
MODEL RISK ADJUSTED HEDGE RATIOS

CAROL ALEXANDER
ANDREAS KAECK*
LEONARDO M. NOGUEIRA

Most option pricing models assume all parameters except volatility are fixed; yet they almost invariably change on re-calibration. This article explains how to capture the model risk that arises when parameters that are assumed constant have calibrated values that change over time and how to use this model risk to adjust the price hedge ratios of the model. Empirical results demonstrate an improvement in hedging performance after the model risk adjustment. © 2009 Wiley Periodicals, Inc. *Jrl Fut Mark* 29:1021–1049, 2009

INTRODUCTION

Model risk refers to the inappropriate use of models for a variety of reasons, such as: unrealistic assumptions, incorrect solution or implementation, inappropriate context, unreliable data, and so forth. For a review of model risk modelling, refer to Derman (1996), Green and Figlewski (1999), Branger and Schlag (2004), Cont (2006), and Psychoyios and Skiadopoulos (2006).

We thank an anonymous referee for excellent comments on an earlier version of this article.

*Correspondence author, ICMA Centre, Henley Business School, University of Reading, Whiteknights, Reading RG6 6BA, UK. Tel: +44-0-118-378-8239, Fax: +44-0-118-931-4741, e-mail: a.kaeck@icmacentre.ac.uk

Received July 2007; Accepted March 2009

-
- *Carol Alexander is at ICMA Centre, The University of Reading, Berkshire, UK.*
 - *Andreas Kaeck is at ICMA Centre, The University of Reading, Berkshire, UK.*
 - *Leonardo M. Nogueira is at Banco Central do Brasil, Brazil and ICMA Centre, The University of Reading, Berkshire, UK.*

The majority of option pricing models allow time variation in only one of the price process parameters, the volatility. The other model parameters are assumed fixed over time and the fact that these parameters usually change on re-calibration is simply ignored in practice. This is one of the main sources of model risk in option pricing models. In this article we show that if the calibrated parameters of a diffusion process change over time in a systematic way then we can use this information to improve the hedging performance of the option pricing model. In other words, we capture and hedge model risk.

After presenting the general framework we consider a specific example within the class of scale-invariant deterministic volatility models.¹ We take a log-normal mixture diffusion as representative of this class and compare its hedging performance, before and after the model risk adjustment, with four other models, each chosen to represent a different hedging alternative. These are the Black–Scholes–Merton model (BSM hereafter, see Black & Scholes, 1973; Merton, 1973), the BSM model with an ad hoc correction for stochastic volatility; the Heston (1993) model with minimum-variance (MV) hedge ratios; and the SABR model (see Hagan, Kumar, Lesniewski, & Woodward, 2002), also with MV hedge ratios.² Without the model risk adjustment the lognormal mixture diffusion performs poorly, but after the adjustment its performance is much improved, although stochastic volatility models still perform slightly better.

The outline of this article is as follows: the section “Related Literature” reviews the related literature on hedging; the section “Adjusting Hedge Ratios for Model Risk” derives the model risk adjustment to the price hedge ratios of a general diffusion model and illustrates the theorem in the context of the log-normal mixture diffusion; the section “Empirical Results” presents the empirical results, which are based on out-of-sample hedging performance of nearly 30,000 observations on S&P 500 index options expiring in December 2007, March 2008, and June 2008; the last section summarizes and concludes.

RELATED LITERATURE

Two main strands of research have been developed in a prolific literature on pricing options: stochastic volatility as in Hull and White (1987), Heston (1993),

¹A price process is scale-invariant (or space-homogeneous) if and only if the marginal distribution of returns is independent of the price scale. This is a reasonable property for *any* option pricing model for tradable assets (Hoogland & Neumann, 2001). When the price of any claim with a pay-off that is homogeneous in the space dimension is observable Alexander and Nogueira (2007a) prove that the partial derivative price hedge ratios—delta, gamma, and higher order—from all scale-invariant models should be identical: they can be expressed in terms of sensitivities that can be readily obtained from the market data by interpolation or some other smoothing technique.

²Minimum-variance hedge ratios differ from the standard partial derivatives of the option price as they take into account any correlation between further state variables, such as stochastic volatility, and the hedging instrument (see e.g. Frey, 1997; Schweizer, 2001).

and many others, where the variance or volatility of the price process is itself stochastic; and local volatility as in Dupire (1994), Derman and Kani (1994), and Rubinstein (1994) where arbitrage-free forward volatilities are “locked-in” by trading options today. Local volatility models are commonly used to hedge over-the-counter options consistently with the observed prices of standard European options through static replication (see e.g. Carr, Ellis, & Gupta, 1998; Derman, Ergener, & Kani, 1995). Following Dumas, Fleming, and Whaley (1998), McIntyre (2001), Brigo and Mercurio (2002), Alexander (2004), and many others, the term “local volatility” has subsequently been extended to cover any deterministic volatility model where forward volatilities are a function of time and asset price. For a detailed review of these models, see Skiadopoulos (2001), Bates (2003), and Fengler (2005).

In comparison with the great proliferation of research on option pricing models there are relatively few studies on their hedging performance. Until recently, research has focused on specific models, such as the Heston (1993) model. Nandi (1998) investigates the importance of the correlation coefficient in the Heston model and concludes that the model’s delta–vega hedging performance is significantly improved when the correlation coefficient is not constrained to be zero. That article also finds that, after taking into account the transactions costs (bid–ask spreads) in the index options market and using S&P 500 futures to hedge, the stochastic volatility model outperforms the Black–Scholes model only if correlation is not constrained to be zero.

Bakshi, Cao, and Chen (1997) compare the hedging performance of models that include stochastic volatility, jumps, and/or stochastic interest rates. They find that a stochastic volatility model such as Heston (1993) is adequate for price hedging. In fact, once stochastic volatility is modelled, the inclusion of jumps leads to no discernible improvement in hedging performance. It is conjectured that this is because the likelihood of a jump during the hedging period is too small, at least when the hedge is rebalanced frequently. Bakshi, Cao, and Chen (2000) consider MV delta hedging and “delta-neutral” hedging (using as many hedging instruments as there are sources of risk, except for jump risk). They find that the inclusion of stochastic interest rates can improve the hedging of long-dated out-of-the-money options, but for other options stochastic volatility is the most important factor to model.

Dynamic delta hedging is possible with local volatility only if the assumption of a fixed local volatility is valid, but this is not justified empirically. The local volatility surface changes considerably on re-calibration and perfect hedging becomes very complex because it involves new sources of randomness. Recognizing this problem, Kani, Derman, and Kamal (1997) propose using “volatility gadgets”, i.e. small portfolios of traded options combined in such a way that it is possible to hedge any specific region from the local volatility

surface. Then, by combining these gadgets, a multitude of perfect hedges is available to the volatility trader. However, in this article we do not focus on perfect hedging.

A few studies of hedging performance with local volatility models have been reported in the academic literature, but their results are controversial. Dumas et al. (1998) test several parametric and semi-parametric forms of the local volatility function, and conclude that BSM deltas appear to be more reliable than any of the local volatility deltas that they tested. McIntyre (2001) reaches a similar conclusion. Hagan et al. (2002) claims that local volatility models produce poor hedge ratios because they predict the wrong dynamics for implied volatility. However, “true” local volatilities have a static smile.³

Finally, Engelmann, Fengler, and Schwendner (2006) test the hedging performance of barrier options under different assumptions for the smile dynamics. They find that delta hedging alone does not lead to satisfying results and that sticky strike models perform best. However, these findings are contradicted by Coleman, Kim, Li, and Verma (2001), who show that BSM deltas are too large for S&P 500 index options and that the average hedging error using a parameterization of the “true” local volatility is smaller than it is using the BSM model. More recently Crépey (2004), Vähämaa (2004), and Alexander and Nogueira (2007a) support this argument and conclude that, on average, “true” local volatility deltas are indeed more effective than BSM deltas.

We should be cautious about these findings because they convey distinct types of model risk. Choosing a good functional form for the local volatility is not straightforward. Several articles base results on floating smile parameterizations of the local volatility surface, and floating smiles are representative of the class of scale-invariant models but not of “true” local volatility. Thus, the parametric or semi-parametric forms used in some of these tests have been inappropriate. When the model’s local volatility surface is a function of the current asset price S_0 as in Hagan et al. (2002) this violates a key assumption of local volatility models, because the surface cannot be static. Secondly, following a result by Coleman et al. (2001), both Crépey (2004), Vähämaa (2004) replace the implied volatility derivative to price by its derivative to strike, i.e. the slope of the smile curve. Yet neither examine the assumptions behind this approximation. Although this “rule-of-thumb” is not necessarily incorrect, neither is its validity obvious given the result of Bates (2005), who proves that for all scale-invariant models the two sensitivities have opposite signs. Thirdly, and this is the point addressed in this article, virtually all these studies implicitly assume

³By “true” local volatility we mean the local volatility surface that is calibrated to option prices using Dupire’s equation. Equivalently, the local variance is the conditional expectation of a stochastic volatility, given the underlying price at some time in the future.

the model is correct and that all the parameters except volatility are constant. However, if the calibrated values of parameters that are assumed constant in the model change when the model is re-calibrated, their hedge ratios are inconsistent with the actual dynamics of the model's local volatility surface. All together, these three key points could explain the contradictory conclusions outlined above.

ADJUSTING HEDGE RATIOS FOR MODEL RISK

In this section we consider an extension of the stochastic local volatility (SLV) model introduced by Alexander and Nogueira (2004). The formulation used in this article allows one to compare hedge ratios before and after an adjustment for uncertainty in calibrated parameters. We first introduce the general theory and then we illustrate our theorem by deriving explicit formulae for the model risk adjusted hedge ratios in a lognormal mixture diffusion.

General Framework

Many models with vanilla call and put prices that are consistent with observed market prices have been proposed in the literature. Yet both the "true" model and its parameters are unknown. It is a common practice to pick a particular pricing model, to recalibrate its parameters frequently, and to ignore the fact that assumed constant parameters change at each recalibration. In the following we specify a general diffusion process with constant drift and non-constant volatility, distinguishing between parameters such as spot volatility, which are time-varying, and parameters that are assumed to be constant.

Assume the option pricing model used specifies the asset price process as a continuous Itô process under the risk-neutral measure:

$$\frac{dS_t}{S_t} = (r - q)dt + \sigma(t, S_t, \xi_t, \lambda)dB_t, \quad (1)$$

where the instantaneous volatility $\sigma(t, S_t, \xi_t, \lambda)$ may be stochastic or deterministic. The risk-free rate r and dividend yield q are assumed constant. The $(h \times 1)$ vector ξ_t collects all time-varying parameters, whereas the $(n \times 1)$ vector λ consists of all parameters that are assumed constant within the model. For instance, standard stochastic volatility models only allow the spot volatility (or equivalently variance) to move over time, whereas all other diffusion parameters of the volatility process are assumed constant.

In this framework, model risk arises when the calibrated values of λ change over time, so we capture this risk by modelling the evolution of the vector λ .

We assume that the risk-neutral dynamics of the parameters $\lambda_t = (\lambda^1, \lambda^2, \dots, \lambda^n)_t'$ are governed by general diffusion processes that can be correlated with the asset price S_t and with each other:

$$d\lambda_t^i = \alpha^i(t, \lambda_t^i)dt + \beta^i(t, \lambda_t^i)dW_t^i, \quad (2)$$

$$dW_t^i = \rho^{i,S}dB_t + \sqrt{1 - (\rho^{i,S})^2}dZ_t^i, \quad (3)$$

with the following quadratic covariations being zero

$$\langle dZ_t^i, dZ_t^j \rangle = 0 \quad \forall i \neq j, \quad \langle dZ_t^i, dB_t \rangle = 0 \quad \forall i. \quad (4)$$

Here $\rho^{i,S}$ is the correlation between variations in λ_t^i and S_t . Note that our assumptions imply that the correlation between parameter i and j equals $\rho^{i,S}\rho^{j,S}$ and thus the correlation structure between parameters in λ is indirectly introduced via their dependence on S . The dynamics (1) and (2) are assumed to satisfy the regularity conditions for an Itô process. Jumps in the instantaneous volatility, if any, could be modelled by adding Poisson jumps to the dynamics of the parameters of the local volatility model. However, in this article we shall confine our discussion to diffusion models only.

With this formulation of stochastic parameters, recalibration over time is in accordance with the model and will reveal time-series values for ξ_t and λ_t . These parameters can be used to obtain vanilla option prices in model (1), regardless of the values for α^i and β^i .⁴ However, the dynamics of calibrated parameters are essential for hedging even for vanilla options. If calibrated parameters in λ_t are time-varying the hedge ratios derived from a model that assumes they are constant will be incorrect.

In the following we shall consider how to adjust the price hedge ratios, i.e. those that cancel the exposure to dS and $(dS)^2$, when some assumed constant model parameters are in fact stochastic. Then the delta and gamma of a general claim, derived under the constant parameter assumption, require an adjustment according to the following theorem.⁵

Theorem 1: In the model (1) with parameter dynamics (2)–(4) the MV delta $\delta(t, S_t, \xi_t, \lambda_t)$ and gamma $\gamma(t, S_t, \xi_t, \lambda_t)$ of a contingent claim with price $f \equiv f(t, S_t, \xi_t, \lambda_t)$ are given by

⁴This is similar to knowing the value for the implied volatility of a standard European option. We can simply use implied volatilities in the BSM pricing formula to obtain the option price, irrespective of any stochastic dynamics for implied volatilities.

⁵We use subscripts to indicate partial derivatives except for subscript t , which is used as the time index for stochastic processes.

$$\delta(t, S_t, \xi_t, \lambda_t) = f_S + \sum_{j=1}^h f_{\xi^j} X^{\xi^j, S} + \sum_{i=1}^n f_{\lambda^i} X^{\lambda^i, S} = \delta(t, S_t, \xi_t, \lambda) + \sum_{i=1}^n f_{\lambda^i} X^{\lambda^i, S} \quad (5)$$

$$\begin{aligned} \gamma(t, S_t, \xi_t, \lambda_t) &= \gamma(t, S_t, \xi_t, \lambda) + \sum_{j=1}^h \left\{ X^{\xi^j, S} \sum_{l=1}^n f_{\xi^j \lambda^l} X^{\lambda^l, S} + f_{\xi^j} \sum_{l=1}^n X_{\lambda^l}^{\xi^j, S} X^{\lambda^l, S} \right\} \\ &+ \sum_{i=1}^n \left\{ X^{\lambda^i, S} \left(2f_{\lambda^i S} + \sum_{k=1}^h f_{\lambda^i \xi^k} X^{\xi^k, S} + \sum_{l=1}^n f_{\lambda^i \lambda^l} X^{\lambda^l, S} \right) \right\} \\ &+ \sum_{i=1}^n \left\{ f_{\lambda^i} \left(X_S^{\lambda^i, S} + \sum_{k=1}^h X_{\xi^k}^{\lambda^i, S} X^{\xi^k, S} + \sum_{l=1}^n X_{\lambda^l}^{\lambda^i, S} X^{\lambda^l, S} \right) \right\}, \end{aligned} \quad (6)$$

where

$$\begin{aligned} X^{\lambda^i, S} &= \frac{\langle d\lambda_t^i, dS_t \rangle}{\langle dS_t, dS_t \rangle} = \frac{\beta^i(t, \lambda_t^i) \rho^{i, S}}{\sigma(t, S_t, \xi_t, \lambda_t) S_t}, \\ X_S^{\lambda^i, S} &= -\frac{\beta^i(t, \lambda_t^i) \rho^{i, S} [\sigma(t, S_t, \xi_t, \lambda_t) + \sigma_S(t, S_t, \xi_t, \lambda_t) S_t]}{(\sigma(t, S_t, \xi_t, \lambda_t) S_t)^2}, \\ X_{\lambda^i}^{\lambda^i, S} &= \frac{\beta_{\lambda^i}^i(t, \lambda_t^i) \rho^{i, S} \sigma(t, S_t, \xi_t, \lambda_t) S_t - \beta^i(t, \lambda_t^i) \rho^{i, S} \sigma_{\lambda^i}(t, S_t, \xi_t, \lambda_t) S_t}{(\sigma(t, S_t, \xi_t, \lambda_t) S_t)^2}, \end{aligned}$$

and $\delta(t, S_t, \xi_t, \lambda)$ and $\gamma(t, S_t, \xi_t, \lambda)$ are the sensitivities derived from the standard view in which λ is assumed constant.

In effect, each hedge ratio now consists of two parts: the sensitivity derived from the standard view (i.e. calibrated to option prices at a fixed point in time) and an adjustment factor that depends on the dynamics of the parameters λ . Note that any stochastic movements of parameters in ξ are already taken into consideration in the derivation of $\delta(t, S_t, \xi_t, \lambda)$. For instance, in the stochastic volatility model of Heston (1993) an adjustment of the hedge ratio due to stochastic volatility is already accounted for and we only need to consider the effect of time-varying parameters in λ such as speed of mean reversion or long-term variance.

The adjustment is driven by the variance of the calibrated parameters and their correlation with the underlying. Considering the special case of $\beta^i(t, \lambda_t^i) = 0$, i.e. deterministic time-varying parameters, we recover the hedge ratios for the standard model. Hence, an adjustment is only necessary when the calibrated parameters are stochastic. Likewise, for the case $\rho^{i, S} = 0$ no adjustment is necessary. This is because we cannot hedge any movements in the parameters if they are not correlated with the hedging instrument. Hence, an adjustment will only be fruitful if the calibrated parameters and the underlying exhibit a systematic correlation structure. If not, standard hedge ratios will be efficient even in the presence of this type of model risk.

A major advantage of the formulation of the model is that its calibration can be broken into two parts: a calibration of the model, assuming constant parameters in λ , to market prices of standard European calls and puts, and a volatility and correlation calibration based on a discrete time-series analysis of calibrated parameters. In the first step, the calibration of the model will provide values of the parameters at the time of calibration t_0 , which are a priori unknown as they are assumed to vary stochastically over time. In the second step we can use the calibrated values at m points prior to t_0 to estimate β^i and the correlations $\rho^{i,S}$ between the underlying S and the parameters λ^i . This way, we include information in the calibrated parameters to account for any systematic behavior in their time series.

Hence, we assume (2) and (4) also hold over a small time-step Δt , thus

$$\begin{aligned} \text{Var}(\Delta\lambda^i) &= \beta^i(t, \lambda_t^i)^2 \Delta t, \\ \text{Cov}(\Delta\ln(S_t), \Delta\lambda^i) &= \sigma(t, S_t, \xi_t, \lambda_t) \beta^i(t, \lambda_t^i) \rho^{i,S} \Delta t \end{aligned}$$

for $1 \leq i \leq n$. For example if the dynamics of the parameters are assumed to be driven by arithmetic Brownian motions we obtain

$$\beta^i = \sqrt{\frac{\text{Var}(\Delta\lambda^i)}{\Delta t}} \quad \text{and} \quad \rho^{i,S} = \frac{\text{Cov}(\Delta\ln(S_t), \Delta\lambda^i)}{\sigma(t, S_t, \xi_t, \lambda_t) \sqrt{\text{Var}(\Delta\lambda^i)}}.$$

Thus, to calibrate the model at time t_0 we utilize the sample covariance matrix of the parameter matrix \mathbf{M} , where $\mathbf{M} = [\Delta\ln(S), \Delta\lambda^1, \Delta\lambda^2, \dots, \Delta\lambda^n]$ is the $m \times (n + 1)$ matrix of variations in each of the risk factors. This procedure has the benefit of being determined by the stability of the calibration. That is, the more stable the model parameters, the more accurate these approximations. Another possibility, not applied in this article, is to use ARIMA models to find the best discrete model for each parameter, and take the limit as the time interval tends to zero to obtain the continuous-time dynamics.⁶

Example: Lognormal Mixture Diffusion Model

In this section we consider a lognormal mixture diffusion with two constant volatility components and constant non-zero means as an example. This particular model is a special case of the lognormal mixture diffusion of Brigo, Mercurio, and Sartorelli (2003) and we have chosen it to illustrate the model risk adjustment because it is representative of the class of scale-invariant deterministic volatility models.

⁶See Hafner (2004) for an application of this technique to the parameters of the DAX implied volatility surface.

The model assumes that the log of the asset price has a marginal density at some future time T given by

$$g^S(x_T) = p\varphi(x_T, (r - q + s_1 - 0.5\sigma_1^2)T, \sigma_1^2T) + (1 - p)\varphi(x_T; (r - q + s_2 - 0.5\sigma_2^2)T, \sigma_2^2T),$$

where $x_T = \ln(S_T)$ and $\varphi(x; \mu, \sigma^2)$ is the normal density. In effect this model has four parameters: a volatility for each of the two states (σ_1 and σ_2), a shift in the drift of the first component s_1 and a mixing law defined by $p \in [0, 1]$. The shift in the second component s_2 is not a free parameter, rather it is fully determined by the choice of $\lambda = \{p, \sigma_1, \sigma_2, s_1\}$. This is because in order to preclude arbitrage opportunities, the drift of the asset price must be equal to the risk-free rate and this condition imposes a restriction on s_2 . Brigo et al. (2003) prove that the dynamics implied by the mixture of two lognormal densities are unique and can be written as

$$dS_t = (r - q)S_t dt + \sigma(t, S_t, \lambda)S_t dB_t, \tag{7}$$

where $\sigma(t, S_t, \lambda)$ is a deterministic function of time, the underlying and the model parameters, which are assumed constant.⁷ None of the model parameters is assumed to move stochastically, therefore ξ is empty and can be omitted.

The lognormal mixture price of a standard option is a linear combination of BSM prices and is therefore straightforward to implement. Following the theory of the last section, we can first calibrate this model to quoted call and put prices at several consecutive points in time. The frequent recalibration will lead to time series of calibrated parameters, which are subsequently used to estimate β^i and $\rho^{i,S}$ for all model parameters. The calculation of the instantaneous volatility function requires the examination of the limit of $\sigma(t, S_t, \lambda)$ as t approaches t_0 , the time of calibration. The result and its partial derivatives are stated in the next corollary.⁸

Corollary: The instantaneous volatility at calibration time t_0 in model (7) is given by

$$\sigma(t_0, S_{t_0}, \lambda) = \sqrt{\frac{p\sigma_1 + (1 - p)\sigma_2}{p\sigma_1^{-1} + (1 - p)\sigma_2^{-1}}}. \tag{8}$$

Furthermore, its partial derivatives follow directly

⁷For the exact functional form see Brigo et al. (2003) or the Appendix for its formulation in terms of futures prices.

⁸Note that (8) does not depend on the shifts s_1 and s_2 . In fact, the instantaneous volatility at time t_0 is defined solely in terms of the variances of the component processes, regardless of their means and of the value of the asset price S_{t_0} .

$$\begin{aligned}\sigma_p(t_0, S_{t_0}, \lambda) &= \eta^{-1}[\sigma_1\sigma_2(\sigma_1^2 - \sigma_2^2)], \\ \sigma_{\sigma_1}(t_0, S_{t_0}, \lambda) &= \eta^{-1}[p\sigma_2(\sigma_1^2 + \sigma_2^2 - p\sigma_1^2 - \sigma_2^2p + 2p\sigma_1\sigma_2)], \\ \sigma_{\sigma_2}(t_0, S_{t_0}, \lambda) &= \eta^{-1}[(1-p)\sigma_1(p\sigma_2^2 - 2p\sigma_1\sigma_2 + 2\sigma_1\sigma_2 + p\sigma_1^2)],\end{aligned}$$

where

$$\eta = 2\sqrt{\frac{(p\sigma_1 + (1-p)\sigma_2)\sigma_1\sigma_2}{p\sigma_2 + (1-p)\sigma_1}}(p\sigma_1 - (1-p)\sigma_2)^2.$$

Proof: The proof for the limit of the instantaneous volatility in Equation (8) is given in Nogueira (2006). The derivatives immediately follow from (8).

It remains to derive first and second order partial derivatives of the claim price with respect to all model parameters. These derivatives are straightforward to calculate as most follow directly from BSM hedge ratios.⁹ This information can now be used to adjust the hedge ratios according to (5) and (6). Typically, for equity index options, the adjusted delta will be lower than the unadjusted delta. Later we show that this is due to the fact that the sum of products $f_{\lambda}X^{\lambda,S}$ is negative. For example, the derivative f_{σ_1} resembles the BSM vega, which is positive for a long position, and combined with the negative correlation $\rho^{i,S}$ that we find empirically in equity markets, the adjustment term becomes negative. Similar arguments hold for the other parameters. The adjusted gammas are usually higher for low moneyness options and lower for high moneyness options. The typical shapes for fixed maturity equity call option deltas and gammas, as a function of moneyness are depicted in Figure 1.

EMPIRICAL RESULTS

In this section we set up a hedging race between popular option pricing models and compare their performance with the performance of the model risk adjusted hedge ratios proposed in the last section. To this end, we first describe our data set and the pricing models that are used in this study.

Data

We have obtained data from Bloomberg on December 2007, March 2008, and June 2008 European put and call options on the S&P 500 spot index from June 1, 2007 until December 18, 2007 (142 business days).¹⁰ Our data set includes

⁹The results are given in the Appendix, and are again phrased in terms of the futures price process for simplicity.

¹⁰This sample period is interesting because it includes the turbulent phase of the sub-prime mortgage crisis. During the sample period, the spot index was bounded between 1,407 and 1,565 index points with an annualized standard deviation (volatility) of approximately 20%.




FIGURE 1

Delta and Gamma with and without model risk adjustment. This figure depicts the differences between deltas and gammas of the unadjusted and adjusted bivariate normal mixture model.

numerous strikes and the maturities of the options range from three days up to almost one year. In order to assure that all options in the sample have a reasonable level of liquidity only options with moneyness $M = K/F$ (where K denotes the strike and F the futures price) between 0.8 and 1.2 are utilized. Also all options whose prices lie outside their arbitrage bounds are discarded. After applying these filters, our sample consists of 27,580 observations. Time series of daily USD Libor rates were downloaded from the British Bankers Association (BBA) website for several maturities and used as a proxy for the risk-free rate. Linear interpolation was applied to produce a continuous function of the Libor rate with respect to time to maturity. This procedure was repeated for every trading day in the sample.

Index options typically face the challenge of estimating the future dividend yield over the life of the option. This estimation is circumvented by interpreting the spot options as options on the futures. As S&P 500 spot options are European style, their payoff at expiration equals the payoff of a European futures option, where the future has the same maturity as the option. Consequently, assuming the futures price F is at its fair value we can use it as the underlying process in our empirical study. Employing futures option pricing formulae also has the advantage of directly providing hedge ratios for the futures. As the S&P spot index is not tradeable, we cannot set up a realistic hedge portfolio by buying or selling the spot asset and hedging arguments ought to be based on the futures. Corresponding futures prices are also obtained from Bloomberg.

Models

We utilize six different pricing models in our empirical hedging study. We first describe the models before discussing their price hedge ratios. All models are expressed in terms of the futures prices F and their dynamics and densities are given under the risk-neutral measure.

(i) *BSM*: The Black–Scholes–Merton model (see Black & Scholes, 1973; Merton, 1973). We use the formulation for futures options due to Black (1976) where futures follow a geometric Brownian motion.

$$dF_t = \sigma F_t dW_t.$$

(ii) *BSM-A*: The Black–Scholes–Merton model where implied volatility is correlated with the underlying asset (see the hedge ratios below for further description).

(iii) *Heston*: The stochastic volatility model of Heston (1993):

$$\begin{aligned} dF_t &= \sqrt{V_t} F_t dW_t, \\ dV_t &= a(m - V_t)dt + b\sqrt{V_t} dZ_t, \end{aligned}$$

with price-variance covariation given by $\langle dW_t, dZ_t \rangle = \rho dt$.

(iv) *SABR*: The “stochastic- $\alpha\beta\rho$ ” model of Hagan et al. (2002):

$$\begin{aligned} dF_t &= \alpha_t F_t^\beta dW_t, \\ d\alpha_t &= \nu \alpha_t dZ_t \end{aligned}$$

with price-volatility covariation given by $\langle dW_t, dZ_t \rangle = \rho dt$.

(v) *NM*: The lognormal mixture diffusion of Brigo et al. (2003), in which the log price of the futures has marginal density at some future time T given as

$$g^F(x_T) = p\varphi(x_T, (s_1 - 0.5\sigma_1^2)T, \sigma_1^2 T) + (1 - p)\varphi(x_T, (s_2 - 0.5\sigma_2^2)T, \sigma_2^2 T),$$

where $x_T = \ln(F_T)$, $\varphi(x; \mu, \sigma^2)$ is the normal density and the shift

$$s_2 = T^{-1}[\ln(1 - pe^{s_1 T}) - \ln(1 - p)]$$

is given by no-arbitrage conditions.

(vi) *NM–SLV*: The adjusted version of the NM model as described in the previous section. This model allows the parameters σ_1 , σ_2 , p , and s_1 to be stochastic.

We could have included a variety of other models but we believe that the models outlined above are sufficiently representative of the main types of

option pricing models: the BSM model as a benchmark, a scale-invariant stochastic volatility model (Heston), a potentially non-scale-invariant stochastic volatility model (SABR),¹¹ and a scale-invariant deterministic volatility model and its adjusted version (NM and NM–SLV). Models (iii)–(v) have been selected because they have a good fit to observed option prices of individual maturities and are tractable, with closed form solutions for vanilla options and hedge ratios. Model (vi) is the only model that accounts for the parameter instability addressed in this article.

The main reason for choosing a finite mixture model as the starting point for the adjustment of the hedge ratios is an apparent contradiction in the literature: recent empirical evidence on lognormal mixture diffusions has found that these models are good for forecasting prices but not for delta hedging, and even the BSM delta may perform better (Gemmill & Saflekos, 2000; Kim & Kim, 2003; Wilkens, 2005). However, no theoretical explanation has been presented to justify their poor hedging performance. This section demonstrates that the hedging performance of these models may be substantially improved using the adjusted hedge ratios proposed in the section “Adjusting Hedge Ratios for Model Risk.”

The delta and gamma hedge ratios for each model are discussed below. For expositional clarity, calibration time is set to $t_0 = 0$.

(i) *BSM*: The BSM model deltas and gammas are obtained using the market implied volatility $\theta(K, T)$ of each vanilla option with strike K and maturity T , and setting $\omega = \{1 \text{ for calls, } -1 \text{ for puts}\}$:

$$\delta^{BSM}(K, T) = \omega e^{-rT} \Phi(\omega d_1(0, \sigma)),$$

$$\gamma^{BSM}(K, T) = \frac{e^{-rT} \varphi(d_1(0, \sigma))}{F_0 \theta(K, T) \sqrt{T}},$$

where φ and Φ are the standard normal density and cumulative density function, respectively, and

$$d_1(\mu, \sigma) = \frac{\ln(F_0/K) + (\mu + 0.5\sigma^2)T}{\sigma \sqrt{T}}. \tag{9}$$

(ii) *BSM-A*: If implied volatility is itself a function of the forward price F , then the first two derivatives of the option price with respect to F can be written as

$$\delta^{BSM-A}(K, T) = \delta^{BSM}(K, T) + \nu^{BSM}(K, T)\theta_F,$$

$$\gamma^{BSM-A}(K, T) = \gamma^{BSM}(K, T) + \nu^{BSM}(K, T)\theta_{FF}$$

$$+ 2\nu a^{BSM}(K, T)\theta_F + \nu o^{BSM}(K, T)(\theta_F)^2,$$

¹¹The SABR model is scale-invariant only if $\beta = 1$.

where ν^{BSM} is the BSM vega, $\nu\sigma^{BSM}$ is the second partial derivative of the option price with respect to implied volatility (volga) and $\nu\alpha^{BSM}$ is the second cross derivative with respect to spot and volatility (vanna). Following Coleman et al. (2001) we replace θ_F with θ_K to obtain an ad hoc correction of the BSM model. As the smile in equity markets is usually downward sloping this will adjust the BSM delta downward and therefore this ad hoc approach is consistent with the negative correlation, which is often observed between implied volatilities and the underlying (see Vähämaa, 2004).

(iii) *Heston*: For the Heston model we use MV hedge ratios. For the derivations of these hedge ratios see Alexander and Nogueira (2007a).

$$\delta^{Heston}(K, T) = f_F + f_V \frac{\rho b}{F_0},$$

$$\gamma^{Heston}(K, T) = f_{FF} + \frac{\rho b}{F_0} \left(2f_{FV} + \frac{\rho b}{F_0} f_{VV} - \frac{1}{F_0} f_V \right),$$

with $f \equiv f^{Heston}(K, T)$

(iv) *SABR*: The MV hedge ratios are (see Alexander & Nogueira, 2007b)¹²

$$\delta^{SABR}(K, T) = f_F + f_\alpha \frac{\rho v}{F_0^\beta},$$

$$\gamma^{SABR}(K, T) = f_{FF} + \frac{\rho v}{F_0^\beta} \left(2f_{F\alpha} + \frac{\rho v}{F_0^\beta} f_{\alpha\alpha} - \frac{\beta}{F_0} f_\alpha \right),$$

with $f \equiv f^{SABR}(K, T)$

(v) *NM*: The NM deltas and gammas are linear combinations of BSM deltas and gammas, each based on a lognormal density with different means and variances (see Brigo et al., 2003):

$$\delta^{NM}(K, T) = pe^{(s_1-r)T} \omega \Phi(\omega d_1(s_1, \sigma_1)) + (1 - p)e^{(s_2-r)T} \omega \Phi(\omega d_1(s_2, \sigma_2)),$$

$$\gamma^{NM}(K, T) = pe^{(s_1-r)T} \frac{\varphi(d_1(s_1, \sigma_1))}{\sigma_1 \sqrt{TF_0}} + (1 - p)e^{(s_2-r)T} \frac{\varphi(d_1(s_2, \sigma_2))}{\sigma_2 \sqrt{TF_0}},$$

where d_1 is given in Equation (9).

(vi) *NM-SLV*: As described in Theorem 1 and the section “Example: Lognormal Mixture Diffusion Model” after, setting $f \equiv f^{NM}(K, T)$ and $\lambda = \{s_1, p, \sigma_1, \sigma_2\}$.

¹²Unlike the Heston model, the SABR model is not scale-invariant and therefore its partial derivative price hedge ratios can be different from those of the Heston or NM model. However, as in the Heston model, minimum-variance hedge ratios are likely to reduce hedging errors because indirect effects from changes in the volatility are taken into account.

Model Calibration

All models are calibrated daily by minimizing the root-mean-squared error (RMSE) between model and market prices for each maturity. As the adjusted BSM hedge ratios are based on the first and second derivative of the implied volatility with respect to strike K , we fit a cubic polynomial to the observed implied volatility function and use the estimated parameters to calculate partial derivatives with respect to K . The calibration of the polynomial function is performed on prices rather than directly on implied volatilities in order to ensure that the calibration objective is the same for all models under consideration (for a similar application and a discussion on different objective functions see Christoffersen & Jacobs, 2004). The SABR model is calibrated with the restriction $\beta = 0.5$. As already remarked by Hagan et al. (2002), β cannot be adequately determined by fitting the model to market prices and its value should be chosen before calibration. Our motivation of using $\beta = 0.5$ is two-fold. Firstly, the model becomes non-scale-invariant and can therefore lead to partial derivative deltas and gammas that are different from the NM and Heston models, and secondly the stochastic process for the futures price remains non-negative. In addition, we fix the long-term volatility m in the Heston model to 14%. Our choice for the long-term volatility is motivated by the comprehensive study conducted by Eraker, Johannes, and Polson (2003).

All models produce an adequate fit to our market data. With an average RMSE of 0.57\$ the NM model exhibits the highest error throughout and the Heston model tends to fit market prices best (with an average RMSE of 0.27\$). The RMSE for the cubic polynomial is 0.28\$ and we obtain 0.38\$ for the SABR model. The main advantage of the stochastic volatility models over the NM model stems from their ability to reproduce more accurate market prices for far out-of-the-money and far in-the-money options. Contrary to the other models, the model implied volatility for the NM model tends to “flatten out” in the extremes, where the influence of the low volatility component becomes negligible. For at-the-money options their fit is nearly indistinguishable from that of the other models. Comparing the small pricing errors with an average bid–ask spread of 2.55\$ we deduce that the fit of all models to market prices is similar enough to preclude differences in the fit to have a significant effect on the models’ hedging performance.¹³

The dynamics of the calibrated parameters for the NM model are shown in Figure 2. The parameters for different maturities tend to move in the same direction, but for short maturity options the calibrations become unstable in

¹³Our empirical results on the performance of the partial derivative delta and gamma for the NM and the Heston model (unreported) show that their hedging performance is almost indistinguishable. As both models are scale-invariant, differences in their hedging performance can only be attributed to their fit of market prices. Therefore, we can safely conclude that calibration error does not work in favor of the Heston model.




FIGURE 2

Evolution of the model parameters for the Normal Mixture model. The black line depicts the December 2007 options, the light gray line the March 2008 options and the dark gray line shows the June 2008 options.

December 2007 when options are within less than three weeks to expiry. This finding is consistent with previous empirical evidence that diffusion-only models may be insufficient to explain short-term option prices (Andersen, Benzoni, & Lund, 2002; Bakshi et al., 1997).

Figure 3 shows that not only the NM parameters themselves, but also their correlations with the underlying futures prices are relatively stable over time. A positive relationship is observed between the variations in the calibrated values of p and F and a negative relationship between the futures price and all other parameters (namely σ_1 , σ_2 , and s_1) is evident.¹⁴ The NM model parameter correlations with the futures price are shown in the upper graph. For these, statistical tests on the whole sample reject the null hypothesis of zero correlation for all parameters and therefore our results suggest that we can extract further information from the time-series properties of the calibrated parameters to adjust the price hedge ratios in the NM model.

By contrast, for the Heston model parameter correlations with the futures price, shown in the lower graph of Figure 3, exhibit no clear pattern, and they

¹⁴These correlations are based on the previous four weeks of calibrated parameters.




FIGURE 3

Evolution of the correlations for the Normal Mixture (upper plot) and Heston model parameters (lower plot) with the underlying futures.

exhibit both positive and negative values within the sample period. The same tests only indicate a significant relationship between the spot volatility and the changes in the futures. However, as this correlation is already incorporated into the Heston model no adjustment is necessary.

Hedging Study

In order to examine the hedging performance we consider both delta and delta–gamma-hedging strategies. Our objective for the delta hedge is to remove any exposure of the hedged portfolio with respect to movements in the stock price dS . This hedge is also locally risk minimizing, because we only admit strategies based on the underlying and the risk-free investment. The gamma hedge aims to remove all quadratic exposure, i.e. all risk arising from terms of order $(dS)^2$. A delta–gamma hedge is effected by hedging with another option on the same underlying and then delta hedging the portfolio. This is desirable because delta hedging in discrete time requires frequent rebalancing and

adding a gamma hedge will diminish the need to trade too frequently. However, our delta–gamma hedging is not locally risk-minimizing as we focus on price hedging only; uncertainties from other risk factors are only taken into account if they exhibit correlations with the hedging instrument. One could consider alternative strategies that include trading another option on the same underlying. For example a delta–vega hedge in a stochastic volatility model locally removes all the risk from the portfolio.¹⁵ However defining a vega is not straightforward for all the models used in this study and therefore we focus on price hedging only.¹⁶

The delta-hedge strategy consists of a delta-hedged short position on each option (including both puts and calls), rebalanced daily. That is, one option on each strike is sold at every trading day in the sample and hedged by buying an amount $\delta(K, T)$ of the respective futures contract, where $\delta(K, T)$ is determined by both the model's and the option's characteristics. In addition we form a zero-cost delta hedge portfolio by investing in the risk-free interest rate.¹⁷ On the next trading day, we determine the P&L of the hedge of each option individually. The delta–gamma-hedge strategy again consists of a shorted option, but this time a second option has to be bought to make the portfolio both delta and gamma neutral. For the gamma hedge we use the option closest to ATM and form a zero-cost portfolio as before. This option-by-option strategy on a large and complete data set of liquid options allows one to assess the effectiveness of hedging by maturity or moneyness of the option, and day-by-day as well as over the whole period.

Table I reports the sample statistics of the daily P&L for each model aggregated over all options and all days in the hedging sample. Our main focus lies in the standard deviation of the daily P&L, because to minimize this is the prime objective of hedging. Another important performance criterion is that the P&L be uncorrelated with the underlying asset. In our case over-hedging would result in a significant positive correlation between the hedge portfolio and the S&P 500 index return. We have therefore performed a linear regression, based on all 27,580 P&L data points where the P&L of each option for the day is explained by a quadratic function of the S&P 500 futures returns x_t :

$$P \& L_t(K, T) = a_0 + a_1 x_t + a_2 x_t^2 + \varepsilon_t \quad (10)$$

with $\varepsilon_t \sim \Phi(0, \sigma_\varepsilon^2)$. The R^2 from this regression should be small as we hope to remove linear and quadratic exposure by using a delta and delta–gamma hedge, respectively. The results are reported in the last column for each model.

¹⁵For the BSM model, delta–gamma and delta–vega hedging produce the same hedge portfolio; however, this property is not shared by other pricing models.

¹⁶For example, in the NM model, the definition of vega is not straightforward as there is no single volatility parameter on which a vega could be defined.

¹⁷This strategy is also used by Bakshi et al. (1997) and Coleman et al. (2001).

TABLE I
Hedging P&L

	<i>Mean</i>	<i>Std. Dev.</i>	<i>Skewness</i>	<i>XS Kurtosis</i>	<i>R²</i>
<i>Delta hedging</i>					
SABR $\beta = 0.5$ (MV)	-0.010	1.343	-0.156	1.675	0.161
BSM-A	-0.016	1.359	-0.247	1.895	0.148
Heston (MV)	-0.047	1.400	-0.554	1.964	0.208
Normal mixture (Adj)	-0.031	1.620	-0.557	2.304	0.208
BSM	-0.072	1.784	-0.924	2.280	0.472
Normal mixture	-0.118	2.891	-0.902	2.525	0.676
<i>Delta-gamma hedging</i>					
Heston (MV)	0.017	0.796	-0.140	4.580	0.017
SABR $\beta = 0.5$ (MV)	0.021	0.807	0.060	5.810	0.006
BSM-A	0.024	0.843	-0.117	5.874	0.037
BSM	0.009	0.860	-0.261	3.342	0.034
Normal mixture (Adj)	0.012	1.070	-0.626	9.206	0.041
Normal mixture	-0.022	1.661	-0.100	4.200	0.009

Note. This table reports descriptive statistics for the hedging P&L for all models used in this study. The statistics are based on the whole sample and includes 27,580 observations.

In the delta-hedge strategy, shown in the upper part of Table I, the standard deviation of the hedged portfolio P&L is the lowest for the SABR model, followed by the adjusted BSM, the Heston, the adjusted NM model, and the BSM model. A similar ranking of the models is obtained by using the R^2 measure, only now the adjusted BSM is outperforming the SABR model according to this criterion. The unadjusted NM model is by far the worst model, with a standard deviation considerably higher than the standard deviation of all the other models (and in particular the adjusted NM model).

The poor hedging performance of the NM model indicates that although the model is capable of fitting option prices well, it does not capture the correct dynamics of the underlying and the volatility surface. The adjustment to the delta hedge corrects for these inappropriate dynamics, and it improves the hedging performance considerably. As a result, the standard deviation of the adjusted delta-hedged portfolio is almost one-half of the standard deviation of the unadjusted delta-hedged portfolio.

For delta-gamma hedging the ranking remains more or less the same as in the delta-hedging strategy. The performance of the NM model is still by far the worst and again the adjustment to the hedge ratios leads to a remarkable reduction in the standard deviation of the hedging error. The BSM model shows considerable improvement when a gamma hedge is added and its performance is close to the adjusted hedge ratios. This is in line with earlier research finding that it is difficult to beat the BSM model once an option is included to hedge




FIGURE 4

Delta hedge average P&L (X axis) versus S&P 500 futures returns (Y axis). The gray line marks the two standard deviation band of the regression, where the standard deviation is obtained from the estimation of σ_{ε}^2 in Equation (10).

either gamma or vega risk (see for example Bakshi et al., 1997). However, the R^2 for both the adjusted and unadjusted BSM and the adjusted NM model remains higher than the R^2 for all the other models, indicating that these hedge ratios are not fully capturing the gamma effects.

The P&L quadratic regression results depicted in Figures 4 and 5 agree with the findings of Table I. As we only rebalance daily, we expect the delta hedge average P&L, shown in Figure 4, to exhibit a quadratic relationship with returns due to the well-known gamma effect in delta-hedged portfolios

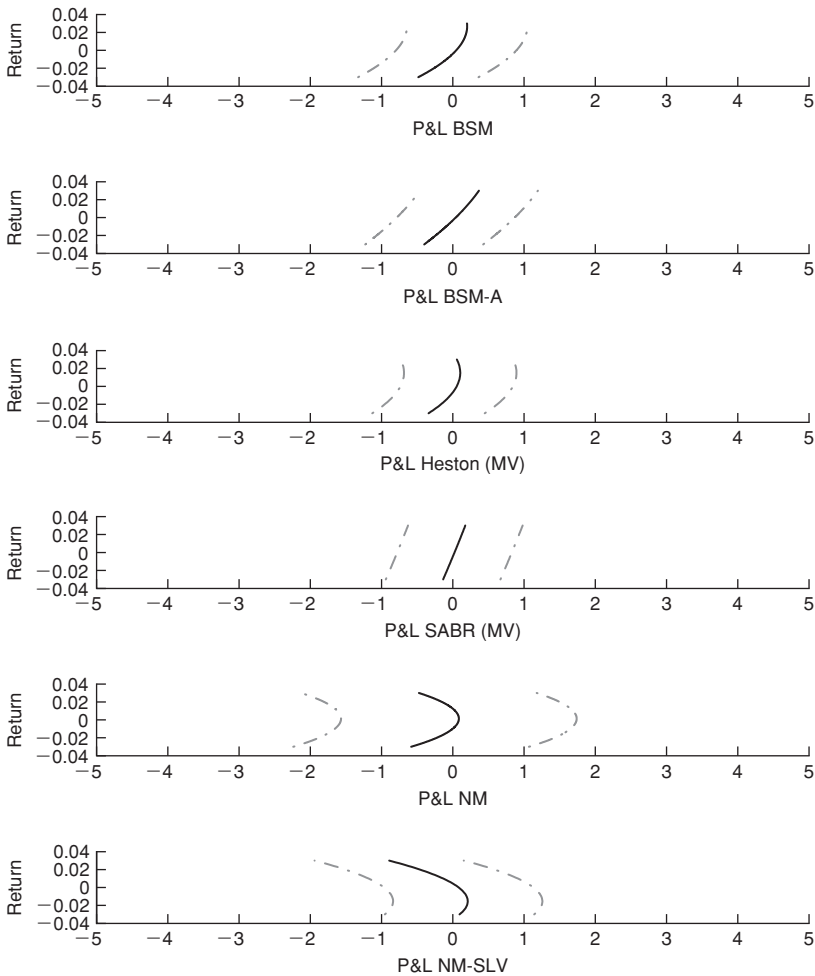


FIGURE 5

Delta–gamma hedge average P&L (X-axis) versus S&P 500 futures returns (Y-axis). The gray line marks the two standard deviation band of the regression, where the standard deviation is obtained from the estimation of σ_g^2 in Equation (10).

with discrete rebalancing. Indeed, in our regressions the coefficient for the squared returns is significant and stable across different models. Yet if the delta hedge had succeeded, the fitted values in Figure 4 would be a symmetric quadratic function of the futures return, and this is not observed in any of the models. There is a significant and positive relationship with returns in all models, which is evidence of over-hedging, and this is especially pronounced in the NM model before the model risk adjustment. This strong positive relationship is no longer evident in Figure 5, and a comparison of the graphs for different models suggests that the SABR model is the best at hedging gamma effects.

TABLE II
Standard Deviation of Hedging Error

	<i>All</i> (%)	<i>0.80–0.90</i> (%)	<i>0.90–0.97</i> (%)	<i>0.97–1.03</i> (%)	<i>1.03–1.10</i> (%)	<i>1.10–1.20</i> (%)
<i>Delta</i>						
SABR $\beta = 0.5$ (MV)	75.29	80.47	76.74	78.58	69.20	63.48
BSM-A	76.18	77.74	75.93	80.26	72.45	69.10
Heston (MV)	78.48	83.97	80.99	79.42	73.28	68.62
Normal mixture (Adj)	90.77	97.63	91.56	93.22	87.14	75.94
Normal mixture	162.04	136.45	169.99	176.97	156.36	140.13
<i>Delta–gamma</i>						
Heston (MV)	92.60	92.51	96.25	98.87	89.89	83.95
SABR $\beta = 0.5$ (MV)	93.84	93.12	98.68	99.62	89.68	86.63
BSM-A	98.10	94.28	98.87	99.70	95.91	102.71
Normal mixture (Adj)	124.43	116.90	106.09	101.91	127.37	166.92
Normal mixture	193.22	158.35	214.09	142.40	225.83	209.55

Note. This table reports the standard error of all hedging models across different moneyness buckets. The error is given as a percentage of the BSM error.

TABLE III
Hedging P&L

	<i>All</i> (%)	<i>1–90</i> (%)	<i>91–180</i> (%)	≥ 181 (%)
<i>Delta</i>				
SABR $\beta = 0.5$ (MV)	74.58	78.62	72.08	69.12
BSM-A	75.35	79.34	72.50	70.31
Heston (MV)	76.98	81.68	72.80	71.56
Normal mixture (Adj)	95.50	90.13	96.61	103.52
Normal mixture	166.76	152.98	181.97	176.99
<i>Delta–gamma</i>				
Heston (MV)	91.65	88.75	91.07	95.23
SABR $\beta = 0.5$ (MV)	92.82	87.13	93.56	98.69
BSM-A	96.88	92.02	96.24	102.69
Normal mixture (Adj)	132.09	101.17	139.04	158.45
Normal mixture	204.72	203.42	217.17	197.97

Note. This table reports the standard error of all hedging models across different maturity buckets for the last three months of the sample. The error is given as a percentage of the BSM error.

A more detailed analysis of these results is presented in Tables II and III. Table II reports the delta and delta–gamma hedge portfolio P&L standard deviations (as a percentage of the BSM standard deviation) averaged by moneyness K/F and over all days in the sample period. We have highlighted in bold the smallest standard deviation for each moneyness interval. For delta hedging, the adjusted BSM model performs best for low moneyness options, whereas the SABR model outperforms the other models for at-the-money options and

high-moneyness options. The NM model is again the poorest model for delta hedging, regardless of the moneyness of the options. The results for the delta–gamma–hedging performance in the lower part of the table are not so straightforward. The most interesting result is that the NM model seems to be completely unable to hedge adequately, but the model risk adjustment almost halves the standard deviation of the hedging error for all moneyness buckets. However, the model is still outperformed by the other three. Even these three show no appreciable improvement on the simple BSM model for ATM options.

Table III reports the standard deviation of the hedging error across different maturity buckets. For this table we restrict our sample to the last three months (September 17, 2007 to December 17, 2007) to guarantee that all maturity buckets contain hedging errors from the same time period.¹⁸ For the delta hedge we find that the hedging performance of all models is nearly independent of the maturity of the options. The ranking of the models remains the same for all maturity buckets, with the SABR model producing the hedging P&L with the smallest standard deviation. For delta–gamma hedging, the SABR model performs best for short-term options, whereas for medium- and long-term options the Heston model performs best. A possible explanation is that the mean-reversion in the variance process in the Heston model might be a more accurate description for the variance process compared with the geometric Brownian motion that governs the SABR variance. Although the SABR model translates short-term volatility effects into the prices of long-term options, the Heston model mitigates these effects by assuming the variance eventually converges toward a long-term value.

Summarizing the empirical results as they relate to our theoretical analysis, we have shown that adjusting the NM deltas and gammas for SLV provides hedge ratios that dramatically improve the performance of the NM model. For delta hedging we obtain a hedging performance that is similar to the SABR and Heston MV hedge ratios. This shows that the dynamics of the parameters of a lognormal mixture diffusion model have important information that can be used for hedging purposes.

CONCLUSIONS

One of the main sources of model risk in option pricing models is parameter instability. The vast majority of pricing models assume that all parameters except spot volatility are constant, yet they change every time the model is calibrated. We have derived hedge ratios that are adjusted to account for this type of model risk. The price hedge ratios are equal to the standard hedge ratios plus

¹⁸We also carried out the same analysis for the whole sample period and found no qualitative difference between those results and the results reported here.

an adjustment factor that depends on the degree of uncertainty in the model parameters and on their correlation with the underlying asset price. Our theoretical and empirical analysis has shown that the hedging performance of a scale-invariant deterministic volatility model improves dramatically if we augment the model by assuming that all the parameters could be stochastic. Although we have only studied the effect of stochastic parameters on one particular scale-invariant option pricing model (the lognormal mixture diffusion) it is likely that the use of stochastic parameters and the adjusted hedge ratios that we have derived would significantly improve the hedging performance of other option pricing models.

APPENDIX

Proof of Theorem 1

The hedge ratio that minimizes the variance of the delta-hedged portfolio and thus removes all dS -risk is given by (see Bakshi et al., 1997; Lee, 2001)

$$\delta = \frac{\langle df, dS \rangle}{\langle dS, dS \rangle}.$$

Applying Itô's lemma to the claim price yields

$$df = (\dots)dt + f_S dS + \sum_{j=1}^h f_{\xi^j} d\xi^j + \sum_{i=1}^n f_{\lambda^i} d\lambda^i$$

and using this in the definition above yields the result for delta, where $X^{\lambda^i, S} \equiv \langle d\lambda^i, dS \rangle / \langle dS, dS \rangle$ and $X^{\xi^j, S} \equiv \langle d\xi^j, dS \rangle / \langle dS, dS \rangle$. The MV gamma is defined as

$$\gamma = \frac{\langle d\delta, dS \rangle}{\langle dS, dS \rangle}.$$

Using Itô's lemma for δ gives

$$\begin{aligned} d\delta = & df_S + \sum_{j=1}^h (X^{\xi^j, S} df_{\xi^j} + f_{\xi^j} dX^{\xi^j, S} + dX^{\xi^j, S} df_{\xi^j}) \\ & + \sum_{i=1}^n (X^{\lambda^i, S} df_{\lambda^i} + f_{\lambda^i} dX^{\lambda^i, S} + dX^{\lambda^i, S} df_{\lambda^i}) \end{aligned}$$

where the dynamics of f_S , f_{λ^i} , f_{ξ^j} , $X^{\lambda^i, S}$, and $X^{\xi^j, S}$ are all dependent on the risk factors S , λ , and ξ , i.e. their dynamics follow from Itô's lemma, e.g.

$$dX^{\xi^j, S} = (\dots)dt + X_S^{\xi^j, S} dS + \sum_{k=1}^h X_{\xi^k}^{\xi^j, S} d\xi^k + \sum_{i=1}^n X_{\lambda^i}^{\xi^j, S} d\lambda^i.$$

Substituting the dynamics for f_S , f_{λ^i} , f_{ξ^j} , $X^{\lambda^i,S}$, and $X^{\xi^j,S}$ into the expression for gamma in Equation (11) yields

$$\begin{aligned} \gamma = & f_{SS} + \sum_{j=1}^h \left\{ X^{\xi^j,S} \left(2f_{S\xi^j} + \sum_{k=1}^h f_{\xi^j\xi^k} X^{\xi^k,S} + \sum_{i=1}^n f_{\xi^j\lambda^i} X^{\lambda^i,S} \right) \right\} \\ & + \sum_{j=1}^h \left\{ f_{\xi^j} \left(X_S^{\xi^j,S} + \sum_{k=1}^h X_{\xi^k}^{\xi^j,S} X^{\xi^k,S} + \sum_{i=1}^n X_{\lambda^i}^{\xi^j,S} X^{\lambda^i,S} \right) \right\} \\ & + \sum_{i=1}^n \left\{ X^{\lambda^i,S} \left(2f_{\lambda^i S} + \sum_{k=1}^h f_{\lambda^i\xi^k} X^{\xi^k,S} + \sum_{l=1}^n f_{\lambda^i\lambda^l} X^{\lambda^l,S} \right) \right\} \\ & + \sum_{i=1}^n \left\{ f_{\lambda^i} \left(X_S^{\lambda^i,S} + \sum_{k=1}^h X_{\xi^k}^{\lambda^i,S} X^{\xi^k,S} + \sum_{l=1}^n X_{\lambda^l}^{\lambda^i,S} X^{\lambda^l,S} \right) \right\}. \end{aligned}$$

Note that the partial derivative of β^i with respect to λ^j is zero unless $i=j$. $X^{\xi^j,S}$ and its partial derivatives follow directly from the assumption on the stochastic differential equation of ξ .

Normal Mixture Diffusion

This appendix provides further details on the lognormal mixture diffusion model for options on futures. In particular it states the pricing formulae and all partial derivatives, which are used in this article. The price of a European call or put ($\omega = 1$ for a call, $\omega = -1$ for a put) option on the futures price with maturity $\tau = T-t$ at time t is given by

$$\begin{aligned} f(\omega, F_t, K, \tau, r, s_1, \sigma_1, \sigma_2, p) = & p \underbrace{P(\omega, F_t, K, \tau, r, s_1, \sigma_1)}_{f^1} \\ & + (1 - p) \underbrace{P(\omega, F_t, K, \tau, r, s_2(\tau, p, s_1), \sigma_2)}_{f^2}, \end{aligned}$$

where

$$P(\omega, F_t, K, \tau, r, s_i, \sigma_i) = e^{-r\tau} (\omega F_t e^{s_i\tau} \Phi(\omega d_1^i) - \omega K \Phi(\omega d_2^i)),$$

with

$$d_1^i = \frac{\ln(F_t/K) + (s_i + 0.5\sigma_i^2)\tau}{\sigma_i\sqrt{\tau}},$$

$$d_2^i = d_1^i - \sigma_i\sqrt{\tau},$$

$$s_2(\tau, p, s_1) = \frac{1}{\tau} (\ln(1 - pe^{s_1\tau}) - \ln(1 - p)).$$

Brigo et al. (2003) prove that this model implies an SDE for the underlying under the risk-neutral measure of the form

$$dF_t = F_t \sigma(t, F_t) dW_t$$

where

$$\sigma(t, y)^2 = \frac{p\sigma_1^2 g_1(y) + (1 - p)\sigma_2^2 g_2(y)}{pg_1(y) + (1 - p)g_2(y)} + \frac{2F_t(ps_1 e^{m_1 n_1} + (1 - p)s_2 e^{m_2 n_2})}{y^2(pg_1(y) + (1 - p)g_2(y))},$$

with

$$g_i(y) = \frac{1}{y v_i \sqrt{2\pi}} \exp\left(-\frac{1}{2v_i^2} \left[\ln\left(\frac{y}{F_t}\right) - m_i + \frac{1}{2} v_i^2\right]^2\right),$$

$$n_i = \Phi\left(v_i^{-1} \left(\ln\left(\frac{F_t}{y}\right) + m_i + \frac{1}{2} v_i^2\right)\right),$$

$$m_i = s_i \tau, v_i = \sigma_i \sqrt{\tau}.$$

The first partial derivatives of the claim price with respect to F and the model parameters mainly follow immediately from well-known results in BSM model. In addition, the gamma is a simple linear combination of BSM gammas. The exact formulae are as follows:

$$f_F = p e^{(s_1 - r)\tau} \omega \Phi(\omega d_1^1) + (1 - p) e^{(s_2 - r)\tau} \omega \Phi(\omega d_1^2),$$

$$f_{FF} = p e^{(s_1 - r)\tau} \frac{\varphi(d_1^1)}{\sigma_1 \sqrt{\tau} F_t} + (1 - p) e^{(s_2 - r)\tau} \frac{\varphi(d_1^2)}{\sigma_2 \sqrt{\tau} F_t},$$

$$f_{\sigma_1} = p F_t e^{(s_1 - r)\tau} \varphi(d_1^1) \sqrt{\tau},$$

$$f_{\sigma_2} = (1 - p) F_t e^{(s_2 - r)\tau} \varphi(d_1^2) \sqrt{\tau},$$

$$f_p = f^1 - f^2 + \left(\frac{e^{(s_2 - r)\tau} \omega F_t \Phi(\omega d_1^2) (1 - e^{s_1 \tau})}{(1 - p e^{s_1 \tau})}\right),$$

$$f_{s_1} = \omega F_t \tau p e^{(s_1 - r)\tau} [\Phi(\omega d_1^1) - \Phi(\omega d_1^2)].$$

The first partial derivatives with respect to the delta are

$$f_{F\sigma_1} = -p e^{(s_1 - r)\tau} \varphi(d_1^1) \frac{d_2^1}{\sigma_1},$$

$$f_{F\sigma_2} = -(1 - p) e^{(s_2 - r)\tau} \varphi(d_1^2) \frac{d_2^2}{\sigma_2},$$

$$f_{Fp} = e^{(s_1 - r)\tau} \omega \Phi(\omega d_1^1) + \frac{e^{(s_2 - r)\tau} (1 - e^{s_1 \tau})}{(1 - p e^{s_1 \tau})} \omega \Phi(\omega d_1^2) \frac{\varphi(d_1^2)}{\sigma_2 \sqrt{\tau}} - e^{(s_2 - r)\tau} \omega \Phi(\omega d_1^2),$$

$$f_{Fs_1} = p e^{(s_1 - r)\tau} \omega \tau \left(\Phi(\omega d_1^1) + \frac{\varphi(d_1^1) \omega}{\sigma_1 \sqrt{\tau}} \right) - (1 - p) \frac{p e^{s_1 \tau} e^{(s_2 - r)\tau}}{1 - p e^{s_1 \tau}} \omega \tau \left(\Phi(\omega d_1^2) + \omega \frac{\varphi(d_1^2)}{\sigma_2 \sqrt{\tau}} \right).$$

The second partial derivatives can be obtained as

$$\begin{aligned}
 f_{\sigma_1\sigma_1} &= pF_t e^{(s_1-r)\tau} \sqrt{\tau} \varphi(d_1^1) \frac{d_1^1 d_2^1}{\sigma_1}, \\
 f_{\sigma_1\sigma_2} &= 0, \\
 f_{\sigma_1 p} &= F_t e^{(s_1-r)\tau} \sqrt{\tau} \varphi(d_1^1), \\
 f_{\sigma_1 s_1} &= pF_t \sqrt{\tau} e^{(s_1-r)\tau} \left(-d_1^1 \varphi(d_1^1) \frac{\sqrt{\tau}}{\sigma_1} + \tau \varphi(d_1^2) \right), \\
 f_{\sigma_2\sigma_2} &= (1-p)F_t e^{(s_2-r)\tau} \sqrt{\tau} \varphi(d_1^2) \frac{d_1^2 d_2^2}{\sigma_2}, \\
 f_{\sigma_2 p} &= F_t \sqrt{\tau} \varphi(d_1^2) e^{(s_2-r)\tau} \left(-1 + \frac{1 - e^{s_1\tau}}{1 - pe^{s_1\tau}} - \frac{d_1^2 (1 - e^{s_1\tau})}{\sigma_2 \sqrt{\tau} (1 - pe^{s_1\tau})} \right), \\
 f_{\sigma_2 s_1} &= pF_t \tau \varphi(d_1^2) e^{(s_1-r)\tau} \left(\frac{d_1^2}{\sigma_2} - \sqrt{\tau} \right), \\
 f_{s_1 s_1} &= p\omega F_t \tau e^{(s_1-r)\tau} \left(\tau \Phi(\omega d_1^1) + \omega \varphi(d_1^1) \frac{\sqrt{\tau}}{\sigma_1} \right) \\
 &\quad - \frac{p\omega F_t \tau^2 e^{(s_1-r)\tau}}{1 - pe^{s_1\tau}} \left(\Phi(\omega d_1^2) - pe^{s_1\tau} \left(\Phi(\omega d_1^2) + \frac{\omega \varphi(d_1^2)}{\sigma_2 \sqrt{\tau}} \right) \right), \\
 f_{s_1 p} &= \omega F_t \tau e^{(s_1-r)\tau} \left[\Phi(\omega d_1^1) - \Phi(\omega d_1^2) - \frac{\omega p \varphi(d_1^2) (1 - e^{s_1\tau})}{\sigma_2 \sqrt{\tau} (1 - pe^{s_1\tau}) (1 - p)} \right], \\
 f_{pp} &= -\frac{e^{(s_2-r)\tau} (1 - e^{s_1\tau})^2 F_t \varphi(d_1^2)}{(1 - p) (1 - pe^{s_1\tau})^2 \sigma_2 \sqrt{\tau}}.
 \end{aligned}$$

BIBLIOGRAPHY

Alexander, C. (2004). Normal mixture diffusion with uncertain volatility: Modelling short and long term smile effects. *Journal of Banking and Finance*, 28, 2957–2980.

Alexander, C., & Nogueira, L. M. (2004). Stochastic local volatility. *Proceedings of the second international IASTED conference on financial engineering and applications* (pp. 136–141). Cambridge: MIT.

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

Alexander, C., & Nogueira, L. M. (2007b). Price hedge ratios for homogeneous claims on tradable assets. *Quantitative Finance*, 7, 473–479.

Andersen, T. G., Benzoni, L., & Lund, J. (2002). An empirical investigation of continuous-time equity return models. *The Journal of Finance*, 57, 1239–1284.

- Bakshi, G., Cao, C., & Chen, Z. (1997). Empirical performance of alternative option pricing models. *Journal of Finance*, 52, 2003–2049.
- Bakshi, G., Cao, C., & Chen, Z. (2000). Pricing and hedging long-term options. *Journal of Econometrics*, 94, 277–318.
- Bates, D. S. (2003). Empirical option pricing: A retrospection. *Journal of Econometrics*, 116, 387–404.
- Bates, D. S. (2005). Hedging the smirk. *Financial Research Letters*, 2, 195–200.
- Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, 3, 167–179.
- Black, F., & Scholes, M. S. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81, 637–54.
- Branger, N., & Schlag, C. (2004). Model risk: A conceptual framework for risk measurement and hedging (working paper).
- Brigo, D., & Mercurio, F. (2002). Lognormal-mixture dynamics and calibration to market volatility smiles. *International Journal of Theoretical & Applied Finance*, 5, 427–446.
- Brigo, D., Mercurio, F., & Sartorelli, G. (2003). Alternative asset-price dynamics and volatility smile. *Quantitative Finance*, 3, 173–183.
- Carr, P., Ellis, K., & Gupta, V. (1998). Static hedging of exotic options. *Journal of Finance*, 53, 1165–1190.
- Christoffersen, P., & Jacobs, K. (2004). The importance of the loss function in option valuation. *Journal of Financial Economics*, 72, 291–318.
- Coleman, T., Kim, Y., Li, Y., & Verma, A. (2001). Dynamic hedging with a deterministic local volatility function model. *Journal of Risk*, 4, 63–89.
- Cont, R. (2006). Model uncertainty and its impact on the pricing of derivative instruments. *Mathematical Finance*, 16, 519–547.
- Crépey, S. (2004). Delta-hedging vega risk? *Quantitative Finance*, 4, 559–579.
- Derman, E. (1996). Model risk. *Risk*, 9, No. 5, 34–37.
- Derman, E., Ergener, D., & Kani, I. (1995). Static options replication. *Journal of Derivatives*, 2, 78–95.
- Derman, E., & Kani, I. (1994). Riding on a smile. *Risk*, 7, 32–29.
- Dumas, B., Fleming, J., & Whaley, R. E. (1998). Implied volatility functions: Empirical tests. *Journal of Finance*, 53, 2059–2106.
- Dupire, B. (1994). Pricing with a smile. *Risk*, 7, 18–20.
- Engelmann, B., Fengler, M. R., & Schwendner, P. (2006). Better than its reputation: An empirical hedging analysis of the local volatility model for barrier options (working paper), unpublished.
- Eraker, B., Johannes, M., & Polson, N. (2003). The impact of jumps in volatility and returns. *Journal of Finance*, 53, 1269–1300.
- Fengler, M. (2005). *Semiparametric modeling of implied volatility*. Berlin: Springer.
- Frey, R. (1997). Derivative asset analysis in models with level-dependent and stochastic volatility. *CWI Quarterly*, 10, 1–34.
- Gemmill, G., & Saflekos, A. (2000). How useful are implied distributions? Evidence from stock index options. *Journal of Derivatives*, 7, 83–98.
- Green, T., & Figlewski, S. (1999). Market risk and model risk for a financial institution writing options. *Journal of Finance*, 54, 1465–1499.
- Hafner, R. (2004). *Stochastic implied volatility. Lecture notes in economics and mathematical systems*. Berlin: Springer.

- Hagan, P., Kumar, D., Lesniewski, A., & Woodward, D. (2002). Managing smile risk. *Wilmott Magazine*, 1, 84–108.
- 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, 327–343.
- Hoogland, J., & Neumann, C. (2001). Local scale invariance and contingent claim pricing. *International Journal of Theoretical and Applied Finance*, 4, 1–21.
- Hull, J. C., & White, A. D. (1987). The pricing of options on assets with stochastic volatilities. *Journal of Finance*, 42, 281–300.
- Kani, I., Derman, E., & Kamal, M. (1997). Trading and hedging local volatility. *The Journal of Financial Engineering*, 6, 1233–1268.
- Kim, I. J., & Kim, S. (2003). On the usefulness of implied risk-neutral distributions—Evidence from the Korean KOSPI 200 index options market. *Journal of Risk*, 6, 93–110.
- Lee, R. (2001). Implied and local volatilities under stochastic volatility. *International Journal of Theoretical and Applied Finance*, 4, 45–89.
- McIntyre, M. (2001). Performance of Dupire's implied diffusion approach under sparse and incomplete data. *Journal of Computational Finance*, 4, 33–84.
- Merton, R. C. (1973). Theory of rational option pricing. *Bell Journal of Economics*, 4, 141–183.
- Nandi, S. (1998). How important is the correlation between returns and volatility in a stochastic volatility model? Empirical evidence from pricing and hedging in the S&P 500 index options market. *Journal of Banking and Finance*, 22, 589–610.
- Nogueira, L. M. (2006). Hedging options with local and stochastic volatility models. Ph.D. thesis, ICMA Centre, The University of Reading.
- Psychoyios, D., & Skiadopoulos, G. (2006). Volatility options: Hedging effectiveness, pricing, and model error. *Journal of Futures Markets*, 26, 1–31.
- Rubinstein, M. (1994). Implied binomial trees. *Journal of Finance*, 49, 771–818.
- Schweizer, M. (2001). A guided tour through quadratic hedging approaches. In E. Jouini, J. Cvitanic, M. Musiela (Eds.), *Option pricing, interest rates and risk management* (pp. 538–574). Cambridge, UK: Cambridge University Press.
- Skiadopoulos, G. (2001). Volatility smile consistent option models: A survey. *International Journal of Theoretical and Applied Finance*, 4, 403–437.
- Vähämaa, S. (2004). Delta hedging with the smile. *Financial Markets and Portfolio Management*, 18, 241–255.
- Wilkins, S. (2005). Option pricing based on mixtures of distributions: Evidence from the Eurex index and interest rate futures options market. *Derivatives Use, Trading Regulation*, 11, 213–231.