REGIME-DEPENDENT SMILE-ADJUSTED DELTA HEDGING

CAROL ALEXANDER*
ALEXANDER RUBINOV
MARKUS KALEPKY
STAMATIS LEONTSINIS

We introduce several regime-dependent smile-adjusted deltas and compare their efficiency with the smile-adjusted deltas that are popular with option traders. Using $16\frac{1}{2}$ years of daily option prices, out-of-sample hedging performance tests for options of all moneyness and maturities and daily, weekly, or fortnightly rebalancing show that even the simplest regime-dependent smile-adjustment consistently outperforms implied BSM delta hedging and local volatility and minimum variance smile-adjustments. Markov-switching deltas offer the best performance, with delta-hedging errors often half the size of implied BSM hedging errors. During volatile markets risk reduction from regime-dependent delta hedging is much greater than during tranquil periods. © 2011 Wiley Periodicals, Inc. Jrl Fut Mark

*Correspondence author, ICMA Centre, Henley Business School at Reading, Whiteknights, Reading RG6 6BA, UK. Tel: +44 (0)1183 786431, Fax: +44 (0)20 8668 6050, e-mail: c.alexander@icmacentre.ac.uk

Received October 2010; Accepted January 2011

- *Carol Alexander is Chair of Financial Risk Management the ICMA Centre, Henley Business School at Reading, United Kingdom.*
- *Alexander Rubinov was a masters student at the Technical University Munich, now at Ergo Asset Management, Munich, Germany.*
- *Markus Kalepky was a masters student at the Technical University Munich, now at Allianz Investment Management, Stuttgart, Germany.*
- *Stamatis Leontsinis was a PhD student at the ICMA Centre, Henley Business School at Reading, now at Fulcrum Asset Management LLP, London, United Kingdom.*

The Journal of Futures Markets, Vol. 00, No. 00, 1–27 (2011)

DOI: 10.1002/fut.20517

1. INTRODUCTION

The majority of option-pricing research focuses on the construction and calibration of models used to price over-the-counter (OTC) illiquid derivatives. Yet exchange trading on standard vanilla options is much more active than trading on OTC products. Market makers on exchanges can make large profits from a high-volume business.[1] The prices of liquid, exchange-traded vanilla options are determined by supply and demand, with market makers setting bid–ask spreads to make a small profit on the round trip after accounting for expected hedging costs. The more accurate their hedging the lower its cost, so the more competitive their spread and the greater their trading volume. Hence, the "daily bread" of the vast majority of option traders depends on the accurate hedging of vanilla options.

Compared with the veritable explosion of research on pricing OTC options with path-dependent and/or exotic pay-offs there has been relatively little research on hedging vanilla options. The market is incomplete when the price process has stochastic volatility, so there is no perfect hedge. Instead, traders should choose a hedge ratio to minimize the hedging costs over the life of the option. Such a delta is commonly termed "locally risk minimizing" or "minimum variance" because it is the delta which minimizes the instantaneous variance of the hedging error: see Schweizer (1991), Bakshi et al. (1997, 2000), Frey (1997), Lee (2001), and others. Poulsen et al. (2009) demonstrate, using simulations and empirical data, that the particular stochastic volatility model used does not seem to matter that much—the important factor for accurate delta hedging is to use a minimum variance rather than a standard (partial derivative) delta.

Bakshi et al. (1997) found no evidence that adding stochastic interest rates and/or jumps in prices improves the hedging performance of the Heston (1993) stochastic volatility model for S&P 500 options. Bakshi et al. (2000) confirm these results even when hedging long-term options. Kim and Kim (2004) find that the Heston model outperforms other stochastic volatility models for Kospi 200 options, confirming that jumps offer no significant improvement. Alexander et al. (2009) explain how hedge ratios should be adjusted to account for the model risk stemming from time-varying calibrated parameters in a model that assumes the parameters are constant. The motivation for these investigations is that traders will use the same model for hedging vanilla options as they do for pricing complex options.

However, in practice many traders base vanilla option hedge ratios on simple adjustments to the standard lognormal model of Black and Scholes (1973)

[1]For instance, at the end of December 2009, banks trading vanilla options on the Chicago Mercantile Exchange held about 35 million contracts. See http://www.cftc.gov/dea/bank/deadec09o.htm—accessed May 2010

and Merton (1973)—henceforth referred to as BSM. The simplest of these is the "implied BSM" approach where each option's price hedge ratios are evaluated at their own implied volatility. Bakshi et al. (1997, 2000), Christoffersen and Jacobs (2004), Alexander et al. (2009), and others have demonstrated how effective this hedging approach is compared with minimum variance hedging in various stochastic volatility models, especially for hedging at-the-money and high-strike options.

Smile-adjustments to the implied BSM delta are derived from the slope of the implied volatility smile. An option-pricing model's calibration usually involves setting the model price equal to the BSM price based on the option's implied volatility. Differentiating both sides of this equality yields the smile-adjusted delta as the implied BSM delta plus an adjustment term equal to the BSM vega times the implied volatility sensitivity to the underlying price.[2] Derman (1999) introduced three different ways of adjusting the BSM model to account for an equity index implied volatility smile, each being relevant in a different market regime. They imply different types of "stickiness" for the volatility in the process for the underlying price F.[3]

According to Derman, the *sticky strike* model applies during stable, trending equity markets. Here each option may be priced using a single, constant-volatility geometric Brownian motion (GBM) for F, but each option uses a different GBM in which the volatility is taken to be the implied volatility of the option. In this regime, the implied BSM delta-hedges are computed using the option's own implied volatility.

There are also multiple GBMs for F in the *sticky moneyness* model, which applies when the market is in a range-bounded regime. But here the (constant) volatility of each GBM depends on the option's moneyness. Hence, to price the option one has to employ multiple GBMs, jumping between them as F evolves over time to use the price process with the relevant moneyness of the option, given the price level at the time. For this reason, sticky moneyness is also referred to as the *floating smile* model, and we need an additional adjustment to the implied BSM price hedge ratio, which depends on how the smile moves when F changes.

The only hedging model with a unique process for F that is used to price every option is the *sticky tree* model, which applies during excessively volatile markets. The volatility is not constant in the process; it is a deterministic function of price and time. The forward volatility conditional on the price level at

[2]The vega is the option's model price sensitivity to the option's implied volatility. Although a smile-adjusted delta incorporates a vega effect, there is no additional option in the hedge, i.e. the hedge is affected just by trading the futures in the amount specified by the smile-adjusted delta.

[3]For equity index options we use the futures price F for the underlying rather than the spot, since they are delta-hedged using the futures.

that time is equivalent to the *local volatility* of Dupire (1994) and Derman and Kani (1994). Derman et al. (1996), Coleman et al. (2001), and Crepey (2004) all show that the local volatility delta-hedge ratio is a smile-adjusted delta where the implied volatility-price sensitivity is equal to the slope of the smile in the strike dimension.

There is conclusive evidence that deltas based on the sticky moneyness model provide worse hedges than the implied BSM delta. The reason for this is that its price hedge ratios are theoretically equivalent to the standard (partial derivative) price hedge ratios derived from *any* proper model for pricing options on a tradable asset, which are well-known to perform poorly relative to the implied BSM. See Alexander and Nogueira (2007) for further details.[4] By contrast, local volatility (sticky tree) delta-hedges are usually more effective than implied BSM deltas, except possibly for hedging high-strike options with frequent rebalancing.[5]

Vähämaa (2004) and Crepey (2004) find that the effectiveness of the local volatility smile-adjustment has regime dependence; in fact, their results confirm Derman's intuition that the local volatility delta is superior only during excessively volatile periods. Alexander et al. (2009) demonstrate that the local volatility deltas can outperform Heston minimum variance deltas for S&P 500 options, and for FTSE 100 options Alexander and Kaeck (2010) show that this is true for options across the entire range of moneyness and for all maturities, over both daily and weekly rebalancing frequencies, and during both volatile and tranquil periods. They also confirm that the main advantage of using the sticky tree/local volatility delta rather than the simple implied BSM delta is for hedging during volatile market conditons.

Given the evidence reviewed above, this study asks an obvious question: can we improve on the implied BSM and local volatility delta-hedges by explicitly modeling regime-dependence in a smile-adjusted delta? We test the performance of several regime-dependent smile-adjusted deltas using $16\frac{1}{2}$ years of daily closing

[4]Alexander and Nogueira (2007) extended and formalized the concept of scale invariance, first introduced by Merton (1976). An option-pricing model is scale invariant if and only if the relative change of the futures price is independent of the price level. All proper models for pricing options on tradable assets must be scale invariant, for if they were not then option prices relative to the futures price would not be invariant under a change of numeraire. Subsequently they showed that all models for pricing homogeneous options on tradable assets have the same theoretical delta and gamma. More precisely, for every option with homogeneous pay-off (be it a standard European, a barrier or an American corridor) it does not matter which model is used for hedging. The standard price hedge ratios (which are partial derivatives of the option price with respect to the futures price) are theoretically identical, and the only empirical difference between them is due to calibration error.

[5]Several empirical studies support this: for instance, see Dumas et al. (1998), Coleman et al. (2001), Vähämaa (2004), Crepey (2004), Alexander and Nogueira (2007) and others.

prices on the FTSE 100 index options, thus covering several business cycles with both volatile and tranquil market conditions.[6] We consider two simple, ad hoc regime-dependent smile-adjustments and then specify a formal model based on Markov switching (MS) of price-volatility sensitivities. In the following: Section II defines the smile-adjusted deltas and their implicit assumptions about implied volatility dynamics; Section III describes the data; Section IV presents the empirical results for the MS model; Section V defines the hedging strategies and presents our results; and Section VI summarizes and concludes.

2. SMILE-ADJUSTMENTS TO THE BSM DELTA

We begin by describing the various smile-adjustments that have already been used in the literature, making explicit their assumptions about the sensitivity of implied volatility to the futures price. Then we propose new smile-adjustments that allow this sensitivity to depend on the current market regime, suggesting two ad hoc regime-dependent smile-adjustments and one that is derived in the context of the MS model.

Suppose a standard European option is written on one index futures contract, with price F. The purpose of delta-hedging is to build a portfolio that neutralizes the sensitivity of the option to changes in F, and the number of futures contracts that should be bought is the delta of the option. This is the derivative of the option price $f(K,T|F)$ with respect to F, where K denotes the strike of the option and T its maturity. Note that f also depends on σ, the volatility of the futures price and r, the discount rate. The option's implied volatility $\sigma(K,T|F)$ is that volatility which equates the BSM price to the market price of the option. It depends on K and T because the implied volatility surface is not flat and it also depends on F because the entire surface moves when the futures price changes. Taking the dependence of implied volatility on F into account requires the use of the chain rule when differentiating the option price with respect to F. This yields the smile-adjusted delta

$$\delta_{adj}(K,T|F) = \frac{\partial f}{\partial F} + \frac{\partial f}{\partial \sigma}\frac{\partial \sigma}{\partial F} = \delta_{BSM}(K,T|F) + \nu_{BSM}(K,T|F)\sigma_F(K,T|F) \qquad (1)$$

with the implicitly defined notation for δ_{BSM}, ν_{BSM}, and σ_F. In other words, to account for dynamics in implied volatility, the implied BSM delta δ_{BSM} should

[6]We remark that all the empirical delta-hedging studies reviewed above utilize not more than 4 years of daily data, with the exception of Dumas et al. (1998), who use $5\frac{1}{2}$ years. In fact, the majority of option-hedging studies are based on only 1 or 2 years of data. The options studied are usually either S&P 500 or FTSE 100 index options, although Crepey (2004) also uses a 2001 sample of DAX 30, SMI, and DJIA index options data, and Poulsen et al. (2009) use a 2004–2005 sample of Eurostoxx 50 and EUR/USM options, in addition to S&P 500 options.

be adjusted by an amount which depends on the option vega ν_{BSM} and the volatility-price sensitivity σ_F.

2.1. Sticky Models

A variety of different terminologies have been applied to these smile-adjusted deltas and our terminology is as follows: the "sticky strike" (SS) delta denotes the implied BSM delta;[7] the "sticky tree" (ST) delta is the local volatility delta; and the "sticky moneyness" (SM) delta denotes the partial derivative delta derived from any proper price process for a tradable asset.[8] In each of these hedging models $\sigma_F(K,T|F) = k\sigma_K(K,T|F)$ with

$$k = \begin{cases} 0 & (SS) \\ 1 & (ST) \\ -K/F & (SM). \end{cases} \tag{2}$$

Now (2) defines the following very different skew dynamics when F changes:

- (SS): the skew does not move;
- (ST): the skew tilts by an amount equal to the slope of the skew at that point;
- (SM): the skew tilts in the opposite direction to the ST model, by an amount equal to K/F times slope of the skew at that point.

If the skew has zero slope at some point, as it does if it has a typical hockey stick shape, then in the ST and SM models it will pivot about this point.[9] For an increase in F, it tilts downward at lower strikes and upward at higher strikes in the ST model, and in the opposite direction in the SM model.

Our empirical study utilizes data over a very long period, with a great variety of option strikes being traded depending on the level of the futures price. Hence, we change from the strike to the moneyness metric, defining moneyness as $m = K/F$. The implied volatility as a function of moneyness is written $\theta(m,T|F) = \sigma(mF,T|F)$, so the ATM volatility is $\theta(1,T|F) = \sigma(F,T|F)$ and the smile-adjusted delta (1) may be written as follows:

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

[7]Alternative terms include the "practitioner Black–Scholes" delta (e.g. Christoffersen & Jacobs, 2004), "the volatility-by-strike" model (e.g. Rosenberg, 2000), and the "absolute smile" method (e.g. Jackwerth & Rubinstein, 2001).

[8]The sticky moneyness model is also referred to as the "sticky delta" model. Its delta is the scale-invariant delta defined by Alexander and Noguiera (2007), which is a partial price derivative hedge ratio that is model-free in the class of scale-invariant models. Hence, SM deltas include, for example, the partial derivative deltas of the Heston (1993) model and of any deterministic volatility function (e.g. a normal mixture diffusion, which has sticky moneyness implied volatility dynamics).

[9]Typically, the equity index smile slopes downward at low strikes and very often also slopes downward at and above the at-the-money (ATM) strike; it is usually only at very high strikes that the smile may be upward sloping, thus having a "hockey stick" shape, and this is often termed the equity volatility skew, or "smirk."

Note that $\sigma_K(K,T|F) = \mathrm{F}^{-1}\theta_m(m,T|F)$ and $\sigma_F(K,T|F) = \theta_F(m,T|F)$, the latter being just a reparameterization of the former. Thus, the sticky models have

$$\theta_F(m,T|F) = \kappa\,\theta_m(m,T|F) \tag{4}$$

where

$$\kappa = \begin{cases} 0 & (SS) \\ 1/F & (ST) \\ -\,m/F & (SM). \end{cases} \tag{5}$$

Equations (3)–(5) define the smile-adjusted deltas $\theta_m(m,T|F,\,t)$ that are derived from the slope of the smile in the moneyness metric. We shall estimate this slope by fitting a cubic polynomial of the form $a_0(t) + a_1(t)m + a_2(t)m^2 + a_3(t)m^3$ to the market-implied volatility smile on each day t, setting

$$\hat{\theta}_m(m,T|F,\,t) = \hat{a}_1(t) + 2\hat{a}_2(t)m + 3\hat{a}_3(t)m^2.$$

2.2. Approximate Minimum Variance Delta

Lee (2001), Section 7.2, derives an approximate locally risk minimizing or "minimum variance" (MV) delta which, in the moneyness metric, is given by

$$\delta_{mv}(m,T|F) = \delta_{BSM}(m,T|F) + \frac{m}{F}\,\nu_{BSM}(m,T|F)\,\theta_m(m,T|F). \tag{6}$$

As such it may be viewed as a smile-adjusted delta with

$$\theta_F(m,T|F) = (m/F)\theta_m(m,T|F).$$

We remark that the MV delta is based on exactly the opposite implied volatility dynamics to those of the SM delta.

2.3. Ad hoc Regime-Dependent Smile-Adjustments

So far we have considered four possible assumptions for implied volatility dynamics, each leading to a different value for $\theta_F(m,T|F)$ in (3) which is derived entirely from the current smile $\theta(m,T|F)$ and the current value of the underlying F. However, Derman (1999) hypothesized that the implied volatility dynamics depend on the market regime, as specified in the introduction, where a regime is defined by the characteristics of returns and volatility that prevails over a period of time. If this hypothesis is correct, then a recent period of historical data on the smile and the underlying will contain information that is relevant for determining the implied volatility dynamics.

In this sub-section, we specify the dynamics by allowing θ_F to depend on recent historical data, in addition to the current smile and underlying price.

To this end, we set $\kappa = \kappa(m,T|F, t)$ in (4) to that multiple which minimizes the ex post standard deviation of the delta-hedging error over the past year. In other words, for each option, at each time t and fixing $\kappa(m,T|F, t)$ we compute the delta-hedging error on every day during the past year and calculate its standard deviation. Then we choose that value for $\kappa(m,T|F, t)$ which minimizes this standard deviation. This way, for every option, a time series of optimal multipliers κ in (4) is derived, whose values reflect the current regime by taking account of the hedging error over the previous year. These regime dependent (RD) deltas have an ad hoc regime dependence because the choice of a one-year sample is arbitrary; a more refined approach might seek to optimize this period.

All the deltas defined so far are restricted by the limitation of the relationship (4) for the volatility–price sensitivity; when $\theta_m(m,T|F) = 0$ the implied volatility will not change with F. That is, the smile can only pivot about the point where it has zero slope. We now consider the sticky tree model (because it offers the best hedging performance of all sticky models) and but now to augment its smile dynamics by adding the possibility for a shift to the smile's pivoting movements, assuming

$$\theta_F(m,T|F) = \alpha(T|F, t) + \frac{1}{F}\,\theta_m(m,T|F). \tag{7}$$

In the empirical results we shall estimate $\alpha(T|F, t)$ by considering the historical relationship between ATM volatility and the underlying price during the recent past. That is, we use the ATM volatility–price relationship to derive an additional shift in the smile, so that implied volatility changes even at a point where $\theta_m(m,T|F) = 0$. As described above, we arbitrarily choose one year of historical data leading up to the time t that the hedge is taken, and perform a simple linear regression of daily changes in ATM volatility on the daily log return on the futures price. Then we set $\hat{\alpha}(T|F, t)$ equal to the estimated slope coefficient, divided by F (since the regressor is the log return rather than changes in price). We call this delta the "regime-dependent shift" (RS). Although ad hoc, the RD and RS deltas are very easy to implement.

2.4. Markov Switching Smile-Adjustment

To determine regime-dependent values for $\theta_F(m,T|F, t)$ in the context of a formal statistical model, we introduce the "fixed moneyness spread" (FMS) at time t as follows:

$$\theta^{fms}(m,T|F, t) = \theta(m,T|F, t) - \theta(1,T|F, t).$$

We do this because Derman (1999) notes that it is the relationship between the ATM volatility $\theta(1,T|F, t)$ and the underlying price that appears to be regime dependent, and Alexander (2001) shows that the FMS has the same relationship with the futures price in all regimes, i.e.

$$\theta^{fms}(m,T|F,t) = -b(t,T)F(t,T)(m(t,T) - 1),$$

where $b(t,T)$ is not regime dependent.

The model is constructed in two stages: (a) the estimation of a MS model for the ATM volatility–price relationship, which determines a regime-dependent value for $\theta_F(1,T|F,t)$; and (b) a principal component analysis (PCA) of the FMS which is used to extrapolate the regime-dependent ATM sensitivity $\theta_F(1,T|F,t)$ to regime-dependent fixed-moneyness sensitivities $\theta_F(m,T|F,t)$.

Hamilton (1989) provided the first formal statistical representation of the idea that economic recessions and expansions influence the behavior of economic variables. Since then, Hamilton's MS techniques have been further developed and widely applied to a variety of disciplines—see Fruhwirth-Schnatter (2006) for a review. Statistical tests are used to identify the optimal MS model structure and the number of distinct regimes. The nonstandard distributions for such test statistics require a computationally intensive bootstrapping framework, such as that recommended by Ryden et al. (1998). In addition, we performed some powerful inference tests based on the comparison of unconditional densities proposed by Ait-Sahalia (1996) and applied to MS models by Breunig et al. (2003). A variety of different model structures were tested, including nonlinear dependence and different numbers of states. The results of these tests, which are summarized in section IV, indicate that the optimal MS model for the ATM volatility-futures price relationship is, for all maturities, a two-state MS model with the following specification

$$y(t) = \alpha^s + \beta_1^s x(t) + \beta_2^s x(t-1) + \beta_3^s y(t-1) + \varepsilon^s(t), \tag{8}$$

where $s = 1, 2$ denotes the state, $y(t)$ denotes the daily change in the T-maturity ATM volatility, and $x(t)$ denotes the daily log return on the futures and $\varepsilon^s \sim NID(0, \sigma_s^2)$. The state variable s is assumed to follow a first-order Markov chain with constant state-transition probabilities $\Pr(s_t = j|s_{t-1} = i) = \pi_{ij}$. Since $\sum_j \pi_{ij} = 1$, the transition matrix may be written as follows:

$$\pi_{ij} = \begin{pmatrix} \pi_{11} & \pi_{12} \\ \pi_{21} & \pi_{22} \end{pmatrix} = \begin{pmatrix} \pi_{11} & 1 - \pi_{11} \\ 1 - \pi_{22} & \pi_{22} \end{pmatrix} \tag{9}$$

Maximum likelihood estimation provides two sets of parameters, their standard errors, the transition matrix and a time series of conditional probabilities $\pi(t)$ for state 1 to be the ruling state at time t. We shall label the states so that $\hat{\sigma}_1 > \hat{\sigma}_2$ and therefore identify state 1 with the "volatile" regime and state 2 with the "tranquil" regime.

In the above, for brevity of notation, we have supressed the dependence of (8) and (9) on the maturity T. Recalling this dependence, (8) provides we have a regime-dependent *and* maturity-dependent value for the sensitivity of T-maturity ATM volatility to the daily change in the T maturity futures price at time t. There are two possible values: the instantaneous sensitivity $\beta^s(T)/F(t,T)$ and the steady-state sensitivity $(\beta_1^s(T) + \beta_2^s(T))/((1 - \beta_3^s(T))F(t,T))$.[10]

Next we apply PCA to the covariance matrix $\mathbf{V}(T)$ of the daily changes $\Delta\theta^{fms}(m,T|F,t)$ in the T-maturity FMS, and use the approximation

$$\Delta\theta^{fms}(m,T|F, t) = \sum_{i=1}^{3} \omega_i(m,T)p_i(t,T) \tag{10}$$

where $\omega_i(m,T)$ is the mth element of the ith eigenvector of $\mathbf{V}(T)$ and $p_i(t,T)$ is the value at time t of the ith principal component. Then, for each component we assume a time-varying but linear relationship with the log return on the futures

$$p_i(t,T) = \gamma_i(t,T)\Delta \log F(t,T) + \varepsilon_i(t,T) \tag{11}$$

where $\varepsilon_i(t,T)$ are i.i.d. processes. By construction principal components have zero unconditional correlation. Thus to specify a common factor $\Delta \log F(t,T)$ with a *constant* coefficient $\gamma_i(T)$ in Equation (11) would require all but one of the $\gamma_i(T)$, $i = 1, 2, 3$ to be zero. However, typically the conditional correlation between principal components is not zero so all three of the time-varying coefficients $\gamma_i(t,T)$ in (11) could be simultaneously different from zero.

Substituting (11) into (10) yields the FMS sensitivity to changes in the futures price

$$\theta_F^{fms}(m,T|F, t) = \frac{\lambda(t,T)}{F(t,T)}. \tag{12}$$

where $\lambda(t,T) = \sum_{i=1}^{3}\omega_i(m,T)\gamma_i(t,T)$. Writing

$$\theta_F(m,T|F, t) = \theta_F(1,T|F, t) + \theta_F^{fms}(m,T|F, t)$$

and substituting (12) plus the regime-dependent sensitivity for the T-maturity ATM volatility in the above, yields:

$$\theta_F^s(m,T|F, t) = \frac{\beta^s(T) + \lambda(t,T)}{F(t,T)}. \tag{13}$$

Substituting this regime-dependent value for $\theta_F(m,T|F, t)$ into (3) gives the MS deltas $\delta_{MS}^1(m,T|t)$ and $\delta_{MS}^2(m,T|t)$ which depend on whether the market is in the volatile state ($s = 1$) or the tranquil state ($s = 2$) at the time the hedge is taken.

[10]In the following, we use the instantaneous sensitivity. In terms of their empirical hedging performance, we find very little difference between the two.

It remains only to decide on the signal to be used as an indication of the current market regime. To this end, we employ the conditional regime probability $\pi(t,T)$ that is estimated from the MS model (8) and (9). For our hedging study, we start by setting $\delta^1_{MS}(m,T|t)$ if $\pi(t,T) > \frac{3}{4}$ and $\delta^2_{MS}(m,T|t)$ otherwise.[11] Then we investigate whether recent data have any useful information for the regime-switching signal, by setting $\delta^1_{MS}(m,T|t)$ if $\pi(t,T) > \tau(m,T|t)$ and $\delta^2_{MS}(m,T|t)$ otherwise, where $\tau(m,T|t)$ is chosen to minimize the ex post standard deviation of the hedging error at moneyness m during the past year. In other words, the selection of the threshold $\tau(m,T|t)$ for regime switching is based on a similar optimization to the selection of the multiplier in Section 2.4.

3. DATA

The option data used for our empirical study, provided by NYSE Euronext, consist of closing prices and implied volatilities of European FTSE 100 index options from October 2, 1992 to March 17, 2009, a period of over 4,000 days with over 2.3 million option prices in the raw data set. At the start of the sample, the price of the FTSE 100 futures was obtained using put–call parity because the options and futures markets closed at different times. Finally, the time series of GBP LIBOR rates were downloaded from the British Banker's Association's homepage and yield curve fitting was performed using Hermite splines.

The raw data were filtered for several reasons. In a few cases, the put and call implied volatility for a given maturity and strike differed, and here we only used the implied volatility of the out-of-the-money (OTM) option, this typically being the more liquid of the two. Second, we deleted all "cabinet" option prices and volatilities from the data set. Cabinet options are those which are very deep OTM and hence are nearly worthless. They are rarely traded, mostly just to close open positions. The exchange is obliged to quote a price for these options while open interest is positive, so they quote the minimum price of 50 pence, but the real value of these options is typically much lower. Hence to include these option prices in the database could distort our results. A third filter was based on the observation that, before 2008, the xx75 strike options were much less liquid than the xx25 strike options, and would therefore be liable to inaccurate stale prices. So we also deleted the xx75 strike options, before 2008, from the data set.

A fourth filter was used to obtain consistent time series of prices and volatilities covering the entire sample period. For this, we first constructed constant maturity option and futures prices with a fixed maturity of 30, 60, . . . , 180 days. We also constructed prices and volatilities at maturities 1, 5, and 10 days

[11]Following standard practice, to avoid excessive responses to occasional outliers the regime probability is smoothed and the regime probability trigger for entry to the high-volatility regime is above $\frac{1}{2}$.

less than these, because they are required for the hedging study. To this end, for every combination of strike and date we fitted arbitrage-free term structures of implied volatility, and evaluated these term structures at the fixed maturities. The best results were obtained using shape-preserving Hermite splines to interpolate and extrapolate the values of the implied volatility term structure, and we checked for no-arbitrage in the result.[12]

Just like cubic splines, Hermite splines consist of piecewise cubic polynomials. However, the second order condition at the knot points is replaced by a shape-preserving condition which is only on the slopes of the cubic polynomials at the knot points. Shape-preserving Hermite splines provide realistic extrapolated values, if these are necessary, and they are much less sensitive to data errors than cubic splines. Using Hermite splines we constructed roughly 100,000 term structures, so it is not possible to examine every term structure to verify that it is realistic. Instead we evaluated the result of Hermite spline interpolation and extrapolation across maturities by investigating the volatility skews for constant maturities of 30, 60, . . . , 180 days.

To avoid calendar arbitrage, option prices of any given moneyness must be nondecreasing with respect to maturity. A necessary but not sufficient condition for this is that the total implied variance increases with term. However, there are no simple no-arbitrage constraints for the shape of a fixed maturity implied volatility skew, except that the wings of the implied variance functions at deep in-the-money (ITM) and deep OTM strikes should be asymptotically linear in log-moneyness (see Lee, 2004). Hence, we transformed the constant maturity implied volatilities into the deltas of a constant maturity option, using constant maturity time series of discount rates and futures prices, and ensured that these deltas were monotonically decreasing with respect to strike; this way we ensured no-arbitrage with respect to strike. Then, to construct constant moneyness contracts, on each date in the sample and for each of the constant maturities considered, we fitted a cubic polynomial to the implied volatility skew. This approach, first used by Malz (1997), has also been applied by Weinberg (2001), Bliss and Panigirtzoglou (2002), and Lynch and Panigirtzoglou (2008).

Of course, constant maturity and moneyness contracts are artificial, nontraded instruments. Nevertheless, they allow one to examine hedging performance of a single contract over the entire sample, which would not have been possible without this construction. To focus on liquid options, we only examine the hedging of 30, 60, . . . , 180-day options with moneyness 0.95, 0.96, . . . , 1.04, 1.05.

[12]Other common approaches to fitting volatility term structures are cubic polynomials or cubic splines, and when these are used it is possible to impose the no-arbitrage constraint on the optimization as explained by Fengler (2009). However, a cubic polynomial does not fit the raw data exactly and a cubic spline is too flexible to be of much use (especially for extrapolation) and it is very sensitive to any misquotes or stale prices that may have been left in the data after the filtering.

This way, our final arbitrage-free data set consists of $16\frac{1}{2}$ years of daily prices on 66 different options, which is much larger than any other dataset previously used to study option hedging.

4. ESTIMATION OF MS SMILE DYNAMICS

The rolling sample analysis reported later in this section entails about 20,000 model estimations. Hence, to reduce complexity, we shall use the same functional form for all maturities and estimate this on all the rolling samples. The specification of our model in the form (8) is justified as follows. Starting with ordinary least squares estimation of single-regime models over the entire data period, we tested a variety of linear, log-linear, and nonlinear relationships between ATM volatility and futures price changes or returns. The best combination of high modified R^2 and coefficient t-ratios was based on a linear relationship between daily changes in ATM volatility and the daily log return on the futures, and for options up to 90 days maturity the fit is slightly improved by adding one lag of the dependent and independent variable. Then we estimated a regime-switching version of this model and applied the specification test of Breunig et al. (2003) to determine the number of states. This procedure compares the empirical density of daily changes in ATM volatility, $\hat{f}(y(i,T))$ (which we fit with a Gaussian kernel) and the unconditional density generated by the MS model, denoted $f_{MS}(y(i,T)\,|\,\hat{\theta})$ where $\hat{\theta}$ are the model parameters. The comparison of the two densities can be made through a distance metric, for instance the Kolmogorov–Smirnov (KS) statistic

$$\max|f_{MS}(y(i,T)\,|\,\hat{\theta}) - \hat{f}(y(i,T))\,|_{i=1}^{N}$$

and/or the mean square error (MSE)

$$N^{-1}\sum_{i=1}^{N}[f_{MS}(y(i,T)\,|\,\hat{\theta}) - \hat{f}(y(i,T))]^2 \qquad ,$$

where N is the number of observations. The smaller the distance between the two densities, the better specified the MS model.

For brevity, we only present the results of these model specification tests for the 30-day maturity. Table I reports the MSE and KS statistics for simple linear models with one, two, and three states and for the two-state model (8) having one lag of both dependent and independent variables. According to the KS statistic, the best model is the two-state model with lagged variables, but according to the MSE this model is only second best, the best being the two-state model without lagged variables. Nevertheless, the hypothesis that (8) is the preferred specification is supported by other statistical tests. Starting with a

TABLE I
Density Tests for Model Specification

Model	*MSE*	*KS*
1 State	0.481	0.726
2 States	0.362	0.607
2 States (lagged)	0.425	0.352
3 States	0.667	0.872

MSE, mean square error.

linear model having no lagged variables we employ the likelihood ratio (LR) and Wald (W) tests specified by Ryden et al. (1998) as follows:[13]

- H_0: 1 State vs. H_1: 2 States. Reject H_0 at 5%; $LR = -33.21$, $W = -4.1267$ (5% critical values: $LR = -40.08$, $W = -4.1277$);
- H_0: 3 States vs. H_1: 2 States. Reject H_0 at 5%; $LR = 1.82$, $W = 18.52$ (5% critical values $LR = -6.05$, $W = -56.01$);
- H_0: 2 States vs. H_1: 2 States (lagged). Reject H_0 at 5%; $LR = -24.76$, $W = 25.45$ (5% critical values $LR = -45.01$, $W = -42.55$).

This justifies the specification (8) and (9) of our two-state MS model for daily changes in ATM volatility and daily log returns on the futures of the same maturity.

4.1. Static Analysis

To analyze the empirical performance of the regime-switching deltas derived from the MS-PCA model, both the MS model and the PCA must be implemented in a dynamic (rolling sample) framework. But first, in order to provide some intuition behind the MS model's estimation and a discussion of the PCA results, we report the empirical results based on the entire sample period, again focusing on the 30-day smile for illustration.

Table II reports the parameter estimates and their standard errors when the MS model (8) and (9) is estimated using daily data between October 1992 and March 2009, for the maturity $T = 30$ days. The two regimes are clearly identified as a high-volatility regime (regime 1) in which the standard error is

[13]Since there are unidentified parameters under the null hypothesis of regime switching, standard tests do not converge to their usual distribution. Therefore, we obtain an empirical distribution of the test statistic under the null by bootstrapping and fitting the model and comparing this with the actual statistic from our data set. Since this procedure requires the estimation of a switching model for every iteration, it is very computationally intensive. We therefore set a limit of 200 simulations and used 40 randomized starting values for each simulation.

TABLE II
Two-State MS Model Estimation

Regime	*Const.*	$x(t)$	$x(t-1)$	$y(t-1)$	$\pi(i, j)$	*St. Error*
High vol.						
Est.	−0.0007	−1.2076	−0.2906	−0.2096	0.94	0.0179
SE	0.0006	0.0303	0.0488	0.0322		0.0004
Low vol.						
Est.	−0.0001	−0.6238	−0.0822	−0.1223	0.97	0.0052
SE	0.0001	0.0129	0.0199	0.0249		0.0001

MS, Markov switching.

almost four times higher than it is in the low-volatility regime (regime 2). The changes in FTSE futures log returns and lagged log returns and lagged changes in ATM volatilities are highly significant and have a negative impact on ATM volatility in both regimes, but the estimated parameters are significantly different. The size of all regression coefficients is much higher in regime 1, indicating that ATM volatility–price sensitivity is much greater when markets are volatile than when they are tranquil.[14] The volatile regime is the least persistent of the two, as it has the lower transition probability.

Figure 1 depicts the corresponding time series $\pi(t, 30)$ of the conditional regime probabilities. This identifies major periods of market turmoil such as the Asian property crash of 1997, the LTCM crisis in 1997, the bursting of the dot-com bubble in 2001–2002 and the global banking crisis of 2008–2009 which led to a global recession.

Turning now to the principal component analysis part of the model, Figure 2 depicts the first three eigenvectors of the covariance matrix of the daily changes in the FMS at maturity 30 days, as a function of moneyness. This shows that the first eigenvector, which captures almost 70% of the variation, represents a tilt in FMS with moneyness <0.99, with the other FMS remaining virtually static; the second eigenvector, which captures about 23% of the variation, represents a tilt in FMS with moneyness >0.99, with the other FMS remaining almost static; and the third eigenvector, which captures about 7% of the variation, represents a change in convexity in FMS with moneyness <1, at the same time as an almost equal and opposite change in covexity of in FMS with

[14]The ATM sensitivity to changes in the futures price, as defined in (13), is $\theta^1(30) = \beta_1(30)/F(30) = -1.2076/3{,}856.8 = -3.13 \times 10^{-4}$in the volatile regime and $\theta^2(30) = -1.62 \times 10^{-4}$ in the tranquil regime. Clearly, the reaction of ATM volatility to futures returns is very different in volatile and tranquil periods. For instance, in the volatile regime a 5% fall in 30-day FTSE futures is associated with about a 6% rise in ATM volatility (i.e. 5×-1.2076); but in the tranquil regime a similar fall in the FTSE futures precipitates little more than 3% rise in ATM volatility ($5 \times -0.6238 = 3.12\%$).

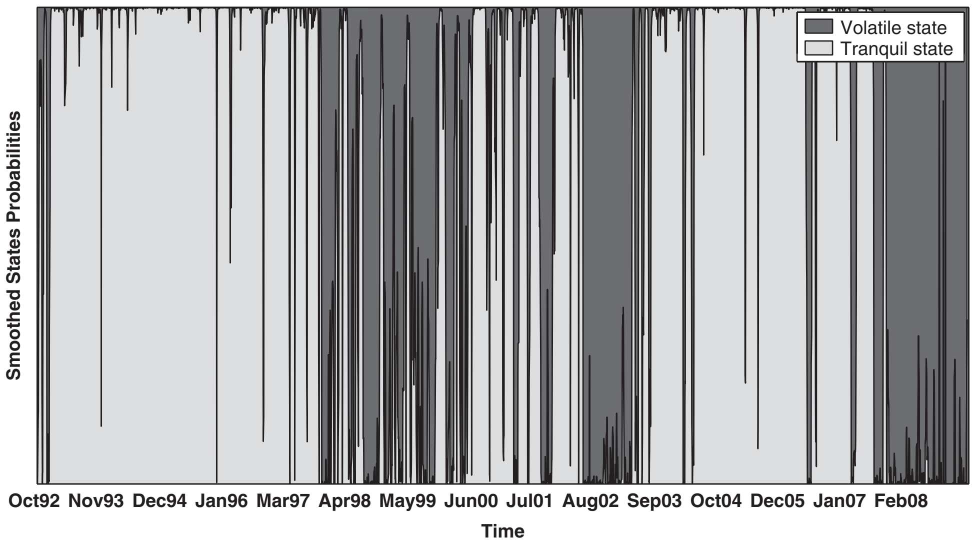

FIGURE 1
Conditional regime probabilities for model in Table I.

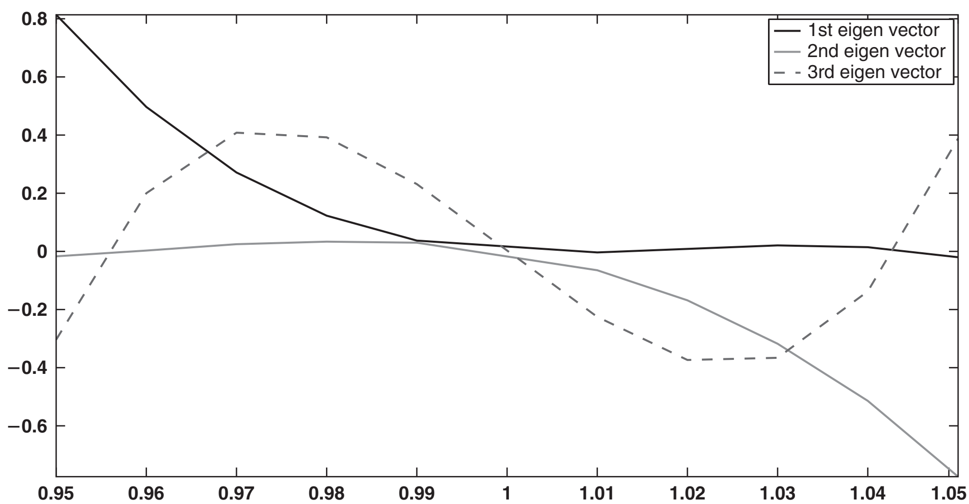

FIGURE 2
Eigenvectors at $T = 30$ days, based on the entire sample.

moneyness >1. A total of 99.75% of the variation over the entire period is captured using just three principal components. The eigenvectors have a similar shape at other maturities, although the tilt and convexity movements are less pronounced; but again virtually 100% of the variation may be captured using a three-component representation.

4.2. Dynamic Analysis

Having used the entire sample to determine which model is best specified, we now implement this model in a rolling window framework. This is to avoid "data snooping." That is, in our hedging study, we must determine the delta using data only up to the point in time when the hedge is implemented. A fairly large sample is required for the specification of MS model parameters and for this reason we fixed a sample size of 1,000 days (approximately 4 years) and estimated the models in a 4-year rolling window framework. That is, starting with the period October 1992–October 1996, for each of six different maturities $T = 30, 60, \ldots, 180$ days, we estimated the MS-PCA model using 1,000 observations on each of the variables (daily changes in ATM volatility, daily changes in FMS and daily log returns on the futures). Then each set of 1,000 observations was rolled forward one day and the MS-PCA models were re-estimated. This was repeated until all $16\frac{1}{2}$ years of data were exhausted. This way we estimated approximately 20,000 MS-PCA models.

The model is completed by estimating the coefficients $\gamma_i(t, 30)$ in (11) for the three principal components. In an advanced implementation of the model, a Kalman filter or bivariate GARCH model could be applied, but to avoid this added complexity (and since we have to about 20,000 model estimations in the dynamic setting below) we shall employ a simple exponentially weighted moving average variance and covariance estimator, with a smoothing constant of 0.94.[15] Thus, for $i = 1, 2$ and 3, we set

$$\gamma_i(t, 30) = \frac{Cov(p_i(t, 30), \Delta\, log\, F(t, 30))}{V(\Delta \log F(t, 30))}.$$

Now, taking our estimate of $\gamma_i(t, 30)$ and the 30-day futures price, both at the end of the sample period, we can apply (13) with $\hat{\beta}^1(T)$ and $\hat{\beta}^2(T)$ to obtain the regime-dependent volatility-price sensitivities $\theta_F^1(m,T|F,\ t)$ and $\theta_F^2(m,T|F,\ t)$.

For brevity, we only summarize the 30-day results here. Full details on the rolling window results are available from the authors on request. Figure 3 depicts the evolution of the percentage of variation explained by each of the first three principal components.[16] This shows that the three components capture virtually all the movements in the FMS, with the total variation explained by the three components being in excess of 99.75% for the vast majority of samples.

As expected, the volatility–price sensitivities $\theta_F^s(m,T|F,\ t)$ are very highly correlated across moneyness, for both $s = 1$ and $s = 2$. They are also very highly

[15]This value for the smoothing constant is commonly applied in the industry, being the value adopted by RiskMetrics for their daily Value-at-Risk calculations.

[16]The percentage of variation explained by the ith component is the ith largest eigenvalue of the covariance matrix of the daily changes of the FMS, based on a rolling window of 1,000 days, divided by the sum of all the eigenvalues.

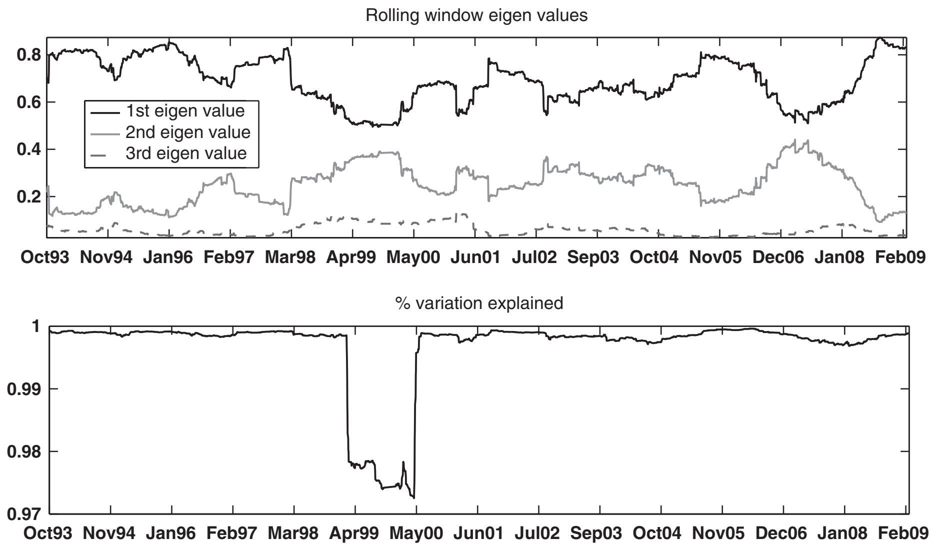

FIGURE 3
Percentage of variation explained in rolling PCA ($T = 30$ days).

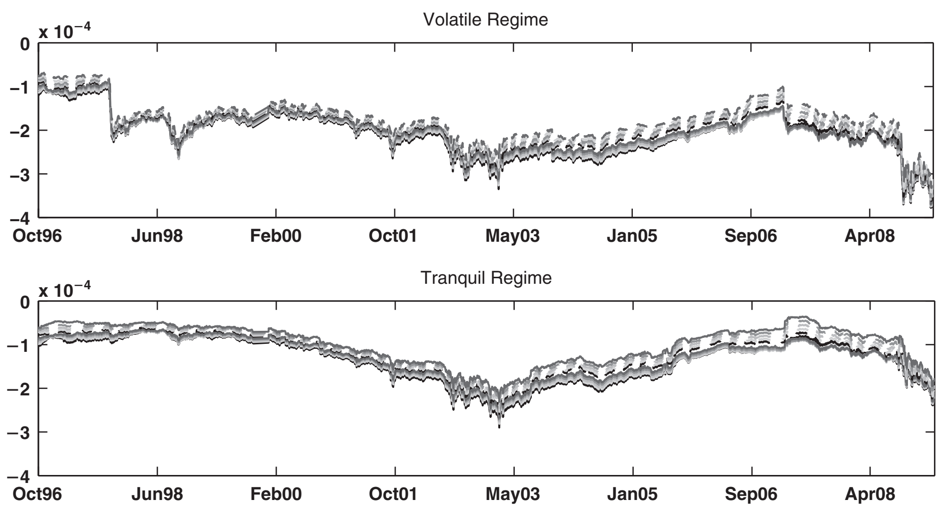

FIGURE 4
Volatility-price sensitivities over time ($T = 30$ days).

correlated across maturity, and they become more negative and less dependent on moneyness during a highly volatile period. Time series of these sensitivities are depicted in Figure 4.[17]

[17]There is very little difference between the steady-state and instantaneous sensitivities, so we only display the latter here.

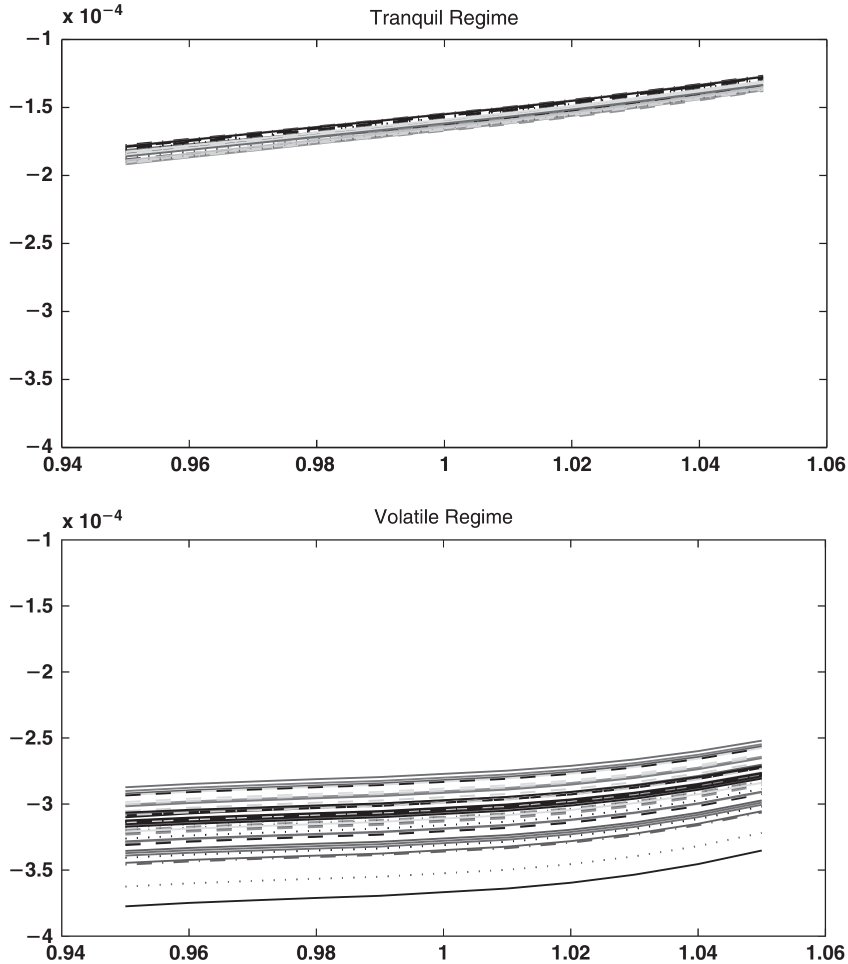

FIGURE 5
Volatility-price sensitivities: July 2005 (above) and October 2008 (below).

To gain further insight to the behavior of the sensitivities during different market regimes, we take 40 consecutive days during a tranquil period (July 2005) and plot the cross-section of $\theta_F(m, 30|F, t)$ as a function of moneyness. Figure 5 compares this with the sensitivities taken around one of the most volatile periods in the sample (October 2008). These sensitivities exhibit features that are common to other volatile and tranquil sub-samples: during crash periods they are much greater and more variable from day to day than the sensitivities during normal market circumstances. Since these sensitivities are the source of difference between various smile-adjusted deltas, we expect that the MS-PCA model will produce hedging results that are very different from those based on deltas which are not regime-dependent, especially during a market crash.

5. HEDGING RESULTS

In our delta-hedging strategies, at each time t a short position in a call option of moneyness m and maturity T is complemented by an investment in $\delta_{adj}(m,T|F,\ t)$ futures contracts of maturity T, where $\delta_{adj}(m,T|F,\ t)$ is one of the smile-adjusted delta-hedges specified in Section II. Defining $V(m,T|F,\ t)$ as the option value when the futures price is $F(t,T)$ at time t, the smile-adjusted delta is used to construct a self-financing portfolio by adding an investment of $B(m,T|F,\ t) = V(m,T|F,\ t) - \delta_{adj}(m,T|F,\ t)F(t,T)$ in a bond of maturity T that returns the risk-free rate. The initial value $\Pi(t)$ of this portfolio is zero, but if the hedge is held from time t to time $t + \Delta t$ the value of the portfolio $\Pi(t + \Delta t)$ at the end of the hedging period, i.e. the hedging error, is:

$$\Pi(t + \Delta t) = rB(m,T|F,t)\Delta t + \delta_{adj}(m,T|F,t)\Delta F(t,T) - \Delta V, \tag{14}$$

where $\Delta F(t,T) = F(t + \Delta t,T) - F(t,T)$ and $\Delta V = V(m,T|F, t + \Delta t) - V(m,T|F,\ t)$.

We use the standard deviation of the hedging errors to measure the performance of each of the seven smile-adjusted hedges defined by (3), which are summarized as follows:

$$\theta_F(m,T|F,\ t) = \begin{cases} 0 & \text{(SS)} \\ \dfrac{1}{F}\,\theta_m(m,T|F,\ t) & \text{(ST)} \\ -\dfrac{m}{F}\,\theta_m(m,T|F,\ t) & \text{(SM)} \\ \dfrac{m}{F}\,\theta_m(m,T|F,\ t) & \text{(MV)} \\ \kappa(m,T|F,\ t)\theta_m(m,T|F,\ t) & \text{(RD)} \\ \alpha(T|F,\ t) + \beta(m,T|F,\ t)\theta_m(m,T|F,\ t) & \text{(RS)} \\ F(t,T)^{-1}(\beta^s(T) + \lambda(t,T)) & \text{(MS).} \end{cases} \tag{15}$$

In the MS delta, the state indicator $s = 1, 2$ is determined by the conditional regime probability $\pi(t,T)$ at time t. We consider two variants of the MS delta: MS1, in which $s = 1$ if $\pi(t,T) > \frac{3}{4}$ and $s = 2$ otherwise; and MS2, in which $s = 1$ if $\pi(t,T) > \tau(m,T|t)$ and $s = 2$ otherwise. The parameters $\kappa(m,T|F,\ t)$, $\alpha(T|F,\ t)$, $\beta(m,T|F,\ t)$, and $\tau(m,T|t)$ are estimated using an in-sample period of 250 days, as described in Section II.

For every $(m,T) \in R$, with $R = \{m,T|m = 0.95, 0.96, \ldots, 1.05; T = 30, 60, \ldots, 180 \text{ days}\}$ and for $\Delta t = 1, 5, 10$ days we now compute the standard deviation of the out-of-sample hedging error (14) over the entire out-of-sample

TABLE III
Standard Deviation of Hedging Error, One-Day Rebalancing

	Moneyness										
Maturity	*0.95*	*0.96*	*0.97*	*0.98*	*0.99*	*1.00*	*1.01*	*1.02*	*1.03*	*1.04*	*1.05*
30-day											
SS	**12.62**	**12.86**	**13.22**	**13.53**	**13.69**	**13.60**	**13.21**	**12.54**	**11.65**	**10.65**	**9.67**
SM (%)	124	127	128	130	131	131	131	131	130	129	128
ST (%)	81	79	77	76	75	75	75	75	76	77	78
MV (%)	82	79	77	76	75	75	75	75	75	76	77
RD (%)	77	72	69	68	67	67	68	72	79	78	93
RS (%)	77	71	68	66	65	65	64	65	66	68	72
MS1 (%)	73	70	68	67	66	66	65	65	66	66	67
MS2 (%)	76	68	66	65	64	63	63	63	63	64	66
90-day											
SS	**15.78**	**15.88**	**16.07**	**16.24**	**16.32**	**16.32**	**16.22**	**15.99**	**15.64**	**15.20**	**14.76**
SM (%)	131	133	134	135	135	136	136	136	136	136	135
ST (%)	74	72	71	70	70	69	69	69	69	70	71
MV (%)	75	73	72	71	70	69	69	69	69	69	70
RD (%)	62	58	56	55	54	54	54	54	55	57	61
RS (%)	65	60	56	54	52	51	50	50	50	52	55
MS1 (%)	62	59	57	56	55	54	53	53	53	53	54
MS2 (%)	59	56	53	52	51	50	50	49	49	50	51
180-day											
SS	**20.86**	**20.61**	**20.62**	**20.65**	**20.62**	**20.53**	**20.39**	**20.20**	**19.92**	**19.59**	**19.36**
SM (%)	124	126	128	129	130	131	132	132	133	133	132
ST (%)	79	77	75	74	73	73	72	72	72	73	74
MV (%)	80	78	76	75	74	73	72	72	72	72	73
RD (%)	62	58	56	55	55	54	53	52	53	58	66
RS (%)	65	58	54	52	51	51	50	49	49	50	53
MS1 (%)	64	61	59	58	57	56	56	56	55	55	56
MS2 (%)	59	55	53	51	50	49	48	48	48	48	49

SS, sticky strike; SM, sticky moneyness; ST, sticky tree; MV, minimum variance; RD, regime dependent; RS, regime-dependent shift.

period. Because 66 options are delta-hedged and rebalanced over three different hedging horizons, and eight different deltas are compared, a total of $66 \times 3 \times 8 = 1{,}584$ standard deviations are computed. For brevity, only the results for hedging 30, 90, and 180-day options with daily and 10-day rebalancing are summarized in Tables III and IV. Examination of the other results, which are available on request, would not change any of our conclusions. The first row of each table shows the out-of-sample hedging error of the implied BSM delta-hedge (labeled SS) and the remaining rows report the other delta-hedging errors as a percentage of this.

Considering first the implied BSM delta-hedge, we find that the performance is best for daily rebalancing of short-term (30-day) options, with OTM calls being easier to hedge than ITM calls, and the hedging errors of ATM options

TABLE IV
Standard Deviation of Hedging Error, 10-Day Rebalancing

	Moneyness										
Maturity	*0.95*	*0.96*	*0.97*	*0.98*	*0.99*	*1.00*	*1.01*	*1.02*	*1.03*	*1.04*	*1.05*
30-day											
SS	**43.21**	**45.05**	**46.60**	**47.62**	**47.78**	**46.81**	**44.56**	**41.22**	**37.19**	**33.00**	**29.01**
SM (%)	117	118	120	121	122	123	124	124	125	125	125
ST (%)	87	86	84	83	82	81	80	80	79	79	80
MV (%)	88	86	85	83	82	81	80	79	79	79	79
RD (%)	83	79	76	74	72	69	66	65	69	78	97
RS (%)	77	74	71	69	67	65	63	60	57	56	56
MS1 (%)	78	77	75	74	72	70	69	67	66	65	64
MS2 (%)	76	74	72	71	69	67	65	63	61	60	59
90-day											
SS	**45.07**	**46.13**	**46.89**	**47.32**	**47.38**	**47.06**	**46.34**	**45.23**	**43.74**	**41.94**	**39.96**
SM (%)	127	128	129	129	130	131	131	131	132	132	132
ST (%)	77	77	76	76	76	75	75	75	75	75	76
MV (%)	78	78	77	76	76	75	75	75	74	74	75
RD (%)	67	66	64	64	63	63	63	64	65	67	69
RS (%)	68	65	64	62	61	61	60	60	60	61	62
MS1 (%)	67	66	65	65	64	63	62	62	62	62	62
MS2 (%)	61	60	59	58	58	57	56	57	57	58	59
180-day											
SS	**55.00**	**55.71**	**56.32**	**56.73**	**56.88**	**56.80**	**56.46**	**55.87**	**55.01**	**53.95**	**52.88**
SM (%)	124	126	127	129	130	131	132	132	132	132	132
ST (%)	79	77	76	75	74	74	73	73	73	74	75
MV (%)	80	78	77	75	74	74	73	73	73	73	74
RD (%)	57	56	57	58	60	61	61	62	65	71	81
RS (%)	66	61	58	56	54	54	54	54	56	60	70
MS1 (%)	62	61	60	59	58	57	57	56	56	56	56
MS2 (%)	57	55	53	52	51	50	50	49	50	51	52

SS, sticky strike; SM, sticky moneyness; ST, sticky tree; MV, minimum variance; RD, regime dependent; RS, regime-dependent shift.

having larger standard deviations than both OTM and ITM calls. The hedging performance deteriorates as the maturity of the option increases and especially as the hedging horizon increases (e.g., compare the daily rebalancing results in Table III with the 10-day rebalancing results in Table IV). Over daily and longer hedging horizons, OTM calls remain easier to hedge than other options, whatever their term. There is little or no improvement in the ITM call option's hedging performance over that for ATM options for the longer term options.

Now consider the seven other deltas, whose performance in Tables III and IV is presented as a percentage of the implied BSM delta's performance. In agreement with the previous literature, the SM delta performs much worse than the implied BSM delta-hedge for all options, irrespective of the rebalancing frequency. However, all other smile-adjusted hedges improve on the implied BSM hedge and the extent of this improvement increases with the maturity of the options. The ST and MV deltas perform very similarly, as expected, with the ST delta being marginally better for ITM call options and the MV delta being marginally better for OTM call options—for ATM options, the deltas are the same. Much the best performance, for all options, is found with regime-dependent deltas, particularly when using with the RS and MS deltas.

For options with maturity greater than 30 days, the use of the full MS model with an optimized threshold (the MS2 delta) has a clear advantage: for instance the standard deviation of the hedging errors is roughly half of that from the implied BSM delta-hedge at both daily and 10-day hedging horizons, for ATM and OTM call options. This is a really substantial improvement on implied BSM hedging. The RS delta, which is much easier to implement than the MS deltas, is also very effective. In fact, with 30-day options and 10-day rebalancing the RS delta outperforms even the MS2 delta for near ATM and OTM options.

We now investigate how delta-hedging performance varies over time, as the market experiences volatile and tranquil periods. Is delta hedging more effective during tranquil periods? Does the choice of delta-hedging model matter more during volatile periods than during tranquil periods? The answer to both these questions is most definitely yes.

For 30-day call options with moneyness 0.95, 1.00, and 1.05, Figure 6 depicts the standard deviations of the daily-rebalanced hedging errors computed out-of-sample over the previous 250 days, rolled over daily. The standard deviation of the MV delta-hedged portfolios was extremely close to those of the ST delta-hedged portfolios (and are identical for ATM options) so we have omitted them, for greater clarity in the figures. The results for hedging options with different moneyness and with maturity from 60 to 180 days are not shown, for brevity, but are available from the authors on request. They exhibit very similar features to Figure 6, but with the general level of the standard deviation increasing with the options' maturity. The hedging error standard deviations are closely related to the market volatility, as is evident on comparing Figure 6 with the 3-month ATM volatility series depicted in Figure 7. Clearly, the uncertainty in the hedging error tends to increase during volatile periods.

According to the criterion of minimizing the hedging error standard deviation, all delta-hedges except SM improve on the BSM, with two exceptions: (1) hedging OTM calls (moneyness 1.05) during the very tranquil period between April 2004 and May 2006, where any delta hedging is particularly effective and it matters very little which model is used; and (2) the year from July 2006 to July 2007, where the simple regime dependent (RD) delta led to an unexpectedly large hedging error standard deviation for the OTM calls.

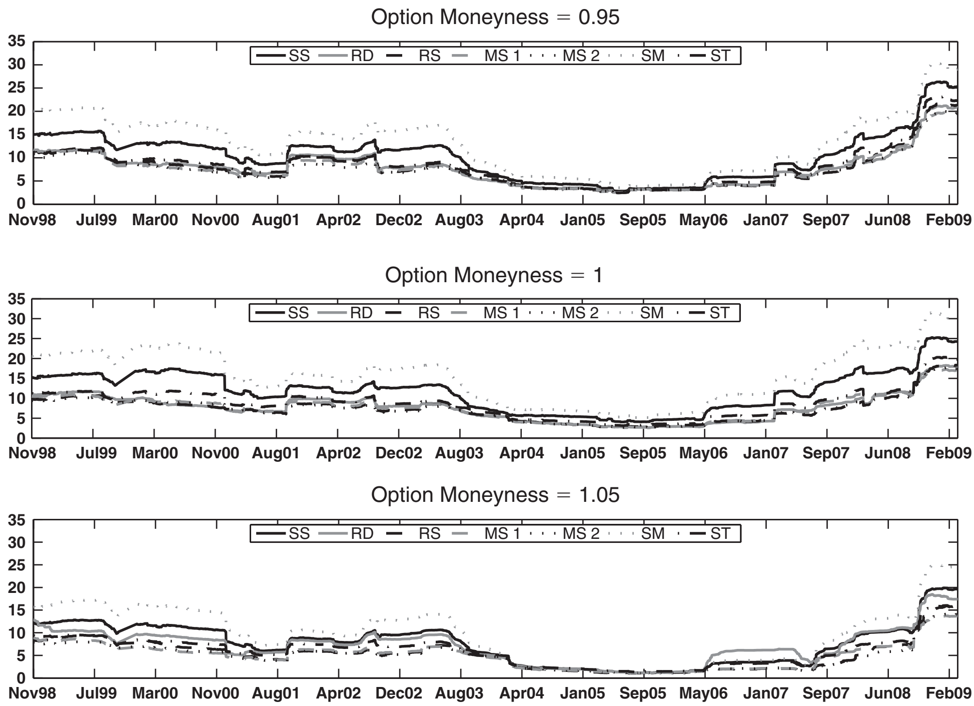

FIGURE 6
Standard deviation of hedging error over previous year (daily rebalancing of delta-hedged 30-day calls).

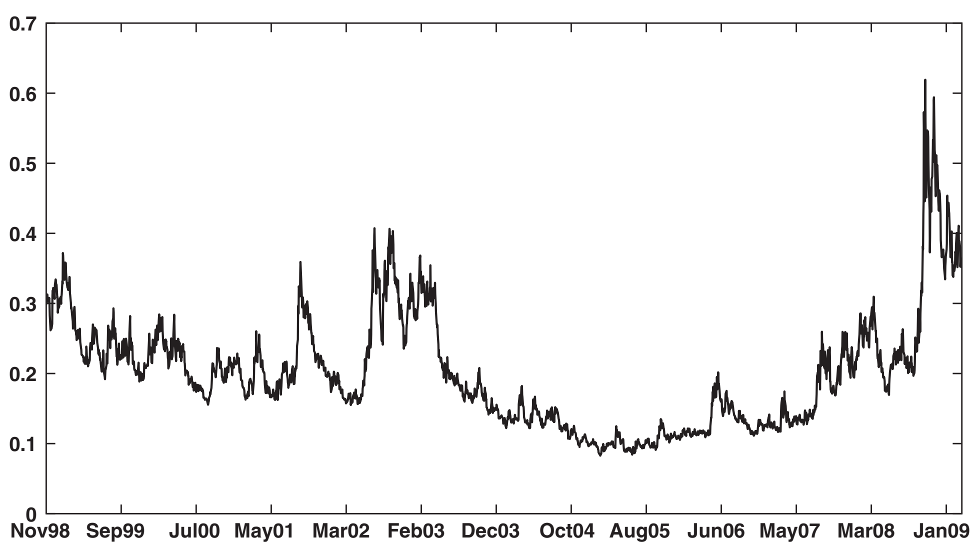

FIGURE 7
FTSE 100 at-the-money volatility.

Finally, we note that the RS and MS hedges perform almost identically when hedging ATM options, but as the option moves away from ATM the formal MS model that we have introduced, especially with an optimized threshold delineating volatile and tranquil regimes, gives a clear advantage to traders who seek to reduce their spreads by using the most accurate delta-hedges possible.

6. SUMMARY AND CONCLUSIONS

We have employed a very much larger sample than any other hedging study to date ($16\frac{1}{2}$ years of daily data on FTSE 100 index options) to calculate the efficiency gains, relative to the implied BSM delta, from utilizing several different smile-adjusted deltas. Considering the standard deviation of out-of-sample delta-hedging errors from 1-, 5- and 10-day rebalancing of delta-hedged vanilla call options, with moneyness 0.95, 0.96, . . . , 1.05 and maturities 30, 60, . . . , 180 days, we provide very strong evidence that tailoring the delta-hedge to the market regime makes a huge difference to hedging performance. On average during the entire sample period, the regime-dependent delta-hedges produce errors that are almost half the size of implied BSM delta-hedging errors, whatever the rebalancing period, with only slightly less of an efficiency gain even when hedging very short-term (30-day) options. For all options and all rebalancing periods, regime-dependent deltas are especially effective during volatile market periods.

Considering seven different smile-adjustments to the implied BSM delta, four of which are regime-dependent, we show that the essential property to capture is that there is a *shift* of the smile accompanying a change in the underlying price. Furthermore, the size of this shift should depend on the current market regime. Three of our deltas have this property, and they each offer a substantial improvement over implied BSM, local volatility, and minimum variance delta hedging. One of these regime-dependent deltas is rather ad hoc, but it is very simple to implement using only the current smile and a short history of ATM volatility. Nevertheless, our results show that the most effective regime-dependent hedging of vanilla FTSE 100 index options is based on a MS relationship between the ATM volatility and the underlying future price, which has a straightforward extension to options of other moneyness via standard principal component analysis.

BIBLIOGRAPHY

Ait-Sahalia, Y. (1996). Nonparametric pricing of interest rate derivative securities. Econometrica, 64, 527–560.

Alexander, C. (2001). Principles of the skew. Risk, 14(1), S29–S32.

Alexander, C., & Kaeck, A. (2010). Does model fit matter for hedging? Evidence from FTSE 100 options. ICMA Centre Discussion Papers in Finance, DP2010-05. Available from http://www.carolalexander.org.
Alexander, C., Kaeck, A., & Nogueira, L. (2009). Model risk adjusted hedge ratios. Journal of Futures Markets, 29(11), 1021–1049.
Alexander, C., & Nogueira, L. (2007). Model-free hedge ratios and scale-invariant models. Journal of Banking and Finance, 31(6), 1839–1861.
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, 277–318.
Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. Journal of Political Economy, 81(3), 637–654.
Bliss, R., & Panigirtzoglou, N. (2002). Testing the stability of implied probability density functions. Journal of Banking and Finance, 26, 381–422.
Breunig, R., Najarian, S., & Pagan, A. (2003). Specification testing of Markov switching models. Oxford Bulletin of Economics and Statistics, 65(1), 703–725.
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(1), 63–89.
Crepey, S. (2004). Delta-hedging vega risk? Quantitative Finance, 4, 559–579.
Derman, E. (1999). Volatility regimes. Risk, 14:2, 55–59.
Derman, E., & Kani, I. (1994). Riding on a smile. Risk, 7(2), 32–39.
Derman, E., Kani, I., & Zou, J. (1996). The local volatility surface: Unlocking the information in index option prices. Financial Analysts Journal, 52, 25–36.
Dumas, D., Pan, J., & Whaley, R. (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.
Fengler, M. (2009). Arbitrage-free smoothing of the implied volatility surface. Quantitative Finance, 9(4), 417–428.
Frey, R. (1997). Derivative asset analysis in models with level-dependent and stochastic volatility. CWI Quarterly, 10(1), 1–34.
Fruhwirth-Schnatter, S. (2006). Finite mixture and Markov switching models. Springer Series in Statistics.
Hamilton, J. (1989). A new approach to the economic analysis of non-stationary time series and the business cycle. Econometrica, 57(2), 357–384.
Heston, S. (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.
Jackwerth, J., & Rubinstein, M. (2001) Recovering stochastic processes from option prices. Journal of Finance, 52(3), 1236 (abstract).
Kim, I., & 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, 117–142.
Lee, R. (2001). Implied and local volatilities under stochastic volatility. International Journal of Theoretical and Applied Finance, 4(1), 1178–1192.

Lee, R. (2004). The moment formula for implied volatility at extreme strikes. Mathematical Finance, 14(3), 469–480.

Lynch, D., & Panigirtzoglou, N. (2008). Summary statistics of option-implied probability density functions and their properties (Working Paper No. 345). Bank of England.

Malz, A. (1997). Estimating the probability distribution of the future exchange rate from option prices. Journal of Derivatives, 5(2), 18–36.

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

Merton, R. (1976). Option pricing when underlying stock returns are discontinuous. Journal of Financial Economics, 3, 125–144.

Poulsen, R., Schenke-Hoppe, K., & Ewald, C-O. (2009). Risk minimization in stochastic volatility models: Model risk and empirical performance. Quantitative Finance, 9(6), 693–704.

Rosenberg, J. (2000). Implied volatility functions: A reprise. Journal of Derivatives, 7, 51–64.

Ryden, T., Teräsvirta, T., & Asbrink, S. (1998). Stylized facts of daily return series and the hidden Markov model. Journal of Applied Econometrics, 13, 217–244.

Schweizer, M. (1991). Option hedging for semimartingales. Stochastic Processes and their Applications, 37, 339–363.

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

Weinberg, S. (2001). Interpreting the volatility smile: An examination of the information content of option prices, Discussion Paper No. 706, Board of Governors of the Federal Reserve System.