Rev Deriv Res (2012) 15:57–79
DOI 10.1007/s11147-011-9069-2

Calibration risk: Illustrating the impact of calibration risk under the Heston model

Florence Guillaume · Wim Schoutens

Published online: 8 July 2011

Abstract It is already well documented that model risk is an important issue regarding the pricing of exotics (see Schoutens et al., in A perfect calibration! Now what?, Wilmott Magazine, March 2004: pp 66–78, 2004). Arguments have been made to put this into the perspective of bid-ask pricing using the theory of conic finance and pricing to acceptability (Cherny and Madan Review of Financial Studies, 22: 2571–2606, 2009). In this paper we show also the presence and importance of calibration risk. More particularly, we point out that a variety of plausible calibration methods lead again to serious price differences for exotics and different distributions of the P&L of the delta-hedging strategy. This is illustrated under the popular Heston stochastic volatility model, which is used among practitioners to price all kinds of exotic and structured products. This paper shows that it is prudent to take some additional safety margin into account for the pricing of these structured notes.

Keywords Heston model · Calibration · Model risk · Calibration risk · Exotic options

JEL classification C63 · G17

F. Guillaume (✉)
Department of Mathematics, T.U.Eindhoven, Eurandom, P.O.Box 513, 5600 MB Eindhoven, The Netherlands
e-mail: guillaume@eurandom.tue.nl

W. Schoutens
Department of Mathematics, K.U.Leuven, Celestijnenlaan 200 B, 3001 Leuven, Belgium
e-mail: Wim@Schoutens.be

1 Introduction

The multiplicity of financial models has inevitably given rise to what is nowadays commonly referred to as model uncertainty or model risk. Indeed the fair price of any financial instrument typically strongly depends on the model which is used to describe the dynamics of the financial market. Moreover, different models yielding the same precision for benchmark instruments turn out to lead to a substantial uncertainty once we are faced with the pricing of more exotic products (see for instance Cont 2006 or Schoutens et al. 2004). In particular, if several trading desks are asked to price a set of exotic derivatives from similar market quotes for the most liquid derivatives, they will probably come up with significantly different prices for the more exotic products. This relatively new kind of risk, due to the choice of the model itself, has significantly increased this last decade given the rapid growth of the derivative market and has led to substantial losses arising from the misvaluation of financial derivatives. In particular, in January 2004 the national Australian Bank reported losses on currency options amounting to more than 280 million U.S. dollars, partially due to misvaluation (see Stulz 2004). The tremendous losses due to model uncertainty suffered by many financial institutions have led the Basel Committee on banking supervision to adopt a directive which compels financial institutions to take into account the uncertainty of the model valuation in the mark-to-model valuation of structured products. In particular the Basel Committee distinguished two types of model risk: the one arising by making use of a *possibly incorrect valuation* procedure and the one due to *unobservable and thus possibly incorrect calibration parameters* (see Basel Committee on Banking Supervision 2009).

Once a particular class of models has been selected, financial institutions are faced with the construction of the risk-neutral pricing measure $\mathbb{Q}$. This calibration procedure consists of the determination of the model parameter set which is compatible with the observed market prices of liquidly traded derivatives. Trading desks thus select the parameters $\mathbf{p}^\star$ which minimize the discrepancy $f(\{P_i\}, \{\hat{P}_i\}; \mathbf{p})$ between the model prices $\hat{P}_i$ and the market prices P_i of benchmark instruments. In particular, they are faced with choosing a particular objective function f which links the model prices with the market quotes. The current industry practice is to use the root mean square error leading to a least-squares problem

$$\mathrm{RMSE} = \sqrt{\sum_{j=1}^{N} \frac{(P_j - \hat{P}_j)^2}{N}}, \tag{1}$$

where N stands for the number of benchmark instruments. However, there exist other alternatives just as suitable such as the average absolute error as a percentage of the mean price (APE) or the average relative error (ARPE):

$$\mathrm{APE} = \frac{1}{\mathrm{mean}_j \hat{P}_j} \sum_{j=1}^{N} \frac{\left|P_j - \hat{P}_j\right|}{N} \quad \text{and} \quad \mathrm{ARPE} = \frac{1}{N} \sum_{j=1}^{N} \frac{\left|P_j - \hat{P}_j\right|}{\hat{P}_j}.$$

The optimal parameter set $\mathbf{p}^\star$ typically strongly depends on the choice of the objective function, leading to significantly different prices for the more exotic options such as one-touch barrier or cliquet options. This particular type of model risk which is inherent to the calibration procedure was introduced by Detlefsen and Härdle 2007 under the terminology calibration risk and formally defined as the different optimal parameter sets arising from the different specifications of the functional f. In this paper, we extend the concept of calibration risk to include not only the choice of the objective function but also the calibration methodology and illustrate its impact under the Heston stochastic volatility model by pricing a wide range of exotics and by looking at the P&L distribution of the delta hedging strategy.

2 Calibration of the Heston model

The recent emergence of a liquid market for volatility derivatives has given rise to new calibration procedures. Before, equity models were calibrated on the basis of the implied volatility surface only by the minimization of a particular objective function. This standard calibration often leads to optimal parameters which are instable or set at extreme levels. However, given the substantial liquidity of the volatility market, practitioners might resort to time series or market quotes to determine some of the model parameters beforehand and perform therefore a calibration on a reduced parameter set. In particular, the spot variance of the Heston model can be inferred from the spot value of the VIX volatility index whereas the long run variance can be determined either from the VIX time series or from the VIX option price surface. The remaining parameters are then calibrated from the stock option price surface by minimizing a particular functional. The different calibration procedures will typically lead to different optimal parameter sets and hence to different exotic prices and different hedge ratios. This gives evidence that it is prudent to consider some additional safety margin for the pricing of these exotics, as it has been advised by the Basel Committee.

2.1 The Heston stochastic volatility model

In (Heston 1993), Heston extends the Black-Scholes model by making the volatility parameter σ stochastic. More particularly, the squared volatility is modeled by a CIR process, which is coherent with the positivity and mean-reverting characteristics of the empirical volatility (Whaley 2009). The stock price process follows the well-known Black-Scholes stochastic differential equation:

$$\frac{dS_t}{S_t} = (r-q)dt + \sqrt{v_t}dW_t, \quad S_0 \geq 0 \tag{2}$$

and the squared volatility process follows the CIR stochastic differential equation:

$$dv_t = \kappa\,(\eta - v_t)\,dt + \lambda\sqrt{v_t}d\tilde{W}_t, \quad v_0 = \sigma_0^2 \geq 0, \tag{3}$$

where $W = \{W_t, t \geq 0\}$ and $\tilde{W} = \left\{\tilde{W}_t, t \geq 0\right\}$ are two correlated standard Brownian motions such that $\mathrm{Cov}\left(dW_t, d\tilde{W}_t\right) = \rho dt$. v_0 is the initial variance, $\kappa > 0$ the mean reversion rate, $\eta > 0$ the long run variance, $\lambda > 0$ the volatility of variance and ρ the correlation. The variance process (3) is always positive and can not reach zero if $2\kappa\eta > \lambda^2$. The model parameters can be determined either by matching data or by calibration. In practice, parameters calibrated on the implied volatility surface might turn out to be unstable and often unreasonable (see Wilmott 2006). Hence, we will propose an alternative methodology which consists of a mixture of matching data and calibration in order to enhance the stability of the model parameters and decrease the calibration computation time.

2.2 Calibration sets

We start with describing three different ways of estimating/calibrating the long run variance parameter η. The first two estimates of the long run variance are determined on the basis of time series, in the example the VIX index, and are therefore historical estimates. The final estimate is a market-implied one that allows us to take into account additional market information in the calibration methodology. In the illustration, it is derived from the VIX option surface which is not a calibration instrument in the standard calibration.

- The **moving window** (MW) estimate
  The moving window estimate is computed as the mean of the variance of the stock price process over a time series window which moves forward through time:

$$\eta^{\mathrm{MW}} = \frac{1}{T^{\mathrm{VIX}}} \int_{t_0 - T^{\mathrm{VIX}}}^{t_0} \left(\frac{\mathrm{VIX}(t)}{100}\right)^2 dt = \underset{t_0 - T^{\mathrm{VIX}} \leq t \leq t_0}{\mathrm{mean}} \left(\frac{\mathrm{VIX}(t)}{100}\right)^2. \tag{4}$$

  For the numerical study, we will consider a length of the VIX time series window equal to 6 months, 3 or 5 years. Since the choice of the length of the time series window to fix the long run variance strongly depends on the market volatility regime (see Guillaume and Schoutens 2010), it might be interesting to consider estimates of η which are independent of the length of the time series window, such as the exponentially weighted moving average estimate.
- The **exponentially weighted moving average** (EWMA) estimate
  The exponentially weighted moving average estimate of the long run variance is given by

$$\eta^{\mathrm{EWMA}} = (1-\alpha) \sum_{i=1}^{N} \alpha^{N-i} \left(\frac{\mathrm{VIX}(t_i)}{100}\right)^2 \tag{5}$$

  where $\alpha \in (0, 1)$, $t_i = t_0 - (N-i)\Delta t$ and where $N \to \infty$ is the number of data in the time series. The most recent the VIX quote, the highest the corresponding weight. In particular, the weight α^{N-i} decreases exponentially as we move back through time. The parameter α determines how responsive the estimate η^{EWMA} is to the most recent daily percentage change of the VIX: a low value of α corresponds to a highly volatile estimate. For the numerical study, we consider a parameter α equal to 0.94, which is also the value used by *JP Morgan* for the *RiskMetrics* database (see Hull 2006).

– The **market-implied** (MI) estimate
Finally, we propose a robust way of computing a risk-neutral long run volatility estimate, referred to as long VIX, or LVIX and inferred from the market price of long term European vanilla options on the VIX; reflecting therefore the expectations of the investors.
The Chicago Board Options Exchange (CBOE) launched trading of VIX option contracts on the 24th of February 2006. After 9 months, the VIX option trading volume almost reached 4.5 million contracts, making it the most successful new product launched in CBOE history. Hence, given the substantial liquidity of VIX options, we expect to infer an accurate estimate of the long run volatility from their market quotes.
From the non-arbitrage principle, the long run volatility should be equal to the at-the-money strike for long term VIX options. Indeed, from the put-call parity, we have that

$$P(K,T) - C(K,T) = \exp(-rT)(K - \text{VIX}_T)$$

and in particular $K = \text{VIX}_T$ iff $P(K,T) = C(K,T)$. Hence, the long run variance can be approximated by the at-the-money strike of long term options on the VIX:

$$\eta^{\text{MI}} = \left(\frac{K^{\text{ATM}}}{100}\right)^2 = \left(\frac{\text{LVIX}(t_0)}{100}\right)^2. \tag{6}$$

For the numerical study, we will consider the options with the longest available quoted maturity. The at-the-money strike is obtained by interpolation of the call-put spread.

We will compare six different calibration performances of the Heston model: a fully free parameter set $\{v_0, \kappa, \eta, \lambda, \rho\}$ and five reduced parameter sets $\{\kappa, \lambda, \rho\}$, using the market data to fix v_0 and η, where v_0 is set equal to the square of the spot price of the VIX index expressed in units:

$$v_0 = \left(\frac{\text{VIX}(t_0)}{100}\right)^2;$$

and η is estimated either

– on the basis of the empirical VIX index by Eq. (4) with three different time series windows (0.5, 3 and 5 years);
– on the basis of the empirical VIX index by Eq. (5);
– using the VIX option market quotes by Eq. (6).

For the numerical study, we consider daily S&P 500 and VIX market quotes for a period extending from the 24th of February 2006 until the 31st of October 2009, including therefore the credit crunch. The different calibrations are performed on the whole set of quoted vanilla options (i.e. for both call and put options on the whole strike and time to maturity ranges). The vanilla option prices are computed by using the Carr-Madan formula (see Carr and Madan 1998 and Albrecher et al. 2007 for the closed-form expression of the Heston characteristic function).

3 Calibration performance and evolution of the model parameters through time

Except for a period extending from mid-October 2008 until mid-December 2008, the RMSE functional (1) obtained by considering the market implied estimate of η is pretty close to the RMSE functional of the full calibration procedure (see Fig. 1[1]). In particular, η^{MI} leads to a better fit of the S&P 500 implied volatility surface than the time series estimates obtained by the MW or the EWMA technique; which is coherent with the fact that η^{MI} reflects the future expectations of market participants whereas η^{MW} and η^{EWMA} reflect their past expectations.

Moreover, the moving window calibration performance depends on the length of the time series window. In particular, we can distinguish two periods: the first one extending from February 2006 until July 2007 and the second one from August 2007 until October 2009. For the first period, the optimal moving window calibration is obtained by considering the widest time series window, i.e. T^{VIX} equal to 5 years whereas, for the second period, the narrowest window (i.e. T^{VIX} equal to 6 months) leads to the best calibration of the S&P 500 option surface.

To have some insight into the precision of the different calibration settings, it is interesting to have a look at the evolution through time of the two parameters which are inferred beforehand from the VIX quotes, i.e. the initial variance v_0 and the long run variance η.

From the evolution of the spot VIX (see upper part of Fig. 2), it is clear that the two periods previously mentioned correspond to two different volatility regimes. More precisely, the transition in between the two periods coincides with the beginning of the credit crisis, characterised by a substantial increase of the VIX index. Moreover, the market implied spot variance, as well as the calibrated v_0 to a smaller extent, exhibit a sharp increase from mid-September 2008. This date coincides with the bankruptcy of *Lehman Brothers* which is therefore the trigger of the panic wave occurring between October and December 2009, where the market spot variance reached some exceptional level of more than 65%, which corresponds to a volatilty near 80%.

By comparing Figs. 1 and 2, we note that the significant difference between the RMSE objective function of the standard calibration and this of the market implied setting over the period ranging from mid-October to mid-December 2008 might be explained by the significant difference between the calibrated v_0 and the square of the spot VIX during the panic wave. Nevertheless, except for this 2 months period, the square of the spot VIX appears to be close to the initial variance parameter obtained by

[1] Given the significant difference in the magnitude order of the different quantities of interest, the total period will be typically split into two parts.

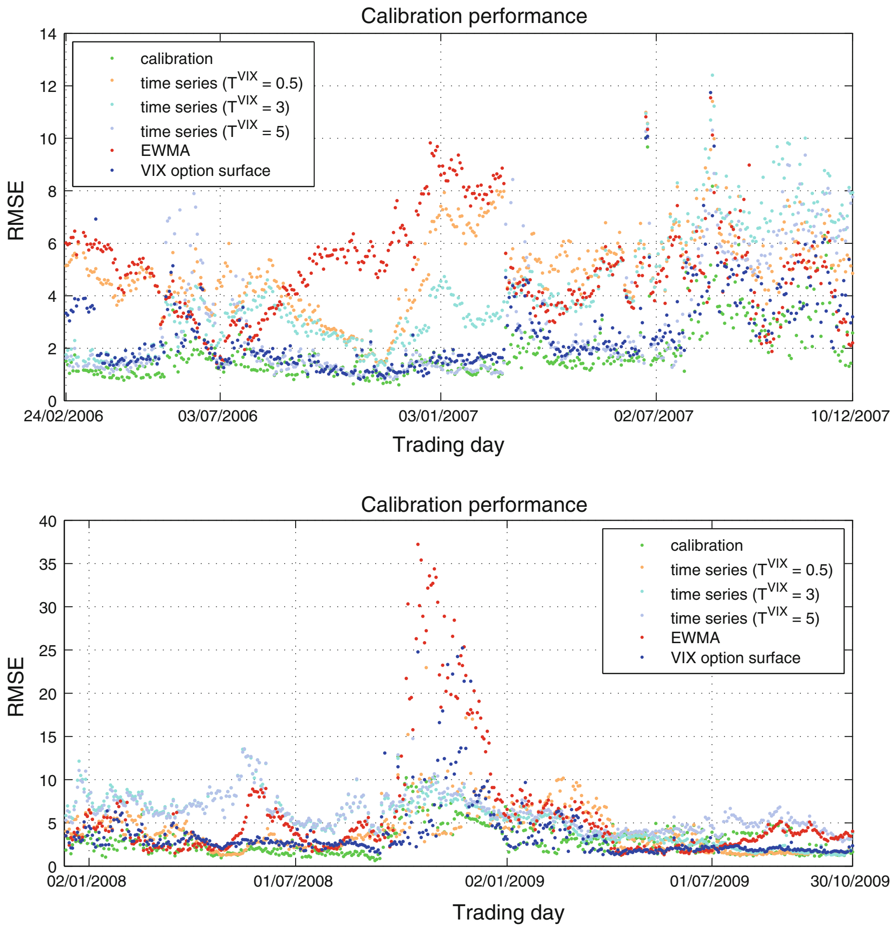

Fig. 1 Evolution of the RMSE through time for the different calibration procedures

calibrating the whole parameter set on the S&P 500 volatility surface and both exhibit a similar trend.

On the other hand we note that the estimate of the long run variance η turns out to be significantly different from one calibration procedure to the other. In particular, the time series estimates and the market implied estimate of η are typically lower than the value of η resulting from the standard calibration. However, we notice some exception for the EWMA and the 6 months MW estimate from October 2008 until May 2009 and for the market implied estimate from October 2008 until December 2008, which might be explained by the relatively high value of the market implied spot variance in comparison to the calibrated value of v_0 during the investors' fear wave of the end of 2008. Furthermore the moving window estimates exhibit a significantly smoother trend than the calibrated parameter, especially for a wide time series window whereas both the market implied and the EWMA estimate of η roughly follow the same trend as the calibrated parameter. Moreover, the market implied estimate and to a larger

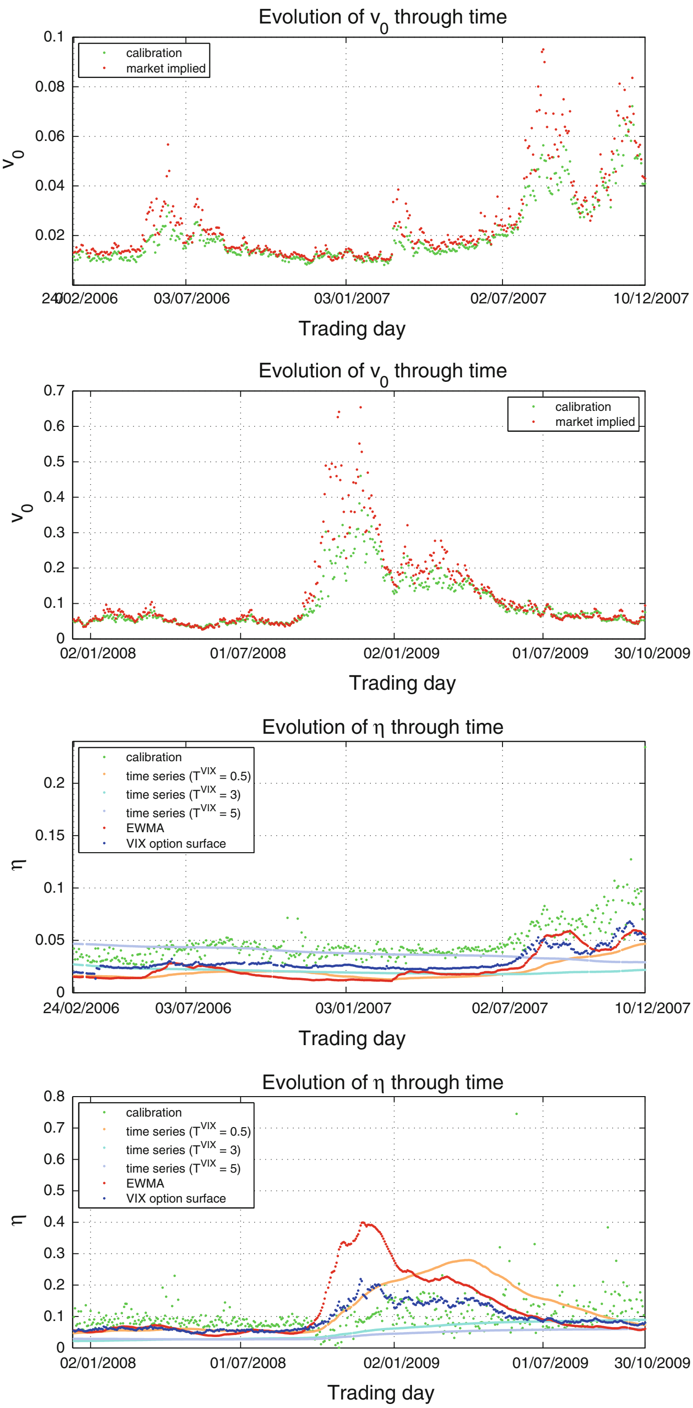

Fig. 2 Evolution of the parameter v_0 (*upper*) and η (*lower*) through time for the different calibration procedures

extent the EWMA estimate exhibit the same, although clearly smoother, trend as the spot variance (or equivalently as the square of the spot VIX).

The calibrated parameter usually turns out to be closer to η^{MI} than to η^{EWMA} or η^{MW} which explains the better fit of the S&P 500 option price surface obtained by inferring the long run variance from VIX option quotes. For the first period, the MW long run variance computed for a period T^{VIX} of 5 years leads to an accurate fit of the S&P 500 option surface since then the average of the calibrated parameter η turns out to be of the same order of magnitude than η^{MW}. On the other hand, for the second period the 6 months window leads to the smallest difference between the calibrated and moving window estimates of the long run variance.

From the lower part of Fig. 2, it is clear that both η^{MI} and η^{EWMA} immediately react to the switch in the volatility regime, which is consistent with their definition. Indeed, the computation of η^{MI} is directly inferred from current market quotes and η^{EWMA} is assessed by associating more weight to the spot variance than to historical variances. On the other hand, the reflection of the market trend in the moving window estimate is delayed by a period which increases with the length of the time series window. In particular, the sixth months MW estimate reflects almost immediately the switch from one volatility regime to the other, although we clearly observe a delay in between the reaction time of this estimate and the reaction time of both the EWMA and MI estimates. On the other hand, the three and 5 years estimates need more than 1 year to reflect the switch of volatility regime. Moreover, these two estimates are not able to reflect the peak in the VIX index occurring during the panic wave, which might be explained by the (too) wide time series window.

The matching of the spot variance and the long run variance results in some adjustment of the other parameters calibrated on the S&P 500 implied volatility surface. Figure 3 shows the evolution through time of these parameters, i.e. λ, κ and ρ under both the standard and the reduced settings. In particular, the parameters λ and κ are typically set to a higher value than the optimal ones whereas the correlation ρ is typically set to a lower value.

4 Pricing of exotic options

Both the standard and the market implied reduced calibration procedure lead to a pretty good fit of the whole set of liquid S&P 500 options except during the panic wave period (see Fig. 1 or Guillaume and Schoutens 2010 for more evidence). Hence, we can usually hardly discriminate between the two calibration methods from the sole replication of the market price of benchmark instruments. We will next compare the calibration procedures by pricing several exotic options ranging from one-touch barrier options, lookback options and cliquet options.

The path dependent nature of exotic options requires the use of the Monte Carlo procedure to simulate sample paths of the underlying index and its volatility, or equivalently its variance. The stock price process (2) is discretised by using a first order Euler scheme and the variance process (3) using a second order Milstein scheme. The Monte Carlo simulation is performed by considering one million scenarios and 252 trading days a year.

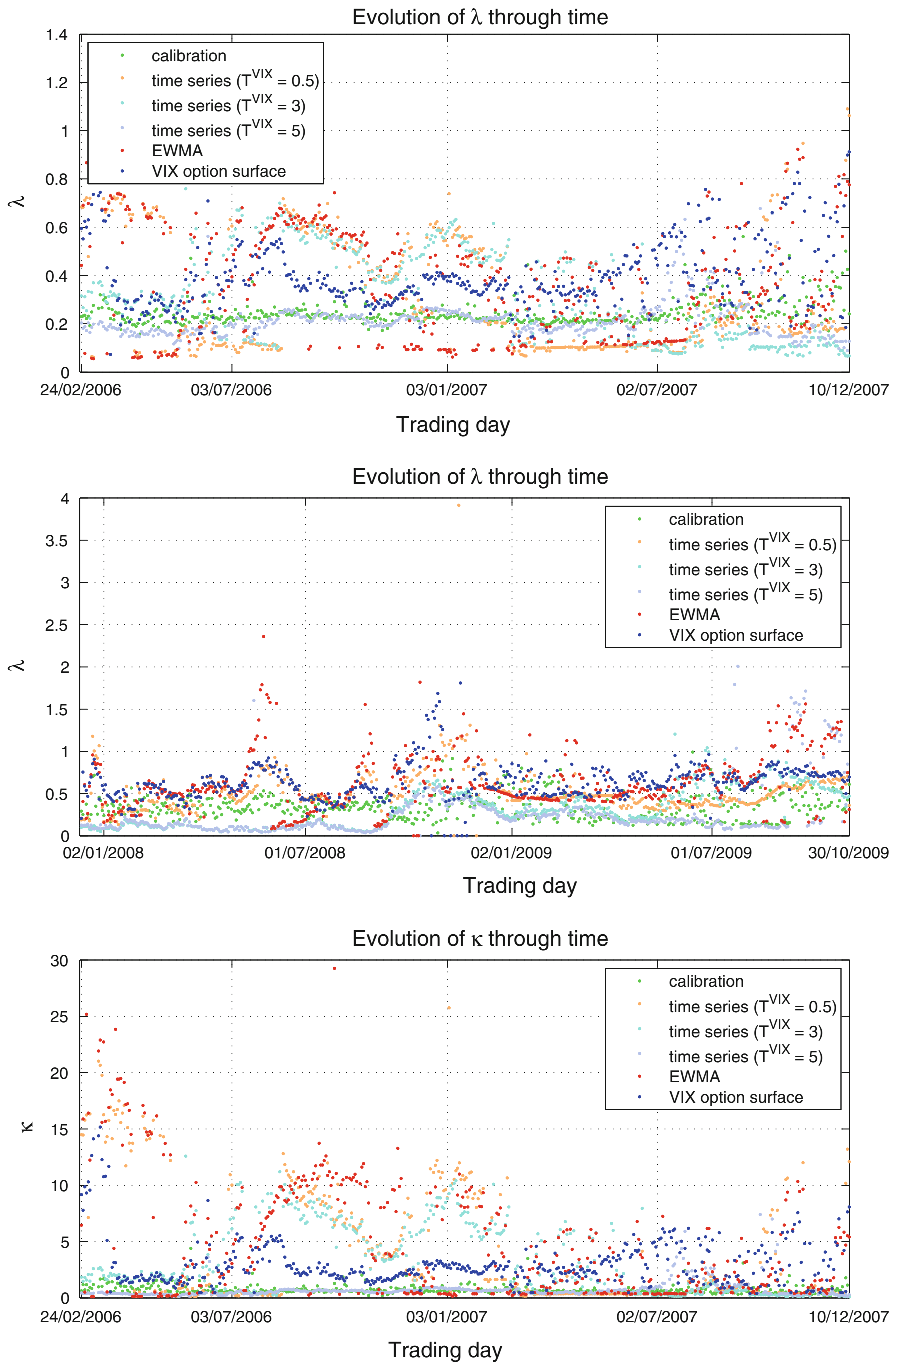

Fig. 3 Evolution of the parameters λ (*upper*), κ (*center*) and ρ (*lower*) through time for the different calibration procedures

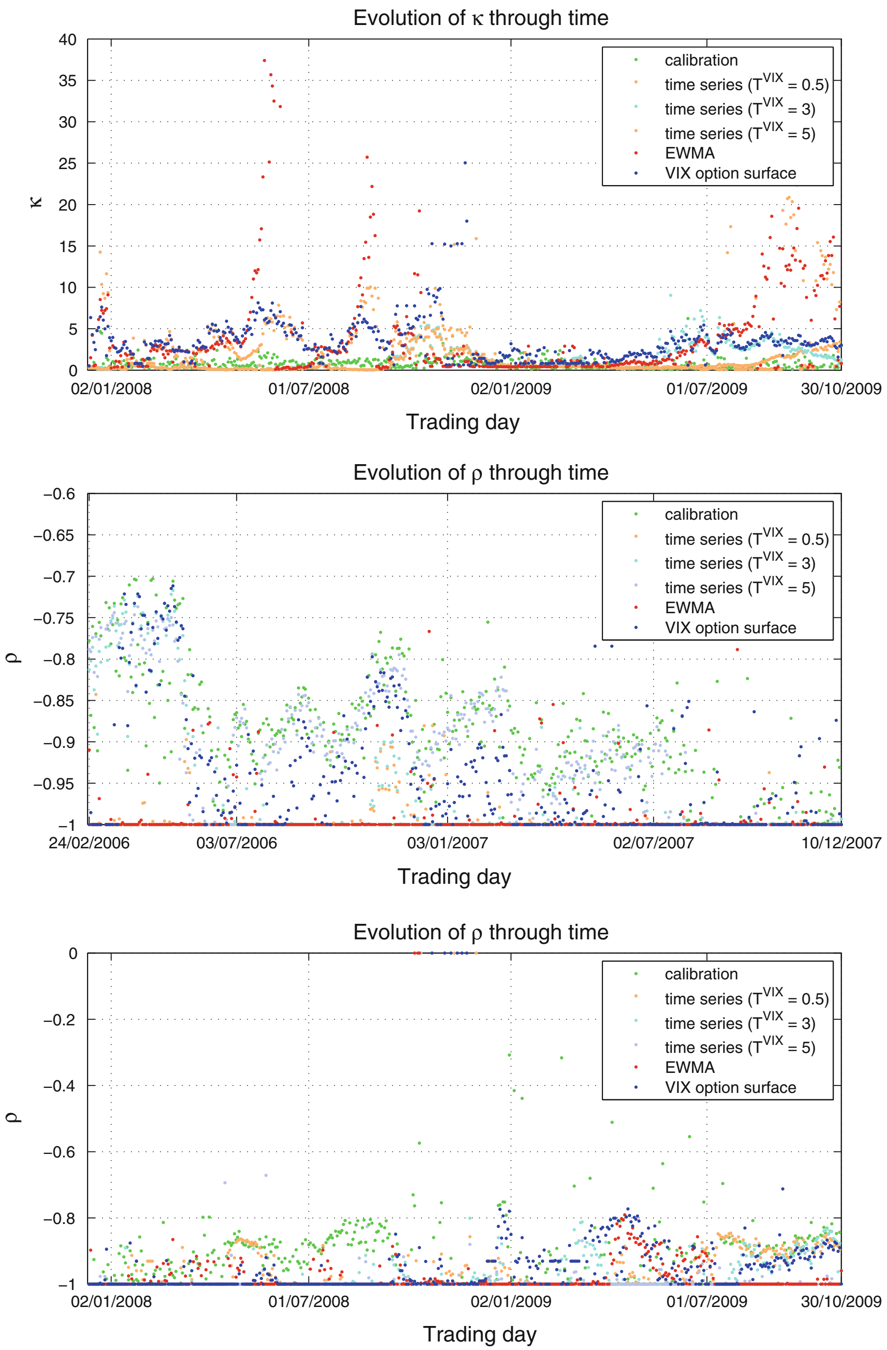

Fig. 3 continued

Table 1 Call and put lookback prices

Calibration	31/10/2006		18/09/2008		19/10/2009	
	LC	LP	LC	LP	LC	LP
calibration	247.1450	144.2566	320.0110	277.1198	296.2825	287.0761
MW - $T^{\text{VIX}} = 0.5$	243.3814	134.8871	291.4523	268.1978	290.9191	285.8290
MW - $T^{\text{VIX}} = 3$	243.3555	135.1626	290.4874	245.2662	294.0654	285.7428
MW - $T^{\text{VIX}} = 5$	246.2371	141.4669	292.0366	243.4269	269.9777	283.7615
EWMA	212.9100	107.4609	313.6819	280.8053	268.9217	283.8554
VIX	246.4530	139.3560	304.7779	281.0082	286.9627	285.9993

The payoff of lookback call and put options corresponds to the call and the put vanilla payoff where the strike is taken equal to the lowest and highest levels the stock has reached during the option lifetime, respectively. The initial price of the lookback call and put options is given by

$$LC = \exp(-rT)\mathbb{E}_{\mathbb{Q}}\left[(S_T - m_T^S)^+\right] \quad \text{and} \quad LP = \exp(-rT)\mathbb{E}_{\mathbb{Q}}\left[(M_T^S - S_T)^+\right],$$

respectively where m_t^X and M_t^X denote the minimum and maximum processes of the process $X = \{X_t, 0 \leq t \leq T\}$, respectively:

$$m_t^X = \inf\{X_s, 0 \leq s \leq t\} \quad \text{and} \quad M_t^X = \sup\{X_s, 0 \leq s \leq t\}.$$

The payoff of a one-touch barrier option depends on whether the underlying stock price reaches the barrier H during the lifetime of the option. We illustrate the findings by looking at the down-and-in put and the up-and-in call price:

$$\text{DIBP} = \exp(-rT)\mathbb{E}_{\mathbb{Q}}\left[(K - S_T)^+ \mathbf{1}\left(m_T^S \leq H\right)\right]$$

and

$$\text{UIBC} = \exp(-rT)\mathbb{E}_{\mathbb{Q}}\left[(S_T - K)^+ \mathbf{1}\left(M_T^S \geq H\right)\right].$$

The payoff of a cliquet option depends on the sum of the stock returns over a series of consecutive time periods $\left[t_i, t_{i+1}\right]$; each local performance being first floored and/or capped. Moreover the final sum is usually further floored and/or capped to guarantee a minimum and/or maximum overall payoff such that cliquet options protect investors against downside risk while allowing them for significant upside potential:

$$\text{Cliquet} = \exp(-rT)\mathbb{E}_{\mathbb{Q}}\left[\min\left(\text{cap}^G, \max\left(\text{floor}^G, \sum_{i=1}^{N}\min\left(\text{cap}^L, \max\left(\text{floor}^L, \frac{S_{t_i}-S_{t_{i-1}}}{S_{t_{i-1}}}\right)\right)\right)\right)\right].$$

As it can be seen from Table 1, Figs. 4 and 5, the price of the different exotic options under the Heston model turns out to be sensitive to the calibration method. In

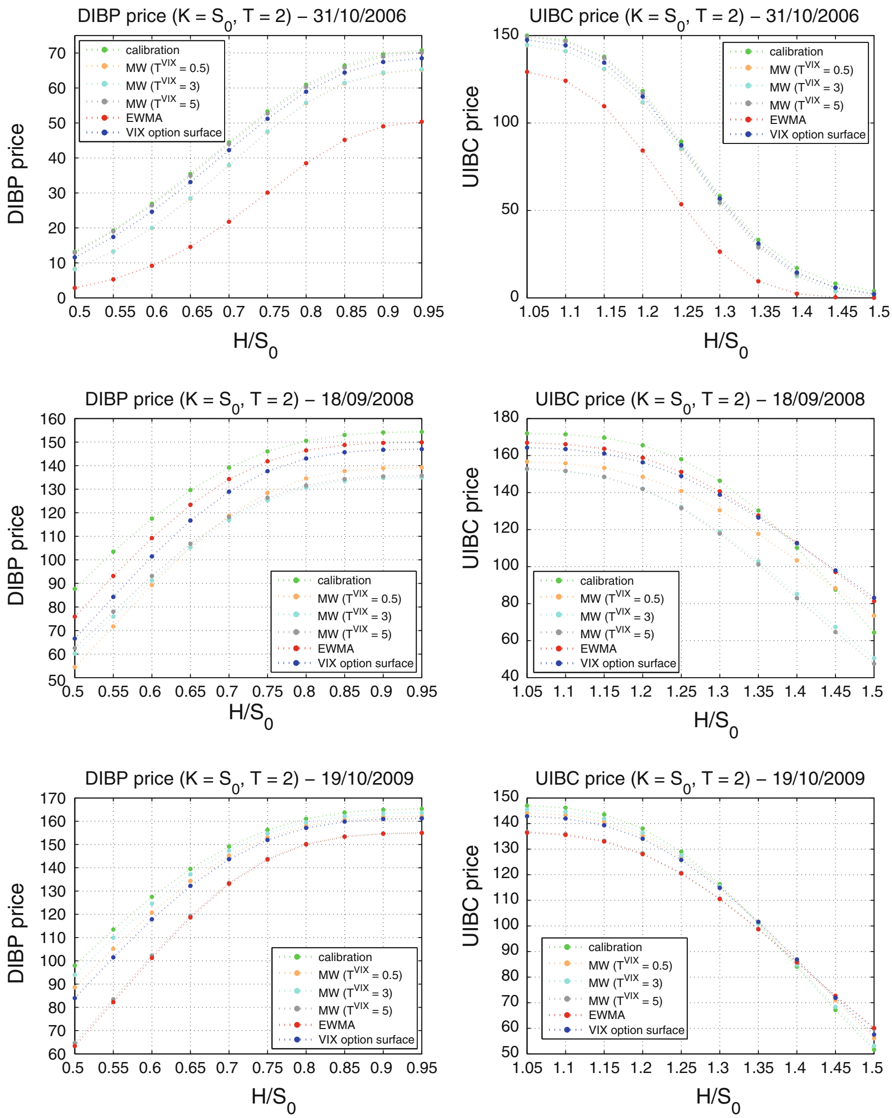

Fig. 4 Down-and-in put (*left*) and up-and-in call (*right*) option prices

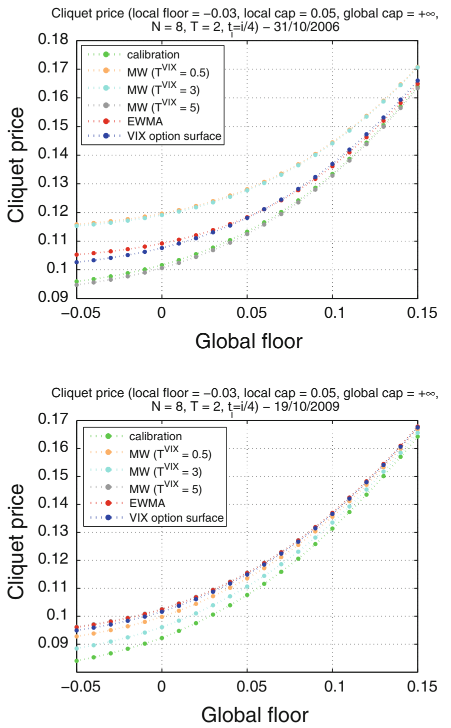

Fig. 5 Cliquet option prices

particular, the reduced calibration which leads to the lowest objective function does not necessarily lead to exotic option prices which are the closest to the prices obtained by considering the standard optimal parameter set (for more details see Guillaume and Schoutens 2010). It is also clear that the calibration risk depends on the contract specifications such as the barrier level or the cap and floor levels. In particular, the calibration risk of cliquet options decreases with the global floor level and it turns out to be significant for the widely traded capital protected cliquets (i.e. cliquets with a global floor equal to zero).[2]

Table 2 shows an estimate of the impact of the global calibration risk (i.e. arising from the choice of both the calibration methodology and the objective function) on short and long term exotic option prices. We measure this risk by

$$\frac{\max_{i,j} \hat{P}_{i,j} - \min_{i,j} \hat{P}_{i,j}}{\sum_{i=1}^{M} \sum_{j=1}^{N} \frac{\hat{P}_{i,j}}{NM}} \tag{7}$$

[2] For more exotic prices, see Guillaume and Schoutens (2010).

Table 2 Global calibration risk (7) for different exotic options maturities

Quoting date/T	0.25	0.5	1	2
Lookback call				
31/10/2006	0.099454	0.114375	0.138969	0.191766
18/09/2008	0.202491	0.176726	0.210603	0.262799
11/12/2008	0.201495	0.251864	0.307942	0.341169
19/10/2009	0.082971	0.102074	0.126567	0.166206
DIBP ($K = S_0$, $H = 0.75\quad S_0$)				
31/10/2006	1.960904	0.874819	0.543727	0.629034
18/09/2008	0.656278	0.397121	0.366175	0.181164
11/12/2008	0.362686	0.351415	0.383152	0.416484
19/10/2009	0.519131	0.246416	0.176765	0.191923
UIBC ($K = S_0$, $H = 1.25\quad S_0$)				
31/10/2006	9.118238	3.831980	2.355436	0.823998
18/09/2008	1.753460	0.939958	0.427795	0.352885
11/12/2008	0.313618	0.330268	0.371631	0.390901
19/10/2009	3.645171	0.817110	0.212333	0.241628
Cliquet ($\text{floor}^L = -0.03$, $\text{cap}^L = 0.05$, $\text{floor}^G = 0$, $\text{cap}^G = +\infty$, $N = 6$, $t_i = T/6$)				
31/10/2006	0.084757	0.123370	0.161060	0.146625
18/09/2008	0.136895	0.194299	0.276360	0.226304
11/12/2008	0.109767	0.177890	0.253856	0.269136
19/10/2009	0.188438	0.329906	0.324200	0.175856

where $\hat{P}$ is the Monte Carlo price of the exotic and where N denotes the number of calibration settings and M the number of objective functions. The number of prices taken into account in the global calibration risk estimate (7) amounts thus to 18.

The impact of the calibration risk increases with the exotic option maturity for lookback and cliquet options. On the other hand, for barrier options, it decreases with the time to expiration, except during highly volatile periods (i.e. for the 11th of December 2008) where the impact of calibration risk is roughly the same whatever the option maturity.

4.1 Calibration risk: the choice of the calibration procedure

In order to quantify the calibration risk, it is interesting to have a look at the evolution of the exotic prices through time under the different calibration settings (see Figs. 6, 7 & 8). We computed the price of different exotic options every 2 weeks by the Monte Carlo simulation for the six sets of model parameters, the number of computation days amounting then to 96. We first notice that the calibration risk is higher during the panic wave period, especially for the lookback and knock-in barrier options. Moreover, except during the investors' fear wave, the lookback call and put prices are not significantly different from one calibration procedure to the other; indicating that the

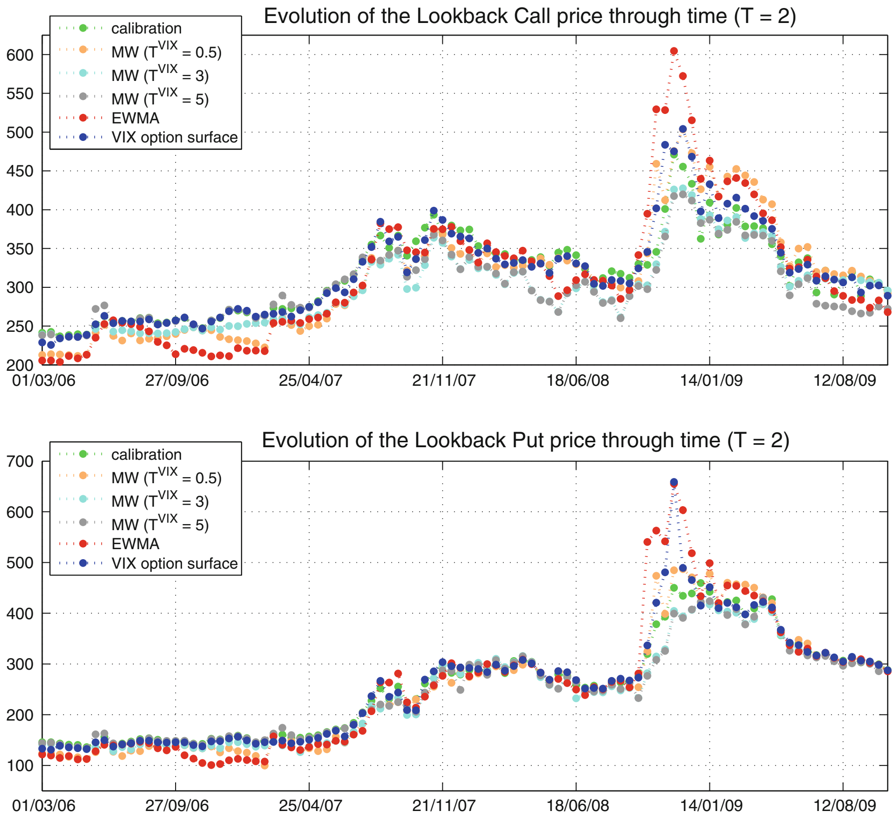

Fig. 6 Evolution of the lookback call (*upper*) and put prices (*lower*) through time for the different calibration procedures

calibration risk is not predominant for this kind of exotic options. The biggest difference between the exotic option prices under the different settings corresponds to the cliquet options, followed by the barrier options.

4.2 Calibration risk: the choice of the objective function

Figures 9, 10 and 11 show the evolution of the different exotic prices through time obtained with the different objective functions under the full and market implied calibration procedures. It is clear that the calibration risk turns out to be more significant during the panic wave period. Moreover, the difference of the exotic prices obtained by considering different functionals is more marked under the standard calibration setting, which might be explained by the fact that the number of degrees of freedom is higher for this calibration procedure. Indeed, for the reduced settings, the parameters v_0 and η are directly inferred from volatility market data and are thus constant with respect to the choice of the objective function.

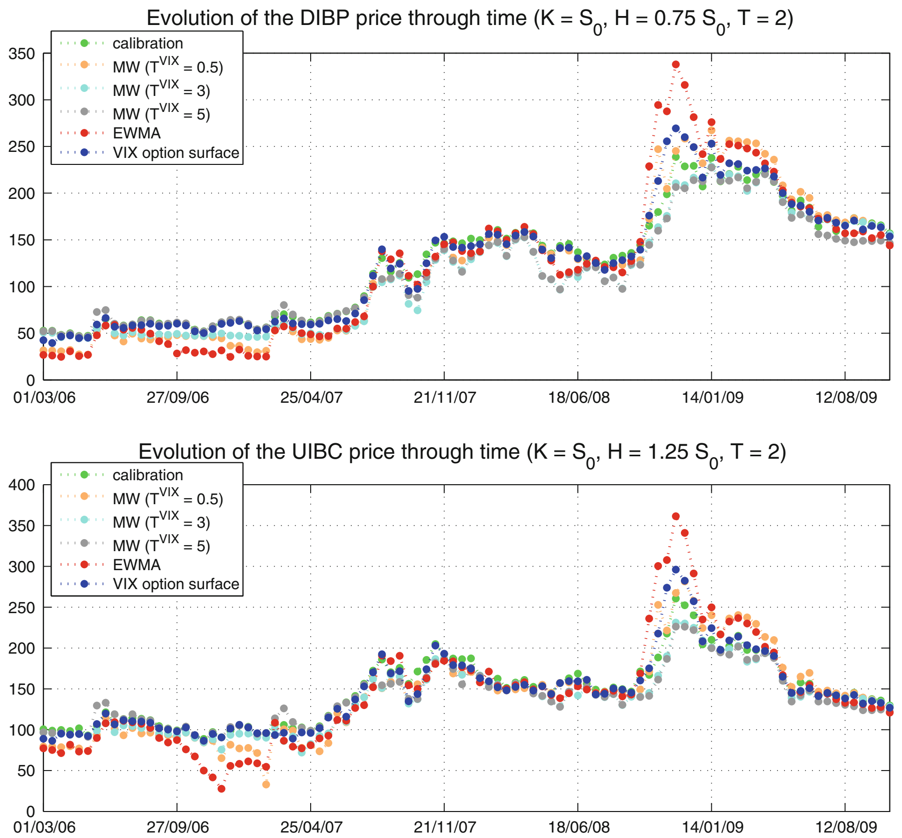

Fig. 7 Evolution of the down-and-in put (*upper*) and up-and-in call prices (*lower*) through time for the different calibration procedures

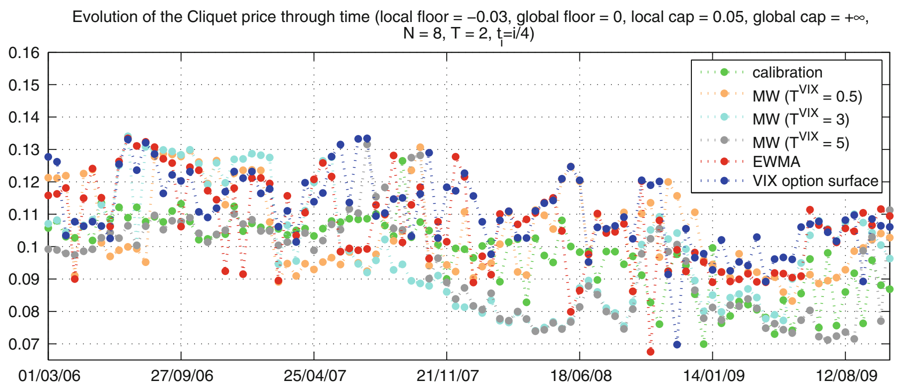

Fig. 8 Evolution of the cliquet price through time for the different calibration procedures

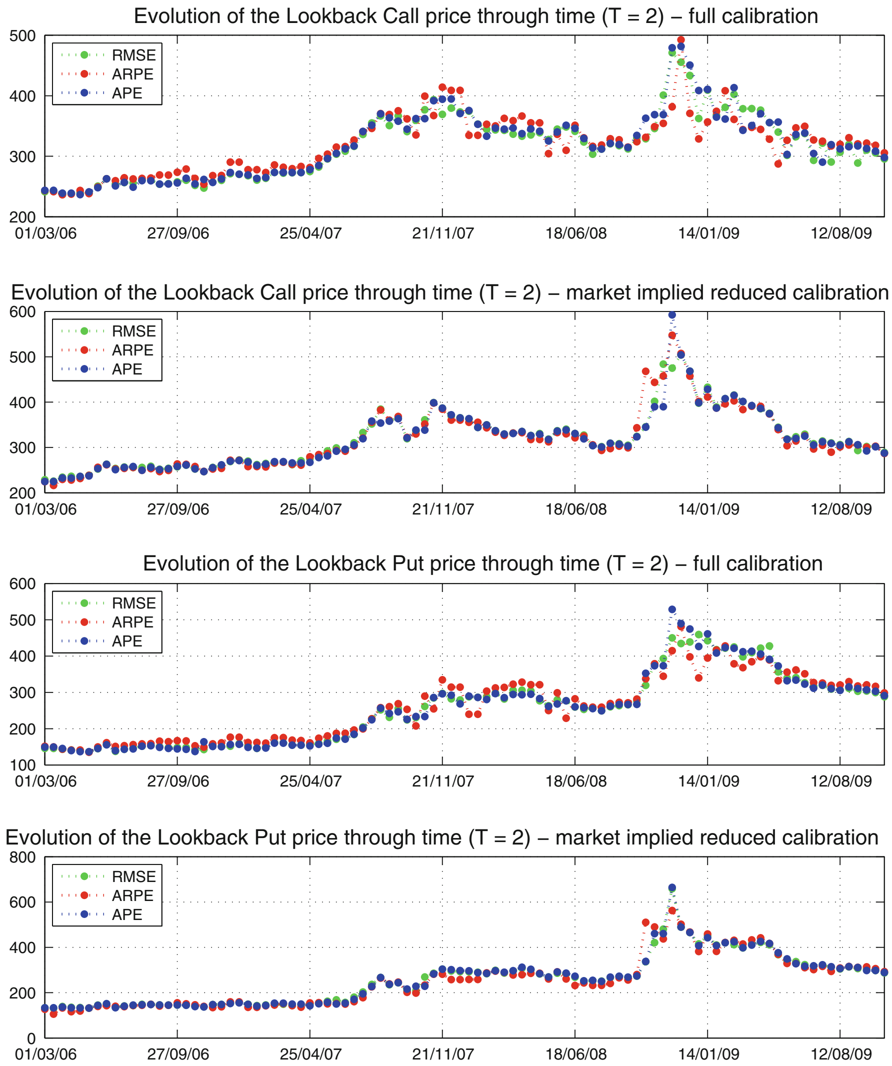

Fig. 9 Evolution of the lookback call (*upper*) and put prices (*lower*) through time for the different objective functions

5 Delta hedging

This section features the impact of the choice of the calibration methodology and of the objective function on the P&L of the delta hedging strategy. More particularly, we hedge a short position in 3 months call options with different strike prices by adjusting daily our position in the underlying stock. We consider four typical trading days and for each of them we draw 100,000 Monte Carlo sample paths of the underlying

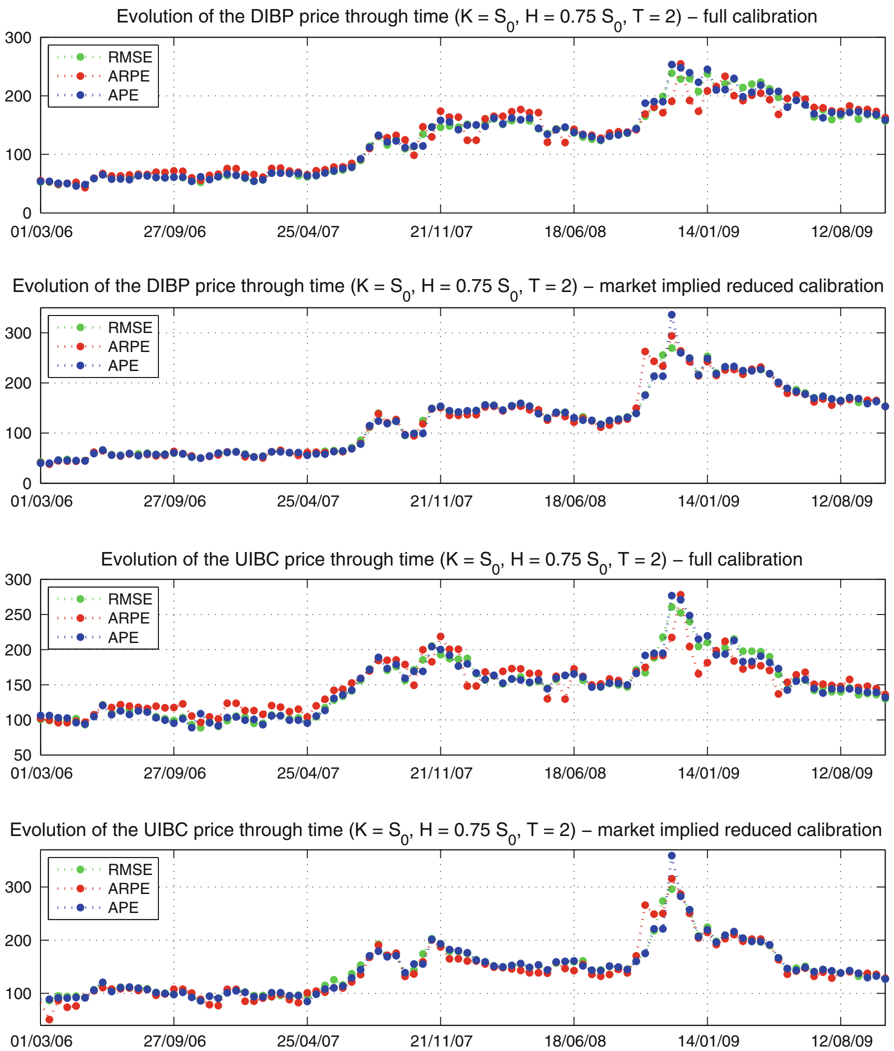

Fig. 10 Evolution down-and-in put (*upper*) and up-and-in call (*lower*) prices through time for the different objective functions

process for the different optimal parameter sets. We assess the impact of the calibration risk on the delta hedge P&L both qualitatively and quantitatively by looking at the empirical P&L distribution (see Figs. 12 & 13) and by computing the Sharpe ratio Sharpe (1994) and the Gain loss ratio Bernardo and Ledoit (2000). Table 3 shows an estimate of the relative variability of these two standard risk measures RM with respect to the calibration methodology and the specification of the functional. This estimate is computed by

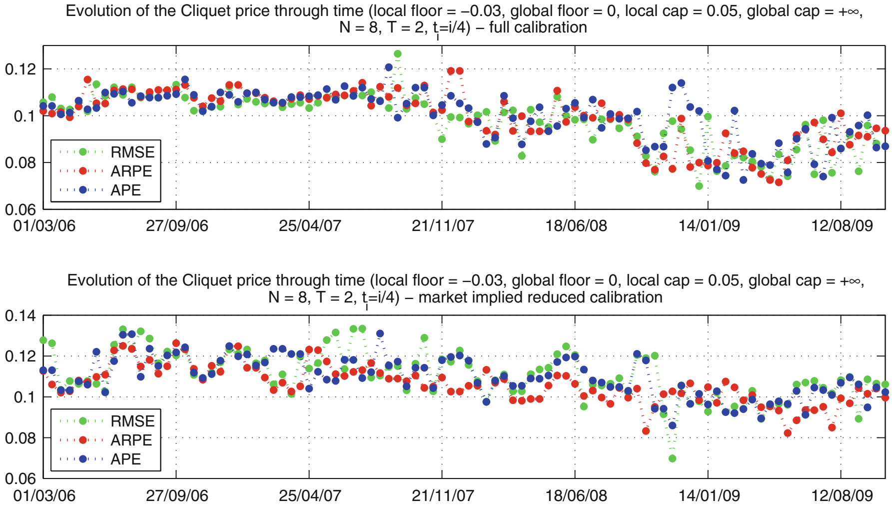

Fig. 11 Evolution of the cliquet price through time for the different objective functions

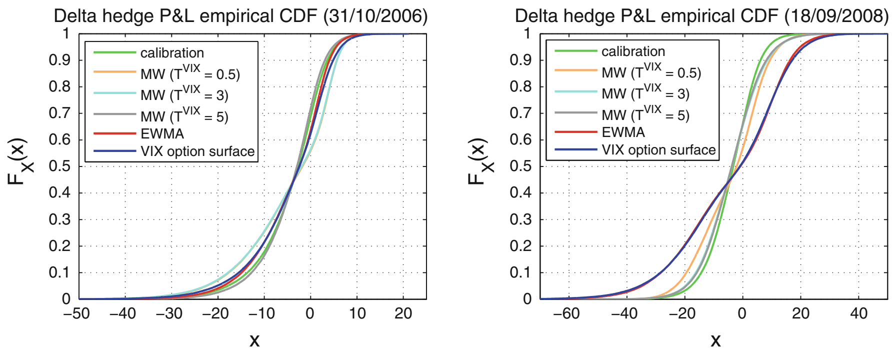

Fig. 12 Delta hedge P&L cumulative distribution function for the different calibration procedures

$$\frac{\max_{i,j} \hat{\mathrm{RM}}_{i,j} - \min_{i,j} \hat{\mathrm{RM}}_{i,j}}{\sum_{i=1}^{M} \sum_{j=1}^{N} \frac{\hat{\mathrm{RM}}_{i,j}}{NM}} \tag{8}$$

where $\hat{\mathrm{RM}}$ is the Monte Carlo Sharpe ratio or Gain loss ratio of the delta hedge profit and loss.

Both the choice of the objective function and the choice of the calibration methodology are reflected in the P&L of the delta hedging strategy, whatever the volatility regime. Nevertheless, the impact is more marked during the low volatility period (i.e. for the 31st of October 2006) and for deep in-the-money call options.

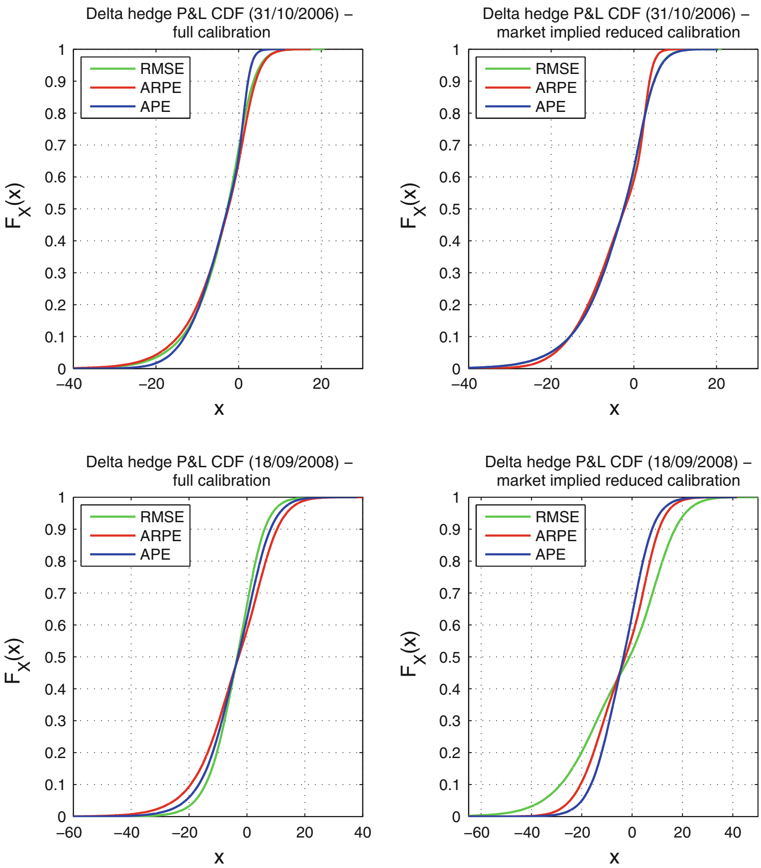

Fig. 13 Delta hedge P&L cumulative distribution function for the different objective functions

6 Conclusion

We propose a market implied estimate of the long run variance of stochastic volatility models such as the Heston model which is directly obtained from the put-call parity of long maturity vanilla options on the VIX index. We show that this estimate follows the same trend as the long run variance parameter η obtained by the standard calibration of the Heston model on the option price surface, although it is typically lower. We also perform a detailed study of the calibration performance of the Heston model, considering either the common calibration on the whole parameter set or a reduced calibration on the set $\{\kappa, \lambda, \rho\}$ where the parameter v_0 is inferred from the spot VIX and the parameter η either from the VIX time series or from the VIX option surface. The optimal reduced calibration procedure is obtained by considering the market implied estimate of the long run variance, since then the different objective functions turn out to be pretty close to the optimal ones, except during the period characterized by huge investors fear.

Table 3 The impact of calibration (8) on the delta hedge P&L

Quoting date/K	0.9	0.95	1	1.05	1.1
Sharpe ratio					
31/10/2006	−0.631981	−0.586003	−0.516548	−0.330157	−0.329874
18/09/2008	−0.540852	−0.521437	−0.522575	−0.531354	−0.517259
11/12/2008	−0.793438	−0.729650	−0.656095	−0.709454	−0.750170
19/10/2009	−0.466888	−0.434629	−0.401373	−0.342886	−0.280010
Gain loss ratio					
31/10/2006	3.193305	2.696743	1.012186	0.200702	0.082658
18/09/2008	1.513670	1.003081	0.673916	0.465949	0.309048
11/12/2008	0.275325	0.206421	0.158969	0.143276	0.125561
19/10/2009	0.821799	0.422809	0.242792	0.141269	0.074425

Although the market implied reduced calibration leads to a fit of the option surface of a similar quality than the standard calibration, the price of a wide range of exotic options (one touch barrier, lookback and cliquet options) turns out to be significantly different under the two calibration settings, the calibration risk being predominant for the cliquet and barrier options. Moreover, the delta hedge P&L turns out to depend on both the calibration methodology and the objective function. This might be explained by the fact that the two calibration procedures, as well as, to a larger extent, the different specifications of the objective function lead to significantly different optimal parameter sets. Hence, even within a particular model, model risk and calibration risk are present. We are thus faced with choosing a route to price exotic products out of the liquid vanilla option prices.

References

Albrecher, H., Mayer, P., Schoutens, W., & Tistaert, J. (2007). The little Heston trap. *Wilmott Magazine*, Jannuary 2007, 83–92.

Basel Committee on Banking Supervision (2009). Revisions to the Basel II market risk framework. *Bank for International settlements*.

Bernardo, A. E., & Ledoit, O. (2000). Gain, loss, and asset pricing. *Journal of Political Economy, 108*(1), 144–172.

Carr, P., & Madan, D. B. (1998). Option valuation using the fast Fourier transform. *Journal of Computational Finance, 2*, 61–73.

Cherny, A., & Madan, D. B. (2009). New measures of performance evaluation. *Review of Financial Studies, 22*, 2571–2606.

Cont, R. (2006). Model uncertainty and its impact on the pricing of derivative instruments. *Mathematical Finance, 16*(3), 519–547.

Detlefsen, K., & Härdle, W. K. (2007). Calibration risk for exotic options. *The Journal of Derivatives, 14*(4), 47–63.

Guillaume, F., & Schoutens, W. (2010). Use a reduced Heston or reduce the use of Heston? *Wilmott Journal, 2*(4), 171–192.

Guillaume, F., & Schoutens, W. (2010). Model and calibration risk for the Heston model. *Eurandom report*.

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

Hull, J. (2006). *Options, futures, and other derivatives* (6th ed.). New Jersey: Pearson Prentice Hall.

Schoutens, W., Simons, E., & Tistaert, J. (2004). A Perfect calibration! Now what?, *Wilmott Magazine*, March 2004, pp 66–78.

Sharpe, W. F. (1994). The Sharpe ratio. *The Journal of Portfolio Management, 21*(1), 49–58.

Stulz, R. M. (2004). Should we fear derivatives?. *The Journal of Economic Perspectives, 18*(3), 173–192.

Whaley, R. E. (2009). Understanding the VIX. *The Journal of Portfolio Management, 35*, 98–105.

Wilmott, P. (2006). *Paul Wilmott on quantitative finance*. London: Wiley.