

EURANDOM PREPRINT SERIES
2010-026

**Model and calibration risks
for the Heston model**

F. Guillaume, W. Schoutens
ISSN 1389-2355

Model and calibration risks for the Heston model

Florence Guillaume*

Wim Schoutens[†]

June 10, 2010

Abstract

Parameters of equity pricing models, such as the Heston's stochastic volatility model, have to be calibrated every day to new market data of European vanilla options by minimizing a particular functional. Hence, the optimal parameter set might turn out to vary significantly on a daily basis, depending on the quality of the initial guess and therefore on the local minima which is reached by the local optimizer method. However, thanks to the emergence of market data for volatility derivatives, 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 v_0 of the Heston model can be inferred beforehand from the spot value of the volatility index whereas the long run variance η can be estimated either from the time series of the volatility index or from the VIX option surface.

This paper provides a market-implied estimate of η which is inferred from the Put-Call parity for long maturity options on the VIX. We then compare the such obtained mark-to-market estimate with the η parameter obtained from the calibration and with the estimate inferred from the VIX time series by using either a moving window or the exponentially weighted moving average technique. Moreover, this paper features a detailed calibration performance study of the Heston model for the two calibration procedures, i.e. the standard calibration on the whole parameter set and the different reduced calibrations on the parameter set $\{\kappa, \lambda, \rho\}$. We also investigate the calibration risk which arises by considering different objective functions and/or different calibration methodologies and price a wide range of exotic options for the different calibration settings. For the numerical study, we consider daily S&P500 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.

*T.U.Eindhoven, Department of Mathematics, Eurandom, P.O.Box 513 5600 MB Eindhoven, the Netherlands.
E-mail: guillaume@eurandom.tue.nl

[†]K.U.Leuven, Department of Mathematics, Celestijnenlaan 200 B, B-3001 Leuven, Belgium. E-mail: Wim@Schoutens.be

1 Introduction

This paper features the calibration performance of the Heston model for two calibration procedures. The first one consists of the common calibration on the whole parameter set $\{v_0, \kappa, \eta, \lambda, \rho\}$ whereas the second one consists of a calibration on the reduced parameter set $\{\kappa, \lambda, \rho\}$ where the spot variance v_0 is inferred beforehand from the spot VIX and the long run variance η is determined beforehand on the basis of either VIX time series or the VIX option surface. In [5], Guillaume and Schoutens have investigated the fit of the implied volatility surface under the Heston model for a period extending from the 24th of February 2006 until the 30th of September 2008 for both the standard calibration and a particular reduced calibration where η is inferred from a moving window of the VIX time series. In particular they have shown that, for the period considered, both calibration methods lead to a fit of the option surface of a similar quality for a well chosen length of the time series window. In this paper we perform a similar study but for an extended time period (including the whole credit crisis) and for additional reduced calibration procedures as well. Indeed, in [5], it has been shown that the choice of the length of the time series window to fix the long run variance strongly depends on the market volatility regime. Hence it might be interesting to find an estimate of η which is independent of any time series window. The first matching method we propose to approximate the long run variance is to use the exponentially weighted moving average (or EWMA) technique. However, both the EWMA and the moving window methodologies rest on time series of the VIX index and not on current market data which reflect the future expectations of the investors. Hence, we will propose a market implied estimate of η inferred from the VIX option surface, or more precisely from VIX options with a long maturity.

The Chicago Board Options Exchange launched trading of VIX option contracts on the 24th of February 2006. After nine months, the VIX option trading volume almost reaches 4.5 million contracts, making it the most successful new product launched in CBOE history. Figure 1 shows the monthly volume of VIX options which had experienced a substantial increase until the mid of 2007. Since then, the monthly volume of VIX option contracts has oscillated between one and three millions, with an average amounting to roughly two millions. Hence, given the substantial liquidity of VIX options, it is possible to infer an accurate estimate of the long run volatility from their market quotes.

2 Calibration of the Heston model

2.1 The Heston's stochastic volatility model

In [6], Heston proposed a model which extends the Black-Scholes model by making the volatility parameter σ stochastic. In particular, the squared volatility is modeled by a CIR process, which is coherent with the positivity and mean-reverting characteristics of the empirical volatility (see for instance [10]). The stock price process follows the well-known Black-Scholes stochastic differential equation:

$$\frac{dS_t}{S_t} = (r - q)dt + \sqrt{v_t}dW_t, \quad S_0 \geq 0 \quad (2.1)$$

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

$$dv_t = \kappa(\eta - v_t)dt + \lambda\sqrt{v_t}d\tilde{W}_t, \quad v_0 = \sigma_0^2 \geq 0, \quad (2.2)$$

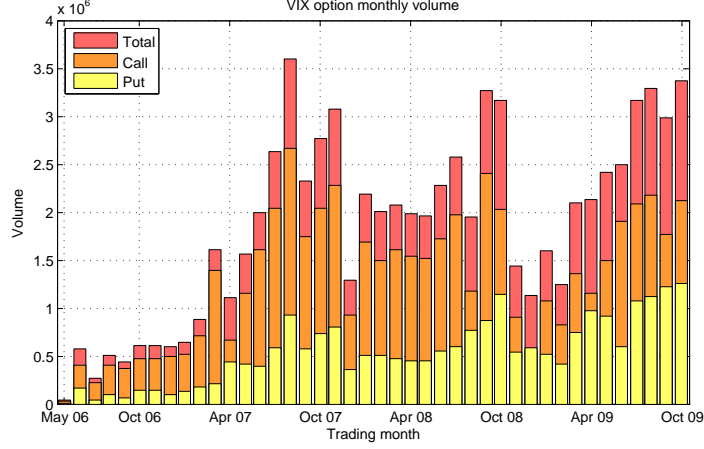


Figure 1: Monthly volume of VIX Options.

where $W = \{W_t, t \geq 0\}$ and $\tilde{W} = \{\tilde{W}_t, t \geq 0\}$ are two correlated standard Brownian motions such that $\text{Cov}(dW_t, d\tilde{W}_t) = \rho dt$ and where 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 is always positive and can not reach zero if $2\kappa\eta > \lambda^2$. Moreover, the deterministic part of the CIR process is asymptotically stable if $\kappa > 0$ and tends towards the equilibrium point $v_t = \eta$. The model parameters can be determined either by matching data or by calibration. In practice, calibrated parameters turn out to be unstable and often unreasonable (see [11]). Hence, we will propose a method which consists of a mixture of matching data and calibration in order to increase the stability of the model parameters and decrease the calibration computation time.

2.2 Pricing of vanilla Options using Characteristic Functions

The vanilla option prices are computed by using the Carr-Madan formula (see [2]):

$$C(K, T) = \frac{\exp(-\alpha \log(K))}{\pi} \int_0^{+\infty} \exp(-iv \log(K)) \varrho(v) dv,$$

where

$$\begin{aligned} \varrho(v) &= \frac{\exp(-rT) \mathbb{E}[\exp(i(v - (\alpha + 1)i) \log(S_T))]}{\alpha^2 + \alpha - v^2 + i(2\alpha + 1)v} \\ &= \frac{\exp(-rT) \phi(v - (\alpha + 1)i, T)}{\alpha^2 + \alpha - v^2 + i(2\alpha + 1)v}, \end{aligned}$$

where α is a positive constant such that the $(1 + \alpha)$ th moment of the stock price process exists and ϕ is the risk neutral (i.e. under Q) characteristic function of the log stock price process at maturity T :

$$\phi(u; T) = \mathbb{E}_Q[\exp(iu \log(S_T)) | S_0, \sigma_0^2].$$

Under the Heston model, the characteristic function is given by (see [1] or [9]):

$$\begin{aligned}\phi(u; T) = & \exp(iu(\log S_0 + (r - q)T)) \exp\left(\frac{\eta\kappa}{\lambda^2} \left((\kappa - \rho\lambda iu - d)T - 2\log\left(\frac{1 - ge^{-dt}}{1 - g}\right)\right)\right) \\ & \exp\left(\frac{\sigma_0^2}{\lambda^2} (\kappa - \rho\lambda iu - d) \frac{1 - e^{-dt}}{1 - g}\right),\end{aligned}$$

where

$$\begin{aligned}d &= \sqrt{(\rho\lambda iu - \kappa)^2 + \lambda^2(iu + u^2)} \\ g &= \frac{\kappa - \rho\lambda iu - d}{\kappa - \rho\lambda iu + d}.\end{aligned}$$

Using the Carr-Madan formula and combining it with numerical techniques that evaluate the integral involved in an efficient way (based on the Fast Fourier Transform (FFT)) lead to an extremely fast pricing algorithm of the entire option surface (see [2]). Indeed the algorithm generates in one run all prices for a fine strike grid and all the given maturities.

2.3 Calibration sets

The goodness of fit of the Heston model is determined by the root mean square error objective function:

$$\text{RMSE} = \sqrt{\sum_{j=1}^N \frac{(P_j - \hat{P}_j)^2}{N}} \quad (2.3)$$

where N stands for the number of quoted options and P and \hat{P} denote the market and model price, respectively.

Minimizing the objective function (2.3) is clearly a nonlinear programming problem with the nonlinear constraint $2\kappa\eta > \lambda^2$. Unfortunately the root mean square error is far from being convex and it turns out that there exist usually many local minima (see for instance [8]).

In practice, the model parameters have to be recalibrated every day to new market data. Hence, the optimal parameter set might turn out to vary significantly on a day-to-day basis, depending on the quality of the initial guess and therefore on the local minima which is reached by the local optimization method. In order to avoid such a fluctuation of the model parameters, practitioners might resort to time series or market quotes to determine some of the model parameters. In particular, the parameters v_0 can be fixed beforehand from the spot value of the volatility index whereas η can be estimated either from the historical time series of the volatility index by the moving window or the exponentially weighted moving average technique or from the market quotes of VIX options.

- 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:

$$\eta^{\text{MW}} = \frac{1}{T^{\text{VIX}}} \int_{-T^{\text{VIX}}}^0 \left(\frac{\text{VIX}(t)}{100}\right)^2 dt = \text{mean}_{-T^{\text{VIX}} \leq t \leq 0} \left(\frac{\text{VIX}(t)}{100}\right)^2. \quad (2.4)$$

For the numerical study, we will consider a length of the VIX time series window equal to six months, three years or five years.

- The **exponentially weighted moving average** (EWMA) estimate

The exponentially weighted moving average estimate of the long run variance is given by

$$\eta^{\text{EWMA}} = (1 - \alpha) \sum_{i=1}^N \alpha^{N-i} \left(\frac{\text{VIX}(t_i)}{100} \right)^2 \quad (2.5)$$

where $\alpha \in (0, 1)$, $t_i = t_0 - (N - i)\Delta t$ and where $N \rightarrow \infty$ is the number of data in the time series. The most recent the VIX quote, the highest the corresponding weight. In particular, the weights α^{N-i} decrease 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 the value used by *JP Morgan* for the *RiskMetrics* database (see [7]). The MW and EWMA estimates of the long run variance are determined on the basis of time series and are therefore historical estimates.

- The **market-implied** (MI) estimate

We propose a robust way of computing a risk-neutral long run volatility estimate, referred as long VIX, or LVIX and inferred from the market price of long term European vanilla options on VIX; reflecting therefore the expectations of the investors. 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 volatility might be approximated by the at-the-money strike of long term options on VIX:

$$\eta^{\text{MI}} = \left(\frac{K^{\text{ATM}}}{100} \right)^2 = \left(\frac{\text{LVIX}(t_0)}{100} \right)^2. \quad (2.6)$$

For the numerical study, we will consider the options with the quoted maturity which is the closest to one year. The at-the-money strike is obtained by interpolation of the call-put spread. The market-implied estimate of η allows us to take into account additional market information in the calibration methodology since it is derived from the VIX option surface which is not a calibration instrument in the standard calibration procedure.

We will thus compare the calibration performance of the Heston model by using

1. a fully free parameter set $\{v_0, \kappa, \eta, \lambda, \rho\}$;
2. a reduced parameter set $\{\kappa, \lambda, \rho\}$, using the market data to fix v_0 and η :
 - v_0 is set equal to the square of the spot price of the VIX index expressed in units:

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

- η is estimated on the basis of either the empirical VIX index (by equation (2.4) or (2.5)) or the VIX option market quotes (by equation (2.6)).

For the numerical study we use S&P500 and VIX quotes. The different calibrations are performed on the whole set of quoted vanilla options (i.e. both call and put options on the whole strike and time to maturity ranges); which includes on average 914.6485 options, the minimum and maximum number of options amounting to 437 and 1567, respectively.

3 Calibration performance and Evolution of the model parameters through time

This section features the evolution of the RMSE functional and the different model parameters through time for a period extending from the 24th of February 2006 until the 30th of October 2009 for the full calibration and the different reduced calibrations considered in Section 2.3¹.

As mentioned in [5], the moving window calibration performance depends on the length of the moving 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 five years whereas, for the second period, the narrowest window (i.e. T^{VIX} equal to six months) leads to the best calibration of the option surface.

Moreover, except for a period extending from mid-October 2008 until mid-December 2008, the RMSE functional obtained by considering the market implied estimate of η is pretty close to the RMSE functional of the full calibration procedure. In particular, η^{MI} leads to a better fit of the implied volatility surface than the time series estimate 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 the past expectations of investors. We also notice that the EWMA functional is usually closer to the six months MW objective function than to the other MW objective functions.

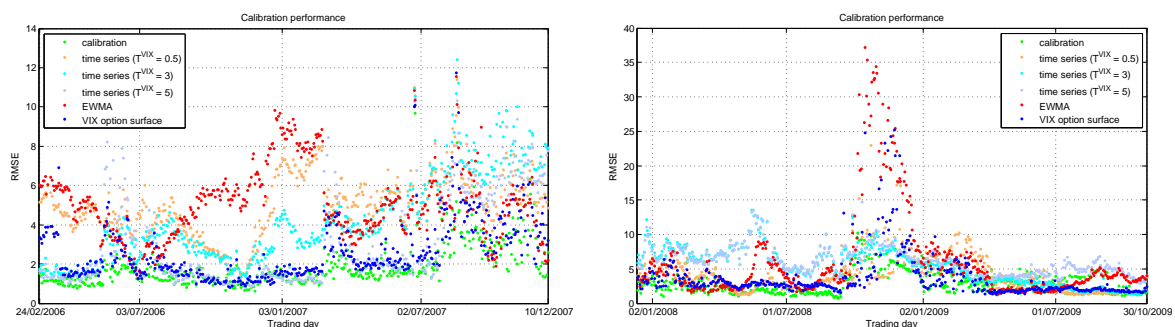


Figure 2: Evolution of the RMSE through time for the different calibration procedures.

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 .

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

From the evolution of the spot VIX (see upper part of Figure 3), it is clear that the two periods previously mentioned correspond to two different volatility regimes. More precisely, as mentioned in [5], 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 η in a smaller extent, exhibits 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 reaches some exceptional level of more than sixty-five percents.

By comparing Figure 2 and Figure 3, we note that the significant difference between the RMSE objective function of the full calibration and this of the mark-to-market 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 two months period, the square of the spot VIX appears to be close to the initial variance parameter obtained by calibrating the whole parameter set 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 mark-to-market estimate of η are typically lower than the value of η resulting from the whole calibration. However we notice some exception for the EWMA and the six months MW estimates 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 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 mark-to-market and the EWMA estimate of η roughly follow the same trend as the calibrated parameter. Moreover, the mark-to-market estimate and in a larger extend the EWMA estimate exhibit the same, although clearly smoother, trend as the spot variance (or equivalently as the square of the spot VIX). We also observe that, although smoother, the six months MW estimate follows a similar trend as the EWMA estimate, which explains the small difference of the RMSE under these two settings.

The calibrated parameter usually turns out to be closer to η^{MI} than to η^{EWMA} or η^{MW} which explains the better fit 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 five years leads to an accurate fit of the option surface since then the average of the calibrated parameter η turns out to be of the same order of magnitude than the parameter inferred from the VIX time series. On the other hand, for the second period the six months window leads to the smallest difference between the calibrated and moving window estimate of the long run variance.

From the left lower part of Figure 3, 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 the market quotes and the 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 five years estimates need more than one 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.

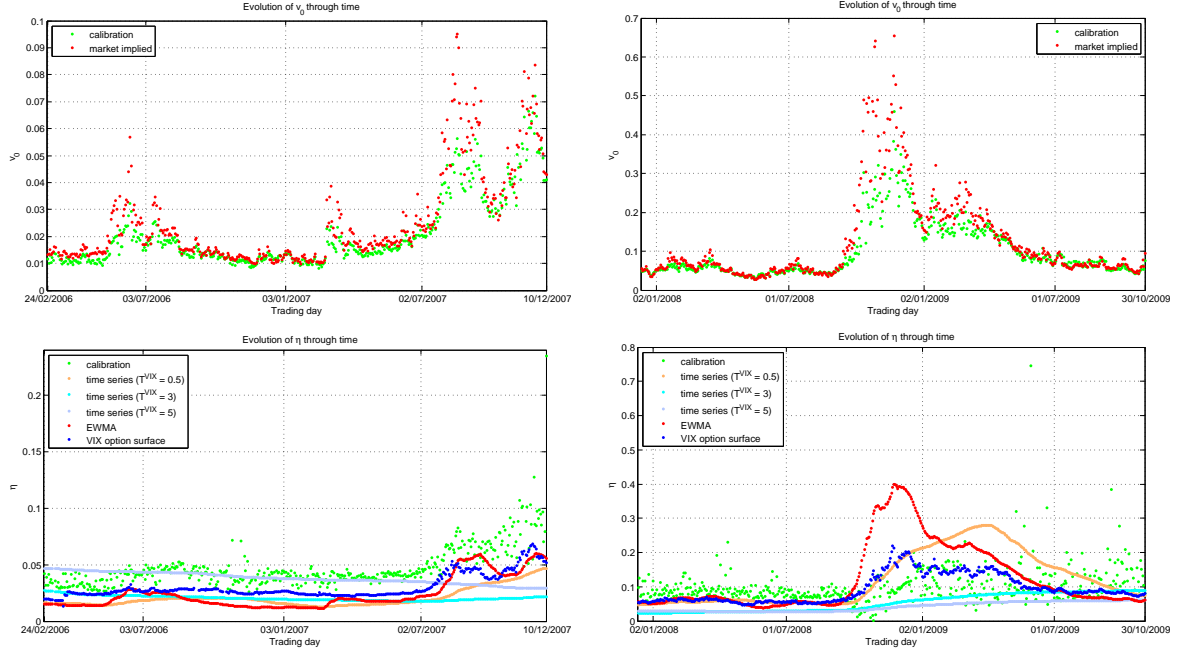


Figure 3: Evolution of the parameter v_0 (upper) and η (lower) through time for the different calibration procedures.

The matching of the spot value and the long run variance results in some adjustment of the other parameters calibrated on the S&P500 implied volatility surface. Figure 4 shows the evolution through time of these parameters, i.e. λ , κ and ρ under both the full and the reduced settings. In particular, the parameters λ and κ are typically set to a higher value than the optimal one whereas the correlation ρ is typically set to a lower value. We also observe that the parameter ρ is from time to time touching its boundary value of -1 and that the parameter κ might reach some exceptionally high level.

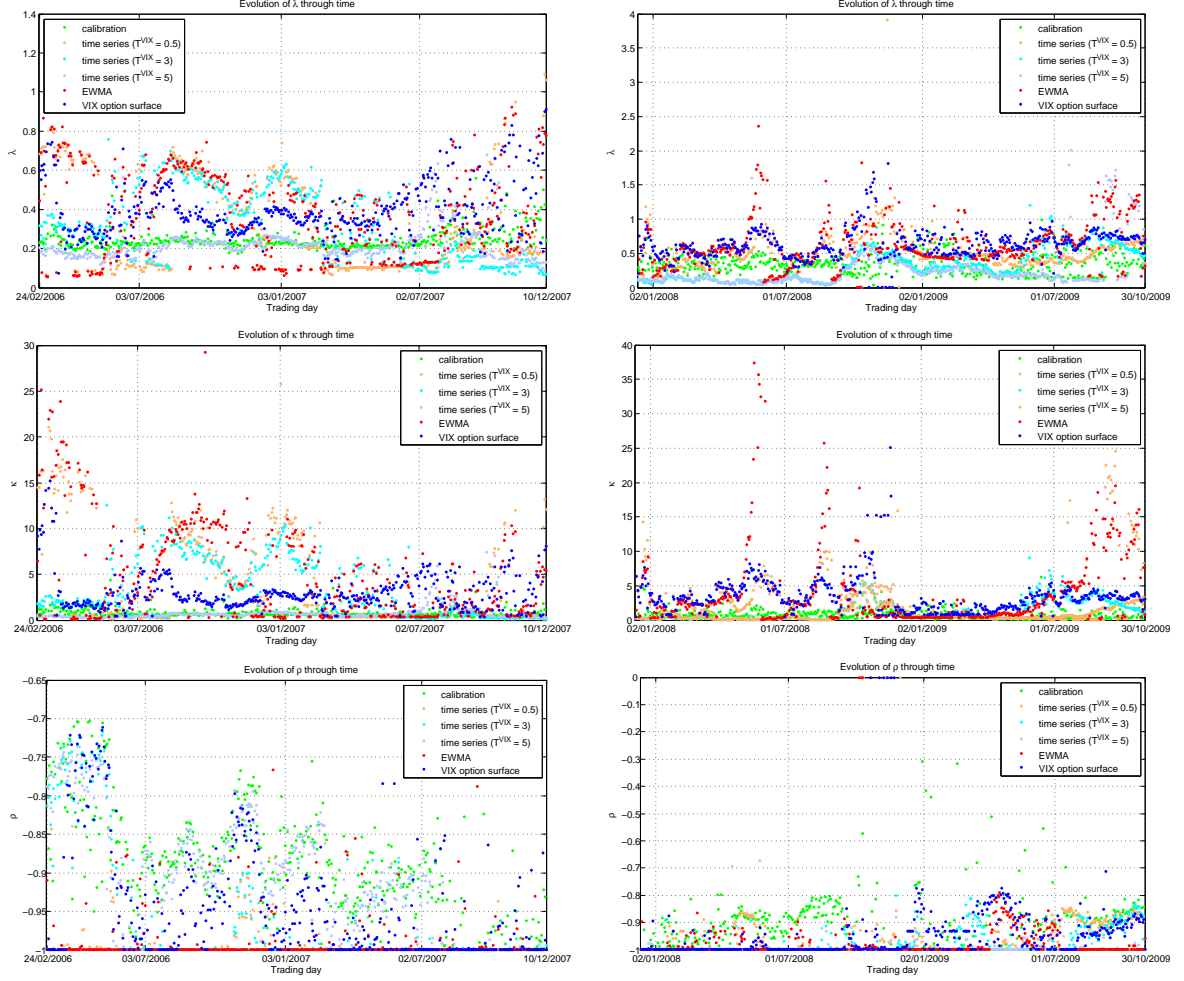


Figure 4: Evolution of the parameters λ (upper), κ (center) and ρ (lower) through time for the different calibration procedures.

The goodness of fit of the full and reduced calibrations is shown on Figure 5, Figure 6 and Figure 7 for the 31st of October 2006, the 18th of September 2008 and the 11th of December 2008, respectively (i.e. for three typical quote dates, the first one belonging to the low volatile regime period, the second one to the crisis period and the third one to the panic wave period)². For each quote date, we consider the time series estimate of η which leads to the smallest root mean square error. Except for the panic wave period, the model vanilla prices obtained with and without fixing beforehand the parameters v_0 and η barely differ from each other and are pretty close to the

²Except during the panic wave period, the biggest distance between the model and market option prices occurs for the shortest maturity and might be explained by the fact that the dividend yield we consider is fixed and not maturity dependent.

corresponding market prices.

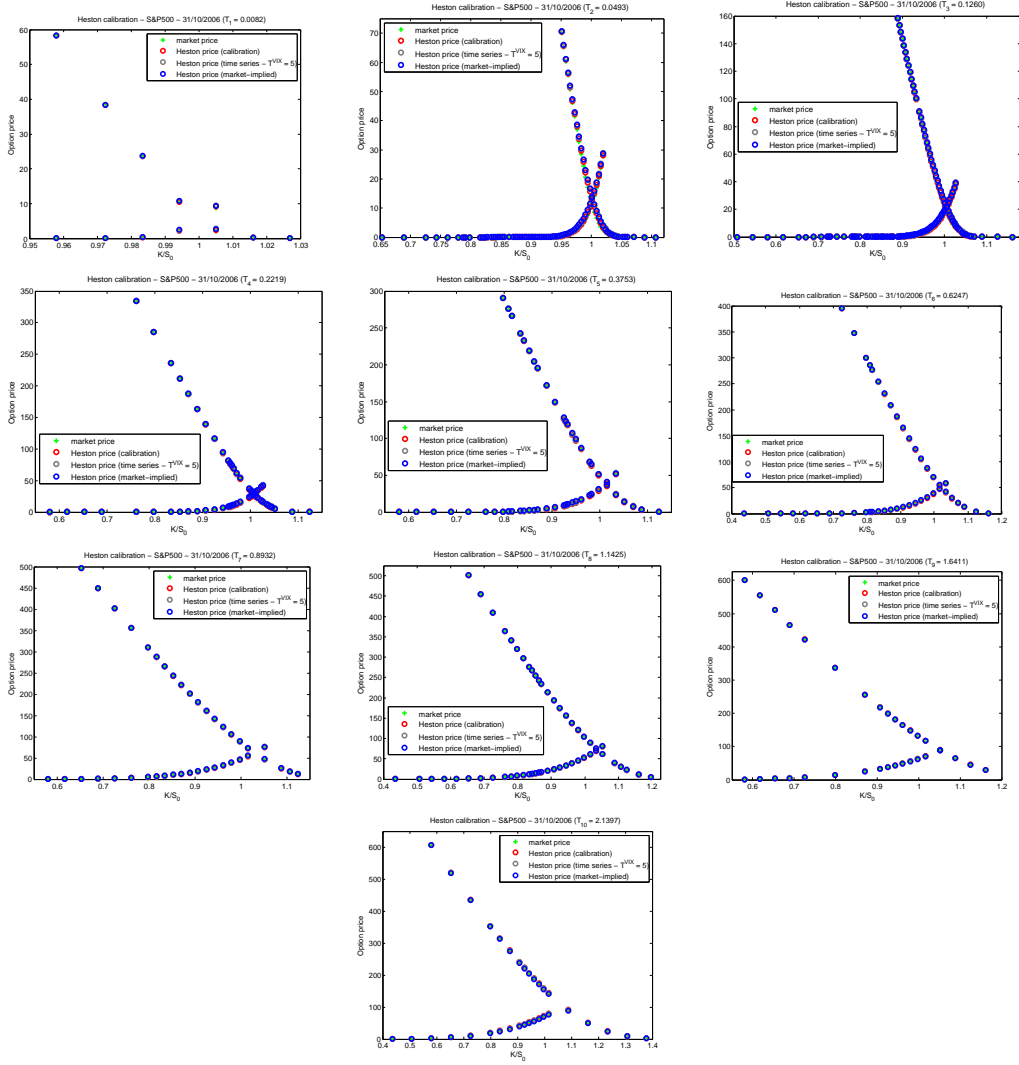


Figure 5: calibration of the Heston model - 31/10/2006

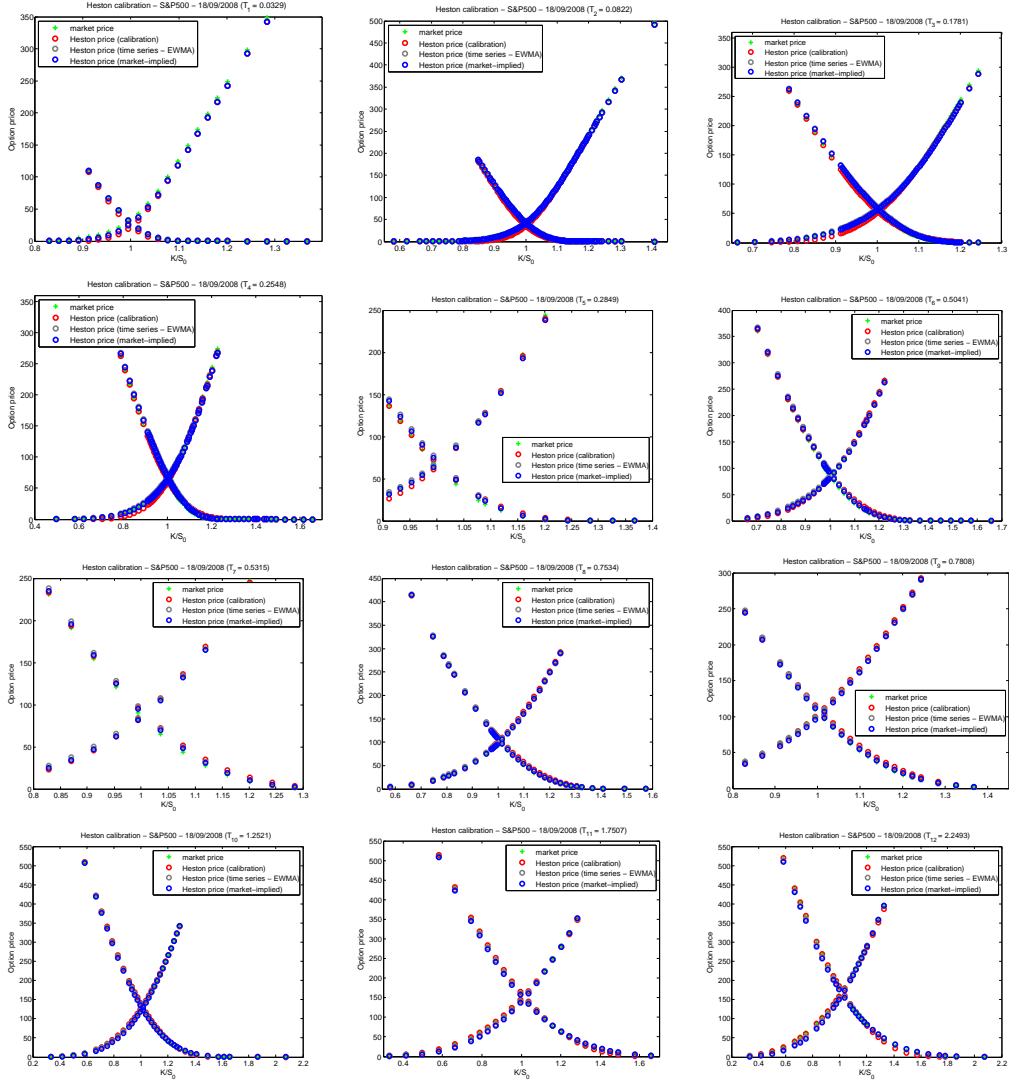


Figure 6: calibration of the Heston model - 18/09/2008

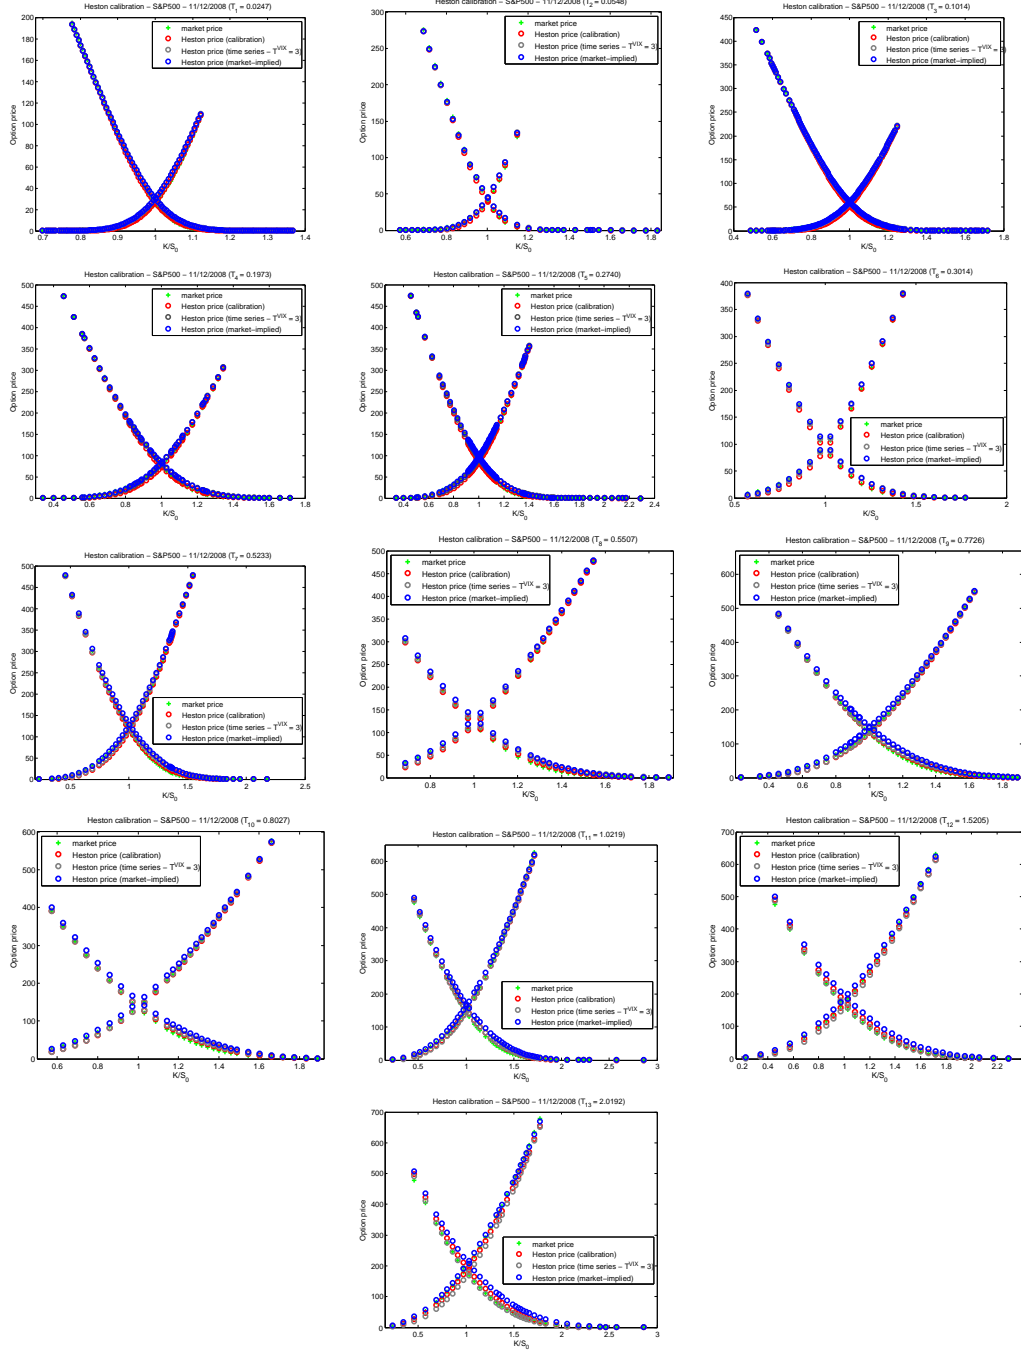


Figure 7: calibration of the Heston model - 11/12/2008

3.1 Calibration with respect to different objective functions

This subsection features a comparison of the fit of vanilla option market prices under the different calibration procedures and for different functionals: the root mean square error (RMSE), the average absolute error as a percentage of the mean price (APE) and the average relative error (ARPE) where

$$\text{APE} = \frac{1}{\text{mean}_j \hat{P}_j} \sum_{j=1}^N \frac{|P_j - \hat{P}_j|}{N}$$

and

$$\text{ARPE} = \frac{1}{N} \sum_{j=1}^N \frac{|P_j - \hat{P}_j|}{\hat{P}_j}.$$

Table 1 shows the average precision and daily variation of the different model parameters for the different calibration procedures and objective functions.

RMSE						
	calibration	MW ($T^{\text{VIX}} = 0.5$)	MW ($T^{\text{VIX}} = 3$)	MW ($T^{\text{VIX}} = 5$)	EWMA	market-implied
average precision	2.439650	4.502293	4.834824	4.654399	5.530451	3.099928
average Δv_0	0.006793	0.009460	0.009460	0.009460	0.009460	0.009460
average $\Delta \eta$	0.022588	0.000567	8.864555e-05	5.796916e-05	0.001152	0.001986
average $\Delta \lambda$	0.080715	0.082411	0.050275	0.045892	0.134313	0.098315
average $\Delta \kappa$	0.413670	1.264941	0.631513	0.570401	2.030938	1.179841
average $\Delta \rho$	0.052782	0.014224	0.012173	0.015022	0.017561	0.038741
ARPE						
	calibration	MW ($T^{\text{VIX}} = 0.5$)	MW ($T^{\text{VIX}} = 3$)	MW ($T^{\text{VIX}} = 5$)	EWMA	market-implied
average precision	0.235445	0.358760	0.359595	0.348677	0.324854	0.305351
average Δv_0	0.009716	0.009460	0.009460	0.009460	0.009460	0.009460
average $\Delta \eta$	0.044100	0.000567	8.864555e-05	5.796916e-05	0.001152	0.001986
average $\Delta \lambda$	0.111124	0.062548	0.049030	0.051594	0.079298	0.082063
average $\Delta \kappa$	0.684585	0.634506	0.578065	0.522025	0.686207	0.691249
average $\Delta \rho$	0.053700	0.035944	0.048639	0.038836	0.034663	0.038442
APE						
	calibration	MW ($T^{\text{VIX}} = 0.5$)	MW ($T^{\text{VIX}} = 3$)	MW ($T^{\text{VIX}} = 5$)	EWMA	market-implied
average precision	0.019201	0.032250	0.035712	0.034686	0.038651	0.022346
average Δv_0	0.007801	0.009460	0.009460	0.009460	0.009460	0.009460
average $\Delta \eta$	0.027242	0.000567	8.864555e-05	5.796916e-05	0.001152	0.001986
average $\Delta \lambda$	0.094447	0.078535	0.059901	0.049769	0.126614	0.101588
average $\Delta \kappa$	0.537634	1.198459	0.931575	0.708374	1.879946	1.139020
average $\Delta \rho$	0.060817	0.030710	0.021987	0.016877	0.037403	0.043229

Table 1: Average precision and average day-to-day variation of the model parameters for the different calibration procedures and objective functions.

Globally speaking, the lost of precision due to the reduction of the number of calibrated parameters is not really significant if the long run variance is inferred from market data (i.e. from the VIX option surface). On the other hand if η is estimated from the time series of the VIX, the objective function is significantly higher than the optimal one, especially under the EWMA setting and for the RMSE and APE functionals. Moreover, fixing beforehand the parameter η , whatever

the technique, leads to a significant reduction of the day-to-day variation of the long run variance and, in a lower extend, of the correlation. On the other hand, the parameter κ turns out to vary more from one day to the other on average when v_0 and η are determined beforehand. Hence, the gain in stability for the long run variance is somehow offset by the reduction of stability in the mean reversion rate. The daily variation of the parameters v_0 and λ is either slightly larger or slightly smaller if we consider a reduced calibration, depending on the choice of the objective function and the particular technique used to estimate η . Similar conclusions can be drawn from the histogram of the daily variation of the model parameters (see Appendix A for the RMSE objective function case). We notice that the high value of the average daily change of κ under the different reduced calibration settings is due to a few values of $\Delta\kappa$, the huge majority of the values being of magnitude order one or below.

The value of the different objective functions for the full and reduced calibration procedures is given in Table 2 for four particular trading days: the 31st of October 2006 (low volatility regime period), the 18th of September 2008 (credit crisis period), the 11th of December 2008 (panic wave period) and the 19th of October 2009 (end of the credit crunch).

31/10/2006 : the low volatility regime period						
	calibration	MW ($T^{\text{VIX}} = 0.5$)	MW ($T^{\text{VIX}} = 3$)	MW ($T^{\text{VIX}} = 5$)	EWMA	market-implied
RMSE	0.867766	1.705515	1.680931	1.119678	5.497030	1.150935
ARPE	0.251314	0.322358	0.321602	0.259621	0.334561	0.245063
APE	0.009164	0.026294	0.025964	0.012442	0.049938	0.012713
18/09/2008 : the credit crisis period						
	calibration	MW ($T^{\text{VIX}} = 0.5$)	MW ($T^{\text{VIX}} = 3$)	MW ($T^{\text{VIX}} = 5$)	EWMA	market-implied
RMSE	3.817187	5.386691	8.079560	8.387793	3.866353	3.847507
ARPE	0.176251	0.376801	0.549708	0.563411	0.412158	0.356739
APE	0.032360	0.042608	0.065999	0.069585	0.032711	0.101225
11/12/2008 : the panic wave period						
	calibration	MW ($T^{\text{VIX}} = 0.5$)	MW ($T^{\text{VIX}} = 3$)	MW ($T^{\text{VIX}} = 5$)	EWMA	market-implied
RMSE	5.006764	9.074369	6.377103	6.785686	14.718441	9.193661
ARPE	0.246217	0.214005	0.629620	0.704610	0.317759	0.214418
APE	0.026347	0.047523	0.040497	0.043182	0.084877	0.045216
19/10/2009 : the end of the credit crisis						
	calibration	MW ($T^{\text{VIX}} = 0.5$)	MW ($T^{\text{VIX}} = 3$)	MW ($T^{\text{VIX}} = 5$)	EWMA	market-implied
RMSE	1.280985	1.503507	1.368586	3.035212	3.120949	1.716362
ARPE	0.292831	0.255591	0.251697	0.290182	0.279091	0.306004
APE	0.011769	0.012445	0.011580	0.020253	0.020669	0.013665
Calibration time						
	calibration	MW ($T^{\text{VIX}} = 0.5$)	MW ($T^{\text{VIX}} = 3$)	MW ($T^{\text{VIX}} = 5$)	EWMA	market-implied
RMSE	136.1719	40.8281	41.1719	22.6406	66.8281	46.1563
ARPE	140.2031	55.2656	53.9063	29.8125	32.9844	56.3906
APE	92.2969	36.1563	37.3438	22.9531	35.1563	35.8750

Table 2: Objective functions (upper) and Computation time in seconds (below) for the full and reduced calibration procedures.

The matching technique leading to the lowest value of the objective function is given in bold type for each objective function. For the first quote date, both the 5 years moving window and the

market implied techniques outperform the other reduced calibration procedures. For the second quote date either the EWMA or the MI estimates of η leads to an objective function of the same magnitude order as the optimal one. During the panic wave period, the Heston model usually leads to a poor fit of the option surface, especially for the widely used RMSE objective function. On the 19th of October 2009, the two narrowest moving windows and the market-implied techniques lead to an accurate fit of the option surface. As a conclusion, considering the market-implied estimate of the spot variance and of the long run variance leads to a distance between the model and market vanilla option prices pretty close to the optimal distance for each objective function, except during the investors fear wave of the end of 2008. Note that it might turn out that the value of the objective function under one of the reduced calibration settings is lower than the value obtained by considering the standard calibration procedure, which means that the full optimization technique reaches a local minimum. Moreover the computation time of the calibration procedure is reduced by at least almost three times for each functional if the calibration is performed on a reduced set which might turn out to be a significant advantage for practitioners.

Table 3 shows the optimal parameter set under the RMSE objective function for the four particular trading days mentioned in Table 2. It is clear that the optimal parameter set depends a lot on the choice of the calibration procedure.

31/10/2006 : the low volatility regime period			
	calibration	MW ($T^{\text{VIX}} = 0.5$)	MW ($T^{\text{VIX}} = 3$)
$\{v_0^*, \kappa^*, \eta^*, \lambda^*, \rho^*\}$	$\{0.0107, 0.5672, 0.0395, 0.2116, -0.8037\}$	$\{0.0123, 3.6445, 0.0202, 0.3834, -0.9315\}$	$\{0.0123, 3.5697, 0.0203, 0.3804, -0.9248\}$
	MW ($T^{\text{VIX}} = 5$)	EWMA	market-implied
$\{v_0^*, \kappa^*, \eta^*, \lambda^*, \rho^*\}$	$\{0.0123, 0.4602, 0.0396, 0.1908, -0.8302\}$	$\{0.0123, 3.2988, 0.0133, 0.2964, -0.8802\}$	$\{0.0123, 1.4575, 0.0250, 0.2702, -0.8335\}$
18/09/2008 : the credit crisis period			
	calibration	MW ($T^{\text{VIX}} = 0.5$)	MW ($T^{\text{VIX}} = 3$)
$\{v_0^*, \kappa^*, \eta^*, \lambda^*, \rho^*\}$	$\{0.0685, 0.2878, 0.0951, 0.2339, -0.9420\}$	$\{0.1096, 4.6785, 0.0501, 0.4705, -0.9999\}$	$\{0.1096, 1.5314, 0.0317, 0.2861, -0.9141\}$
	MW ($T^{\text{VIX}} = 5$)	EWMA	market-implied
$\{v_0^*, \kappa^*, \eta^*, \lambda^*, \rho^*\}$	$\{0.1096, 1.3190, 0.0282, 0.2726, -0.9106\}$	$\{0.1096, 5.1096, 0.0633, 0.8040, -0.9100\}$	$\{0.1096, 6.7472, 0.0600, 0.8998, -0.8489\}$
11/12/2008 : the panic wave period			
	calibration	MW ($T^{\text{VIX}} = 0.5$)	MW ($T^{\text{VIX}} = 3$)
$\{v_0^*, \kappa^*, \eta^*, \lambda^*, \rho^*\}$	$\{0.2403, 0.5527, 0.1271, 0.3748, -0.9803\}$	$\{0.3111, 0.8404, 0.1946, 0.5719, -0.9305\}$	$\{0.3111, 1.1906, 0.0579, 0.3713, -0.9876\}$
	MW ($T^{\text{VIX}} = 5$)	EWMA	market-implied
$\{v_0^*, \kappa^*, \eta^*, \lambda^*, \rho^*\}$	$\{0.3111, 1.1608, 0.0431, 0.3165, -0.9835\}$	$\{0.3111, 0.4115, 0.3757, 0.5561, -0.9999\}$	$\{0.3111, 0.8404, 0.2003, 0.5802, -0.9305\}$
19/10/2009 : the end of the credit crisis			
	calibration	MW ($T^{\text{VIX}} = 0.5$)	MW ($T^{\text{VIX}} = 3$)
$\{v_0^*, \kappa^*, \eta^*, \lambda^*, \rho^*\}$	$\{0.0429, 1.4953, 0.0914, 0.5227, -0.8311\}$	$\{0.0462, 2.4491, 0.0783, 0.6191, -0.8492\}$	$\{0.0462, 1.4713, 0.0889, 0.5115, -0.8335\}$
	MW ($T^{\text{VIX}} = 5$)	EWMA	market-implied
$\{v_0^*, \kappa^*, \eta^*, \lambda^*, \rho^*\}$	$\{0.0462, 11.003, 0.0603, 1.1522, -0.9754\}$	$\{0.0462, 11.7220, 0.0598, 1.1841, -0.9842\}$	$\{0.0462, 3.4363, 0.0729, 0.7080, -0.8666\}$

Table 3: Optimal parameter set $\{v_0^*, \kappa^*, \eta^*, \lambda^*, \rho^*\}$ for the RMSE objective function

The calibration risk was defined by Detlefsen and Hardle in [4] as the difference in the value of the calibrated parameters arising from the different specifications of the objective function. Figures 8 and 9 show the maximum absolute value of the difference between the optimal parameter $p^* \in \{v_0^*, \kappa^*, \eta^*, \lambda^*, \rho^*\}$ obtained with the different objective functions,

$$\max(|p_{\text{RMSE}}^* - p_{\text{APE}}^*|, |p_{\text{RMSE}}^* - p_{\text{ARPE}}^*|, |p_{\text{APE}}^* - p_{\text{ARPE}}^*|)$$

for the different calibration procedures³.

³The period is split in between the low volatility period and the credit crisis period given the significant difference in the calibration risk magnitude order during these two periods.

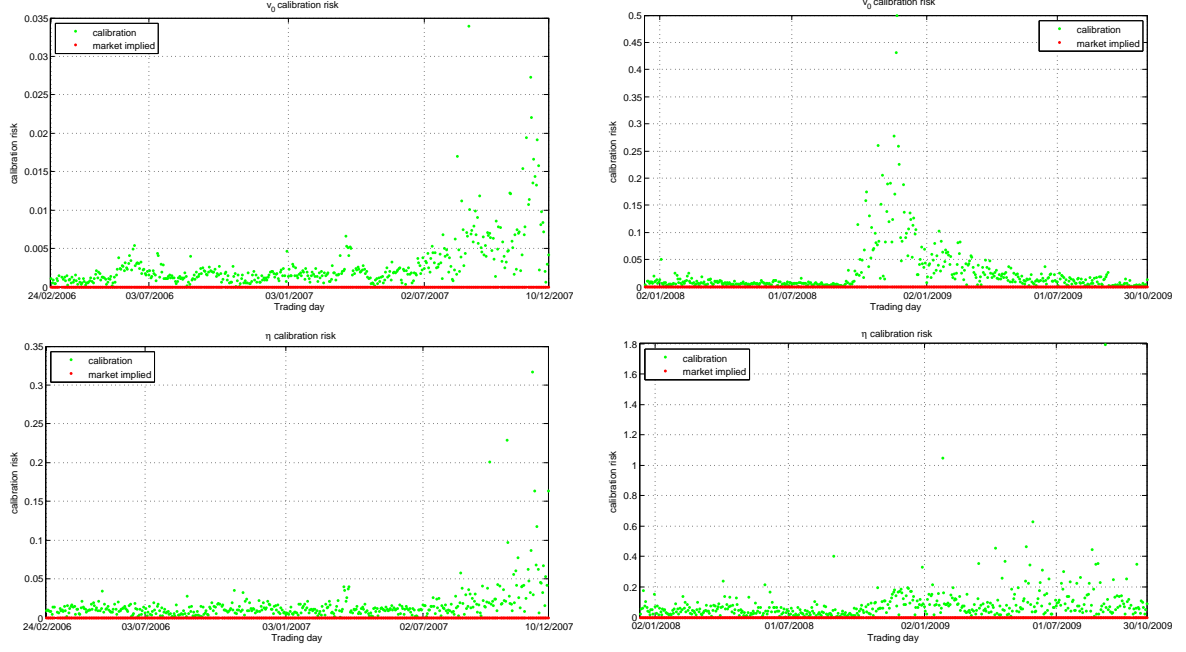


Figure 8: Evolution of the calibration risk for the parameter v_0 (upper) and η (lower) through time for the different calibration procedures.

As mentioned in [5], v_0 and η do not vary with the functional for the reduced calibrations since they are directly inferred from the VIX spot and the VIX historical prices or the VIX option surface, respectively. Hence each reduced calibration procedure allows to eliminate the calibration risk arising from the determination of these two parameters. Moreover, the optimal correlations ρ 's obtained by considering one or another functional are usually closer to each other in the reduced calibration settings than in the full calibration setting, leading to a lower calibration risk. The contrary holds for the mean reversion rate κ whereas the calibration risk of the volatility of variance λ is globally speaking of the same magnitude order under the standard and reduced calibration procedures.

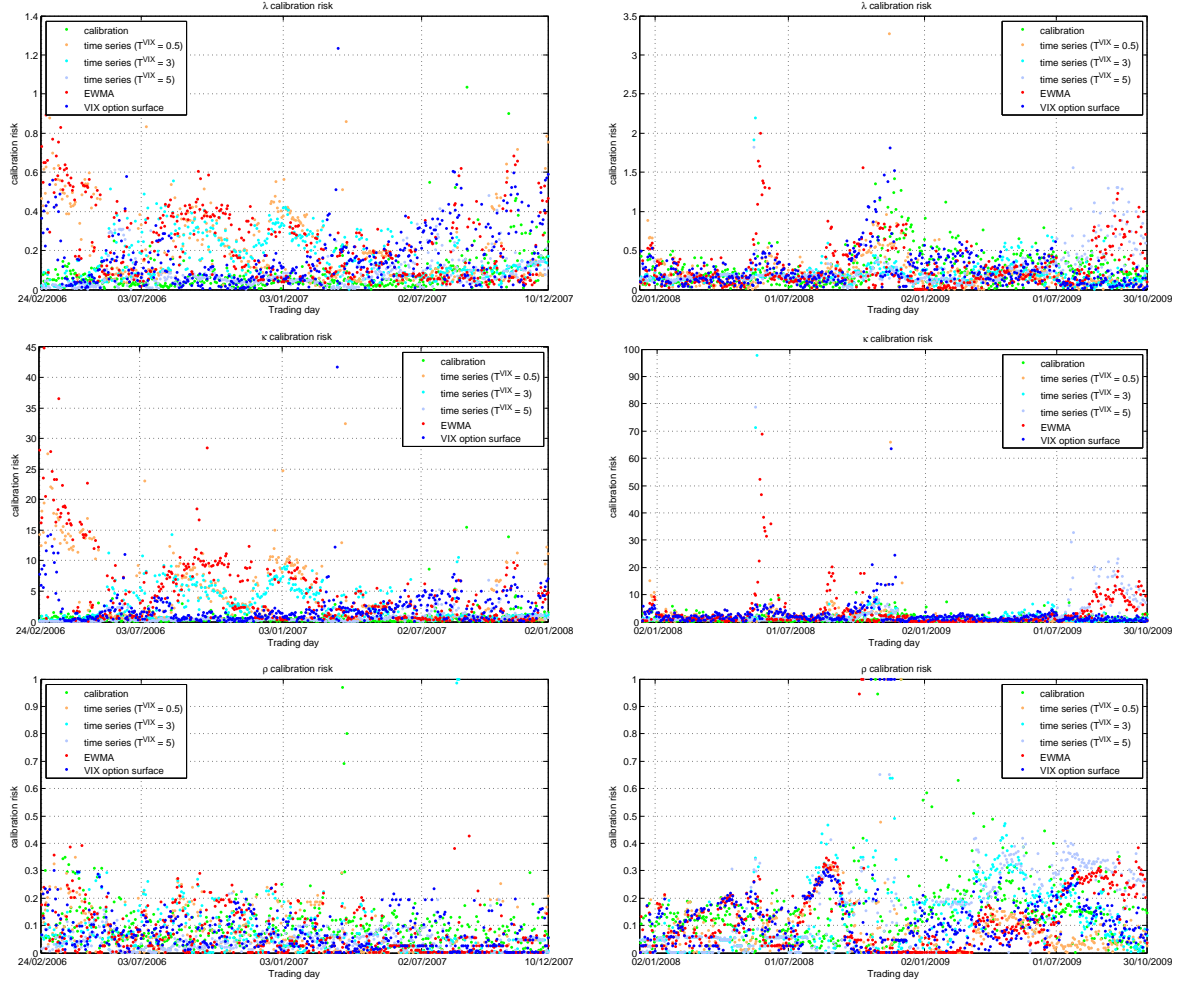


Figure 9: Evolution of the calibration risk for the parameter λ (upper), κ (center) and ρ (lower) through time for the different calibration procedures.

4 Pricing of Exotic Options

This section features a comparison of the price of different exotic options under the different calibration procedures for the three particular quoting dates defined in Section 3.1, considering the RMSE objective function. As it can be seen from Figure 5, Figure 6 and Figure 7 as well as from Table 2, both the standard and the market-implied reduced calibration procedures lead to a pretty good fit of the whole set of S&P500 liquid options except during the panic wave period. Hence, we can usually hardly discriminate between the two calibration methods. We will then further compare the calibration procedures by pricing several exotic options ranging from Asian options, 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 (2.1) is discretised by using a first order Euler scheme and the variance process (2.2) using a second order Milstein scheme. The Monte Carlo simulation is performed by considering one million scenarios and 252 trading days a year.

The payoff of Asian options depends on the arithmetic average of the stock price from the emission to the maturity date of the option. The time $t_0 = 0$ Asian call and put price is given by

$$AC = \exp(-rT) \mathbb{E}_{\mathbb{Q}}[(\text{mean}_{0 \leq t \leq T} S_t - K)^+]$$

and

$$AP = \exp(-rT) \mathbb{E}_{\mathbb{Q}}[(K - \text{mean}_{0 \leq t \leq T} S_t)^+],$$

respectively.

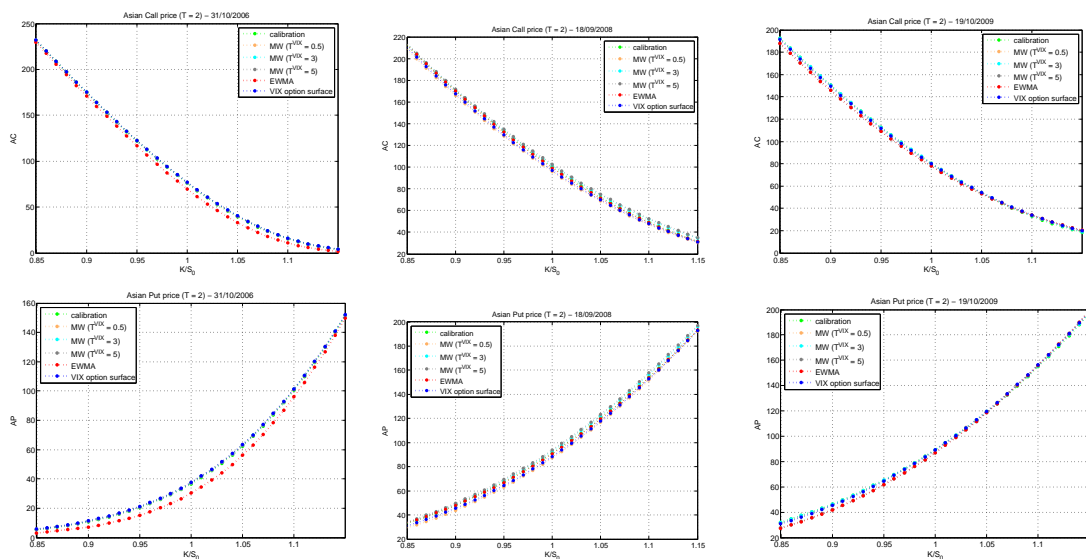


Figure 10: Asian call (upper) and put (lower) option prices

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

$$LC = \exp(-rT) \mathbb{E}_{\mathbb{Q}}[(S_T - m_T^S)^+]$$

and

$$LP = \exp(-rT) \mathbb{E}_{\mathbb{Q}}[(M_T^S - S_T)^+],$$

respectively where m_t^X and M_t^X denotes the minimum and maximum process 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 lookback call and put prices are given in Table 4.

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

Table 4: Call and put Lookback prices.

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. There exist two different types of barrier options:

- the knock-out barrier options which cease to exist when the underlying asset price reaches the threshold H . The down-and-out barrier is worthless unless the minimum of the stock price remains above some low barrier H whereas the up-and-out barrier is worthless unless the maximum of the stock price remains below some high barrier H :

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

and

$$\text{UOBC} = \exp(-rT) \mathbb{E}_{\mathbb{Q}}[(S_T - K)^+ \mathbf{1}(M_T^S < H)];$$

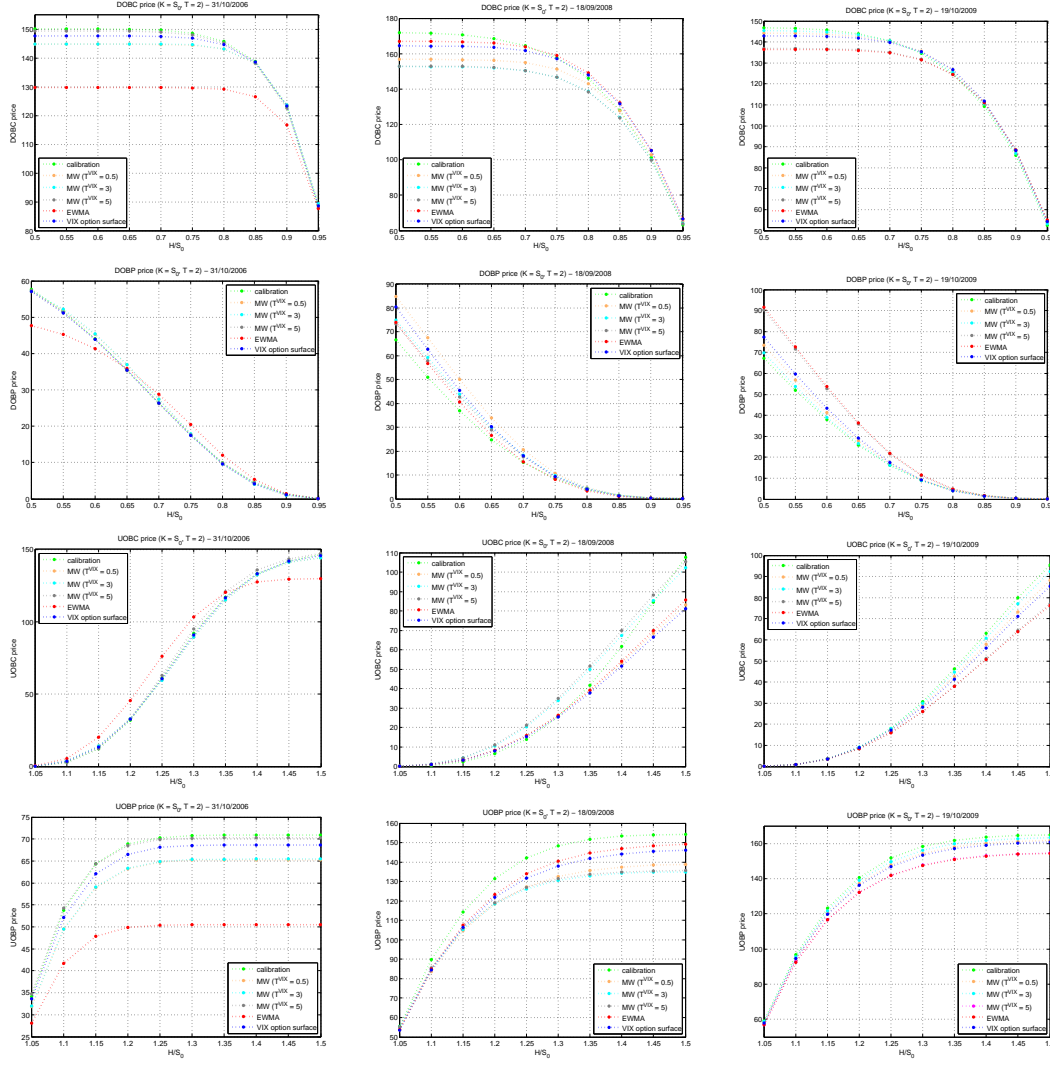


Figure 11: Down-and-out (upper) call and put and Up-and-out (lower) call and put option prices

- the knock-in barrier options which come into existence only when the asset price reaches the barrier H . The down-and-in barrier is worthless unless the minimum of the stock price hits some low barrier H whereas the up-and-in barrier is worthless unless the maximum of the stock price crossed some high barrier H :

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

and

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

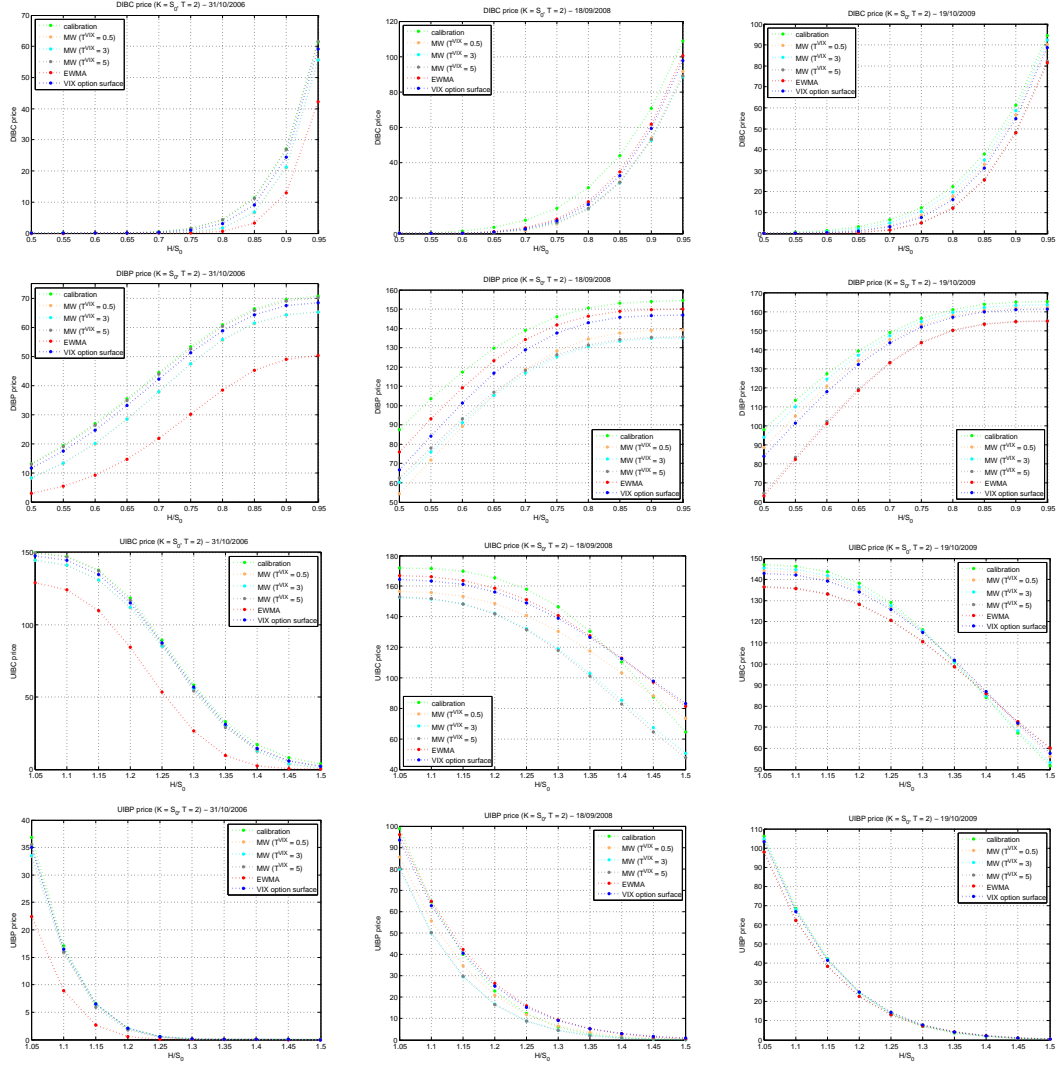


Figure 12: Down-and-in (upper) call and put and Up-and-in (lower) call and put option prices

The payoff of a cliquet option depends on the sum of the stock return over a series of consecutive time periods $[t_i, t_{i+1}]$; 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 the cliquet options protect the investor against downside risk while allowing him for significant upside potential:

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

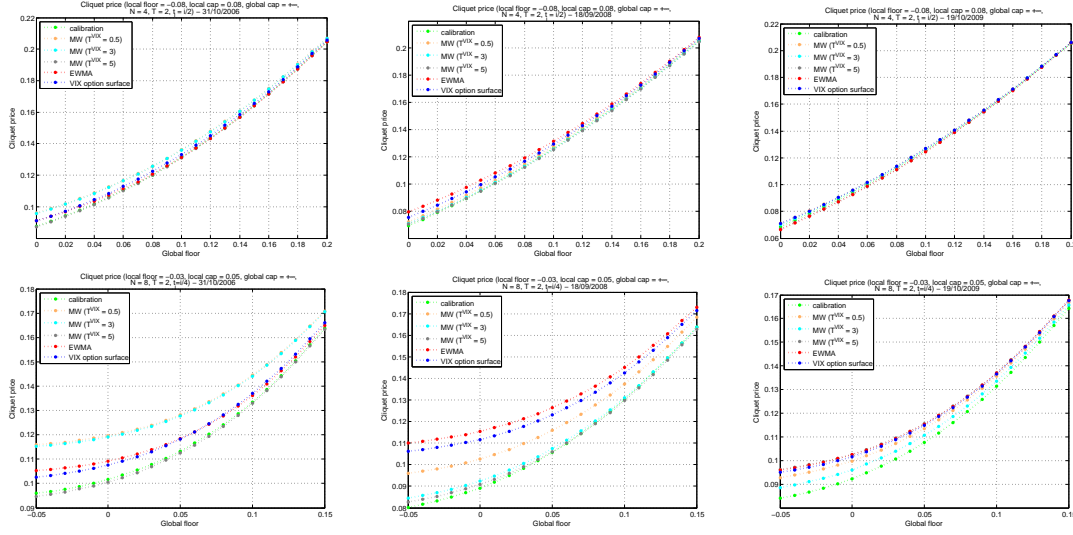


Figure 13: Cliquet option prices

As it can be seen from Figure 10 to Figure 13, except for the Asian options, the Heston price of the different exotic options turns out to be sensitive to the calibration method. For the 31st of October 2006, the exotic prices obtained by the reduced calibration procedure under the market-implied setting, and in a larger extend under the 5 years moving window setting, barely differ from the prices obtained under the optimal parameter set⁴. On the other hand, the EWMA reduced calibration leads to significantly different exotic prices. For the 18th of September 2008, the closest prices to the prices computed with the optimal parameter set are usually obtained by considering the EWMA technique, although this depends on the particular exotic option. Indeed, for the cliquet options the distance between the standard calibration and reduced calibration prices is minimum if we consider the five years window and maximum for the EWMA technique. For the 19th of October 2009, the closest prices are usually obtained by considering the three years window parameter set. As a conclusion 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. 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, for the knock-out barrier options the calibration risk significantly increases with the option prices, i.e. increases and decreases with the ratio $\frac{H}{S_0}$ for up-and-out and down-and-out barrier options, respectively. Note that for the knock-in barrier options the influence of the barrier level on the calibration risk is less marked. Moreover, the calibration risk of cliquet options decreases with the global floor level. In particular, the calibration risk of the widely traded capital protected cliquets (i.e. cliquets with a global floor equal to zero) turns out to be significant.

⁴Nevertheless, we notice an exception for the cliquet options under the MI setting.

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 Figure 14 to Figure 18). We computed the price of 14 different exotic options every two weeks by the Monte Carlo simulation for the six sets of model parameters obtained under the different calibration procedures, 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 Asian, Lookback and knock-in barrier options. Moreover, except during the investors' fear wave, the Asian, and in a smaller extend the Lookback call and put prices are not significantly different under the different calibration procedures indicating that the calibration risk is not predominant for these two kinds 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.

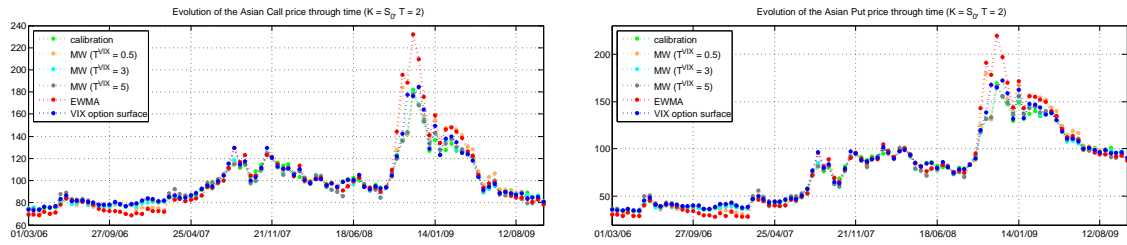


Figure 14: Evolution of the Asian call (left) and put price (right) through time for the different calibration procedures.

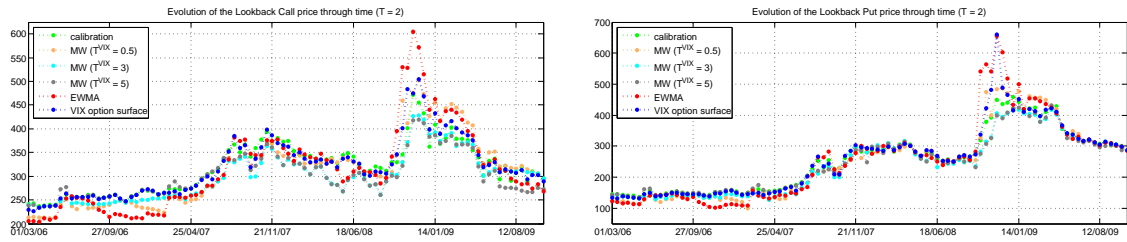


Figure 15: Evolution of the Lookback call (left) and put price (right) through time for the different calibration procedures.

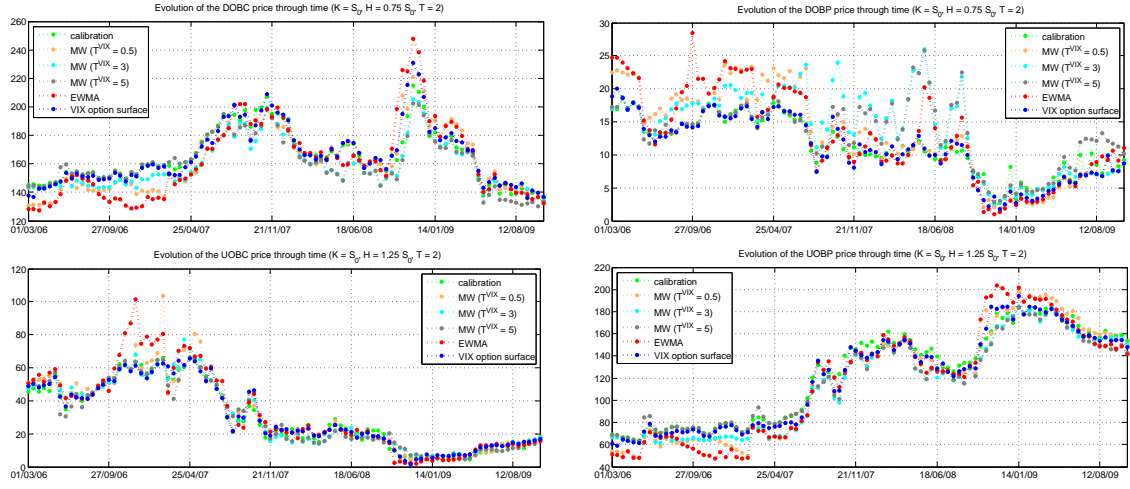


Figure 16: Evolution of the Down-and-out (upper) call (left) and put and Up-and-out (lower) call (left) and put price (right) through time for the different calibration procedures.

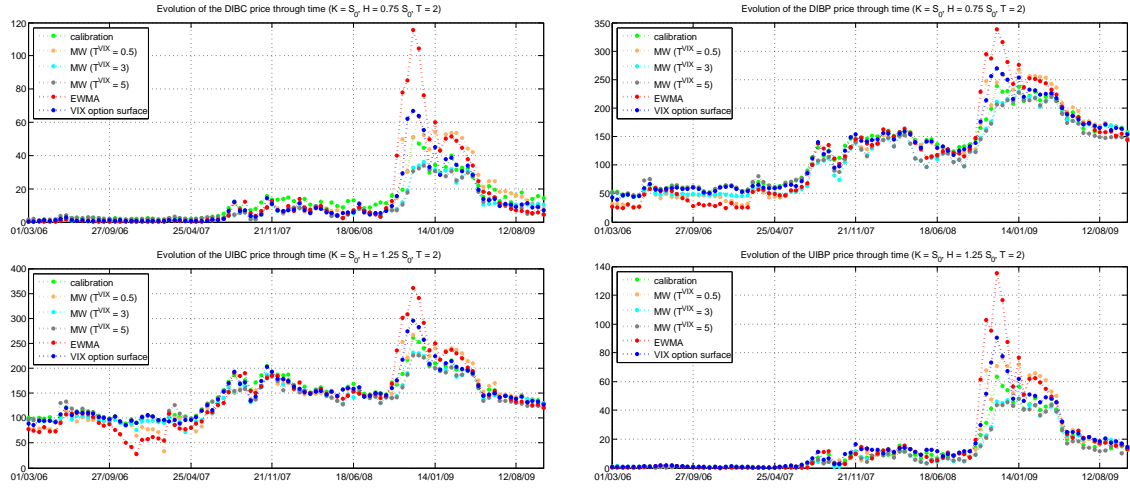


Figure 17: Evolution of the Down-and-in (upper) call (left) and put and Up-and-in (lower) call (left) and put price (right) through time for the different calibration procedures.

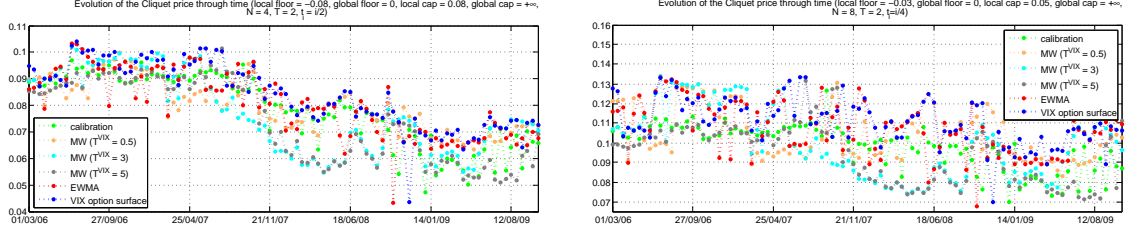


Figure 18: Evolution of the cliquet price through time for the different calibration procedures.

The Box plot of the ratio between the reduced and full exotic option prices, $\frac{P^{RC}}{P^{FC}}$ where P^{RC} and P^{FC} denotes the reduced and full calibration prices, respectively, is given in Annex B. For the Asian options, the ratio $\frac{P^{RC}}{P^{FC}}$ is included between 0.9 and 1.1 for almost all the days and for all the reduced calibration settings (see Figure 27). Moreover, the ratio closest to one is obtained by considering the MI parameter set, followed by the three and five years MW parameters. The greatest difference between P^{RC} and P^{FC} corresponds to the EWMA setting. Furthermore, roughly fifty percents of the ratios are above one under each calibration setting, indicating that none of the reduced calibration leads to typically higher or lower Asian option prices.

Except for the EWMA parameter set, the empirical value range of the ratio $\frac{P^{RC}}{P^{FC}}$ is higher for the Lookback options than for the Asian options, indicating an higher calibration risk. Moreover, all the reduced Lookback prices, except the MI ones, are typically lower than the prices corresponding to the full calibration procedure. The box plots of the DOBC, UOBP, DIBP and UIBC barrier options have a similar shape (i.e. both the boxes as well as the whole plot have a similar length) to this of the Lookback options, indicating that the calibration risk is of the same magnitude order for these exotic options. On the other hand, both the length of the box and the total value range are significantly higher for the DOBP, UOBC, DIBC and UIBP barrier options, highlighting the predominance of the calibration risk for these exotic structures. We also notice that for almost all the barrier options, the reduced price P^{RC} is usually lower than the full price P^{FC} , the opposite trend being observed for the DOBP and UOBC barrier options only as well as for the cliquet options. The extend of the cliquet box plots is located in between this of the Lookback box plots and this of the riskiest barrier option box plot, indicating a significant model risk for the cliquet options too. It is also clear that the smallest value range of the ratio $\frac{P^{RC}}{P^{FC}}$ is obtained under the MI setting, followed by the five and three years MW procedures whereas the biggest difference between any full and reduced exotic prices results from the EWMA and the six months MW calibration procedures.

4.2 Calibration risk: the choice of the objective function

Figure 19 to Figure 23 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 a different functional is more marked under the full 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 market data and are thus constant with respect to the choice of the objective function.

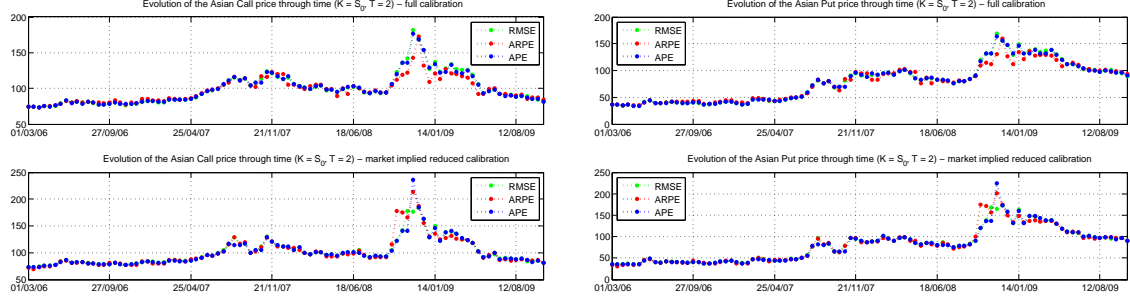


Figure 19: Evolution of the Asian call (left) and put price (right) through time for the different objective functions.

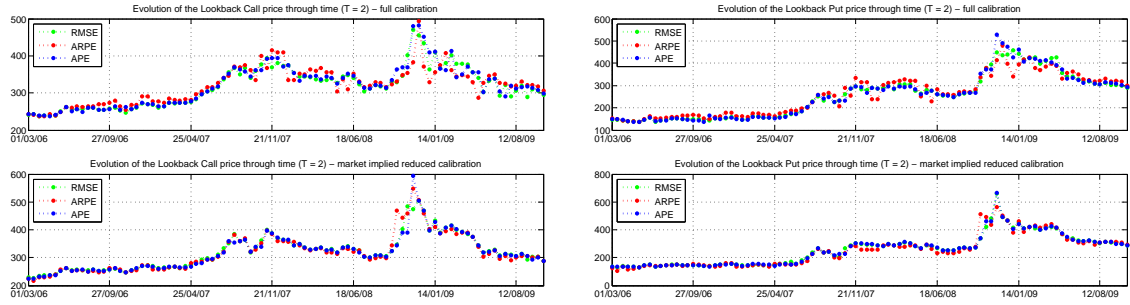


Figure 20: Evolution of the Lookback call (left) and put price (right) through time for the different objective functions.

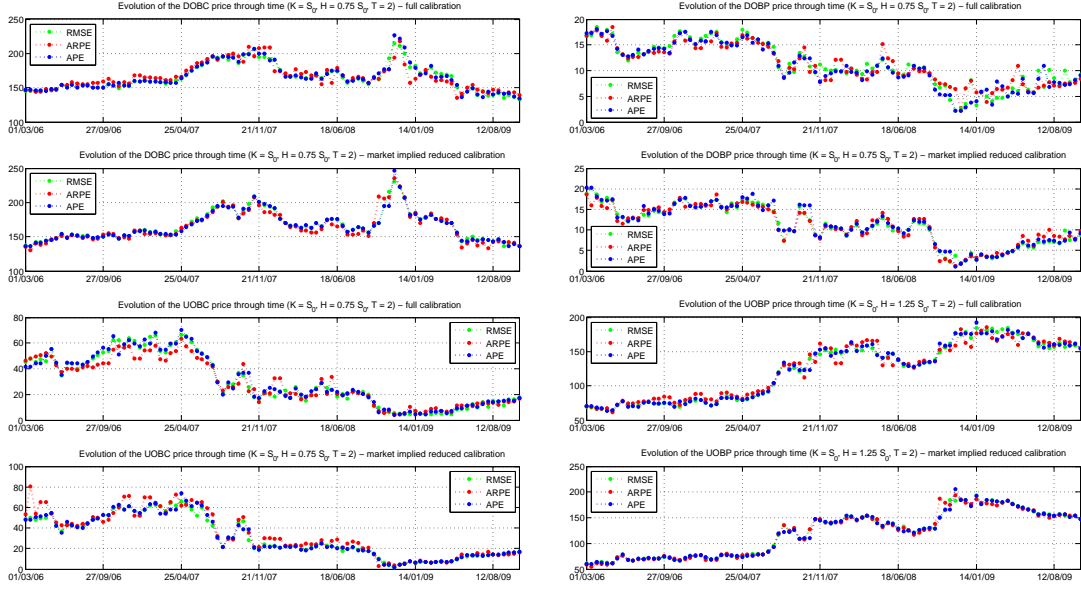


Figure 21: Evolution of the Down-and-out (upper) call (left) and put and Up-and-out (lower) call (left) and put price (right) through time for the different objective functions.

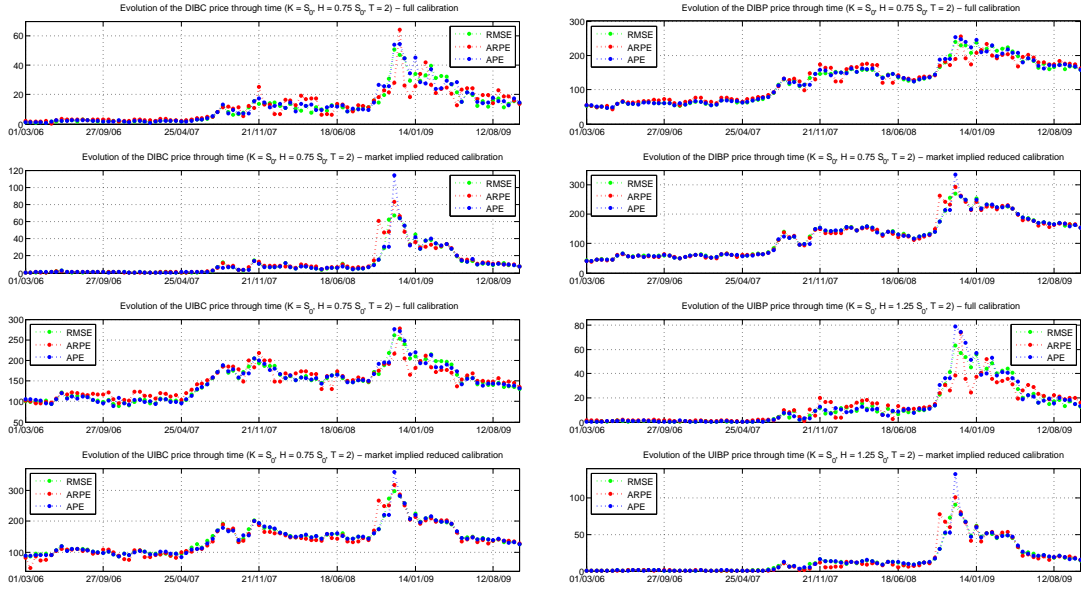


Figure 22: Evolution of the Down-and-in (upper) call (left) and put and Up-and-in (lower) call (left) and put price (right) through time for the different objective functions.

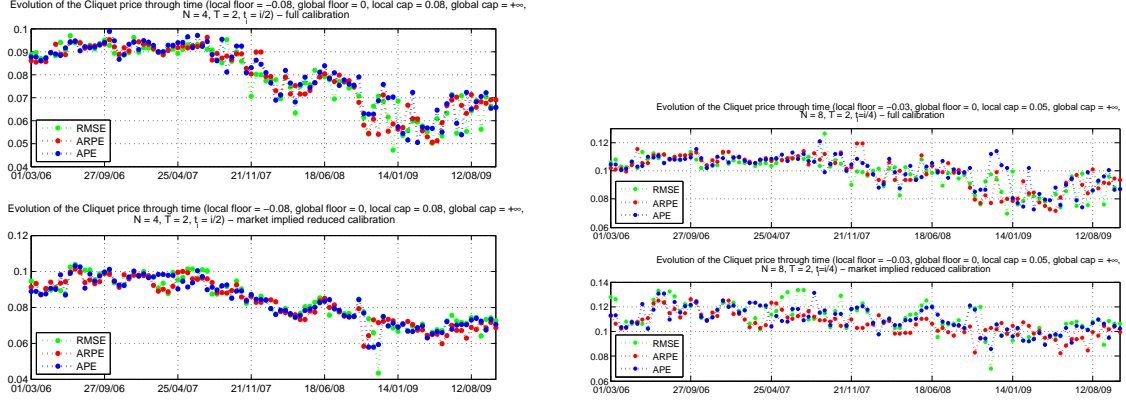


Figure 23: Evolution of the cliquet price through time for the different objective functions.

Figure 24 shows the evolution of the risk which arises from both the calibration setting and the functional specification. We measure this risk by

$$\frac{\sqrt{\text{Var}(\hat{P}_{i,j})}}{\sum_{i=1}^M \sum_{j=1}^N \frac{\hat{P}_{i,j}}{NM}} \quad (4.1)$$

where N denotes the number of calibration settings and M the number of objective functions (see [3]). The number of prices taken into account in the global model risk estimate amounts thus to 18.

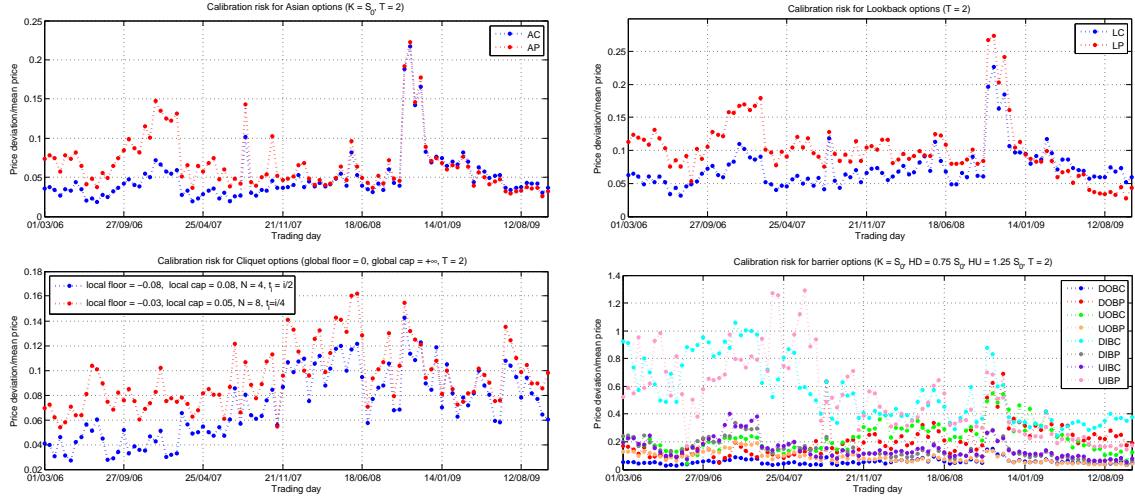


Figure 24: Average model risk for the different exotic options.

It is clear that the global model risk is higher for the cliquet and some of the barrier options (i.e. the DIBC and UIBP, and in a smaller extend the DOBP and UOBC options, see also Table 5). Moreover, the model risk is significantly higher during the panic wave period for the Asian, lookback and some barrier options (i.e. the DOBP and UOBC options).

As it can be seen from Table 5, the calibration risk arising from the specification of the objective function is significantly higher on average (often of one magnitude order) than the model risk due to the choice of the calibration procedure. Nevertheless, we notice an exception for the Asian call option, although the calibration and model risks then turn out to be of the same magnitude order. Moreover, taking into account the different calibration procedures and not only the different objective functions allows to reduce the model risk in the sense that its measure (4.1) turns out to be on average lower.

calibration	AC	AP	LC	LP	Cliquet 1	Cliquet 2	DOBC	DOBP	UOBC	UOBP	DIBC	DIBP	UIBC	UIBP
Calibration risk arising from the calibration procedure														
RMSE	0.033908	0.048530	0.060000	0.064051	0.085506	0.112690	0.039281	0.198357	0.135771	0.067489	0.524566	0.123583	0.087288	0.357493
APE	0.038819	0.054284	0.066131	0.080649	0.071326	0.095247	0.045615	0.185915	0.152006	0.075266	0.528947	0.130241	0.118739	0.465904
ARPE	0.066253	0.088669	0.096187	0.140164	0.047125	0.063327	0.077479	0.2089320	0.254173	0.115150	0.642981	0.182209	0.212589	0.688647
Calibration risk arising from the objective function														
calibration	0.020271	0.225636	0.520212	0.425640	1.727657	1.726586	0.274359	1.243682	0.814589	0.134273	1.302822	0.199536	0.217512	1.318924
MW ($T^{\text{VIX}} = 0.5$)	0.022378	0.252357	0.506769	0.380589	1.727545	1.726371	0.256693	1.205416	0.747263	0.168450	1.374050	0.285221	0.207119	1.318154
MW ($T^{\text{VIX}} = 3$)	0.037969	0.256160	0.492920	0.378976	1.727664	1.726532	0.246357	1.153557	0.692015	0.149320	1.486095	0.242757	0.172117	1.427167
MW ($T^{\text{VIX}} = 5$)	0.039331	0.232118	0.497650	0.399492	1.727786	1.726739	0.251869	1.173680	0.731435	0.112777	1.470301	0.168577	0.181764	1.442246
EWMA	0.014861	0.246158	0.506096	0.381877	1.727508	1.726304	0.256144	1.196937	0.760032	0.167349	1.377761	0.284868	0.224092	1.306372
market-implied	0.018896	0.229911	0.513259	0.407050	1.727500	1.726191	0.266475	1.240836	0.780326	0.139903	1.390347	0.222046	0.187905	1.296268
Calibration risk arising from both the calibration procedure and the objective function														
ALL	0.049028	0.213804	0.434454	0.352613	1.453121	1.452149	0.226264	1.011796	0.644223	0.147501	1.174243	0.227528	0.211867	1.135486

Table 5: Calibration Risk.

5 Conclusion

We have proposed 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 have shown that this estimate follows the same trend as the long run variance parameter η obtained by the standard calibration of the Heston model, although it is typically lower. We also performed 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 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 one, except during the period characterized by huge investors fear.

Moreover each reduced calibration procedure leads to a significant reduction of both the calibration computation time and the calibration risk (except for the mean reversion rate κ) which arises by specifying different measures for the error between the model and market prices. Furthermore, considering the reduced calibration set leads to more stable long run variance η and correlation ρ through time, which has to be counterbalanced by the lost of stability of the mean reversion rate κ .

Although the market implied reduced calibration leads to a fit of the option surface of a similar quality to 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. This might be explained by the fact that the two calibration procedures, as well as, in a larger extend, 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 way to price exotic products out of the liquid vanilla option prices.

A Day-to-day variation of the model parameters

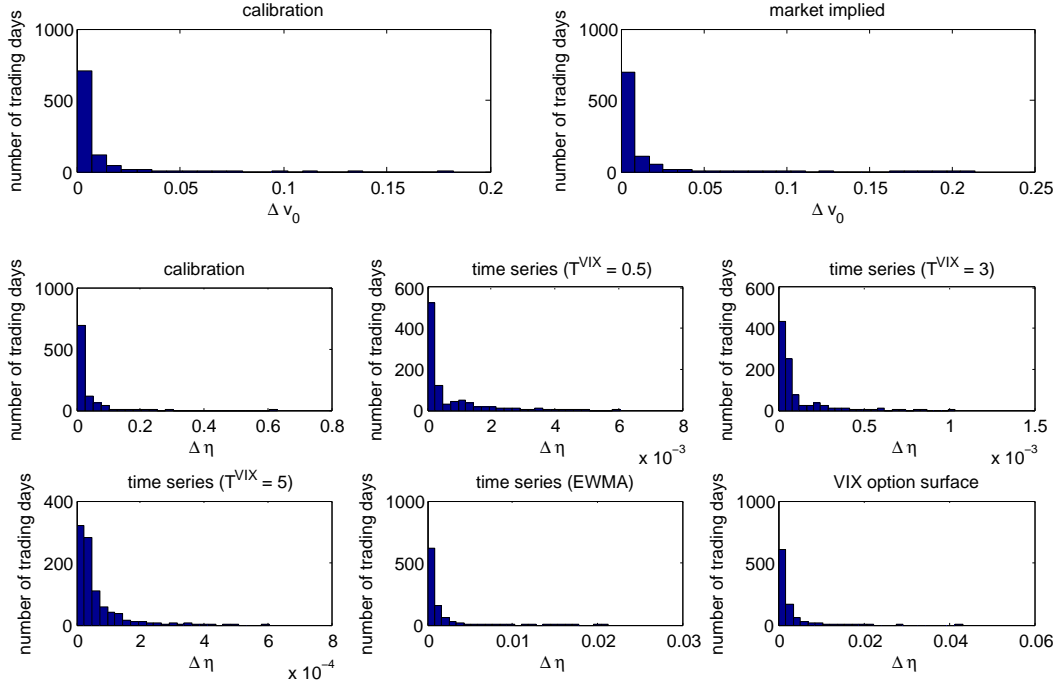


Figure 25: Histogram of the day-to-day variation of v_0 (upper) and η (lower) for the different calibration procedures.

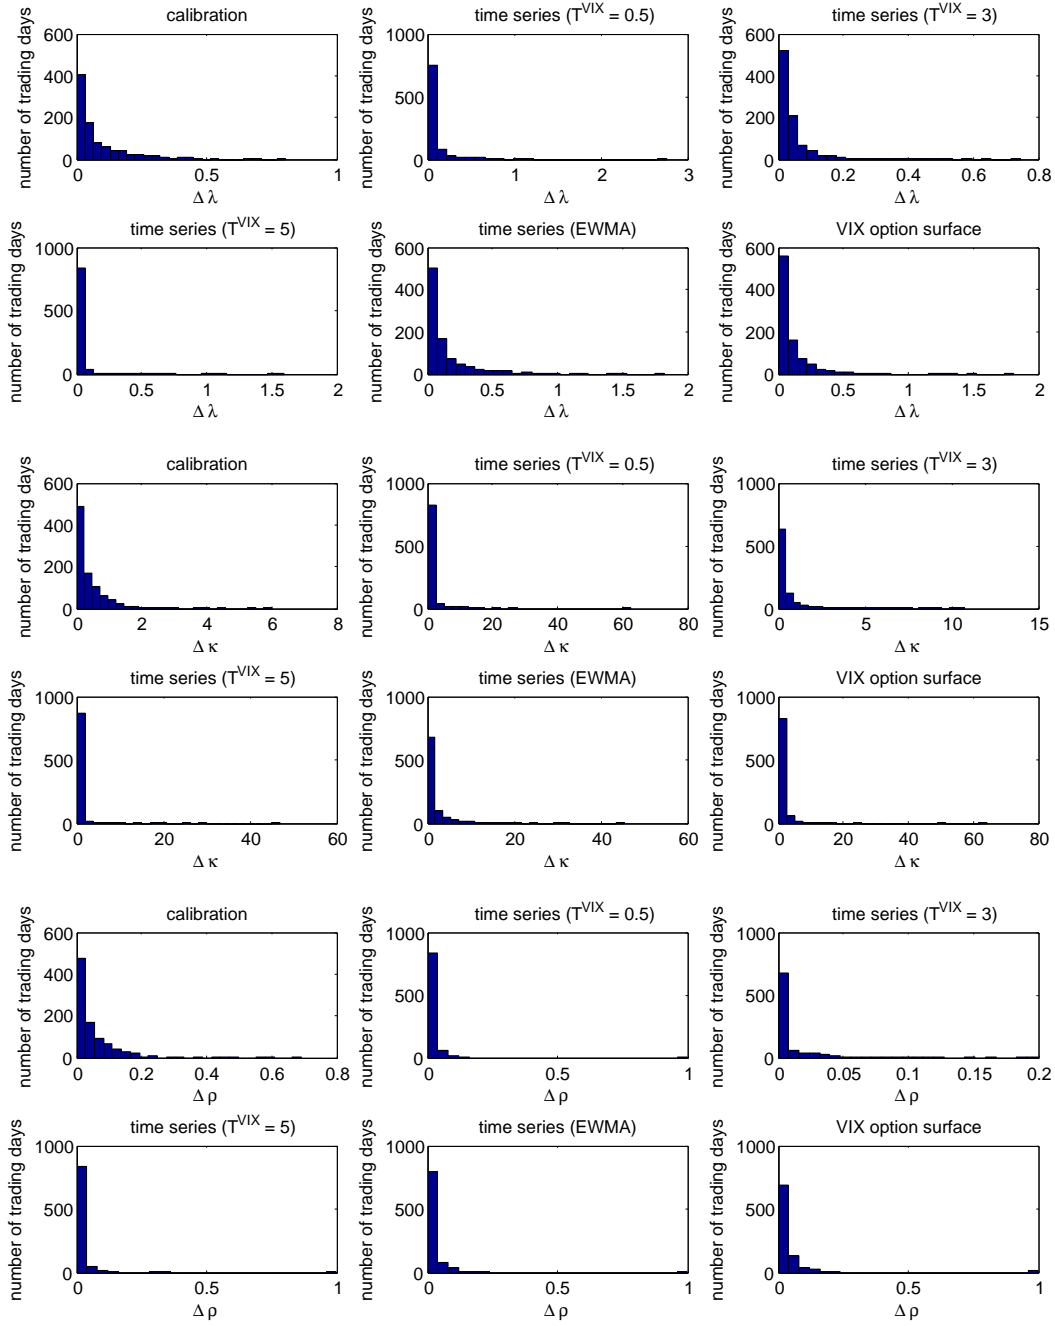


Figure 26: Histogram of the day-to-day variation of λ (upper), κ (center) and ρ (lower) for the different calibration procedures.

B Box plot of the ratio between the reduced and the full exotic prices

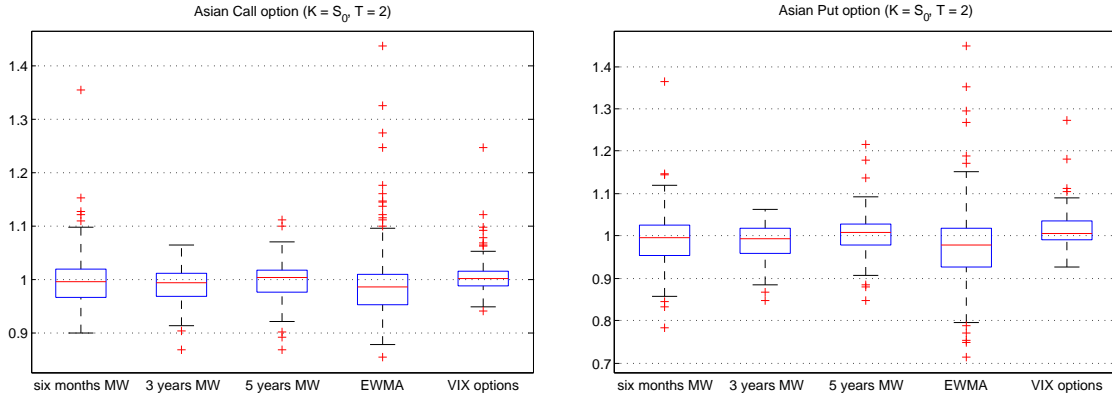


Figure 27: Box plot of the ratio between the reduced and full Asian call (left) and put price (right) for the different reduced calibration procedures.

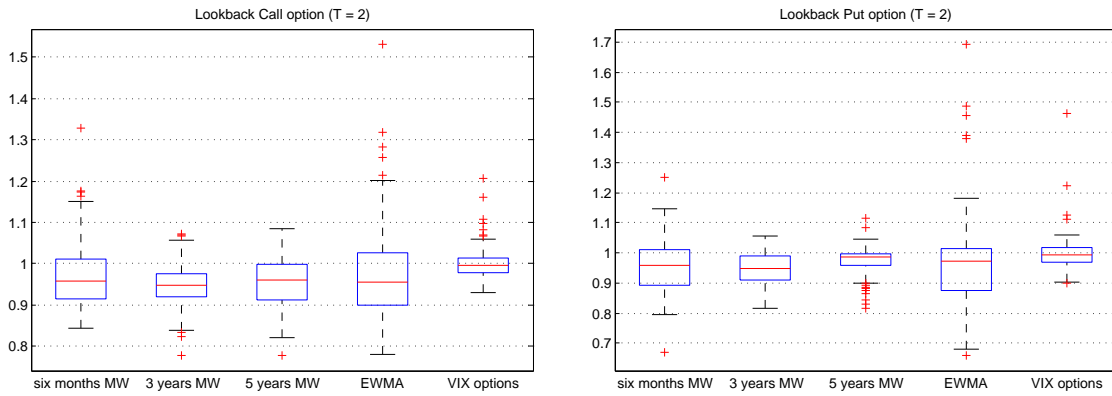


Figure 28: Box plot of the ratio between the reduced and full Lookback call (left) and put price (right) for the different reduced calibration procedures.

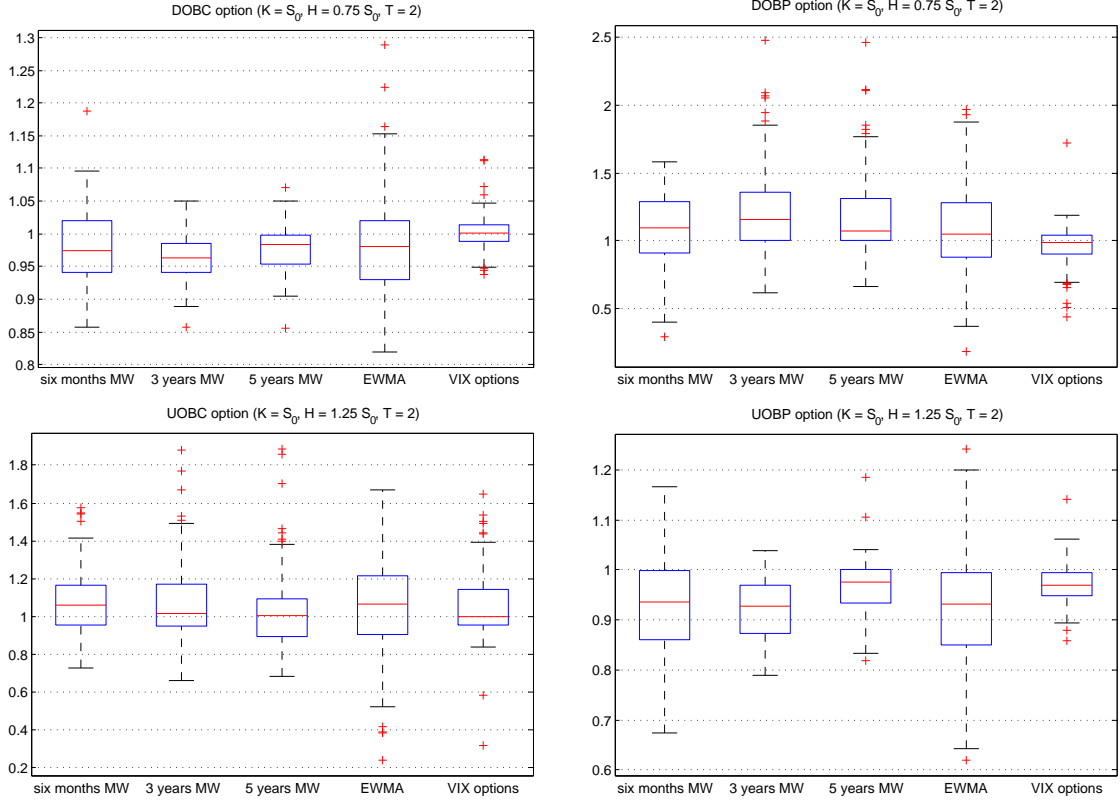


Figure 29: Box plot of the ratio between the reduced and full Down-and-out (upper) call (left) and put and Up-and-out (lower) call (left) and put price (right) for the different reduced calibration procedures.

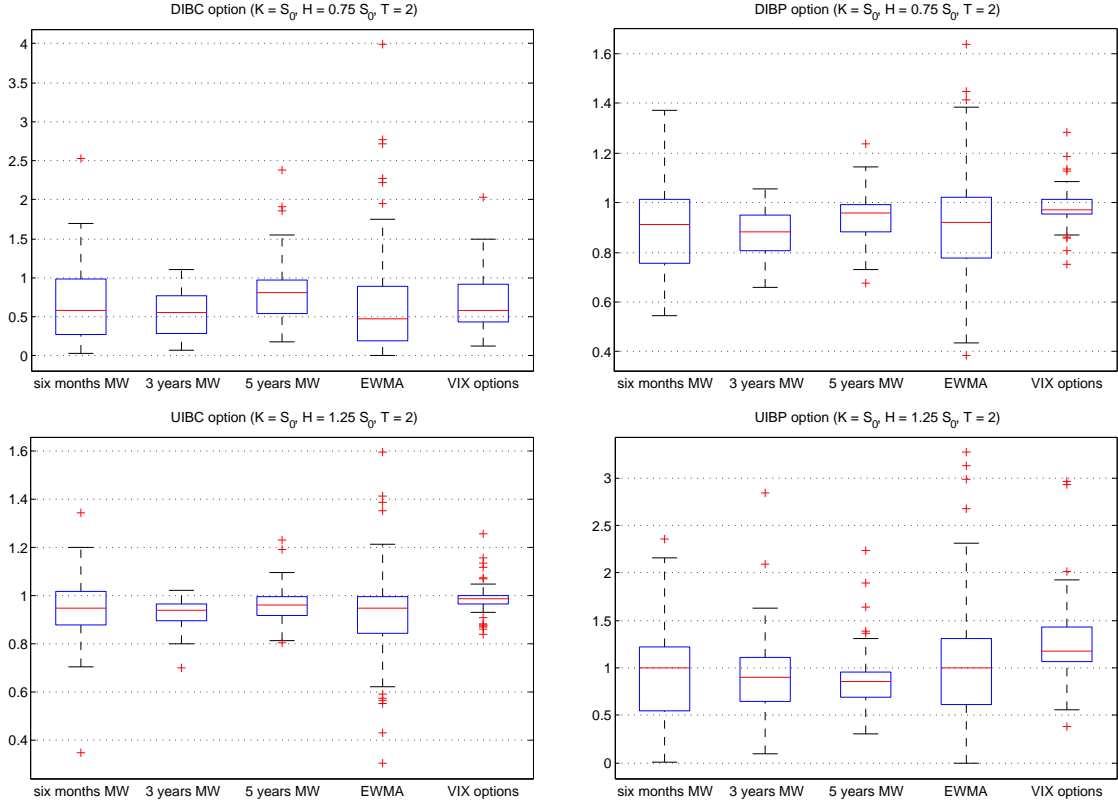


Figure 30: Box plot of the ratio between the reduced and full Down-and-in (upper) call (left) and put and Up-and-in (lower) call (left) and put price (right) for the different reduced calibration procedures.

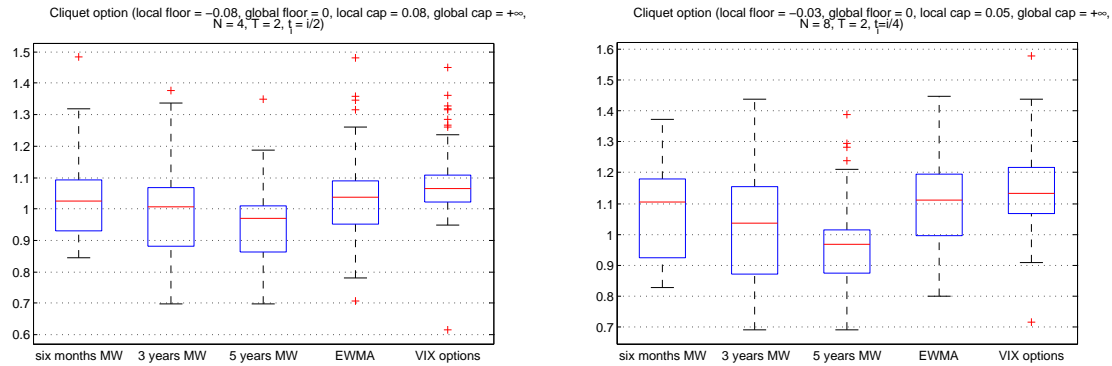


Figure 31: Box plot of the ratio between the reduced and full cliquet price for the different reduced calibration procedures.

References

- [1] Albrecher, H., Mayer, P., Schoutens, W. and Tistaert, J. (2007) The little Heston trap. *Wilmott Magazine*, 83-92.
- [2] Carr, P. and Madan, D.B. (1998) Option valuation using the fast Fourier transform. *Journal of Computational Finance* **2**, 61-73.
- [3] Cont, R. (2006) Model Uncertainty and its Impact on the Pricing of derivative Instruments. *Mathematical Finance* **16-3**, 519-547.
- [4] Detlefsen, K. and Hardle W. K. (2007) Calibration Risk for Exotic Options. *The Journal of Derivatives* **14**, 47-63.
- [5] Guillaume, F. and Schoutens W. (2010) Use a reduced Heston or reduce the use of Heston? *Wilmott Journal*, to appear.
- [6] Heston, S. (1993) A closed-form solution for options with stochastic volatility. *Review of Financial Studies* **6**, 327-243.
- [7] Hull, J. (2006) *Options, Futures, and Other Derivatives (sixth edition)*. Pearson Prentice Hall.
- [8] Mikhailov, S. and Nögel, U. (2003) Heston's Stochastic Volatility Model Implementation, Calibration and Some Extensions. *Wilmott magazine*, July 2003, 74-79.
- [9] Schoutens, W., Simons, E. and Tistaert, J. (2004) A Perfect calibration ! Now what ? *Wilmott Magazine*, March 2004, 66-78.
- [10] Whaley, R. E. (2009) Understanding the VIX. *The Journal of Portfolio Management*, **35**, 98-105.
- [11] Wilmott, P. (2006) *Paul Wilmott on Quantitative Finance*. Wiley.