

EURANDOM PREPRINT SERIES
2010-003

**Use a reduced Heston or
reduce the use of Heston?**

F. Guillaume, W. Schoutens
ISSN 1389-2355

Use a reduced Heston or reduce the use of Heston?

Florence Guillaume*

Wim Schoutens[†]

January 12, 2010

Abstract

This paper features the calibration performance of the Heston model for two different calibration procedures. The first consists of the standard calibration on the whole parameter set and the second one, called reduced calibration consists of a calibration on the reduced parameter set $\{\kappa, \lambda, \rho\}$ where the spot variance v_0 and the long run variance η are inferred beforehand from the time series of the VIX volatility index. It is shown that both calibration procedures lead to an accurate fit of the vanilla option surface. Furthermore both the computation time and the calibration risk of the reduced calibration procedure turns out to be significantly lower, which might turn out to be a considerable advantage for practitioners. This paper also features a comparison of the price of different exotic options for the two calibration procedures.

*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 a calibration performance of the Heston model under two calibration settings: the first consists of the standard calibration procedure 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 initial variance of the S&P500 stock price process v_0 and its long run variance η are fixed beforehand from the time series of the VIX CBOE volatility index. First, we calibrate the Heston model on liquid options for the whole strike and maturity range for the two procedures and for three particular windows of the VIX time series. In particular, it is shown that the two methodologies lead to a pretty good fit of the implied volatility surface for a well chosen window to estimate the long run variance. Moreover the alternative calibration procedure turns out to be promising for practitioners purposes since it allows to significantly reduce both the calibration computation time and the calibration risk while leading to objective functions of the same order of magnitude

Next, we price several exotic options in the two settings for three particular quoting dates and show that the model prices might differ significantly from one procedure to the other. In particular the length of the time series window to estimate η has a significant influence on a wide range of exotic option prices.

2 Calibration of the Heston model

2.1 The Heston's stochastic volatility model

In [5], 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 [8]). The stock price process follows the well-known Black-Scholes stochastic differential equation:

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

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

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

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 [9]). 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.

2.2 Pricing of vanilla Options using Characteristic Functions

The vanilla option prices are computed by using the Carr-Madan formula (see [4]):

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

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 [3] or [7]):

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

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 [4]). 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 [6]). 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 to determine some of the model parameters. In particular, both the parameters v_0 and η can be fixed beforehand from the spot value and the historical time series of the volatility index. 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 the empirical VIX index:

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

considering a time series window T^{VIX} equal to six months, three years or five years.

The model calibration is done on the whole set of quoted vanilla options¹ on the S&P500 index, which includes a minimum of 437 options and a maximum of 1143 options, the average number amounting to 743.7423. The numerical study uses the S&P500 data together with the VIX quotes. For other underlying stocks, one can compute a volatility index out of the option quotes (see [1] for details on the VIX computation).

2.4 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 to the 30th of September 2008 for the full calibration and the reduced calibration with a parameter η inferred from three different windows of the VIX time series, i.e for a period T^{VIX} equal to six months, three years and five years.

From Figure 1, it is clear that the VIX time series window has a significant influence on the quality of the fit of market prices. In particular we can clearly distinguish two periods, the first one ranging from February 2006 until July 2007 and the second one from August 2007 until September 2008. For the first period, the wider the slot for the computation of the long run variance η , the better the fit of market data. In particular T^{VIX} equal to five years leads to a pretty good fit of the option surface and under that particular setting the loss of precision due to the reduced calibration set is not really significant with respect to the gain of computation time (see Table 2). On the other hand, for the second period of time, the best fit is obtained by considering a narrow window, i.e. T^{VIX} equal to six months.

¹i.e. for both call and put options on all the quoted strikes and times to maturity.




Figure 1: 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 of the two parameters which are inferred beforehand from the VIX time series, i.e. v_0 and η . The square of the spot VIX appears to be close to the initial variance parameter obtained by calibrating the whole parameter set (see Figure 2) and both exhibit a similar trend. On the other hand we note that the long run variance η computed on the basis of time series turns out to be significantly different from this resulting from the whole calibration. In particular, the time series estimate and the calibrated parameter exhibit a quite different trend over time : the time series estimate turns out to be much smoother and typically lower than the calibrated parameter. Furthermore the calibrated long run variance significantly varies from one trading day to the other from July 2007 on (i.e. for the second period). We particularly note that its day-to-day variation appears to reach some exceptional level of more than 15 percents during the time period considered which is quite inconsistent with its intrinsic nature. For the first period, the best fit is obtained by the backward long run variance computed for a period T^{VIX} of five years 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. On the other hand, for the second period the six months window leads to the smallest difference between the calibrated and VIX implied estimate of the long run variance.

The transition between the two periods turns out to coincide with the beginning of the credit crunch. From Figure 2, it is clear that the market has adjusted the long run variance η at the beginning of the credit crisis. Indeed the calibrated η slightly oscillates around 0.04 until July 2007 before exhibiting a sharp increase during July and August 2007. From then on the average long run variance reaches an higher level (around 0.08) and the daily calibrated value of η oscillates around this level until the end of September 2008. The replacement in the VIX time series window of the most distant value by the current one leads to a long run variance estimate which does not reflect immediately the switch of the long run variance regime observed in the market, especially for a wide time series window. Indeed, the wider the time series window, the longer the delay between the switch of the long run volatility regime directly observed in the market and the switch resulting from the VIX time series. This can be directly seen from the curve corresponding to the three years window where the transition period between the low and high long run variance regime lasts more

than one year. Moreover, the average of the market long run variance has to stay at the same level during a period of time comparable with the length of the time series window in order to totally capture the market long run variance trend in the time series estimate of η . Actually, for the time period considered, only the six months time series window turns out to capture the change in the long run volatility observed in the market.

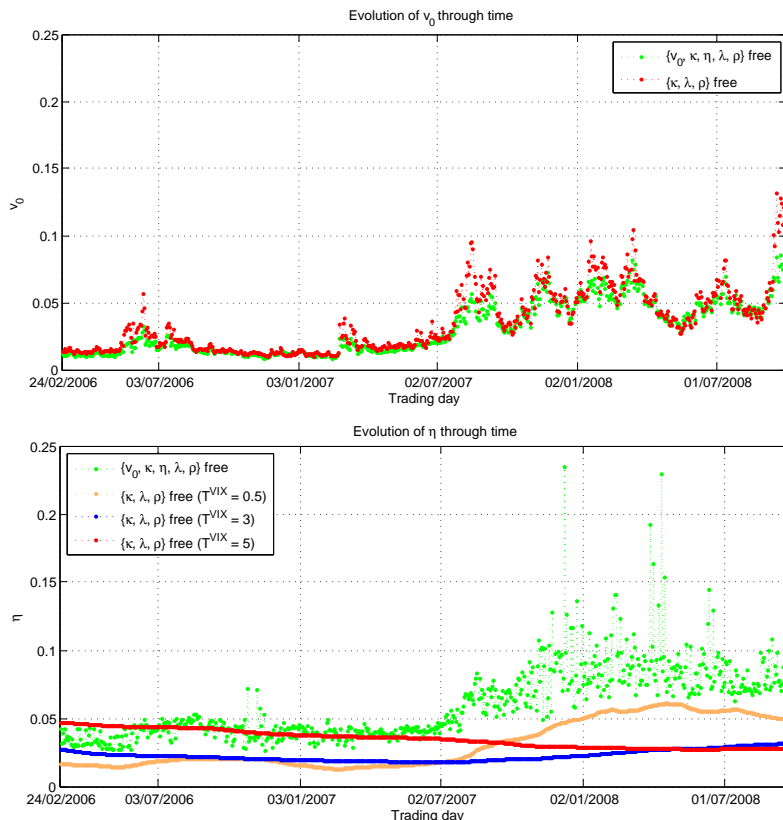


Figure 2: Evolution of the parameter v_0 (upper) and η (lower) through time for the different calibration procedures.

The matching of the spot value and the long run variance of the volatility index results in some adjustment of the other parameters calibrated on the implied volatility surface. Figure 3 shows the evolution through time of these parameters, i.e. λ , κ and ρ under both the full and reduced settings. As expected, until July 2007, the three parameters follow the same trend under the full calibration and the reduced calibration for a slot T^{VIX} equal to five years since then both v_0 and η barely differ from one calibration setting to the other. Moreover a time series long run variance significantly lower than the optimal one results in a significant adjustment of the three free parameters. In particular, the mean reversion rate κ is usually set to an higher value whereas the correlation ρ is usually set to a lower value than the optimal one. We also observe that the parameter ρ is from

time to time touching its boundary value of -1 (see also Table 3).

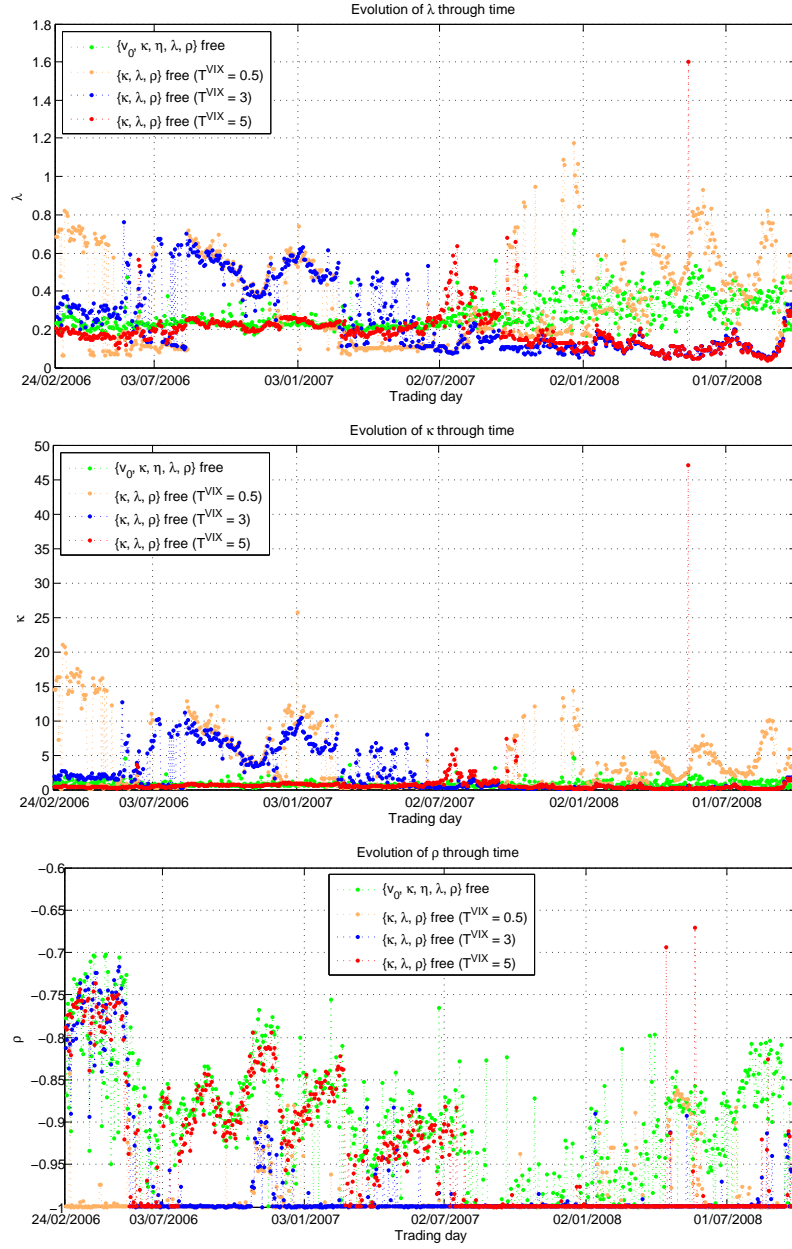


Figure 3: Evolution of the parameters λ (upper), κ (center) and ρ (lower) through time for the different calibration procedures.

The goodness of fit of the full and reduced calibrations is shown on Figure 4 and Figure 5 for the 24th of February 2006 and the 24th of September 2008, respectively². The model vanilla prices obtained with and without fixing beforehand the parameters v_0 and η barely differ from each others when the time window T^{VIX} is adequately chosen.

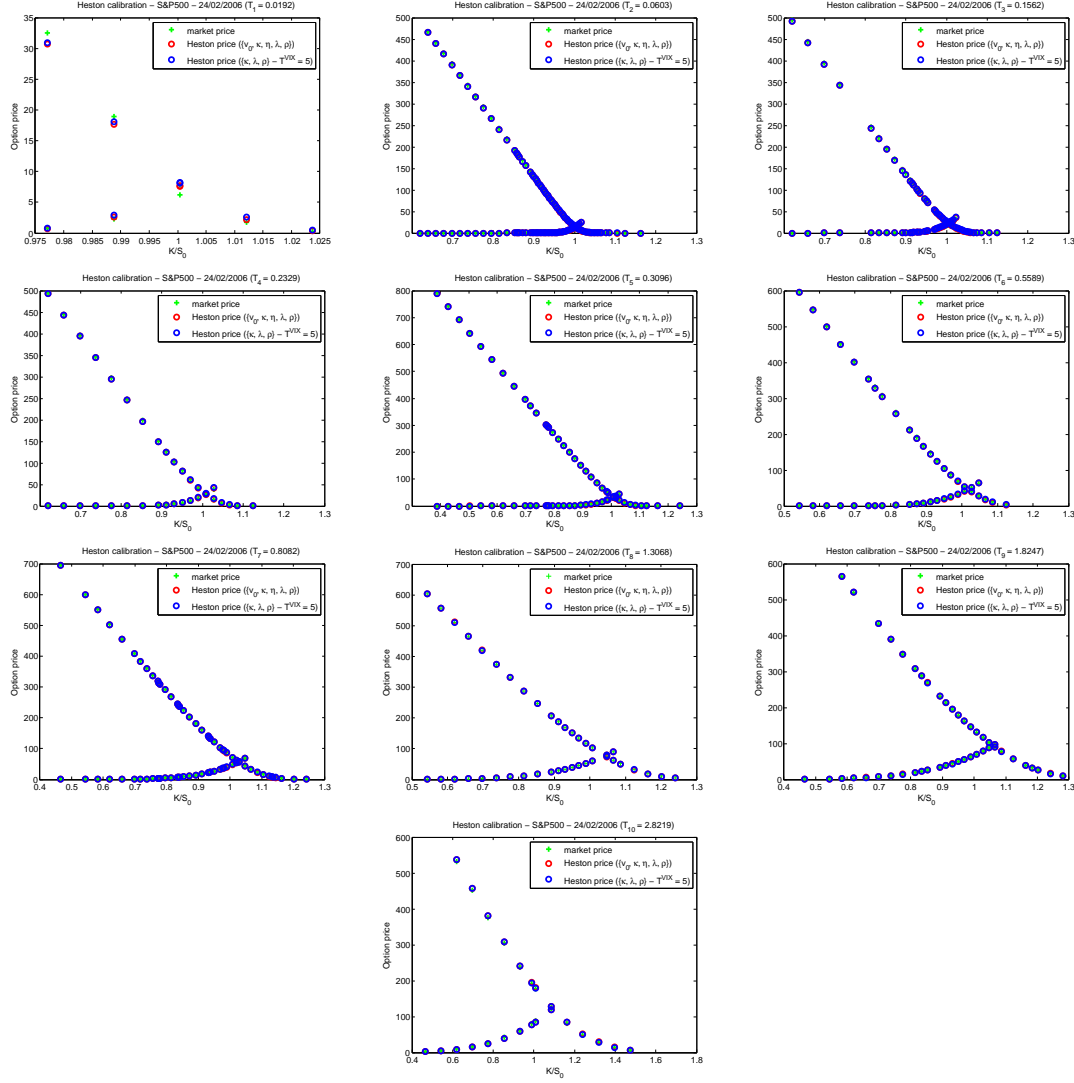


Figure 4: calibration of the Heston model - 24/02/2006

²The 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.

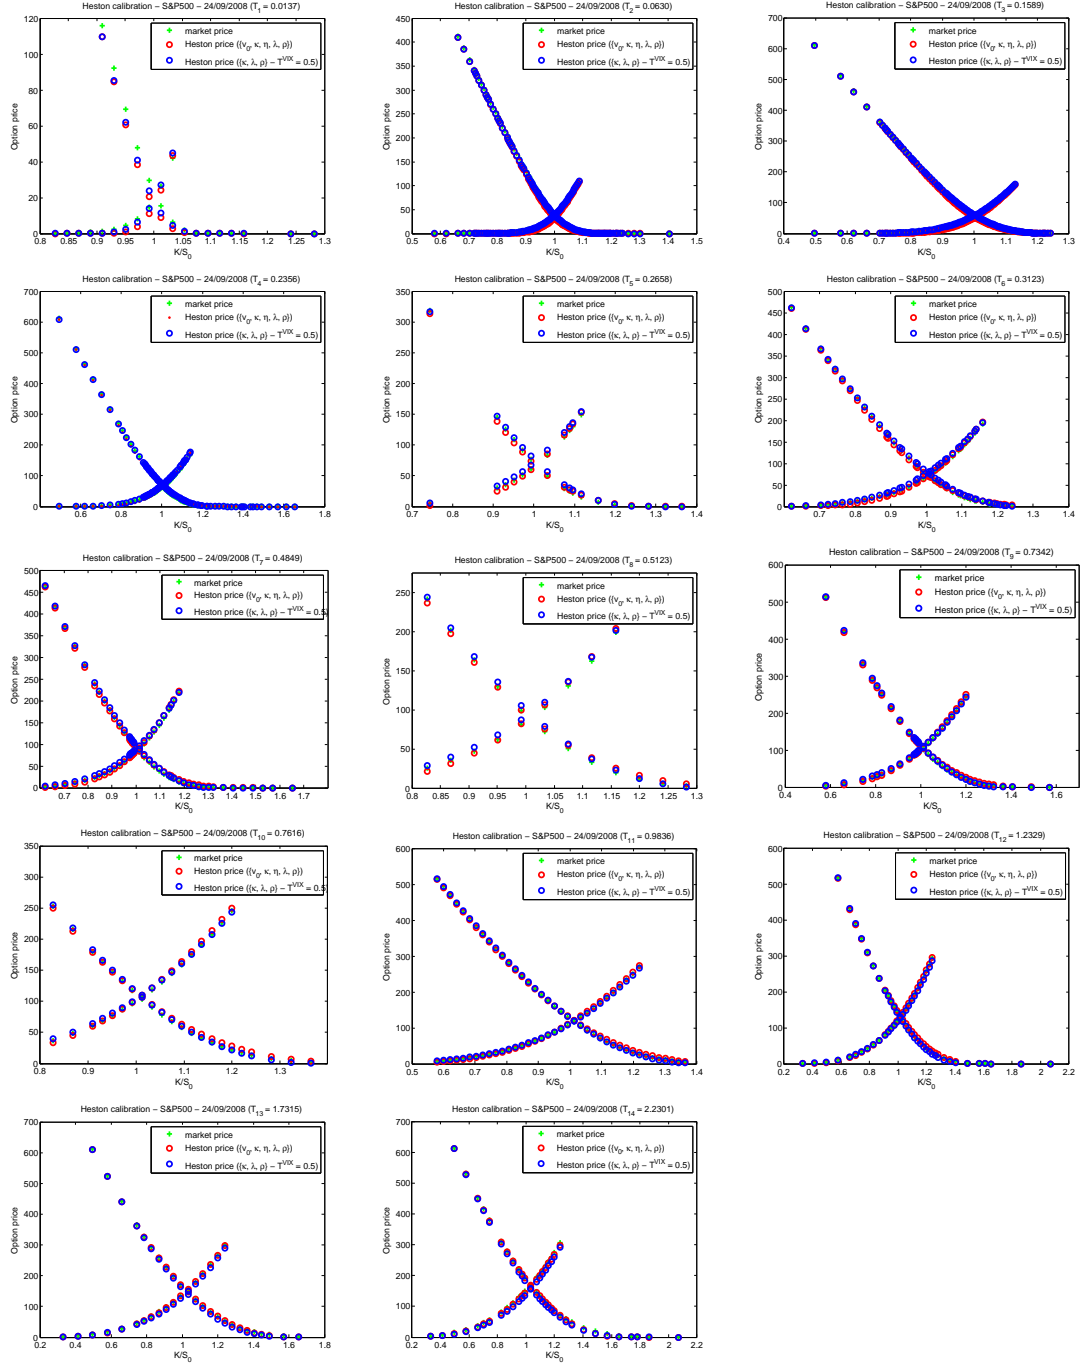


Figure 5: calibration of the Heston model - 24/09/2008

Figure 12 and Figure 13 show the histogram of the daily variation of the model parameters for the full calibration and the reduced calibration (see appendix A). It is particularly obvious that the day-to-day variation of η is lower when this parameter is directly inferred from the VIX index than when it is calibrated on the S&P500 option prices. In particular the average daily variation can be reduced by 99.70 percents (see Table 1) for a five years window. Moreover, Table 1 indicates that, generally speaking, the optimal window is obtained by considering a time slot of five years, although the average precision relative to the six months window is of the same range. Furthermore, the calibration on the reduced set $\{\kappa, \lambda, \rho\}$ with $T^{\text{VIX}} = 5$ leads to a significant decrease of the day-to-day variation of λ and ρ , and to a smaller extent of the daily change of κ , and consequently to more stable parameters.

	$\{v_0, \kappa, \eta, \lambda, \rho\}$	$\{\kappa, \lambda, \rho\} (T^{\text{VIX}} = 0.5)$	$\{\kappa, \lambda, \rho\} (T^{\text{VIX}} = 3)$	$\{\kappa, \lambda, \rho\} (T^{\text{VIX}} = 5)$
average precision	1.9729	4.5405	5.1539	4.3239
average Δv_0	0.002786	0.004157	0.004157	0.004157
average $\Delta \eta$	0.01094	0.000136	3.85E-05	3.31E-05
average $\Delta \lambda$	0.05135	0.077394	0.048333	0.027986
average $\Delta \kappa$	0.36694	1.3696	0.71801	0.31591
average $\Delta \rho$	0.041399	0.007477	0.010409	0.015182

Table 1: Average RMSE and average day-to-day variation of the model parameters for the different calibration sets.

2.5 Calibration with respect to different objective functions

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

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

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

and

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

The value of the different objective functions for the full and reduced calibration procedures is given in Table 2 for three particular trading days: the 24th of February 2006 (beginning of the first period), the 18th of July 2007 (end of the first period) and the 24th of September 2008 (end of the second period). Fixing the spot variance and the long run variance beforehand leads to a distance between the model and market vanilla prices pretty close to the optimal distance for each objective function for a well chosen time series window T^{VIX} . The time series window leading to the lowest

value of the objective function is given in bold type. For the first two quoting dates the optimal window is obtained by considering a period of 5 years whereas for the last quoting date, the lowest value of the functional is obtained with $T^{\text{VIX}} = 0.5$. Note that it might turn out that the value of the objective function under the reduced calibration setting is lower than the value obtained by considering the full calibration procedure, which means that the full optimization technique reaches a local minimum. Moreover the computation time of the calibration procedure is reduced by more than 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.

functional	RMSE	APE	AAE	ARPE
full calibration - 24/02/2006	1.51008	0.01090	1.07454	0.20763
reduced calibration ($T^{\text{VIX}} = 0.5$) - 24/02/2006	5.14267	0.02984	2.94191	0.28295
reduced calibration ($T^{\text{VIX}} = 3$) - 24/02/2006	1.75493	0.01223	1.20507	0.20478
reduced calibration ($T^{\text{VIX}} = 5$) - 24/02/2006	1.61958	0.01210	1.19238	0.22230
full calibration - 18/07/2007	1.93539	0.01256	1.19435	0.21991
reduced calibration ($T^{\text{VIX}} = 0.5$) - 18/07/2007	8.15286	0.06056	5.75795	0.33036
reduced calibration ($T^{\text{VIX}} = 3$) - 18/07/2007	8.29926	0.06188	5.88320	0.38100
reduced calibration ($T^{\text{VIX}} = 5$) - 18/07/2007	3.03135	0.01458	1.38618	0.23617
full calibration - 24/09/2008	3.77414	0.03298	3.04660	0.12662
reduced calibration ($T^{\text{VIX}} = 0.5$) - 24/09/2008	4.17867	0.02791	2.57776	0.19769
reduced calibration ($T^{\text{VIX}} = 3$) - 24/09/2008	6.18813	0.04758	4.39502	0.32388
reduced calibration ($T^{\text{VIX}} = 5$) - 24/09/2008	6.55433	0.05036	4.64949	0.33456
cpu	RMSE	APE	AAE	ARPE
full calibration	123.96875	81.828125	93.421875	130.484375
reduced calibration	20.0625	20.984375	22.984375	26.375

Table 2: Objective functions (upper) and Computation time in seconds (below) for the full and reduced calibration procedures (S&P500).

Table 3 shows the optimal parameter set under the RMSE objective function for the three particular trading days mentioned in Table 2.

	24/02/2006	18/07/2007	24/09/2008
full calibration	{0.0115, 0.6265, 0.0422, 0.2299, -0.7779}	{0.0225, 0.5780, 0.0518, 0.2446, -0.8872}	{0.0746, 0.33690, 0.05510, 0.1927 - 1.0000}
reduced calibration ($T^{\text{VIX}} = 0.5$)	{0.0131, 14.4819, 0.0162, 0.6840, -1.0000}	{0.0232, 0.1788, 0.0189, 0.0822, -0.9998}	{0.1077, 2.0829, 0.0521, 0.4657, -1.0000}
reduced calibration ($T^{\text{VIX}} = 3$)	{0.0131, 1.8388, 0.0271, 0.3155, -0.8129}	{0.0232, 0.1590, 0.0178, 0.0752, -1.0000}	{0.1077, 1.3660, 0.0324, 0.2975, -0.9720}
reduced calibration ($T^{\text{VIX}} = 5$)	{0.0131, 0.4423, 0.0467, 0.2033, -0.7883}	{0.0232, 1.3256, 0.0342, 0.3012, -0.8829}	{0.1077, 1.2888, 0.0285, 0.2710, -0.9173}

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 [2] as the difference in the value of the calibrated parameters arising from the different specifications of the objective function. Figures 6 and 7 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 two calibration procedures. Note that v_0 and η do not vary with the functional for the reduced calibration since they are directly inferred from the VIX spot and historical prices. Hence the reduced calibration procedure allows to eliminate the calibration risk arising from the determination of these two parameters. Moreover, for each other model parameter, the optimal values obtained by considering one or another functional are usually closer to each other in the reduced calibration setting than in the full calibration setting, leading to a lower calibration risk.

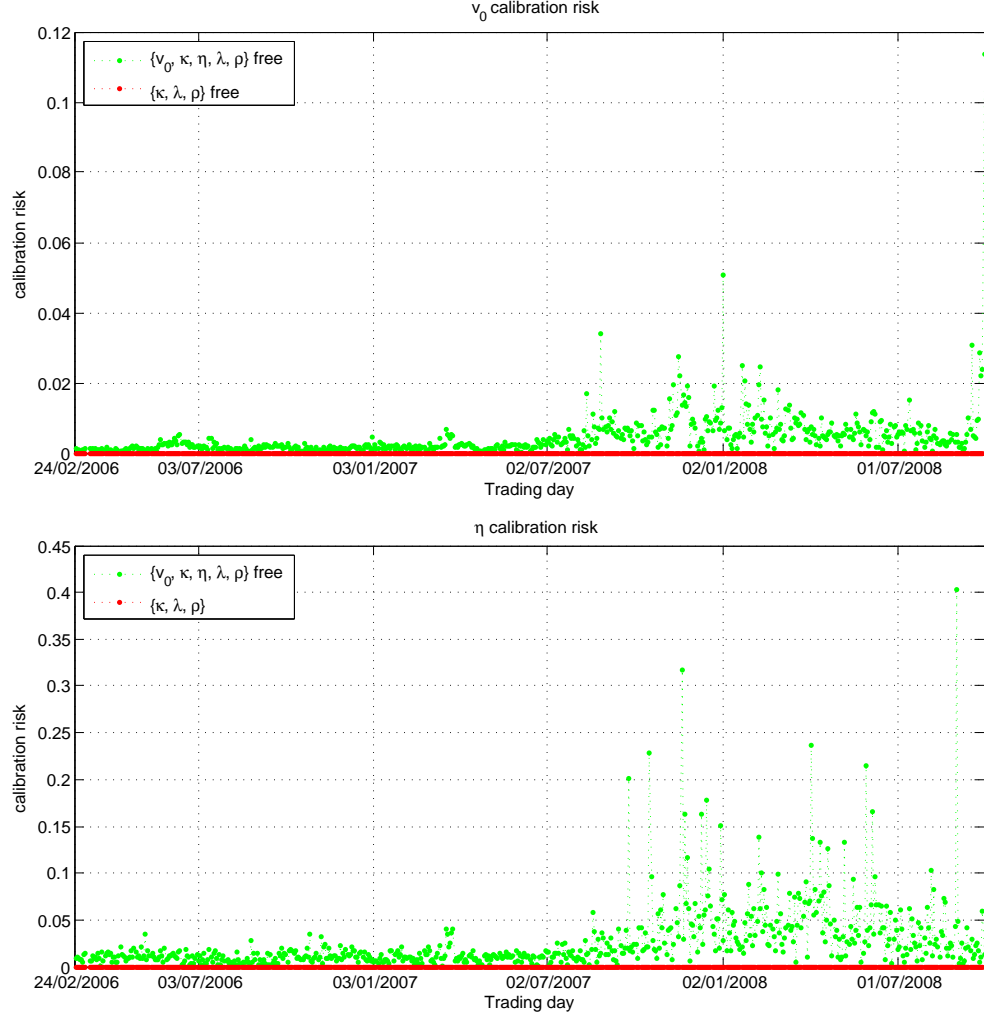


Figure 6: Evolution of the calibration risk for the parameter v_0 (upper) and η (lower) through time for the different calibration procedures.

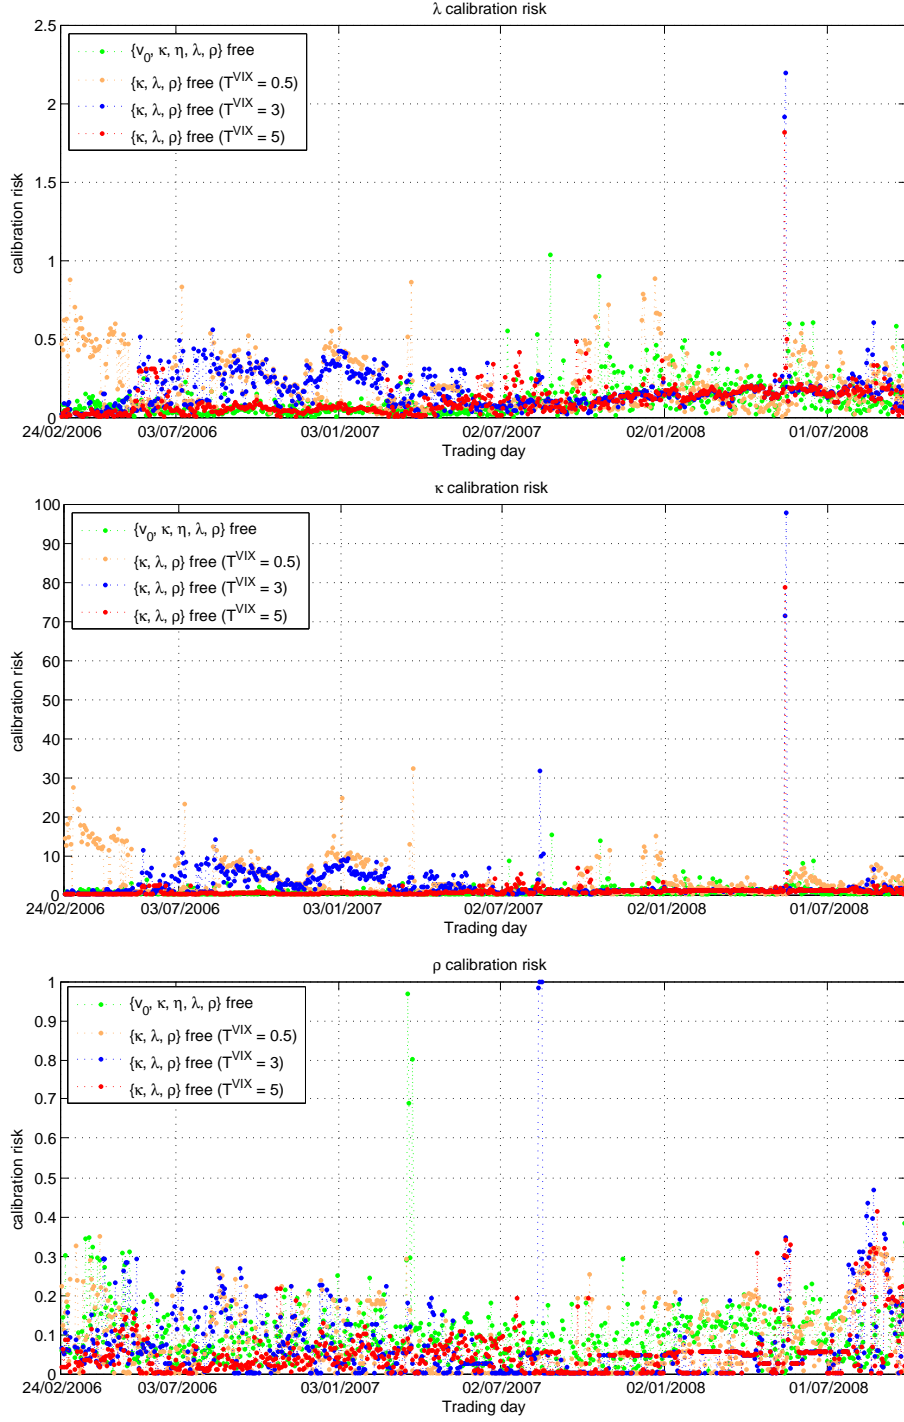


Figure 7: Evolution of the calibration risk for the parameter λ (upper), κ (center) and ρ (lower) through time for the different calibration procedures.

3 Pricing of Exotic Options

This section features a comparison of the price of different exotic options under the two calibration procedures for the three particular quoting dates defined in Section 2.5, considering the RMSE objective function. As it can be seen from Figure 4 and Figure 5, both the full and reduced calibration procedures leads to a pretty good fit of the whole set of liquid options for a well chosen window T^{VIX} . Hence, we can 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

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

and

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

respectively.

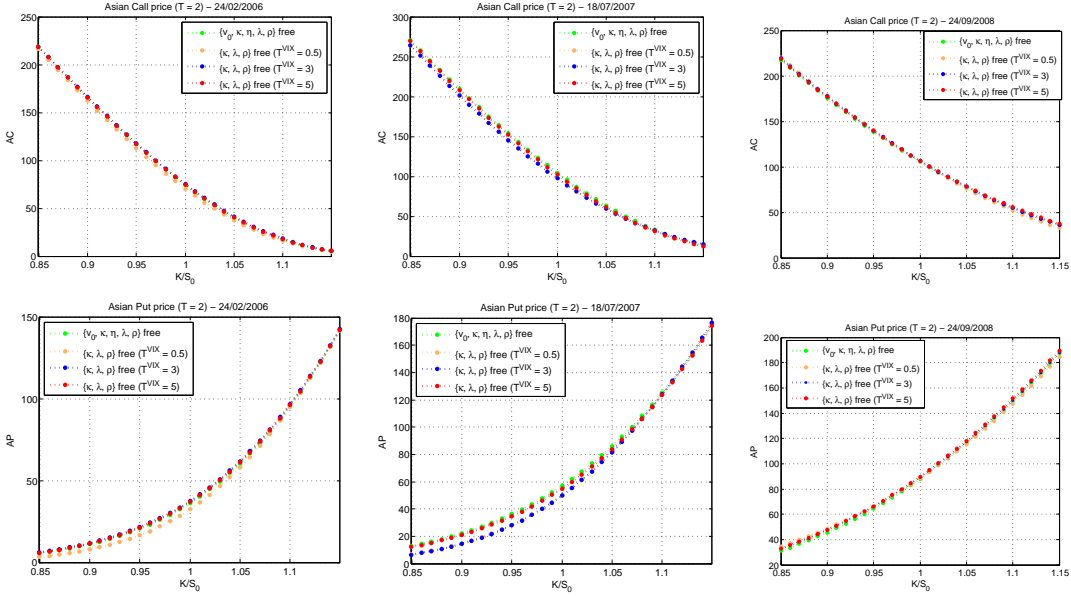


Figure 8: Asian call (upper) and put (lower) option prices

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 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 denote 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	24/02/2006		18/07/2007		24/09/2008	
	LC	LP	LC	LP	LC	LP
full	241.1438	143.9279	334.5474	201.2216	322.6543	263.8103
reduced - $T^{\text{VIX}} = 0.5$	213.4778	122.9003	289.6962	171.1820	319.5498	244.3405
reduced - $T^{\text{VIX}} = 3$	241.5010	140.3158	288.6279	173.3066	306.1144	231.4793
reduced - $T^{\text{VIX}} = 5$	240.3859	142.9716	323.5411	188.9741	302.2986	234.5406

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 :

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

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

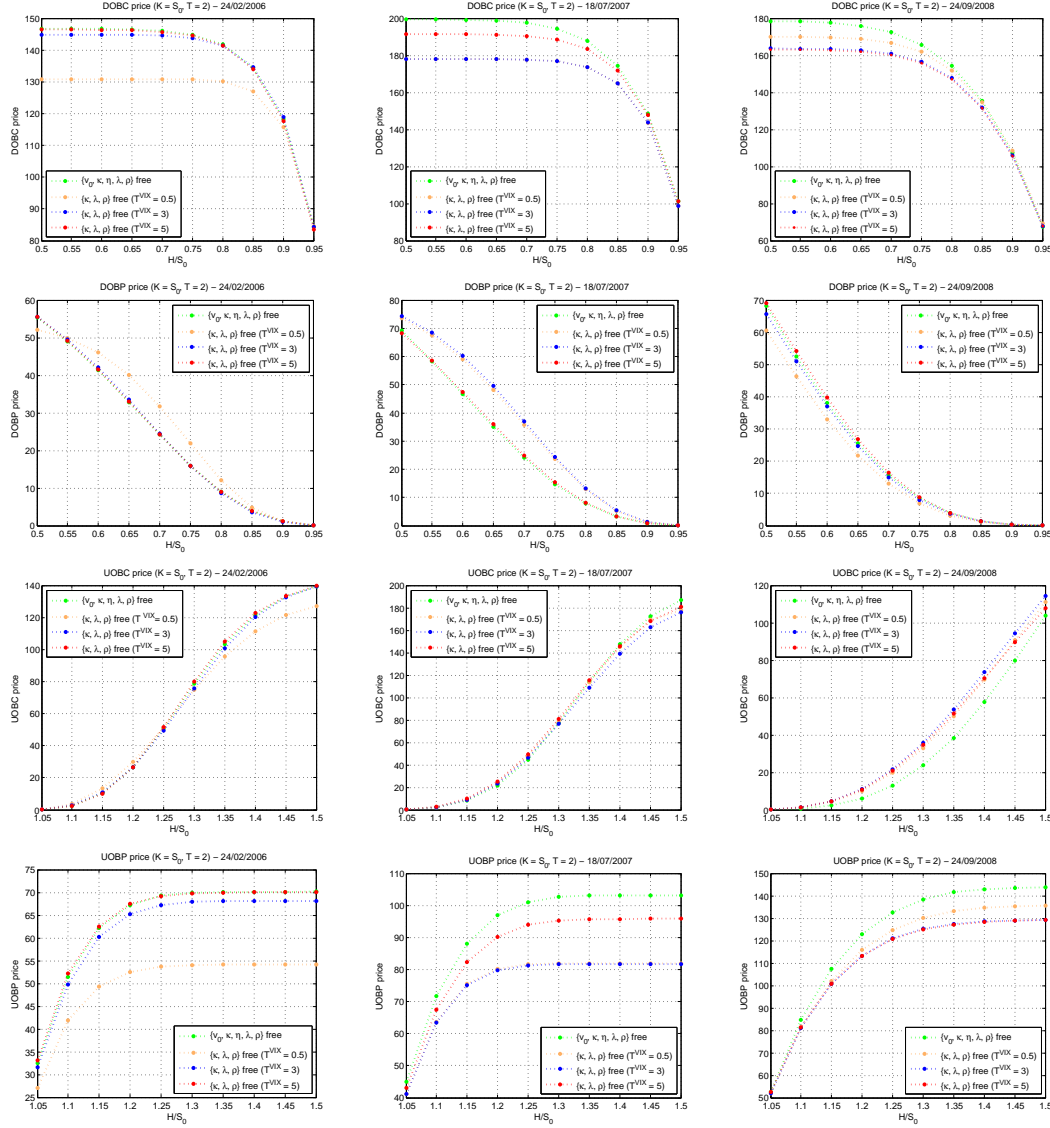


Figure 9: Down-and-out (upper) call and put and Up-and-out (lower) call and put option prices

- 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 :

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

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

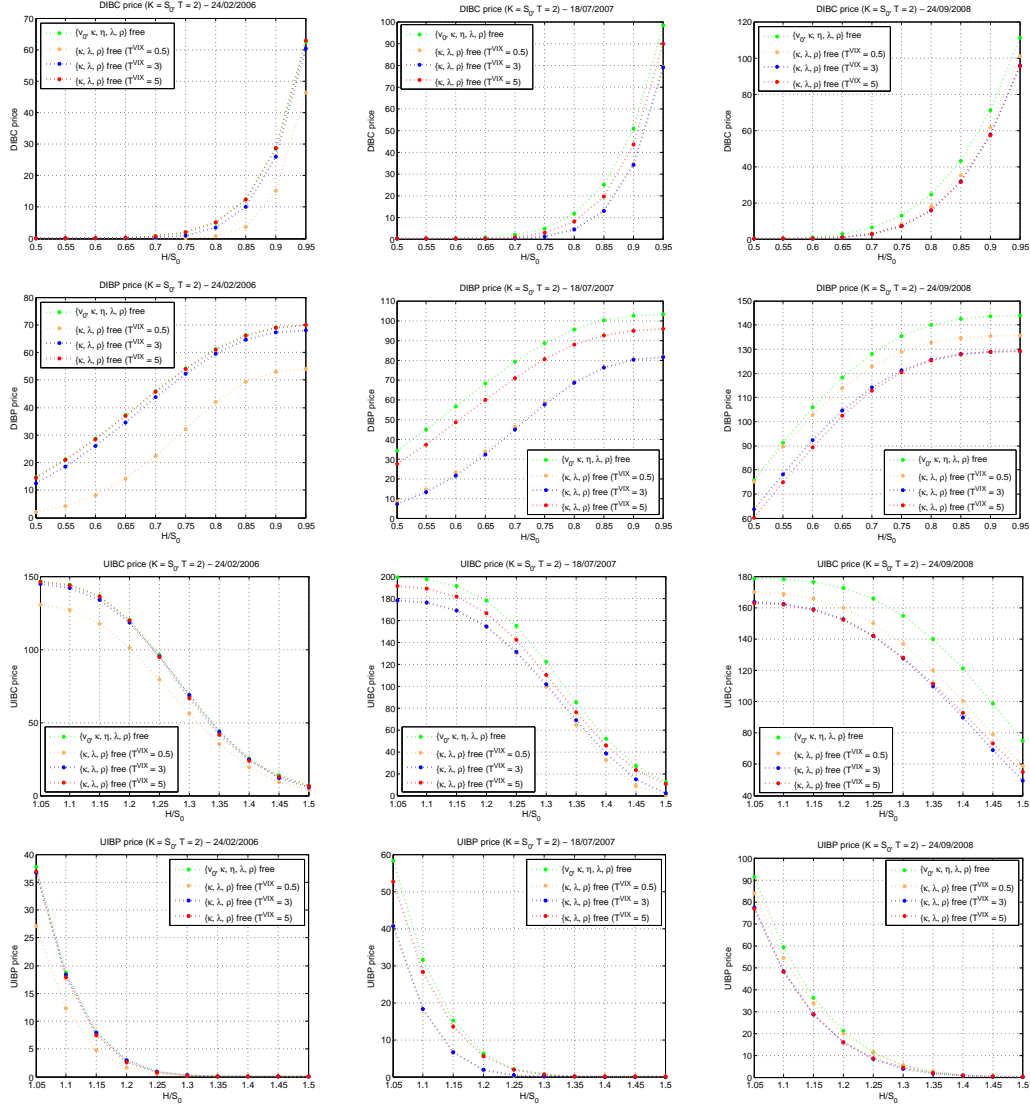


Figure 10: Down-and-in (upper) call and put and Up-and-in (lower) call and put option prices

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.

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

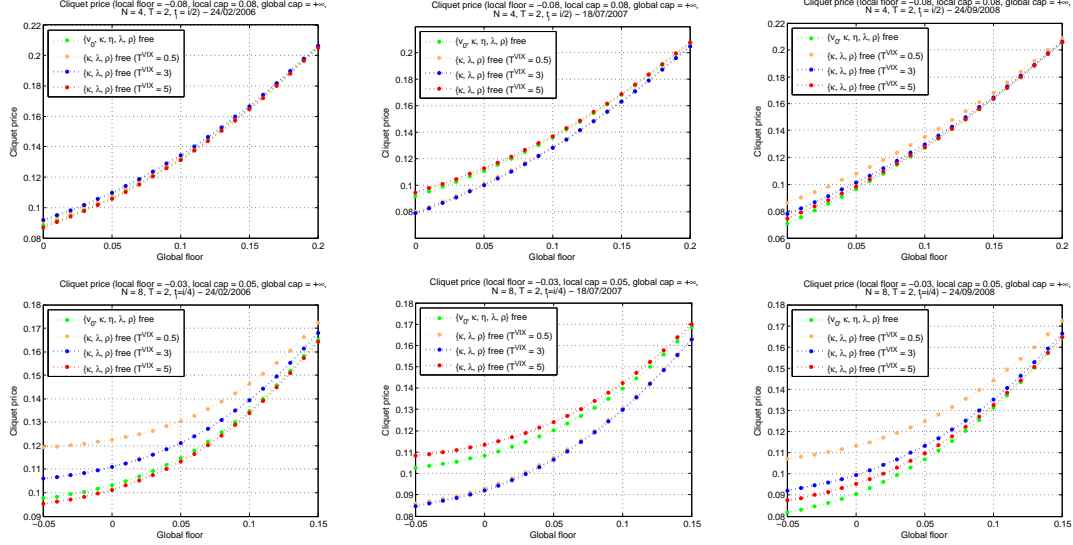


Figure 11: Cliquet option prices

As it can be seen from Figure 8 to Figure 11, except for the Asian options, the Heston price of the different exotic options turns out to be sensitive to the calibration method, and especially to the time series window. For the 24th of February 2006, the exotic prices obtained by the reduced calibration procedure with $T^{\text{VIX}} = 5$ barely differ from the prices obtained under the optimal parameter set. Moreover, the three years window leads to slightly different exotic prices whereas the 6 months window is characterized by significantly different exotic prices which might be explained by the particularly high value of the parameter κ (see Table 3). For the 18th of July 2008, each time series window leads to significantly different exotic prices under the reduced calibration setting than under the standard one. The closest prices are obtained by considering a window of five years. For the 24th of September 2008, the closest price is obtained with the six months window except for the cliquet options and the DOBP options. Hence the time series window 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.

4 Conclusion

This paper features a detailed comparison of the calibration performance of the Heston model, considering either the standard calibration on the whole parameter set $\{v_0, \kappa, \eta, \lambda, \rho\}$ or a calibration

on a reduced parameter set $\{\kappa, \lambda, \rho\}$, fixing beforehand v_0 and η from the time series of the VIX index. We have shown that both procedures lead to a fit of vanilla options of similar quality for well chosen time series window whereas the reduced procedure is characterized by both a significantly lower computation time and calibration risk which arises by considering different measures for the error between the observed market prices and the corresponding model prices. Moreover fixing beforehand the initial variance and the long run variance leads to more stable parameters of the Heston model through time.

We have then shown that the price of a wide range of exotic options (one touch barrier, lookback and cliquet options) turns out to be sensitive to the calibration procedure, and in particular to the length of the time series window.

As a general conclusion, we have indicated that even within one model, model risk is present. More precisely, different objective functions, the full or the reduced calibration procedure and the length of the time series window to estimate η , all lead to different optimal parameters. This is on itself not very surprising, but what is important is that optimal parameters can significantly differ leading to different exotic prices. We are therefore faced, even within one model, with choosing a route to price the exotic structures out of an existing vanilla option surface. We have described several alternatives, with their pros and cons. The analysis does not put one particular method to the fore. We have furthermore noticed that sometimes the Heston model is not really able to fit the option surface with a realistic set of optimal parameters: we have observed very high values of κ and the correlation is often reaching its lower boundary value.

A Day-to-day variation of the model parameters

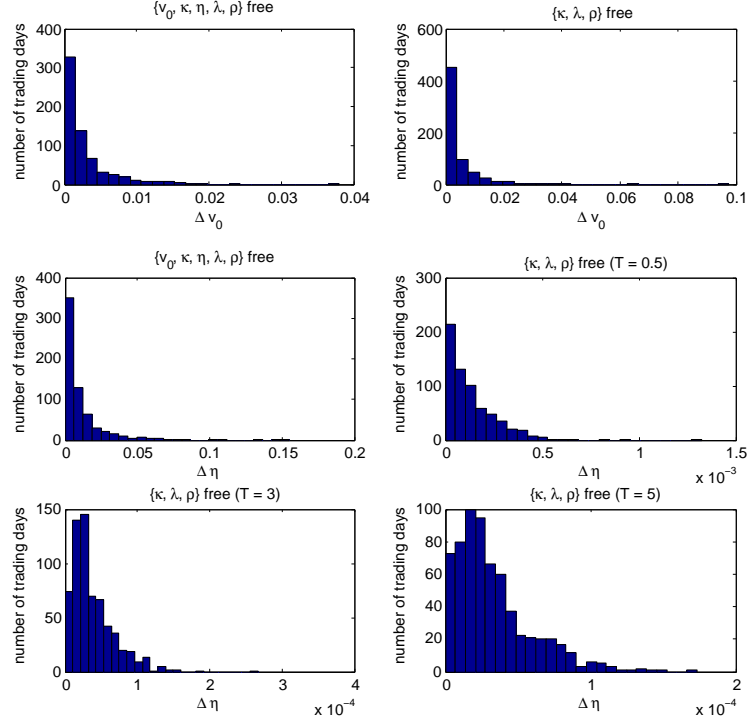


Figure 12: Histogram of the day-to-day variation of v_0 (upper) and η (lower) for the different calibration procedures.

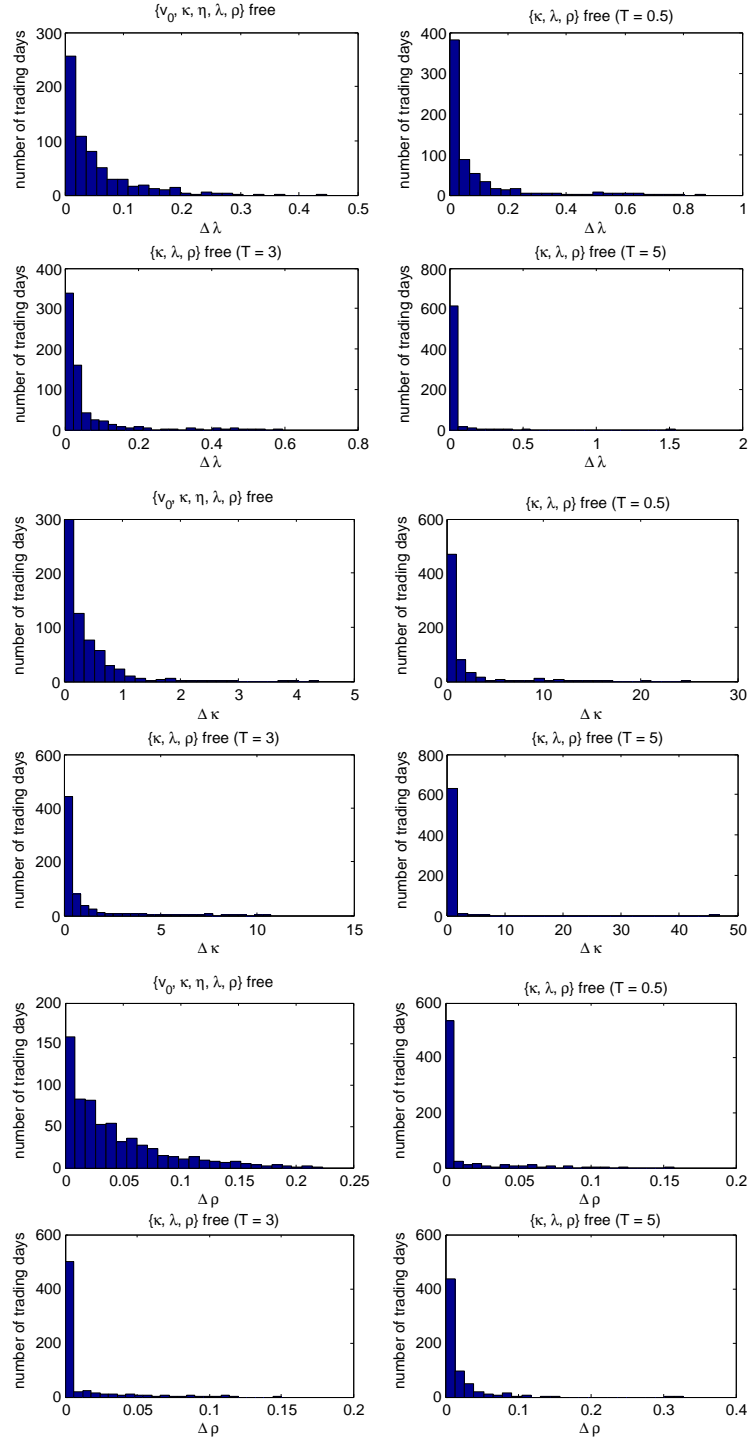


Figure 13: Histogram of the day-to-day variation of λ (upper), κ (center) and ρ (lower) for the different calibration procedures.

References

- [1] Chicago Board Options Exchange (2003) VIX: CBOE volatility index. *Working paper, Chicago*.
- [2] Detlefsen, K. and Hardle W. K. (2007) Calibration Risk for Exotic Options. *The Journal of Derivatives* **14**, 47-63 .
- [3] Albrecher, H., Mayer, P., Schoutens, W. and Tistaert, J. (2007) The little Heston trap. *Wilmott Magazine*, 83-92.
- [4] Carr, P. and Madan, D.B. (1998) Option valuation using the fast Fourier transform. *Journal of Computational Finance* **2**, 61-73.
- [5] Heston, S. (1993) A closed-form solution for options with stochastic volatility. *Review of Financial Studies* **6**, 327-243.
- [6] Mikhailov, S. and Nögel, U. (2003) Heston's Stochastic Volatility Model Implementation, Calibration and Some Extensions. *Wilmott magazine*, 74-79.
- [7] Schoutens, W., Simons, E. and Tistaert, J. (2004) A Perfect calibration ! Now what ? *Wilmott Magazine*, March 2004
- [8] Whaley, R. E. (2009) Understanding the VIX. *The Journal of Portfolio Management*, **35**, 98-105.
- [9] Wilmott, P. (2006) *Paul Wilmott on Quantitative Finance*. Wiley.