# Hedging Options in the Presence of Microstructural Noise

David Horn[‡] Eva Schneider[‡] Grigory Vilkov[‡]

This version: September 26, 2007

**Abstract**

In order to use an option pricing model for dynamic hedging an investor will have to calibrate it to a cross-section of option prices. Microstructural noise in option prices results in a set of indistinguishable parametrizations which may give rise to different hedging errors. In our simulation study for the Heston (1993) model, we identify the parameters most important for hedging and show which set of strikes and time to maturity is relevant for the identification of certain parameters. In our empirical study we show that different but indistinguishable parametrizations w.r.t. prices may induce large differences in hedging performance.

**Keywords:** parameter risk, hedging, microstructural noise, stochastic volatility

**JEL:** G11, G12

[‡]Finance Department, Goethe University, Mertonstr. 17-21/Uni-Pf 77, D-60054 Frankfurt am Main, Germany. E-mail: horn@finance.uni-frankfurt.de, schneider@finance.uni-frankfurt.de, vilkov@finance.uni-frankfurt.de

# 1 Introduction

An investor setting up a dynamic hedging strategy will first have to identify a stochastic process for the underlying. Having chosen the appropriate model, the investor will have to calibrate this model to a cross-section of empirical option prices. With the calibrated model the investor can then proceed to calculate the weights of the replicating portfolio. The identification and calibration of the model are of great importance, since a wrong model will result in wrong hedging weights and larger and more volatile hedging errors. When using empirical data to identify and calibrate the model at least two problems arise. First, there is only limited data available. Usually only for very liquid index options we have a broad range of option prices available. For other types of underlyings, the limited number of available option prices at a particular point in time can make the identification and calibration very difficult. Second, the prices we observe on option markets are noisy. When we speak of noise in option prices we mean that we believe in one true (frictionless) option price but this price cannot be observed due to the bid- and ask-spread. Several reasons for noisy option prices have been suggested: bid-ask spread in the underlying, different interest rates for borrowing and lending, non-synchronous observations or rounding to the nearest tick size (see Hentschel (2003)).

The existence of a bid-ask spread complicates the identification and calibration of a model in the following ways. Identification becomes more difficult since two structurally different models that yield option prices, which differ by less than the bid-ask spread, cannot be distinguished anymore (see Dennis and Mayhew (2004)). The same problem applies to the calibration. Even if we knew the true model, two different parametrizations could not be distinguished from another if they produce option prices differing by less than the bid-ask spread. Not only can two parametrizations result in different option prices, but they can also result in different hedging weights in a hedge portfolio. They will thus also yield different hedge errors. In the present paper we want to focus on this second issue. Hence, this paper tries to answer the following research questions: What are the most sensitive parameters for calibration and hedging? Which parameters are the most important to identify a model correctly? Which subset of prices (with respect to strike price and time to maturity) is especially relevant for the identification of certain parameters? Can we rely on cross-sectional information or do we have to use time series information? And finally: Can we verify the theoretical results for the previous questions empirically?

In order to answer these questions we will perform both a simulation analysis in the stochastic volatility framework of Heston (1993) and an empirical study. In the simulation we will generate European option prices in the Heston (1993) model under a fixed parametrization which represents the true model. Then we will identify the parameter range leading to an indistinguishable cross-sectional fit for an investor who calibrates the Heston model to the generated dataset. This means that we ignore model risk and solely focus on parameter risk. In the empirical part of this paper we will calibrate the model of Heston (1993) to option prices of several single stocks and the S&P 100 and identify the indistinguishable parametrizations for the observed bid-ask spread. We find that a parametrization with a good cross-sectional fit does not necessarily have to have good hedging properties and that there are economically important differences in the hedging performance of these indistinguishable parametrizations.

We contribute to the existing literature in the following ways: First, we identify the parameters which prices and greeks are most sensitive to. Second, we show which information (number of strikes, number of maturities) is important for the correct identification of certain parameters. Third, we show which parameters are the most important to identify correctly for the purpose of hedging. Fourth, we measure parameter risk using empirical option prices.

## 2 Literature Review

Besides the papers looking at the sensitivities of option prices to model parameters, there are only few which consider the impact of parameter risk. These papers mainly focus on volatility mis-estimation. Hentschel (2004) provides a thorough analysis of the error in implied volatility estimation induced by measurement errors in the input variables and by truncation, i.e. when low option prices are missing in the calibration. He works in the model of Black and Scholes (1973) and shows that the resulting error is the larger the more the options are OTM or ITM and that due to the truncation error, the confidence interval may not even contain the correct volatility. Figlewski (2004) estimates the volatility by using a simple average of squared returns or via an exponentially-weighted moving average. He then performs a simulation analysis of estimation errors for the computation of the value at risk. For the data generating process, he assumes the dynamics given in Black and Scholes (1973), Heston (1993), Bakshi, Cao, and Chen (1997) and Eraker, Johannes, and Polson (2003). The standard estimation technique of the value at risk is shown to lead to significant errors in the estimation of tail probabilities.

In contrast to these papers, we estimate parameters in the true model and do not take volatility as a simple estimate from time series but calibrate the model to option prices. As our base model, we use the empirically well supported stochastic volatility model of Heston (1993).

Another strand of literature related to our work is concerned with model mis-specification. Here, the mis-specification does not relate to the model parameters but to the whole model class. Dennis and Mayhew (2004) study the impact of noise in option prices on parameter or model estimation and try to answer the question if, given the noise in option prices, one can distinguish between different models. Schoutens, Simons, and Tistaert (2003) calibrate several models to a cross section of European options and compare the resulting prices of exotic options. An and Suo (2003) test several models by looking at the hedging errors of strategies for exotic derivatives.

The closest paper to ours is the paper by He, Kennedy, Coleman, Forsyth, Li, and Vetzal (2006). The authors assume a jump-diffusion model with a local volatility function as the true data-generating process. When calibrating the model, jump parameters are harder to estimate than the local volatility function since a large surface of parameters yields sufficiently small estimation errors. However, the hedging performance is not largely affected by the estimation problems. Both dynamic variance-minimizing hedges and semi-static hedges perform well for the set of instrument options. In our paper, we focus on mis-estimation of stochastic volatility parameters (not jump parameters) and in addition to a simulation study, we perform an empirical analysis.

# 3 Motivation

## 3.1 Parameter Risk in the Heston Model

As stated in the introduction, we want to focus on parameter risk. We therefore assume that the structural type of the model is known. Also, to keep the analysis as simple as possible, we will examine a stochastic volatility model and exclude additional sources of risk, such as multiple volatility components or jumps in the stock price or in volatility. We choose stochastic volatility since this seems to be the most important improvement over the model of Black and Scholes (1973). In particular we will use the model of Heston (1993). This model assumes the following risk-neutral dynamics for the underlying stock $S$ and its local variance $V$:

$$\begin{aligned} dS_t &= rS_t dt + \sqrt{V_t} S_t dW_t^S \\ dV_t &= \kappa(\theta - V_t)dt + \sigma_v \sqrt{V_t} dW_t^V \end{aligned}$$

where $r$ is the riskless rate of return, $\kappa$ is the speed of mean reversion, $\theta$ is the long run mean of local variance $V$, $\sigma_v$ is the 'volatility of variance' parameter and the correlation between the two Wiener processes is described by the parameter $\rho\, dt = E[dW^S dW^V]$. The Heston model is essentially an incomplete market model. A closed form solution for option prices can be derived via Fourier inversion (see Heston (1993)). For hedging purposes all claims can be replicated with a stock, the money market account and one instrument option.

At date $t$, the interest rate $r$ and the stock price $S$ can be observed. The speed of mean-reversion $\kappa$, the long-run mean of variance $\theta$, and the volatility of variance $\sigma_v$ all have to be estimated from a cross section of option prices. The local level of variance $V$ plays a special role since it is theoretically assumed to be a state variable but for practical purposes is usually estimated from empirical option prices like a model parameter.

The starting point for our analysis is the fact that several parameter combinations can provide a similar fit to a given cross section of option prices. Microstructural noise in option prices can then make two parametrizations virtually indistinguishable. If the focus of the investor is the pricing of plain vanilla options, each of these indistinguishable parametrizations might be acceptable. However, when the pricing of exotic options or hedging is considered, different parametrizations might lead to significantly different outcomes. That is, even if the investor knows the true model he will not only incur a hedging error due to discrete trading, but also an additional hedging error due to the fact that he picks a parametrization which is not the true one. This of course raises the question of what the true parametrization is in the presence of a bid-ask spread in option prices. For the rest of the paper we will assume that the true parametrization is the one which would persist if there were no bid-ask-spread, i.e. the parametrization in a world without noise in option prices.

To get a feel for the importance of parameter risk for hedging consider a dynamic delta-vega hedge. For the case of continuous trading and noiseless option prices it is possible to devise a strategy with zero hedging error. The investor trying to hedge a short position in a call then has to solve the following problem: Find the quantities $w_s$, $w_m$, $w_i$ of the stock $S$, the money market account $M$, and the instrument option $C^I$ subject to the constraints

$$\begin{aligned} C^T &= w_s \cdot S + w_m \cdot 1 + w_i C^I \\ C^T_S &= w_s \cdot 1 + w_m \cdot 0 + w_i C^I_S \\ C^T_V &= w_s \cdot 0 + w_m \cdot 0 + w_i C^I_V, \end{aligned}$$

where subscripts denote partial derivatives. The first condition makes the replicating portfolio self-financing, the second one delta-neutral and the third one vega-neutral. Note that the weights of this hedging strategy will depend on the derivatives of the instrument and the target option with respect to $S$ and $V$, which are model and parameter dependent. The correct estimation of the sensitivities will thus be crucial for the performance of the hedging strategy.

## 3.2 Parameter Sensitivity of Option Prices and Greeks

The purpose of this section is to give a first intuition about the relative importance of different parameters for pricing and hedging. In particular we will ask two questions: First, which parameters have the largest impact on option prices? Since the calibration of option pricing models is usually done by fitting observed prices, a parameter with little influence on the option price can easily be misestimated. Second, which parameters have the largest impact on the greeks of the option price? The performance of a dynamic hedge depends heavily on the use of the correct sensitivities.

To answer these two questions we plot the price, delta and vega in the Heston model as a function of moneyness for different parameter levels. Differences in the reaction of prices and greeks to the change of a parameter can then be used to draw some first conclusions regarding the effect of misestimation on hedging performance. Since we expect the effects to be strongest for long-term options we analyze an option with a maturity of 0.9 years.

We start with the analysis of the effect of the parameter changes on prices and on the volatility smile. The sensitivities of the option price to its parameters are depicted in Figure 1. We see that a variation of $\kappa$ and $\sigma_v$ has little effect, while the impact of $\rho$ is a little stronger. Note that there is an inverse effect for ITM and OTM options. The parameters with the strongest influence on the option price are $\theta$ and $V_0$. Except for the case of $\rho$ the sensitivities are highest for ATM options. The plots of the sensitivities of the smile give essentially the same findings. However, it becomes clearer that $\kappa$ and $\sigma_v$ seem to have influence only on ATM options. Again the plots show that $\theta$ and $V_0$ have a strong effect for all moneyness levels and that the influence of $\rho$ is weakest for ATM options and becomes stronger the deeper the options are ITM or OTM.

This is in line with intuition, since local variance and long-term mean of variance are the most important drivers of option prices, while the other parameters exhibit only a second-order impact. Due to its asymmetric payoff structure, an OTM option increases in value for more positive stock-variance correlation where simultaneous increases or decreases of the stock price and its local variance are more likely to occur. On the other hand, its price decreases the more negative the correlation. Overall we can say that mis-estimation is most likely for the parameters $\kappa$ and $\sigma_v$. Mis-estimation is less likely if ATM options are used. An exception here is the estimation of $\rho$ for which we would advise to use OTM or ITM options for calibration. Since ITM options often lack liquidity, OTM options seem to be most appropriate for the calibration.

Figure 3 shows the sensitivity of the Heston delta with respect to the various parameters. Again we find only little influence of $\kappa$ and $\sigma_v$. Except for $\rho$ the sensitivities are highest for OTM and ITM options, and the effect is reversed when moving from ITM to OTM options. For $\rho$, the influence is strongest for ATM options and becomes weaker the deeper the option is ITM or OTM.

The sensitivity of vega is depicted in Figure 4. Here we see that $\rho$ and $\sigma_v$ have little, $\theta$ and $V_0$ moderate influence on the Heston vega. In contrast to delta we see that $\kappa$ has a very strong effect on vega. In total we find that delta reacts most strongly to changes in $\theta$, $\rho$ and $V_0$, while vega reacts most strongly to changes in $\kappa$, $\theta$ and $V_0$.

To sum up the results of this preliminary analysis, we find that when parameters are estimated through a calibration to observed prices, a mis-estimation of $\kappa$ is likely. This can result in strong differences in vega. In contrast, although a mis-estimation of $\sigma_v$ is likely, this will not affect delta or vega severely. For $\rho$, a mis-estimation is likely when calibrating the model only to ATM options. This simple analysis already shows that if several parametrizations provide a satisfactory fit in the presence of microstructural noise, these different parametrizations can lead to very different deltas and vegas, which will in turn affect the composition of our hedging portfolio. The question arises to what extent this carries over to hedging performance.

In the following sections we will analyze the effect of microstructural noise on the hedging performance in more detail.

# 4 Design of the Study

In this section the general design of our simulation analysis is explained. An investor setting up a dynamic delta-vega hedge for a short position in a target call will first estimate the parameters of his model from a cross section of option prices. The objective of his calibration routine will be to find a parametrization $\Psi$ which fits observed option prices best. In the Heston model the investor will have to estimate the parameters vector $\Psi = (\kappa, \theta, \sigma_v, \rho, V_t)$. Several approaches to the calibration of a model are possible. For example, one could employ the minimization of the relative or absolute squared pricing error or the minimization of the squared implied volatility difference. In this paper we will assume that the investor minimizes the sum of squared absolute pricing errors as in Bakshi, Cao, and Chen (1997). This means his minimization problem is of the form

$$\min_{\Psi} \sum_i \left( C_i^{Market} - C_i^{Model}(\Psi) \right)^2,$$

where $i = 1, \ldots N$ is the index for the options used in the calibration. Furthermore, we will assume that option prices are noisy. One implication of this assumption is that a perfect identification of the parametrization is no longer possible. In fact all parametrizations having a pricing error smaller than the bid-ask spread may be true. To find parametrizations which cannot be distinguished in the presence of noisy option prices we implement the following procedure.

For our simulation study, we first assume one benchmark parametrization ($\Theta$) to be the true parametrization, i.e. the parametrization which would persist in a world without noisy option prices. Based on this parametrization, we then compute a set of observable 'market' option prices.[1] Then, we compute the surface of pricing errors for all conceivable parametrizations. We restrict our analysis to the misestimation of two parameters at a time while all other parameters are correctly identified. This procedure has the advantage that it is a conservative approach in that all but two parameters are correctly estimated. We denote by $\widetilde{\Theta}_{(x,y)}$ the parametrization where all but two parameters $(x, y)$ are equal to the parameters in $\Theta$. So for each possible parameter combination $(x, y)$ we compute the sum of squared pricing errors between the true options prices $C_i(\Theta)$ and prices $C_i(\widetilde{\Theta}_{(x,y)})$ $(i = 1, ...N)$ calculated in the Heston model using the parametrization $\widetilde{\Theta}_{(x,y)}$. By plotting the results we obtain a three-dimensional surface, where one parameter is varied along the x-axis, the other one along the y-axis, and the objective function is plotted in z-direction. An often used measure for the total calibration error is the root mean squared error (RMSE). We denote the RMSE by $\Xi_{(x,y)}$:

$$\Xi_{(x,y)} = \sqrt{\frac{1}{N} \sum_i \left( C_i(\Theta) - C_i(\widetilde{\Theta}_{(x,y)}) \right)^2}.$$

An investor calibrating a model will have not be able to distinguish between certain

[1]The benchmark parametrization is taken from Bakshi, Cao, and Chen (1997), Table III, 'All Options', SV-Model. The interest rate is set to 5%, the stock price is 100. The initial variance equals its long-run mean. For the benchmark case 'all options', we consider a cross section of $N = 28$ option prices with maturities of 1 month, 3 months, 6 months and 12 months and moneyness levels in steps of 0.05 from 0.85 to 1.15. For 'ATM options' we use $N = 28$ options with some maturity, but moneyness levels in steps of 0.01 from 0.97 to 1.03, for 'short-term options' $N = 26$ options with maturities of 1 month and 2 months and moneyness levels in steps of 0.02 from 0.88 to 1.12.

parametrization if a bid-ask spread ($BA$) is present. In our study the set of indistinguishable parametrizations $M_{(x,y)}$ is defined as follows.

$$M_{(x,y)} = \left\{ \widetilde{\Theta}_{(x,y)} \middle| \; \Xi_{(x,y)} \leq \frac{BA}{2} \right\}.$$

We thus assume that all parametrizations which result in an RMSE smaller than the upper bound $BA/2$ are not distinguishable from each other given the observed option prices. In the graphs shown in the appendix all parametrizations in the set $M_{(x,y)}$ lie inside the marked bounds. The identification of these sets of indistinguishable parametrizations now allows us to assess the impact on hedging performance caused by the misestimation of parameters.

After calibrating the model to empirical data, the investor sets up a hedge portfolio. Our aim is to identify the difference in performance between the optimal hedge (based on the true parametrization $\Theta$) and the hedges based on all parametrizations within the set $M_{(x,y)}$ for each parameter combination $(x, y)$. The hedging error based on parametrization $\widetilde{\Theta}_{(x,y)}$ is defined as the deviation of the hedge portfolio from the value of the target option after the next time step of length $\Delta t$:

$$\epsilon(\widetilde{\Theta}_{(x,y)}) = C^T_{t+\Delta t} - w_s(\widetilde{\Theta}_{(x,y)}) S_{t+\Delta t} - w_m(\widetilde{\Theta}_{(x,y)}) e^{r\Delta t} - w_i(\widetilde{\Theta}_{(x,y)}) C^I_{t+\Delta t},$$

where $C^T_{t+\Delta t}$ and $C^I_{t+\Delta t}$ are the realized market prices of the target and the instruments option, respectively, at time $t + \Delta t$. For our simulation under the P-measure, we assume that all risk premia are zero and thus the P-measure equals the Q-measure. This makes the average hedging error of all strategies equal to zero and allows to focus solely on the distribution.

The hedging performance under parametrization $\widetilde{\Theta}_{(x,y)}$ is measured by the standard deviation of hedging errors over all observations $M$.

$$\sigma(\widetilde{\Theta}_{(x,y)}) = \sqrt{\frac{1}{M-1} \sum_{i=1}^{M} \epsilon^2(\widetilde{\Theta}_{(x,y)})}.$$

The larger this value the worse the hedging performance. The hedging errors for the true parametrization are taken as a benchmark. Due to discrete trading the hedging errors for this strategy will not be equal to zero and the standard deviation of the hedging error follows as $\sigma(\Theta)$. We correct the standard deviation of hedging errors under parameter risk by this unavoidable standard deviation due to discretization and measure the performance loss due to parameter risk as:

$$\Upsilon_{(x,y)} = \frac{\sigma(\widetilde{\Theta}_{(x,y)}) - \sigma(\Theta)}{C^T_0}.$$

We perform a one-day delta-vega hedge. We hedge an OTM target call with a moneyness (strike/$S_t$) of 1.1 and a maturity of 0.2 years. The hedge portfolio consists of the stock, the money-market account and an instrument option. The instrument option is an ATM call with a maturity of 0.2 years. We perform 10,000 simulation runs with 10 steps per day. After performing the simulations we again draw a three-dimensional surface, where one parameter is varied along the x-axis, the other one along the y-axis and the percentage hedging error is plotted in z-direction. We thus have two plots for each parameter combination: one for the objective function and one for the hedging errors. In order to highlight the effect of microstructural noise on the hedging performance in each plot we mark the subarea of objective function values and hedging errors for the set of indistinguishable parametrizations $M_{(x,y)}$. The hedging errors within the marked subarea result from parametrizations indistinguishable from a noisy cross-section of option prices. All hedging errors in this set are hedging errors which have to be added to the discretization error and stem from the parameter risk induced by the bid-ask-spread in the option price.

# 5 Simulation Results

## 5.1 Pricing and Hedging Performance

In this section we will present and interpret the results of our simulation study. Figures 5 to 8 in the appendix show for each combination of two parameters $(x, y)$ of the Heston model the area of indistinguishable parametrizations $M_{(x,y)}$ when the model is calibrated to noisy option prices (marked area in left surface plots). We also show graphs illustrating how different the resulting hedging performance for each set of parametrizations can be. That is we show the area of possible performance losses for indistinguishable parametrizations (marked area in right surface plots). We perform this analysis for a calibration to all options ('all'), for a calibration only to ATM options ('atm') and for a calibration only to short maturity options ('short').[2] The latter two analyses are done to show how a 'thin' market aggravates the results found for the case of a deeper market where a wide spectrum of moneyness levels and maturities is available. The underlying problem here is that on the one hand parameter identification is easier when the whole range of strikes and maturities is used and that a mis-estimation of parameters is more likely when only short maturity or only ATM options are used. On the other hand, the use of highly liquid short maturity or ATM options may be preferable due to a lower bid-ask spread and the resulting better identification of the true model with more precise option prices.

The numerical results are based on a bid-ask spread of 10% which lies well within the empirical spread observation. For options on the S&P100 individual stocks the average bid-ask spreads between 1996 and 2003 sorted by delta buckets are depicted in Table 1. They are lowest with about 3% for deep ITM options and largest with up to 42% for deep OTM options. For ATM options the average bid-ask spread is around 7%.

We start our interpretation of the graphs with the parameter combination ($\sigma_V$,$\kappa$) in Figure 5. For the case 'all' (upper row) we find that for a given bid-ask spread the identification of $\kappa$ is more difficult than the identification of $\sigma_V$. The marked area in the

[2] Since for all parameter combinations the results for the cases 'all' and 'short' are almost the same, we henceforth only report the results for the case 'all'.

graph shows that $\kappa$ can vary roughly between 0 and 5, while $\sigma_V$ can only deviate by a relatively small amount from its true value. For the 'atm' (lower row) case we immediately see that identification of the true parameter combination becomes extremely difficult. Almost all parameter combinations of $\sigma_V$ and $\kappa$ are possible.

For the parameter combination ($\sigma_V$,$\theta$) in Figure 6 we see that both parameters are quite easy to identify in the 'all' (upper row) case. The area of indistinguishable parametrizations also increases in the 'atm' (lower row) case although not as drastically as in the previous case. An interesting result is that in the 'atm' case the form and curvature of the surface of the objective function is very different from the surface of the hedging performance. In particular the smallest values of the objective function are on a straight line along the true value of $\theta$ while the smallest values of the performance loss function are on a straight line along the true value of $\sigma_V$. Also note that the performance loss function is strongly increasing in $\sigma_V$ from the point of the minimum of $\Upsilon_{(\sigma_V,\theta)}$ onwards. This means that in the calibration procedure a mis-estimation of $\sigma_V$ is likely to be compensated by a correct identification of $\theta$, but this could result in large performance losses.

The parameter combination ($V_0$,$\sigma_V$) is probably the one easiest to identify (see Figure 7). The area of indistinguishable parameters is extremely small. Again the identification becomes worse when only ATM options (lower row) are used. It seems that $V_0$ is more important for prices and hence slightly easier to identify. In the plots for the combination ($\theta$,$\kappa$) in Figure 8 we see that for small values of $\kappa$ the identification of $\theta$ becomes very difficult. The higher $\kappa$, the easier the identification of $\theta$.

Since similar results hold for the other parameter combinations, the graphs are not shown. Note, however, that in accordance with the results in Section 3 it is always very important to use OTM options to identify $\rho$.

In summary, we find that in general pricing performance is a good proxy for hedging performance. A notable exception here is the parameter combination ($\sigma_V$,$\theta$) where the surface of the objective function is very different from the surface of the hedging performance in the 'atm' case. Furthermore, $V_0$ is the most important parameter for pricing and hedging as expected. A correct identification should be easy, but a mis-estimation will have serious consequences for the hedging performance. We find that $\sigma_V$ and $\rho$ are of particular importance: For $\sigma_V$ an increase in hedging error is steeper than the increase in the objective function. This means that a mis-estimation of this parameter has a large impact on hedging errors. For $\rho$ we find that a calibration is very difficult if only ATM options are used. This could result in very high performance losses. Consequently in order to identify $\rho$, it is of first order importance to use OTM options. We also find a cross effect between $\theta$ and $\kappa$. For very low values of $\kappa$ the identification of $\theta$ is very difficult and the closer $\theta$ is to its true value, the more difficult identification of $\kappa$ becomes. Finally, most effects are more pronounced when calibration is done using ATM options only. That is, the performance loss can become a lot higher when using only ATM options. Although ATM options may have lower bid-ask spreads and prices are therefore more precise, we find that the less liquid OTM options contain important information (concerning the tails of the distribution) necessary to identify $\rho$.

## 5.2 Hedging and Microstructural Noise

After having identified the qualitative effects in our simulation study, we next want to show the quantitative effects. In particular we want to answer the following questions: How large is the maximum possible size of the performance loss for different sizes of the bid-ask spread? Although ATM options have a lower bid-ask spread and the price information is therefore more precise, does a calibration to these options, only, result in significantly worse hedging performance? In order to answer these questions we plot the maximum performance loss $\max \Upsilon_{(x,y)}$ as a function of the mean bid-ask spread over options used for calibration for three different cases: all options, atm options and short-maturity options. These graphs give a feel for the quantitative impact of parameter risk on hedging in a worst case scenario. This means we show how large the performance loss can possibly get when the investor cannot distinguish between different parametrizations due to a given bid-ask spread.

The results are depicted in Figure 9. First of all note that since we restrict the parameter range to reasonable values[3], $\Upsilon$ seems to converge to an upper bound in some graphs. However without the restriction on the parameter set, it could increase even further. For example in the case ($\kappa$,$\rho$) the upper bound on $\Upsilon$ in the case 'atm' is obtained for $\kappa = 10$ and $\rho = 0$. Without the restricted parameter range, both parameters could possibly be estimated larger and then $\Upsilon$ would be larger. Also note that the stepwise behavior of the functions is a result of the parameter grid we chose for the calibration.

When comparing the graphs for the different calibrations, we make the following observations: First, for an average bid-ask spread of 5% the investor may suffer from a performance loss of more than 2% (case $\sigma_V$,$\rho$) when the investor calibrates to 'all' options and up to 11% (case $\sigma_V$,$\rho$) when the investor calibrates to 'atm' options. For an average bid-ask spread of 30%, the maximum performance loss lies in a range between 7% (case $\theta$,$\rho$) and 16% (case $\sigma_V$,$V_0$) when the investor calibrates to 'all' options. In the case of 'atm' options the maximum performance loss becomes even higher and is within a range of 11% (case $\kappa$,$\rho$) and 30% (cases $\rho$,$V_0$ and $\sigma_V$,$\rho$). So we see that the lowest performance loss is obtained when calibration is done using 'all' options, while the largest performance loss results for 'atm' options and the performance loss for 'short'-maturity options is in between. The hedging performance for 'short'-maturity options is, however, only marginally worse than the hedging performance for the calibration to 'all options'. This means that for dynamic hedging it is of huge importance to use a wide range of moneyness levels while the range of maturities seems to be less relevant. The use of OTM options for calibration allows the investor to accept a higher bid-ask spread and still achieve a similar hedging performance. As can be seen in the graphs, in most cases a calibration to ATM options with a bid-ask spread of 5% results in a similar hedging performance as a calibration to 'all' options or 'short'-maturity options with a bid-ask spread of more than 20%.

Regarding the relative importance of the particular parameters of the Heston model

[3]The parameters were restricted to be in the following intervals: $\kappa \in [0, 10]$, $\sigma_V \in [0, 1]$, $\theta \in [0, 0.1]$, $\rho \in [-1, 0]$ and $V_0 \in [0, 0.1]$.

we find the following numeric results: performance losses resulting from a calibration to 'all' options are largest when $\sigma_V$ is one of the parameters subject to mis-estimation. i.e. identification is difficult and mis-estimation may lead to large performance losses, because $\sigma_V$ has only a small impact on prices, but has a large impact on hedging errors. A similar result holds for the parameter $\rho$ for which already our introductory analysis has indicated that OTM options are extremely important for the calibration. As can be seen in the graphs, when $\rho$ is one of the parameters subject to mis-estimation, the maximum performance loss increases sharply when only 'atm' options are used for calibration.

# 6 Empirical Illustration

## 6.1 Design of the Study

To support the economic relevance of our simulation results, we assess the impact of parameter risk on hedging errors empirically.

Theoretically, indistinguishable parametrizations are those for which the calibrated option prices lie within the bid-ask spread. Finding empirically the worst case scenario, i.e. the parametrization which yields the worst hedging performance among all indistinguishable parametrizations is computationally infeasible. To keep our analysis as simple as possible we therefore take a conservative approach and perform three different calibrations: to the bid-, ask- and mid-prices. Additionally, we again calibrate our prices to the Heston model which seems to provide a sufficient complexity to reflect our empirical data, while being analytically tractable at the same time. Of course, there may be other risk factors like jumps in the true stock dynamics, especially for the index options. However, this is not the focus of our analysis, we do not want to state that the Heston model is the true model to describe the data. Rather, we intend to show that assuming one model there may be several equally well parametrizations which still yield different hedging performances.

At each trading day in our data set, we perform the following steps. First, we calibrate the Heston model to the empirically observed (bid, ask and mid) option prices. These are all call options traded on that day on the specific underlying available in the Option Metrics database. Second, we build the hedge portfolio. We set up the Heston delta-vega and, for comparison, also the Heston delta hedge for each of the three parametrizations. This is done for all call options traded on that and the following day. Third, to assess the differences in hedging performance, we compute the hedging errors after one day. For both the construction and the evaluation of the hedge portfolio, we use mid-prices and thus do not focus on either short- or long-positions.

The hedging performance is then judged according to the standard deviation of hedging errors, differences between the parametrizations are statistically tested with an $F$-Test and the non-parametric test of Ansari-Bradley (AB).

## 6.2 Data Description

We use Option Metrics to select all call options on the S&P100 index and two individual stocks (Cisco, and AT&T) in the time frame from 2000 to 2004.

First, we take all calls with times to maturity from 14 to 180 days. We eliminate option price observations for dates with zero open interest, with zero bid prices or with missing implied volatility or delta. We also take out deep ITM and far OTM options (with moneyness smaller then 0.8 and larger than 1.2), as the lack of liquidity in those options may distort the results.

After applying these filters we are left with the number of observations as shown in Table 2. From these numbers we can see that we have on average 15 option prices observed on each date for individual securities and 64 options for S&P100.

In the following we treat each underlying separately. On each date we select the option to be used as the instrument for hedging. To do so we first select the options with remaining time to maturity closest to the average time to maturity for all options observed on a given date. From this group of options we then select the one closest to the ATM level.

## 6.3 Results

Table 3 summarizes the results. In addition to the results from the Heston model, we include the hedging errors obtained from a Black-Scholes delta hedge based on implied volatility.

In the table the standard deviation of the hedging errors (relative to the initial value of the option) is depicted as a measure of hedging performance. The point of interest for our analysis is when the hedging performance is significantly different if the hedge is based on different parametrizations. To confirm the visible differences, we perform a statistical test for the difference in variance of hedging errors based on the ask- and the bid-calibration. Under the assumption of normally distributed hedging errors, a two-sided $F$-Test is performed for the null of equal variances. The non-parametric AB test tests the null that hedging errors of bid- and ask-calibration come from the same distribution against the alternative that they come from distributions with same median but different variances. To increase the power of the AB-test, we first normalize our data by subtracting the medians. The results are sorted by the moneyness of the target option.

The general results are as expected. The standard deviations of the hedging errors in case of the delta-vega hedge are much smaller than in case of the simple delta hedge. They are largest for OTM options and smallest for ITM options. Surprisingly, but in line with practitioners' experience is the very good performance of the Black-Scholes delta hedge. Often it outperforms the delta-vega hedge, although this hedge uses an additional instrument. For both statistical tests, the variances of Black-Scholes delta hedging error are significantly different from Heston delta hedging errors. This does not necessarily mean that the Heston hedge is generally worse, but reveals a mistake often arising in the measurement of hedging performance. There is a mismatch between the objective of our hedge and the performance measure. Whereas in our objective we want to set delta and vega of our portfolio equal to zero, the performance measure prefers hedges with low standard deviations. The Black-Scholes hedge portfolio as built in our example often seems to be closer to the minimum-variance hedge than the Heston delta-vega hedge.

Looking at the results for Cisco in Panel A, the standard deviation of relative hedging errors is larger for the Black-Scholes hedge than for the Heston hedge, in particular for the delta-vega hedge. When the Heston model is used, a delta-vega hedge decreases the standard deviation of about 0.02 over all options compared to the simple Heston delta hedge. For different moneyness levels, the hedging performance is best for ITM options and worst for OTM options. In case of OTM options, a simple Black-Scholes hedge has a lower standard deviation than the Heston delta-vega hedge. The differences in hedging performance of the calibrations may be of important size: It is more than 0.1 for all options and more than 0.2 for OTM options. The differences in results of the bid- and ask-calibrations are statistically significant in case of the F-Test (to the 5% level), but not significant in case of the AB-Test. The differences are largest for OTM options where the hedging performance for different calibrations differs in more than 20% of option value.

In case of the S&P100 index in Panel B, the standard deviations are much larger than for the other securities, especially for OTM options. As before, the standard deviation of relative hedging errors is largest for OTM options and smallest for ITM options. On average over all options, the performance of the Black-Scholes delta hedge is better than the performance of the Heston hedge. This is mainly due to the bad performance of the Heston hedge for OTM options. Between the parametrizations, differences of about 0.01 for ATM up to 0.05 for OTM options are observed. They are in most cases statistically significant. Especially for OTM options, the differences in the standard deviation between the parametrizations arise to 5 percentage points.

The results for AT&T are depicted in Panel C of Table 3. For the Heston model, the delta-vega hedge is much better than the simple delta hedge and the standard deviation of hedging errors decreases of about 0.03 on average over all options. As before, the performance of the hedge is best for ITM and worst for OTM options. The differences between the parametrizations may arise up to 0.01 for ATM and OTM options. They are in many cases statistically significant.

In summary, the differences in hedging performance of the parametrizations are economically and statistically significant. Especially for OTM options, the already large standard deviation of the hedging errors of about 24% of the option value may nearly double to about 45%. For ITM and ATM options the standard deviation differs in less than 0.01 which is small in absolute value but relative to the total standard deviation of on average 0.05 still significant.

Compared to the controlled simulation analysis, this empirical study is influenced by many external factors. Most importantly, since our study is based on the Heston model, it is certainly subject to model risk - the fact that our assumed stochastic volatility model is not the true data-generating process. The more surprising is the fact that even in this simple study parameter risk is shown to significantly impact the hedging performance.

# 7 Conclusion

The presence of microstructural noise in option prices makes the calibration of an option pricing model difficult. Even if the investor is sure about the structural type of the model, he will not be able to identify the true parametrization. As we have shown for the case of the Heston model, this is in particular true for the parameters of second-order importance like $\kappa$ or $\sigma_v$. Another parameter difficult to estimate is $\rho$ for which it is extremely important to have OTM option prices at hand. For calibration, a wide range of moneyness levels is therefore much more important than a wide range of maturities.

In a second step, however, finding the correct parametrization is not an objective in itself. The investor always has to take into account for what purpose he needs a calibration. In this paper we have focussed on the impact of different parametrizations on the hedging performance. We have also shown that $\sigma_v$ and $\rho$ are not only difficult to identify, but also that a mis-estimation of these parameters may have severe consequences for the hedging performance. Further, the link of the size of the average bid-ask spread to the average hedging performance allowed us to put the need of OTM option prices for calibration in numbers: To have only ATM options with a bid-ask-spread of 5% available is as good as having options with a wide range of moneyness levels with a bid-ask-spread of 20%!. The most important parameters to identify correctly for hedging purposes are again $\sigma_v$ and $\rho$.

Already our simple empirical study could confirm the theoretical results of the simulation analysis. In terms of standard deviation of hedging errors, the differences between the parametrization arose to 20% relative to the target option value for OTM options. The results of our paper illustrate the difficulties arising from microstructural noise and highlight the importance of parameter risk for hedging. Further research in this field seems necessary.

# References


An, Y., and W. Suo, 2003, The Performance of Option Pricing Models on Hedging Exotic Options, Working Paper.

Bakshi, G., C. Cao, and Z. Chen, 1997, Empirical Performance of Alternative Option Pricing Models, *Journal of Finance* 52, 2003–2049.

Black, F., and M. Scholes, 1973, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy* 81, 637–654.

Dennis, P., and S. Mayhew, 2004, Microstructural Biases in Empirical Tests of Option Pricing Models, EFA 2004 Maastricht Meetings Paper No. 4875.

Eraker, B., M. Johannes, and N. Polson, 2003, The Impact of Jumps in Volatility and Returns, *Journal of Finance* 58, 1269–1300.

Figlewski, S., 2004, Estimation Error in the Assessment of Financial Risk Exposure, Working Paper.

He, C., J.S. Kennedy, T. Coleman, P.A. Forsyth, Y. Li, and K. Vetzal, 2006, Calibration and Hedging under Jump Diffusion, *Review of Derivatives Research* 9, 1–35.

Hentschel, L., 2003, Errors in Implied Volatility Estimation, *Journal of Financial and Quantitative Analysis* 38, 779–810.

Hentschel, L., 2004, Option Hedging in the Presence of Measurement Errors, Working Paper.

Heston, S.L., 1993, A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options, *Review of Financial Studies* 6, 327–343.

Schoutens, W., E. Simons, and J. Tistaert, 2003, A Perfect Calibration! Now What?, Working Paper.

| $\Delta^{BS}$ interval | Bid-Ask Spread |
|---|---|
| [-1.0,-0.8) | 0.03 |
| [-0.8,-0.6) | 0.05 |
| [-0.6,-0.4) | 0.07 |
| [-0.4,-0.2) | 0.14 |
| [-0.2, 0.0] | 0.42 |

Table 1: Bid-Ask Spreads on S&P100 Individual Stock' Put Options

The table shows the average percentage spreads of put options on S&P100 individual stocks from 1996 to 2003 as a function of the $\Delta^{BS}$ interval. $\Delta^{BS}$ is the Black-Scholes delta based on implied volatility.

| Underlying | Number of Observations | Number of Options | Number of Days | Avg. Number of Observations per Day |
|---|---|---|---|---|
| Cisco | 13,850 | 434 | 1,248 | 11 |
| AT&T | 10,128 | 337 | 1,244 | 8 |
| S&P100 | 80,301 | 2,141 | 1,248 | 64 |

Table 2: Data Description

| Panel A: Cisco | | | | | |
|---|---|---|---|---|---|
| | | Moneyness Range | | | |
| | Calibration | 0.8 to 0.95 | 0.95 to 1.05 | 1.05 to 1.2 | All Options |
| Delta (rel.) | Bid | 0.0963 | 0.2728 | 0.5701 | 0.1570 |
| | Mid | 0.0952 | 0.2684 | 0.5629 | 0.1557 |
| | Ask | 0.0945 | 0.2679 | 0.5530 | 0.1499 |
| Delta-Vega (rel.) | Bid | 0.0503 | 0.0844 | 0.2509 | 0.1638 |
| | Mid | 0.0506 | 0.0827 | 0.2419 | 0.1583 |
| | Ask | 0.0521 | 0.0905 | 0.4543 | 0.2874 |
| BS-Hedge (rel.) | | 0.0915 | 0.2369 | 0.3853 | 0.2694 |
| F-Test (p-value) | Delta | 19.82% | 36.80% | 3.91% | 0.17% |
| | Delta-Vega | 1.21% | 0.05% | 0.00% | 0.00% |
| AB-Test (p-value) | Delta | 3.19% | 0.67% | 16.48% | 0.41% |
| | Delta-Vega | 51.16% | 34.95% | 74.17% | 88.38% |

| Panel B: S&P100 | | | | | |
|---|---|---|---|---|---|
| | | Moneyness Range | | | |
| | Calibration | 0.8 to 0.95 | 0.95 to 1.05 | 1.05 to 1.2 | All Options |
| Delta (rel.) | Bid | 0.0295 | 0.1343 | 0.6946 | 0.4523 |
| | Mid | 0.0294 | 0.1257 | 0.5975 | 0.3903 |
| | Ask | 0.0294 | 0.1306 | 0.6420 | 0.4188 |
| Delta-Vega (rel.) | Bid | 0.0188 | 0.0721 | 0.5133 | 0.3317 |
| | Mid | 0.0195 | 0.0683 | 0.4630 | 0.2995 |
| | Ask | 0.0190 | 0.0713 | 0.5013 | 0.3240 |
| BS-Hedge (rel.) | | 0.0370 | 0.0897 | 0.3011 | 0.2010 |
| F-Test (p-value) | Delta | 66.09% | 0.00% | 0.00% | 0.00% |
| | Delta-Vega | 0.00% | 0.00% | 0.00% | 0.00% |
| AB-Test (p-value) | Delta | 82.91% | 0.00% | 0.00% | 0.00% |
| | Delta-Vega | 18.29% | 0.13% | 14.88% | 6.82% |

| Panel C: AT&T | | | | | |
|---|---|---|---|---|---|
| | | Moneyness Range | | | |
| | Calibration | 0.8 to 0.95 | 0.95 to 1.05 | 1.05 to 1.2 | All Options |
| Delta (rel.) | Bid | 0.0349 | 0.0968 | 0.2411 | 0.1570 |
| | Mid | 0.0349 | 0.0926 | 0.2330 | 0.1517 |
| | Ask | 0.0355 | 0.0924 | 0.2234 | 0.1461 |
| Delta-Vega (rel.) | Bid | 0.0283 | 0.0681 | 0.1997 | 0.1287 |
| | Mid | 0.0286 | 0.0688 | 0.1913 | 0.1238 |
| | Ask | 0.0292 | 0.0667 | 0.1882 | 0.1218 |
| BS-Hedge (rel.) | | 0.0398 | 0.0870 | 0.1959 | 0.1300 |
| F-Test (p-value) | Delta | 32.76% | 5.55% | 0.00% | 0.00% |
| | Delta-Vega | 7.50% | 40.14% | 0.08% | 0.00% |
| AB-Test (p-value) | Delta | 85.64% | 11.86% | 3.10% | 8.25% |
| | Delta-Vega | 94.32% | 45.90% | 89.91% | 60.71% |

Table 3: Standard Deviations of Hedging Errors

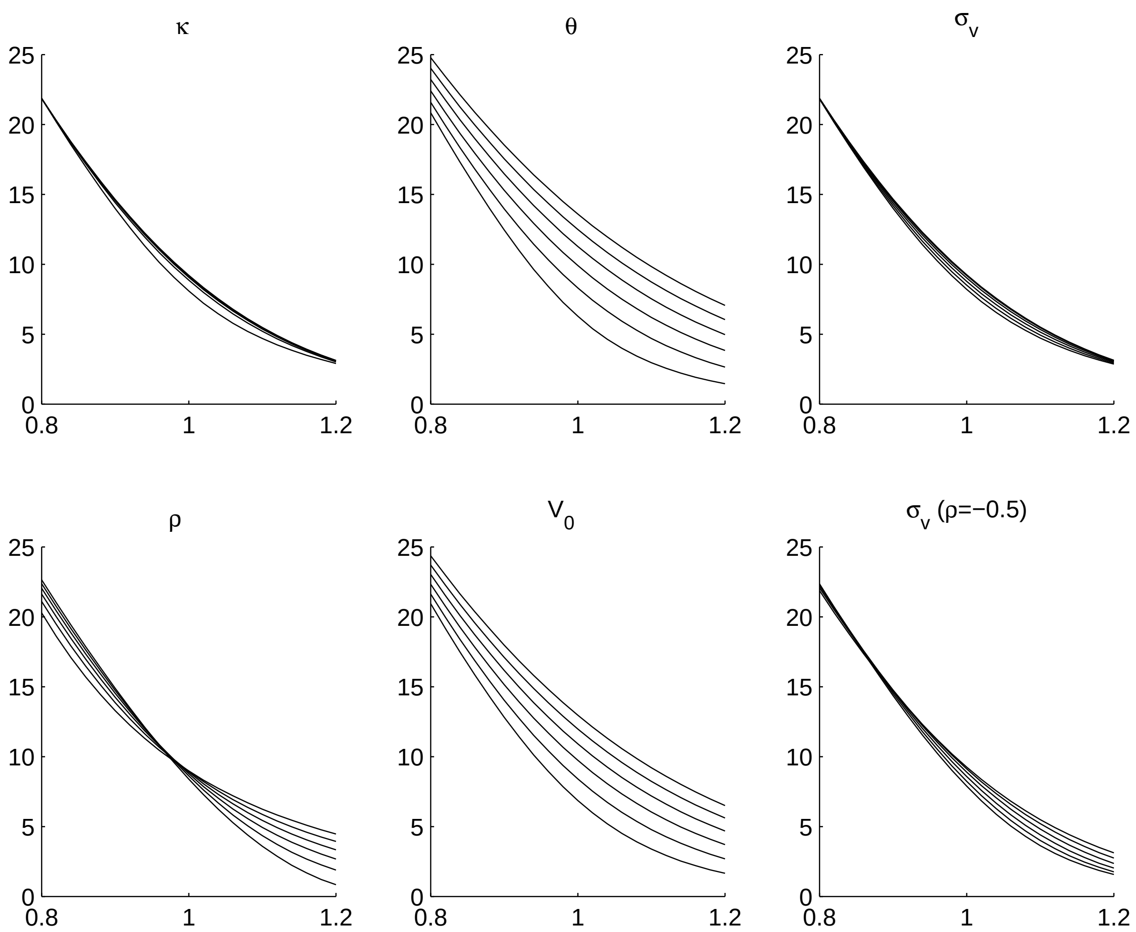


Figure 1: Sensitivity of Call Price

The graphs show the sensitivities of a call price in the Heston (1993) model as a function of moneyness. In each graph one parameter is varied within a certain range while all others are held constant. The base case and the parameter ranges were chosen as follows: $\kappa = 2.0$ $[0.1; 10.0]$, $\theta = 0.06$ $[0.01; 0.2]$, $\sigma_v = 0.5$ $[0.01; 0.9]$, $\rho = 0.0$ $[-1.0; 1.0]$, $V_0 = 0.06$ $[0.01; 0.2]$. The current stock price equals 100 and the interest rate is zero.

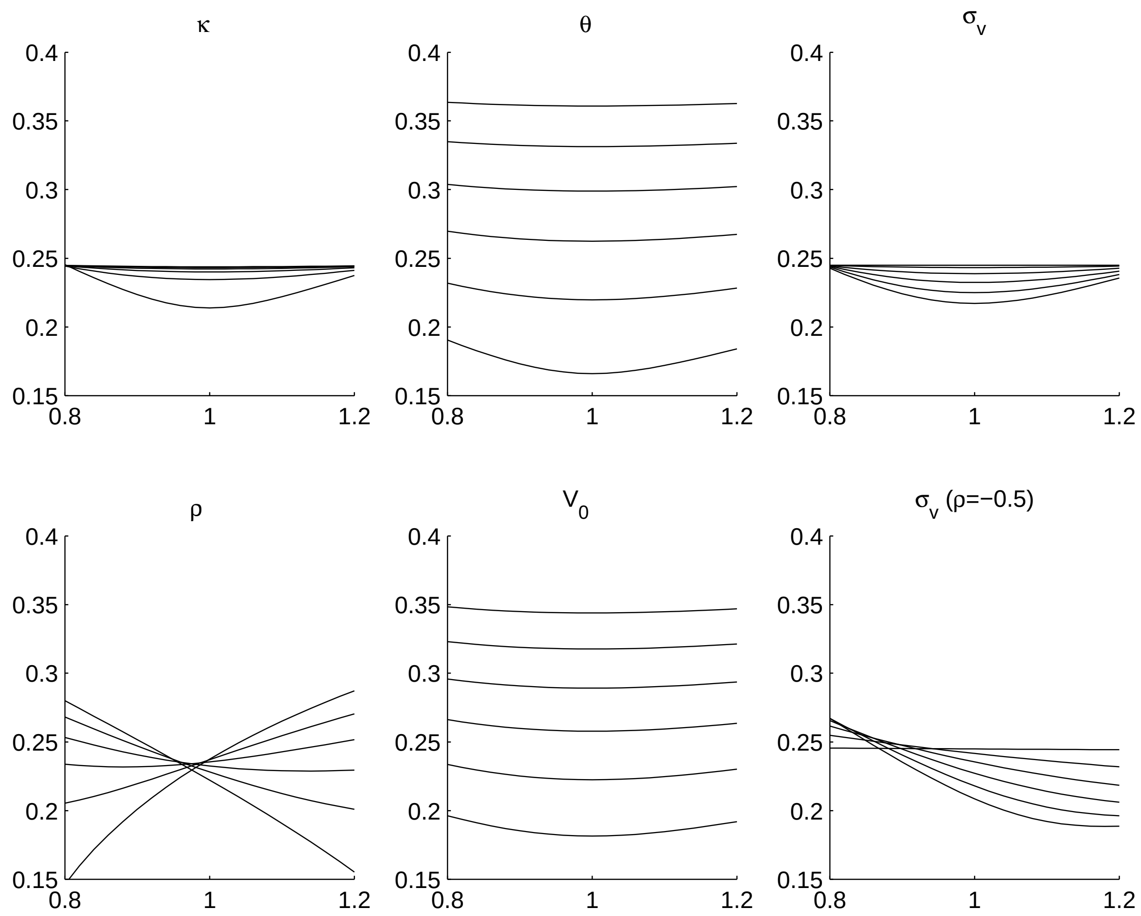


Figure 2: Sensitivity of IV-Smile

The graphs show the sensitivities of the IV-Smile in the Heston (1993) model. In each graph one parameter is varied within a certain range while all others are held constant. The base case and the parameter ranges were chosen as follows: $\kappa = 2.0$ $[0.1; 10.0]$, $\theta = 0.06$ $[0.01; 0.2]$, $\sigma_v = 0.5$ $[0.01; 0.9]$, $\rho = 0.0$ $[-1.0; 1.0]$, $V_0 = 0.06$ $[0.01; 0.2]$. The current stock price equals 100 and the interest rate is zero.

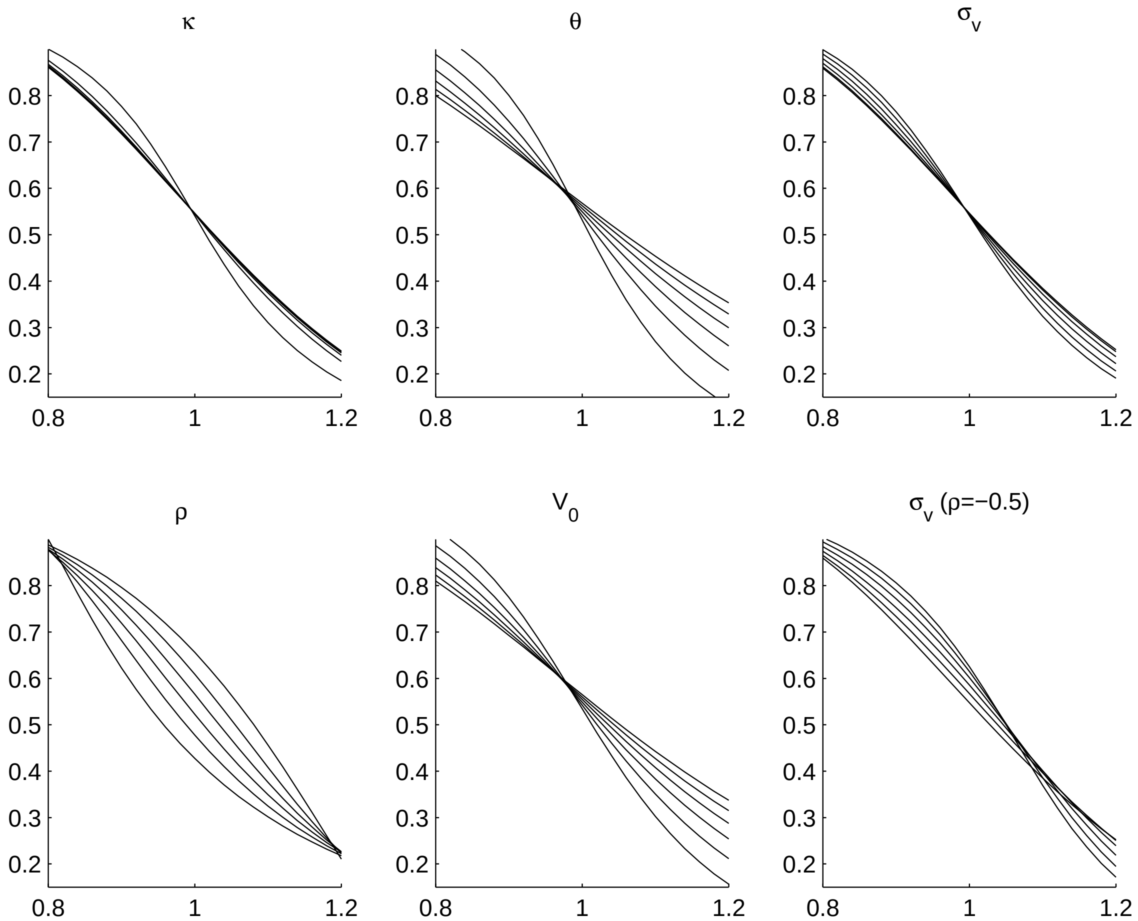


Figure 3: Sensitivity of Option Delta

The graphs show the sensitivities of the option delta in the Heston (1993) model. In each graph one parameter is varied within a certain range while all others are held constant. The base case and the parameter ranges were chosen as follows: $\kappa = 2.0$ $[0.1; 10.0]$, $\theta = 0.06$ $[0.01; 0.2]$, $\sigma_v = 0.5$ $[0.01; 0.9]$, $\rho = 0.0$ $[-1.0; 1.0]$, $V_0 = 0.06$ $[0.01; 0.2]$. The current stock price equals 100 and the interest rate is zero.

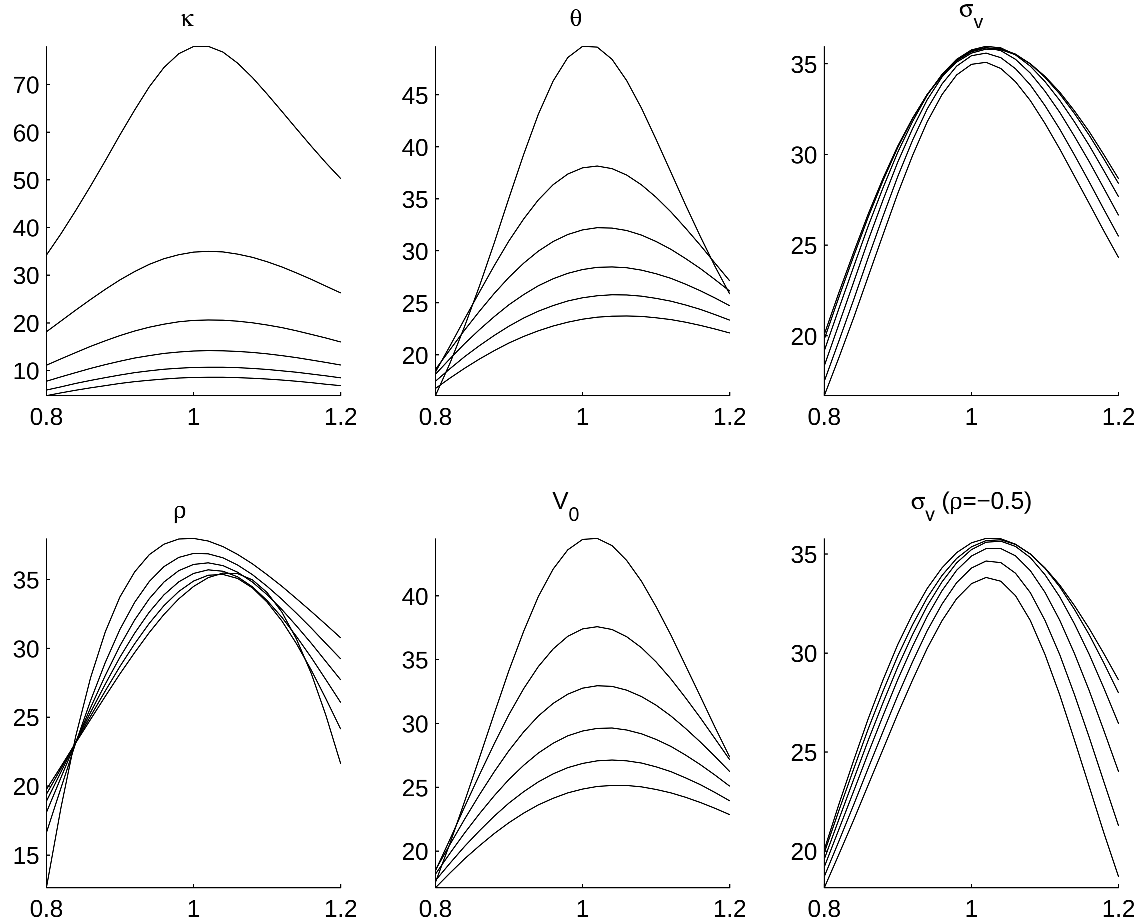


Figure 4: Sensitivity of Option Vega

The graphs show the sensitivities of the option vega in the Heston (1993) model. In each graph one parameter is varied within a certain range while all others are held constant. The base case and the parameter ranges were chosen as follows: $\kappa = 2.0$ $[0.1; 10.0]$, $\theta = 0.06$ $[0.01; 0.2]$, $\sigma_v = 0.5$ $[0.01; 0.9]$, $\rho = 0.0$ $[-1.0; 1.0]$, $V_0 = 0.06$ $[0.01; 0.2]$. The current stock price equals 100 and the interest rate is zero.

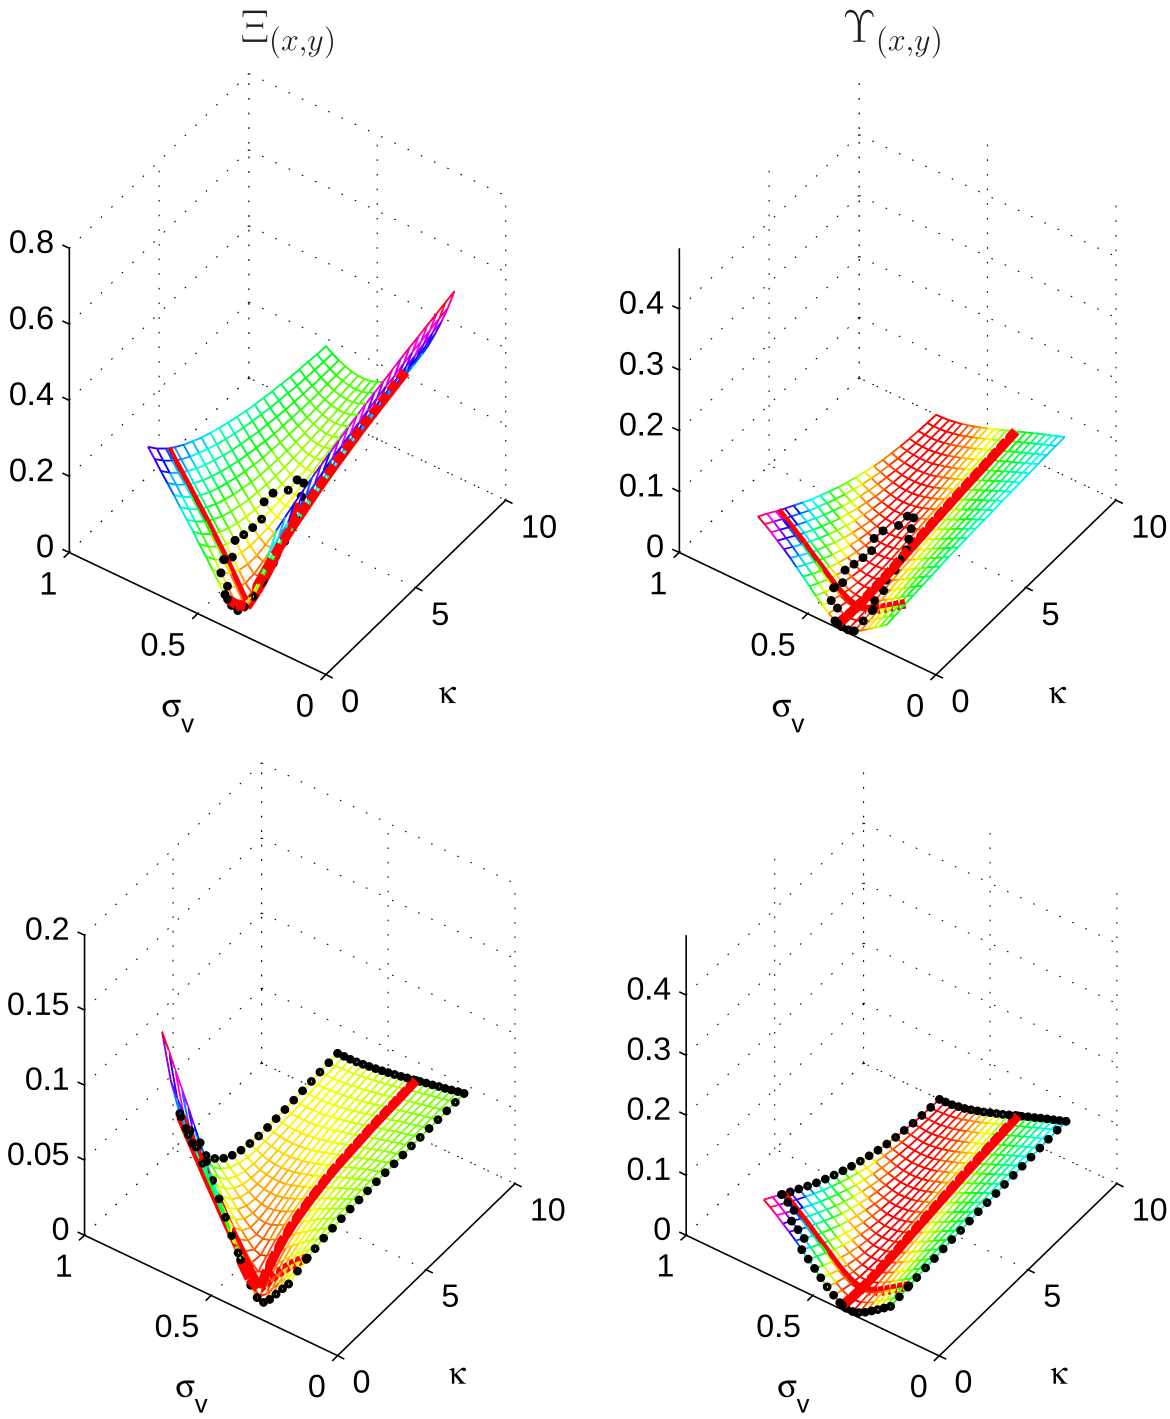


Figure 5: Objective Function and Standard Deviation of Hedging Errors for $\sigma_v$ and $\kappa$

The graphs plot the objective function of calibration $\Xi_{(x,y)}$ (left) and the standard deviation of hedging errors (as percentage of the option price) $\Upsilon_{(x,y)}$ (right) as a function of the parameters $\sigma_v$ and $\kappa$. In the upper graphs the calibration is done for 'all' options, in the lower graphs for 'atm' options only.

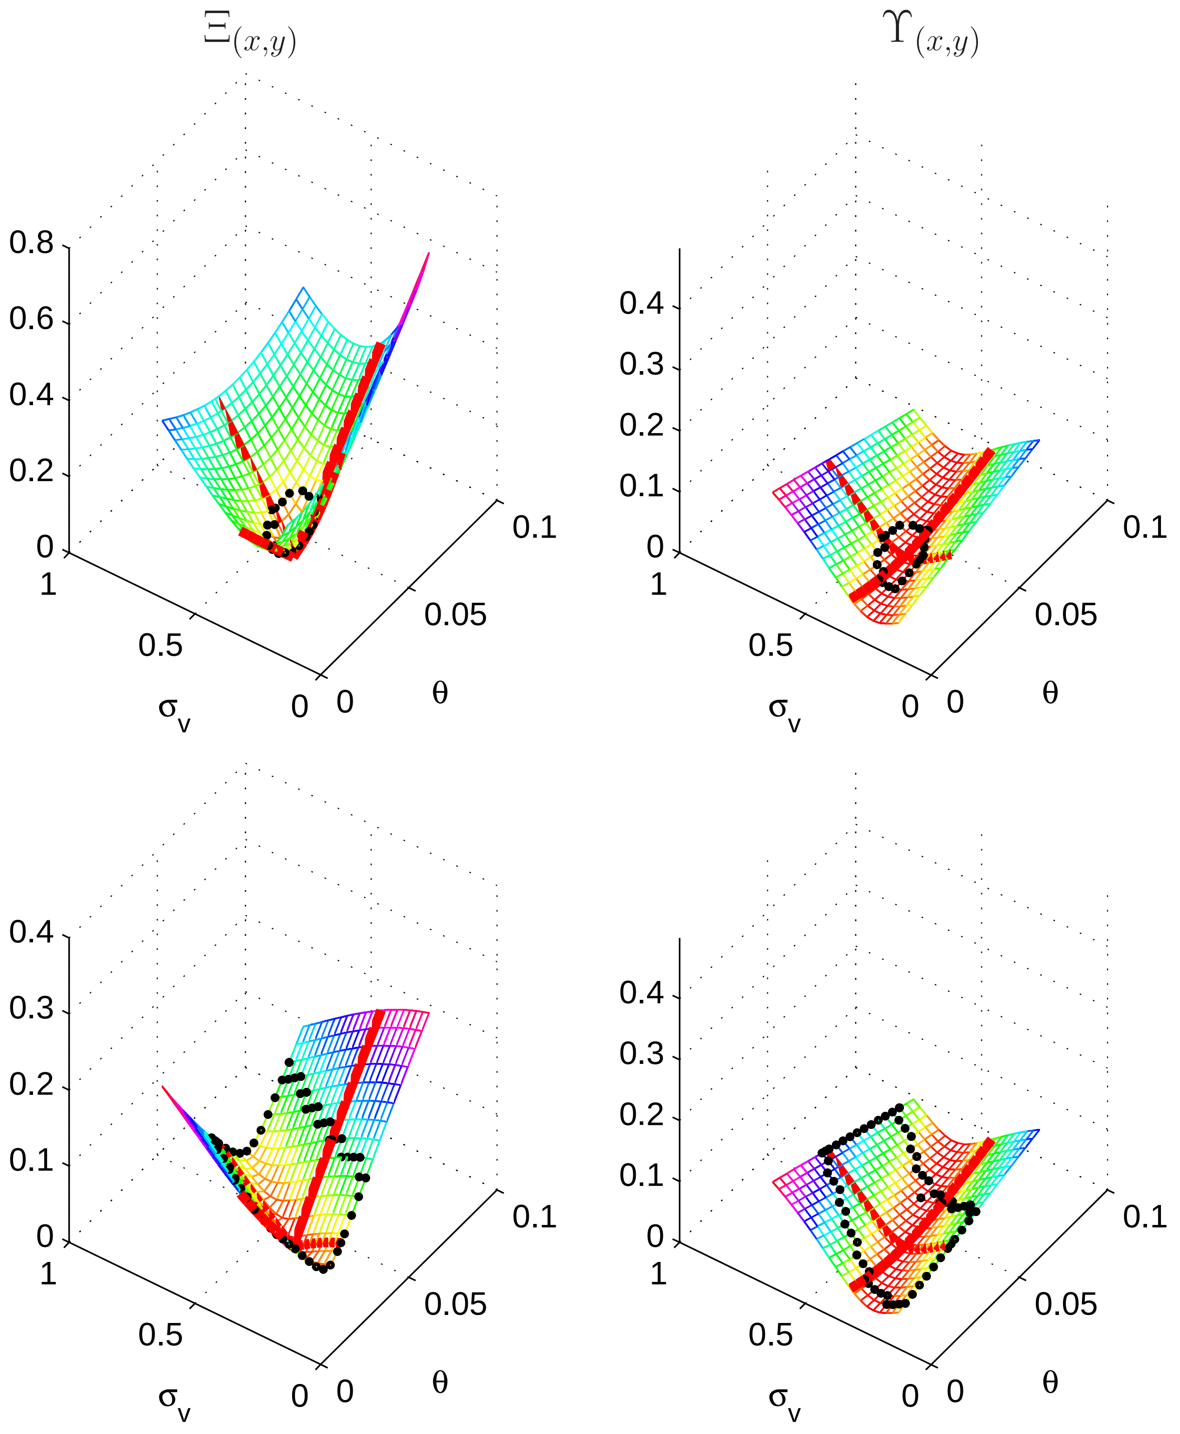


Figure 6: Objective Function and Hedge Errors for a Variation of $\sigma_v$ and $\theta$

The graphs plot the objective function of calibration $\Xi_{(x,y)}$ (left) and the standard deviation of hedging errors (as percentage of the option price) $\Upsilon_{(x,y)}$ (right) as a function of the parameters $\sigma_v$ and $\theta$. In the upper graphs the calibration is done for 'all' options, in the lower graphs for 'atm' options only.

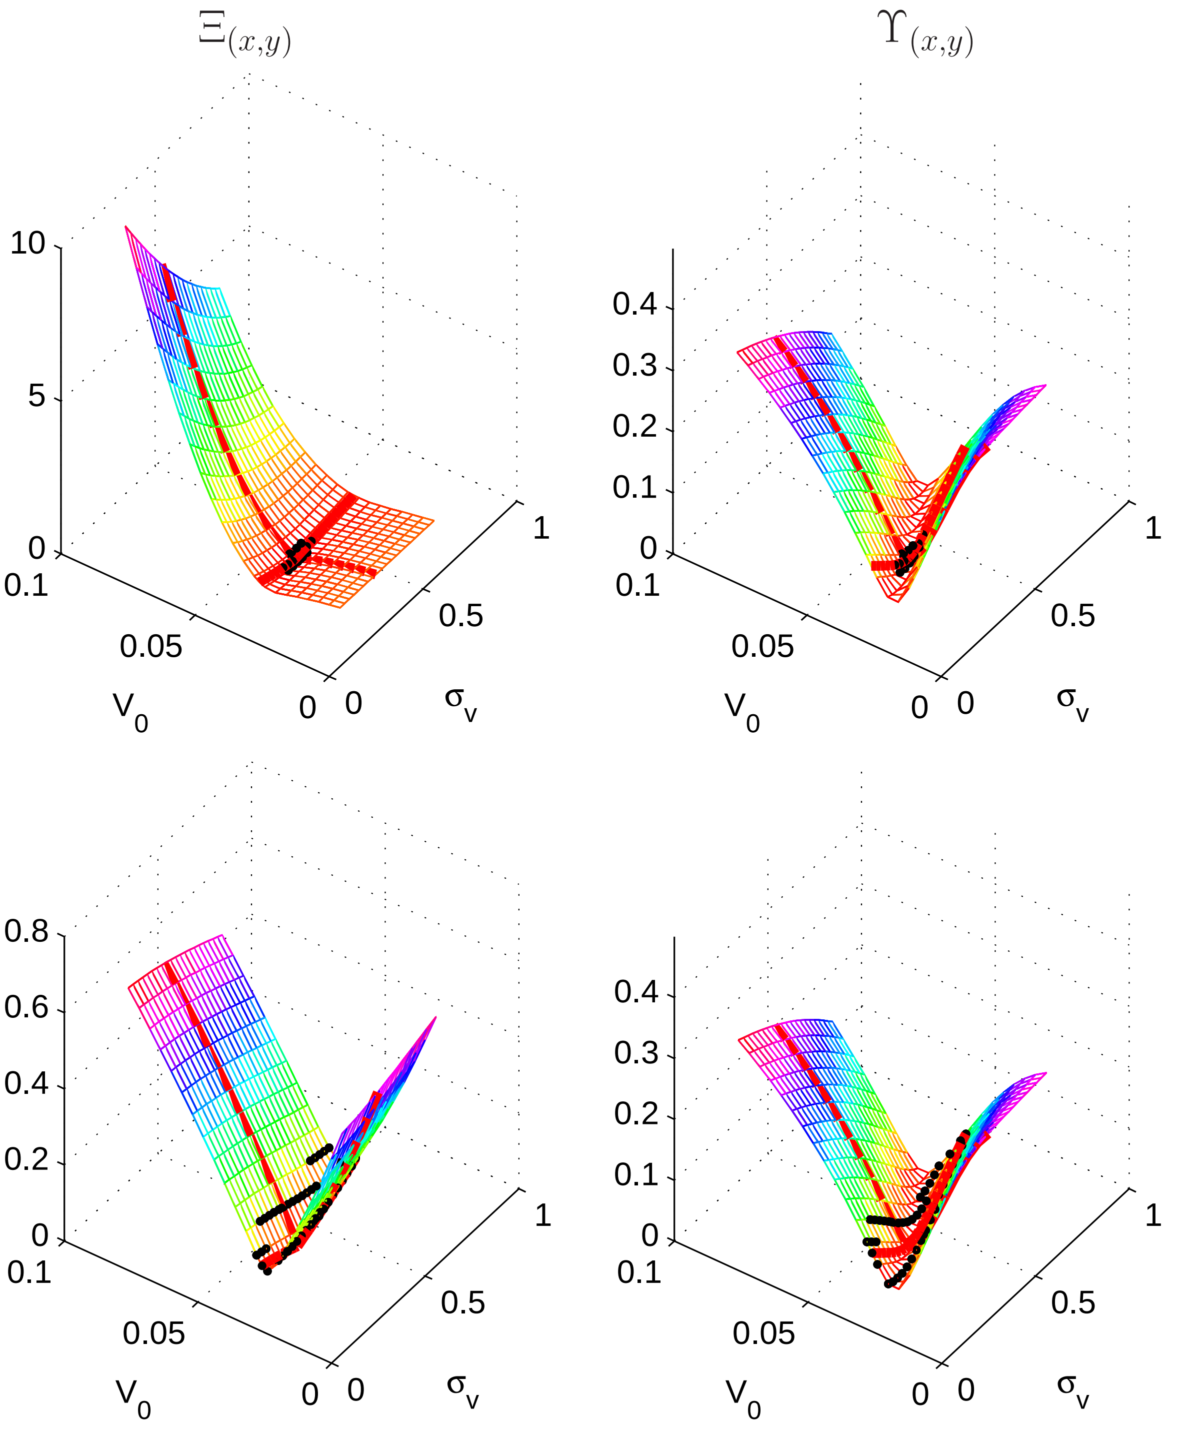


Figure 7: Objective Function and Hedge Errors for a Variation of $V_0$ and $\sigma_v$

The graphs plot the objective function of calibration $\Xi_{(x,y)}$ (left) and the standard deviation of hedging errors (as percentage of the option price) $\Upsilon_{(x,y)}$ (right) as a function of the parameters $\sigma_v$ and $V_0$. In the upper graphs the calibration is done for 'all' options, in the lower graphs for 'atm' options only.

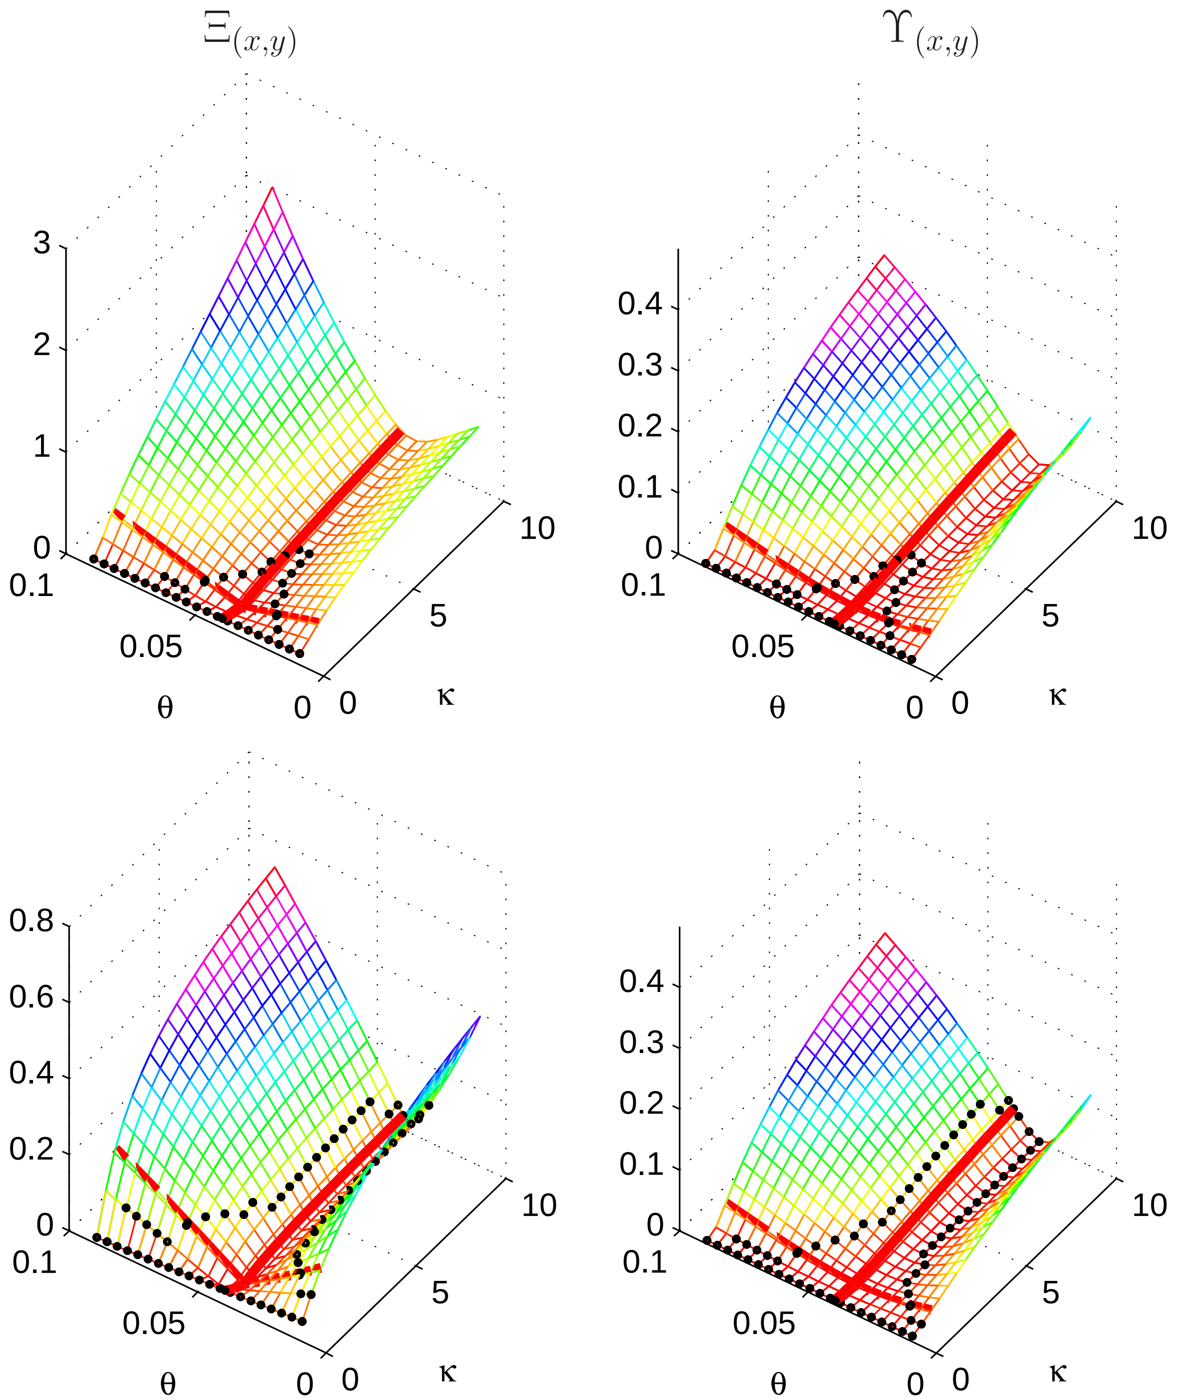


Figure 8: Objective Function and Hedge Errors for a Variation of $\theta$ and $\kappa$

The graphs plot the objective function of calibration $\Xi_{(x,y)}$ (left) and the standard deviation of hedging errors (as percentage of the option price) $\Upsilon_{(x,y)}$ (right) as a function of the parameters $\sigma_v$ and $\theta$. In the upper graphs the calibration is done for 'all' options, in the lower graphs for 'atm' options only.

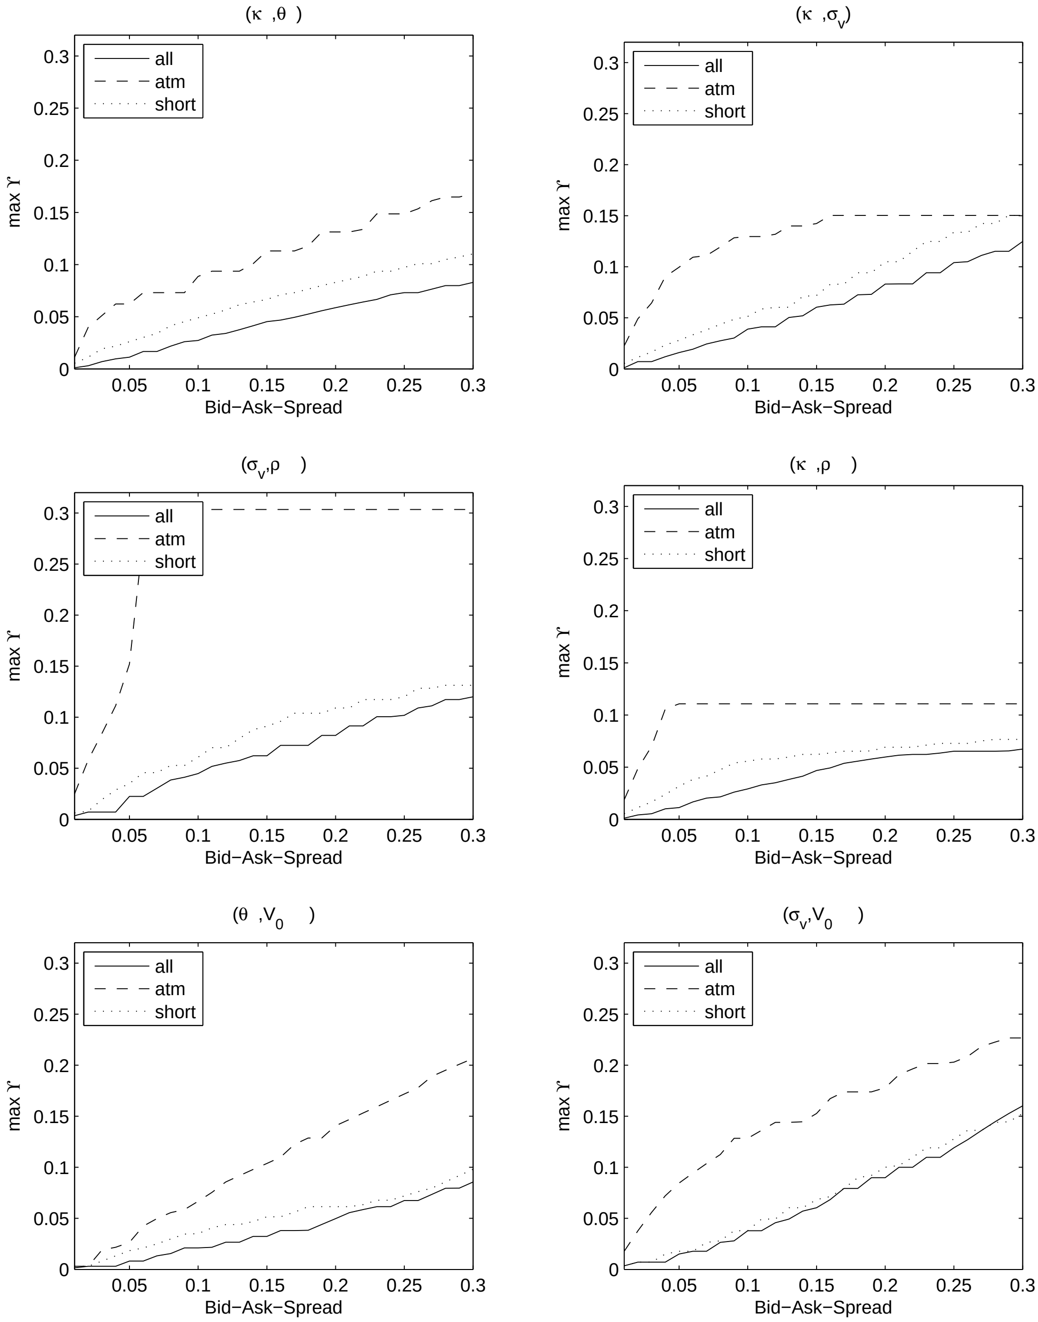


Figure 9: Bid-Ask Spread vs. Standard Deviation of Hedging Error

The graphs show the maximum of $\Upsilon_{(x,y)}$, i.e. the standard deviation of the hedging error, as a function of the mean bid-ask spread over all options used for calibration for different parameter combinations $(x, y)$. The solid lines represent the case where the model is calibrated to 'all' options, the dotted lines where it is calibrated to 'short' maturity options and the dashed lines where it is calibrated to 'atm' options only.