Master Degree in Economics and Finance

Master Project

Option Pricing in the Heston Stochastic Volatility Model: An Empirical Evaluation

Authors
Patrick Altmeyer
Jacob Daniel Grapendal
Makar Pravosud
Gand Derry Quintana

Supervisor
Prof. Eulàlia Nualart

June 2018, Barcelona

Option Pricing in the Heston Stochastic Volatility Model: An Empirical Evaluation

Barcelona GSE Master Project

Hit.Knit, a.k.a.;
Patrick Altmeyer
Jacob Daniel Grapendal
Makar Pravosud
Gand Derry Quintana

31 May, 2018

Abstract

This paper evaluates the calibration method of the Heston model presented by Alòs, De Santiago, and Vives (2015). We propose a slightly more efficient configuration of the optimisation procedure by introducing initial bounds on parameters. Using an extensive simulation study we generally obtain parameter estimates in agreement with true values. In an empirical application carried out onto high-frequency option data the calibration method fails to generate satisfactory results. Nonetheless, we conclude that similarly simple calibration methods as the one examined here should be used in combination with more sophisticated option pricing models.

1 Introduction

There exists a substantial body of literature concerned with the calibration of the Heston model for pricing financial options under stochastic volatility, many of which rely on computationally expensive algorithms. Alòs, De Santiago, and Vives (2015) propose a simple calibration method which considers only the three most critical regions of the implied volatility surface. Although their procedure is parsimonious and very easily implemented, their paper contributes to a model whose empirical applicability is contested. Here we therefore evaluate their model in an extensive numerical exercise as well as an application to real data. The paper is organized as follows. In Section 2 we briefly outline the theoretical background. Section 3 tests the results presented in Alòs, De Santiago, and Vives (2015) using a numerical exercise that is larger in scale and scope than the one in the original paper. Finally, in Section 4 we extend this framework to the empirical level using high-frequency option data. We find that whilst the calibration method has solid theoretical foundations and produces satisfactory estimation results within the Heston universe, inherent limitations of the latter disqualify the calibration for practical use.

2 Theoretical Framework

The framework discussed in this paper belongs to a school of pricing models that addresses a caveat of the seminal method of Black and Scholes (1973), namely its assumption that the volatility parameter of a financial asset remains constant across time. Although various configurations of stochastic volatility option pricing models exist in the literature, the common denominator is the introduction of a second nested Brownian motion that governs the movement of volatility in addition to the conventional stock price driving Brownian motion. The Heston model discussed here encompasses a mean-reverting stochastic volatility movement that is proportional to the square root of the underlying asset's variance. This section will present the mathematical framework of the Heston model, the closely related Lewis Approximation Formula as well as the approximation formula presented in Alòs, De Santiago, and Vives (2015).

2.1 The Heston Model

The stock price under the risk neutral measure evolves according the following dynamics

$$dS_t = S_t r dt + S_t \sigma_t \left(\rho dW_t + \sqrt{1-\rho} dB_t \right), \ t \in [0, T] \tag{1}$$

$$d{\sigma_t}^2 = \kappa \left(\theta - {\sigma_t}^2 \right) dt + v\sqrt{\sigma_t} dW_t \tag{2}$$

$$S_0 > 0 \tag{3}$$

$$\sigma_0 > 0 \tag{4}$$

where S_t is the stock price at time t, $r \geq 0$ is the constant instantaneous interest rate, σ_t is the current level of stochastic volatility which evolves according to the Cox, Ingersoll, Ross process, $\rho \in (-1, 1)$ is the correlation between the shocks to movements in stock prices and volatility, W_t and B_t are two independent Brownian motion processes, θ is the long term volatility, ν is the volatility of volatility, and κ is the mean reverting parameter term. The parameters need to satisfy that $2\kappa\theta > \nu^2$.

The price C_t of the European Call struck at K is defined as

$$C\left(S_t, v_t, \tau\right) = e^{-r\tau} \underset{t}{E}\left[\left(S_T - K\right)^+\right] = e^{x_t} P_1\left(x_t, v_t, \tau\right) - e^{-r\tau} K P_2\left(x_t, v_t, \tau\right) \tag{5}$$

where $\tau = T - t$ and $x_t = \ln(S_t)$. The second part of equation (5) is similar to the Black Scholes formula, where the terms P_1 and P_2 are the conditional probability of call option expiring at the money, given the current value of the stock price and current volatility.

Specifically, Heston (1993) defines P_1 and P_2 through Fourier Inverse Transformation as

$$P_j(x, v, T; \ln(K)) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \mathrm{Re}\left[\frac{e^{-\mathrm{i}\phi \ln(K)} f_j(x, v, T; \phi)}{\mathrm{i}\phi}\right] d\phi, \quad \text{for } j = 1, 2 \tag{6}$$

where the characteristic functions $f_j()$ take the following form:

$$f_j(x, v, T; \phi) = \exp[C(\tau; \phi) + D(\tau; \phi)v + \mathrm{i}\phi x] \tag{7}$$

with

$$C(\tau; \phi) = r\mathrm{i}\phi\tau + \frac{a}{\sigma^2}\left\{\left(b_j - \rho\sigma\phi\mathrm{i} + d\right)\tau - 2\ln\left[\frac{1 - ge^{d\tau}}{1 - g}\right]\right\} \tag{8}$$

$$b_1 = \kappa + \lambda - \rho\sigma, b_2 = \kappa + \lambda \tag{9}$$

$$D(\tau; \phi) = \frac{b_j - \rho\sigma\phi\mathrm{i} + d}{\sigma^2}\left[\frac{1 - ge^{d\tau}}{1 - g}\right] \tag{10}$$

$$g = \frac{b_j - \rho\sigma\phi\mathrm{i} + d}{b_j - \rho\sigma\phi\mathrm{i} - d} \tag{11}$$

$$d = \sqrt{\left(\rho\sigma\phi\mathrm{i} - b_j\right)^2 - \sigma^2\left(2u_j\phi\mathrm{i} - \phi^2\right)}, \; u_1 = \frac{1}{2}, \; u_2 = -\frac{1}{2} \tag{12}$$

where $\lambda = \lambda(x, v, t)$ represents the price of volatility risk. This parameter can be computed outside the model.

2.2 Lewis' Approximation Formula

The analytical solution to the Heston model is not easily implemented due to numerical issues. We therefore resort to one of the most commonly used approximation formulae to compute call prices. According with Lewis (2009) the price of the European call can be computed as

$$V(S,v,\tau) = S - Ke^{-r\tau}\frac{1}{\pi}\int_{\frac{i}{2}}^{\infty+\frac{i}{2}} e^{-iKX}\frac{\hat{H}(k,v,\tau)}{k-ik}\Bigg]dk \tag{13}$$

where

$$X = \ln\left(\frac{S}{K}\right) + r\tau \tag{14}$$

$$\hat{H}(\kappa,v,\tau) = \exp\left(\frac{2\kappa\theta}{\sigma^2}\left[qg - \ln\left(\frac{1-he^{-\xi q}}{1-h}\right) + \nu g\left(\frac{1-e^{-\xi q}}{1-he^{-\xi q}}\right)\right]\right) \tag{15}$$

$$g = \frac{b-\xi}{2} \tag{16}$$

$$h = \frac{b-\xi}{b+\xi} \tag{17}$$

$$q = \frac{\sigma^2\tau}{2} \tag{18}$$

$$\xi = \sqrt{b^2 + \frac{4\left(k^2-ik\right)}{\sigma^2}} \tag{19}$$

$$b = \frac{2}{\sigma^2}(ik\rho\sigma + \kappa) \tag{20}$$

Mrázek and Pospíšil (2017) show that formulae (13) and (5) are equivalents. They point out that the great advantage of Lewis' formula is that it is well behaved and only one numeric integration is required. We will use the Lewis formula to compute the call prices in the numerical exercise.[1]

[1]Note that in their calibration exercise Alòs, De Santiago, and Vives (2015) use an online tool that implements the analytical solution using a program made available by T. Kluge (2002). For a more extensive simulation study using this online calculator becomes infeasible. Using Lewis (2009) formula for pricing instead, we obtain calibrated parameters very close to those in Alòs, De Santiago, and Vives (2015).

2.3 Alòs, De Santiago, and Vives (2015) Approximation Formula

Alòs, De Santiago, and Vives (2015) show that the Heston formula for the European call price (Equation 5) can be approximated by

$$\tilde{C}_t = BS\left(t, X_t, \tilde{v}_t\right) + H\left(t, X_t, \tilde{v}_t\right) U_t + K\left(t, X_t, \tilde{v}_t\right) R_t \tag{21}$$

where $BS\left(\tau, X_t, \tilde{v}_t\right)$ is the Black Scholes formula using the usual inputs: current stock price S_t, time to maturity τ, risk-free interest rate r, $X_t = \ln\left[\frac{S_t}{K}\right]$, $x_t = \ln\left[S_t\right]$, and volatility $\tilde{v}_t$. This last term denotes the future expected volatility, defined as:

$$\sigma = \tilde{v}_t = \sqrt{\frac{1}{\tau}\int_t^T \underset{t}{E}\left[{v_s}^2\right]\, ds} = \sqrt{\theta + \frac{{v_t}^2 - \theta}{\kappa\tau}\left(1 - e^{-\kappa\tau}\right)} \tag{22}$$

The remaining terms are denoted as:

$$H(t, x, \sigma) = \frac{e^x}{\sigma\sqrt{2\pi\tau}}\exp\left[-\frac{{d_+}^2}{2}\right]\left(1 - \frac{{d_+}^2}{\sigma\sqrt{\tau}}\right) \tag{23}$$

$$U_t = \frac{\rho\sigma}{2\kappa^2}\left\{\theta\kappa\tau - 2\theta + {v_t}^2 + e^{-\kappa\tau}\left(2\theta - {v_t}^2\right) - \kappa\tau\mathrm{e}^{-\kappa\tau}\left({v_t}^2 - \theta\right)\right\} \tag{24}$$

$$K(t, x, \sigma) = \frac{e^x}{\sigma\sqrt{2\pi\tau}}\exp\left[-\frac{{d_+}^2}{2}\right]\frac{{d_+}^2 - \sigma d_+\sqrt{\tau} - 1}{\sigma\sqrt{\tau}} \tag{25}$$

$$d_+ = \frac{1}{\sigma\sqrt{\tau}}\left[\ln\left(\frac{S_t}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)\tau\right] \tag{26}$$

$$d_- = d_+ - \sigma\sqrt{\tau} \tag{27}$$

Alòs, De Santiago, and Vives (2015) prove that the approximation error with respect to the real value of the option is less than $\mathcal{O}\left(\sigma^2(\rho + \theta)\right)$. Therefore, the approximation is accurate for smaller values of the volatility. Furthermore, this semi-closed-form solution is cheaper in terms of computational costs and is used for tractable calibration of Heston parameters.

3 Numerical Exercise

3.1 Methodology

The calibration procedure used in this paper is similar to the one used by Alòs, De Santiago, and Vives (2015) with the slight difference that the optimisation used in this project imposes stricter parameter constraints as discussed further below. In general, the five unknown parameters of the Heston model are:

$$\Theta = \{\sigma_0, \kappa, \theta, \nu, \rho\}$$

Alòs, De Santiago, and Vives (2015) Taylor approximate implied volatility $I(T,K)$ and analyse its asymptotic behaviour. They show that when the option is at the money we have:

$$I(T,K) \approx \sigma_0 + \frac{3\sigma_0^2\rho\nu - 6\kappa(\sigma_0^2 - \theta) - \nu^2}{24\sigma_0}T \tag{28}$$

This implies that for $T = 0$ we have that $I(0,K) \approx \sigma_0$. Another of their contributions is that for short times to maturity we have:

$$I(T,K) \approx \sigma_0 - \frac{\rho\nu}{4\sigma_0}(x - \log K) + \frac{\nu^2}{24\sigma_0^3}(x - \log K)^2 \tag{29}$$

Finally, for long times to maturity the implied volatility of at-the-money European call options is given by:

$$I(T,K) \approx \sqrt{\theta}\left(1 + \frac{\nu\rho}{4\kappa} - \frac{\nu^2}{32\kappa^2}\right) + \left(\frac{\sigma_0^2 - \theta}{2\kappa\sqrt{\theta}} + \nu\rho\frac{\sigma_0^2 - 2\theta}{4\kappa^2\sqrt{\theta}} - \nu^2\frac{\sigma_0^2 - \frac{5}{2}\theta + 4\kappa}{32\sqrt{\theta}\kappa^3}\right)\frac{1}{T} \tag{30}$$

These three limiting regions are considered as the most critical areas to recover the whole volatility surface and lay the foundation for the calibration procedure. In order to estimate Θ, we follow three steps which will be briefly outlined in the following.

3.1.1 Step 1: OLS Regressions

We run OLS regressions based on Equation 28 to Equation 30 adding some restrictions on the data as shown below:

OLS equations	Restrictions on simulated data
$\beta_{1,1} + \beta_{1,2}\mathrm{T} + \epsilon_1$	T$\in$[0.01,0.2] and S_0=K e^{-rT}
$\beta_{1,1} + \beta_{2,2}\ln(\frac{S_0}{K}e^{-rT}) + \beta_{2,3}\ln(\frac{S_0}{K}e^{-rT})^2 + \epsilon_2$	T=0.01 and $\ln(\frac{S_0}{K}e^{-rT}) \in [0.8, 1.2]$
$\beta_{3,1} + \beta_{3,2}\frac{1}{T} + \epsilon_3$	T$\geq$3.5 and S_0=K e^{-rT}

3.1.2 Step 2: Matching

We exploit the relationship between the theoretical parameters in Equation 28 to Equation 30 and the estimated coefficients of the Step 1 as outlined below:

	Calibration steps	Theoretical value	Empirical value
(i)	Estimate σ_0	σ_0	$\hat{\beta}_{1,1}$
(ii)	Estimate $\nu\rho$	$-\frac{\nu\rho}{4\sigma_0}$	$\hat{\beta}_{2,2}$
(iii)	Given (i) and (ii), solve the system for $\{\kappa,\theta,\nu\}$	$\frac{-\kappa\left({\sigma_0}^2-\theta\right)}{4\sigma_0}+\frac{\nu\rho}{8}\sigma_0-\frac{\nu^2}{24\sigma_0}$ $\sqrt{\theta}\left(1+\frac{\nu\rho}{4\kappa}-\frac{\nu^2}{32\kappa^2}\right)$ $\frac{\sigma_0^2-\theta}{2\kappa\sqrt{\theta}}+\nu\rho\frac{\sigma_0^2-2\theta}{4\kappa^2\sqrt{\theta}}-\nu^2\frac{\sigma_0^2-2.5\theta+4\kappa}{32\kappa^3\sqrt{\theta}}$	$\hat{\beta}_{1,2}$ $\hat{\beta}_{3,1}$ $\hat{\beta}_{3,2}$
(iv)	Estimate ρ from (ii), given $\hat{\nu}$ and $\hat{\sigma_0}$	ρ	$-\frac{\hat{\beta}_{2,2}4\hat{\sigma}_0}{\hat{\nu}}\in$(-1,1)

3.1.3 Step 3: Optimisation

In the matching step we have three equations and three unknown parameters $\{\kappa,\theta,\nu\}$. However, this nonlinear system of equations has multiple solutions and optimisation is sensitive to the estimated coefficients, the resulting initial guesses as well as to the numerical method used. This introduces some complications with regard to some specifications of the true model parameters Θ. On the basis of this we have decided to introduce a few modifications and found that this has slightly improved the calibration.[2] In the following, we will look at the above-mentioned issues one by one.

3.1.4 Loss function

Alòs, De Santiago, and Vives (2015) proposed to optimise sum of squared residuals of the system of equations given in (iii). They use Solver to do so, given an initial guess for $\{\kappa,\theta,\nu\}$. In our case, instead of optimising the sum of square errors, we impose the following loss function to optimise the system:

$$L(\kappa,\theta,\nu)=\left(f^1(\kappa,\theta,\nu)+f^2(\kappa,\theta,\nu)+f^3(\kappa,\theta,\nu)\right)^2 \tag{31}$$

where

$$\begin{aligned}
f^1(\kappa,\theta,\nu)&=\frac{-\kappa\left(\hat{\sigma_0}^2-\theta\right)}{4\hat{\sigma_0}}+\frac{\hat{\nu\rho}}{8}\sigma_0-\frac{\nu^2}{24\hat{\sigma_0}}-\hat{\beta}_{1,2}\\
f^2(\kappa,\theta,\nu)&=\sqrt{\theta}\left(1+\frac{\hat{\nu\rho}}{4\kappa}-\frac{\nu^2}{32\kappa^2}\right)-\hat{\beta}_{3,1}\\
f^3(\kappa,\theta,\nu)&=\frac{\hat{\sigma_0}^2-\theta}{2\kappa\sqrt{\theta}}+\hat{\nu\rho}\frac{{\sigma_0}^2-2\theta}{4\kappa^2\sqrt{\theta}}-\nu^2\frac{\hat{\sigma_0}^2-2.5\theta+4\kappa}{32\sqrt{\theta}\kappa^3}-\hat{\beta}_{3,2}
\end{aligned} \tag{32}$$

can best be thought of as residuals. The advantage of the loss function given by Equation 31 is that it penalizes not only squared residuals, but also the cross products between them. Nonetheless, that function is still not convex and local minima can occur. Figure 1 illustrates this issue using simulated data: from left to right we fix one parameter at a time and observe that in two out of three cases finding a minimum for κ is not trivial given the shape of the objective function. As a consequence, bounds on this parameter should be quite strict.

[2]For a comparison of the approximation errors under weaker and stricter constraints please refer to Table 3 in the appendix

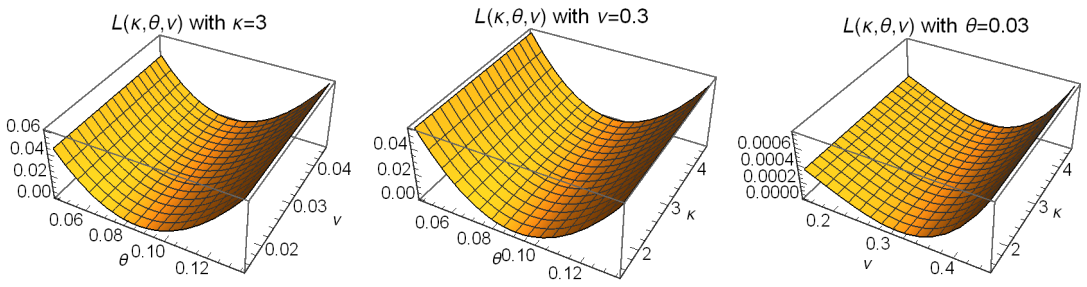

Figure 1: Loss functions of simulated data fixing one parameter at a time.

3.1.5 Initial Guess

We generally define initial values for $\{\kappa, \theta, \nu\}$ in the same way as Alòs, De Santiago, and Vives (2015), although we introduce a slight modification for ν_0: an initial high value for ν_0 is imposed whenever the estimated coefficient $\hat{\beta}_{2,3}$ is negative, since otherwise the optimisation becomes infeasible. We have encountered this issue for some parameter specifications used in the simulation study below. For the initial guesses we have therefore defined

$$\begin{aligned} \nu_0 &= \begin{cases} \sqrt{24\hat{\sigma}_0^3 \hat{\beta}_{2,3}} \text{ if } \hat{\beta}_{2,3} \geq 0 \\ 0.5 \text{ otherwise} \end{cases} \\ \theta_0 &= \hat{\beta}_{3,1}^2 \\ \kappa_0 &= \text{Solve}\left[\frac{-\kappa\left(\hat{\sigma_0}^2 - \theta_0\right)}{4\hat{\sigma_0}} + \frac{\hat{\nu\rho}}{8}\hat{\sigma_0} - \frac{{\nu_0}^2}{24\hat{\sigma_0}} = \hat{\beta}_{1,2}, \{\kappa\}\right] \end{aligned} \tag{33}$$

3.1.6 Bounds

There is no explicit need to impose parameter bounds in the optimisation. However, for certain values of Θ it fails to converge. Because of this we have imposed bounds on the search of optimal values based on initial guesses for parameters. Specifically we have defined:

$$\begin{aligned} \theta_0 &\in [0.85\theta_0, 1.05\theta_0] \\ \kappa &\in [\kappa_0, 1.05\kappa_0] \end{aligned} \tag{34}$$

These bounds are arbitrary and purely based on a guess and verify process. Furthermore, we have deemed it reasonable to impose $\rho \in [-1, 1]$. From this it follows:

$$\nu_0 \geq 4\hat{\sigma_0}\hat{\beta}_{2,2} \tag{35}$$

To conclude, the optimisation process can be summarized as solving

$$\left\{\hat{\kappa}, \hat{\theta}, \hat{\nu}\right\} = \text{argmin}\left\{L(\kappa, \theta, \nu)\right\} \tag{36}$$

subject to Equation 34 and Equation 35 using as initial guess values given by Equation 33[3].

3.2 Calibration Results

To gain a first impression of the calibration's accuracy we have computed absolute percentage deviations of the estimates from the true model parameters as in Alòs, De Santiago, and Vives (2015). Table 1 displays true and calibrated parameters as well as their absolute percentage deviations for different parameter specifications. In general, we observe that errors are highest for the volatility of volatility ν and lowest for σ_0, the initial volatility. We have split the simulation into three subsets where firstly we look at the non-correlated case, then we let only ρ vary and finally we let all parameters vary. Overall, average errors are roughly similar across these three subsets. Rows 3 and 8 correspond to the parameter choices in Alòs, De Santiago, and Vives (2015). A few relationships between the different parameters and their corresponding errors appear to emerge: firstly, for both ν and σ the calibration improves as the true supplied values increase; secondly, errors for θ and κ increase with the magnitude of their true values; finally, errors for ρ are highest around $\rho = 0$. Decreases in accuracy are likely driven by moving into complex regions of the loss functions in which numerical methods fail. Overall these dependencies on true parameters demonstrate that in reality the accuracy of the calibration may be compromised.

Besides looking at approximation errors of the calibrated parameters it should also be enlightening to analyse more closely in what specific areas the estimated implied volatility surface deviates from the true surface. For each of the parameter combinations in Table 1 we have firstly computed call prices under the Heston model over the strike range $K \in [95, 105]$ and maturity range $M \in [0.01, 4]$ using true parameters and then calibrated ones. Values for the implied volatility are then derived by inverting the Black-Scholes formula. We calculate the absolute percentage deviation of the true volatility surface under the Heston model and the surface obtained from using the calibrated parameter values. These approximation errors are plotted for each parameter specification in Figure 2. The dashed rectangle indicates where the option is at the money.

Note that deviations are generally minimized for small times to maturity and near the money. Even more so, for some specification errors also decrease for long times to maturity. These findings are intuitive, since the regressions used for the calibration primarily focus on these areas of the volatility surface and it should therefore not be surprising to obtain accurate results here. Finally, a satisfying result is that the overall percentage errors are small. Out of all simulations the maximum average error across the entire surface is 0.999 percent for specification 15. The mean error across simulations is 0.481 percent. Overall these results indicate that if we assume that the Heston model accurately describes option prices under stochastic volatility, the calibration by Alòs, De Santiago, and Vives (2015) works well despite depending on the given parameters to some extent.

[3]We have used Mathematica to perform the optimisations.

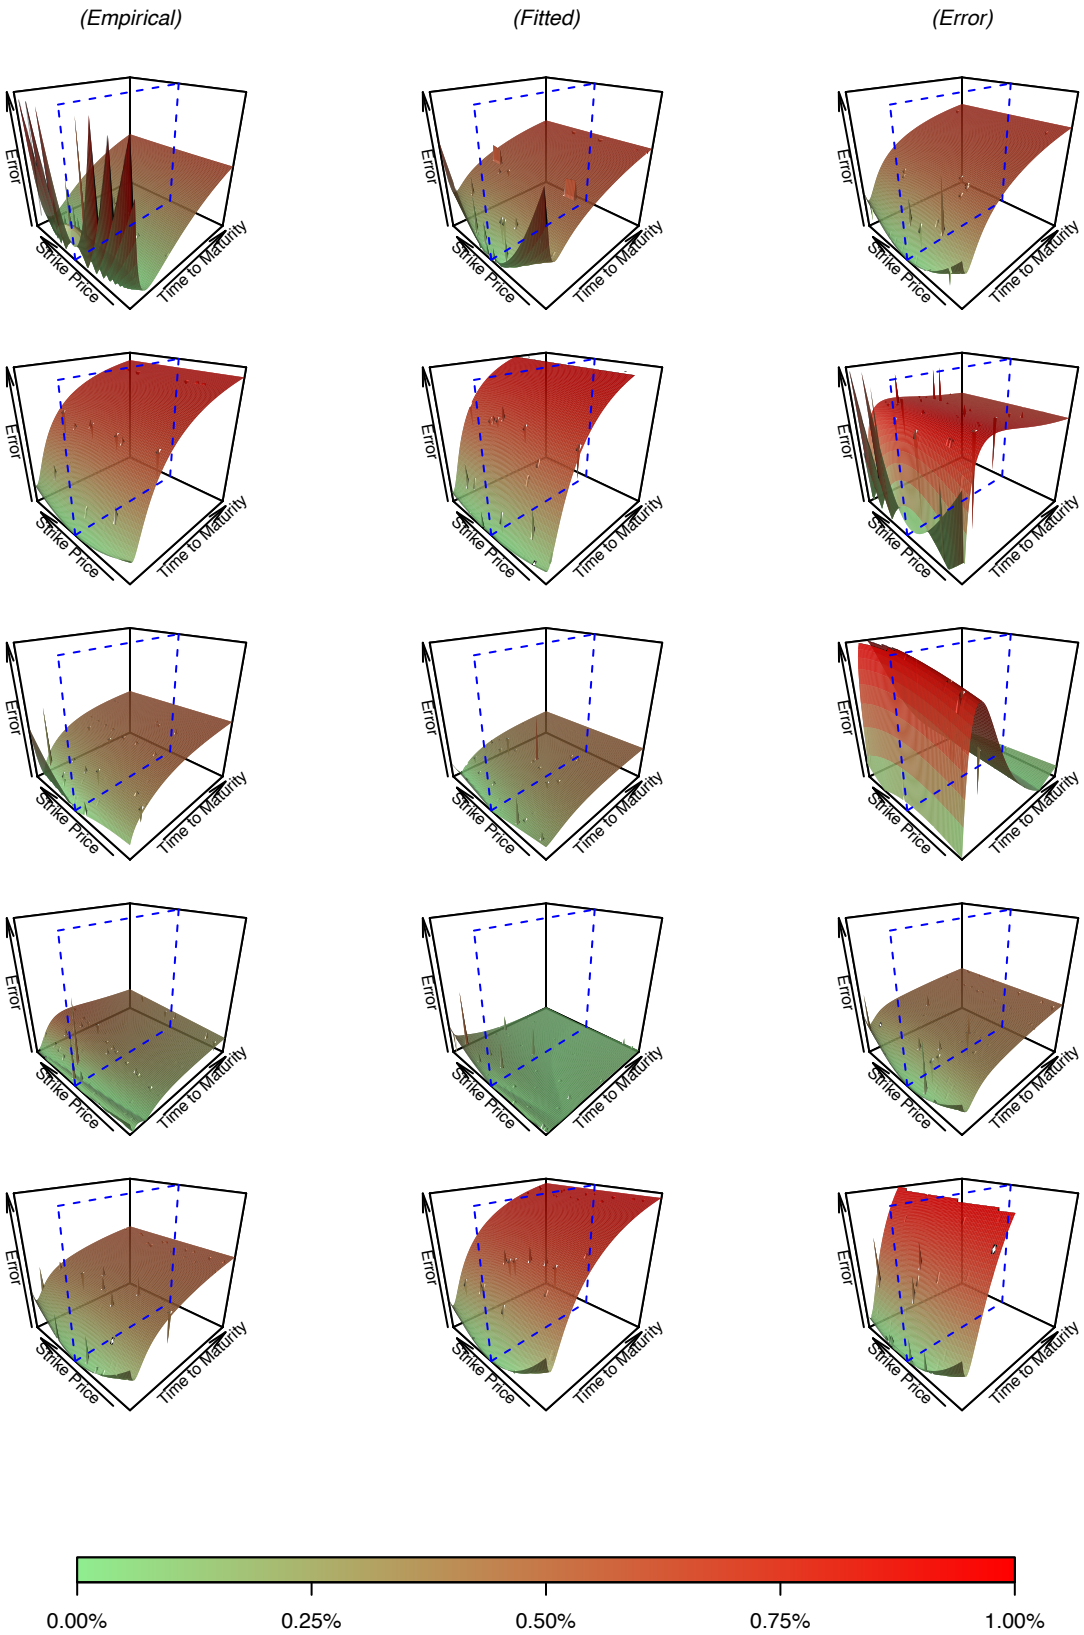

Figure 2: Absolute percentage difference between implied volatility surfaces derived from Heston call prices using true and calibrated parameters, respectively.

True parameters					Calibrated parameters					Absolute percentage errors				
ν	κ	θ	ρ	σ_0	$\hat{\nu}$	$\hat{\kappa}$	$\hat{\theta}$	$\hat{\rho}$	$\hat{\sigma}_0$	ε_ν	ε_κ	ε_θ	ε_ρ	ε_{σ_0}
0.200	2.000	0.080	0.000	0.100	0.290	2.001	0.081	0.000	0.100	45.244	0.053	1.787	0.000	0.225
0.250	2.500	0.085	0.000	0.150	0.326	2.533	0.086	0.000	0.150	30.438	1.308	1.648	0.000	0.023
0.300	3.000	0.090	0.000	0.200	0.380	3.056	0.092	0.000	0.200	26.728	1.882	1.891	0.000	0.007
0.350	3.500	0.095	0.000	0.250	0.433	3.576	0.097	0.000	0.250	23.805	2.173	2.225	0.000	0.001
0.400	4.000	0.100	0.000	0.300	0.483	4.125	0.102	0.000	0.300	20.842	3.121	2.479	0.000	0.002
0.200	2.000	0.080	-0.500	0.100	0.237	1.998	0.079	-0.410	0.100	18.316	0.083	1.121	17.982	0.216
0.250	2.500	0.085	-0.500	0.150	0.260	2.530	0.084	-0.464	0.150	3.905	1.187	1.087	7.212	0.025
0.300	3.000	0.090	-0.500	0.200	0.298	3.021	0.089	-0.491	0.200	0.569	0.705	0.627	1.874	0.008
0.350	3.500	0.095	-0.500	0.250	0.411	3.066	0.096	-0.418	0.250	17.497	12.386	1.297	16.414	0.002
0.400	4.000	0.100	-0.500	0.300	0.336	3.439	0.100	-0.588	0.300	16.099	14.036	0.078	17.538	0.001
0.300	3.000	0.090	-0.400	0.200	0.331	3.057	0.090	-0.353	0.200	10.381	1.901	0.062	11.630	0.007
0.300	3.000	0.090	-0.200	0.200	0.370	3.085	0.091	-0.158	0.200	23.202	2.827	0.965	20.852	0.007
0.300	3.000	0.090	-0.100	0.200	0.377	3.077	0.091	-0.078	0.200	25.643	2.568	1.400	22.397	0.007
0.300	3.000	0.090	0.100	0.200	0.378	3.023	0.092	0.077	0.200	25.869	0.759	2.373	22.537	0.007
0.300	3.000	0.090	0.300	0.200	0.356	2.918	0.093	0.246	0.200	18.774	2.735	3.383	17.895	0.007

Table 1: Calibration results for varying parameter values with $S_0 = 100$ and $r = 0$.

4 Empirical Application

Whilst the literature contains a range of studies aiming to validate the Heston model using empirical data, the general consensus of the model's efficacy remains contested. For instance, Drăgulescu and Yakovenko (2002) celebrate that their analytical solution of the Heston model using the Fokker-Planck equation tracks the Dow Jones index with considerable precision. However, this study is challenged by Daniel, Joseph, and Brée (2005), who in turn provide evidence that the accuracy of these results was driven by questionable data preprocessing. Furthermore, Daniel, Joseph, and Brée (2005) find that the Heston model only outperformed the constant volatility framework for small time horizons. Equivalently, Fiorentini, Leon, and Rubio (2002) apply the model to Spanish IBEX-35 options and conclude that Heston estimates only marginally outperform the the standard Black-Scholes framework, while similar results are presented by Bakshi, Cao, and Chen (1997) and Chernov and Ghysels (2000).

Bearing in mind the heterogeneity observed across the literature with regard to the practical utility of the model, the following section continues with an empirical extension of the calibration method presented by Alòs, De Santiago, and Vives (2015).

4.1 Data Collection

The regression structures in the calibration method used by Alòs, De Santiago, and Vives (2015) imply that an empirical extension requires a multidimensional option dataset spanning across both strike prices and maturities. Alike any quantitative study of equity options, we encountered stringent constraints in the supply of readily available data. In fact, sources of time-series data on different option contracts are restricted to either financial institutions or costly commercial platforms such as Bloomberg Terminal or Thomson Reuters. In order to obtain a rich dataset without spending thousands of dollars or compromising our limited connections in the banking industry, we confront this issue by employing a data mining strategy that is presented below.

The R library *quantmod* contains the useful function *getOptionChain* that queries the Yahoo Finance API for spot option prices on a selected ticker for all available maturities and contracts. Our configuration of this function, *optionQuery*, outputs a data table with an enriched variable structure including precise time to maturity (in fractions of years) as well as the prevailing spot price of the underlying (using quantmod's *getQuote* function). The output of our function is highly tractable for aggregation, as showcased below in Table 2.

ID	Strike	Timetomat	Last	Bid	Ask	Vol	Spot	TimeStamp
AAPL180720C00002500	2.50	0.15	185.13	185.75	186.45	1261	188.58	2018-05-26 14:03:09
AAPL181019C00002500	2.50	0.40	187.49	185.50	186.85	4	188.58	2018-05-26 14:03:09
AAPL190118C00042500	42.50	0.65	146.09	145.80	146.95	40	188.58	2018-05-26 14:03:09
AAPL200117C00050000	50.00	1.64	140.00	136.50	141.00	30	188.58	2018-05-26 14:03:10
AAPL180921C00075000	75.00	0.32	115.07	113.25	114.65	65	188.58	2018-05-26 14:03:09
AAPL190621C00085000	85.00	1.07	84.50	82.00	86.30	4	188.58	2018-05-26 14:03:09
AAPL200619C00085000	85.00	2.07	100.65	101.00	105.00	2	188.58	2018-05-26 14:03:10
AAPL180615C00100000	100.00	0.05	89.15	88.35	89.00	2	188.58	2018-05-26 14:03:09
AAPL180817C00100000	100.00	0.23	89.91	88.40	89.80	1	188.58	2018-05-26 14:03:09
AAPL181116C00110000	110.00	0.47	65.80	73.10	75.35	2	188.58	2018-05-26 14:03:09
AAPL180608C00130000	130.00	0.03	36.90	53.40	54.70	1	188.58	2018-05-26 14:03:09
AAPL180706C00180000	180.00	0.11	10.34	10.15	10.35	11	188.58	2018-05-26 14:03:09

Table 2: optionQuery function output.

The isolated use of *optionQuery* helps to little avail, however, seeing as it only provides spot prices. In order to create the actual time-series database, we have created an automated script that calls *optionQuery* every minute throughout NASDAQ's opening times (i.e. 09:30 to 16:00, U.S. Eastern Time). Data has been sampled between Tuesday, May 01 and Friday, May 18, 2018. We have selected Facebook Inc. (FB) seeing as it is normally traded at high volume. Furthermore, the volatility of this ticker should be of particular interest owing to Mark Zuckerberg's recent appearances in the U.S. Senate and the European Parliament, which occurred during the scraping period. Ultimately, the aggregated output of these sequentially produced data files generates a rich dataset that we use in the calibration of Heston's model. In total the data panel consists of 3720 minute-wise observations across 531 option contracts. For the computation of minute-wise implied volatilities one also needs information about the instantaneous interest rate. Historical data for the U.S. 10-year treasury note (TNX) is readily available at a daily frequency. To match it with the high-frequency option data we have interpolated the daily observations using a cubic spline.[4][5]

4.2 Calibration and results

To perform the calibration on empirical data, the general approach is the same as for simulated data. However, a number of complicating factors with respect to real data should be pointed out. Recall that these regressions rely on options falling into narrowly defined ranges for strike prices and times to maturities. The accumulated option chain data for Facebook call options contains 531 unique contracts with unique combinations of K and T, but for each regression only subsets of these contracts fall into the specified ranges. A principal concern was therefore to avoid running into small sample issues.

Consider for example the first regression for which Alòs, De Santiago, and Vives (2015) use simulated implied volatilities for at-the-money options with times to maturity $T \leq 1$. In practice, it is not possible to find options that are exactly at-the-money. Since at any given point t the spot price $S_t \in \mathbb{R}_+$ can take infinitely many possible values, the probability that it just matches the strike price $S_t = K$ is zero. To deal with this issue we have specified a certain range of options which are considered as being at-the-money. In particular, we have defined options with $K \in [S_t - 4, S_t + 4]$ as at-the-money options, which for the underlying Facebook data leads to including options with strike prices roughly 2 percent left and right of S_t on average. In terms of expiration times we have included all options with maturities falling into the smallest quartile across all expiration times within the option chain, so $T \leq T_{Q1}$. The second regression involves regressing implied volatilities on log-moneyness for short times to maturity. In order to preserve enough data we have included all options for which $\ln(\frac{S_t}{K}) \in [-0.3, 0.3]$, so thirty percent left and right of $S_t = K$ and only those which had the shortest time to maturity in our sample. For the third regression we need at-the-money options with long times to expiration. The filtration of at-the-money option was done in the same fashion as in the first regression. Long times to maturity have been defined as $T \geq T_{Q3}$, so we have included all options for which time to maturity falls into the third quarter of all times to maturity across the sample. Following these steps we have been able to obtain subsets large enough for

[4] We are happy to share our data which also contains data for Apple Inc. (AAPL) with the readers of this paper. The website https://griipen.shinyapps.io/bgse/ allows the user to browse and download daily aggregates of our data, and the full minute-wise dataset is available on request.

[5] If readers want to use our code for future analysis, our GitHub depository https://github.com/HitKnit/BGSE2018/tree/HitKnit-optionscraping contains scripts for minute-wise option scraping (prices.R) and subsequent data compiling (dataCompiler.R).

regression analysis. Below we firstly turn to intra-daily data before analysing the model's output across days. This should give us an idea of the calibration's robustness across time.

Similarly as in subsection 3.2 in this section we look at differentials between the empirical and calibrated implied volatility surfaces in order to evaluate the model's applicability in practice. For the intraday analysis we have computed hourly averages of the minute-wise data for Wednesday, May 16, 2018, that is we have divided the trading day into seven sessions. Secondly, we applied the calibration model to each interval and retrieved the estimated volatility surface.

In Figure 3 we depict in each row from left to right the empirical and calibrated volatility surface as well as their absolute percentage deviation. We present results for the opening session (10:00:00 to 11:00:00), the midday session (12:00:00 to 13:00:00) and the closing session (15:00:00 to 16:00:00).[6] Since the underlying stock price moves across the trading, so do the calibration and plot ranges for the strike price and times to maturity. The first observation one can make is that errors are generally larger for options close to maturity, a feature has previously figured in the literature (Janek et al. (2011)). In particular, for the opening and lunch session of this specific day the errors drastically surge for small times to maturity. The kinks observed on these plots coincide with times to maturity of around 6 months. Secondly, the error peaks appear to display inverse smile shapes, indicating that the the calibration exaggerates the smile around this region. However, for larger times to maturity errors are more homogeneous across strikes on this particular day. Furthermore, note that the empirical surfaces exert volatility spikes for in-the-money options close to maturity. This feature is persistent and is only partially captured by the estimated surfaces. In general though, the overall shape of the empirical surfaces is roughly reproduced by the calibration, although it should be noted that absolute percentage errors throughout the day lie between 0.017 and 45.589 percent. Finally, with respect to the robustness of the calibration, one can observe that the shape and magnitude of the error surfaces do not seem to depend on the time of day.

For the calibration at daily frequencies we have aggregated averages of all variables in the minute-wise data set for each trading day between Tuesday May 01, 2018 and Friday May 18, 2018. This yields data for 14 trading days across three weeks. In Figure 4 we have plotted volatility and error surfaces in the same way as above for Wednesday of each week.[7] Once again ranges for the strike price and times to maturity are slightly different for each plot to accommodate spot price movements. First of all, note that the error surfaces display larger heterogeneity across days vis-a-vis hours. Recall that for Wednesday, May 16, 2018 we already observed above that errors across strike prices are largely homogeneous. While this is still the case for aggregated data of this day, not all days exert this feature. Especially for May 09 but also May 16 the empirical surfaces dip for out-of-the money options close to expiration, a feature that is not captured well by the calibrated surfaces. We have observed that within each week these patterns are actually fairly consistent (see appendix).

Two general patterns emerge across all days in the sample. Firstly, errors are larger for small times to maturity: calibrated surfaces deviate on average around 16.488 percent from empirical surfaces for short times to maturity ($M < 0.5$) while for longer times to maturity errors are only on the order of 7.632 percent. It should be pointed out that this is exactly the opposite of what we observe for simulated data. Secondly, fitted surfaces consistently overestimate the smiles close to maturity: for at-the-money options ($K \in [0.98S_t, 1.02S_t]$) close to maturity ($M < 0.5$) the average estimation error is 18.291 percent, while for corresponding in- and out-the-money options the average estimation errors are only 14.152 and 16.261 percent, respectively.

[6] Plots for every single session of the day can be found in Figure 5 in the appendix.

[7] Plots for every single day can be found in Figure 6 to Figure 8 in the appendix.

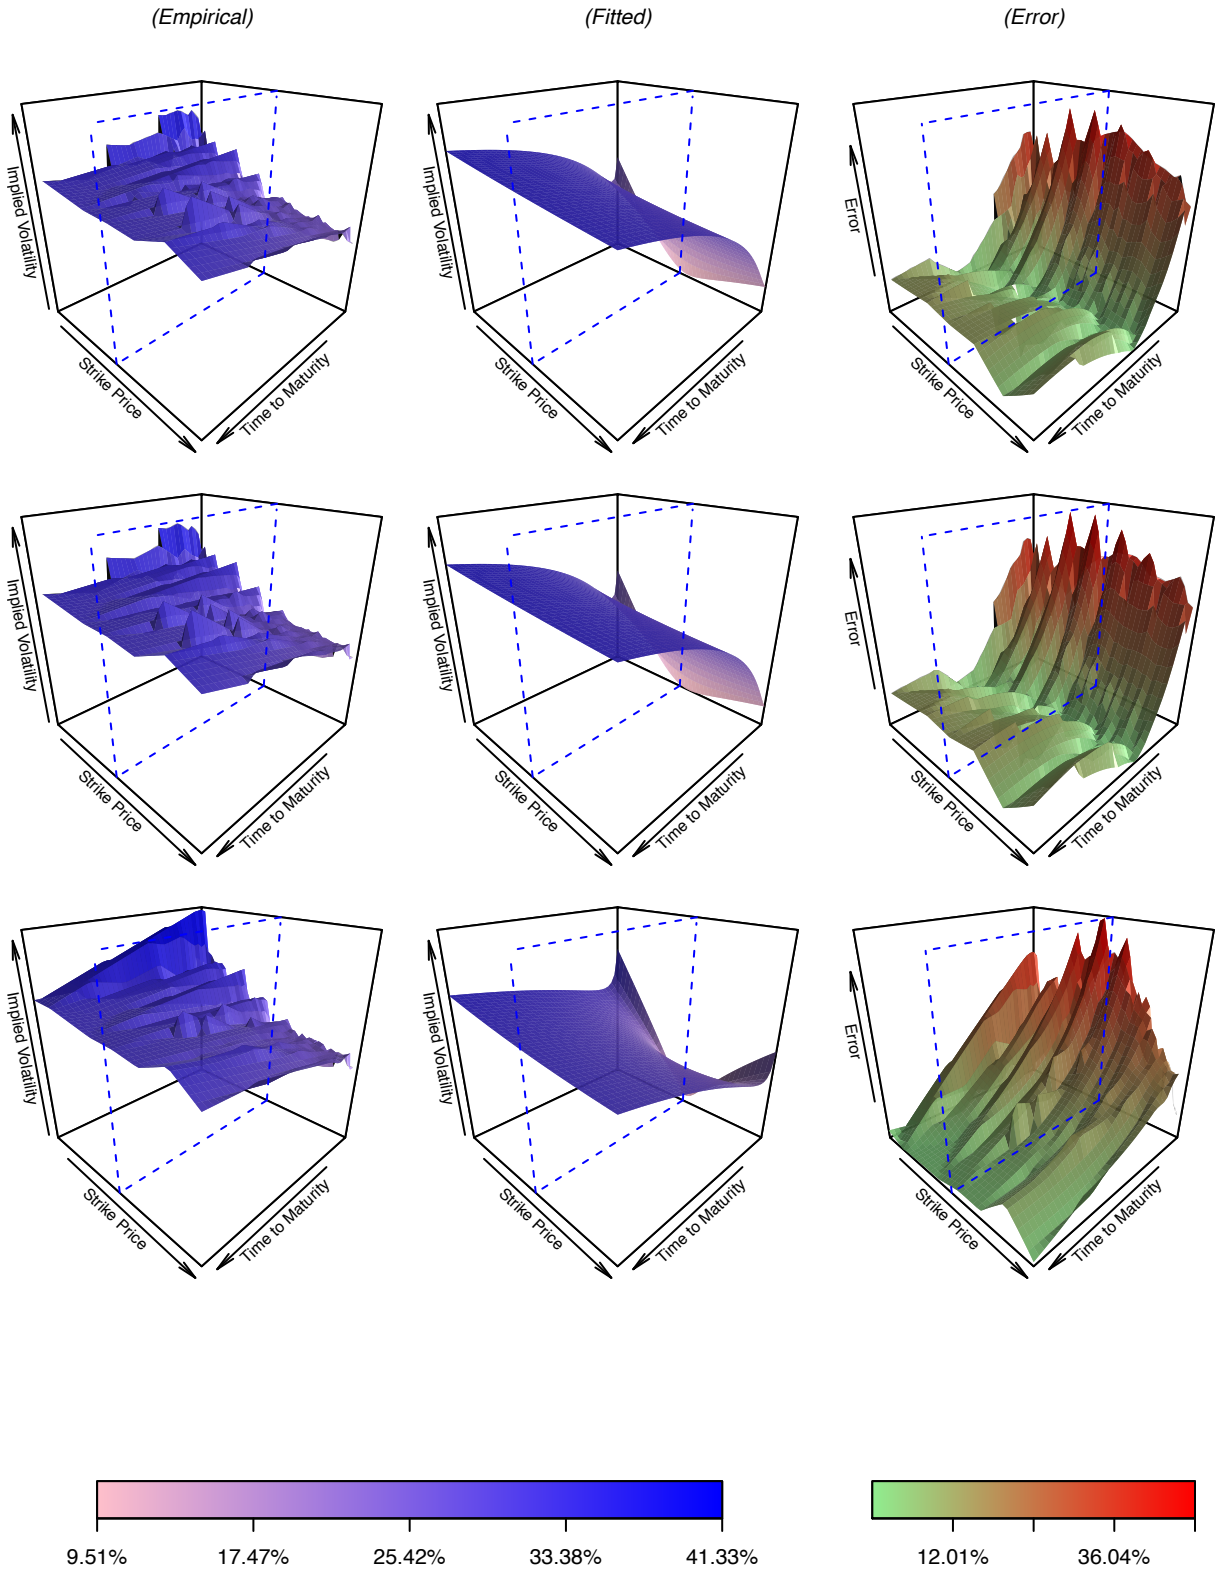

Figure 3: Comparison of empirical implied volatility surface and surface optained from using calibrated parameters for Facebook. From top to bottom volatilities are shown for the trading sessions from 10:00:00 to 11:00:00, from 12:00:00 to 13:00:00 and from 15:00:00 to 16:00:00 on May 16, 2018. The dashed frame indicates the spot price, hence options surrounding this frame are near the money.

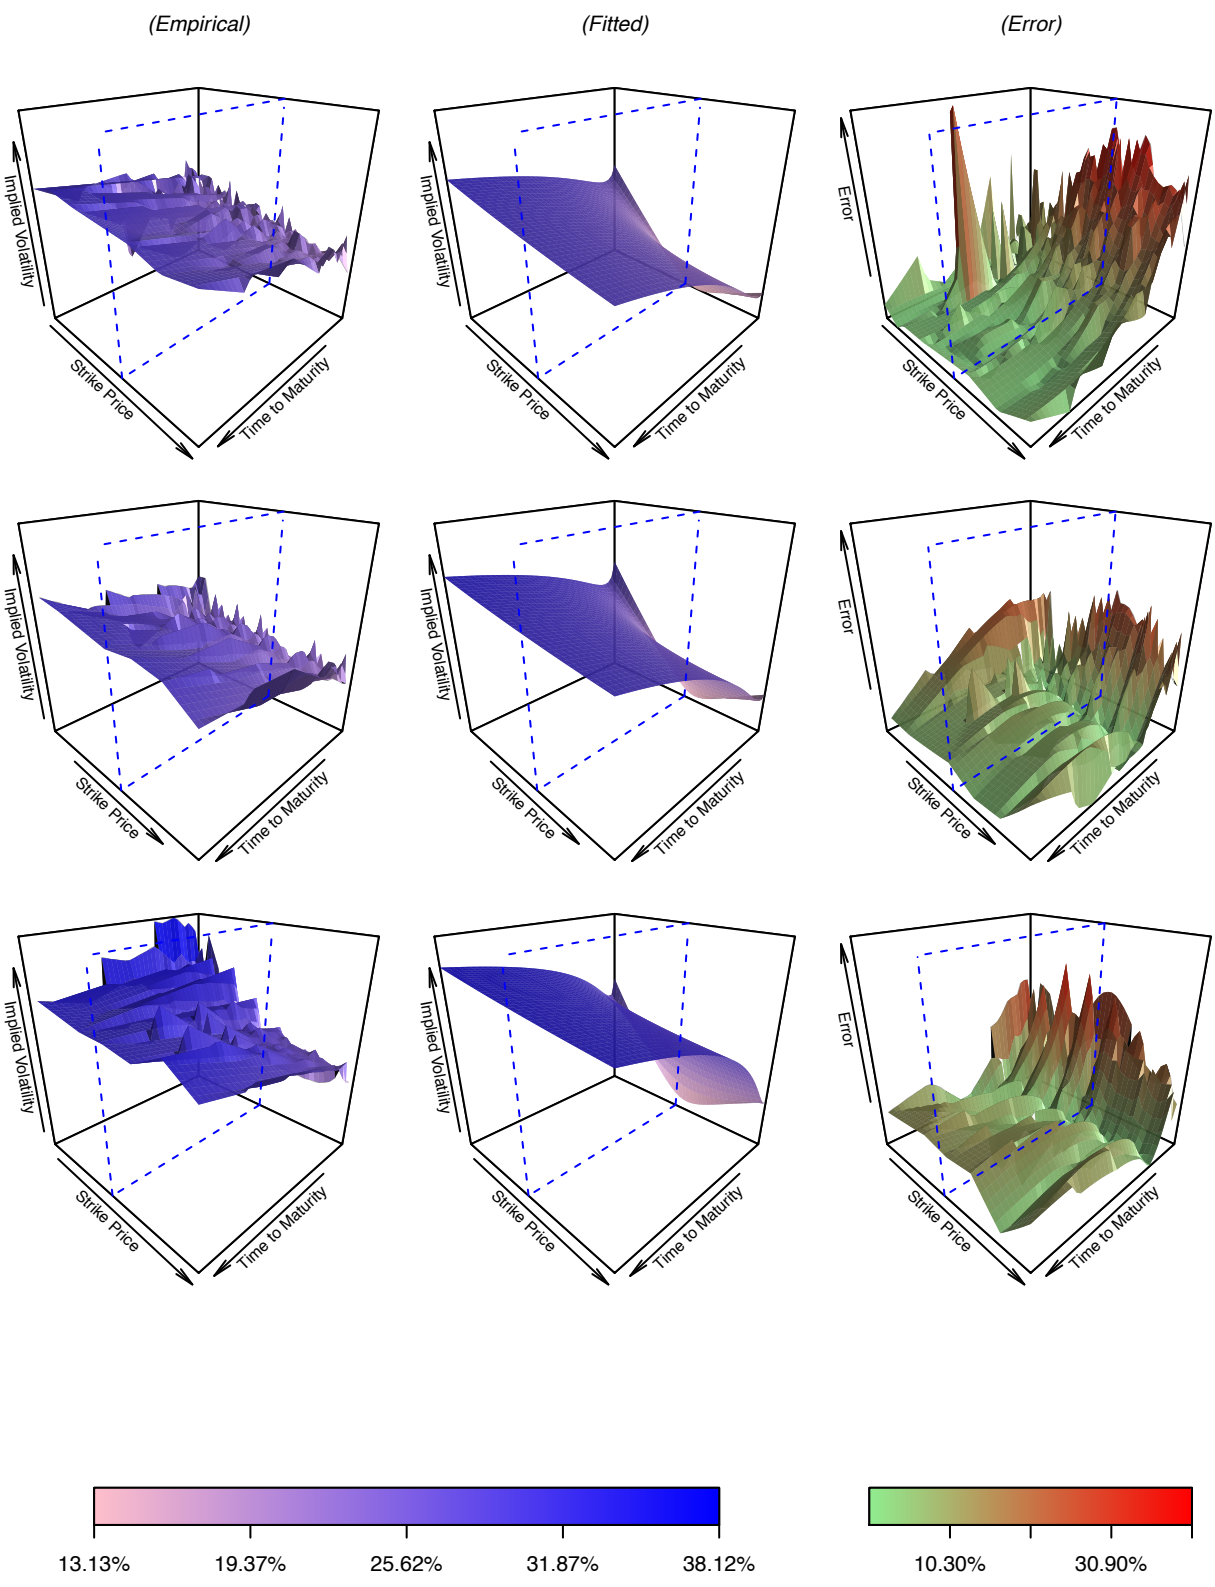

Figure 4: Comparison of empirical implied volatility surface and surface optained from using calibrated parameters for Facebook. From top to bottom volatilities are shown for May 02, May 09 and May 16, 2018. The dashed frame indicates the spot price, hence options surrounding this frame are near the money.

Overall the findings of this section demonstrate the practical shortfalls of both Heston model and ultimately also the calibration method proposed by Alòs, De Santiago, and Vives (2015), which relies on the validity of the former. Not only do we find approximation errors that are generally quite substantial, but also persistent across time. The fact that the approximation consistently generates error of roughly the same magnitude for the same options demonstrates that there is scope for improvement.

5 Conclusion

One primary aim of this paper has been to explore the Heston calibration method introduced by Alòs, De Santiago, and Vives (2015). For this purpose we have firstly extended the scale and scope of their simulation study. We have found that the performance of the calibration to some extend depends on the initial set of Heston parameters. This issue has been extensively discussed in previous literature and is likely due to the five-dimensional space in which calibration for the Heston model takes place, as this environment concedes little promise of a well-behaved objective function (see Cui, Baño Rollin, and Germano (2017) for a brief review). A comparison of the true and calibrated implied volatility surfaces based on simulated data within the Heston framework has further shown that the calibration is least accurate in those areas of the volatility universe where it does not actually operate, that is far away from the limits. Overall, however, the magnitude of errors is negligible implying that relying on a computationally simple calibration at the limits is not penalized.

A second goal of this paper has consisted in testing the calibration's potential value for practitioners. For this purpose we have gathered substantial amounts of recent real data on prices of numerous Facebook call options. Using the same methodology as for simulated data, we have found that unfortunately average approximation errors are large and disqualify the calibration from any reasonable prospect of practical application. In particular we observe that the shape of the differential surfaces is roughly consistent across time and hence approximation errors are predictable. This indicates that adaptations of the Heston model could avoid these errors. Methods that have already been used to address this include employing fractional stochastic volatility jump diffusion models (Pospíšil and Sobotka (2016)) and rough volatility models in general.

Overall, we conclude from our empirical study that Alòs, De Santiago, and Vives (2015) calibration inherits flaws from the Heston model. Nonetheless, the first part of our analysis has demonstrated that the idea of using rather simple calibration methods that rely on only a limited number of regressions is a fruitful ground for further research. Similar methods as the one proposed by Alòs, De Santiago, and Vives (2015) could be used to accommodate more involved extensions of the Heston model which more accurately describe the behaviour of option prices under stochastic volatility.

References

Alòs, Elisa, Rafael De Santiago, and Josep Vives. 2015. "Calibration of Stochastic Volatility Models via Second-Order Approximation: The Heston Case." *International Journal of Theoretical and Applied Finance* 18 (06). World Scientific: 1550036.

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen. 1997. "Empirical Performance of Alternative Option Pricing Models." *The Journal of Finance* 52 (5). Wiley Online Library: 2003–49.

Black, Fischer, and Myron Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81 (3). The University of Chicago Press: 637–54.

Chernov, Mikhail, and Eric Ghysels. 2000. "A Study Towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purpose of Options Valuation." *Journal of Financial Economics* 56 (3). Elsevier: 407–58.

Cui, Yiran, Sebastian del Baño Rollin, and Guido Germano. 2017. "Full and Fast Calibration of the Heston Stochastic Volatility Model." *European Journal of Operational Research* 263 (2). Elsevier: 625–38.

Daniel, Gilles, Nathan L Joseph, and David S Brée. 2005. "Stochastic Volatility and the Goodness-of-Fit of the Heston Model." *Quantitative Finance* 5 (2). Taylor & Francis: 199–211.

Drăgulescu, Adrian A., and Victor M. Yakovenko. 2002. "Probability Distribution of Returns in the Heston Model with Stochastic Volatility." *Quantitative Finance* 2 (6): 443–53.

Fiorentini, Gabriele, Angel Leon, and Gonzalo Rubio. 2002. "Estimation and Empirical Performance of Heston's Stochastic Volatility Model: The Case of a Thinly Traded Market." *Journal of Empirical Finance* 9 (2). Elsevier: 225–55.

Heston, Steven L. 1993. "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options." *The Review of Financial Studies* 6 (2). Oxford University Press: 327–43.

Janek, Agnieszka, Tino Kluge, Rafał Weron, and Uwe Wystup. 2011. "FX Smile in the Heston Model." In *Statistical Tools for Finance and Insurance*, 133–62. Springer.

Kluge, T. 2002. "Option Pricer Based on Heston's Stochastic Vol Model." http://kluge.in-chemnitz.de/tools/pricer/.

Lewis, Alan L. 2009. *Option Valuation Under Stochastic Volatility Ii*. Finance Press, Newport Beach, CA.

Mrázek, Milan, and Jan Pospíšil. 2017. "Calibration and Simulation of Heston Model." *Open Mathematics* 15 (1). De Gruyter Open: 679–704.

Pospíšil, Jan, and Tomáš Sobotka. 2016. "Market Calibration Under a Long Memory Stochastic Volatility Model." *Applied Mathematical Finance* 23 (5). Taylor & Francis: 323–43.

Appendix

True parameters					Strict constraints					Weak constraints				
ν	κ	θ	ρ	σ_0	ε_ν	ε_κ	ε_θ	ε_ρ	ε_{σ_0}	ε_ν	ε_κ	ε_θ	ε_ρ	ε_{σ_0}
0.200	2.000	0.080	0.000	0.100	15.539	0.334	1.120	0.000	0.225	66.269	1.473	5.903	0.000	0.225
0.250	2.500	0.085	0.000	0.150	4.745	1.021	0.928	0.000	0.023	42.236	0.215	3.894	0.000	0.023
0.300	3.000	0.090	0.000	0.200	0.245	1.615	1.003	0.000	0.007	35.763	1.069	3.490	0.000	0.007
0.350	3.500	0.095	0.000	0.250	4.104	1.938	0.950	0.000	0.001	30.908	1.600	3.373	0.000	0.001
0.400	4.000	0.100	0.000	0.300	7.435	2.998	0.866	0.000	0.002	2962.991	2307.081	246.570	0.000	0.002
0.200	2.000	0.080	-0.500	0.100	10.891	0.184	1.750	12.490	0.216	29.266	1.477	1.197	24.929	0.216
0.250	2.500	0.085	-0.500	0.150	0.190	1.134	1.401	3.771	0.025	4.529	0.171	0.359	7.765	0.025
0.300	3.000	0.090	-0.500	0.200	4.612	0.649	0.995	2.284	0.008	2.784	0.011	0.498	0.362	0.008
0.350	3.500	0.095	-0.500	0.250	14.888	12.400	0.918	14.516	0.002	6480.393	3208.703	764.502	98.508	0.002
0.400	4.000	0.100	-0.500	0.300	20.280	14.060	0.389	23.702	0.001	3414.517	2837.738	267.414	97.194	0.001
0.300	3.000	0.090	-0.400	0.200	6.155	1.834	1.703	3.941	0.007	13.185	1.119	0.634	13.819	0.007
0.300	3.000	0.090	-0.200	0.200	0.920	2.570	1.359	3.378	0.007	30.553	1.982	2.307	25.309	0.007
0.300	3.000	0.090	-0.100	0.200	0.976	2.312	1.223	3.439	0.007	34.462	1.741	2.951	27.487	0.007
0.300	3.000	0.090	0.100	0.200	2.780	0.490	0.698	0.290	0.007	34.549	0.025	3.925	27.535	0.007
0.300	3.000	0.090	0.300	0.200	11.702	3.003	0.221	10.443	0.007	23.955	3.432	4.429	21.327	0.007

Table 3: Approximation errors using strict and weak constraints with $S_0 = 100$ and $r = 0$.

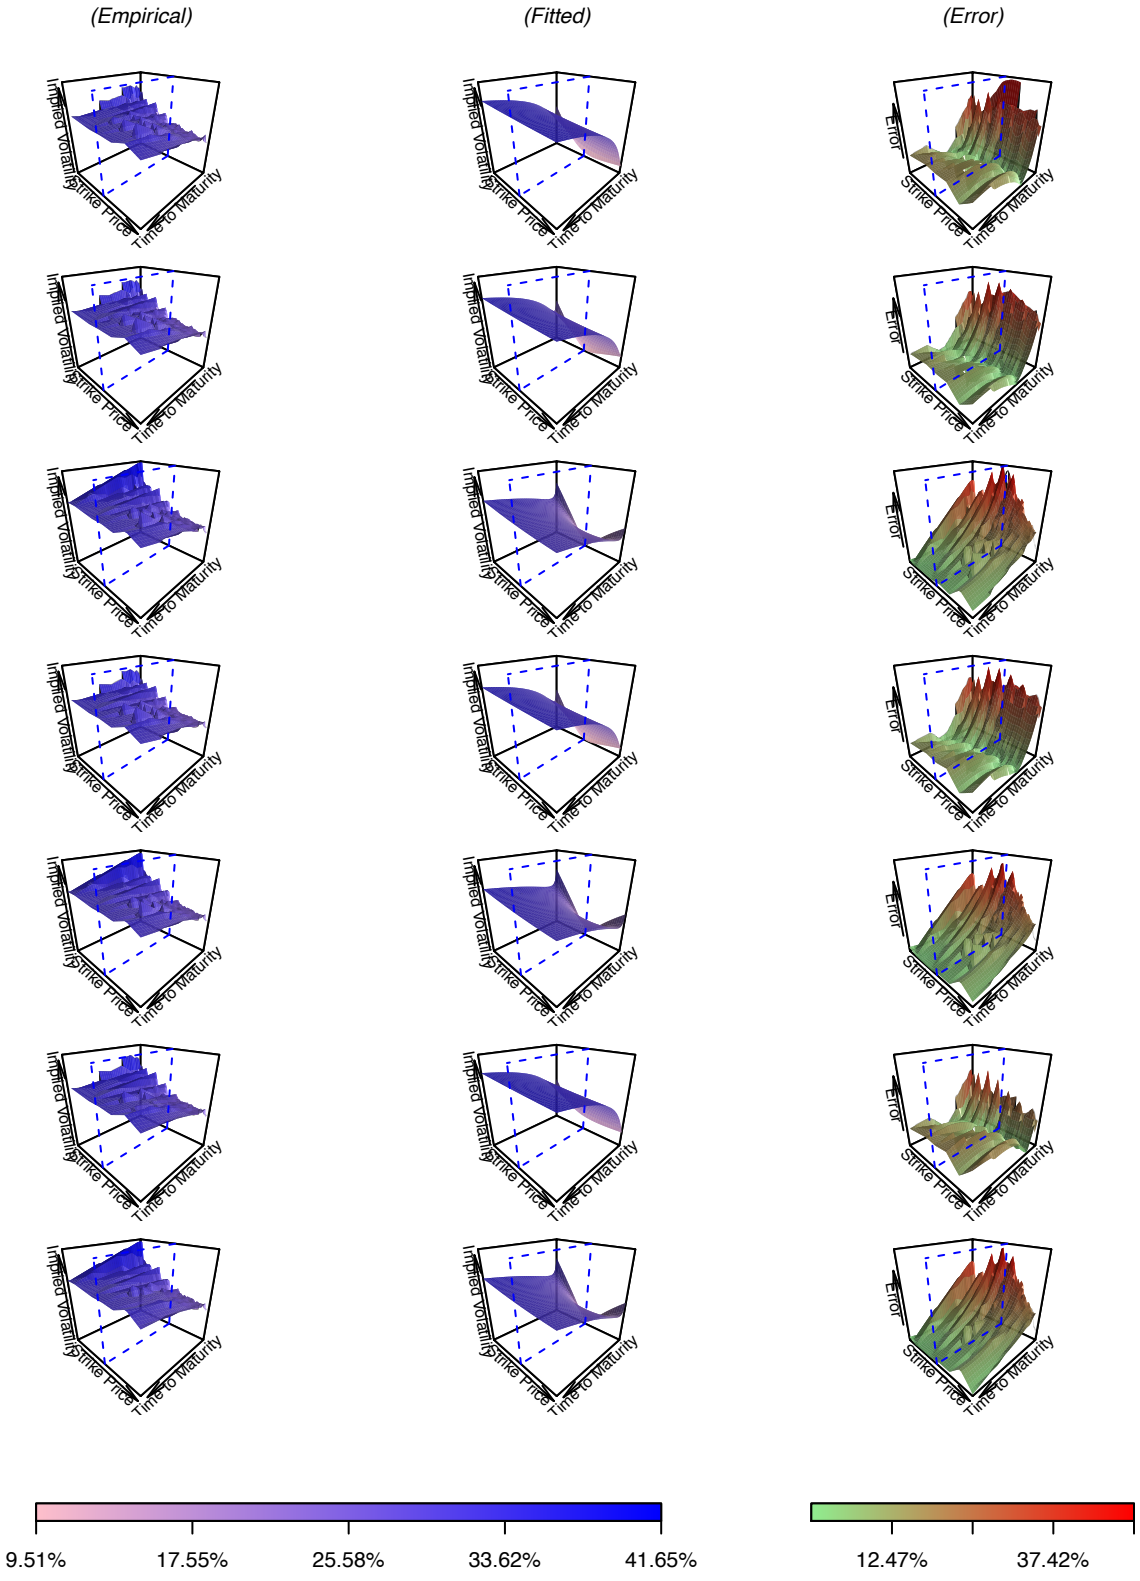

Figure 5: Volatility surfaces for all seven trading intervals on Wednesday, May 16, 2018.

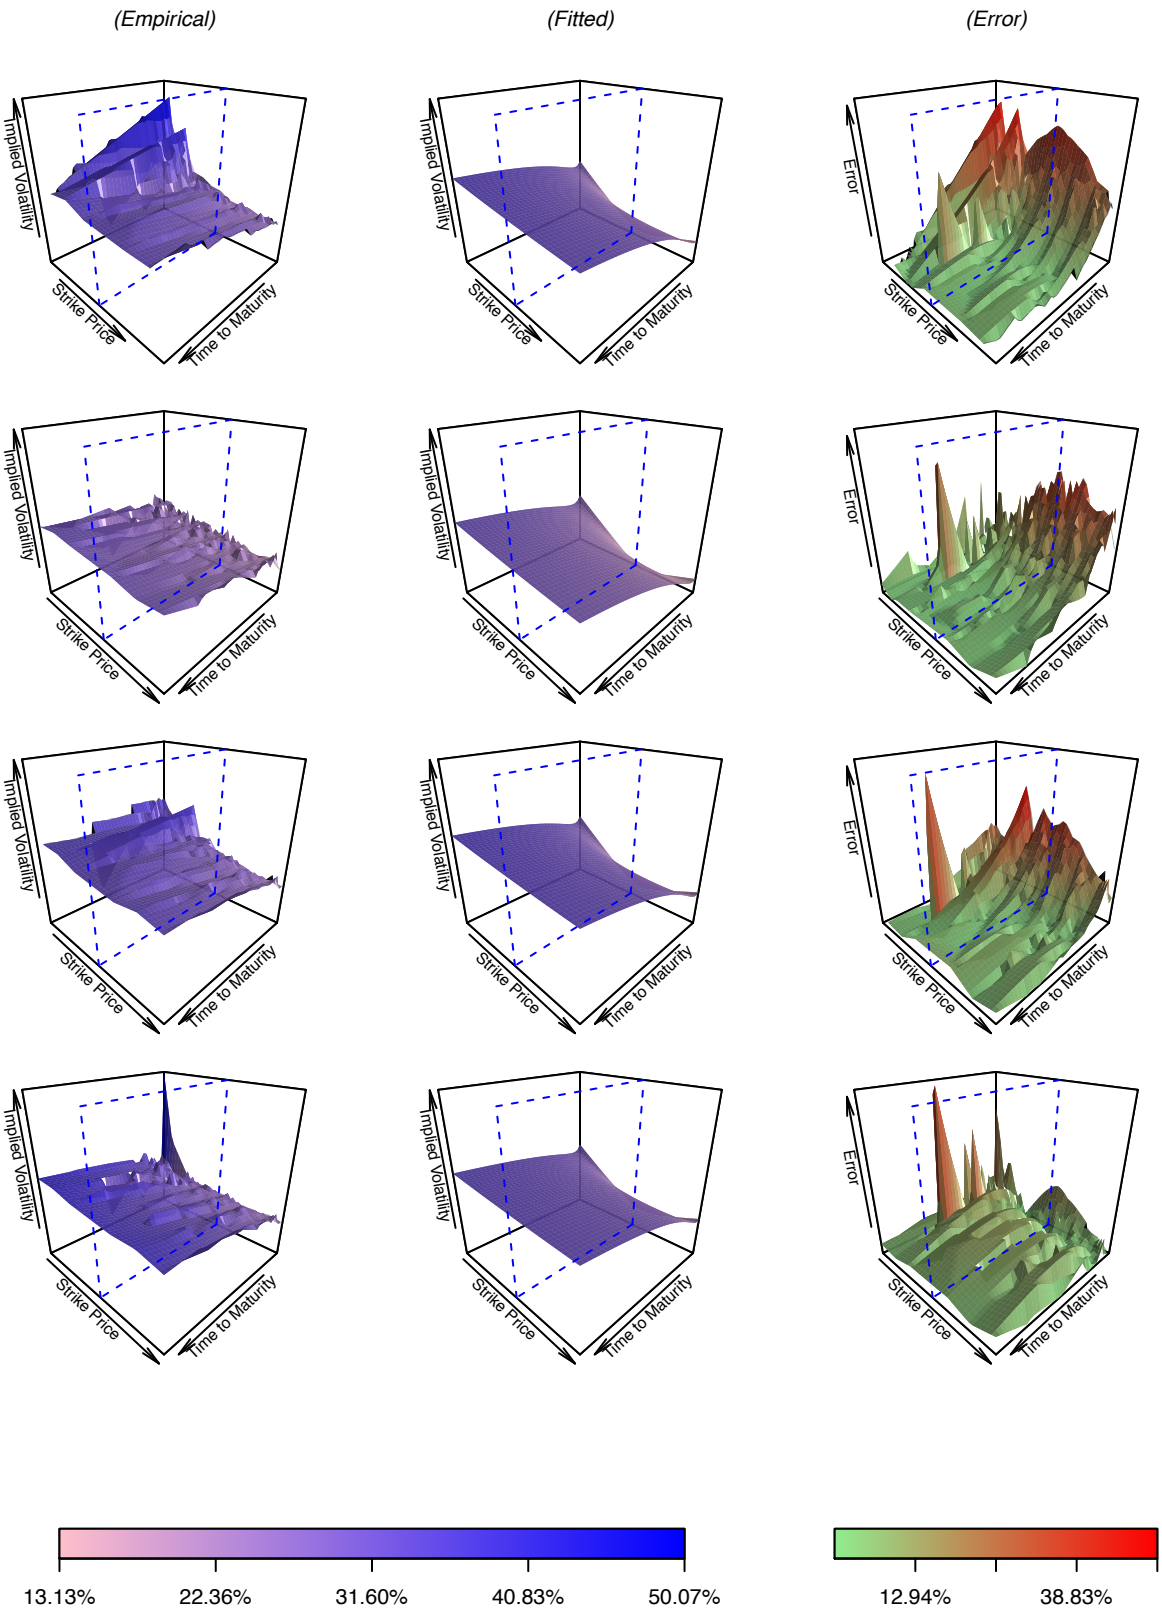

Figure 6: Comparison of empirical implied volatility surface and surface optained from using calibrated parameters for the first week from May 01, 2018 to May 04, 2018.

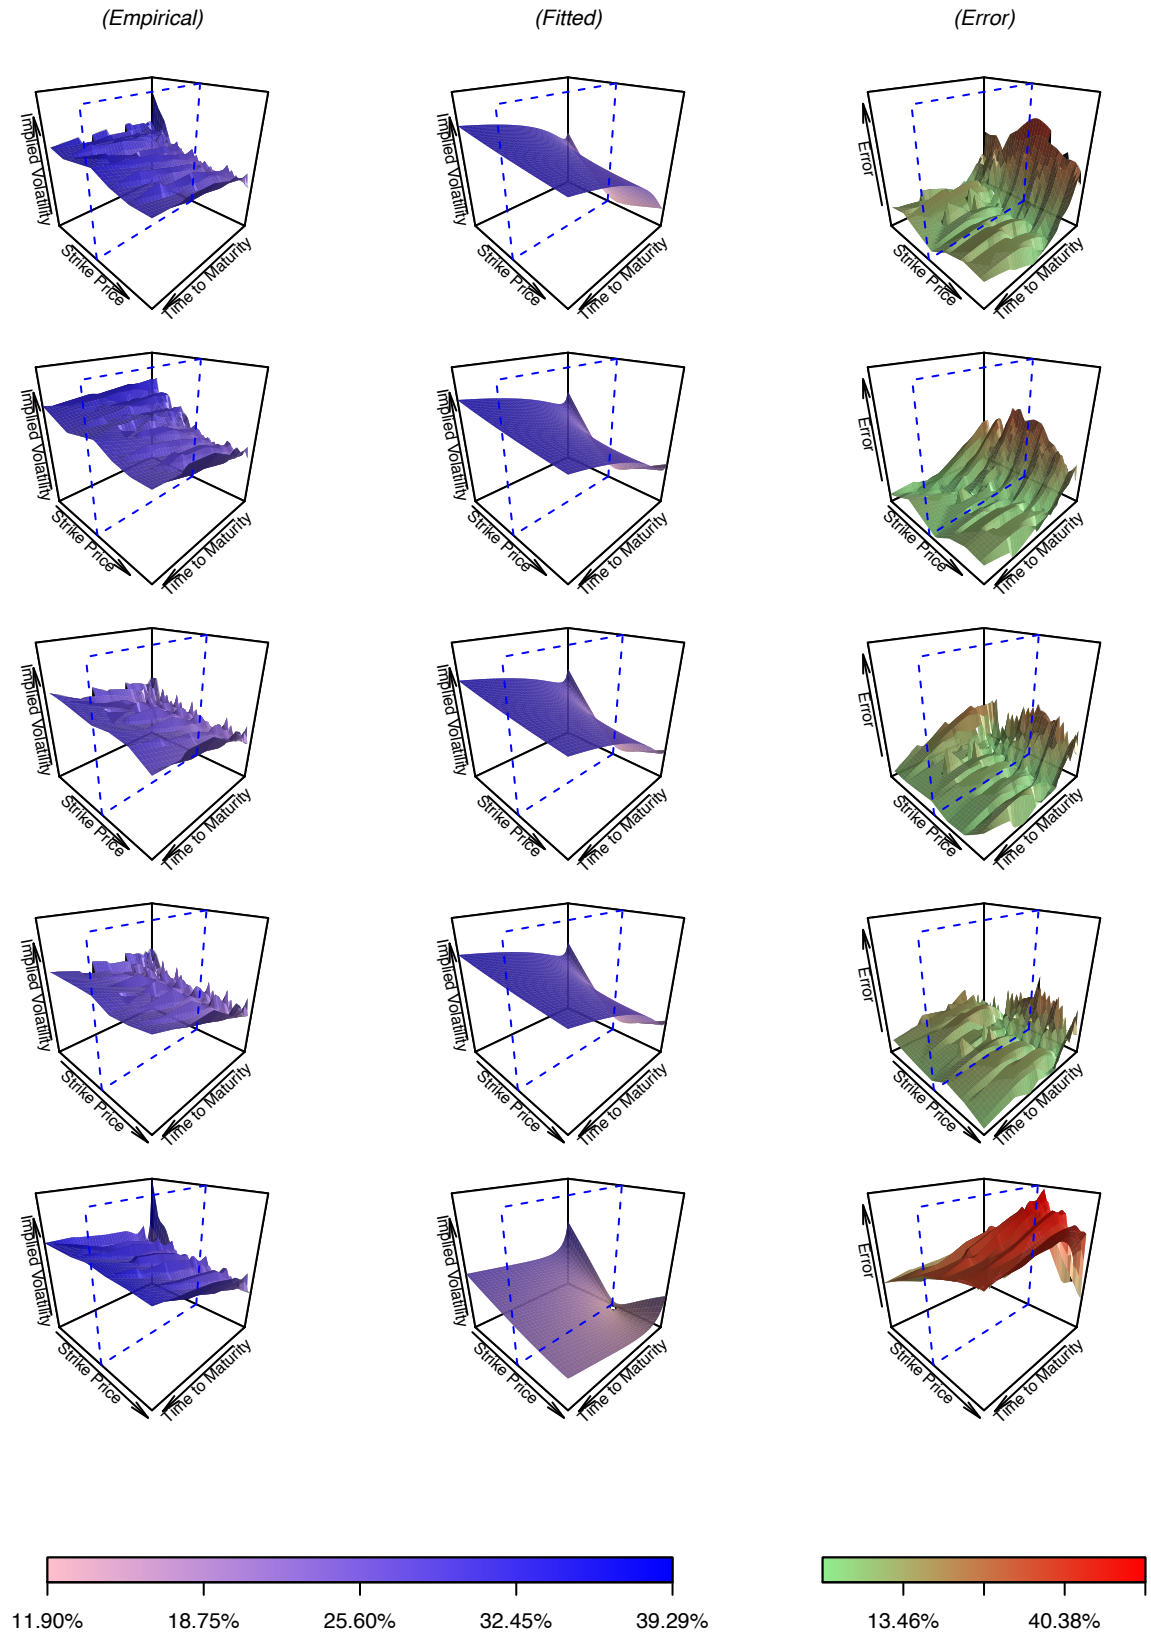

Figure 7: Comparison of empirical implied volatility surface and surface optained from using calibrated parameters for the second week from May 07, 2018 to May 11, 2018.

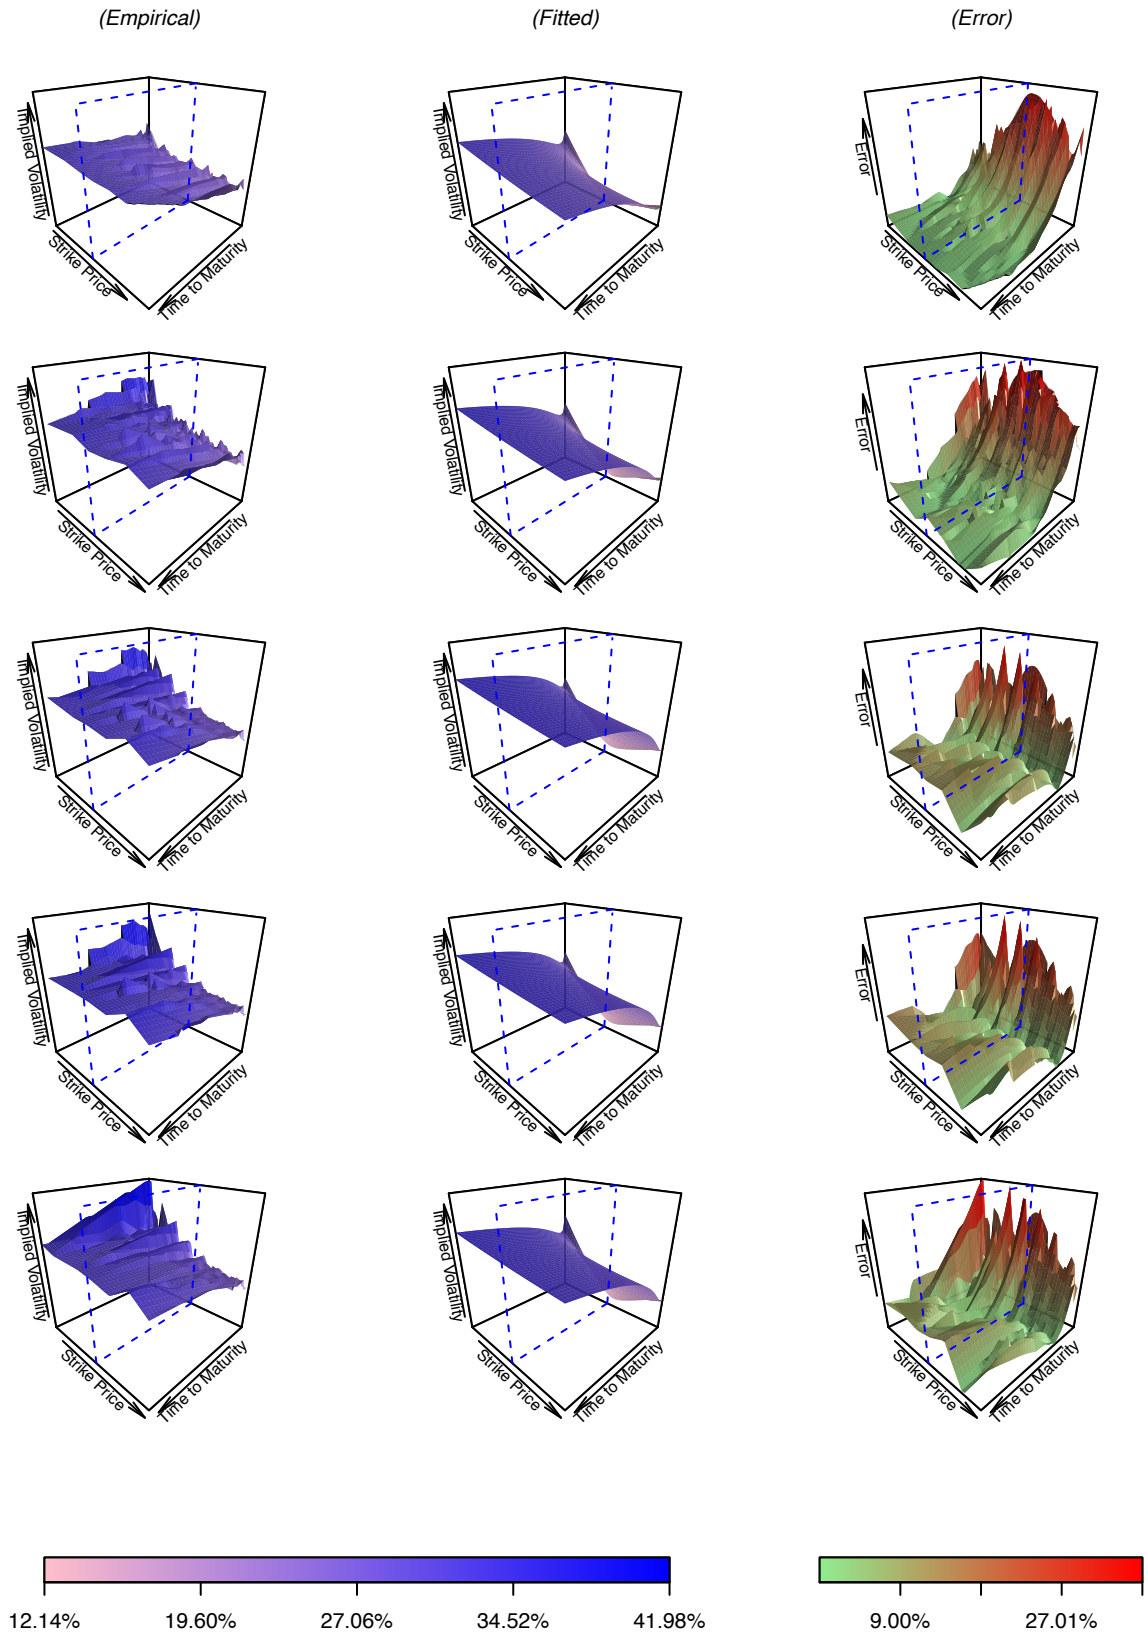

Figure 8: Comparison of empirical implied volatility surface and surface optained from using calibrated parameters for the third week from May 14, 2018 to May 18, 2018.