Estimation of Risk-Neutral Density Surfaces*

A. M. Monteiro† R. H. Tütüncü‡ L. N. Vicente§

Abstract

Option price data is often used to infer risk-neutral densities for future prices of an underlying asset. Given the prices of a set of options on the same underlying asset with different strikes and maturities, we propose a nonparametric approach for estimating risk-neutral densities associated with several maturities. Our method uses bicubic splines in order to achieve the desired smoothness for the estimation and an optimization model to choose the spline functions that best fit the price data. Semidefinite programming is employed to guarantee the nonnegativity of the densities. We illustrate the process using synthetic option price data generated using log-normal and absolute diffusion processes as well as actual price data for options on the S&P500 index.

We also used the risk-neutral densities that we computed to price exotic options and observed that this approach generates prices that closely approximate the market prices of these options.

1 Introduction

After the appearance of the Black-Scholes (BS) model in 1973, the option pricing theory has undergone a strong and sustained development. Given option market prices one can invert the BS formula in order to obtain the option implied volatility of underlying securities, that is, the BS formula can be used as an implied volatility calculator. While this approach is still commonly used in practice, due to the discrepancies observed between the constant volatility assumption of the BS model and the implied volatility observed from the market prices, namely the smile or skew effect, or the existence of a volatility term structure, there has been a constant stream

*Support has been provided by FCT under grants POCI/MAT/59442/2004 and PTDC/MAT/64838/2006.

†Faculdade de Economia, Universidade de Coimbra, Av. Dias da Silva, 165, 3004-512 Coimbra, Portugal (amonteiro@fe.uc.pt).

‡Goldman Sachs Asset Management (reha.tutuncu@gs.com).

§CMUC, Department of Mathematics, University of Coimbra, 3001-454 Coimbra, Portugal (lnv@mat.uc.pt).

of generalized models that try to better explain options prices. Continuous or jump diffusion models as well as stochastic volatility models have been developed to incorporate new evidence from the markets and gave rise to significant advances in the option pricing theory.

The estimation of the risk-neutral density or the implied volatility from option prices has its principal use in the pricing and hedging of other options, and as such it is mainly an interpolation tool. Several approaches have been presented in the literature to estimate risk-neutral densities. The existing approaches can be classified as parametric or non-parametric depending on the degrees of freedom and number of parameters needed to define the models (see the surveys [5, 8, 27]).

Market information such as the series of daily prices is commonly used to infer the set of parameters of an option pricing model. One strategy often used is to choose parameters that best fit the market prices. A common measure for the fitting is the least-squares distance between the market price and the theoretical price predicted by the model, given by $\sum_{i \in \mathcal{C}} (C_i(y) - C_i)^2$, where $C_i(y)$ is the option price computed by the model which depends on the parameters y, and C_i is the market price for option i. The parameters of the model can be determined by minimizing the measure used for the fitting, leading to a parameter estimation or inverse problem. The objective function of such optimization problems, as a function of the parameters y, is usually nonconvex (see, e.g., Hamida and Cont [22]) which can pose difficulties to optimization solvers.

Several authors proposed regularization methods in order to enforce well-posedness and overcome numerical difficulties, using smooth functions to minimize the distance between market prices and theoretical prices. Avellaneda et al. [4] recovered the implied volatility surface extending the one-period entropy minimization, presented by Butchen and Kelly [11] and Stutzer [34], to a multi-period model. Crépey [15] considered a minimization problem with a regularization parameter based on smoothness norms for volatility functions. This approach was also followed by Lagnado and Osher [29], and also by Jackson, Suli, and Howison [26] who used splines for the regularization effect. Coleman, Verma, and Li [12] considered a minimization procedure based on bicubic splines to ensure the desired smoothness of the implied volatility surface.

The method we propose here is closest to the approach outlined by Coleman, Verma, and Li [12]. Instead of trying to model the implied volatility surface we propose to directly invert the risk-neutral density surface. As in [12], we employ bicubic splines to obtain the desired density smoothness, and consider a least-squares objective function where the optimization variables are the parameters defining the bicubic splines used to describe

the risk-neutral density surface. Our optimization problem has a convex objective function and convex constraints.

A major difficulty in implied volatility or risk-neutral density estimation problems is the enforcement of the non-negativity of the estimated volatilities or densities. We explore two approaches. In the first case, we impose the nonnegativity of the recovered density functions only at the spline knots. This results in a problem with linear constraints, leading to a quadratic programming (QP) approach. In the second model, following an argument suggested for the single maturity case [33], we propose a semidefinite programming (SDP) approach, where, due to appropriate semidefinite and second order cone type constraints, one can rigorously ensure the nonnegativity of the recovered risk-neutral densities everywhere.

Exotic options are derivative contracts that offer nonstandard payoffs and allow a wider range of market strategies. Binary (or digital) options are an example of exotic options with discontinuous payoffs and are notoriously difficult to hedge. Cash-or-nothing calls, in particular, pay a fixed amount if the asset price, at maturity, is at or above the strike price, and nothing otherwise. We apply our estimation method for the risk-neutral density to price exotic options recently traded at the CBOE. The numerical results show that the prices of exotic options recovered from our risk-neutral estimation closely approximate the market prices of these securities and differ significantly from those obtained using the Black-Scholes model and implied volatilities. Finally, we remark that estimating the risk-neutral density surface does not provide a tool for pricing path-dependent options.

The rest of the paper is organized as follows. Section 2 describes the formulation of the estimation problem (the objective function, the constraints, and the QP and SDP approaches). In Section 3 we report our numerical experiments for the estimation of risk-neutral densities. Then, in Section 4, we use our risk-neutral estimation to price binary options. Some concluding remarks are made in Section 5.

2 Problem formulation using bicubic splines

2.1 Basic formulation of the problem

We start by presenting a basic optimization model to determine the risk-neutral density surface of an underlying asset from the known prices of options on this asset. Our goal is to compute a twice continuously differentiable bicubic spline function $p(\omega, u)$ representing the risk-neutral density function at state ω and time u. The inputs to the problem are a set of option exercise dates $\mathcal{T}$, for which one knows the sets of corresponding option prices

($\mathcal{C}_T$ and $\mathcal{P}_T$, $T \in \mathcal{T}$, for calls and puts respectively) and the corresponding sets of strike prices $\mathcal{K}_T$, $T \in \mathcal{T}$. We are given also interest rates related to the period from the initial time to the exercise date, which might vary across maturities. Based on the values for the strikes and on the range of terminal values for the underlying asset, we choose an interval $[a, b]$ which should contain all such values. We also select an interval $[c, d]$ containing the maturities. The set $[a, b] \times [c, d]$ forms the domain of the function $p(\omega, u)$ we are trying to determine.

The decision variables in our optimization model are the coefficients of the spline functions used to describe $p(\omega, u)$. To properly formulate the optimization problem we first need to specify a "super-structure" for the spline functions, namely, the number and location of the bicubic spline knots. We consider $n_s + 1$ asset knots in $[a, b]$, $a = \omega_1 < \omega_2 < \cdots < \omega_{n_s+1} = b$, and $n_t + 1$ temporal knots in $[c, d]$, $c = u_1 < u_2 < \cdots < u_{n_t+1} = d$. Asset knots are dependent on strikes or underlying asset prices but do not need to coincide with them. Similarly, temporal knots do not need to coincide with maturities. Thus, we partition the domain $[a, b] \times [c, d]$ into a $n_s \times n_t$ grid. For each rectangular region R^{st}, $s = 1, \ldots, n_s$, $t = 1, \ldots, n_t$ (defined by these two partitions), the corresponding bicubic spline element has 16 parameters. Therefore, the total number of spline parameters is equal to $16n_sn_t$. These parameters, denoted by $y \in \mathbb{R}^{16n_sn_t}$, amount to all optimization variables in the basic formulation. Let $p_y(\omega, u)$ represent the twice continuously differentiable bicubic spline. One must impose a number of constraints to represent all spline requirements (continuity up to second derivatives at the boundaries of R^{st}, $s = 1, \ldots, n_s$, $t = 1, \ldots, n_t$ and natural spline conditions). These constraints are linear in y and can be represented as follows:

$$(a^k)^\top y \;=\; b^k,\; k = 1, \ldots, n_c, \tag{1}$$

for appropriate vectors $a^k \in \mathbb{R}^{16n_sn_t}$ and scalars b^k, $k = 1, \ldots, n_c$. The number n_c, the total number of spline constraints, equals $(42n_s - 18)n_t + 21n_s - n_t - 17$.

Since $p_y(\omega, u)$ approximates a probability density function, it must satisfy

$$p_y(\omega, u) \;\geq\; 0,\; \forall (\omega, u) \in [a, b] \times [c, d]$$

and

$$\int_a^b p_y(\omega, u) d\omega \;=\; 1,\;\; \forall u \in [c, d].$$

These requirements can be imposed at the temporal knots by setting, respectively,

$$p_y(\omega, u_t) \;\geq\; 0,\; \forall \omega \in [a, b],\; t = 1, \ldots, n_t + 1 \tag{2}$$

and

$$\int_a^b p_y(\omega, u_t)d\omega \;=\; 1,\; t = 1, \ldots, n_t + 1. \tag{3}$$

The constraints (3) are linear in the variables of the problem. As in [33], where the authors studied the single maturity case, we will consider in the next two subsections two approaches to deal with the nonnegativity constraints (2).

Considering $p_y(\omega, u)$ as the risk-neutral density, the discounted expected value for future prices for an option corresponding to maturity $T \in \mathcal{T}$ and strike $K \in \mathcal{K}_T$ is given, respectively for calls and puts, by

$$\begin{aligned} C_{T,K}(y) &= e^{-r(T-t_0)} \int_a^b p_y(\omega, T)(\omega - K)^+ d\omega, \\ P_{T,K}(y) &= e^{-r(T-t_0)} \int_a^b p_y(\omega, T)(K - \omega)^+ d\omega. \end{aligned}$$

Thus,

$$[C_{T,K} - C_{T,K}(y)]^2 \quad \text{and} \quad [P_{T,K} - P_{T,K}(y)]^2$$

are the squared distances between the market and recovered prices. Summing up for all maturities and strikes, we obtain the following least-squares residual objective function

$$E(y) = \sum_{T \in \mathcal{T}} \left\{ \sum_{K \in \mathcal{C}_T} \theta_{T,K}[C_{T,K} - C_{T,K}(y)]^2 + \sum_{K \in \mathcal{P}_T} \mu_{T,K}[P_{T,K} - P_{T,K}(y)]^2 \right\}. \tag{4}$$

The weights $\theta_{T,K}$, $\mu_{T,K}$ may serve the purpose of balancing the relevance of each contribution under the presence of market information such as transaction volumes.

Leaving aside the nonnegativity constraints (2) momentarily, our basic formulation consists of the following quadratic programming problem (quadratic objective function and linear constraints):

$$\min_y \;\; E(y) \;\text{ s.t. }\; (1),\; (3). \tag{5}$$

We have pointed out before that all spline constraints (1) are linear. It is also easy to see that each term in the objective function is quadratic. In fact, let us look, for instance, at the formula for $C_{T,K}(y)$ for given $T \in \mathcal{T}$ and $K \in \mathcal{C}_T$. Assuming that $\omega_\ell \le K < \omega_{\ell+1}$, we get

$$\begin{aligned} & e^{r(T-t_0)} C_{T,K}(y) \\ &= \int_a^b p_y(\omega, T)(\omega - K)^+ d\omega \end{aligned}$$

$$= \sum_{s=\ell}^{n_s} \int_{\omega_s}^{\omega_{s+1}} p_y(\omega, T)(\omega - K)^+ d\omega$$
$$= \int_{K}^{\omega_{\ell+1}} p_y(\omega, T)(\omega - K) d\omega + \sum_{s=\ell+1}^{n_s} \int_{\omega_s}^{\omega_{s+1}} p_y(\omega, T)(\omega - K) d\omega.$$

Since $p_y(\omega, T)$ is a piecewise cubic polynomial whose coefficients are the variables y, one can clearly see that this expression is linear with respect to these variables. Once the integrals are calculated, one can see that these terms involve powers of the type ω^5 which will generate large values and thus introduce ill-conditioning in the Hessian matrix of the objective function $E(y)$. We will address this issue by scaling the variables as discussed in Section 3.

The nonnegativity constraint (2) for the density surface can either be imposed heuristically (without guarantees) or rigorously. These two alternatives lead to two different optimization problems that we explore in the next two subsections.

2.2 Quadratic programming approach

Our first approach to enforce nonnegativity of the estimated risk-neutral densities consists of imposing nonnegativity at all spline knots:

$$p_y(\omega_s, u_t) \geq 0, \ s = 1, \ldots, n_s + 1, \ t = 1, \ldots, n_t + 1. \tag{6}$$

This gives rise to a new QP formulation:

$$\min_y \ E(y) \ \text{ s.t. } (1), \ (3), \ (6). \tag{7}$$

The objective function is the same and the new constraints are linear inequalities in the optimization variables y. So, we are still in the presence of a convex QP, a relatively simple problem to solve. The disadvantage of the QP formulation is the possibility that a cubic polynomial function can be nonnegative at the end-points of an interval but negative at some of the inner points; this would not be appropriate for the density function we intend to estimate. Although many of our experiments with this formulation have resulted in nonnegative densities, we have observed in some instances that the densities recovered exhibited small negative values between the knots.

2.3 Semidefinite programming approach

Since our objective is to obtain nonnegative estimated risk-neutral densities, we suggest another approach based on semidefinite programming. The nonnegativity of the densities is based on a characterization of nonnegativity

for polynomials in real intervals, given by Bertsimas and Popescu [6]. To apply this characterization at the lines defined by the spatial and temporal knots, respectively, $\omega_s,\ s = 1, \ldots, n_s + 1$, and $u_t,\ t = 1, \ldots, n_t + 1$, we introduce new variables $X^\ell = [x^\ell_{ij}]_{i,j=0,\ldots,3}, \ell = 1, \ldots, \bar{n} = 2n_s n_t + n_s + n_t$. In what follows the symbol $X \succeq 0$ denotes that the matrix X is symmetric and positive semidefinite. According to the abovementioned characterization of nonnegativity for polynomials (see also [33]), the piecewise bicubic polynomial $p_y(\omega, u)$ is nonnegative in $[a, b] \times \{u_t\}$, for all $u_t,\ t = 1, \ldots, n_t + 1$, and in $\{\omega_s\} \times [c, d]$, for all $\omega_s,\ s = 1, \ldots, n_s + 1$, if and only if

$$\begin{aligned} & H^\ell_k \bullet X^\ell \ = \ 0,\ k = 1, 2, 3,\ \ell = 1, \ldots \bar{n}, \\ & (g^\ell_k)^\top y + H^\ell_k \bullet X^{st} = 0,\ k = 4, \ldots, 7,\ \ell = 1, \ldots \bar{n}, \\ & X^\ell \succeq 0,\ \ell = 1, \ldots \bar{n}, \end{aligned} \tag{8}$$

for appropriately chosen vectors $g^\ell_k,\ k = 4, \ldots, 7$, and matrices $H^\ell_k,\ k = 1, \ldots, 7$, for $\ell = 1, \ldots \bar{n}$, where the symbol $\bullet$ denotes the trace matrix inner product.

Note that because X^ℓ is positive semidefinite, we have that x^ℓ_{00} is nonnegative which then means, in the context of our notation, that the bicubic polynomial defined at the knots (ω_s, u_t) is always nonnegative. Thus, this approach includes the linear inequality constraints (6) imposed for the QP problem, yielding a stronger guarantee of nonnegativity. We can now pose a semidefinite program to recover the risk-neutral density surface, in the form:

$$\min_{y, X^1, \ldots, X^{\bar{n}}} \ E(y) \ \text{ s.t. } (1),\ (3),\ (8). \tag{9}$$

The constraints in (8) are linear in the optimization variables, except for the positive semidefiniteness constraints. Since all these constraints are convex and the objective function is quadratic the problem (9) is a convex semidefinite programming problem.

Semidefinite and conic optimization software typically solves problems with linear objective functions and thus we must provide a reformulation of problem (9) to be able to use available software options. It is easy to see that (9) is equivalent to the following problem:

$$\min_{t, y, X^1, \ldots, X^{\bar{n}}} \ t \ \text{ s.t. } t \geq E(y),\ (1),\ (3),\ (8). \tag{10}$$

Since $E(y)$ is a convex quadratic function of y, the constraint $t \geq E(y)$ can be reformulated as a second order cone constraint (see [31]). The resulting problem is still more complex than the QP but it can be efficiently solved with the current conic and semidefinite software. In our numerical experiments, we used the solver SDPT3 [35].

3 Numerical experiments

We now present a number of experiments for the approaches introduced in this paper to dynamically estimate the risk-neutral densities. To solve the convex QP problem (7) we used the MATLAB QUADPROG solver. To solve the SDP problem (9) (reformulated using the second order cone constraints) we applied the interior-point code SDPT3 [35]. Since the Hessian matrix is extremely ill-conditioned due to the magnitude of the powers of ω, we had to scale the QP and SDP problems, choosing the average value of the components of the vector of spatial knots as a scaling factor. We then solved scaled versions of problems (7) and (10) by scaling first each spline knot component by the average scaling value ω_{avg}. This led to new scaled variables related linearly to the original ones. The objective function was scaled by $1/\omega_{avg}^2$. The MATLAB version used was 7.5.0.342 (R2007b).

3.1 Black-Scholes data

In this first set of tests, we study the ability of our method to recover a known risk-neutral density based on the prices of options that are computed using this density. For this purpose, we generated option prices using the Black-Scholes model which corresponds to the choice of a log-normal risk-neutral density function. Using the function BLSPRICE, from the Financial Toolbox of MATLAB, we generated call option prices considering 7 maturities $\mathcal{T} = \{4/12, 5/12, 6/12,\ 8/12, 9/12, 11/12, 12/12\}$ and, for each maturity, 11 call options whose strikes are equally spaced between 20 and 100. We set $S_0 = 50$ as the initial market price for the underlying asset, 0.1 as the risk-free interest rate, a volatility of 0.2, and considered no dividend rate. The log-normal risk-neutral density that corresponds to these settings is plotted for $t \in (4/12 : 0.002 : 1)$ (see the plot at the bottom of Figure 2).

Following [3, 12] we considered a region of the domain centered around S_0 that is appropriate for the choice of the spline knots. Our spline knots formed a non-uniform rectangular mesh. There were 40 asset knots varying from 20 to 100. Since the risk-neutral density is more nonlinear (has greater curvature) around the current prices, we prefer to put more knots around S_0, to capture this nonlinearity with better precision. Thus, the number of asset knots near the at-the-money locations was higher and corresponded to half of the total number of asset knots, to account for the importance of the at-the-money options. These asset knots were constant in time and did not coincide with strike prices. The temporal knots set chosen was $\{4/12, 10/12, 12/12\}$.

We solved the scaled formulation for problems (7) and (10) and obtained the risk-neutral densities (see Figure 2). For this data set, we did not consider any weights $\theta_{T,K}$ and $\mu_{T,K}$. Since the Hessian of $E(y)$ is strongly

rank-deficient (and despite the fact that it is theoretically positive semidefinite), we observed numerically the presence of small negative eigenvalues. The MATLAB QP solver encountered some difficulties due to these negative eigenvalues. The scaling of the Hessian reduced the ill-conditioning of the matrix, but we still had to perturb the Hessian by adding a term of the form ξI, where $\xi = |\lambda_{min}|$, which guaranteed a numerically positive semidefinite Hessian.

The estimated call prices, computed using the recovered risk-neutral density, very accurately estimate the "true" Black-Scholes prices for all seven maturities (see Figure 1). The residual $E(y)$, obtained for the sum of the seven maturities, and the average absolute error per option are similar for QP and SDP approaches as we can see from Table 1. However, we found the relative errors averaged over the whole set of strike prices substantially higher than the averages for at-the-money options mostly because of the difficulty of accurately estimating prices of very cheap, far-out-of-the-money options. With very small true prices, even small price deviations cause large relative errors. Imposing a floor of 0.01 on theoretical and estimated prices brings the overall averages down to values comparable to those achieved by at-the-money options. Thus, the average relative error per option (reported in the third line of Table 1) was computed as

$$\frac{1}{|\cup_{T\in\mathcal{T}}\mathcal{C}_T|}\left(\sum_{T\in\mathcal{T}}\sum_{K\in\mathcal{C}_T}\frac{|\max(C_{T,K},0.01)-\max(C_{T,K}(y),0.01)|}{\max(C_{T,K}(y),0.01)}\right). \quad (11)$$

The last line of Table 1 lists the average relative error per option considering only those options for which the corresponding strike prices belong to the interval $[0.9S_0, 1.1S_0]$. In addition, in Table 2, we broke down per maturity the average absolute and relative errors per option reported before in Table 1. Due to the imposition of the floor of 0.01, there are no significant differences among the numbers across different maturities.

In terms of the recovery of the risk-neutral densities, both the QP and SDP approaches performed very well for all maturities (see the upper plots of Figure 2). Since we know the log-normal distributions represent the "true" risk-neutral densities for this experiment, we plotted them against the recovered risk-neutral densities for the QP and SDP approaches (see the plots at the middle of Figure 2). While the fit is very good, the recovered density deviates slightly from the true density for the first maturity, which, again, is due to the very cheap, far-out-of-the-money options.

The dimensions of the QP problem solved are the following: 1248 variables, 2424 equality constraints, and 120 inequality constraints. The SDP problem, reformulated using the second order cone (SOC) constraints, had 1248 linear variables, 197 SDP variables, 1249 SOC variables, 2424 linear

BS data set	QP	SDP
residual $E(y)$	3.30×10^{-3}	2.73×10^{-3}
average abs. error per option	3.94×10^{-3}	3.43×10^{-3}
average rel. error per option	1.37×10^{-2}	1.03×10^{-2}
aver. r.e. (options at-the-money)	3.83×10^{-3}	3.42×10^{-3}

Table 1: Residual for option prices for all maturities and absolute and relative errors per option (Black-Scholes data set). The average relative errors were computed according to (11).

constraints, 352 constraints involving only SDP variables, 476 constraints involving linear and SDP variables, and 1248 constraints involving linear and SOC variables.

Finally, we point out that the problem is relatively robust with respect to the number of asset knots; a similar performance is found if we consider 20 or 30 asset knots or vary its locations, or change the temporal knots.

3.2 Absolute diffusion data

To further demonstrate the effectiveness of our approach in recovering a known risk-neutral density surface for the underlying asset, we generated another set of option prices, this time assuming that the underlying assets follow an absolute diffusion process of the following form (see [12]):

$$dS_t \;=\; (r-d)S_t dt + 15 dW_t, \;\; t \in [0,\tau], \tau > 0. \tag{12}$$

Above, W_t is a standard Brownian motion and τ is a fixed trading horizon. The choice of the parameters followed [12]: initial value for the asset $S_0 = 100$, risk-free interest rate $r = 0.05$, and dividend rate $d = 0.02$. There were 7 maturities $\mathcal{T} = \{5/12, 7/12, 8/12, 9/12, 10/12, 11/12, 12/12\}$, and for each maturity 15 call prices. The call option prices were computed using the analytic formula for pricing European call options [14] when the underlying asset follows an absolute diffusion process.

We considered 30 equally spaced asset spline knots between 50 and 150. The set of temporal spline knots was $\{5/12, 12/12\}$. The Hessian perturbation is the same as in the Black-Scholes case. We did not incorporate any weights $\theta_{T,K}$ and $\mu_{T,K}$.

As we can see from Figures 3 and 4, the shapes of the theoretical and recovered surfaces are very similar. There is a good fitting when we plot, at the maturities, the theoretical densities against the recovered ones. We only observed a slight discrepancy at the recovered risk-neutral density for

Maturity	QP	SDP
1		
average abs. error per option	7.15×10^{-3}	6.38×10^{-3}
average rel. error per option	1.09×10^{-2}	9.42×10^{-3}
aver. r.e. (options at-the-money)	1.42×10^{-2}	1.27×10^{-2}
2		
average abs. error per option	1.65×10^{-3}	1.61×10^{-3}
average rel. error per option	4.74×10^{-3}	4.19×10^{-3}
aver. r.e. (options at-the-money)	3.52×10^{-4}	7.54×10^{-4}
3		
average abs. error per option	5.07×10^{-3}	4.34×10^{-3}
average rel. error per option	2.51×10^{-2}	2.41×10^{-2}
aver. r.e. (options at-the-money)	5.31×10^{-3}	5.06×10^{-3}
4		
average abs. error per option	5.45×10^{-3}	4.18×10^{-3}
average rel. error per option	1.85×10^{-2}	1.34×10^{-2}
aver. r.e. (options at-the-money)	4.46×10^{-3}	3.79×10^{-3}
5		
average abs. error per option	2.40×10^{-3}	2.47×10^{-3}
average rel. error per option	8.83×10^{-3}	8.41×10^{-3}
aver. r.e. (options at-the-money)	6.92×10^{-4}	2.82×10^{-4}
6		
average abs. error per option	1.50×10^{-3}	1.70×10^{-3}
average rel. error per option	4.69×10^{-3}	2.91×10^{-3}
aver. r.e. (options at-the-money)	5.81×10^{-4}	6.71×10^{-4}
7		
average abs. error per option	4.34×10^{-3}	3.37×10^{-3}
average rel. error per option	2.32×10^{-2}	9.48×10^{-3}
aver. r.e. (options at-the-money)	1.16×10^{-3}	7.28×10^{-3}

Table 2: Absolute and relative errors per option and per maturity (Black-Scholes data set). The average relative errors were computed according to (11).

abs. diff. data set	QP	SDP
residual $E(y)$	6.03×10^{-3}	6.03×10^{-3}
average abs. error per option	5.19×10^{-3}	5.21×10^{-3}
average rel. error per option	2.32×10^{-2}	2.28×10^{-2}
aver. r.e. (options at-the-money)	1.69×10^{-3}	1.69×10^{-3}

Table 3: Residual for option prices for all maturities and absolute and relative errors per option (absolute diffusion data set). The average relative errors were computed according to (11).

the first maturity. The residual $E(y)$ for the option prices for all maturities and the errors per option are also good, for both QP and SDP approaches, as one can see from Table 3.

Once again, our methodology appears robust to the selection of the number and location of the spline knots. For instance, if we change the number of asset knots to any number belonging to $[20, 40]$, we obtain essentially identical results. Similarly, if we change the location of the temporal knots, the estimation results remain the same.

In this case, the QP problem solved had 464 variables, 594 equality constraints, and 60 inequality constraints. The SDP problem, reformulated using the second order cone (SOC) constraints, exhibited 464 linear variables, 88 SDP variables, 465 SOC variables, 594 linear constraints, 264 constraints involving only SDP variables, 352 constraints involving linear and SDP variables, and 464 constraints involving linear and SOC variables.

3.3 Heston data

In order to test the ability of our method to recover densities obtained from different processes followed by the underlying, we considered a stochastic volatility model, namely, the Heston stochastic volatility model.

The sets of option prices were generated assuming that the stock price S_t follows a process of the form [16, 24]:

$$dS_t = \mu S_t dt + \sigma_t S_t d\omega_t^1. \tag{13}$$

Here μ is the drift parameter, ω_t^1 is a standard Brownian motion and σ_t is the time-dependent volatility. Introducing the new variable $\upsilon_t = \sigma_t^2$, called variance, we assume that it follows a mean-reverting process

$$d\upsilon_t = k[\theta - \upsilon_t]dt + \gamma\sqrt{\upsilon_t}d\omega_t^2. \tag{14}$$

The factor $k[\theta - \upsilon_t]$ ensures mean-reversion of the variance towards the long-time mean θ and k is the speed of the adjustment. Also, ω_t^2 is a standard

Heston data set	QP	SDP
residual $E(y)$	1.19×10^{-1}	1.21×10^{-1}
average abs. error per option	1.76×10^{-2}	1.75×10^{-2}
average rel. error per option	1.87×10^{-1}	1.86×10^{-1}
aver. r.e. (options at-the-money)	3.81×10^{-2}	3.63×10^{-2}

Table 4: Residual for option prices for all maturities and absolute and relative errors per option (Heston data set). The average relative errors were computed according to (11).

Brownian motion and γ is the variance noise. The processes ω_t^1 and ω_t^2 are correlated with correlation λ. We generated option prices for 7 maturities: T=$\{0.1, 0.2, 0.25, 0.3, 0.4, 0.45, 0.5\}$. For each maturity we generated 20 call option prices. The strikes are equally spaced between $0.8 \times S_0$ and $1.4 \times S_0$. We set $S_0 = 50$ as the initial price for the underlying asset. There are 28 asset spline knots, equally spaced between $0.8 \times S_0$ and $1.4 \times S_0$. The temporal knots are $\{0.1, 0.5\}$. The Hessian perturbation is the same as in the Black-Scholes case and there are no weights. Figures 5 and 6 show the results for the pdf estimation.

The residual $E(y)$ for the option prices for all maturities and the errors per option, for QP and SDP approaches, are given in Table 4.

In order to stress the validity of our method, we fit the risk-neutral density surface to a set of option prices related to four maturities, namely $\{0.1, 0.25, 0.4, 0.5\}$. After that, we generate out-of-sample option prices, based on the estimated risk-neutral density, for another set of four maturities $\{0.2, 0.3, 0.35, 0.45\}$. The strike prices used in the out-of-sample estimation are different from the ones used in the risk-neutral density estimation. The fitting of the out-of-sample prices against theoretical prices is good as we can see from Figure 7. The average relative value per option, for out-of-sample prices, is 6.31×10^{-2} (of similar order as for in-sample prices).

3.4 S&P500 data

In this next set of experiments, we test the effectiveness of our approach on recovering risk-neutral densities from traded option price data. The first data set is the one used in [3, 12] and refers to European call options on the S&P500 index traded in October 1995. In the original data, there are 10 maturities and 10 strikes for each maturity. We considered only 7 of these 10 maturities, as in [12], for easier comparison of the results. The specifications of the data are as follows: the spot asset price is 590, the risk-free interest

S&P 500 (1995) data set	QP	SDP
residual $E(y)$	12.89	12.99
average abs. error per option	2.88×10^{-1}	2.89×10^{-1}
average rel. error per option	3.61	3.61
aver. r.e. (options at-the-money)	4.14×10^{-2}	4.31×10^{-2}

Table 5: Residual for option prices for all maturities and absolute and relative errors per option (S&P 500 (1995) data set). The average relative errors were computed according to (11).

rate is 0.06, and the dividend rate is 0.0262. The location of the spline knots was chosen to be around the initial underlying asset price. We considered 28 equally spaced asset knots belonging to $[472, 826]$. The set of temporal knots is $\{0.175, 2\}$. The modification of the Hessian of the objective function is the same as before and we did not consider any weights $\theta_{T,K}$ and $\mu_{T,K}$.

As illustrated in Figures 8 and 9, both the QP and SDP approaches gave similar results and deviated clearly from log-normality. It is interesting to observe the uncertainty associated with longer maturities. The cumulative residual $E(y)$ for QP and SDP approaches as well as the errors per option are reported in Table 5.

One can compare the obtained average, absolute and relative errors per option to the error reported in [12] which seems to be around 0.0076, slightly better than ours, although it is difficult to establish rigorous comparisons without knowing exactly how the error in [12] was computed.

The dimensions of the QP problem solved are the following: 432 variables, 552 equality constraints, and 56 inequality constraints. The SDP problem, reformulated using the second order cone (SOC) constraints, had 432 linear variables, 54 SDP variables, 433 SOC variables, 552 linear constraints, 162 constraints involving only SDP variables, 216 constraints involving linear and SDP variables, and 432 constraints involving linear and SOC variables.

As we did with simulated prices generated using the Heston model, we generate out-of-sample estimates of some of the option prices in this SPX set. First, we estimate the density surface using four maturities. Then, based on the estimated risk-neutral density we generate out-of-sample prices for the remaining maturities. The fitting for the out-of-sample prices against market prices is acceptable as we can see from Figure 10. The average relative error per option, for out-of-sample prices, is 1.87.

Our second S&P500 data set was collected from CBOE on October 4, 2006 with maturities in October, November, and December (the set of ma-

S&P500 (2006) data set	QP	SDP
residual $E(y)$	2.74	3.03
average abs. error per option	0.728	0.761
average rel. error per option	0.126	0.130
aver. r.e. (options at-the-money)	0.126	0.130

Table 6: Residual for option prices for all maturities and absolute and relative errors per option (S&P500 (2006) data set). The average relative errors were computed according to (11).

turities was $\mathcal{T} = \{16/360, 44/360, 72/360\}$). The strikes differ from one maturity to the other and are, respectively,

$$[1320, 1325, 1330, 1335, 1340, 1345, 1350, 1360, 1370, 1375, 1380, 1385]\,,$$
$$[1320, 1330, 1335, 1345, 1350, 1360, 1370, 1375, 1380, 1385]\,,$$
$$[1320, 1325, 1330, 1335, 1340, 1345, 1350, 1360, 1375, 1380, 1385]\,.$$

We considered as option prices the average of the bid and ask prices. The data selection followed [2]. Occasionally, quoted option prices may exhibit arbitrage opportunities, often as a result of illiquidity in certain strikes and maturities. Since a risk-neutral density surface consistent with observed prices can not exist when the prices contain arbitrage opportunities, we use the following strategy to eliminate arbitrage in the data (see also [33]). Using put-call parity we translated the put prices into call prices, then checked for monotonicity and convexity (as a function of the strike price), and then removed prices that violate these properties from the estimation data.

The temporal spline knots corresponded to the first and last maturities. The placement of the asset knots was made as before around the initial underlying asset price. In this case there were 16 equally spaced asset knots between 1281 and 1675. (If we choose the number of knots too small (less than 10) the recovered risk-neutral densities exhibit less smoothness than in the cases where the number of knots exceeds 10.) The Hessian was modified as for the other data sets. The overall residual $E(y)$ found for the three maturities and the errors per option is given in Table 6.

Figure 11 depicts the recovered densities using the QP and SDP approaches for this data set. The fitting for the first maturity is worse than for the other two, for both QP and SDP approaches (see Figure 12).

For this data set, the QP problem solved had 288 variables, 363 equality constraints, and 32 inequality constraints. The SDP problem, reformulated using the second order cone (SOC) constraints, had the following dimensions: 288 linear variables, 46 SDP variables, 289 SOC variables, 363 linear

CPU time (sec.)	QP	SDP
BS data set	53.6	161.4
abs. diff. data set	3.2	40.7
S&P500 (1995) data set	2.9	18.8
S&P500 (2006) data set	1.4	4.9

Table 7: CPU times for QP and SDP approaches.

constraints, 158 constraints involving only SDP variables, 184 constraints involving linear and SDP variables, and 288 constraints involving linear and SOC variables.

3.5 Advantages of the SDP formulation and CPU times

We have observed, especially for nearly all market data tested in this paper, that the QP approach yielded risk-neutral densities exhibiting negative values. The SDP approach, in turn, always returned positive densities (along the spatial and temporal lines of the discretization). In our previous paper, for the static case, the SDP formulation showed multimodality in the recovered densities, a feature not detected by the solution of the QP formulation.

We also point out that guaranteeing nonnegativity of the risk-neutral densities inside the discretization rectangles (in the dynamic case and by interpolation) was never an issue once this was the case along the spatial and temporal lines, a feature guaranteed by the SDP approach.

The SDP approach is computationally more expensive but not as much as one would think, as reported in the Table 7 (the tests were made in an Intel(R) Core(TM)2 Duo CPU T7100, 1.8 GHz, 2GB RAM)

4 Pricing binary options

The CBOE has started to trade binary option contracts on the S&P500 and the VIX, in July, 2008. These options are of European style. CBOE binary call options pay out $C = 100$USD if the settlement value S_T, at maturity T, is at or above the strike price K, and pay nothing if the settlement value is below the strike price.

In this section we apply our method for estimating the risk-neutral density to price binary options. The price of the above binary options can be calculated as the discounted value of C times the risk-neutral probability that the underlying S_T is in the money at maturity. Using our risk-neutral density estimation, the price of a binary call option, with strike price K and

maturity T, is thus given by

$$B_{p_y}(K,T) \;=\; e^{-r(T-t_0)}C\int_K^{+\infty} p_y(\omega,T)d\omega.$$

Since we chose a range $[a,b]$ for possible terminal values for the underlying asset we have, instead,

$$B_{p_y}(K,T) \;=\; e^{-r(T-t_0)}C\int_K^{b} p_y(\omega,T)d\omega. \tag{15}$$

For binary options of this type, it is known that there exists a closed pricing formula using the Black-Scholes formula (see, e.g., [25]):

$$B(K,T,\sigma_{imp}) \;=\; Ce^{-r(T-t_0)}N(d_2(\sigma_{imp})), \tag{16}$$

where

$$d_2(\sigma_{imp}) \;=\; \frac{\log(S_0/K)+(r-d-\sigma_{imp}^2/2)(T-t_0)}{\sigma_{imp}\sqrt{T-t_0}}$$

and where d is the dividend rate for the S&P500 index and σ_{imp} is an approximation for the implied volatility.

In order to compare these two different pricing strategies, we considered a set of European vanilla call option prices on the S&P500 index collected on July 15, 2008, with 3 maturities: July, August, and September. For each maturity we have 15 strike prices. We collected also a set of binary call option prices on July 15, 2008 and considered the market binary call prices given by (bid price+ask price)/2.

For our computations, we first estimated the risk-neutral density for the three maturities according to the SDP approach given in Section 2.3. Then, the numerically estimated risk-neutral density was used in (15) to compute the corresponding binary call prices. Finally, we calculated the binary option prices from (16), setting σ_{imp}, for each binary option, to the implied volatility computed from the corresponding European vanilla call option price on the S&P500 index.

Figure 13 depicts the binary call prices using SDP risk-neutral estimation against the BS implied binary call prices and the market binary prices. We observe that the difference between BS implied binary prices and SDP recovered binary prices is significant and that market binary prices are much closer to those recovered using the SDP formulation. The errors between the estimated prices and the market data are given in Table 8.

S&P500 (2008) binary data set	(14) using SDP	(15)
overall residual	1.27×10^{-4}	1.31×10^{-3}
average abs. error per option	1.31×10^{-2}	4.84×10^{-2}
average rel. error per option	1.59×10^{-1}	2.16×10^{-1}
aver. r.e. (options at-the-money)	1.16×10^{-1}	1.79×10^{-1}

Table 8: Residual for option prices for all maturities and absolute and relative errors per option (S&P500 (2008) binary data set). The average relative errors were computed according to (11).

5 Concluding remarks

We have proposed a nonparametric method for estimating the dynamics of the risk-neutral density. Bicubic splines are used to generate the surface of densities, ensuring the appropriate smoothness. The optimization problems formulated (QP and SDP) proved to be effective in recovering the densities from a variety of data, including data generated from the Black-Scholes model, an absolute diffusion model, and the Heston stochastic volatility model, and data collected from S&P500 index. The fitting of the option prices for the generated models data was extremely good, and the SDP approach ensured nonnegativity of the risk-neutral densities everywhere. The fitting for the market data respected, in general, the interval for the bid-ask spread. While our QP and SDP approaches performed well on the 1995 S&P500 index data, the performance was not so good for the 2006 S&P500 index data, especially in the fitting for the first maturity. We should stress that the 2006 data set is smaller than the 1995 one, with strikes more irregularly located.

Finally, we used the proposed approach to price exotic options and observed a departure from the prices computed using the Black-Scholes implied volatility formula and a better fitting to the market exotic prices.

References

[1] K. Abadir and M. Rockinger. Density functionals, with an option-pricing application. *Econometric Theory*, 19:778–811, 2003.

[2] Y. Aït-Sahalia and A. Lo. Nonparametric estimation of state-price densities implicit in financial asset prices. *The Journal of Finance*, 53:499–547, 1998.

[3] L. Andersen and R. Brotherton-Ratcliffe. The equity option volatility smile: an implicit finite-difference approach. *The Journal of Computational Finance*, 1:5–32, 1998.

[4] M. Avellaneda, C. Friedman, R. Holmes, and D. Samperi. Calibrating volatility surfaces via relative-entropy minimization. *Applied Mathematical Finance*, 4:37–64, 1997.

[5] B. Bahra. Implied risk-neutral probability density functions from options prices: Theory and application. Technical report, Bank of England, 1997.

[6] D. Bertsimas and I. Popescu. On the relation between option and stock prices: A convex programming approach. *Operations Research*, 50:358–374, 2002.

[7] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81:637–659, 1973.

[8] O. Bondarenko. Estimation of risk-neutral densities using positive convolution approximation. *Journal of Econometrics*, 116:85–112, 2003.

[9] D. Breeden and R. Litzenberger. Prices of state-contingent claims implicit in options prices. *Journal of Business*, 51:621–651, 1978.

[10] B. Brunner and R. Hafner. Arbitrage-free estimation of the risk-neutral density from implied volatility smile. *The Journal of Computational Finance*, 7:75–106, 2003.

[11] P. W. Buchen and M. Kelly. The maximum entropy distribution of an asset inferred from option prices. *Journal of Financial and Quantitative Analysis*, 31:143–159, 1996.

[12] T. Coleman, Y. Li, and A. Verma. Reconstructing the unknown local volatility function. *The Journal of Computational Finance*, 2:77–102, 1999.

[13] R. Cont and J. Fonseca. Dynamics of implied volatility surfaces. *Quantitative Finance*, 2:45–60, 2002.

[14] J. Cox and S. Ross. The valuation of options for alternative stochastic processes. *Journal of Financial Economics*, 3:145–166, 1976.

[15] S. Crépey. Calibration of the local volatility in a trinomial tree using tikhonov regularization. *Inverse Problems*, 19:91–127, 2003.

[16] A. Drăgulescu and V. Yakovenko. Probability distribution of returns in the Heston model with stochastic volatility. *Quantitative Finance*, 2:443–453, 2002.

[17] D. Duffie. *Dynamic Asset Pricing Theory*. Princeton University Press, Princeton, NJ, 2001.

[18] B. Dumas, J. Fleming, and R. Whaley. Implied volatility functions: Empirical tests. *The Journal of Finance*, 53:2059–2106, 1998.

[19] W. Feller. Diffusion processes in genetics. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 227–246, 1951.

[20] W. Feller. Two singular diffusion problems. *Annals of Mathematics*, 54:173–182, 1951.

[21] W. Feller. *An Introduction to Probability Theory and its Applications*. John Wiley & Sons, New York, 1971.

[22] S. B. Hamida and R. Cont. Recovering volatility from option prices by evolutionary optimization. *The Journal of Computational Finance*, 8:43–76, 2005.

[23] S. Herzel. Arbitrage opportunities on derivatives: A linear programming approach. *Dynamics of Continuous, Discrete and Impulsive Systems, Series B: Applications and Algorithms*, 12:589–606, 2005.

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

[25] J. C. Hull. *Options, Futures, and Other Derivatives*. Prentice-Hall, NJ, fifth edition, 2003.

[26] N. Jackson, E. Süli, and S. Howison. Computation of deterministic volatility surfaces. *The Journal of Computational Finance*, 2:5–32, 1999.

[27] J. C. Jackwerth. Option-implied risk-neutral distributions and implied binomial trees: A literature review. *The Journal of Derivatives*, 7:66–82, 1999.

[28] J. C. Jackwerth and M. Rubinstein. Recovering probabilities distributions from options prices. *The Journal of Finance*, 51:1611–1631, 1996.

[29] R. Lagnado and S. Osher. A technique for calibrating derivative security pricing models: Numerical solution of an inverse problem. *The Journal of Computational Finance*, 1:13–25, 1997.

[30] D. Lamberton and B. Lapeyre. *Introduction to Stochastic Calculus Applied to Finance*. Chapman et Hall, London, 1996.

[31] M. S. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret. Applications of second-order cone programming. *Linear Algebra and Its Applications*, 284:193–228, 1998.

[32] S. Mayhew. On estimating the risk-neutral probability distribution implied by option prices. Technical report, Purdue University, 1995.

[33] A. M. Monteiro, R. H. Tütüncü, and L. N. Vicente. Recovering risk neutral probability density functions from options prices using cubic splines. *European Journal of Operational Research*, 187:525–542, 2007.

[34] M. Stutzer. A simple nonparametric approach to derivative security valuation. *The Journal of Finance*, 51:1633–1652, 1996.

[35] R. H. Tütüncü, K. C. Toh, and M. J. Todd. Solving semidefinite-quadratic-linear programs using SDPT3. *Mathematical Programming*, 95:189–217, 2003.

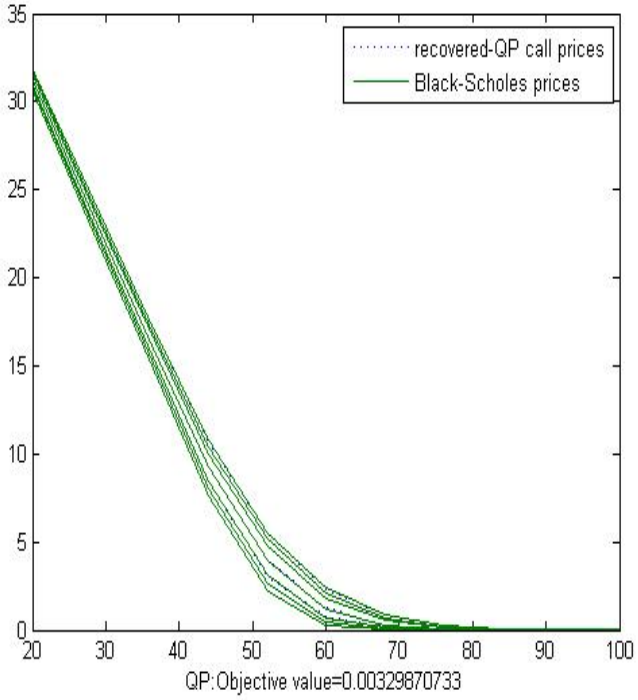

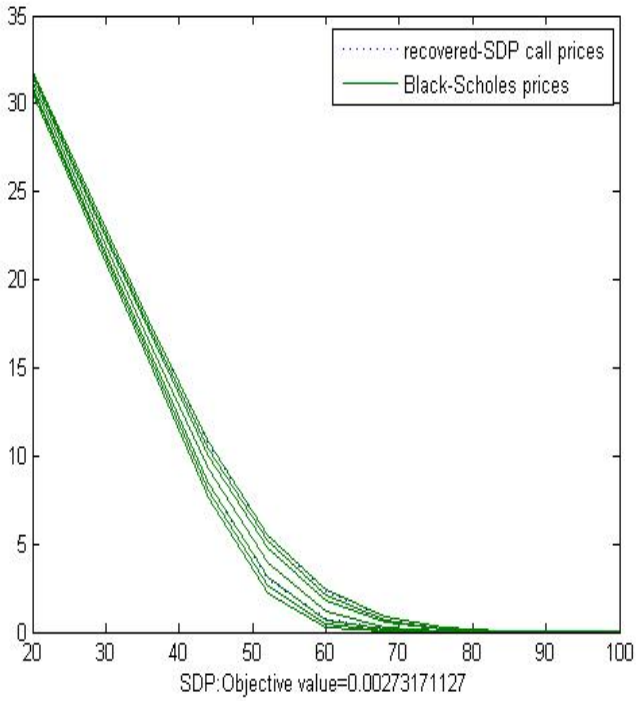

Figure 1: Recovered expected prices plotted against Black-Scholes prices for the seven maturities (QP and SDP approaches).

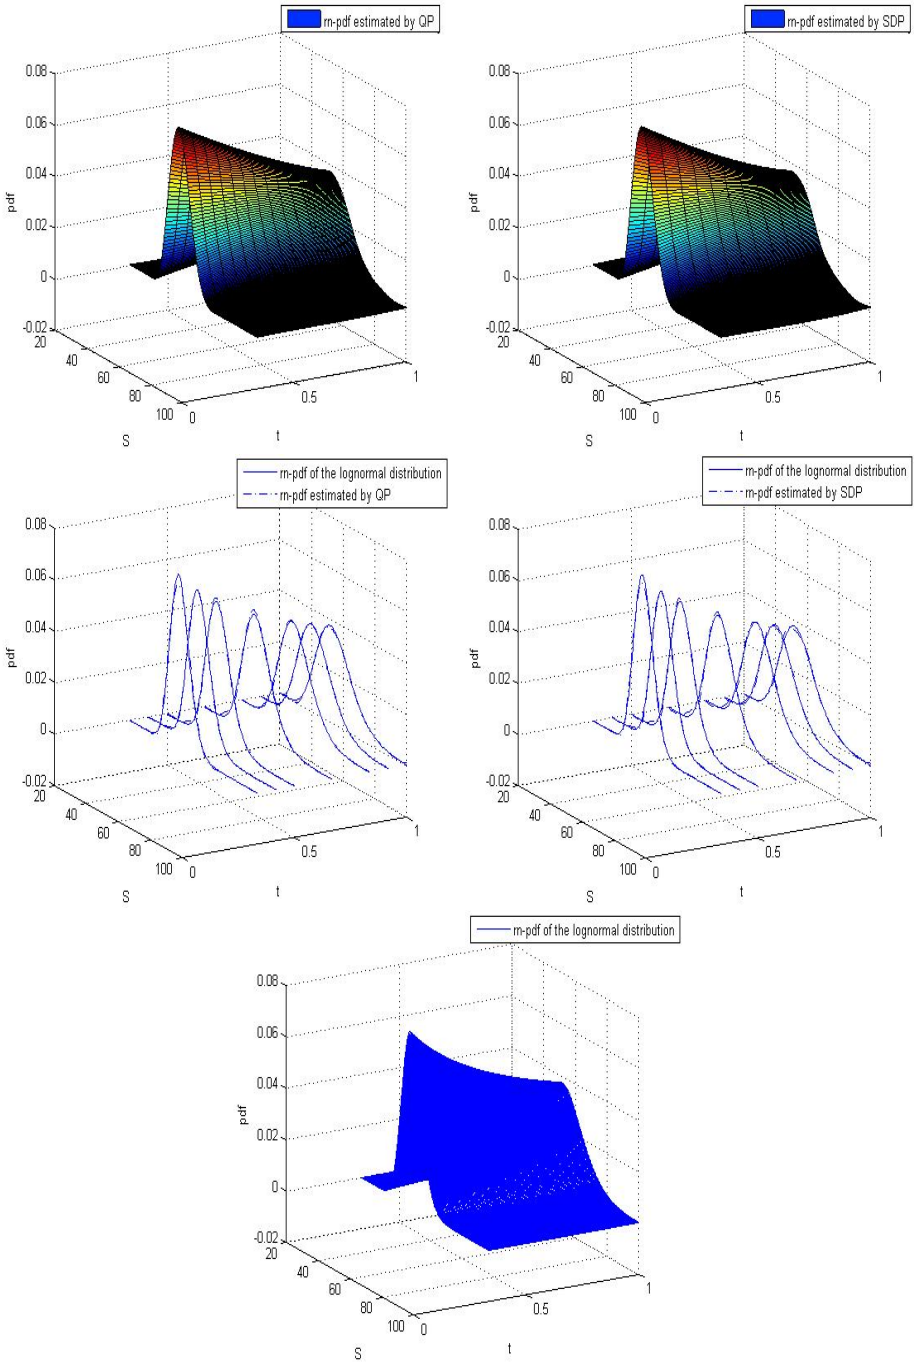

Figure 2: Recovered risk-neutral density surface from data generated by a Black-Scholes model using QP and SDP approaches, plotted against the log-normal densities at the maturities.

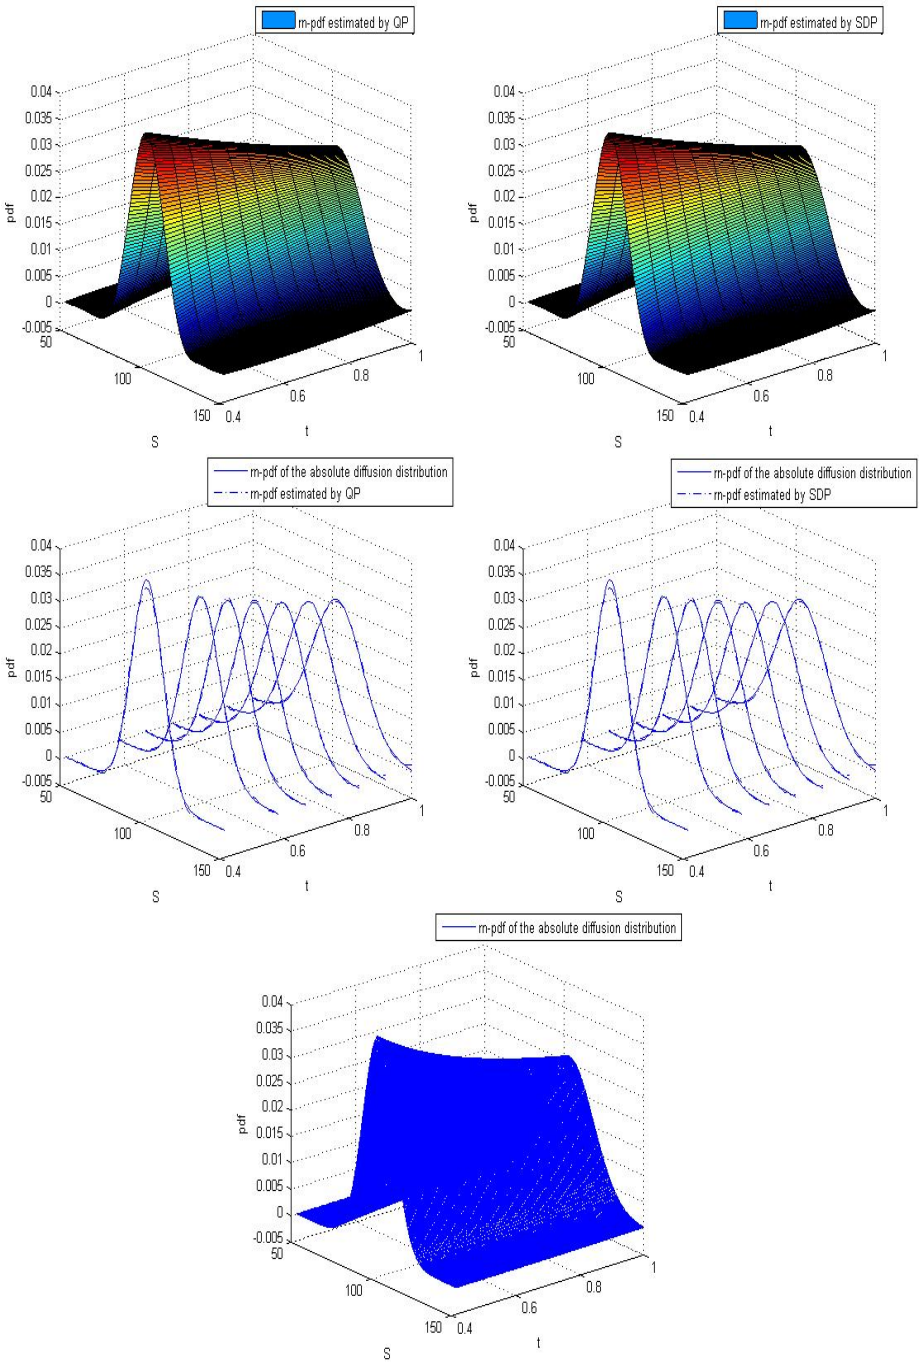

Figure 3: Recovered risk-neutral density surface from data generated by an absolute diffusion model using QP and SDP approaches, plotted against the absolute diffusion densities at the maturities.

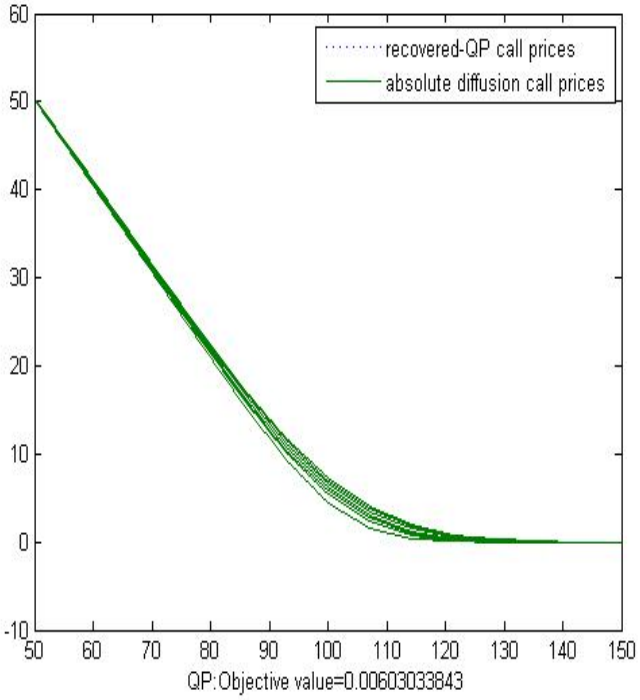

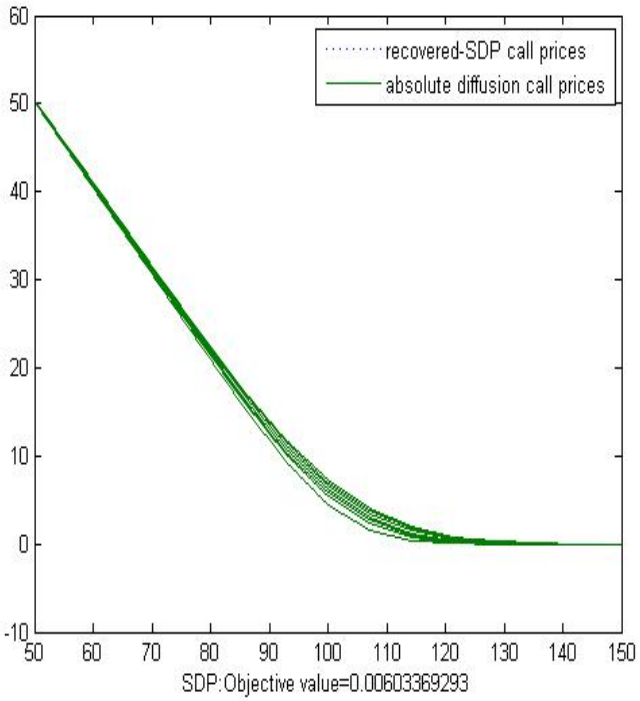

Figure 4: Recovered expected prices plotted against the absolute diffusion prices for the seven maturities (QP and SDP approaches).

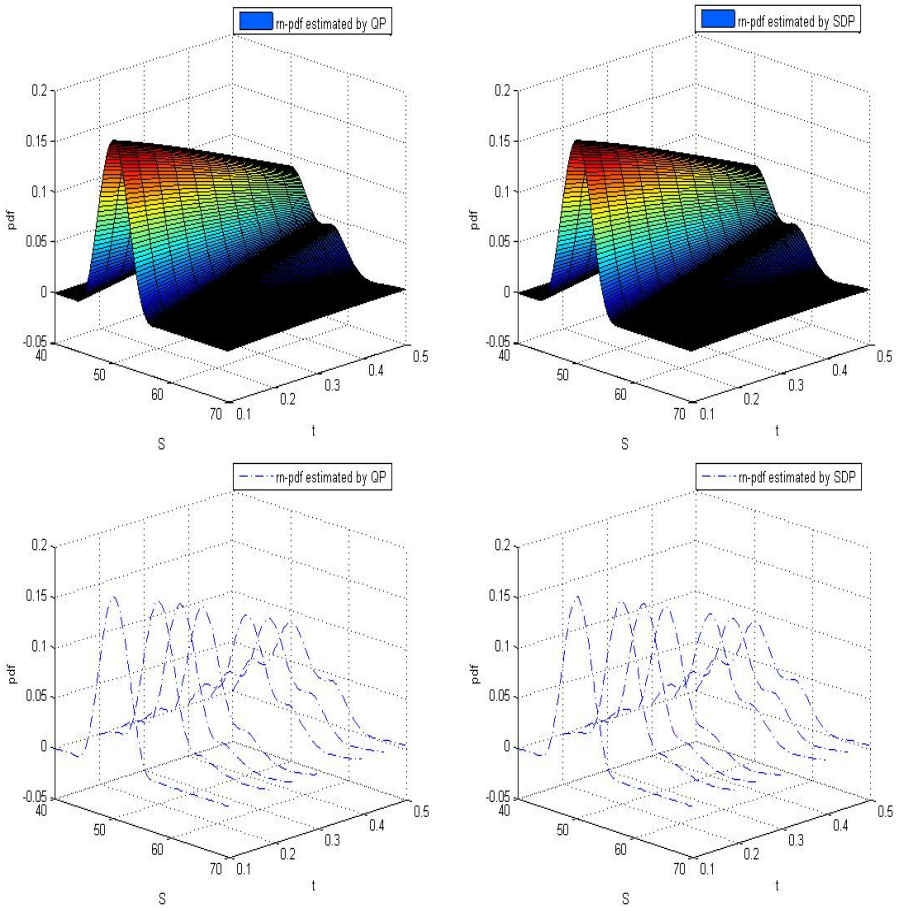

Figure 5: Recovered risk-neutral density surface from data generated by the Heston stochastic volatility model using QP and SDP approaches.

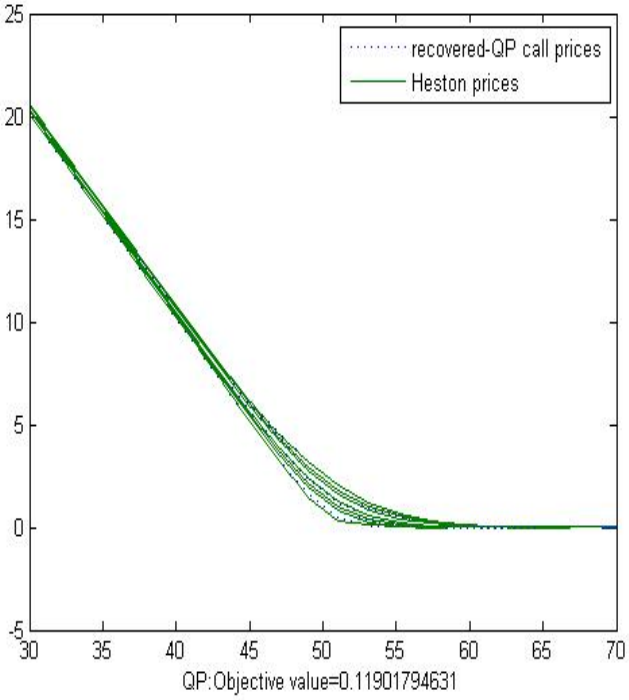

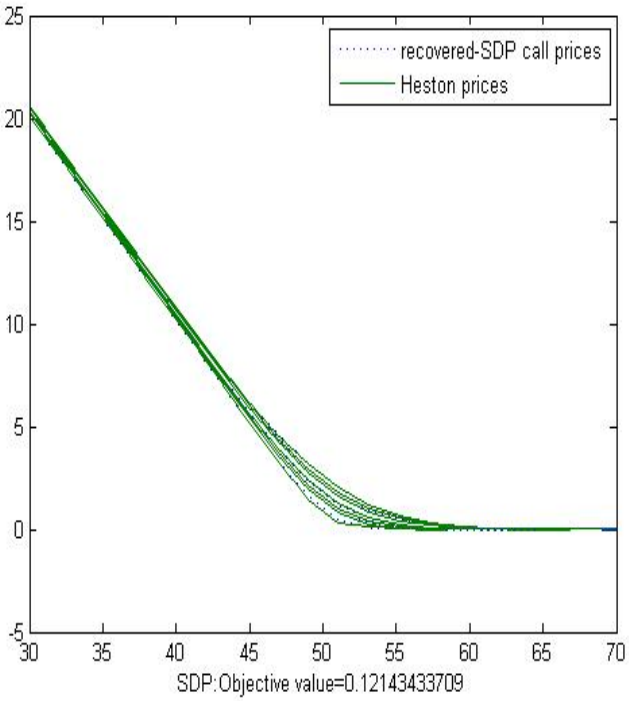

Figure 6: Recovered expected prices plotted against the Heston prices for the seven maturities (QP and SDP approaches).

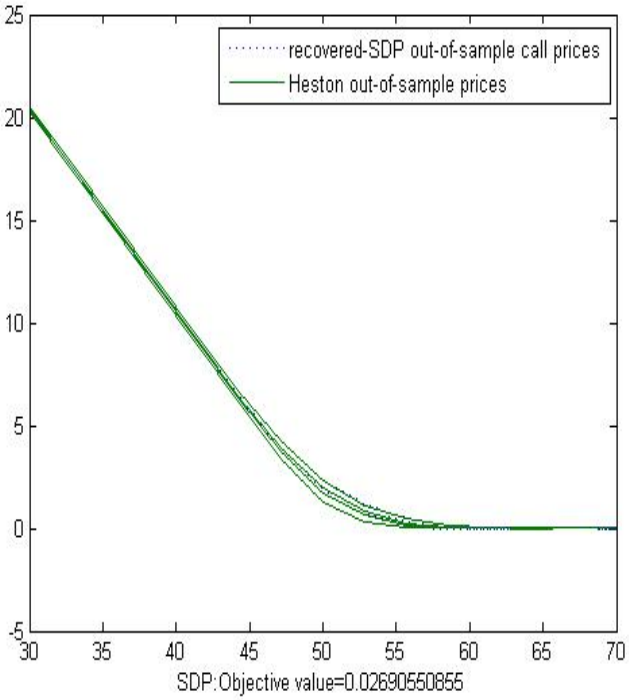

Figure 7: Recovered expected out-of-sample prices plotted against Heston prices for the four out-of-sample maturities (SDP approach).

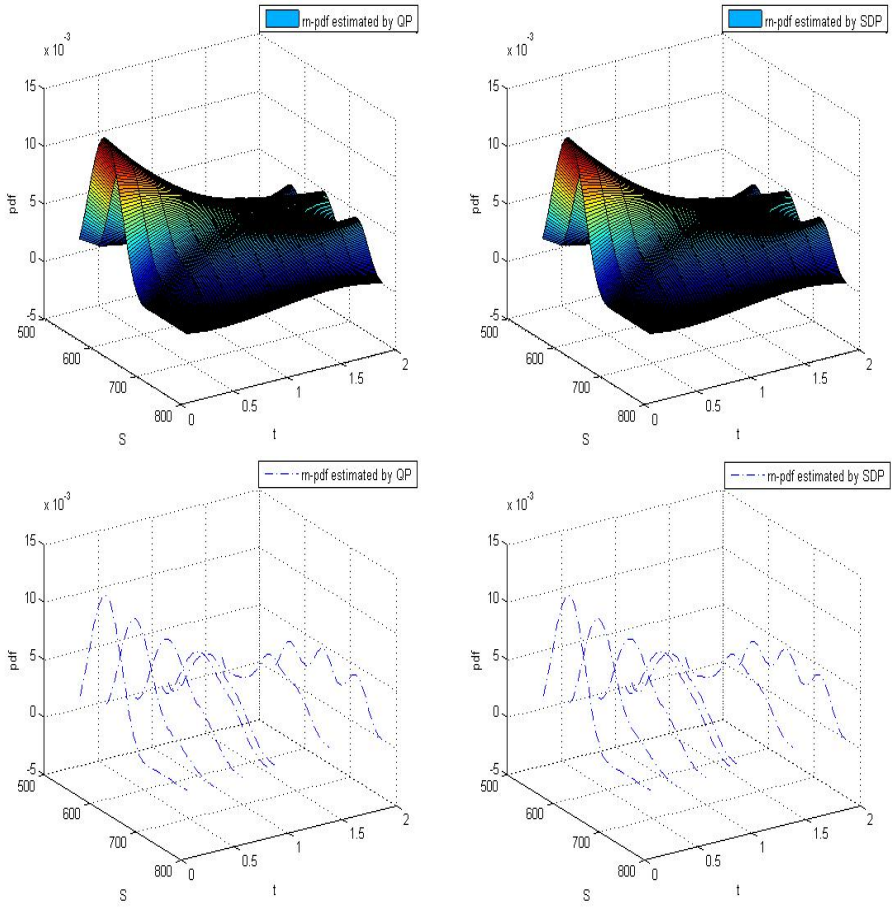

Figure 8: Recovered risk-neutral density surface and risk-neutral densities at the maturities, from S&P500 1995 data, using QP and SDP approaches.

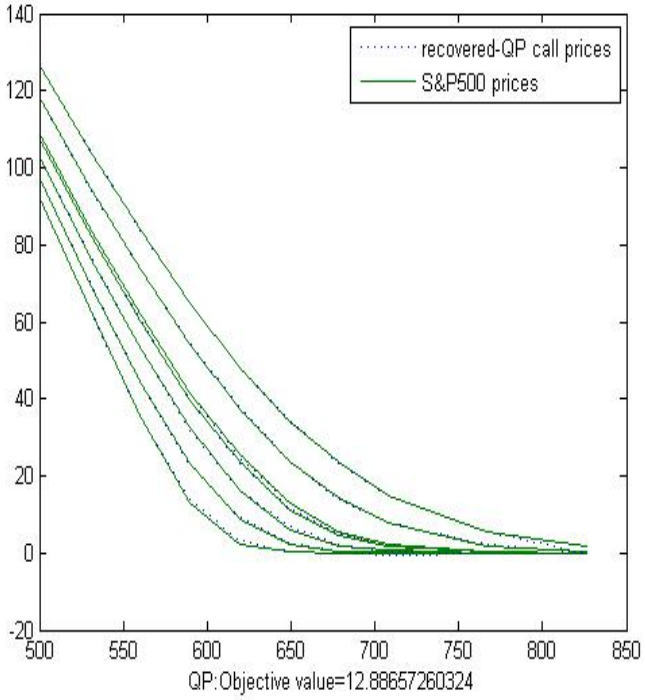

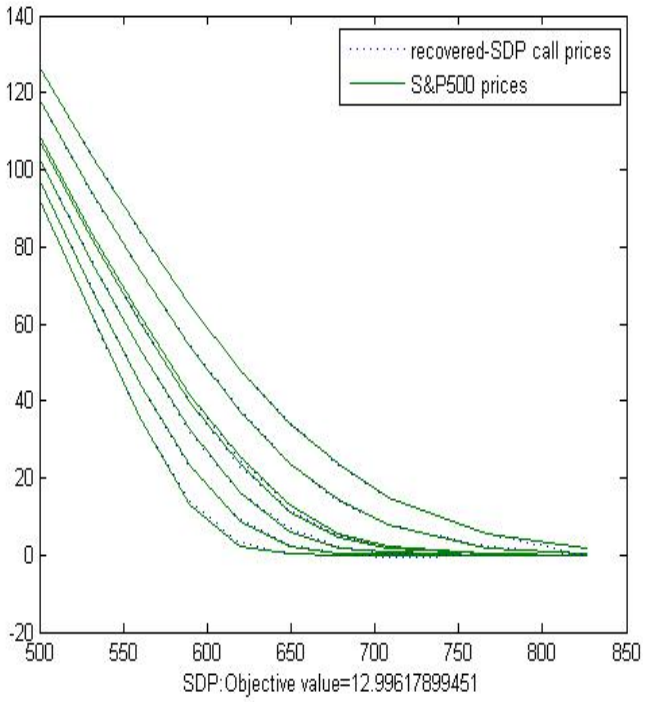

Figure 9: Recovered expected prices plotted against the S&P500 1995 prices for the seven maturities (QP and SDP approaches).

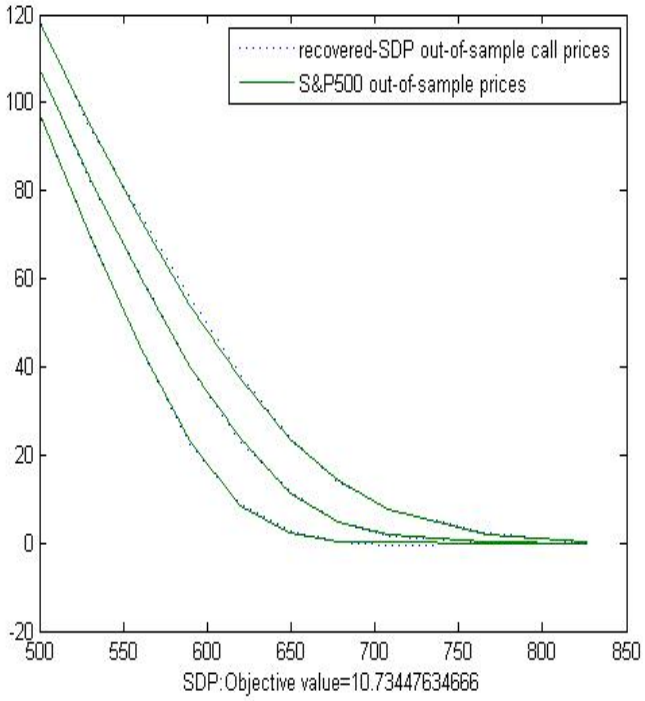

Figure 10: Recovered expected out-of-sample prices plotted against S&P500 1995 prices for the four out-of-sample maturities (SDP approach).

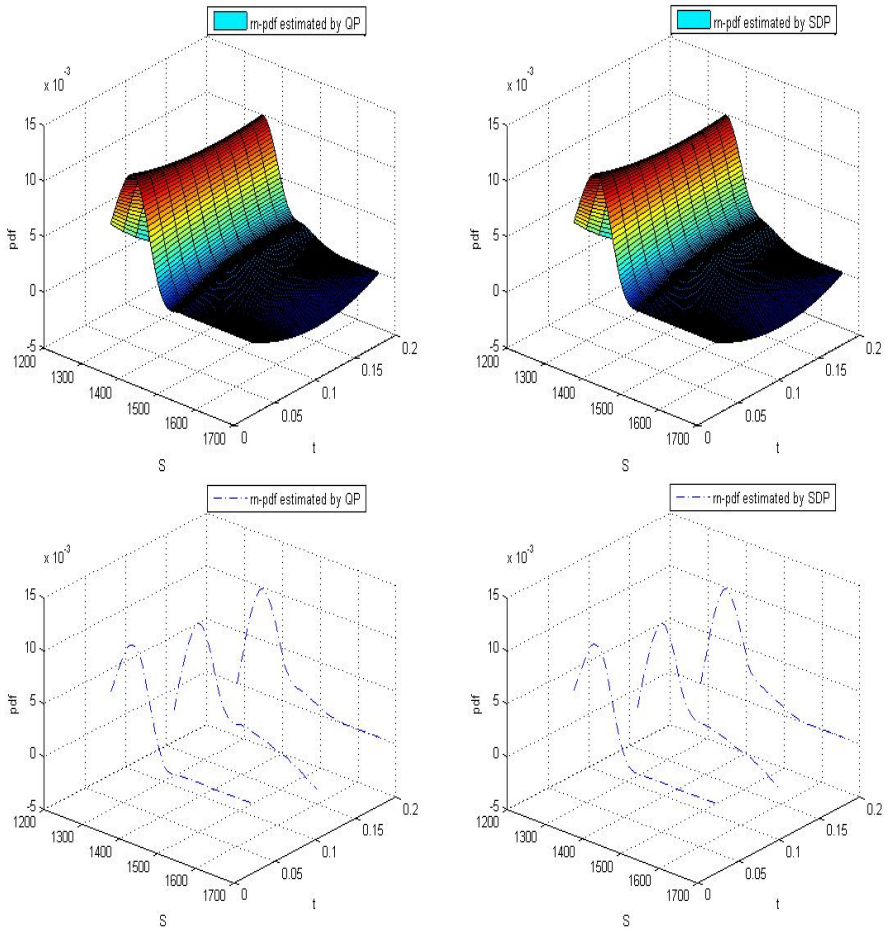

Figure 11: Recovered risk-neutral density surface and risk-neutral densities at the maturities, from S&P500 2006 data, using QP and SDP approaches.

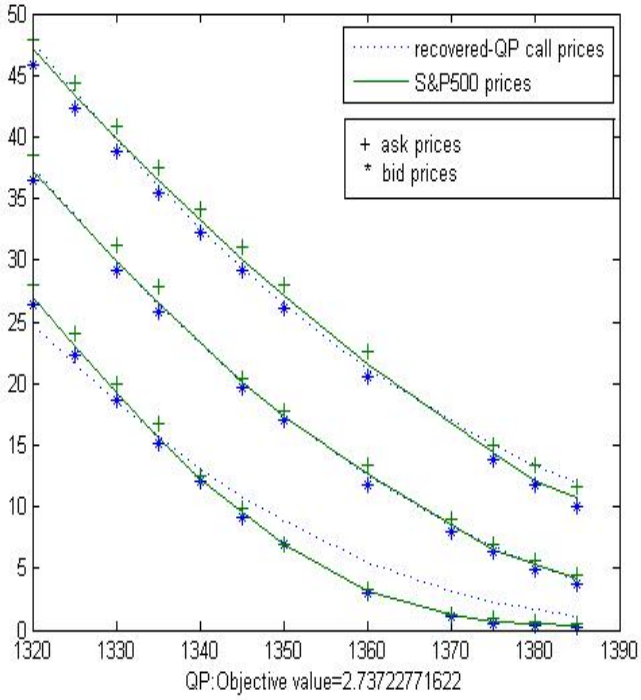

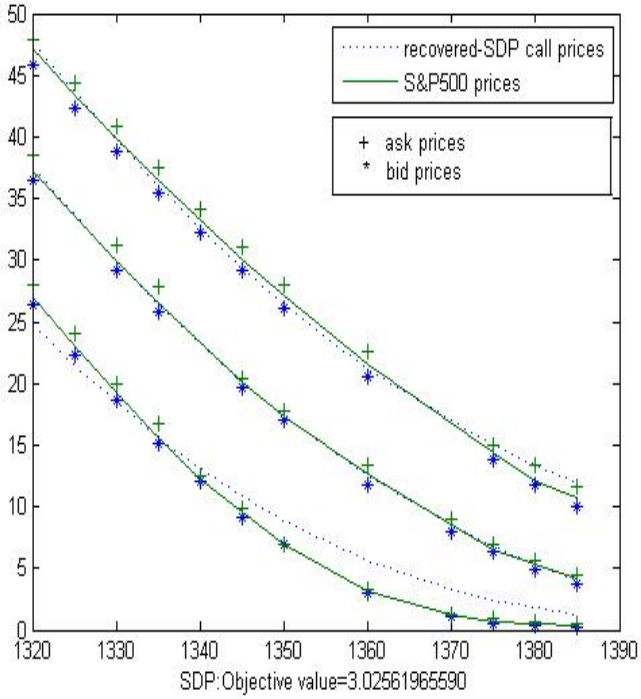

Figure 12: Recovered expected prices plotted against the S&P500 2006 prices for the three maturities (QP and SDP approaches).

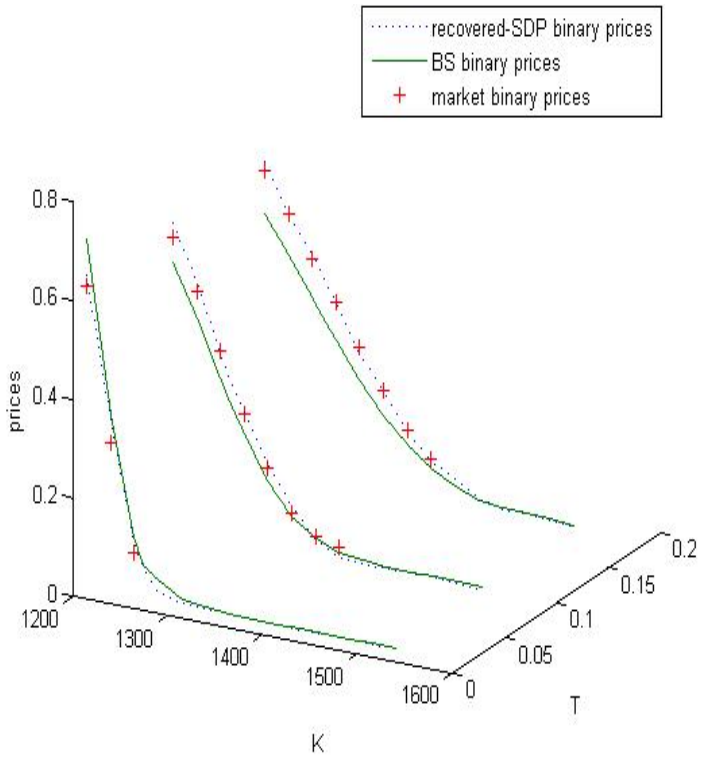

Figure 13: Recovered binary call prices from S&P500 data using the SDP approach against Black-Scholes binary prices and S&P500 2008 binary prices.