

A New Method of Employing the Principle of Maximum Entropy to Retrieve the Risk Neutral Density

by

Leonidas S. Rompolis*

Abstract

This paper suggests a new method of implementing the principle of maximum entropy to retrieve the risk neutral density of future stock, or any other asset, returns from European call and put prices. Instead of option prices, the method employs risk neutral moments as constraints. These moments can be retrieved from market option prices in a first step. Compared to other existing methods of retrieving the risk neutral density based on the principle of maximum entropy, the benefits of the method that the paper suggests is the use of all the available information provided by the market more sufficiently. To evaluate the performance of the suggested method, the paper presents simulation and empirical evidence.

Keywords: Maximum entropy, risk neutral density, risk neutral moments.

JEL: C40, C61, G13.

The author would like to thank Thanassis Stengos and Elias Tzavalis for their useful comments and suggestions in an earlier version of the paper.

*Department of Public and Business Administration, University of Cyprus. 9-11 Larnakos street, Nicosia, Cyprus. *Email:* rompolis@ucy.ac.cy. Leonidas S. Rompolis acknowledges financial support from the Iraklitos Fellowships of Research of the Athens University of Economics and Business.

1 Introduction

The estimation of the risk neutral density (RND), an essential tool for valuing derivatives, remains one of the most crucial issues in finance. Since the advent of the Black-Scholes (BS) option pricing formula based on the rather restricted assumption that log-returns follow the normal distribution, several methods were proposed to circumvent the empirical failures of this model associated with this assumption (see Garcia, Ghysels and Renault (2003), for a review). Most of the theoretical and empirical studies, which are aimed to improve the performance of the BS model, have focused on recovering the correct RND implied by option prices. These studies can be classified into two broad categories. The first assumes that the prices of the underlying asset follow a specific parametric model (see, for example, Heston (1993) and Bates (1996)). The second category employs nonparametric or density approximation methods to retrieve the RND from a set of observed option prices (see, for example, Rubinstein (1994), Ait-Sahalia and Lo (1998) and Madan and Milne (1994)).

Among the second class of studies, several authors have suggested the application of the principle of maximum entropy to recover the RND from market option prices. In general, the principle of maximum entropy is a Bayesian method of statistical inference that estimates a distribution from partial information in the form of a finite number of moments. The maximum entropy density (MED) is the least prejudiced estimate, compatible with the given price information in the sense that it will be maximally noncommittal with respect to missing, or unknown, information (see Jaynes (1979)). Thus, it can provide a feasible approximation of the RND implied by the available option prices data, without having to specify any underlying theoretical structure. By construction, this method produces strictly positive probabilities which can be thought of as one of its main advantages compared to other approaches of retrieving the RND like, for example, the standard Edgeworth expansion which implies negative risk neutral probabilities (see Jarrow and Rudd (1982) and Corrado and Su (1996)).

There are a few different applications of the principle of maximum entropy in the literature to recover the RND from option prices. Buchen and Kelly (1996) were among the first who have relied on this principle to recover the RND from cross-sectional sets of European option price data at any point in time imposing the constraints that the observed option prices are the expectations of their future payoffs. On the other hand, Stutzer (1996) has

used time series observations (historical data) of the underlying asset to predict the probability distribution of asset returns. He then used the principle of maximum entropy to estimate the RND subject to the constraint that the underlying asset is a martingale as well as any other derivative whose payoff occurs at this date.

This paper suggests a new methodology for implementing the principle of maximum entropy to retrieve the RND from option prices. The crucial feature, which differentiates our method from the aforementioned ones, is the way that the information contained in option prices is exploited. Instead of relying on the values of option prices, we employ the non-central moments of the RND implied by a cross-sectional set of European option prices as constraints to solve the MED problem. These moments can be retrieved from option prices in a model-free manner based on Bakshi, Kapadia and Madan (2003) and Rompolis and Tzavalis (2005, 2007) formulas. Our new methodology of implementing the principle of maximum entropy has several advantages compared to the previously mentioned ones. First, the entire cross-sectional data sets of option prices can be used in the estimation of the RNM and, hence, the RND. This is computationally impossible in the other existing methods of estimating the MED, which rely on option price constraints. This is due to the well known ill-conditioned problem of the Hessian matrix appeared in the estimation of the MED as the number of option price constraints increases (see Buchen and Kelly (1996)). Second, by construction the MED retrieved by our method will be a smooth density function, compared to the exponential piecewise linear function implied by the other existing methods. Third, our method can easily accommodate the recently proposed method of Wu (2003) to control for the ill-conditioned problem of the Hessian matrix in the estimation of the MED, if this appears in practice. In line with Wu's method, we can sequentially update the number of moment constraints in the MED estimation, starting from lower to higher order moments.

The paper is organized as follows. Section 2 presents our new methodology of implementing the principle of maximum entropy to recover the RND and discuss some of its features in more detail. Sections 3 and 4 evaluate the efficiency of our method by conducting two exercises. The first is based on a Monte Carlo simulation study, while the second empirically assesses its ability to successfully predict one-day ahead option prices. Section 5 concludes the paper.

2 Maximum entropy density implied by option prices

Let r be the instantaneous riskless interest rate, $C_t(S_t, K)$ and $P_t(S_t, K)$ denote the price of a European call and put at time t , respectively, with maturity date T , where S_t stands for the current-time t price of the underlying stock and K is the strike price. Define the $(T - t)$ -period log-return of the underlying stock as $X = \ln\left(\frac{S_T}{S_t}\right)$. Let μ_i be the i th-order non-central moment of X . Then, the maximum entropy density (MED), denoted as $p(x)$, can be obtained by maximizing Shannon's (1948) entropy measure:

$$W = - \int_D p(x) \ln p(x) dx \quad (1)$$

subject to the $m + 1$ moment constraints,

$$\int_D x^i p(x) dx = \mu_i, \quad i = 0, 1, \dots, m, \quad (2)$$

where D is the interval support of the RND and $\mu_0 = 1$. The above definition of the MED guarantees that the candidate estimator gives a probability density function which will always have positive values. This is due to the first constraint and the logarithmic nature of the function appeared in relationship (1). The second constraint, i.e. for $i = 1$, concerning the mean of the RND, guarantees that the underlying asset price is a martingale under the risk neutral measure. This is the main constraint assumed by Stutzer's (1996) method. For $i = 2, 3, 4$ the other three constraints are related to the variance, skewness and kurtosis coefficients of the RND.

To obtain the MED by solving the above problem, Rockinger and Jondeau (2002) (see also Agmon et. al. (1979a, b)) have proved that the maximization of Shannon's entropy, satisfying the moment conditions given by equation (2), is equivalent to the unconstrained minimization of the following function,

$$Q(\lambda_1, \dots, \lambda_m) = \int_D \exp\left(\sum_{i=1}^m \lambda_i (x^i - \mu_i)\right) dx,$$

with respect to the Lagrange multipliers $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_m)'$. The solution to the above

problem, known as the dual problem, takes the form:

$$p(x) = \frac{1}{Q(\lambda_1, \dots, \lambda_m)} \exp \left(\sum_{i=1}^m \lambda_i (x^i - \mu_i) \right). \quad (3)$$

The last relationship shows that the MED will belong to the Personian family. For small values of m , it is possible to obtain explicit solutions of (3). If $m = 0$, meaning that no information is given beyond the fact that one seeks a density, $p(x)$ becomes the density function of a uniform distribution over the interval support D . As one adds the first and then the second moment, then the exponential and normal density (implying the BS model) can be derived from equation (3), respectively. The knowledge of the third and higher moments does not yield a density in closed form. In this case only numerical solutions of (3) can be provided. Relation (3) also implies that the MED is an exponential polynomial function providing a smooth probability density function, compared to the exponential piecewise linear function implied by the existing MED methods of Buchen and Kelly (1996) and Stutzer (1996).

As equation (3) indicates, to retrieve the RND we can rely on prior estimates of moments μ_i . These can be directly obtained from out-of-the-money (OTM) European call and put prices employing the formulas of the risk neutral moments (RNM) suggested by Bakshi, Kapadia and Madan (2003) for $i = 1, 2, 3, 4$ and, more recently, extended by Rompolis and Tzavalis (2005, 2007) to any order i .¹ These formulas use the entire cross-sectional option data set available in the market at any point in time. The RNM estimated by the

¹These formulas are given as:

$$\begin{aligned} \mu_1 &= e^{r(T-t)} \left(1 - \int_{S_t}^{+\infty} \frac{1}{K^2} C_t(S_t, K) dK - \int_0^{S_t} \frac{1}{K^2} P_t(S_t, K) dK \right) - 1, \\ \mu_i &= e^{r(T-t)} \left\{ \int_{S_t}^{+\infty} \frac{i}{K^2} \left[\ln \left(\frac{K}{S_t} \right) \right]^{i-2} \left[i - 1 - \ln \left(\frac{K}{S_t} \right) \right] C_t(S_t, K) dK \right. \\ &\quad \left. + \int_0^{S_t} \frac{i}{K^2} \left[\ln \left(\frac{K}{S_t} \right) \right]^{i-2} \left[i - 1 - \ln \left(\frac{K}{S_t} \right) \right] P_t(S_t, K) dK \right\}, \text{ for } i \geq 2. \end{aligned}$$

As these formulas employ integrals of continuous functions to retrieve the values of the risk neutral moments based on them, we can employ cubic splines to interpolate the implied by our option prices volatilities between two different points of the data (see Campa et. al. (1998)). Due to the lack of option prices at 0 and $+\infty$, we can extrapolate the implied volatilities by a linear function over the intervals $(0, K_{\min}]$ and $[K_{\max}, +\infty)$, where K_{\min} and K_{\max} is the minimum and maximum strike prices given by our data, respectively. The extrapolation is truncated at strike prices, denoted as K_0 and K_∞ , which correspond to put and call prices which are very close to zero (e.g. smaller than 10^{-3}). These strike prices, define the lower and upper bounds of the integrals, respectively. They also determine the support of the estimated moments given by $[l, u] \equiv \left[\ln \left(\frac{K_0}{S_t} \right), \ln \left(\frac{K_\infty}{S_t} \right) \right]$.

above method reflect all the market available information for the RND at any point in time. Thus, the MED estimate of the RND retrieved by the above approach can be thought of as incorporating all the market information sufficiently. This is practically infeasible by the other existing MED methods of estimating the RND in which option prices are directly imposed as constraints in the maximization of the entropy measure. As aptly first noticed by Buchen and Kelly (1996), this is due to the well known ill-conditioning feature of the Hessian matrix. This problem becomes apparent in the estimation of the MED, if a large number of option prices are used as constraints. This does not allow us to use all the available market option price information in the estimation of the MED. It is most likely to happen when exact (theoretical) option prices are corrupted by a noise term, as often happen in practice (see Buchen and Kelly (1996)).

The existence of a solution to the MED problem (1)-(2) requires that Frontini's and Tagliani (1994) sufficient condition to hold. This theorem states that the necessary and sufficient condition for the existence of a MED given the $m + 1$ moments is that the Hankel matrix, defined as

$$\mathbf{H}_{2j} = \begin{bmatrix} \mu_0 & \mu_1 & \dots & \mu_j \\ \mu_1 & \mu_2 & \dots & \mu_{j+1} \\ \vdots & \vdots & \dots & \vdots \\ \mu_j & \mu_{j+1} & \dots & \mu_{2j} \end{bmatrix}$$

$\forall j = 0, 1, \dots, n$, where $m = 2n$, is positive definite. Given that the RND exists (which is true under the no arbitrage condition) and $\{\mu_i\}_{i=0}^{\infty}$ given by the formulas in footnote 1 constitutes by construction the sequence of its moments, the Hankel matrix \mathbf{H} will be positive definite by the converse of Hamburger's theorem.² This proves that Frontini's and Tagliani condition holds, which in turn means that problem (1)-(2) has a solution. The uniqueness of this solution is guaranteed by the positive definiteness of the Hessian matrix denoted \mathbf{G} , which has the following elements:

$$\mathbf{G}_{ij} = \frac{\partial^2 Q}{\partial \lambda_i \partial \lambda_j} = \int_D (x^i - \mu_i)(x^j - \mu_j)p(x)dx.$$

As long as $p(x)$ is a density function, \mathbf{G} defines a variance-covariance matrix which will be

²Hamburger's theorem states that a necessary and sufficient condition for the existence of a distribution with infinite support having moments $\{\mu_i\}_{i=0}^{\infty}$ is that the Hankel matrix \mathbf{H}_{2j} is positive definite, $\forall j > 0$ (see Shohat and Tamarkin (1943), for a survey).

always positive definite, which proves the uniqueness of the solution (3).

Numerically, the dual problem (3) can be solved using the well known Newton's method. This iteratively updates the following relationship:

$$\boldsymbol{\lambda}^{(k)} = \boldsymbol{\lambda}^{(k-1)} - \mathbf{G}^{-1} \frac{\partial Q}{\partial \boldsymbol{\lambda}}. \quad (4)$$

To use (4), we need to construct Q , which involves the computation of an integral. To this end, we will first map the interval support $D = [l, u]$ to $[-1, 1]$ and then, following Rockinger and Jondeau (2002), we will approximate the integral by a Gauss-Legendre quadrature.³ To control for a possible ill-conditioned problem of the Hessian matrix \mathbf{G} when the number of moments μ_i increases, we suggest that we employ the recently developed method of Wu (2003). This method proposes sequentially updating the moment constraints into the estimation process (4) from lower to higher order. The estimated coefficients based on the lower order moments are used as initial values to update the density estimates, when additional higher order moments are considered. This procedure can control for the ill-conditioning feature of the Hessian matrix \mathbf{G} , as the difference between the values of the moments predicted by the MED based on the lower order moments and the actual moments approaches to zero as the number of constraints increases. As a final, note that in implementing the MED the optimal order of the expansion of the density function $p(x)$ (i.e. the number of moment constraints m) is required. This can be chosen based on one of the well known information criteria, such as the Akaike (AIC) and the Schwarz (SC) (see also Wu (2003)). These can be calculated based on the sum of squared errors of the observed option prices from their predicted values obtained through the estimation of the MED.⁴

³Theoretically, the support of the log-return distribution and consequently of its moments is given by the interval $(-\infty, +\infty)$. The domain $[l, u]$ is necessary for the application of the maximum entropy method. This can be considered as an approximation of the true one.

⁴Specifically, these criteria choose the order of truncation of the MED to minimize:

$$\log(SSE_m/N) + (m+1)C(N)/N,$$

where SSE_m is the sum of squared option pricing errors for a given order m of the MED, N is the size of the sample (total number of option prices) and $C(N)$ is 2 for the AIC and $\log(N)$ for the SC criteria.

3 Simulation study

In this section we conduct a simulation study with the aim to assess if our MED based methodology proposed in the previous section can be successfully employed to retrieve the RND of the log-return (denoted X , with values x) from European option prices.

In our simulation study, we assume that option prices are generated from the stochastic volatility with jumps (SVJ) model suggested by Bates (1996). This model considers that the risk neutralized stochastic processes of the stock price S_t and its volatility, V_t , are given as follows:

$$\begin{aligned} dS_t &= S_t \left[(r - \lambda \bar{\mu}) dt + \sqrt{V_t} dW_{t,1} + \mu dz_t \right] \\ \text{with } V_t &= \kappa(\theta - V_t) dt + \sigma \sqrt{V_t} dW_{t,2} \\ \text{prob}(dz_t = 1) &= \lambda dt, \ln(1 + \mu) \sim N \left(\ln(1 + \bar{\mu}) - \frac{1}{2} \delta^2, \delta^2 \right), \end{aligned} \quad (5)$$

where $W_{t,1}$ and $W_{t,2}$ are two correlated Brownian motions with correlation coefficient ζ , λ is the annual frequency of jumps, μ is the random percentage jump conditional on a jump occurring and z is a Poisson counter with intensity λ . The SVJ model is frequently used in practice to improve upon the pricing performance of the BS model as it implies high levels of negative skewness and kurtosis of the RND. These two features of the RND are consistent with the pattern of the implied volatility, across different strike prices (see Bakshi, Cao and Chen (1997), and Chernov and Ghysels (2000)).

To generate empirically plausible option prices from model (5), we are based on values of its structural parameters found in the empirical literature (see Bakshi, Cao and Chen (1997)). These are set to the following levels: $\vartheta = 0.01$, $\sigma = 0.4$, $\kappa = 3.93$, $\zeta = -0.52$, $\lambda = 0.61$, $\bar{\mu} = -0.09$ and $\delta = 0.14$. The strike prices K are taken to span the closed interval $[K_{\min} = 820, K_{\max} = 1260]$ at every 20 points, where K_{\min} and K_{\max} denote the lower and upper values of K . These values of K_{\min} and K_{\max} are the average levels of the minimum and maximum strike prices of a sample of short term option prices written on the S&P 500 index from January 1996 to December 2002, which is used in our empirical study. To be consistent with the magnitude of them, the values of the two state variables S_t and V_t are taken to be the mean of the S&P 500 index and the variance of its log-return over the sample period. These are given as $S_t = 1080$ and $V_t = 0.026$, respectively. For the interest

rate r , we used the average level of the three-month US Treasury bill of the above period, given by $r = 0.05$.

Based on the above values of the SVJ option pricing model, we generated a set of $N_1 = 22$ call option prices with a short-term to maturity interval of $T - t = 2.5$ months (0.21 years), where the biggest failures of the BS model are observed. From this set, next we produced another one consisting of 22 put option prices using the put-call arbitrage relationship, thus bringing the total number of generated option prices to 44. This should be considered as a cross-sectional sample of option prices comparable to those often used to estimate the RND in practice. Following Ait-Sahalia and Duarte (2003), we added a noise term to our generated set of implied volatilities drawn from the uniform distribution, with interval $[-0.025, 0.025]$. This term can be taken to reflect random effects of the bid-ask spread and the different degree of liquidity on option call/put prices.⁵ From these prices, we then derived their implied volatilities which were used to obtain estimates of the RNM formulas given in fn 1. Given these estimates, we then obtained the RND based on the MED given by (3), following our methodology described in the previous section. These estimates of the RND will be henceforth denoted as MED-RNM- (N_1) . The optimal order m considered in our estimation of the MED was chosen to be 14 based on the SC information criterion.

To assess the performance of our method, we have also estimated the RND based on the MED methodology suggested by Buchen and Kelly (1996), henceforth denoted as MED-BK, using option prices as constraints. The number of option prices data used in this estimation (which is equal to the order of the expansion) was determined by the SC criterion. This was found to be on average equal to $N_2 = 7$.⁶ To investigate if our method performs well enough even for a such very small data set of $N_2 = 7$ option prices, we have implemented it to this data set, too. These estimates of our method will be henceforth denoted as MED-RNM- (N_2) .

In Table 1 we report the mean values of two well known metrics of density accuracy calculated over the 1000 iterations for our Monte Carlo exercises. These are the L^1 and L^2

⁵Note that the perturbed option prices does not violate the arbitrage conditions, i.e. monotonicity and convexity.

⁶We have also tried for a larger number of option prices as constraints (that is $N_2 > 7$), but we have found that the problem of ill-conditioned Hessian matrix appears.

metrics, defined as

$$L^1 = \int |p(x) - p^*(x)| dx \quad \text{and} \quad L^2 = \int (p(x) - p^*(x))^2 dx,$$

respectively, where $p^*(x)$ denotes the true RND function and $p(x)$ the estimated ones. The values of the RND $p^*(x)$ are obtained by inverting the theoretical characteristic function implied by the SVJ option pricing model.

Table 1: L^1 and L^2 measures

	MED-RNM-(N_1)	MED-RNM-(N_2)	MED-BK
L^1	0.062	0.096	0.103
L^2	0.015	0.040	0.049

The results of Table 1 clearly indicate that the MED-RNM method provides a very accurate estimation of the true RND. This is supported by both measures L^1 and L^2 , reported in the table. As was expected, the larger the sample size, the closest the MED to the true density is. The results of the table also show that MED-RNM approach slightly outperforms the MED-BK method even for the very small data set of $N_2 = 7$ option prices.⁷ These results are also supported by the inspection of Figure 1, which graphically presents the different estimates of the MED against the true density. As this figure shows, the MED-RNM method provides a smooth approximation of the true RND compared to the exponential piecewise linear density given by the MED-BK. This is true for the different sizes of option data set considered.

4 Empirical exercise

In this section, we conduct an empirical exercise with the aim of examining if there are any price valuation (prediction) gains in employing the MED-RNM method of retrieving the RND to provide accurate out-of-sample-forecasts of market option prices, implying small valuation errors.⁸ In our analysis, we use a sample of European call and put option prices

⁷Obviously, this can be attributed to the fact that the estimates of the RNM obtained in the first step are close to the true ones, even for a very small data set considered (see also Rompolis and Tzavalis (2005)).

⁸This is a standard exercise carried out in the literature to evaluate the empirical performance of many option pricing models or their implied risk neutral densities by option prices (see Bakshi, Cao and Chen (1997), and Heston and Nandi (2000)).

written on the S&P 500 index. The maturity interval of our options is close to 22 trading days. Our sample covers the period from January 1996 to December 2002. In order to use the same market information for both methods and to overcome the illiquid feature of the in-the-money (ITM) calls, our initial cross-sectional data set consists of out-of-the money (OTM) calls and puts and at-the-money (ATM) calls. The ITM calls, which are used by the MED-BK method, can then be calculated using the OTM puts through the put-call parity (see also Ait-Sahalia and Lo (1998)).

To study the distributional features of the RND implied by our data, our analysis starts with presenting the estimates of the RND over the whole sample period based on the two maximum entropy estimation methods: the MED-RNM and MED-BK (see Figure 2). This is a crucial step in our analysis as it can indicate potential sources of option pricing prediction failures of the above two methods, across different moneyness levels. As was expected, both methods' estimates reveal that our set of option prices imply a RND with high degree of kurtosis and negative skewness. However, the MED-RNM method gives fatter left and right tails than the MED-BK. The latter seems to assign higher probabilities for returns around the mean of the RND.

The above differences in the estimates of the RND can be translated into significant valuation gains in employing the MED-RNM method of estimating the RND to predict future option prices. To investigate if this is true, in Table 2 we present the values of some very well known performance evaluation measures of an out-of-sample forecasting exercise. This exercise was set up as follows. On the third Wednesday of each month, we retrieved estimates of the RND based on the MED-RNM and MED-BK methods, respectively, using our set of option prices at each point in time over the whole sample period. These estimates were then used to predict one-period ahead prices of call and put options (i.e. for the third Thursday of each month), with one period less to maturity horizon. For both methods, the predicted values of the option prices were derived by numerically evaluating the standard Cox-Ross (1976) option pricing formula. The values of the evaluation measures that the table reports are the mean absolute error (MAE) and the mean percentage error (MPE) between the predicted and the actual option prices, over the whole sample period.

The results of Table 2 clearly indicate that the MED-RNM method considerably improves upon the performance of the MED-BK method to predict future option prices, especially for the OTM calls and puts. This can be supported by both values of the forecasting

performance evaluation metrics reported in the table. The MED-BK method seems to slightly outperforms the MED-RNM only for ITM calls. The higher benefits of the MED-RNM than the MED-BK method which are observed across the different moneyness levels can be obviously attributed to the fact that, by exploiting more sufficiently the sample information of all available option prices at hand, it can more accurately estimate the true RND, especially around its tails.

5 Conclusions

This paper suggests a new method of implementing the principle of maximum entropy to retrieve the risk neutral density (RND) of future asset return implied by European call and put prices. Instead of employing option prices as constraints, the suggested method relies on estimates of risk neutral moments (RNM) implied by option prices in retrieving the maximum entropy density (MED). Compared with the earlier applications of the principle of maximum entropy to option pricing, our method merits several interesting features. First, it can use all the market available information to estimate the RND sufficiently, as it relies on estimates of risk neutral moments as constraints. Second, it can always provide a smooth probability density function by construction. Third, based on moment constraints it can more easily employ Wu's (2003) sequential updating procedure in increasing the number of constraints in the estimation of the MED. Thus it can better control for the well documented problem of ill-conditioned Hessian matrices faced by the other existing methods of implementing the principle of maximum entropy in practice, when the number of constraints rises.

To assess the performance of the suggested method, the paper conducts a simulation study generating option prices from the stochastic volatility with jumps model which is frequently used in practice to improve upon the pricing performance of the Black-Scholes model. It also carries out an empirical exercise investigating if there are any price valuation gains in employing our method to value option prices in practice. The results of both of these exercises clearly support the use of the new method of implementing the principle of maximum entropy that the paper suggests in option pricing.

References

- Agmon, N., Y. Alhassid and R. D. Levine, 1979a, An algorithm for finding the distribution of maximal entropy, *Journal of Computational Physics* 30, 250-259.
- Agmon, N., Y. Alhassid and R. D. Levine, 1979b, The maximum entropy formalism, *Maximum Entropy Formalism*, R. D. Levine and M. Tribus, eds. Cambridge, MA: MIT Press, 207-209.
- Ait-Sahalia, Y. and J. Duarte, 2003, Nonparametric option pricing under shape restrictions, *Journal of Econometrics* 116, 9-47.
- Ait-Sahalia, Y. and A.W. Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance* 53, 499-547.
- Bakshi, G., C. Cao and Z. Chen, 1997, Empirical performance of alternative option pricing models, *Journal of Finance* 52, 2003-2049.
- Bakshi, G., N. Kapadia and D. Madan, 2003, Stock return characteristics, skew laws, and the differential pricing of individual equity options, *The Review of Financial Studies* 16, 101-143.
- Bates, D. S., 1996, Jumps and stochastic volatility: exchange rate processes implicit in Deutsche Mark options, *Journal of Financial Studies* 9, 69-107.
- Buchen, P. W. and M. Kelly, 1996, The maximum entropy distribution of an asset inferred from option prices, *Journal of Financial and Quantitative Analysis* 31, 143-159.
- Campa, J. M., P.H K. Chang and R. L. Reider, 1998, Implied exchange rate distributions: Evidence from OTC option markets, *Journal of International Money and Finance* 17, 117-160.
- Chernov, M. and E. 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, 407-458.
- Corrado, C. and T. Su, 1996, Skewness and kurtosis in S&P 500 index returns implied by option prices, *Journal of Financial Research* 19, 175-192.
- Cox, J. and S. Ross, 1976, The valuation of options for alternative stochastic processes, *Journal of Financial Economics* 3, 145-166.

- Frontini, M. and A. Tagliani, 1994, Maximum entropy in the finite Stieltjes and Hamburger moment problem, *Journal of Mathematical Physics* 35, 6748–6756.
- Garcia, R., E. Ghysels and E. Renault, 2003, The econometrics of option pricing, forthcoming in *Handbook of Financial Econometrics*, Yacine Ait-Sahalia and Lars Peter Hansen eds, Elsevier-North Holland, Amsterdam.
- Heston, S.L., 1993, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* 6, 327-343.
- Heston, S. L. and S. Nandi, 2000, A closed-form GARCH option valuation model, *Review of Financial Studies* 13, 585-625.
- Jarrow, R. and A. Rudd, 1982, Approximate valuation for arbitrary stochastic processes, *Journal of Financial Economics* 10, 347-369.
- Jaynes, E. T., 1979, Where do we stand on maximum entropy?, *Maximum Entropy Formalism*, R. D. Levine and M. Tribus, eds. Cambridge, MA: MIT Press, 115-118.
- Madan, D. B. and F. Milne, 1994, Contingent claims valued and hedged by pricing and investing in a basis, *Mathematical Finance* 4, 223-245.
- Rockinger, M. and E. Jondeau, 2002, Entropy densities with an application to autoregressive conditional skewness and kurtosis, *Journal of Econometrics* 106, 119-142.
- Rompolis, L. S. and E. Tzavalis, 2005, Retrieving risk neutral moments from option prices, Mimeo, Department of Accounting and finance, Athens University of Economics and Business.
- Rompolis, L. S. and E. Tzavalis, 2007, Recovering risk neutral densities from option prices: A new approach, *Journal of Financial and Quantitative Analysis*, forthcoming.
- Rubinstein, M., 1994, Implied binomial trees, *Journal of Finance* 49, 771-818.
- Shannon, C. E., 1948, The mathematical theory of communication, *Bell Systems Technical Journal* 27, 379-423.
- Shohat, J. A. and J. D. Tamarkin, 1943, *The problem of moments*, Providence, RI: AMS Mathematical Survey.

Stutzer, M., 1996, A simple nonparametric approach to derivative security valuation, *Journal of Finance* 51, 1633-1652.

Wu, X., 2003, Calculation of maximum entropy densities with application to income distribution, *Journal of Econometrics* 115, 347-354.

Table 2: Out-of-Sample Option Pricing Forecasting Performance

	OTM puts	ATM puts	ITM puts	All puts	ITM calls	ATM calls	OTM calls	All calls	All options
MAE									
MED-RNM	1.34	5.93	10.05	3.98	9.66	5.78	1.59	5.19	4.54
MED-BK	1.93	7.16	10.65	4.69	9.53	6.82	2.21	5.85	5.22
MPE (%)									
MED-RNM	32.39	23.27	16.52	27.33	11.82	22.13	45.53	28.71	27.81
MED-BK	63.45	31.18	18.22	47.68	11.62	28.66	67.48	39.98	44.09

Notes: This table presents the mean absolute error (MAE) and the mean percentage error (MPE) for our empirical exercise forecasting out-of-sample option prices. This is done across different moneyness levels defined as follows. A call (put) option is said to be in-the-money (ITM) if $S_t/K \geq 1.03$ ($S_t/K \leq 0.97$), at-the-money (ATM) if $S_t/K \in (0.97, 1.03)$ and out-of-the-money (OTM) if $S_t/K \leq 0.97$ ($S_t/K \geq 1.03$).

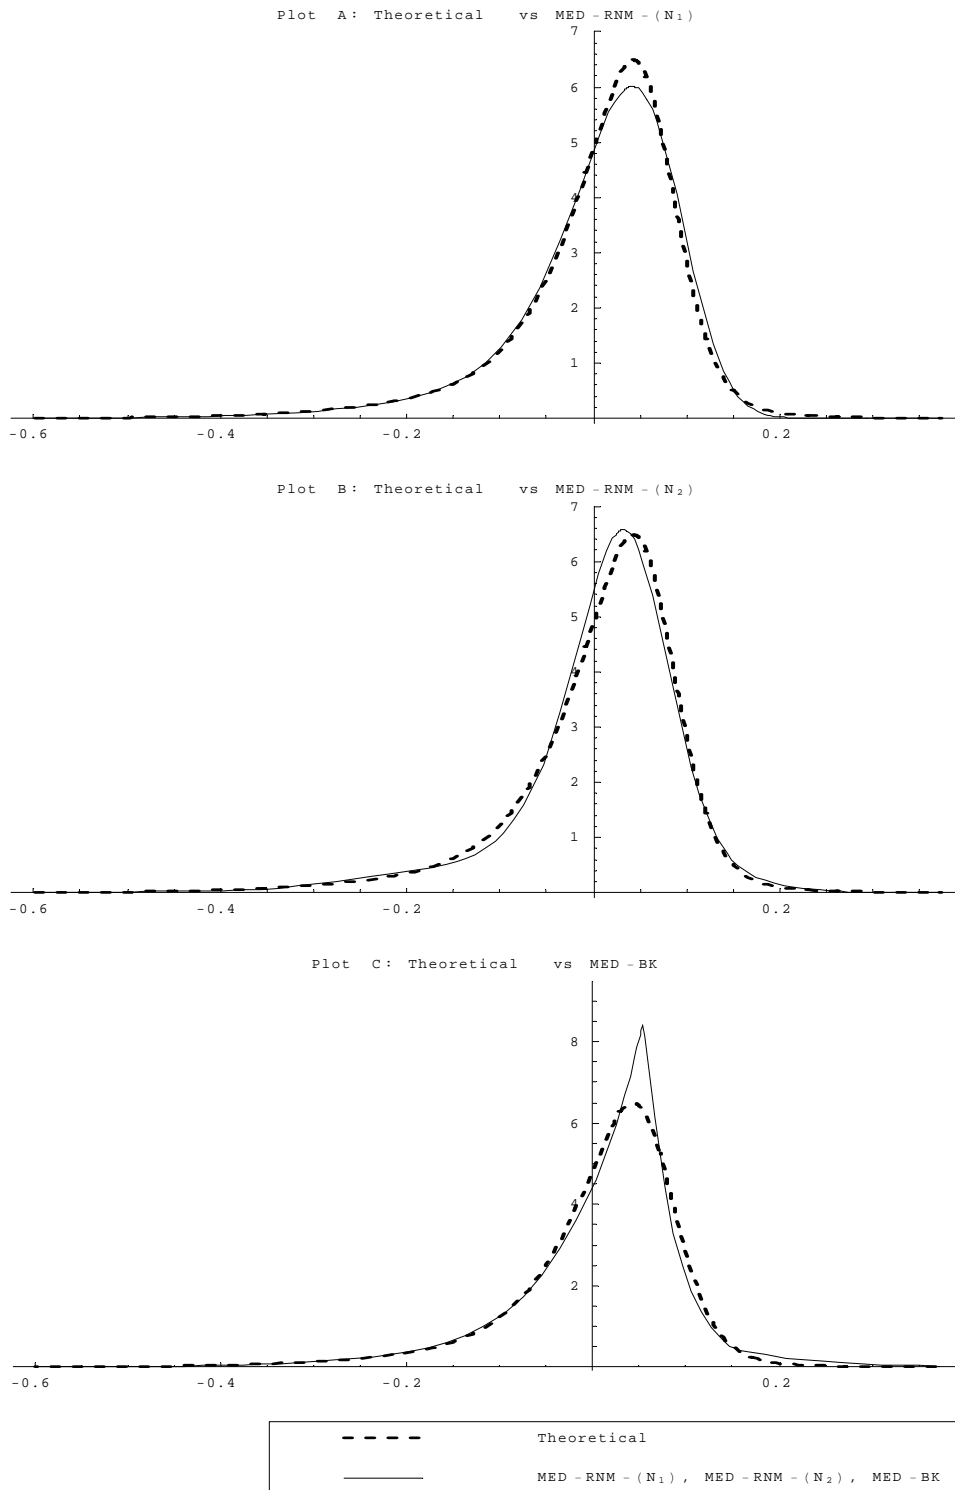


Figure 1: MED estimates of the RND implied by the SVJ model




Figure 2: MED estimates of the RND