The entropy-based implied volatility and its information content*

Xiao Xiao†[1,2,3] and Chen Zhou ‡[1,2,4]

[1]Erasmus School of Economics, Erasmus University Rotterdam

[2]Tinbergen Institute, The Netherlands

[3]Duisenburg School of Finance

[4]De Nederlandsche Bank

June 4, 2015

Abstract

This paper investigates the maximum entropy method for extracting the implied volatility. Numerical examples show that the maximum entropy method outperforms the Black-Scholes model and the model-free method in extracting the implied volatility. This phenomenon is more pronounced if the risk neutral distribution of the underlying asset deviates from the normal distribution, or if the number of available options is limited. In addition, the maximum entropy method allows for constructing confidence intervals around the implied volatility and extracting the implied skewness and kurtosis. We apply the maximum entropy method to the S&P500 index options for predicting future realized volatility. We find that the entropy-based implied volatility subsumes all information in the Black-Scholes implied volatility and historical volatility and has more predictive power than the model-free implied volatility. Lastly, entropy-based variance risk premium performs better than other alternatives in predicting future monthly market returns.

JEL Classification: C14, G13, G17

*The authors thank Casper de Vries, Wolfgang Hardle, John Crosby, Matthias Fengler, Rex Wang, Isabel Casas, Frank Liu, Andrei Lalu, Yang Liu and Erkki Slide for helpful discussions and comments. This paper was circulated under the title "Option implied risk measures: a maximum entropy approach". Views expressed are those of the authors and do not necessarily reflect official positions of De Nederlandsche Bank.

†Department of Economics, Erasmus University Rotterdam, P.O.Box 1738, 3000DR, The Netherlands. Email: xiao@ese.eur.nl.

‡Department of Econometrics, Erasmus University Rotterdam, P.O.Box 1738, 3000DR, The Netherlands. Email: zhou@ese.eur.nl. Economics and Research Division, De Nederlandsche Bank, 1000AB Amsterdam, The Netherlands. Email: c.zhou@dnb.nl.

Keywords: Nonparametric estimation, risk neutral distribution, skewness, kurtosis

1 Introduction

Investors use options to hedge their positions against unfavorable future movements of asset prices. Consequently, option prices reflect investors' perceptions on such movements. Risk measures implied by option prices, which are the characteristics of the risk neutral distribution of the underlying asset, can therefore be informational superior to their historical counterparts in forecasting the risk of the underlying asset. With a large literature emphasizing on the information content of option implied risk measures, little efforts have been devoted to examine whether these measures actually capture the characteristics of the risk neutral distribution. The situation is even more in doubt when the method of extracting option implied risk measures is based on certain parametric assumptions without empirical validation. An example reflecting this critique is the working horse methodology in practice, the Black-Scholes formula. Extracting the implied volatility by the Black-Scholes formula (Black and Scholes (1973)) using options with different strik prices results in the well-known volatility smile or smirk. This is against the uniquely defined volatility in the model. Furthermore, Neumann and Skiadopoulos (2013) showed that the implied skewness calculated from S&P500 index options is consistently negative and the implied kurtosis is always higher than three during the period from 1996 to 2010. The empirical evidence points to the fact that the risk neutral distribution observed in the financial market is inconsistent with the Gaussian assumption in the Black-Scholes model. Therefore, whether the Black-Scholes implied volatility captures the actual volatility of the risk neutral distribution is very much in doubt. This critique may hold true for any parametric method for extracting the implied volatility.

In this paper, we investigate a non-parametric method, the maximum entropy (ME) method , for extracting the option implied risk measures, in particular the implied volatility. We show at least four advantages of the entropy-based implied volatility (EBIV). First, the

ME method does not rely on any parametric model while allowing the data to determine the shape of the risk neutral distribution. Second, different from the model-free method in Britten-Jones and Neuberger (2000) and Bakshi et al. (2003), the ME method does not require a large number of options with strike prices covering the entire support of the return distribution. Even with limited number of options, this method can produce more accurate estimates than the Black-Scholes implied volatility (BSIV) and the model-free implied volatility (MFIV). Third, the ME method also allows for calculating implied skewness and implied kurtosis. Last but not least, the ME method allows for constructing confidence intervals around the implied volatility by utilizing a nonparametric analog of likelihood ratio statistics proposed by Kitamura and Stutzer (1997).

Using non-parametric methods to extract risk measures of the risk neutral distribution has been studied extensively in the literature, in particular the so-called model-free method. This stream of literature started from the pioneer work of Britten-Jones and Neuberger (2000) and Bakshi et al. (2003), with following-up works in Dennis and Mayhew (2002), Jiang and Tian (2005), Bali and Murray (2013),Neumann and Skiadopoulos (2013), and DeMiguel et al. (2014) . The model-free method derives the risk measures of the risk neutral distribution entirely from no-arbitrage conditions. The implied voliatlity can be approximated by a linear combination of European call and put option prices with strikes spanning the full range of possible values for the underlying asset at maturity. However, the approximation turns to be an exact relation only if the underlying risk neutral distribution follows a diffusion process. By contrast, Jiang and Tian (2005) showed the truncation error and the discretionary error of the model-free method under the stochastic volatility and random jump (SVJ) model. Although the estimation error is admissible under certain conditions, it tends to be larger when the underlying distribution is more negatively skewed, when the available number of options is limited and when the market is more volatile. These are exactly the situations in which the ME method performs more accurate and robust.

This paper also contributes to the literature on testing the information content of implied risk measures. Several studies find that the implied volatility is superior to the historical volatility of the underlying asset in predicting future realized volatility; see Day and Lewis (1992), Canina and Figlewski (1993), Lamoureux and Lastrapes (1993), Christensen and

Prabhala (1998), Fleming (1998), Blair et al. (2001) and Busch et al. (2011). In addition, DeMiguel et al. (2014) showed that using the implied risk measures can improve the selection of mean-variance portfolios which leads to a better out-of-sample performance. We test the information content of the EBIV, and compare it with that of the BSIV and MFIV.

The ME method for extracting the option implied risk measures is closely related to the principle of maximum entropy proposed in Buchen et al. (1996). Buchen et al. (1996) found that given simulated option prices at different strikes, estimating the risk neutral distribution by maximizing the entropy can accurately fit the true risk neutral density. In this paper, we apply this method to obtain the characteristics of the risk neutral distribution, such as the EBIV and other implied risk measures. Different from Buchen et al. (1996), we focus on the implied risk measures rather than the full risk neutral distribution. In addition, the empirical goal is to compare the estimation error and forecasting ability of the EBIV to the other alternatives such as the BSIV and MFIV. Lastly, we provide a novel methodological contribution constructing confidence intervals around the EBIV based on Kitamura and Stutzer (1997). This study is also related to Stutzer (1996). Stutzer (1996) investigated options pricing using the maximum entropy method. By contrast, we conduct the reverse procedure to extract information from option prices.

We have three main findings. First, we propose a set of estimators for the volatility, skewness and kurtosis of the risk neutral distribution. We show that they are more accurate than their counterparts using the Black-Scholes or the model-free methods. In particular, when the risk neutral distribution exhibits heavy tail and negative skewness, the EBIV is more accurate than the BSIV and MFIV. If the number of available options is reduced or the true volatility increases, the estimation error of MFIV becomes more salient while the EBIV remain robust under different specifications. Second, to the best of our knowledge, this paper is the first to construct confidence intervals around implied volatility using the ME method. We construct the confidence interval around the EBIV and calculate the coverage ratio under different distributions. The coverage ratios are found to be close to the confidence levels. Third, using the prices of S&P500 index options, we provide both in-sample and out-of-sample evidence that the EBIV performs better than the BSIV or MFIV in forecasting future realized volatility, and the variance risk premium derived from the EBIV performs better that that derived from the EBIV or MFIV in forecasting future stock returns. More specifically, we find that the EBIV subsumes all information in the BSIV and the historical volatility in predicting future realized volatility. In addition, the EBIV has a higher forecasting power than the MFIV. In the out-of-sample analysis, the EBIV continues to provide superior forecasts to the BSIV and MFIV. More importantly, the EBIV performs the best in high volatility regimes.

The remainder of the paper proceeds as follows. Section 2 discusses the estimation of the option implied risk measures using the ME method. Section 3 compares the accuracy of different implied risk measures and shows the coverage ratio of the confidence intervals around the EBIV. The information content of different implied volatilities are compared in Section 4. Section 5 concludes.

2 The entropy-based implied volatility

In this section, we first introduce the ME method for extracting the risk neutral distribution from option prices. The implied volatility is consequently calculated from the extracted risk neutral distribution. An important feature of this method is that it allows for constructing the confidence interval around the implied volatility as explained in Section 2.2

2.1 The maximum entropy method

The absence of arbitrage guarantees the existence of a risk neutral probability measure under which the price of any security equals to the expectation of its discounted payoffs. In the reminder of the paper, all probability measures refers to the risk neutral probability measure, unless otherwise specified.

Let X_t be a random variable that represents the gross return of a stock at the expiry time t in the future. Denote S_0 as the current price of the stock. At time 0, the value of a call option with strike price K equals to the expectation of its discounted payoff at time t as follows

$$C = \mathbb{E}[\max(S_0 X_t - K, 0)]/r_t, \tag{1}$$

where r_t is the gross risk free rate from time 0 to t. In a discrete state setting, we assume that there are n possible states for X_t as X_{t1},...,X_{tn}, with probabilities q_1, ..., q_n respectively.

In addition, we require $q_i > 0$ and $\sum_{i=1}^{n} q_i = 1$. The pricing equation (1) can be rewritten as:

$$C = \sum_{i=1}^{n} q_i(\max(S_0 X_{ti} - K, 0))/r_t.$$

A similar pricing equation can be correspondingly established for put options.

The number of possible states is usually much larger than the number of available options. Consequently, the pricing equations on available options are not sufficient to uniquely determine the underlying risk neutral distribution. Buchen et al. (1996) show that if the pricing equations are regarded as constraints on the risk neutral distribution, by maximizing the entropy, defined as

$$\ell_{ET} = -\sum_{i=1}^{n} q_i log(q_i),$$

a unique optimal distribution can be obtained. Since the entropy measures the amount of missing information, the optimal distribution is the least prejudiced distribution compatible with the given constraints. From the viewpoint of statistical inference, there is no reason to prefer any other distribution, if the only available information is the pricing equations (Buchen et al. (1996)).

More specifically, suppose there are k_1 call options with strike price $K_c(j)$ and option price $C(j), \, , j = 1, ..., k_1$. In addition, there are k_2 put options with strike price $K_p(j)$ and option price $P(j), \, , j = 1, ..., k_2$. Then the constraints based on the call and put options are:

$$C(j) = \sum_{i=1}^{n} q_i(\max(S_0 X_{ti} - K_c(j), 0))/r_t, \; j = 1, ..., k_1 \tag{2}$$

$$P(j) = \sum_{i=1}^{n} q_i(\max(K_p(j) - S_0 X_{ti}, 0))/r_t, \; j = 1, ..., k_2 \tag{3}$$

$$\sum_{i=1}^{n} q_i = 1, \; q_i > 0,$$

To present the constraints in a concise manner, we express the $k_1 + k_2$ constraints in equation

(2) and (3) as:

$$\sum_{i=1}^{n} q_i g_j(X_{ti}) = 0,\ j = 1, ..., k, \tag{4}$$

where $k = k_1 + k_2$. The Lagrange function associated with the constrained optimization problem is:

$$\mathcal{L} = \sum_{i=1}^{n} q_i log(q_i) + \gamma(\sum_{i=1}^{n} q_i - 1) + \lambda'(\sum_{i=1}^{n} q_i g(X_{ti})),$$

where $\gamma \in \mathbb{R}$ and $\lambda \in \mathbb{R}^m$ are the Lagrange multipliers, $g(X_{ti}) = (g_1(X_{ti}), ..., g_k(X_{ti}))^T$. The first order conditions for $\mathcal{L}$ are solved by:

$$\hat{q}_i = \frac{\exp(\hat{\lambda}' g(X_{ti}))}{\sum_{i=1}^{n} \exp(\hat{\lambda}' g(X_{ti}))},\ i = 1, ..., n \tag{5}$$

$$(\hat{\lambda}_1, ..., \hat{\lambda}_k) = \arg\min \sum_{i=1}^{n} \exp(\lambda' g(X_{ti})), \tag{6}$$

where λ_j is the Lagrange multiplier of the jth constraint in equation (4). The value of $(q_1, ..., q_n)$ that maximizes ℓ_{ET} under the k constraints is also called the exponential tilting estimator. Notice that the estimated $\hat{q}_i$ is presented as a function of the Lagrange multipliers which are uniquely solved from minimizing a strictly convex function.

After estimating the risk neutral probabilities associated to the predetermined states, the EBIV is calculated as:

$$EBIV = V^Q = \sqrt{\sum_{i=1}^{n} \hat{q}_i (log(X_{ti}) - \mu^Q)^2},\ \mu^Q = \sum_{i=1}^{n} \hat{q}_i log(X_{ti}).$$

We choose to calculate the EBIV of the continuously compounded returns rather than the discrete return in order to compare it later with the BSIV, because the BSIV is also based on the continuously compounded return.

The entropy is also named as the Kullback-Leibler divergence measure, which is a member of the Cressie-Read divergence family. Correspondingly, maximizing any member in the Cressie-Read divergence family results in a non-parametric method for extracting a probability distribution under given constraints. Such methods form the class of generalized empirical likelihood methods. A notable example of such a method is the so-called empirical likelihood (EL) method, in which the Euclidean divergence is maximized. However, there are at least two reasons why the ME method is preferred over the EL method. First, the ME method provides a robust performance with respect to the variation in the possible states. Regardless whether we simulate states from a certain distribution, or enforce a series of equally distanced values as states, the extracted risk neutral distribution do not change provided that the chosen states cover the range of the strike prices. On the contrary, the result following the EL method depends largely on the choice of the states[1]. Second, Theorem 1 in Schennach (2007) shows that the EL method suffers from a dramatic degradation of its asymptotic properties under even the slightest amount of misspecification.

2.2 Constructing the confidence interval of the EBIV

An important feature of the ME method is that it facilitates the construction of a confidence interval around the EBIV as follows. First, given the level of the mean μ_Q, consider a hypothesis testing problem as $H_0: \ V^Q = V_0^Q$, where V_0^Q is a given level of volatility to be tested. Kitamura and Stutzer (1997) proposed the following likelihood ratio testing statistics:

$$LR_T = 2n[\log M(V_0^Q) - \log M(\hat{V}^Q)] \xrightarrow{d} \chi_1^2$$

where $M(\hat{V}^Q) = \frac{1}{n}\sum_{i=1}^n \exp(\hat{\lambda}' g(X_{ti}))$ is the minimized function (6) under the k constraints in (4), and $M(V_0^Q) = \frac{1}{n}\sum_{i=1}^n \exp(\tilde{\lambda}' g(X_{ti}) + \tilde{\lambda}_{k+1} g_{k+1}(X_{ti}))$ is the achieved minimal value in the optimization problem

$$(\tilde{\lambda}_1, ..., \tilde{\lambda}_k, \tilde{\lambda}_{k+1}) = \frac{1}{n} \arg\min \sum_{i=1}^n \exp(\lambda' g(X_i) + \lambda_{k+1} g_{k+1}(X_i)),$$

with an additional $(k+1)$-th constraint as $g_{k+1}(X_{ti}) = (X_{ti} - \mu_Q)^2 - (V_0^Q)^2$.

For each V_0^Q, the null is rejected or not rejected under a given confidence level α. Then, one may vary the value of the V_0^Q around the estimated implied volatility $\hat{V}^Q$. Since LR_T

[1]Simulation results on the comparison of the two methods are upon request.

is increasing for $V_0^Q > \hat{V}^Q$ and decreasing for $V_0^Q < \hat{V}^Q$, given the confidence level α, there exists two values V_L^Q and V_H^Q such that H_0 is rejected for $V_0^Q < V_L^Q$ and for $V_0^Q > V_H^Q$ while H_0 is not rejected for $V_0^Q \in [V_L^Q, V_H^Q]$. Then the interval $[V_L^Q, V_H^Q]$ is regarded as the confidence interval of V^Q with the given confidence level α. Obviously, we have that $\hat{V}^Q \in [V_L^Q, V_H^Q]$.

In this procedure, we assume that the mean of the continuous compounded return μ^Q is fixed when varying the constraint based on V_0^Q. Such an assumption is partially supported by the fact that the mean of the discrete return is fixed provided that both at-the-money call and put option prices are available. According to the put-call parity, the mean of the discrete stock return is derived as:

$$\sum_{i=1}^{n} q_i X_{Ti} = \frac{(C_{atm} - P_{atm}) r_t + 1}{S_0},$$

where C_{atm} and P_{atm} are at-the-money call and put option prices. Approximately, we regard the mean of the continuously compounded return as fixed at $\sum_{i=1}^{n} \hat{q}_i log(X_{Ti})$ when we vary the V_0^Q. In the simulation, we do observe that the means of the discrete return and the continuously compounded return are close with the difference at a negligible magnitude.

3 The performance of the EBIV: a numerical study

In this section, we compare the performance of the ME method, the model-free method, and the Black-Scholes (BS) model for backing out implied volatility from option prices. The ME method is shown to be more accurate than the other two methods when there are less number of options available and when the underlying distribution is heavy-tailed with non-zero skewness.

3.1 Three methods for backing out implied volatility

The most conventional method for backing out implied volatility from option prices is to use the BS model. We first calculate the implied volatilities from all available option prices, and then take the average as the estimate. This results in the BSIV.

When the underlying return distribution deviates from log-normal, taking the average of the BS implied volatilities may not be an efficient way to aggregate information across different strike prices. Britten-Jones and Neuberger (2000) and Bakshi et al. (2003) propose the MFIV which is independent of the pricing models. It is derived entirely from no-arbitrage conditions and can be considered as a linear combination of European call and put option prices with strikes spanning the full range of possible values for the underlying asset at maturity. The MFIV is defined as follows:

$$MFIV = e^{rT}V - \mu^2,$$
$$V = \int_{S_0}^{\infty} \frac{2(1-\ln[\frac{K}{S_0}])}{K^2} C(K,T)dK + \int_0^{S_0} \frac{2(1+\ln[\frac{S_0}{K}])}{K^2} P(K,T)dK,$$

where $C(K,T)$ ($P(K,T)$) is the call (put) option price with strike price K and maturity T and μ is the mean of the risk neutral return, which can also be replicated by a option portfolio. The details for calculating μ are given in Appendix 6.1. In a discrete setting, the term V can be approximated as:

$$V \approx \sum_{i=1}^{m} \frac{2(1-\ln[\frac{K_i}{S_0}])}{K_i^2} C(K_i,T)(K_{i+1}-K_i) + \sum_{j=m+1}^{n} \frac{2(1+\ln[\frac{S_0}{K_j}])}{K_j^2} P(K_j,T)(K_j-K_{j+1}),$$

where $K_1 > K_2 > ... > K_m > S$ and $K_{m+1} < K_{m+2} ... < K_n < S$ are the strike prices of the available options. Practically, the number of available options is limited. Therefore, we apply a curve-fitting method to interpolate and extrapolate the prices of the unavailable options as follows. Available option prices are first mapped to implied volatilities using the BS model. For options with strike prices within the available range, following Bates (1991) and Jiang and Tian (2005), we use cubic spines to interpolate their implied volatilities. For options with strike prices beyond the available range, we use the end-point implied volatility as their implied volatilities. Then we use the BS model to transform the obtained implied volatilities for unavailable options back to option prices. Eventually we have the option prices with moneyness ranging from 0.35 to 1.65 with interval 0.002. All the prices of these options are used for calculating the MFIV.

Since the model-free method is derived under the diffusion assumption, it only holds approximately when the underlying distribution is skewed or heavy-tailed. Jiang and Tian (2005) show that the estimation error of MFIV is admissible under certain conditions. By contrast,the ME method is a completely nonparametric method which is not derived under conditions. Hence, it is expected that the EBIV can perform better than the MFIV when the underlying distribution is more negatively skewed or heavy-tailed.

3.2 The underlying risk neutral distributions in the numerical study

We consider four data generating processes to generate the underlying continuously compound returns. To begin with, we consider the case that stock price follows a geometric Brownian motion under the risk neutral measure:

$$dS_t = rS_t dt + \sigma S_t dw_t,$$

where S_t is the stock price at time t, r is the risk-free rate, σ is the constant instantaneous volatility of the process and dw_t is the increment in a standard Wiener process. Throughout the section, we employ an annual risk-free rate r at 5%, an annual volatility σ at 20% (or 40%) and the initial stock price S_0 at 100. Under this model, the risk neutral distribution of the continuously compounded T-year returns $\ln(R_T)$ is normally distributed:

$$\ln(R_T) \sim N((r - \frac{1}{2}\sigma^2)T, \sigma^2 T). \tag{7}$$

The mean of the risk neutral distribution, $(r - \frac{1}{2}\sigma^2)T$, ensures that the expectation of R_T is e^{rT} under the risk neutral measure. Note that the BS model and the model-free method are derived based on this assumption, these two methods should provide accurate estimates for the implied volatility in this case.

We also consider the underlying distributions deviating from the normal distribution, in particular, the Student-t and the skewed Student-t distributions. More specifically, the continuously compound return is given as

$$\ln(R_T) \sim (r - \frac{1}{2}\sigma^2)T + \sigma\sqrt{T}\,\epsilon. \tag{8}$$

Firstly, we consider that ϵ follows standardized Student-t distribution with degree of freedom 5. Then, we consider the case that ϵ follows the skewed Student-t distribution ($skewt(\eta, \lambda)$) suggested in Hansen (1982) with degree of freedom η and skewness parameter λ. We use two sets of parameters: $\eta = 5$, $\lambda = -0.3$ and $\eta = 5$, $\lambda = -0.7$. The latter is more negatively skewed than the former.

Notice that although the mean return is comparable with that in (7), the pricing equation, $ER_T = e^{rT}$, does not hold if ϵ follows the Student-t or the skewed Student-t distributions, though it remains approximately true.

3.3 Results

Based on the risk neutral distributions specified in section 3.2, we calculate the call and put option prices using numerical integration for several moneyness with one month to expiration. Results for other expiration horizons are provided in the robustness check in Section 3.4. The call and put option prices with strike price K and maturity T are calculated by numerical integration:

$$C(K,T) = \int_{K/S_0}^{\infty} (S_0 R_T - K) f(R_T) dR_T / r_T \tag{9}$$

$$P(K,T) = \int_{0}^{K/S_0} (K - S_0 R_T) f(R_T) dR_T / r_T \tag{10}$$

where $f(R_T)$ is the density function of R_T. The density function when $\ln(R_T)$ follows the skewed Student-t distribution is provided in the Appendix 6.2.

Following Bakshi et al. (2003), we only consider out-of-the-money (OTM) options and at-the-money (ATM) options due to the liquidity reason in the option market. Moreover, in-the-money call and put option prices can be derived from put-call parity under the no-arbitrage condition. Consequently, they do not provide additional information for extracting the implied volatility.

We consider different ranges of strike prices which result in different estimation accuracy. In the first case, we specify the moneyness (K/S_0) of call options from 1 to 1.15 with equal interval 0.025, and the moneyness of put options from 0.85 to 1 with the same interval. There are 14 options in total. In the second case, we reduce the number of available options and only consider six options: call options with moneyness 1, 1.05, 1.1 and put options with moneyness 0.95, 0.975, 1. By comparing the two cases, we evaluate the performance of the three methods with different number of available options. The calculated option prices with the chosen moneynesses under different distributions are reported in Table 1.

Table 2 reports the estimated implied volatilities under different distributions using the three methods. The first row shows the true volatility of the underlying distribution and the second row shows the number of options. Under the normal distribution, the all three methods provides very accurate estimates, while for other risk neutral distributions, all methods have some estimation errors. We calculate the relative improvements of the MFIV and EBIV compared to the BSIV in the parenthesis, which is defined as $\frac{|XXIV - TrueVolatility|}{|BSIV - TrueVolatility|}$, where XXIV is either MFIV or EBIV.

From Table 2, we find that when the underlying distribution is heavy-tailed or negatively skewed, the EBIV estimates are closer to the true value than both the BSIV and MFIV. The improvement is substantial.

Although the MFIV performs better than the BSIV under heavy-tailed or skewed distributions, the estimation error increases when the underlying distribution is more negatively skewed. However, the estimation improvement of the EBIV remains robust across different specifications. When we decrease the available number of options or increase the true volatility from 0.2 to 0.4, the better performance of EBIV compared to MFIV becomes more evident. For example, an estimation error of the MFIV under the $skewt(5, -0.7)$ distribution with 6 options and true volatility 0.4 is even larger than that of the BSIV, while the EBIV has an estimation error as low as 31% of that of the BSIV. The unrobust performance of the MFIV may be attributed to the increase of the truncation error and extrapolation error, whereas the ME method does not suffer from such a problem.

An additional advantage of the ME method is that one may construct the confidence interval around the estimate of the implied volatility. Table 3 shows the coverage rate of the EBIV confidence interval under the four distributions with six option prices. To be more specific, we simulate 10000 states from the true distribution for 100 times[2]. For each set of

[2] When simulating the states, kurtosis of the simulated sample might differ from the true value in many

simulated states, we first calculate option prices, and then back out the EBIV using the ME method from the calculated option prices. Then we construct the confidence interval around the EBIV and examine whether the true volatility falls into the constructed interval. Across the 100 simulations, we count the number of "coverage", i.e. when the true volatility falls into the confidence interval, and divide that number by 100. From Table 3, we find that the coverage rates of the confidence interval are close to the confidence levels under all four distributions.

3.4 Robustness and discussion

Results from Section 3.3 show that the EBIV aggregates information in the option prices more efficiently than the MFIV, when the underlying distribution exhibits heavy tail and when the number of options is limited. We now conduct robustness check to ensure the generality of our findings.

In Table 4 and Table 5, we provide the estimated implied volatilities based on option prices with three-month and one-year maturity, respectively. In addition to the 14 or 6 options with symmetric moneyness, we consider asymmetric moneyness with 3 call options (moneyness ranging from 1 to 1.05) and 7 put options (moneyness ranging from 0.85 to 1), with equal interval 0.025. This is to reflect the situation that there are more available put options than call options. Column 5 an 8 in Table 4 and Table 5 report the results of asymmetric moneyness. Results in these two tables show that the EBIV still has a better estimation accuracy than the MFIV and BSIV for longer maturity options.

The second generalization is to consider more complex data generating process for the risk neutral distribution, for example, the SVJ model. Further, we may consider estimating higher moments such as implied skewness and kurtosis. The SVJ model has been applied for pricing options in Bakshi et al. (1997) and for illustrating truncation errors of model-free

simulations. The reason is that sample kurtosis is very sensitive to extreme observations. If there are no extreme observations in the simulated sample, sample kurtosis is downward biased. To alleviate this bias, we only consider the simulated sample if the sample kurtosis is higher than 80% of the true kurtosis.

implied volatility in Jiang and Tian (2005). The model is specified as:

$$\frac{dS_t}{S_t} = \sqrt{V_t}dW_t + J_t dN_t - \mu_J \lambda dt,$$
$$dV_t = (\theta_v - \kappa_v V_t)dt + \sigma_v \sqrt{V_t} dW_t^v,$$
$$dW_t dW_t^v = \rho dt,$$

where N_t follows a homogeneous Poisson process with jump intensity λ and $\ln(1+J_t)$ follows a normal distribution $N(\ln(1+\mu_J) - \frac{1}{2}\sigma_J^2, \sigma_J^2)$. If $\lambda = 0$, the model reduces to the Heston (1993) model. We choose the parameters as $\kappa_v = 1$, $\sigma_v = 0.25$, $\rho = 0$, $\lambda = 0.5$, $\theta_v = V_0\kappa_v$, $V_0 = 0.1854^2$. In addition, to evaluate the impact of the jump process, we choose two sets of parameters for the jumps: (1) $\mu_J = -1.75$, $\sigma_J = 0.5$, (2) $\mu_J = -0.075$, $\sigma_J = 2.5$, correspondingly. The volatility is 0.203 for the first set of parameters and 0.453 for the second set of parameters. Since the unconditional return distribution are unknown here, we opt to use simulations to get option prices and the true skewness and kurtosis. First, we simulate 21 daily returns to get one monthly return, and repeat this for 100,000 times. Second, option prices, the true volatility, skewness and kurtosis are calculated based on the simulated monthly return. Third, we estimate the implied volatility by the BS model, the model-free method and the ME method, and compare the estimates with the true volatility. In addition, we calculate the implied skewness and kurtosis using the model-free method and the ME method. The results for options with one-month maturity are reported in Table 6. In Table 6, the second column for each parameter set reports the simulated moments and the column 3-5 reports the implied moments estimated from different methods. The results show that the ME method gives the most accurate estimation in all cases, and is especially robust when the unconditional distribution has higher volatility, more negative skewness and higher kurtosis.

Lastly, we check the robustness of the ME method by focusing on its fundamental step: the estimated risk neutral distribution. For each risk neutral distribution, we compare the sample of the simulated distribution (blue bars) and the estimated density produced by the ME method (red lines) in Figure 1. More specifically, the estimated risk neutral densities in these figures are estimated from 14 options with one year maturity. The figures show that the risk neutral density estimated by the ME method matches the true density in all four cases. Option prices with different moneynesses essentially provide information on different parts of the distribution.

To conclude, the ME method provides more accurate estimates of option implied volatility than the BS model and the model-free method. Our main findings are robust to the choice of different number of options, maturities and the data generating process. In addition, the ME method can also be applied for other higher moments of the risk neutral distribution, because it provides an accurate estimation for the risk neutral distribution.

4 The information content of EBIV

In this section, we conduct an empirical analysis to explore the information content of the EBIV using the S&P500 index option traded in the Chicago Board Options Exchange (CBOE). We first review the basic statistical properties of the various volatility measures and then investigate their relative performance as predictors of the subsequent realized volatilities of the underlying S&P500 index. We also analyze the forcasting power of the variance risk premia derived from different implied volatility on the subsequent returns of the S&P500 index.

4.1 Data

Our sample period covers from January 1996 to August 2014. We get the S&P500 index price data from The Center for Research in Security Prices database. We obtain the S&P500 index options data from the Ivy DB database of OptionMetrics. Continuously-compounded zero-coupon interest rates are also obtained from OptionMetrics as a proxy for the risk-free rate. From the CBOE, we get daily levels of the newly calculated VIX index[3] and match them with the trading days on which options with one month expiration are traded.

Our analysis is conducted based on call and put options quoted on the S&P500 index with 30 days expiration. We choose one-month maturity because the options with one month

[3] Although the CBOE changed the methodology for calculating the VIX in September 2003, they have backdated the new index using the historical option prices.

to expire are more actively traded than with other maturities. From January 1996 to February 2007, in each month, there is only one day on which options with 30 days expiration are traded. From March 2007, there are several such days in each month. To avoid the overlapping problem described in Christensen and Prabhala (1998), Christensen et al. (2001) and Jiang and Tian (2005), we select one date in each month from 2006 to 2014, such that the time intervals between any two adjacent dates are the closest to 30 days. With this procedure, there are 222 selected dates in total. Midpoints of the bid-ask spread are used as the option prices instead of the actual trade prices. This follows Jackwerth (2000) who demonstrates that measurement of risk neutral distribution is not sensitive to the existence of spreads.

Table 7 presents the descriptive statistics of the out-of-the-money call and put options. We apply several filters to select the options. First, option quotes less than 3/8 are excluded from the sample. Such low prices may not reflect the true option value due to proximity to tick size. Second, options with zero open interest are excluded from the sample. Third, following Aït-Sahalia and Lo (1998) and Bakshi et al. (2003), we exclude in-the-money options, because they have less liquidity than out-of-the-money options.

We use all available option prices with moneyness between 0.85 to 1.15 to calculate the MFIV. Similar as in Section 3.1, we use cubic spines to interpolate the implied volatility of the options with strike price within the range. For options with strike prices beyond the range, we use the end-point implied volatility to extrapolate. Then we use BS model to transform the obtained implied volatilities back to option prices. Eventually we have option prices with moneyness ranging from 0.35 to 1.65. All the prices of these options are used for calculating the MFIV.

By contrast, we do not use all option prices in the ME method. To calculate the EBIV, we select options with moneyness that are closest to the moneyness ranging from 0.85 to 1.15 with equal interval 0.025. The reason is that if there are many option prices as constraints, the ME method may run into numerical difficulties. This is most likely to happen if the covariance matrix of the constraints in equation (4), $cov(g_i(X_t), g_j(X_t))$, is close to singular. In that case, the numerical solution for the Lagrange multipliers becomes unstable (Buchen et al. (1996)). For BSIV, we calculate the mean of the Black-Scholes implied volatility using all available option prices after the filters procedure.

On each selected trading day, we also calculate the historical realized volatility (RV) in the previous month. Following Christensen and Prabhala (1998), we adopt the realized volatility over the 30 calendar days proceeding the current observation dates as the lagged realized volatility RV_t. It is computed as the sample standard deviation of the daily index returns:

$$RV_t = \sqrt{\frac{1}{30}\sum_{i=1}^{30}(r_{t,i} - \bar{r}_t)^2},$$

where $\bar{r}_t = \frac{1}{30}\sum_{k=1}^{30} r_{t,k}$, and $r_{t,i}$, $i = 1, ..., 30$, are the log index returns on the 30 days preceding to the selected trading day t. All of the volatility measures are expressed in annual terms to facilitate interpretation. Finally, to analyze the predictability of variance risk premium on future stock return, we calculate the log monthly returns of the S&P500 index.

4.2 Descriptive statistics of different volatility measures

Table 8 reports descriptive statistics of the five measures of volatility: RV, VIX, BSIV, MFIV and EBIV. Table 9 shows the correlation matrix of these measures. We first observe that the mean of the four implied volatility measures, VIX, BSIV, MFIV and EBIV, are comparable. All of them exceed the mean of the realized volatility measure RV by about 24%, which is in line with the positive volatility risk premium. Second, the four implied volatility measures are highly correlated, with all correlation coefficients above 0.99. VIX is more correlated with MFIV and EBIV than with BSIV. This may be a consequence of the fact that the three shares the same nonparametric feature by construction.

In Figure 2, we plot the estimated MFIV and EBIV from 1996 to 2014. Figure 1a shows their strong comovement during the period. From Figure 1b, we observe that in most of the time, the differences between the EBIV and the MFIV are negative and small. Occasionally, the differences can be positive and large. This is consistent with the results in Table 4 and 5: when the number of the put options is more than the number of call options, the MFIV may overestimate the true volatility under the low volatility regime ($\sigma = 20\%$). Conversely, the underestimation occurs under the higher volatility regime ($\sigma = 40\%$). Figure 1c plots the spread between the EBIV and the MFIV against the BSIV. By regarding the BSIV as an indicarion of high and low volatility regime, this scatter plot further demonstrates that in a higher volatility regime, the spread is also higher. We attribute this phenomenon to the fact that the model-free method produces less accurate estimates of implied volatility when the market condition becomes more volatile.

Figure 2a presents the confidence interval for the estimated EBIV and Figure 2b shows the length of the confidence interval over time. Lastly, in Figure 4, we provide the estimated risk neutral distributions on four example dates using the ME method. From the figures, we observe that the estimated risk neutral distributions may differ across different market environment.

4.2.1 Forecasting the stock market volatility

Prior research has extensively analysed the information content of the BSIV on predicting the future realized volatility. In particular, recent studies seem to agree on the informational superiority of the BSIV compared to historical volatility. Further, Jiang and Tian (2005) investigated the information content of the MFIV. They find that the MFIV subsumes all information contained in the BSIV and historical volatility. In this paper, we assess the predictive power of the EBIV, and compare it to the other implied volatility measures.

To nest previous research within our framework, we use five competing volatility measures: RV_t, $BSIV_t$, $MFIV_t$, VIX_t and $EBIV_t$ to forecast the realized volatility in the next period RV_{t+1}. To explore the predictive ability of the implied volatility measures, we first include each of them within an in-sample regression separately. We run the following regression

$$RV_{t+1} = \alpha_i + \beta_i x_{i,t} + \epsilon_{i,t+1},$$

with different predictors $x_{i,t} \in I = \{RV_t, BSIV_t, MFIV_t, EBIV_t, VIX_t\}$. R^2 of these regressions captures the proportion of total variation in the ex-post realized volatility explained by the predictors. We also employ encompassing regressions to investigate whether one implied volatility measure subsumes the information in another one. The regression is specified as follows:

$$RV_{t+1} = \alpha + \beta_i x_{i,t} + \beta_j x_{j,t} + \epsilon_t,\ i \neq j,$$

where $x_{i,t}$ and $x_{j,t}$ are two different elements from I.

Table 10 summarizes both univariate and encompassing regressions for the realized volatility in the next month. First, from the estimation results of univariate regressions in Table 10, we observe that the adjusted R^2 is the highest when using the EBIV compared to using other implied volatility measures. Second, in the encompassing regressions with lagged historical volatility RV_t, BSIV, MFIV and EBIV subsume all information in historical volatility, because the coefficient of RV_t becomes statistical insignificant at 5% level if any of the three implied volatility measures enters the regression. Third, the EBIV subsumes all information in the BSIV, because the coefficient of the BSIV is not statistically significant at 5% level when considering both EBIV and BSIV as regressors. Fourth, if using both MFIV and EBIV as regressors, we find that the EBIV subsumes all information in the MFIV. Having a significantly positive sign of the regression coefficients on the EBIV and an insignificantly negative sign on the MFIV indicates that the EBIV plays a dominant role in explaining the variations of the future realized volatility. In summary, the evidence suggests that, among all the implied volatility measures, the EBIV explains the most variation in the next month realized volatility with the highest in-sample fit. It is also notable that even if the MFIV uses more options as inputs, its information content does not overweight that of EBIV.

We then turn to the out-of-sample evidence reported in right columns of Table 10. We use moving window of 100 observations preceding to the period to be forecasted as the estimation window in the regression. Consequently, the remaining 122 months are the forcasting period. The overall measure of forecasting accuracy is the relative RMSE. If we denote $\widehat{RV}_t$ as a forecast for RV_t, the RMSE is formally defined as,

$$RMSE = \frac{\sqrt{\frac{1}{122}\sum_{t=101}^{222}(\widehat{RV}_t - RV_t)^2}}{\frac{1}{122}\sum_{t=101}^{222} RV_t}.$$

The column labelled “All days” covers the full forecasting period. When using one implied volatility measures as a sole regressor, we find that the EBIV provides the best performance among all volatility measures. Forecast precision deteriorates as we move down along the following order: the MFIV, the VIX, the BSIV and the RV. Moreover, the RMSE corresponding to using the EBIV only, 0.477, is lower than using other combinations of volatility measure. This suggests the forecast superiority of the EBIV.

The last three columns in Table 10 report this out-of-sample results in three subperiods. We divide the monthly forecast of future volatility into three subsamples by sorting the BSIV preceding to the forecasted month in ascending order. The results for the months with the low, middle and high BSIV are reported in the “Low”, “Middle” and “high” columns, respectively. First, from the results using a sole regressor, the EBIV performs better than other implied volatility measures in the high volatility regime. All measures exhibit comparable performance in the low volatility regime. Second, the results of the encompassing regressions show that when combining the RV with another IV measure, the combination with the EBIV leads to the best out-of-sample performance. Lastly, it is notable that using the combination of the RV and the EBIV provides the best in-sample fit, and its out-of-sample predictive performance is only slightly worse than using the EBIV only.

As a further robustness check, we provide regression results based on the log realized volatility in Table 11. In this specification, we find that the EBIV performs as good as the MFIV, and they both perform better than other implied volatility measures. While the EBIV subsumes information in the BSIV in both specifications, the MFIV subsumes all information in the BSIV only in the log volatility regression. Table 11 shows that the out-of-sample RMSE when using the EBIV is lower than using other combination of volatility measures.

Finally, we conduct a robustness check by using a broader choice of strike prices, i.e. moneyness ranging from 0.5 to 1.5. The quantitative results remain valid[4]. An additional observation is that the adjusted R^2 when using the BSIV as a sole regressor becomes smaller when we incorporate more options. The BSIV is thus not efficient to integrate the information in a large number of option prices. By contrast, the adjusted R^2 of using the EBIV as the sole regressor increases in this case. Therefore, the ME method can better integrate the information contained in multiple option prices.

[4]Regression results and out-of-sample analysis are available upon request.

4.3 Forecasting stock market returns

At last, we investigate the relation between variance risk premium (VRP) and future market return. The theoretical model in Bollerslev et al. (2009) suggest that VRP may serve as a useful predictor for the future returns. VRP is defined as the difference between market's risk-neutral expectation of future return variation and the current return variation. i.e. the spread between the implied variance and the realized variance. In this paper, we intend to compare the performance of VRP using different measure of implied variance. We use univariate regressions to examine the in-sample fit and out-of-sample forcasting performance.

Denote the ex-post return for month $t+1$ as R_{t+1}, the regressions take the form:

$$R_{t+1} = \alpha_i + \beta_i x_{i,t} + \epsilon_{i,t+1}, \tag{11}$$

where $x_{i,t}$ is one of the item in $I = \{VRP_{BS,t}, VRP_{MF,t}, VRP_{EB,t}\}$, $VRP_{BS,t} = BSIV_t^2 - RV_t^2$ and $VRP_{MF,t}$ and $VRP_{EB,t}$ are calculated based on $MFIV_t$ and $EBIV_t$ in a similar way.

Table 12 reports the results of predicting future monthly returns. In all regressions, the estimated slope coefficients associated with the VRP measures are significant at 5% confidence level. In addition, VRP_{EB} variable explains more variations in future monthly returns than VRP_{BS} and VRP_{MF} with the highest R^2 at 7.1%. The out-of-sample setup is similar to that in Section 4.2.1. The only difference is that the variable to be forecasted here is the monthly stock return R_t instead of the realized volatility RV_t. In the out-of-sample results, VRP_{EB} performs as good as VRP_{MF}. The VRP_{EB} and the VRP_{MF} both perform better than VRP_{BS} and they forecast more accurately in the Medium and the High volatility regime.

All our empirical results point to the direction that the EBIV performs at least comparable with the MFIV in different specifications. In many cases, its information content is of the highest among all available volatility measures, both in terms of in-sample fit and out-of-sample predictive power.

5 Conclusion

This paper provides the first comprehensive investigation on the option implied volatility estimated by the ME method. The ME method extracts the risk neutral distribution of an asset, given a set of option prices at different strikes. The EBIV is then calculated based on the estimated risk neutral distribution. Compared to parametric methods such as the Black-Scholes model, the ME method does not depend on any parametric assumptions. Compared to the MFIV, proposed by Bakshi et al. (2003), the ME method does not require many options with strike prices spanning the full range of possible values for the underlying asset at expiry. Therefore, the ME method combines the advantages in the model-free and the parametric methods: on the one hand, it aggregates information in multiple options with different strikes efficiently; on the other hand, it produces accurate estimates even if the number of options is limited. Lastly, it allows for constructing confidence interval around the estimated implied volatility thanks to a nonparametric analog of likelihood ratio statistics.

With numerical examples, we show that the EBIV has a lower estimation error than the BSIV and the MFIV, particularly when the underlying distribution exhibits heavy tail and non-zero skewness. With limited number of available options or under high volatility level, the accuracy of the EBIV remains robust while the estimation error of the MFIV gets higher. The confidence interval around the EBIV has a coverage ratio that is close to the correct confidence level across various numerical examples. These findings are robust to the choice of different number of options, maturities and the data generating process.

We remark that the ME method also yields estimators for other higher moments of the risk neutral distribution, such as skewness and kurtosis. The estimators perform better than their counterparts when using the model-free method. A potential explanation is that the ME method provides an accurate estimation for the risk neutral distribution.

Using the S&P500 index options, we empirically test the information content of the EBIV in predicting future monthly realized volatilities and index returns. Our in-sample regression results show that the EBIV subsumes all information in the BSIV and the lagged realized volatility. In addition, it has a better explanation power on future realized volatility than the MFIV. In the out-of-sample analysis, the EBIV provides superior forecasts, particularly during high volatility regime. Lastly, the entropy-based variance risk premium explains more variations in future monthly return. To conclude, empirical evidence supports that the EBIV is a more accurate estimator of investors' expectation on future return variation.

A potential drawback of the ME method is that the tail region of the estimated risk neutral distribution largely depends on the options with the highest and lowest strike prices. Given limited number of available options, the estimated density can be less accurate for that part. This may be the reason why estimating implied skewness and kurtosis is less accurate than implied volatility. Improving the estimation of the tail region of the risk neutral density and the option implied skewness and kurtosis is left for further research.

6 Appendix

6.1 Calculation of the Model-free implied moments

The calculation of the model-free option implied moments follows from Bakshi et al. (2003). Let the t-period continuous compounded return be given by: $R_t = \ln[S_t] - \ln[S_0]$. The fair values of the mean, volatility, cubic and quartic contract at time 0 are defined as:

$$M(0,t) = E[e^{-rt}R_t], \;\; V(0,t) = E[e^{-rt}R_t^2], \;\; W(0,t) = E[e^{-rt}R_t^3], \;\text{and}\; X(0,t) = E[e^{-rt}R_t^4].$$

To simplify the notations, we ignore the time period information in the parenthesis in the following equations, for instance $V = V(0,t)$. Further, under the risk neutral measure, the values M, V, W and X can be replicated by the option prices as,

$$\begin{aligned}
M &= 1 - e^{-rt} - \frac{1}{2}V - \frac{1}{6}W - \frac{1}{24}X, \\
V &= \int_S^\infty \frac{2(1-\ln[\frac{K}{S_0}])}{K^2}C(K,t)dK + \int_0^S \frac{2(1+\ln[\frac{S_0}{K}])}{K^2}P(K,t)dK, \\
W &= \int_S^\infty \frac{6\ln[\frac{K}{S}] - 3(\ln[\frac{K}{S_0}])^2}{K^2}C(K,t)dK - \int_0^S \frac{6\ln[\frac{K}{S}] + 3(\ln[\frac{S_0}{K}])^2}{K^2}P(K,t)dK, \\
X &= \int_S^\infty \frac{12(\ln[\frac{K}{S}])^2 - 4(\ln[\frac{K}{S_0}])^3}{K^2}C(K,t)dK - \int_0^S \frac{12(\ln[\frac{K}{S})^2] + 4(\ln[\frac{S_0}{K}])^3}{K^2}P(K,t)dK.
\end{aligned}$$

The t-period risk neutral return mean μ, volatility $MFIV$, skewness $MFIS$, and kurtosis $MFIK$ are given as

$$\begin{aligned}
\mu &= e^{rt}M, \\
MFIV &= \sqrt{e^{rt}V - \mu^2} \\
MFIS &= \frac{e^{rt}W - 3\mu e^{rt}V + 2\mu^3}{(e^{rt}V - \mu^2)^{3/2}}, \\
MFIK &= \frac{e^{rt}X - 4\mu e^{rt}W + 6e^{rt}\mu^2 V - 3\mu^4}{(e^{rt}V - \mu^2)^2}.
\end{aligned}$$

6.2 Density function of log Skewed Student-t distribution

Hansen (1994) suggests the Skewed Student-t distribution, which follows the density function:

$$f(x) = \begin{cases} bc(1 + \frac{1}{\eta-2}(\frac{bx+a}{1-\lambda})^2)^{-(1+\eta)/2} & \text{if } x < -a/b, \\ bc(1 + \frac{1}{\eta-2}(\frac{bx+a}{1+\lambda})^2)^{-(1+\eta)/2} & \text{if } x \geq -a/b, \end{cases}$$

where $2 < \eta < \infty$, and $-1 < \lambda < 1$. The constants a, b and c are given by:

$$a = 4\lambda c(\frac{\eta - 2}{\eta - 1}), \;\; b^2 = 1 + 3\lambda^2 - a^2, \;\; \text{and} \;\; c = \frac{\Gamma(\frac{\eta+1}{2})}{\sqrt{\pi(\eta-2)}\Gamma(\eta/2)}.$$

Assume that the density function of $\ln(R_t)$ is $f(x)$, then the density function of the discrete return $y = R_t$ can be expressed as:

$$g(y) = f(\ln y)\frac{1}{y} = \begin{cases} \frac{b}{y}c(1 + \frac{1}{\eta-2}(\frac{b\ln y+a}{1-\lambda})^2)^{-(1+\eta)/2} & \text{if } y < e^{-a/b}, \\ \frac{b}{y}c(1 + \frac{1}{\eta-2}(\frac{b\ln y+a}{1+\lambda})^2)^{-(1+\eta)/2} & \text{if } y \geq e^{-a/b}. \end{cases}$$

Table 1: Option prices under different risk neutral distributions

(a) Panel A: $\sigma = 0.2$

	Moneyness(K/S_0)	lognormal	Student t	Skewt1	Skewt2
Call	1.15	0.020	0.078	0.020	0.000
	1.125	0.057	0.125	0.038	0.001
	1.1	0.148	0.210	0.080	0.003
	1.075	0.349	0.373	0.188	0.022
	1.05	0.744	0.691	0.474	0.237
	1.025	1.435	1.292	1.146	1.002
	1	2.512	2.333	2.336	2.310
Put	0.85	0.003	0.029	0.062	0.093
	0.875	0.015	0.054	0.105	0.149
	0.9	0.061	0.107	0.184	0.242
	0.925	0.193	0.222	0.329	0.402
	0.95	0.504	0.469	0.598	0.675
	0.975	1.106	0.979	1.086	1.137
	1	2.096	1.917	1.922	1.899

(b) Panel B: $\sigma = 0.4$

	Moneyness(K/S_0)	lognormal	Student t	Skewt1	Skewt2
Call	1.15	0.730	0.811	0.394	0.041
	1.125	1.049	1.063	0.598	0.131
	1.1	1.479	1.410	0.921	0.400
	1.075	2.046	1.882	1.417	0.947
	1.05	2.774	2.521	2.140	1.779
	1.025	3.688	3.368	3.123	2.886
	1	4.805	4.456	4.367	4.247
Put	0.85	0.353	0.396	0.580	0.702
	0.875	0.613	0.602	0.814	0.944
	0.9	1.003	0.913	1.138	1.266
	0.925	1.552	1.369	1.582	1.691
	0.95	2.289	2.017	2.180	2.248
	0.975	3.231	2.897	2.967	2.966
	1	4.390	4.036	3.976	3.879

Note: This table reports call and put option prices with different moneynesses under different risk neutral distributions. Panel A reports results for $\sigma = 0.2$ and Panel B reports for $\sigma = 0.4$. Risk neutral distributions of the continuously compounded stock returns follow the normal, the Student-t or two skewed Student-t distributions as specified in (7) and (8). The degree of freedom of the Student-t and the two skewed Student-t distributions is 5. For the two skewed Student-t distributions, the skewness parameters are -0.3 and -0.7. The risk-free rate is 5%, K is the strike price, S_0 is the initial stock price 100, and the standard deviation of the underlying risk neutral distributions σ is 0.2.

Table 2: Comparison across the three methods, one-month maturity

	Volatility σ	0.2	0.2	0.4	0.4
	Option No.	14	6	14	6
Normal	BSIV	0.200	0.200	0.400	0.400
	MFIV	0.200	0.200	0.400	0.400
	EBIV	0.200	0.202	0.402	0.413
Student-t	BSIV	0.211	0.192	0.385	0.373
	MFIV	0.198	0.195	0.387	0.374
		(0.209)	(0.628)	(0.896)	(0.979)
	EBIV	0.199	0.196	0.393	0.393
		(0.139)	(0.466)	(0.444)	(0.264)
Skewt1	BSIV	0.206	0.192	0.374	0.368
	MFIV	0.197	0.197	0.383	0.366
		(0.433)	(0.342)	(0.657)	(1.082)
	EBIV	0.198	0.196	0.391	0.391
		(0.322)	(0.558)	(0.361)	(0.269)
Skewt2	BSIV	0.195	0.187	0.350	0.359
	MFIV	0.196	0.196	0.375	0.355
		(0.818)	(0.334)	(0.501)	(1.106)
	EBIV	0.197	0.193	0.384	0.387
		(0.668)	(0.541)	(0.308)	(0.310)

Note: This table reports the estimated implied volatility calculated from 14 or 6 options with one-month maturity by the Black-Scholes formula (BSIV), the model-free method (MFIV) and the maximum entropy method (EBIV) under different risk neutral distributions. The first row provides the true volatilities of the underlying risk neutral distribution. The second row presents number of options used in calculating implied volatilities. The moneynesses of the 14 options range from 0.85 to 1.15 and the moneynesses of the 6 options range from 0.95 to 1.05, both with equal interval 0.025. The degree of freedom of the Student-t and the two skewed Student-t distributions is 5. For the two skewed Student-t distributions, the skewness parameters are -0.3 and -0.7, for "Skew1" and "Skew2" respectively. The estimation improvements of the MFIV and the EBIV compared to the BSIV are presented in parenthesis, which is defined as $\frac{|XXIV - TrueVolatility|}{|BSIV - TrueVolatility|}$, where XXIV is either EBIV or MFIV.

Table 3: Coverage rates of the confidence intervals around the EBIV

volatility		Normal	Student t	Skewt1	Skewt2
0.2	95%	91.21%	91.40%	92.30%	93.10%
	90%	85.00%	85.50%	86.60%	92.30%
0.4	95%	92.39%	91.60%	93.40%	92.50%
	90%	88.78%	86.70%	87.60%	84.50%

Note: This table reports the coverage rates of confidence intervals under different risk neutral distributions. The upper panel is for $\sigma = 0.2$ under 95% and 90% confidence levels and the lower panel is for $\sigma = 0.4$. The degree of freedom of the Student-t and the two skewed Student-t distributions is 5. For the two skewed Student-t distributions, the skewness parameters are -0.3 and -0.7, for "Skew1" and "Skew2" respectively. The details of constructing the confidence interval of the EBIV is in Section 2.2.

Table 4: Comparison across the three methods: 3-month maturity

	Volatility	0.2	0.2	0.2	0.4	0.4	0.4
	Option No.	14	6	3+7	14	6	3+7
Normal	BSIV	0.200	0.200	0.200	0.400	0.400	0.400
	MFIV	0.200	0.200	0.200	0.400	0.400	0.400
	EBIV	0.201	0.202	0.200	0.404	0.417	0.405
student t	BSIV	0.195	0.192	0.200	0.384	0.376	0.379
	MFIV	0.194	0.193	0.204	0.386	0.376	0.382
		(1.086)	(0.908)	(8.784)	(0.862)	(0.995)	(0.854)
	EBIV	0.197	0.196	0.197	0.395	0.397	0.393
		(0.515)	(0.453)	(6.329)	(0.328)	(0.131)	(0.343)
skewt1	BSIV	0.192	0.188	0.199	0.368	0.361	0.370
	MFIV	0.194	0.191	0.208	0.379	0.363	0.378
		(0.746)	(0.762)	(10.454)	(0.664)	(0.941)	(0.714)
	EBIV	0.197	0.196	0.197	0.391	0.393	0.391
		(0.343)	(0.359)	(3.486)	(0.274)	(0.185)	(0.302)
skewt2	BSIV	0.184	0.177	0.190	0.341	0.341	0.354
	MFIV	0.191	0.186	0.208	0.367	0.346	0.369
		(0.534)	(0.581)	(0.763)	(0.563)	(0.902)	(0.659)
	EBIV	0.195	0.193	0.195	0.384	0.384	0.385
		(0.290)	(0.310)	(0.471)	(0.281)	(0.262)	(0.329)

Note: This table reports the estimated implied volatility calculated from 14 or 6 options with three-month maturity by the Black-Scholes formula (BSIV), the model-free method (MFIV) and the maximum entropy method (EBIV) under different risk neutral distributions. The first row provides the true volatilities of the underlying risk neutral distribution. The second row presents number of options used in calculating implied volatilities. The moneynesses of the 14 options range from 0.85 to 1.15 and the moneynesses of the 6 options range from 0.95 to 1.05, both with equal interval 0.025. The degree of freedom of the Student-t and the two skewed Student-t distributions is 5. For the two skewed Student-t distributions, the skewness parameters are -0.3 and -0.7, for "Skew1" and "Skew2" respectively. The estimation improvements of the MFIV and the EBIV compared to the BSIV are presented in parenthesis, which is defined as $\frac{|XXIV - TrueVolatility|}{|BSIV - TrueVolatility|}$, where XXIV is either EBIV or MFIV.

Table 5: Implied volatilities estimated using the three methods: 1-year maturity

	Volatility	0.2	0.2	0.2	0.4	0.4	0.4
	Option No.	14	6	3+7	14	6	3+7
Normal	BSIV	0.200	0.200	0.200	0.400	0.400	0.400
	MFIV	0.200	0.200	0.200	0.400	0.400	0.400
	ETIV	0.204	0.202	0.200	0.407	0.423	0.407
student t	BSIV	0.188	0.192	0.198	0.392	0.386	0.388
	MFIV	0.190	0.192	0.199	0.390	0.381	0.387
		(0.850)	(0.933)	(0.598)	(1.363)	(1.328)	(1.037)
	EBIV	0.197	0.197	0.197	0.401	0.406	0.399
		(0.256)	(0.402)	(1.181)	(0.101)	(0.452)	(0.072)
skewt1	BSIV	0.191	0.193	0.201	0.366	0.363	0.368
	MFIV	0.193	0.193	0.203	0.382	0.368	0.382
		(0.844)	(0.977)	(3.253)	(0.536)	(0.867)	(0.572)
	EBIV	0.200	0.199	0.201	0.399	0.404	0.398
		(0.040)	(0.086)	(0.651)	(0.036)	(0.112)	(0.049)
skewt2	BSIV	0.188	0.183	0.193	0.331	0.334	0.342
	MFIV	0.190	0.187	0.200	0.366	0.346	0.369
		(0.840)	(0.810)	(0.054)	(0.486)	(0.812)	(0.546)
	EBIV	0.193	0.196	0.198	0.389	0.393	0.390
		(0.542)	(0.253)	(0.316)	(0.163)	(0.102)	(0.175)

Note: This table reports the estimated implied volatility calculated from 14 or 6 options with one-year maturity by the Black-Scholes formula (BSIV), the model-free method (MFIV) and the maximum entropy method (EBIV) under different risk neutral distributions. The first row provides the true volatilities of the underlying risk neutral distribution. The second row presents number of options used in calculating implied volatilities. The moneynesses of the 14 options range from 0.85 to 1.15 and the moneynesses of the 6 options range from 0.95 to 1.05, both with equal interval 0.025. The degree of freedom of the Student-t and the two skewed Student-t distributions is 5. For the two skewed Student-t distributions, the skewness parameters are -0.3 and -0.7, for "Skew1" and "Skew2" respectively. The estimation improvements of the MFIV and the EBIV compared to the BSIV are presented in parenthesis, which is defined as $\frac{|XXIV-TrueVolatility|}{|BSIV-TrueVolatility|}$, where XXIV is either EBIV or MFIV.

Table 6: Implied moments for the stochastic volatility and jump (SVJ) model)

Parameter set I				Parameter set II			
	volatility	skewness	kurtosis		volatility	skewness	kurtosis
true value	0.203	-0.401	3.179	true moments	0.453	-0.554	3.53
BS	0.2	-	-	BS	0.445	-	-
MF	0.205	-0.274	2.924	MF	0.427	-0.059	1.76
ET	0.203	-0.404	3.207	ET	0.454	-0.543	3.435

Note: This table provides a comparison between the three methods, the Black-Scholes formula (BS), the model-free (MF) method and the entropy-based (EB) method, when the option prices are simulated from the stochastic volatility and jump (SVJ) model (details in Section 3.4). For the two sets of parameters given in Section 3.4, we present the true moments of the underlying distribution in the first row and the implied moments calculated from different methods in the second to fourth rows. To determine the true moments of the risk neutral distributions with the two parameter sets, $100,000$ monthly returns are simulated by aggregating each 21 daily returns simulated from the SVJ model. Option prices and the true moments are calculated based on the simulated monthly return. There are 14 options in total, with one-month maturity. The moneynesses of the options range from 0.85 to 1.15 with equal interval 0.025.

Table 7: Descriptive statistics of the S&P500 index options with 1 month expiration

(a) Panel A: Call options

K_c/S	1	1.025	1.05	1.075	1.1	1.125	1.15
Mean	24.62	11.38	4.78	2.23	1.28	0.93	0.77
Variance	80.02	65.66	33.11	15.39	7.36	4.35	2.28
Skewness	1.11	1.66	3.15	5.26	7.17	8.66	9.73
Maximum	7.69	1.28	0.40	0.40	0.40	0.40	0.40
Minimum	64.85	53.25	45.20	36.80	29.30	24.50	19.30
obs.	222	222	222	196	123	57	33

(b) Panel B: Put options

K_p/S	0.85	0.875	0.9	0.925	0.95	0.975	1
Mean	1.96	2.69	3.91	5.78	8.92	14.38	24.10
Variance	9.06	13.37	20.66	30.97	47.59	66.61	80.86
Skewness	4.57	4.12	3.41	2.83	2.25	1.66	0.99
Maximum	0.40	0.40	0.40	0.53	1.33	3.65	8.25
Minimum	25.20	30.05	34.45	40.35	47.35	55.25	62.65
obs.	222	190	216	220	221	221	222

Note: This table reports descriptive statistics of the S&P500 call and put options with 1 month expiration from January 1996 to August 2014. The options are selected by the procedure illustrated in Section 4.1. The first row in Panel 7a (Panel 7b) shows moneynesses of the out-of-the-money call (put) options, where K_c (K_p) are exercise prices of the call (put) options, S is the current price of the S&P500 index. The last row labelled "obs" shows the number of observations in each moneyness category.

Table 8: Descriptive statistics of different measures of volatility

	Mean	Median	Std. Dev.	Skewness	Kurtosis	Maximum	Minimum
RV	0.172	0.145	0.101	2.739	14.158	0.784	0.055
VIX	0.217	0.200	0.093	2.460	13.122	0.809	0.102
BSIV	0.219	0.204	0.078	2.994	17.341	0.765	0.130
MFIV	0.209	0.195	0.085	2.460	13.349	0.758	0.103
EBIV	0.210	0.194	0.089	2.257	11.343	0.745	0.097

Note: This table reports the descriptive statistics for volatility measures RV, VIX, BSIV, MFIV and EBIV. RV is the realized volatility of the preceding 30 days defined in 11. VIX is the volatility index provided by CBOE. BSIV is the average Black-Scholes implied volatility of all available option prices after filter procedure. MFIV is calculated based on Appendix 6.1. The details of calculating EBIV is in 2.1 and 4.1. Statistics are reported for the full sample from January 1996 to August 2014. In all tables and figures, the volatility measures are annualized and given in decimal form.

Table 9: Correlation matrix of different measures of volatilities

	RV	VIX	BSIV	MFIV	EBIV
RV	1.000	0.734	0.735	0.735	0.735
VIX	0.734	1.000	0.994	0.998	0.998
BSIV	0.735	0.994	1.000	0.996	0.993
MFIV	0.735	0.998	0.996	1.000	0.998
EBIV	0.735	0.998	0.993	0.998	1.000

Note: This table reports the correlation coefficients across volatilities measures: RV, VIX, BSIV, MFIV and EBIV. RV is the realized volatility of the preceding 30 days defined in 11. VIX is the volatility index provided by CBOE. BSIV is the average Black-Scholes implied volatility of all available option prices after filter procedure. MFIV is calculated based on Appendix 6.1. The details of calculating EBIV is in 2.1 and 4.1. The sample period is from January 1996 to August 2014.

Table 10: Predict the realized volatility: volatility regressions

	In-Sample Estimation					Out-of-sample RMSE			
	α	β_1	β_2	β_3	$adj.R^2$	All days	Low	Medium	High
RV	0.050***	0.680***	-	-	0.486	0.506	0.483	0.815	1.142
	(3.681)	(7.544)	-	-					
BSIV	-0.024	0.917***	-	-	0.554	0.485	0.295	0.819	1.125
	(-1.428)	(9.982)	-	-					
VIX	-0.003	0.805***	-	-	0.554	0.486	0.284	0.824	1.127
	(-0.221)	(9.929)	-	-					
MFIV	-0.006	0.823***	-	-	0.559	0.484	0.283	0.824	1.120
	(-0.419)	(9.879)	-	-					
EBIV	0.005	0.787***	-	-	0.564	0.477	0.286	0.823	1.096
	(0.397)	(10.401)	-	-					
RV+BSIV	-0.014	0.157	0.740***	-	0.562	0.486	0.316	0.812	1.126
	(-1.234)	(1.073)	(5.102)	-					
RV+MFIV	0.000	0.151	0.672***	-	0.566	0.484	0.306	0.816	1.118
	(-0.047)	(1.123)	(5.687)	-					
RV+EBIV	0.008	0.135	0.658***	-	0.571	0.478	0.305	0.815	1.101
	(0.945)	(1.109)	(5.956)	-					
BSIV+MFIV	0.004	-0.527	1.294	-	0.562	0.491	0.283	0.843	1.132
	(0.216)	(-0.422)	(1.110)	-					
BSIV+EBIV	0.032	-0.790	1.457*	-	0.570	0.484	0.292	0.822	1.118
	(0.975)	(-0.787)	(1.660)	-					
MFIV+EBIV	0.010	-0.299	1.070*		0.567	0.496	0.305	0.818	1.165
	(0.806)	(-0.463)	(1.700)						

Note: This table reports regression results of forecasting future realized volatility using different measures of volatility. The sample period extends from January 1996 to August 2014. All regressions are based on monthly nonoverlapping observations. The dependent variable is the realized volatility in the next month defined in equation (11). Robust t-statistics are reported in parentheses taking into account the heteroscedastic and autocorrelated error structure. The out-of-sample analysis is conducted based on a moving window of 100 observations preceding to the period to be forecasted. Besides the results for "All Days" in the forecasting period, the whole sample is further splitted into three sub-samples by sorting the BSIV in the month preceding to the forecasting period in ascending order. "Low", "Medium" and "High" represent different volatility regime.

Table 11: Predict the realized volatility: log volatility regressions

	In-Sample Estimation					Out-of-sample RMSE			
	α	β_1	β_2	β_3	$adj.R^2$	All days	Low	Medium	High
RV	-1.216***	1.394***	-	-	0.447	0.200	0.181	0.171	0.146
	(-5.156)	(11.760)	-	-					
BSIV	-0.119	2.284***	-	-	0.576	0.175	0.140	0.161	0.136
	(-0.554)	(18.128)	-	-					
VIX	-0.542**	2.025***	-	-	0.578	0.175	0.136	0.163	0.137
	(-2.544)	(16.639)	-	-					
MFIV	-0.496**	2.049***	-	-	0.582	0.175	0.136	0.162	0.137
	(-2.328)	(16.776)	-	-					
EBIV	-0.565***	1.970***	-	-	0.582	0.174	0.135	0.162	0.135
	(-2.801)	(17.397)	-	-					
RV+BSIV	-0.128	0.056	2.213***	-	0.579	0.176	0.142	0.161	0.136
	(-0.617)	(0.288)	(8.471)	-					
RV+MFIV	-0.494**	0.038	2.007***	-	0.585	0.175	0.139	0.162	0.137
	(-2.316)	(0.196)	(8.566)	-					
RV+EBIV	-0.561***	0.049	1.917***	-	0.585	0.175	0.138	0.162	0.135
	(-2.760)	(0.262)	(8.678)	-					
BSIV+MFIV	-0.696**	-1.134	3.059*	-	0.586	0.175	0.138	0.163	0.135
	(-2.264)	(-0.611)	(1.787)	-					
BSIV+EBIV	-0.699**	-0.645	2.522*	-	0.585	0.174	0.135	0.162	0.136
	(-2.235)	(-0.383)	(1.688)	-					
MFIV+EBIV	-0.521**	1.255	0.764		0.585	0.175	0.139	0.161	0.137
	(-2.399)	(0.451)	(0.284)						

Note: This table reports regression results of forecasting future realized log volatility using different measures of volatility. The sample period extends from January 1996 to August 2014. All regressions are based on monthly nonoverlapping observations. The dependent variable is the realized log volatility in the next month defined in equation (11). Robust t-statistics are reported in parentheses taking into account the heteroscedastic and autocorrelated error structure. The out-of-sample analysis is conducted based on a moving window of 100 observations preceding to the period to be forecasted. Besides the results for "All Days" in the forecasting period, the whole sample is further splitted into three sub-samples by sorting the BSIV in the month preceding to the forecasting period in ascending order. "Low", "Medium" and "High" represent different volatility regime.

Table 12: Predict the monthly returns using variance risk premium

	In-Sample Estimation			Out-of-sample RMSE			
	α	β_1	$adj.R^2$	All days	Low	Medium	High
VRP_{BS}	0.002	0.353***	0.051	0.982	1.010	1.016	0.961
	(0.427)	(3.916)					
VRP_{MF}	0.000	0.411***	0.065	0.971	1.045	1.010	0.939
	-(0.040)	(4.430)					
VRP_{EB}	0.000	0.425***	0.071	0.971	1.045	1.009	0.941
	(0.046)	(4.273)					

Note: This table presents predictive regression results of variance risk premium on future monthly return. VRP_{BS} is the variance risk premium calculated by the difference between $BSIV^2$ and realized variance in the last month RV^2, which is defined in equation (11). VRP_{MF} and VRP_{EB} are variance risk premium calculated based on MFIV and EBIV. The sample period extends from January 1996 to August 2014. All regressions are based on monthly nonoverlapping observations. The dependent variable is the S&P500 index return in the next month. Robust t-statistics are reported in parentheses taking into account the heteroscedastic and autocorrelated error structure. The out-of-sample analysis is conducted based on a moving window of 100 observations preceding to the period to be forecasted. Besides the results for "All Days" in the forecasting period, the whole sample is further splitted into three sub-samples by sorting the BSIV in the month preceding to the forecasting period in ascending order. "Low", "Medium" and "High" represent different volatility regime.

Figure 1: Comparing the estimated risk neutral distribution to the true distribution

(a) Standard normal

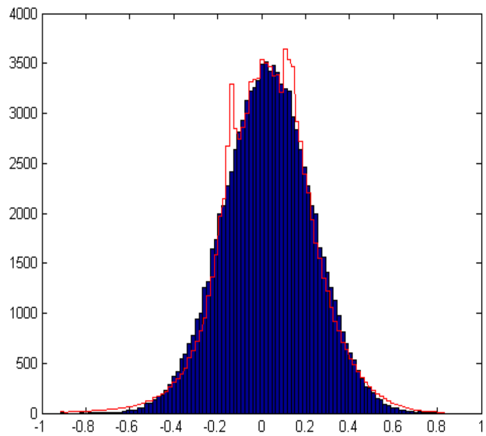

(b) Student t

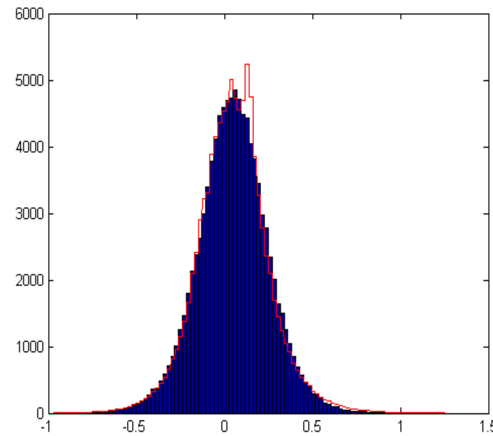

(c) Skewt1

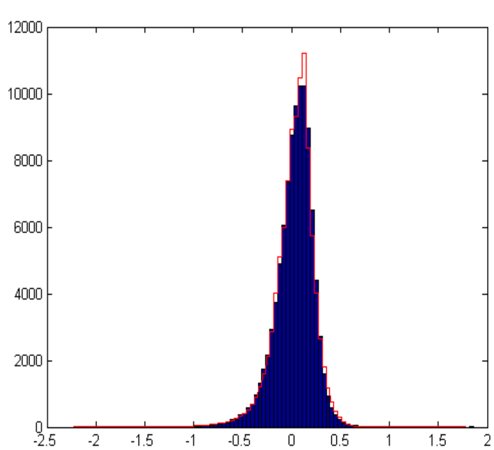

(d) Skewt2

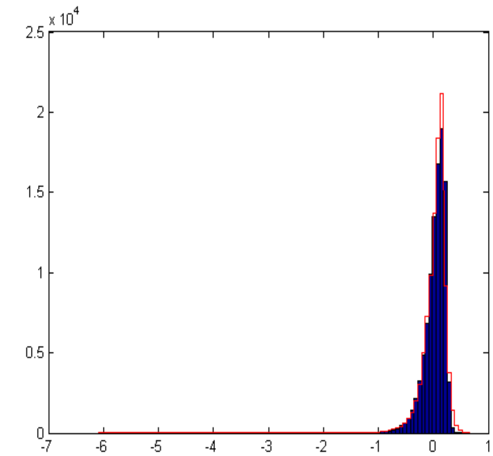

Note: The blue bars are histograms of the given risk neutral distributions of the continuous compounded return. The red lines are the estimated risk neutral densities using the Maximum Entropy method. The distributions are estimated from 14 options. Their moneynesses range from 0.85 to 1.15 with an equal interval 0.025. The degree of freedom of the Student t and the two skewed Student-t distributions is 5. For the two skewed Student-t distributions, the skewness parameters are -0.3 and -0.7, for "Skew1" and "Skew2" respectively.

Figure 2: Comparison between the MFIV and the EBIV

(a) EBIV and MFIV

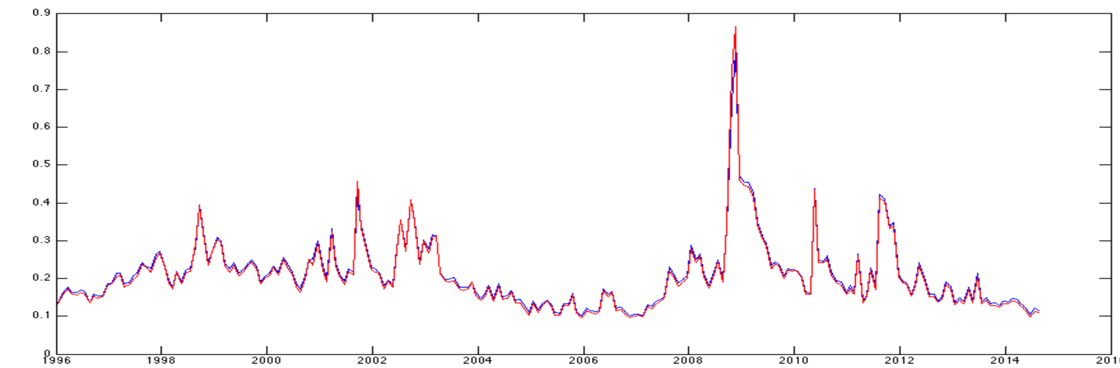

(b) Difference between EBIV and MFIV

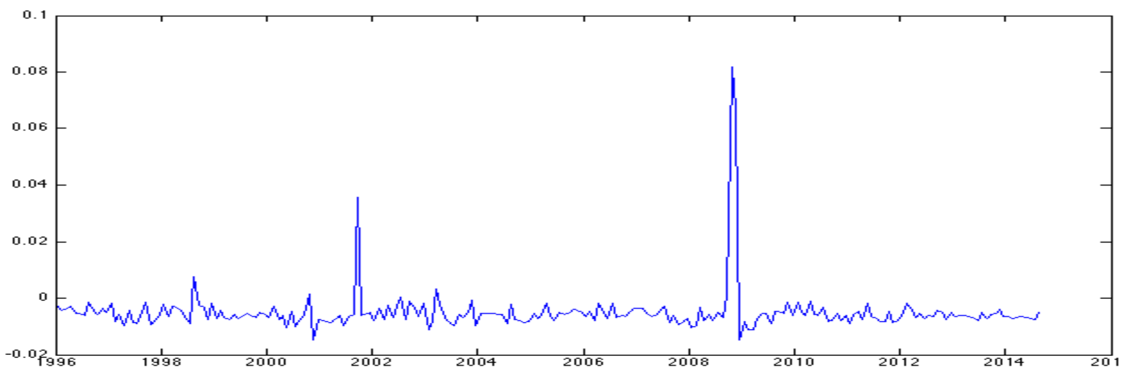

(c) BSIV and MFIV-EBIV

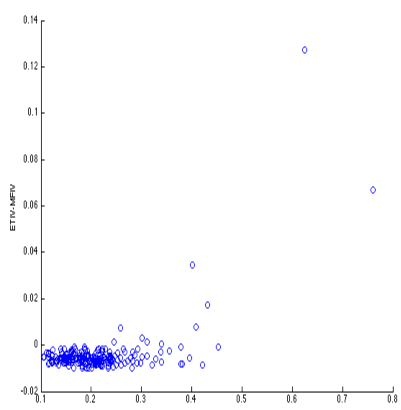

Note: This figure shows the comparison between the MFIV and the EBIV from January 1996 to August 2014. The blue line in 1a represents the MFIV and the red line represents the EBIV. Figure 1b shows the time series of the spread between EBIV and MFIV. Figure 1c shows the scatter plot of the BSIV (horizontal axis) against the spread between MFIV and EBIV (vertical axis).

Figure 3: Confidence interval using the maximum entropy method

(a) Confidence interval for EBIV

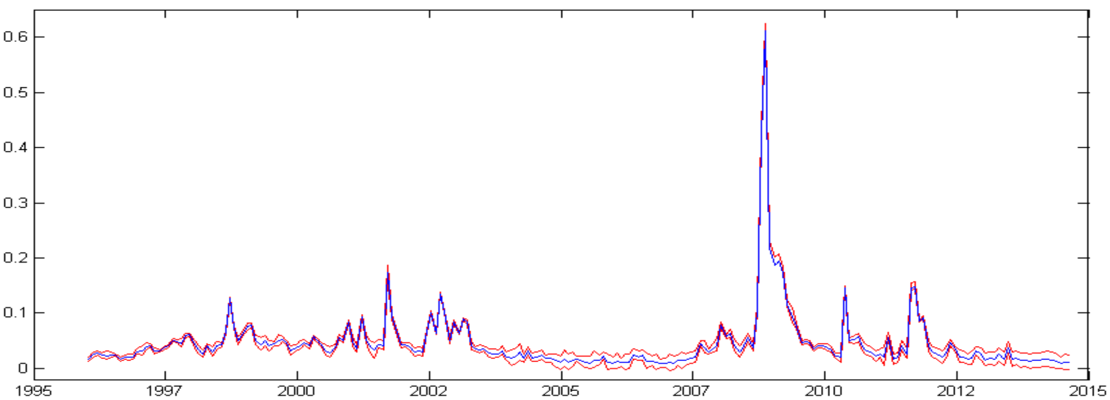

(b) Length of the confidence interval

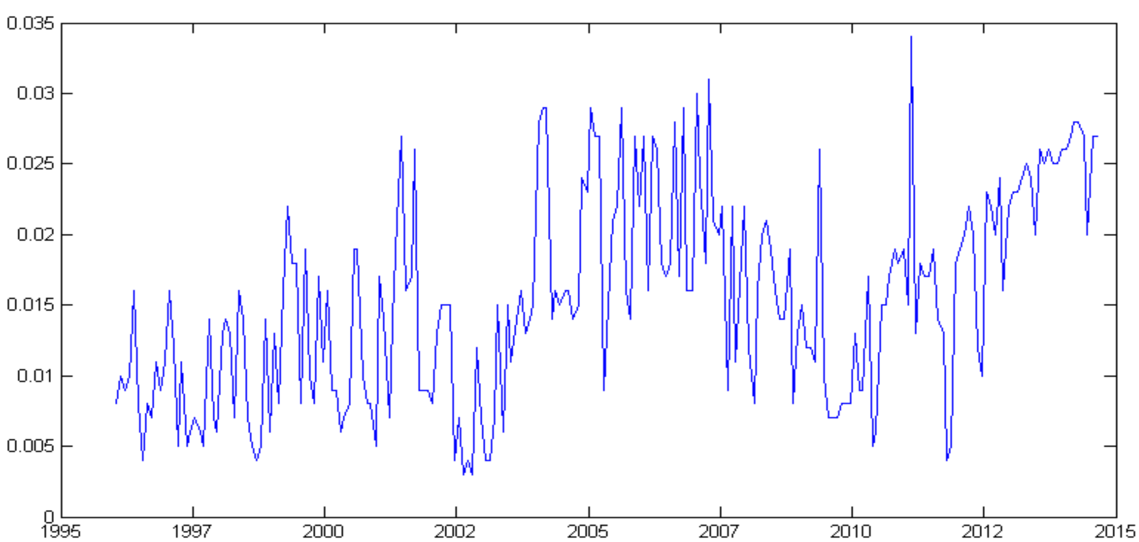

Note: The red lines in Figure 2a are the upper and lower bounds of the confidence interval around the estimated EBIV with a confidence level 95%. The blue line in Figure 2a represents the point estimate of the EBIV. Figure 2b shows the time series of the lengths of the confidence intervals.

Figure 4: Estimated risk neutral distributions on four selected dates

(a) June 18th, 1998

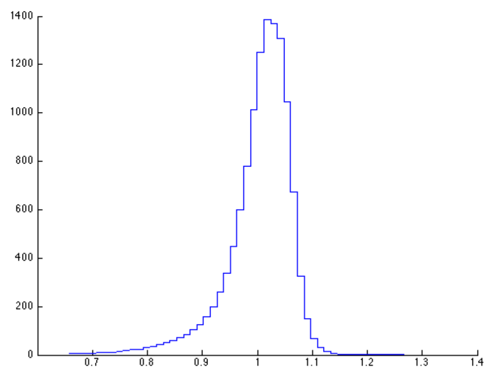

(b) September 20th, 2001

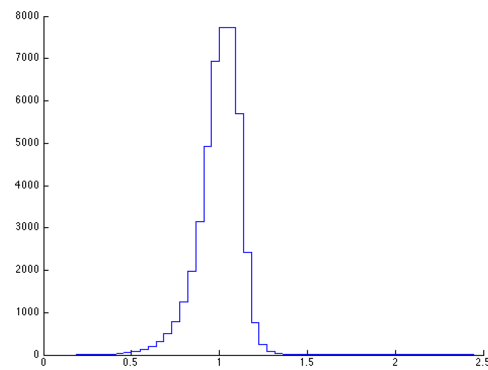

(c) October, 23rd, 2008

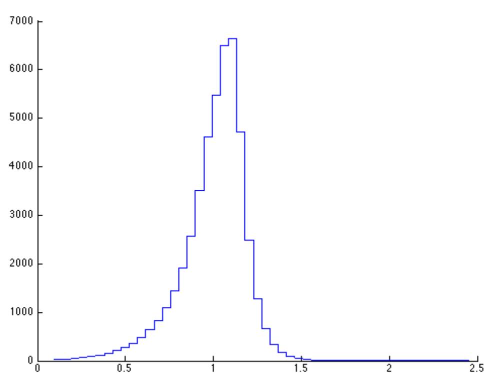

(d) June 19th, 2014

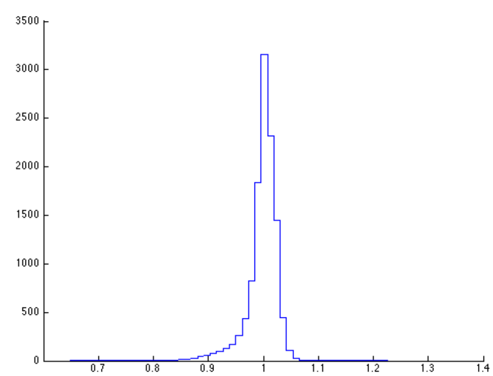

Note: This figure shows the estimated risk neutral distributions using the maximum entropy method on four selected dates.

References

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

Bakshi, G., Cao, C., and Chen, Z. Empirical performance of alternative option pricing models. *Journal of Finance*, 52(5):2003–2049, 1997.

Bakshi, G. S., Kapadia, N., and Madan, D. Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies*, 16:101–143, 2003.

Bali, T. G. and Murray, S. Does Risk-Neutral Skewness Predict the Cross-Section of Equity Option Portfolio Returns? *Journal of Financial and Quantitative Analysis*, pages 1145–1171, 2013.

Bates, D. S. The crash of'87: Was it expected? the evidence from options markets. *Journal of Finance*, 46(3):1009–1044, 1991.

Black, F. and Scholes, M. The pricing of options and corporate liabilities. *Journal of Political Economy*, pages 637–654, 1973.

Blair, B. J., Poon, S. H., and Taylor, S. J. Modelling S&P 100 volatility: The information content of stock returns. *Journal of Banking and Finance*, 25:1665–1679, 2001.

Bollerslev, T., Tauchen, G., and Zhou, H. Expected Stock Returns and Variance Risk Premia. *Review of Financial Studies*, 22:4463–4492, 2009.

Britten-Jones, M. and Neuberger, A. Option Prices, Implied Price Processes, and Stochastic Volatility. *Journal of Finance*, 55:839–866, 2000.

Buchen, P. W., Url, S., Kelly, M., Journal, T., Analysis, Q., Archive, T. J., and Archive, T. The Maximum Entropy Distribution of an Asset Inferred from Option Prices. *Journal of Financial and Quantitative Analysis*, 31:143–159, 1996.

Busch, T., Christensen, B. J., and Nielsen, M. Ø. The role of implied volatility in forecasting future realized volatility and jumps in foreign exchange, stock, and bond markets. *Journal of Econometrics*, 160:48–57, 2011.

Canina, L. and Figlewski, S. The informational content of implied volatility. *Review of Financial Studies*, 6:659–681, 1993.

Christensen, B. J. and Prabhala, N. R. The relation between implied and realized volatility. *Journal of Financial Economics*, 50:125–150, 1998.

Christensen, B. J., Hansen, C. S., and Prabhala, N. R. The telescoping overlap problem in options data. *Working Paper*, 2001.

Day, T. E. and Lewis, C. M. Stock market volatility and the information content of stock index options. *Journal of Econometrics*, 52:267–287, 1992.

DeMiguel, V., Plyakha, Y., Uppal, R., and Vilkov, G. Improving Portfolio Selection Using Option-Implied Volatility and Skewness. *Journal of Financial and Quantitative Analysis*, pages 1–57, 2014.

Dennis, P. and Mayhew, S. Risk-Neutral Skewness: Evidence from Stock Options. *Journal of Financial and Quantitative Analysis*, 37:471–493, 2002.

Fleming, J. The quality of market volatility forecasts implied by S&P 100 index option prices. *Review of Financial Studies*, 5:317–345, 1998.

Hansen, L. P. Large Sample Properties of Generalized Method of Moments Estimators. *Econometrica*, 50:1029–1054, 1982.

Jackwerth, J. C. Recovering Risk Aversion from Option Prices and Realized Returns. *Review of Financial Studies*, 13:433–451, 2000.

Jiang, G. J. and Tian, Y. S. The model-free implied volatility and its information content. *Review of Financial Studies*, 18:1305–1342, 2005.

Kitamura, Y. and Stutzer, M. An Information-Theoretic Alternative to Generalized Method of Moments Estimation. *Econometrica*, 65:861–874, 1997.

Lamoureux, C. G. and Lastrapes, W. D. Forecasting stock-return variance: toward an understanding of stochastic implied volatilities. *Review of Financial Studies*, 6:293–326, 1993.

Neumann, M. and Skiadopoulos, G. Predictable dynamics in higher-order risk-neutral moments: Evidence from the s&p 500 options. *Journal of Financial and Quantitative Analysis*, 48(03):947–977, 2013.

Schennach, S. M. Point estimation with exponentially tilted empirical likelihood. *Annals of Statistics*, 35:634–672, 2007.

Stutzer, M. A Simple Nonparametric Approach to Derivative Security Valuation. *Journal of Finance*, 51:1633 – 1652, 1996.