Available online at www.sciencedirect.com
ScienceDirect
Physica A 382 (2007) 121–128
PHYSICA A
www.elsevier.com/locate/physa

Non-parametric extraction of implied asset price distributions

Jerome V. Healy, Maurice Dixon*,[1], Brian J. Read, Fang Fang Cai

CCTM, London Metropolitan University, 31 Jewry Street, London EC3N 2EY, UK

Available online 1 March 2007

Abstract

We present a fully non-parametric method for extracting risk neutral densities (RNDs) from observed option prices. The aim is to obtain a continuous, smooth, monotonic, and convex pricing function that is twice differentiable. Thus, irregularities such as negative probabilities that afflict many existing RND estimation techniques are reduced. Our method employs neural networks to obtain a smoothed pricing function, and a central finite difference approximation to the second derivative to extract the required gradients.

This novel technique was successfully applied to a large set of FTSE 100 daily European exercise (ESX) put options data and as an Ansatz to the corresponding set of American exercise (SEI) put options. The results of paired t-tests showed significant differences between RNDs extracted from ESX and SEI option data, reflecting the distorting impact of early exercise possibility for the latter. In particular, the results for skewness and kurtosis suggested different shapes for the RNDs implied by the two types of put options. However, both ESX and SEI data gave an unbiased estimate of the realised FTSE 100 closing prices on the options' expiration date. We confirmed that estimates of volatility from the RNDs of both types of option were biased estimates of the realised volatility at expiration, but less so than the LIFFE tabulated at-the-money implied volatility.

Keywords: Option pricing; Risk neutral density; Risk management; Neural nets; Econophysics

1. Introduction

Extracting the risk neutral density (RND) function from option prices is well defined in principle, but is very sensitive to errors in practice. For risk management, knowledge of the entire RND provides more information for value-at-risk (VaR) calculations than implied volatility alone [1]. Many asset pricing models used in finance, including the Black–Scholes (BS) model for option prices, rely on the conventional assumption that the statistical distribution of asset returns is normal, and the price distribution is log-normal [2]. This assumption is consistent with geometric Brownian motion as the underlying mechanism driving price movements. It is now well known that historical asset price distributions exhibit fat tails. That is, they are slightly smaller near the mean and larger at extreme values. This has important implications for financial risk management, as it suggests that large price movements occur more frequently than they would for a normal

*Corresponding author.
E-mail address: M.Dixon@Londonmet.ac.uk (M. Dixon).
[1]Visitor: e-Science, RAL, Didcot, Oxon, OX11 0QX, UK.

0378-4371/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.physa.2007.02.013

distribution with the same variance. It also suggests that the underlying price process does not follow a geometric Brownian motion. Option prices represent a rich source of information on the statistical properties of the underlying asset. Exchange traded options are now available on financial assets including stock indices and futures; often these are heavily traded, so are very liquid. While daily time series of asset prices contain just one observation per day, there is a set of option prices available for each maturity date. These option prices reflect traders' expectations regarding future price movements of the underlying asset, so they allow alternative approaches to estimating financial risk.

In this paper we present a simple yet effective method for extracting non-parametric estimates of the complete distribution for the value of the underlying asset at maturity of an option, known as the RND, from sets of daily option prices. We have recently reported the application of data mining techniques using neural nets (NNs) to model European style FTSE 100 index options [3]. We extended that approach by introducing prediction intervals [4]. Here, we demonstrate how NNs can be used to obtain estimates of the expected value and standard deviation (volatility), as well as higher moments, of the implied distribution of the asset price at the expiration of an option written on it.

2. Risk neutral distributions

RNDs have many practical applications. They are used by central banks to assess market expectations regarding future stock prices, commodity prices, interest rates, and exchange rates in connection with setting monetary policy [2]. They are useful to market practitioners as an aid to investment decisions. RNDs extracted from exchange traded options can be used to price exotic options. For risk management, they can provide measures of VaR [1].

The prices of European exercise options can be expressed as the expected value of their payoffs, discounted at the risk-free interest rate.

$$C(X,t,T) = \mathrm{e}^{-r(T-t)} \int_X^{\infty} \rho(S_T)(S_T - X)\,\mathrm{d}S_T, \tag{1a}$$

$$P(X,t,T) = \mathrm{e}^{-r(T-t)} \int_{-\infty}^{X} \rho(S_T)(X - S_T)\,\mathrm{d}S_T. \tag{1b}$$

In Eq. (1) $C(X,t,T)$ and $P(X,t,T)$ are the prices of calls and puts trading at time t for expiration at some later time T. X is the strike price, and r is the risk-free interest rate. $\rho(S_T)$ is the distribution for the value of the underlying asset S at T predicted from time t. Given an assumption about the functional form of $\rho(S_T)$, options can be priced for any value of exercise price X. Conversely, given a series of synchronous market prices observed at some time t, for options expiring at some later time T, this calculation can be inverted and an estimate of $\rho(S_T)$ extracted. Breeden and Litzenberger [5] showed that the cumulative density function (negatively signed) for the value of the underlying asset S at time t is given by the first partial differential with respect to X of Eq. (1), while the RND is obtained by differentiating equation (1) twice with respect to X.

$$\frac{\partial f(X,t,T)}{\partial X} = -\mathrm{e}^{-r(T-t)} \int_X^{\infty} \rho(S_T)\,\mathrm{d}S_T \quad \text{and} \quad \frac{\partial^2 f(X,t,T)}{\partial X^2} = \mathrm{e}^{-r(T-t)}\rho(S_T). \tag{2}$$

In Eq. (2), $f(X,t,T)$ represents the call or put European option pricing functions. In reality, X is not continuous and options are only available for a limited number of exercise prices at discrete intervals. It has been shown [5], that for discrete data, finite difference methods can be used to obtain a numerical solution to Eq. (2). Neuhaus [6] has shown how the RND can be obtained via Eq. (2) using finite differences on the first derivative.

2.1. Recovering RNDs: Existing methods

Bondarenko has noted that there are many alternative methods in the literature for estimating RNDs from option prices [7]. Among parametric methods he considers those for the stochastic process, those for the implied volatility, and those directly modelling the RND. He identifies the advantage of the parametric models as being parsimonious in the number of parameters. He concludes however that "unfortunately, despite theoretical advances, the popular parametric models seem unable to fully explain the data". Bondarenko contrasts these with the non-parametric methods; these are data driven while making no strong assumptions about the data generating process. Among these he considers the methods of implied trees, smoothing techniques, maximum entropy, and neural networks. He concludes their disadvantage is that they are not effective on small samples. Two extensive tabulations with critical discussion of methods for extracting RNDs have been given in Jackwerth's monograph [8]. One tabulation was for parametric methods for which he used three classes: expansion methods, generalised distribution methods, and mixture methods. The other tabulation was for non-parametric methods in which he classified maximum entropy methods, kernel methods, and curve fitting methods. Jackwerth noted problems with all existing methods. He observed that non-parametric models oscillate wildly because of over fitting when small changes in market prices arising from noise in the data are fitted by the models. Jackwerth argues that where there are a reasonable number of observed options the choice between non-parametric methods does not really matter because it does not change the results much. Jackwerth also has concerns about the statistics of RNDs, observing that bootstrapping led to much lower confidence. Having found existing methods unsatisfactory both Bondarenko and Jackwerth suggest new methods called, respectively, positive convolution approximation (PCA) and fast and stable method (FSM). Bondarenko used Monte Carlo simulation to compare favourably the PCA with seven other leading parametric and non-parametric candidates. Both Bondarenko and Jackworth constrain their solutions such that the RND is positive across the range of strikes, while normalisation of the integral to one is by construction for Bondarenko and by scaling for Jackwerth.

Our own survey found that two techniques were popular with practitioners. They are the mixture of lognormals method and the smoothed implied volatility smile. The first of these is an example of a parametric method which works with Eq. (1); the second is an example of a non-parametric method which works with Eq. (2).

The mixture of lognormals technique originated with Ritchey [9]. Ritchey used a mixture of just two lognormal densities to minimise the number of parameters to be estimated, and this has become the standard procedure. Given the parameters and the observed option prices, the implied RND can then be constructed. The method is prone to over fitting since option price series frequently have 20 or fewer observed prices corresponding to different exercise prices with the same t and T.

Rather than trying to fit a smooth function to the prices Shimko [10] found it better to interpolate the corresponding implied volatilities using the BS formula as a transformer. The RND can be obtained directly from Eq. (2) provided the option pricing function $f(X, t, T)$ is observable. Unfortunately, only a relatively small number of option prices corresponding to discrete exercise prices are observable for a given time t. The resulting RND cannot be guaranteed to be positive because this is not a constrained optimisation.

Bliss and Panigirtzoglou [11] found that the smoothed implied volatility smile estimation method produced more stable RNDs than the mixture of lognormals method. By contrast, the mixture of lognormals method was found to be sensitive to computational problems.

2.2. Recovering RNDs from index options with NNs

There now exist a substantial number of studies dealing with the application of NNs to option pricing; see for example Bennell and Sutcliffe [12]. A close reading of this literature reveals that, surprisingly, there have been few empirical studies of their use for extracting RNDs from market traded option prices. It is commonly mentioned as feasible since NNs are twice differentiable. However, empirical results are presented only exceptionally for example [13,14]. A version of the parametric mixture of lognormals method was implemented by Schittenkopf and Dorffner [13] using mixture density networks for the FTSE 100. Herrmann and Narr [14] differentiated NN option pricing functions fitted directly to prices of options on the German DAX index to obtain RNDs. They used average values for some input variables when training their models. The resulting RNDs were compared graphically with corresponding lognormal RNDs obtained using the BS formulae. No statistical tests were performed, and only goodness-of-fit and error measures were provided. Garcia and Gencay [15] report the conditions under which NNs with nodes of logistic functions can achieve any desired degree of accuracy. Very recently Andreou et al. [16] reported a substantial improvement on pricing fits by combining two techniques. One is to create a hybrid by differencing with the parametric model. The other is to use a Huber function to implement a form of robust regression that minimises the 'leverage' effect of outliers. In assessing Andreou's results, allowance must also be made for the fact that confidence intervals for robust estimates are wider than least squares, even asymptotically.

Despite an extensive literature search, for extractions from NNs no studies were found other than [13,14]. The non-parametric extraction of RNDs from option prices is an example of an ill-posed problem, in that small changes in option prices can lead to large changes in the estimated density function. (Market prices of ESX and SEI options are only quoted at a granularity of 0.5 index points.) It requires the use of methods that are robust to slight changes in the option prices. NNs have been shown to be suitable for directly fitting option prices, avoiding the need to work in implied volatility space as in the smoothed implied volatility method.

2.3. *American put options*

Most exchange traded options are American options and these can be exercised at any time prior to maturity. It is expected that this feature, which is reflected in their price, affects any extracted RND. The theory underlying RNDs is only applicable to ESX options, and cannot be applied to SEI options without modification. Dupont [17] discussed this and suggested that the early exercise correction is not significant in practice except for deep in-the-money options. It is an open question whether RNDs extracted from American options are significantly different empirically from those extracted from corresponding European options. The American style SEI option is based directly upon the FTSE 100 index while the European ESX option is based upon an implied future on the FTSE 100 index. The underlyings converge when time approaches maturity date. Only put options are considered here; in the absence of dividend payments early exercise of American call options is never optimal, so they can be priced as European. Here, RNDs from American Ansatz and European put option pricing functions are extracted and compared.

3. Data and method

For this work, smoothed prices corresponding to each exercise price in a daily price series were estimated through directly fitting a neural net to create an option pricing function. Then RNDs were extracted by twice partially differentiating the functions numerically with respect to exercise price. A training set of 13,790 FTSE 100 ESX put options was prepared along with another training set which comprised 14,619 FTSE 100 SEI put options. To ensure prices were liquid, only options with positive values for contract volume and open interest were selected. A disjoint test set of 60 daily option price series, containing data on a total of 1238 (European) put options on the FTSE 100 index, was created. Option and underlying asset prices for the American puts were added for the same time, risk-free interest rate, and interleaved exercise price. The resulting test set had the following two special features: (1) included options were traded for one of 60 consecutive monthly expirations; (2) the options had a maturity of one calendar month (17 or 18 trading days). This was because any longer maturity results in overlapping data for some variables.

In creating the test set, the objective was to obtain a set of option price series with constant maturities, which was non-overlapping. The latter feature was required to avoid serial dependence between successive observations, which might bias statistical results for the RNDs [18]. Pricing models for European and American put options were separately trained, using logistic functions in a 5-11-1 architecture. The inputs were the five BS variables, and the targets were market prices of European and American put options, respectively. Once trained, the pricing models were applied to the test set to generate series of smoothed European and American option prices, taking care to use the correct values of the underlying asset as inputs to each model. Each generated price series was then numerically differentiated to estimate $\partial^2 f(X,t,T)/\partial X^2$ using symmetric central finite differences. Making f correspond to the neural net, the following formula was applied:

$$\frac{\partial^2 f(X,S,t,r,\sigma)}{\partial X^2} = \frac{f(X+\varepsilon,S,t,r,\sigma) - 2f(X,S,t,r,\sigma) + f(X-\varepsilon,S,t,r,\sigma)}{\varepsilon^2}, \tag{3}$$

where $f(X, S, t, r, \sigma)$ is a neural net option pricing function with the five standard BS input variables defined in [3] using at-the-money volatility, and ε is a small increment. If $\hat{\rho}(X_i)$ is the RND for a strike price interval $[X_i, X_{i+1}]$, then

$$\hat{\rho}(X_i) \approx \frac{\partial^2 f(X, S, t, r, \sigma)}{\partial X^2} \varepsilon \, e^{-r(T-t)}, \tag{4}$$

where ε is the interval between adjacent values of X_i. Eq. (4) was used to obtain point estimates of the RND for the interval corresponding to each X_i. The median values were calculated. Also the tail probabilities were calculated by constructing the cumulative distribution function as suggested by Neuhaus [6], using Eq. (2).

$$Mean = \sum_{i=1}^{n} \frac{X_i + X_{i+1}}{2} \hat{\rho}(X_i), \tag{5}$$

$$Stdev = \sqrt{\sum_{i=1}^{n} \left(\frac{X_i + X_{i+1}}{2} - Mean \right)^2 \hat{\rho}(X_i)}, \tag{6}$$

$$Skewness = \sum_{i=1}^{n} \left(\frac{X_i + X_{i+1}}{2} - Mean \right)^3 \frac{\hat{\rho}(X_i)}{Stdev^3}, \tag{7}$$

$$Kurtosis = \sum_{i=1}^{n} \left(\frac{X_i + X_{i+1}}{2} - Mean \right)^4 \frac{\hat{\rho}(X_i)}{Stdev^4}. \tag{8}$$

The first four moments of each recovered RND were estimated from Eqs. (5) to (8). The annualised percentage implied volatility was derived from Eq. (6) as $100 * (\text{Stdev}/\text{Mean}) * v^{0.5}$ where $v = (\text{time-for-year}/\text{time-to-expiry})$.

4. Results: RNDs of European and American put options

The fit of the NN models to the market prices of the option in the test set gave an R^2 of 0.988 and 0.984 for the European and American put options, respectively. The calculated t-statistics were −0.09 and −1.32, respectively, compared with a critical value of 1.96 for a two-tailed test. These results confirmed that both NN models were unbiased estimators of market prices. Tables 1 and 2 compare summary statistics for RNDs recovered from the test data of non-overlapping, constant maturity (1 month as 17 or 18 trading days), sets of European and American exercise put options on the FTSE 100 Index. Table 1 gives results for direct comparisons of the time series of moments. The results of the paired t-tests indicate significant differences, for each statistic tested, between RNDs for European and American put options in spite of $R^2 = 0.9996$ for mean and median. In particular, the results for skewness and kurtosis suggested the shapes of the RNDs were different for each type of put option. The F-test is a test to determine whether two samples are from populations with different variances; here it indicates significant differences for annualised IV and for kurtosis. Further tests were carried out to assess the practical effects of these differences on the predictive properties of

Table 1
European vs American FTSE 100 Index Put Options: Comparison of RNDs

SummaryStat	R^2	F-stat[a]	t-stat (paired)[a]	$H0$:[a]
Mean	0.9996	1.01	9.09	Reject
Median	0.9996	1.00	5.07	Reject
Stdev	0.640	1.45	5.35	Reject
Ann.I.V.%	0.820	1.78	4.29	Reject
Skewness	0.932	1.09	19.01	Reject
Kurtosis	0.779	1.96	−13.04	Reject

[a] $F_{crit} = 1.54$. $t_{crit}(2\text{ tail}) = 2.00$. $H0$:= no difference in SummaryStat means at 95% confidence.

Table 2
Performance: Comparison of market value with RNDs' one month estimate

Option	Parameter	R^2	F–stat[a]	t-stat[a]	$H0$:[a]
ESX	T.FTSE100 vs Median	0.955	1.01	−0.02	Accept
ESX	Realised vol. vs Ann.I.V.%	0.379	2.01	−3.47	Reject
SEI	T.FTSE100 vs Median	0.955	1.01	0.05	Accept
SEI	Realised vol. vs Ann.I.V.%	0.318	3.58	−2.55	Reject
Market Outcome	Realised vol. v ATMIV (LIFFE)%	0.426	1.46	−4.50	Reject

[a] $F_{crit} = 1.54$. $t_{crit}(2\text{ tail}) = 2.00$. $H0$:= no difference in means at 95% confidence.

Table 3
Skewness and Kurtosis: Estimated RNDs vs Normal

Option	Skewness	Normal	t-stat (paired)[a]	Kurtosis	Normal	t-stat (paired)[a]
ESX	−0.43	0	−10.1	2.82	3	−4.4
SEI	−0.65	0	−14.7	3.19	3	3.4

[a] Reject $H0$:= no difference in means at 95% confidence if $abs(t - \text{test}) > t_{crit}(2\text{ tail}) = 2.00$.

RNDs from each type of put option. Results of these tests are presented in Table 2. The median and annualised percentage implied volatility from each RND are compared with the actual traded FTSE 100 closing price (T.FTSE 100) and realised volatility on the expiration date of the option. This is a test of the one month (17 or 18 trading days) predictive characteristics of the median and implied volatility from the estimated RNDs. The t-test results in Table 2 indicate that the medians of RNDs from both SEI and ESX options provide an unbiased estimate of FTSE 100 closing prices on the expiration date of the options, one month later. The annualised implied volatilities of the RNDs from both types of options, on the other hand, are biased estimates of the actual realised volatilities at expiration of the option. However, the volatility estimates from the RNDs compare favourably with LIFFE tabulated at-the-money implied volatility, which gives an even more biased estimate of realised volatility. These results suggest unbiased estimates of future asset prices can be obtained from RNDs from both European and American put options on those assets. Table 3 shows the average departure of the ESX and SEI RNDs from normality. The t-stats from paired t-tests of the series of values of skewness and kurtosis estimated from the RNDs, and those for the normal distribution, lead to rejection of the null hypothesis of no difference at the 95% level. It is important to appreciate that RNDs extracted from option prices are distinct from the historical distribution of the underlying asset returns (prices), since one is risk adjusted and the other is not. However, their shapes should be the same. Thus, if asset returns were normally distributed we would not expect to obtain negatively skewed, leptokurtic, or platykurtic RNDs. The results above are therefore consistent with other empirical findings suggesting that asset returns are not normally distributed.

5. Discussion and conclusions

5.1. Summary of findings

In this empirical study, we applied our method of estimating RNDs to a large set of FTSE 100 European style put option daily data, and then as an Ansatz to a corresponding set of American style options on the same underlying asset. Our results in Table 1 suggest that the RNDs obtained from each style of option are significantly different, reflecting the distorting influence of the early exercise possibilities for the American put options. We confirmed that estimates of volatility from the RNDs from both types of option were biased estimates of the actually realised volatilities at expiration, suggesting that prices for the options tend to over estimate volatility. However, caution is necessary in interpreting the latter results, as it is difficult to reliably

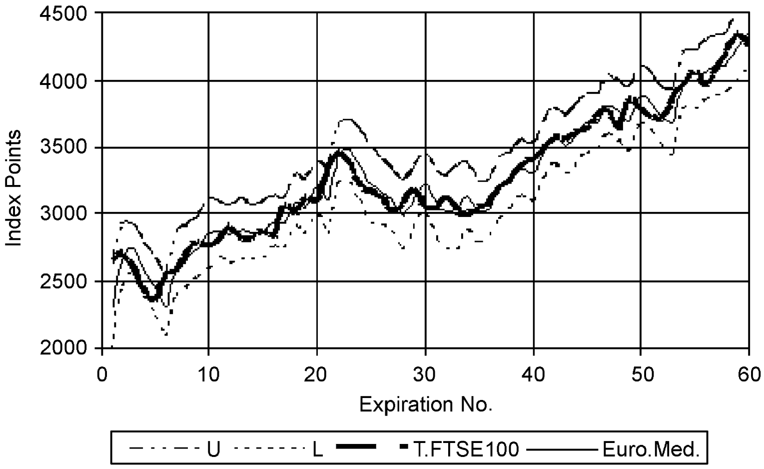

Fig. 1. ESX put option RNDs: Euro.Med is the one month estimate of realised FTSE 100 (T.FTSE100) showing it lies within U and L the 95% confidence band of the model.

estimate realised volatility from at most 18 daily observations of returns. The values of skewness and kurtosis obtained also suggest that the underlying pricing process departs from geometric Brownian motion. The results presented in Table 2 are surprising. They suggest that in practice, RNDs from both European and American put options can be used interchangeably to obtain estimates of future asset prices. This holds although the theory underlying RNDs outlined in Section 2 applies only to European style exercise, and despite the existence of significant model differences revealed by the (albeit more powerful) paired t-test results presented in Table 1. The median represents the centre of the distribution so is insensitive to extrapolations and errors in the tails of the RND. In addition, the standard deviation of RNDs for the American put options is smaller at 107 index points compared with 113 on average for the European options. The economic interpretation of the finding that the variance for the American put is smaller is that it is less risky for the purchaser than the European put because of the early exercise feature; this is reflected in the price of the former leading to the different RND. Reassuringly, Fig. 1 shows for ESX how in all cases the actual FTSE 100 closing prices lie within U and L, the ± 2 standard deviation confidence intervals constructed from the estimated standard deviations of each separate RND. The same applies to SEI.

5.2. Features and novelty of our approach to recovering RNDs

A fully non-parametric application of NNs was used to extract the RND from market prices of traded options. The use of a single NN model per option type series means that the set of RNDs are consistent in the sense of being predicted from the same model. Moreover, the combination of NNs and the methods of Breeden and Litzenberger [5] and Neuhaus [6] is straightforward to implement for practitioners, and does not require special programming. Unlike Herrmann and Narr [14] we do not substitute constants for some input variables, and we present results based on formal hypothesis tests, not just graphical comprisons and global goodness-of-fit criteria. Also, the analytical differentiation used by Herrmann and Narr is difficult to implement and requires special programming. Unlike Schittenkopf and Dorfner [13], we do not use a mixture of densities model, and thus avoid the limitations of that approach. Since we apply the method of Neuhaus to compute the probabilities in the tails of the distribution, the problem of accurately computing higher moments alluded to in Bondarenko is somewhat ameliorated. However, we did make use of the median as the centre of the distribution rather than the mean, which is only central in symmetrical distributions. The first key difference from both Bonderenko [7] and from Jackwerth [8] is that we construct a model by fitting to an extensive set of the market option prices that excludes the set for which we are attempting to extract the RND, whereas both Bonderenko and Jackwerth include the target set in the modelling; in Jackwerth's case via the implied volatility. The integrals of our RNDs are not constrained but approximate very closely to one while only deviating in the fifth decimal digit. However, not all of our RNDs are positive over their full strike range and we have noted at high moneyness (large X for puts) the tail can be negative, although not to the extent reported by Herrmann and Narr. Unlike Jackwerth, we use only an at-the-money market volatility in predicting the RNDs from our model.

5.3. Further work

Overall, our results suggest that NNs provide a promising method for use in extracting RNDs from option prices; this merits further investigation. To evaluate fully the potential and limitations of the approach we describe here, an empirical comparison of the statistical characteristics of RNDs from the PCA, FST, double lognormal, from a smoothed implied volatility smile, and from our approach using NNs is required. In particular, further consideration needs to be given to the observation of Dupont [17] on deep in-the-money American style options.

References

[1] Y. Ait-Sahalia, A.W. Lo, Nonparametric risk management and implied risk aversion, J. Econometrics 94 (2000) 9–51.
[2] B. Bahra, Implied Risk-Neutral Probability Density Functions From Option Prices: Theory and Applications, Bank of England, London, EC2R 8AH, 1997.
[3] J. Healy, M. Dixon, B. Read, F.F. Cai, A data centric approach to understanding the pricing of financial options, Euro. Phys. J. B 27 (2002) 219–227.
[4] J.V. Healy, M. Dixon, B.J. Read, F.F. Cai, Confidence limits for data mining models of option prices, Physica A 344 (2004) 162–167.
[5] D.T. Breeden, R.H. Litzenberger, Prices of state-contingent claims implicit in option prices, J. Bus. 4 (1978) 621–651.
[6] H. Neuhaus, The information content of derivatives for monetary policy, Discussion Paper 3/95, Economic Research Group of the Deutsche Bundesbank, 1995.
[7] O. Bondarenko, Estimation of risk-neutral densities using positive convolution approximation, J. Econometrics 116 (2003) 85–112.
[8] J.C. Jackwerth, Option-Implied Risk-Neutral Distributions and Risk Aversion, ISBN 0-943205-66-2 published by Research Foundation of the AIMR, 2004.
[9] R.J. Ritchey, Call option valuation for discrete normal mixtures, J. Finan. Res. 13 (1990) 285–296.
[10] D. Shimko, Bounds of probability, Risk 6 (4) (1993) 33–37.
[11] R. Bliss, N. Panigirtzoglou, Testing the stability of implied probability density functions, J. Bank. Finance 26 (2002) 381–422.
[12] J. Bennell, C.M.S. Sutcliffe, Black–Scholes versus artificial neural networks in pricing FTSE 100 options, Intell. Syst. Acc. Finance Manage. 12 (4) (2004) 243–260.
[13] C. Schittenkopf, G. Dorffner, Risk neutral density extraction from option prices: improved pricing from mixture density networks, IEEE Trans. Neural Networks 12 (2001) 716–723.
[14] R. Herrmann, A. Narr, Risk Neutrality, Risk 10, Technology Supplement, 1997, pp. 23–29.
[15] R. Garcia, R. Gencay, Pricing and hedging derivative securities with neural networks and a homogeneity hint, J. Econometrics 94 (2000) 93–113.
[16] P.C. Andreou, C. Caralambous, S.H. Martzoukos, Robust artificial neural networks for pricing of European options, Comput. Econ. 27 (2006) 329–351.
[17] D.Y. Dupont, Extracting Risk-neutral Probability Distributions from Option Prices Using Trading Volumes as a Filter, 2001. ⟨http://www.ihs.ac.at/publications/eco/es-104.pdf⟩.
[18] M.M. Dacorogna, R. Gencay, U. Muller, R.B. Olsen, O.V. Pictet, An Introduction to High Frequency Finance, Academic Press, New York, 47–51. 2001.