From skews to a skewed-t:
Modelling option-implied returns by a skewed Student-t

Cyriel de Jong and Ronald Huisman

ERIM REPORT SERIES *RESEARCH IN MANAGEMENT*	
ERIM Report Series reference number	ERS-2000-12-F&A
Publication status / version	draft / version March 2000
Number of pages	10
Email address first author	Cjong@fbk.eur.nl
URL Location (electronic Version)	http://www.eur.nl/WebDOC/doc/erim/erimrs20000518152023.pdf
Address	Erasmus Research Institute of Management (ERIM) Rotterdam School of Management / Faculteit Bedrijfskunde Erasmus Universiteit Rotterdam PoBox 1738 3000 DR Rotterdam, The Netherlands Phone: # 31-(0) 10-408 1182 Fax: # 31-(0) 10-408 9020 Email: info@erim.eur.nl Internet: www.erim.eur.nl

Bibliographic data and classifications of all the ERIM reports are also available on the ERIM website:
www.erim.eur.nl

ERASMUS **R**ESEARCH **I**NSTITUTE OF **M**ANAGEMENT

REPORT SERIES
RESEARCH IN MANAGEMENT

BIBLIOGRAPHIC DATA AND CLASSIFICATIONS		
Abstract	In this paper we present a new methodology to infer the implied risk-neutral distribution function from European-style options. We introduce a skewed version of the Student-t distribution, whose main advantage is that its shape depends on only four parameters, of which two directly control for the levels of skewness and kurtosis. We can thus easily vary parameters to compare different distributions and use the parameters as inputs to price other options. We explain the method, provide some empirical results and compare them with the results of alternative models. The results indicate that our model provides a better fit to market prices of options than the Shimko or implied tree models, and has a lower computation time than most other models. We conclude that the skewed Student-t method provides a good alternative for European-style options.	
Library of Congress Classification (LCC)	5001-6182	Business
	4001-4280.7	Finance Management, Business Finance, Corporation Finance
	6021	Futures, Options
Journal of Economic Literature (JEL)	M	Business Administration and Business Economics
	G 13 C 4 C40	Contigent pricing; Futures pricing Econometric and Statistical Methods: - Special Topics General
European Business Schools Library Group (EBSLG)	85 A	Business General
	220 A	Financial Management
	220 R	Option markets
Gemeenschappelijke Onderwerpsontsluiting (GOO)		
Classification GOO	85.00	Bedrijfskunde, Organisatiekunde: algemeen
	85.30	Financieel management, financiering
	85.30	Financieel management, financiering
Keywords GOO	Bedrijfskunde / Bedrijfseconomie	
	Financieel management, bedrijfsfinanciering, besliskunde	
	Optiehandel, Verdelingen (statistiek)	
Free keywords	Options, implied volatility, implied distribution, Student-t, skewness	
Other information		

From skews to a skewed-t:

Modelling option-implied returns by a skewed Student-t

Cyriel de Jong[a] and Ronald Huisman[b]

[a] Rotterdam School of Management, Erasmus University, cjong@fbk.eur.nl, +31 10 408 2300
[b] Finedge Rotterdam and Rotterdam Rotterdam School of Management, Erasmus University, ronaldhuisman@finedge.nl

Abstract

In this paper we present a new methodology to infer the implied risk-neutral distribution function from European-style options. We introduce a skewed version of the Student-t distribution, whose main advantage is that its shape depends on only four parameters, of which two directly control for the levels of skewness and kurtosis. We can thus easily vary parameters to compare different distributions and use the parameters as inputs to price other options. We explain the method, provide some empirical results and compare them with the results of alternative models. The results indicate that our model provides a better fit to market prices of options than the Shimko or implied tree models, and has a lower computation time than most other models. We conclude that the skewed Student-t method provides a good alternative for European-style options.

Keywords

options, implied volatility, implied distribution, Student-t, skewness

1 Introduction

One of the fundamental assumptions in the Black and Scholes (1973) model is that the risk-neutral expected returns on the underlying asset are drawn from a normal distribution. Volatility, measured as the standard deviation of the expected returns, determines the exact shape of this distribution. A good insight in volatility is therefore crucial to calculate proper option values. Since the 1987 crash however, implied volatilities calculated from option market prices have varied over the strike price (or moneyness) and time to maturity, instead of being constant as assumed by the Black-Scholes model. The variation indicates that risk-neutral expected returns are not normally distributed. Since 1987, implied volatility has been a convex function of strike price and referred to as a skew or smile, depending on its exact shape (Rubinstein - 1994, Derman - 1999). Skews and smiles thus refer to non-normal characteristics of implied risk-neutral return distributions. The exact shape of the implied distribution gives important information that can be used for pricing other options on the same underlying asset, for comparing options on different assets and for closely monitoring changes in the markets perception.

Breeden and Litzenberger (1978) were the first to show how the implied risk-neutral distribution function could be derived from option prices: the probabilities are equal to the second order derivatives of option prices with respect to the strike price. Shimko (1993) offers a practical application of this general idea. He proposes to model the volatility smile as a quadratic function of moneyness, and then to calculate the second order derivative numerically. This approach is simple and fast, but inaccurate outside the range of traded strike prices. Other methods construct implied binomial (Rubinstein - 1994) or trinomial trees (Nagot and Trommsdorff - 1999), or estimate the end-of-term distribution non-parametrically (Aït-Sahalia and Lo - 1995, Jackwerth and Rubinstein – 1996). In this paper we present a different methodology to infer the implied risk-neutral distribution function from European-style options. We introduce a skewed version of the Student-t distribution, which is known to provide a good fit to historical returns on many financial assets (see Huisman - 1999, Huisman, Koedijk, and Pownall – 1999 among others). The skew or smile pattern of implied volatility as a function of strike is a direct indication of skewness and excess kurtosis of the implied risk-neutral return distribution. A smile implies fat tails; a skew implies both fat tails and skewness. The skewed Student-t distribution is able to capture these characteristics. The advantage of our method (see the appendix) is that the whole distribution depends on only four parameters, of which two directly control for the levels of skewness and kurtosis. Moreover, the skewed Student-t nests the normal distribution. We can thus easily vary parameters to compare different distributions and use the parameters as inputs to price other options. The disadvantage is that we pre-impose a structure on the implied distribution and therefore give up some of the flexibility of non-parametrical methods. Moreover, unlike tree models, our model cannot be applied directly to American-style options.

In the following sections we explain the method, provide some empirical results and compare them with the results of alternative models. The results indicate that our model provides a better fit to market prices of options than the Shimko or implied tree models. Keeping in mind the computational speed of our model and its direct parameterisation of skewness and kurtosis, we conclude that the skewed Student-t method provides a good alternative for European-style options.

2 Results

In order to gain insight in the performance of our Student-t based model we start with an example that is presented both by Shimko (1993) and Nagot and Trommsdorff (1999). We compare our results with theirs, as well as the classical lognormal distribution. We obtain the four parameters of the skewed Student-t by minimising the sum of squared deviations between the observed and fitted prices.

Table 1 Shimko example

Strike	325	345	360	365	375	385
Price	66.500	46.000	33.000	27.750	20.125	13.500
Strike	390	395	400	405	410	425
Price	9.625	7.250	5.375	3.375	1.875	0.250

Consider the following European-style call options on the S&P 500 index. The prices are from October 21, 1991. The index value is 390.02, the interest rate 5.03%, the continuous dividend yield 3.14%, and the time to maturity 0.16. The prices of the call options are listed in table 1. The optimisation of the skewed-t method leads to the following estimates for the skewness (g) 0.50, degrees of freedom (a) 167.20, and normalising constants ($\tilde{m}$) 0.066 and ($\tilde{s}$) 0.044. The skewness parameter indicates pronounced negative skewness (γ equals one for a symmetric distribution). The degrees of freedom are relatively high (a level of 3 to 6 is not uncommon in financial markets – see Huisman - 1999), indicating limited excess kurtosis.

Table 2 Results Shimko example

This table presents the implied distribution characteristics from the example presented in table 1 for four methods. Normal refers to the lognormal implied distribution, Shimko refers to Shimko's method, ITT is the trinomial tree introduced by Nagot and Trommsdorff, Skewed-t is our method. The RMSE is the root mean squared error between the actual prices and the fitted prices of the option.

	RMSE	Mean	St. Deviation	Skewness	Kurtosis
Normal	0.74	391.2	510.0	0.17	3.05
Shimko	0.49	391.4	554.5	-0.96	5.84
ITT	0.34	391.7	515.3	-0.57	3.39
Skewed-t	0.32	391.8	502.9	-0.67	3.19

In table 2 we compare the fit as well as the first four distributional moments of the lognormal, the Shimko, the implied trinomial tree (ITT), and our method. The fit of the skewed-t method is somewhat better than the ITT, as indicated by the root mean squared error between the actual and fitted prices (0.32 versus 0.34). Both methods clearly improve upon Shimko and the lognormal method. The main differences appear in the estimates for skewness and kurtosis. By construction, the lognormal distribution has very low levels of positive skewness and excess kurtosis. Due to the instability of the tails, the levels of skewness and kurtosis in Shimko's approach are rather unreliable. For completeness, we show the implied volatilities and the density function of every method in Figure 1 and 2. The similarity in reported statistics between the skewed-t and trinomial tree approach is clearly visible in the shape of the density function and the implied volatility smile. They coincide at practically every price. It seems as if the evolution of the trinomial tree results in a skewed-t end-of-term distribution.

Figure 1 Implied distributions, Shimko example

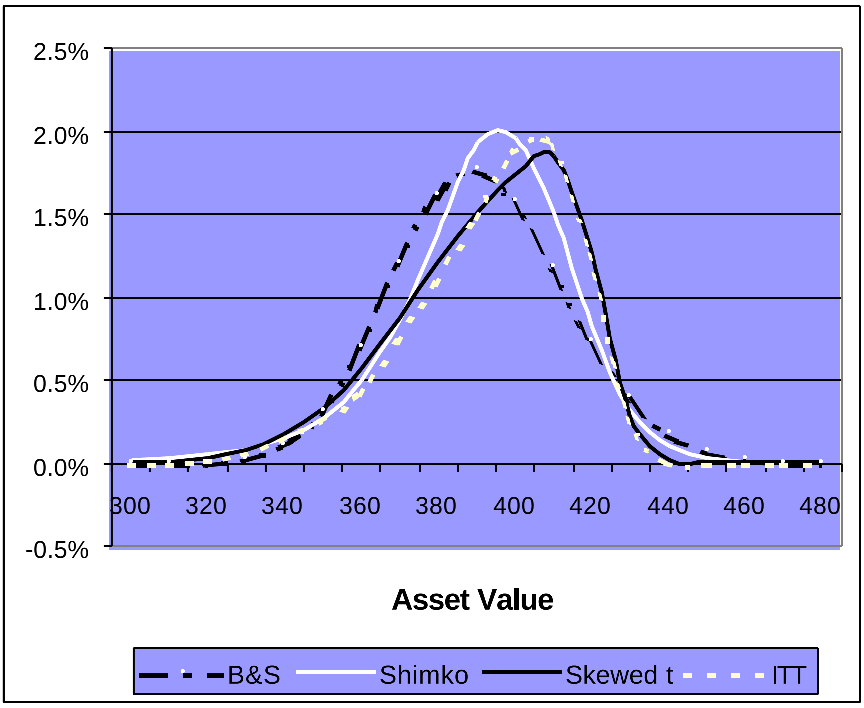

Figure 2 Implied volatilities, Shimko example

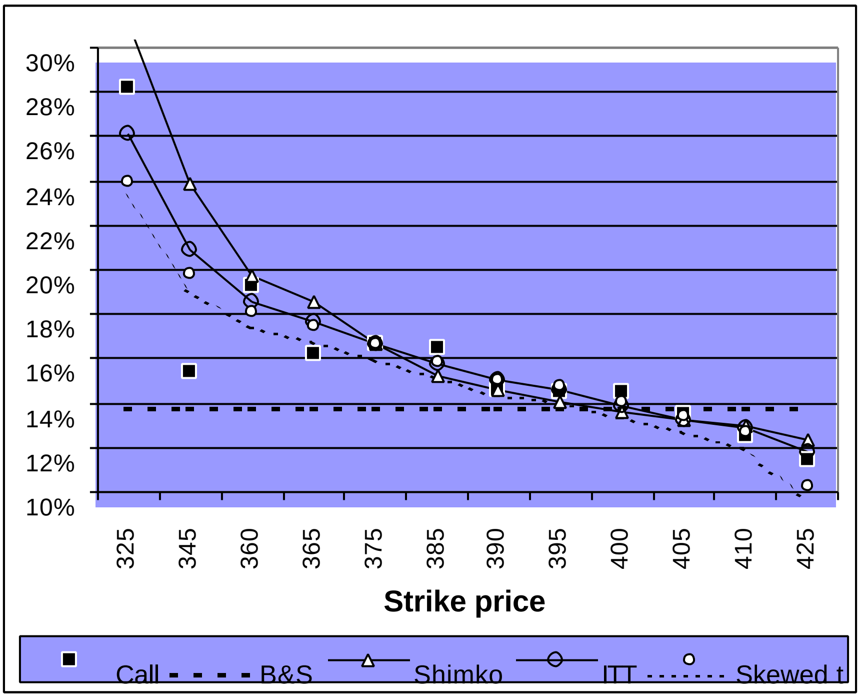

Based on the Shimko example we conclude that the skewed-t method shows a slightly better performance than the ITT method; the advantages lie in the computational speed and the direct control of two parameters for skewness and kurtosis. In the following empirical example we again compare the performance of the four methods, but with more recent data. We take the midprices of call options on the FTSE 100 Index from 10-9-1999, 10:17 a.m. (table 3).

Table 3 FTSE 100 example

Strike	5875	5925	5975	6025	6075	6125
Price	464.5	421.0	379.5	339.5	301.0	265.0
Strike	6175	6225	6275	6325	6375	6425
Price	231.0	200.0	170.0	143.0	118.0	96.5
Strike	6475	6525	6575	6625	6675	
Price	75.5	58.5	43.5	31.5	23.0	

The expiration date is 15-10-1999, the spot 6278.30, the continuous interest rate 5.19%, and the continuous dividend yield 1.88% (obtained from the Bloomberg TRMS sheet at approximately the same time as the price quotes). The implied volatilities can be described as a skew: a clear straight and downward sloping line (see figure 3). This is already a first indication of negative skewness and only limited tail fatness in the implied distribution.

The estimation results confirm this indication (table 4). Apart from the lognormal structure, skewness is apparent in all implied distributions, ranging from –0.50 for the ITT, -0.55 for the skewed-t, to –0.83 for Shimko. As could be expected from the implied volatilities, the implied distributions contain limited excess kurtosis (estimates of 3.10 for ITT and 3.21 for the skewed-t)[1]. The RMSE of the skewed-t method equals 0.94 versus 1.06 for ITT, 2.40 for Shimko, and 11.73 for the lognormal method, stressing the strength of our skewed-t. The parameter estimates for the skewed-t are skewness (γ) 0.64, degrees of freedom (α) 34.86, and normalising constants ($\tilde{m}$) -0.016 and ($\tilde{s}$) 0.195.

[1] Shimko's method yields a kurtosis of 4.14, but is rather unreliable, due to the instability of the tail probabilities. The same holds for the skewness estimate of –0.83, though to a lesser extent.

Figure 3 Implied volatilities, FTSE100 example.

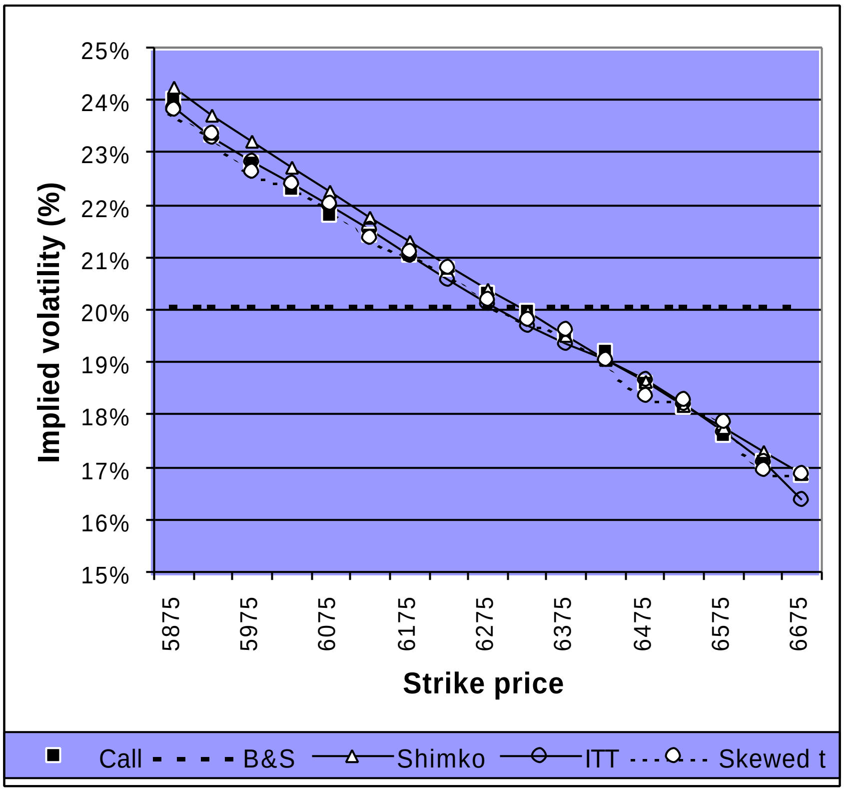

Table 4 Results FTSE 100 example

This table presents the implied distribution characteristics from the example presented in table 3 for four methods. Normal refers to the lognormal implied distribution, Shimko refers to Shimko's method, ITT is the trinomial tree introduced by Nagot and Trommsdorff, Skewed-t is our method. The RMSE is the root mean squared error between the actual prices and the fitted prices of the option.

	RMSE	Mean	St. Deviation	Skewness	Kurtosis
Normal	11.73	6298.4	395.1	0.18	3.04
Shimko	2.40	6298.8	402.5	-0.83	4.14
ITT	1.06	6312.3	370.3	-0.55	3.10
Skewed-t	0.94	6332.8	369.8	-0.50	3.21

3 Conclusion

The two empirical examples show that the skewed-t approach is an attractive alternative to infer the risk-neutral density from European-style option prices. The skewed-t method is based on estimating four parameters, of which two directly control for skewness and tail fatness. This makes a comparison over different options or assets very easy. The skewed-t approach is fast and provides a fit that is slightly better than the implied trinomial tree of Nagot and Trommsdorff and much better than Shimko's approach. Unlike the trinomial tree, there is no convincing way for the skewed-t model to account for the early exercise premium in American-style options, which makes our method only applicable to European-style options.

References

Aït-Sahalia Y and A Lo, 1995, *Nonparametric estimation of state-price densities implicit in financial asset prices*, Working paper series 5351, National Bureau of Economic Research

Black F and M Scholes, 1973, *The pricing of options and corporate liabilities*, Journal of Political Economy 81, pages 637-659

Breeden D and R Litzenberger, 1978, *Prices of state-contingent claims implicit in option prices*, Journal of Business 51, pages 621-652

Derman E, 1999, *Regimes of volatility*, Risk April, pages 55-59

Huisman R, 1999, *Adventures in international financial markets*, non-published manuscript

Huisman R, C Koedijk, and R Pownall, 1998, *VAR-x: fat tails in financial risk management*, Journal of Risk 1

Jackwerth J and M Rubinstein, 1996, *Recovering probability distributions from option prices*, Journal of Finance 5, pages 1611-1631

Nagot I and R Trommsdorff, 1999, *The tree of knowledge*, Risk August, pages 99-102

Rubinstein M, 1994, *Implied binomial trees*, Journal of Finance 49, pages 771-818

Shimko D, 1993, *Bounds of probability*, Risk April, pages 33-37

Appendix

The skewed-t method assumes that the expected risk-neutral return distribution implied in option prices equals a skewed Student-t distribution. In this appendix we introduce this distribution.

Most are familiar with the central Student-t distribution. Therefore, let us first consider the probability density function $f(x \mid \alpha)$ of the central Student-t distribution with α degrees of freedom. It reads:

$$f(x \mid \alpha) = c(\alpha) \cdot \left(1 + \frac{x^2}{\alpha}\right)^{-\frac{\alpha+1}{2}} \qquad \text{(A.1)}$$

where $c(\alpha)$ a constant that exclusively depends on α (Bain and Engelhardt, 1992, p. 274)

The central Student-t distribution is symmetric with mean equal to zero. For values of α larger than two, the variance is defined and equals $\frac{\alpha}{\alpha - 2}$. The parameter α is called the number of degrees of freedom and controls the level of tail fatness. The smaller the degrees of freedom, the fatter the tails. The Student-t distribution nests the normal distribution: if α approaches infinity the Student-t converges to the normal distribution.

The Student-t distribution can be made skewed by inverse scaling of the probability density function (p.d.f.) both sides of the mode. This does not affect unimodality and enables control of the probability mass both sides of the mode with a single parameter. The p.d.f. $g(x \mid \alpha, \gamma)$ of this skewed-t reads:

$$g(x \mid \alpha, \gamma) = \begin{cases} f(x \cdot \gamma \mid \alpha) & x < 0 \\ f(x/\gamma \mid \alpha) & x \geq 0 \end{cases} \qquad \text{(A.2)}$$

with γ: skewness parameter that assigns the distribution of mass both sides of the mode

The p.d.f. is symmetric for $\gamma = 1$, negatively skewed for $\gamma < 1$ and positively skewed for $\gamma > 1$. We obtain the standard normal distribution if $\gamma = 1$ and α approaches infinity. Note that γ not only influences the skewness, but also the mean of the distribution. In order to obtain the desired mean and variance, one normalises the observations as follows:

$$x = \frac{\ln\left(\frac{S_T}{S_0}\right) - \tilde{m}}{\tilde{s}},$$

where: S_T and S_0 the asset values at expiry and now respectively

$\tilde{m}$ and $\tilde{s}$ normalising constants

The skewed-t distribution is therefore parameterised by four parameters: the skewness parameter g, the kurtosis parameter a, and the normalising constants $\tilde{m}$ and $\tilde{s}$. There are several methods that can be used to obtain parameter estimates. Unlike the Black-Scholes formula, there is no closed form for the option price under the skewed-t distribution. Therefore we divide the density function into 200 intervals and analytically derive the option price. In the optimization algorithm (that finds the best-fitting parameters) we start with 10 intervals and then progressively increase the number of intervals. Convergence is independent of the final number of intervals.

ERASMUS **R**ESEARCH **I**NSTITUTE OF **M**ANAGEMENT

REPORT SERIES
RESEARCH IN MANAGEMENT

Publications in the Report Series Research* in Management

Impact of the Employee Communication and Perceived External Prestige on Organizational Identification
Ale Smidts, Cees B.M. van Riel & Ad Th.H. Pruyn
ERS-2000-01-MKT

Critical Complexities, from marginal paradigms to learning networks
Slawomir Magala
ERS-2000-02-ORG

Forecasting Market Shares from Models for Sales
Dennis Fok & Philip Hans Franses
ERS-2000-03-MKT

A Greedy Heuristic for a Three-Level Multi-Period Single-Sourcing Problem
H. Edwin Romeijn & Dolores Romero Morales
ERS-2000-04-LIS

Integer Constraints for Train Series Connections
Rob A. Zuidwijk & Leo G. Kroon
ERS-2000-05-LIS

Competitive Exception Learning Using Fuzzy Frequency Distribution
W-M. van den Bergh & J. van den Berg
ERS-2000-06-LIS

Start-Up Capital: Differences Between Male and Female Entrepreneurs, 'Does Gender Matter?'
Ingrid Verheul & Roy Thurik
ERS-2000-07-STR

The Effect of Relational Constructs on Relationship Performance: Does Duration Matter?
Peter C. Verhoef, Philip Hans Franses & Janny C. Hoekstra
ERS-2000-08-MKT

Marketing Cooperatives and Financial Structure: a Transaction Costs Economics Analysis
George W.J. Hendrikse & Cees P. Veerman
ERS-2000-09-ORG

A Marketing Co-operative as a System of Attributes: A case study of VTN/The Greenery International BV,
Jos Bijman, George Hendrikse & Cees Veerman
ERS-2000-10-ORG

---

* ERIM Research Programs:
LIS Business Processes, Logistics and Information Systems
ORG Managing Relationships for Performance
MKT Decision Making in Marketing Management
F&A Financial Decision Making and Accounting
STR Strategic Renewal and the Dynamics of Firms, Networks and Industries

ERASMUS **R**ESEARCH **I**NSTITUTE OF **M**ANAGEMENT

REPORT SERIES

RESEARCH IN MANAGEMENT

Evaluating Style Analysis
Frans A. De Roon, Theo E. Nijman & Jenke R. Ter Horst
ERS-2000-11-F&A