M De Araújo and E Maré*

Examining the volatility skew in the South African equity market using risk-neutral historical distributions

1. INTRODUCTION

1.1 Volatility skews and the notion of fairness

The volatility skew is a function relating the implied volatility of an option to its strike price. For most equities and their associated indices, a plot of this function yields a decreasing curve, which indicates that, as the strike price of an equity option increases, so its implied volatility tends to decrease. Implied volatilities also differ across maturities. A plot of implied volatility as a function of strike and term to maturity is called a volatility surface (see Figure 1).

Volatility skews have been discussed in the derivatives literature for more than a decade. Derman and Kani (1994), Dupire (1994), Rubinstein (1994) and Derman (1999) are a few of the many published contributions on the topic.

It has been argued that the skew exists to correct inaccuracies in the assumptions of the conventional Black and Scholes model (Black and Scholes, 1973; Jackwerth, 2000). In particular, the assumptions of normally distributed logarithmic asset returns and zero transaction costs result in a mispricing of options. Chen, Palmon and Wald (2003) have recently attempted to quantify the extent to which each of these factors may impact on the skew. They find that most of the so-called smile bias is removed when correcting for the non-realistic Black and Scholes assumptions.

The implied volatility skew is, however, only evidenced after the equity market crash of 1987. This is notably the case for the S&P 500 Index (Derman 1999). Prior to 1987, the market seems to have priced implied volatility independently of strike level. It may be fair then to argue that market participants quote implied volatilities for buying/selling options based on their aversion to market crashes, their specific views on future volatility levels and on supply and demand considerations. These arguments are supported by Rubinstein (1994), Bates (1996), Clews, Panigirtzoglou and Proudman (2000) and Jackwerth (2000).

Some market participants view the volatility skew as a means by which market makers may extract more profitability from options trades. It should be remembered that implied volatility quotes contain a bid-offer spread which reflects profit potential to the option market maker, while the skew reflects potential risk and should not be construed as a risk margin. Notwithstanding this, it is still valuable to measure the richness or cheapness of the observed volatility skew to determine just how fair it actually is. This is especially necessary for illiquid traded options where price transparency proves to be problematic as is the case in the South African equity derivatives market.

To determine if the observed volatility skew is fair in some sense, we need to derive the skew independently of model assumptions and subjective market expectations. We utilise nonparametric option pricing techniques along the same lines as Stutzer (1996) and Duan (2002) to estimate a nonparametric volatility skew based only on the historical returns of the underlying asset (see section 2). Our approach is documented in section 3 below. We apply these techniques to the FTSE/JSE Top40 index futures options' markets and present our results in section 4. We conclude with a discussion of our results and their relevance to the South African market.

2. NONPARAMETRIC ESTIMATION OF THE VOLATILITY SKEW

In liquid markets, the volatility skew is typically estimated from a sample of traded option prices. The sample is used to calibrate the relationship between option prices and the underlying asset price. For many South African counters however, there are few or no option trades, resulting in a lack of price transparency. For these counters, utilising observed option prices to calibrate the volatility skew might be unfeasible or impossible. An alternative approach has been to utilise nonparametric option pricing techniques to estimate the volatility skew (Stutzer, 1996 and Duan, 2002).

The Black and Scholes (1973) option pricing model and other parametric models make several assumptions that aim to simplify model building and make their solutions analytically tractable. The assumptions of normally distributed logarithmic asset returns, constant volatility and interest rates, costless trading and zero dividends all belie reality. Simplifying assumptions often lead to models that are not a true reflection of the dynamics of the underlying asset, resulting in mis-estimation of option values as was argued in the section above.

* Department of Mathematics and Applied Mathematics, University of Pretoria, Pretoria 0002, Republic of South Africa.
Email: mark.dearaujo@absa.co.za; eben.mare@up.ac.za

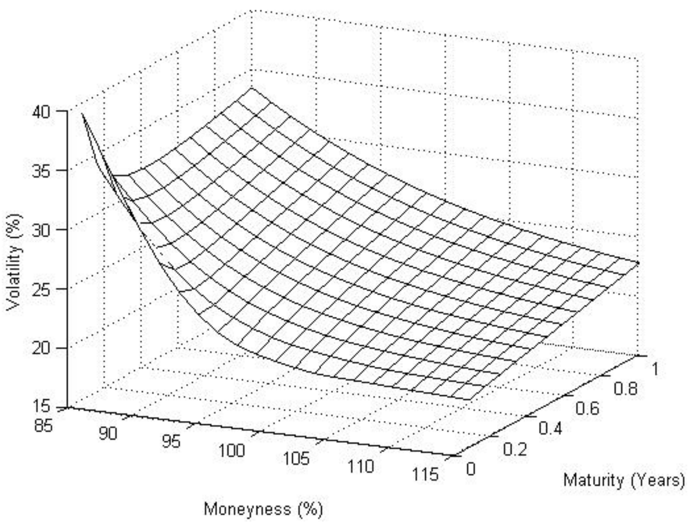

Figure 1: Volatility surface for the Top40 index for the period from January 1996 to September 2005. Data: INet-Bridge

Nonparametric option pricing techniques utilise spot market observed security prices in order to determine the probability distribution of the underlying asset. Derman and Kani (1994) and Rubinstein (1994) utilise implied binomial tree approaches to find risk-neutral distributions which result in estimated option prices that match observed option prices. Buchan and Kelly (1996) and the canonical valuation method of Stutzer (1996) estimate option prices by recovering the risk-neutral distribution from the empirical distribution obtained from a set of stock prices. Stutzer and Chowdhury (1999) apply this technique to the pricing of options on bond futures. Duan (2002) expands the canonical valuation method of Stutzer to allow for the pricing of path-dependent options. It is Duan's nonparametric formulation that forms the basis of this investigation.

These nonparametric approaches impose no assumptions about the nature of the probability distribution of the asset. This leaves little room for mis-estimation of option pricing resulting from an incorrect choice of the probability distribution of the underlying asset. Thus nonparametric methods do not suffer from the smile bias that is evidenced in parametric models.

Using nonparametric methods to estimate the price of an option and reversing this price through the Black and Scholes formula will yield the conventional implied volatility. Following this procedure for a range of possible strikes will result in the so-called nonparametric volatility skew (Derman and Kani, 1994; Duan, 2002). The nonparametric volatility skew is free of any model-specific assumptions and, as it is based only on observed asset data, it is a minimally subjective estimate of the volatility skew. Thus, we will contrast nonparametrically derived volatility skews against market observed implied volatility skews to determine the fairness of observed skews within the South African market.

3. METHODOLOGY

Options on index futures trade in reasonable volumes in the South African market on both the South African Futures Exchange (SAFEX) and over-the-counter (OTC) markets. Within the South African market, FTSE/JSE Top40 index (Top40) futures options are the most actively and liquid traded derivatives.

SAFEX options are traded on the basis of implied volatility and the option price is calculated using the Black formula (Black 1976) for options on futures. In order to establish whether the market observed volatility skews are fair within this market we compare observed volatility skews to nonparametric volatility skews.

To estimate nonparametric volatility skews we implement a modified version of Stutzer's canonical valuation method (Stutzer, 1996) that was proposed by Duan (2002). The method is based only on the price data of the underlying asset and does not require observed option price data to be calibrated.

We outline the methodology used here; the technical aspects have been relegated to the Appendix.

Firstly, we obtain a nonparametric distribution function for the normalised continuously compounded return of the underlying asset. The aim is to characterise the risk-neutral distribution directly from the normalised empirical distribution.

Secondly, we need to find the risk-neutral distribution corresponding to the normalized asset return, i.e., we find the density function that minimizes the relative entropy subject to the condition that the expected return of the risk-neutral distribution is equal to the risk-free rate less the dividend yield.

Once we have found the risk-neutral distribution, we perform a Monte Carlo simulation to generate the asset price at maturity and compute the expected value of the option.

We vary the maturity and strike parameters and repeat the nonparametric option pricing procedure outlined above. The option values produced in this way are then converted to implied volatilities using the Black and Scholes formula (1973) and the implied volatilities are plotted against strike and term to maturity, yielding the volatility surface. We document the results of our analysis as applied to the Top40 index futures market in the section below.

4. RESULTS OF FTSE/JSE TOP40 VOLATILITY SKEW ANALYSIS

Applying the method described in Section 3 we calculate nonparametric volatility skews for the Top40 index for various maturities. For this purpose we obtain Top40 closing prices and dividend yields from INet-Bridge for the period 1 January 1996 to 16 September 2005. This data period yields 2480 returns. The prevailing continuously compounded interest rate on the 16 September 2005 is assumed to be $r = 0{,}0709$. We estimate the parameters $\mu = 0{,}0764$ and $\sigma = 0{,}2104$ from the entire data sample. We assume 252 trading days per year.

The resulting volatility surface is shown in Figure 1.

In Figures 2 and 3 we show the calculated skews for three and six months respectively. For comparative purposes we also show the market implied volatility skew for options traded on SAFEX.

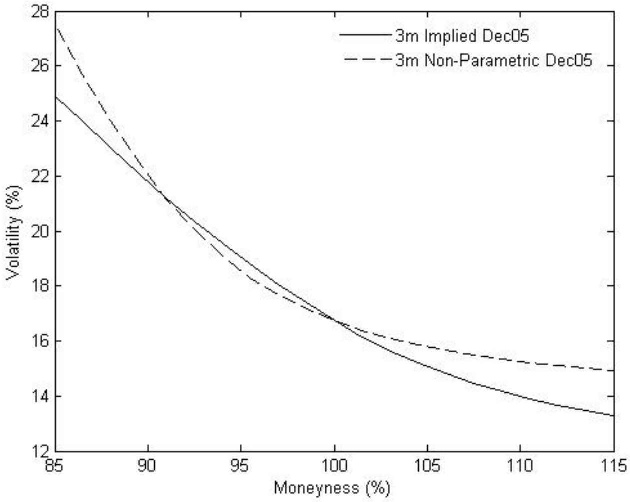

Figure 2: Three month Top40 nonparametric skew vs. implied skew for the December 2005 expiry as measured and observed (respectively) on 16 September 2005. Data: INet-Bridge and Peregrine Securities

It is important to note that the implied volatility skew is dependent on supply and demand factors, the open interest of all option strikes for a given maturity, bid/offer spreads as well as traders' views on the market. The nonparametric skew on the other hand is calculated using only historical data and is therefore a reflection of the historical cost of replicating an option with a specific strike and maturity. It is therefore not expected that the two skews will be identical.

In both figures we obtain a reasonable resemblance between the market skew and the nonparametric skew. A noticeable difference appears for the shorter dated skew to the right-hand side of the graph (see, Figure 2). This area of the skew is relevant to the pricing of deep out-of-the-money call options. To understand this divergence it is worthwhile noting the workings of the South African index derivative market in more detail.

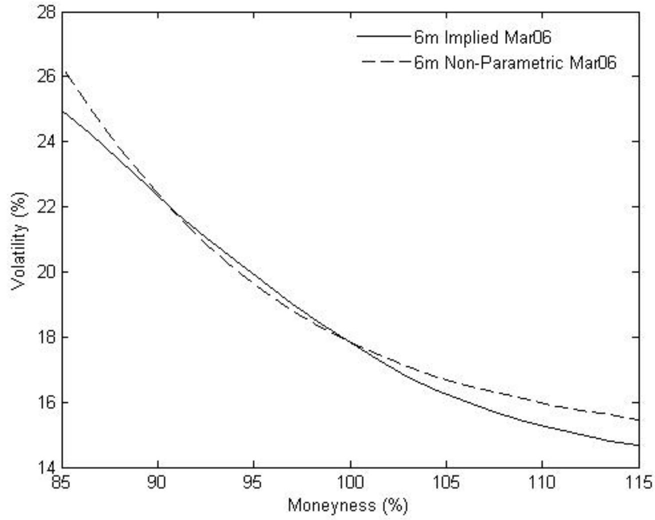

Figure 3: Six month Top40 nonparametric skew vs. Implied skew for the March 2006 expiry as measured and observed (respectively) on 16 September 2005. Data: INet-Bridge and Peregrine Securities.

Typically, the larger banks act as option market makers, and are responsible for providing and maintaining liquidity within the options market. Their aim is to earn the bid-offer spread by trading options or by dynamically replicating long or short option positions. The option market maker therefore manages a book of option positions with different maturities and strikes by dynamically trading in the underlying. This is frequently called Delta hedging.

Institutions, on the other hand, frequently maintain buy-sell and hold strategies. This means that an institution would buy or sell an option and retain this position until its expiry. Institutions frequently buy put options for protection and sell call options to fund the purchase of the put options or perform so-called out-of-the-money call overwriting to earn the option premium. The option market maker will therefore end-up with the opposite position, i.e., frequently long of call options which are out-of-the-money. As the options near their maturities, the market maker would then attempt to sell these options in the market, which leads to an over-supply. Other market makers would bias their bid prices to take account of the over-supply. This explains why the implied volatility skew, especially for short-dated options, is cheap relative to the nonparametrically derived skew.

5. CONCLUSIONS

We have demonstrated how a nonparametric option pricing method could be used in the context of the South African equity derivatives market. This methodology allows investors to glean information about the volatility surface of a specific underlying without the necessity of a liquid derivative market in that underlying.

Nonparametrically derived skews for the SAFEX index derivative market allow us to establish whether options are trading expensive or cheap relative to their historically calculated replication costs. In a comparison of estimated nonparametric skews to observed market skews for the FTSE/JSE Top40 index, it is noted that implied volatilities for far out-of-the-money options are cheap relative to the historical costs of replicating these options. This may be attributable to the trading behaviour of market makers closer to contract expiries.

The nonparametric approach for estimating volatility skews is also clearly valuable in the South African context where investors would want to obtain meaningful indications of where derivative prices should trade on counters with illiquid derivative contracts.

ACKNOWLEDGEMENTS

The authors wish to extend their appreciation to the comments and suggestions of an anonymous referee which improved significantly on an earlier version of this paper.

REFERENCES

Bates DS. 1996. Testing option pricing models. In Maddala GS and Rao CR, eds., *Handbook of statistics*. Amsterdam: North-Holland. 15:567-611.

Black F. 1976. The pricing of commodity contracts. *Journal of Financial Economics*, 3:167-179.

Black F and Scholes M. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81:637-654.

Buchen PW and Kelly M. 1996. The maximum entropy distribution of an asset inferred from option prices. *Journal of Financial and Quantitative Finance,* 31:143-159.

Chen R, Palmon O and Wald J. 2003. What is behind the smile? Transaction costs or fat tails. *Rutgers University Paper*.

Clews R, Panigirtzoglou N and Proudman J. 2000. Recent developments in extracting information from options markets. *Bank of England Quarterly Bulletin*, 40(1):50-57.

Cover TM and Thomas JA. 1991. *Elements of information theory*. New York: Wiley and Sons.

Derman E. 1999. Regimes of volatility: Some observations on the variation of S&P 500 implied volatilities. Quantitative Strategies Research Notes, Goldman Sachs.

Derman E and Kani I. 1994. Riding on a smile. *Risk*, 7:32-39.

Duan JC. 2002. Nonparametric option pricing by transformation. Working paper, University of Toronto.

Dupire B. 1994. Pricing with a smile. *Risk*, 7:18-20.

Jackwerth J. 2000. Option implied risk neutral distributions and implied binomial trees: A literature review. *Journal of Derivatives*, 7:66-82.

Rubinstein M. 1994. Implied binomial trees. *Journal of Finance*, 69:771-818.

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

Stutzer M and Chowdhury M. 1999. A simple non-parametric approach to bond futures option pricing. *Journal of Fixed Income*, 8(4):67-76.

Zou J and Derman E. 2001. Strike-adjusted spread: A new metric for estimating the value of equity option. In C. Alexander ed., *Mastering Risk, Vol. 2: Applications*. Financial Times/Prentice Hall. 39-56.

APPENDIX: NONPARAMETRIC OPTION PRICING METHOD OF DUAN

The nonparametric option pricing approach described in this appendix follows from the work of Duan (2002). Buchen and Kelly (1996) and Stutzer (1996) lay the foundations for the approach and ought to be read in conjunction. (See also Zou and Derman, 2001.)

The aim is to characterise the risk-neutral distribution directly from the normalised empirical distribution of the underlying asset, following from which, the option price may be estimated.

Given a sample of asset returns $\{R_t, t = 1,...,n\}$, with sample mean $\bar{\mu}$ and sample standard deviation $\bar{\sigma}$, define the normalized asset return $\{Z_t, t = 1,...,n\}$ by $Z_t = \Phi^{-1}\left(G\left(\frac{R_t - \bar{\mu}}{\bar{\sigma}}\right)\right)$ where $\Phi(\cdot)$ is the standard normal distribution function and $G(\cdot)$ is the distribution function of $\left\{\frac{R_t - \bar{\mu}}{\bar{\sigma}}, t = 1,\ldots,n\right\}$. The normalised asset return has a standard normal density $\phi(x)$.

To estimate $G(\cdot)$, we first determine the empirical distribution function of the sample $\left\{\frac{R_t - \bar{\mu}}{\bar{\sigma}}, t = 1,\ldots,n\right\}$, given by $\hat{G}(x) = \frac{1}{n}\sum_{t=1}^{n} I\{(R_t - \bar{\mu})/\bar{\sigma} \le x\}$, where $I\{\bullet\}$ is the indicator function. Since $\hat{G}(\cdot)$ is a step function, it is smoothed to generate $G(\cdot)$. Smoothing will allow for inversion of the distribution function. Details of the smoothing technique applied can be found in Appendix 5.1 in Duan (2002).

To determine the risk-neutral distribution corresponding to the normalised asset return, we find the density function $f(x)$ that minimizes the relative entropy between itself and $\phi(x)$ subject to the condition that the expected return of the risk-neutral distribution is equal to the risk-free rate less the dividend yield.

Using the relative entropy principle, the risk-neutral density for the normalized return $\{Z_t\}$ is given by minimizing the integral

$$\int_{-\infty}^{\infty} f(x)\ln\left(\frac{f(x)}{\phi(x)}\right)dx,$$

with respect to $f(x)$, subject to the usual density function conditions

$$\int_{-\infty}^{\infty} f(x)dx = 1,$$

and

$$\int_{-\infty}^{\infty} xf(x)dx = c_t.$$

The constant c_t reflects that the risk-neutral density may be a function of time.

The above problem has the solution (Cover and Thomas, 1991)

$$f(x,\lambda_t) = \frac{\phi(x)\exp(\lambda_t x)}{\int_{-\infty}^{\infty}\phi(x)\exp(\lambda_t x)dx} = \phi(x - \lambda_t),$$

The value λ_t corresponds to a given value of c_t. Hence the density function is fully parameterized by λ_t, which is determined by the constraint that the expected return of risk-neutral density must be equal to the risk-free rate, r, less the dividend yield, d.

The sequence of asset returns can be expressed as $\{R_t, t = 1,...,n\} = \{\bar{\sigma}G^{-1}(\Phi(Z_t)) + \bar{\mu}, t = 1,...,n\}$. Due to this, we use the bisection search method to find the value of λ_t^* that solves

$$\int_{-\infty}^{\infty} \exp\left(\sigma G^{-1}\left(\Phi(x)\right) + \mu\right)\phi\left(x - \lambda_t^*\right)dx = \exp(r - d).$$

Note that it is not necessary to have $\mu = \bar{\mu}$ and $\sigma = \bar{\sigma}$. The risk-neutral density function of the normalized returns $\{Z_t\}$ is then given by $\phi(x - \lambda_t^*)$.

The risk-neutral asset value dynamics become

$$S_t = S_{t-1}\exp\left(\sigma G^{-1}\left(\Phi(x)\right) + \mu\right),$$

where Z_t is a normal random variable with mean λ_t^* and variance 1. The asset price until the maturity may then be modelled by using a Monte Carlo simulation from which the expected value of the option payoff is computed.