

Testing the stability of implied probability density functions

Robert R. Bliss ^{a,*}, Nikolaos Panigirtzoglou ^{b,1}

^a *Federal Reserve Bank of Chicago, 230 South La Salle Street, Chicago, IL 60604-1413, USA*

^b *Monetary Instruments and Markets Division, Bank of England, Threadneedle Street, London EC2R 8AH, UK*

Abstract

This paper examines the absolute and relative robustness of two of the most common methods for estimating implied probability density functions (PDFs) – the double-lognormal approximating function and the smoothed implied volatility smile methods – using short sterling futures options and the FTSE 100 index options. The changes resulting from randomly perturbing quoted prices by no more than a half tick provide a lower bound on the confidence intervals of the summary statistics derived from the estimated PDFs. Our tests show that the smoothed implied volatility smile method dominates the double lognormal as a technique for estimating implied PDFs. © 2002 Elsevier Science B.V. All rights reserved.

JEL classification: G13; C13; C15

Keywords: Implied density functions; Stability of estimates; Confidence intervals

1. Introduction

Implied probability density functions (PDFs) estimated from cross-sections of observed option prices are gaining increasing attention. They are used to

* Corresponding author. Tel.: +1-312-322-2313; fax: +1-312-322-2357.

E-mail addresses: robert.bliss@chi.frb.org (R.R. Bliss),
nikolaos.panigirtzoglou@bankofengland.co.uk (N. Panigirtzoglou).

¹ Tel.: +44-171-601-5440; fax: +44-171-601-5953.

price complex derivatives. A number of authors have used implied PDFs as indicators of market sentiment to examine whether options markets anticipated major economic events.² Central banks, in particular, have been interested in using implied PDFs to assess market participants' expectations of future changes in interest rates, stock prices, and exchange rates.³ A number of methods have been developed in the literature for estimating implied PDFs. To date however, little attention has been paid to the robustness of these estimates or to the confidence that users can place in the summary statistics, for example skewness or the 99th percentile, derived from these fitted PDFs.

This paper begins to address these questions by examining the absolute and relative robustness of two common methods for estimating implied PDFs – the double-lognormal approximating function (DLN), and the smoothed implied volatility smile (SML) methods – to small errors in recorded option prices. We do this by randomly perturbing prices by no more than plus or minus one half of the quotation tick size. The half-tick size represents the minimum irreducible uncertainty associated with option prices.

Tests are conducted using short sterling futures options and FTSE 100 index options contracts, both trading on the London financial futures and options exchange (LIFFE). Our results show that the double-lognormal method for estimating implied PDFs is systematically less stable than the smoothed implied volatility smile method, even when the latter is calibrated to have the same goodness-of-fit. We conclude that the smoothed implied volatility smile method dominates the double lognormal as a technique for estimating implied PDFs.

The remainder of the paper is organized as follows: Section 2 provides an overview of PDF estimation techniques and their applications, together with a discussion of the potential sources of error in the underlying options prices. Technical details of the two implied PDF estimation methods used in this paper are given in Appendix A. Section 3 discusses the data and empirical tests to be carried out. Details of the underlying options contracts are given in Appendix B. Section 4 discusses the empirical results and Section 5 concludes.

² See for example Butler and Davies (1998), Campa et al. (1998), Coutant et al. (1999), Gemmill and Saffekos (1999), Leahy and Thomas (1996), Malz (1996), McCauley and Melick (1996), McManus (1999), Nakamura and Shiratsuka (1999), and Söderlind (1999).

³ Examples of central bank research in the use of implied PDFs include Banco de España: Manzano and Sanchez (1998); Bank of Canada: McManus (1999); Bank of England: Bahra (1996, 1997) and Butler and Davies (1998); Bank of International Settlements: Galati and Melick (1999); Bank of Japan: Nakamura and Shiratsuka (1999) and Shiratsuka (1999); Banque de France: Coutant et al. (1999) and Coutant (1999); European Central Bank: Hördahl (1999); Federal Reserve Bank of Atlanta: Abken et al. (1996); Federal Reserve Bank of New York: Malz (1996, 1997a,b); Federal Reserve Board: Leahy and Thomas (1996) and Melick and Thomas (1997); and Deutsche Bundesbank: Neuhaus (1995).

2. Implied PDF estimation

2.1. Literature review

Methods for estimating implied PDFs fall into five groups: stochastic process methods, implied binomial trees, PDF approximating function methods, finite difference methods, and implied volatility smoothing methods.

Stochastic process methods for estimating PDFs begin by assuming a model for the stochastic process driving the prices of the underlying security, usually one for which it is possible to obtain an analytical solution to the implied PDF for a given horizon.⁴ After estimation, the parameters of the stochastic process are plugged into the analytical formula for the PDF. For instance, Malz (1996) fits a lognormal-jump diffusion process to over-the-counter foreign exchange derivative prices and then analytically computes risk-neutral realignment probabilities around the time of the 1992 Exchange Rate Mechanism crisis. The stochastic process approach can be used in the absence of options prices (the other approaches cannot). For instance, Hördahl (1999) applied the Longstaff–Schwartz model to Swedish interest rates. The Longstaff–Schwartz model has analytic solutions for both the term structure of interest rates observed at any point in time and the future distribution of short rates at any given horizon. Hördahl used the observed term structure of Swedish interest rates to estimate the Longstaff–Schwartz model parameters. These estimated parameters were then substituted into the analytically derived PDF function to produce the estimated implied PDF.

The implied binomial tree method was developed in Rubinstein (1994) and Jackwerth and Rubinstein (1996). The method seeks to build a binomial tree for the value of the underlying asset. The tree is constructed so as to minimize deviations from a lognormal process subject to the tree fitting the observed options prices. The implied binomial tree is thus a non-parametric Bayesian technique related to stochastic process methods in that its focus is on modeling the evolution of the underlying asset’s price.

Approximating function methods begin with the option pricing relation in Cox and Ross (1976) who show that the price of an option is the discounted risk neutral expected value of the payoffs:

$$C(t, T, K) = e^{-r(T-t)} \int_K^\infty w(S_T)(S_T - K) dS_T, \tag{1}$$

$$P(t, T, K) = e^{-r(T-t)} \int_{-\infty}^K w(S_T)(K - S_T) dS_T, \tag{2}$$

⁴ Such analytic tractability is not necessary. Monte Carlo methods could also be used to generate the PDF for intractable stochastic processes.

where $C(t, T, K)$ and $P(t, T, K)$ are the prices of European calls and puts observed at time t having expiries at T and strike prices of K , r is the riskless rate of interest, and $w(S_T)$ is the risk-neutral probability density function for the value of the underlying asset S at time T . Parametric approximating function methods assume that $w(S_T)$ has a particular functional form, chosen to allow for a variety of possible shapes. Parameter values are found by minimizing some function of the fitted price errors. Examples of the approximating functions that have been used include: mixtures of lognormals, developed by Melick and Thomas (1997);⁵ Hermite polynomials, developed by Madan and Milne (1994); and a Burr III distribution, used by Sherrick et al. (1996). Alternatively, non-parametric methods can be used. Examples include the kernel estimator of Ait-Sahalia and Lo (1998) and maximum entropy methods developed by Buchen and Kelly (1996).

The mixtures of lognormals and the related mixtures of normals applied to the logprice are the most widely used methods for estimating implied PDFs. This paper uses a double lognormal as one method for estimating implied PDFs. Details of the method are presented in Appendix A.

The finite difference methods begin with the observation, made by Breeden and Litzenberger (1978), that differentiating Eqs. (1) or (2) once with respect to K produces the cumulative density function (less 1)

$$\frac{\partial C(t, T, K)}{\partial K} = -e^{-r(T-t)} \int_K^{\infty} w(S_T) dS_T \quad (3)$$

while differentiating twice yields the probability density function

$$\frac{\partial^2 C(t, T, K)}{\partial K^2} = e^{-r(T-t)} w(K). \quad (4)$$

Breeden and Litzenberger (1978) show that one can use finite difference methods to approximate Eq. (4) using strikes where bond prices are observed. Neuhaus (1995) applied finite difference methods to Eq. (3) instead.

The smoothed implied volatility smile method was originally developed by Shimko (1993). The method is an approximating function method applied to the volatility smile rather than to the PDF. Option prices are first converted to implied volatilities using the Black–Scholes options pricing formula. A continuous approximating (smoothing) function is then fitted to the implied volatilities and the associated strike prices (on the X -axis). This continuous implied volatility function is converted back into a continuous call price function and then Eq. (4) is used to obtain the PDF. The Black–Scholes model is used here simply as a transformation or mapping from one measurement space to another. The smoothed implied volatility smile method does not as-

⁵ A variant is to model the logprice as a mixture of normals, as was done in Söderlind and Svensson (1997) and Söderlind (1999).

sume that the underlying price process is lognormal. Malz (1997a,b) used delta, $\Delta \equiv \partial C / \partial F$, rather than strike price as the X -axis variable when fitting the implied volatility smile smoothing function. Both Shimko and Malz used low-order polynomial functional forms to fit the implied volatility smile. Campa et al. (1997) introduced the use of smoothing splines to fit the implied volatility function; in their case as a function of the strike price. The second method examined in this paper is a variant of the smoothed implied volatility smile method developed by Panigirtzoglou at the Bank of England, combining the innovations of Malz (1997a,b) and Campa et al. (1997). The method uses a natural spline, as described in Appendix A, to fit Black–Scholes implied volatilities as a function of the deltas of the options in the sample.

A number of papers have compared different implied PDF estimation methods. Campa et al. (1997) compared binomial tree, smoothed implied volatility smile and mixtures of lognormal methods. Comparing various moments of the implied distributions they concluded that all methods produced similar results. They chose to use the binomial tree method in their subsequent analysis. Coutant et al. (1999) compared single lognormal, mixtures of lognormals, Hermite polynomials and maximum entropy methods. Again results were broadly similar, although they noted that the maximum entropy method ran into convergence problems. They chose to use the Hermite polynomial approach in their subsequent analysis. Hördahl (1999) compared implied PDFs derived from the Longstaff–Schwartz stochastic process with PDFs derived using the double-lognormal method. He concluded that the PDFs implied by the two methods were similar and therefore the Longstaff–Schwartz stochastic process method could reliably be used where options data was non-existent. McManus (1999) compared two stochastic process methods, using Black and jump diffusion processes, and four approximating function methods, double lognormal, fourth and sixth order Hermite polynomials, and maximum entropy. Using comparisons of in-sample fit, he concluded that the double-lognormal method was best. Sherrick et al. (1996) compared two PDF approximating function approaches using double-lognormal and Burr III functions. Based on in-sample goodness-of-fit, they concluded that the Burr III approximating function produced the better results.

Like all statistics estimated from finite data samples, implied PDFs and their summary statistics are point estimates, subject to estimation error. However, while many papers have estimated and interpreted implied PDFs, surprisingly few have considered the reliability of estimated implied PDFs and their associated summary statistics. Two methods have been used in previous papers to examine the stability of implied PDFs: working with the parameter variance–covariance matrix and perturbing pseudo-prices generated from known PDFs.

Söderlind and Svensson (1997) assumed that the distribution of estimated parameters was multivariate normal. The PDF confidence intervals were derived analytically using the delta method applied to the heteroskedastic-consistent

estimator. Melick and Thomas (1998) used the Hessian at the maximum likelihood solution as the estimated parameter variance–covariance matrix, again assuming the parameters were multivariate normals. They then used a Monte Carlo simulation to randomly perturb the parameters, recomputing the implied PDF for each simulation. Both papers applied their methods to a single cross-section of option prices, which were then analyzed visually by plotting the value of the PDF and the estimated 5–95% confidence intervals.⁶ Melick and Thomas also used a second method for obtaining the distribution of the implied PDF. This was to bootstrap their original sample of option prices and re-estimate the PDF for each resampling.

Both Söderlind and Svensson (1997) and Melick and Thomas (1998) found that confidence intervals based on the theoretical distributions of the parameters at the solution appeared to be quite narrow. However, when Melick and Thomas resampled the data the confidence intervals were much wider. This disparity suggests that the assumptions underlying the maximum likelihood estimation of the implied PDFs were perhaps violated in some way.

Söderlind (1999) and Cooper (1999) both began with known PDFs. The PDF was used to generate fitted prices, which were then perturbed. The resulting pseudo-prices were then used to estimate the implied PDF. Söderlind (1999) estimated implied PDFs from actual prices and then applied Monte Carlo methods to the fitted option prices. Two error distributions were examined: in the first experiment Söderlind used normally distributed perturbations with variance equal to the observed variance of the actual fitted price errors; in the second experiment Söderlind resampled from the actual fitted price errors. To examine the resulting distributions Söderlind plotted the time series of means, 5th and 95th percentiles of the distribution each day for five months, and four different contract types. The confidence intervals were deemed to be narrow.⁷ Cooper (1999) generates PDFs from an assumed Heston stochastic volatility process and then generates pseudo-prices from the PDFs. Cooper then applies the test methodology developed in this paper to the pseudo-prices.

Advocates of the pseudo-prices approach used by Söderlind (1999) and Cooper (1999) argue that by beginning with a known PDF one can compare the fitted implied PDFs to the ‘true’ PDF to examine how well the estimated PDFs fit the original PDFs. This is correct, but may be of limited usefulness.

⁶ It is somewhat difficult to interpret these error bands. For each value of X the confidence intervals represent the confidence band for the PDF at that single point. However, taken together the lower bound necessarily integrates to less than unity, and the upper bound integrates to more than unity. Thus, unlike confidence intervals for an estimated parameter, which represent possible values for that parameter, the confidence intervals that Söderlind and Svensson (1997) and Melick and Thomas (1998) estimate do not represent possible PDFs.

⁷ The confidence intervals were 2–6% wide for interest rates that vary \approx 2–3% over the entire sample period. Others might reasonably consider this range to be wide.

Goodness-of-fit results may not be generalized to PDFs outside the set examined. A double-lognormal implied PDF estimation method may do well when the assumed PDF has a double-lognormal functional form, but may do less well when the assumed PDF is another distribution. Neither Söderlind nor Cooper consider this issue. Since we cannot know the true distribution underlying actual option prices, it is difficult to extrapolate from such experiments to practical applications.

The robustness results of both parameter variance–covariance matrix and pseudo-price approaches are apt to be misleading for another reason. Stability of an estimated PDF has two components: the theoretical stability at the solution, and the stability of the convergence to a solution. Söderlind and Svensson (1997) and Melick and Thomas (1998) examined only the stability at the solution. These studies rely on assumed distributions for estimated parameters and on estimated variance–covariance matrices. This approach is open to the criticism that actual parameter distributions may be far different from the assumed distribution.

The other component of stability is the stability of the convergence to the original solution. Methods such as Söderlind (1999) and Cooper (1999) that use data created from idealized PDFs – either fitted values or simulated from assumed stochastic processes – are imposing a degree of smoothness in the simulated data that may not be congruent with reality. Perturbing fitted, rather than actual, prices may result in quite different convergence behavior of the optimizing algorithm. Actual fitted-price errors may be larger than the small perturbations used in Cooper (1999). Söderlind (1999), by resampling from actual fitted price errors, offers a somewhat better methodology. However, where actual fitted price errors are not homoskedastic and independent across strikes, his method may still mislead. In both cases, the convergence of the optimization may be influenced by the well-behaved nature of the assumed functional form of the PDFs used to generate pseudo-prices.

Only by perturbing actual option prices can we examine the robustness of estimated implied PDFs in an environment that approximates the real world. This is the first paper to offer systematic comparison of the absolute and relative robustness of implied PDF estimation methods to measurement errors in actual option prices.⁸

2.2. Sources of error in option prices

The prices used as inputs for estimating implied PDFs are subject to various errors that cause the observed prices to deviate from those we expect would

⁸ The resampling approach of Melick and Thomas (1998) is a plausible alternative where there are sufficient usable strikes available in each cross-section of prices. Melick and Thomas limited their study to one estimation method and a single cross-section of prices.

obtain in a frictionless world, the world envisioned in the models we invert in order to estimate distributions from prices. These include:

- Data errors – mistakes in the recording and reporting of prices.
- Non-synchronicity – arising from the need to use multiple simultaneous prices (option and underlying values) as inputs to the model.
- Liquidity premia – arising from the potential impact of differential liquidity on prices.
- Discreteness – arising from quoting, trading, and reporting of prices in discrete increments.

It is frequently possible to obtain evidence suggestive of pricing errors, though it is not always possible to determine whether there is in fact an error, or of what type. Suspicious circumstances would include a series of option prices which violate basic arbitrage restrictions such as monotonicity (call prices should decrease as exercise prices increase) or convexity (prices of option triplets should be convex in their exercise prices). Another basic no-arbitrage relation is put–call parity. This may be verified for individual strikes using observed values for the underlying asset’s price and the riskless interest rate to test put–call parity for each pair of puts and strikes. Alternatively, if some doubt exists as to the appropriate values for the underlying asset and risk-free rate (Treasury bills may not be a good proxy), a cross-section of puts and calls may be tested simultaneously by finding the underlying price and interest rate values that minimize put–call parity violations across all put–call pairs, and then examining the magnitude of violations given these ‘best fit’ values. When violations of these no-arbitrage restrictions occur, it is unclear whether it is due to data errors, non-synchronicity, or liquidity premia.

The data used in this study consist of settlement prices, which are used to mark positions to market at the end of each day’s trading. Settlement prices are set by the exchange at the end of trading. However, as most option strikes trade infrequently and with great variations in time-of-last-trade, the market information used by the exchange when setting settlement prices is likely to be non-synchronous. Unless the LIFFE actively corrects for non-synchronicity, the problem will be transferred to simultaneously determined settlement prices.

To the degree that liquidity is reflected in options prices, it represents a misspecification of the model that we use to infer unobservables such as implied volatility and PDFs from option prices. There is abundant evidence of differential liquidity across options with different strikes for the same expiry. Unfortunately, there is no option pricing model (that we are aware of) that incorporates liquidity into pricing equations. Even if there were, liquidity is time varying and difficult to measure. Thus, the potential impact of differential liquidity on the values derived from options prices is a currently unresolved problem. The problem can however be mitigated by using only the most liquid

strikes – implicitly assuming that there is no premium for liquid options, only discounts for illiquid ones. Doing this has the added advantage of reducing the potential severity of non-synchronicity problems. This approach is practical when computing implied volatilities, when we are interested in representative values for time-series application, rather than the cross-section of implied volatilities. However, when computing implied PDFs, restricting the estimation to the most liquid issues limits the range of available strikes and thus, to the extent that there is information in the illiquid strike quotes, limiting the information incorporated into the estimated implied PDF. Furthermore, as only the four or five nearest-the-money strikes trade reasonably often, restricting our sample to these few would preclude application of implied PDF estimation methods with more than four or five parameters.

It is worth noting that option prices can provide information about the underlying density function only at their strike prices. The shape of the density function between strikes may be constructed by smoothing. If the strikes are not too widely spaced and the PDF not too ill-behaved, this smoothness assumption is likely to be innocuous. However, a cross-section of options can only tell us the total probability mass above the highest strike and below the lowest strike, and that imperfectly. The shape of the tails beyond the range of included strikes is entirely an artifact of the PDF estimation method used. Unfortunately, estimates of higher moments such as skewness and kurtosis are sensitive to small variations in the tails of the distribution. It is thus desirable to use as wide a range of strikes as possible so as to reduce the reliance on unverifiable assumptions about the functional form of the PDF.

The discreteness with which prices are quoted imposes an irreducible level of uncertainty as to the underlying ‘true’ or equilibrium price of an option. Even if no data errors occur in the reporting process and there are no non-synchronicity and liquidity errors, it remains the case that we cannot know to an accuracy of less than one half a tick what price the option would have traded at if prices were quoted on a continuum of positive real numbers.

2.3. Weighting

In fitting an implied PDF, regardless of the method used, the objective is to minimize some function of the distance between the observed call and put prices, C_i , and P_i , $i = 1, \dots, N$, and the fitted prices derived from the estimated PDF, \widehat{C}_i and \widehat{P}_i . In a maximum likelihood framework, where the errors attached to the observed prices are assumed to be normally distributed with mean zero and variances η_i^2 , we would have

$$\min_{\Phi} \sum_{i=1}^{N_c} \frac{(C_i - \widehat{C}_i(\Phi))^2}{\eta_i^2} + \sum_{i=N_c+1}^{N_c+N_p} \frac{(P_i - \widehat{P}_i(\Phi))^2}{\eta_i^2}, \tag{5}$$

where Φ is the vector of parameters that define the fitted prices, including the PDF, and N_c and N_p are respectively the numbers of call and put prices to be fitted. Defining $w_i \equiv 1/\eta_i^2$ we see that this objective function is just weighted least squares:

$$\min_{\Phi} \sum_{i=1}^{N_c} w_i (C_i - \widehat{C}_i(\Phi))^2 + \sum_{i=N_c+1}^{N_c+N_p} w_i (P_i - \widehat{P}_i(\Phi))^2. \quad (6)$$

While in this paper we do not use maximum likelihood or impose the normality assumption, we do retain the weighted squared fitted price error loss function. Unfortunately, η_i is not known and must be inferred. The determination of η_i depends in turn on which sources of error in quoted or fitted prices we wish to consider. Errors in the inputs to the fitted price computations, as well as errors in the observed prices, all contribute to η_i .

Fitted prices are functions of the strike price, K , the underlying asset's current price, S , the time to expiry, τ , the riskless rate r , and the risk-neutral distribution of values of the underlying asset at expiry, the PDF. The strike price is a contractual parameter and is known with certainty, as is the expiry date. Uncertainty regarding precise time-of-quote is generally a tiny fraction of time-to-expiry (minutes or hours rather than weeks or months) and option prices are not sensitive to small changes in time-to-expiry. Thus, for all practical purposes we may consider time-to-expiry as known with certainty. The riskless rate is more problematic. Proxies, such as an equivalent-maturity T-Bill rate, may be affected by market microstructure factors unrelated to the rates at which market participants can borrow and lend (see Duffee, 1996) and other money market rates may embed non-equivalent default premia. Fortunately, like time-to-expiry, small variations in discount rates have a negligible effect on option prices.⁹

Uncertainty regarding the value of the underlying asset is an important factor in determining the uncertainty regarding fitted option values. Changes in the value of the underlying asset have a large impact on the theoretical value of all but deep out-of-the-money options. Uncertainty regarding the value of the underlying asset arises from uncertainty as to the exact time at which the price of the option was determined and hence which intra-day value of the underlying asset should be used. Lastly, uncertainty regarding the probability distribution of the value of the underlying asset is also an important component of the uncertainty regarding the fitted option values.

Thus, in determining η_i we should ideally consider three potential sources of error: non-synchronicity, uncertainty regarding the distribution of future values of the underlying asset, and uncertainty regarding the actual equilibrium

⁹ For options such as short sterling that have a pay-at-exercise feature, both time-to-expiry and the riskless rate drop out of the pricing equations.

price arising from quote discreteness. We may safely ignore other factors. Unfortunately, there is no simple or generally accepted manner for modeling all of these effects.

Errors arising from non-synchronicity affect the values of the underlying asset, S . These in turn are related to the call (and put) prices through delta, $\Delta \equiv \partial C / \partial S$. Thus, an error of ϵ_S in measuring the price of the underlying asset results in an error for the option price, ϵ_{C_i} of

$$\epsilon_{C_i} = \frac{\partial C_i}{\partial S} \epsilon_S = \Delta_i \epsilon_S. \tag{7}$$

The value of Δ increases from zero for deep out-of-the-money options to ≈ 0.5 for at-the-money options and then 1.0 for deep in-the-money options. Translating uncertainty about the current value of the underlying asset, which is the same for all strikes, into uncertainty about the option price leads to inverse- Δ weighting ($w_i = 1/\Delta_i^2$). However, this has the disadvantage of the weights becoming excessively large as $\Delta \rightarrow 0$ for deep out-of-the-money options, which are also the most illiquid.

Errors arising from uncertainty about the distribution of futures values of the underlying asset relate the unknown PDF directly to the option prices. In the context of the Black–Scholes pricing model, the uncertainty concerning the PDF lies only in the unobservable volatility parameter σ , as it is assumed that the other parameters are observable and that the functional form of the distribution is lognormal. The relation between volatility, σ , and call (and put) price is termed v (vega),

$$v \equiv \frac{\partial C}{\partial \sigma}. \tag{8}$$

The value of v approaches zero for deep out-of-the-money and in-the-money options and reaches a maximum for at-the-money options. This is because the value of far-away-from-the-money options is almost entirely determined by the intrinsic value, and the time value, which depends on σ , is vanishingly small. If we assume that uncertainty regarding implied volatility is the same across strikes, translating this (homoskedastic) uncertainty into uncertainty about the option price leads to inverse- v weighting ($w_i = 1/v_i^2$). Equal weighting ($w_i = 1$) when fitting smoothing functions to implied volatilities is the same thing. Both weightings produce the non-sensical result of giving the greatest weight to options with the lowest v , those farthest-away-from-the-money, which are also the least liquid contracts and most susceptible to non-synchronicity errors. The alternative of v weighting ($w_i = v_i^2$) is intuitively appealing as it would place the most weight on near-the-money options, and corresponding lesser weight on away-from-the-money options. Nonetheless, without a known or assumed structure for option pricing errors, v -weighting is ad hoc.

Equal weighting of fitted price errors is appropriate where the sources of price measurement error are homoskedastic. In this paper we focus on price errors resulting from the discrete tick size, which is the same for all options regardless of moneyness. Our maintained hypothesis is that the discreteness with which options are quoted imposes a homoskedastic uncertainty on the observed prices unrelated to the determinants of their fundamental value. For this reason we use equal weighting in the DLN method and set $w_i = 1, \forall i$. In the SML method we are minimizing not fitted price errors, but fitted implied volatility errors. The ν -weighting of the fitted implied volatility errors is equivalent to equally weighting the fitted price errors of the options from which the target implied volatilities are derived under the Black–Scholes model. Fig. 1 shows this graphically by plotting, for one contract, the option prices with error bars corresponding to plus and minus one half of a tick (difficult to see due to the small tick size) together with the corresponding implied volatilities for the original option prices and error bars for the implied volatilities for option prices one half tick above and below the original prices. The figure shows how small equal-sized price perturbations can produce variable-size implied volatility perturbations with the size of the change increasing as strikes move further from at-the-money. Thus, the SML estimation method uses ν -weighting when fitting the volatility smile. However, when we compare the two methods we do so on the basis of their equally weighted fitted price errors.

It would be difficult to devise an ideal weighting scheme that was not completely ad hoc; one that could account for all sources of pricing error. Such a weighting scheme would require an asymmetric function that placed greatest weight on near-the-money options and decreasing weight on away-from-the-money options, but with weights falling off faster for in-the-money options than for out-of-the-money options. The choice of weighting scheme is likely to be less important if fitted price errors are small. Fortunately, about 90% of the fitted price errors in our estimations are less than one half of a tick. To ensure that the full sample results are not dependent on choice of weighting scheme we test several weighting schemes using a subset of the data used in this study.

2.4. Mean-forward price equality

Option theory dictates that the mean of the risk-neutral PDF should equal the currently observed forward price of the underlying asset. In the DLN procedure, it is possible to impose the forward–mean equality as a constraint using the futures price as a proxy for the forward price, thus reducing the free parameters from five to four. However, this theoretical relation is not required by the mathematics underlying the DLN method, it follows from related, but

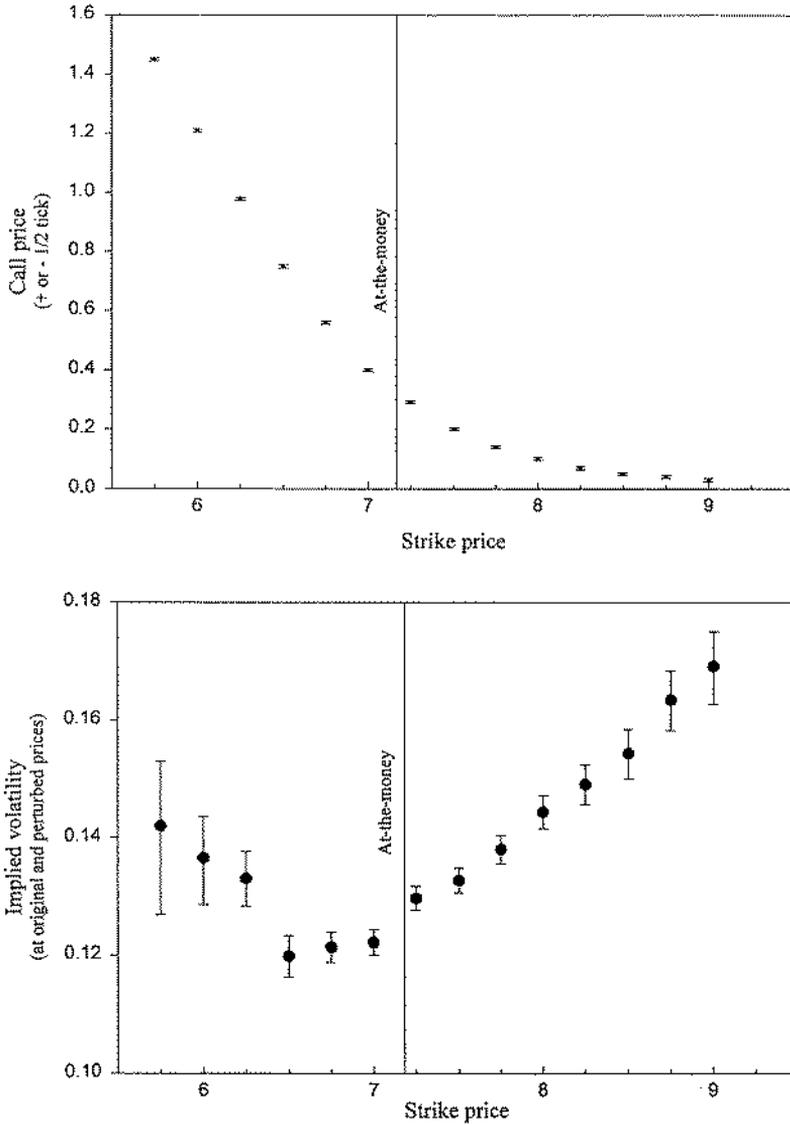


Fig. 1. The effect of price perturbations on implied volatility (note: tick size = 0.01).

separate, arbitrage arguments. This constraint will usually be binding and will degrade the goodness-of-fit. Not imposing the constraint allows us to see how closely the estimated PDF conforms to the theoretical restriction on the mean;

in effect, how well the underlying conditions for no-arbitrage hold. The choice is a matter of taste.¹⁰

By construction, the SML method prices a zero-strike call to be equal to the value of the underlying asset. The value of a zero-strike call is just the expected value of the underlying asset at expiry.¹¹ For the short sterling options used in this study, the underlying is the futures price. The FTSE 100 options, while options on an index, can be thought of as an option on the futures on the index, as the futures contract expires at the same time as the option and so will have the same value as the index at option expiration. Thus, the SML method naturally enforces the forward–mean constraint, abstracting from forward–futures differences.

3. Stability test methodology

The stability of an estimated function, as used in this paper, is a measure of how much estimates are likely to be affected by data imperfections or computational problems. There are several methods of assessing stability. For simple linear models we can examine the conditioning of the data matrix. It is well known that an ill-conditioned problem leads to unstable estimates.¹² However, no simple equivalent of the condition number of a data matrix exists for more complicated estimation procedures such as various methods of estimating the PDFs implicit in a cross-section of option prices.

In this paper, we therefore rely on repeated-estimation methods. Bootstrap and jack-knife methods are one possibility already discussed. In the present context, this would take the form of repeatedly selecting a subset (with or without replacement) of the option prices available in a cross-section, estimating the PDF on this sample, and repeating the procedure numerous times to build a distribution of estimated PDFs. These methods work best where there is a large number of option strikes from which to sample. In practice, this is not always the case.

An alternative re-estimation method is to slightly perturb the inputs and then re-estimate. This can be done any number of times with even a small number of strikes. Perturbing the data simulates the effects of measurement error between the ‘true’ option value representing the underlying economic

¹⁰ In an earlier version of this paper, using approximately 175 option cross-sections and 30 Monte Carlo simulations per cross-section, we imposed the mean–forward constraint. The results obtained were not qualitatively different from those we present in this paper.

¹¹ This is true for options on futures and for pay-upon-exercise deferred premium options. For normal options on positive-investment underlying assets such as stocks, the value of a zero-strike call would equal the present value of the expected value of the underlying asset at expiry.

¹² See Belsley et al. (1980) for a discussion.

factors we seek to uncover (the market-clearing risk-neutral distribution of future values of the underlying asset) and the observed quotes that add to this information noise from the various sources of error discussed above. Furthermore, if the perturbations are calibrated to the size of the possible measurement errors, the distribution of simulated PDF summary statistics provides a confidence region for assessing the summary statistics, and their period-to-period changes, estimated from the original unperturbed data.

In this paper, we introduce the price-perturbation method for assessing the stability of PDFs estimated from options prices. We apply this technique to two methods for estimating PDFs and to options on two important underlying assets.

3.1. Data

The data used in this study are the daily settlement prices published by LIFFE. These prices are based on quotes and transactions during the day and are used to mark options and futures positions to market. Two contract types are used to ensure the results are not contract specific. These are the FTSE 100 index options and the short sterling futures options. Details of the contracts are given in Appendix B. Fig. 2 presents the distribution of daily trading for FTSE 100 and short sterling puts and calls during the 1997 sample period as a function of moneyness. The Tukey (box-and-whiskers) plot corresponding to a moneyness index of 1 is the distribution of daily number of trades for strikes at or immediately in-the-money, measured against the contemporaneous (intraday) value of the underlying asset. The extreme dots denote the 5th and 95th percentiles of the distributes; the “T”s indicate the 10th and 90th percentiles; the box itself spans the 25th to 75th percentiles with the median indicated by the horizontal line within the box. Thus, immediately in-the-money FTSE 100 call strikes trade at least once on 90% of the sample days, with a median of 5 trades per day. Fig. 2 shows the large differences in liquidity between the two contracts used in this study and between in- and out-of-the-money FTSE 100 contracts. Only 4 short sterling call strikes and 5 short sterling put strikes trade on at least 25% of days. The FTSE 100 trading activity shows a strong market interest in out-of-the-money puts and calls, while in-the-money trading drops off quickly as strikes move away from the money. This pattern is less evident for short sterling contracts, with only a slight hint of deep out-of-the-money put trading. Summary statistics for the full sample and final sample (following various filters described below) are presented in Table 1. The original data set covered all observed option cross-sections (a set of put and call prices with identical expiries observed on a given quotation date) for all available expiries for these two contract types during 1997: 1506 FTSE 100 option cross-sections and 1000 short sterling option cross-sections are used.

For short sterling options, put–call parity always holds exactly, so puts and calls for the same strike are redundant. So, for the short sterling portions of the

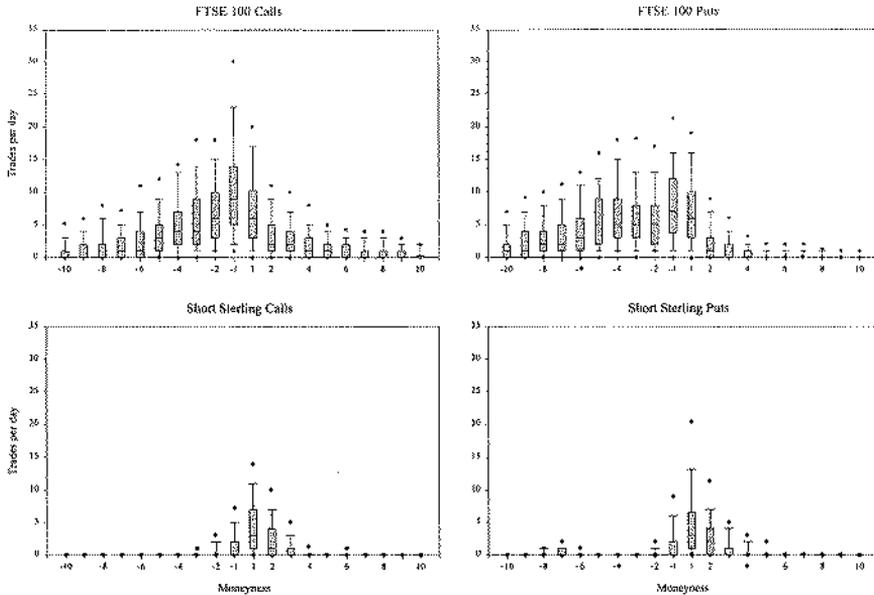


Fig. 2. Distribution of daily trading frequency of options contracts by degree of moneyness (*note*: moneyness value of 1 indicates strike at or immediately in-the-money; 2 indicates next in-the-money strike; -1 is first strike strictly out-of-the-money, etc.).

study we use only call prices. For FTSE 100 options put–call parity does not always hold and so FTSE 100 put and call prices are not redundant. Rather than include both, we seek to use the most liquid strikes. A related investigation by Bliss and Xu at the Bank of England looked at daily trading and quotation activity for both short sterling and FTSE 100 options contracts as a function of moneyness. Except for the four or five nearest-the-money strikes and with expirations of less than six months, most option strikes are not quoted or traded on most dates. That study also confirmed the general understanding that out-of-the-money calls tend to be more liquid than puts of the same strike, and similarly for out-of-the-money puts and in-the-money calls. Thus we use only out-of-the-money options in the FTSE 100 portions of this study.

The data were filtered to exclude option cross-sections with less than seven days to expiry or less than five ‘good’ option strikes. A minimum of five strikes is required to estimate the five parameter double-lognormal function. Good strikes are defined as those with positive put and call prices¹³ for which it is possible to compute a Black–Scholes implied volatility that is strictly greater

¹³ Away-from-the-money short sterling options are frequently quoted at their intrinsic value regardless of time to expiry.

Table 1
Summary statistics of option cross-section samples

| Original sample | FTSE 100 | Short sterling |
|--------------------------------------|-------------|----------------|
| Number of cross-sections | 1,506 | 1,000 |
| Strikes per cross-section | | |
| Range | 13–55 | 13–22 |
| Average | 26.4 | 17.6 |
| Time to expiry (years) | | |
| Range | 0.004–0.996 | 0.004–0.996 |
| Average | 0.374 | 0.501 |
| <i>Filter losses</i> | | |
| Short expiry or too few good strikes | 60 | 206 |
| DLN failures on unperturbed data | 8 | 21 |
| Simulation failures | 5 | 5 |
| DLN spikes | 18 | 57 |
| <i>Final sample</i> | | |
| Number of cross-sections | 1,415 | 721 |
| Strikes per cross-section | | |
| Range | 13–55 | 5–18 |
| Average | 26.1 | 10.8 |
| Time to expiry (years) | | |
| Range | 0.024–0.996 | 0.163–0.996 |
| Average | 0.389 | 0.628 |

than zero. These two filters reduced the sample sizes to 1446 FTSE 100 option cross-sections and 794 short sterling option cross-sections.

3.2. Comparability of estimation methods

When comparing implied PDF estimation methods it is important to ensure that inputs of the two methods are as similar as possible. The goodness-of-fit of the SML method can be controlled while that of the DLN method cannot. The DLN method sometimes fails entirely, thus producing no output for option cross-sections for which the SML method is successful. In ensuring comparability of results we adjust for both factors.

There is a natural tension between goodness-of-fit and stability. While not invariably true, one expects a method that fits the data accurately to be less stable to perturbations of the data. In this paper our focus is on stability and so we abstract from goodness-of-fit considerations. The SML method involves a smoothing parameter, λ , which controls the trade-off between smoothness and goodness-of-fit. The DLN method has no such degree of control. In this paper, λ was selected so that the two goodness-of-fit measures, as measured by the mean squared fitted option price error across all option cross-sections and

strikes (those included in the estimations), were approximately equal in the unperturbed data sets. In this way, we are able to compare the stability of two PDF estimation methods that fit the data equally well. In practice, the λ required to accomplish this is too ‘loose’ and occasionally produces improbably contorted PDFs, just as the DLN method sometimes produces improbably spiked PDFs. A tighter λ , while still fitting the option prices to well within a half tick in the vast majority of cases, will produce PDFs that are more plausible (see Appendix A.2). This smoother SML PDF would be more stable than the SML PDF calibrated to match the goodness-of-fit of the DLN PDFs.

The DLN method failed to converge to a solution on the original unperturbed data for a number of option cross-sections. This never occurred with the SML method. The middle panel of Table 1 tracks the resulting adjustments to the samples. To ensure comparability of tests across PDF estimation methods, we excluded option cross-sections for which it was not possible to compute both DLN and SML solutions. This reduced the sample to 1438 FTSE 100 option cross-sections and 783 short sterling option cross-sections. Similar convergence failures occurred during price-perturbation simulations. Again, failures in either method resulted in the option cross-section being excluded from the sample, reducing the sample sizes to 1433 FTSE 100 option cross-sections and 778 short sterling option cross-sections. Deleting cross-sections where the DLN solutions produced evidence of a spiked PDF (when volatility constraint was binding or when the mode of the PDF had an extremely high value) reduced the final samples to 1415 FTSE 100 option cross-sections and 721 short sterling option cross-sections. The quotation dates and times to expiry of the surviving option cross-sections are plotted in Fig. 3. The short-expiry short sterling cross-sections invariably had too few usable strikes. This problem occurred less frequently with FTSE 100 options.

3.3. Monte Carlo simulations

To test the relative effects of measurement error on the stability of estimated PDFs, we take observed option prices, perturb them and re-estimate the PDFs repeatedly. To obtain each simulated price we add a uniformly distributed random price perturbation of between plus and minus one half of the contract’s tick size. The tick size for the short sterling contract is 0.01, and for the FTSE 100 contract, 0.5. As the simulated prices lie within a half tick of the original data they are observationally equivalent to the original data.¹⁴ For each set of simulated prices, we estimate both DLN and SML PDFs. This process is then repeated 100 times for each option cross-section.

¹⁴ For example, short sterling option prices of 1.796 and 1.804 would both be quoted as 1.800 and hence are ‘observationally equivalent’.

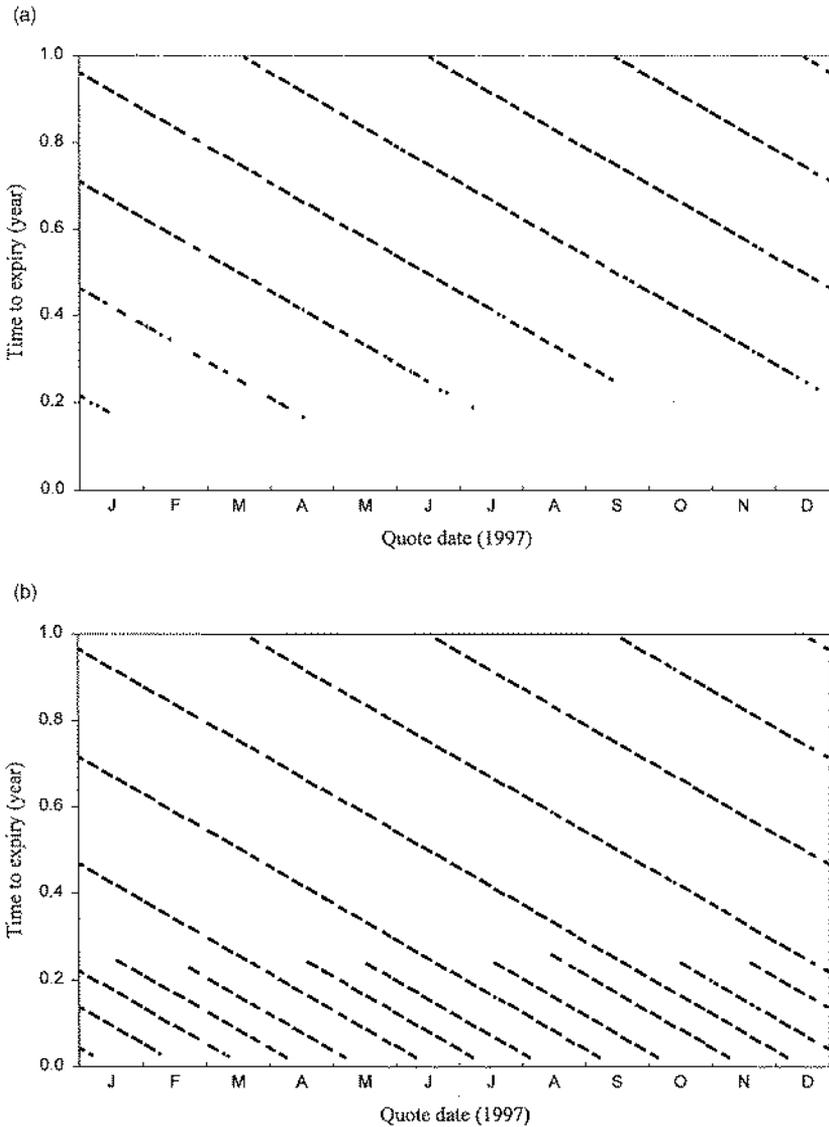


Fig. 3. (a) Short sterling option cross-sections, (b) FTSE 100 option cross-sections.

Numerous DLN and some SML failures occurred during the price-perturbation simulations. DLN failures were due either to convergence failures or degeneracy resulting in PDFs not integrating to unity. SML failures occurred when Black–Scholes implied volatilities could not be computed for all of the

perturbed prices in the option cross-section.¹⁵ When failures occurred for either DLN or SML, that set of simulated prices was discarded and another random sample was drawn. If 50 such simulation failures occurred before the target of 100 successful solutions was reached, then the entire option cross-section was discarded from the simulations sample.

The result of this process was a set of 100 DLN PDFs and 100 SML PDFs for each option cross-section, and their associated mean-squared fitted option price errors (VOFs), estimated on identical sets of perturbed prices for each option cross-section. Fig. 4 plots selected results of the perturbation exercise for both SML and DLN methods for a sample contract, presenting the PDFs estimated using the original unperturbed data together with the perturbed-data PDFs producing the highest and lowest values of skewness for each method. It is readily apparent for this contract, at least, that the DLN estimation method produced a wider variation of fitted PDFs and derived statistics over the identical set of input data, than did the SML estimation method.

The perturbed-data PDFs were then filtered to delete instances when the DLN method arrived at a corner solution, usually the lower bound on one of the component lognormal variances (indicative of a possible spike). In the event of a DLN corner solution the corresponding SML PDF was also deleted, thus maintaining exact comparability of the perturbed-data sets used in the subsequent analysis. This reduced the number of usable FTSE 100 simulations from 141,500 to 140,610 and the number of short sterling simulations from 68,000 to 63,611.

The several filters applied to the unperturbed data PDFs and to the simulation results exclude most ill-behaved DLN solutions. Because there are no corresponding problems with the SML method, the filtering favors the DLN method.

It is difficult to compare more than a few PDFs in their entirety (for example by overlaying graphs). Therefore, we analyze the perturbed-price PDFs by examining the distributions of twelve PDF summary statistics. For a number of applications, such as inferring asymmetries of market expectations, estimated PDFs are used as an intermediate step to computing measures of asymmetry or skewness. For such purposes it is the stability of the derived statistic that is of interest. For applications, such as pricing other derivatives, the entire PDF is needed. However, the stability of the moments derived from estimated PDFs, when taken together, provides insight into the stability of the entire PDF. Thus, focusing on implied PDF summary statistics provides a practical method for assessing the absolute and relative stability of estimated PDFs.

This study examines the following PDF summary statistics:

¹⁵ By construction Black–Scholes implied volatilities can be estimated for all the unperturbed prices.

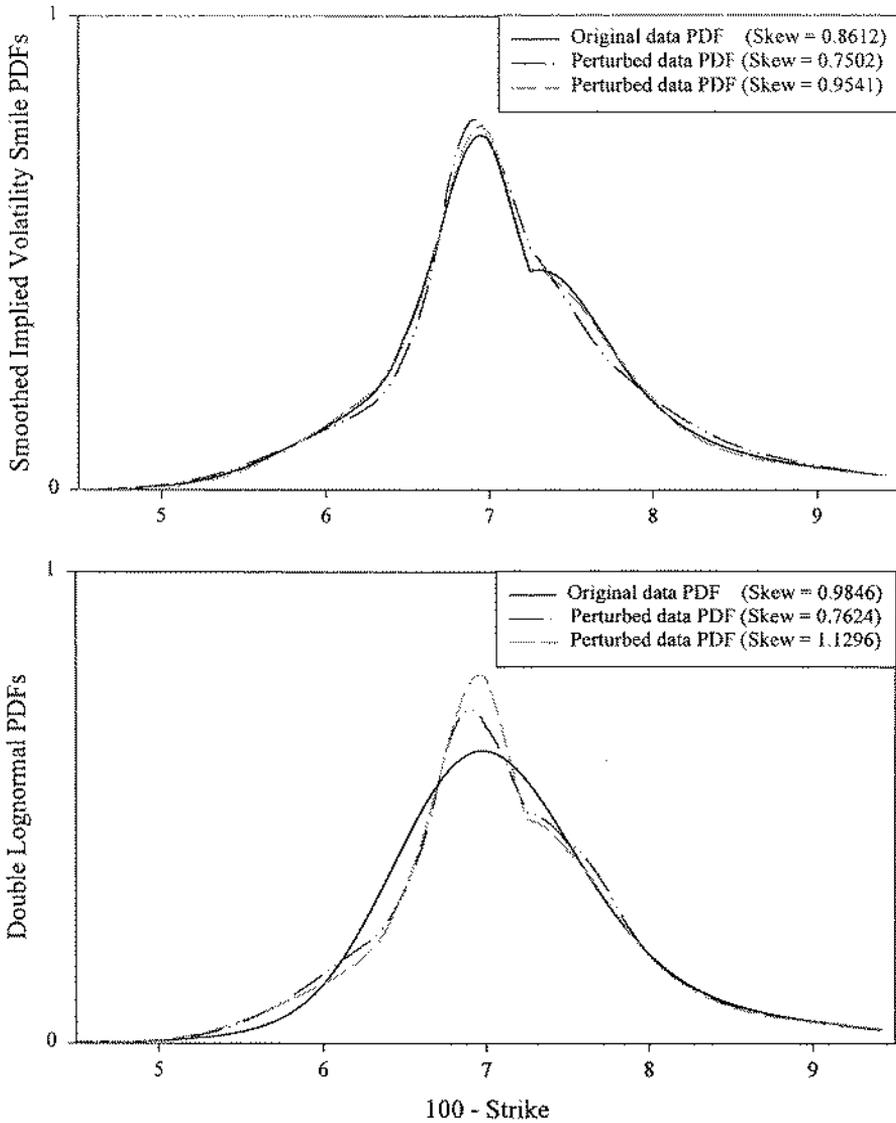


Fig. 4. Original and perturbed-data PDFs estimated from a single short sterling contract (note: contract expiry 17-Sep-97; quoted 3-Mar-97; 14 usable strikes).

- $\hat{\mu}$: Mean.
- $Skew_1$: The skewness coefficient; defined as the third central moment normalized by the cube of the standard deviation:

$$\text{Skew}_1 = \frac{\hat{m}^3}{\hat{\sigma}^3}, \quad (9)$$

where m^3 is the third central moment about the mean. This is the most commonly used measure of skewness.

- *Skew*₂: The Pearson mode-based skewness measure, defined as

$$\text{Skew}_2 = \frac{\hat{\mu} - \widehat{\text{mode}}}{\hat{\sigma}}. \quad (10)$$

- *Skew*₃: The Pearson median-based skewness measure, defined as

$$\text{Skew}_3 = \frac{\hat{\mu} - \widehat{X}_{50}}{\hat{\sigma}}, \quad (11)$$

where X_n is the n th percentile of the PDF, in this case 50th percentile or median.

- *Skew*₄: A measure of asymmetry defined by

$$\text{Skew}_4 = \frac{\widehat{X}_{75} - \widehat{X}_{50}}{\widehat{X}_{50} - \widehat{X}_{25}}. \quad (12)$$

When computing sample statistics, this measure is robust to the presence of outliers.¹⁶ In the context of PDF functions, this measure should be robust to fluctuations in the tails of the distribution, where there is no underlying options data.

- *Kurt*: The kurtosis coefficient; defined as the fourth central moment normalized by the square of the variance:

$$\text{Kurt} = \frac{\hat{m}^4}{\hat{\sigma}^4}, \quad (13)$$

where m^4 is the fourth central moment about the mean.

- X_n : Tail percentiles X_{01} , X_{05} , X_{95} , X_{99} . These are important in risk management.

To compute the above moments we first compute the value of the PDF at 10,000 points spanning a range of values sufficient to ensure that the PDF integrates to approximately unity.¹⁷ We then numerically integrate the appropriate function of the PDF to estimate the moments, numerically integrate the CDF to estimate the percentiles, and find the maximum value of the PDF to estimate the mode. For the simulations, we then compute the deviations of the VOFs and summary statistics from their unperturbed values.

¹⁶ See Barnett and Lewis (1984, p. 81).

¹⁷ If the PDF integrated to less than 0.90, the solution is deemed a failure. This usually occurred for DLN PDFs were the PDF contained a spike, rather than because the range of integration was insufficient. When the PDF integrated to a value between 0.90 and 1.00, the PDF was normalized by dividing by that value.

Table 2
Distribution of PDF summary statistics for unperturbed data

| Summary statistic | Statistic means | | | | Statistic standard deviations | | | |
|-------------------|-----------------|---------|----------------|--------|-------------------------------|--------|----------------|--------|
| | FTSE100 | | Short sterling | | FTSE100 | | Short sterling | |
| | DLN | SML | DLN | SML | DLN | SML | DLN | SML |
| VOF | 0.0191 | 0.0199 | 0.8630 | 0.8863 | 0.1169 | 0.0330 | 0.9329 | 0.8387 |
| $\hat{\mu}$ | 4,752.7 | 4,755.1 | 7.205 | 7.214 | 366.4 | 368.0 | 0.317 | 0.314 |
| $\hat{\sigma}$ | 501.8 | 498.0 | 0.662 | 0.649 | 259.1 | 252.5 | 0.205 | 0.213 |
| Skew ₁ | -0.544 | -0.683 | 0.657 | 0.766 | 0.359 | 0.363 | 0.349 | 0.279 |
| Skew ₂ | -0.136 | -0.239 | 0.172 | 0.202 | 0.241 | 0.175 | 0.295 | 0.204 |
| Skew ₃ | -0.087 | -0.119 | 0.085 | 0.102 | 0.051 | 0.042 | 0.093 | 0.038 |
| Skew ₄ | -0.105 | -0.131 | 0.505 | 0.165 | 0.110 | 0.151 | 3.630 | 0.194 |
| Kurt | 4.345 | 3.962 | 5.178 | 4.626 | 4.734 | 1.579 | 2.542 | 0.878 |
| X_{01} | 3,406.4 | 3,422.3 | 5.791 | 5.897 | 522.6 | 461.1 | 0.498 | 0.500 |
| X_{05} | 3,815.3 | 3,808.6 | 6.231 | 6.260 | 432.2 | 399.3 | 0.404 | 0.415 |
| X_{95} | 5,500.7 | 5,466.7 | 8.367 | 8.411 | 684.1 | 663.6 | 0.493 | 0.502 |
| X_{99} | 5,841.5 | 5,759.5 | 9.161 | 9.166 | 875.4 | 812.1 | 0.703 | 0.701 |
| <i>N</i> | 1,415 | | 721 | | 1,415 | | 721 | |

4. Empirical results

The means and standard deviations of the unperturbed data VOFs and summary statistics for DLN and SML are presented in Table 2. The VOFs, means, standard deviations, and tail percentiles are quite close in their means and, except for the FTSE 100 VOFs, in their standard deviations.¹⁸ Differences in means of the several skewness measures sometimes appear moderately large (for example FTSE 100 Skew₂ and short sterling Skew₄). The standard deviations of the DLN skewness measures are generally larger than the SML skewness measures, in several cases by a factor of two or more. Similarly, the DLN method produces much greater variation in estimated kurtosis than does the SML method of both FTSE 100 and short sterling. However, the mean kurtosis is comparable across methods. The unperturbed data results suggest that the DLN and SML methods are similar in performance.

Fig. 5 plots the DLN results against the SML results for selected statistics. The range of the DLN axes were shortened to exclude outliers. Goodness-of-fit is correlated for the two methods for the short sterling cross-sections, but not for the FTSE 100 cross-sections. There are varying degrees of correlation for the other statistics. It is noteworthy that particularly large values of DLN skewness and kurtosis are not associated with particularly large SML values.

Table 3 presents measures of the distribution of day-to-day changes in the various PDF summary statistics for the unperturbed option cross-sections.

¹⁸ By construction the mean VOFs are comparable.

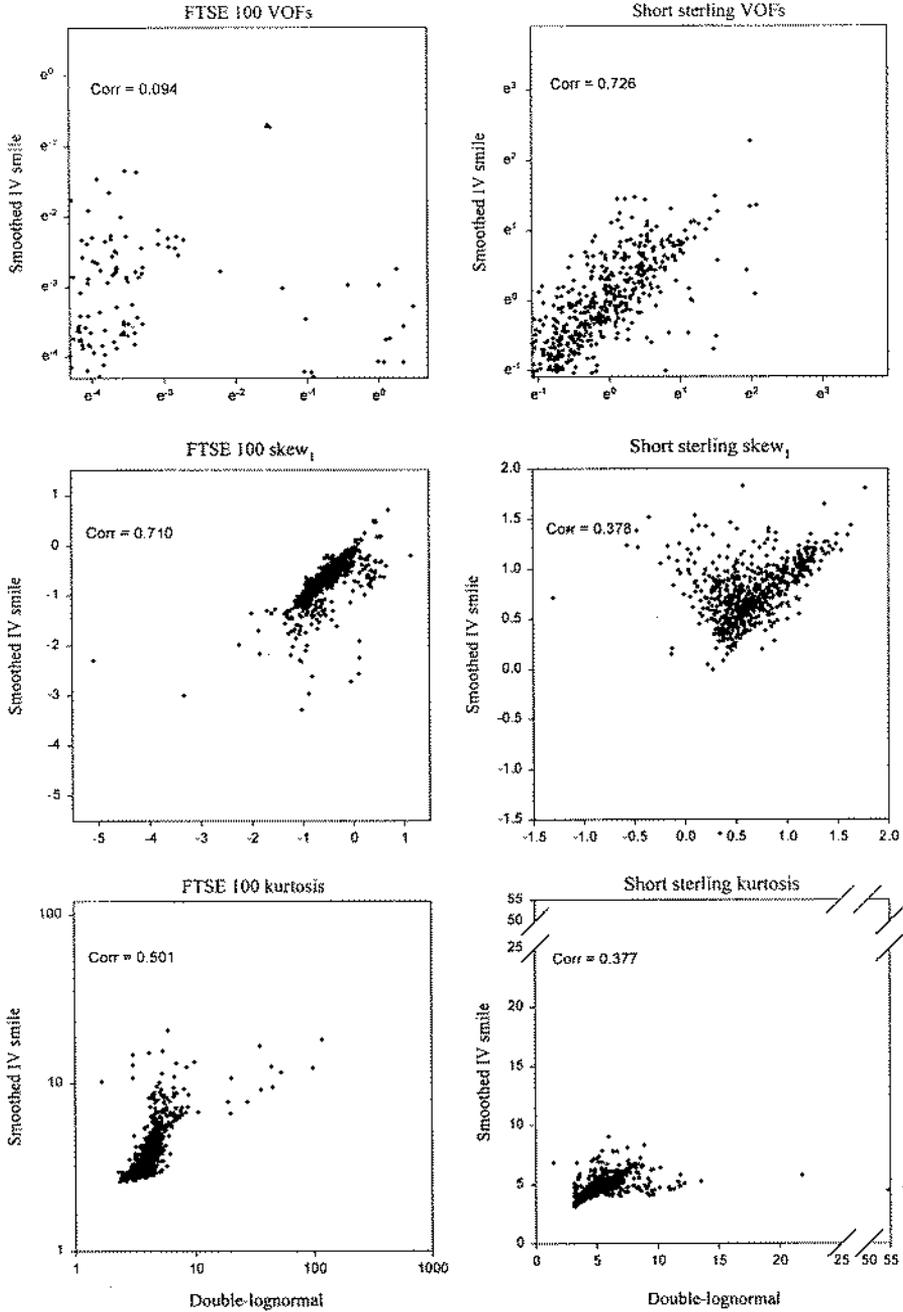


Fig. 5. Comparison of SML and DLN summary statistics.

Table 3
Unperturbed data PDF summary statistics day-to-day changes

| Summary statistic | Mean absolute | | | | Standard deviation | | | |
|-------------------|---------------|-------|----------------|-------|--------------------|-------|----------------|-------|
| | FTSE100 | | Short sterling | | FTSE100 | | Short sterling | |
| | DLN | SML | DLN | SML | DLN | SML | DLN | SML |
| $\hat{\mu}$ | 40.8 | 39.9 | 0.042 | 0.037 | 57.5 | 55.2 | 0.062 | 0.049 |
| $\hat{\sigma}$ | 11.4 | 9.6 | 0.030 | 0.016 | 23.4 | 18.5 | 0.042 | 0.021 |
| Skew ₁ | 0.118 | 0.044 | 0.226 | 0.151 | 0.321 | 0.079 | 0.324 | 0.201 |
| Skew ₂ | 0.088 | 0.041 | 0.132 | 0.162 | 0.311 | 0.068 | 0.391 | 0.231 |
| Skew ₃ | 0.020 | 0.008 | 0.048 | 0.027 | 0.061 | 0.014 | 0.082 | 0.036 |
| Skew ₄ | 0.046 | 0.026 | 0.510 | 0.147 | 0.115 | 0.050 | 3.448 | 0.193 |
| Kurt | 0.680 | 0.126 | 1.360 | 0.538 | 5.352 | 0.313 | 3.406 | 0.753 |
| X_{01} | 80.4 | 48.2 | 0.184 | 0.059 | 159.9 | 73.0 | 0.271 | 0.106 |
| X_{05} | 62.1 | 49.9 | 0.085 | 0.041 | 103.5 | 72.8 | 0.124 | 0.066 |
| X_{95} | 43.9 | 36.9 | 0.069 | 0.055 | 63.8 | 52.9 | 0.095 | 0.078 |
| X_{99} | 74.3 | 39.3 | 0.163 | 0.113 | 133.8 | 59.6 | 0.242 | 0.170 |

| Summary statistic | Inter-quartile range | | | | X_{05} to X_{95} range | | | |
|-------------------|----------------------|-------|----------------|-------|----------------------------|-------|----------------|-------|
| | FTSE100 | | Short sterling | | FTSE100 | | Short sterling | |
| | DLN | SML | DLN | SML | DLN | SML | DLN | SML |
| $\hat{\mu}$ | 59.5 | 59.0 | 0.067 | 0.060 | 182.3 | 180.5 | 0.163 | 0.146 |
| $\hat{\sigma}$ | 10.4 | 8.2 | 0.043 | 0.023 | 51.9 | 43.1 | 0.137 | 0.068 |
| Skew ₁ | 0.063 | 0.046 | 0.302 | 0.229 | 0.514 | 0.209 | 1.108 | 0.678 |
| Skew ₂ | 0.045 | 0.047 | 0.134 | 0.210 | 0.207 | 0.178 | 0.455 | 0.773 |
| Skew ₃ | 0.017 | 0.009 | 0.058 | 0.042 | 0.071 | 0.038 | 0.238 | 0.123 |
| Skew ₄ | 0.034 | 0.026 | 0.131 | 0.230 | 0.182 | 0.119 | 1.165 | 0.645 |
| Kurt | 0.238 | 0.093 | 1.320 | 0.724 | 1.528 | 0.593 | 5.951 | 2.566 |
| X_{01} | 74.2 | 61.1 | 0.228 | 0.035 | 372.9 | 212.1 | 0.902 | 0.394 |
| X_{05} | 76.2 | 65.0 | 0.113 | 0.044 | 267.0 | 220.0 | 0.378 | 0.216 |
| X_{95} | 62.5 | 55.0 | 0.104 | 0.090 | 189.2 | 161.8 | 0.301 | 0.245 |
| X_{99} | 79.3 | 54.6 | 0.185 | 0.078 | 355.5 | 174.9 | 0.774 | 0.558 |

While the average values of these statistics and their standard deviations did not differ systematically across PDF estimation methods, their changes do. In virtually every instance the day-to-day changes in the SML-derived statistics had a smaller dispersion than the corresponding DLN-derived numbers; in some cases by a factor of three or more. Thus while average values do not differ, there is much less short-run stability in the DLN numbers. Changes in the underlying options prices from one day to the next appear to have a larger effect on DLN PDFs than on SML PDFs.

Fig. 6 plots the day-to-day changes in selected DLN statistics against the corresponding SML values. The correlations are weak at best. Thus the DLN and SML methods are apt to give conflicting signals of changes in the underlying distribution. The X_{01} results provide evidence of the hazards of extrapolating the tails of the PDF beyond the range of available data. Recall that

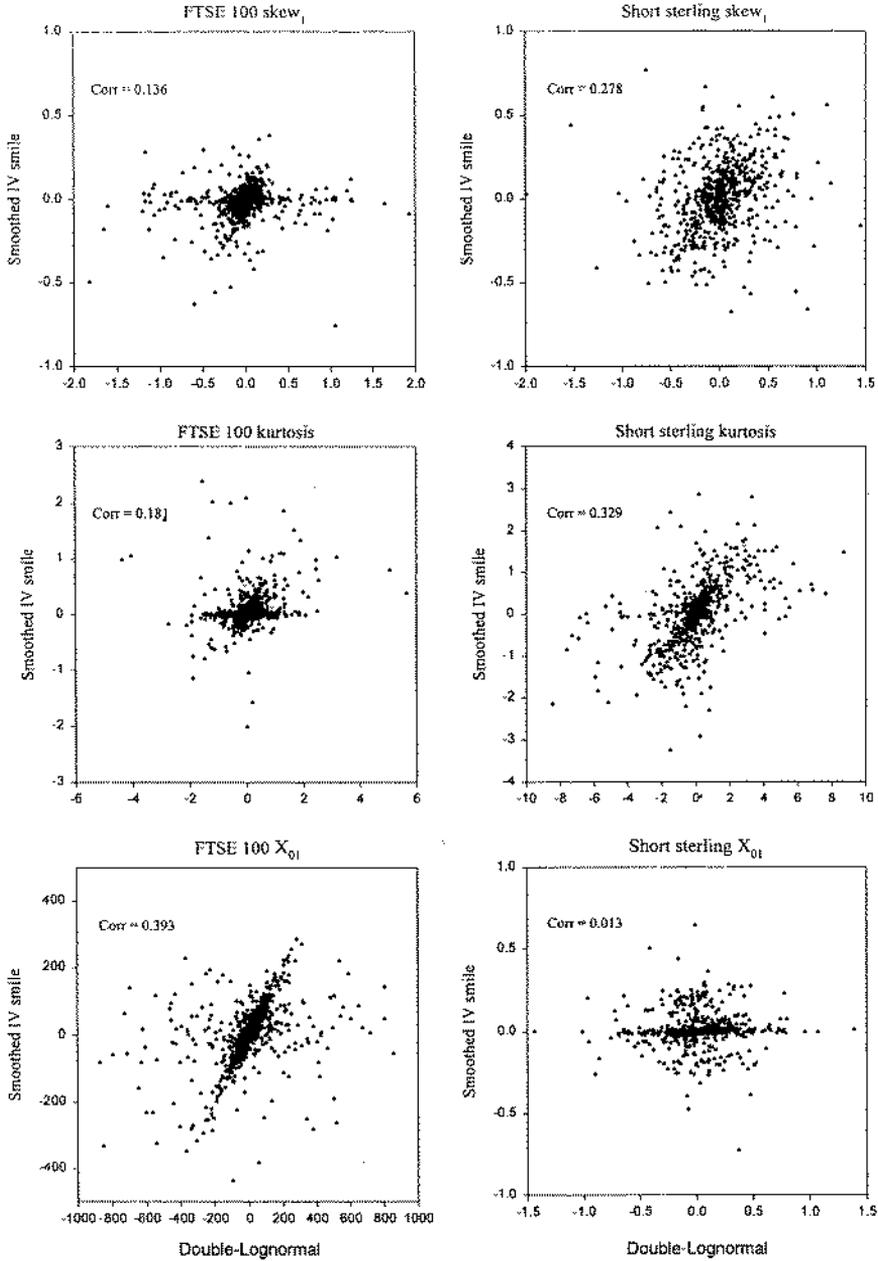


Fig. 6. Comparison of day-to-day changes in SML and DLN summary statistics.

the number and range of usable FTSE 100 strikes in each cross-section is greater than in most short sterling cross-sections. Thus, there is more information about away-from-the-money prices in the FTSE 100 cross-sections than in the short sterling cross-sections, leading to relatively high positive correlation for the FTSE 100 X_{01} statistics, in contrast to the complete lack of correlation for the same short sterling statistic.

This greater day-to-day stability of the SML method when fitting cross-sections of options cannot be simply be due to the SML PDFs under-fitting the data and thus not responding to actual changes in the underlying risk-neutral distributions. The SML PDFs are calibrated to have the same average goodness-of-fit as the DLN PDFs. It is likely that the relative instability of the DLN PDFs is due to the parametric nature of the DLN method which may result in local price changes affecting the entire distribution, while the non-parametric SML method is more likely to isolate the effects of small price changes to a smaller portion of the estimated PDF. Alternatively, the problem could be due to numerical instability in the DLN method, as suggested by the failure of this method to obtain a solution in a number of cases (see Table 1), a failure not observed with the SML method.

To test this conjecture, we turn to our price-perturbation tests. By keeping the perturbed data close to the original data (within half a tick) we are holding the underlying distribution nearly constant for each round of simulations. Table 4 presents various measures of dispersion in the perturbed-price implied PDF summary statistics. The upper panel presents the statistics for various subsets of the results. Use of percentile ranges is robust as long as the range is not so broad as to include ill-behaved tail outcomes. For every range we observe that the dispersion of SML results is almost always smaller than those derived from the DLN method, often by an order of magnitude. As the number of simulations included expands, the ranges expand as well. The X_{01} to X_{99} DLN results show a particularly precipitous decline in the ability of the DLN method to produce stable PDFs. The exception is in the SML $\hat{\mu}$ -measures. SML PDFs naturally fit the PDF mean to the futures price. The DLN means on the other hand vary considerably.

Table 5 presents the standard deviations of the perturbed-price PDF summary statistics. The ‘before filtering’ sample includes the entire set of option cross-sections included in the analysis (after filtering as described in the previous section). These statistics show a marked disparity in the stability of the DLN results when compared to the SML results. In some cases, such as the FTSE 100 $\hat{\mu}$, $\hat{\sigma}$, Skew_1 , Skew_2 , and Kurt , the differences are approximately two orders of magnitude (100 times).

Standard deviation estimates are not robust to outliers. To test the robustness of these results we filter the simulated data once again. The $X_{0.5}$ to $X_{99.5}$ range of each summary statistic for each estimation method was first computed. Any simulations that contained a summary statistic that fell outside

Table 4
Distribution of PDF summary statistics for perturbed data

| Summary statistic | Inter-quartile range | | | | X_{10} to X_{90} range | | | |
|-------------------|----------------------|-------|----------------|-------|----------------------------|-------|----------------|-------|
| | FTSE100 | | Short sterling | | FTSE100 | | Short sterling | |
| | DLN | SML | DLN | SML | DLN | SML | DLN | SML |
| $\hat{\mu}$ | 0.37 | 0.00 | 0.009 | 0.000 | 0.83 | 0.00 | 0.024 | 0.000 |
| $\hat{\sigma}$ | 0.32 | 0.20 | 0.019 | 0.006 | 0.86 | 0.40 | 0.048 | 0.011 |
| Skew ₁ | 0.006 | 0.003 | 0.168 | 0.084 | 0.024 | 0.009 | 0.462 | 0.176 |
| Skew ₂ | 0.005 | 0.004 | 0.079 | 0.100 | 0.013 | 0.009 | 0.194 | 0.264 |
| Skew ₃ | 0.002 | 0.001 | 0.033 | 0.023 | 0.005 | 0.002 | 0.096 | 0.045 |
| Skew ₄ | 0.004 | 0.002 | 0.069 | 0.136 | 0.011 | 0.006 | 0.354 | 0.274 |
| Kurt | 0.026 | 0.006 | 0.608 | 0.299 | 0.114 | 0.025 | 2.093 | 0.624 |
| X_{01} | 4.19 | 1.47 | 0.126 | 0.059 | 11.11 | 3.23 | 0.370 | 0.106 |
| X_{05} | 1.59 | 0.79 | 0.061 | 0.027 | 5.01 | 1.50 | 0.155 | 0.051 |
| X_{95} | 0.93 | 0.81 | 0.033 | 0.020 | 2.37 | 1.55 | 0.085 | 0.036 |
| X_{99} | 2.13 | 1.64 | 0.088 | 0.067 | 7.22 | 3.21 | 0.215 | 0.123 |
| N | 140,610 | | 68,374 | | 140,610 | | 68,374 | |

| Summary statistic | X_{05} to X_{95} range | | | | X_{01} to X_{99} range | | | |
|-------------------|----------------------------|-------|----------------|-------|----------------------------|-------|----------------|-------|
| | FTSE100 | | Short sterling | | FTSE100 | | Short sterling | |
| | DLN | SML | DLN | SML | DLN | SML | DLN | SML |
| $\hat{\mu}$ | 1.67 | 0.00 | 0.043 | 0.000 | 14.77 | 0.01 | 0.115 | 0.000 |
| $\hat{\sigma}$ | 3.00 | 0.54 | 0.082 | 0.014 | 33.48 | 0.92 | 0.202 | 0.020 |
| Skew ₁ | 0.091 | 0.018 | 0.771 | 0.249 | 0.848 | 0.057 | 1.732 | 0.444 |
| Skew ₂ | 0.037 | 0.014 | 0.296 | 0.451 | 0.283 | 0.032 | 0.724 | 1.219 |
| Skew ₃ | 0.011 | 0.002 | 0.160 | 0.062 | 0.086 | 0.004 | 0.371 | 0.099 |
| Skew ₄ | 0.033 | 0.011 | 0.985 | 0.369 | 0.324 | 0.029 | 9.541 | 0.595 |
| Kurt | 0.339 | 0.056 | 4.023 | 0.862 | 2.287 | 0.236 | 12.330 | 1.413 |
| X_{01} | 32.37 | 4.94 | 0.615 | 0.130 | 320.10 | 10.31 | 1.503 | 0.164 |
| X_{05} | 12.46 | 1.99 | 0.244 | 0.064 | 107.88 | 3.11 | 0.506 | 0.081 |
| X_{95} | 7.48 | 1.99 | 0.131 | 0.045 | 78.92 | 2.89 | 0.253 | 0.062 |
| X_{99} | 37.61 | 4.14 | 0.354 | 0.150 | 586.06 | 6.73 | 0.902 | 0.190 |
| N | 140,610 | | 68,374 | | 140,610 | | 68,374 | |

this range was then discarded. Applied to a single summary statistic this would have reduced the sample size by 1%. Applied to all 22 summary statistics (11 statistics, 2 estimation methods), this reduced the sample sizes to 127,491 simulations for FTSE 100 and 59,915 for short sterling.

Filtering out the potential outliers somewhat mitigates the extreme disparities between the two methods, but does not change the results. In virtually all cases the SML method is more stable than the DLN method, usually by an order of magnitude. The sole exception is the short sterling Skew₂ results. Taken together, these test results show that DLN-estimated PDFs are extremely sensitive to small changes in the data as compared to the SML-derived PDFs.

Table 5
Standard deviation of PDF summary statistics for perturbed data

| Summary statistic | Before filtering | | | | After filtering | | | |
|-------------------|------------------|-------|----------------|-------|-----------------|-------|----------------|-------|
| | FTSE100 | | Short sterling | | FTSE100 | | Short sterling | |
| | DLN | SML | DLN | SML | DLN | SML | DLN | SML |
| $\hat{\mu}$ | 12.50 | 0.01 | 0.046 | 0.000 | 1.00 | 0.00 | 0.011 | 0.000 |
| $\hat{\sigma}$ | 12.23 | 0.18 | 0.040 | 0.004 | 2.14 | 0.15 | 0.020 | 0.004 |
| Skew ₁ | 0.174 | 0.009 | 0.287 | 0.079 | 0.050 | 0.005 | 0.192 | 0.068 |
| Skew ₂ | 0.168 | 0.009 | 0.330 | 0.170 | 0.017 | 0.005 | 0.078 | 0.146 |
| Skew ₃ | 0.035 | 0.001 | 0.079 | 0.019 | 0.005 | 0.001 | 0.040 | 0.017 |
| Skew ₄ | 0.076 | 0.005 | 9.427 | 0.115 | 0.022 | 0.003 | 0.987 | 0.104 |
| Kurt | 3.661 | 0.039 | 3.311 | 0.274 | 0.165 | 0.018 | 1.199 | 0.231 |
| X_{01} | 62.08 | 1.69 | 0.229 | 0.039 | 18.45 | 1.37 | 0.169 | 0.038 |
| X_{05} | 24.62 | 0.74 | 0.093 | 0.019 | 6.85 | 0.58 | 0.068 | 0.018 |
| X_{95} | 19.76 | 0.61 | 0.042 | 0.014 | 5.35 | 0.58 | 0.035 | 0.013 |
| X_{99} | 79.16 | 1.34 | 0.141 | 0.046 | 43.15 | 1.18 | 0.101 | 0.043 |
| N | 140,610 | | 68,374 | | 127,491 | | 59,915 | |

Note: The filtered sample was constructed by deleting all observations that contained an outlier for any of the eleven summary statistics. An outlier is defined as any value outside the 0.5 to 99.5 percentiles of the respective empirical distribution.

The perturbed-price simulations may also be used to provide rough confidence intervals for the estimated unperturbed data PDF summary statistics. For example, if we are interested in a 90% confidence interval we may use the X_{05} to X_{95} ranges of the summary statistic deviations from the perturbation results. Table 6 presents these ranges as percentages of the corresponding mean statistic value and the mean absolute day-to-day change in the statistic. The FTSE 100 DLN Skew₁ estimates have a rough 90% confidence interval of 0.091 or 17% of the average estimated skewness (−0.544) and 77% of the mean absolute day-to-day change in Skew₁ (0.118). The FTSE 100 SML Skew₁ measure has a confidence interval of only 0.018 or 3% of the average skewness value (−0.683) and 41% of the mean absolute day-to-day change in Skew₁ (0.044). For short sterling the confidence regions are much larger in relation to the measured statistics for both PDF estimation methods. This is because the average number of strikes per option cross-section is much smaller. For the short sterling DLN Skew₁, the 90% confidence interval (0.771) is actually 17% larger than the average value of the statistic (0.657) and more than three times the mean absolute day-to-day change in the statistic (0.226). For short sterling SML Skew₁, the confidence interval (0.249) is 33% of the average value (0.766) and 165% of the mean absolute day-to-day change (0.151). The problem is similarly severe for other skewness measures and for kurtosis. The tail percentiles also show relatively larger confidence intervals for DLN and smaller ones for SML. In almost all cases, the confidence intervals for short sterling higher moment and tail summary statistics exceed the corresponding mean

Table 6
90% confidence intervals as percentages

| Summary statistic | 90% confidence intervals as percentages of: | | | | | | | |
|----------------------|---|------------|----------------|------------|---------------------------------|------------|----------------|------------|
| | Mean statistic | | | | Mean absolute day-to-day change | | | |
| | FTSE100 | | Short sterling | | FTSE100 | | Short sterling | |
| | DLN (%) | SML (%) | DLN (%) | SML (%) | DLN (%) | SML (%) | DLN (%) | SML (%) |
| $\hat{\mu}$ | 0 | 0 | 1 | 0 | 4 | 0 | 102 | 0 |
| $\hat{\sigma}$ | 1 | 0 | 12 | 2 | 26 | 6 | 273 | 88 |
| Skew ₁ | 17 | 3 | 117 | 33 | 77 | 41 | 341 | 165 |
| Skew ₂ | 27 | 6 | 172 | 223 | 42 | 34 | 224 | 278 |
| Skew ₃ | 13 | 2 | 188 | 61 | 55 | 25 | 333 | 230 |
| Skew ₄ | 31 | 8 | 195 | 224 | 72 | 42 | 193 | 251 |
| Kurt | 8 | 1 | 78 | 19 | 50 | 44 | 296 | 160 |
| X_{01} | 1 | 0 | 11 | 2 | 40 | 10 | 334 | 220 |
| X_{05} | 0 | 0 | 4 | 1 | 20 | 4 | 287 | 156 |
| X_{95} | 0 | 0 | 2 | 1 | 17 | 5 | 190 | 82 |
| X_{99} | 1 | 0 | 4 | 2 | 51 | 11 | 217 | 133 |

Note: The 90% confidence interval is defined as the difference between the 5th and the 95th percentiles of the perturbed statistics.

absolute day-to-day changes. The exceptions are the short sterling Skew₂ and Skew₄ statistics. The confidence intervals for $\hat{\mu}$ and $\hat{\sigma}$ are much smaller. For FTSE 100 DLN PDFs, the 90% confidence intervals for $\hat{\mu}$ and $\hat{\sigma}$ (1.67 and 3.00) are 4% and 26% of their respective mean absolute day-to-day changes (40.8 and 11.4). For FTSE 100 SML PDFs the corresponding numbers are 0 and 6% respectively. The short sterling DLN $\hat{\mu}$ confidence interval is likewise approximately equal to the day-to-day variation, while short sterling DLN and SML $\hat{\sigma}$ confidence intervals are roughly three times and almost equal to their mean absolute day-to-day changes respectively.

These results show that, except for measures of location and dispersion ($\hat{\mu}$ and $\hat{\sigma}$) where cross-section sizes are large (FTSE 100), the effect of price measurement error introduces a significant degree of uncertainty to the accuracy of implied PDF-derived summary statistics. It should be remembered that these are minimal confidence intervals – they account for only one source of pricing uncertainty and probably the smallest one at that.

Lastly, we examine the impact that the choice of weighting scheme might have on our results. We examined three weighting schemes: equal weighting; vega weighting; and an ad hoc weighting that placed weights of 1 on the four strikes closest-to-the-money, weights of 1/2 for the remaining out-of-the-money strikes, and 1/4 for the remaining in-the-money strikes. These fitted price error weights were applied to the DLN estimations. For the SML estimations the fitted price error weights were converted to fitted implied volatility error weights by multiplying them by the appropriate vegas. The test sample was

created by finding the first 120 cross-sections for which it was possible to find DLN and SML solutions (again, only the former were problematic) for all three pairs of estimations. For each pair of estimations the λ was calibrated so that the SML method had approximately the same overall goodness-of-fit as the DLN method. For each cross-section, 100 perturbed-price simulations were run, as before.

Table 7 presents selected results. Again the DLN method is less stable overall than is the SML method, regardless of the weighting scheme used. In only three cases for each estimation pair does the DLN show less variability in the PDF summary statistic: Skew₂ and Skew₄ in the inter-quartile range results, and Skew₂ in the 90% confidence interval results. There is no variation in this

Table 7
Distribution of PDF summary statistics for perturbed data and various weightings

| Summary statistic | Inter-quartile range | | | | | |
|-------------------|----------------------------|--------|----------------|--------|----------------|--------|
| | Equal weighting | | Vega weighting | | Step weighting | |
| | DLN | SML | DLN | SML | DLN | SML |
| $\hat{\mu}$ | 0.0085 | 0.0000 | 0.0091 | 0.0000 | 0.0082 | 0.0000 |
| $\hat{\sigma}$ | 0.0175 | 0.0056 | 0.0182 | 0.0049 | 0.0174 | 0.0057 |
| Skew ₁ | 0.1804 | 0.0705 | 0.2014 | 0.0569 | 0.1691 | 0.0660 |
| Skew ₂ | 0.0557 | 0.0770 | 0.0636 | 0.0761 | 0.0600 | 0.0797 |
| Skew ₃ | 0.0231 | 0.0158 | 0.0263 | 0.0152 | 0.0256 | 0.0193 |
| Skew ₄ | 0.0471 | 0.0917 | 0.0579 | 0.0785 | 0.0543 | 0.1103 |
| Kurt | 0.7687 | 0.2753 | 0.8125 | 0.2071 | 0.7078 | 0.2630 |
| X_{01} | 0.1352 | 0.0508 | 0.1343 | 0.0349 | 0.1150 | 0.0423 |
| X_{05} | 0.0572 | 0.0264 | 0.0647 | 0.0210 | 0.0586 | 0.0243 |
| X_{95} | 0.0436 | 0.0175 | 0.0403 | 0.0161 | 0.0399 | 0.0166 |
| X_{99} | 0.0963 | 0.0559 | 0.1064 | 0.0452 | 0.0952 | 0.0540 |
| N | 11,443 | | 11,615 | | 11,635 | |
| Summary statistic | X_{05} to X_{95} range | | | | | |
| | Equal weighting | | Vega weighting | | Step weighting | |
| | DLN | SML | DLN | SML | DLN | SML |
| $\hat{\mu}$ | 0.0360 | 0.0000 | 0.0380 | 0.0000 | 0.0350 | 0.0000 |
| $\hat{\sigma}$ | 0.0702 | 0.0137 | 0.0715 | 0.0129 | 0.0653 | 0.0139 |
| Skew ₁ | 0.7056 | 0.1927 | 0.7995 | 0.1629 | 0.7088 | 0.1890 |
| Skew ₂ | 0.2322 | 0.2992 | 0.2321 | 0.2666 | 0.2113 | 0.3594 |
| Skew ₃ | 0.1107 | 0.0423 | 0.1083 | 0.0411 | 0.1062 | 0.0498 |
| Skew ₄ | 0.3026 | 0.2382 | 0.5344 | 0.2052 | 0.4933 | 0.2774 |
| Kurt | 3.0000 | 0.7647 | 3.7220 | 0.5572 | 3.0960 | 0.7263 |
| X_{01} | 0.6667 | 0.1166 | 0.7041 | 0.0829 | 0.6043 | 0.0998 |
| X_{05} | 0.2797 | 0.0620 | 0.3054 | 0.0488 | 0.2792 | 0.0555 |
| X_{95} | 0.1384 | 0.0397 | 0.1404 | 0.0365 | 0.1370 | 0.0386 |
| X_{99} | 0.2791 | 0.1363 | 0.3453 | 0.1068 | 0.3197 | 0.1317 |
| N | 11,443 | | 11,615 | | 11,635 | |

pattern across weighting methods. Nor does the choice of weighting scheme much affect the size of the individual statistics. This is consistent with the hypothesis advanced earlier that because the fitted price errors are generally very small, the weights used to multiply them have little impact on the estimation. These results provide evidence that the full-sample results presented above are not artifacts of the somewhat arbitrary choice of weighting used in the tests.

5. Summary and conclusions

This paper has developed and applied a methodology for assessing the relative and absolute stability of implied PDFs estimated from options prices to small errors in prices. The technique consists of re-estimating PDFs on randomly perturbed sets of options prices constructed so as to be observationally equivalent to the original data. The perturbed data PDFs are then evaluated by examining the distribution of a number of summary statistics estimated from each PDF.

The test methodology introduced in this paper is applied to the DLN and to the SML PDF estimation methods. These PDF estimation methods are tested on two sets of options data: FTSE 100 and short sterling.

The results in this paper provide strong evidence of the superior stability of the smoothed implied volatility smile implied PDF estimation method over the double-lognormal method. Furthermore, we show that higher-order statistics, such as skewness, cannot always be estimated with precision; for either PDF estimation method, the confidence intervals can be so large as to make the estimates useless. This is particularly true of the short sterling data. This suggests that day-to-day or month-to-month changes in skewness and kurtosis should not be over-analyzed. Regardless of the data set or estimation method used, confidence regions are too large to reliably measure day-to-day changes in summary statistics. Studies which look at announcements on estimated implied PDFs should take this uncertainty into consideration. Only persistent changes in these measures, of an order exceeding the minimal confidence intervals derived using the methods developed in this paper, are likely to contain useful information.

The problems with the DLN method arise from its sensitivity to computational problems. The statistical analysis presented here ignores the frequent convergence failures encountered in applying the DLN implied PDF estimation method and the tendency of the method to produce spurious spikes in the estimated PDFs when one of the component lognormal density functions collapses. It may be possible to restart the DLN solution by repeatedly using a grid of initial values and then visually inspecting the PDFs for abnormalities to arrive at a plausible solution. This is impracticable for large numbers of option

cross-sections and the solution thus laboriously achieved is apt still to be more sensitive to data problems than solutions from more robust methods such as the smoothed smile used in this paper.¹⁹

The smoothed implied volatility smile implied PDF estimation method is remarkably free of computational problems. Applied to a large cross-section of strikes, such as we find with FTSE 100 options, the estimated implied PDFs are reasonably insensitive to small measurement errors. However, for small cross-sections, such as we find with short sterling options, the estimated SML implied PDFs, while much more stable than the DLN implied PDFs estimated on the same data set, are nonetheless not so stable that derived statistics can safely be viewed as precise measures. As always in statistics, parameters estimated from sample data are estimates with associated uncertainties. In the case of higher-moment statistics derived from options prices, this cannot be ignored.

It is clear from this analysis that the double-lognormal method, while widely used, should not be. In the future, other alternatives may well prove superior to the smoothed implied volatility smile. It is likely that the smoothed implied volatility smile method used in this paper, chosen to have the same goodness-of-fit as the DLN method, may itself be improved upon by trading off a small degree of fit for a more stable implied PDF estimate. The methods used in this paper both fit the data to a high precision, which the nature of the quotation process cannot justify. There is considerable room to trade off fit and stability. The smoothed implied volatility smile method permits this fine tuning. The double-lognormal method does not.

Acknowledgements

The authors acknowledge the helpful comments and suggestions of Roger Clews, Eric Edmond, Gordon Gemmill, Holger Neuhaus, James Proudman, Steven Smith, Paul Söderlind, and Gary Xu and seminar participants at the Bank of England, the Federal Reserve Bank of Chicago, and the University of Wisconsin. Genevieve Pham-Kantor and Michael Rorke provided valuable programming assistance. Any remaining errors are the authors' responsibility. The views expressed herein are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Chicago or the Bank of England.

¹⁹ It is conceivable that a solution arrived at by the labor intensive measures mentioned would be more stable than the straightforward application of the DLN estimation method.

Appendix A. Methods for estimating implied PDFs

A.1. Double-lognormal approximating function

The double-lognormal approximating function method for estimating implied PDFs from option prices is based on the following theoretical pricing relations for European calls and puts:

$$C_t(K) = e^{-rt} \int_K^\infty (S_T - K)df(S_T) \tag{A.1}$$

and

$$P_t(K) = e^{-rt} \int_{-\infty}^K (K - S_T)df(S_T), \tag{A.2}$$

where C and P are the call and put prices observed at time t ; r is the riskless rate; τ is the time to expiry; K is the exercise price; and $df(S_T)$ is the risk-neutral probability density function for the value of the underlying asset, S , at expiry, $T = t + \tau$. The double-lognormal method approximates this density function with a mixture of two lognormal density functions:

$$df(S_T) = \theta L(S_T|\mu_1, \sigma_1, S_t) + (1 - \theta)L(S_T|\mu_2, \sigma_2, S_t), \tag{A.3}$$

$$L(S_T) = \frac{1}{S_T \sigma \sqrt{2\pi\tau}} \exp \left\{ \frac{-[\log S_T - \log S_t - (\mu - \frac{1}{2}\sigma^2)\tau]^2}{2\sigma^2\tau} \right\}, \tag{A.4}$$

where S_t is the current value of the underlying asset and $\{\mu_1, \sigma_1, \mu_2, \sigma_2, \theta\}$ are the unknown parameters that define the double-lognormal density function; $\theta \in [0, 1]$.²⁰ Thus, the fitted value for a call price, given parameters $\{\mu_1, \sigma_1, \mu_2, \sigma_2, \theta\}$, is given by

$$\begin{aligned} \widehat{C}_t(K|\mu_1, \sigma_1, \mu_2, \sigma_2, \theta) = e^{-rt} \left\{ \theta \int_K^\infty (S_T - K)L(S_T|\mu_1, \sigma_1, S_t)dS_T \right. \\ \left. + (1 - \theta) \int_K^\infty (S_T - K)L(S_T|\mu_2, \sigma_2, S_t)dS_T \right\} \end{aligned} \tag{A.5}$$

with an equivalent expression for the value of put option. Given observations of call and put prices, the parameters, $\{\mu_1, \sigma_1, \mu_2, \sigma_2, \theta\}$, of the implied double-

²⁰ See Melick and Thomas (1997) and Bahra (1997) for development of this model.

lognormal PDF can be estimated using non-linear optimization methods to minimize the weighted sum of fitted price errors:

$$\min_{\{\mu_1, \sigma_1, \mu_2, \sigma_2, \theta\}} \left\{ \sum_{i=1}^{N_c} w_i [C_t(K_i) - \widehat{C}_t(K_i | \mu_1, \sigma_1, \mu_2, \sigma_2, \theta)]^2 + \sum_{j=1}^{N_p} w_j [P_t(K_j) - \widehat{P}_t(K_j | \mu_1, \sigma_1, \mu_2, \sigma_2, \theta)]^2 \right\} \tag{A.6}$$

subject to

$$\sum_{i=1}^{N_c} w_i + \sum_{j=1}^{N_p} w_j = 1 \tag{A.7}$$

and

$$w_i, w_j \geq 0 \quad \forall i, j, \tag{A.8}$$

and where N_c and N_p and the number of calls and put contracts in the estimation sample for a given pair of observations and expiry dates $\{t, T\}$ and the w_i, w_j are the weights placed on each option.

A.2. Smoothed implied volatility smile

The smoothed implied volatility smile implied PDF estimation method is based on a few simple ideas. Breeden and Litzenberger (1978) showed that the PDF for distribution of the value of the underlying asset at option expiry, $f(S_T)$, is related to call (or put) prices through

$$f(S_T) = \frac{\partial^2 C}{\partial K^2}, \tag{A.9}$$

so that if we observed the call price function (as a function of strike) we could differentiate twice to obtain the PDF. However, we only observe call prices for a relatively few discretely spaced strikes.

The obvious solution is to interpolate (and extrapolate), or alternatively smooth, the observed prices by fitting a function to the observed prices. However, several technical reasons make this approach undesirable. Small fitted price errors can have large effects on the PDFs, particularly in the tails, and otherwise advantageous smoothing splines are not well suited to fitting a fundamentally exponential functional form that asymptotes to zero slope for deep out-of-the-money and slope of one for deep in-the-money.

Shimko (1993) proposed as an alternative that the observed option prices be first converted to implied volatilities using the Black–Scholes option pricing formula. The implied volatility function could then be fitted and the continuum

of fitted implied volatilities converted back to a continuum of fitted option prices and thence to a fitted PDF. Shimko chose to use a simple quadratic polynomial smoothing function within the span of available strikes and with flat linear extrapolations outside the range of available strikes. While this method relies on the Black–Scholes pricing model as a method for transforming from one space (price/strike) to another (BS-implied volatility/strike) this is purely a computational convenience and does not presume that the Black–Scholes model is true.

Malz (1997a,b) modified Shimko’s technique by transforming the original data from price/strike space to implied volatility/delta space, where

$$\text{delta} \equiv \Delta = \frac{\partial C}{\partial S}. \quad (\text{A.10})$$

Using delta, Δ , rather than strike, K , as the function argument, groups away-from-the-money implied volatilities more closely together than near-the-money implied volatilities. This has the effect of permitting greater “shape” near the center of the distribution where the data is more reliable (frequently traded), without using a variable smoothness penalty across the length of the spline.²¹ Malz (1997a,b) followed Shimko in using a low-order polynomial as the smoothing function.

Campa et al. (1997) introduced the use of a smoothing spline for fitting implied volatility curves. They applied this to smoothing the implied volatility/strike function. Use of a natural spline, rather than a low order polynomial, permits the user to control the smoothness of the fitted function. The spline is also less restrictive of the shapes the fitted function can assume.

The smoothed smile PDF estimation method used in this paper was developed by Panigirtzoglou in previously unpublished work at the Bank of England. The method follows Malz (1997a,b) in smoothing in implied volatility/delta space and Campa et al. (1997) in using a natural spline to smooth the function. In addition, the fitted implied volatility errors are weighted by the option vegas ($v \equiv \partial C / \partial \sigma$) to account for presumed homoskedastic pricing errors in the underlying raw price data.

A natural spline is a piece-wise cubic polynomial. For a natural spline the points on the x -axis corresponding to each of the data points define the “knot points”.²² Between the knot points the function is a simple cubic polynomial. However, the function is constrained so that it is continuous at the knot points and has continuous first and second derivatives. The smoothness of the spline is controlled by a smoothness penalty, λ , which multiplies a measure of the degree

²¹ See Waggoner (1997) for the application of variable smoothness penalties to fitting natural splines.

²² Some applications use a smaller set of knot points at predetermined intervals or selected to span approximately equal number of data points.

of curvature in the function – the integral of the squared second derivative of the function over its range. The objective function to be minimized is thus

$$\min_{\Theta} \sum_{i=1}^N w_i (IV_i - \widehat{IV}_i(\Delta_i, \Theta))^2 + \lambda \int_{-\infty}^{\infty} f''(x; \Theta)^2 dx, \tag{A.11}$$

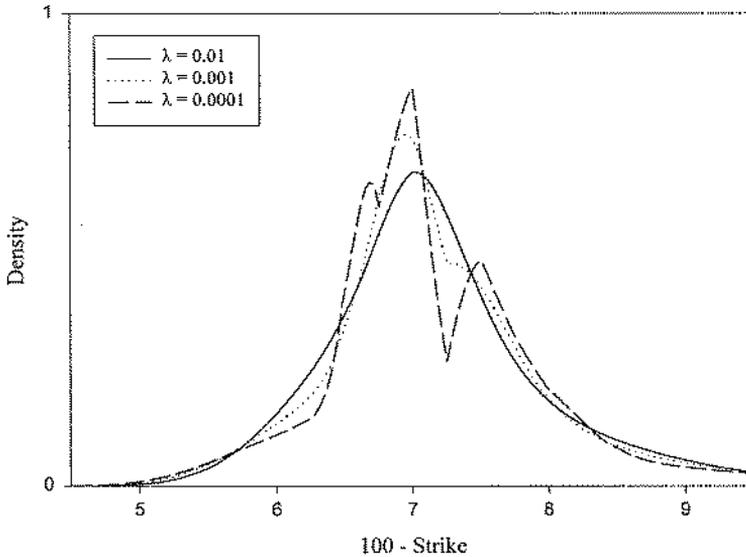
where Θ is the matrix of parameters of the cubic spline (knot points and component polynomial parameters), $f(\Theta)$ is the cubic spline function, and $\widehat{IV}_i(\Delta_i, \Theta)$ is the fitted implied volatility at Δ_i , given the spline parameters Θ . Implementing the natural spline is both simple and computationally efficient.

The cubic smoothing spline has the property that it becomes linear outside the range of available data. If the slope at the extreme knot points is negative, it is possible for extrapolate values of the fitted density to be negative. It is possible to alter the spline function or to constrain it to avoid this problem, though at some cost the computation efficiency. In this study, the negative fitted-density problem did not occur.

Once the natural spline is fitted, a large number of equally Δ -spaced points on the function are computed. These are then converted into equally K -spaced values in price/strike space. These in turn are used to compute the PDF using numerical methods.

The presence of the smoothing parameter λ in the objective function being minimized allows the user to control the trade-off between goodness-of-fit and smoothness of the estimated PDF. Fig. 7 illustrates the effects of varying λ when fitting PDFs to a single contract. The PDF corresponding to $\lambda = 0.01$ produces a noticeably smoother PDF than does the PDF for $\lambda = 0.0001$ however at the cost of a markedly poorer goodness-of-fit (0.6454 vs. 0.0804). However, this variation of goodness-of-fit must be put into context: The maximum fitted price error for $\lambda = 0.01$ is approximately equal to the option's tick size of 0.005. The maximum error for the $\lambda = 0.001$ PDF is approximately one half the tick size – effectively fitting the data to within the precision of option-price measurement. Thus, while the looser $\lambda = 0.0001$ PDF has a complex shape, this appears to result from over-fitting the data to a degree that exceeds the precision of the underlying data. The less complex $\lambda = 0.001$ PDF is appears to fit the data adequately, while the smooth $\lambda = 0.01$ PDF may be somewhat over-smoothing the data. Varying λ in this case has little effect on derived summary statistics except for kurtosis.²³

²³ By construction the smoothed implied volatility smile method will always fit the mean of the implied PDF to the future's price, regardless of the value of λ .



| λ | 0.01 | 0.001 | 0.0001 |
|-----------|--------|--------|--------|
| VOF | 0.6454 | 0.1780 | 0.0804 |
| Mean | 6.7600 | 6.7600 | 6.7600 |
| Std. Dev. | 0.5887 | 0.5887 | 0.5890 |
| Skewness | 0.4633 | 0.4531 | 0.4457 |
| Kurtosis | 3.5729 | 3.6728 | 2.7364 |
| Max FPE | 0.0048 | 0.0027 | 0.0019 |

Fig. 7. PDFs estimated from a single short sterling contract using different smoothing parameters (λ) (note: Contract expiry 17-Sep-97; quoted 3-Mar-97; 14 usable strikes, VOF is the value of the objective function, the sum of squared fitted volatility errors, FPE is the fitted price errors).

Appendix B. Data description

B.1. Short sterling

The LIFFE short sterling options contract is an American-style option on the three month sterling interest rate futures contract. Options and the underlying futures contract expire simultaneously on the third Wednesday of the nearest March, June, September, and December (i.e. Shiratsuka, 1999). No options are listed with expiries beyond one year. The short sterling option exercise prices are in intervals of 0.25 (25 basis points) except for the near contract which is quoted in strike intervals of 0.125. The minimum tick size is 0.01.²⁴ The underlying futures contract cash settles to 100 minus the 3-month London interbank offer rate (LIBOR) rate on the futures contract expiry date. Option premia are not paid at time of purchase. Option positions are marked-to-market daily and the

²⁴ After April 1998 the tick size was reduced to 0.005 for contracts trading at less than 0.03.

premium is paid only if the option is exercised. The combined effect of the deferred premium and the (initially costless) underlying asset results in the following put–call parity relation for the short sterling contract:²⁵

$$C - P = S - K, \quad (\text{B.1})$$

where C is the call price for an option with strike K , P is the corresponding put price, and S is the current value of the underlying asset. This condition is not violated for any contracts in the 1997 data set, indicating that the condition is imposed when short sterling option settlement prices are determined. It also means that put and call prices are redundant and we may use either. In this study we use call prices.

In 1997 the average near-the-money²⁶ short sterling call price was 0.219 (range 0.010–0.600).²⁷ The tick size is therefore 4.6% of the average near-the-money short sterling option premium. For away-from-the-money options the tick size would be a correspondingly higher percentage of the time value (the intrinsic value portion of the option price is irrelevant in this context). If we are interested in the time series of changes in derived statistics rather than in their absolute levels, the more relevant benchmark for evaluating the tick size is the daily price change. In 1997, the daily near-the-money short sterling call price changes ranged between -0.240 and 0.280 , with a mean absolute value of 0.155 and a standard deviation of 0.141 . Thus, the tick size is 6.6% of the mean absolute daily price change and 7.1% of the standard deviation of daily near-the-money short sterling call price changes and correspondingly greater for away-from-the-money strikes.

B.2. FTSE 100

The LIFFE FTSE 100 option contract used in this study is an European-style option on the FTSE 100 stock index.²⁸ Options are listed for the nearest four months and for the nearest June and December. FTSE 100 options expire

²⁵ This assumes all cash flows occur at expiration and ignores mark-to-market effects. Evidence from studies of price differences for interest rate futures and interest rate forwards indicates that the mark-to-market distortion is minimal. Strict equality theoretically only holds for European-style options. However, the absence of potential “dividends” from holding the underlying asset makes early exercise of the call sub-optimal, and the premium-payment-on-exercise raises the threshold at which it is optimal to exercise the put early. In practice the short sterling contract trades as if it were a European-style option.

²⁶ Defined as the strike immediately above and the strike at or immediately below the futures price.

²⁷ Statistics are based on all options with at least seven days to expiry. Price change statistics are based on consecutive quotations over intervals of no more than three days. The same applies to the FTSE 100 options statistics below.

²⁸ LIFFE also lists a American-style FTSE 100 contract. This contract is less liquid than the European-style contract.

on the third Friday of the expiry month. The FTSE 100 option strikes are in intervals of 50 or 100 points depending on time-to-expiry. The minimum tick size is 0.5. FTSE 100 options cash settle to the daily settlement price, which is determined by taking the average level of the FTSE 100 index sampled every 15 seconds between 16:20 and 16:30 (London time). For expiring contracts the sampling is done between 10:10 and 10:30 on the last trading day. The option premium is payable in full at the time of purchase (next business day).

Because the FTSE 100 futures contract expires on the same date as the option and therefore will have the same value as the index when the option expires, the European-style FTSE 100 contract may be viewed as an option on the futures. The put–call parity relation for the FTSE 100 contract is therefore

$$C - P = (S - K)e^{-r_f\tau}, \quad (\text{B.2})$$

where C is the call price for an option with strike K , P is the corresponding put price, S is the current value of the underlying asset, in this case a futures price, τ is the time to expiry, and r_f is the risk-free interest rate. Again, this theoretical relation ignores mark-to-market effects.

Put–call parity is frequently violated for FTSE 100 options contracts in the 1997 data set as a result put and call prices are not redundant.

In 1997 the average near-the-money FTSE 100 call price was 188.4 (range 9.50 to 583.5). The tick size is therefore 0.3% of the average near-the-money FTSE 100 option premium. In 1997, the daily price changes ranged between –164.0 and 186.0, with a mean absolute value of 42.44 and a standard deviation of 49.78. Thus the tick size is 1.2% of the mean absolute daily price change and 1.0% of the standard deviation of daily price changes. One would expect therefore that tick-size effects are less of a problem in the FTSE 100 option market than in the short sterling market.

References

- Abken, P.A., Madan, D.B., Ramamurtie, S., 1996. Estimation of risk-neutral and statistical densities by hermite polynomial approximation: With an application to Eurodollar futures options. Working paper no. 96-5, Federal Reserve Bank of Atlanta.
- Ait-Sahalia, Y., Lo, A.W., 1998. Nonparametric estimation of state-price densities implied in financial asset prices. *Journal of Finance* 53, 499–547.
- Barnett, V., Lewis, T., 1984. *Outliers in Statistical Data*, second edition. Wiley, New York.
- Bahra, B., 1996. Probability distributions of future asset prices implied by options prices. *Bank of England Quarterly Bulletin* 36 (3) 299–311.
- Bahra, B., 1997. Implied risk-neutral probability density functions from option prices: Theory and application. Working paper no. 66, Bank of England.
- Belsley, D.A., Kuh, E., Welch, R.E., 1980. *Regression Diagnostics: Identifying Influential Data and Sources of Collinearity*. Wiley, New York.

- Breeden, D.T., Litzenberger, R.H., 1978. Prices of state-contingent claims implicit in option prices. *Journal of Business* 51, 621–651.
- Buchen, P.W., Kelly, M., 1996. The maximum entropy distribution of an asset inferred from option prices. *Journal of Finance and Quantitative Analysis* 31, 143–159.
- Butler, C., Davies, H., 1998. Assessing market views on monetary policy: The use of implied risk-neutral probability distributions. Paper presented at BIS/CEPR conference on Asset Prices and Monetary Policy.
- Campa, J.M., Chang, P.H.K., Reider, R.L., 1997. Implied exchange rate distributions: Evidence from OTC option markets. NBER Working paper no. 6179.
- Campa, J.M., Chang, P.H.K., Refalo, J.F., 1998. An options-based analysis of emerging market exchange rate expectations: Brazil's Real plan, 1994–1997. Working paper, New York University.
- Cooper, N., 1999. Testing techniques for estimating implied RNDs from the prices of European and American options. Unpublished working paper, Bank of England.
- Coutant, S., 1999. Implied risk aversion in option prices. Working paper, Banque de France.
- Coutant, S., Jondeau, E., Rockinger, M., 1999. Reading interest rate and bond futures options' smiles: How PIBOR and notational operators appreciated the 1997 French snap election. Working paper NER no. 54, Banque de France.
- Cox, J., Ross, S., 1976. The valuation of options for alternative stochastic processes. *Journal of Financial Economics* 3, 145–166.
- Duffee, G.R., 1996. Idiosyncratic variation of Treasury bill yields. *Journal of Finance* 51, 527–551.
- Galati, G., Melick, W., 1999. Perceived central bank intervention and market expectations: An empirical study of the yen/dollar exchange rate 1993–96. Working paper no. 77, Bank of International Settlements.
- Gemmill, G., Saflekos, A., 1999. How useful are implied distributions? Evidence from stock-index options. Working paper, City University Business School.
- Hördahl, P., 1999. Estimating the implied distribution of the future short term interest rate using the Longstaff–Schwartz model. Working paper, European Central Bank.
- Jackwerth, J.C., Rubinstein, M., 1996. Recovering probability distributions from options prices. *Journal of Finance* 51, 1611–1631.
- Leahy, M.P., Thomas, C.P., 1996. The sovereignty option: The Quebec referendum and market views on the Canadian dollar. International Finance Discussion Paper no. 555, Board of Governors of the Federal Reserve System.
- Madan, D.B., Milne, F., 1994. Contingent claims valued and hedged by pricing and investing in a basis. *Mathematical Finance* 4, 223–245.
- Malz, A.M., 1996. Using option prices to estimate realignment probabilities in the European monetary system: The case of sterling–mark. *Journal of International Money and Finance* 15, 717–748.
- Malz, A.M., 1997a. Option-implied probability distributions and currency excess returns. Staff Reports no. 32, Federal Reserve Bank of New York.
- Malz, A.M., 1997b. Estimating the probability distribution of the future exchange rate from options prices. *Journal of Derivatives* 5 (2) 18–36.
- Manzano, M.C., Sanchez, I., 1998. Indicators of short-term interest rate expectations: The information contained in the options market. Documento de Trabajo no. 9816, Servicio de Estudios, Banco de Espana.
- McCauley, R., Melick, W.R., 1996. Propensity and density. *Risk* 9, 52–54.
- McManus, D.J., 1999. The information content of interest rate futures options. Working paper no. 99-15, Bank of Canada.
- Melick, W.R., Thomas, C.P., 1997. Recovering an asset's implied PDF from option prices: An application to crude oil during the Gulf crisis. *Journal of Financial and Quantitative Analysis* 32, 91–115.

- Melick, W.R., Thomas, C.P., 1998. Confidence intervals and constant maturity series for probability measures extracted from option prices. Paper presented at the Bank of Canada Conference on Information Contained in Prices of Financial Assets.
- Nakamura, H., Shiratsuka, S., 1999. Extracting market expectations from option prices. Working paper, Bank of Japan.
- Neuhaus, H., 1995. The information content of derivatives for monetary policy-implied volatilities and probabilities. Economic Research Group Discussion Paper no. 3/95, Deutsche Bundesbank.
- Rubinstein, M., 1994. Implied binomial trees. *Journal of Finance* 49, 771–818.
- Sherrick, B.J., Garcia, P., Tirupattur, V., 1996. Recovering probabilistic information from option markets: Tests of distributional assumptions. *Journal of Futures Markets* 16, 545–560.
- Shimko, D.C., 1993. Bounds of probability. *Risk* 6, 33–37.
- Shiratsuka, S., 1999. Information content of implied probability distributions: Empirical studies of Japanese stock price index options. Working paper, Bank of Japan.
- Söderlind, P., 1999. Market expectations in the UK before and after the ERM crisis. Working paper no. 210, SSE/EFI Working Paper Series in Economics and Finance, Stockholm School of Economics.
- Söderlind, P., Svensson, L.E.O., 1997. New techniques to extract market expectations from financial instruments. *Journal of Monetary Economics* 40, 383–429.
- Waggoner, D.F., 1997. Spline methods for extracting interest rate curves from coupon bond prices. Working paper no. 97-10, Federal Reserve Bank of Atlanta.