

Workshop on estimating and interpreting probability density functions 14 June 1999

Background note

P H Kevin Chang and William R Melick

Starting in the late 1980s, financial and economic researchers became increasingly sophisticated in their attempts to analyze market expectations embedded in option prices. Moving beyond the study of implied Black-Scholes volatilities, this body of work has focused on the recovery of either the stochastic process followed by the underlying asset price or the density function from which the asset price at expiration will be drawn. The workshop was meant to share information and results on the latter exercise, the estimation of terminal (at expiration) probability density functions (PDFs) implied by option prices. Toward that end, this note is meant to provide some context reading the papers presented at the workshop. The first section of the note provides a brief overview (taxonomy) of the various methods used to estimate PDFs. The second section discusses issues of interpretation, providing an initial exploration of possible lines for future research.

I. Estimation of PDFs

In other surveys of PDF recoveries (in particular see Bahra (1997), techniques have been classified as falling into one of four areas: I) recovery of the stochastic process for the price of the underlying asset, with the PDF obtained as a by-product of the exercise, II) a functional form for the PDF is assumed with the parameters for the function estimated by minimizing the difference between actual and predicted option prices, III) smoothing techniques that relate option prices in some fashion to only exercise prices, allowing for the recovery of the PDF through differentiation, and IV) non-parametric techniques.

For purposes of this workshop; however, it may be more useful to follow a slightly different tack. Setting aside the methods that focus on the stochastic process¹, it is possible to classify the remaining techniques into two broad categories, based on the risk-neutral valuation equation and its second derivative. It is hoped that this classification scheme, although not perfect, will shed more light on the methodology behind the various techniques.

¹ Readers interested in the recovery of the stochastic process should consult the articles by Bates (1991 and 1996a and 1996b) and Malz (1996). Estimates of deterministic local volatility are viewed as essentially recovering the stochastic process followed by the asset price. See Bodurtha and Jermakyan (1999), Levin (1998), and Levin, Mc Manus, and Watt (1999) for more on this technique.

For the simplest case of European options, and as shown by Cox and Ross (1976), the price of the call option with strike X can be written in terms of the risk-neutral PDF for the underlying price ($f(S)$) by

$$c[X] = e^{-rT} \cdot \int_X^{\infty} (S - X) \cdot f(S) dS \quad (1)$$

where e^{-rT} is the relevant discount factor. As shown by Breeden and Litzenberger (1978), $f(S)$ can be isolated by differentiating equation (1) twice, yielding²

$$\frac{\partial^2 c[X]}{\partial X^2} = e^{-rT} \cdot f(S) \quad (2)$$

These two equations provide a convenient means for classifying the different techniques used for recovering PDFs.³ Roughly half of the techniques essentially work with equation (1) - using assumptions about the form or family of the PDF, and evaluations of the integral in equation (1), to estimate the parameters of the PDF such that predicted option prices best fit the observed option prices. The remaining techniques exploit equation (2), using a variety of means to generate the function $c[X]$ and then differentiating the function (either numerically or analytically) to obtain the PDF. Some of the very first PDF recoveries were based on equation (2), therefore they are discussed first.

A. Methods Based on Equation (2)

Making use of equation (2), Shimko (1993) was one of the first to recover the risk-neutral PDF. This technique uses the Black-Scholes formulae to translate a scatter-plot of option prices against strike prices into a scatter plot of implied volatilities against strike prices (the smile relationship). The points in the scatter plot are then used to fit a quadratic equation relating the volatility to the strike price. This then allows for the Black-Scholes equation to be written in terms of only the strike price (rather than the strike price and volatility), giving an equation that relates the option's price to only the strike price. This equation is then differentiated twice to obtain the PDF. For ranges outside of the observed strike prices, lognormal distributions are grafted on to the tails of the PDF, using the condition that the PDF has to integrate to one to pin down the parameters of the PDF.

The technique of Malz (1997) also makes use of equation (2). Using data taken from the over-the-counter (OTC) foreign exchange options market, Malz obtains, without any translation, a scatter plot of implied volatility against delta (a measure of moneyness). He then fits a particular functional form to this scatter-plot, such that each point lies on the line. Like Shimko (1993), this then allows the

² By similar reasoning, the cumulative distribution function (CDF) can be obtained by differentiating a single time. This technique is used by Neuhaus (1995).

³ Ignoring any complications introduced by American options.

option pricing equation to be written only in terms of the strike price. This equation is then differentiated (numerically) to obtain the CDF and PDF. Unlike Shimko (1993), Malz does not make special allowances for the tails, instead allowing the fitted curve to cover the entire range of possible deltas, hence the entire support of the density function.

Neuhaus (1995) also makes use of equation (2), although he chooses to differentiate a single time to recover the CDF. The derivatives are numerical and discrete in that he uses only the available strikes. Unlike Shimko and Malz he does not construct a smooth equation relating the option's price to the strike price. This can be seen as an advantage or disadvantage, however, it only allows for probability calculations at and between strike prices.

Jackwerth and Rubinstein (1996) propose a maximum smoothness criteria that essentially uses a butterfly spread variant of equation (2) to minimize the curvature in the resulting implied PDF.

Finally, Aït-Sahalia and Lo (1998) use a non-parametric method to generate a relationship between the option price and the strike price. In a data-intensive method that makes use of a cross-sectional time-series of option prices and strikes (rather than just a single day's cross-section of option prices and strikes as is used in all other studies considered herein), they utilize the Nardaraya-Watson non-parametric kernel regression to estimate the functional form that relates the call price to the strike price. The second derivative of this function then gives the PDF.

B. Methods Based on Equation (1)

The other studies considered in this note essentially make use of equation (1), typically using an nonlinear optimization method to find the exact form of the PDF that produces predicted option prices that are "close" to the observed option prices. These techniques differ in the amount of structure they place on the PDF to be derived.

Sherrick, Garcia and Tirupattur (1996), using prices for options on soybean futures contracts, specify a Burr III PDF. As in several other studies, they estimate the parameters of the density by minimizing the sum of squared option pricing errors.

Melick and Thomas (1997), using a model proposed by Ritchey (1990), specify that the PDF is to be a mixture of lognormal densities, providing bounds on American options on futures that are similar to equation (1) in that they are written in terms of the PDF. In their application to the crude oil market, they use a mixture of three lognormals, other analysts applying their technique have often used a mixture of two lognormals.⁴ Mizrach (1996) also specifies that the density is a mixture of lognormals. In a similar vein, Söderlind and Svensson (1998) stipulate that the discount factor and the underlying asset price be drawn from a mixture of bivariate normal distributions.

Several techniques have been proposed that place less structure on the functional form of the PDF. Rubinstein (1994), using a lognormal assumption, pre-specify prior terminal nodes (probabilities) for a binomial tree. Posterior terminal node probabilities are then calculated by deviating as little as possible from the prior nodes such that the predicted option prices fall between the observed bid/ask option prices and that other arbitrage possibilities are eliminated.

In a related approach, Buchen and Kelley (1996) propose a maximum entropy estimate of the distribution. They also use the lognormal density as a prior and find that the resulting PDF will be the product of piece-wise uniform-exponential distributions.

Finally, Madan and Milne (1994) propose a finite Hermite polynomial expansion to estimate the PDF, generalizing on the Black-Scholes assumption of a single lognormal PDF. Although they do not recover the complete PDF, Corrado and Su (1996) follow a very similar approach to recover implied measures of skewness and kurtosis.

C. Existing comparisons⁵

Several studies have compared some of the methods for recovering the PDF. Bahra (1997) essentially directly implements equation (2) using butterfly spread prices as well as using the methods of Shimko (1993) and Melick and Thomas (1997). In the study he uses options prices from LIFFE (FTSE 100, Long Gilt, Euromark, Bund and Short-sterling) and the PHLX (exchange rates). However, a formal comparison of the techniques was beyond the scope of his paper.

Campa, Chang and Reider (1998) compare a modified version of the Shimko (1993) technique and the techniques of Rubinstein (1994) and Melick and Thomas (1997) using data from the OTC currency markets. They find that the three methods produce “remarkably similar PDFs” and note that each approach has its strengths and weaknesses. For the balance of the paper they report results using the method of Rubinstein (1994).

Jondeau and Rockinger (1997) compare the techniques of Melick and Thomas (1997), Madan and Milne (1994), Corrado and Su (1996) along with the stochastic process techniques of Malz (1996) and Heston (1993). They use data for two dates from the OTC market for options on the FF/DM exchange rate. They find the fit of the Melick and Thomas approach to be very good, but prefer the stochastic process approach of Malz (1996) for ease of interpretation.

Coutant, Jondeau and Rockinger (1998) compare the techniques of Melick and Thomas (1997), Madan and Milne (1994) and Buchen and Kelley (1996). They find that the three methods yield similar PDFs,

⁴ See Bahra (1997) and Campa, Chang and Reider (1998).

⁵ This note only discusses papers issued prior to the workshop. The paper presented at the workshop by Des McManus (contained in this volume) also compares several of the estimation techniques.

although they prefer the Hermite expansion of Madan and Milne (1994) for reasons of robustness and ease of calculation.

II. Interpreting PDFs

A. Why use PDFs to study asset prices?

A central motivation for computing PDFs from observed options data is to understand how the distribution implied by market prices differs from a theoretical distribution assumed *a priori*, which in finance is usually the lognormal. Typically, this divergence from the lognormal is meant to be of some qualitative importance and not merely superficial. For example, computation of the implied PDF may reveal that market expectations are in fact characterized by multiple modes, or a degree of skewness and kurtosis significantly different from that found in the lognormal distribution.

Computation of a PDF usually occurs in the context of a more general fundamental economic question, often relating to a possible change in regime, or other phenomena that would affect expectations before showing up in time series data. In certain papers [e.g. Rubinstein (1994), Ait-Sahalia and Lo (1995), Jackwerth and Rubinstein (1996)], a PDF is computed in order to characterize expected returns in the stock market, especially the probability of a crash or correction. In Melick and Thomas (1997), the PDF of future oil prices reveals the effects of the Gulf war in 1991 on the expected price of oil. Leahy and Thomas (1996) derive the PDF of the Canadian dollar-U.S. dollar exchange rate during the October 1995 referendum on Quebec independence. In these last two papers, the PDF is sometimes characterized by two modes corresponding to two political outcomes—war vs. peace, or independence vs. national unity. Campa, Chang, and Reider (1997) compute PDFs on key cross rates within the “Exchange Rate Mechanism” of the European Monetary System in order to determine the size of ERM bandwidths consistent with market expectations of exchange rate convergence/divergence. Campa, Chang, and Reider (1998) derive a number of exchange rate PDFs in order to study the relation between skewness and spot, with implications for whether exchange rates follow implicit target zones.

B. Data Limitations in Analyzing Large Potential Price Changes

A PDF-based approach to forecasting is especially useful for important potential qualitative changes in asset prices (e.g. market crashes or booms, currency devaluations, exchange rate or interest rate regime changes) that have influenced expectations without necessarily being detectable in time series data directly. Often these qualitative regime changes imply a price change of a magnitude that is large relative to day-to-day variation in the underlying asset price.

Yet, sometimes, the questions of greatest economic or policy interest—i.e. market expectations of large price movements associated with a regime shift, are in fact the most difficult to answer because of data limitations. In particular, one usually has the fewest observations of deep out-of-the-money or deep in-the-money options. Usually at-the-money options are the most actively traded, while options away-from-the-money are more illiquid and of less reliability. Synchronous observation of the option price and the spot price (e.g. at the day's close) become more difficult to coordinate, and idiosyncracies in supply and demand conditions (e.g. a large market participant suddenly needing to liquidate a position) play a greater role. For out-of-the-money options, bid-ask spreads become a higher percentage of the option premium, and thus distort the underlying economic price. Thus, data quality typically diminishes precisely in the regions of the most interest.

For certain markets, there are very few (if any) options trading with strike prices in the region of greatest economic or policy interest. When this occurs, the PDF in these regions of analysis become more art than science. Often, one's inferences in these outer regions of the distribution, unfortunately, depend more on one's choice of estimation method or smoothing technique rather than on the data itself. When estimating the PDF at any given point, observed prices of options with a higher and lower strike price provide a "reality check" on how one assigns probability to that region. For strike prices above the highest observed strike or below the lowest observed strike, it is impossible to use additional option data to verify empirically whether one's construction of the PDF is realistic. At best, one constructs a PDF that is consistent with available data, internally consistent, and economically sensible. Then, one must recognize the limitations of any inferences derived from the PDF in the regions based on extrapolation rather than interpolation of data.

C. Evaluating PDFs Empirically

One way of judging a PDF derived from option data is its empirical performance—i.e. its ability to predict realizations in the returns of the underlying asset. This is intrinsically a difficult proposition because a PDF represents the range of possible realizations, and in reality, there will be only a single realization of returns for any given forecast horizon. As long as that realization has positive probability in the PDF, then one cannot reject that PDF.

D. Stationarity and Aggregation over Time

One approach to increasing the number of observations is to aggregate over time, using multiple realizations to evaluate a given PDF. For example, one could compare a year's worth of one-month exchange rate returns to the distribution of returns implied by a cross-section of one-month options observed at an instant in time, say on January 1. Of course, a major drawback to this approach is the

implicit assumption of stationarity in the distribution. The density function observed on January 1 need not apply for the whole year, as expectations are almost certain to change—sometimes significantly—over time. Clearly, the briefer the time horizon of a given PDF, the more independent realizations of returns one will observe over time. Yet, options of very short time time-to-expiration (e.g. a few days) rarely trade with any liquidity, presumably because bid-ask spreads would be disproportionately large relative to the option premium.

E. Risk-neutral vs. Actual PDFs

Another difficulty in both the interpretation and evaluation of risk-neutral PDFs derived from option prices is the impossibility of distinguishing between actual probability, in a purely statistical sense, and the risk-neutral probability. A state that may have a relatively high probability in the risk-neutral density may in fact have a relatively low statistical probability of actually occurring but simply have a high valuation. For example, a stock market crash may have low statistical probability, but a dollar in that state of nature may be very highly valued (relative to a dollar in other states of nature). This will be reflected in the pricing kernel that transforms statistical probability into the risk-neutral probability. From a research point of view, however, it is impossible to disentangle the contribution of statistical probability and the contribution of relative marginal utility of different states.

In the absence of a full-scale economic model in which marginal utilities under different states of nature are made explicit, one can at best use economic intuition to identify qualitatively how the risk-neutral and actual distributions should be expected to differ. Most would agree that a stock market crash is a scenario in which a one dollar payout would be relatively highly valued. With other assets, it may be less clear. Should a given payout be more valuable when there has been a major dollar devaluation or a major yen devaluation? Would a dollar be worth more when interest rates have suddenly risen or suddenly fallen by 100 basis points? Answers to such questions will probably depend on a number of factors, including: wealth asymmetries (payoffs presumably have greater utility in lower-wealth states), policy asymmetries (one should derive higher utility from payoffs occurring when policy is less accommodative or stimulatory), or risk-appetite asymmetries (higher utility associated with payoffs in states of lower investor risk tolerance).

F. Peso Problems.

The presence of “peso problems” also complicates the empirical evaluation of PDFs. A small probability, large-magnitude event can have an important influence on the shape of a PDF, even if this event is not observed over a given finite sample period. The absence of that rare event in the data would not necessarily invalidate the PDF. The effect of peso problems is even stronger if the low-

probability event is associated with relatively high marginal utility (as may be the case with a stock market crash, or an exchange rate or interest rate regime change), causing the risk-neutral probability to appear higher than the statistically expected probability.

G. Key Advantages of PDF-Based Analysis

In spite of these complications, an analytical approach based on PDFs derived from option prices still has much to recommend over alternatives based on time series data. Most important, since they are based on market prices of options, PDFs are forward-looking. Thus, they are capable of incorporating a wide range of future eventualities that simply are not captured using historical data. They do not require a long historical time series in order to be estimated accurately, and furthermore are instantly capable of reflecting a change in market sentiment. A sudden shift in beliefs due to a political announcement or economic news could be immediately captured in option prices and the implied PDF. Second, they are well-suited to capturing the uncertainty inherent in financial markets, that of “multiple scenarios.” The shape of a potential distribution will depend on market data across multiple strike prices rather than on a mathematical function of the standard errors of an econometric regression. Third, PDFs are relatively free of mathematical priors imposed by a specific economic model or structure. While some parameterization is needed in order to map a finite number of data points into a smooth PDF, it has been shown that for certain regions representing a large percentage of total probability, key characteristics of the derived PDFs are relatively independent of the methodology used, suggesting some robustness to this approach. While there are dangers to over-inference from the derived PDF, they can be applied to virtually any financial market and still provide these key advantages over methods based solely on time series data.

References

- Aït-Sahalia, Yacine and Andrew Lo (1998): "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices". *Journal of Finance*, 53(2) (April) pp. 499-547.
- Bahra, Bhupinder (1997): "Implied Risk-Neutral Probability Density Functions from Option Prices: Theory and Application". *Bank of England Working Paper Series*, No. 66 (July).
- Bates, David S. (1991): "The Crash of '87: Was It Expected? The Evidence From The Options Markets". *Journal of Finance*, Vol. 46 (July) 1009-1044.
- Bates, David S. (1996a): "Dollar Jump Fears, 1984-1992: Distributional Abnormalities Implicit in Currency Futures Options". *Journal of International Money and Finance* 15(1) pp. 65-93.
- Bates, David S. (1996b): "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options". *Review of Financial Studies* 9(1) (Spring) pp. 69-107.
- Bodurtha, James N. and Martin Jermakyan (1999): "Non-Parametric Estimation of an Implied Volatility Surface". forthcoming *Journal of Computational Finance*
- Breeden, D. and R. Litzenberger (1978): "Prices of State Contingent Claims Implicit in Options Prices". *Journal of Business*, Vol. 51 (October) 621-651.
- Buchen, Peter W. and Michael Kelley (1996): "The Maximum Entropy Distribution of an Asset Inferred from Option Prices". *Journal of Financial and Quantitative Analysis* 31(1) (March) pp. 143-159.
- Campa, Jose M, P.H. Kevin Chang and Robert L. Reider (1997): "ERM Bandwidths for EMU and After: Evidence from Foreign Exchange Options". *Economic Policy*, 24 pp. 55-89
- Campa, Jose M, P.H. Kevin Chang and Robert L. Reider (1998): "Implied Exchange Rate Distributions: Evidence from OTC Option Markets". *Journal Of International Money And Finance*, 17(1) (February) pp. 117-160
- Corrado, Charles J. and Tie Su (1996): "Skewness and Kurtosis in S&P 500 Index Returns Implied by Option Prices". *Journal of Financial Research* 19(2) (Summer) pp. 175-192.
- Coutant, Sophie and Eric Jondeau and Michael Rockinger (1998): "Reading Interest Rate and Bond Futures Options' Smiles: How PIBOR and Notional Operators Appreciated the 1997 French Snap Election". Banque de France Working Paper (January).
- Cox, John C. and Stephan A. Ross (1976): "The Valuation of Options for Alternative Stochastic Processes". *Journal of Financial Economics* 3 pp. 145-166.
- Heston, Steven L. (1993): "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options". *Review of Financial Studies*, 6(2), pp. 327-345.
- Jackwerth, Jens C. and Mark Rubinstein (1996): "Recovering Probability Distributions from Options Prices". *Journal of Finance*, 51(5), (December) pp. 1611-1631.
- Jondeau, Eric and Michael Rockinger (1997): "Reading the Smile: The Message Conveyed By Methods Which Infer Risk Neutral Densities" Banque De France Working Paper (June).
- Leahy, Michael and Charles Thomas (1996): "The Sovereignty Option: The Quebec Referendum and Market Views on the Canadian Dollar". Board of Governors of the Federal Reserve System, International Finance Discussion Paper 555.
- Levin, Alexander (1998): "Recovering Implied Volatility and Distribution for CAD/USD Foreign Exchange Rate Using Regularization Method". Bank of Montreal (in preparation).
- Levin, Alexander, Des Mc Manus and David Watt (1999): "The Information Content of Canadian Dollar Futures Options", in *Information in Financial Asset Prices – Proceedings of a Conference*, May 1998, Bank of Canada pp. 229-275.

- Madan Dilip B. and Frank Milne (1994): "Contingent Claims Valued and Hedged by Pricing and Investing in a Basis". *Mathematical Finance* 4 pp. 223-245.
- Malz, Allan (1996): "Using Option Prices to Estimate Ex Ante Realignment Probabilities in the European Monetary System: The Case of Sterling-Mark". *Journal of International Money and Finance*, 15(5), (May) pp. 717-748.
- Malz, Allan (1997): "Estimating the Probability Distribution of the Future Exchange Rate from Option Prices". *The Journal of Derivatives* (Winter) pp. 20-36.
- Melick, William R. and Charles P. Thomas (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(1) March pp. 91-115.
- Mizrach, Bruce (1996): "Did Option Prices Predict the ERM Crises?" Rutgers University Department of Economics Working Paper #96-10 (April).
- Neuhaus, Holger (1995): "The Information Content of Derivatives for Monetary Policy". Discussion Paper 3/95, Economic Research Group of the Deutsche Bundesbank.
- Ritchey, Robert J. (1990): "Call Option Valuation for Discrete Normal Mixtures". *Journal of Financial Research*, 13(4) (Winter) pp. 285-296
- Rubenstein, Mark (1994): "Implied Binomial Trees". *Journal of Finance*, 49(3) (July) pp. 771-818.
- Sherrick, Bruce J., Philip Garcia and Viswanath Tirupattur (1996): "Recovering Probabilistic Information From Option Markets: Tests of Distributional Assumptions". *Journal of Futures Markets* 16(5) pp. 545-560.
- Shimko, David (1993): "Bounds of Probability", *Risk*, Vol. 6 (4) (April), pp. 33-37.
- Söderlind, Paul and L.E.O. Svensson (1998): "New Techniques to Extract Market Expectations From Financial Instruments". *Journal of Monetary Economics*, Vol. 40 (2) (October) pp. 383-429.