PROCEEDINGS

November 1999

ESTIMATING AND INTERPRETING PROBABILITY DENSITY FUNCTIONS

Proceedings of the workshop held at the BIS on 14 June 1999

BANK FOR INTERNATIONAL SETTLEMENTS
Monetary and Economic Department
Basel, Switzerland

Note:

The papers included in this volume are to be considered as working papers. They should be cited as working papers and considered preliminary drafts of any subsequent publication. They are reproduced here to make them easily available to anyone having an interest in the subject of the workshop. Although all papers have been screened for relevance to the subject matter of the workshop, they have not been subject to a rigorous refereeing process nor edited for form or content by the Bank for International Settlements.

Copies of publications are available from:

Bank for International Settlements
Information, Press & Library Services
CH-4002 Basel, Switzerland

Fax: +41 61 / 280 91 00 and +41 61 / 280 81 00

This publication is available on the BIS website (www.bis.org).

ISBN 92-9131-074-3

PROCEEDINGS

November 1999

ESTIMATING AND INTERPRETING PROBABILITY DENSITY FUNCTIONS

Proceedings of the workshop
held at the BIS on 14 June 1999

BANK FOR INTERNATIONAL SETTLEMENTS
Monetary and Economic Department
Basel, Switzerland

Table of contents

Participants in the meeting

Name	Affiliation	E-mail
Ms J Aguilar	Bank of Sweden	javiera.aguilar@riksbank.se
Mr J Berk	Netherlands Bank	j.m.berk@bportugal.pt
Mr R Bliss	Federal Reserve Bank of Chicago	robert.bliss@chi.frb.org
Mr V Brousseau	European Central Bank	vincent.brousseau@ecb.int
Mr J Campa	New York University	jcampa@stern.nyu.edu
Mr N Cooper	Bank of England	n.cooper@bankofengland.co.uk
Ms S Coutant	Bank of France	scountant@banque-france.fr
Mr T Flury	Swiss National Bank	flury.thomas@snb.ch
Mr G Gemmill	City University Business School	g.gemmill@city.ac.uk
Mr M Loretan	Board of Governors of the Federal Reserve System	loretanm@frb.gov
Mr J Barros Luís	Bank of Portugal	jhcluis@bportugal.pt
Mr J Mahoney	Federal Reserve Bank of New York	jim.mahoney@ny.frb.org
Mr A Malz	RiskMetrics	allan.malz@riskmetrics.com
Ms M C Manzano	Bank of Spain	cmanzano@bde.es
Mr D McManus	Bank of Canada	dmcmanus@bank-banque-canada.ca
Mr W Melick	Kenyon College	melickw@kenyon.edu
Mr H Neuhaus	European Central Bank	holger.neuhaus@ecb.int
Mr S Shiratsuka	Bank of Japan	shigenori.shiratsuka@boj.or.jp
Mr P Söderlind	Stockholm School of Economics	paul.soderlind@hhs.se
Mr C Upper	Deutsche Bundesbank	christian.upper@bundesbank.de
Mr R Violi	Bank of Italy	viol.4108@interbusiness.it
Mr R Wouters	Belgian National Bank	rafael.wouters@nbb.be
BIS: Mr R Filosa		renato.filosa@bis.org
Mr G Galati		gabriele.galati@bis.org
Mr E Remolona		eli.remolona@bis.org

Foreword

In recent years, central banks have increasingly used option markets to construct measures of market conditions and market participants' expectations. Most recently, techniques have been developed that use option prices to estimate or recover the entire expected distribution (probability density function, PDF) of future financial asset prices such as interest rates, exchange rates and equity prices. These PDFs allow for a more complete characterisation of the state of market expectations.

There are a number of different techniques currently used to estimate PDFs from option prices, and in some cases they have produced different results. There is at this point no consensus on which technique should be used in which situation. Moreover, opinion differs as to how PDFs should be interpreted. As the use of estimated PDFs has become increasingly popular in the central banking community, the BIS decided to organise a one-day workshop on estimation and interpretation of PDFs. The workshop was held in Basel on 14 June 1999 and organised by Gabriele Galati of the BIS and William Melick of Kenyon College, Ohio; it brought together experts from central banks, academia and the investment community. The background paper by Kevin Chang and William Melick provides an overview of the issues involved in the estimation and interpretation of PDFs.

The workshop was divided into two sessions. The first session addressed issues related to the estimation techniques. Before the workshop, participants received a common data set of settlement price data for options on eurodollar futures between 1 September 1998 and 30 November 1998. Participants were asked to estimate an implied PDF using their own technique for each trading day in the data set and to provide a standard set of summary statistics. The results, which are summarised in a note by William Melick, were discussed at the beginning of the first session.

Three papers were then presented. The paper by Neil Cooper (Bank of England) compared the accuracy of alternative estimation techniques using simulated distributions and applying a Monte Carlo test. The paper by Sophie Coutant (Bank of France) provided a theoretical framework for separating the risk-neutral density function from a risk-aversion function. She then estimated these functions using data on CAC options. Des McManus (Bank of Canada) applied alternative techniques to estimate PDFs from eurodollar futures and used different statistics to evaluate their performance.

The second session focused on applications of PDFs and on issues related to their economic interpretation. Gordon Gemmill (City University Business School) analysed the behaviour of PDFs estimated from FT-SE 100 option prices around a number of "crash episodes" and several election dates in the United Kingdom and evaluated their predictive content. José Campa (New York University), Kevin Chang (Credit Suisse First Boston) and James Refalo (New York University) used option data from Brazil to describe expectations of the real/dollar exchange rate and analysed the credibility of different exchange rate regimes. Finally, the paper by Jorge Barros Luís and Bernardino Adão (Bank of Portugal) used PDFs to evaluate expectations regarding interest rate convergence in Europe in the run-up to monetary union.

The BIS hopes that circulation of these conference proceedings will stimulate further discussion and research on implied PDFs. As the papers in this volume make clear, these PDFs offer an important new means of characterising market expectations. However, the papers also make clear the difficulties and caveats involved in interpreting these characteristics.

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 process[1], 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

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

where $e^{-r\cdot T}$ is the relevant discount factor. As shown by Breeden and Litzenberger (1978), $f(S)$ can be isolated by differentiating equation (1) twice, yielding[2]

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

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, the 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 comparisons[5]

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 Financ*e 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 Notianal 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.

William R Melick

"Results of the estimation of implied PDFs from a common dataset"

Results of the Estimation of Implied PDFs from a Common Dataset

William R. Melick

A part of the June 14, 1999 BIS workshop involved estimation of implied probability density functions (pdfs) by a number of participants using a common dataset. In early April, each participant received settlement price data for options on Eurodollar futures. These are American options that trade on the Chicago Mercantile Exchange. The data covered the 61 trading days from September 1, 1998 through November 30, 1998 for the December 1998 contract. The option strikes and the futures prices were subtracted from 100 (with calls redefined as puts and puts redefined as calls) in order that the probability density functions were estimated in terms of the more intuitive short-term interest rate as opposed to the discount price. These data were chosen for two reasons 1) this particular options market is among the most active in the world 2) the conference was attended by many central bank economists who for monetary policy reasons are eager to learn more about movements in short-term interest rates, and 3) the period was an active one, with the Federal Reserve lowering the federal funds rate by a total of 75 basis points from September through November.[1]

It is believed that the data are of a fairly high quality. Any option that had no open interest, exercises, or volume on a given day was excluded from the dataset. The remaining options were checked to ensure that they satisfied arbitrage restrictions involving monotonicity, slope, concavity, and put-call parity (within ranges that would result from the transactions costs involved in eliminating the arbitrage possibility). For the 61 trading days there was an average of a bit more than 25 option settlement prices per day with an average of roughly 18 unique strikes prices per day.

Each workshop participant was asked to estimate an implied PDF by whatever technique they desired for each of the 61 trading days. The participants were asked to provide a standard set of results, namely the mean, standard deviation, and 11 percentiles.[2] A total of 19 workshop participants submitted estimates, with 14 of those containing all percentiles for the 42 trading days from September 1, 1998 through October 30, 1998. These 14 "complete" submissions were used in the analysis to follow.

Within the 14 submissions, a variety of techniques were used to recover the PDF. Using the taxonomy developed in the background note, 5 of the submissions used some variant of Equation (2), either smoothing the volatility smile and then differentiating twice to recover the PDF or using finite difference methods on the option prices directly to recover the cumulative density function (CDF). Most of the

[1] The target federal funds rate was lowered by 25 basis points on three occasions - following the regularly scheduled FOMC meetings held on September 29, 1998 and November 17, 1998 and following a conference call meeting of the FOMC on October 15, 1998.

[2] The percentiles were 0.005, 0.001, 0.050, 0.100, 0.250, 0.500, 0.750, 0.900, 0.950, 0.990, 0.995 where percentile x is defined as the eurodollar rate such that there is an x chance of the eurodollar rate falling below that rate.

remaining submissions used a method based on Equation (1), with the great majority of these using the mixture of lognormals assumption. Finally, one submission recovered the PDF as a by-product of specifying that the interest rate futures price follow a jump diffusion process.

I. Dispersion Across Percentiles

Given this variety of techniques, a natural first question to consider is the extent to which the different techniques (and different estimation algorithims for a given technique) produce different results. For each of the 42 trading days, the median estimate across the 14 participants for each of the 11 percentiles was calculated. These median percentiles for each trading day were then subtracted from each participants' percentile estimates for that day to create a standard measure of dispersion for each of the 11 percentiles that could be meaningfully aggregated across time. Therefore, for each of the 11 percentiles a total of 588 (14 participants 42 days) deviations were calculated, providing a measure of the dispersion across the estimates. Chart 1 shows the plots for the 11 percentiles and is somewhat discouraging.[3] If all the techniques yielded identical estimates then there would be no deviation from the median estimate of the 11 percentiles for any of the 14 participants on any of the 42 trading days. This would result in 588 zeroes being plotted for each of the percentiles - Chart 1 would show just 11 points that would form a horizontal line at zero on the vertical axis.

Obviously the actual results are nowhere near the ideal of zero dispersion, but a pattern does emerge, namely the dispersion is greater in the lower and upper percentiles than around the 0.500 percentile. As an example, on one trading day one participants' estimate of the 0.005 percentile was almost 1.8 percentage points (180 basis points) away from the median estimate of that percentile. That is, if the median estimate for the 0.005 percentile were a 3-month Eurodollar futures rate of 3.80 percent, one of the participants' estimates for the 0.005 percentile was 2.0 percent Dispersion around the 0.500 percentile was much smaller, amounting to roughly 12 basis points below and 52 basis points above the median. For example, even on the trading day with the largest dispersion, all of the participants estimates for the 0.500 percentile fell within a range of 64 basis points. (To provide a sense of the magnitude of this dispersion the average estimate of the 0.500 percentile across the 42 trading days at roughly 5.1 percentage points.)

However, a closer examination of the estimates indicated that almost all of the large outliers from the 11 median percentiles came from a single participant.[4] Chart 1a re-plots the data from Chart 1 excluding this participant. As can be seen (the scales on the two charts are identical) the dispersion for the 13 remaining participants is much lower. The largest deviation from any of the median percentiles (again at the 0.005

[3] The plots for the 0.005 and 0.010 percentiles are very close together, as are the plots for the 0.990 and 0.995 percentiles.
[4] The participant was not using an unusual technique, in fact the participant was using the mixture of lognormals assumption, the most popular technique among the 14 complete submissions.

percentile) now amounts to 96 basis points compared to 180 basis points in Chart 1. The range of deviations from the 0.500 percentile now amounts to roughly 25 basis points compared to 64 basis points in Chart 1.

However, the question remains whether the dispersion in Chart 1a is significant in any sense. The answer surely depends on the purpose to which the PDF estimation is being applied. The results shown in Chart 1a indicate that between the 0.100 and 0.900 percentiles there is not that much difference across the techniques. That is, practitioners can have some confidence that the results they report are not overly sensitive to the particular method they use to estimate the PDF. Outside of these percentiles, the sensitivity to the technique increases dramatically. This increase can be a problem for some but not all applications. For example, an analysis for policy-making purposes that uses PDF estimation to provide a 90% confidence interval for market expectations of the future short-term interest rate will not be all that sensitive to the choice of PDF estimation technique. On the other hand, an analysis for a value-at-risk calculation that used PDF estimation to provide a measure of the future short-term interest rate below which there is less than a 1% chance of falling will be quite sensitive to the choice of PDF estimation technique.

Sensitivity of the tail percentiles to the choice of estimation technique is not surprising, given that these regions of the density have few, if any, actively traded options with strike prices in the region. As discussed in Melick and Thomas (1997), outside of the lowest and highest available strike prices there is an infinite variety of probability mass that can be consistent with the observed option prices. Put a little more precisely, for example, below the lowest strike option prices only reveal information about the combination of $\Pr[f < X_L] \cdot E[f \mid f < X_L]$, where f and X_L are the underlying price and lowest strike respectively. As the option price constrains only the value of the product there can be significant variation across the techniques in the two terms of the product. The reported percentiles are only related to $\Pr[f < X_L]$, just one term in the product, so a large dispersion in the tail percentiles across techniques might well be expected. That is, the observed option price provides information about the product, not the two terms, hence two methods could provide very different estimates of one of the terms so long as there were offsetting differences in the estimates for the other term.

A final source of variation relates to the difficult nature of the estimation problem. Those techniques that make use of equation (1) typically involve a nonlinear, constrained optimization. The particular solution algorithm (and the parameters involved in the algorithm such as step size and convergence tolerance) can have dramatic impacts on the estimated PDF. That is, two researchers who both use a mixture of lognormals assumption for the form of the PDF may arrive at much different conclusions depending on the optimization algorithm used. As discussed in McCullough and Vinod (1999) the variation across algorithms can be large, especially for the difficult estimations involved in a PDF recovery.[5] The same sort of problems obtain for those using techniques based on equation (2). The choice of analytic versus numerical derivatives will create variation in the percentiles, even if each researcher is generating a strike price option price mapping (e.g. volatility smile) in exactly the same way.

II. Effects of Large Events

The dataset also provides the opportunity to assess whether the dispersion in the estimated percentiles increases during periods of large changes in economic conditions. Over the period September 1, 1998 through October 30, 1998 the FOMC lowered the target federal funds rate by 25 basis points on two occasions, September 29 and October 15. The latter cut came as a great surprise to financial market participants. Charts 2 and 3 plot the range of estimates for each of the 11 percentiles on each of the 42 trading days with vertical lines indicating the dates of the changes in the target federal funds rate. Although there is an increase in the dispersion of the percentiles following each of the changes, the increases are not large relative to other increases that do not coincide with FOMC policy changes.[6]

Conclusion

A very preliminary analysis of the submissions to the common dataset exercise suggests that measures of tail probabilities are quite sensitive to the technique used to estimate a PDF from options prices. However, within the 10th and 90th percentiles, sensitivity to technique is much less of an issue. Finally, a shock to the underlying market does appear to increase the dispersion of the estimates of the PDF percentiles, although the increase is similar to increases seen on other dates where shocks are not readily identified.

References

McCullough, B.D. and H.D. Vinod (1999), “he Numerical Reliability of Econometric Software”. *Journal of Economic Literature* 37(2) June pp. 633-665.

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.

---

5 Future work on the estimates from this common dataset will compare the percentiles of a subset of the participants who are known to be using the same technique (pdf assumption) but different optimization packages and algorithms.

6 These charts include all of the participants, like Chart 1 but unlike Chart 1a.

Chart 1
Scatter Plot of Scaled Percentiles - Sept. 1 - Oct. 30
(Estimated Percentile - Median of Estimated Percentile)

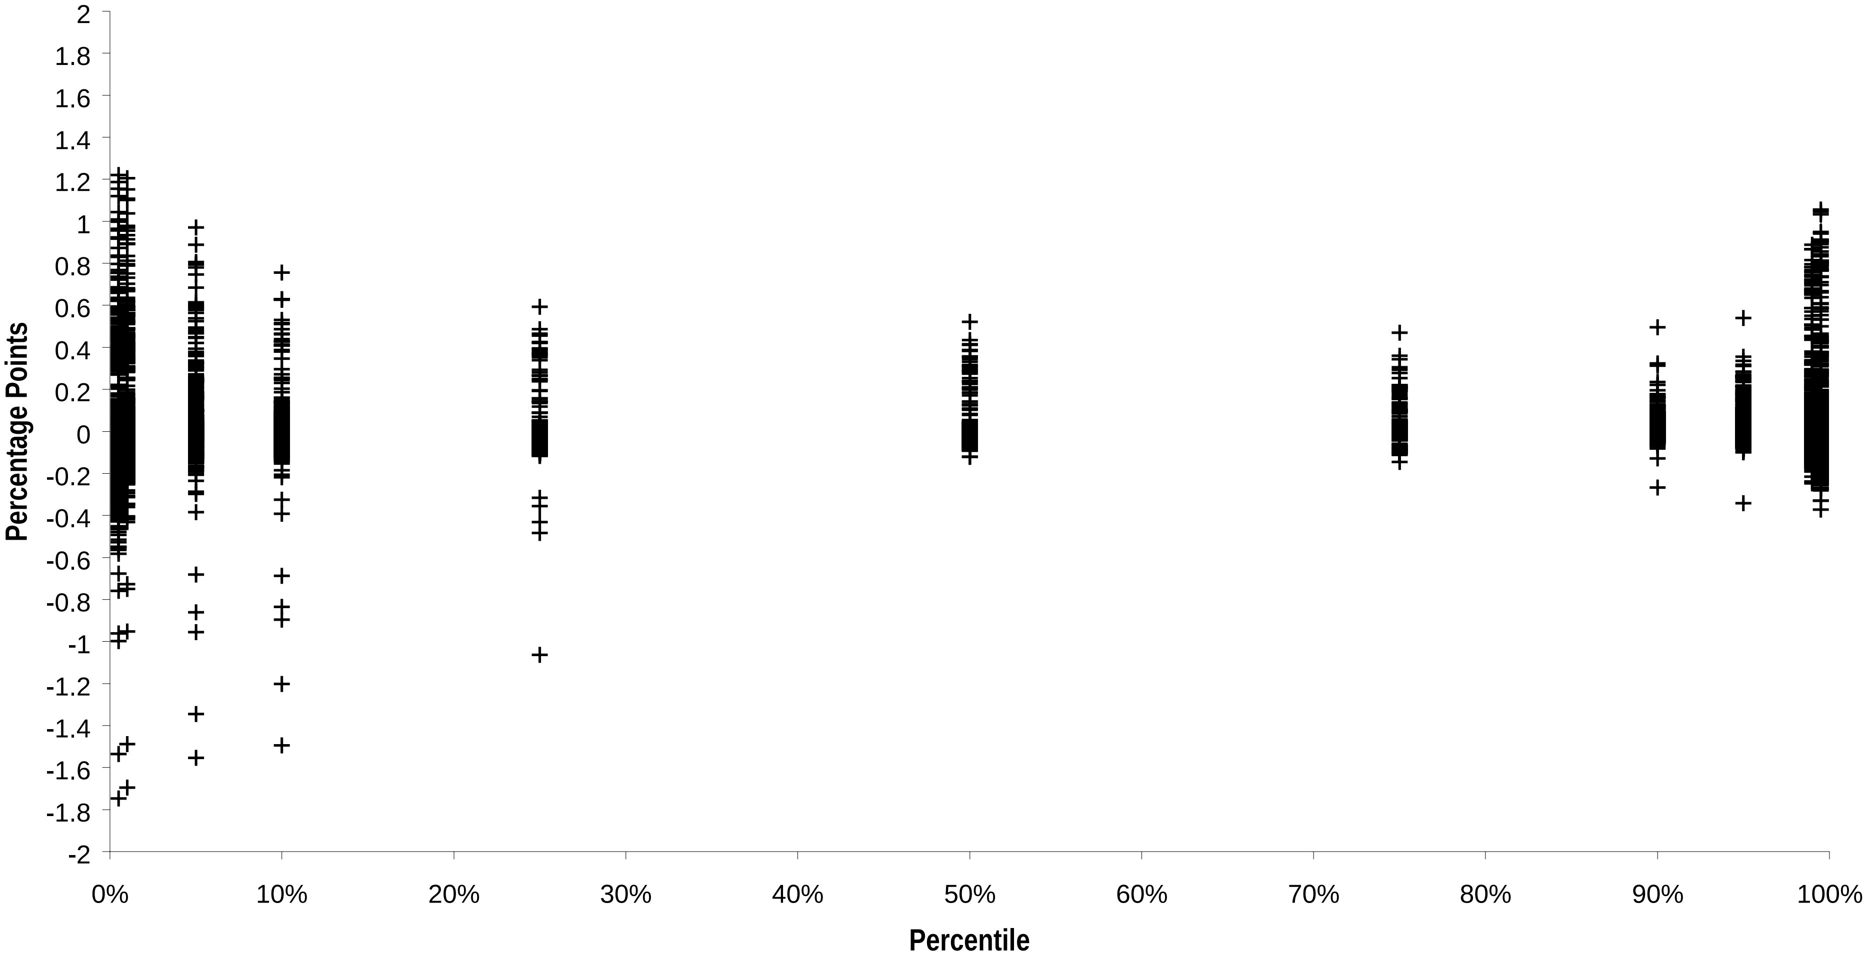

Chart 1a
Scatter Plot of Scaled Percentiles - Sept. 1 - Oct. 30 - Excluding One Participant
(Estimated Percentile - Median of Estimated Percentile)

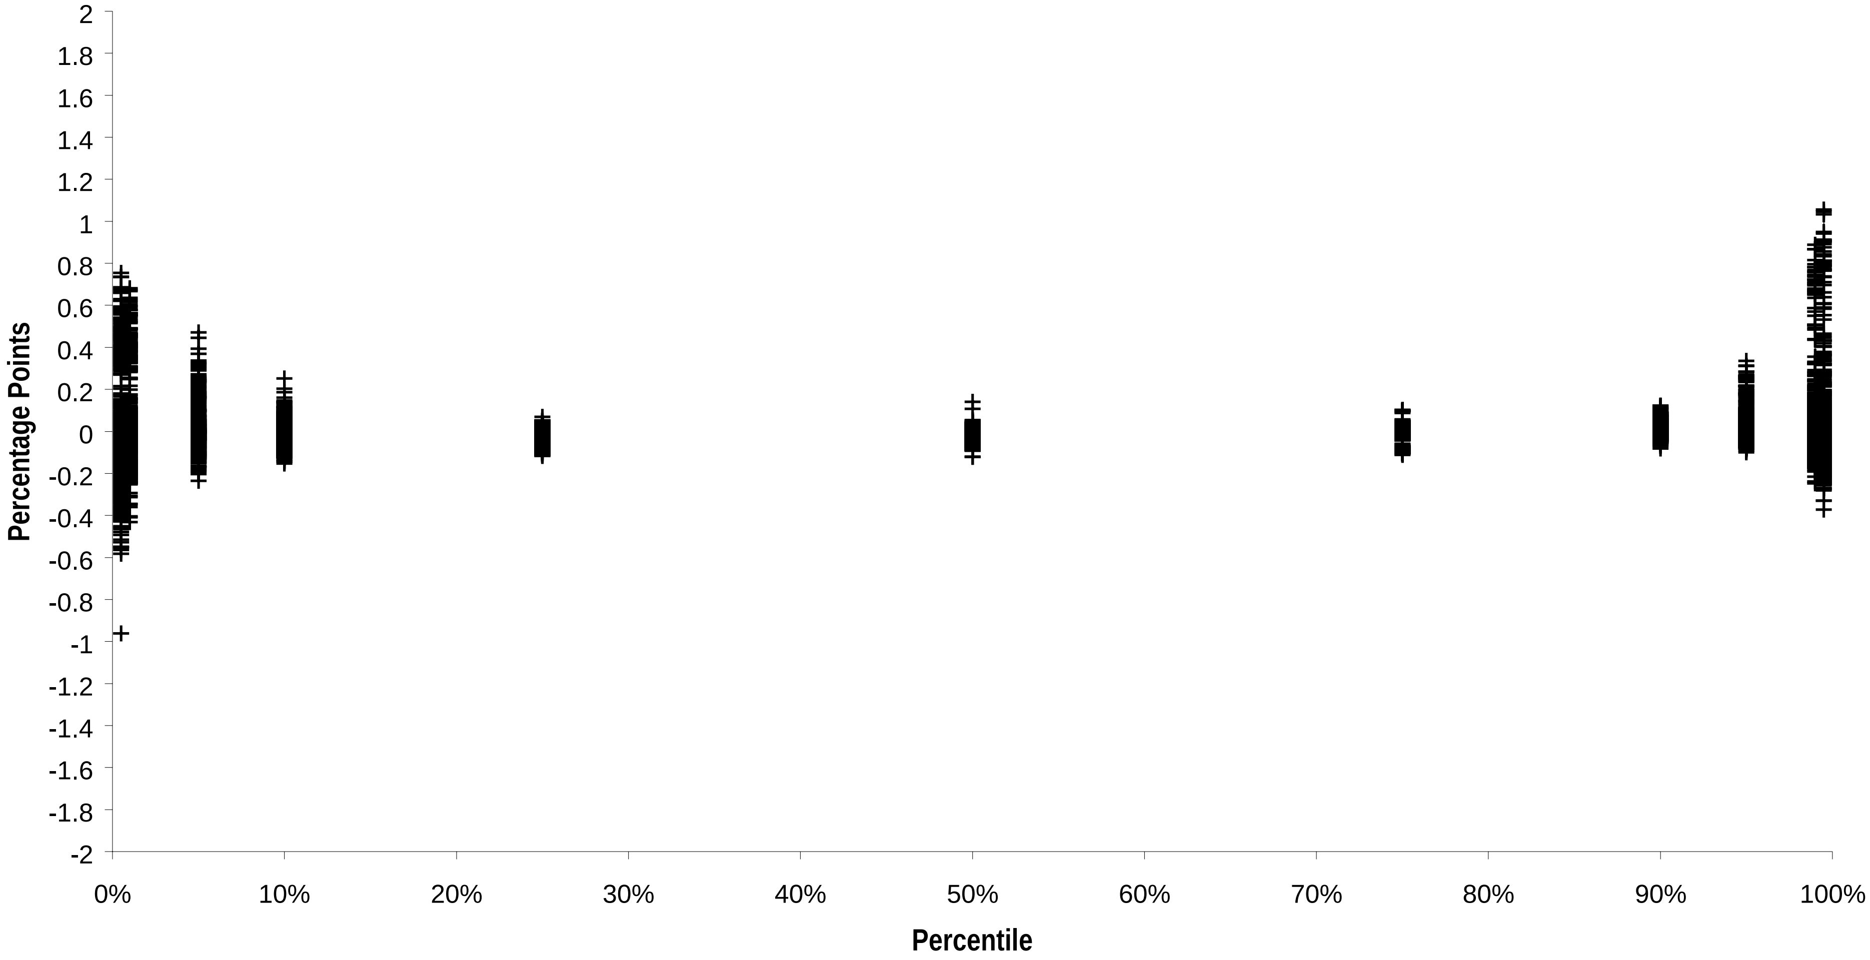

Chart 2
Range of Estimated Percentiles

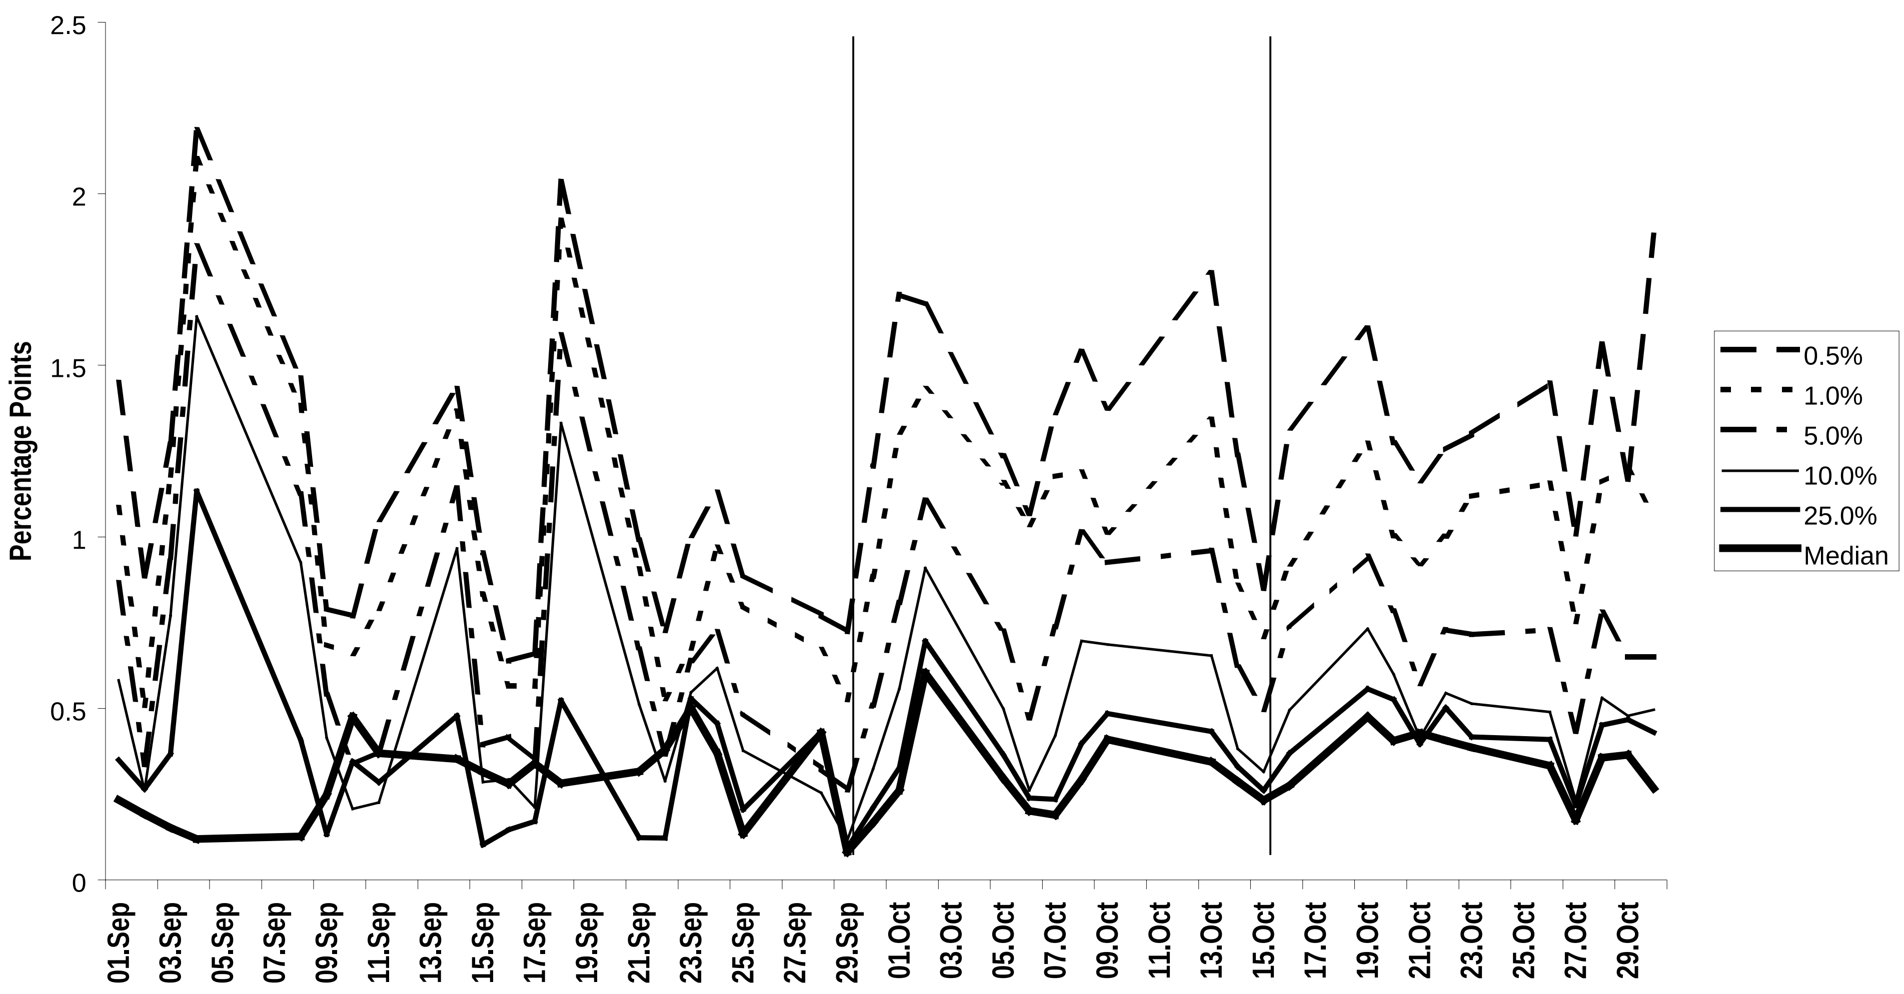

Chart 3
Range of Estimated Percentiles

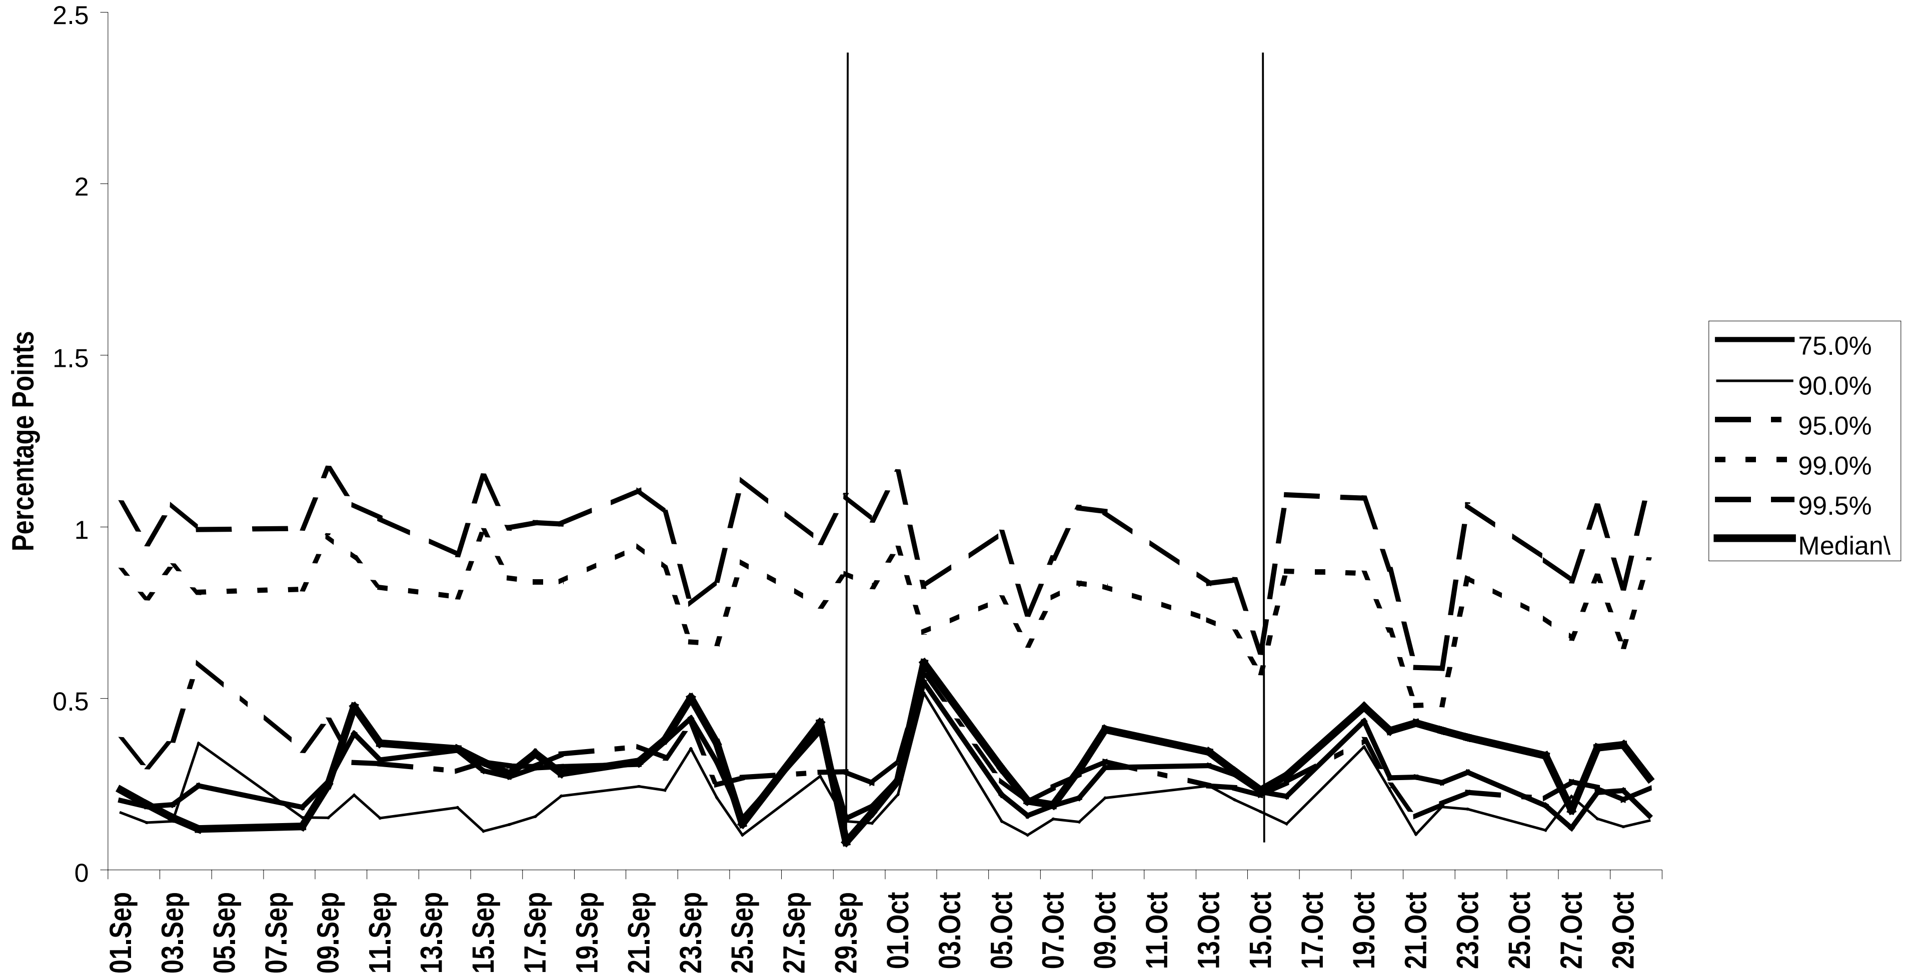

Neil Cooper

"Testing techniques for estimating implied RNDs from the prices of European-style options"

Discussants:

Holger Neuhaus

Jan Marc Berk

TESTING TECHNIQUES FOR ESTIMATING IMPLIED RNDS FROM THE PRICES OF EUROPEAN-STYLE OPTIONS

Abstract:

This paper examines two approaches to estimating implied risk-neutral probability density functions from the prices of European-style options. It sets up a monte carlo test to evaluate alternative techniques' ability to recover simulated distributions based on Heston's (1993) stochastic volatility model. The paper tests both for the accuracy and stability of the estimated summary statistics from RNDs. We find that a method based on interpolating the volatility smile out-performs the commonly used parametric approach that uses a mixture of lognormals.

July 2nd 1999

Neil Cooper
Bank of England

The author would like to thank Robert Bliss, Nikolaos Panigirtzoglou, James Proudman, Paul Soderlind, Gary Xu and participants (especially the discussants Jan Marc Berk and Holger Neuhaus) at the BIS workshop on implied PDFs for helpful comments. The views expressed in this paper are those of the author and do not necessarily reflect those of the Bank of England.

1. Introduction

In the last five years, there has been great interest amongst policy-makers in extracting information from the prices of financial assets. Options prices, in particular, have proved to be be a particular rich source of information since they enable the extraction of a complete implied risk-neutral probability density function (RNDs) for the assets, interest rates and commodity prices upon which they trade. These RNDs have proven particularly useful in interpreting the market's assessment of the balance of risks associated with future movements in asset prices.

Reflecting this interest, a relatively large number of papers have been published that set out alternative techniques for the estimation of implied RNDs with examples of their application to particular markets. Despite this wide range of papers, nearly all are based on one of three basic approaches:

- estimating the parameters of a particular stochastic process for the underlying asset price from options prices and constructing the implied RND from the estimated process - see Malz(1995) and Bates(1996) for examples that incorporate jump processes;

- fitting a particular parametric functional form for the terminal asset price, for example a mixture of lognormals directly to options prices - see Bahra (1996,1997) and Melick and Thomas (1997);

- interpolating across the the call pricing function or the volatility smile, following Shimko (1993), and employing the Breeden and Litzenberger (1978) result that the implied distribution may be extracted by calculating the second partial derivative of that function with respect to the strike price.

The first approach has the disadvantage that it is based on a particular stochastic process: we cannot observe whether the assumed process can capture the density functions that are implicit within options' prices. In this paper we focus on the second and third approaches which are more flexible since by trying to estimate the density function directly they are consistent with many different stochastic processes.

Given these alternative techniques, a natural question is: "Which technique performs the best?" A key concern is the accuracy and stability of the estimated RNDs. Suppose we observe an estimated RND that displays bi-modality or "spikes."[1] Should we interpret this as reflecting actual expectations or estimation errors? If we believe it to be the latter then the value of using implied RNDs is seriously diminished.

This paper attempts to address these concerns. It examines the empirical performance of two approaches to RND estimation by testing the ability of alternative techniques' ability to recover the implied density function from a set of simulated prices. The simulated prices are generated from a quite general stochastic volatility model set out in Heston (1993). By using simulated prices, rather than actual prices, we can compare estimated RNDs against the "true" RND implied by the underlying price process. We test not just the stability of estimated RNDs and their robustness to small errors as in Bliss and Panigirtzoglou (1999), but also their ability to closely recover the summary statistics from the true density function given sufficient data.

The paper is organised as follows. Section two sets out the two estimation techniques that we compare. Section three sets out the approach we will use for assessing the performance of the alternative methodologies. Section four presents results for European-style options and section five concludes.

2. Alternative Techniques for Estimating Implied RNDs from Options' Prices

2.1 Underlying Economics

In this section we examine the two estimation approaches that are tested within this paper. Both may be derived from the Cox and Ross (1976) pricing model. This model gives current time t European-style call option prices as the risk-neutral expected payoff of the option at expiry T, discounted back at the risk-free rate:

$$C(S, X, t) = e^{-rt} \int_{X}^{\infty} \left(S_T - X\right) g(S_T) dS_T \qquad (1)$$

[1] See Bahra(1996) for examples of such spiked distributions when using the mixture of lognormals approach.

where S_T is the terminal underlying asset price at T, $g(S_T)$ is its RND, X is the strike price and r and τ=T-t are the risk-free rate and the maturity of the option respectively. The put price can be recovered either through put-call-parity or by replacing the payoff of the call (S_T-X) with the payoff of the put (X-S_T) in the above formula and by integrating from zero to the strike price.

The first estimation approach tested in this paper involves specifying a particular parametric functional form for the RND $g(S_T)$ and fitting this distribution to the observed range of strike prices via non-linear least squares. Although a range of functional forms have been suggested, the most commonly used is a mixture of two lognormals[2]. The form chosen should be sufficiently flexible to capture the features of distributions that we might expect to find implicit within the data - excess kurtosis, either positive or negative skewness, and perhaps bi-modality. The mixture of lognormals is parsimonious because it matches these criteria with just five parameters to be estimated.

The mixture lognormal is given by:

$$g(S_T) = qL(a_1, b_1) + (1-q)L(a_2, b_2) \qquad (2)$$

where q, a_1, a_2, b_1, b_2 are the parameters to be estimated. The fitted call and put prices are given by[3]:

$$\hat{C}(S, X_i, t) = e^{-rt} \int_{X_i}^{\infty} (S_T - X_i)\big(qL(a_1, b_1) + (1-q)L(a_2, b_2)\big) dS_T \qquad (3)$$

$$\hat{P}(S, X_i, t) = e^{-rt} \int_{0}^{X_{i1}} (X_i - S_T)\big(qL(a_1, b_1) + (1-q)L(a_2, b_2)\big) dS_T \ .$$

To fit the parameters of the RND we minimise the following:

$$\min_{a_1, b_1, a_2, b_2, q} \sum_{i=1}^{m} \left(\hat{C}_{i,t} - C_{i,t}\right)^2 + \sum_{j=1}^{n} \left(\hat{P}_{i,t} - P_{i,t}\right)^2 \qquad (4)$$

[2] See Bahra (1996,1997) and Melick and Thomas (1997)
[3] As explained in Bahra (1997) for the futures options traded at LIFFE that have futures-style margining, the discount factor disappears.

The second approach to estimating implied RNDs that we test here which we term the “smile” approach, exploits the result derived by Breeden and Litzenberger (1978) that the RND can be recovered by calculating the second partial derivative of the call-pricing function with respect to the strike price. This result can be derived simply by taking the second partial derivative of equation (1) with respect to the strike price to get:

$$\frac{\partial^2 C}{\partial X^2} = e^{-rt} g\left(S_T\right) \qquad (5)$$

So we just have to adjust up the second partial derivative by exp($r\tau$) to get the RND g(S_T). In practice we only have a discrete set of strike prices. So to obtain an estimate of the continuous call-pricing function we need to interpolate across the discrete set of prices. Following Shimko(1993) this interpolation can be done by interpolating across the volatility smile and using Black-Scholes to transform this back to prices. The reason for doing this rather than interpolating the call-pricing function directly is that it is difficult to fit accurately the shape of the latter. And since we are interested in the convexity of that function any small errors will tend to be magnified into large errors in the final estimated RND.

Shimko (1993) used a quadratic functional form to interpolate across the volatility smile. Instead, we follow Bliss and Panigirtzolglou (1999) and use a cubic smoothing spline to interpolate in a similar way to Campa and Chang (1998). This is a more flexible non-parametric curve that gives us control on the amount of smoothing of the volatility smile, and hence the smoothness of the estimated RND. But following Malz (1997), Bliss and Panigirtzoglou (1999) also first calculate the Black-Scholes deltas of the options and use delta rather than strike to measure the money-ness of options. In practice this makes interpolation of the volatility smile even easier, since it becomes a simpler shape to approximate in “delta-space”. Finally, to generate the implied RND we calculate the second partial derivative with respect to strike price numerically as for (5) and adjust for the effect of the discount factor.

So summarising, estimation via the smile-based approach proceeds by:

- calculating implied volatilities of the call and put options;
- calculating the Black-Scholes deltas of the options using those implied volatilites;

- constructing the volatility smile by joining the implied volatilities for out-of-the-money calls with those of the out-of-the-money puts[4];
- interpolating across the volatility smile in "delta-space" via a cubic smoothing spline;
- transforming back to a price function using the Black-Scholes model;
- taking the second partial derivation of that function with respect to strike and adjusting for the discount factor within equation (5) to generate the final estimated RND.

3. A Monte Carlo Approach to Testing PDF Estimation Techniques

This section of the paper explains the testing procedures we will use to assess the performance of the two estimation approaches set out above. One approach to testing these techniques is to examine how closely they fit actual options data (for example see the approaches taken by Campa and Chang (1998), Jondeau and Rockinger (1998) and Bliss and Panigirtzoglou (1999). But in doing so it is difficult to assess which of the estimated RNDs most closely match the true risk-neutral density since this is unobservable. In the absence of knowledge of what the true density function is, it is difficult to judge this.

Instead we use simulated artificial options price data. We can simulate options prices that correspond to a given risk-neutral density function and see whether the estimation techniques can recover the RND. In addition following Bliss and Panigirtzoglou (1999), we also test whether the estimation technique is robust to small errors in prices that might result in the real world from the existence of discrete tick size intervals.

Any good RND estimation technique should be able to recover the true RND under a wide range of market conditions: that is conditions of high and low volatility; where the true density function has either positive or negative skews; and where we use options across the full range of maturities that are encountered in practice - anything from one week out to a year. So we need a way of generating options data that match this range of conditions.

[4] We use out-of-the-money options because traded volumes concentrate on at-the-money and out-of-the-money options. Also the out-of-the-money option value is composed enitirely of the time value of the option rather than its intrinsic value as for in-the-money options. It is the time value of the options only that reflects the shape of the RND.

To generate sufficiently interesting “true” risk-neutral densities that incorporate the features discussed above, we use Heston’s (1993) stochastic volatility model to generate prices. For European options, this model has a closed form solution. Under Heston’s model, the underlying asset price dynamics are described by the following stochastic differential equations:

$$\begin{aligned} dS &= mSdt + \sqrt{v_t}\,Sdz_1 \\ dv_t &= k(q - v_t)dt + s_v\sqrt{v_t}\,dz_2 \end{aligned} \tag{12}$$

Here the volatility of the underlying asset $\sqrt{v_t}$ is also stochastic. The conditional variance v_t follows a mean reverting process such that the volatility mean-reverts to a long run of $\sqrt{q}$ at a rate dictated by k. The term s_v sets the volatility of the volatility. Finally, the two Wiener processes dz_1 and dz_2 have a correlation given by r. By changing the correlation parameter we can generate skewness in asset returns. Suppose we have a negative correlation between shocks to the asset price and volatility. This means that as we get negative shocks to the price, volatility will tend to increase. This increase in volatility then increases the chance that we can get further large downwards movements. Thus a negative correlation can generate negative skewness in the unconditional distribution of returns. This will be reflected in a downwards volatility smile in the options generated under these parameters. A positive correlation between volatility and the asset price has the opposite effect[5].

Heston shows that for European call options[6] on assets that behave according to (12) it is possible to calculate prices with the following formula:

$$\begin{aligned} C(S, v_t, X) &= SP_1 - Xe^{-rt}P_2 \\ P_j(y, v_t, T; \ln(X)) &= \frac{1}{2} + \frac{1}{p}\int_0^\infty \mathrm{Re}\left[\frac{e^{-i f \ln(X)} f_i(y, v, T; f)}{i f}\right] d f \end{aligned} \tag{13}$$

where X is the strike price, y=ln(S), $i = \sqrt{-1}$,

[5] See Das and Sundaresan (1998) for more details on the relation between conditional skewness and kurtosis and the parameters of this stochastic volatility model.
[6] Put prices can be generated simply via put-call-parity.

$$f_j\left(y, v_t, t; f\right) = e^{C(t,f)+D(t;V)v_t + i fy},$$

$$C(t;f) = r\,fi\,t + \frac{a}{s_v^{\,2}}\left\{\left(b_j - r s_v\,fi + d\right)t - 2\ln\left[\frac{1-ge^{dt}}{1-g}\right]\right\},$$

$$D(t;f) = \frac{b_j - r s_v\,fi + d}{s_v^{\,2}}\left[\frac{1-e^{dt}}{1-ge^{dt}}\right],$$

$$g = \frac{b_j - r s_v\,fi + d}{b_j - r s_v\,fi - d},$$

$$d = \sqrt{\left(r s_v\,fi - b_j\right)^2 - s_v^2\left(2u_j\,fi - f^2\right)}.$$

and a=κθ, b_1=κ + λ - ρσ_v, b_2=κ + λ.

To generate the true density function and its associated summary statistics we simply apply equation (5) to (13). Figures 1 and 2 show the effect of changing ρ on the terminal asset price distribution and on the volatility smile for options generated under this model with current and long run volatility of 30%, mean reversion κ=2 and volatility of volatility σ_v of 40%. We can see that the Heston model can generate the sorts of shapes of both the volatility smile and the underlying asset distributions that can be observed in the real world.

Figure 1: Implied RNDs Under Alternative Correlation Parameters

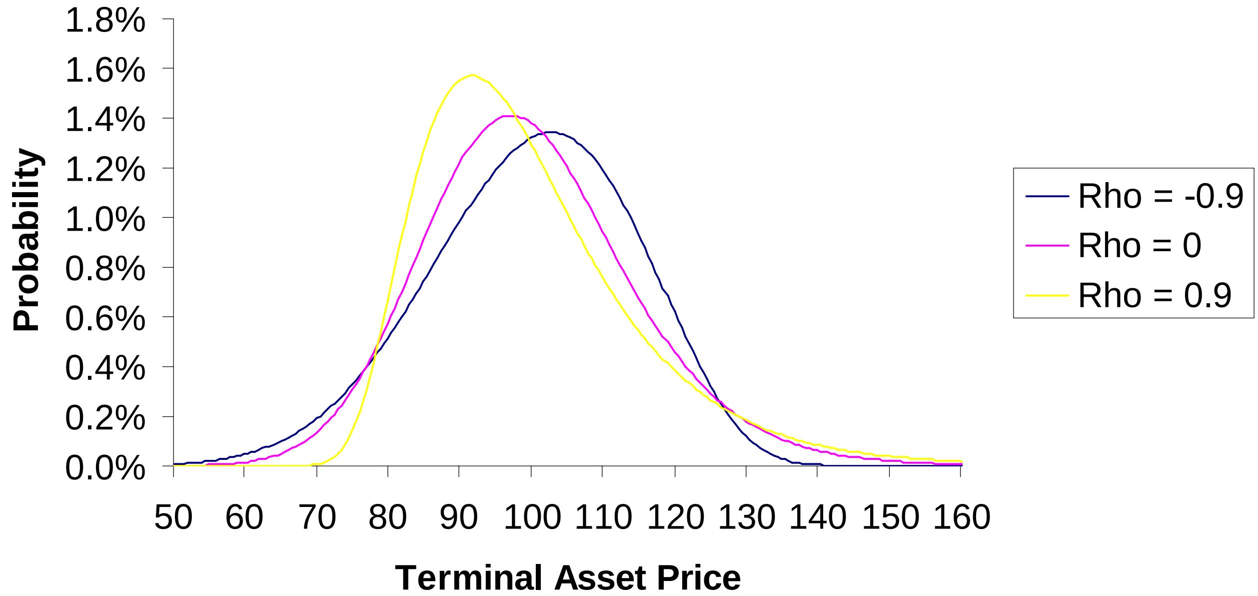

Figure 2: Volatility Smile Under Alternative Correlation Parameters

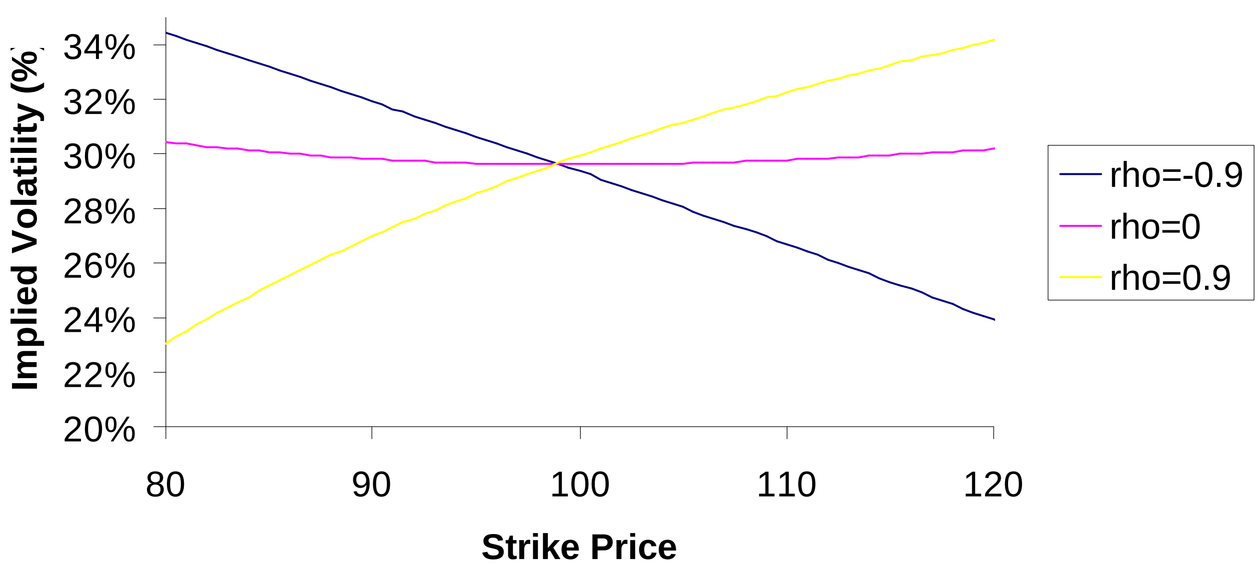

An additional feature of the real world that we want to incorporate is the existence of errors that are the result of discrete tick size intervals (and possibly any small violations of arbitrage within the settlement prices used for estimation). We want our estimation methodology to be robust to these small errors in the prices. So we perform the following test of the two RND estimation techniques.

We first establish a set of six scenarios corresponding to low and high volatility and three levels of skewness. For each scenario we generate a set of options prices with strikes ranging from 30% out-of the-money to 40% in-the-money. Then for each combination of scenario and maturity we use the approach developed by Bliss and Panigirtzoglou (1999) to first shock each price by a random number uniformly distributed from -1/2 to +1/2 a "tick size". This tick size was chosen as 0.05 to reflect the sorts of tick sizes that are typically found for exchange-traded options. Given these shocked prices we fit RNDs using the two techniques described in section two and calculate the summary statistics. We repeat this procedure of shocking the prices and then fitting the RNDs 100 times for each scenario and maturity combination. Finally we calculate in each case the mean and standard deviation of the calculated summary statistics and the squared pricing errors. In essense, this technique simply amounts to a monte carlo test of the finite sample properties of the two estimators of the sort that is commonly used within standard econometrics - see Greene (1997) Ch.5 or Davidson and McKinnon (1993) Ch. 18.

We then assess the two techniques by comparing the mean estimated summary statistics with the true summary statistic. We are looking for a technique that has both mean estimates of the statistics that are close to the true ones and one that has small standard deviations for the calculated statistics in the presence of the small errors within the options prices used i.e. it is stable. We also want an estimation procedure that performs well across the range of scenarios and maturities. The next section performs these tests for European-style options.

4. Results

This section includes the results of the tests that we described above for the two estimation approaches. As described above, we test performance across a range of six scenarios. The Heston model parameters used for each scenario are set out in table 1 below. These were chosen to generate true RNDs that corresponded to situations of negative skewness, and weak and strong positive skewness in the terminal asset price and also conditions of low and high volatility. To generate these differing levels of skewness in the terminal asset price distributions, we use three different levels of the correlation parameter -0.9, 0 and 0.9. The long run volatilities of 30% for the high volatility scenarios were chosen on the basis of the levels of implied volatility typically seen within equity markets. The low volatility (10%) scenarios can be thought of as consistent with levels often seen within FX and interest rate markets[7].

Table 1: Model parameters used under each scenario

	Strong Negative Skew		Strong Positive Skew
Low Volatility	Scenario 1 κ=2, θ^2=0.1, σ_v=0.1, ρ=-0.9	Scenario 2 κ=2, θ^2=0.1, σ_v=0.1, ρ=0	Scenario 3 κ=2, θ^2=0.1, σ_v=0.1, ρ=0.9
High Volatility	Scenario 4 κ=2, θ^2=0.3, σ_v=0.4, ρ=-0.9	Scenario 5 κ=2, θ^2=0.3 σ_v=0.4, ρ=0	Scenario 6 κ=2, θ^2=0.3, σ_v=0.4, ρ=0.9

We test the performance of the two estimation techniques under each of these scenarios across four different maturities - 2 weeks, 1 month, 3 months and 6 months. For each scenario and maturity pairing we first generate the true RND and calculate their summary statistics - their mean, standard deviation, skewness (the third central moment) and kurtosis (the fourth central moment). Table 2 sets out the true summary statistics for all the maturity and scenario combinations.

[7] We also assume that the market price of volatility risk is zero and that the time t conditional volatility is equal to the long run volatility.

Table 2: True Summary Statistics

	Scenario	2 weeks	1 month	3 month	6 month
Mean	1	100.000	100.000	100.000	100.000
	2	100.000	100.000	100.000	100.000
	3	100.000	100.000	100.000	100.000
	4	100.000	99.999	99.999	99.994
	5	99.999	100.000	99.999	99.998
	6	99.988	99.997	99.995	99.961
Std Dev	1	1.958	2.878	4.956	6.966
	2	1.962	2.888	5.004	7.081
	3	1.965	2.898	5.052	7.201
	4	5.849	8.555	14.524	20.099
	5	5.888	8.677	15.093	21.485
	6	5.921	8.799	15.687	22.900
Skewness	1	-0.198	-0.280	-0.418	-0.474
	2	0.060	0.089	0.159	0.231
	3	0.318	0.459	0.743	0.956
	4	-0.166	-0.228	-0.301	-0.265
	5	0.180	0.272	0.504	0.756
	6	0.523	0.776	1.346	1.840
Kurtosis	1	3.037	3.081	3.178	3.221
	2	3.038	3.086	3.221	3.355
	3	3.160	3.344	3.930	4.599
	4	2.984	2.962	2.878	2.748
	5	3.119	3.269	3.809	4.616
	6	3.426	4.036	6.272	9.026

For all combinations the futures price has been set at 100, so the true mean of the distributions are equal to 100[8]. As we would expect, the standard deviation of the true RNDs increases with maturity and as volatility is increased. Scenarios 1 and 4 which have a negative correlation between the underlying asset price and volatility display negative skewness in the terminal asset price distribution. Except for scenario four[9], the kurtosis of the terminal asset price distribution is greater than three and increases with maturity.

[8] The slight differences from 100 are caused by error in the numerical integration used to calculate the summary statistics.

[9] For this scenario the combination of a high volatility of volatility and the negative correlation between volatility and the asset price appears to reduce the probabilities attached to extreme outcomes.

Given these true summary statistics, we proceed to test the two estimation approaches to recover them from simulated options prices. We use equation (13) to generate European-style call and put futures options for all the scenario and maturity pairings for strikes ranging from 70 to 140 with strikes spaced at intervals of one apart from each other. For each pairing we then test the two techniques by shocking the prices for the tick size errors as described above and estimating RNDs and their summary statistics. This is repeated a hundred times for each pairing. Then for each pairing, and each summary statistic we calculate a measure of the mean estimate and measures of the estimates' stability (the standard deviation, and the distance between the five and ninety five percentiles) from the sets of a hundred estimated summary statistics under each technique. A good technique should have a mean estimate of each of the summary statistics that is close to the true ones and so may be said to be unbiased. A low standard deviation of the estimated summary statistics across the full range of scenarios and maturities indicates that the estimation technique is stable in the presence of small errors wthin the prices.

Table 3 below gives the mean estimated summary statistics for both approaches across all the different scenario and maturity pairings. To assess the unbiasedness of the two estimators, however, we are interested in the difference between the true summary statistics and the mean estimates from each approach. So table 4 calculates the difference between the true and the mean estimated statistics as a percentage of the true value of the summary statistic: (true-mean)/true.

Examining the top two panels of table 4 we can see that the mean of the smile approach is always almost always exactly equal to the true mean. The reason for this accuracy is that when we transform from the volatility smile to the pricing function using Black-Scholes this constrains the mean of the RND to be equal to the forward rate. The mixture lognormal estimation as described by equation (4) does not explicitly constrain the mean of the RND to be equal to the forward rate, so we get small errors between the true and actual means. This could be eliminated if we added an extra constraint to (4) to ensure that the mean equalled the forward rate as described in Bahra (1997), but this may come at the cost of extra instability in the fitted RNDs.

Table 3: Mean Estimated Summary Statistics

Smile Technique

	Scenario	2 weeks	1 month	3 month	6 month
Mean	1	100.0000	100.0000	100.0000	100.0000
	2	100.0000	100.0000	100.0000	100.0000
	3	100.0000	100.0000	100.0000	100.0000
	4	100.0000	100.0000	100.0000	100.0000
	5	100.0000	100.0000	100.0000	100.0000
	6	100.0000	100.0000	100.0000	100.0000
Std Dev	1	1.9595	2.8787	4.9588	6.9676
	2	2.0615	2.9891	5.0970	7.1651
	3	1.9685	2.9035	5.0581	7.2073
	4	5.8485	8.5588	14.5087	20.0582
	5	5.8905	8.6778	15.0693	21.3463
	6	5.9316	8.8035	15.5988	22.6372
Skewness	1	-0.1848	-0.2654	-0.3915	-0.4466
	2	0.1968	0.1898	0.2081	0.2548
	3	0.3132	0.4421	0.7130	0.9128
	4	-0.1521	-0.2076	-0.2578	-0.1960
	5	0.1784	0.2692	0.4798	0.6825
	6	0.5124	0.7505	1.2121	1.5861
Kurtosis	1	3.0163	3.0363	3.0701	3.0744
	2	3.8713	3.6523	3.4221	3.3638
	3	3.1615	3.2812	3.6934	4.1546
	4	2.9720	2.9415	2.8054	2.6707
	5	3.0781	3.1813	3.5137	3.9345
	6	3.3829	3.8291	5.1522	6.8006

Mixture Lognormal Technique

	Scenario	2 weeks	1 month	3 month	6 month
Mean	1	100.1680	100.0450	100.0312	100.0525
	2	99.9950	99.8531	99.8428	100.0006
	3	101.9543	100.4185	100.0332	99.9818
	4	99.8328	99.6520	100.7073	99.8851
	5	100.0027	100.0113	100.0024	100.0000
	6	100.2045	98.8045	99.9303	99.8592
Std Dev	1	2.6334	2.8864	4.9503	6.9544
	2	1.9732	2.9732	5.0552	7.1094
	3	7.4278	3.7705	5.0581	7.1893
	4	5.8645	8.5775	15.0537	20.1169
	5	5.8861	8.6659	15.0867	21.4762
	6	5.9116	9.4515	15.5956	22.5237
Skewness	1	-0.0585	-0.2709	-0.3695	-0.4148
	2	0.0594	0.1232	0.2165	0.2408
	3	0.1855	0.4236	0.6167	0.8436
	4	0.1764	0.2543	-0.2479	-0.1381
	5	0.1776	0.2603	0.4995	0.7504
	6	0.1821	0.7642	1.1701	1.4949
Kurtosis	1	3.0942	3.0672	3.0629	3.0668
	2	3.0720	3.0711	3.1870	3.4901
	3	2.3902	3.2004	3.5323	4.0006
	4	3.0554	3.1175	2.7957	2.7313
	5	3.0848	3.1811	3.7677	4.5488
	6	3.2202	3.6502	4.9617	6.0907

The second set of panels gives the results for the estimated standard deviations. For scenarios 4 to 6 and for all scenarios with maturities of three months and above the mean estimates are close to the true standard deviations for both techniques. In most of these cases the mean errors are less than 1%. For scenarios 1 and 3 and for the two week and one month maturities however, the mixture lognormal appears to perform significantly worse. For low times to maturity and low volatility, the mixture lognormal over-estimates the true standard deviation on average.

The results for the higher moments are much more variable. The absolute size of the mean errors as a proportion of the true statistic are much higher than for the first two moments. For skewness, these figures partly over-state the problems, however, because the true skewness is close to zero for at least scenarios 2 and 5. Compared to the mixture lognormal, the smile-based technique has less biassed results for skewness under scenarios 1, 3, 4 and 6 - those that display more extreme levels of skewness. But for scenarios 2 and 5 the smile-based approach does better for some maturities but worse than the mixture lognormal for others. Like skewness, the mean errors of the kurtosis estimates are larger and more variable than for the mean or standard deviation. Broadly, the mixture lognormal mean estimates are poorest for scenarios 3 and 6 in which skewness is strongest. The smile-based approach has the poorest results for scenarios 5 and 6 when maturity is three months or above.

Table 4: Difference Between True and Mean Estimated Summary Statistics (as % of the True)

	Smile Technique					Mixture Lognormal Technique			
	Scenario	2 weeks	1 month	3 month	6 month	2 weeks	1 month	3 month	6 month
Mean	1	0.00%	0.00%	0.00%	0.00%	-0.17%	-0.05%	-0.03%	-0.05%
	2	0.00%	0.00%	0.00%	0.00%	0.00%	0.15%	0.16%	0.00%
	3	0.00%	0.00%	0.00%	0.00%	-1.95%	-0.42%	-0.03%	0.02%
	4	0.00%	0.00%	0.00%	-0.01%	0.17%	0.35%	-0.71%	0.11%
	5	0.00%	0.00%	0.00%	0.00%	0.00%	-0.01%	0.00%	0.00%
	6	-0.01%	0.00%	0.00%	-0.04%	-0.22%	1.19%	0.07%	0.10%
Std Dev	1	-0.05%	-0.04%	-0.05%	-0.03%	-34.47%	-0.31%	0.12%	0.16%
	2	-5.09%	-3.51%	-1.87%	-1.19%	-0.58%	-2.96%	-1.03%	-0.40%
	3	-0.18%	-0.19%	-0.12%	-0.09%	-277.99%	-30.11%	-0.12%	0.16%
	4	0.01%	-0.05%	0.10%	0.20%	-0.26%	-0.27%	-3.65%	-0.09%
	5	-0.04%	0.00%	0.16%	0.65%	0.03%	0.13%	0.04%	0.04%
	6	-0.18%	-0.05%	0.56%	1.15%	0.15%	-7.42%	0.59%	1.64%
Skewness	1	6.83%	5.39%	6.30%	5.81%	70.51%	3.44%	11.55%	12.51%
	2	-229.59%	-113.88%	-30.89%	-10.39%	0.51%	-38.86%	-36.15%	-4.36%
	3	1.57%	3.72%	3.97%	4.51%	41.69%	7.73%	16.95%	11.75%
	4	8.51%	8.82%	14.22%	26.03%	206.12%	211.70%	17.50%	47.91%
	5	0.90%	1.20%	4.84%	9.77%	1.35%	4.45%	0.93%	0.80%
	6	1.93%	3.24%	9.93%	13.82%	65.16%	1.46%	13.05%	18.78%
Kurtosis	1	0.68%	1.44%	3.41%	4.54%	-1.89%	0.43%	3.64%	4.78%
	2	-27.42%	-18.34%	-6.23%	-0.27%	-1.11%	0.49%	1.07%	-4.03%
	3	-0.04%	1.89%	6.02%	9.65%	24.36%	4.30%	10.11%	13.00%
	4	0.40%	0.70%	2.53%	2.81%	-2.39%	-5.24%	2.86%	0.60%
	5	1.31%	2.68%	7.75%	14.76%	1.09%	2.69%	1.08%	1.45%
	6	1.27%	5.14%	17.86%	24.66%	6.02%	9.57%	20.90%	32.52%

On the basis of these tests for the ability on average to estimate the true summary statistics, it is not immediately obvious that one of the techniques is better than the other. The smile-based approach does appear to do marginally better in estimating the first two moments particularly at short maturities with low volatility. For the third and fourth moments, however, neither technique obviously out-performs the other.

But when we look at the stability of the estimates, the story is far more clear cut. Table 5 sets out the standard deviations of the estimated summary statistics. High standard deviations of the summary statistics are indicative of instability in the estimated RNDs. For nearly all the scenarios the mixture lognormal has much higher standard deviations of the estimates for all statistics than for the smile-based approach. This mirrors Bliss and Panigirtzoglou's (1999) findings that the mixture lognormal is unstable using actual options data for FTSE and 3 month sterling interest rates. In particular, the mixture lognormal appears to perform badly under scenarios 1 and 3 and to some extent scenario 2, when maturities of the options are one month or less. This suggests that the mixture lognormal approach is unstable when volatility is low, the maturity of the options is low and when there is a strong negative or positive skew.

Table 5: Standard Deviation of Summary Statistics

Smile Technique

	Scenario	2 weeks	1 month	3 month	6 month
Mean	1	0.0000	0.0000	0.0000	0.0000
	2	0.0000	0.0000	0.0000	0.0000
	3	0.0000	0.0000	0.0000	0.0000
	4	0.0000	0.0000	0.0000	0.0000
	5	0.0000	0.0000	0.0000	0.0000
	6	0.0000	0.0000	0.0000	0.0000
Std Dev	1	0.0123	0.0110	0.0088	0.0091
	2	0.0144	0.0137	0.0112	0.0094
	3	0.0139	0.0123	0.0112	0.0100
	4	0.0093	0.0095	0.0062	0.0063
	5	0.0104	0.0080	0.0075	0.0065
	6	0.0097	0.0079	0.0080	0.0068
Skewness	1	0.0204	0.0192	0.0130	0.0085
	2	0.0201	0.0234	0.0104	0.0068
	3	0.0191	0.0166	0.0106	0.0080
	4	0.0096	0.0064	0.0030	0.0021
	5	0.0091	0.0061	0.0035	0.0027
	6	0.0102	0.0066	0.0038	0.0028
Kurtosis	1	0.0175	0.0156	0.0141	0.0100
	2	0.0645	0.0333	0.0163	0.0101
	3	0.0517	0.0296	0.0215	0.0189
	4	0.0092	0.0065	0.0035	0.0022
	5	0.0100	0.0078	0.0076	0.0069
	6	0.0157	0.0150	0.0143	0.0139

Mixture Lognormal Technique

	Scenario	2 weeks	1 month	3 month	6 month
Mean	1	4.2206	0.2585	0.0198	0.0176
	2	0.1927	0.9256	0.7257	0.0149
	3	11.9560	6.2271	0.3363	0.0198
	4	0.0133	0.0390	5.7981	2.4340
	5	0.0121	0.0137	0.0185	0.0082
	6	0.0859	4.2933	0.0153	0.0093
Std Dev	1	4.3470	0.0730	0.0102	0.0107
	2	0.0693	0.5684	0.2252	0.0126
	3	12.2235	7.5644	0.0994	0.0116
	4	0.0099	0.0093	5.4027	0.5017
	5	0.0117	0.0092	0.0135	0.0276
	6	0.0111	2.4101	0.0088	0.0088
Skewness	1	0.2348	0.1663	0.0248	0.0142
	2	0.1721	0.2341	0.2684	0.0165
	3	0.7527	0.1899	0.1975	0.0130
	4	0.0003	0.0458	0.0644	0.2577
	5	0.0015	0.0055	0.0115	0.0167
	6	0.0423	0.1839	0.0049	0.0037
Kurtosis	1	1.5428	0.1002	0.0532	0.0335
	2	0.0832	0.1835	0.0827	0.1190
	3	0.9664	0.3374	0.2073	0.0565
	4	0.0002	0.0185	0.1992	0.2575
	5	0.0239	0.0274	0.0677	0.1572
	6	0.0271	0.3777	0.0249	0.0227

Table 5 indicates that the “smile”-based estimation is far more stable than the mixture lognormal approach. But how does the instability of the mixture lognormal estimates manifest itself? To see this we look in detail at the estimated RNDs under scenario 3, one of the scenarios in which mixture lognormal is most unstable, and compare them to the true RND. Figures 3,4 and 5 compare the true RND with estimated RNDs using the two techniques for maturities of two weeks, 1 month and 3 months respectively. In each figure, the top panel displays the true RND, while the bottom two panels each display thirty of the RNDs estimated from the previous tests for the smile-based approach and the mixture lognormal technique.

It is immediately clear that the smile-based RND estimates are far more stable than the mixture lognormal RNDs. The former match the shapes of the true RND closely, particularly for the longer maturities. At the two week maturity there is greater variation in the fitted RNDs but this is to be expected given that the tick size errors that are added to the prices will have a greater proportionate impact on the time value of these shorter maturity options.

The mixture lognormal distributions are highly unstable at the two week and one month maturities. The most common - but not only - cause of this instability is the existence of “spikes” in the distribution. The spikes occur when the variance of one of the distributions collapses. The mixture of two distributions then looks like a single lognormal distribution with a spike, usually towards the centre. Clearly such a spike is not contained within the true distribution and reflects estimation errors. As the maturity of the options increases, the proportion of these spiked distributions falls. In addition to these spiked cases, there are a few mixture lognormal RNDs which are not spiked but which display skewness which is quite different to the true RND.

Examining further the one month mixture lognormal distributions, we can see that for a larger proportion of cases, the mixture lognormal technique manages to get a close fit to the true distribution than at two weeks. As we move to three months we get a higher proportion still of fitted mixture lognormal RNDs that are close to the true distribution. At this maturity, the optimisation used to fit the RND appears to be flipping between between two minima - one which closely matches the true RND and one which contains a spike and hence severely mis-estimates the RND. It is the existence of these spiked distributions that causes the increased standard deviation of the estimated summary statistics compared with the smile-based approach.

What appears to be the key difference between the two estimation approaches is that the small errors in the prices cause only small local errors in the estimated RNDs under the smile approach, while for the mixture lognormal non-linear least squares estimation, the errors can be sufficient for the minimisation to reach very different parameter estimates with large changes in the shape of the estimated RND as a result. The end result of this is that while the bias of the mixture lognormal estimator does not appear to be much larger than the smile-based estimator (at least for the third and fourth moments) it is far more unstable. Since in practice we are often concerned with changes in the

PDF from one day to another this instability is a concern and reduces the value of the mixture lognormal technique as a practical tool.

Figure 3

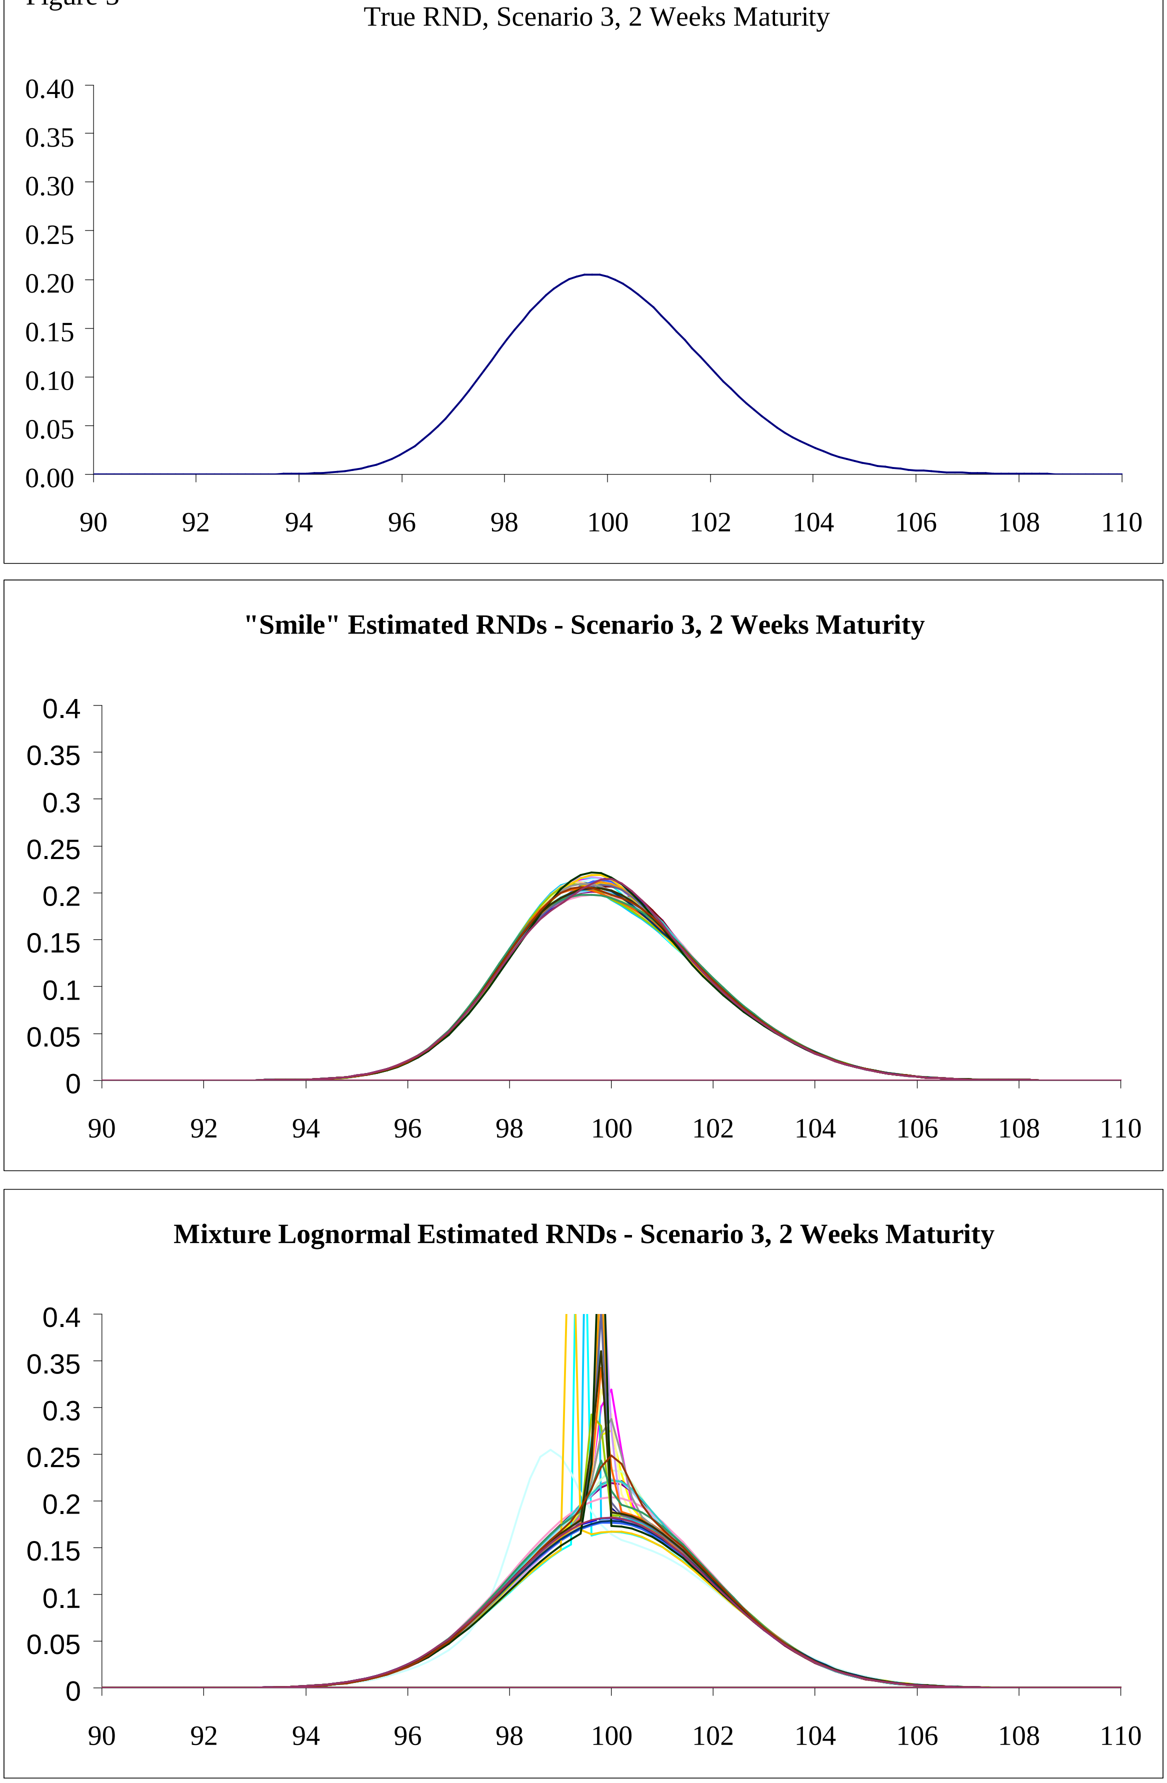
True RND, Scenario 3, 2 Weeks Maturity
0.40
0.35
0.30
0.25
0.20
0.15
0.10
0.05
0.00
90
92
94
96
98
100
102
104
106
108
110
"Smile" Estimated RNDs - Scenario 3, 2 Weeks Maturity
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
90
92
94
96
98
100
102
104
106
108
110
Mixture Lognormal Estimated RNDs - Scenario 3, 2 Weeks Maturity
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
90
92
94
96
98
100
102
104
106
108
110

Figure 4

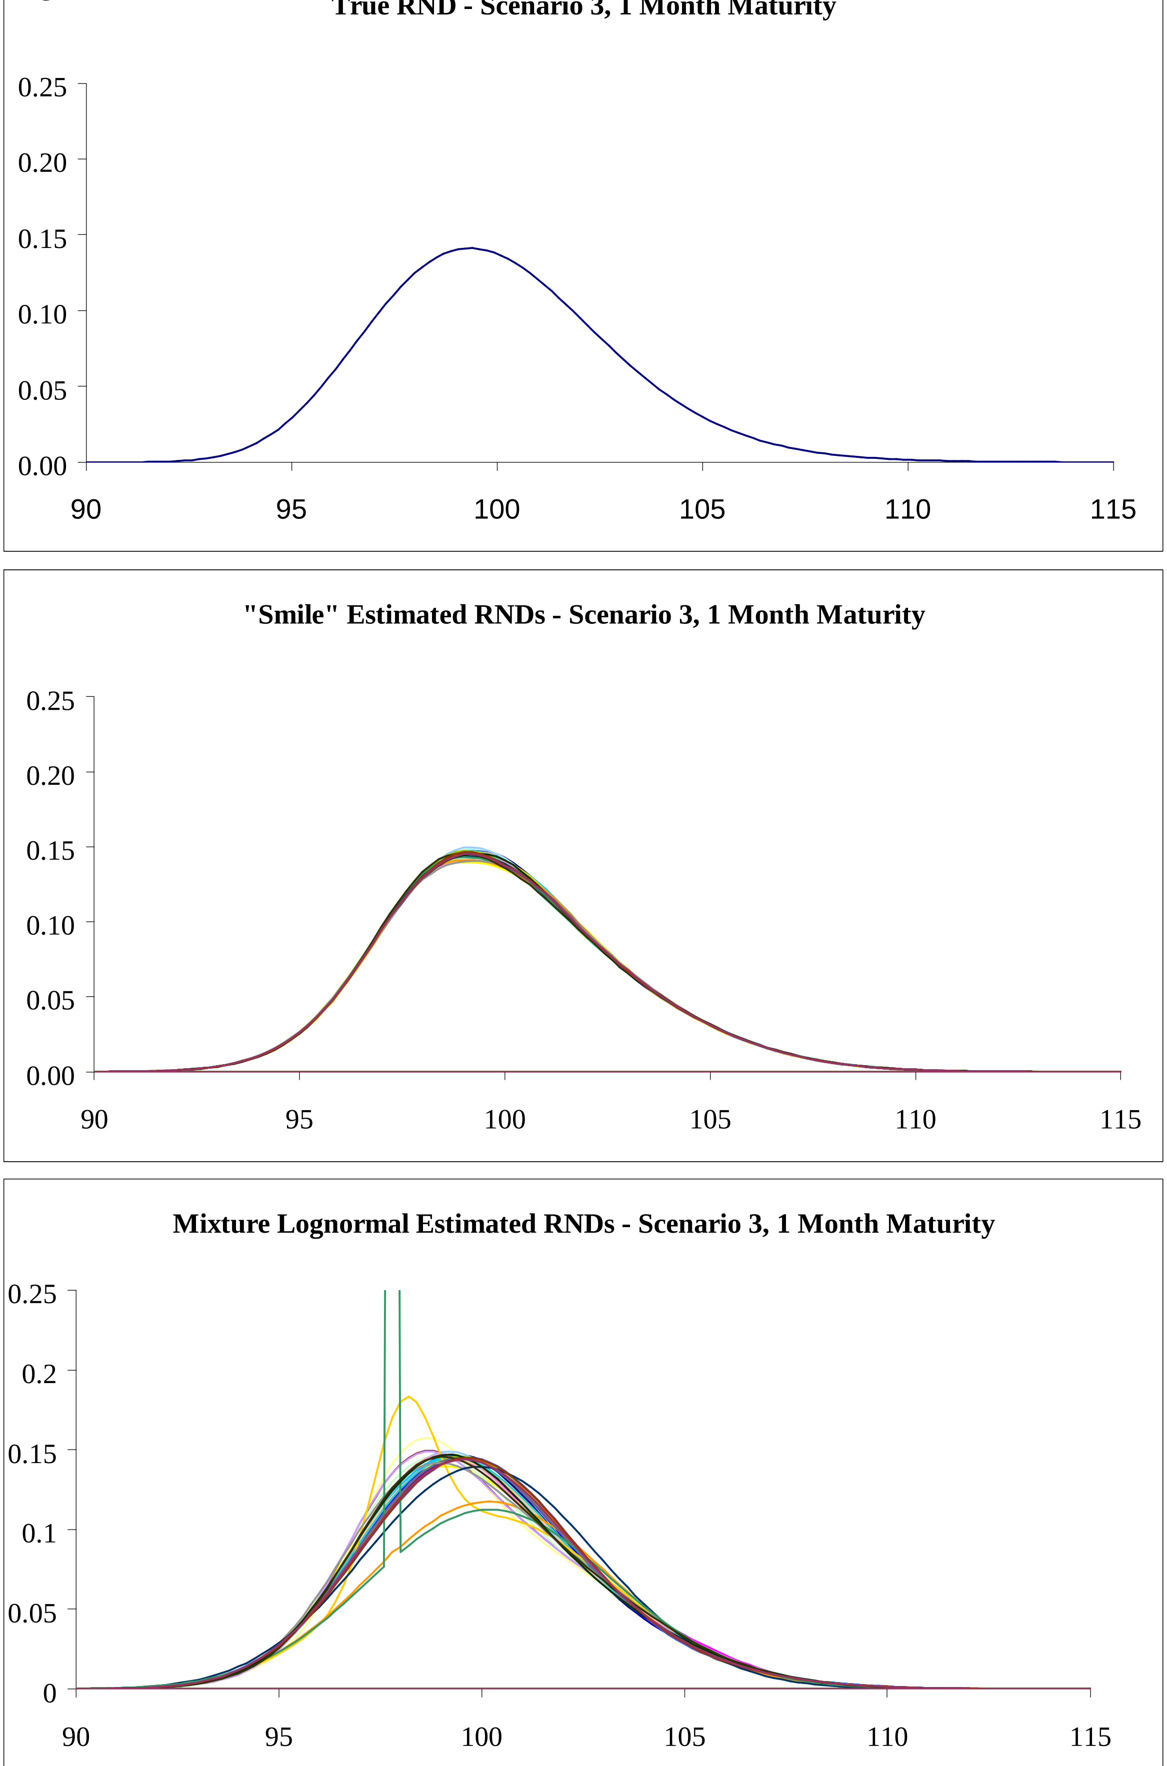
True RND - Scenario 3, 1 Month Maturity
0.25
0.20
0.15
0.10
0.05
0.00
90
95
100
105
110
115
"Smile" Estimated RNDs - Scenario 3, 1 Month Maturity
0.25
0.20
0.15
0.10
0.05
0.00
90
95
100
105
110
115
Mixture Lognormal Estimated RNDs - Scenario 3, 1 Month Maturity
0.25
0.2
0.15
0.1
0.05
0
90
95
100
105
110
115

Figure 5

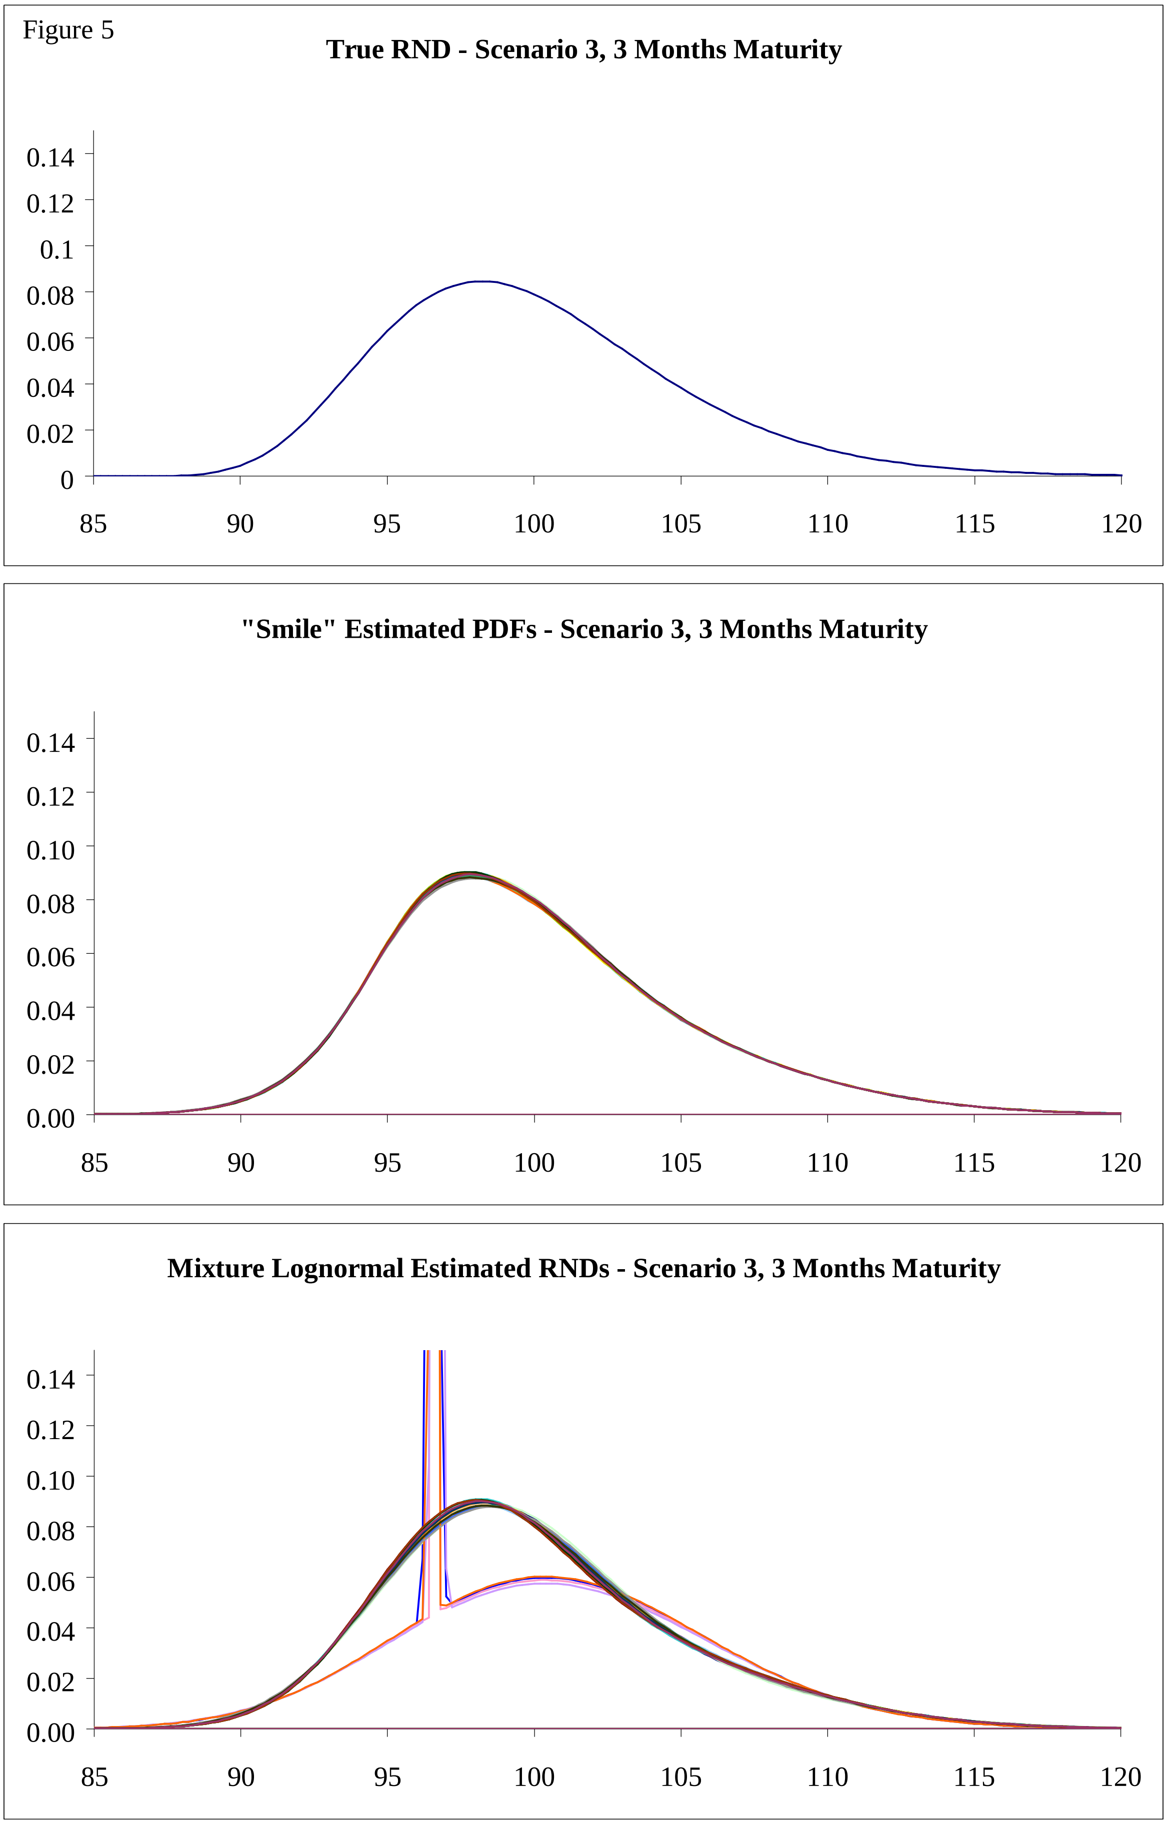
True RND - Scenario 3, 3 Months Maturity
0.14
0.12
0.1
0.08
0.06
0.04
0.02
0
85
90
95
100
105
110
115
120
"Smile" Estimated PDFs - Scenario 3, 3 Months Maturity
0.14
0.12
0.10
0.08
0.06
0.04
0.02
0.00
85
90
95
100
105
110
115
120
Mixture Lognormal Estimated RNDs - Scenario 3, 3 Months Maturity
0.14
0.12
0.10
0.08
0.06
0.04
0.02
0.00
85
90
95
100
105
110
115
120

5. Assessing Alternative Approaches

This paper has examined two alternative approaches to estimating implied RNDs from European-style options. The first was the commonly used mixture lognormal approach which uses non-linear least squares estimation to fit a parametric form to observed options prices. The second approach interpolated across the volatility smile using a cubic smoothing spline and then employed the Breeden and Litzenberger result to recover the RND by calculating the second partial derivative of the call pricing function with respect to the strike price.

The monte carlo tests of the two estimators in section four demonstrated that the second "smile"-based approach performed a little better in terms of its ability to match the first two moments of the true RND. We also saw that the higher moments appear to be much more difficult to estimate accurately with both techniques often resulting in estimates that are on average quite a long way from the true ones.

But we also observed that the smile-based technique was far more stable than the mixture lognormal approach. The latter technique has severe mis-estimation problems when using options on low volatility assets or when using low maturity (less than three months) options. This mis-estimation most often shows up as a "spiked" distribution when one of the lognormal distribution's estimated variance falls to a very low level. In contrast the "smile"-based estimation appears to perform well across all scenarios and maturities (although the existence of discrete tick size errors does create increased instability at maturities below 1 month). These results suggest the use of the smile-based approach over the mixture lognormal by practitioners and researchers alike. They also suggest that where the mixture lognormal is still used that the results have to be interpreted with great caution given the tendency of the estimation approach to severely mis-estimate the true RND.

Future work at the Bank will use the monte carlo tests set out here to examine the empirical performance of RND estimators that use American-style options. As well as examining the accuracy and stability of the techniques, this forthcoming work will examine how important it is to take account of the early exercise premium when using these options.

Bibliography

Bates, D (1996), “Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options”, *Review of Financial Studies*, Vol 9, No. 1 pp 69-107

Bahra, B. (1996), “Probability Distributions of Future Asset Prices Implied By Option Prices”, *Bank of England Quarterly Bulletin*, August 1996, pages 299-311.

Bahra, B. (1997), “Implied Risk-Neutral Probability Density Functions from Options Prices: Theory and Application”, *Bank of England Working Paper* No.66

Bliss, R and Panigirtzoglou, N (1999). “Testing the Stability of Implied PDFs”, *Bank of England mimeo*

Breeden, D.T., and R.H. Litzenberger (1978), “Prices of State-Contingent Claims Implicit in Option Prices”, *Journal of Business*, 51, 4, 621-651.

Campa, J., Chang, K. and Reider, R. (1998) “Implied Exchange Rate Distributions: Evidence From OTC Option Markets”, *Journal of International Money and Finance*

Cox, J.C. and Ross, S.A (1976) “The Valuation of Options for Alternative Stochastic Processes”, *Journal of Financial Economics* 3, 145-66

Das, S and Sundaresan, R. (1998), “ Of Smiles and Smirks: A Term Structure Perspective”, *Working Paper*

Davidson and McKinnon (1993) “Estimation and Inference in Econometrics” Oxford University Press

Heston, S (1993), “A Closed-Form Solution for Options with Stochastic Volatility With Applications to Bond and Currency Options”, *The Review of Financial Studies*, Vol 6, No 2, pp 327-343.

Greene, W. (1997) “Econometric Analysis”, Prentice Hall

Jondeau, E. and Rockinger, M (1998), "Reading the Smile: The Message Conveyed By Methods Which Infer Risk Neutral Densities", CEPR *Discussion Paper* No 2009.

Kurpiel, A. and Roncalli, T. (1999) "Hopscotch Methods For Two-State Financial Models", *FERC Working Paper, City University*

Malz, A (1995), " Using Options Prices to Estimate ReAlignment Probabilities in the European Monetary System", *Federal Reserve Bank of New York Staff Reports*, No. 5

Malz, A (1997), "Estimating the Probability Distribution of Future Exchange Rates from Options Prices", *Journal of Derivatives*

Melick and 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*, Vol 32 No1 March 1997

Shimko (1993), "Bounds of Probability", RISK, Vol 6, No.4.

Discussion of Neil Cooper's paper:

TESTING TECHNIQUES FOR ESTIMATING IMPLIED RNDS FROM THE PRICES OF EUROPEAN-STYLE OPTIONS

by Holger Neuhaus[1]

What I have to say requires fifteen minutes. I have ten. So, fasten your seat belts, we are about to take off.

The motivation for Neil Cooper's paper is comparably old: in 1995 when I was just finishing my research paper I was talking a lot to Bhupi Bahra at the Bank of England and he was uncertain whether to implement Shimko's smile approach or Melick's and Thomas' mixture of lognormals to derive implied probability density functions. Well - you all know what the Bank eventually decided, but now Neil Cooper tries to investigate in an objective way which of the two methods is the better one and he does that by the following methodology:

Fig. 1 True risk-neutral density function

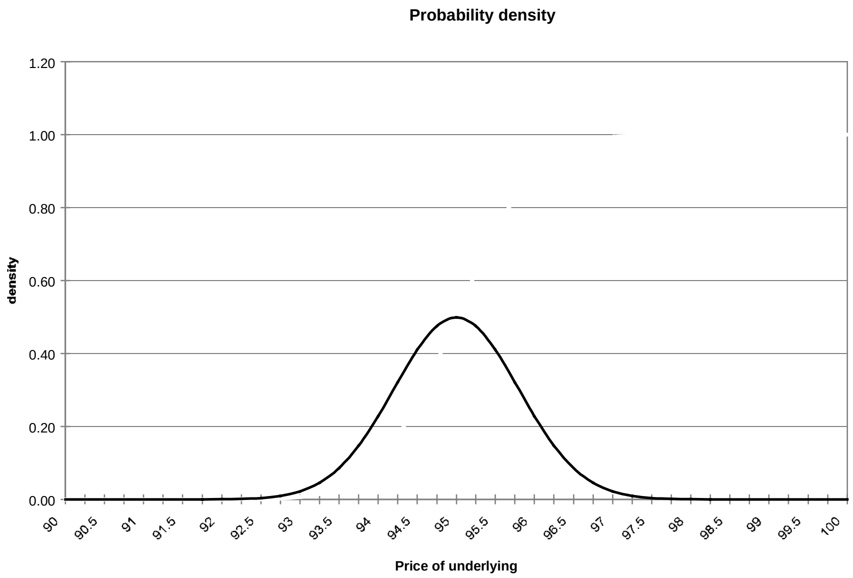

Neil imposes probability density distributions like the one in figure 1 – although his are fancier, of course. He then uses this function to derive corresponding option prices which, in turn, serve as input for the models for backing out the *implied* probability distribution. Eventually, the results can be compared with the imposed distributions he started with. Moreover Neil shocks the option prices by "half a tick size" to find out how sensitive the models are to inaccurate option prices, the inaccuracy being caused by discrete tick-sizes.

[1] The views expressed represent exclusively the opinion of the author and do not necessarily correspond to those of the European Central Bank.

All in all, this is a promising approach as it is the only method that allows a comparison of the estimates with a known probability density function. However, a relevant question is whether the implementation of the method is close enough to reality to allow for a fair comparison of the different approaches for backing out implied probabilities.

In this context, one observation frequently made in the real world is that one may not always have a sufficient number of strikes and option prices to cover the whole distribution – as illustrated in figure 2.

Fig. 2 True risk-neutral density function with missing strikes

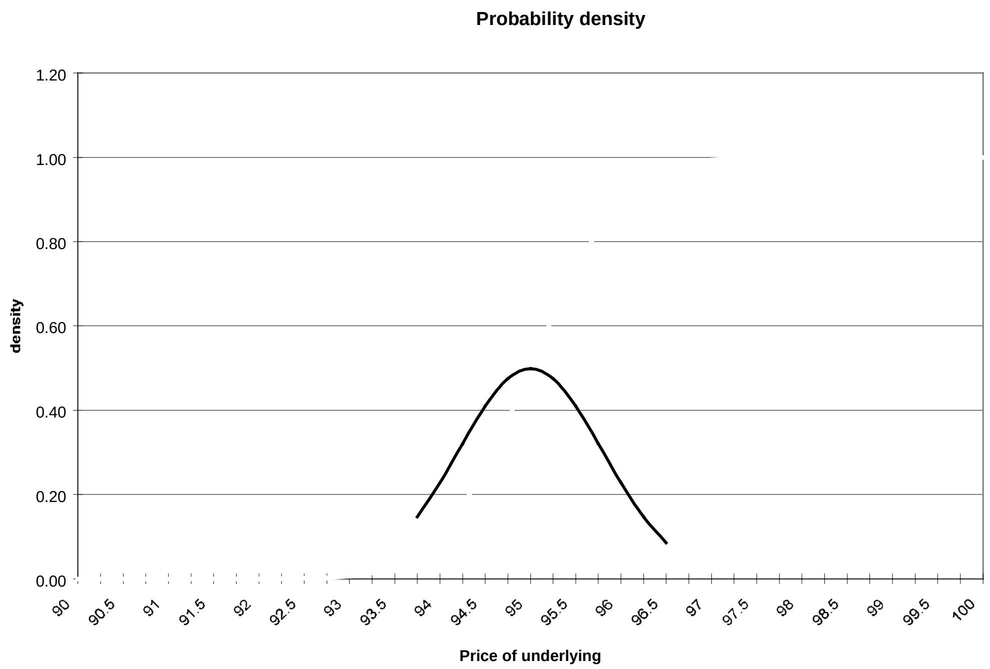

This presentation is of course an exaggeration, but it drives home my point that if only a part of the density function is backed out, it can be difficult to correctly allocate the missing probability mass. When looking at the smile technique, should one, for example, extrapolate the smile or keep the implied volatility constant at the tails?

Some models have more problems with these errors than others do and Neil's dataset is “too” complete and thus leaves out an important aspect in the evaluation of different approaches chosen. As a matter of fact, this is one issue where, as we have already discussed earlier today, the same model can yield (slightly) different results when implemented in a (slightly) different manner.

In this context I also keep mentioning that one means for deciding on how to allocate the missing probability mass is not to estimate the probability density function but to estimate the first derivative of the option price with respect to the strike price, i.e. the implied (*cumulative*) probability distribution. If the lowest and highest cumulative probability derived from the data are, say, 1% and 98% respectively, it can be inferred that 1% of the probability mass is missing at the lower end and 2% at the right tail. This could, in particular, be valid for the smile technique.

Fig. 3 True risk-neutral density and cumulative distribution function with missing strikes

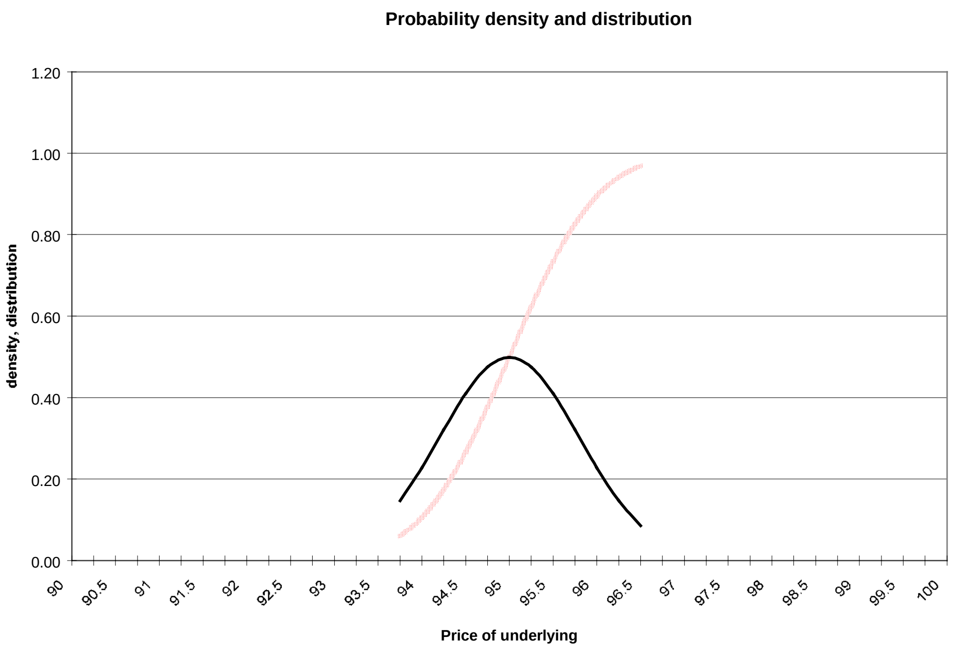

When comparing different approaches to recover implied probabilities, it is also important to mimic realistically another feature of the relevant option market: the number of and distance between strike prices. On the foreign exchange market, for example, few strike prices exist, while, say, for short-term interest rates, some derivative exchanges provide a large number of options. The number of, and distance between, strike prices are determined by certain rules aiming at striking a balance between having a choice between a large number of strike prices and liquid option contracts. Looking at Neil's paper, I have the impression that the intervals chosen between the strike prices are comparably close, which could, in my view, favour the smile approach.

Anyway, looking at Neil’s paper more closely, in particular at the summary statistics he generated after he had backed out the (shocked) implied probability density 100 times for each of his scenarios/maturities, one can conclude the following. For facilitating the comparison, I categorised the results as g or b. G is green and good, b is brown and bad[2], depending on the relative performance of the different methods' estimates.

Table 1: Mean of estimates

Table 3: Mean Estimated Summary Statistics

	Scenario	Smile Technique 2 weeks	1 month	3 month	6 month	Mixture Lognormal Technique 2 weeks	1 month	3 month	6 month
Mean	1	100.000	100.000	100.000	100.000	100.168	100.045	100.031	100.053
	2	100.000	100.000	100.000	100.000	99.995	99.853	99.843	100.001
	3	100.000	100.000	100.000	100.000	101.954	100.419	100.033	99.982
	4	100.000	100.000	100.000	100.000	99.833	99.652	100.707	99.885
	5	100.000	100.000	100.000	**100.000**	100.003	100.011	100.002	**100.000**
	6	100.000	100.000	100.000	100.000	100.204	98.804	99.930	99.859
Std Dev	1	1.9595	2.8787	4.9588	6.9676	2.6334	2.8864	4.9503	6.9544
	2	2.0615	2.9891	5.0970	7.1651	1.9732	2.9732	5.0552	7.1094
	3	1.9685	2.9035	5.0581	7.2073	7.4278	3.7705	5.0581	7.1893
	4	5.8485	8.5588	14.5087	20.0582	5.8645	8.5775	15.0537	20.1169
	5	5.8905	8.6778	15.0693	21.3463	5.8861	8.6659	15.0867	21.4762
	6	5.9316	8.8035	15.5988	22.6372	5.9116	9.4515	15.5956	22.5237
Skewness	1	-0.1848	-0.2654	-0.3915	-0.4466	-0.0585	-0.2709	-0.3695	-0.4148
	2	0.1968	0.1898	0.2081	0.2548	0.0594	0.1232	0.2165	0.2408
	3	0.3132	0.4421	0.7130	0.9128	0.1855	0.4236	0.6167	0.8436
	4	-0.1521	-0.2076	-0.2578	-0.1960	0.1764	0.2543	-0.2479	-0.1381
	5	0.1784	0.2692	0.4798	0.6825	0.1776	0.2603	0.4995	0.7504
	6	0.5124	0.7505	1.2121	1.5861	0.1821	0.7642	1.1701	1.4949
Kurtosis	1	3.0163	3.0363	3.0701	3.0744	3.0942	3.0672	3.0629	3.0668
	2	3.8713	3.6523	3.4221	3.3638	3.0720	3.0711	3.1870	3.4901
	3	3.1615	3.2812	3.6934	4.1546	2.3902	3.2004	3.5323	4.0006
	4	2.9720	2.9415	2.8054	2.6707	3.0554	3.1175	2.7957	2.7313
	5	3.0781	3.1813	3.5137	3.9345	3.0848	3.1811	3.7677	4.5488
	6	3.3829	3.8291	5.1522	6.8006	3.2202	3.6502	4.9617	6.0907

Indeed, the smile outperforms the mixture of lognormals at most occasions. However, the difference in quality between the models is not that clear-cut. Sometimes the smile yields better results but only marginally and the question has to be addressed: how big is big, i.e. when are results significantly different. In this context, I would like to mention that some of the percentage mistakes presented in the paper are misleading, in particular those where the benchmark value is close to zero. A point of greater substance is that, as we had already discussed today, the measurement of the skewness and kurtosis by calculating the 3rd and 4th moment is not recommendable but could be replaced by other measures.

Looking at the expected value of the smile technique, I would like to note that, even when using the shocked option prices, the smile leads to more reliable estimates of the true distribution than the figures provided in what Neil calls the "true" summary statistics of the imposed probability distributions. Because of the numerical integration method used, the "true" mean is slightly different from 100, while the smile always generates an expected value of 100. That looks very - or even too - robust to me.

And indeed, looking at the standard deviation of the estimated summary statistics (mean, standard deviation, skewness, and kurtosis), the smile looks very robust and fares better than the mixture of lognormals approach.

---

2 In black and white copies the brown fields are the darker ones.

Table 2: Standard deviation of estimates

Table 5: Standard Deviation of Summary Statistics

Smile Technique

	Scenario	2 weeks	1 month	3 month	6 month
Mean	1	0.0000	0.0000	0.0000	0.0000
	2	0.0000	0.0000	0.0000	0.0000
	3	0.0000	0.0000	0.0000	0.0000
	4	0.0000	0.0000	0.0000	0.0000
	5	0.0000	0.0000	0.0000	0.0000
	6	0.0000	0.0000	0.0000	0.0000
Std Dev	1	0.012303	0.010959	0.008757	0.009063
	2	0.014392	0.013679	0.011181	0.009371
	3	0.013898	0.012333	0.011182	0.010014
	4	0.009302	0.009502	0.006178	0.006323
	5	0.01037	0.008047	0.007502	0.006537
	6	0.009686	0.007899	0.008014	0.006796
Skewness	1	0.020438	0.019151	0.012986	0.008499
	2	0.020113	0.023413	0.010409	0.006759
	3	0.019066	0.016583	0.010591	0.008017
	4	0.009557	0.006367	0.00304	0.002144
	5	0.00911	0.006055	0.003462	0.002719
	6	0.0102	0.006642	0.003844	0.002834
Kurtosis	1	0.017536	0.015553	0.014076	0.010039
	2	0.064452	0.033344	0.016313	0.010056
	3	0.051657	0.029554	0.021507	0.018863
	4	0.009222	0.006492	0.003471	0.00217
	5	0.010033	0.007826	0.00757	0.006854
	6	0.015728	0.015041	0.014292	0.013867

Mixture Lognormal Technique

	Scenario	2 weeks	1 month	3 month	6 month
Mean	1	4.220623	0.258529	0.019781	0.01758
	2	0.19266	0.925638	0.725677	0.014939
	3	11.95599	6.227077	0.336293	0.019803
	4	0.013332	0.03904	5.798117	2.433972
	5	0.012113	0.01368	0.01845	0.008227
	6	0.085873	4.29325	0.015288	0.009329
Std Dev	1	4.347047	0.072974	0.010168	0.010725
	2	0.069271	0.568389	0.225205	0.012629
	3	12.22352	7.564441	0.09944	0.011641
	4	0.009903	0.009298	5.402745	0.501706
	5	0.011721	0.009154	0.013471	0.027603
	6	0.011146	2.410094	0.008808	0.008811
Skewness	1	0.234815	0.166297	0.024775	0.014154
	2	0.172144	0.234108	0.268442	0.016485
	3	0.752712	0.189888	0.1975	0.012989
	4	0.000301	0.045771	0.064426	0.25774
	5	0.00151	0.005481	0.011528	0.016678
	6	0.042258	0.183944	0.004942	0.003734
Kurtosis	1	1.542771	0.10017	0.053217	0.03349
	2	0.083219	0.183497	0.082698	0.119001
	3	0.966391	0.337422	0.207318	0.056505
	4	0.00019	0.018513	0.199159	0.257488
	5	0.023942	0.027396	0.067714	0.157168
	6	0.027087	0.377703	0.024928	0.022667

Again, the issue is: how big is big? And: why is the smile approach looking so good? One reason is, of course, its design: option prices are translated into implied volatilities, a smooth function is calculated that fits the smile (in delta space) and is used to back out the density function. When the prices are shocked by small amounts, in this case by half a tick, the shocks have a small impact on the implied volatility, which is then smoothed away by the curve fitted to the smile. The mixture of lognormals is more sensitive with respect to these shocks as the option prices are directly used to estimate the density function. That is why I had a look at the option prices used, by processing some of Neil's data with my own model. A further reason – to be honest – was that I had marvellous results as regards the mean and standard deviation for my own model (but less promising estimates of skew and kurtosis).

The equation I used is an approximation of the first derivative of the option price with respect to the strike, which yields the implied cumulative probability distribution at that strike (K_i), which I approximate with a simple difference quotient, which in this audience does not require a lot of explanation.[3]

[3] Cf. Holger Neuhaus (1995), The information content of derivatives for monetary policy – implied volatilities and probabilities, Deutsche Bundesbank Economic Research Group, Discussion paper 3/95 (July 1995).

C_i, K_i and F_T are the price of option i, its strike price and the value of the futures at the expiry of the option. To be precise, the option in this equation should be either margined or C should already be adjusted for the discount factor (as is the case here).

(1) $C = \int_{-\infty}^{+\infty} w(F_T) \max(0, F_T - K) dF_T$

(2) $C_K = -\int_{K}^{+\infty} w(F_T) dF_T$

(3) $p(F_T > K_i) = \dfrac{C_{i-1} - C_{i+1}}{K_{i+1} - K_{i-1}}$

Even if one does not recognise the quotient (consisting of the difference between two option prices divided by the difference in their strike prices) as a cumulative probability distribution, it is clear that deep in-the-money options going deeper into the money by one unit will increase in value by the same amount, i.e. the ratio should be exactly unity. Is this always the case in the paper? No: as one can see in table 3, there are small problems with the true prices (having taken into account the discount factor)

Table 3: True prices and implied cumulative probability distribution values

Strike	C	C compounded	Probability distribution
71	28.94429	29.000015	1.0000
72	27.94621	28.000015	1.0000
73	26.94813	27.000015	1.0000
74	25.95006	26.000015	1.0000
75	24.95198	25.000015	1.0000
76	23.9539	24.000015	1.0000
77	22.95582	23.000015	1.0000
78	21.95774	22.000015	1.0000
79	20.95966	21.000015	1.0000
80	19.96159	20.000015	1.0000
81	18.96351	19.000015	1.0000
82	17.96543	18.000015	1.0000
83	16.96735	17.000015	1.0000005009626
84	15.96927	16.000014	1.0000
85	14.97119	15.000015	1.0000
86	13.97311	14.000014	1.0000
87	12.97504	13.000015	1.0000
88	11.97696	12.000014	1.0000005009626
89	10.97888	11.000014	1.0000001502888
90	9.980799	10.000014	0.9999999499037
91	8.98272	9.000014	1.0000001502888

in particular, have an influence on the results of the mixture of lognormals approach. However, does it imply that this approach is indeed too sensitive? Not necessarily, as the prices used allow for arbitrage and are thus “wrong”, albeit to a limited extent only. A way forward in the research would be to just shock prices of options that are not *deep* in or out of the money.

Table 4: Some shocked prices and implied cumulative probability distribution values

	Strikes				
Shock	**78**	**79**	**80**	**81**	**82**
1	0.9963	0.993189	1.002543	1.011356	0.99394
2	0.999621	1.024057	1.008464	0.977107	0.993295
3	1.007501	1.004483	0.994608	1.008015	1.012959
4	0.999495	1.017243	1.005898	1.003024	0.994951
5	0.995052	0.990734	0.990985	1.001388	0.99782
6	1.010791	1.004175	0.990104	0.983198	0.992493
7	0.997608	0.984138	1.00091	1.019346	1.001663
8	1.000668	0.998463	0.984069	1.000081	1.001444
9	1.019999	0.991685	0.998495	1.014912	1.00225
10	1.011092	1.018889	0.981598	0.996102	1.011387
11	1.002029	1.01328	0.999109	0.993804	0.985691
12	1.016841	0.989879	0.990558	1.013711	1.003615
13	1.001403	0.98645	1.00682	0.994389	0.986779
14	1.004532	1.015066	0.98888	0.993103	1.015303
15	1.017755	0.984908	0.997246	0.996538	0.999667
16	1.003631	1.003485	1.002598	1.004908	0.995692
17	0.992376	0.998081	0.998523	1.004808	1.015937
18	1.00167	0.987514	1.00775	1.011708	0.989414
19	0.98171	1.0117	1.009924	0.978308	1.00168
20	0.999093	1.002282	1.01776	1.002007	1.00061
21	0.981344	0.997791	1.006633	1.002047	1.009586
22	1.013006	1.000001	0.983824	0.998456	1.02053
23	1.006047	0.985566	0.996316	1.008509	0.984099
24	0.993576	1.001908	1.003176	0.985739	0.988618
25	0.988323	1.005294	0.988197	1.00032	1.013966
26	0.993232	1.003898	1.010038	1.010592	0.998971
27	0.99807	0.990528	1.01003	1.006907	0.987676
28	1.021555	0.991778	0.981613	1.002087	1.009585
29	0.977072	1.000689	1.013373	0.996852	0.98623
30	0.996151	0.999982	0.997678	1.019486	1.003256

Conclusions

To summarise what I found out in the short time I had to look at the paper, I think that it is interesting and a step in the right direction. Nevertheless, I have some food for thought:

- Is there "too much" data?
  - Some models are designed to generate results with a very limited data set (strike prices) only, some require more. The choice of model is likely to depend on the market to be monitored.
  - The number of and interval between strike prices used in the comparison should reflect the features of the market that should be monitored.
- Should the probability *distribution* be estimated rather than the *density*?
- Is the smile technique "too" stable?
- Is the mixture lognormals too sensitive as regards errors in prices?
- Data must not allow for arbitrage.
  - This also holds true for "shocked" data.
  - Only the prices of at-the-money options but not of far in-the-money or out-of-the-money options should be shocked.
- How big is big?
  - The size of the errors has to be put into perspective. Criteria used should involve also, for example, computational costs in a broad sense (computer and software requirements, robustness of the estimates).
  - In particular, some percentage errors shown may be misleading.
- The use of the third and fourth moment is debatable.

Discussion of Neil Cooper's paper:

Testing techniques for estimating implied RNDs from the prices of European-style options

Discussant: Jan Marc Berk

BIS, 14 June 1999

1. My contribution is structured as follows. I will start by giving a brief summary of the paper. This is then followed by some comments, and I conclude by sketching some paths for future work on the subject at hand.

Summary

2. The paper aims to compare two methods for calculating PDFs. Both methods are applied on both European and American-style options. The performance of both methods is tested by means of Monte Carlo analysis, although the current version of the paper only deals with comparing the methods applied to European options. Innovative aspects of the paper include the variant used for calculating volatility-smile-based PDFs, and the Monte Carlo experiment.

3. The methods used to construct PDFs for *European* options are the well-known mixture of lognormals approach (MLN), as documented by, for example, Bahra (1997), and a method based on interpolation of the volatility smile (IVS), as introduced by Shimko (1993). The paper slightly amends the Shimko approach, as it uses cubic splines in stead of quadratic forms, and interpolates in delta space in stead of volatility space. These amendments are in line with, for example, Malz (1997).

4. PDFs for *American*-style options are the MLN variant introduced by Melick and Thomas (1997), and the early exercise premium is taken into account within the IVS method by using the approximation of Barone-Adesi and Whaley (1987).

5. The performance of both methods is compared in a Monte Carlo experiment, using simulated artificial data in stead of observed prices. However, by using the stochastic volatility model of Heston (1993), the author generates quite realistic data, whilst retaining the advantage of knowing the 'true' PDF.

6. The results form the Monte Carlo analysis, in the current version of the paper applied only to European options, are that, on average, there is no clear winner between MLN and IVS. However, the latter method provides far more stable estimates. The instability of MLN estimates are due to spikes, and reflect estimation errors. Instability increases with volatility and skewness, and decreases with

time to maturity. Based on these Monte Carlo analysis, the author expresses a preference of IVS over MLN.

Comments

7. I find the paper of Neil Cooper very interesting, and as it is work in progress, I suspect it will become even more interesting. The paper reflects my own experiences, or should I say frustations, with the (in)stability of the MLN method. Without meaning to detract from the quality of the paper, there are some points that, in my view, deserve some further consideration. Given the time constraint, I will only briefly touch upon them here:

8. Whilst the application of the MLN method in the paper is fairly standard, the version of the IVS method employed is more innovative. The paper could benefit form a more extensive discussion on the effects of the amendments *vis-à-vis* the Shimko approach.

9. In a similar vein, no mention is made in the paper of possible drawbacks of the IVS method, such as the problem of fitting the tails of the PDF (ie outside the observed range of strikes), and negative probabilities.

10. The comparison of both methods is based on Monte Carlo analysis. Yet, given the results of Melick and Thomas (1998), who find widely different results for simulations based on Monte Carlo and bootstrap methods, and given the the assumptions underlying the MC-method (independent errors, regularity conditions) vis-à-vis actual options prices, some attention to the validity of Monte Carlo as a tool for comparison seems in order.

11. The comparison of both methods uses artificial data, so there is no distinction between exchange-traded and otc data. Campa, Chang and Reider (1997) compare MLN and IVS methods using otc data, and find that they yield similar results. Could or should the choice of method (MLN versus IVS) be dependent on the type of data used?

12. The focus of the paper is primarily technical and not economic, which is understandable given its objective. However, more attention to the economic aspects would seem in order, as it could provide an answer to the question as to how important the instability of MLN based PDFs is. Clearly, this answer depends on the purpose of the analysis using PDFs.

Way to proceed

13. Based on my, admittedly limited, knowledge of the estimation and use of PDFs, there are two basic questions which in my view remain to be answered in a convincing way. First, regarding the method used to calculate PDFs, do we really need to impose so much structure? Second, regarding the estimation of PDFs, do the data allow us to impose so much structure?

14. My personal opinion on these questions is that we should use different methods for different purposes, also taking the amount and types of data into account. As an economist, I would tend to say that economic considerations should govern the purpose of the analysis, as well as that the results of the analysis should be useful to economists. As an economist working in a monetary policy department, I will go even one step further and state that the results of the analysis should be useful to policy makers. Given the fairly technical nature of work involving PDFs, it is my own experience that translating the results of PDF based analyses to policy messages is by no means an easy task.
15. Data considerations are also of importance in the choice of method. I already touched upon the difference between exchange-traded and otc data and possible implications for the choice of calculation method. Moreover, we all encounter situations when only a few data points are available, or that only a limited subset of a larger set of prices reflect sufficient liquidity. In these situations, I found entropy-based (Bayesian) methods for calculating PDFs useful. Moreover, the field of maximum entropy econometrics has a firm statistical foundation, and provides a natural metric for evaluating different methods.Finally, when not even a limited set of data on options prices is available, it may be still possible to extract a PDF, using alternative methods (Hördahl, 1999).

References

Bahra, B (1997): “Implied risk-neutral probability density functions from option prices: theory and application”. *Working Paper no 66*, Bank of England, London.

Barone-Adesi, G and R E Whaley (1987): “Efficient analytic approximation of American option values”. *Journal of Finance*, 42, pp. 301-20.

Campa, J.M., P H K Chang and R L Reider (1997): “Implied exchange rate distributions: evidence from otc option markets”. *Working Paper no 6179*, NBER, Cambridge (MA).

Heston, S (1993): “A closed-form solution for options with stochastic volatility with applications to bond and currency options”. *The Review of Financial Studies*, 6, pp. 327-43.

Hördahl, P (1999): “Estimating the implied distribution of the future short term interest rate using the Longstaff-Schwartz Model”. Paper presented at the ECB Workshop on yield curve modeling, Frankfurt, March 1999.

Malz, A. (1997): “Estimating the probability distribution of the future exchange rate from option prices”. *The Journal of Derivatives*, Winter 1997, pp. 20-36.

Melick, W R and C 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, pp. 91-115.

Melick, W R and C P Thomas (1998): “Confidence intervals and constant maturity series for probability measures extracted from option prices”. Paper presented at the conference ‘Information Contained in Prices of Financial Assets’, Bank of Canada.

Shimko, D (1993): “Bounds of probability”. *Risk*, 6, pp. 33-37.

Sophie Coutant

"Implied risk aversion in options prices using Hermite polynomials"

Discussant:

Robert Bliss

Implied Risk Aversion in Options Prices Using Hermite Polynomials

Sophie Coutant[1]

BIS Workshop at the BIS on 14 June 1999

Abstract

The aim of this paper is to construct a time-varying estimator of the investors' risk aversion function. Jackwerth (1996) and Aït-Sahalia and Lo (1998) show that there exists a theoretical relationship between the Risk Neutral Density (RND), the Subjective Density (SD), and the Risk Aversion Function. The RND is estimated from options prices and the SD is estimated from underlying asset time series. Both densities are estimated on daily French data using Hermite polynomials' expansions as suggested first by Madan and Milne (1994). We then deduce an estimator of the Risk Aversion Function and show that it is time varying.

Résumé

Nous construisons dans ce papier un estimateur variant avec le temps de la fonction d'aversion au risque d'un investisseur. Jackwerth (1996) et Aït-Sahalia et Lo (1998) montrent qu'il existe une relation théorique entre la densité neutre au risque, la densité subjective et la fonction d'aversion au risque. On estime la densité neutre au risque à partir des prix d'options et la densité subjective à partir d'une série chronologique du sous-jacent. Chaque densité est estimée en données quotidiennes sur le marché français, en utilisant à la suite de Madan et Milne (1994) des expansions en polynômes d'Hermite; on en déduit alors un estimateur de la fonction d'aversion au risque pouvant varier dans le temps.

Keywords: Risk aversion function, Index option's pricing, Risk neutral density, Statistical density, Hermite polynomials.

---

[1] The author is from the Banque de France, Centre de Recherche and from the CEREG at University Paris IX Dauphine. Please forward comments to author's adress: Banque de France, 41-1391 DGE-DEER-Centre de Recherche, 31 rue Croix-des-Petits-Champs, 75049 Paris, France, tel.: 01-42-92-91-67, email: scoutant@banque-france.fr. The author aknowledges helpful comments of Sanvi Avouyi-Dovi, Jean Cordier, Helyette Geman, Eric Jondeau, Michael Rockinger and Pierre Sicsic. Special thanks goes to Robert Bliss, Mico Loretan and Alain François-Heude and other seminar participants at the Bank of
International Settlements and at the Université Catholique de Louvain.

1 Introduction

An important area of recent research in finance is devoted to the information content in options prices that can be obtained in estimating implied Risk Neutral Densities (RND). Whereas this density gives information about market-makers expectations concerning the future behaviour of the underlying asset, it does not allow to infer anything related to investors' risk aversion. In return, there exists a relationship between the risk neutral density, the subjective density (SD) and the risk aversion function.

Although this theoretical relationship is well known, few works have been interested in the topic in an empirical framework. To our knowledge, the two major studies which deal with are those from Jackwerth (1996) and Aït-Sahalia and Lo (1998). On the one hand, they estimate the RND from options prices and on the other hand they estimate the SD from time series of the underlying asset. By comparing both densities, they conclude that risk aversion is time varying.

Following Jackwerth, and Aït-Sahalia and Lo, we extract both densities (RND and SD) and show that investors' risk aversion function is time varying. The contribution of this study is twofold: first, we investigate French dataset, and second we estimate the model at a daily frequency.

With regard to the RND, in addition to seminal work on options pricing by Black and Scholes (1973) and Merton (1973), we may cite Breeden and Litzenberger (1978) who first found a relationship between options prices and the risk neutral density. Nevertheless their method requires a big range of strike prices; over the past few years, a whole literature has looked into the problem of estimating the RND of the option's underlying asset. We may mention stochastic volatility models such as Hull and White (1987), Chesney and Scott (1989) or Heston (1993); to the latter Bates (1991 and 1996) adds a jump process in the asset return diffusion. Madan and Milne (1994) and Jarrow and Rudd (1982) respectively approximate the RND by Hermite and Edgeworth expansions. Rubinstein (1994), Dupire (1994) and Derman et Kani (1994) suggest to use implied binomial trees. Bahra (1996), and Melick and Thomas (1997) assume lognormal mixture for the RND. Aït-Sahalia (1998) uses kernels estimators of the RND. Lastly we refer to Campa, Chang and Reider (1997), Jondeau and Rockinger (1998) or Coutant, Jondeau and Rockinger (1998) for a comparison of several methods of extracting the RND from options prices on a particular event.

Section 2 first presents a brief review of the investment's theoretical foundations in an economy with a single consumption good, second it describes the traditional Black and Scholes

model and explains why this model is too far from reality. Section 3 describes the model that used: Hermite polynomials approximations and shows how we estimate the risk neutral density using options and the subjective density using underlying time series. Finally Section 4 first describes the dataset and analyses statistical properties, second explains which optimisation's proceeds are used to estimate the models and third studies results on French daily
dataset. Section 5 concludes. Technical results are detailed in the Appendix.

2 Methodology

2.1 Implied risk aversion

The basic investment choice problem for an individual is to determine the optimal allocation of his wealth among the available investment opportunities. We stand in a standard investment theory (see Lucas (1978)). There is a single physical good *S* which may be allocated to consumption or investment and all values are expressed in term of units of this good; there is a risk-free asset, i.e. an asset whose return over the period is known with certainty. Any linear combination of these securities which has a positive market value is called a portfolio. It is assumed that the investor chooses at the beginning of a period the feasible portfolio allocation which maximises the expected value of a Von Neumann-Morgenstern utility function for the end-of-period wealth. The only restriction is the budget constraint. We denote this utility function by $U(.)$, and by W_T the terminal value of the investor's wealth at time T. It is further assumed that U is an increasing strictly concave function of the range of feasible values for W, and that U is twice-continuously differentiable. The only information about the assets that is relevant to the investor's decision is the density probability of W_T.

In addition, it is assumed that:

Hypothesis 1: *Markets are frictionless: there are no transactions costs nor taxes, and all securities are perfectly divisible*.

Hypothesis 2: *There are no-arbitrage opportunities in the markets. All risk-free assets must have the same return between t and T. This return will be denoted by $r_t\,(T)$ and is assumed to be known and constant.*

Hypothesis 3: *There are no institutional restrictions on the markets. Short-sales are allowed without restriction.*

As Aït-Sahalia and Lo (1998) write it, the equilibrium price of the risky asset S_t at date t with a T-liquidating payoff $Y(W_T)$ is given by:

$$S_t = E\left[\Psi(W_T)M_{t,T}\right], (1)$$

$$M_{t,T} = \frac{U'(W_T)}{U'(W_t)}, (2)$$

under the true probability, where $M_{t,T}$ is the stochastic discount factor between consumption at dates t and T.

In equilibrium, investor optimally invests all his wealth in the risky stock for all $t<T$ and then consumes the terminal value of the stock at T, $W_T = S_T$.

If we notice by *p(.)* the *subjective density* (SD) of W_T, we may rewrite (1) as:

$$\begin{aligned} S_t &= \int_0^\infty \Psi(W_T)\frac{U'(W_T)}{U'(W_t)}p(W_T)dW_T \\ &= e^{-r_t(T)(T-t)}\int_0^\infty \Psi(W_T)q(W_T)dW_T \\ &= e^{-r_t(T)(T-t)}E_t^*\left[\Psi(W_T)\right] \end{aligned}$$

with

$$q(W_T) = \frac{M_{t,T}}{\int_0^\infty M_{t,T}\,p(W_T)dW_T}\,p(W_T)\,(3)$$

is called the *state-price density* or *risk neutral density* (RND) which is the equivalent in a continuous-time world of the Arrow-Debreu state-contingent claims in a discrete-time world[2].

A way to specify the preference ordering of all choices available to the investor is the risk-aversion function. A measure of this risk-aversion function is the *absolute risk-aversion function A(.)* of Pratt and Arrow (see Pratt (1964)) given by:

$$A(S) = -\frac{U''(S)}{U'(S)}. (4)$$

By the assumption that U is increasing (*U'(S)>0*) and strictly concave (*U"(S)<0*), function *A(.)* is positive; such investors are called *risk-averse*. An alternative, but related measure of risk aversion is the *relative risk-aversion function*:

[2] Recall that Arrow-Debreu contingent claims pay $1 in a given state and nothing in all other states.

$$R(S) = -\frac{U''(S)}{U'(S)}S. (5)$$

From (3), we can deduce than the ratio *q/p* is proportional to $M_{t,T}$ and we can write:

$$V(S_T) = \frac{q(S_T)}{p(S_T)} = qM_{t,T} = q\frac{U''(S_T)}{U'(S_T)}. (6)$$

where q is a constant independent of the level of *S*.

Differentiating (6) with respect to S_T leads to:

$$V'(S_T) = q\frac{U''(S_T)}{U'(S_T)}$$

and

$$-\frac{V'(S_T)}{V(S_T)} = -\frac{U''(S_T)}{U'(S_T)} = A(S_T)$$

We then may calculate *A(.)* as a function of *p(.)* and *q(.)* and we easily obtain an estimator of the absolute risk-aversion function, which does not depend on the parameter q:

$$A(S_T) = \frac{p'(S_T)}{p(S_T)} - \frac{q'(S_T)}{q(S_T)}. (7)$$

At this stage, we need to specify a general form for the utility function and we add the following hypothesis:

Hypothesis 4: *We stand in a state in which investors have preferences characterised by Constant Relative Risk Aversion (CRRA) utility functions (see Merton (1969, 1971)). Those functions have the following*

general form:

$$U(S) = \frac{S^{1-l}}{1-l},\ if\ \ l \neq 1 (8)$$

$$A(S) = \frac{l}{S}, (9)$$

$$\begin{aligned} &U(S) = \ln(S),\ if\ \ l = 1 \\ &A(S) = \frac{1}{S}, \end{aligned} \qquad (10)$$

where l be a nonnegative parameter representing the level of investor's risk aversion.

An estimation of the parameter l will directly give us an idea on the investors' risk aversion level. Once one has supposed a form for the utility function, he must specify a model to extract subjective density p and risk neutral density q. In order to study investor's reactions across time, the risk aversion is to be time-varying. So we replace all previous notations by p_t, q_t, A_t and l_t where t denotes all dates of our dataset. In the next section, first we give an example using the traditional Black-Scholes model, second we explain why Black-Scholes model does not correspond to reality and third we present an extension of Black-Scholes model: Hermite polynomials model which allows for more properties of the data.

2.2 Hermite polynomials expansion vs Black-Scholes

Now, we wish to develop the method for a traditional option pricing model. We have to keep in mind that we need to estimate subjective density p_t and risk-neutral density q_t at each date and then extract parameter l from these estimations.

A large part of the literature concerning options pricing is based on the Black and Scholes (1973) model. Assets returns are lognormally distributed with known mean and variance. The underlying asset S_t, $t \leq T$ follows a Brownian diffusion:

$$dS_t = m_t S_t dt + s_t S_t dW_t \,, (11)$$

where W_t is a Brownian motion under the subjective probability, m_t is the rate of return of S under the SD and s_t is the volatility; both are supposed to be constant for a certain date t. Harrison and Kreps (1979) show that when hypotheses (1) to (3) hold, there exists a unique risk neutral probability equivalent to the subjective one, under which discounted prices of any asset are martingales. Under this equivalent probability, the underlying asset price S_t is distributed as following:

$$dS_t = (r_t - d_t) S_t dt + s_t S_t dW_t^* \,, (12)$$

where W_t^* is a standard Brownian motion under the risk neutral probability, d_t denotes the implied dividend at time t and s_t is the volatility which appears to be the same than under the true probability. In the Black-Scholes model, asset price S_t follows a lognormal under both probability[3]. Risk Neutral

---

[3] Applying Ito's formula to ln(S_t) and (11) gives us:

$$d\ln(S_t) = \frac{dS_t}{S_t} + \frac{1}{2}\left(-\frac{1}{S_t^2}\right)\mathrm{var}(dS_t)dt = \left(m - \frac{1}{2}s^2\right)dt + s\,dW_t$$

Density (RND), $q_t^{BS}(S, s_t)$ and Subjective Density (SD), $p_t^{BS}(S, s_t, m_t)$ only differ in mean and are given by:

$$\mathrm{p}_{\mathrm{t}}^{\mathrm{BS}}(S, s_t, m_t) = \frac{1}{s_t\sqrt{T-t}\sqrt{2pS}}\exp\left[-\frac{\left(\ln(S)-m_t(m_t)\right)^2}{2s_t^2(T-t)}\right], (13)$$

$$\mathrm{q}_{\mathrm{t}}^{\mathrm{BS}}(S, s_t) = \frac{1}{s_t\sqrt{T-t}\sqrt{2pS}}\exp\left[-\frac{\left(\ln(S)-m_t(r_t-d_t)\right)^2}{2s_t^2(T-t)}\right] (14)$$

where

$$m_t(x) = \ln(S_t) + \left(x - \frac{1}{2}s_t^2\right)(T-t).$$

By replacing (13) and (14) and under hypothesis (4) we directly obtain:

$$A_t^{\mathrm{BS}}(S) = \frac{m_t - (r_t - d_t)}{s_t^2 S}. (15)$$

An estimation of parameters m_t and s_t allows us to estimate absolute risk aversion function when the underlying follows (11).

Black and Scholes is based on the fundamental hypothesis that volatility is deterministic, skewness and excess kurtosis are zero. Those hypotheses have been widely reconsidered on the last few years, owing to the fact that option price at maturity is very sensitive to the underlying asset's distribution specifications. Figure 2 shows typical volatility smiles for two dates, May 1995, 5^{th}, date that we can call agitated, and July 1996, 25^{th}, date that we can call flat: we observed that implied volatility at date t is constant neither in strike price neither in maturity; volatility is higher for small strikes, which means that market makers will pay more for a call option on a smaller strike: this feature will appear in the density with a presence of asymmetry; volatility smile for the second date is very U-shape: we will notice a kurtosis effect in the density.

We impose another model for the underlying which allows for skewness and kurtosis. Following Madan and Milne (1994) and Abken, Madan and Ramamurtie (1996), we adopt an Hermite polynomials approximation for the density. Their model operates as follows.

First, we add the following hypotheses to hypotheses (1)-(4):

Hypothesis 5: *The set of all contingent claims is rich enough to form a Hilbert space that is separable and for which an orthonormal basis exists as a consequence. The markets are assumed to be complete.*

Hypothesis 6: *Abken, Madan and Ramamurtie suppose that under a reference measure, the asset price evolves as (11), i.e. as a geometric Brownian motion. Then* S_t *can be written as:*

$$S_T = S_t \exp\left[\left(m_t - \frac{1}{2} s_t^2\right)(T-t) + s_t \sqrt{T-t} z\right] \quad (16)$$

where z follows a N(0,1).

Madan and Milne (1994) assume than SD and RND may be written as a product of a change of measure density and reference measure density *n(z)*:

$$\tilde{p}_t^{\mathrm{HER}}(z) = n_t(z) n(z) \quad (17)$$

$$\tilde{q}_t^{\mathrm{HER}}(z) = u_t(z) n(z) \quad (18)$$

where $\tilde{p}_t^{\mathrm{HER}}(z)$ and $\tilde{q}_t^{\mathrm{HER}}(z)$ are respectively subjective and risk neutral densities. In our particular case *n(z)* will be a Gaussian distribution of zero mean and unit variance. A basis for the Gaussian reference space may be constructed by using Hermite polynomials which form an orthonormal system for the Hilbert space[4].

As we have carried out for the benchmark model, we wish to estimate time-varying risk aversion function when supposing an Hermite polynomials expansion for the density; therefore, we need to estimate both risk-neutral and subjective densities. Next section is divided in two parts. In a first part, we give the way to estimate risk-neutral model from options prices, and in a second part we show how to use these estimated parameters as observed data to estimate subjective model and extract l_t.

3 Models' specifications

3.1 Risk Neutral Model

To estimate implied volatilities risk neutral parameters we use options prices. A call option (put option) is the right to buy (to sell) the option's underlying asset at some future date -the

[4] Hermite polynomial of order *k* is defined as follows:

$$f_k(z) = \frac{(-1)^k}{\sqrt{k!}} \frac{¶^k n(z)}{¶z^k} \frac{1}{n(z)} \text{ with } < f_k, f_j >= \int_{-\infty}^{+\infty} f_k(z) f_j(z) n(z) dz = \begin{cases} = 0 \; if \, j \neq k \\ = 1 \; if \, j = k \end{cases}$$

expiration date- at a prespecified price -the striking price. This right has a price today that is a function of the option's specifications. Since under the risk neutral probability discounted prices are martingales, the current option's price may be written as the discounted end-of-period option's payoff expectation. If we denote by $C_e(t,S,K,T)$, a European call price of exercise price K and maturity T, we have:

$$C_e(t,S_t,K,T) = e^{-r_t(T)(T-t)} \int_0^{\infty} \max(S_T - K, 0) q_t(S_T) dS_T . (19)$$

As CAC 40 options are American style options, we introduce the approach developed by Melick and Thomas (1997) to price American options. They show that the option's price could be flanked by two bounds representing minimal and maximal value of the price. This method can be applied to any stochastic process if we know the shape of the future underlying's distribution. If we can bound the option's price, we will be able to write it as a weighted sum of the bounds. The idea of the method comes from the martingale's hypothesis of the underlying asset under the risk neutral probability. Low and high bounds for an option call are given by:

$$C_t^u = \max[E_t(S_T) - K, r_t(1) C_e(t,S_t,K,T)] (20)$$

$$C_t^l = \max[E_t(S_T) - K, r_t(T) C_e(t,S_t,K,T)], (21)$$

then the price $C_a(t,S_t,K,T)$ of an American call can be written as:

$$C_a(t,S_t,K,T) = \begin{cases} w_1 C_t^u + (1-w_1) C_t^l \ \ if E_t(S_T) \geq K \\ w_2 C_t^u + (1-w_2) C_t^l \ \ if E_t(S_T) < K \end{cases} . (22)$$

Let $C_e^{\mathrm{HER}}(t,S_t,K,T,s_t,q_t^*)$ be the price of a European call of strike K and maturity T where q_t^* denotes the vector of parameters that describes the risk neutral density. Under hypotheses(1)-(6), $C_e^{\mathrm{HER}}(t,S_t,K,T,s_t,q_t^*)$ is given by:

$$\begin{aligned} C_e^{\mathrm{HER}}(t,S_t,K,T,s_t,q_t^*) &= e^{-r_t(T)(T-t)} \int_0^{\infty} (S_T - K)^+ \tilde{q}_t^{\mathrm{HER}}(z, s_t, q_t^*) dz \\ C_e^{\mathrm{HER}}(t,S_t,K,T,s_t,q_t^*) &= e^{-r_t(T)(T-t)} \sum_{k=0}^{+\infty} a_{k,t} b_{k,t} \end{aligned} \quad (23)$$

where S_T is given by (16) and by definition of a basis:

$$a_{k,t} = \int_{-\infty}^{+\infty} \left(S_t \exp\left[(m_t - \frac{1}{2} s_t^2)(T-t) + s_t \sqrt{T-t} z \right] - K \right)^+ f_k(z) n(z) dz \ (24)$$

and $b_{k,t}$, $k=1, 2,...$ represent the implicit price of Hermite polynomial risk $f_k(z)$[5] which needs to be estimated so that $q_t^*=(b_{0,t}, b_{0,t},...)$.

The derivation of expression (23) can be found in Appendix.

Replacing in (18) gives the RND of z:

$$\tilde{q}_t^{\mathrm{HER}}(z, s_t, q_t^*) = \sum_{k=0}^{+\infty} b_{k,t} f_k(z) n(z). (25)$$

For a practical purpose, the sum is truncated up to an arbitrary order L_b. When the sum is truncated up to an order L_b, then the density (25) may lead to some negative values for some given $b_{k,t}$, $k=1, 2,..L_b$. Balistkaia and Zolotuhina (1988) give the positivity constraints when $L_b=6$ and Jondeau and Rockinger (1999) give an ingenious way to implement positivity's constraints when $L_b=4$. For simplifications reasons and since we only need moments up to the fourth order, we restrict our model to $L_b=4$. Madan and Milne (1994) then show that the risk neutral density of the future underlying asset can be written as:

$$q_t^{\mathrm{HER}}(S, s_t, q_t^*) = q_t^{\mathrm{BS}}(S, s_t) P_H(h), (26)$$

where

$$P_H(h) = \left[b_{0,t} - \frac{b_{2,t}}{\sqrt{2}} + \frac{3b_{4,t}}{\sqrt{24}} + (b_{1,t} - 3\frac{b_{3,t}}{\sqrt{6}})h + (\frac{b_{2,t}}{\sqrt{2}} - \frac{6b_{4,t}}{\sqrt{24}})h^2 + \frac{b_{3,t}}{\sqrt{6}}h^3 + \frac{b_{4,t}}{\sqrt{24}}h^4 \right] (27)$$

$$h = \frac{\ln(S) - \left[\ln(S_t) + (r_t - d_t - \frac{1}{2} s_t^2)(T-t) \right]}{s_t \sqrt{T-t}}, (28)$$

and $q_t^{\mathrm{BS}}(S, s_t)$ is given by (14).

One can choose to estimate $b_{k,t}$, $k=1,...,4$ or follow Abken, Madan and Ramamurtie (1996) by imposing $b_{0,t}=1$ $b_{1,t}=0$, $b_{2,t}=0$ and estimate s_t, $b_{3,t}$ and $b_{4,t}$ only (See Appendix for technical details on restrictions on $b_{0,t}$, $b_{1,t}$ $b_{2,t}$ and positivity constraints on $b_{3,t}$ and $b_{4,t}$).

We wish in the next section to estimate the subjective density, in order to compute the absolute risk aversion function (7).

---

[5] The Hermite polynomials through the fourth order are:

$$f_0(z) = 1,\ f_1(z) = z,\ f_2(z) = \frac{1}{\sqrt{2}}(z^2 - 1)$$

$$f_3(z) = \frac{1}{\sqrt{6}}(z^3 - 3z),\ f_4(z) = \frac{1}{\sqrt{24}}(z^4 - 6z^2 + 3)$$

3.2 Subjective Model

To estimate the SD, we discretize equation (16) after applying Ito's lemma which straight gives us:

$$x_{(k+1)\Delta t} = x_{k\Delta t} + \left(m_{k\Delta t} - \frac{1}{2} s_{k\Delta t}^{2} \right)\Delta t + s_{k\Delta t}\sqrt{\Delta t}\, e_{(k+1)\Delta t}, (29)$$

where $x_{k\Delta t} = \ln(S_{k\Delta t})$ and Dt is a time discretization step (Dt=1/260 for daily data), kDt, $k=1,...N$, are the dates of discretization with $t=NDt$. For example, if data are daily, t will equal one year. After a change

of probability $e_{(k+1)Dt}$ will have the following distribution $\tilde{p}_{k\Delta t}^{\mathrm{HER}}(z)$:

$$\tilde{p}_{k\Delta t}^{\mathrm{HER}}(z) = n(z)\left[1 + \frac{3\hat{b}_{4,k\Delta t}}{\sqrt{24}} - 3\frac{\hat{b}_{3,k\Delta t}}{\sqrt{6}} z - \frac{6\hat{b}_{4,k\Delta t}}{\sqrt{24}} z^{2} + \frac{\hat{b}_{3,k\Delta t}}{\sqrt{6}} z^{3} + \frac{\hat{b}_{4,k\Delta t}}{\sqrt{24}} z^{4}\right] (30)$$

The general idea of the method is that parameters s_{kDt}, $b_{3,kDt}$ and $b_{4,kDt}$ are the same than those estimated in the previous section for the date $t=kDt$ because they are invariant when we switch from risk neutral world to real world. So we can consider them as observed variables. The only parameter to estimate is the drift m_{kDt}; to allow this latter to vary across time, we can write it as:

$$m_{(k+1)\Delta t} = a_{0} + a_{1} m_{k\Delta t} + b_{1} e_{(k+1)\Delta t}, (31)$$

where a_0, a_1 and b_1 are to be estimated.

Once we have estimated m_{kDt}, the subjective density $p_{k\Delta t}^{\mathrm{HER}}(S,q)$, where q denotes the vector of parameters to be estimated, that is a_0, a_1 and b_1, of S_{kDt} is known and is given by:

$$p_{k\Delta t}^{\mathrm{HER}}(S,q) = p_{k\Delta t}^{\mathrm{BS}}(S, \hat{s}_{k\Delta t}, m_{k\Delta t})\left[1 + \frac{3\hat{b}_{4,k\Delta t}}{\sqrt{24}} - 3\frac{\hat{b}_{3,k\Delta t}}{\sqrt{6}} h - \frac{6\hat{b}_{4,k\Delta t}}{\sqrt{24}} h^{2} + \frac{\hat{b}_{3,k\Delta t}}{\sqrt{6}} h^{3} + \frac{\hat{b}_{4,k\Delta t}}{\sqrt{24}} h^{4}\right],$$

where h is given by

$$h = \frac{\ln(S) - \left[\ln(S_{k\Delta t}) + (m_{k\Delta t} - \frac{1}{2}\hat{s}_{k\Delta t}^{2})\Delta t\right]}{\hat{s}_{k\Delta t}\sqrt{\Delta t}},$$

and $p_{k\Delta t}^{\mathrm{BS}}(S, \hat{s}_{k\Delta t}, m_{k\Delta t})$ is given by (13).

The risk aversion function for Hermite polynomials model is then given by:

$$A_t^{\mathrm{HER}}(S) = \frac{p_t^{\mathrm{HER}'}(S)}{p_t^{\mathrm{HER}}(S)} - \frac{q_t^{\mathrm{HER}'}(S)}{q_t^{\mathrm{HER}}(S)} = \frac{l_t}{S}. \quad (32)$$

Analytic form of those functions are given in Appendix.

4 Results

4.1 Data description

We consider the case of the CAC 40 index[6] and short time-to-maturity CAC 40 options[7].

The whole database has been provided by the SBF-Bourse de Paris (Société des Banques Françaises) which produces monthly CD-ROMs including tick-by-tick quotations of the CAC 40 caught every 30 seconds, and all equities options prices quoted on the MONEP tick-by-tick. The database includes time quotation, maturity, strike price, closing and settlement quotes for all calls and puts and volume from January 1995 through June 1997. Short maturity CAC 40 options prices need to be adjusted for dividends. Aït-Sahalia and Lo (1998a) suggest to extract an implied forward underlying asset F_t using the call-put parity on the at-the-money option which requires that the following equation holds:

$$C_{\mathrm{atm}} - P_{\mathrm{atm}} = e^{-r_t(T)(T-t)}(F_t - K_{\mathrm{atm}}) \quad (33)$$

where C_{atm}, P_{atm} and K_{atm} respectively denote the price of the call, the price of the put and the strike at-the-money. Once we have obtained F_t, we may deduce the implied dividend $d_t(T)$ at time t for a maturity T using the arbitrage relation between F_t and S_t

$$F_t = e^{(r_t(T)-d_t(T))(T-t)} S_t \quad (34)$$

Since CAC 40 options data contains many misspriced prices, once needs to filter the data very carefully. First following Aït-Sahalia and Lo (1998a), we drop options with price less than 1/8.

[6] CAC 40 index leans on the major shares of Paris Stock Market. It is constructed from 40 shares quoted on the monthly settlement market and selected in accordance with several requirements (capitalization, liquidity,...). CAC 40 is computed by taking the arithmetical\ average of assets quotations which compose it, weighted by their capitalization.

[7] CAC 40 options are traded on the MONEP (Marché des Options Négociables de Paris). They are american type and there are four expiration dates for each date: 3 months running and a quarterly maturity among March, June, September or December. Two consecutives strike prices are separated by a standard interval of 25 basis points.

Second, for our study, we kept the most liquid maturity which usually appears to be the closest to 30 days yield-to-maturity.

Table 1 shows summary statistics of the CAC 40 index return historical distribution. Negative skewness and positive excess kurtosis show nonnormality of historical distribution, implying a leptokurtic and skewed distribution. Statistic *W* used by Jarque and Bera (1980) to construct a normality test allows to reject normality at 95%.

The Ljung-Box (1978) statistic LB(20) to test heteroskedasticity rejects the homoskedasticity for the square returns. The Ljung-Box (1978) statistic *LB(20)* corrected for heteroskedasticity computed with 20 lags allows to detect autocorrelation returns. Diebold (1988) suggests a Ljung-Box statistic corrected for heteroskedasticity LB_c. We notice that autocorrelation of squared returns is significantly higher than autocorrelation of returns, which implies than large changes tend to be followed by large changes, of either sign.

Table 1: Descriptive statistics of the CAC 40 daily index return for the period from January 1995 to June 1997. Table 1 shows several statistics describing returns series: mean, standard deviation, skewness and excess kurtosis. *LB(20)* is the Ljung-Box statistic to test heteroskedasticity. ρ(h) is the autocorrelation of order *h*. $LB_c(20)$ is the Ljung-Box statistic corrected for heteroskedasticity for the nullity test of the 20 first autocorrelations of returns. Under nullity hypothesis, this statistic is distributed as $\chi^2(2)$ with 20 degrees of freedom. *W* is the Jarque and Bera (1980) statistic that allows to test for normality[8].

[8]Jarque and Bera's statistic is based on empirical skewness, *sk* and kurtosis *kt* given by:

$$sk = \frac{1}{N}\sum_{t=1}^{N}\frac{(x_t - \overline{m})^3}{\overline{s}^3} \text{ et } kt = \frac{1}{N}\sum_{t=1}^{N}\frac{(x_t - \overline{m})^4}{\overline{s}^4}$$

where $\overline{m}$ and $\overline{s}$ represent respectively the empirical mean and empirical standard deviation.
We note by t_1 and t_2 the following statistics:

$$t_1 = \sqrt{N\frac{sk^2}{6}},\ t_2 = \sqrt{N\frac{(kt-3)^2}{24}}.$$

Under the nul hypothesis of normality, the Jarque and Bera's statistic $W = t_1^2 + t_2^2$ asymptotically follows a $\chi^2(2)$.

	x_t	x_t^2
Number of observations	650	650
Mean	$0.60\ 10^{-3}$	$0.99\ 10^{-4}$
Standard Deviation	$1.00\ 10^{-2}$	$0.16\ 10^{-3}$
Skewness	-0.163	3.913
Excess kurtosis	0.928	21.205
LB(20)	26.740	49.166
r(1)	-0.010	0.010
r(5)	-0.081	-0.031
r(10)	-0.033	0.092
r(20)	-0.022	0.070
LB_c(20)	26.692	25.182
W	26.178	13836.563

4.2 Estimations' procedures

A non-linear least squares method is implemented to estimate risk neutral parameters. At each date t, the non-linear least squared estimator (NLLSE) $\hat{b}^*_{NLLSE} = \left\{ s_t, b_{3,t}, b_{4,t}, w_{1,t}, w_{2,t} \right\}$ is obtained so that it minimises the distance between observed and theoretical implied volatilities computed with Hermite polynomials' model (s_i^{BS} for observed ones and s_i^{HER} for theoretical ones):

$$b^*_{NLLSE} = \underset{b^* \in \Theta^*}{\arg\min} \sum_{i=1}^{m_c} \left(s_i^{\mathrm{BS}} - s_i^{\mathrm{HER}}(b^*) \right)^2, (35)$$

where m_c denotes the number of observed call options at date t, $\Theta^* = \left(R^+, D, [0,1], [0,1] \right)$ where D is the domain of ($b_{3,t}$, $b_{4,t}$) for which (25) remains positive for all z (see figure 3).

Subjective model (29)-(31) is estimated by maximum likelihood method. The log-likelihood function L of $x=(x_1,\ldots,x_{NDt})'$ is given by:

$$L(x; b) = \sum_{k=1}^{N} L_{k\Delta t}(x_{k\Delta t}) \ (36)$$

where L_{kDt} is the log-likelihood function of x_{kDt}.

The maximum likelihood estimator (MLE) $\hat{b}_{MLE} = \{a_0, a_1, b_1\}$ is obtained so that it maximises the following optimisation problem[9]:

$$\hat{b}_{MLE} = \underset{b \in \Theta}{\arg\max}\left[L(x; b)\right], (37)$$

where $\Theta^* = (R, R, R)$.

To find the implied coefficient of risk aversion l_t, one can solve:

$$l_t = \underset{l \in R^+}{\arg\min} \sum_{r=1}^{M} \left(\frac{p_t^{\mathrm{HER}\prime}(S_r)}{p_t^{\mathrm{HER}}(S_r)} - \frac{q_t^{\mathrm{HER}\prime}(S_r)}{q_t^{\mathrm{HER}}(S_r)} - \frac{l}{S_r} \right), (38)$$

where M is a constant and S_r, $r=1,...,M$ is a range of points around the underlying at date t, S_t.

4.3 Empirical results

In this section, we analyse empirical results.

In figure (4a)-(4b), we show two estimated risk neutral densities for the dates May 1995, 5th with maturity of 56 days and July 1996, 25th with maturity 36 days. The first one corresponds to a so called agitated date during French Presidential Elections and the second corresponds to a quiet date. We notice that asymmetry is higher for the first one. The daily time series for the estimates of the parameters in a risk neutral world are shown in figure (5a)-(6b).

We notice that implied volatilities given by Hermite polynomials' model in figure (5a) appear to be larger than those obtained from Black and Scholes model in figure (5b) which seems to imply that Black-Scholes volatilities are undervalued. The different picks at the beginning of the period may come from the fact that CAC 40 options are much less liquid during 1995 than 1996. We turn to market prices of skewness b_3 in figure (6a); this latter is significantly different from zero during the whole period. Parameter b_3 gives some information about the skewness of the distribution when parameter b_4 gives information about the excess kurtosis which is significantly positive. The skewness appears to be negative along almost all the period which indicates that investors anticipate a decrease more often than an increase in the underlying index. We notice four agitated sub-periods. The first one corresponds to French presidential elections of May 1995. The second one and the

[9] Estimations have been done with the software GAUSS using Optmum routine.

third one respectively in May 1996 and February 1997 are not as so clear and may be due to perturbation in U.S. market. The latest is the French snap elections of May 1997.

During these period, market seems to be agitated which can be seen in the kurtosis. It gives an idea about extreme events.

Figure 7 shows Mean Square Errors (MSE) of parameters[10]. All MSE appear to be less than $8\ 10^{-2}$, that is quite satisfying and confirms the choice of the method. Other properties of the method is that it is computationally fast and it may take into account possible dirty data. Empirical results of these properties can be found in Coutant, Jondeau and Rockinger (1998).

In order to show the consistence of the model, we show in table 2 estimated parameters of the model under the true probability when parameters are supposed to be constant. Volatility parameter is higher than average volatility estimated in a risk neutral world. Parameter b_4 is significantly different from zero which is not the case of b_3.

In table 3, estimation of model (29)-(30) is presented. All parameters appear to be significant and the daily time series of estimated drift $m_{k\Delta\tau}$, $k=1,...,N$ from (31) with values of table 3 is given by figure 8.

Table 2: Estimation of the model (29)-(30) when parameters $(\mu_t, s_t, b_{3,t}, b_{4,t}) = (\mu, s, b_3, b_4)$ are supposed to be constant:

	μ	σ	b_3	b_4
BS	0.192 (1.898)	0.161 (29.684)		
Hermite	0.186 (1.708)	0.159 (29.247)	-0.003 (-0.191)	0.177 (3.199)

[10]MSE at date t is calculated as follow:

$$\mathrm{MSE}_t = \frac{1}{m_c - m_b} \sum_{i=1}^{m_c} \left(s_i^{\mathrm{BS}} - s_i^{\mathrm{HER}}(\hat{b}_t^*) \right)^2,$$

with the notations used in (35), m_b is the number of parameters to estimate and $\hat{b}_t^*$ is the vector of estimated parameters at date t.

Table 3: Estimation of the time varying drift μ_t in (31):

α_0	α_1	β_1
-0.113	0.481	0.296
(-3.163)	(2.177)	(5.067)

Figure (8) shows Absolute Risk Aversion functions for several days. First date is 28 February 1995 and CAC 40 moderately rose during this month: implied risk aversion coefficient l_t=4.999 is rather high. Second date is 28 April 1995, the index improved since mid-March and l_t =1.051. The date 15 July 1996, sees a short drop of the CAC 40, l_t =11.404 is very high. Finally last date takes place on 13 November 1996, during a significant growth of the underlying and l_t =3.103. We may conclude from these observations that investor's risk aversion substantially depends on the index's evolution. When CAC 40 goes up, investors have a moderate risk aversion, even they are nearly risk neutral for 13 November 1996.

Figure (9) represents the risk aversion level obtained with (38).

5 Conclusion

In this paper, we have empirically investigated investors' risk aversion coefficient implied in options prices. We showed that this latter could be estimated by the knowledge of a combination of information under risk neutral and subjective probabilities.

We have focused on CAC 40 index options, and we have supposed CRRA utility functions and an Hermite polynomial expansion for risk neutral and subjective densities. This model has the advantage to give directly the skewness and the kurtosis in addition to numerical properties. We first estimated Hermite polynomials' model under a risk neutral probability using options prices, and second injected risk-neutral parameters obtained in an equivalent discretized model under a subjective probability. We then used time series of the CAC 40 index to estimate the subjective density. A relation between densities and their derivatives allowed us to compute all absolute risk aversion functions on the period from 1995 to 1996. Risk aversion function appeared to be time varying and investors' risk aversion is very sensitive to the way underlying asset evolutes. Risk

Aversion coefficient is a good tool to test market-makers reactions to particular events or announcements.

Some future studies could turn on comparing results from several investor's preferences choices and another kind of risk, so that volatility risk for example. In a future research, we will focus on modelling the risk aversion coefficient in order to forecast the true density.

Appendix

Compute derivatives $p_t^{\mathrm{HER}\,\prime}(S)$ and $q_t^{\mathrm{HER}\,\prime}(S)$:

$$\mathrm{q}_{\mathrm{t}}^{\mathrm{HER}\,\prime}(S) = \frac{1}{s_t^2 (T-t)\sqrt{2pS}} \exp\left[-\frac{1}{2}\left(\frac{\ln(S) - m_t^*}{s_t\sqrt{T-t}}\right)^2\right]\widetilde{P}_H(h)$$

$$\widetilde{P}_H(h) = -\frac{P_H(h)}{S} - \frac{P_H(h)}{Ss_t\sqrt{T-t}}\,\frac{\ln(S) - m_t^*}{s_t\sqrt{T-t}} + \frac{Q_H(h)}{Ss_t\sqrt{T-t}}$$

$$m_t^* = \ln(S_t) + \left(r_t - d_t - \frac{1}{2}s_t^2\right)(T-t),$$

where $P_H(.)$ is given by (27), h is given by (28) and

$$Q_H(h) = \frac{3b_{3,t}}{\sqrt{6}}(-1 + h^2) + \frac{4b_{3,t}}{\sqrt{24}}(h^3 - 3h).$$

To obtain $\mathrm{p}_{\mathrm{t}}^{\mathrm{HER}\,\prime}(S)$ just replace m_t^* by:

$$m_t = \ln(S_t) + \left(m_t - \frac{1}{2}s_t^2\right)(T-t)$$

European call in the Hermite polynomials basis:

The price of a European call is given by (19)

$$C_e^{\mathrm{HER}}(t, S_t, K, T, s_t, q_t^*) = e^{-r_t(T)(T-t)}\int_{-\infty}^{+\infty}(S_T - K)^+\,\widetilde{q}_t^{\mathrm{HER}}(z, s_t, q_t^*)dz$$

$$= e^{-r_t(T)(T-t)}\int_{-\infty}^{+\infty}\left(S_t\exp\left[(r_t - d_t - \frac{1}{2}s_t^2)(T-t) + s_t\sqrt{T-t}z\right] - K\right)^+\widetilde{q}_t^{\mathrm{HER}}(z, s_t, q_t^*)dz$$

$$= e^{-r_t(T)(T-t)}\int_{-\infty}^{+\infty}\left(S_t\exp\left[(r_t - d_t - \frac{1}{2}s_t^2)(T-t) + s_t\sqrt{T-t}z\right] - K\right)^+ u_t(z, q_t^*)n(z)dz.$$

All functions can be expressed in terms of the basis so that:

$$\left(S_t \exp\left[(r_t - d_t - \frac{1}{2} s_t^2)(T-t) + s_t \sqrt{T-t}z \right] - K \right) = \sum_{k=0}^{+\infty} a_{k,t} f_k(z)$$

$$u_t(z, q_t^*) = \sum_{j=0}^{+\infty} b_{j,t} f_j(z)$$

then

$$C_e^{\mathrm{HER}}(t, S_t, K, T, s_t, q_t^*) = e^{-r_t(T)(T-t)} \int_{-\infty}^{+\infty} \sum_{k=0}^{+\infty} a_{k,t} f_k(z) \sum_{j=0}^{+\infty} b_{j,t} f_j(z) n(z) dz$$

$$= e^{-r_t(T)(T-t)} \sum_{k=0}^{+\infty} \sum_{j=0}^{+\infty} a_{k,t} b_{j,t} \int_{-\infty}^{+\infty} f_k(z) f_j(z) n(z) dz$$

$$= e^{-r_t(T)(T-t)} \sum_{k=0}^{+\infty} a_{k,t} b_{k,t}$$

Parameters $a_{k,t}$:

Coefficients $a_{k,t}$ for the call price are given by:

$$a_{k,t} = a(k, S_0, x, m, s, t) = \left. \frac{\P\Phi(u, S_0, x, m, t)}{\P u^k} \right|_{u=0} \frac{1}{\sqrt{k!}} \quad (39)$$

$$\Phi(u, S_0, x, m, s, t) = S_0 \exp(mt + s\sqrt{t}z) N(d_1(u)) - xN(d_2(u)) \quad (40)$$

Explicitly $a_{k,t}$ are given as follows:

$$d_1 = \frac{1}{s\sqrt{T-t}}\ln(F_t) + \frac{1}{2}s\sqrt{T-t},\ d_2 = d_1 - s\sqrt{T-t},$$
$$a_{0,t} = F_t N(d_1) - KN(d_2),$$
$$a_{1,t} = s\sqrt{T-t}F_t N(d_1) + F_t n(d_1) - Kn(d_2),$$
$$a_{2,t} = \frac{1}{\sqrt{2}}\left[\left(s\sqrt{T-t}\right)^2 F_t N(d_1) + 2s\sqrt{T-t}F_t n(d_1) + F_t n'(d_1) - Kn'(d_2)\right],$$
$$a_{3,t} = \frac{1}{\sqrt{6}}\left[\left(s\sqrt{T-t}\right)^3 F_t N(d_1) + 3\left(s\sqrt{T-t}\right)^2 F_t n(d_1) + 3s\sqrt{T-t}F_t n'(d_1) + F_t n''(d_1) - Kn''(d_2)\right],$$
$$a_{4,t} = \frac{1}{\sqrt{24}}\begin{bmatrix}\left(s\sqrt{T-t}\right)^4 F_t N(d_1) + 4\left(s\sqrt{T-t}\right)^2 F_t n(d_1) + 6\left(s\sqrt{T-t}\right)^2 F_t n''(d_1) \\ +4s\sqrt{T-t}F_t n''(d_1) + F_t n^{(3)}(d_1) - Kn^{(3)}(d_2)\end{bmatrix}$$

where *n(.)* and *N(.)* are the normal and cumulative normal densities.

Restrictions on parameters in Hermite's model:

Let h be:

$$h = \frac{\ln(S_T) - \left[\ln(S_t) + (r_t - d_t - \frac{1}{2}s_t^2)(T-t)\right]}{s_t\sqrt{T-t}},$$

the risk neutral distribution of h is:

$$\tilde{q}_t^{\mathrm{HER}}(z) = n(z)P_H(z),$$

where $P_H(z)$ is given by (27) and *n(z)* is the Gaussian distribution with mean 0 and variance 1.

$\tilde{q}_t^{\mathrm{HER}}(z)$ must satisfy:

$$\int_{-\infty}^{+\infty}\tilde{q}_t^{\mathrm{HER}}(z)dz = 1,.$$

which implies that

$$\int_{-\infty}^{+\infty}n(z)\left[b_{0,t} - \frac{b_{2,t}}{\sqrt{2}} + \frac{3b_{4,t}}{\sqrt{24}} + (b_{1,t} - 3\frac{b_{3,t}}{\sqrt{6}})z + (\frac{b_{2,t}}{\sqrt{2}} - \frac{6b_{4,t}}{\sqrt{24}})z^2 + \frac{b_{3,t}}{\sqrt{6}}z^3 + \frac{b_{4,t}}{\sqrt{24}}z^4\right]dz = 1,$$
$$\left[b_{0,t} - \frac{b_{2,t}}{\sqrt{2}} + \frac{3b_{4,t}}{\sqrt{24}} + (\frac{b_{2,t}}{\sqrt{2}} - \frac{6b_{4,t}}{\sqrt{24}}) + 3\frac{b_{4,t}}{\sqrt{24}}z^4\right] = 1,$$
$$b_{0,t} = 1$$

for all t.

We also want to impose that the future underlying asset's expectation equals the current future price, that is:

$$E_t(S_T) = S_t e^{(r_t - d_t)(T-t)} \Leftrightarrow E_t(h) = 0,$$

$$\int_{-\infty}^{+\infty} z \tilde{q}_t^{\text{HER}}(z) dz = 0,$$

which gives the restriction for parameter $b_{1,t}$:

$$\int_{-\infty}^{+\infty} z n(z) \left[b_{0,t} - \frac{b_{2,t}}{\sqrt{2}} + \frac{3b_{4,t}}{\sqrt{24}} + (b_{1,t} - 3\frac{b_{3,t}}{\sqrt{6}})z + (\frac{b_{2,t}}{\sqrt{2}} - \frac{6b_{4,t}}{\sqrt{24}})z^2 + \frac{b_{3,t}}{\sqrt{6}} z^3 + \frac{b_{4,t}}{\sqrt{24}} z^4 \right] dz = 0,$$

$$\left[(b_{1,t} - 3\frac{b_{3,t}}{\sqrt{6}}) + 3\frac{b_{3,t}}{\sqrt{6}} \right] = 0,$$

$$b_{1,t} = 0$$

for all t.

Finally, third restriction comes from variance which is imposed to be the same under the transformed measure than under the reference measure:

$$\int_{-\infty}^{+\infty} z^2 \tilde{q}_t^{\text{HER}}(z) dz = 1,$$

$$\int_{-\infty}^{+\infty} z^2 n(z) \left[b_{0,t} - \frac{b_{2,t}}{\sqrt{2}} + \frac{3b_{4,t}}{\sqrt{24}} + (b_{1,t} - 3\frac{b_{3,t}}{\sqrt{6}})z + (\frac{b_{2,t}}{\sqrt{2}} - \frac{6b_{4,t}}{\sqrt{24}})z^2 + \frac{b_{3,t}}{\sqrt{6}} z^3 + \frac{b_{4,t}}{\sqrt{24}} z^4 \right] dz = 1,$$

$$\left[b_{0,t} - \frac{b_{2,t}}{\sqrt{2}} + \frac{3b_{4,t}}{\sqrt{24}} + 3(\frac{b_{2,t}}{\sqrt{2}} - \frac{6b_{4,t}}{\sqrt{24}}) + 15\frac{b_{4,t}}{\sqrt{24}} \right] = 1,$$

$$b_{2,t} = 0$$

for all t.

Positivity's constraints on parameters $b_{3,t}$ and $b_{4,t}$:

Let g and g be the skewness and excess kurtosis respectively. A straight calculus leads to:

$$g_1 = \int_{-\infty}^{+\infty} z^3 \tilde{q}_t^{\text{HER}}(z) dz = \sqrt{6} b_3, (41)$$

$$g_2 = \int_{-\infty}^{+\infty} z^4 \tilde{q}_t^{\text{HER}}(z) dz - 3 = \sqrt{24} b_4. (42)$$

Then (25) can be rewritten in terms of g_1 and g_2:

$$\tilde{q}_t^{\text{HER}}(z, s_t, q_t^*) = n(z)\left[1 + \frac{g_1}{6} H_3(z) + \frac{g_2}{24} H_4(z)\right],$$

where $H_j(z) = \sqrt{j!}\, f_j(z)$ is the non standardised Hermite polynomial of order *j*.

Density (25) remains positive when

$$P_H(z) = 1 + \frac{g_1}{6} H_3(z) + \frac{g_2}{24} H_4(z) \geq 0.$$

Jondeau and Rockinger (1999) explain that this is the case if a couple (g_1, g_2) lies within the envelope generated by the hyperplane $P_H(z)=0$, with $z \in R$. This envelope is given by the system

$$\begin{cases} P_H(z) = 0, \\ P_H'(z) = 0, \end{cases}$$

with

$$P_H'(z) = \frac{g_1}{6} H_2(z) + \frac{g_2}{24} H_3(z).$$

They find that solving the problem gives explicitly g_1 and g_2 as a function of *z*:

$$\begin{cases} g_1(z) = -24 \dfrac{H_3(z)}{d(z)}, \\ g_2(z) = 72 \dfrac{H_2(z)}{d(z)}, \end{cases}$$

with

$$d(z) = 4H_3^2(z) - 3H_2(z)H_4(z).$$

After some demanding calculus, Jondeau and Rockinger (1999) find numerically and analytically that the authorised domain for g_1 and g_2 is a steady, continuous and concave curve. The domain for b_3 and b_4 is given by figure 3.

Captions

Figure 1a: Daily CAC 40 index over the period January 1995 to July 1997.

Figure 1b: Daily CAC 40 index returns over the period January 1995 to July 1997.

Figure 2a: CAC 40 volatility smile for the date 05/05/1995 and the maturity 56 days.

Figure 2b: CAC 40 volatility smile for the date 25/07/1996 and the maturity 36 days.

Figure 3: Domain authorised by the skewness and the kurtosis for positivity constraint of an Hermite polynomials' density

Figure 4a: Risk neutral density for the CAC 40 computed with Hermite polynomials for the date 05/05/1995 and the maturity 56 days.

Figure 4b: Risk neutral density for the CAC 40 computed with Hermite polynomials for the date 25/07/1996 and the maturity 36 days.

Figure 5a: Estimation of parameter s_t in Hermite's model under the risk neutral probability.

Figure 5b: Estimation of implied Black's volatilities under the risk neutral probability.

Figure 6a: Estimation of parameter $b_{3,t}$ in Hermite's model under the risk neutral probability.

Figure 6b: Estimation of parameter $b_{4,t}$ in Hermite's model under the risk neutral probability.

Figure 7: Mean Squares Errors for the estimation of risk neutral parameters in Hermite's model.

Figure 8: Graphs of implied absolute risk aversion functions for the dates 28/02/1995, 28/04/1995, 15/07/1996 and 13/11/1996.

Figure 9: Implied risk aversion's coefficients for the period January 1995 to July 1997.

References

Abken, P.A., D. Madan and S. Ramamurtie (1996), "Estimation of Risk Neutral Densities by Hermite Polynomial Approximation: With an Application to Eurodollar Future Options", Federal Reserve Bank of
Atlanta, *Working Paper* 96-5.

Aït-Sahalia, Y. et A.W. Lo (1998a), "Nonparametric Estimation of State-price Densities Implicit in Financial Asset Prices", *Journal of Finance,* 53, 499-547.

Aït-Sahalia, Y. et A.W. Lo (1998b), "Nonparametric Risk Management and Implied Risk Aversion", Working paper, University of Chicago, forthcoming in *Journal of Econometrics*.

Bahra, B., (1996), "Probability Distributions of Futures Asset Prices Implied by Option Prices", *Bank of England Quarterly Bulletin*, 299-311.

Balitskaia, E.O. and L.A. Zolotuhina (1988), "On the Representation of a Density by an Edgeworth Series", *Biometrika*, 75(1), 185-187.

Bates, D. S. (1991), "The Crash of '87: Was it Expected? The Evidence From Options Markets", *Journal of Finance*, 46(3), 1009-1044.

Bates, D. S. (1996b), "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options", *Review of Financial Studies*, 9(1), 69-107.

Black, F. and M. Scholes (1973), "The Pricing of Options and Corporate Liabilities", *Journal of Political Economy*, 81(3), 637-654.

Breeden, D.T., and R.H. Litzenberger (1978), "Prices of State-contingent Claims Implicit in Option Prices",
Journal of Business, 51(4), 621-651.

Campa, J., P.H.K. Chang and R.L. Reider (1997b), "Implied Exchange Rate Distributions: Evidence from OTC Options Markets", *NBER* Working Paper N°6179.

Chesney, M. and L. Scott (1989), "Pricing European Currency Options: a Comparison of the Modified Black-Scholes Model and a Random Variance Model", *Journal of Financial and Quantitative*
Analysis, 24, 267-284.

Cox, J.C., J.I. Ingersoll and S.A. Ross (1985b), "A Theory of the Term Structure of Interest Rates", *Econometrica*, 53(2), 385-407.

Coutant, S., E. Jondeau and M. Rockinger (1998), "Reading Interest Rate and Bond Future Options' Smiles Around the 1997 French Snap Election", *CEPR* n°2010.

Derman, E., and I. Kani (1994), "Riding on a Smile", *RISK*, 7, 32-39.

Diebold, F.X. (1988), *Empirical Modeling of Exchange Rate Dynamics*, Springer-Verlag, New-York.

Dupire, B. (1994), "Pricing with a Smile", *RISK*, 7, 18-20.

Engle, R.F. (1982), "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of the United Kingdom Inflation", *Econometrica*, 50, 987-1008.

Harrison, J.M. and D. Kreps (1979), "Martingales and Arbitrage in Multiperiod Securities Markets", *Journal of Economic Theory*, 20, 381-408.

Hull, J.C., and A.D. White (1987), "The Pricing of Options on Assets with Stochastic Volatilities", *Journal of Finance*, 42(2), 281-300.

Heston, S.L. (1993), "A Closed-form Solution for Options on Assets with Stochastic Volatility with Application to Bond and Currency Options", *Review of Financial Studies*, 6(2), 327-343.

Jackwerth, J.C. (1996), "Recovering Risk Aversion from Option Prices and Realized Returns", Working paper, London Business School.

Jarque, C.M. and A.K. Bera (1980), "Efficient Test of Normality, Homoscedasticity and Serial Independence of Regression Residuals", *Economic Letters*, 6, 255-259.

Jarrow, R. and A. Rudd (1982), "Approximate Valuation for Arbitrary Stochastic Processes ", *Journal of Financial Economics*, 10(3), 349-369.

Jondeau, E. and M. Rockinger (1998), "Reading the Smile: The Message Conveyed by Methods which Infer Risk Neutral Densities.", *CEPR* n°2009.

Jondeau, E. and M. Rockinger (1999), "Estimating Gram-Charlier Expansions with Positivity Constraints", Banque de France, NER n°56.

Lucas, R. (1978), "Asset Prices in Exchange Economy", *Econometrica*, 46, 1429-1446.

Ljung, G. and G. Box (1978), "On a Measure of Lack of Fit in Time Series Models", *Biometrika*, 66, 67-72.

Madan, D.B., and F. Milne (1994), "Contingent Claims Valued and Hedged by Pricing and Investing in a Basis", *Mathematical Finance*, 4, 223-245.

Melick, W.R. and C.P. Thomas (1997), "Recovering an Asset's Implied PDF from Options Prices: an Application to Crude Oil During the Gulf crisis", *Journal of Financial and Quantitative Analysis*, 32, 91-116.

Merton, R. (1969), "Lifetime Portfolio Selection Under Uncertainty: the Continuous-time Case", *Review of Economics and Statistics*, 51, 847-857.

Merton, R. (1971), "Optimum Consumption and Portfolio Rules in a Continuous Time Model", *Journal of Economic Theory*, 3, 373-413.

Pratt, J.W. (1964), "Risk Aversion in the Small and in the Large", *Econometrica*, 32, 122-136.

Rubinstein, M. (1994), "Implied Binomial Trees", *Journal of Finance*, 69, 771-818.

Shimko, D. (1993), "Bounds of Probability", *Risk*, 6(4), 34-36.

Figure 1a: daily CAC 40 index
over the time period January 1995 to July 1997

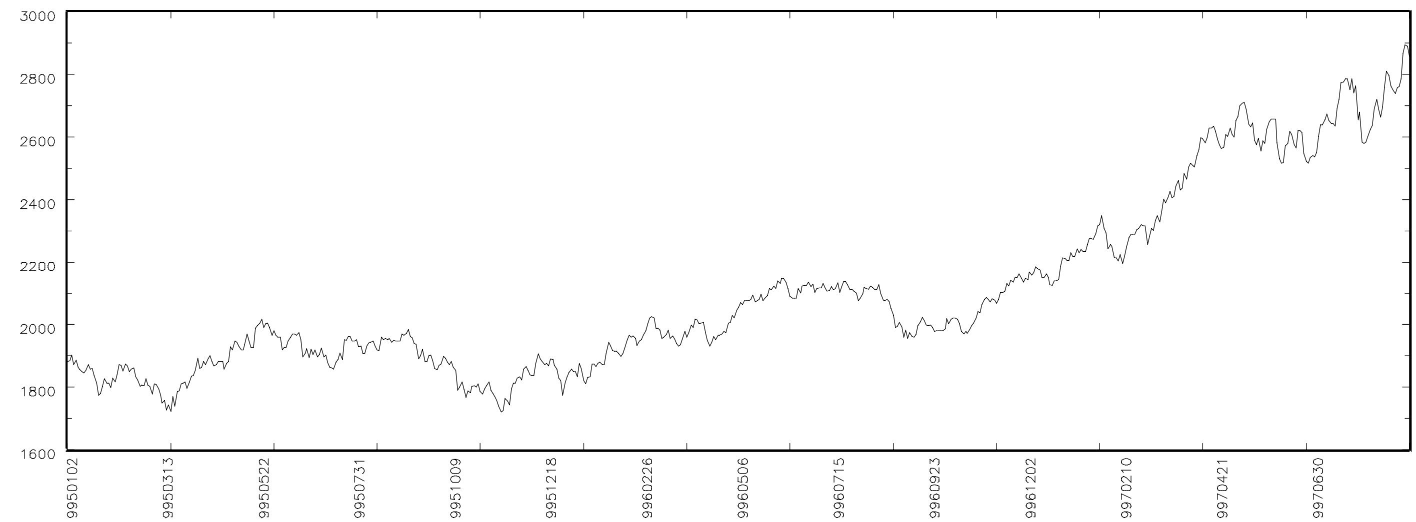

Figure 1b: daily returns on the CAC 40 index
over the time period January 1995 to July 1997

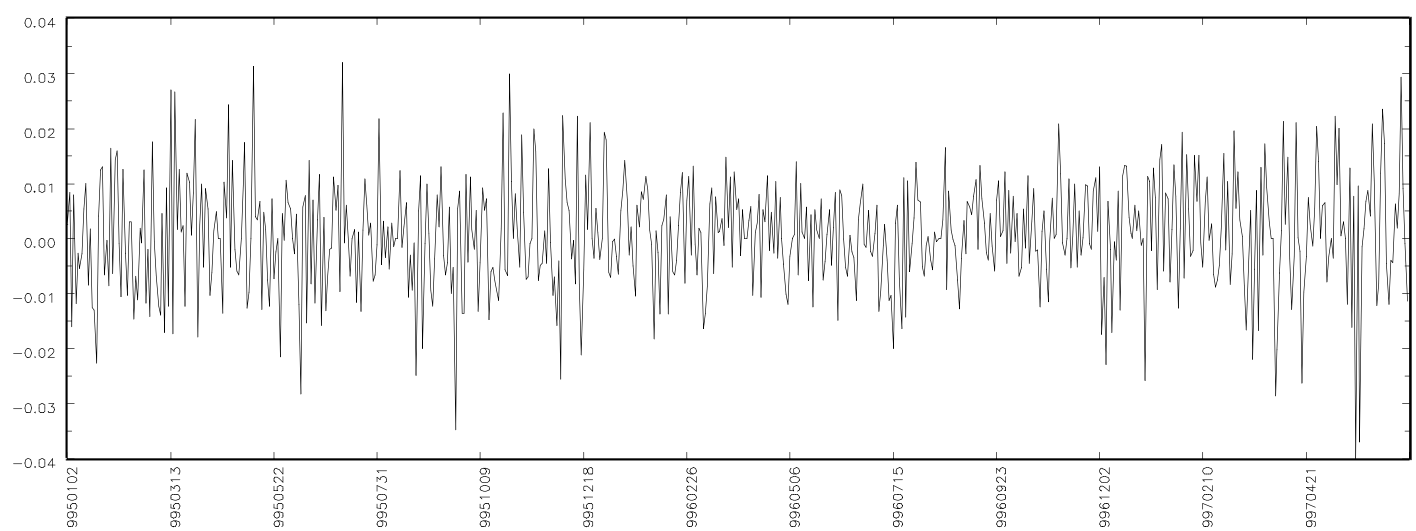

Figure 2a: volatility smile for the date 19950505

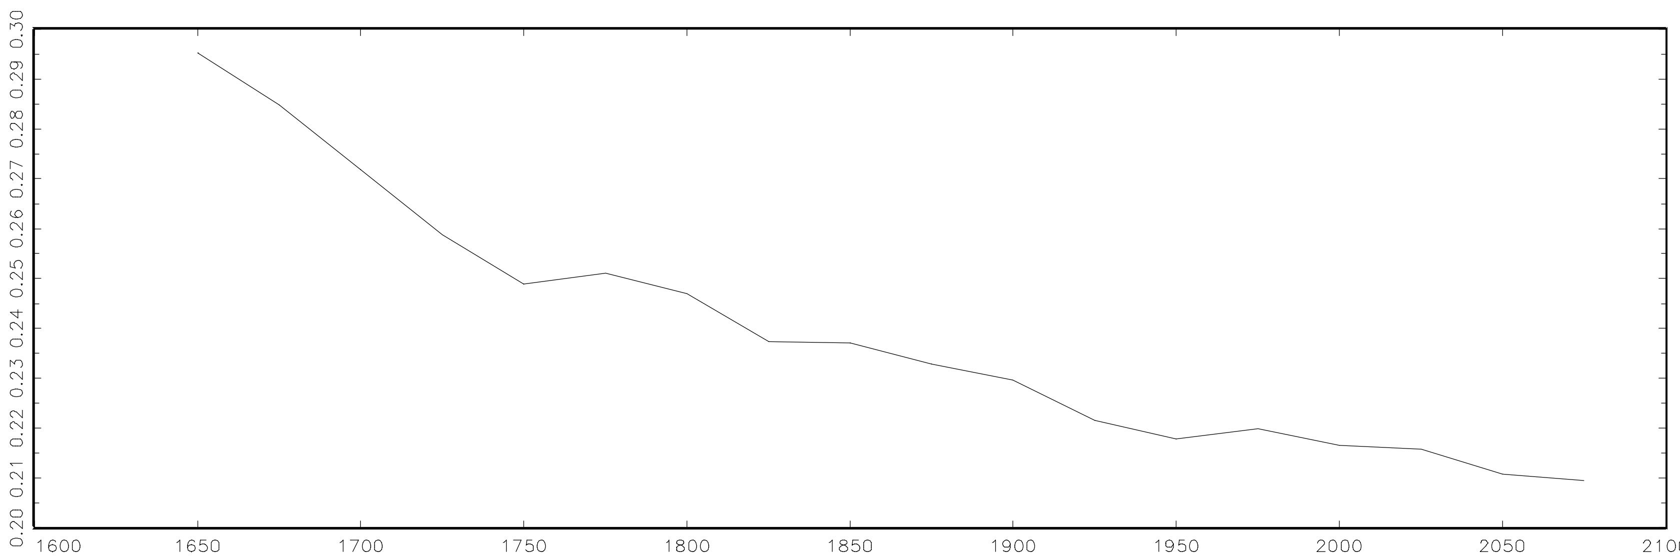

Figure 2b: volatility smile for the date 19960725

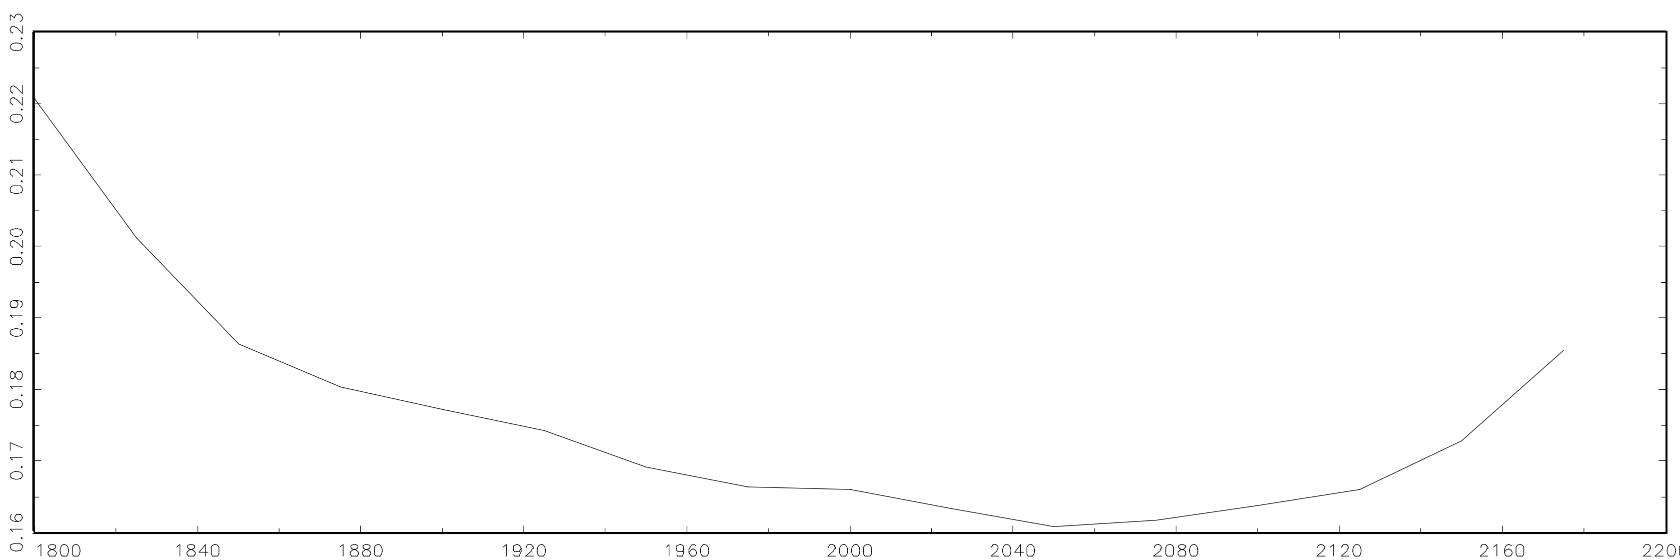

Figure 3: authorized domain for skewness and kurtosis

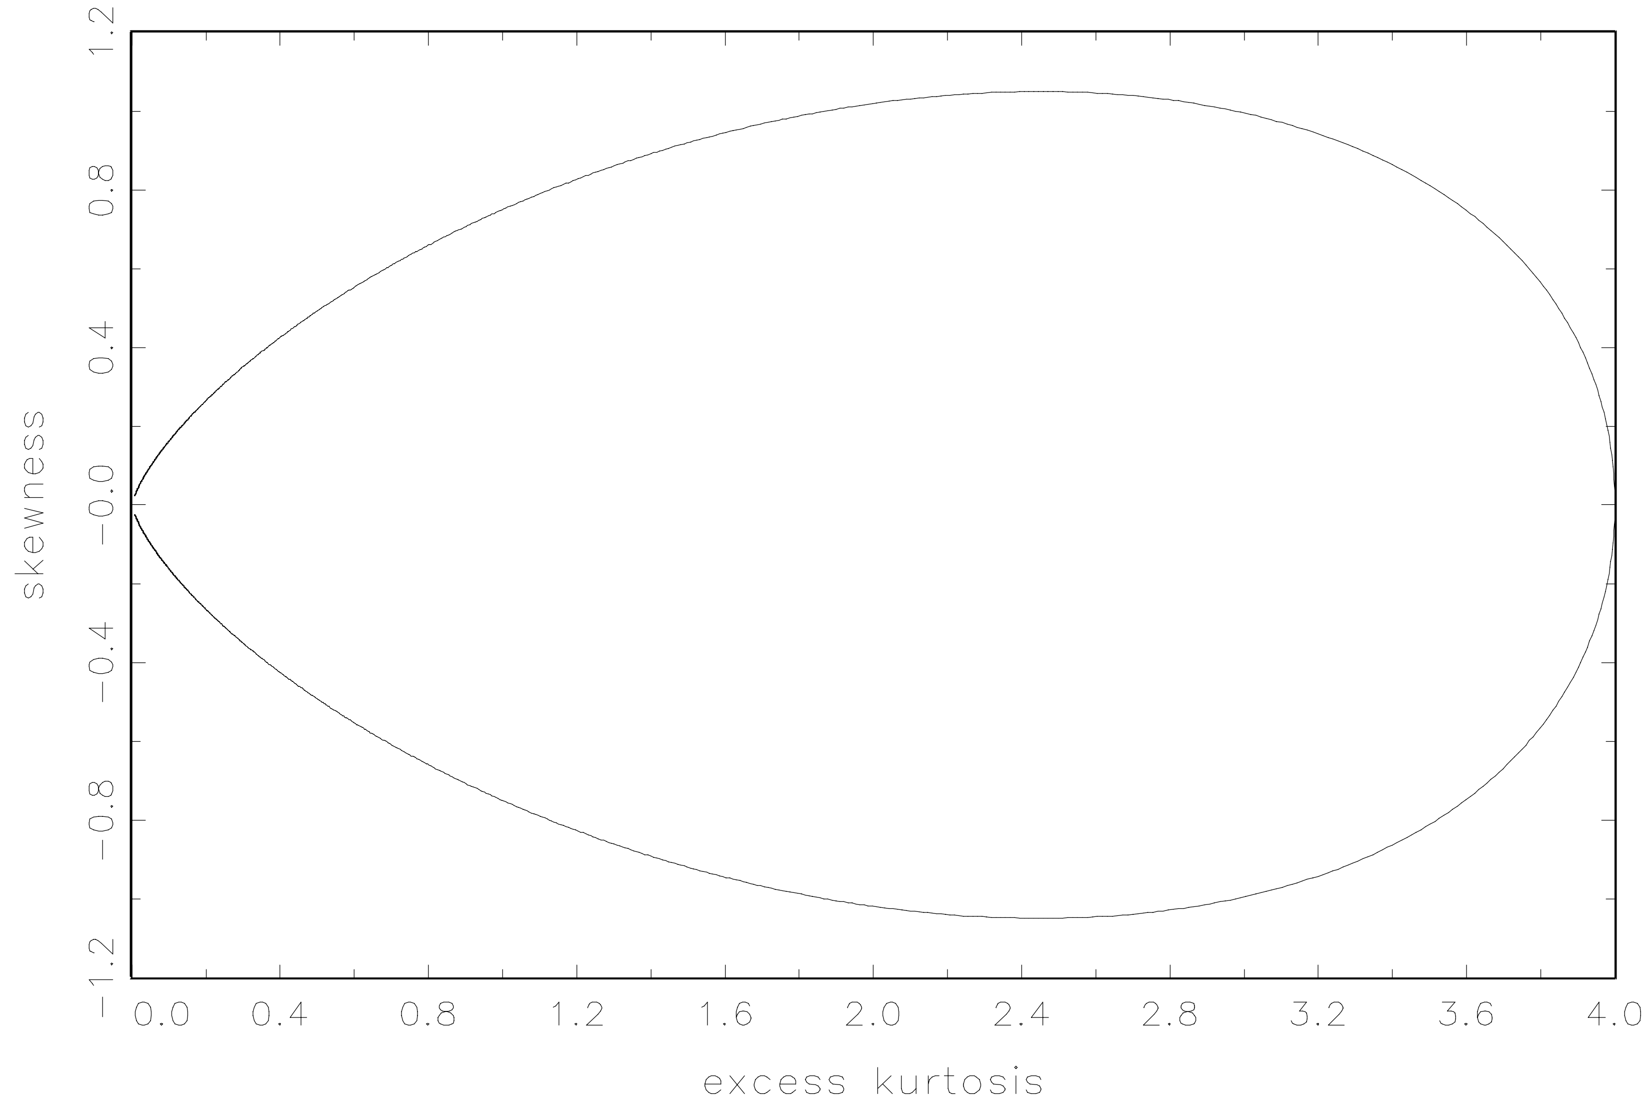

Figure 4a: RND for the date 05/05/1995 and maturity of 56 days

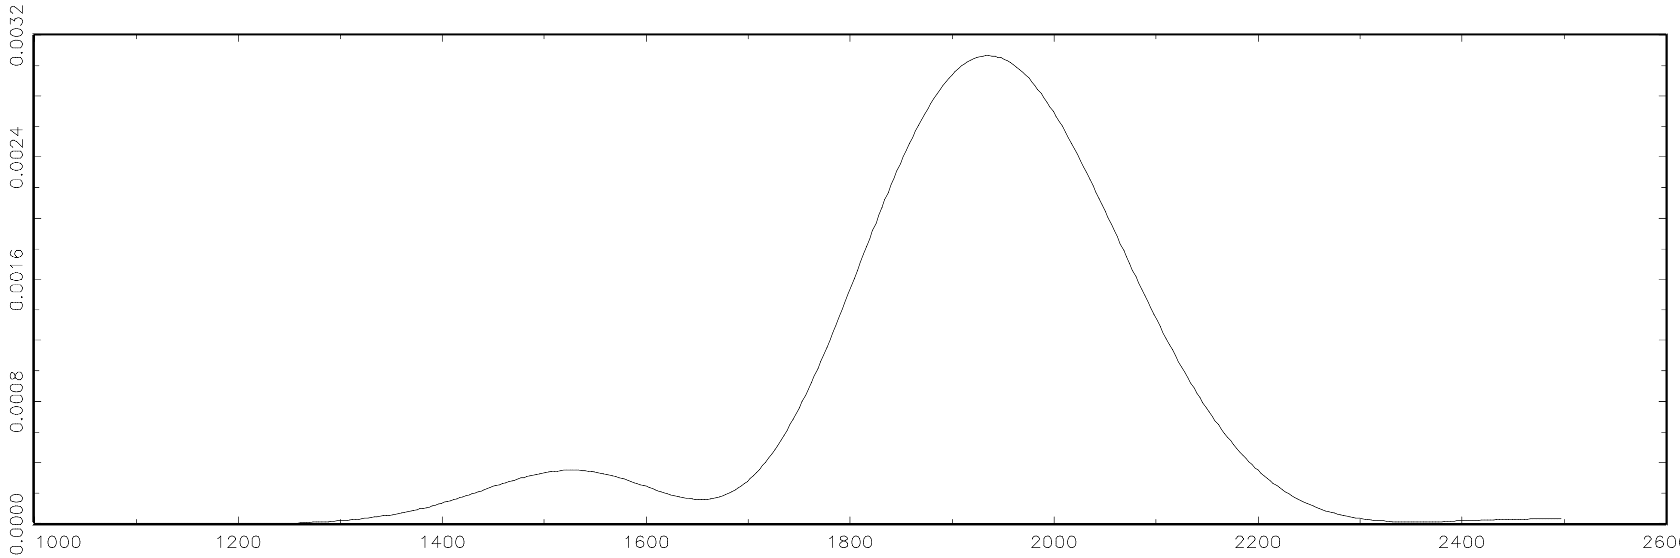

Figure 4b: RND for the date 25/07/1996 and maturity of 36 days

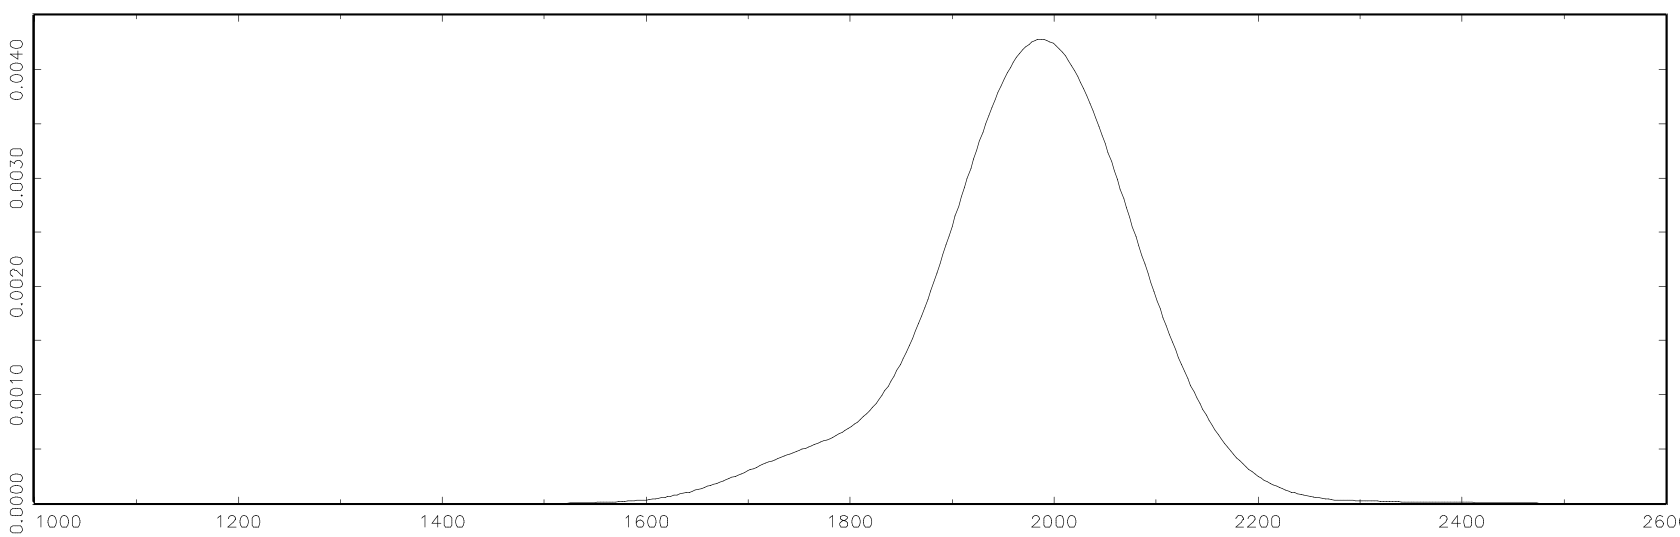

Figure 5a: parameter σ estimated in Hermite polynomials' model over the period 01/01/1995 to 30/06/1997

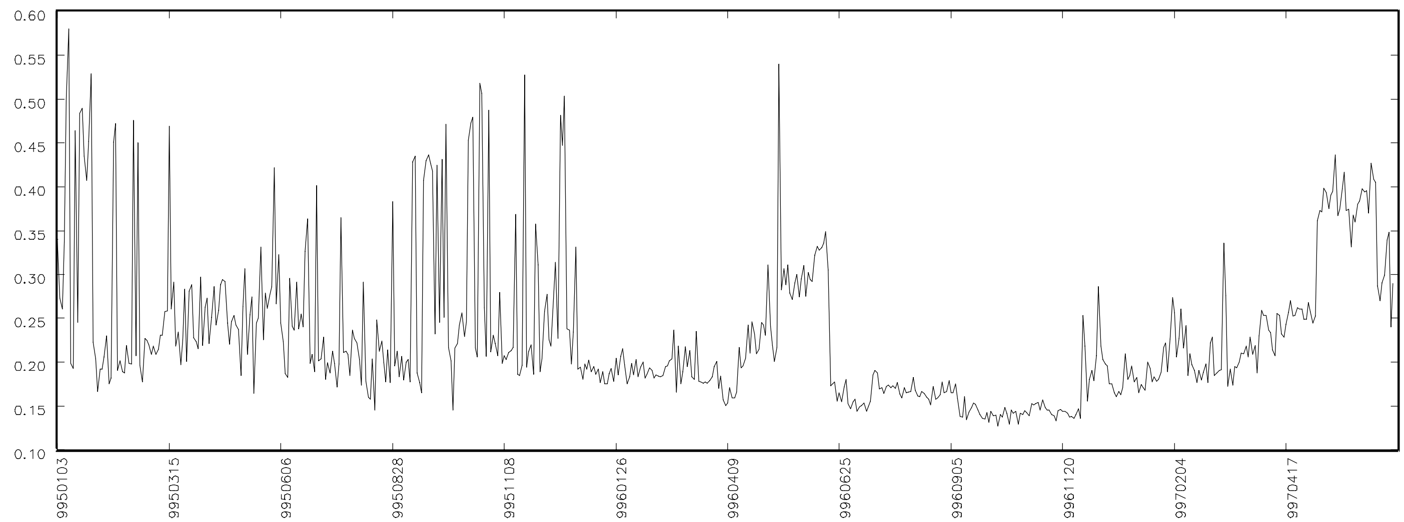

Figure 5b: Black_Scholes volatilities of the CAC 40 index over the period 01/01/1995 to 30/06/1997

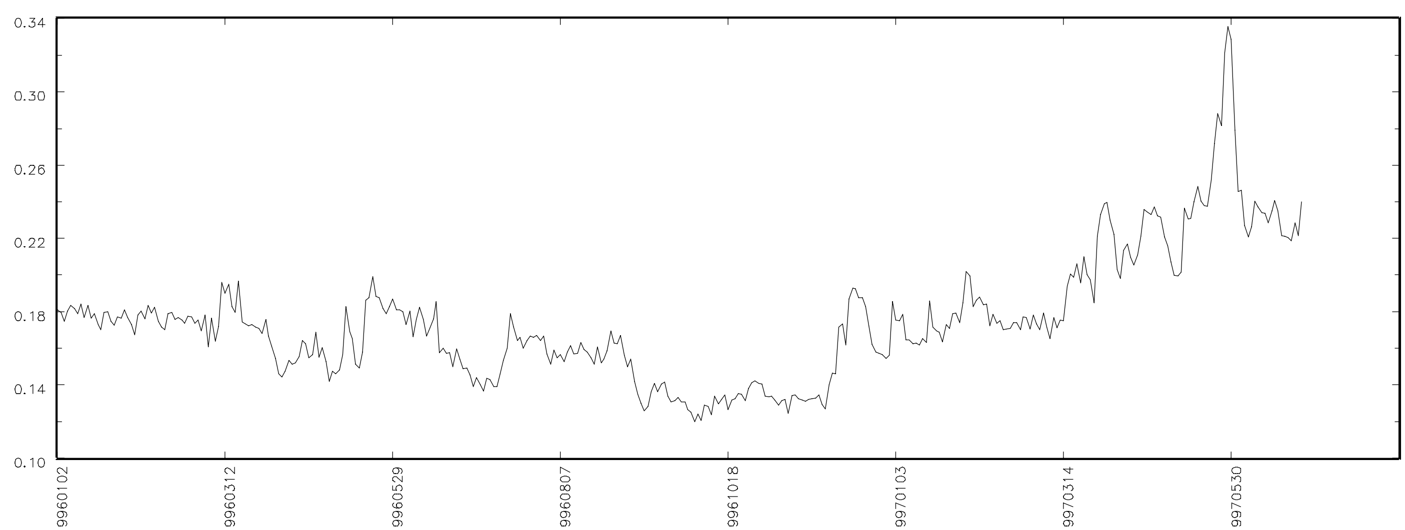

Figure 6a: parameter b3 estimated in Hermite polynomials' model over the period 01/01/1995 to 30/06/1997

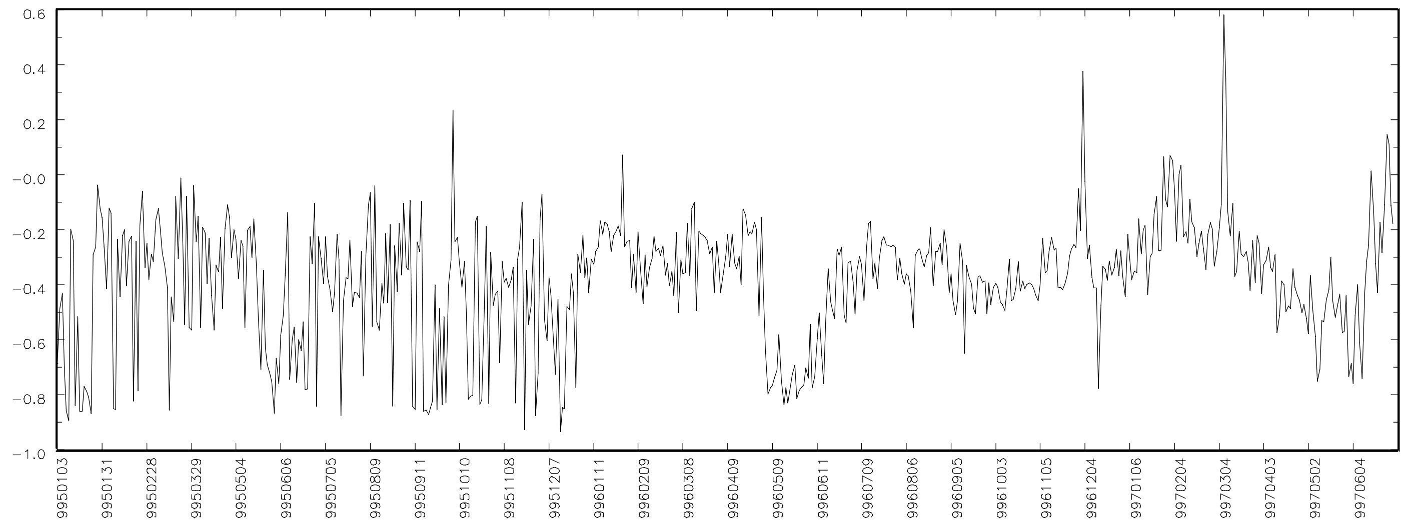

Figure 6b: parameter b4 estimated in Hermite polynomials' model over the period 01/01/1995 to 30/06/1997

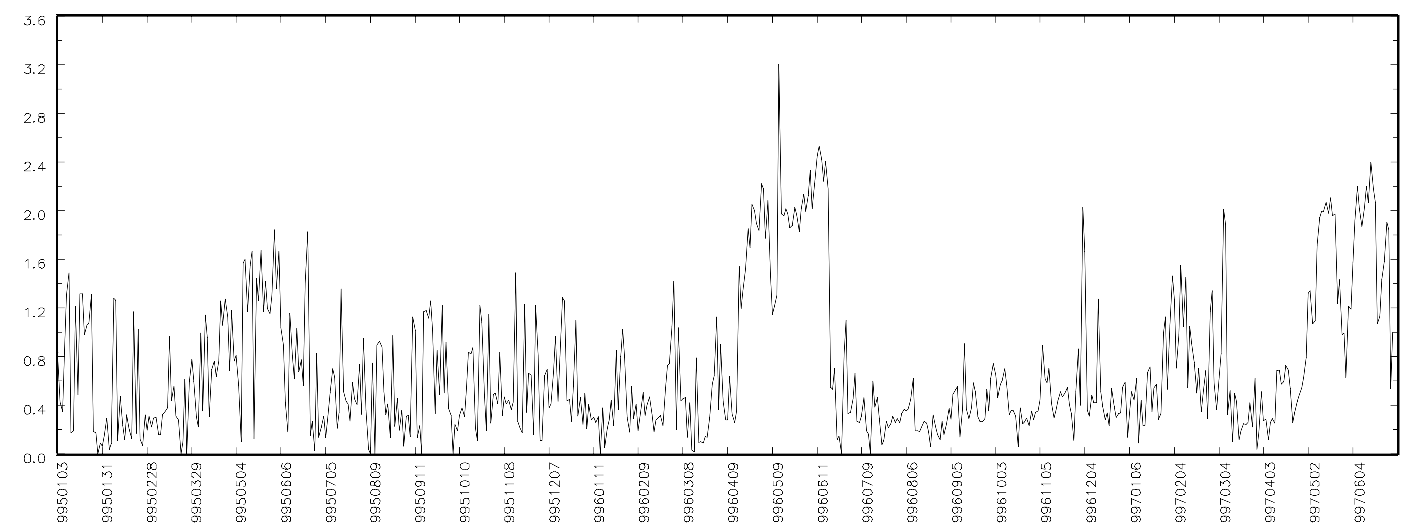

Figure 7: MSE of parameters in Hermite polynomials' model over the period du 01/01/1995 au 30/06/1997

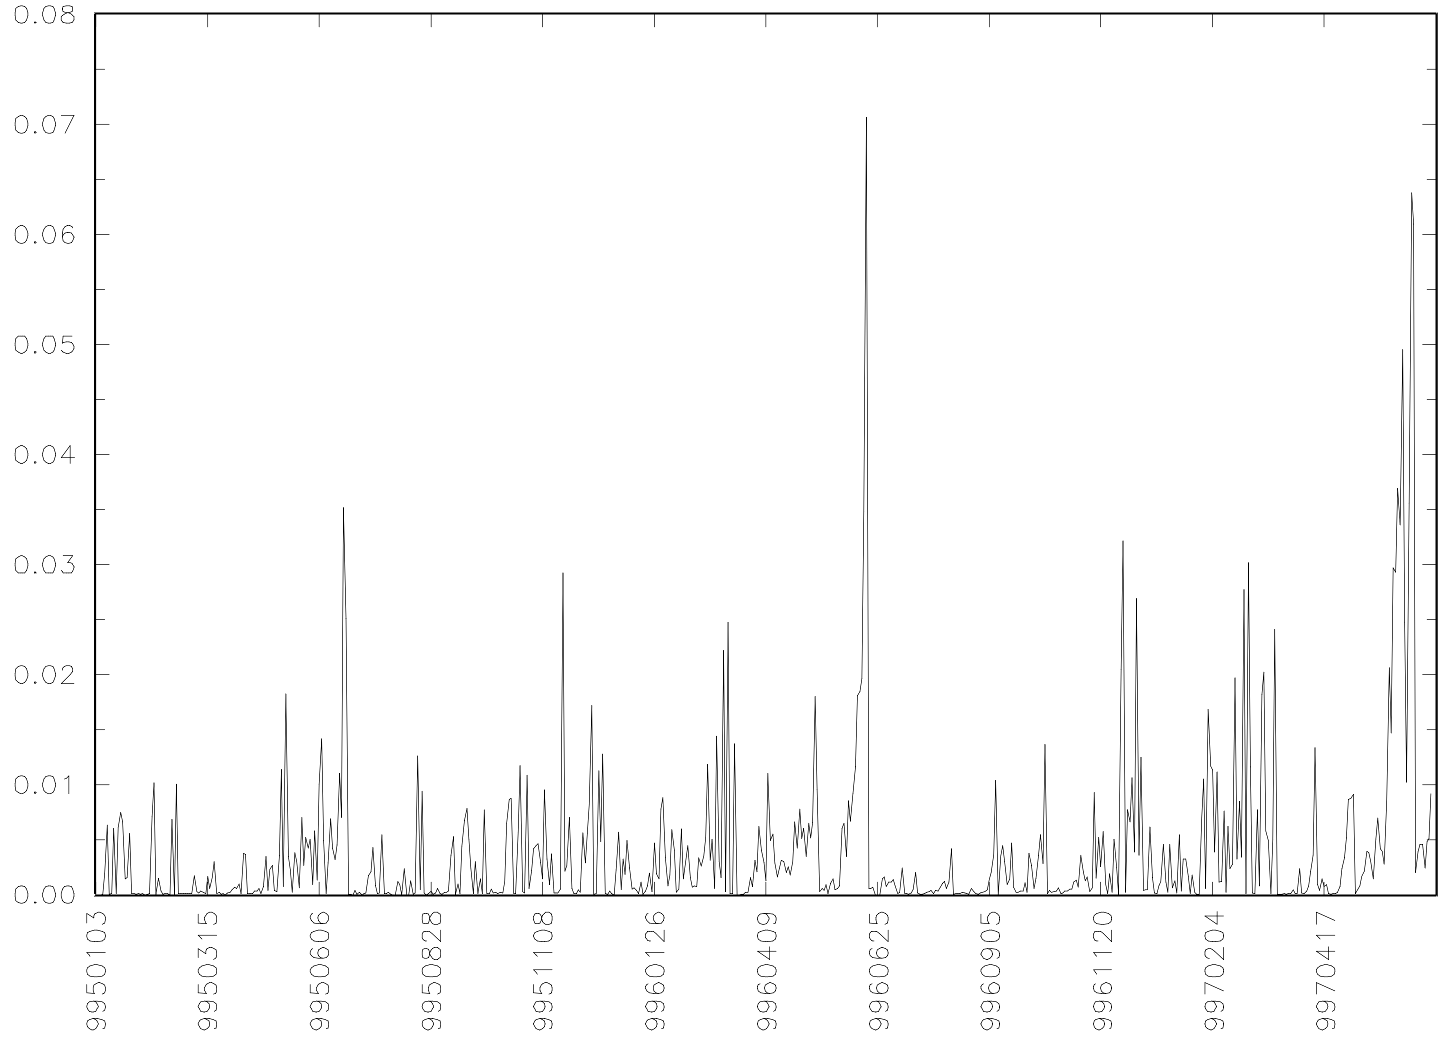

Figure 8a: absolute Risk Aversion function for the date 19950228

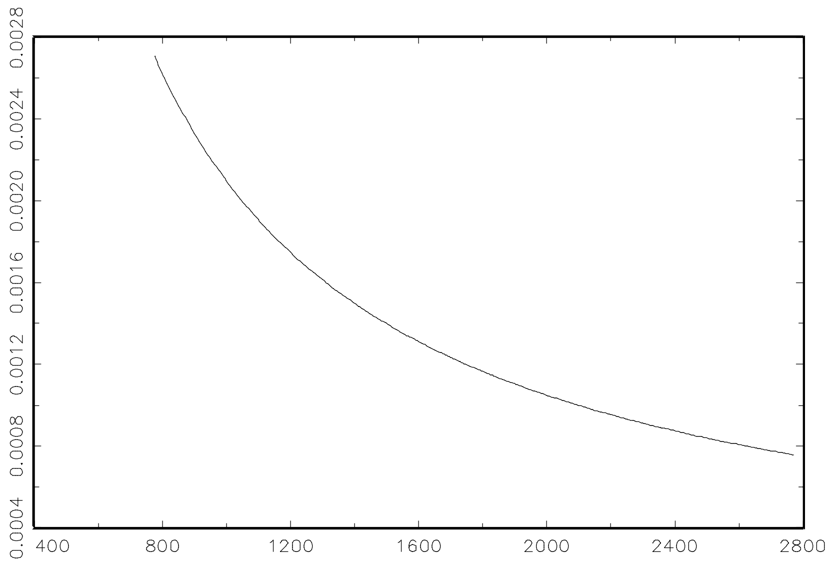

Figure 8b: absolute Risk Aversion function for the date 19950428

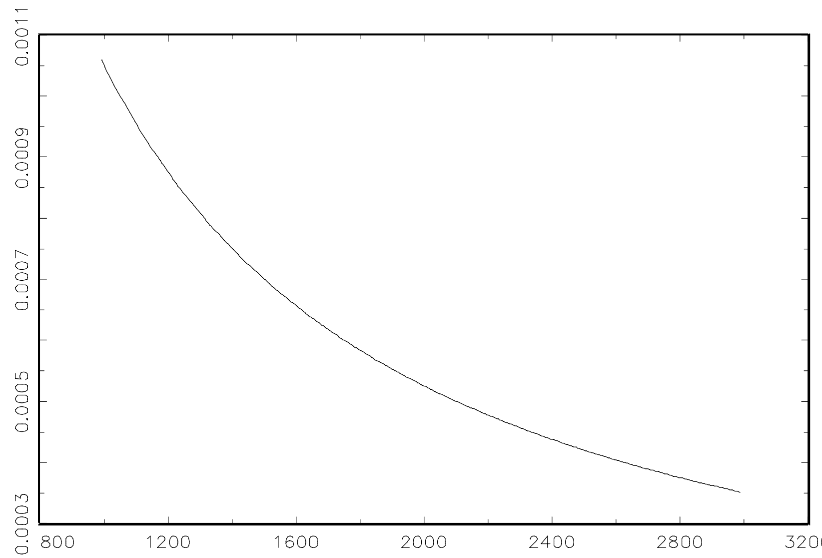

Figure 8c: absolute Risk Aversion function for the date 19960715

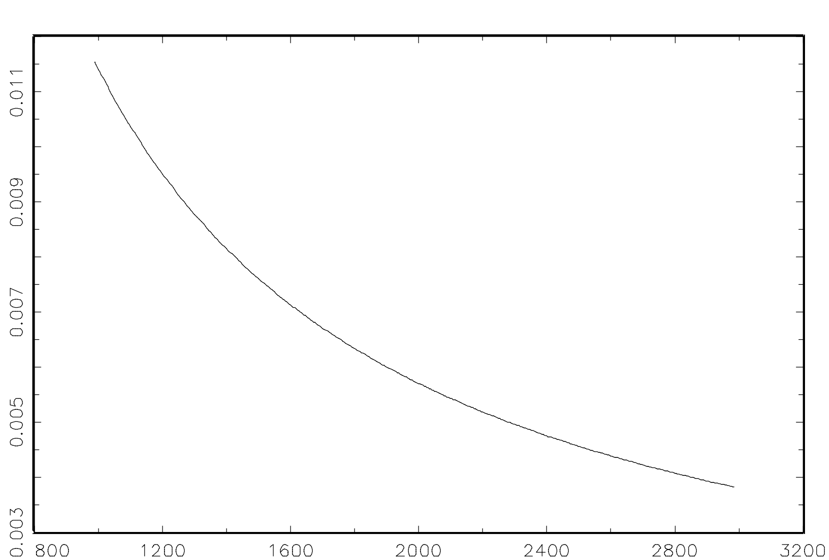

Figure 8d: absolute Risk Aversion function for the date 19961113

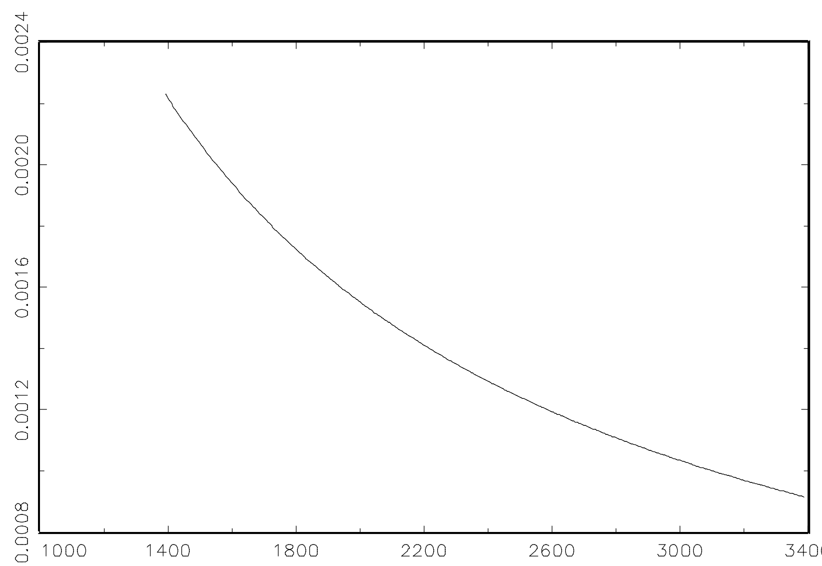

Figure 9: λ estimated with Hermite polynomials over the period du 01/01/1995 au 30/06/1997

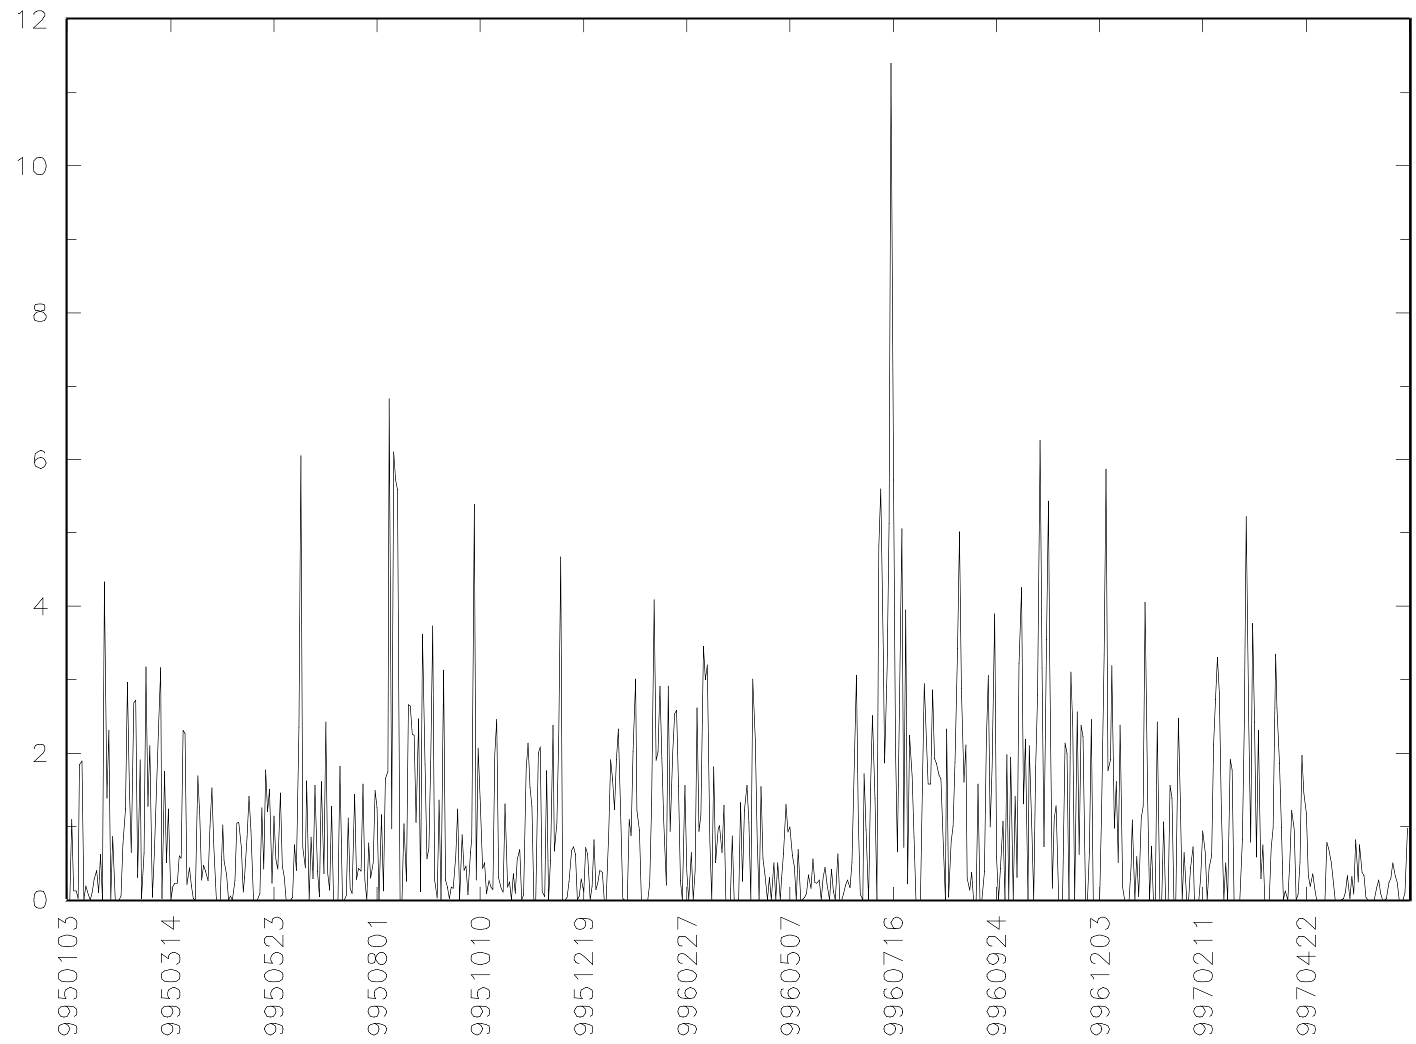

Discussion of the paper by Sophie Coutant, Banque de France:

Implied Risk Aversion in Options Prices

Discussant: Robert Bliss

- Starting point is equation (7) $A(S_T) = \frac{p'(S_T)}{p(S_T)} - \frac{q'(S_T)}{q(S_T)}$.
- If we know two functions, we can estimate third.
- We can estimate RND $q(S_T)$ from options prices.
- If we want SD $p(S_T)$ we must specify risk aversion funtion $A(S_T)$.
  - Some simply assume investors are risk neutral:
  - $A(S_T) = 0$.
- There is considerable evidence that investors are NOT risk neutral.
- If we can estimate SD (from past data), we can learn about risk aversion function $A(S_T)$.
- This paper makes strong assumptions about both $A(S_T), q(S_T)$, and $p(S_T)$
  - $$\begin{aligned} p_t(S_T) &= v_t(S_T)n(S_T) \\ q_t(S_T) &= \lambda_t(S_T)n(S_T) \end{aligned}$$
  - Hermite polynomial representation (for $v(S_T)$ and $\lambda(S_T)$?)
    - Sums of Hermite polynomials are general approximating functions.
    - Paper truncates sum at 4th order
      - Is 4th order precise enough? (No discussion here or in Abken *et al.*)
    - Restrict $b_{0,t} = 1,\ b_{1,t} = 0,\ b_{2,t} = 0$.
      - Because Abken *et al.* do (to match RND and SD mean and variance).
      - What is motivation? Abken *et al.* do not explain.
    - Under these restrictions summed Hermite polynomials are no longer a general approximating function.
  - $$\begin{aligned} x_{(k+1)\Delta\tau} &= x_{k\Delta\tau} + \mu_{k\Delta\tau}\Delta\tau + \sigma_{k\Delta\tau}\sqrt{\Delta\tau} e_{(k+1)\Delta\tau} \\ \mu_{(k+1)\Delta\tau} &= \alpha_0 + \alpha_1 \mu_{k\Delta\tau} + \beta_1 e_{(k+1)\Delta\tau} \end{aligned}$$
    - This is a non-mean reverting process: $x_{k\Delta\tau} \to \pm\infty$.
  - $A(S_T) = \frac{\lambda_{t?}}{S_T}$ follows from CRRA: $U(S) = \frac{S^{1-\lambda}}{1-\lambda}$
    - $\Rightarrow A(S_T) = \frac{p'^{Hermite}(S_T)}{p^{Hermite}(S_T)} - \frac{q'^{Hermite}(S_T)}{q^{Hermite}(S_T)} = \frac{\lambda_t}{S_T}$.

- Given $p^{Hermite}(S_T)$ and $q^{Hermite}(S_T)$, λ_t is estimated w/ least squares.
- Given $p^{Hermite}(S_T)$ and $q^{Hermite}(S_T)$, why impose a functional form on $A(S_T)$?
  - Just compute $A(S_T)$.
  - Or let data suggest appropriate functional form.
    (To test CRRA assumption)

Conclusion

- Paper addresses a difficult but important problem.
  - Critical for using RDNs to assess SDs and market expectations
  - Approach is imaginative.

- Methodology used makes numerous strong, structural assumptions.
  - Restricts possible solution space to particular parsimonious function.
  - If structural assumptions are correct, answers are useful.
  - If structural assumptions are wrong, what do we have?

- Recommendations
  - Provide empirical support for structural assumptions.
  - Or better yet, use methodology to study $A(S_T)$.

Des J McManus

"The information content of interest rate futures options"

Discussants:

James M Mahoney

Roberto Violi

Bank of Canada Banque du Canada

The Information Content of Interest Rate Futures Options*

Des J. Mc Manus

Research and Risk Management
Financial Markets Department
Bank of Canada, Ottawa, Ontario, Canada

Abstract

Option prices are being increasingly employed to extract market expectations and views about monetary policy. In this paper, Eurodollar options are monitored to examine the evolution of market sentiment over the possible future values of Eurodollar rates. Risk-neutral probability functions are employed to synopsize the information contained in the prices of Eurodollar futures options. Several common methods of estimating risk-neutral probability density functions are examined. A method based on a mixture of lognormals density is found to rank first and a method based on a Hermite polynomial approximation is found to rank second. Several standard summary statistics are also examined, namely volatility, skewness and kurtosis. The volatility measure is fairly robust across methods, while the skewness and kurtosis measure are model-sensitive. As a concrete example, the days surrounding the September 1998 Federal Market Open Committee are examined.

*The author would like to thank David Watt, Paul Gilbert, Toni Gravelle, Peter Thurlow, and Mark Zelmer for their helpful comments. Special thanks goes to Michael Rockinger for data that enabled computer code testing. The views expressed in this paper are those of the author and should not be attributed to the Bank of Canada.

Introduction and overview

Timely information is crucial to central banks for formulating and implementing monetary policy. There are of course many sources of information. Macroeconomic data releases, regional industry visits and surveys, and financial market data are all examples of sources that central Banks use. This paper focusses on the latter source—in particular, the derivative markets sector of financial markets, which has gained prominence as a source of information.

Derivative markets have the desirable property of being forward-looking in nature and thus are a useful source of information for gauging market sentiment about future values of financial assets. Indeed, several studies have used option prices to extract market expectations and views about monetary policy [Bahra (1996), Söderlind and Svensson (1997), Söderlind (1997), Butler and Davies (1998), and Levin, Mc Manus, and Watt (1998)]. In particular, Bahra noted that option prices may prove to be useful to monetary authorities as valuable sources to (i) assess monetary conditions, (ii) assess monetary credibility, (iii) assess the timing and effectiveness of monetary operations, and (iv) identify market anomalies.

In this paper, eurodollar futures options are monitored to examine the evolution of market sentiment over the possible future values of eurodollar rates. The key tool used to synopsize the information contained in the prices of eurodollar futures options is the risk-neutral probability density function (PDF). Risk-neutral PDFs provide the probabilities attached by a risk-neutral agent to particular outcomes for future values of eurodollar rates. In addition, changes in the shape and location of the risk-neutral PDF can point to changes in the tone of the market.

Many methods exist to extract risk-neutral PDFs from option prices. This paper compares several common methods of estimating risk-neutral PDFs with the aim of determining which method most accurately prices observed market options. Encouragingly, the mixture of lognormals method ranked first—this method is now used at the Bank for examining the information content of foreign exchange futures options.[1] However, the mixture of lognormal method can occasionally run into problems. When it does, an alternative method called the Hermite polynomial method is more appropriate. The Hermite method ranked second and yielded similar results to the mixture of lognormal method.

Several standard summary statistics can be derived from the risk-neutral PDFs, namely volatility, skewness, and kurtosis. Invariably, these statistics are always quoted in conjunction with the risk-neutral PDF estimates.

A second objective of the paper is to ascertain the robustness and usefulness of these statistics. The volatility measure was found to be fairly robust across the different risk-neutral PDFs. However, the estimates of skewness and kurtosis were found to be model-dependent. The skewness measure for the exchange rate is

1. Foreign exchange futures options are examined to monitor the evolution of the markets' sentiment over future Canadian dollar exchange rates.

now quoted weekly at the Bank. The results of this paper show that further research needs to be conducted on an appropriate measure of market sentiment asymmetry.

As a concrete example, the days surrounding the September 1998 Federal Open Market Committee (FOMC) meeting are examined using the risk-neutral PDF methodology. Risk-neutral PDFs are used to monitor the response of market sentiment over the future levels of the eurodollar rates to the 29 September FOMC statement. The risk-neutral PDFs indicated an increase in market uncertainty prior to the 29 September meeting date, a lessening of uncertainty on the meeting date, and a renewed increase in uncertainty the day after the meeting. The risk-neutral PDFs clearly suggest a bearish market sentiment for the eurodollar rate, both prior to and after the FOMC meeting. Thus, some market participants expected the Fed easing and also anticipated further rate cuts would follow before mid-December 1998.

This paper is organized as follows: Section 1 reviews exchange-traded interest rate futures and interest rate futures options. Section 2 presents the general theory behind the pricing of interest rate futures options. Section 3 gives an overview of several of the common methods that are used to extract risk-neutral PDFs. (Those readers not interested in the technical details of the various option-pricing models may wish to skip section 3.) Section 4 describes the data. Section 5 compares the risk-neutral PDFs from the various estimation methods. Section 6 presents a study of the September 1998 FOMC meeting, focusing on the response of the risk-neutral PDF to the meeting. Section 7 concludes the paper and discusses possible further work.

The work in the present paper closely follows the work and methodologies of Jondeau and Rockinger (1998), and Coutant, Jondeau, and Rockinger (1998).

1. The instruments

The primary focus of this paper is exchange-traded interest rate futures and interest rate futures options. In the United States and Canada, the main exchanges for interest rate products are the Chicago Merchantile Exchange (CME) and the Montreal Exchange (ME). The CME lists a host of contracts on short-term U.S. and foreign securities. For example, both futures and futures options are listed for 3-month eurodollars, 1-month LIBOR, 13-week Treasury bills, euroyen and eurocanada. On the other hand, the ME lists relatively few interest rate futures, namely, 1-month Canadian bankers' acceptance futures (BAR), 3-month Canadian bankers' acceptance futures (BAX), 5-year Government of Canada bond futures (CGF), and 10-year Government of Canada bond futures (CGB). Futures options are listed for the 3-month Canadian bankers' acceptance futures (OBX) and the 10-year Government of Canada bond futures (OGB). Options are also listed for a small selection of Government of Canada bonds.

According to the CME, the eurodollar futures (ED) are "the most liquid exchange-traded contracts in the world when measured in terms of open interest" (Chicago Mercantile Exchange 1999). For example, a snapshot of the futures market on 15 January 1999 reveals that the March 99 ED contract had a trading volume of 76,109 and an open interest of 465,398. The eurodollar futures options (ZE) on this contract, March 99 ZE, had a combined trading volume of 27,939 and a combined open interest of 748,664. The numbers for the eurocanada futures contract pale in comparison; on 14 January 1999 the March 99 futures contract had zero trading volume and an open interest of only 190.

Statistics from the ME reveal that the BAX contract is the most actively traded contract at that exchange. The average daily volume and open interest for all BAX contracts for 1998 was 27,104 and 171,354, respectively. In comparison, the OBX futures options had an average daily volume and open interest of 840 and 15,505, respectively. The OBX volume and open interest are minuscule compared with the figures for the ZE contracts, especially considering the fact that the OBX data is aggregated across all maturity dates trading while the ZE data refers to a single maturity date. Thus, for the remainder of the paper, only CME futures and futures option data will be used.

ED contracts are listed for the quarterly cycle of March, June, September, and December, and also for the two nearest serial (non-quarterly) months. ED futures contracts are traded using a price index. The futures interest rate is calculated by subtracting the futures price from 100. For example, a ED price of 95.80 corresponds to a futures interest rate of 4.20 per cent. Thus if investors expect short-term interest rates to decline (increase), they would go long (short) the futures contract. ED contracts have a contract size of U.S.$1 million. They also feature a minimum allowable price move or tick size of 0.01, with the single exception of when a futures contract is in its expiration month, in which case the minimum tick size is reduced to 0.005. A tick value of 0.01 corresponds to a value of U.S.$25 (Contract size ¥ Tick Value ¥ Maturity of the underlying futures contract = 1,000,000 ¥ 0.01/100 ¥ 3/12). Futures contracts cease trading at 11:00 am London time on the second London business day prior to the third Wednesday of the contract month.

The ZE contract cycle, maturity date, and minimum tick size are the same as those of the underlying ED contract. The ZE contract size is simply one futures contract. Eurodollar futures options consist of American-style[2] call and put[3] options written on the underlying ED futures contract. A 3-month ED futures call option gives the holder the right but not the obligation to buy a 3-month ED futures contract. Now, investors who expect U.S. short-term interest rates to decline would also be expecting the price of the futures contract to increase. Thus, they might be inclined to purchase a 3-month ED futures call option to speculate on their belief. Hence, an exchange-listed interest rate futures call option is equivalent to a put option on the futures interest rate because of the inverse relationship between prices and interest rates, and the fact that exchange-listed interest rate futures options are quoted in units of price rather than percentage interest rates.

---

2. An American option allows the holder to exercise the option on any date up to and including the maturity date—the maturity date is also referred to as the expiration date or the exercise date. European options only allow exercise on the expiration date. American options are always more expensive than European options with the same characteristics because of the added feature of early exercise. In general, the early exercise feature of American options makes these options more difficult to price than European options.
3. A call option gives the holder the right but not the obligation to buy the underlying asset at a predetermined strike price. A put option gives the holder the right but not the obligation to sell the asset at the strike price.

For notational convenience, exchange-listed call (put) options that are quoted in units of price are converted to put (call) options that have units of interest rate, that is, to percentage interest rates.

2. General theory

The valuation of interest rate futures options is best illustrated by first considering the pricing of European-style options. Let $\tilde{r}(t)$ denote the futures interest rate at time *t*—recall $\tilde{r}(t) = 100 - \tilde{p}(t)$, where $\tilde{p}(t)$ is the listed futures price at time *t*. Let *X* and *T* denote the strike price and the time to maturity of the option, respectively. Note that the strike price of a call option on the futures interest rate is equal to 100 minus the listed strike price of an interest rate futures put option. First, note that on their maturity dates the price of a call and put option will be

$$\begin{aligned} \tilde{C}(T, X) &= \max\{0, \tilde{r}(t) - X\} \equiv (\tilde{r}(t) - X)^+ \\ \tilde{P}(T, X) &= \max\{0, X - \tilde{r}(t)\} \equiv (X - \tilde{r}(t))^+ \end{aligned} . \tag{1}$$

Prior to maturity, European options are priced by taking the expectation of the discounted future cash flows. In this case, the future cash flows are the possible payouts of the options at maturity; see equation (1). The cash flows are discounted using the future values of the instantaneous risk-free rate. Thus, the value of European call and put options prior to maturity are given by the following formulae, respectively:

$$\begin{aligned} C(0, X) &= E_0\left[\exp\left\{-\int_0^T \tilde{r}_i(\tau)\ d\tau\right\}\ \tilde{C}(T, X)\right] \\ P(0, X) &= E_0\left[\exp\left\{-\int_0^T \tilde{r}_i(\tau)\ d\tau\right\}\ \tilde{P}(T, X)\right] \end{aligned} , \tag{2}$$

where E_0 represents the risk-neutral expectation, as opposed to the true or actual expectation, and $\tilde{r}_i(\tau)$ refers to the continuously compounded instantaneous interest rate. To simplify matters, the instantaneous rate is taken to be a fixed risk-free interest rate r_f. Strictly speaking, this assumption is incorrect, however it is common practice among market participants and academics alike.

Thus, the value of the European call and put options can then be expressed as:

$$\begin{aligned} C(0, X) &= \exp\{-r_f T\}\ E_0[(\tilde{r}(T) - X)^+] \\ P(0, X) &= \exp\{-r_f T\}\ E_0[(X - \tilde{r}(T))^+] \end{aligned} . \tag{3}$$

2.1 American-style interest rate futures options

Exchange- traded interest rate futures options are typically American-style options. Thus, the above pricing formulae for European-style options needs to be adjusted to account for the possibility of early exercise. Explicit formulae for American-style options are generally not available. However, Melick and Thomas (1997), Leahy and Thomas (1996), and Söderlind (1997) have shown that the following bounds can be placed on the prices of American-style currency futures options:

$$\begin{aligned}\bar{C}_A(0, X) &= E_0[\max\{0, \tilde{r}(T) - X\}]\\ \underline{C}_A(0, X) &= \max\{E_0[\tilde{r}(T)] - X, \exp(-r_f T)E_0[\max\{0, \tilde{r}(T) - X\}]\}\\ \bar{P}_A(0, X) &= E_0[\max\{0, X - \tilde{r}(T)\}]\\ \underline{P}_A(0, X) &= \max\{X - E_0[\tilde{r}(T)], \exp(-r_f T)E_0[\max\{0, X - \tilde{r}(T)\}]\}\end{aligned} \quad . \tag{4}$$

American-style options can then be priced as a weighted average of the upper and lower bounds, namely:

$$\begin{aligned}C_\theta(0, X) &= \omega_i \bar{C}_A(0, X) + (1 - \omega_i)\underline{C}_A(0, X)\\ P_\theta(0, X) &= \omega_i \bar{P}_A(0, X) + (1 - \omega_i)\underline{P}_A(0, X)\end{aligned} \qquad \text{where } i = 1,2 \text{ and } 0 \le \omega_i \le 1 . \tag{5}$$

Following Melick and Thomas (1997), the weights applied will depend on whether the particular option is in-the-money[4] or out-of-the-money. That is, by convention, $i = 1$ for in-the-money call or put options, and $i = 2$ for out-of-the-money call or put options.

2.2 General methodology

The formulae for the prices of European options, (3), can be written explicitly in terms of the risk-neutral PDF, $q[\tilde{r}(T)]$, as follows:

$$\begin{aligned}C(0, X) &= \exp\{-r_f T\} \int_X^\infty \{\tilde{r}(T) - X\}\ q[\tilde{r}(T)]\ d\tilde{r}(T)\\ P(0, X) &= \exp\{-r_f T\} \int_0^X \{X - \tilde{r}(T)\}\ q[\tilde{r}(T)]\ d\tilde{r}(T)\end{aligned} \quad . \tag{6}$$

The risk-neutral PDF for the interest rate, $q[\tilde{r}(T)]$, provides the probabilities attached by a risk-neutral agent today (that is, time $t = 0$) to particular outcomes for future interest rates[5] that could prevail on the maturity date of the option contract.

Various methodologies have been proposed to obtain the risk-neutral PDF from observed futures option prices.[6] The techniques used in this paper—a full discussion follows later—all allow the risk-neutral

4. A European interest rate call (put) option is in-the-money if the futures interest rate is above (below) the strike interest rate, out-of-the-money if the futures interest rate is below (above) the strike interest rate, and at-the-money if the futures interest rate equals the strike interest rate.
5. In the context of this paper, the future interest rate refers to the 3-month eurodollar rate.

PDF to be expressed in a parametric form. Thus, it is helpful to introduce the following notation: let θ denote the parametric vector for the risk-neutral PDF—of course the makeup of this vector will vary depending on the technique being used. Now, let $C_\theta(0, X)$, and $P_\theta(0, X)$ be the theoretical call and put futures option prices with exercise price X [the theoretical prices are calculated from equation (5) with the aid of equations (4) and (6)]. Also, let $C(X)$ and $P(X)$ be the observed call and put futures option prices with exercise price X. Finally, let the theoretical interest rate futures price derived from the option-pricing model under risk-neutral density, $q[\tilde{r}(T)]$, be given by $F_\theta(0, T)$ (= $E_0[\tilde{r}(T)]$), and let the observed interest rate futures price be given by $F(0, T)$.

The parameters of the risk-neutral PDFs, θ, are estimated by minimizing the squared pricing errors associated with the call futures option prices, the put futures options prices, and the interest rate futures price. The minimization problem is:

$$\min_\theta \left[\sum_{i=1}^{n} [C(X_i) - C_\theta(0, X_i)]^2 + \sum_{j=1}^{m} [P(X_j) - P_\theta(0, X_j)]^2 + [F(0, T) - F_\theta(0, T)]^2 \right] \tag{7}$$

where the number of call and put options are allowed to differ.

3. Overview of some specific techniques

As mentioned earlier, many techniques exist to extract risk-neutral PDFs from option prices. In this section, the theory behind some of the more common techniques is reviewed. In general, the techniques considered in this paper fall, with one exception, into two broad categories: a stochastic process for the evolution of the short-term interest rate is specified, or a parametric form for the risk-neutral PDF over the interest rate on the maturity date of the option is specified. The former category contains Black's model and a jump-diffusion model. The latter category contains methods based on a mixture of lognormal density functions and a Hermite polynomial expansion. The single exception is the method of maximum entropy.

3.1 Black's model

Black's model (1976) is the baseline model for pricing futures options. The model is very similar to the Black–Scholes model (1973). The futures interest rate, $\tilde{r}(t)$, is assumed to follow a lognormal process

$$d\tilde{r}(t) = \sigma \tilde{r}(t) dW(t) \tag{8}$$

6. There are four main methods of extracting risk-neutral PDFs from option prices: (i) specify a generalized stochastic process for the price of the underlying asset, (ii) specify a parametric form for the risk-neutral PDF, (iii) smooth the implied volatility function, and (iv) use non-parametric techniques. For a broad review of these techniques see Levin, Mc Manus, and Watt (1998).

where σ is the volatility of the futures interest rate, and dW is a Wiener process, that is *W(t)* is a geometric Brownian motion process in a risk-neutral world. For such a process, the risk-neutral PDF is a lognormal density:

$$q[\tilde{r}(T)] = \frac{1}{\sqrt{2\pi}\sigma\sqrt{T}\ \ \tilde{r}(T)} \exp\left\{-\frac{1}{2}\left(\frac{\log(F(0,T)/\tilde{r}(T)) - \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}\right)^2\right\} \quad , \tag{9}$$

where $F(0,T)$ is the interest rate futures rate. Furthermore, in Black's model the theoretical prices of European call and put futures options are given by

$$C_\theta(0, X) = \exp\{-r_f T\}\,[F(0,T)N(d_1) - XN(d_2)] \tag{10}$$

$$P_\theta(0, X) = \exp\{-r_f T\}\,[XN(-d_2) - F(0,T)N(-d_1)]\,, \tag{11}$$

where

$$d_1 = \frac{\log\{F(0,T)/X\}}{\sigma\sqrt{T}} + \frac{1}{2}\sigma\sqrt{T} \quad \text{and} \quad d_2 = d_1 - \sigma\sqrt{T} \tag{12}$$

and $N(x)$ represents the standardized cumulative normal probability distribution function evaluated at x.

At this point, it is worthwhile giving an example of how interest rate futures options are priced using Black's model. Consider the March 1999 ED futures and futures options listed on the CME on January 29, 1999. The March 1999 three-month ED futures contract had a listed settlement price of 95.04. The ZE call contract with strike price 95.00 had a settlement price of 0.060 and the ZE put contract with the same strike price had a settlement price of 0.020. The other inputs required for Black's model are the time-to-maturity of the contracts, the risk-free rate, and the instantaneous volatility. There are 45 days until the expiration of the contracts on March 15. Thus, the time to maturity is T = 0.125 (= 45/360). The risk-free rate is 4.97 per cent, which was calculated as weighted average of 30-day and 60-day eurodollar spot rates. The volatility is 6.02 per cent. First, convert the futures price and the strike price to interest rates. Thus $F(0,T)$ = 4.96 per cent (=100 – 95.04) and X = 5.00 per cent (= 100 – 95.00). Recall that a price call is equivalent to an interest rate put. Hence, the listed call can be priced by using equation (11) to yield a theoretical price of 0.065. The listed put can be priced using equation (10) to yield a theoretical price of 0.025. The theoretical prices are fairly close to the listed prices. Note that the discrepancies in the theoretical and listed price increase as the strike price moves away from the futures price. Table 1 compares the listed and theoretical option prices for a few different strike prices.

Table 1: Listed and theoretical prices of eurodollar futures options

Strike	CME call price	CME put price	Theoretical call price	Theoretical put price
94.875	0.170	0.005	0.167	0.003
95.000	0.060	0.020	0.065	0.025
95.125	0.020	0.105	0.012	0.097

The option contracts refer to March 1999 3-month eurodollar futures options. The CME prices are settle prices for these options for 29 January 1999. The settlement price for 3-month eurodollar futures contract on that date is 95.04. The theoretical prices are calculated using Black's model with a risk-free interest rate of 4.97 per cent and a volatility of 6.2 per cent.

3.2 Mixture of lognormals

A popular choice for the risk-neutral PDF is that of a weighted sum of independent lognormal density functions, which is referred to as a mixture of lognormals. Levin, Mc Manus, and Watt (1998) used this technique to extract the Canada–U.S. exchange rate from Canadian dollar futures options listed on the CME. The mixture of lognormal distributions is a flexible way to deal with departures from the assumptions underlying Black's model without having to specify a stochastic process for the evolution of the futures rate. As well, the mixture of lognormals has the advantage of retaining Black's model as a special subcase. The number of lognormals is usually dictated by the data constraints. Two lognormals are chosen for the present study.

The risk-neutral PDF with a weighted mixture of two lognormal distributions is given by

$$q[\tilde{r}(T)] = \phi_1 q_1[r(T)] + (1-\phi_1)q_2[\tilde{r}(T)], \tag{13}$$

where $0 < \phi_1 \le 1$ and

$$q_i[\tilde{r}(T)] = \frac{1}{\sqrt{2\pi}\sigma_i\ \tilde{r}(T)}\exp\left\{-\frac{1}{2}\left(\frac{\log(\tilde{r}(T))-\mu_i}{\sigma_i}\right)^2\right\}, \text{ for } i = 1,2.$$

Black's model is given by the special case $\phi_1 = 1$, $\mu_1 = \log F(0, T) - \frac{1}{2}\sigma^2 T$ and $\sigma_1 = \sigma\sqrt{T}$.

The theoretical European call and put prices for the mixture of lognormals are

$$
\begin{aligned}
C_\theta(0, X) &= \phi_1\left[\exp\left(\mu_1 + \frac{1}{2}\sigma_1^2\right)N(d_1) - XN(d_2)\right] \\
&\quad +(1-\phi_1)\left[\exp\left(\mu_2 + \frac{1}{2}\sigma_2^2\right)N(d_3) - XN(d_4)\right] \\
P_\theta(0, X) &= \phi_1\left[-\exp\left(\mu_1 + \frac{1}{2}\sigma_1^2\right)N(-d_1) + XN(-d_2)\right] \\
&\quad +(1-\phi_1)\left[-\exp\left(\mu_2 + \frac{1}{2}\sigma_2^2\right)N(-d_3) + XN(-d_4)\right]
\end{aligned} \tag{14}
$$

where

$$
\begin{aligned}
d_1 &= \frac{1}{\sigma_1}[\mu_1 + \sigma_1^2 - \log(X)]\ ,\ d_2 = d_1 - \sigma_1 \\
\mathrm{d}_3 &= \frac{1}{\sigma_2}[\mu_2 + \sigma_2^2 - \log(X)]\ ,\ d_4 = d_3 - \sigma_2
\end{aligned} . \tag{15}
$$

The theoretical futures price is given by

$$
F_\theta(0, T) = \phi_1 \exp\left(\mu_1 + \frac{1}{2}\sigma_1^2\right) + (1-\phi_1)\exp\left(\mu_2 + \frac{1}{2}\sigma_2^2\right) . \tag{16}
$$

3.3 Jump diffusion

Black's model can be extended to account for asymmetries by adding a jump-diffusion process to Black's basic model. Thus, $\tilde{r}(T)$ is assumed to follow a lognormal jump-diffusion process. The evolution is characterized by two components, a lognormal process and a Poisson jump process,

$$
d\tilde{r}(t) = (\mu - \lambda E[k])\,\tilde{r}(t)\,dt + \sigma_\omega\,\tilde{r}(t)\,dW(t) + k\tilde{r}(t)dq_{0,\,t}, \tag{17}
$$

where $dq_{0,\,t}$ is a Poisson counter on the time interval (0,*t*), λ is the average rate of occurrence of the jumps, and *k* is the jump size. In other words, the probability that one jump occurs within the time interval *dt* is $\mathrm{Prob}[dq_{0,\,dt} = 1] = \lambda\ dt$ and the probability that no jumps occur is $\mathrm{Prob}[dq_{0,\,dt} = 0] = 1 - \lambda\ dt$. For simplicity, *k* is assumed to be constant. In general *k* is stochastic.

Bates (1991) showed that a European call could be priced as

$$
C(0, X) = \exp\{-r_f T\} \sum_{n=0}^{\infty} \mathrm{Prob}\begin{bmatrix} n \text{ jumps} \\ \text{occur} \end{bmatrix} E_0\left[(\tilde{r}(T) - X)^+ \middle|\ \begin{matrix} n \text{ jumps} \\ \text{occur} \end{matrix}\right], \tag{18}
$$

where

$$
\mathrm{Prob}\begin{bmatrix} n \text{ jumps} \\ \text{occur} \end{bmatrix} = \frac{(\lambda T)^n}{n!}e^{-\lambda T} .
$$

A similar formula exists for European puts. For simplicity, assume that at most one jump can occur over the lifetime of the option [see Malz (1996, 1997)]. Ball and Torous (1983, 1985) call this the Bernoulli version of the model. The price of a European call then becomes

$$
\begin{aligned}
C_\theta(0, X) &= \exp\{-r_f T\} \ \text{Prob}\begin{bmatrix}\text{no jumps} \\ \text{occur}\end{bmatrix} \ E_0\left[(\tilde{r}(T) - X)^+ \middle| \begin{array}{c}\text{no jumps} \\ \text{occur}\end{array}\right] \\
&\quad + \exp\{-r_f T\} \ \text{Prob}\begin{bmatrix}\text{1 jump} \\ \text{occurs}\end{bmatrix} \ E_0\left[(\tilde{r}(T) - X)^+ \middle| \begin{array}{c}\text{1 jump} \\ \text{occurs}\end{array}\right] \\
&= (1 - \lambda T) \ \exp\{-r_f T\}\left[\frac{F(0, T)}{1 + \lambda k T} N(d_1) - X N(d_2)\right] \\
&\quad + \lambda T \ \exp\{-r_f T\}\left[\frac{F(0, T)}{1 + \lambda k T}(1 + k) N(d_3) - X N(d_4)\right]
\end{aligned}
\tag{19}
$$

where

$$
\begin{aligned}
d_1 &= \frac{1}{\sigma_\omega \sqrt{T}}\left[\log\left(\frac{F(0, T)}{1 + \lambda k T}\right) - \log(X) + \frac{1}{2}\sigma_\omega^2 T\right], \quad d_2 = d_1 - \sigma_\omega \sqrt{T} \\
d_3 &= \frac{1}{\sigma_\omega \sqrt{T}}\left[\log\left(\frac{F(0, T)}{1 + \lambda k T}(1 + k)\right) - \log(X) + \frac{1}{2}\sigma_\omega^2 T\right], \quad d_4 = d_3 - \sigma_\omega \sqrt{T}
\end{aligned}
\tag{20}
$$

The price of a European call then becomes

$$
\begin{aligned}
P_\theta(0, X) = \ & (1 - \lambda T) \ \exp\{-r_f T\}\left[-\frac{F(0, T)}{1 + \lambda k T} N(-d_1) + X N(-d_2)\right] \\
& + \lambda T \ \exp\{-r_f T\}\left[-\frac{F(0, T)}{1 + \lambda k T}(1 + k) N(-d_3) + X N(-d_4)\right]
\end{aligned}
\tag{21}
$$

Furthermore, the theoretical futures price is $F_\theta(0, T) = F(0, T)$. Note that the future interest rates conditional on no jump occurring and one jump occurring are

$$
\begin{aligned}
E_0\left[\tilde{r}(T) \middle| \begin{array}{c}\text{no jumps} \\ \text{occur}\end{array}\right] &= \frac{F(0, T)}{1 + \lambda k T} \\
E_0\left[\tilde{r}(T) \middle| \begin{array}{c}\text{1 jump} \\ \text{occurs}\end{array}\right] &= \frac{F(0, T)}{1 + \lambda k T}(1 + k)
\end{aligned}
\tag{22}
$$

Thus, the option-pricing formulae consist of a weighted sum of Black's option-pricing formulae where the weights are given by the probability of no jumps occurring and one jump occurring over the lifetime of the option. The option-pricing formulae are very similar to the formula for the mixture of lognormals. Indeed, the jump diffusion is a subcase of the mixture of lognormals PDF. The jump-diffusion PDF is given by equation (13) with $\phi_1 = 1 - \lambda T$, $\mu_1 = \log F(0, T) - \log(1 + \lambda\kappa T) - \frac{1}{2}\sigma_\omega^2 T$, $\mu_2 = \mu_1 + \log(1 + \kappa)$, and $\sigma_1 = \sigma_\omega \sqrt{T} = \sigma_2$.

3.4 Hermite polynomial approximation

Asymmetries in the option data can also be modelled by adding perturbations to Black's baseline model. The Hermite polynomial approximation is a scheme to add perturbations such that successive perturbations are orthogonal. A Hermite polynomial expansion around the baseline lognormal solution is analogous to performing a Fourier expansion. Each additional term in the Hermite polynomial expansion is related to higher moments of the distribution. The general idea is that the Hermite polynomials act as a basis for the set of risk-neutral PDFs. In other words, the risk-neutral PDF can be approximated by a linear summation of Hermite polynomials—the more polynomials the better the approximation; in theory, an infinite series of polynomials gives an almost perfect fit. The technique was developed by Madan and Milne (1994) and later employed to price eurodollar futures options by Abken, Madan, and Ramamurtie (1996).

As a starting point, consider the following lognormal diffusion process:

$$d\tilde{r}(t) = \mu \tilde{r}(t)\,dt + \sigma \tilde{r}(t)\,dW(t)\,, \tag{23}$$

which can be solved to yield

$$\tilde{r}(t) = F(0,T)\exp\left[\left(\mu - \frac{1}{2}\sigma^2\right)T + \sigma\sqrt{T}z\right], \tag{24}$$

where z is distributed as standard normal, that is $z \sim \mathrm{N}(0,1)$. The Hermite polynomial adjustments are constructed with respect to the normalized variable

$$z = \frac{\log[\tilde{r}(T)/\mathrm{F}(0,T)] - \left(\mu - \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}. \tag{25}$$

The risk-neutral PDF for z is denoted by $Q(z)$ and can be written as

$$Q(z) = \lambda(z)n(z)\,, \tag{26}$$

where $n(z)$ is the reference PDF and $\lambda(z)$ captures departures from the reference PDF. The reference PDF is taken as the standardized unit normal PDF, $n(z) = \exp[-z^2/2]/\sqrt{2\pi}$. The departures from normality are captured by an infinite summation of Hermite polynomials, that is:

$$\lambda(z) = \sum_{k=0}^{\infty} b_k\phi_k(z)\,, \tag{27}$$

where b_k are constants and

$$\phi_k(z) = \frac{(-1)^k}{\sqrt{k!}}\frac{1}{n(z)}\frac{d^k n(z)}{dz^k} = \frac{-1}{\sqrt{k}}\frac{d\phi_{k-1}(z)}{dz} + \frac{1}{\sqrt{k}}z\,\phi_{k-1}(z) \tag{28}$$

are an orthogonal system of standardized Hermite polynomials.[7]

The price of any contingent claim payoff $g(z)$ is given by

$$V[g(z)] = \exp\{-r_f T\} \; E_0[g(z)] = \exp\{-r_f T\}\int g(z) \sum_{k=0}^{\infty} b_k \phi_k(z) n(z) dz \qquad (29)$$
$$= \exp\{-r_f T\} \sum_{k=0}^{\infty} g_k \, b_k$$

where $g_k = \int g(z)\phi_k(z)n(z)\,dz$. Now, European call and put options have the contingent claim payoffs, respectively

$$g(z;\text{call}) = \left(F(0,T)\exp\left[\left(\mu - \frac{1}{2}\sigma^2\right)T + \sigma\sqrt{T}z\right] - X\right)^+$$
$$g(z;\text{put}) = \left(X - F(0,T)\exp\left[\left(\mu - \frac{1}{2}\sigma^2\right)T + \sigma\sqrt{T}z\right]\right)^+ \qquad (30)$$

Thus, European call and put prices can be written as

$$C_\theta(0,X) = \exp\{-r_f T\} \sum_{k=0}^{\infty} \alpha_k \, b_k$$
$$P_\theta(0,X) = \exp\{-r_f T\} \sum_{k=0}^{\infty} \beta_k \, b_k \qquad (31)$$

where $\alpha_k = \int g(z;\text{call})\phi_k(z)n(z)\,dz$ and $\beta_k = \int g(z;\text{put})\phi_k(z)n(z)\,dz$. Madan and Milne (1994) show that

$$\alpha_k = \frac{1}{\sqrt{k!}} \left.\frac{\partial^k \Phi(u)}{\partial u^k}\right|_{u=0}, \qquad (32)$$

where the generating function $\Phi(u)$ is given by

$$\Phi(u) = F(0,T)\exp\{\mu T + \sigma\sqrt{T}u\}\, N[d_1(u)] - XN[d_2(u)]$$
$$d_1(u) = \frac{\log\{F(0,T)/X\}}{\sigma\sqrt{T}} + \frac{1}{2}\sigma\sqrt{T} + u \qquad (33)$$
$$d_2(u) = d_1(u) - \sigma\sqrt{T}$$

7. The first four standardized Hermite polynomials are $\phi_0(z) = 1$, $\phi_1(z) = z$, $\phi_2(z) = (z^2 - 1)/\sqrt{2}$, $\phi_3(z) = (z^3 - 3z)/\sqrt{6}$ and $\phi_4(z) = (z^4 - 6z^2 + 3)/\sqrt{24}$. Higher-order Hermite polynomials can be easily calculated using the recurrence relationship $\phi_k(z) = \frac{z}{\sqrt{k}}\phi_{k-1}(z) - \sqrt{\frac{k-1}{k}}\phi_{k-2}(z)$. The polynomials are orthogonal because $\int_{-\infty}^{\infty}\phi_k(z)\phi_j(z)n(z)\,dz$ equals one if $k = j$ and zero otherwise.

and that

$$\beta_k = \begin{cases} \alpha_0 + X - F(0,T)\exp\{\mu T\} & \text{if } k = 0 \\ \alpha_k - \dfrac{\sigma\sqrt{T}}{\sqrt{k!}}\, F(0,T)\exp\{\mu T\} & \text{if } k > 0 \end{cases} \quad . \tag{34}$$

For empirical work, the Hermite polynomial expansion must be truncated at a finite order in *z*. Two approximations are considered in this paper, a fourth-order and a sixth-order approximation. First consider the sixth-order approximation. The risk-neutral PDF for the sixth-order Hermite approximation is given by:

$$\begin{aligned} Q(z) = \frac{1}{\sqrt{2\pi}}\exp\left[-\frac{z^2}{2}\right] \Bigg[& \left(b_0 - \frac{b_2}{\sqrt{2}} + \frac{3b_4}{\sqrt{24}} - \frac{15b_6}{\sqrt{720}}\right) + \left(b_1 - \frac{3b_3}{\sqrt{6}} + \frac{15b_5}{\sqrt{120}}\right)z + \left(\frac{b_2}{\sqrt{2}} - \frac{6b_4}{\sqrt{24}} + \frac{45b_6}{\sqrt{720}}\right)z^2 \\ & + \left(\frac{b_3}{\sqrt{6}} - \frac{10b_5}{\sqrt{120}}\right)z^3 + \left(\frac{b_4}{\sqrt{24}} - \frac{15b_6}{\sqrt{720}}\right)z^4 + \frac{b_5}{\sqrt{120}}z^5 + \frac{b_6}{\sqrt{720}}z^6 \Bigg] \end{aligned} \quad . \tag{35}$$

Under the reference measure, *z* is normally distributed with a mean of 0 and a variance of 1. Under the measure *Q(z)*, *z* has mean $\mathrm{E}_Q[z] = b_1$ and variance $\mathrm{E}_Q[(z - \mathrm{E}_Q[z])^2] = b_0 + \sqrt{2}b_2 - b_1^2$. Furthermore, $\int Q(z)dz = b_0$. Thus, the restriction $b_0 = 1$ must be imposed to insure that the PDF Q integrates to unity. The following restrictions on b_1 and b_2, $b_1 = 0$ and $b_2 = 1$ are imposed to insure that *z* to have mean zero and unit variance with respect to the probability density *Q(z)*. Hence, under the above restrictions the risk-neutral PDF for *z* is

$$\begin{aligned} Q(z) = \frac{1}{\sqrt{2\pi}}\exp\left[-\frac{z^2}{2}\right] \Bigg[& \left(1 + \frac{3b_4}{\sqrt{24}} - \frac{15b_6}{\sqrt{720}}\right) + \left(-\frac{3b_3}{\sqrt{6}} + \frac{15b_5}{\sqrt{120}}\right)z + \left(-\frac{6b_4}{\sqrt{24}} + \frac{45b_6}{\sqrt{720}}\right)z^2 \\ & + \left(\frac{b_3}{\sqrt{6}} - \frac{10b_5}{\sqrt{120}}\right)z^3 + \left(\frac{b_4}{\sqrt{24}} - \frac{15b_6}{\sqrt{720}}\right)z^4 + \frac{b_5}{\sqrt{120}}z^5 + \frac{b_6}{\sqrt{720}}z^6 \Bigg] \end{aligned} \tag{36}$$

and the risk-neutral PDF for $\tilde{r}(T)$ is

$$q[\tilde{r}(T)] = \frac{1}{\sigma\sqrt{T}}\,\frac{1}{\tilde{r}(T)}Q\left[\frac{\log[\tilde{r}(T)/F(0,T)] - \left(\mu - \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}\right]. \tag{37}$$

Finally, the futures price is given by

$$F_\theta(0,T) = F(0,T)\exp\{\mu T\}\left[\sum_{k=0}^{6}\frac{b_k}{\sqrt{k!}}(\sigma\sqrt{T})^k\right] \quad . \tag{38}$$

The fourth-order approximation is simply given by setting $b_5 = 0$ and $b_6 = 0$ in the above equations.

3.5 Method of maximum entropy

The concept of entropy originated in the world of classical thermodynamics as a measure of the state of disorder of a system. Shannon (1948) later introduced the idea to information theory, where entropy was taken as a measure of missing information. Jaynes (1957, 1982) extended the idea to the field of statistical interference using the principle of maximum entropy (PME). Buchen and Kelly (1996) applied the PME to estimating risk-neutral PDFs from option prices. This estimate "will be the least prejudiced estimate, compatible with the given price information in the sense that it will be maximally noncommittal with respect to missing or unknown information."

The PME is a Bayesian method of statistical inference that only uses the price information given and makes no parametric assumptions about the form of the risk-neutral PDF. The method starts with a definition of the entropy of a distribution *q*:

$$S(q) = -\int_0^{\infty} q(x) \log[q(x)]\, dx, \tag{39}$$

which is maximized subject to the constraints

$$\begin{aligned} 1 &= \int_0^{\infty} q(x)\; dx \\ C_i &= \exp\{-r_f T\} \int_0^{\infty} q(x)\; c_i(x)\; dx \qquad \text{where } i = 1 \ldots m, \\ F(0, T) &= \int_0^{\infty} x\, q(x)\; dx \end{aligned} \tag{40}$$

where C_i is the market price of the contingent claim whose payoff at time *T* is given by $c_i(x)$. The risk-neutral PDF is then given by:

$$\begin{aligned} q(x) &= \frac{1}{\mu} \exp\left\{\lambda_0 x + \sum_{i=1}^{m} \lambda_i\, c_i(x)\right\} \\ \mu &= \int_0^{\infty} \exp\left\{\lambda_0 x + \sum_{i=1}^{m} \lambda_i\, c_i(x)\right\}\; dx \end{aligned} \tag{41}$$

Now suppose that the contingent claims consist of European call and put options. Estimating the parameters $\{\lambda_i\}_{i=0}^{m}$ is simplified if only one type of contingent claim is used. Thus, convert the put options to call options using put–call parity. Hence, a futures put option with strike price *X* and observed price *P* is converted to a futures call option with the same strike price and observed price

$C = P - \exp\{-r_f T\} \ [X - F(0, \mathrm{T})]$. For notational convenience, order the resulting set of call options in terms of increasing strike prices, that is, $X_1 < X_2 < \ldots < X_m$.

The futures contract is also considered to be a call option with strike price $X_0 = 0$ and an observed price of $C_0 = \exp\{-r_f T\} \ F(0, T)$. Thus, the constraints for the futures contract and the futures call options can be written as

$$C_i = \exp\{-r_f T\} \int_0^{\infty} q(x)\,(x - X_i)^+ \ dx \qquad \text{where } i = 0 \ldots m \quad . \tag{42}$$

Coutent, Jondeau, and Rockinger (1998) show that the risk-neutral PDF can be written as:

$$q(x) = \begin{cases} \dfrac{1}{\mu} \exp[a_i x + b_i] & \text{for } X_i \le x < X_{i+1} \text{ where } i = 0 \ldots m-1 \\ \dfrac{1}{\mu} \exp[a_m x + b_m] & \text{for } X_m \le x \end{cases} \tag{43}$$

where $a_i = a_{i-1} + \lambda_i$ for $i \ge 1$ with $a_0 = \lambda_0$ and $b_i = b_{i-1} - (a_i - a_{i-1})X_i$ for $i \ge 1$ with $b_0 = 0$. The normalization constant is given by

$$\mu = -\frac{1}{a_m} \exp[a_m X_m + b_m] + \sum_{i=0}^{m-1} \frac{1}{a_i} \{\exp[a_i X_{i+1} + b_i] - \exp[a_i X_i + b_i]\}. \tag{44}$$

Furthermore, the theoretical European call price for strike price X_i is given by C_i where

$$\exp\{r_f T\} \ C_i = \begin{cases} -\left(\dfrac{X_m - X_i}{a_m} - \dfrac{1}{a_m^2}\right) \exp[a_m x + b_m] \\ + \displaystyle\sum_{k=i}^{m-1} \left\{ \left(\dfrac{X_{k+1} - X_i}{a_k} - \dfrac{1}{a_k^2}\right) \exp[a_k X_{k+1} + b_k] - \left(\dfrac{X_k - X_i}{a_k} - \dfrac{1}{a_k^2}\right) \exp[a_k X_k + b_k] \right\} \end{cases} \tag{45}$$

[see Coutent, Jondeau, and Rockinger (1998) for details].

The risk-neutral PDF is characterized by the parameters $\{a_i\}_{i=0}^{m}$, which are estimated by minimizing the squared call pricing errors; see equation (7). The convergence of the estimation process is enhanced by picking initial values for the parameters that are reasonable. Coutent, Jondeau, and Rockinger (1998) suggest choosing parameters so that the risk-neutral PDF, (43), is approximately equal to Black's risk-neutral PDF, (9).[8]

4. Data

The data consist of end-of-the-day settlement prices for American-style eurodollar futures options and eurodollar futures that are traded on the CME, and covers 23 September 1998 to 30 September 1998, inclusively. The dates were chosen to include the Federal Open Market Committee (FOMC) meeting on 29 September 1998. The data also consists of 60- and 90-day spot eurodollar rates. The risk-free rate was constructed by linearly interpolating between these rates and then converting the result to a continuously compounded rate.

The average daily trading volume of the ED futures and the Dec98 ED futures options over the period 23–30 September 1998 was 110,217 and 74,030 contracts, respectively. The average number of Dec98 ED futures options traded was 16. These contracts had a wide range of strike prices, typically from 4.0 per cent to 6.5 per cent (see Table 2 for details).

5. Comparing the models

The various methods outlined in section 3 are compared in this section. First, the models are compared according to their pricing errors—the pricing error is the difference between the theoretical option price and the observed option price. Second, the models are compared using several summary statistics, notably the mean, annualized volatility, skewness, and kurtosis (see the Appendix for further discussion on these quantities). Third, the models are compared by examining the risk-neutral PDFs. This comparison is both graphical and analytic—the analytic analysis consists of comparing the cumulative distribution functions for the various PDFs.

5.1 Metrics for comparison

As mentioned above, the models are compared by examining the pricing errors associated with each model. The pricing error, which is the basic building block, is the difference between the theoretical option price and the observed option price. Thus, the pricing errors for call and put futures options are $C_\theta(0, X_i) - C(X_i)$ and $P_\theta(0, X_j) - P(X_j)$, respectively. These raw pricing errors are illustrated in Figures 1 through 6 [hollow bullets (o) indicate pricing errors for call options and asterisks (*) indicate pricing errors for put options]. Strike prices are marked along the horizontal axis. Black's model clearly gives the highest pricing errors. The method of maximum entropy appears to give the lowest pricing errors. The mixture of lognormals and the Hermite polynomial approximations yield similar pricing errors. Not surprisingly, the mixture of lognormals

---

8. The initial values of the parameters s $\{a_i\}_{i=0}^{m}$ can be estimated as follows. First, generate a data set of interest rates, $\{x\}$, and the corresponding Black's risk-neutral PDF, $\{q_B(x)\}$. Next, estimate the parameters $\{\lambda_i\}_{i=0}^{m}$ for the regression $\log[q_B(x)] = -\log\mu + \lambda_0 x + \sum_{i=1}^{m} \lambda_i (x - X_i)^+ + \varepsilon$. The initial values are then according to the algorithm used for equation (43).

method has smaller pricing errors than the jump-diffusion method, and the sixth-order Hermite polynomial approximation has smaller pricing errors than the fourth-order Hermite approximation. The mixture of lognormals and both the Hermite methods tend to have similar pricing errors.

An alternative to looking at the raw pricing errors is to combine the pricing errors into a single quantity that measures the accuracy of fit. Several measures of accuracy of fit exist in the literature. However, only two measures will be considered in this paper: the mean squared error (MSE) and the mean squared percentage pricing error (MSPE). The choice of measures is motivated by the fact that the loss function (7) is quadratic in the pricing errors. The MSE and the MSPE are calculated as follows:

$$\begin{aligned} \mathrm{MSE} &= \frac{1}{n+m-k}\sum_{i=1}^{n}[C(X_i)-C_\theta(0,X_i)]^2+\frac{1}{n+m-k}\sum_{j=1}^{m}[P(X_j)-P_\theta(0,X_j)]^2 \\ \mathrm{MSPE} &= \frac{1}{n+m-k}\sum_{i=1}^{n}\left[\frac{C(X_i)-C_\theta(0,X_i)}{C(X_i)}\right]^2+\frac{1}{n+m-k}\sum_{j=1}^{m}\left[\frac{P(X_j)-P_\theta(0,X_j)}{P(X_j)}\right]^2 \end{aligned} \tag{46}$$

where n and m are the number of observed call and put prices, and k is the number of independent parameters for the risk-neutral PDF being used, $k = \#\{\theta\}$. The MSE places more weight on larger errors than smaller errors. The MSPE is dimensionless, and thus facilitates comparison across both different methods and different data sets.

Neither the MSE nor the MSPE measures point to a single method that always ranks first. However, averaging the measures over the sample period yields a clear ranking. Both the MSE and the MSPE measures rank the mixture of lognormal method first, the sixth-order Hermite polynomial approximation a close second, and the fourth-order Hermite approximation third (see Table 3). The results may of course be dependent on the ranking scheme employed. However, the other ranking schemes that were considered ranked the mixture of lognormals method first and either one of the Hermite approximations or the method of maximum entropy second. Finally, the results may be dependent on the sample. Only further testing with more diverse data sets will resolve this issue.

5.2 Summary statistics

The models can also be compared according to summary statistics that are calculated with respect to the logarithm of the futures rate. The standard statistics examined are the mean, annualized volatility, skewness, and kurtosis (see the Appendix for a more in-depth explanation). For any given day, the means calculated from each model are practically identical. This result is not too surprising, given that the PDFs are risk-neutral.

The evolution of volatility over the event period follows a fairly consistent pattern. All methods have volatility increasing from 23 September to 24 September, decreasing from 28 September to 29 September, and increasing again from 29 September to 30 September (see Figure 7 and Tables 4a to 9a). The level of volatility from 24 September to 28 September varies across models. On average, the mixture of lognormals yields the highest estimates of volatility and Black's model yields the lowest estimates. Also, the jump model tends to yield higher volatilities than the sixth-order Hermite approximation, the sixth-order Hermite approximation tends to yield higher volatilities than the fourth-order Hermite approximation, and the fourth-order Hermite approximation tends to yield higher volatilities than the method of maximum entropy.

The skewness estimates vary widely across the models (see Figure 7), although all models have negative skewness for each day of the study period. However, no consistent pattern exists for the day-to-day evolution of skewness across methods. For example, from 25 September to 28 September, the mixture of lognormals method measure of skewness becomes more negative while both the Hermite approximations becomes less negative. Likewise, the kurtosis estimates vary dramatically across models. All the models do, however, yield kurtosis numbers greater than 3, indicating fat-tailed (leptokurtotic) distributions.

In summary, the lower moments of the distribution, namely the mean and the volatility, tend to be consistent across models. But the discrepancies between the distributions tend to be exaggerated when higher moments are considered. The skewness and kurtosis measures appear to be very model-dependent, and thus are probably not reliable as indicators of market sentiment.

5.3 The shape of things to come

The risk-neutral PDFs implied by the various models for 23 September to 30 September are illustrated in Figures 1 through 6. The PDFs for the mixture of lognormals method, the jump-diffusion method, and the Hermite polynomial-approximation methods are invariably bimodal. The higher peak is situated almost directly above the futures rate, and in most cases a much lower second peak is situated above a eurodollar rate that is roughly 100 basis points lower than the futures rate (see Figures 1through 6). However, most of the mixture of lognormal risk-neutral PDFs have no lower peak. Instead, they have heavy left tails, indicating negative skewness. The Black risk-neutral PDF is always unimodal. The method of maximum entropy risk-neutral PDF is extremely spiky for all the dates considered. The method of maximum entropy estimates one parameter for every strike price, and thus tends to overfit when there is a large number of strike prices, which is the case in this study. Furthermore, the method of maximum entropy PDFs appear choppy because the first derivative of the PDF is discontinuous at the strike prices.

The cumulative distribution functions (CDFs) are helpful in comparing models. The CDFs are more easily interpreted than the PDFs, since they give the probabilities that the futures rate will be less than a given rate on the maturity date of the futures contract. (Analytic expressions for the CDFs for the various models are in the Appendix.) A selection of the probabilities can be found in Tables 4b through 9b. The CDFs are plotted in Figures 8a through 8d. Black's model consistently underestimates the probabilities in the left tail of the distribution compared with the other models. Not surprisingly, the method of maximum entropy CDF is very different from the other CDFs. The CDFs for the mixture of lognormals method, and the fourth- and sixth-order Hermite polynomial approximation are very close to each other, as can seen both from Tables 4b, 5b, 6b, 7b, 8b, and 9b and from Figure 8d. (For clarity, the aforementioned CDFs are only plotted in Figure 8d).

5.4 General comments on estimation procedures

The method of maximum entropy tends to overfit. This is directly related to the small number of degrees of freedom. Furthermore, the estimation procedure was the slowest to converge. The mixture of lognormals method can also be slow to converge, especially if the true risk-neutral PDF is close to being lognormal. The problem is that there is not a unique set of parameter values that gives a lognormal distribution. Likewise, the jump-diffusion method is plagued by the same problem. The jump-diffusion method works well when there is a reasonable likelihood of a jump occurring. However, as with the mixture of lognormals method, the jump-diffusion method has degenerate parameterizations for lognormal distributions. The Hermite polynomial-approximation methods are quick to converge and do not admit degenerate parameterizations. The Hermite method always converges; the fourth-order approximation converges faster than the sixth-order approximation. The only drawback with the Hermite polynomial-approximation methods is that the estimation of the risk-neutral PDF can occasionally yield negative probability values. These negative probability values can occur because the Hermite method employed is an approximation method that involves truncating an infinite series.

Overall, the mixture of lognormals method and the sixth-order Hermite polynomial-approximation method are probably the best methods to use for extracting risk-neutral PDFs from interest rate option prices. Coutant, Jondeau, and Rockinger (1998) favoured the fourth-order Hermite polynomial-approximation method in their comparison of various methods using French data.

Finally, given the variability of the skewness estimates across methods and the relative consistency of the CDFs, a more accurate measure of skewness could probably be constructed by comparing the tails of the PDFs as opposed to using the third central moment of the distribution. Such a measure exists in the literature: relative intensity [see Campa, Chang, and Reider (1997)] compares the likelihood of large upward movements in the eurodollar rate to large downward movements.

6. The event

As mentioned earlier, the dates of the study were chosen to coincide with the FOMC meeting on 29 September 1998. The FOMC is a 12-member committee, consisting of the seven members of the Board of Governors of the Federal Reserve System, the president of the Federal Reserve Bank of New York, and four of the presidents of the other 11 Reserve Banks; the latter positions rotate yearly.

The FOMC meets eight times a year and has primary responsibility for conducting monetary policy. The committee decides on the desired level of the federal funds rate. Press releases are often posted immediately after meetings, especially if the Fed’s stance on monetary policy has changed. For example, the press release following the 29 September 1998 meeting started: “The Federal Open Market Committee decided today to ease the stance of monetary policy slightly, expecting the federal funds rate to decline 1/4 percentage point to around 5 1/4 per cent.” This reduction was the first of a series of reductions in the Fed fund target rate in 1998. Two later reductions of 25 basis points each occurred on 15 October 1998 and 17 November 1998.

The annualized volatility numbers generally increased over the first half of the period—based on the results of the previous section, the analysis of the present section uses the risk-neutral PDF from the mixture of lognormals method—starting off at 17.82 per cent on 23 September, rising to a high of 19.20 per cent on 28 September, falling to a low of 15.24 per cent on 29 September, and finally starting upwards again on 30 September to 16.92 per cent. Thus, uncertainty, as measured by annualized volatility, initially increased, and peaked the day prior to the FOMC meeting. Uncertainty reduced on the day of the meeting but started to increase again the following day.

The probability of the ED futures rate being below 5.00 per cent on 14 December 1998 rose from 33 per cent to 38 per cent over the period. In addition, the probability of the ED futures rate being below 5.25 per cent rose from 63 per cent to 75 per cent. Furthermore, the skewness numbers remained negative over the entire period, indicating a bearish market tone. Interestingly, skewness became even more negative the day after the Fed easing, indicating that a further Fed easing was expected by some market participants. These findings are consistent with the general market views of the time. Anecdotal evidence suggests that, while market participants anticipated an easing at the 29 September FOMC meeting, some were disappointed by the size of the move (25 basis points) and immediately priced in a further rate reduction by the November meeting.

7. Conclusion

The information content of exchange-traded eurodollar futures options were examined in this paper. Several techniques for extracting risk-neutral PDFs from ED futures option prices were compared. The mixture of lognormals method ranked first, with both the lowest MSE and MSPE. However, this method is occasionally slow to converge due to degeneracies in the parameter space. Typically, the lack of convergence occurs when the risk-neutral PDF appears to be close to a single lognormal distribution. In this case, the alternative sixth-order Hermite polynomial-approximation method yields better results. The Hermite method is quick to converge and gives comparable results to the mixture of lognormals method. However, the method occasionally yields PDFs that have negative probabilities—these negative probabilities are an artifact of the approximation method and are not too worrisome, since they tend to occur near the tails of the distribution.

The higher central moments of the risk-neutral PDFs, namely skewness and kurtosis, are unstable across estimation techniques and thus are probably not overly informative as measures of asymmetry in market sentiment. In contrast, the CDF was found to be stable across the three methods that yielded the lowest MSPEs, namely the mixture of lognormals and the two Hermite polynomial-approximation methods. Thus, measures of skewness based on the CDF are probably more appropriate. One candidate is relative intensity, which compares the likelihood of large upward movements in the ED rate to the likelihood of large downward movements.

Risk-neutral PDFs are useful tools for monitoring market sentiment, as was indicated by the analysis of the 29 September 1998 FOMC meeting. Various methods were used to extract risk-neutral PDFs from ED futures options over the period around the FOMC meeting in order to examine the evolution of market sentiment over the future values of ED rates. Uncertainty grew in the market prior to the meeting and abated on the day of the meeting, only to increase again the following day. Market participants remained bearish on future ED rates both prior to and after the Fed easing, indicating that some of them expected further rate cuts.

Information extracted from option prices can be used to monitor market sentiment. However, the best way to present this information is still up for debate. In particular, work needs to be done on appropriate measures of asymmetry and the predictive power of these measures.

Bibliography

Abken, P.A., D.B. Madan, and S. Ramamurtie. 1996. "Estimation of Risk-Neutral and Statistical Densities By Hermite Polynomial Approximation: With an Application to Eurodollar Futures Options." Federal Reserve Bank of Atlanta Working Paper 96-5.

Bahra B. 1996. "Implied risk-neutral probability density functions from option prices: theory and application." *Bank of England Quarterly Bulletin* (August): 299–311.

Ball C. and W.N. Torous. 1983. "A Simplified Jump Process for Common Stock Returns." *Journal of Financial and Quantitative Analysis* (18): 53–65.

Ball C. and W.N. Torous. 1985. "On Jumps in Common Stock Prices and their Impact on Call Option Pricing." *Journal of Finance* (50): 155–73.

Bates D.S. 1991. "The Crash of '87: Was It Expected? The Evidence from Options Markets." *Journal of Finance* (46): 1009–44.

Black F. and B. Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* (81): 637–57.

Black F. 1976. "The Pricing of Commodity Contract," *Journal of Financial and Quantitative Analysis* (3): 153–67.

Buchen P.W. and M. Kelly. 1996. "The Maximum Entropy Distribution of an Asset Inferred from Option Prices." *Journal of Financial and Quantitative Analysis* (31): 143–59.

Butler C. and H. Davies. 1998. "Assessing Market Views on Monetary Policy: The Use of Implied Risk-Neutral Probability Distributions." In *The role of asset prices in the formulation of monetary policy*, BIS Conference Paper 5. March.

Campa J.M., P.H.K. Chang, and R.L. Reider. 1997. "Implied Exchange Rate Distributions: Evidence from OTC Option Markets." National Bureau of Economic Research Working Paper 6179.

Chicago Mercantile Exchange. 1999. "CME Interest Rate Products." URL: http://www.cme.com/market/interest/howto/products.html.

Coutant S., E. Jondeau, and M. 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 54.

Jaynes E.T. 1957. "Information Theory and Statistical Mechanics." *Physics Reviews* (106): 620–30.

Jaynes E.T. 1982. "On the Rationale of Maximum-Entropy Methods." *Proceedings of the IEEE* (70): 939–52.

Jondeau E. and M. Rockinger. 1997. "Reading the Smile: The Message Conveyed by Methods which Infer Risk Neutral Densities." Centre for Economic Policy Research Discussion Paper 2009.

Leahy M.P. and C.P 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, A., D.J. Mc Manus, and D.G. Watt. 1998. "The Information Content of Canadian Dollar Futures Options." In *Information in Financial Asset Prices: Proceedings of a Conference Held by the Bank of Canada.*

Madan, D.P. and F. Milne. 1994. "Contingent Claims Valued and Hedged by Pricing and Investing in a Basis." *Mathematical Finance* (4): 223–45.

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): 717–48.

Malz, A.M. 1997. "Estimating the Probability Distribution of the Future Exchange Rate From Option Prices." *Journal of Derivatives* (5): 18–36.

Melick W.R. and C.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): 91–115.

Shannon C.E. 1948. "The Mathematical Theory of Communication," *Bell Systems Technical Journal* (27): 379–423.

Söderlind, P. and L.E.O. Svensson. 1997. "New Techniques to Extract Market Expectations from Financial Instruments." *Journal of Monetary Economics* 40 (2): 383–430.

Söderlind, P. 1997. "Extracting Expectations about UK Monetary Policy 1992 from Options Prices." Mimeo.

Table 2: Federal Open Market Committee meeting, September 1998

September 1998	60-day euro-dollar rate	90-day euro-dollar rate	Risk-free rate	Euro-dollar futures rate	Trading volume of euro-dollar futures	Number of different option contracts	Trading volume of euro-dollar futures options
Wednesday 23	5.5313	5.5000	5.3620	5.115	101,026	16	79,626
Thursday 24	5.5000	5.4688	5.3333	5.035	121,205	18	74,215
Friday 25	5.3907	5.3594	5.2306	5.040	124,453	15	96,714
Monday 28	5.3594	5.3282	5.2039	5.060	78,949	15	84,918
Tuesday 29	5.3438	5.3750	5.2217	5.110	142,304	14	50,615
Wednesday 30	5.3594	5.4063	5.2430	5.050	93,363	18	58,089

Note: The day of Federal Open Market Committee meeting is highlighted.

Table 3: Eurodollar futures options:
Pricing errors for call and put futures options, September 1998

Measure	Model	23 Sept.	24 Sept.	25 Sept.	28 Sept.	29 Sept.	30 Sept.	Average	Ranking
Mean squared error ($\times 10^{-5}$)	Black	10.610	8.066	9.771	8.626	7.000	9.230	8.884	6
	MLN[a]	0.777	0.558	0.987	0.863	0.796	1.108	0.848	1
	Jump	1.482	0.604	0.930	0.928	1.983	1.735	1.277	4
	Hermite (4)	0.857	0.591	1.193	1.046	0.945	1.099	0.955	3
	Hermite (6)	0.779	0.608	1.093	0.892	0.667	1.181	0.870	2
	Maximum entropy	2.763	2.974	0.945	0.720	0.609	2.588	1.767	5
Mean squared per-centage pricing error ($\times 10^{-2}$)	Black	7.458	12.754	16.310	18.210	9.302	18.798	13.805	5
	MLN	0.292	0.151	0.201	3.401	0.179	2.735	1.160	1
	Jump	7.071	0.121	0.209	6.914	1.357	5.054	3.454	4
	Hermite (4)	0.249	0.378	2.221	4.808	0.098	2.731	1.748	3
	Hermite (6)	0.896	0.563	0.538	1.887	0.114	3.506	1.251	2
	Maximum entropy	60.805	40.399	4.005	3.146	0.560	38.907	24.637	6
See Section 5.1 for further details.									

a. Mixture of lognormals

Table 4a: Eurodollar futures options, 23 September 1998

23 September	Mean	Volatility	Skewness	Kurtosis
Black	1.629	15.85	0	3
MLN[a]	1.629	17.82	–0.956	5.877
Jump	1.629	17.52	–1.133	5.254
Hermite (4)	1.629	17.66	–0.866	5.438
Hermite (6)	1.629	17.26	–0.719	3.774
Maximum entropy	1.627	16.55	–1.170	5.341

a. Mixture of lognormals

Table 4b: Eurodollar futures options, 23 September 1998
Probabilities for the eurodollar rate on 14 December 1998

23 September	$\text{Prob}[\tilde{r}(T) \leq R]$					
	4.50	**4.75**	**5.00**	**5.25**	**5.50**	**5.75**
Black	0.05	0.17	0.40	0.65	0.84	0.94
MLN[a]	0.08	0.14	0.32	0.63	0.87	0.96
Jump	0.07	0.14	0.33	0.62	0.85	0.96
Hermite (4)	0.08	0.13	0.33	0.63	0.86	0.96
Hermite (6)	0.10	0.14	0.32	0.63	0.87	0.96
Maximum entropy	0.10	0.14	0.30	0.70	0.83	0.99

The probabilities are the risk-neutral probabilities that the market assigns on the given date for the eurodollar rate on 14 December 1998 to be less than the stated *R* value. See the Appendix for details.

a. Mixture of lognormals

Table 5a: Eurodollar futures options, 24 September 1998

24 September 1998	Mean	Volatility	Skewness	Kurtosis
Black	1.613	16.97	0	3
MLN[a]	1.613	18.35	–0.808	3.987
Jump	1.613	18.47	–0.978	4.801
Hermite (4)	1.613	18.37	–0.806	4.524
Hermite (6)	1.613	18.48	–0.998	4.592
Maximum entropy	1.611	17.56	–1.023	4.622

a. Mixture of lognormals

Table 5b: Eurodollar futures options, 24 September 1998
Probabilities for the eurodollar rate on 14 December 1998

24 September 1998	$\text{Prob}[\tilde{r}(T) \leq R]$					
	4.50	**4.75**	**5.00**	**5.25**	**5.50**	**5.75**
Black	0.09	0.25	0.48	0.71	0.87	0.95
MLN[a]	0.09	0.20	0.43	0.70	0.88	0.97
Jump	0.09	0.20	0.43	0.70	0.89	0.97
Hermite (4)	0.10	0.20	0.43	0.70	0.89	0.97
Hermite (6)	0.09	0.20	0.43	0.70	0.88	0.97
Maximum entropy	0.12	0.18	0.47	0.68	0.89	0.99

The probabilities are the risk-neutral probabilities that the market assigns on the given date for the eurodollar rate on 14 December 1998 to be less than the stated *R* value. See the Appendix for details.

a. Mixture of lognormals

Table 6a: Eurodollar futures options, 25 September 1998

25 September 1998	Mean	Volatility	Skewness	Kurtosis
Black	1.614	16.75	0	3
MLN[a]	1.614	19.03	–1.208	5.646
Jump	1.613	19.19	–1.280	6.305
Hermite (4)	1.614	18.26	–0.813	4.735
Hermite (6)	1.614	18.96	–1.227	5.971
Maximum entropy	1.614	18.26	–0.697	3.680

a. Mixture of lognormals

Table 6b: Eurodollar futures options, 25 September 1998
Probabilities for the eurodollar rate on 14 December 1998

25 September 1998	$\text{Prob}[\tilde{r}(T) \leq R]$					
	4.50	**4.75**	**5.00**	**5.25**	**5.50**	**5.75**
Black	0.08	0.24	0.48	0.71	0.87	0.96
MLN[a]	0.08	0.19	0.42	0.69	0.88	0.97
Jump	0.08	0.19	0.43	0.69	0.88	0.97
Hermite (4)	0.09	0.19	0.42	0.70	0.89	0.97
Hermite (6)	0.07	0.19	0.43	0.70	0.88	0.97
Maximum entropy	0.13	0.15	0.49	0.67	0.90	0.96

The probabilities are the risk-neutral probabilities that the market assigns on the given date for the eurodollar rate on 14 December 1998 to be less than the stated *R* value. See the Appendix for details.

a. Mixture of lognormals

Table 7a: Eurodollar futures options, 28 September 1998

28 September 1998	Mean	Volatility	Skewness	Kurtosis
Black	1.619	16.21	0	3
MLN[a]	1.618	19.20	–1.712	10.699
Jump	1.618	18.66	–1.563	7.842
Hermite (4)	1.619	17.55	–0.749	5.017
Hermite (6)	1.618	18.51	–1.168	6.897
Maximum entropy	1.617	17.43	–0.596	4.681

a. Mixture of lognormals

Table 7b: Eurodollar futures options, 28 September 1998
Probabilities for the eurodollar rate on 14 December 1998

28 September 1998	$\text{Prob}[\tilde{r}(T) \leq R]$					
	4.50	**4.75**	**5.00**	**5.25**	**5.50**	**5.75**
Black	0.06	0.21	0.45	0.70	0.87	0.96
MLN[a]	0.06	0.16	0.40	0.69	0.89	0.97
Jump	0.06	0.16	0.40	0.69	0.89	0.97
Hermite (4)	0.08	0.16	0.39	0.69	0.90	0.97
Hermite (6)	0.05	0.16	0.41	0.69	0.89	0.98
Maximum entropy	0.11	0.14	0.45	0.67	0.89	0.98

The probabilities are the risk-neutral probabilities that the market assigns on the given date for the eurodollar rate on 14 December 1998 to be less than the stated *R* value. See the Appendix for details.

a. Mixture of lognormals

Table 8a: Eurodollar futures options, 29 September 1998

29 September 1998	Mean	Volatility	Skewness	Kurtosis
Black	1.629	13.74	0	3
MLN[a]	1.629	15.24	–0.711	6.681
Jump	1.629	15.46	–1.754	9.955
Hermite (4)	1.629	15.07	–0.608	6.026
Hermite (6)	1.629	14.45	–0.952	3.206
Maximum entropy	1.629	14.86	–0.718	5.984

a. Mixture of lognormals

Table 8b: Eurodollar futures options, 29 September 1998
Probabilities for the eurodollar rate on 14 December 1998

29 September 1998	$\text{Prob}[\tilde{r}(T) \le R]$					
	4.50	**4.75**	**5.00**	**5.25**	**5.50**	**5.75**
Black	0.02	0.12	0.37	0.68	0.89	0.97
MLN[a]	0.05	0.11	0.30	0.70	0.92	0.97
Jump	0.03	0.10	0.33	0.67	0.90	0.98
Hermite (4)	0.06	0.09	0.31	0.69	0.92	0.98
Hermite (6)	0.08	0.09	0.30	0.71	0.90	0.95
Maximum entropy	0.06	0.10	0.30	0.71	0.92	0.96

The probabilities are the risk-neutral probabilities that the market assigns on the given date for the eurodollar rate on 14 December to be less than the stated *R* value. See the Appendix for details.

a. Mixture of lognormals

Table 9a: Eurodollar futures options, 30 September 1998

30 September 1998	Mean	Volatility	Skewness	Kurtosis
Black	1.617	14.19	0	3
MLN[a]	1.617	16.92	–1.434	8.794
Jump	1.617	16.88	–1.755	8.810
Hermite (4)	1.618	15.58	–0.848	5.623
Hermite (6)	1.618	15.39	–0.700	5.294
Maximum entropy	1.615	15.09	–1.216	5.216

a. Mixture of lognormals

Table 9b: Eurodollar futures options, 30 September 1998
Probabilities for the eurodollar rate on 14 December 1998

30 September 1998	$\text{Prob}[\tilde{r}(T) \leq R]$					
	4.50	**4.75**	**5.00**	**5.25**	**5.50**	**5.75**
Black	0.04	0.18	0.45	0.74	0.91	0.98
MLN[a]	0.07	0.13	0.37	0.75	0.93	0.98
Jump	0.05	0.14	0.40	0.73	0.92	0.99
Hermite (4)	0.07	0.13	0.38	0.74	0.94	0.99
Hermite (6)	0.08	0.14	0.38	0.74	0.94	0.99
Maximum entropy	0.07	0.14	0.40	0.76	0.95	1.00

The probabilities are the risk-neutral probabilities that the market assigns on the given date for the eurodollar rate on 14 December 1998 to be less than the stated *R* value. See the Appendix for details.

a. Mixture of lognormals

Figure 1: Eurodollar futures options, 23 September 1998

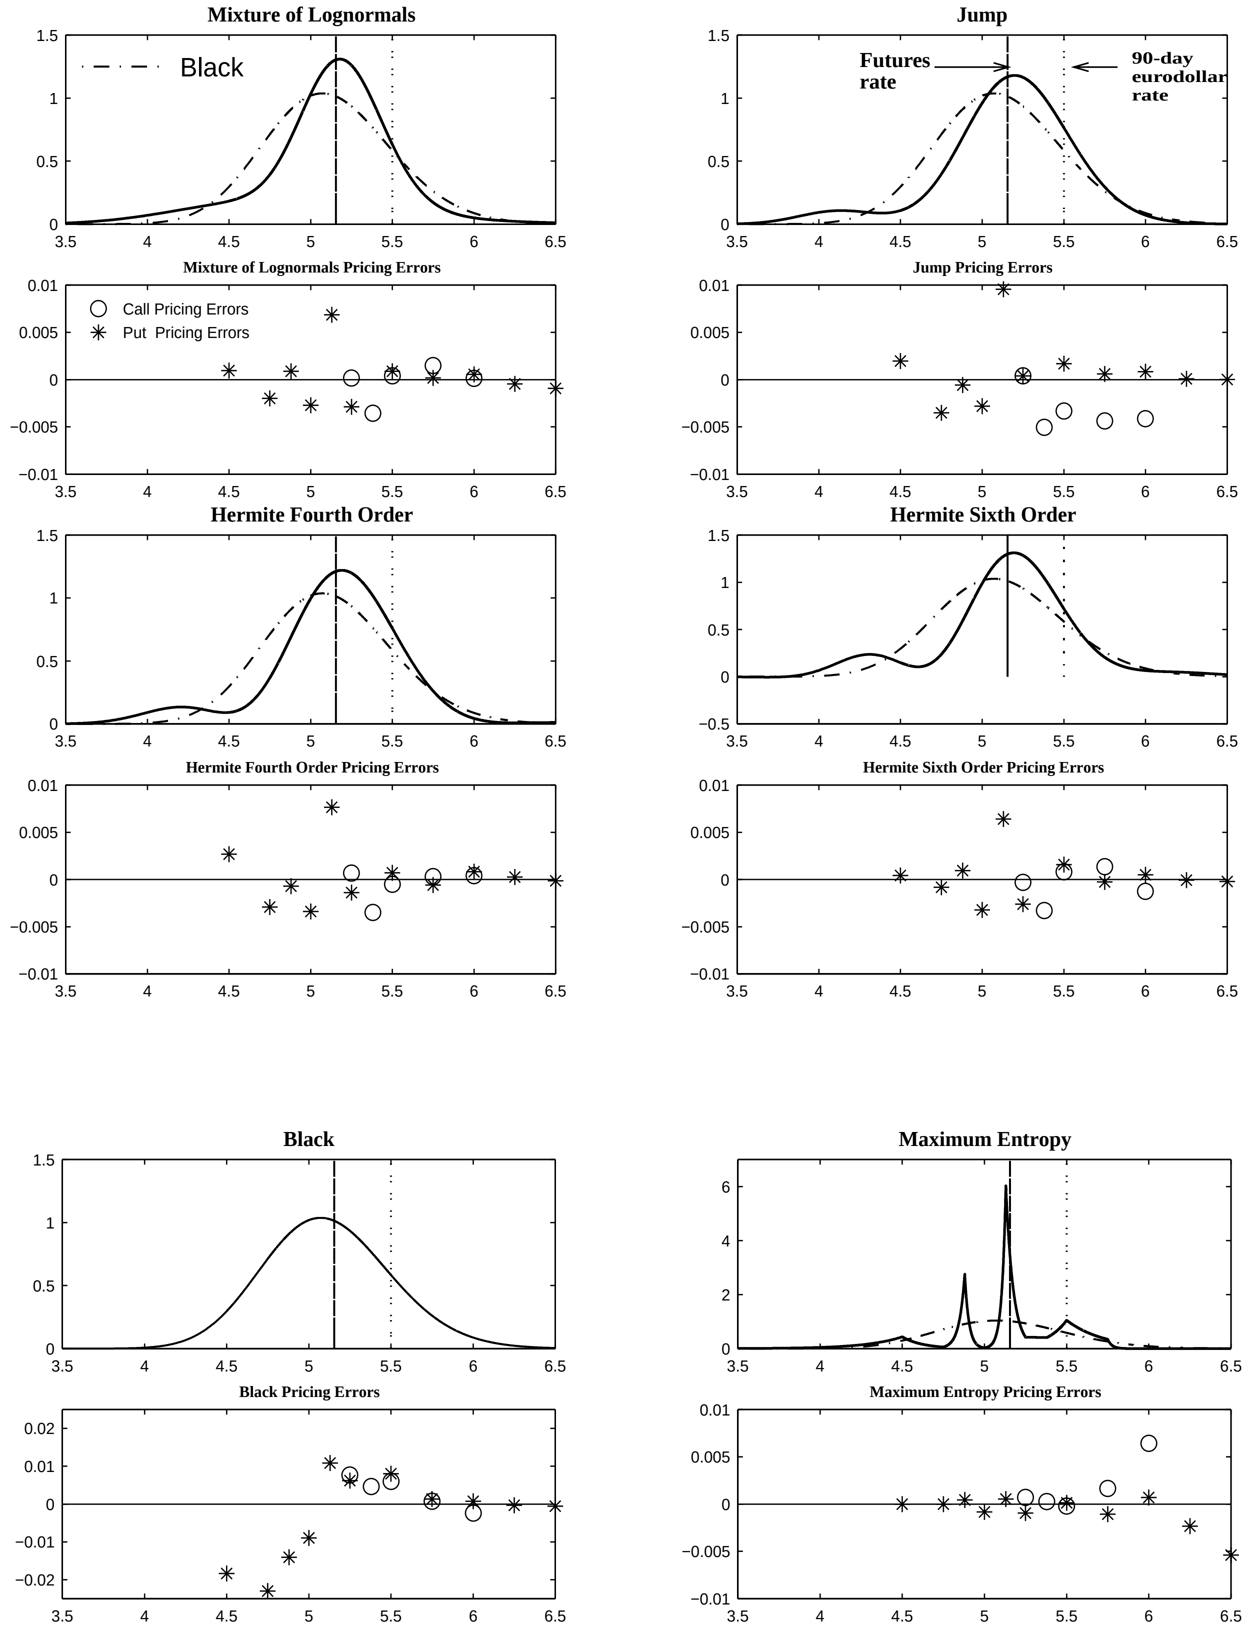

Figure 2: Eurodollar futures options, 24 September 1998

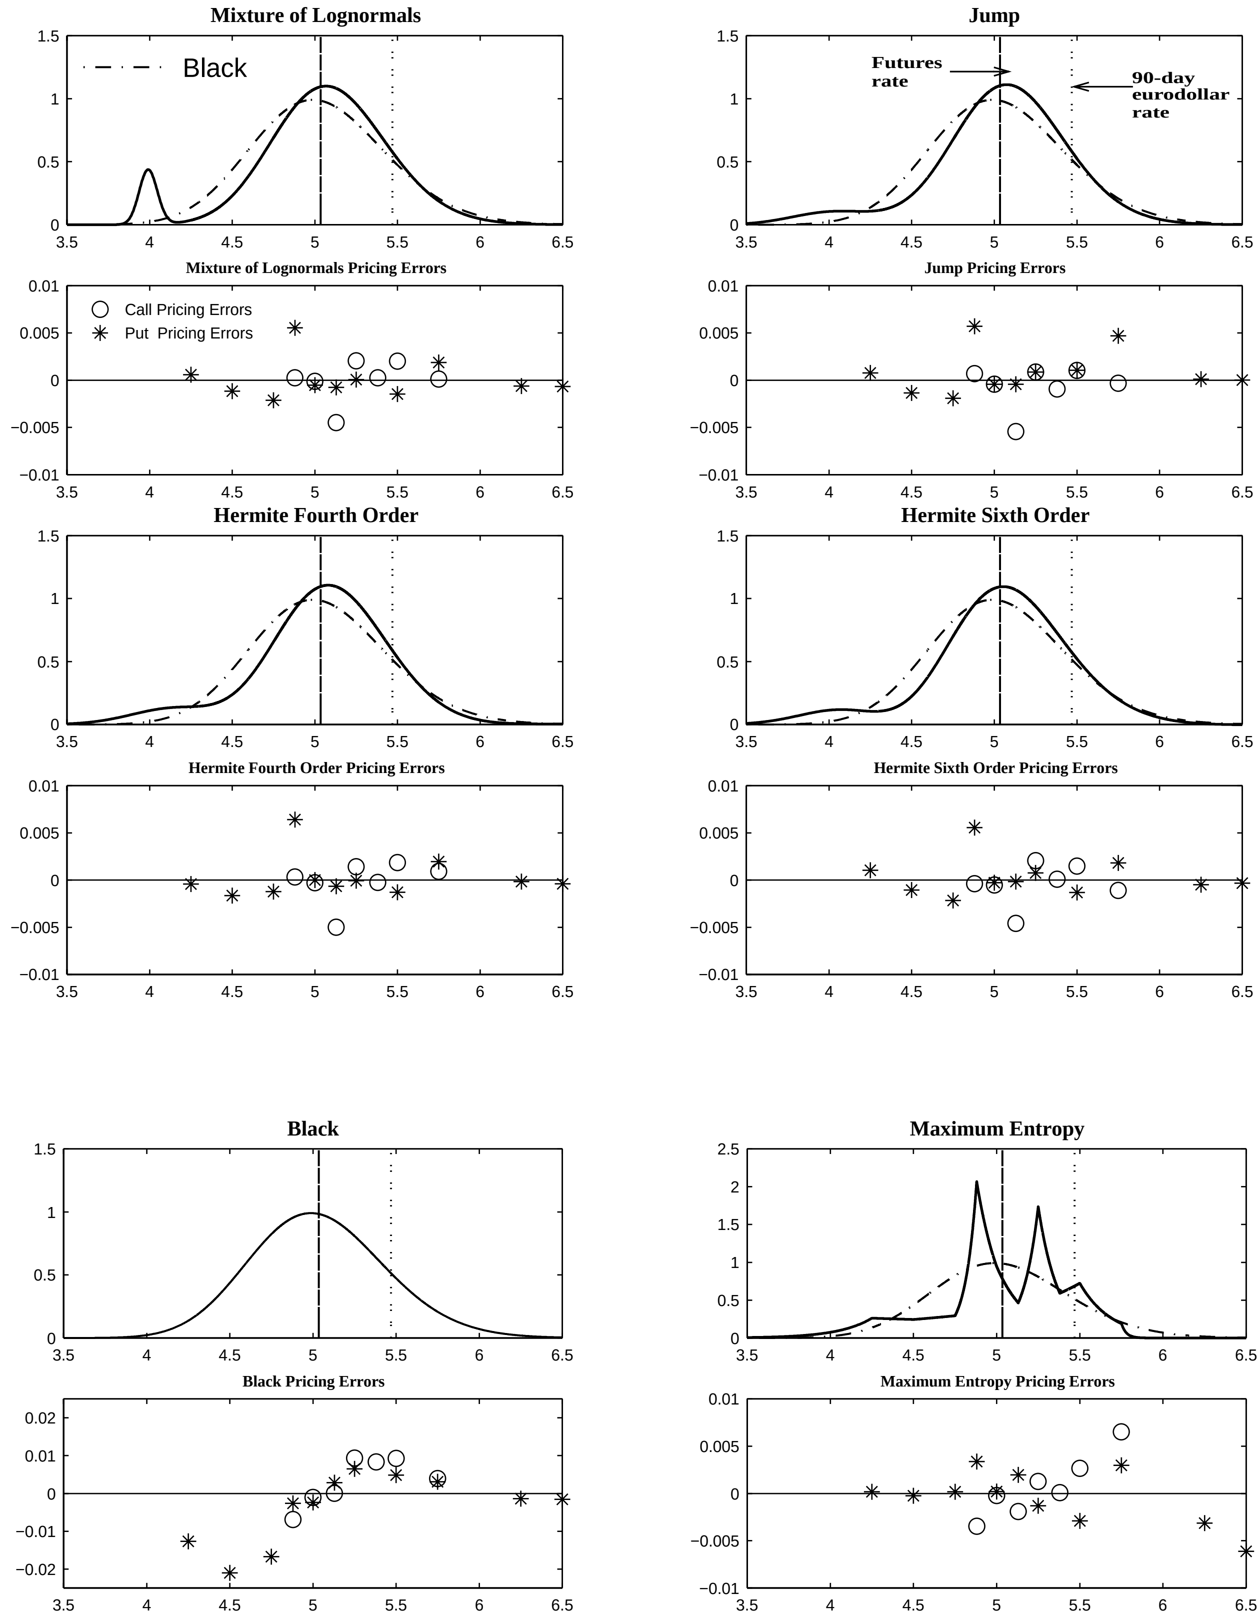

Figure 3: Eurodollar futures options, 25 September 1998

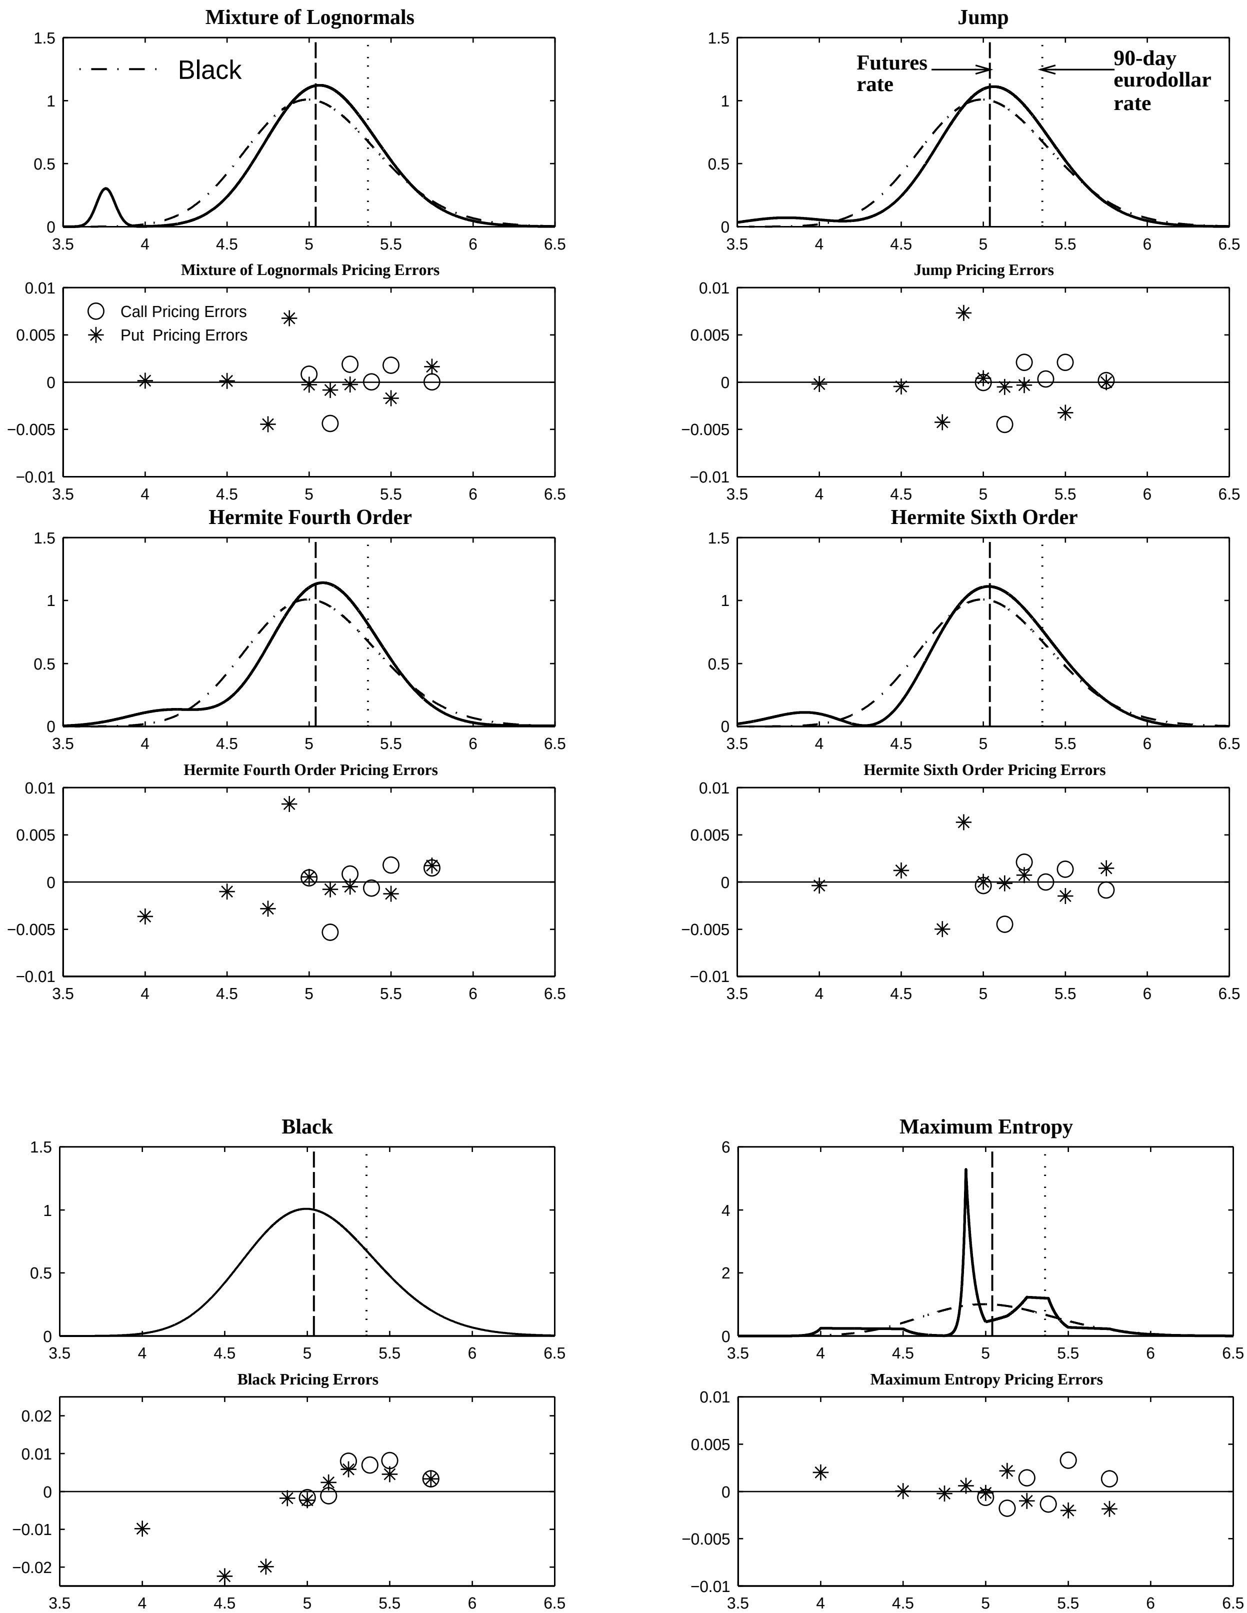

Figure 4: Eurodollar futures options, 28 September 1998

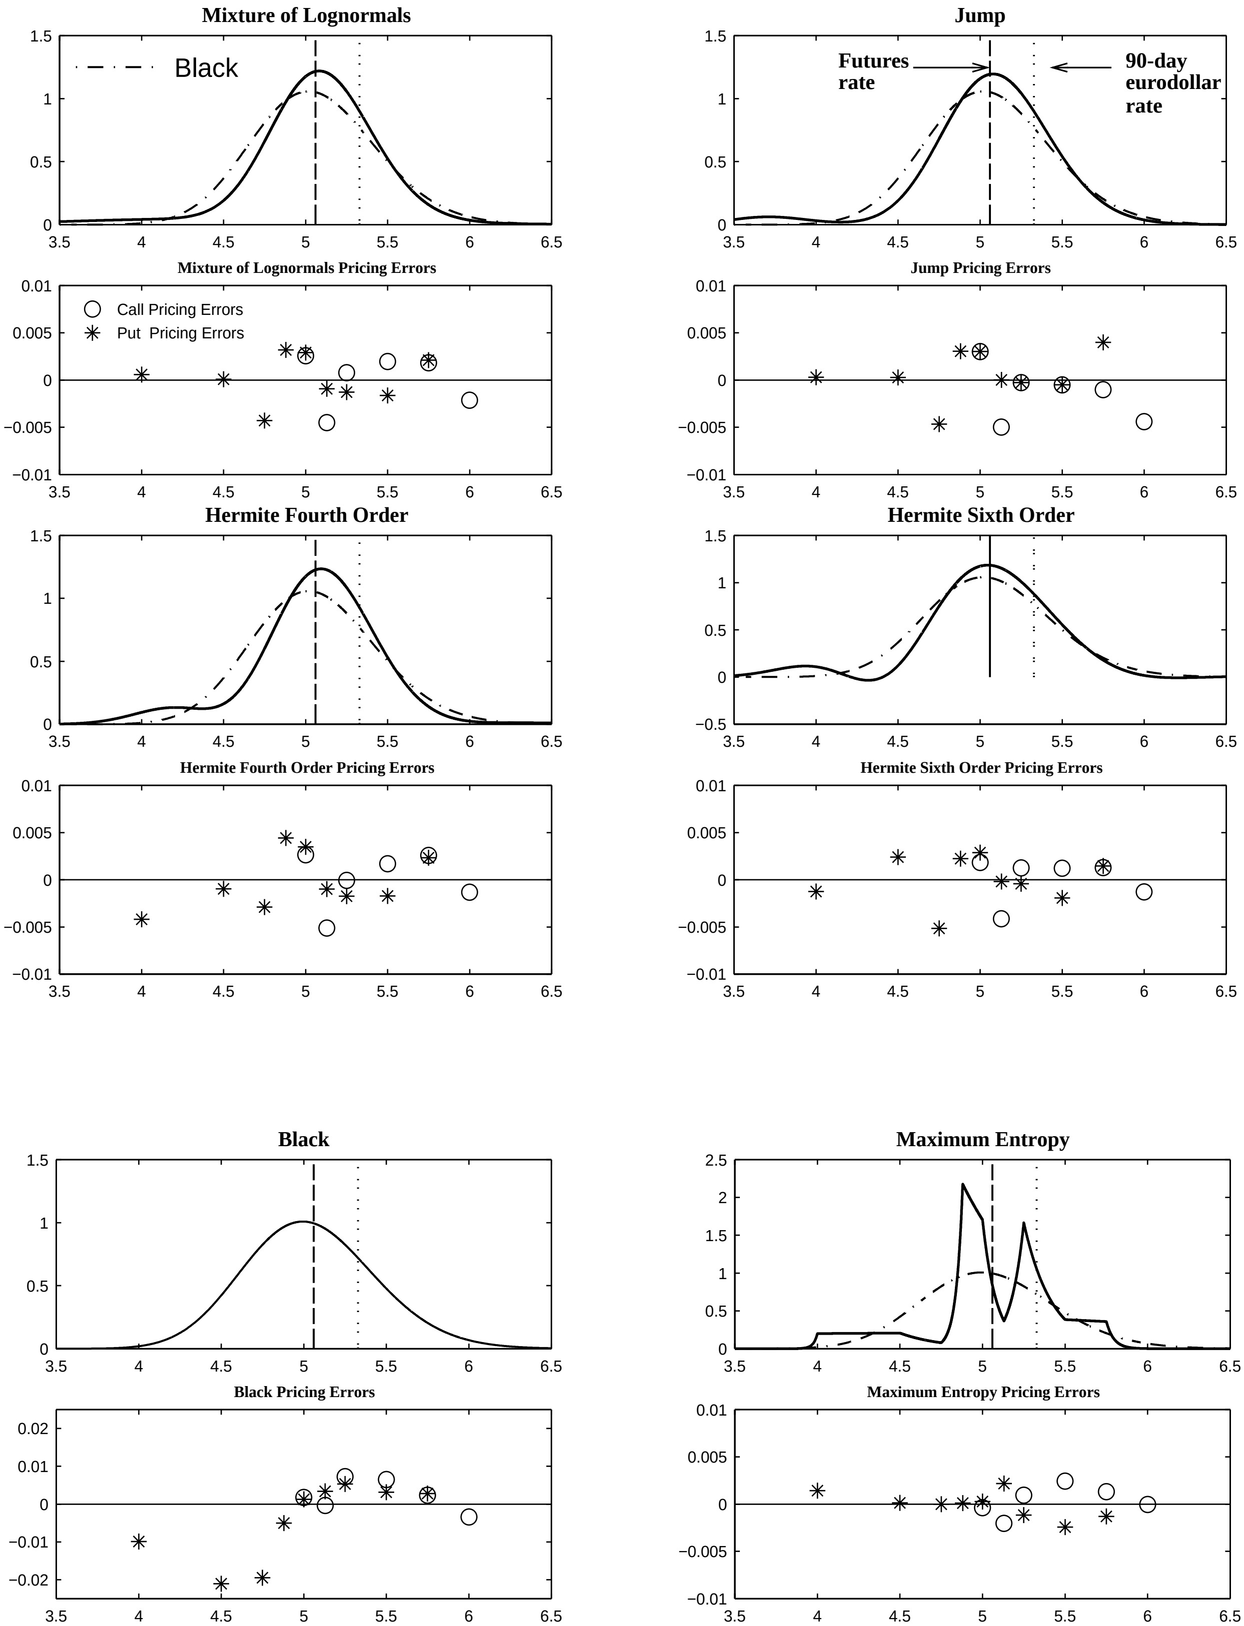

Figure 5: Eurodollar futures options, 29 September 1998

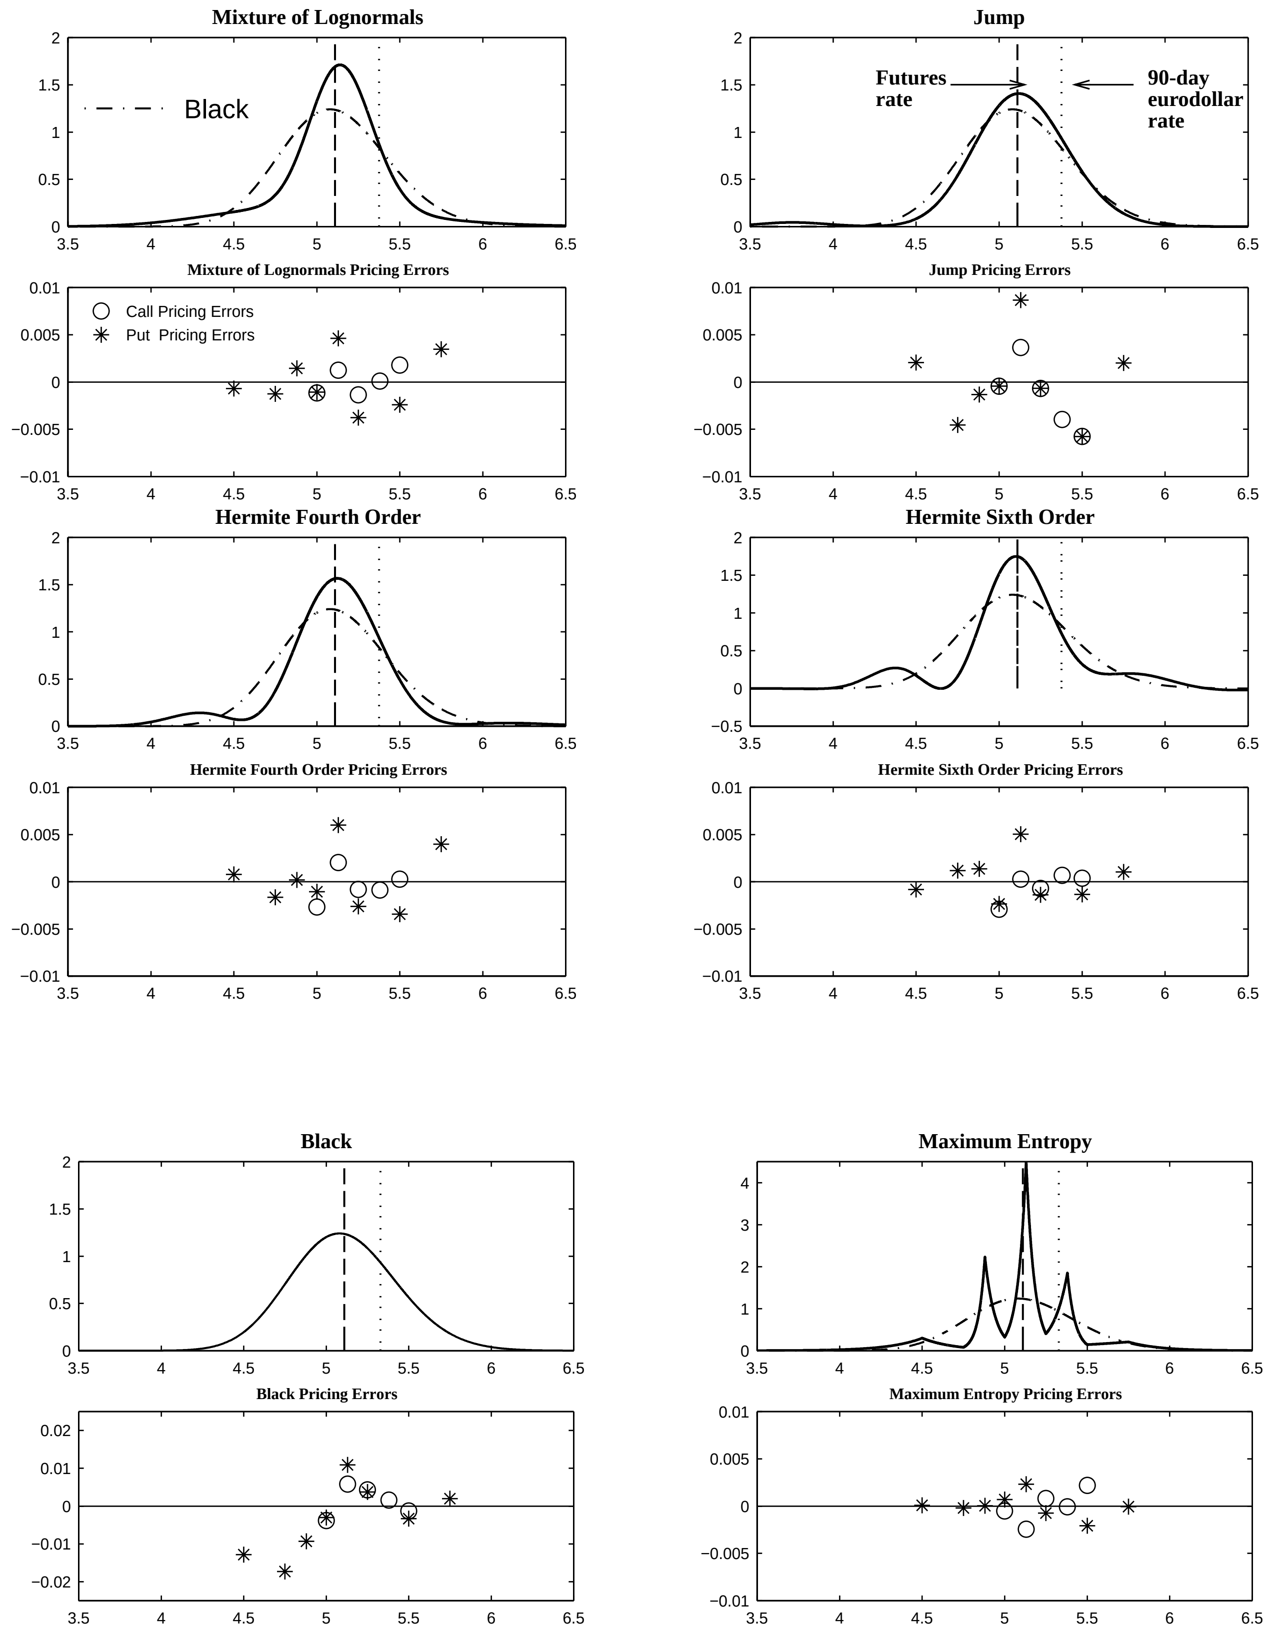

Figure 6: Eurodollar futures options, 30 September 1998

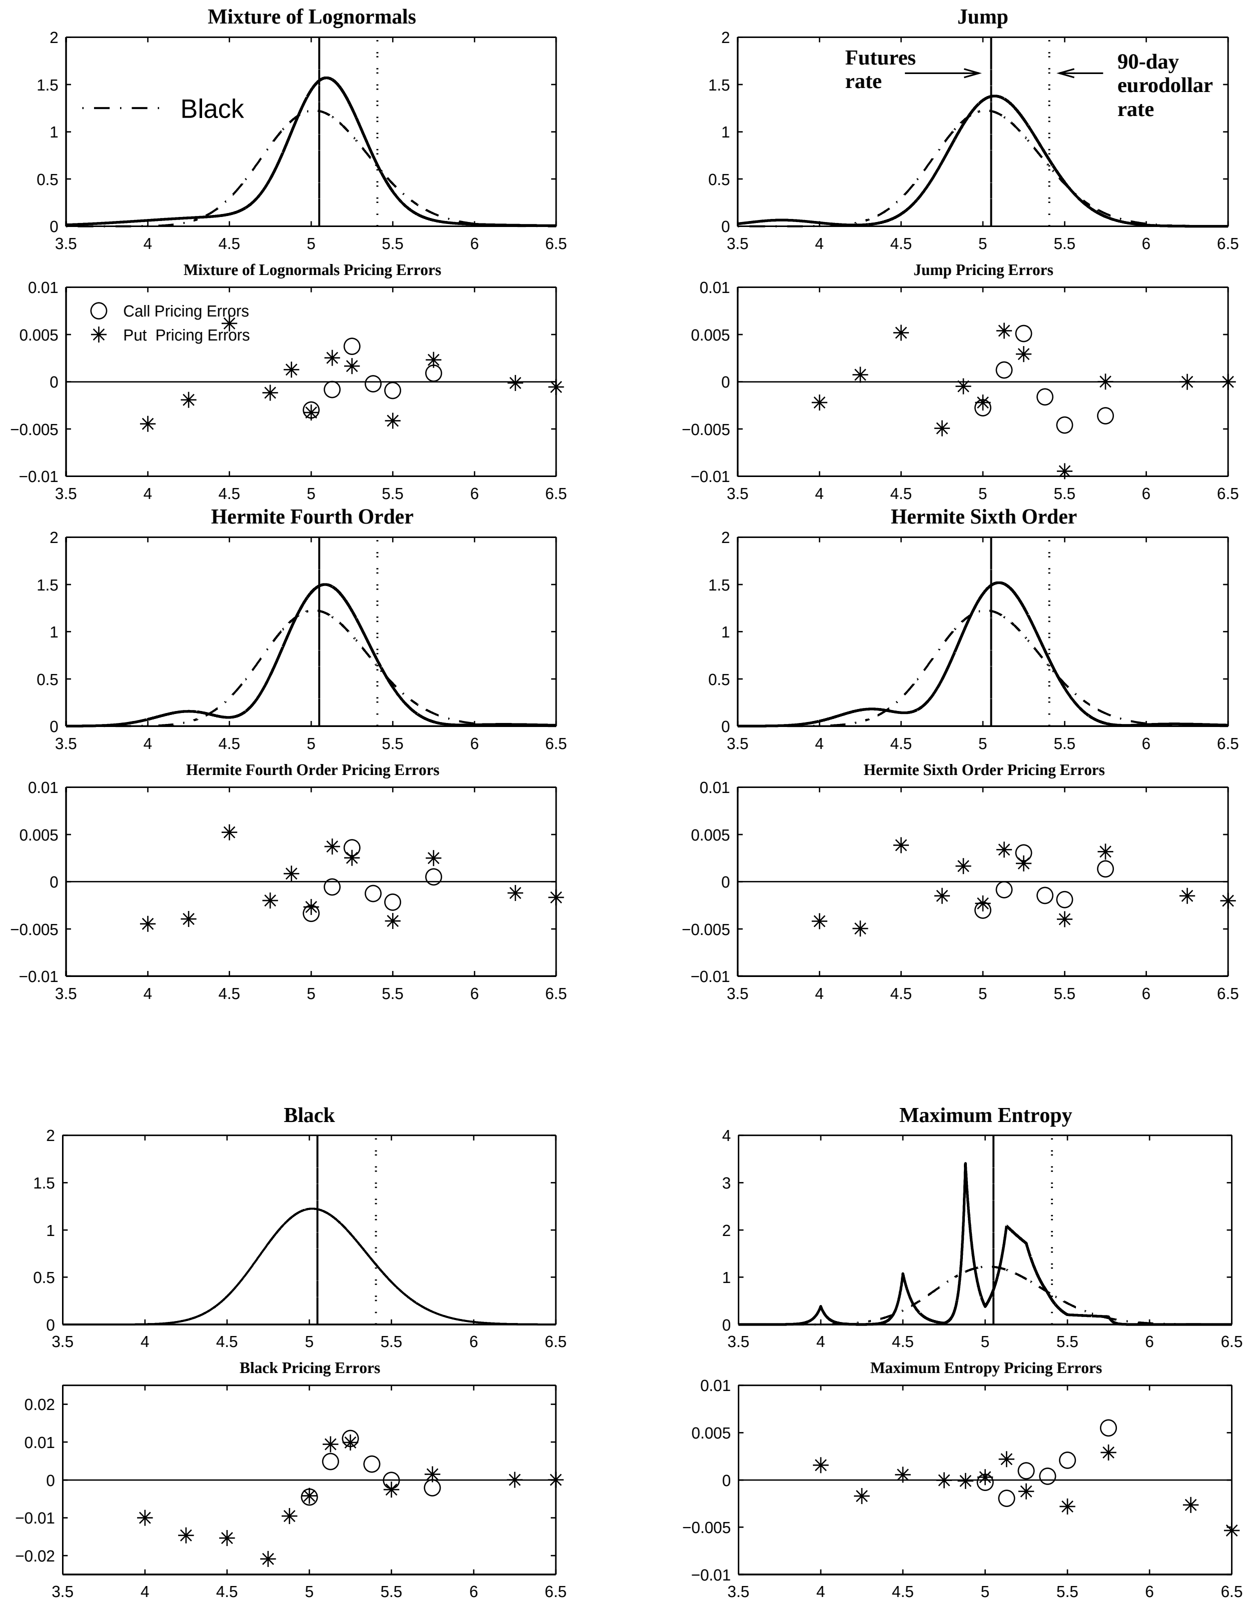

Figure 7: Eurodollar futures options, moments

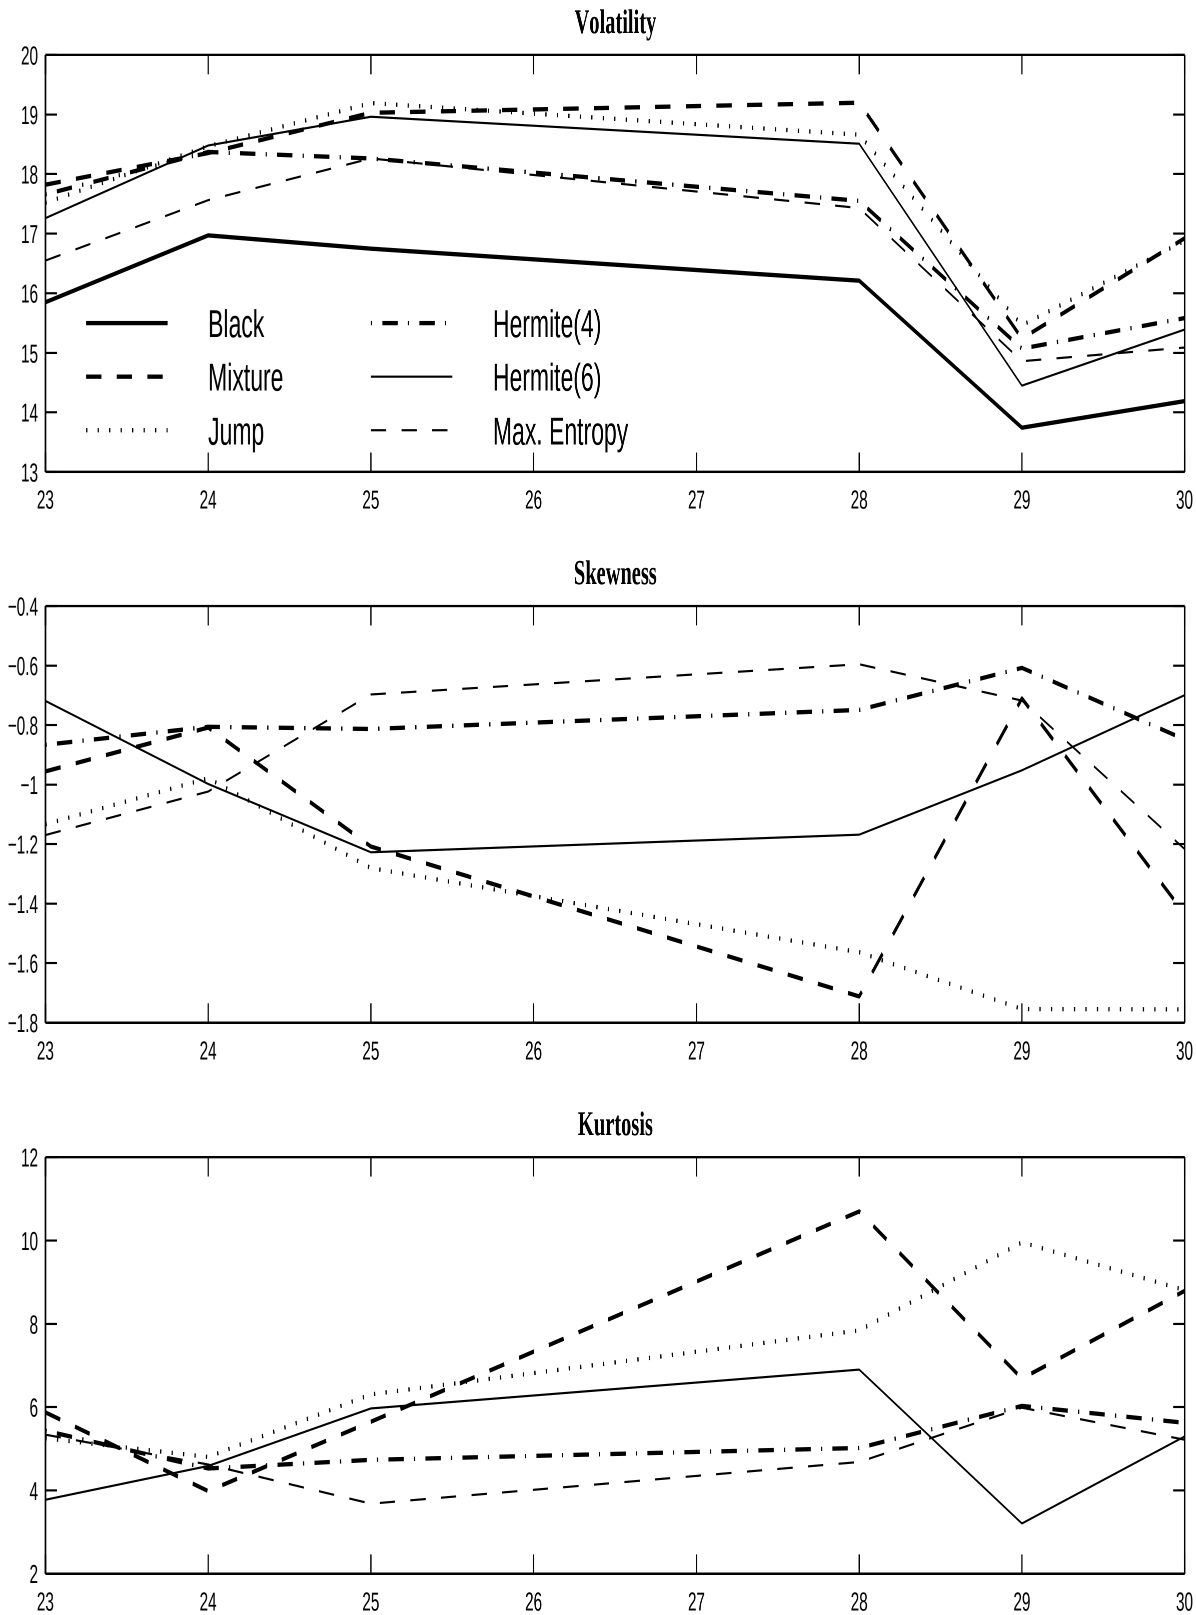

Figure 8a: Cumulative distributions

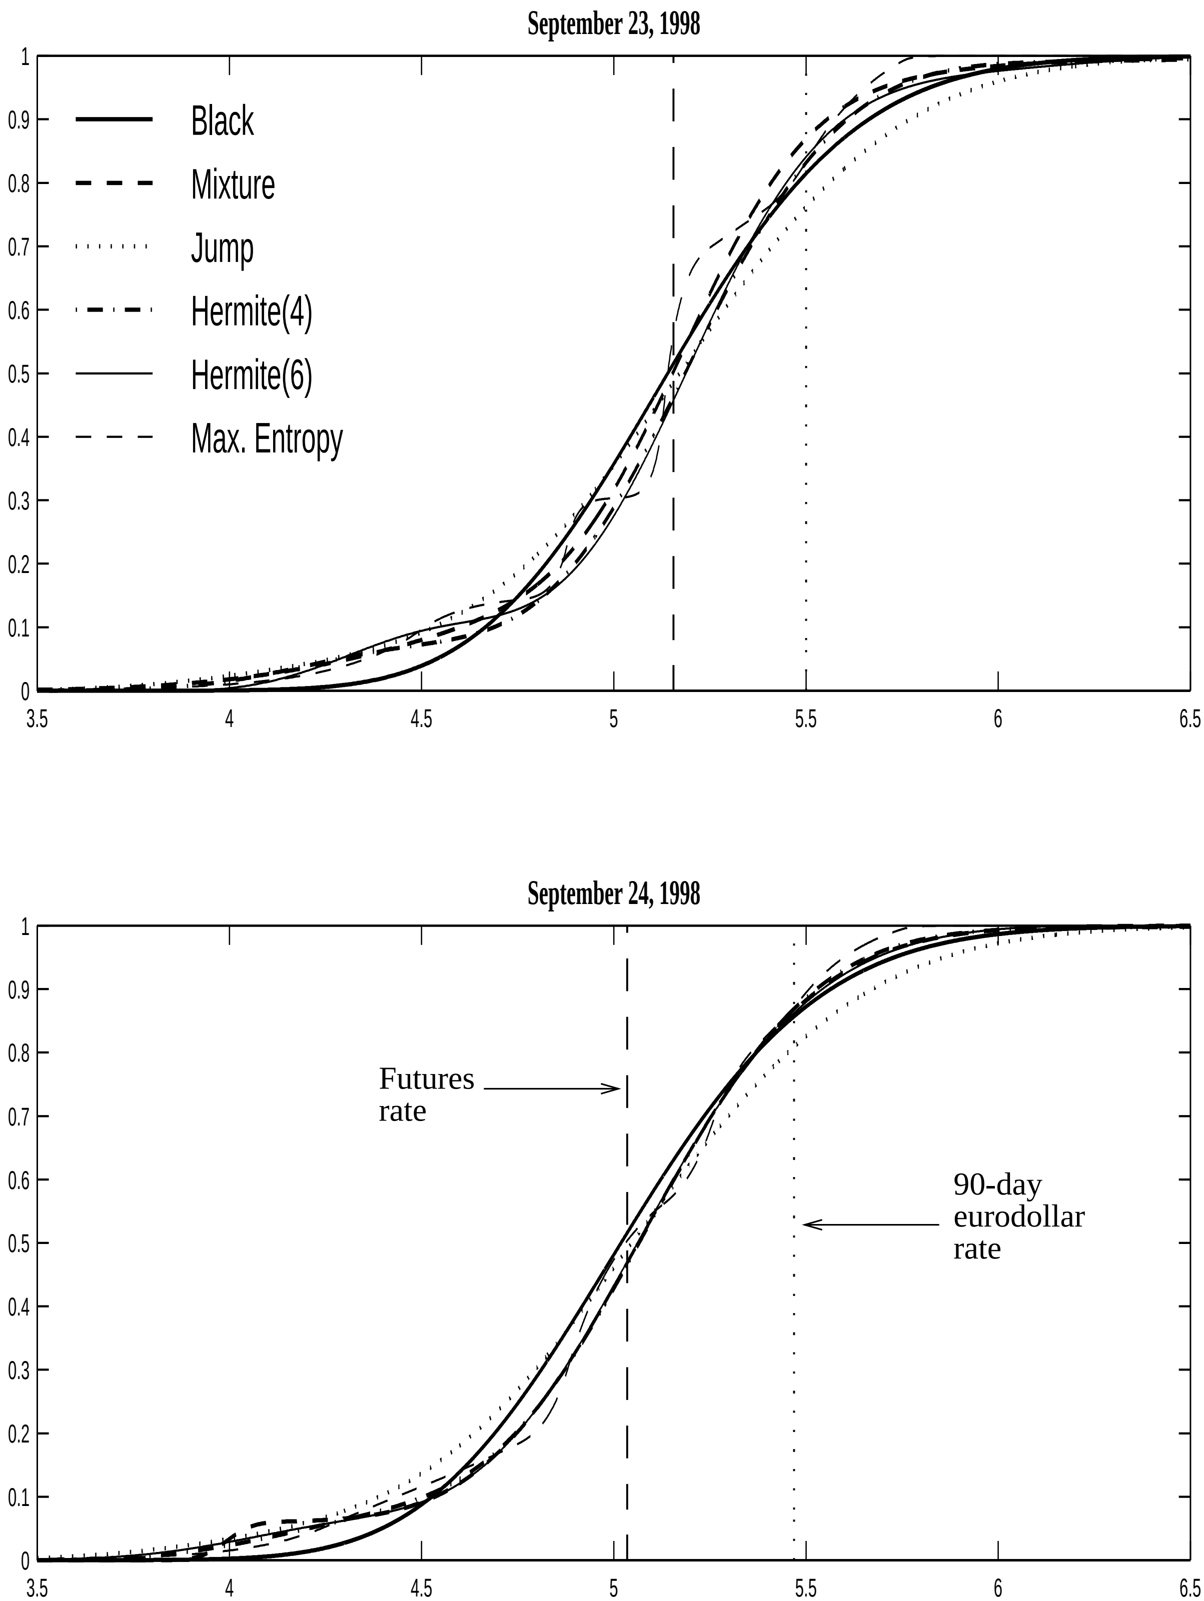

Figure 8b: Cumulative distributions

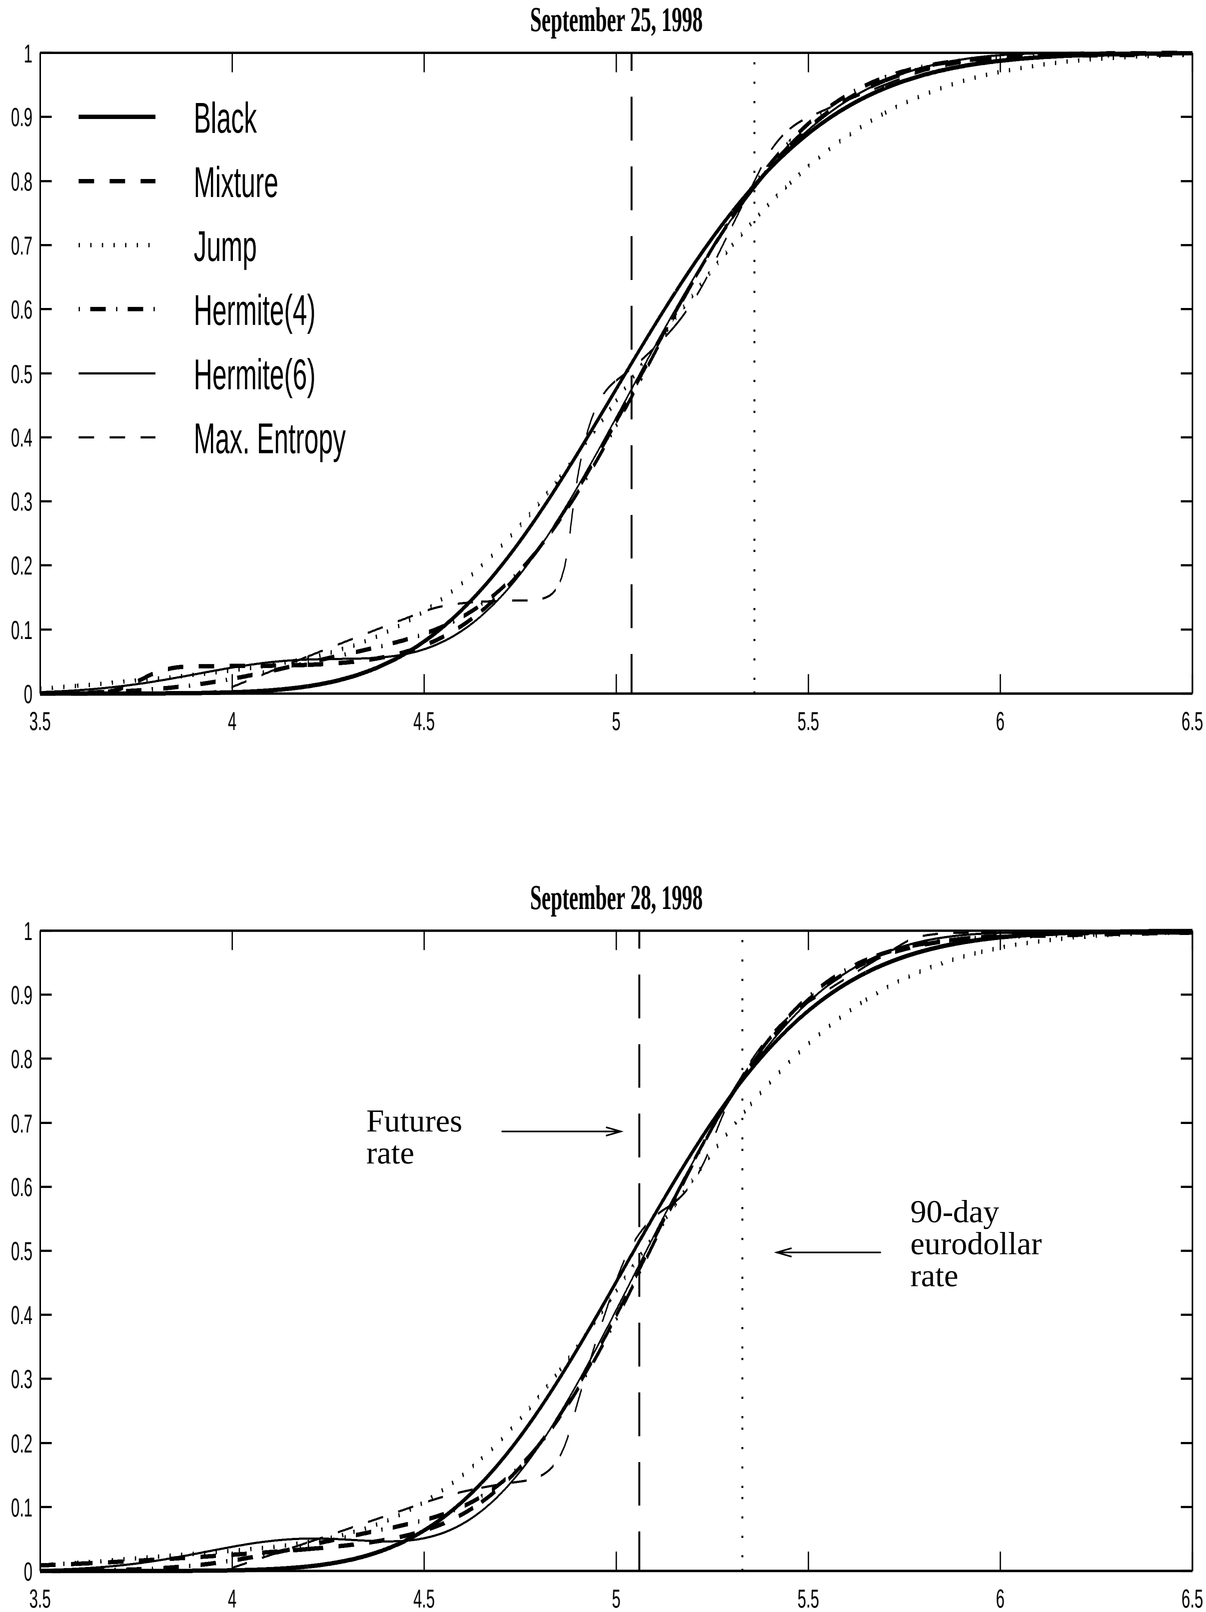

Figure 8c: Cumulative distributions

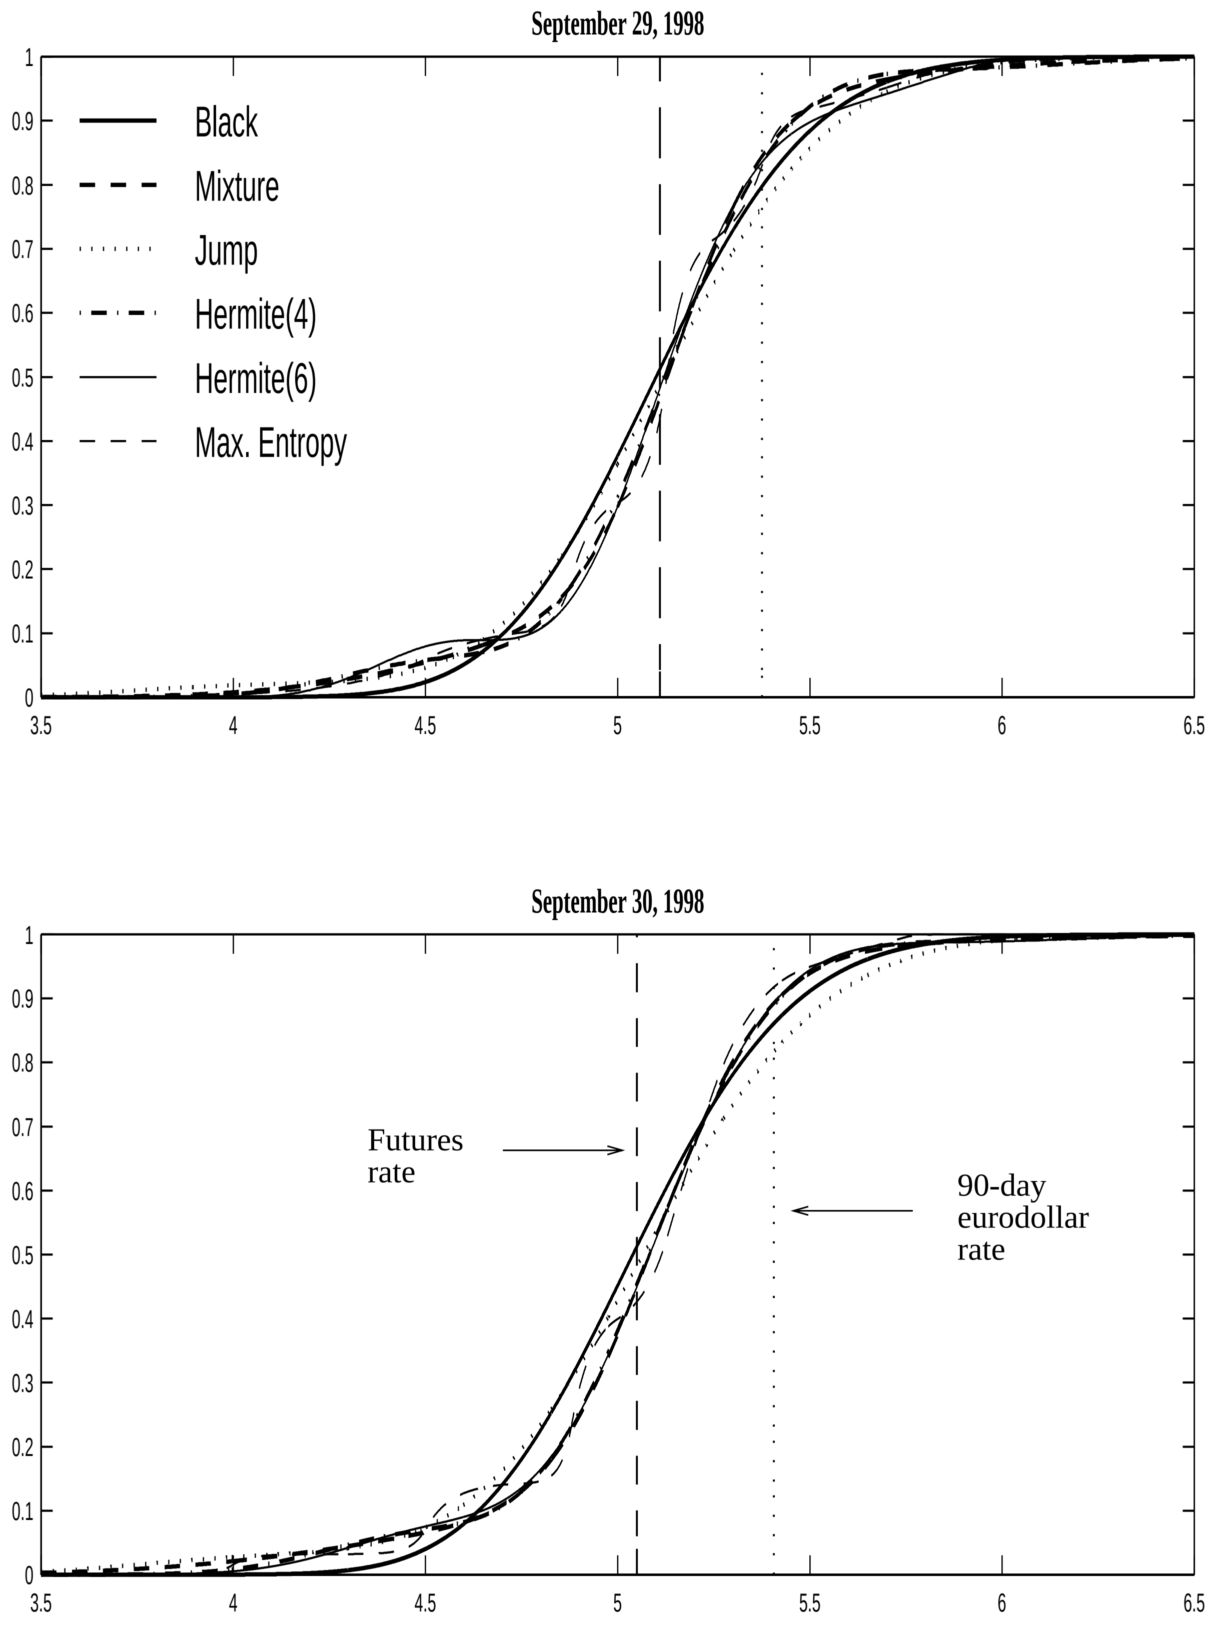

Figure 8d: Cumulative distribution

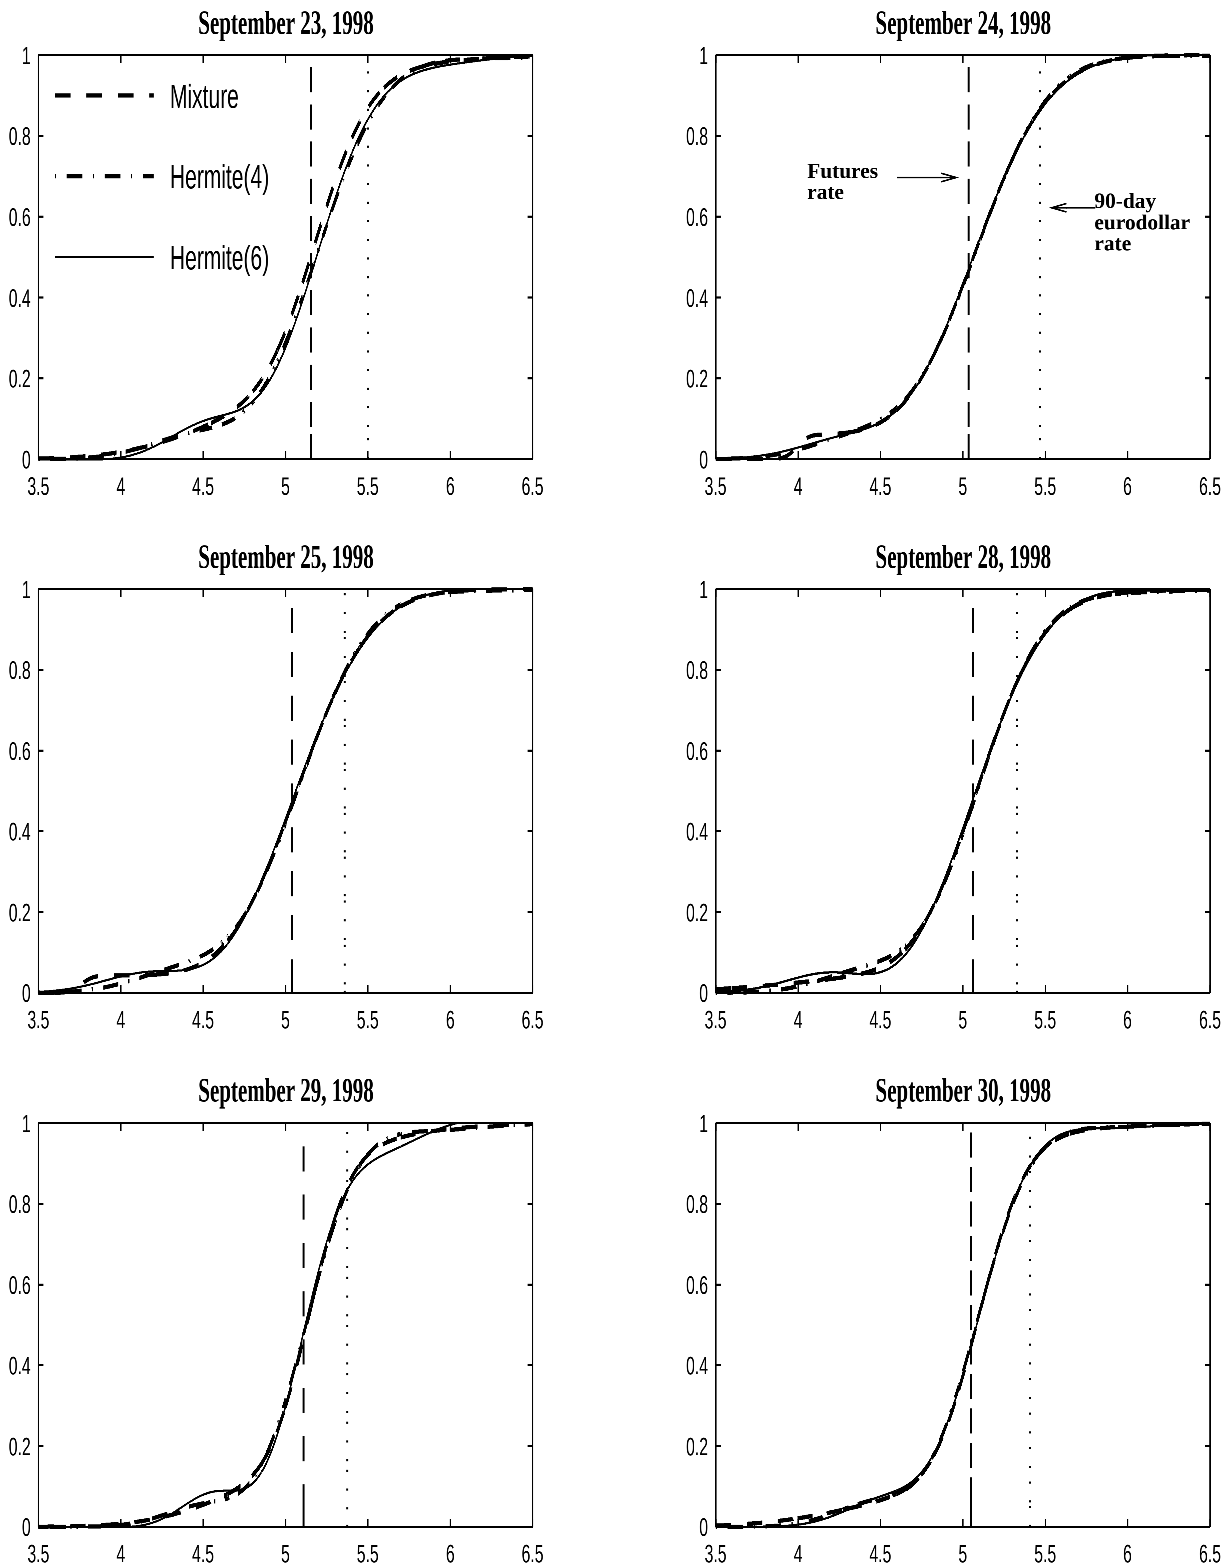

Appendix: PDF summary statistics

The risk-neutral PDF, $q[\tilde{r}(T)]$, synopsizes the information contained in the price of interest rate futures options. Thus, a graphical depiction of the risk-neutral PDF yields market perceptions over the future value of interest rates. In addition, several numerical statistics also yield helpful insights. In particular, the probability that the futures rate will be less than a given rate, R, on the maturity date of the futures contract is insightful, namely

$$\text{Prob}[\tilde{r}(T) \le R] = \int_0^R q[\tilde{r}(T)]\, d\tilde{r}(T)\,. \tag{47}$$

In addition, several summary statistics calculated with respect to the logarithm of the futures rate are useful, such as the mean, annualized volatility, skewness, and kurtosis. The annualized volatility provides an indication of the dispersion of opinion in the market surrounding the future interest rate. The skewness compares the probability of a large upward movement in the futures rate to the probability of a large downward movement. Risk-neutral PDFs are either symmetric, skewed left or skewed right. A skewed left distribution places greater weight on the likelihood the future interest rate will be far below, as opposed to far above, the current futures price on the maturity date of the option. Finally, kurtosis indicates the possibility of large changes in interest rates prior to the maturity of the futures option.

Note that is the futures rate $\tilde{r}(T)$ has a PDF $q[\tilde{r}(T)]$ then the logarithm of the futures rate, $\log[\tilde{r}(T)]$, has a PDF $Q(\log[\tilde{r}(T)]) = \tilde{r}(T)\,\mathrm{q}[\tilde{r}(T)]$. Thus, the mean, variance, skewness, and kurtosis with respect to the logarithm of the futures rate are given, respectively, by

$$\begin{aligned}
\mu &= E_Q[\log\tilde{r}(T)] \\
\text{Var} &= E_Q[(\log\tilde{r}(T) - \mu)^2] \\
\text{Skew} &= \left(E_Q[(\log\tilde{r}(T) - \mu)^3]\right)/\text{Var}^{3/2} \\
\text{Kurt} &= \left(E_Q[(\log\tilde{r}(T) - \mu)^4]\right)/\text{Var}^2
\end{aligned} \tag{48}$$

where E_Q represents expectations with respect to the PDF Q. The annualized volatility is given by $\sigma = \sqrt{\text{Var}/T}$.

The PDF summary statistics for the models outlined in section 3 are as follows.

A.1 Black's model

The cumulative distribution function is

$$\text{Prob}[\tilde{r}(\text{T}) \le R] = N\left(\frac{\log\{R/F(0,T)\}}{\sigma\sqrt{T}} + \frac{1}{2}\sigma\sqrt{T}\right) \tag{49}$$

The variance, skewness, and kurtosis are

$$\text{Var} = \sigma^2 T\text{, Skew} = 0\text{, and Kurt} = 3. \tag{50}$$

A.2 Mixture of lognormals

The cumulative distribution function is

$$\text{Prob}[\tilde{r}(T) \le R] = \phi_1 N\left(\frac{\log R - \mu_1}{\sigma_1}\right) + (1-\phi_1)N\left(\frac{\log R - \mu_2}{\sigma_2}\right) \tag{51}$$

The variance, skewness, and kurtosis are

$$\begin{aligned}
\text{Var} &= \phi_1\sigma_1^{\ 2} + \phi_2\sigma_2^{\ 2} + \phi_1\phi_2(\mu_1-\mu_2)^2 \\
\text{Skew} &= \left\{\phi_1\phi_2(\mu_1-\mu_2)\left[3\left(\sigma_1^{\ 2}-\sigma_2^{\ 2}\right) + (\phi_2-\phi_1)(\mu_1-\mu_2)^2\right]\right\}/\text{Var}^{3/2} \\
\text{Kurt} &= \left\{3\left(\phi_1\sigma_1^{\ 4} + \phi_2\sigma_2^{\ 4}\right) + 6\phi_1\phi_2(\mu_1-\mu_2)^2\left[\phi_2\sigma_1^{\ 2} + \phi_1\sigma_2^{\ 2}\right] + \phi_1\phi_2(\mu_1-\mu_2)^4\left(\phi_1^3+\phi_2^3\right)\right\}/\text{Var}^2
\end{aligned} \tag{52}$$

A.3 Jump diffusion

The cumulative distribution and PDF summary statistics are given by equations (51) and (52) above with $\phi_1 = 1-\lambda T$, $\mu_1 = \log F(0,T) - \log(1+\lambda\kappa T) - \frac{1}{2}\sigma_\omega^2 T$, $\mu_2 = \mu_1 + \log(1+\kappa)$, and $\sigma_1 = \sigma_\omega\sqrt{T} = \sigma_2$.

A.4 Hermite polynomial approximation

The cumulative distribution is

$$\begin{aligned}
\text{Prob}[\tilde{r}(T) \le R] = {} & N(Z) + \frac{1}{\sqrt{2\pi}}\exp\left[-\frac{Z^2}{2}\right]\left[\frac{b_3}{\sqrt{6}}(1-Z^2) + \frac{b_4}{\sqrt{24}}Z(3-Z^2)\right. \\
& \left. + \frac{b_5}{\sqrt{120}}(-3+6Z^2-Z^4) + \frac{b_6}{\sqrt{720}}Z(-15+10Z^2-Z^4)\right]
\end{aligned} \tag{53}$$

where $Z = \dfrac{\log[R/F(0,T)] - \left(\mu - \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}$ and the restrictions $b_0 = 1$, $b_1 = 0$ and $b_2 = 1$ have been imposed. (The fourth-order approximation is given by setting $b_5 = 0$ and $b_6 = 0$.) Under the same restrictions, the variance, skewness, and kurtosis are

$$\begin{aligned} \text{Var} &= \sigma^2 T \\ \text{Skew} &= \sqrt{6}\,b_3 \\ \text{Kurt} &= 3 + \sqrt{24}\,b_4 \end{aligned} \quad . \tag{54}$$

A.5 Method of maximum entropy

The cumulative distribution is given by equation (55) if $X_i \le R < X_{i+1}$

$$\text{Prob}[\tilde{r}(T) \le R] = \left(\begin{array}{l} \dfrac{1}{\mu\, a_i}\{\exp[a_i R + b_i] - \exp[a_i X_i + b_i]\} \\ + \displaystyle\sum_{k=0}^{i-1} \dfrac{1}{\mu\, a_k}\{\exp[a_k X_{k+1} + b_k] - \exp[a_k X_k + b_k]\} \end{array}\right. \tag{55}$$

and by equation (56) if $X_m \le R$:

$$\text{Prob}[\tilde{r}(T) \le R] = \left(\begin{array}{l} \dfrac{1}{\mu\, a_m}\{\exp[a_m R + b_m] - \exp[a_m X_m + b_m]\} \\ + \displaystyle\sum_{k=0}^{m-1} \dfrac{1}{\mu\, a_k}\{\exp[a_k X_{k+1} + b_k] - \exp[a_k X_k + b_k]\} \end{array}\right. . \tag{56}$$

(57)

(58)

No closed-form solutions exist for the mean, variance, skewness, and kurtosis. These statistics are calculated by numerical integration. (59)

Discussion of

"The Information Content of Interest Rate Futures Options"

by Des McManus, Bank of Canada

Discussant: James M. Mahoney

This paper attempts to determine which method of estimating the Eurodollar futures' probability density function (PDF), of six methods under consideration, most accurately fits a cross-section of observed Eurodollar futures options prices. The paper uses various metrics to determine which method is flexible enough to price the array of available options prices found in the market, e.g., with the implicit smiles found in the implied volatility – strike price graphs. The paper is well done and advances our understanding of the relative effectiveness in each of the methodologies. Most of my comments are directed at extending the already-impressive amount of work evident in the paper.

Data issues: Observation weights

In coming up with the parameters of the risk-neutral PDF, the estimation procedure is set up to minimize the equally weighted sum of squared deviations between the modeled prices and the market observed prices. There may be some good reasons to weight these observations differently in the minimization problem. First, weighting the observations by trading volume may make sense if there are differences in bid-ask spreads or non-synchronous trading, which translates in differences in measurement errors, across the observations. Second, there may be reasons to underweight deep out-of-the-money options, as little or no information may be available in options whose market prices are near the exchange-mandated minimum tick size (0.01 for Eurodollar futures and options) and where small measurement errors may significantly alter the results of the minimization problem. And third, although not an issue in the empirical section of this paper due to the small sample window, one needs to decide how close to expiration are the prices of option useful in deriving risk-neutral PDFs.

Metrics for comparison

Two criteria are used as metrics of comparison among the six models. First, the mean squared error (MSE) and the mean squared percent error (MSPE) are used. Again, the issue of how these observations should be weighted arises. Equally weighting the observations may not provide reliable results, as, e.g., a small absolute error may yield large percent errors, which would dominate the MSPE metric.

The second metric of comparison in the paper is based on a set of summary statistics: volatility, skewness, and kurtosis. Volatility measures were consistent across methodologies, while the measures of skewness and kurtosis were erratic. A concern here is the lack of robustness in these summary statistics. For example, the volatility, skewness and kurtosis are very different across six methodologies in Table 7a, but the PDFs across the six methodologies are remarkably similar in Table 7b. In addition to robustness, an additional concern is the difficulty in interpreting the summary statistics. For example, again in Table 7, the volatility and kurtosis of the mixture of log-normal (MLN) methodology are greater than the volatility and kurtosis of the Hermite polynomial of degree 4 (H(4)), which might suggest that large movements in interest rates are more likely for the MLN methodology. Yet, the probability of a large move in rates (with rates rising above 5.75% or falling below 4.5%) is actually greater for the H(4) methodology. This counter-intuitive result (driven by difference in skewness in the PDFs) suggests that the summary statistics may not be the most useful gauges of PDFs, as these summary statistics have non-natural units of measurement that are hard to interpret. Perhaps more useful would be a more extensive evaluation of the PDFs from the perspective of various quintiles (i.e., drop Table 7a and extend Table 7b).

Additional metrics for comparison could be used. For example, the cross-sectional PDF could be estimated using a sub-sample of available options prices, and the out-of-sample cross-sectional fit could be tested. This methodology could also be used to test the robustness of the methodology, in the sense that eliminating one observation should not drastically alter the results of the PDF estimation. For an additional metric of comparison, one could use the time series of subsequent observed Eurodollar interest rates to see how well the models predict out of sample.

All in all, the paper provides useful information to the practitioner and researcher alike on the various strengths and weaknesses in six possible ways to fit the cross-section of observed options prices. This paper also helps lay the groundwork for future research on whether any of these methodologies are useful in forecasting future changes in asset prices.

Discussion of
"The Information Content of Interest Rate Futures Options"
by Des McManus, Bank of Canada
Discussant: Roberto Violi, Banca d'Italia, Research Department

The paper applies and compares several common methods to arrive at estimates of the (risk-neutral) probability distribution of future values of Eurodollar rates; its main aim is to determine which method most accurately prices observed market options. As a concrete example of application, the days surrounding some recent FOMC meetings are examined and estimated PDFs are used as indicator of market sentiment to gauge the uncertainty over the future levels of the Eurodollar rates.

While the results of the paper indicate that further research needs to be conducted, several broad conclusions are drawn from the estimates:

- The mixture of lognormals estimation method ranked first in pricing accurately observed market options;
- The Hermite polynomial estimation method ranked second, yielding similar results to the mixture of lognormals; when the latter method fails, the Hermite method constitutes an appropriate alternative.
- Estimates of skewness and kurtosis, unlike the volatility measure, are found to be unstable across estimation techniques, hence model dependent.

The author provides a very articulated review of the estimation techniques and illustrates several summary statistics, suggesting both graphical and analytical methods, in comparing various models.

Two main issues are left open for future research:

- the appropriate measure of asymmetry for estimated PDFs (over and above those which can be constructed out of the more stable cumulative distribution function: like inter-quartile differences);
- evaluation of the predictive power of the measures of asymmetry.

My comments will be mainly concentrated on the issues left open for future research. The basic motivation behind the PDFs' identification and estimation techniques can be found in the weaknesses of Black-Scholes' (BS) theory of option pricing, namely in the divergence between observed options market price and BS-based theoretical options price. The most common bias is the familiar "smile effect" and other "moneyness" biases are fairly well-known. These biases are understood to be related to measures of asymmetry, for example excess kurtosis. Perhaps less understood biases, I believe, relate to the maturity of the option (see Backus, Foresi, Li and Wu, 1997): the upward slope (on average) of the term structure of implied volatility; the excess kurtosis tendency to decline (on average) with maturity. Both biases have been found in foreign exchange currency options, but clear evidence of a term structure of volatility and kurtosis has also been detected for interest rates. It can be shown, in a suitably defined theoretical environment, that both biases eventually decline with maturity. Hence, for long options, the BS formula can be a good approximation of observed option market price in the foreign exchange.

The tendency for kurtosis to decline with maturity in many models is a consequence of a stronger result: the central limit theorem; this statement clearly doesn’t apply to all theoretical environment (the unit-root volatility model is an obvious counterexample). Similarly, a mean-reverting process of the underlying rate can be a source of changing volatility at different time horizons (as, for example, in Vasicek, 1977). Stochastic volatility models can also be a way to capture leptokurtosis, as in Heston (1993), or simpler time-varying volatility models, such as ARCH process; according to Das and Sundaram (1997) the degree of conditional leptokurtosis for interest rates varies with the time interval between data observation.

To summarise, jump-diffusion models with mean-reverting short rate seem to allow for parameter choices which match conditional skewness and kurtosis at varying maturities of the term structure of interest rates. When jumps are introduced into a pure diffusion model of interest rate, (Gaussian) volatility drops sharply with respect to its prior level and coefficient of mean reversion for the short rate decline substantially (see Das, 1998). This may imply that jumps account for a substantial component of volatility and provide a source of mean reversion.

Testing these propositions for the option markets requires enough depth across the moneyness spectrum and across maturity, perhaps not readily found for exchange-traded options[1]. The Bates (1991) option pricing model[2], which add a jump-diffusion process to the standard Black's basic model, tested by McManus did not fare particularly well in the cross-sectional comparison[3]. Generally, it appears that the more structural model, including the semi-parametric Hermite as well as the semi-non parametric Gram-Charlier (Edgeworth) polynomial expansion technique, have difficulties in coping with the degree of asymmetry found in the data; non-structural approaches, like a mixture of lognormals[4], seem to warrant a higher degree of flexibility in matching data; yet it is less clear how we should interpret such advantage when confronting the observed option market price. One possible interpretation for the mixture of lognormals is that the return on an asset at any given time can be drawn from one, out of several (oftentimes two), normal distribution; each possible draw has a given probability to occur. The benefit of such specification is that it allows for the possibility that

---

1 An attempt along these lines, e.g. incorporating the information embedded in the term structure of implied volatility, can be found in Fornari and Violi (1998).

2 See also Malz (1996) for an application to the estimation of exchange rate realignments probability in the EMS.

3 It is perhaps interesting that Jondeau and Rockinger (1998) estimated a jump-diffusion model of foreign exchange currency option which outperformed all other models at longer maturities; they were able to provide an interpretation for it which should also apply to models of the term structure of interest rates (see also Backus, Foresi and Wu, 1997).

4 This definition is borrowed from Jondeau and Rockinger (1998).

occasionally the return is generated from a distribution with a higher variance, while maintaining the structure of normal densities, conditional on the realisation of a particular draw; this is a simple and effective way to capture fat tails in the return-generating process. Basically, the model allows a jump from one distribution to another, which is similar to the traditional jump-diffusion model allowing the possibility of a jump between an infinite number of normal distributions[5]. In the empirical application of these latter models a simple Bernoulli version is actually used: over the horizon of the option there will be at most one jump of constant size.

Models are estimated at various dates and maturities, yielding a different set of estimates for each date and maturity. This is *prima facie* in sharp contrast with the assumption of constant parameters in the underlying process; the time series of parameters so obtained may also correspond to a process of the underlying asset which might have little to do with historically observed processes. Perhaps more worryingly, estimating parameters tend to display more often than not great instability across dates and maturities. In the literature, and McManus' paper is no exception, the estimates are interpreted as being those perceived to be valid at each point in time by market participants till the expiration of the option. This is clearly inconsistent with the no-arbitrage pricing principle, upon which models are assumed to be based, since parameter changes over time are not treated accordingly within the option pricing model. It also casts serious doubts on the validity of time series information inferred from the volatility and asymmetry measures obtained from the estimated models.

I believe that progress could be made by improving the option price modelling in allowing changing moments over time. As an example of this claim, I have estimated volatility, skewness and excess kurtosis, based on the methodology presented in Fornari and Violi (1998), combining the Soederlind and Svensson (1997) mixture of lognormals model with the Jamshidian (1989) closed form solution for options on discount bonds. As an alternative to Black's model, this model assumes that a single factor model for the term structure of interest rates holds: the overnight rate evolves according to a mean-reverting Gaussian diffusion process, as in Vasicek (1977), with distinct parameters in each of the two regimes; unlike the Black's model, interest rates display a term structure of volatility, though a deterministic one. Only the mixing probabilities of regime switch are allowed to change over time, whereas the parameters of the term structure within each regime are kept constant[6]. No rationale is provided for the changes in the mixing probability; this is clearly a limitation which calls for improvement in future research. Data used are the same as in McManus' paper, comprising CME Eurodollar futures and options contracts with maturity December 1998 for the trading days from September 1, 1998 through November 30, 1998[7].

---

5 See Kon (1984) for details.

6 The TSP 4.4 command LSQ (see the TSP Reference Manual for details) is used to estimate the parameters of the option pricing model; the squared deviation between actual and theoretical option prices is the minimised objective function. No specific handling for the American early exercise premium is adopted.

7 Kindly provided by Gabriele Galati for the BIS workshop exercise on PDF estimation.

Chart 1.

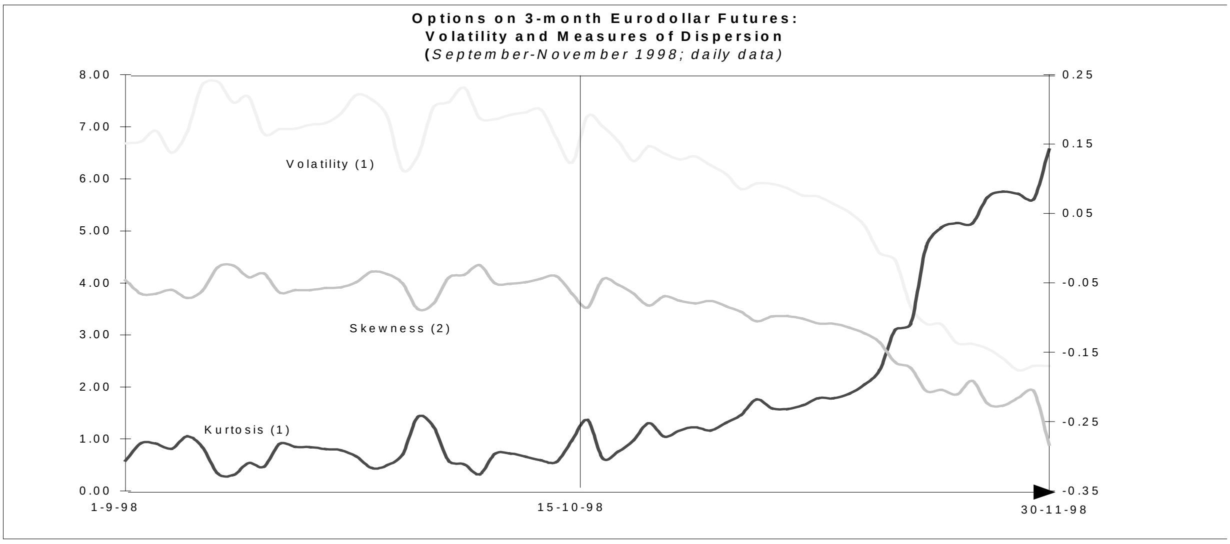

(1) Left-hand scale.
(2) Right-hand scale

Chart 1 corroborates some of the findings reported in the literature; volatility tend to decline and excess kurtosis to rise as maturity approaches; skewness changing of sign relates to the growing perception of FED monetary policy easing bias developing at the US FED in those days.

Pricing errors implied by the estimated model, though larger than the ones reckoned in McManus paper, are still within a reasonable range, as displayed in

Chart 2.

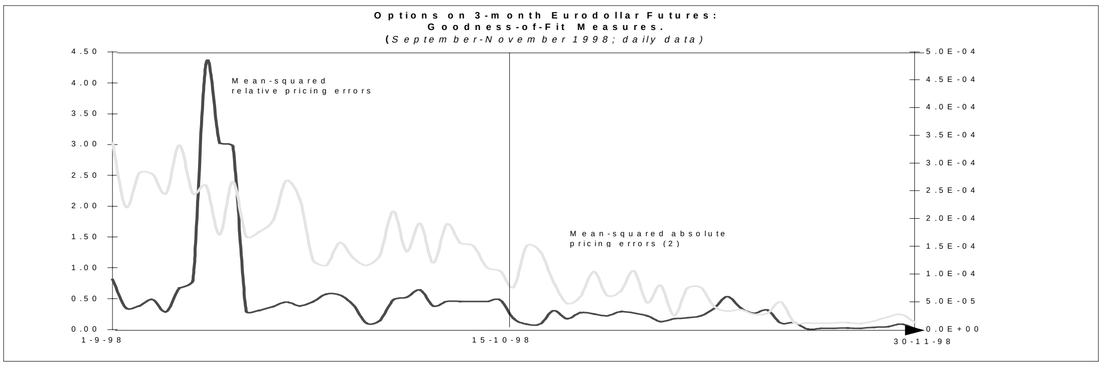

References

Backus, D, S Li Foresi and K and L Wu (1997): “Accounting for Biases in Black-Scholes”. L N Stern School of Business Working Paper S97-34, New York University, New York (NY).

Backus, D, S Foresi. and L Wu (1997): “Macroeconomic Foundations of Higher Order Moments in Bond Yields”. L N Stern School of Business, mimeo, New York University, New York (NY).

Bates, D S (1991): “The Crash of ‘87: Was it Expected? The evidence from Options Markets”. *Journal of Finance*, Vol. 46, pp. 1009-44.

Das S R(1998): “Poisson-Gaussian Processes and the Bond Market”. NBER Working Paper 6631, Cambridge (MA).

Das S R and R Sundaram (1997): “Of Smile and Smirks, A Term Structure Perspective”. NBER Working Paper 5976, Cambridge (MA).

Fornari F and R Violi (1998): “The Probability Density Function of Interest Rates implied in the Prices of Option”. Banca d'Italia, Temi di Discussione, n.339, Rome.

Heston S (1993): “A Closed Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Option”. *Review of Financial Studies*, Vol. 6, pp 327-44.

Jamshidian ,F.(1989): “An Exact Bond Option Formula”. Journal of Finance, Vol. 44, pp. 205-209.

Jondeau E and M Rockinger (1998): “Reading the Smile: The Message Conveyed by Methods which Infer Risk Neutral Densities”. CEPR, Working Paper 2009, London (UK).

Kon, S (1984): “Model of Stock Return: A Comparison”. *Journal of Finance*, Vol. 39, pp. 147-65.

Malz, A (1996): “Using Option Prices to Estimate Realignment Probabilities in the European Monetary System: The Case of Sterling-Mark”. *Journal of International Economics*, Vol. 15, pp.717-48.

Soederlind P and L Svensson (1997): “New Technique to Extract Market Expectations from Financial Instruments”. *Journal of Monetary Economics*, Vol. 40, pp.383-429.

Vasicek, O (1977): “An Equilibrium Characterization of the Term Structure”. *Journal of Financial Economics*, Vol.5, pp. 177-88.

Gordon Gemmill and Apostolos Saflekos

"How useful are implied distributions? Evidence from stock-index options"

Discussants:

Shigenori Shiratsuka

Raf Wouters

How Useful are Implied Distributions?

Evidence from Stock-Index Options

by

Gordon Gemmill and Apostolos Saflekos*

Abstract

Option prices can be used to construct implied (risk-neutral) distributions, but it remains to be proven whether these are useful either in relation to forecasting subsequent market movements or in revealing investor sentiment. We estimate the implied distribution as a mixture of two lognormals and then test its one-day-ahead forecasting performance, using 1987-97 data on LIFFE's FTSE-100 index options. We find that the two-lognormal method is much better than the one-lognormal (Black/Scholes) approach at fitting observed option prices, but it is only marginally better at predicting out-of-sample prices. A closer analysis of four "crash" periods confirms that the shape of the implied distribution does not anticipate such events but merely reflects their passing. Similarly, during three British elections the implied distributions take on interesting shapes but these are not closely related to prior information about the likely outcomes. In short, while we cannot reject the hypothesis that implied distributions reflect market sentiment, we find that sentiment (thus measured) has little forecasting ability.

Keywords: option pricing, implied distribution, volatility smile, market sentiment, crashes, elections.

* The authors are grateful for comments from Robert Bliss and Paul Dawson.

Introduction

Investors, risk-managers and policy-makers all need to forecast the probability distribution of prices if they are to take rational decisions. Conventionally, an estimate of the variance is obtained from recent data on returns. A month of data may give a reasonable estimate of the variance, but observations over several months are required if the skewness and kurtosis are to be measured accurately. Another approach is to use options data to construct implied distributions. These are the so-called risk-neutral distributions (RNDs) which traders are using when they set the prices of the options and which relate to the period until an option expires. One day's options can reveal not only the forecast variance, but also the whole shape of the risk-neutral distribution (including skewness and kurtosis). Using options is therefore an efficient way in which to forecast the whole distribution.[1]

Several alternative methods have been suggested for extracting the risk-neutral distribution from option prices, the main difference between them being the extent to which they constrain its shape. At one extreme, Longstaff (1995) imposes no constraints but the result can be a rather "badly-behaved" or spiky distribution. At the other extreme, Rubinstein (1994) and Jackwerth and Rubinstein (1996) constrain their distributions to be those with the smallest possible deviations from the lognormal. Somewhere in between these two extremes is the assumption of this paper, which is based on the work of Ritchey (1990), Melick and Thomas (1997) and Bahra (1997). We assume that the distribution can take any shape which may be approximated by a mixture of two lognormals.

The first objective of this study is to examine whether an option pricing model, based upon two lognormal distributions, performs well for equity-index options (having previously been applied only to oil futures and interest-rates). The performance of the method is determined not only by measuring the (ex-post) fit of the implied distribution, but also ex-ante by testing how well it forecasts option prices out-of-sample. We find for LIFFE's FTSE-100 index options over the 1987-97 period that although the model fits the data significantly better than the Black/Scholes model, the out-of-sample performance is only marginally better. This is consistent with work on US index options by Dumas, Fleming and Whaley (1998), who found that taking account of volatility smiles did not help in forecasting one-day-ahead option prices.[2]

If implied distributions are of rather limited use in normal periods, it might still be possible that they help to forecast market movements during exceptionally turbulent periods. Our second objective is therefore to examine the performance of the method around crashes (of October 1987, October 1989 and October 1997), British general elections (of May 1987, April 1992 and May 1997), and

[1] Of course, it is important to remember that these implied distributions are different from directly sampled distributions, as they reflect risk-neutral processes (see Harrison and Kreps, 1979). In the absence of transactions costs the shapes should be the same, but the location of the implied distribution reflects only a risk-free rate of return.

[2] There is a one-to-one relationship between the volatility smile and the implied distribution, as demonstrated explicitly by Shimko (1994), so forecasting with the volatility smile is equivalent to forecasting with the implied distribution.

extraordinary events (the sterling currency crisis of September 1992). Our chosen periods are more general than those examined (in a related way) by other researchers. For example, Bates (1991) and Gemmill (1996) have examined volatility smiles for the US and UK around the 1987 crash and Malz (1997), Campa and Chang (1995) have examined options on sterling in the ERM crisis of 1992. Coutant, Jondeau and Rockinger (1998) have examined implied distributions for interest rates at the time of the snap general election in France in 1997. We find that the method does not help to reveal the probability of crashes, because increased left-skewness **follows** rather then preceeds these events. Nevertheless, the method can help to reveal the divergent expectations which arise immediately after crashes and during election campaigns. In other words, the method helps to reveal “market sentiment”, which could be useful for the policy-stance of a central bank (e.g. Federal Reserve, Bank of England) and for investors who may wish to take positions based upon the difference between their forecast of the distribution and the consensus of the market.

The paper is written as follows. The next section provides a theoretical presentation of the main techniques used to derive the implied distribution of future asset prices from option prices. The two-lognormal mixture distribution method is described, as well as the practical issues of the implementation (namely the data used and the selection of the studied periods). Section 3 presents the empirical results from applying the method to FTSE-100 index options, at the same time assessing its usefulness. Conclusions and suggestions for further research are given in the fourth and final section of the paper.

2. Theoretical framework

2.1 Review of literature and development of the model

Option prices reflect forward-looking distributions of asset prices. In the absence of market frictions, it is possible to take a set of option prices, for a single maturity and at various exercise prices, and imply the underlying risk-neutral distribution (RND). Breeden and Litzenberger (1978) first showed how the second partial derivative of the call-pricing function with respect to the exercise price is directly proportional to the RND function. However, since observed option prices are only available at discretely spaced intervals rather than being continuous, some approximation for the second derivative is necessary and more than one implied distribution can be implied depending on the approximation chosen. As Jackwerth and Rubinstein (1996) observe, selecting among the competing distributions then amounts to a choice of how to interpolate or extrapolate option prices across exercise prices.

The most direct way of estimating the implied distribution is by simple application of the Breeden and Litzenberger result to a function relating the call price to exercise prices. This has been done by Longstaff (1995) and Ait-Sahalia and Lo (1995). The former implements a procedure that attributes a probability to an option mid-way between two adjacent exercise prices, then uses this to solve for the next probability, and so on. Ait-Sahalia and Lo first smooth the pricing function with a set of polynomials and then proceed in a similar way.

Shimko (1993, 1994) proposes an alternative approach by interpolating in the implied-volatility domain instead of the call-price domain. He begins by fitting a quadratic relationship between implied volatility and exercise price. The Black/Scholes formula is then used to invert the smoothed volatilities into option prices. At this point he has a continuous spectrum of call prices as a function of the exercise prices and the application of the Breeden and Litzenberger result is straightforward, generating the implied probability distribution.

The main limitation of the above techniques is the need for a relatively wide range of exercise prices. This can be overcome by imposing some form of prior structure on the problem. One such prior (used by Bates (1991, 1996) and Malz (1997)) is to assume a particular stochastic process for the price dynamics of the underlying asset. In their papers the asset price is assumed to follow a jump-diffusion process. In other words, the basic probability distribution is lognormal, but it can jump up or down[3].

Alternatively, the imposed structure may apply to the distribution of the future asset price itself, instead of the asset-price dynamics. This approach proves to be more general than making assumptions about the stochastic process of the underlying asset price, because any given RND function is consistent with many different stochastic processes, whereas a given stochastic price process implies a unique RND function (Melick and Thomas, 1997). The approach of Rubinstein (1994) and Jackwerth and Rubinstein (1996) falls into this category. They employ an optimisation algorithm to find that RND function which is closest to lognormal, taking account of bid/ask bounds on the observed option prices.[4]

The framework used in the current paper follows Ritchey (1990), who notes that a wide variety of shapes may be approximated with a mixture of lognormal distributions. He assumes that the implied density function, $f(S_T)$, of the underlying asset terminal price, S_T, comprises a weighted sum of k individual lognormal density functions:

$$f(S_T) = \sum_{i=1}^{k} \left[\theta_i L(a_i, b_i, S_T) \right] \qquad \textbf{Equation 1}$$

where $L(\alpha_i, b_i, S_T)$ is the i^{th} lognormal density function with parameters α_i, b_i:

$$L(a_i, b_i, S_T) = \frac{1}{S_T b_i \sqrt{2\pi}} e^{\left\{-(\ln S_T - a_i)^2\right\}/2 b_i^2} \qquad \textbf{Equation 2}$$

---

3 Malz assumes that there is either no jump or just one jump over the life of the option, in which case the terminal RND function is a mixture of two lognormal distributions (Bahra, 1997)

4 The distance criteria used in the two papers are, respectively, a quadratic difference and a smoothness function.

$$a_i = \ln S + (\mu_i - \sigma_i^2 / 2)\tau \text{ and } b_i = \sigma_i \sqrt{\tau} \qquad \textbf{Equation 3}$$

In the above equations, S is the spot price of the underlying asset, τ (=T-t) is the time remaining to maturity and μ and σ are the parameters of the normal RND function of the underlying returns. The weights θ_i are positive and sum to unity.

Melick and Thomas (1997) apply this framework to options on crude oil futures, using a mixture of three lognormal distributions. Bahra (1997), Butler and Davies (1998) and Soderlund and Svensson (1997) use a mixture of two lognormals on interest-rate futures. Since our data on FTSE-100 options cover a limited range of exercise prices for each maturity, it seems more appropriate to use two lognormals, which require only five parameters: the mean of each lognormal, α_1, α_2, the standard deviation of each lognormal, b_1, b_2 and the weighting coefficient, θ.

2.2 The two-lognormal mixture distribution method for equity index options

Let the terminal pay-off on a European call maturing at time T be $\max(S_T - X, 0)$, given terminal asset price S_T and exercise price X. Assuming that the risk-free interest rate r is constant, the life of the option is τ and the asset price is S, then the price of the call is the discounted expected payoff (conditional upon finishing in the money) times the probability of finishing in the money:

$$c(X, \tau) = e^{-r\tau} \int_X^{\infty} f(S_T)(S_T - X)\, dS_T \qquad \textbf{Equation 4}$$

where $f(S_T)$ is the risk-neutral probability density function of the terminal asset price at time T. Similarly, the terminal payoff on a European put is $\max(X - S_T, 0)$ and its current price is:

$$p(X, \tau) = e^{-r\tau} \int_0^{X} f(S_T)(X - S_T)\, dS_T \qquad \textbf{Equation 5}$$

Under the assumption that the probability density function is a mixture of two lognormals (with weights θ and $(1-\theta)$), the above equations for call and put prices can be rewritten as:

$$c(X,\tau) = e^{-r\tau} \int_X^{\infty} \left[\theta L(a_1, b_1, S_T) + (1-\theta) L(a_2, b_2, S_T)\right](S_T - X)\, dS_T \quad \textbf{Equation 6}$$

$$p(X,\tau) = e^{-r\tau} \int_0^{X} \left[\theta L(a_1, b_1, S_T) + (1-\theta) L(a_2, b_2, S_T)\right](X - S_T)\, dS_T \quad \textbf{Equation 7}$$

These equations can be used iteratively to minimise the deviation of estimated prices from observed prices, a search being made over the five parameters. We use both puts and calls across five exercise prices and minimise the total sum of squared errors for the ten options:

$$\sum_{i=1}^{n}\left[c(X_i,\tau)-\hat{c}_i\right]^2+\sum_{i=1}^{n}\left[p(X_i,\tau)-\hat{p}_i\right]^2 \quad \textbf{Equation 8}$$

where b_1, $b_2 > 0$, $0 \le \theta \le 1$, subscript i denotes an observation and ^ denotes an estimate.

Bahra (1997) shows that equations 6 and 7 have the following closed-form solutions:

$$c(X,\tau) = e^{-r\tau}\{\theta\left[e^{a_1+b_1^2/2}N(d_1)-XN(d_2)\right]+$$

$$(1-\theta)\left[e^{a_2+b_2^2/2}N(d_3)-XN(d_4)\right]\} \quad \textbf{Equation 9}$$

and

$$p(X,\tau) = e^{-r\tau}\{\theta\left[-e^{a_1+b_1^2/2}N(-d_1)-XN(-d_2)\right]+$$

$$(1-\theta)\left[-e^{a_2+b_2^2/2}N(-d_3)-XN(-d_4)\right]\} \quad \textbf{Equation 10}$$

where

$$d_1 = \frac{-\ln X + a_1 + b_1^2}{b_1}, \quad d_2 = d_1 - b_1$$

$$d_3 = \frac{-\ln X + a_2 + b_2^2}{b_2}, \quad d_4 = d_3 - b_2$$

Equations 9 and 10 have a very simple interpretation: the model prices are just weighted sums of two Black/Scholes solutions, each having its own mean and variance.

2.3 Data sources

The empirical research in this study is based on the daily settlement prices of FTSE-100 calls and puts covering up to five exercise prices and four maturities for each day from January 1st 1987 to December 31st 1997[5]. Our sample contains American-style options for the period to March 1992 and European-style options thereafter, the switch being made because the latter were only thinly traded

[5] We gratefully acknowledge LIFFE for financial assistance in collecting some of the data and for providing the other data.

before then.[6] Exercise prices have been chosen such that one is at-the-money, two are in-the-money and two are out-of-the-money. The interest rates used are UK Sterling 3-month interbank deposit rates, retrieved from Datastream.

2.4 Hypotheses to be tested

Hypothesis 1: ***The two-lognormal model performs better than a one-lognormal (Black/Scholes) model***

This paper is primarily a critical examination of the two-lognormal model. On one day per month for the period 1/87 to 11/97 the distribution is implied and then used to price options on the next day. This allows us to test the method's forecasting performance relative to the Black/Scholes model over quite a long period.[7]

Hypothesis 2: ***The option market anticipates crashes***

One particularly interesting question is whether option markets have any usefulness in predicting extreme events, such as stock market crashes. The crash of October 1987, the mini crash of October 1989 and the market turmoil of October 1997 were chosen as examples of such occasional events. Although other studies have looked at some of these periods (e.g. Bates, 1991, for the US and Gemmill, 1996, for the UK), the use of the two-lognormal mixture is original. We also examine the European monetary crisis of September 1992 (when sterling left the Exchange Rate Mechanism (ERM) and was devalued by more than 10%) as another period of great uncertainty for the British stockmarket.

Hypothesis 3: ***A bimodal distribution is appropriate during elections***

The two-lognormal mixture may prove particularly useful in periods when a market jump is expected but the direction of the jump is unknown. Such is the case during political elections. Assuming that there are only two possible outcomes (for example, Labour victory or Conservative victory) and that investors prefer one to the other, a stock-index option which matures after an election should reflect a bimodal underlying distribution. The mixture of two separate lognormal distributions should therefore fit the observed option prices particularly well at such times. We have included the British elections of 1987, 1992 and 1997 in our analysis in order to test this hypothesis.

---

[6] Strictly speaking the method is only applicable to European options, because we are attributing all of an option's value to the terminal distribution and not to early exercise. However, the value of early exercise on these index options is likely to be small: see Dawson (1994) for an analysis on FTSE-100 options.

[7] Dumas, Fleming and Whaley (1998) use a similar approach on S&P 500 options, but based upon the volatility smile rather than the implied distribution.

3. Results

3.1 Forecasting Performance over the Whole Period, January 1987-December 1997

The options used in this part of the analysis are chosen on one day per month (from the middle of the week) such that they have approximately 45 days to maturity. Table 1 gives a summary of the conventional dispersion and shape statistics of the implied distribution over the whole period. The results leave no doubt about two features. First, the implied distributions have fatter tails than those of lognormal distributions, with kurtosis positive in each subperiod and averaging 1.54. This result is expected, as it is the corollary of the well-known volatility smile which is found for many different options (e.g. Bahra, 1997, on interest rates, Melick and Thomas, 1997, on crude oil futures, Malz, 1996, on foreign exchange).

Second, and more importantly, the generated distributions exhibit consistently negative skewness (that is, they have a more pronounced tail to the left) averaging -0.26 over the whole period and becoming more pronounced over time. Figure 1 plots the monthly results, which indicate that after March 1991 there is no month in which skewness is positive, although variation is quite large. This result differentiates equity-index options from options on other underlying assets and is also well documented from volatility smiles (see e.g. Gemmill, 1996).

The performance of the two-lognormal model can be compared with Black/Scholes in two ways; first, in how well (ex-post) the two models fit observed option prices within sample, and second, in forecasting (ex-ante) the price of an option on day t+1. To do the latter we obtain the model parameters (a_1, b_1, a_2, b_2 and θ) that best fit the option prices observed on day t. Then for forecasting we update the means of the distributions to take account of changes in the stock price.[8]

$$a_i^{t+1} = a_i^t + \ln \frac{S^{t+1}}{S^t} \qquad \textbf{Equation 11}$$

Similarly, we update the variances to take account of one day's less time to maturity:

$$b_i^{t+1} = b_i^t \sqrt{\frac{\tau^{t+1}}{\tau^t}} \qquad \textbf{Equation 12}$$

Table 2 shows the errors obtained by the two-lognormal and Black/Scholes methods, both within sample (ex-post) and out-of-sample (ex-ante). Results are given separately for 1/87 to 2/92, for which

[8] Strictly speaking, the adjustment should take account of the change in the forward price, for which the change in spot price is a good approximation unless an ex-dividend date is straddled (which does not occur in our sample).

American options were used, and for 3/92 to 11/97, for which European options were used. The two-lognormal method has an in-sample performance which is considerably better than that of the Black/Scholes model, being 29% better in terms of root-mean-squared error for the American options in the earlier period and 89% better for the European options in the more recent period. This is to be expected since it uses five parameters (a_1, b_1, a_2, b_2, θ) as compared with the two parameters (a, b) of Black/Scholes.[9] The out-of-sample (forecasting) test shows a root-mean-squared-error improvement of only 11% for the American options (54 out of 61 observations show an improvement), but a more impressive gain of 43% for the European options (all 68 observations show an improvement). However, these *relatively* large RMSE improvements for European options translate into *absolute* gains of about 1-2 index points per option, which are small when compared with a bid/ask spread of at least 2 points. Hence the method gives a consistent but small improvement in forecasting performance on average across the eleven year period.

3.2 Market crashes

The crash of 1987

Even if the method gives only small benefits in most periods, it may fit the data and forecast better than Black/Scholes in periods when there are significant events. On Monday, October 19th 1987 the FTSE-100 fell by 10.9% to 2052. The following day, after the news from Wall Street convinced everybody that this was a global crash, there was a further drop of 12.2%. The decline continued and the index dropped to 1684 on October 26th, 1608 on November 4th and to the 1987 low of 1565 on November 9th. This represented a fall of 32% in three weeks. The London stock market did not recover these losses for more than 18 months.

We have divided our analysis of implied distributions around this time into two distinct periods: the first is the two trading weeks immediately before the crash and the second is the month immediately after the crash. Results on shape and goodness of fit are summarised in Table 3 (first segment) and representative implied distributions are plotted in Figure 2.[10] It should be noted that in this and subsequent figures, the first distribution is plotted as estimated but the other distributions have been adjusted to give the same time to maturity as the first. Without such an adjustment there would be a narrowing of the distributions as maturity approached.[11] Prior to the crash (13th October in Figure 2)

---

[9] In fact, Black/Scholes normally has only one unknown parameter, the volatility of the distribution. However, since we do not use the forward price of the underlying asset as a parameter in the minimisation procedure but imply it as the mean of the distribution, the number of B/S parameters is two and the number of additional parameters of the two-lognormal model is therefore three.

[10] Because averaging across distributions tends to remove deviations from lognormality, we have plotted the distribution on that particular day which has skewness nearest to the period average.

[11] The adjustment to the mean takes account of the change in spot price and the size of the contango (forward price less spot price), treating each distribution separately. We have $\alpha_I^* = \ln S + (\alpha_I - \ln S)(t_1 / t_2)$, where α_I^* is the adjusted value for the ith mean, t_1 is the maturity observed and t_2 is the maturity required for purposes of comparison. Each variance is adjusted as in Equation 12.

the distribution is unimodal, with fat tails and slightly positive skewness. The means of the two component distributions are close together ($a_1 \approx a_2$), but there is a large difference in their standard deviations (b_1, b_2) which generates the fat tails. Option prices on the three trading days from October 19th to October 21st have been excluded, because both methods lead to huge errors when fitted.[12] The two lognormal distributions thereafter move apart and, on some days, give a bimodal composite, as illustrated for November 6th in the middle segment of Figure 2. The representative distribution for the whole month after the crash, as shown by November 26th in Figure 2, is not bimodal but does show quite widely separate means for the two component distributions ($a_1 \neq a_2$). In this period the mode of the composite function exhibits great instability, jumping regularly between 1300 and 1600 which reflects the difficulty which investors had in reaching a new consensus.

From before to after the crash, average volatility jumps from 20.8% to 50.7%[13] and skewness falls from positive (0.36) to negative (-0.26). However, kurtosis falls compared to the pre-crash period (from 1.74 to 0.02). Finally, it is interesting to note that the mean of the implied distribution is below the spot price for much of this period, consistent with the observed backwardation in the futures market.

The mini crash of 1989

On October 16th 1989, two years after the 1987 crash, the FTSE-100 dropped by 74 points (3.3%) to 2163. This was the most dramatic plunge of the index for more than a year, but this time the market reacted in a more muted way than in 1987, avoiding panic and quickly recovering its prior level. Our analysis is divided into a before-crash period of September 11th to October 15th and an after-crash period of October 16th to November 10th.[14] There are only slight changes in the implied distributions from the first to the second period (see Table 3, second segment, and Figure 3), which is in stark contrast to what happened in 1987. Volatility increases from 19.3% to 27.8%, but that is a small change relative to events in October 1987. Left-skewness increases from -0.26 to -0.35, but there is a reduction in kurtosis just as in 1987 (from 1.74 to 0.02). In both the pre-crash and post-crash periods the two component distributions move closely together, preserving the unimodal nature of the composite RND function.

The ERM crisis of 1992

The period of our study is mid-August to mid-October 1992. During the second and third quarter of 1992 a number of European currencies – including sterling – were subject to strong pressures which

---

An alternative to such an approach would have been to do the analysis with options of two maturities and then synthesise an implied distribution for a constant maturity. This is discussed by Butler and Davies (1998).

12 Both methods give large fitting errors from October 22nd to October 28th (daily MSE in the range 3 to 8), but these results are included as the period is particularly important.

13 Volatility is measured empirically by integrating across the composite distribution.

14 November index options were used for the analysis.

eventually pushed them to the ERM floor. The British government tried to hold sterling's value against the D-Mark through a series of interest rates increases. As rates rose, so the stockmarket fell by about 15% from early May to late August. The last effort to defend sterling was an unprecedented 5 percentage point rise in interest rates, announced in two steps, on September 16th 1992. Before the end of that day sterling had left the ERM and the second of the two interest rate rises did not come into effect. The devaluation of sterling and lower expected interest rates then pushed the FTSE-100 up by 7.7% over two days.

Despite the rise in share prices, the shape of the implied distribution is only mildly changed from before to after the devaluation, becoming almost bimodal (see Table 3, third segment, and Figure 4). Average volatility and kurtosis both increase slightly (volatility from 22.9% to 23.5%, kurtosis from 1.61 to 2.50). Skewness, which is extremely negative in both cases, moves from -0.60 to -0.74. It seems that exit from the ERM, an *upward* "shock" for the stockmarket, had a very muted impact on the implied distribution, which contrasts with the impact of the *downward* shocks in 1987 and 1989 which increased left-skewness and volatility. This result is consistent with time-series models of volatility (such as EGARCH) which find significant asymmetry: volatility rises by much more when the market falls than it does when the market rises (see, for example, Crouhy and Rockinger, 1993).

The Asian crash of 1997

Our analysis covers 59 trading days, from 9/9/97 to 28/11/97, and is based on options maturing in December. In the five weeks before the 20th October volatility is relatively high (19.6% – see Table 3, fourth segment). The implied distribution is highly skewed to the left (–0.81) but its kurtosis is of lognormal size (0.03). After October 20th, volatility almost doubles (34.9%) and kurtosis increases significantly (1.05). The implied distribution is still skewed to the left, but on average less so than before (-0.56). The representative plots in Figure 5 show how the distribution changes from being almost bimodal on 16th September, to being stretched and clearly bimodal on 29th October. This suggests that investors hold widely difference views about the potential level of the index at this time.

Our analysis of crashes can be summarised as follows. The largest effect is simply an asymmetric impact on volatility, which responds more to a market fall than to a market rise. Left-skewness rises hugely after the 1987 crash, but does not change much when smaller shocks occur. There is no consistent pattern in kurtosis during these events, but there is a tendency for a bimodal distribution to appear after the event. Finally, there is no pattern in the results to suggest that any of these four events was anticipated by participants in the options market.

3.3 British Elections

The 1987 election

The 1987 election was called on May 11th and held one month later on June 11th. During the whole campaign the Conservatives maintained a clear lead in the opinion polls over Labour and the stockmarket rose by 5.5%. Investors were said to be awaiting a "Japanese wall of money" which would arrive after a Conservative victory (Financial Times, May 30th 1987).

Our study covers 30 trading days from 12/5/87 to 23/6/87 and is based on the prices of the July (American-style) options. During the campaign the distribution is volatile, averaging 26.5% as given in Table 4, and quite left-skewed, averaging -0.62. The representative distribution plotted for 20th May in Figure 6 indicates a nearly bimodal distribution. Taking account of the information available at the time (see Gemmill, 1992) it is reasonable to assert that the mode at 2010 represents a Labour win while the mode at 2300 represents a Conservative win. Relative to the average futures price for this period, these represent an anticipated 3.4% rise if the Conservatives win and a 9.6% fall if Labour wins. After the election the distribution resumes its familiar shape, becoming unimodal and with a more modest dispersion (volatility=21.9%) and skewness (–0.09).

The two-lognormal model proves to be useful during the election campaign in revealing the market's sentiment, confirming that two distinct outcomes are perceived to be possible. However, that perception is itself difficult to explain, given the extremely high probability of a Conservative win which the opinion polls forecast throughout this period (see Gemmill, 1992) and the absence of any upward movement in the market after the Conservative win. The method appears to have captured market sentiment in advance of a known event, but that sentiment also appears to have been a rather misleading forecast of the election outcome.

The 1992 election

The 1992 general election was called on March 11th and held four weeks later, on April 9th. Unlike the previous election, when the Conservatives had been the strong favourites, the outcome of this election was very uncertain. The opinion polls gave Labour a narrow lead over the Conservatives, which, if confirmed on election day, might have led to a hung parliament. Therefore, the possible scenarios were three: a Labour government, a Labour-Liberal Democrat coalition and a Conservative government.

Just as in the previous election, the market's disposition in favour of Conservatives was often heard and the prospect of a change in government caused the stock market to slide. The index eventually fell to its 1992 low a few days before the election. On the day of the election, the index gained considerable ground, since a published opinion poll created "last-minute hopes" for a Conservative victory. The following day the Conservatives returned to power once again and the stock market gained 136 points (5.4%) in one session.

Our analysis covers the period from March 12th to April 30th 1992 and is based on the May (American-style) options. During the campaign volatility is 23%, skewness –0.20 and kurtosis 1.96 (see Table 4). This time we do not observe a bimodal pattern (see Figure 7), even though the polls indicate that the final winner is less clear than in 1987. After the election volatility falls (to 17%), skewness becomes less negative (–0.12) and kurtosis rises to 5.53. In sum, the Conservative win brings forth a smile rather than a sneer, but the options do not forecast that the market will rise on a Conservative victory.

The 1997 election

The election was announced on March 17th and held six weeks later, on May 1st. During the unusually long campaign, the Labour party maintained an estimated lead over the Conservatives which ranged from 28% at the start to 5% a week before voting. Unlike 1992, this time the actual outcome confirmed the predictions of the opinion polls and Tony Blair became the first Labour prime minister in 18 years.

Our analysis covers 50 trading days, from 18/3/97 to 30/5/97, and is based on the June (European-style) options. During the campaign the implied distribution is skewed to the left (–0.53, see Table 4) and kurtotic (+0.53). After the election, volatility decreases (from 15.4% to 13.0%) and both kurtosis and skewness decline (skewness=–0.45, kurtosis=0.29). The representative distributions for 24th April and 15th May in Figure 8 are very similar in shape. In sum, the 1997 election is almost a "non-event" for the index-options market.

What has been learnt from the study of election periods? In principle, they provide an ideal test of the informational content of implied distributions, since a known event is certain to occur on a specific date but its impact has to be forecast. The two-lognormal method gives much smaller root-mean-squared errors (in sample) than does the Black/Scholes model, particularly in 1997. It reveals a rather bimodal distribution in 1987, but not in 1992 or 1997 when such a distribution would have seemed more plausible given the more balanced contests. In sum, the analysis helps to "tell a story" about investors' expectations, but it is not a story which is supported by subsequent outcomes: if investors' expectations are revealed by the implied distributions during election campaigns then those expectations do not seem to have much forecasting power.

4. Conclusions

In this study we have examined whether implied distributions are informative with respect to subsequent stockmarket moves and to what extent they may be used to reveal investor sentiment. To do this we have applied the mixture-of-two-lognormals technique to London's FTSE-100 index options and critically examined the in-sample and out-of-sample performance of this model in a variety of periods.

The analysis was intended to test three general hypotheses: 1) the two-lognormal model performs better than Black/Scholes; 2) implied distributions indicate that the option market anticipates crashes; and 3) the method is particularly useful in periods when a bimodal distribution is to be expected.

We <u>accept</u> the first hypothesis (better than Black/Scholes), but with reservations. The method gives a better in-sample fit to observed option prices and its forecasting performance out-of-sample over 1987 to 1997 is also better but not by enough to be economically useful.

We <u>reject</u> hypothesis 2 (that the options market anticipates crashes). Neither before the large crash of 1987 nor before the much smaller crashes of 1989 and 1997 did the options market become more left-skewed. The upward adjustment of the stockmarket after sterling left the ERM in September 1992 was also not anticipated. Generally, we can say that the index-option market *reacts* to crucial events such as stock market crashes, it does not *predict* them.

We <u>weakly accept</u> hypothesis 3: the method does help to reveal market sentiment during elections. In particular, during the 1987 election the method allows us to reveal the development of a bimodal distribution, reflecting widely different potential outcomes. Nevertheless, while this may help in telling a "market story", it is not one which is consistent with rational expectations: in 1992 and 1997 the election outcome was much more uncertain than in 1987, but a bimodal shape failed to appear. In particular, the market rose on the unexpected Conservative win in 1992, but the options had not shown that a jump was at all likely.

In sum, implied distributions (recovered by using the two-lognormal mixture technique) provide some potential insight into stockmarket sentiment, but their forecasting performance is not markedly better than that of Black/Scholes. Similar conclusions were reached for the US market (using different methods) by Dumas, Fleming and Whaley (1998). These empirical results cast doubt on the view that the shape of the implied distribution is a rational expectation. Fundamentally, what has to be explained is why the implied distribution is so left-skewed and why its shape changes so frequently? The most plausible explanation is portfolio-insuring behaviour (see Grossman and Zhou, 1996) and that does not require implied distributions to be good forecasts: they just need to reflect recent moves in the stockmarket and particular investor preferences.

Table 1: Dispersion and shape statistics of implied distributions for the period 1987-1997

Period	Volatility (%)	skewness	kurtosis
1987-89	22.6	-0.075	1.970
1990-91	21.9	-0.064	2.710
1992-93	17.0	-0.252	1.261
1994-95	15.6	-0.353	0.417
1996-97	14.5	-0.652	1.134
All years	18.7	-0.257	1.544

Notes:

The data are for one day per month, averaged over the periods shown. Data up to (and including) 1993 are for American options, thereafter for European options. The skewness and kurtosis have been measured by deducting the appropriate values for a lognormal distribution, hence the null hypothesis for each is zero. Data run until November of 1997 only.

Table 2: Root-mean-squared errors of the two-lognormal and the Black/Scholes models

	Fitting (ex-post)		*Forecasting (ex-ante)*	
Period	*Two-lognormal*	*Black/Scholes*	*Two-lognormal*	*Black/Scholes*
1/87-2/92 (American-style options)	2.15	3.02	4.39	4.94
3/92-11/97 (European-style options)	0.28	2.46	1.83	3.22

Notes:
The root-mean-squared errors are measured in index points for the option prices. The data are for one day per month, averaged over the periods shown.

Table 3: Period average statistics on implied distributions for four crash periods

Event	Trading Dates	No. of trading days	change in spot	volatility	RMSE of B/S	RMSE of 2-lognormal	skewness relative to lognormal	kurtosis relative to lognormal
crash of 1987	1/10/87-16/10/87	12	-	0.208	2.20	2.00	0.363	1.736
	22/10/87-10/11/87	28	-28.2%	0.507	3.72	3.14	-0.263	0.020
mini-crash of 1989	11/9/89-13/10/89	25	-	0.193	3.65	2.58	-0.264	3.695
	16/10/89-10/11/89	20	-7.0%	0.278	4.33	2.19	-0.350	1.762

Event	Trading Dates	No of trading days	change in spot	volatility	RMSE of B/S	RMSE of 2-lognormal	skewness relative to lognormal	kurtosis relative to lognormal
ERM crisis of 1992	24/8/92-15/9/92	16	-	0.229	2.91	0.70	-0.601	1.615
	16/9/92-15/10/92	22	+9.0%	0.234	2.81	0.75	-0.735	2.504
Asian crash of 1997	9/9/97-17/10/97	29	-	0.196	4.13	0.26	-0.814	0.025
	20/10/97-28/11/97	30	-5.9%	0.349	5.55	0.86	-0.556	1.050

Notes:

Data up to (and including) 2/92 are for American options, thereafter for European options. The skewness and kurtosis have been measured by deducting the appropriate values for a lognormal distribution, hence the null hypothesis for each is zero.

Table 4: Period average statistics on implied distributions for three election periods

Election	Trading Dates	No of trading days	change in spot	volatility	RMSE of B/S	RMSE of 2-lognormal	skewness relative to lognormal	kurtosis relative to lognormal
1987	12/5/87-11/6/87	22	+4.4%	0.265	3.63	2.14	-0.606	0.620
	12/6/87-23/6/87	8	+4.1%	0.219	2.83	1.38	-0.090	0.893
1992	12/3/92-9/4/92	21	-4.1%	0.234	2.80	1.77	-0.200	1.960
	10/4/92-30/4/92	13	+7.6%	0.168	3.84	3.27	-0.118	5.527
1997	18/3/97-1/5/97	31	-1.9%	0.154	4.59	0.16	-0.532	0.527
	2/5/97-30/5/97	19	+7.5%	0.130	2.86	0.29	-0.445	0.295

Notes:
Data up to (and including) 2/92 are for American options, thereafter for European options. The skewness and kurtosis have been measured by deducting the appropriate values for a lognormal distribution, hence the null hypothesis for each is zero.

Figure 1 Skewness of the Implied Distributions

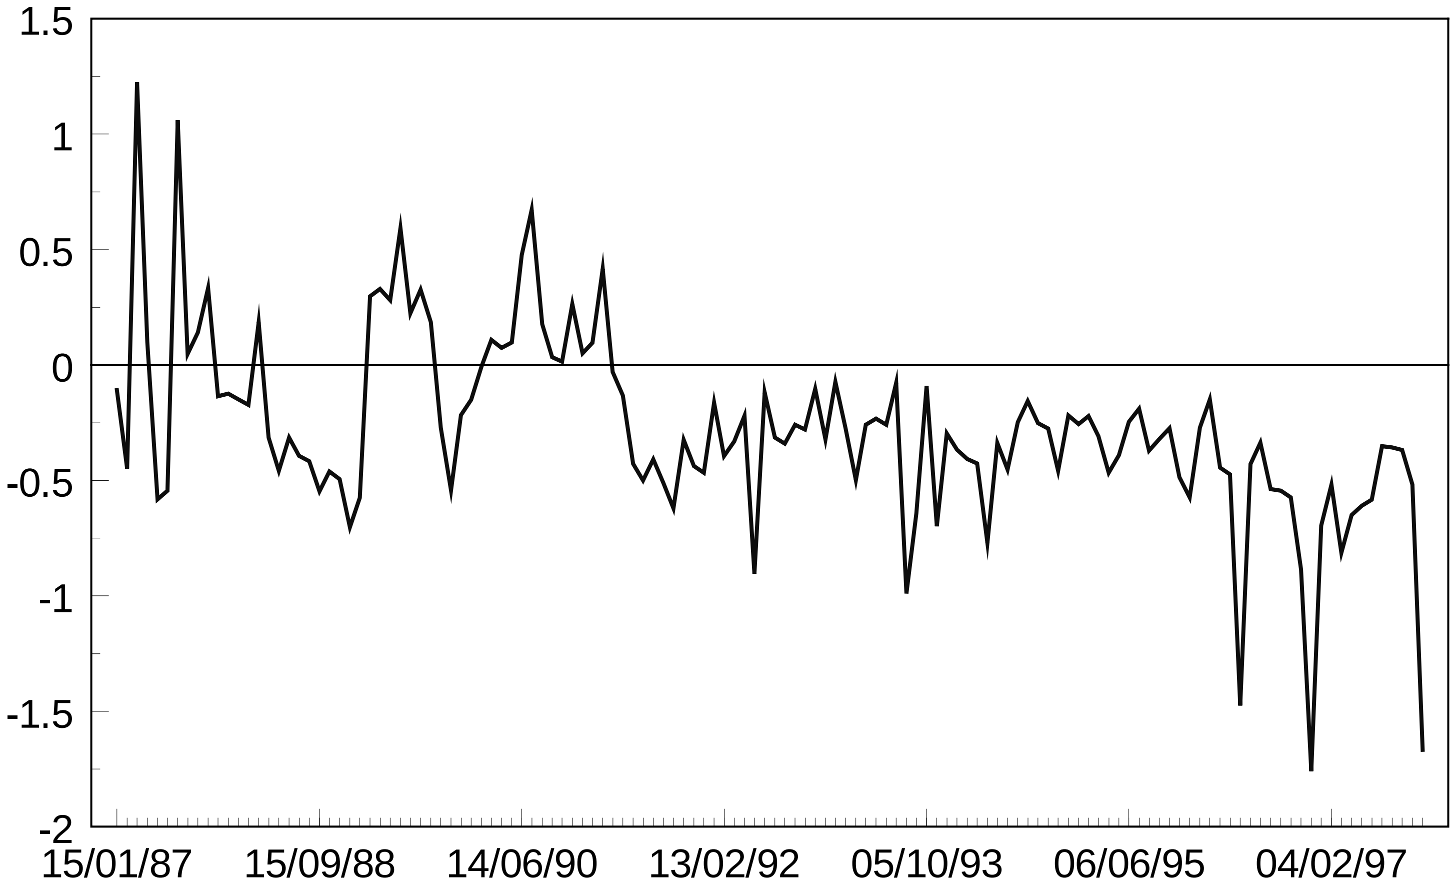

Figure 2   Implied Distributions Around the Crash of 1987

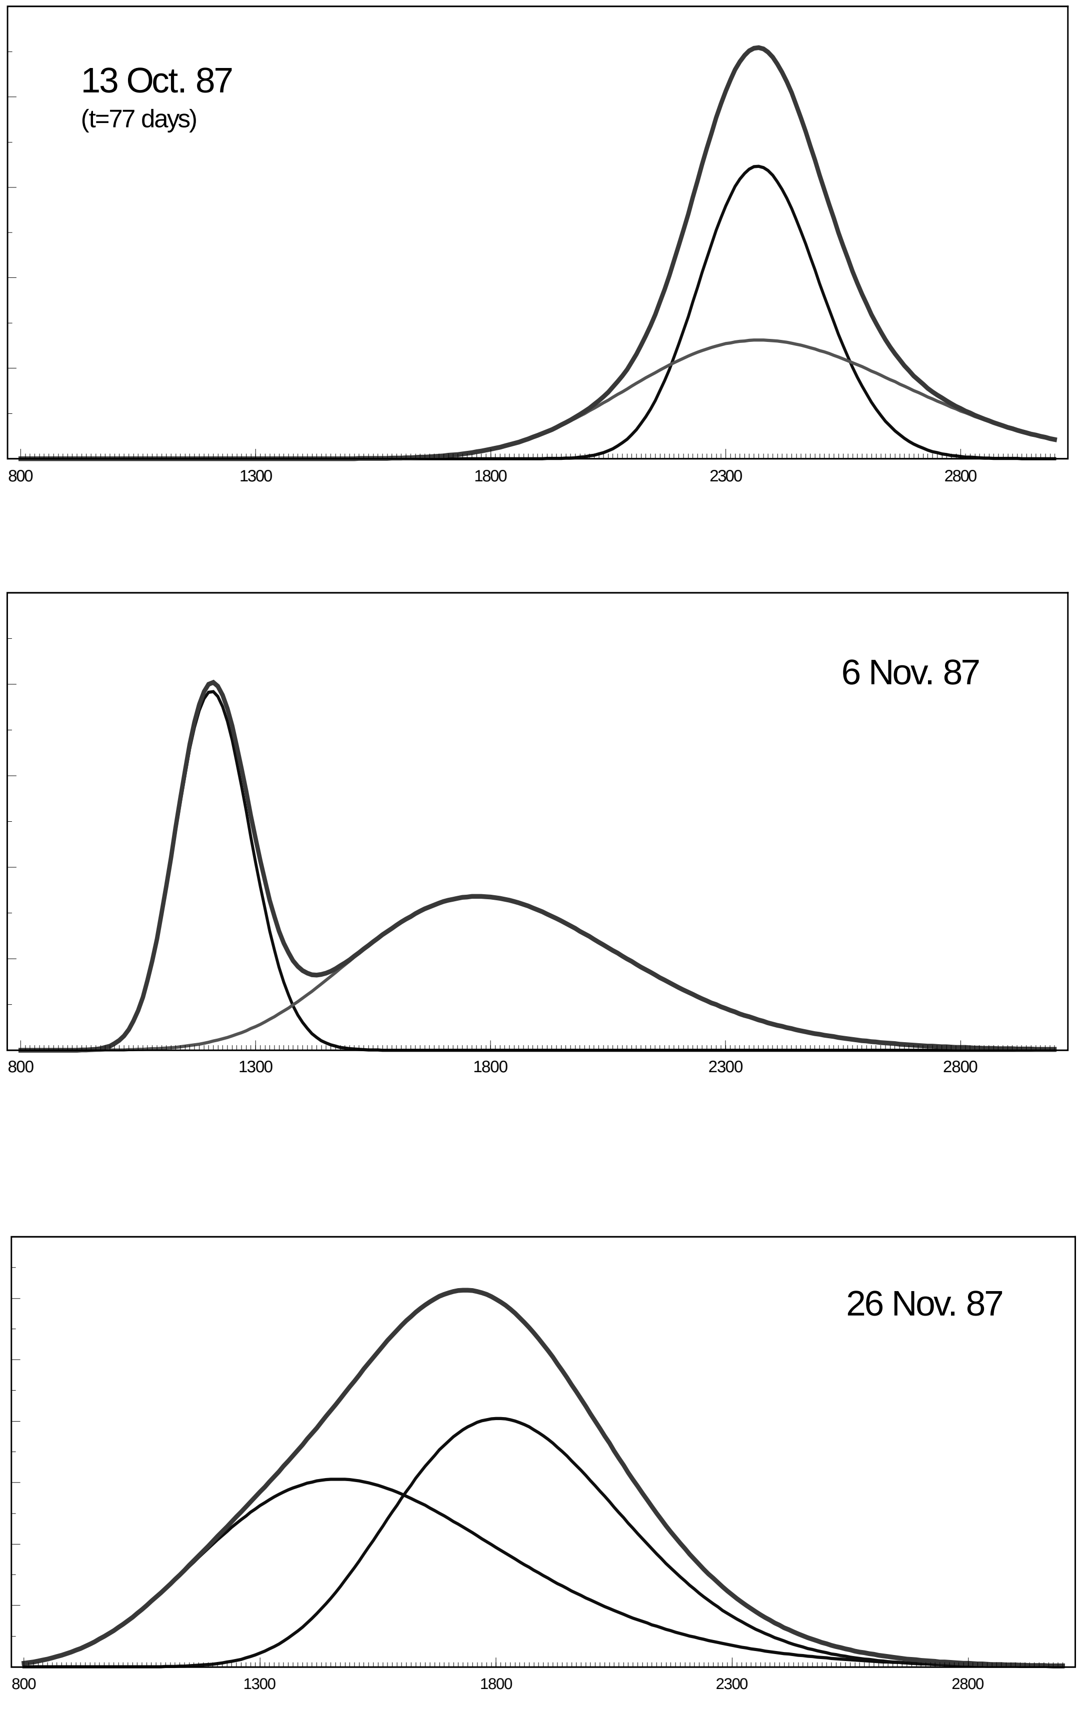

Figure 3   Implied Distributions Around the Crash of 1989

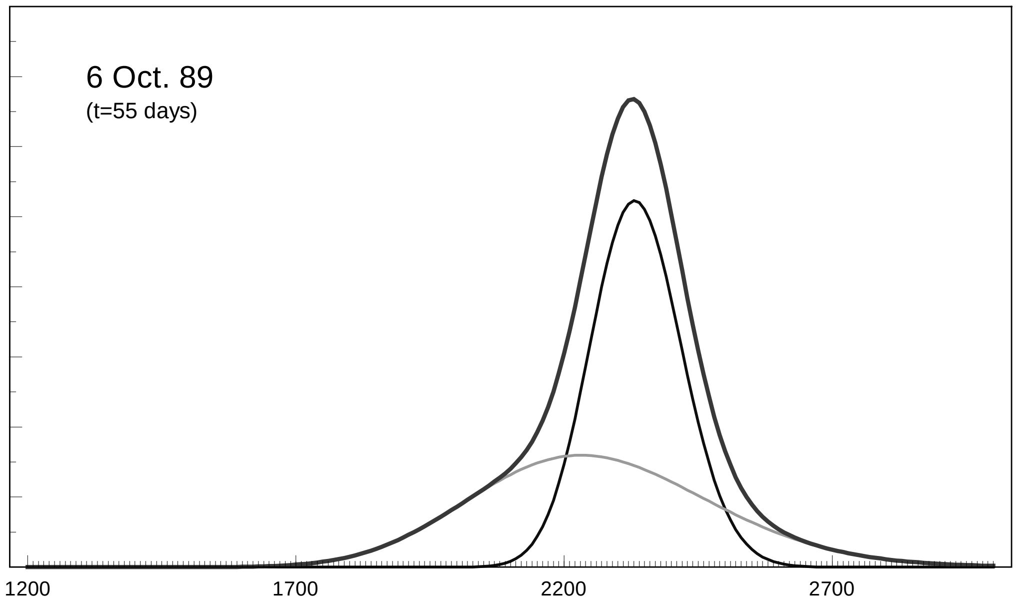

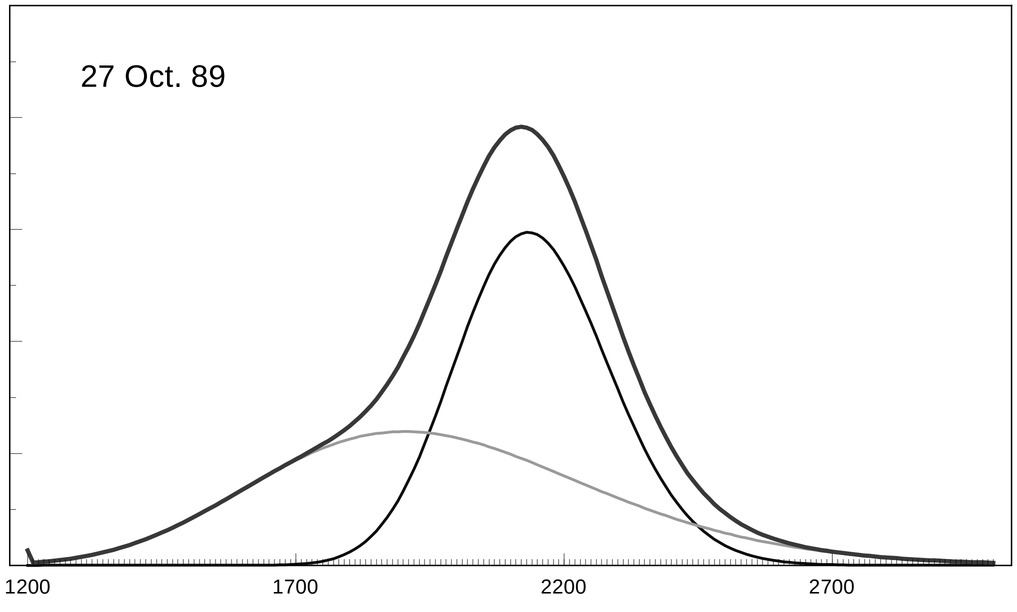

Figure 4   Implied Distributions Around the Sterling Crisis of 1992

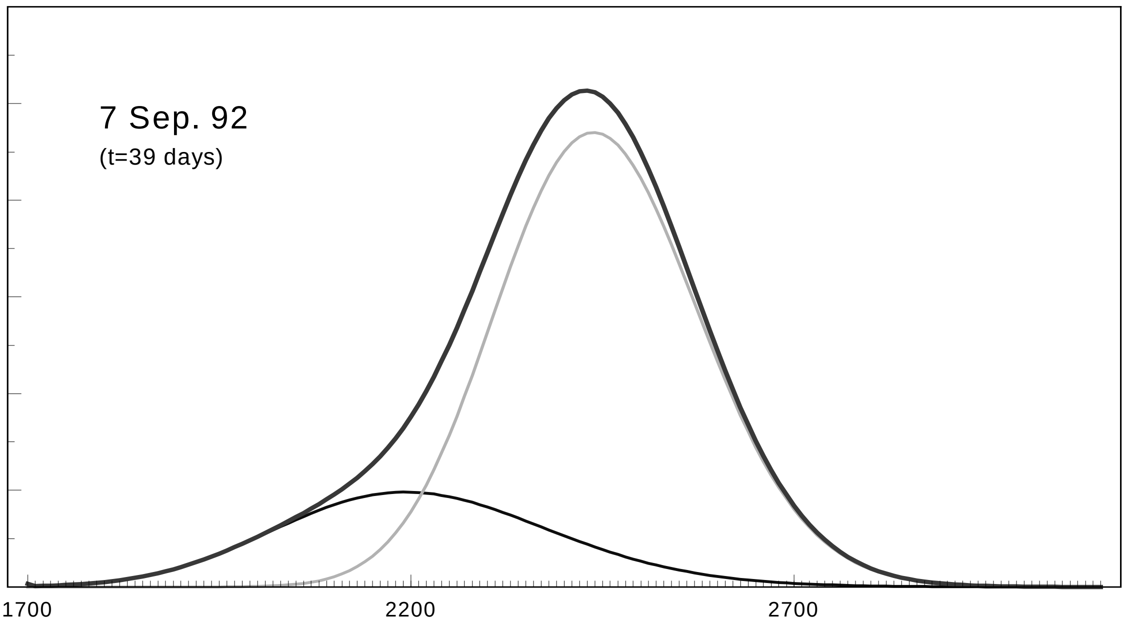

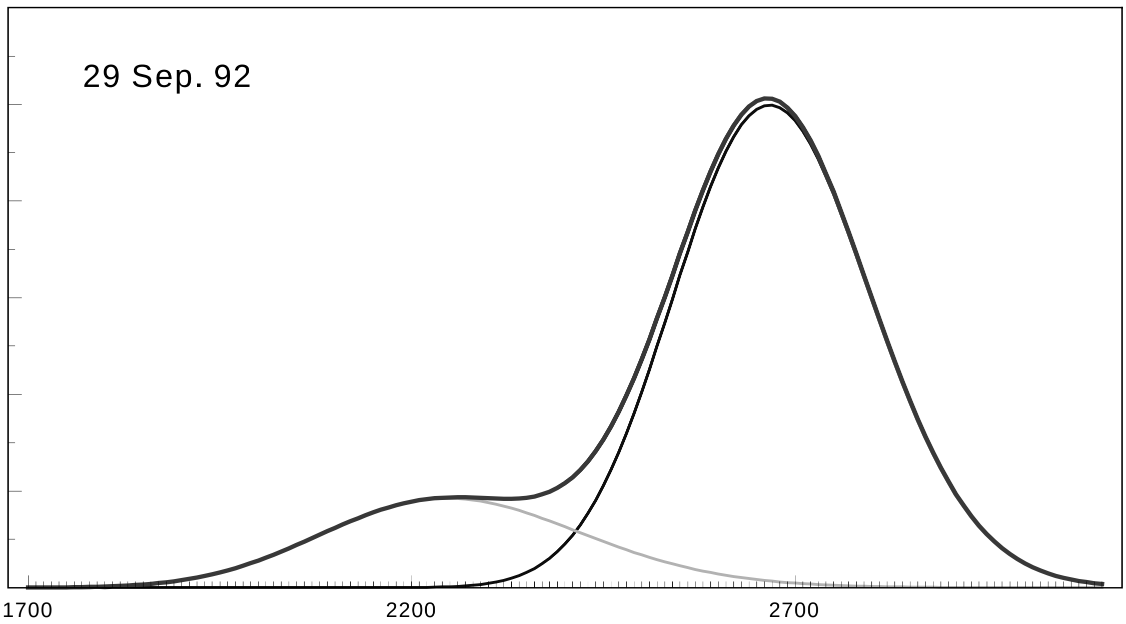

Figure 5 Implied Distributions Around the Asian Crash of 1997

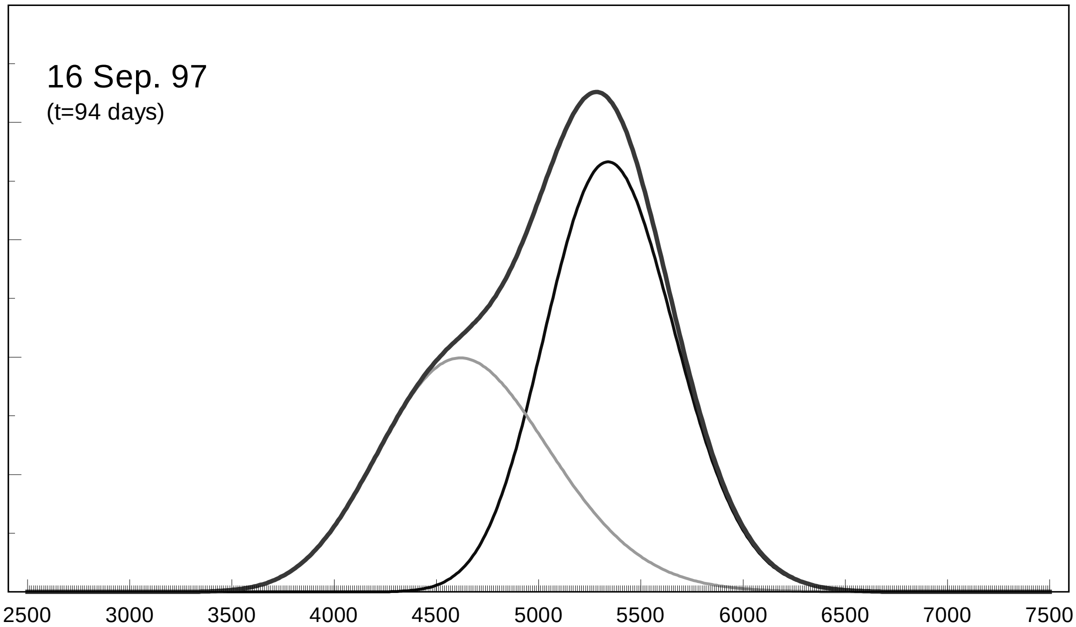

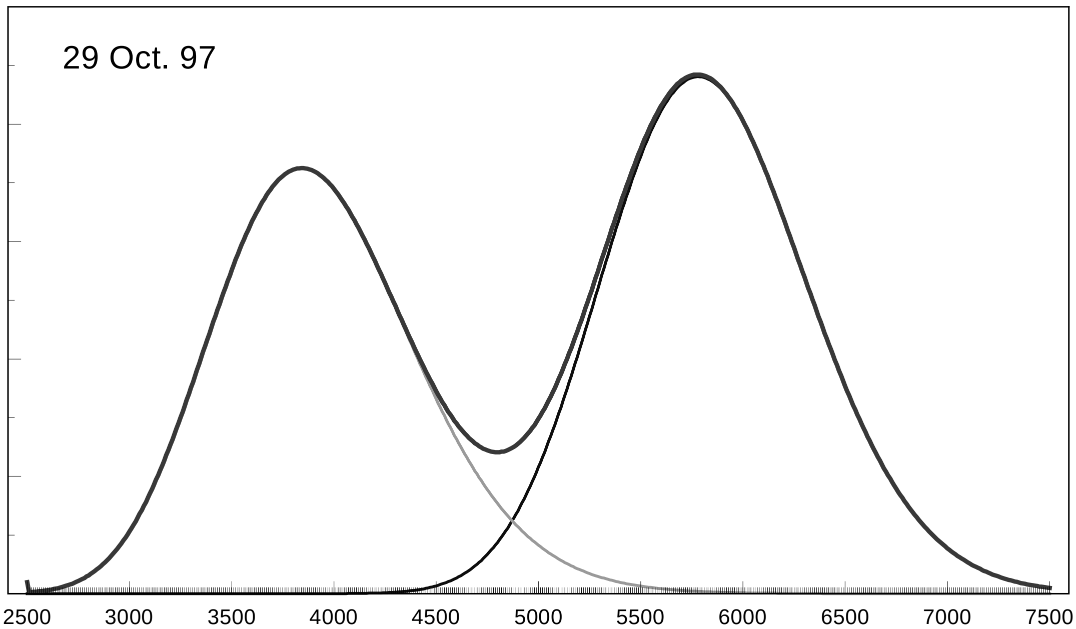

Figure 6   Implied Distributions Around the Election of 1987

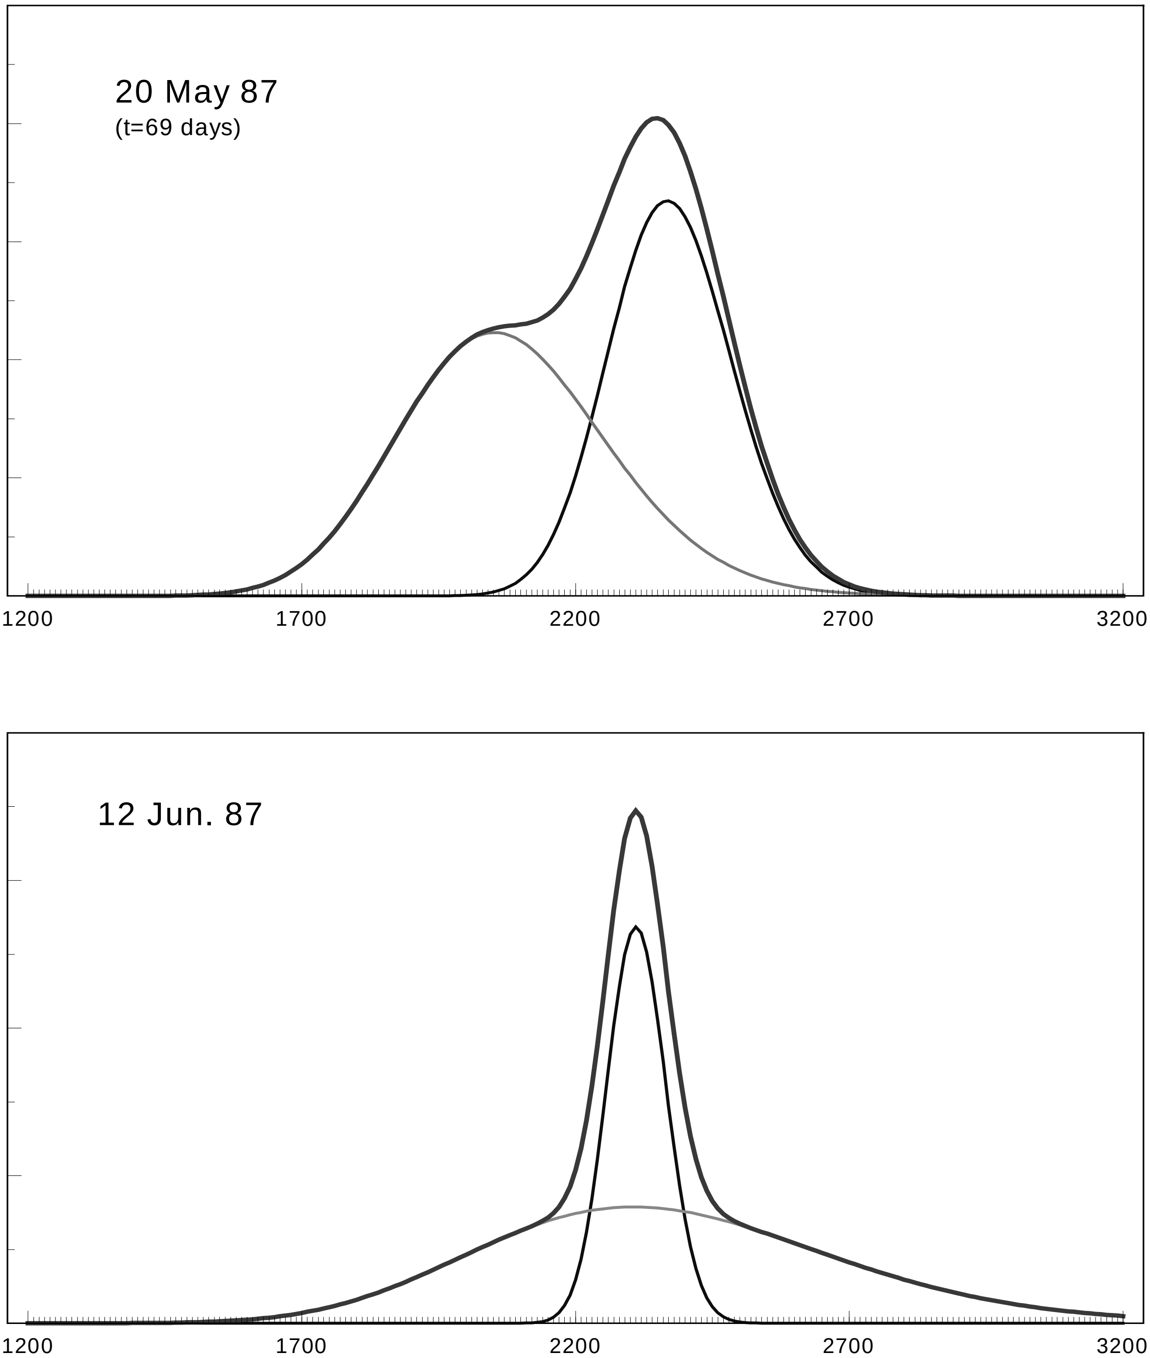

Figure 7   Implied Distributions Around the Election of 1992

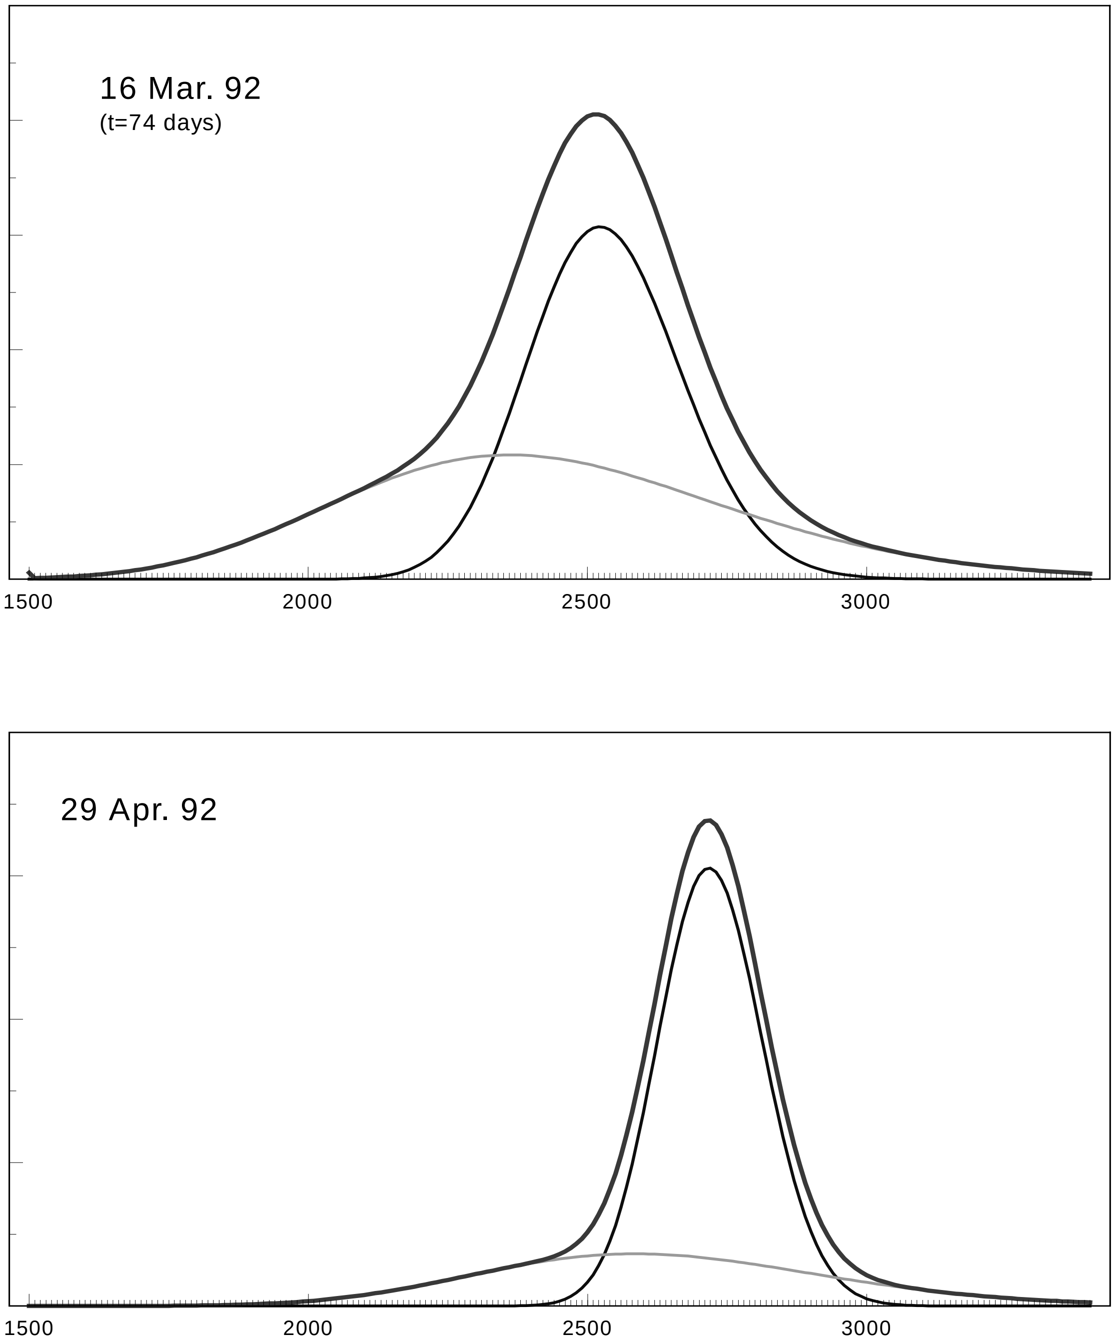

Figure 8 Implied Distributions Around the Election of 1997

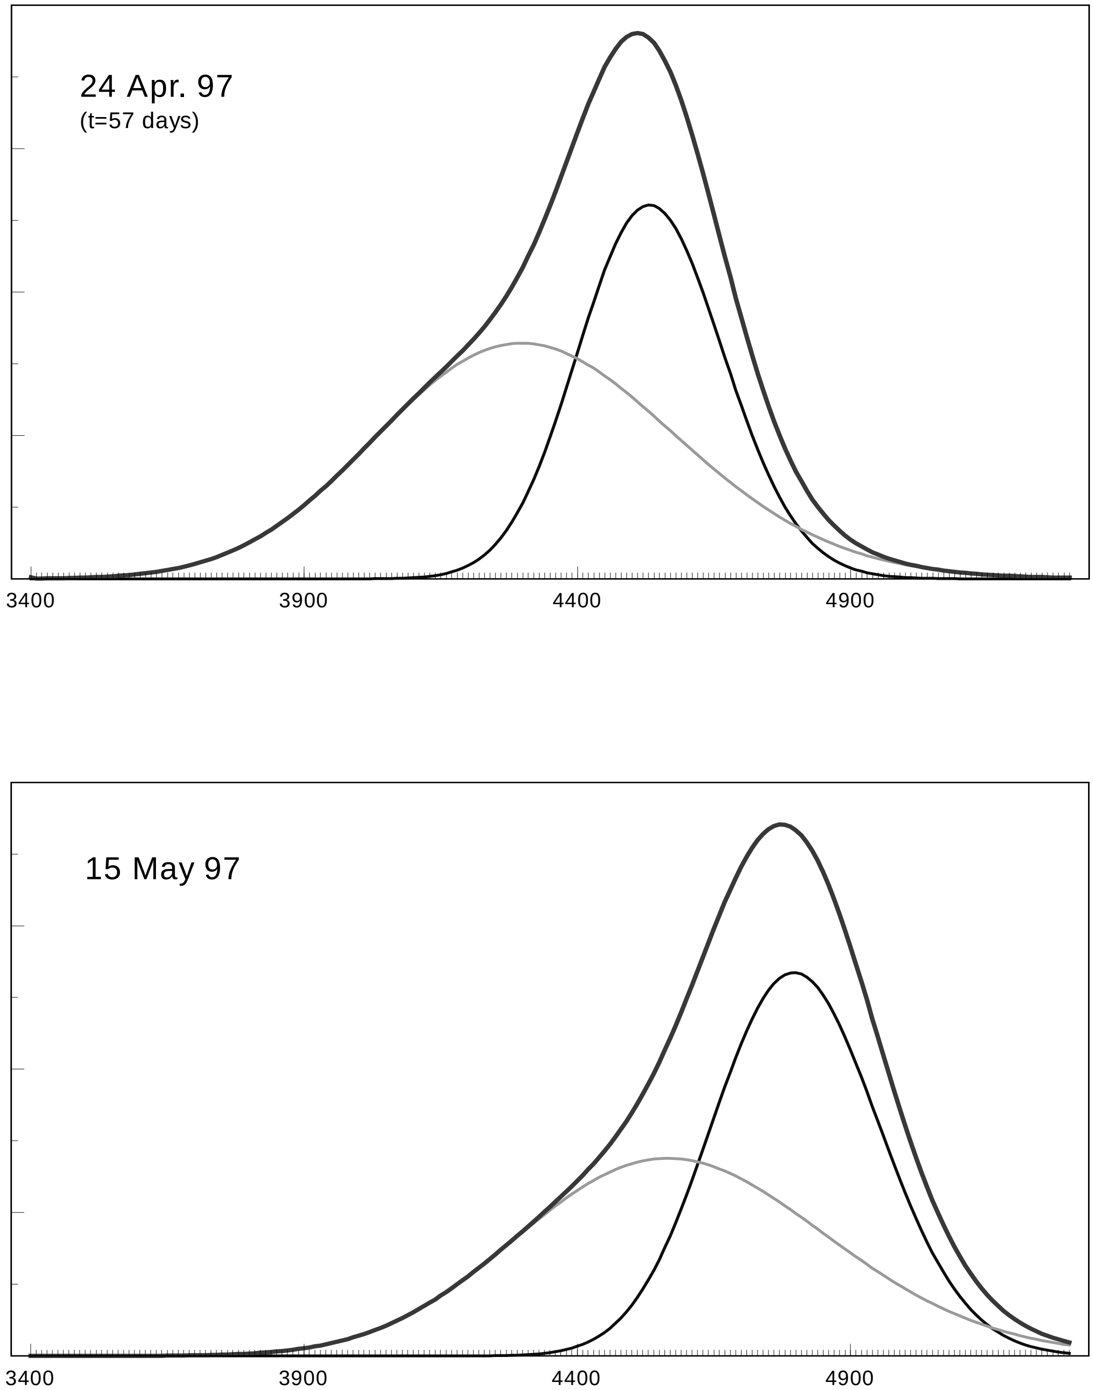

References

Ait-Sahalia, Y and A Lo (1995): “Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices”. NBER Working Paper No. 5351.

Bahra, B (1997): “Implied Risk-Neutral Probability Density Functions From Option Prices: Theory and Application”. Bank of England Working Paper No. 66, July.

Bates, D (1991): “The Crash of ’87: Was it expected? The evidence from Options Markets”. *Journal of Finance*, 46, July, pp 1009–44.

Bates, D (1996): “Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in PHLX Deutschemark Options”, *Review of Financial Studies*, 9, pp 69–107.

Black, F and M Scholes (1973): “The Pricing of Options and Corporate Liabilities”. *Journal of Political Economy*, 81, pp 637–59.

Breeden, D and R H Litzenberger (1978): “Prices of State-Contingent Claims Implicit in Option Prices”. *Journal of Business*, 9 (4), pp 621–51.

Butler, C and H Davies (1998): “Assessing Market Views on Monetary Policy: the Use of Implied Risk Neutral Probability Distributions”. Working Paper, Monetary Instruments and Markets Division, Bank of England.

Campa, J and P Chang (1996): “Arbitrage-Based Tests Of Target-Zone Credibility: Evidence From ERM Cross-Rate Options”. American Economic Review, 86 (4), pp 726–40.

Coutant, S, E Jondeau and M Rockinger (1998): “Reading Interest Rate and Bond Futures Options’ Smiles: How PIBOR and Notional Operators Appreciated the 1997 French Snap Election”. Working Paper, HEC School of Management, Jouy-en-Josas, France.

Crouhy, M and G Rockinger. (1993): “Asymmetric Conditional Heteroskedastic Processes with Hysteresis Effects”. Working Paper, September, HEC, Paris.

Dawson, P (1994): “Comparative Pricing of European and American Index Options: an Empirical Analysis”. *Journal of Futures Markets*, 14 (3), pp 363–78.

Dumas, B, J Fleming. and R Whaley (1998): “Implied Volatility Functions: Empirical Tests”. *Journal of Finance*, 53, (6), pp 2059–2106.

Gemmill, G (1992): “Political risk and market efficiency: Tests based in British stock and options markets in the 1987 election”. *Journal of Banking and Finance*, 16, pp 211–31.

Gemmill, G (1995): “Stockmarket Behaviour and Information in British Elections”, Working Paper, City University Business School.

Gemmill G (1996): “Did Option Traders Anticipate the Crash? Evidence from Volatility Smiles in the UK with US comparisons”. *Journal of Futures Markets*, 16, December, pp 881–97.

Grossman, S and Z Zhou (1996): “Equilibrium Analysis of Portfolio Insurance”. *Journal of Finance*, 51 (4), pp 1379–1403.

Harrison, J M and D M Kreps (1979): “Martingales and Arbitrage in Multi-Period Securities Markets”. *Journal of Economic Theory*, 20, pp 381–408.

Jackwerth, J (1997): “Implied Risk Aversion”. Working Paper, London Business School.

Jackwerth, J and M Rubinstein (1996): “Recovering Probability Distributions from Option Prices”. *Journal of Finance*, 51, December, pp 1611–31.

Longstaff, F (1995): “Option pricing and the Martingale Restriction”. *Review of Financial Studies*, 8 (4), pp 1091–1124.

Malz, A (1997): “Estimating the Probability Distribution of the Future Exchange Rate from Option Prices”. *Journal of Derivatives*, 5 (2), pp 18-36.

Malz, A (1996): "Option-Based Estimates of the Probability Distribution of Exchange Rates and Currency Excess Returns", mimeo, Federal Reserve Bank of New York.

Melick, W R and C 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*, March, pp 91–115.

Ritchey, R J (1990): "Call Option Valuation for Discrete Normal Mixtures". *Journal of Financial Research*, 13 (4), 285–96.

Rubinstein, M (1994): "Implied Binomial Trees". *Journal of Finance*, 49 (3), July, pp 771–818.

Shimko, D (1993): "Bounds of Probability". *Risk*, 6 (4), April, 33–7.

Shimko, D (1994): "A Tail of Two Distributions", *Risk*, 7 (9), September, pp 39–45.

Söderlind, P and L Svensson (1997): "New Techniques to Extract Market Expectations from Financial Instruments". Working Paper No. 5877, NBER.

Discussion of Gemmill and Saflekos paper
"How Useful are Implied Distributions? Evidence from Stock-Index Options"

Discussant: Shigenori Shiratsuka

The purpose of this paper is to estimate implied probability distributions (IPDs) as a mixture of two lognormal distributions by using British stock price index option data. It also empirically tests (1) the performance of pricing formula, (2) the forecastability of crashes, and (3) the revelation of market sentiments. The conclusions are as follows: (1) the two-lognormal method is better than the Black-Scholes model in fitting observed option prices (in-sample estimation), but there is no significant difference in predicting out-of-sample prices; (2) the shape of the IPDs does not indicate the forecastability of market crashes; and (3) the IPDs help to reveal changes in market sentiment.

Since I think that these conclusions are reasonable, I will make comments with a view to improving their robustness. My comments are divided into two parts: one is concerned with the technical issue of empirical procedures and the other is concerned with the more general issue of the case study methodology used to infer market sentiment from the IPDs.

Empirical Procedures

Regarding empirical procedures, I would like to point out two problems in this paper. First, I recommend that the authors use data of whole trading days in the forecasting exercise to compare the performance of different pricing formulas. The current estimations are done with data of one day in each month, probably because of the possible problem of autocorrelation caused by overlapping sample periods. However, considering that a standard method of adjusting for the effects of autocorrelation is available, I think it would be better to test forecasting power by using a larger sample of whole trading days.[1]

Second, I am wondering if the methodology used for estimating the IPDs to examine the forecastability of crashes is appropriate. According to this paper, we can see large estimation errors in market turbulence, suggesting that the approximation of the IPDs by a mixture of two lognormal distributions might be too rigid a restriction for the period of market stress, which is thought to be an important time for case studies. Therefore, I am not sure that the estimation methodology of the IPDs

---

[1] See Newey and West (1987) for details of how to estimate autocorrelation-heteroskedasticity robust standard errors.

The opinions expressed are those of the author and do not necessarily reflect those of the Bank of Japan.

is appropriately selected for examining whether the IPDs might be useful for anticipating crashes and gauging market sentiment? In this case, it might be that a simpler but less restricted methodology would be better for examining market sentiment during the stress period.

Case Study Method

Let me turn next to the method of case studies used in examining market sentiment. Here, I would like to propose another way of deriving market sentiment or expectations from the estimated IPDs, based on my own recent research (Nakamura and Shiratsuka, 1999).

It is important to recognize that there is a risk of misreading the estimation errors as indicators of changes in market expectations, if we look at the shape of the IPD on a specific date. Alternatively, we can avoid this risk by observing the general trend of changes in each summary statistic that is shown by the distribution shape, and focus on examining the relationship between asset price movements and changes in distribution shape.[2] In this case, I would like to emphasize the effectiveness of a time-series plot of summary statistics for the estimated IPDs in the case studies for examining the changes in market sentiment over time.

We estimated the IPDs for the Nikkei 225 Stock Price Index Option from mid-1989 to 1996 on a daily basis (see Figure 1). By observing the time-series movements of underlying asset prices and summary statistics of IPDs, we found the typical patterns in linking large fluctuations of asset prices and response in the shape of IPDs as follows (see also Figure 2):

(1) The standard deviation rises during a period of sharp decline in stock prices (high positive correlation with lagged absolute changes in stock prices), and a shift in level occurred at the end of 1989.

(2) The skewness moves in the opposite direction to changes in market level, reflecting lags in the adjustment of market participants' confidence to the market levels (high negative correlation with simultaneous and lagged changes in stock prices).

(3) The excess kurtosis jumps in the case of extreme price changes (positive correlation with simultaneous absolute changes in stock prices).

By comparing the actual movement of summary statistics with the aforementioned their typical patterns in response to the market fluctuations, the estimated IPDs provide us with a lot of information on the market sentiments and expectations. For example, the magnitude of changes in summary

[2] This paper computed summary statistics on original strike prices, and standardized summary statistics of IPDs for comparison. However, in our paper (Nakamura and Shiratsuka, 1999), we computed summary statistics on log-transformed strike prices, and did not need to standardize computed summary statistics for evaluating the divergence from normal position.

statistics tells us of the impact of external shocks. The length of adjustment periods provides us with information on the smoothness of adjustment of market expectations. Through the case studies of various episodes in Japanese financial markets during the period from 1989 to 1996, we have shown such usefulness of the IPDs as an information variable for monetary policy.

Conclusions

In summary, market participants' expectations are too diverse and informative to be captured merely by using a single summary statistic, i.e., the mean, because the same mean value implies different market expectations and policy implications, depending on the shape of the probability distribution of the expected outcome. In particular, since market participants' confidence in stock prices differs substantially depending on timing, we can expect to capture more market information, both in quality and in quantity, by carefully examining the changes in market participants' expectations that lie behind stock price fluctuations.

In this sense, as the papers and discussions contributed to this workshop suggest, the implied probability distributions extracted from option prices will provide useful information for the conduct of monetary policy. However, studies on how to make use of the information extracted from option prices in policy judgments have only just begun, and further research is necessary in this area.

References

Nakamura, Hisashi, and Shigenori Shiratsuka: "Extracting Market Expectations from Option Prices: Case Studies in Japanese Option Markets." *BOJ Monetary and Economic Studies*, 17 (1), 1999, pp. 1-43.

Newey, Whitney K and Kenneth D.West: "A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." Econometrica, 55 (3), 1987, pp. 703-08.

Figure 1. Stock Prices and Market expectations (Overview)

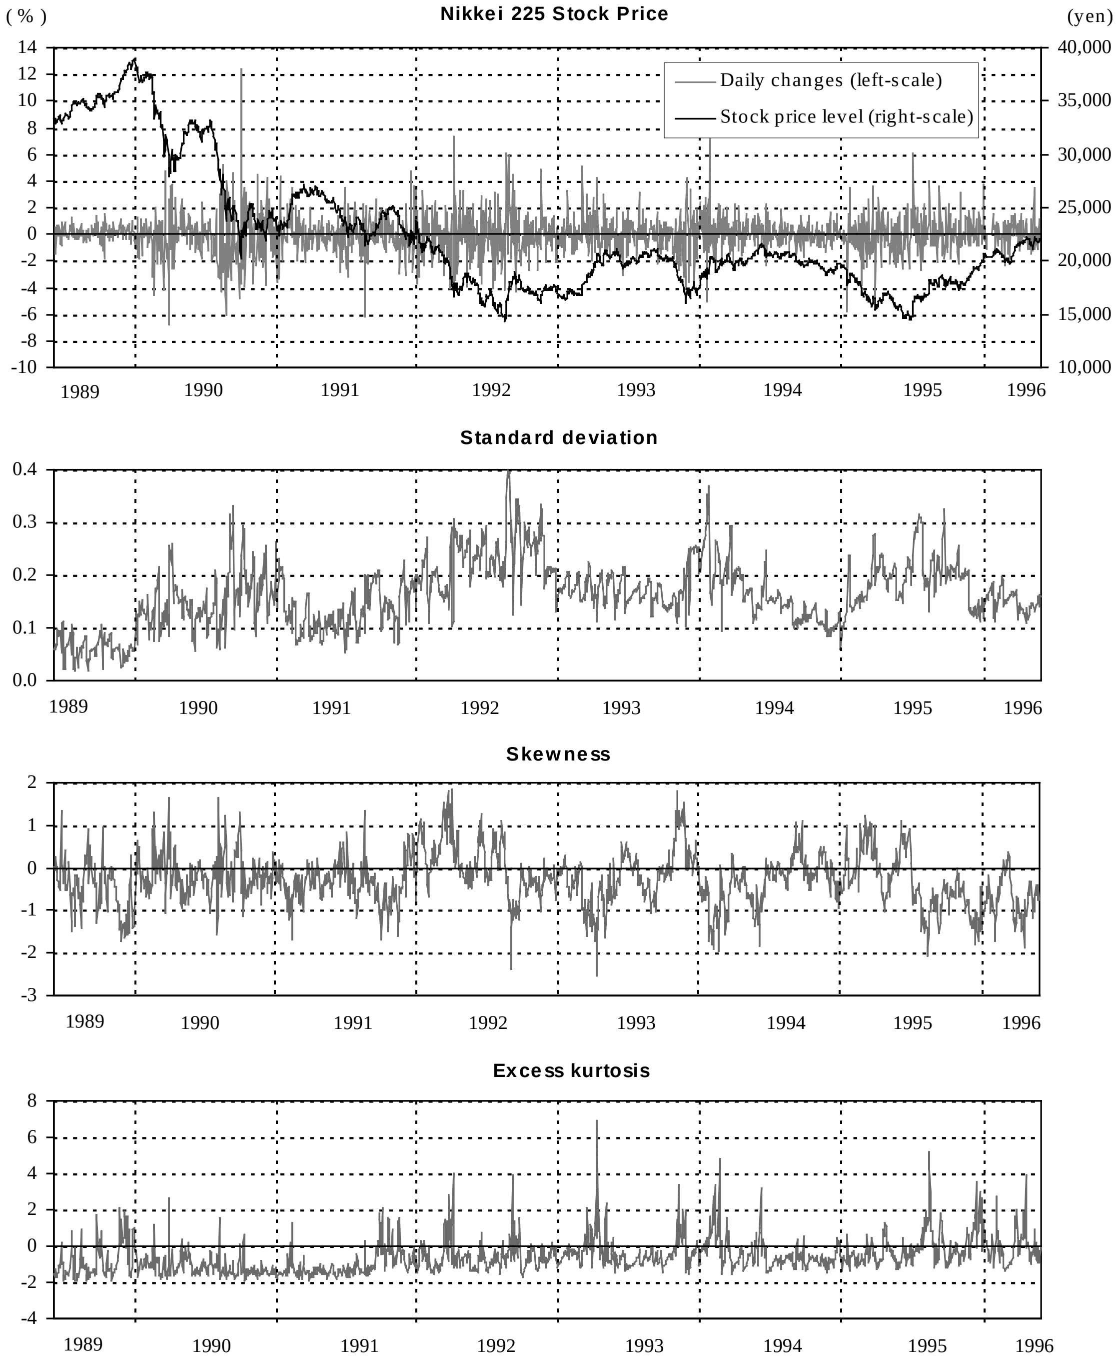

Source: Nakamura and Shiratsuka (1999)

Figure 2. Dynamic Cross-Correlation

(1) Market Changes (t) vs. Summary Statistics (t-k)

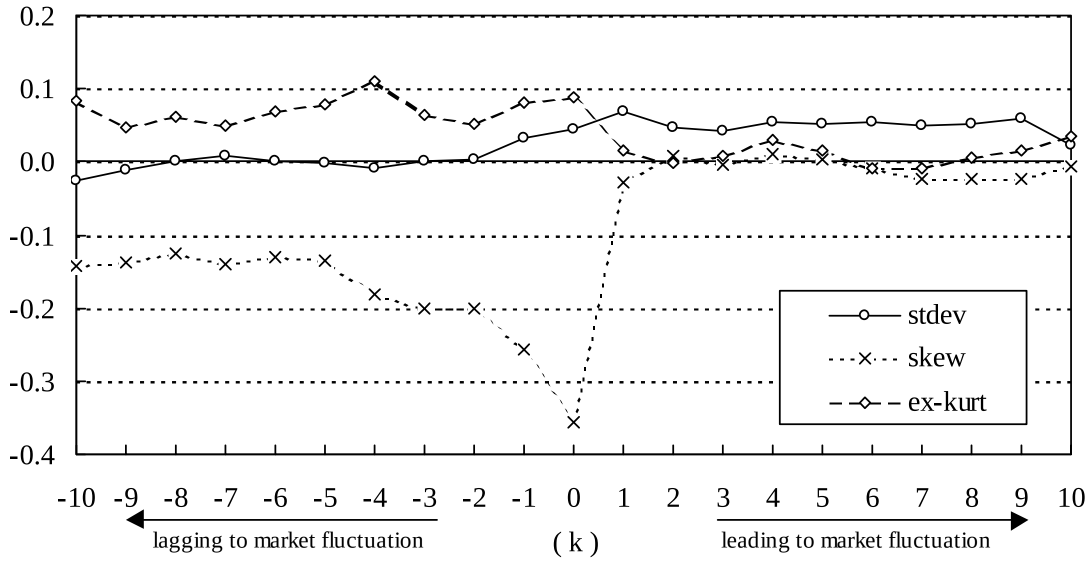

(2) Absolute Market Changes (t) vs Summary Statsitics (t-k)

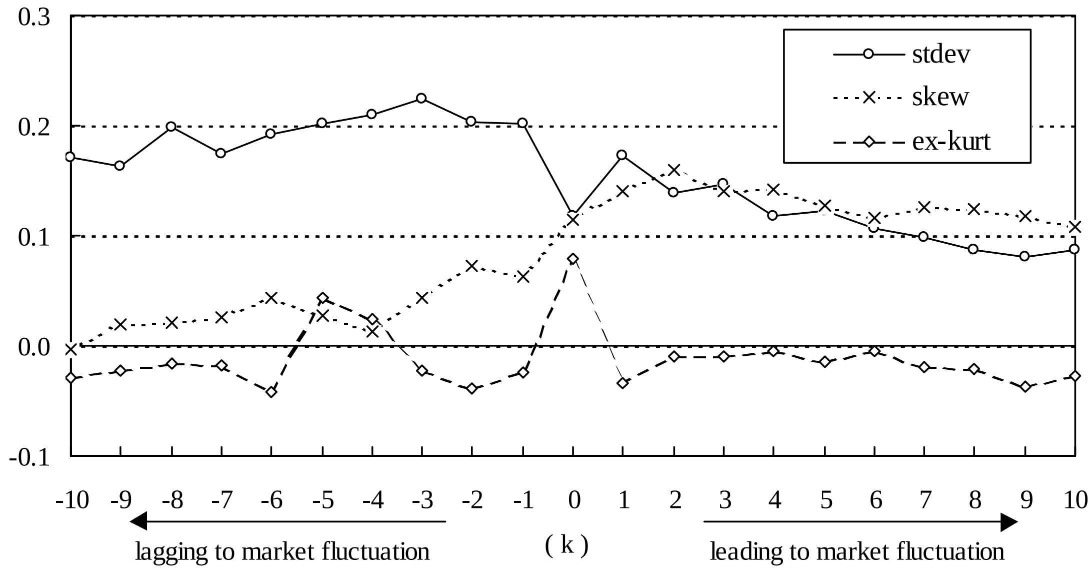

Discussion of
"How useful are implied distributions ? Evidence from stock-index options"
by G. Gemmill and A. Saflekos
Discussant: Raf Wouters

It is a pleasure for me to discuss this paper. The paper covers several different topics :

- the estimation method of the implied distribution ;
- the test of the information content of the implied distribution ;
- the effect of specific events on the implied distribution.

This diversity of topics makes it somewhat easier to discuss the paper as I can choose the topics that look most interesting to me. Since the estimation technique was discussed during the morning session, I will concentrate my remarks on the test of the information content of the implied distribution (i.d.).

The authors use an out-of-sample pricing prediction to test the information quality of the estimated i.d. This is a strong test for the i.d. as it depends on the whole structure of the distribution.

Traditional tests of the information content of option prices have concentrated on the test of the volatility measure only: they analysed whether the implied volatility was an unbiased predictor of the ex-post realisation of the average volatility over the option maturity. The general conclusion was that the implied volatility is a biased estimator because systematically overpredicting the realised volatility. On the other hand, the relation between implied volatility and the realised volatility was proved to be very significant and the implied volatility outperformed the forecasts of time-series based volatility measures such as Garch estimation techniques.

The out of sample pricing test used in this paper is a more general test : by pricing the whole range of options, the error will depend on both the volatility forecast but also on the higher moments describing the whole shape of the i.d.

The authors summarise these results by the RMSE of the forecast and compare this figure with the in sample errors. The mixture of two-lognormals delivers a strong improvement both in sample (but without correction for the number of parameters in the model) and out of sample compared to the B.S. model. These statistical gains are impressive at least for the period where European options are used (indicating perhaps that the method as applied here is not well suited for American style options).

This is a strong result and confirms similar conclusions in the literature (e.g. Dumas a.o. JF 1998). However, the authors minimise this finding: the improvement in value terms is not economic meaningful as it falls within the typical bid/ask spread in the market. They conclude from this that the forecasting performance is not markedly better than B.S. and that these results cast doubt on the hypothesis that the shape of the i.d. is a rational expectation.

In my opinion this conclusion is drawn somewhat to quickly : I would like to see more detailed results and add more tests:

- the average RMSE is too general: I would suggest for instance some statistic on the number of cases and the size of improvements in the fit in excess of the bid/ask spread (the outcome of some investment strategy would even be more conclusive). I would like to add here that by minimising the sample size to 10 observations, one probably makes it more difficult to find important improvements : so why not take the whole set of available prices into account?

- I wonder whether the errors are systematically related to strike prices and whether such systematic errors are higher / lower compared to the B.S. model

- I would also prefer to see predictions over a longer horizon: this can increase the possibility to find economic significant improvements;

- I would also compare the pricing error to a third model, for instance a simple approach based on a martingale assumption in which the pricing in the forecast is based on the observed volatility smile of the previous day.

Whatever the answers to these remarks, the important question remains as to why the prediction error is so much bigger out-of-sample as compared to the in-sample estimation error:

- a first explanation is the misspecification of the model based on the mixture of two-lognormals. However the small in-sample errors, typical for methods based on option prices of one maturity and at one point in time, provide not much hope that further improvements can be achieved within this approach;

- a second answer is the arrival of new information that changes not only the price of the underlying asset but also the volatility and the whole shape of the distribution. Under this hypothesis one can still test whether the forecast is unbiased and rational, but it will be statistically more difficult to find the answer. But more important, under this hypothesis the implied volatility and the i.d. can no longer be considered as constant and one should go in the direction of modelling the behaviour of these variables over time (Examples like Ait-Sahalia JF 1998 give promising results in this respect). But if one accepts that volatility is non-constant over time and that this variability is very important, one arrives quickly at the limit of methods based on a specification of the terminal distribution or other “non-structural” methods to estimate the option pricing function. These methods can never explain the economic logic behind the changes in the volatility or in the higher moments of the distribution, neither can they explain, in economic terms, the observed negative skewness and high kurtosis. Therefore one should move to structural approaches that model explicitly the dynamic process of the asset price and the volatility process. Within such a framework one can look for alternative explanations for these observations in the direction of :

- the role of the leverage effect;

- the role of non-normally distributed shocks;

- the variable risk-aversion, etc;

The test of such a model typically covers a broad range of option prices both cross-section and over time. The existence of higher in-sample errors in such application leaves more room to find economic significant differences in and out of sample;

- a third alternative explanation for the high prediction errors can be found in specific factors or characteristics of the option market. For instance through the existence of :

- bid-ask spreads that generate pricing errors;

- important liquidity problems due to the lower number of transactions in far out or in the money options;

- specific exposure or insurance arguments in the option market that distort the option prices.

This last type of explanations for the pricing errors make the use of the i.d. to investigate their information content problematic. Such distortions make the i.d. less useful or at least more difficult to interpret. Only option investors would still be interested to study these effects.

Now perhaps it is too strong to separate these different explanations underlying high prediction errors : if option market behaviour has a feedback to the underlying asset market, as one should expect in a general equilibrium framework, the different explanations can no longer be separated, and the interpretation becomes difficult in any case.

As long as our knowledge on the mechanisms that drive the movements in the distribution remains limited, we should be careful in using this information, and indeed consider it, as the authors suggest, only as some indication of market sentiment:

- simple volatility measures can be used to give an indication of the uncertainty of investors outlook. The use of implied volatility for analysing credibility of monetary policy typically falls in this category ;

- but the interpretation of the higher moments should be made very cautiously.

This conclusion also follows from the review in the paper of the i.d. behaviour around specific events. The i.d. does not show any systematic behaviour during these periods. But the exercises, and in particular the in-sample statistics, illustrate that the mixture of two lognormals does a reasonable job in describing the information of option prices during periods of high negative skewness, high kurtosis or bimodal distribution.

José M Campa, P H Kevin Chang and James F Refalo

"An options-based analysis of emerging market exchange rate expectations: Brazil's Real Plan, 1994–1997"

Discussants:

Paul Söderlind

Mª Cruz Manzano

An Options-Based Analysis of Emerging Market Exchange Rate Expectations: Brazil's *Real* Plan, 1994-1997

by

José M. Campa
New York University and NBER,

P.H. Kevin Chang
Credit Suisse First Boston, New York,

and
James F. Refalo
New York University

Revised first draft: June 23, 1998
Latest revision: December 3,1998

ABSTRACT

This paper uses currency option data from the BMF, the Commodities and Futures exchange in Sao Paulo, Brazil, to investigate market expectations on the Brazilian Real-U.S. dollar exchange rate from October 1994 through July 1997. Using options data, we derive implied probability density functions (PDF) for expected future exchange rates and thus measures of the credibility of the "crawling peg" and target zone ("maxiband") regimes governing the exchange rate. Since we do not impose an exchange rate model, our analysis is based on either the risk-neutral PDF or arbitrage-based tests of target zones. The paper, one of the first to use options data from an emerging market, finds that target zone credibility was poor prior to February 1996, but improved afterwards. The market anticipated periodic band adjustments, but over time developed greater confidence in the Real. We also test whether devaluation intensities estimated from these option prices can be explained by standard macroeconomic factors.

JEL Codes: F31, G15, G13, F37.

Acknowledgements: For data, helpful comments and discussions, we would like to thank: Rodrigo Azevedo, Axel Blikstad, Silvio Flores. Ricardo Lacerda, Fernando Pinheiro, Marco Suplicy, Marco Aurelio Teixeira Alex Treveza Gyorgy Varga, and seminar participants at the University of Virginia and the International Monetary Fund. We also thank Steve Sedmak for research assistance. This research was partly supported by a CIBEAR grant from the Marshall School of Business, University of Southern California.

An Options-Based Analysis of Emerging Market Exchange Rate Expectations: Brazil's *Real* Plan, 1994-1997

This paper uses a new data set of options on the Brazilian Real / US dollar exchange rate to extract market expectations, as embodied in the risk-neutral probability density function (PDF), of real-dollar exchange rates over horizons of one to three months. Unlike ordinary exchange rate forecasts that provide only a point estimate of the future exchange rate, options-based forecasts, by permitting the derivation of a PDF, describe a range of realizations and the probability attributed to each range.

This PDF-based approach is especially effective for an analysis of the Real/$ exchange rate, which since the June 1994 Real Plan has been characterized by a combination of a crawling peg and a target zone regime. Over short horizons, the exchange rate has followed a crawling peg surrounded by a "miniband," but for long horizons, superimposed on the crawling peg, there has also been an official "maxiband" with a fixed (non-crawling) central rate, floor, and ceiling.

The PDF's derived in this paper enable us to compare market expectations embedded in options with these two concurrent regimes. From the derived PDF's, we can identify any divergences between market expectations and the existing crawling peg—e.g., whether markets in fact anticipated a faster depreciation, and if so, where (relative to the crawling peg) probability was concentrated. Relative to single-point expectation of the future exchange rate, a great advantage of a full PDF is the ability to disentangle magnitude and probability of expected depreciation—e.g., a high probability of a small depreciation *vs.* a low probability of a large depreciation, with presumably very different policy responses. For the longer-horizon fixed target zones, we can perform a similar decomposition of probability and magnitude of depreciation, and moreover, conduct "arbitrage-based tests" of credibility, developed in Campa and Chang (1996), that are virtually assumption-free. Given these target zones, we are also able to determine both "intensities" and probabilities of realignment, and to investigate possible economic determinants of realignment intensity. Thus, a single approach using dollar-Real options permits us to analyze both facets of the post-Real Plan Brazilian exchange rate regime.

This work contributes to the growing literature on the use of options to characterize expected asset returns, and in particular to predict currency crises. Recent empirical work using options to identify the distribution of expected exchange rates includes Malz (1996) and Campa et al. (1997, 1998). Papers specifically focusing on currency crises, especially the 1992 ERM crisis, include Campa and Chang (1996), Malz (1996), and Mizrach (1996). These can be contrasted against measures of devaluation risk *not* based on options, as in Bertola and Svensson (1993), Kaminsky, Lizondo, and Reinhart (1997) and, Svensson (1991).

The motivation for this research is two-fold: first, to use options-based estimates of the PDF to compare and contrast market expectations with the two concurrent exchange rate regimes in the post-Real Plan Brazil; and second, to observe the time path of market perceptions to gauge policy effectiveness over time. Furthermore, this is one of the first options-based tests of exchange rate regime credibility on an emerging market. Within emerging markets, this is also the first paper to deal with the data challenges of exchange-traded options, rather than over-the-counter (OTC) volatility quotes. OTC data are normally subject to less observation error, and are by construction free of arbitrage violations. Hence, OTC data are easier to interpret empirically. Thus, results obtained here may have implications for applying this technique to other emerging markets, including those with only exchange-traded currency options.

The remainder of the paper is structured as follows. Section I describes the theoretical background behind the use of option prices to determine risk-neutral probability density functions (PDF's)—and for target zone regimes, the derivation of re-alignment intensities and probabilities, as well as arbitrage-based tests of credibility. Section II discusses the Real Plan and pertinent historical background, including the "miniband" and "maxiband" regimes. Section III introduces our option data, provides summary statistics, and conducts a preliminary analysis. Section IV investigates the behavior of the PDF over time, and in the context of a crawling peg, describes the probability and magnitude characterizing expected deviations from this regime. Section V addresses the "maxiband" target zones, estimated realignment intensities and probabilities, and arbitrage-based measures of credibility. Section VI explores the empirical relation between estimated intensities and standard macroeconomic factors. Section VII concludes.

I. Options-Based Indicators of Devaluation and Tests of Exchange Rate Band Credibility

Options—whose payoff depends on a limited range of future exchange rates rather than an entire distribution—are able to provide more precise information than other financial indicators about the future exchange rates expected by the market, and the amount of probability attributed to any given realization. In contrast, the forward rate, for example, can indicate only the mean of this distribution. The advantages of an options-based approach will be discussed further below.

A. Options and the Risk-Neutral Distribution

We begin with a few brief definitions. A *call* option gives its holder the right but not the obligation to purchase a fixed amount of foreign currency (in the case of Brazilian Real-US dollar options, $1000 US) at a pre-determined price (referred to as the *strike* price or *exercise* price) in local currency. A *put* option gives the right but not the obligation to conduct the reverse transaction, i.e., to sell a fixed amount of the foreign currency ($1000 US) for a given strike price in local currency.

An *American* option may be exercised at any time before its expiration date; a *European* option, only on its expiration date. Because the European option can be exercised only on a single date, an analytical relationship known as *put-call parity* can be established between the price of a European call and European put of the same strike. This relation, which is derived from arbitrage restrictions, permits the price of a call to be computed from the price of a put with the same strike, and vice-versa. The analysis in this paper focuses exclusively on relationships derived from European call options (though some of the call option data were constructed from European put data via put-call parity).

It was first shown in Breeden and Litzenberger (1978) that the decline in the value of a European call option due to an infinitesimal increase in the strike price equals the discounted risk-neutral probability that the option will finish "in-the-money" (spot exceeding the strike on expiration). Accordingly, the value of a call option (under risk-neutrality) at time T with a strike price K is then given by

$$Call_{K,T} = \frac{1}{1+i_T} \int_K^{\infty} (S_T - K) f(S_T) dS_T \ , \tag{1}$$

where S_T is the spot exchange rate at time T, $f(S_T)$ is the risk-neutral density function for the spot at time T, and i_T is the domestic risk-free rate for an investment maturing at time T. The partial derivative of equation (1) with respect to strike price K is:

$$\frac{\partial Call_{K,T}}{\partial K} = -\frac{1}{1+i_T}\left[1 - F(K)\right], \tag{2}$$

where F(K) is the risk-neutral cumulative density function of the exchange rate at time T, evaluated at strike price K. Taking the second partial derivative of equation (1) with respect to strike yields:

$$\frac{\partial^2 Call_{K,T}}{\partial K^2} = \frac{1}{1+i_T} f(S_T) \, . \tag{3}$$

This then provides a direct relationship between observed European call prices and the value of the corresponding risk-neutral probability density function, i.e. the PDF. Note that the call price is based on the payoff (S_T - K) multiplied by its risk-neutral probability $f(S_T)$, which incorporates both the actual probability of that realization of spot and the value the market places on that state of nature. In other words,

$f(S_T)$ is not necessarily the actual density function, since—because of risk—a dollar in one state may be valued differently from a dollar in a different state.

Equation (3) is important because it provides the method by which the PDF can be extracted from call prices. If a continuous call price function twice-differentiable in strike exists, then the PDF is uniquely determined. In reality, such a continuous call price function is not available, but will be estimated from discrete point observations using a method described in Section IV of this paper.

B. "Intensity" of Devaluation or Realignment

When there are specific reference exchange rates in place, as in the case of target zones, a risk-neutral PDF can be used to indicate the perceived probability of devaluations or "re-alignments" of various sizes beyond that specific reference level. By looking at only that part of the PDF representing a deviation from the reference exchange rates, we can isolate the risk of a change in regime. A summary measure incorporating both probability and magnitude of change from given reference rates, over all possible realizations deviating from these reference rates, can be termed an "intensity" measure. Campa and Chang (1996) define such an intensity as:

$$G(T) = \int_{\bar{S}}^{\infty} (S_T - \bar{S}) f(S_T) dS_T \ , \tag{4}$$

Intuitively, intensity G(T) is a risk-neutral probability-weighted average of all exchange-rate realizations requiring a re-alignment, or under deviation scenarios beyond S-overbar. In other words, the magnitude of realignment is multiplied by the risk-neutral probability of each realization. Comparing equations (4) and (1), the intensity of realignment is simply the future value of a European call with a strike price at the upper-bound. Mathematically,

$$G(T) = Call_{\bar{S},T}(1 + i_T) \ . \tag{5}$$

Though this call with a strike price at the upper-bound does not exist in most cases, its price (and hence the intensity of realignment) is easily calculated once a risk-neutral PDF has been derived.

C. *Minimum* Intensity of Devaluation or Realignment

In the absence of a complete risk-neutral PDF, Campa and Chang (1996) show how one can compute a lower bound on re-alignment intensity using far fewer data points but relying on convexity properties of the call price with respect to strike, and the existence of one credible second reference rate.

This method uses both an observed at-the-money option, as well as the hypothetical price of an option whose strike price is the *credible* side of the target zone. The latter option will *always* end in-the-money, and hence can be evaluated as a bond or forward contract, since there is no uncertainty and no time value. The method then relies on the arbitrage-based condition that call prices are always a convex and non-increasing function of strike price. Therefore, when the call is graphed as a function of strike, the point corresponding to a call with a strike at the upper band (i.e. the realignment intensity) must lie *above* the rightward extension of a line connecting $Call_{\underline{S},T}$ (a call with a strike at the lower band expiring at time T) and any $Call_{K,T}$ (with strike K below the upper band). The following inequality summarizes this:

$$Call_{\overline{S},T}(1+i_T) \geq \left[\frac{\overline{S}-\underline{S}}{K-\overline{S}}\right] + K - \overline{S} \tag{6}$$

D. Arbitrage-Based Tests of Target Zone Credibility

Campa and Chang (1996) also develop two tests of band credibility relying solely on arbitrage or convexity arguments, without assumptions about risk preferences. These tests will be used for analysis and comparison in Section V.

The first test (hereafter referred to as “Test 1”) is based only on a simple no-arbitrage restriction: the maximum future spot rate cannot exceed any credible upper band. At expiration, the payoff of a European call equals, at most, spot minus strike. Therefore, under credibility, the maximum value of the call cannot exceed the present value of the upper band minus the strike. Thus, credibility can be rejected whenever

$$Call_{K,T} > \frac{\overline{S}-K}{1+i_T} \tag{7}$$

Note that this test can be used even when there is only one reference rate.

The second test (“Test 2” from here on) is derived from convexity arguments and also provides an upper bound for the value of a call with a strike between two reference rates, or within the bands of a target zone. The argument is that under credibility, a call with a strike at or below the lower band will always finish in-the-money, and therefore is worth exactly its intrinsic value. This intrinsic value is $S_0/(1 + i_T^*)$-$K/(1 + i_T)$, where i_T^* is the foreign risk-free rate and S_0 is the current spot. Furthermore, a call with a strike greater than the upper band will always finish out-of-the-money, and therefore be worthless. Call value, when mapped against strike, is a convex function passing through these two points. Therefore, a straight line—since we do not know the degree of convexity of the call function, but do know that it cannot be less convex than a line—connecting these two points must provide an upper bound on all points in between. Thus, credibility can be rejected whenever the call value exceeds this upper bound, or

$$Call_{K,T} > \left[\frac{S_0}{1+i_T *} - \frac{\underline{S}}{1+i_T} \right] \frac{\overline{S}-K}{\overline{S}-\underline{S}} . \tag{8}$$

Notice that by rearranging the terms of Test 2, we can show that the RHS of Test 2 is equal to the RHS of Test 1 times a coefficient less than one, provided the forward rate does not exceed the upper band. Therefore, as long as the forward rate is within the band, Test 2 is always at least as restrictive as Test 1. The one advantage of Test 1 is that it does not require the existence of a credible second reference rate, and provides a valid test even in the absence of, for example, a credible lower band. Test 2, in fact, is a test of the joint hypothesis that two reference rates, or the lower and the upper band of a target zone, are credible.

This section has summarized four different (but related) measures of exchange rate band credibility based on options data alone: (1) the PDF-based realignment intensity, when a full PDF can be described; (2) the minimum realignment intensity, given a credible second reference rate, (3) arbitrage test 1 (on one exchange rate band), and (4) arbitrage test 2 (based on lower and upper bands). In terms of how these four measures are related, recall that Test 2 is always more powerful than Test 1, but requires more assumptions. There is also a one-to-one mapping between violation by Test 2 and a positive minimum intensity of realignment, as both are based on convexity properties alone.

II. The Real Plan and Relevant History

Brazil has been subject to high levels of inflation since the early 1980's, and had unsuccessfully attempted to rein in inflation several times prior to 1994’s Real Plan. Economic problems, in part, date back to 1964 when the military overthrew the civilian government, resulting in military control of the economy until 1985. (It was not until 1990 that the first popularly elected president was inaugurated.) During this military-ruled period, Brazil pursued industrialization policies based on import substitution, creating a number of large state-owned enterprises. The government engaged in protectionist trade policies to spur such industrialization and to create economic independence in key industries. As a result, by the 1980's, foreign investment in support of the inefficient industries collapsed, and hyperinflation followed because the high levels of government spending could not be reduced in line with reduced capital inflows. By 1990, hyperinflation had been structured into the economy, through both indexation and expectations, with the concomitant debilitating effects.

Prior to the Real Plan, several attempts were made to contain inflation, usually involving combinations of wage and price controls, tightening of the money supply, tax hikes, and freezing of bank deposits. These all failed as the fundamental problem lay in expectations of high inflation and excessive government spending. Wage and price controls were often ignored by the private sector, as immediate shortages often resulted, creating price pressure. In many industries, cartels also prevailed, reflecting the low degree of market competition. Government spending proved difficult to curtail given large entitlement programs in place, and vested interests resisting spending cutbacks and privatization of government industries, particularly during a recession. Fiscal troubles were compounded by a badly written constitution providing tenure for government employees after only five years (making them virtually impossible to lay off), and guaranteeing individual states a right to share in the federal revenues without restricting state spending. An attempt to introduce a new currency in 1993, the Cruzeiro Real ran into the same problem of inflationary expectations.

The Real Plan, introduced in December 1993 by Finance Minister Fernando Henrique Cardoso, differed from the previous plans in that it directly addressed the problem of inflationary expectations. Cardoso recognized that past inflation was being transmitted into future expectations by indexing and various contract negotiations, as inflation figured into all wage and business contracts. The idea was to break this connection by creating a unit of transactional account in which price and wage contracts would be negotiated and written, and whose value would be kept roughly equal to $1. The official currency, the cruzeiro real, would then be devalued against this unit. The Unit was called the Unit of Real Value (URV), and was introduced in March 1994. At the same time, the constitutional links between revenue and expenditure were circumvented by creating a special fund (Fundo Social de Emergencia - FSE) to eliminate the public sector deficit, thereby addressing a fundamental source of inflationary pressure. (The creation of the FSE was necessary to avoid the structural claims guaranteed by the constitution to the states and to entitlement programs.) Four months after the introduction of the URV, the Real was introduced. The central bank (Banco Central do Brazil) committed not to permit a depreciation beyond 1.00 Real/$, though appreciation would be allowed. Furthermore, a reserve ratio was implemented requiring one American dollar to each Real emitted.

The result of the Real Plan was a reduction in inflation from 50% per month, as of June 1994, to less than 2% per month by the end of the year. Inflation has since then continued to drop, and in May 1998, 12-month inflation was 3.12%, its lowest value since November 1949. The Real Plan has also had positive effects on the rate of economic activity. Brazil's real GDP grew at an average annual rate of 4.0% during the four-year period 1994-1997, compared with an average annual growth rate of -0.2% during the four years prior to the implementation of the plan[1].

Exchange rate stabilization was an integral part of the Real Plan. Upper and lower bands ("maxibands")—as indicators of the maximum possible movement up or down—were established in March 1995, at a rate of .93 and .88 Real/$. Since then, they have been adjusted several times to allow the Real to depreciate at a controlled rate.

While announcing these broader "maxibands," the central bank in practice followed a "crawling peg" system, in which the Real gradually depreciated, but remained within a "miniband" surrounding a depreciating central rate . Under this informal system, the Real's central rate was devalued approximately 0.5%-0.6%/month, and central bank intervention assured that at all times, the spot rate would not deviate by more than 0.25% (half the "miniband width") in either direction. In practice, the central rate was devalued discretely by about 0.10% (although sometimes 0.05% or 0.15%) about 5-7 times per month. Starting April 1997, the government started devaluing the central rate by about 0.7% monthly. To discourage speculation against the system, the actual magnitude and timing of these mini-devaluations was kept slightly irregular. Furthermore, the size of the minidevaluation would be smaller than the width of the miniband itself, so the instantaneous direction of the spot rate could not be known with certainty, discouraging "one way" bets.

While this system of a predictable crawling peg surrounded by a miniband provided short-term stability in the spot rate, the government wished to maintain some flexibility in its commitment to the exchange rate over longer periods such as several months. To commit to a very narrow range, even one surrounding a crawling peg, risked tying the government's hands unnecessarily and inviting outside speculation against the currency. Thus, the government remained free to alter either the rate of devaluation or the width of the miniband. At the same time, the government also wished to provide some indicative levels for medium-term forecasting. This dual objective was reached by instituting wider "maxibands." Though the exchange rate never technically violated these bands, the central bank adjusted the maxibands as markets gradually approached the maximum Real/$ exchange rate, an event that has typically occurred every six to twelve months (Figure 1) .

---

[1] It is worth noting that measures to address the structural problems also appear to be proceeding. Privatization of state industries is continuing, and the Brazilian Congress has agreed to several constitutional reforms. These include: the relaxation of protectionist provisions not mentioned here, reform of the social security system, and provisions relaxing the excessive protections provided to pubic workers.

Since the original maxibands were implemented in March 1995, through the end of our data in July 1997, the bands were changed on three separate occasions: June 22, 1995; January 30, 1996; and February 18, 1997. Since the end of our data, the maxibands were changed on January 19, 1998, to the current upper and lower bands of 1.22 and 1.12 Real/$ respectively. In April of 1998, the government also announced that the lower end of the miniband would depreciate at a rate of 0.65% a month, while the upper band would depreciate at a rate of 0.75%, de facto widening the minibands over time. Despite its short history, the Real Plan appears to have been quite successful in taming Brazilian inflation and establishing a relatively stable currency with a reasonably predictable rate of devaluation. Given this track record, the following sections will seek to investigate issues of exchange rate regime credibility—both the crawling peg and the maxiband system—and how market perceptions of the distribution of Real/$ spot rate have changed over time.

III. Data Description

The data obtained consist of high, low, average, and last transaction prices for every trading day of dollar futures (daily observations of contracts of multiple maturities), calls and puts (daily observations of multiple strikes and expiries), and closing spot rates, from the Commodities & Futures Exchange (Bolsa de Mercadorias & Futuros, known as the BM&F) in Sao Paulo, Brazil. The data cover the period from July 1994, shortly after implementation of the Real Plan, through July 1997. Calls are initially both European and American, until a 1995 shift in convention, making all calls expiring after October 1995 European. All put contracts are European.

The BM&F was formed in July 1985 and began trading in January 1986. Currently, the exchange offers a range of futures and options contracts on the US dollar, the Ibovespa (the Brazilian stock index), sovereign debt instruments, inter-bank deposit rates, US-Brazilian interest rate spreads, gold, cattle, and agricultural commodities. With a total trading volume of 102.3 million contracts and a financial volume of 6.1 trillion US dollars during 1997, BM&F is currently ranked fourth among the world's derivative exchanges[2]. In 1997, 39.7 million contracts traded were US dollar futures, 8.1 million contracts were US dollar call options, and 71,820 contracts were US dollar put options[3]. Total trading volume in foreign exchange contracts has actually declined since 1985, but this is due primarily to increases in contract size; financial volume has more than doubled from 1996 to 1997. US dollar contracts for both futures and

[2] Bolsa de Mercadorias & Futuros 1997 Annual Report.
[3] These figures are for contracts based on commercial US dollar rate. Contracts are also available based on the floating rate, however, these represent less than one percent of total transactions volume.

options apply to the “commercial” (as opposed to financial) exchange rate on a notional amount of $100,000.

In this paper, we focus on European call data, which significantly outnumber put data. Put data were translated using put-call parity and used only to augment the call data if a corresponding call did not exist. The call data consist of 5855 observations from the time period with mixed American and European calls, and 4837 usable (about 200 questionable observations were deleted) observations of the later time period with purely European calls. These call prices were supplemented with put data consisting of 530 observations for the first time period, and 218 from the second time period. This revised data set forms the basis for our subsequent analysis.

In deriving the PDF and conducting credibility tests, we use futures prices as an approximation of the forward rate, as in Bodurtha and Courtadon (1987). US interest rate data are daily Eurodollar rates for 1 day, 1 week, 1 month, 3 month, 6 month, and 1 year, obtained through Datastream . Linear interpolation between the two closest maturities along the yield curve is used to obtain the Eurodollar rate corresponding to the options’ maturity. For example, if an 18-day rate is required then we use a weighted average of the 1-week and 1-month rates. Brazilian interest rate data is computed using covered interest rate parity, using the appropriate futures contract (whose maturity normally coincides with that of the options), spot exchange rate (again the mid-point of bid and ask), and computed Eurodollar rate. Since we have closing spot and U.S. interest rate data, we use the last traded futures contract in each day’s calculations. Finally, exchange rate band information was obtained from the World Bank.

Macroeconomic indicators used in Section V to determine economic explanatory variables are drawn primarily from International Financial Statistics (IFS) by the International Monetary Fund. The choice of variables, follows Rose and Svensson (1994). The “real exchange rate" is constructed from the nominal exchange rate (IFS code ...rf), the US PPI (IFS code 63BB.ZF), and the Brazilian WPI (IFS code 63.Z.CF). “Output" is represented by industrial production (IFS code 66...b for the United States, Data Stream code BRINPRODH for Brazil). “Inflation" is the percentage change in the consumer prices (IFS line 64...x). The “trade balance" is the ratio of exports to imports (IFS line 70 divided by line 71); "Reserves" are foreign exchange excluding gold (IFS code 1l.d) and “Money" is Reserve Money (IFS line 14).

In Table 1a, we report the mean and standard deviation of strike price over three maturity ranges (i.e. 1-30 days, 31-60 days, and 61-90 days) and four time periods corresponding to different exchange rate regimes. Maturities vary because unlike over-the-counter option contracts, which have a fixed time-to-expiration, BM&F standardized options and futures contracts settle on the first business day of the maturity month.[4] Note that especially in the first time period (March 10, 1995 – June 22, 1995) even the mean strike price was often outside the band.

In Table 1b, we report the distribution of these strike prices over time relative to the spot, forward, and upper-band. The concentration of strike prices is important for two reasons. First, it indicates in what exchange rate range market liquidity and interest were greatest. Second, it will affect the reliability of the PDF we extract from these data. Generally, the PDF is most reliable in ranges spanned by the observed strikes. Notice in Table 1b that the distribution of strike prices has become more concentrated over time: the percentage of strikes above the spot is increasing over the four periods, but the percentage above the upper-band is decreasing over the four periods. To the degree that market activity reflects a concentration of expectations (to be verified more formally later in the paper), this pattern suggests that market expectations are exhibiting less dispersion over time, and that the upper-band is becoming increasingly credible (as indicated by decline in the percentage of strikes exceeding the upper-band).

The behavior of the underlying Real-dollar exchange rate also appears to have shifted over these periods. Table 2 reports the standard deviation of daily changes in the spot and forward rates. These standard deviations have decreased over the first three periods, a pattern coincident with less dispersed expectations as suggested by the increased concentration in observed strike prices.

IV. The Implied PDF and Expected Deviations from the Crawling Peg

A. Estimation of the Risk-Neutral PDF (over 15-day periods)

We first use our option data to derive risk-neutral PDF's corresponding to horizons of one, two, and three months. Because of data limitations, this procedure will require certain numerical approximations, but the resulting PDF's provide potentially richer information about expectations than simple point-estimate characterizations of expectations as provided by the forward rate or an econometric model.

---

[4] The exchange does offer a 'flexible' option contract that can be tailored to the issuers needs including style, maturity, dollar value, etc. However, data on these contracts were not available, and in any event,

A common approach to deriving the PDF from option prices characterizes Black-Scholes volatilities ("vols") implied in option prices as a function of the strike price. For any given date and time horizon, one can interpolate and extrapolate from existing implied vols to express implied vol as a continuous function of only the strike price. This function is commonly known as the "volatility smile"[5]. The function is then transformed into a continuous call price function that is twice-differentiable in strike. This approach does not require that the Black-Scholes model hold; indeed, the fact that implied vol varies with strike rather than being constant across all strikes is itself evidence against Black-Scholes assumptions. Note that the numerical technique in this volatility smile-based approach can vary, as discussed in Shimko (1993) and Campa, Chang, and Reider (1997, 1998). In Shimko (1993), the method applied in this paper, the implied volatility smile is fitted as a quadratic function of the strike. In contrast, Campa, Chang, and Reider (1997, 1998) use the method of cubic splines.[6]

Table 3 reports, by option maturity and exchange rate regime, the mean and standard deviation of the Black-Scholes implied volatilities extracted from observed option prices. A number of stylized facts are worthwhile noting. First, in all cases but one, shorter maturities are associated with high mean volatility. When normally calm markets experience occasional periods of high uncertainty expected to be temporary, implied volatility will increase, and most markedly for short-maturity options. Longer-dated options will also show a rise in volatility, but since the high-uncertainty state is not expected to continue throughout the option's remaining life, the implied volatility will reflect both the high-uncertainty period and the normal lower-uncertainty period, thereby diluting the effect of the temporary high-volatility period. Second, by similar reasoning, the standard deviation of short-dated volatility will be the highest, since longer-dated volatility will again reflect an average of high-volatility and low-volatility periods. Third, the mean implied volatility is one to two orders of magnitude greater than realized volatility obtained from the time series of exchange rate changes. This is because implied volatility reflects the presence of "peso problem"—the risk of a rare but substantial exchange rate shock, in this case a devaluation of the Real. Throughout the periods in question, the Real has remained stable, or depreciated only gradually against the dollar. For the most part, in our very brief sample, the Real has avoided the large price movements reflected in options' implied

---

given the potentially unique structure of each contract, each observation would have to be individually evaluated. Furthermore, low liquidity would reduce the reliability of such data.

[5] Volatility plotted as a function of the strike price often resembles a "smile" because Black-Scholes implied volatilities tend to increase as the strike price moves away from the forward rate.

[6] For a number of dates, we also fit a cubic spline (as in Campa, Chang, and Reider (1997, 1998)) to the data, and obtain similar results to the quadratic, suggesting that the results are robust to the method used.

volatility. Of course, in small samples, realized volatility can very easily be substantially below implied volatility in the presence of a "peso problem."

In our attempt to derive a PDF from the Brazilian options data, a significant empirical problem is that, for any given observation date and maturity date, we observe an insufficient range of strike prices to trace out a reasonably complete volatility smile. This prevents us from constructing daily estimates of the PDF on all but a few dates. Also, as mentioned previously, options expire on the first business day of every month, reducing the frequency to only monthly if we wish to compare PDFs with the same time horizon. To overcome these data limitations, we make the assumption that the shape of the volatility smile remains constant for a period of 15 days. For convenience, we assign the period's midpoint as the "observation date" for each 15-day period. For instance, for 60-day call options, implied volatilities are collected for options ranging from 53 to 67 days to expiration. Each volatility corresponds to a strike/forward ratio for the collection period. We convert each strike/forward ratio to an absolute price by multiplying by the forward rate central to the period. The implicit assumption is that during this period, the relationship between volatility and the strike/forward ratio remains constant.

Aggregating option observations over such 15-day periods, we obtain a semi-monthly series of PDF's for 35, 60, and 91-day call options (the 35-day periodicity captured a greater spectrum of strikes than did a 30-day). Many of these PDF's are estimated using over 20 data points on the volatility smile, and most use over 10 data points. Only in one case do we use as few as three options data points. PDF's are discarded if the associated continuous call price function is non-convex, as occurred in two instances. PDF's were also smoothed using an exponential smoothing technique, which removes non-monotonicities or negative values on the posterior and anterior slopes. When this technique is applied, if a non-monotonicity or negative value is detected, the computed PDF at this point is modified to decline from the previous value towards zero at an exponentially decreasing rate.

Figures 2a-c provide three-dimensional time series of risk-neutral PDF's estimated using numerical derivatives, for 35, 60, and 91-day options respectively. The PDF's are presented as a function of the strike/forward ratio. The first observation in any of the graphs is October 3, 1994 though the continuum of observations does not in general start until June 2, 1995. (The dates on the horizontal axes are shown in reverse to facilitate a better view of the fluctuations in the estimated distributions over time.) All time series appear to exhibit increasing skewness and decreasing kurtosis over time. Positive skewness in this context indicates that a large depreciation of the Real is more likely than a large appreciation. The increase in skewness largely stems from the disappearance of a downside tail, in the region of Real appreciation. Kurtosis, on the other hand, reflects "fatter tails," relative to the lognormal distribution. Kurtosis (above that found in the lognormal distribution) denotes a relatively high probability of extreme outcomes—holding volatility constant. These graphical results reinforce our earlier inferences from the distribution of strike prices. Positive skewness confirms that the market perceives a greater probability of a large Real depreciation than a large Real appreciation. Increasing kurtosis indicates the relative increase in very large expected exchange rate changes, and hence conditional on the level of volatility, less total probability of devaluation. Towards the end of the sample, it is striking how the part of the distribution below the forward rate is extremely concentrated in values very close to (but below) the forward rate (i.e. small Real appreciation). In contrast, for values above the forward rate, the distribution quickly drops to zero for points beyond a 2% depreciation from the forward rate. This is consistent with the government's stated policy of constant depreciation over time.

B. Deviations from the Crawling Peg (Miniband) Regime

We now use these PDF's to identify potential divergences between market expectations and the existing crawling peg regime of 0.5%-0.6% per month. We focus on possible Real devaluations of a larger magnitude than the crawling peg, namely 2% and 5% over horizons of 35, 60, and 91 days (approximately 1, 2, and 3 months). All these combinations of devaluations and time horizons represent a rate of Real depreciation at least as fast as under the crawling peg, and usually more so. For example, the existing crawling peg would imply about a 1.5% depreciation over three months.

For each devaluation size (x%) and horizon, we calculate the "probability" of devaluations at least x%. "Probability" denotes the total amount of probability, not weighted by distance, representing devaluations of at least x% from the current spot. Graphically, this corresponds to the area under the curve in the right-hand tail of the risk-neutral PDF beyond an x% devaluation. In contrast, "intensity" denotes the total probability, *weighted by the amount of depreciation beyond x%*, of all devaluations of at least x%.

Table 4 depicts the probability, at the start of each month, of a depreciation of at least 2% or 5% over horizons of 35, 60, and 91-days. A number of points are striking in this table. First, the credibility of the crawling peg regime has improved consistently over time. Late in the sample, probabilities of a given depreciation (from spot) are much lower than early in the sample.

Second, within any of the three exchange rate regimes, the probability of a 2% or 5% depreciation does not change markedly in the months just prior to the maxiband realignment. In the first regime, from August 1995 (when our data begin) through January 1996, depreciation probabilities remain high throughout these six months. Beginning in February 1996, the probability of depreciation drops significantly, and remains low even up to the February 1997 realignment. This does not imply that markets expected no maxiband realignment, as we will see in the following section. Yet, the options data indicate that any anticipated maxiband realignment was not expected to be accompanied by a large spot depreciation.

Third, around times of realignments (the months preceding January 1996 and January 1997) the probability of a 5% depreciation is usually far smaller than that of a 2% depreciation. This indicates that the probability mass of a depreciation of 2% or more arises primarily from expected small depreciations—i.e. between 2% and 5%--rather than expected large depreciations of 5% or more. Thus, even when the crawling peg regime is not perceived as fully credible by the market—i.e. some depreciation beyond the usual 0.5%-0.6% per month is expected, much of the market's "doubt" surrounding the crawling peg regime is in the form of minor rather than major expected depreciations beyond the crawl.

Fourth, our estimates of depreciation risk prove extremely sensitive to news affecting the Brazilian economic and political situation. Probabilities of large depreciations increased considerably in April and May of 1996. This coincided with a humiliating defeat suffered in Congress by the Brazilian government on Social Security Reform, a key part of the structural reforms under the Real Plan. Likewise, in May 1997, bribery accusations against some Congress members resulted in a sharp temporary rise in depreciation probabilities. Table 5 lists certain key economic and political events that may have played a role in the market's perception of depreciation risk over this period.

V. Empirical Findings: Tests of Exchange Rate "Maxi-Band" Credibility

A. Arbitrage-Based Tests (Daily Observations)

We now focus on Brazil's "maxibands" and perform a number of tests, including the arbitrage-based tests of band credibility using Tests 1 and 2 (equations (7) and (8) respectively) discussed in Section I of this paper. We start by focusing simply on the behavior of the spot and three ranges of forward rates (1-30 days, 31-60 days, and 61-90 days) against the band. We see in Figure 1a that there is no violation of the upper-band by the spot. In Figure 1b, for 1-30 day data, an ongoing violation of the upper-band by the forward occurs only in the first target zone regime, although the longer-dated forward prices (Figures 1c-d) do approach and at times cross the upper-band in other regimes just prior to subsequent adjustments.

The options-based test results are graphed in Figures 3a-l. Figures 3a-f report the results from Test 1 for options in three different maturity ranges (1-30 days, 31-60 days, and 61-90 days), while Figures 3g-l report the results of Test 2 for the same maturity ranges. We report two figures for each maturity. The first figure plots the observed price of the call option on a given date minus the corresponding "maximum" consistent with credibility from all the options with the relevant maturity range observed that day. Positive values for Tests 1 and 2 constitute a violation of upper-band credibility. If there are multiple call options observed on a given date, then only the maximum such statistic for each date is reported. On some days, these maxima include some calls whose strikes exceed the upper-band, i.e. automatic violations of the target zone. Recall that since these are arbitrage-based tests, a single option can be sufficient to reject credibility. The second figure reports for each day the percentage of the observed options resulting in a rejection of credibility, indicating the concentration of market liquidity in the non-credible area. This approach does not mix calls with different expiration dates, as these arbitrage-based tests specifically refer to a given band width and time horizon.

Credibility of the exchange rate band is consistently rejected for the initial months of the exchange rate band. During all of 1995 and until the exchange rate realignment of January 30, 1996, options with maturities beyond 30 days were consistently priced higher than their maximum value consistent with credibility. During this period, there also existed a large number of options traded with strike prices larger than the existing upper band, i.e. automatic violations of credibility. Using options with maturities less than 60 days, we find credibility harder to reject from February 1996 until about November 1996, with the exception of a few days around August 1996 coinciding with the turmoil caused by the resignation of the Argentinean Finance Minister. Options with longer maturities (more than 60 days) rejected the credibility of this exchange rate band slightly earlier, starting around mid-summer 1996.

After the realignment of January 30, 1997, credibility of the exchange rate band could still be rejected. Yet, the percentage of traded options whose price was inconsistent with credibility of the new band declined significantly, and remained stable through the end of the sample on July 30, 1997, at around 20% of the traded options.

B. Probability and Realignment Intensities of the Maxibands

As we did with the minibands above, we compute the estimated monthly probabilities and intensities of devaluation (reported in Table 6) implied by the estimated PDF's at the three different horizons. Devaluation probabilities were consistently large at all horizons during the first part of the sample ("Regime II"), until the realignment of January 30, 1996. After that devaluation, probabilities were very close to zero until about November 1996, four months prior to the February 18, 1997 realignment, when the probabilities of devaluation began to steadily increase again.

Realignment intensities in Table 6 are expressed on an annualized basis as a percentage of the existing upper band. These numbers refer the product of the probability of a devaluation and the expected size of the devaluation (measured from the upper band). At the beginning of January 1996, the estimated 35-day devaluation intensity was slightly higher than 10% annually. This suggests, for instance, a 50% probability of a 2% depreciation of the spot rate beyond the upper band over a 35-day horizon. This number seems plausible given the government's policy of aiming for a steady monthly nominal devaluation of the Real of about 0.5-0.6%. The low realignment intensities observed prior to the following realignment on February 19, 1997 corroborates this point. Estimated three-month realignment intensities at the beginning of February 1997 are 2.75% while the estimated probability of the devaluation was almost 98 percent. This again indicates that, although a realignment was widely expected, the expected devaluation of the spot rate from such a realignment was very small and of the same order as the observed depreciations in the previous months.

Like our estimates of expected depreciations beyond the crawling peg, probabilities and intensities of realignments (devaluations beyond the maxibands) also prove sensitive to news affecting the Brazilian economic and political situation. For example, the failure to pass Social Security Reform legislation (April-May 1996) and the Congressional bribery scandal (May 1997) increased both realignment intensity and realignment probability, especially at the 91-day horizons.

VI. Economic Determinants of Realignment Intensity

To ascertain whether variation in realignment intensity can be explained by common macroeconomic variables, we perform regressions whose dependent variable is the monthly estimates of devaluation intensity and its lower bound, as estimated in Section V. The macroeconomic variables used are similar to those in Rose and Svensson (1994). No lagged right-hand-side variables were included, however, because of the limited number of left-hand-side observations available.

The specific equation used is:

$$\text{Intensity}_t = \alpha + \beta_1(\text{RER})_t + \beta_2\,(\text{Infl})_t + \beta_3(\text{Output})_t + \beta_4(\text{Trade})_t + \beta_5(\text{FRES})_t + \beta_6(\text{Money})_t + \varepsilon_t \qquad (9)$$

The explanatory variables on the RHS are:

- the real exchange rate (RER), determined using the nominal monthly average exchange rate, the US PPI, and the Brazilian WPI;
- cumulative inflation (Infl), which is the difference between the Brazilian and US CPI's;
- Brazilian output divided by US output (Output);
- Brazilian trade balance divided by the US trade balance (Trade);
- Brazilian foreign reserves divided by US foreign reserves (FRES);
- and the ratio of Brazilian high-powered money to its US counterpart (Money).

All variables except inflation are expressed in logs. On the left-hand-side, we use devaluation intensity derived from the full estimated PDF.

Results from OLS regressions using equation (9) for the 35-day, 60-day, and 91-day intensity data are reported in Table 7.[7] We should first note the low power of these regressions owing to the small number of observations in our sample. The regression results clearly indicate the low explanatory power of these macroeconomic variables. For the specification using the 90-day realignment intensity, we can not reject the hypothesis that all the coefficients equal zero. None of the indicators is significant in all three regressions. Money is the only variable that has a significant coefficient in more than one regression—with higher money growth associated with higher realignment intensity. The coefficients on Trade and on Reserves do have the expected sign and are significant in the regression of the 60-day intensity. Increases in the Brazilian trade deficit and decreases in its level of reserves appear to increase the intensity of realignment. These economic linkages are *not* confirmed in regressions of the two other horizons' intensity, where the coefficients are insignificant and the sign changes.[8] Given the small number of observations, it is not appropriate to draw general conclusions from these estimates. Nevertheless, the results are consistent with the general conclusions of Svensson and Rose (1994) and Campa and Chang (1998): that macroeconomic variables are largely unable to explain intertemporal movements in realignment risk.

---

[7] The results reported here do not change qualitatively if one replaces the dependent variable (realignment intensities) with either the probabilities of depreciations reported in Table 4 or the probabilities of devaluations reported in Table 6.

[8] We performed similar regressions using the average monthly minimum intensity of realignment computed according to equation (6) and the results were equally unsuccessful.

VI. Conclusion

This paper has used a new data set of exchange-traded options from August 1995 through July 1997 to derive risk-neutral probability density functions for the Real/Dollar exchange rate over horizons ranging from one to three months. The PDF is a superior indicator to a single point estimate of exchange rate expectations, such as a forward rate or survey-based forecast, in that it assigns varying amounts of probability to different possible outcomes. Although we introduce some approximations to compensate for sparse data, we make no assumptions about exchange rate dynamics. The PDF then can be used to analyze both the crawling peg and the maxiband exchange rate regimes. These two overlapping systems have been in operation in Brazil since early 1995, several months after the June 1994 introduction of the Real Plan, designed to combat inflation and currency depreciation.

In assessing market expectations under the crawling peg, we use the risk-neutral PDF to calculate both the intensity and probability of depreciation beyond the crawling peg. A high probability accompanied by a relatively low intensity, for example, indicates that the market anticipates depreciation beyond the peg, but most of this depreciation is concentrated just outside the peg. Empirically, we find that the credibility of the peg has increased over time, and that the occasional spikes in depreciation intensity and probability can usually be explained by identifiable political or economic news in Brazil.

Our evaluation of the maxiband regime consists of two arbitrage-based tests of target zone credibility, as well as a measure of devaluation intensity outside the band. Tests based on arbitrage reject credibility whenever observed option prices are inconsistent with zero probability lying outside the band. When this occurs, devaluation intensity outside the band is positive. The numerical value of this intensity then provides a quantitative indicator of markets' questioning the maxiband regime. Empirically, we are usually able to reject credibility, but find that through our sample ending in July 1997, the intensity of devaluation has fallen over time as the regime became increasingly credible.

This paper also provides a more general methodology for extracting the risk-neutral PDF even when data are limited. In particular, we aggregate observations over several days, normalizing the option price by the contemporaneous forward rate. Our method involves fitting a single volatility smile to these multi-day observation periods. Assuming stationarity of the distribution over each period, this approach results in more precision when relatively few options are observed, a common difficulty with many emerging markets.

Analysis of the shape of the PDFs over time also provides insight into market perceptions. In general, the PDFs appear to exhibit a greater degree of kurtosis and skewness (towards Real devaluation) with time. Increased kurtosis, i.e. fatter tails for a given level of volatility, suggests that increasingly markets believed that *if* a depreciation were to occur, it would be a large depreciation. Holding volatility constant, an increase in kurtosis implies less probability of a devaluation outside the target zone, but a larger expected devaluation if devaluation occurs.

We also run regressions seeking to identify macroeconomic determinants of realignment risk. We find little evidence that standard macroeconomic indicators can explain observed realignment risk, consistent with Rose and Svensson (1994) and Campa and Chang (1998). Our observation of increasing kurtosis over time suggests that devaluation outside the band is increasingly perceived as a rare large event, rather than a more likely but not necessarily large event.

Overall, the paper's findings reinforce earlier work on options' superior ability, relative to macroeconomic or interest-rate based indicators, to anticipate the periodic realignments of the exchange rate bands. By providing a more sensitive indicator of exchange rate risk—either in the form of depreciation beyond the crawling peg or a realignment of the maxibands—we have also documented the steady increase in exchange rate credibility during the first years of Brazil's Real Plan.

References

Abken, P.A. (1995) Using Eurodollar futures options: gauging the market's view of interest rate movements. *Economic Review* **16** (Federal Reserve Bank of Atlanta), 10-30.

Ait-Sahalia, Yacine, and Andrew W. Lo (1995) Nonparametric estimation of state-price densities implicit in financial asset prices. Working Paper.

Bank for International Settlements (1996) *Central Bank Survey of Foreign Exchange Market Activity in April 1995*, Basel: Bank for International Settlements.

Bates, David S. (1991) The Crash of '87: was it expected? *Journal of Finance* **46**, 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**, 65-93.

Bates, David S. (1996b) Testing option pricing models. *Statistical Methods in Finance* **14567-611**.

Bertola, Giuseppe and Ricardo Caballero (1992) Stochastic devaluation risk and the empirical fit of target-zone models. *American Economic Review* **82**, 520-536.

Bertola, Giuseppe and Svennson Lars E.O. (1993) Stochastic Devaluation Risk and the Empirical Fit of Target-Zone Models. *Review of Economic Studies*, July 1993, 60(3), 689-712.

Breeden, Douglas and Robert Litzenberger (1978) Prices of state-contingent claims implicit in option price. *Journal of Business* **51**, 621-651.

Bodurtha, James N. Jr. and Georges R. Courtadon (1987) Tests of an American Option Pricing Model on the Foreign Currency Options Market. *Journal of Financial and Quantitative Analysis* Vol. 22, No.2, 154-167.

BM&F 1997 Annual Report, Bolsa de Mercadoris & Futuros, Sao Paulo, Brazil.

BM&F 1997 Facts and Figures, Bolsa de Mercadoris & Futuros, Sao Paulo, Brazil

Campa, José M. and P.H. Kevin Chang (1996) Arbitrage-based tests of target zone credibility: evidence from ERM cross-rate options. *American Economic Review* **86**, 726-740.

Campa, José M., P.H. Kevin Chang (1998) ERM Realignment Risk and Its' Economic Determinants as Reflected in Cross-Rate Options. *Economic Journal* **108,** 1046-1066.

Campa, José M., P.H. Kevin Chang, and Robert L. Reider (1997) ERM bandwidths for EMU and after: evidence from foreign exchange options. *Economic Policy* **24**, 55-89.

Campa, José 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.**

Chiras, Donald P. and Steven Manaster (1978) The information content of option prices and a test of market efficiency. *Journal of Financial Economics* **10**, 213-234.

Day, T. and C. Lewis (1992) Stock market volatility and the information content of stock index options. *Journal of Econometrics* **52**, 267-287.

Dooley, Michael P., and Jeffrey Shafer (1984) Analysis of short-run exchange rate behavior: March 1973 to November 1981. In: *Floating Exchange Rates and the State of World Trade and Payments*, eds Bigman and Taya, Cambridge, MA: Ballinger Publishing Company.

Economist Intelligence Unit (1997) 1996-1997 Brazil Country Profile. The Economist, London, UK.

Engel, Charles M. and James D. Hamilton (1990) Long swings in the dollar: are they in the data and do the markets know it? *American Economic Review* **80**, 689-713.

Grundy, Bruce D. (1991) Options prices and the underlying asset's return distribution. *Journal of Finance* **46**, 1045-1069.

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)**, 327-343.

Hsieh, David A. (1988) Statistical properties of daily exchange rates: 1974-1983. *Journal of International Economics* **24**, 124-145.

Hsieh, David A. (1989) Modelling heteroscedasticity in daily foreign-exchange rates. *Journal of Business Economics and Statistics* **7**, 307-317.

Hull, John and Alan White (1987) The pricing of options on assets with Stochastic volatilities. *Journal of Finance* **42(2)**, 281-300.

Jackwerth, Jens Carsten and Mark Rubinstein (1996) Recovering probability distributions from contemporaneous security prices. *Journal of Finance* **51(5)**, 1611-1631.

Jorion, Philippe (1995) Predicting volatility in the foreign exchange market. *Journal of Finance* **50(2)**, 507-528.

Kaminsky, Graciela, Saul Lizondo, and Carmen M. Reinhardt (1997) Leading Indicators of Currency Crises. IMF Working Paper, WP/97/79, 43 pages.

Kroner, Kenneth R., Kevin P. Kneafsey, and Stijn Claessens (1993) Forecasting volatility in commodity markets. World Bank Policy Research Working Paper,no. 1226.

Lamoureux, Christopher G. and William D. Lastrapes (1993) Forecasting stock-return variance: toward an understanding of Stochastic implied volatilities. *Review of Financial Studies* **6**, 293-326.

Leahy, Michael P. and Charles P. Thomas (1996) The sovereignty option: market views on the Quebec referendum. International Finance Discussion Papers,no. 555, Federal Reserve Board.

Malz, Allan M. (1996a) 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, Allan M. (1996b) Option-based estimates of the probability distribution of exchange rates and currency excess returns. Federal Reserve Bank of New York Working Paper.

McCauley, Robert and William R. Melick (1996a) Risk reversal risk. *Risk* **9(11)**, 54-57.

McCauley, Robert and William R. Melick (1996b) Propensity and density. *Risk* **9(12)**, 52-54.

Melick, William R. and Charles P. Thomas (1997) Recovering an asset's implied pdf from options prices: an application to crude oil during the Gulf crisis. *Journal of Financial and Quantitative Analysis* **32(1)**, 91-115.

Mizrach, Bruce (1996) Did option prices predict the ERM crisis? Rutgers University Department of Economics Working Paper 96-10.

Newey, Whitney K. and Kenneth D. West (1987) Hypothesis testing with efficient method of moments estimation. *International Economic Review* **28**, 777-787.

Political Risk Services (1993) Political Risk Yearbook Volume III, South America. William D. Coplin and Michael K. Oleary, Syracuse, NY, several years.

Rosenberger, James and Miriam Gasko (1983) Comparing location estimators: trimmed means, medians, and trimean. In: *Understanding Robust and Exploratory Data Analysis*, eds David Hoaglin, Frederick Mosteller, and John Tuckey, New York: John Wiley.

Rubinstein, Mark (1994) Implied binomial trees. *Journal of Finance* **49**, 771-818.

Scott, Louis O. (1992) The information content of prices in derivative security markets. *IMF Staff Papers* **39**, 596-625.

Shimko, David (1993) Bounds of probability. *Risk* **6**, 33-37.

Svensson, Lars E. O. (1991) The Simplest Test of Target Zone Credibility. IMF Staff Papers, September 1991, 38(3), 655-65.

Svensson, Lars E. O. (1992) An interpretation of recent research on exchange rate target zones. *Journal of Economic Perspectives* **6**, 119-144.

Table 1a: Mean and Standard Deviation of Strikes

This table reports the mean and standard deviation (in parenthesis) of strikes for European call data in three maturity ranges (1-30 days, 31-60 days, and 61-90 days) for the period 3/95 to 7/97. The four periods over which these statistics are computed correspond to different exchange rate band regimes: March 10, 1995 through June 22, 1995 (.88-.93 R/$), June 23, 1995 through Jan 30, 1996 (.91-.99 R/$), Jan 31, 1996 through February 18, 1997 (.97-1.06 R/$), and February 19, 1997 through July 30, 1997, the end of data set (1.05-1.14 R/$). The number of observations is listed below each statistic.

	3/10/95 – 6/ 22/95	6/23/95 – 1/30/96	1/31/96 – 2/18/97	2/19/97 – 7/30/97
R/$ Band	.88-.93	.91-.99	.97-1.06	1.05-1.14
1-30 Days, Mean	1.027	1.010	1.040	1.097
Std.Dev.	(.227)	(.138)	(.051)	(.050)
# obs.	11	158	563	386
31-60 Days, Mean	.853	.985	1.047	1.101
Std. Dev.	(.066)	(.040)	(.050)	(.031)
# obs.	14	204	740	437
61-90 Days	.956	.991	1.046	1.107
Std. Dev.	(.172)	(.019)	(.050)	(.027)
# obs.	20	203	608	315

Table 1b: Percentage of Strikes Above to the Spot, Forward, and Upper-Band

This table reports the percentage of strike prices above the spot rate, forward rate and upper-bands, for European call data in three maturity ranges (1-30 days, 31-60 days, and 61-90 days), for the period 3/95 to 7/97. The four periods over which these statistics are computed correspond to different exchange rate band regimes: March 10, 1995 through June 22, 1995 (.88-.93 R/$), June 23, 1995 through Jan 30, 1996 (.91-.99 R/$), Jan 31, 1996 through February 18, 1997 (.97-1.06 R/$), and February 19, 1997 through July 30, 1997, the end of data set (1.05-1.14 R/$).

	3/10/95 – 6/ 22/95	6/23/95 – 1/30/96	1/31/96 – 2/18/97	2/19/97 – 7/30/97
R/$ Band	.88-.93	.91-.99	.97-1.06	1.05-1.14
1-30 Days				
Spot	.64	.85	.94	.99
Forward	.55	.80	.82	.96
Upper-Band	.55	.40	.28	.14
31-60 Days				
Spot	.43	.95	.97	1.00
Forward	.21	.77	.76	.89
Upper-Band	.21	.47	.33	.18
61-90 Days				
Spot	.40	.995	.997	1.00
Forward	.15	.78	.77	.87
Upper-Band	.30	.65	.36	.21

Table 2: Standard Deviation of Changes in the Forward and Spot Rates

This table reports the standard deviation of daily percent changes in the spot rate and three forward rates relative over the four Real/Dollar maxiband regimes during the sample period, 3/95 to 7/97. Observations on the forward rates are separated in three maturity ranges (1-30 days, 31-60 days, and 61-90 days). The four regimes are: March 10, 1995 through June 22, 1995 (.88-.93 R/$), June 23, 1995 through Jan 30, 1996 (.91-.99 R/$), Jan 31, 1996 through February 18, 1997 (.97-1.06 R/$), and February 19, 1997 through July 30, 1997, the end of the data set (1.05-1.14 R/$). Number of observations is provided below each statistic.

	3/10/95 – 6/ 22/95	6/23/95 – 1/30/96	1/31/96 – 2/18/97	2/19/97 – 7/30/97
R/$ Band	.88-.93	.91-.99	.97-1.06	1.05-1.14
Spot				
Std. Dev.	0.0045	0.0010	0.00073	0.00076
# obs.	69	149	259	110
1-30 Day				
Std. Dev.	.0065	.0026	.0015	.0017
# obs.	67	146	247	107
31-60 Day				
Std. Dev.	.0071	.0024	.0016	.0018
# obs.	67	147	251	107
61-90 Day				
Std. Dev.	.0079	.0026	.0018	.0018
# obs.	67	146	240	102

Table 3: Implied Volatilities from Real-U.S. Dollar Options

This table reports the mean and standard deviation of the implied volatilities (in percentage terms) of the options in the sample, 3/95 to 7/97. Observations on the options are separated into three categories defined by maturity (1-30 days, 31-60 days, and 61-90 days). The four periods over which these statistics are computed correspond to different exchange rate maxiband regimes: March 10, 1995 through June 22, 1995 (.88-.93 R/$), June 23, 1995 through Jan 30, 1996 (.91-.99 R/$), Jan 31, 1996 through February 18, 1997 (.97-1.06 R/$), and February 19, 1997 through July 30, 1997, the end of the data set (1.05-1.14 R/$). The number of observations is provided below each statistic.

	3/10/95 – 6/ 22/95	6/23/95 – 1/30/96	1/31/96 – 2/18/97	2/19/97 – 7/30/97
R/$ Band	.88-.93	.91-.99	.97-1.06	1.05-1.14
1-30Day				
Mean	46.64	13.65	5.06	5.75
Std. Dev.	29.20	19.85	6.68	6.43
# obs.	11	151	520	385
31-60Day				
Mean	18.74	4.57	4.13	4.23
Std. Dev.	7.83	5.96	4.48	3.41
# obs.	14	202	702	433
61-90 Day				
Mean	19.85	4.37	3.48	3.59
Std. Dev.	11.74	2.32	4.19	2.56
# obs.	20	201	581	311

Table 4: Probabilities of a 2% and of a 5% depreciation over 35, 60, and 91-day horizons, 8/95-7/97.

This table reports the probability that the expected exchange rate will depreciate by more than 2% and 5% over a given horizon. These probabilities are estimated monthly from implied PDFs at three different horizons (35, 60 and 91-days).

Date	35 Day		60 Day		91 Day	
	2 %	5%	2 %	5%	2 %	5%
Regime II: [.91-.99]						
Aug-95	3.32	0.51	12.38	2.98	21.78	9.47
Sep-95	25.39	19.38	12.00	4.04	12.35	1.90
Oct-95	31.48	24.64	32.43	26.27	9.68	1.82
Nov-95	15.88	12.70	35.18	24.59	8.26	1.57
Dec-95	28.00	18.80	10.79	7.75	2.74	0.24
Jan-96	22.34	14.56	3.76	0.86	4.72	0.56
Feb-96	0.50	0.02	2.12	0.17	3.50	0.37
Regime III: [.97-1.06]						
Mar-96	2.95	2.25	2.76	0.78	4.42	1.04
Apr-96	0.20	0.01	0.03	0.00	11.35	9.16
May-96	0.52	0.06	0.82	0.08	37.15	28.82
Jun-96	0.43	0.02	7.09	5.05	2.01	0.18
Jul-96	0.06	0.00	0.06	0.00	1.98	0.15
Aug-96	0.09	0.00	0.00	0.00 .	.	
Sep-96	0.66	0.04	2.26	0.20	2.42	0.19
Oct-96	0.00	0.00 .	.		0.02	0.00
Nov-96	0.00	0.00	0.30	0.00	0.84	0.01
Dec-96	1.34	0.07	0.12	0.00	0.00	0.00
Jan-97	0.06	0.00	6.28	2.17	2.35	0.09
Feb-97	0.33	0.00	0.89	0.00	0.00	0.00
Mar-97	0.33	0.00	0.01	0.00	1.96	0.22
Regime IV: [1.05-1.14]						
Apr-97	0.43	0.00	0.46	0.00	0.11	0.00
May-97	0.03	0.00	1.81	0.09	65.21	53.23
Jun-97	0.19	0.00	0.83	0.00	1.12	0.00
Jul-97	0.23	0.00	0.23	0.00	0.25	0.00

Table 5: Significant Events Affecting the Real-Dollar Exchange Rate, November 1994 – July 1997

Table presents a list of significant world or Brazilian events that occurred over the period covered by the data.

Date		Event(s)
1994:	March	Unit of Real Value introduced as basis for all Brazilian financial contracts and indices
	July	Real first introduced as official Brazilian currency
	November	U.S. Congress approves GATT
	December	Devaluation of Mexican Peso
1995:	January	Continued depreciation of Mexican Peso
	February	$40 billion bail-out plan for Mexico announced
	March	First Real/Dollar Maxibands introduced [.88-.93]. Mexican Peso continues to tumble Argentina seeks $3 billion in credit lines to counter contagion effects from Mexican crisis
	April	Mexican peso shows steady appreciation / Mexican stocks start to rebound / Four largest Japanese brokerage houses announce $1 billion in losses
	May	Dollar begins to appreciate against yen / trade gap with Japan declines
	June	Realignment of the Real/Dollar Maxibands. New bands [.91-.99]. Dow Jones experiences second largest decline in history
	August	Toyota Invests $150 million in new car manufacturing facility in Brazil
	October	Mexico begins repayment of US loan package.
	November	Delays in approval of constitutional reform for Social Security, Indexation and Taxes
1996:	January	Realignment of the Real/Dollar Maxibands. New bands [.97-1.06]. FEF established by the Brazilian congress to eliminate fiscal deficit, renewed for 18 months
	February	US bond prices tumble, biggest drop in 7 months
	March	Dow experiences 3rd largest decline ever at beginning of month, then reaches record levels March 18
	May	Yields on US 30 year treasuries exceed 7% for first time in months
	June	Brazilian government suffers humiliating defeat on Social Security Reform in Congress
	July	Dow tumbles; 219 point swing in trading
	August	Argentina's finance minister is replaced raising uncertainty in emerging markets
	November	Brazilian municipal elections are held with mixed results for the party in power
	December	Peruvian terrorists seize Lima residence of Japanese Ambassador / 300 point decline and recovery of Dow
1997:	January	Mexico repays final $3.5 billion of US loan package
	February	Realignment of the Real/Dollar Maxibands. New bands [1.05-1.14]. Constitutional amendment for reelection of high officials passes lower house
	May	Scandal on the government buying some congressional votes. Transaction tax is introduced to "cool" the economy Massive speculative attack on the Thai baht
	July	The baht devalues by about 15-20 percent Philippines, Indonesia, Singapore and Malaysia widened or abandoned their existing exchange rate bands

Table 6: Probabilities of Realizations outside the Maxiband and Intensities of Maxiband Realignment, 8/95-7/97

This table reports the total probability of the expected exchange rate realizations outside the maxiband and the annualized expected intensities of realignment (as a % of the upper end of the band) from the estimated PDFs at three different horizons (35, 60 and 91-days).

Date:	Probability (in %) outside the Upper Band			Intensity of Realignment		
(start of Month)	35-Day	60-day	91-Day	35-Day	60-day	91-Day
Regime II: [.91-.99]						
Aug-95	2.01	7.82	25.63	0.19	1.04	3.52
Sep-95	24.50	15.97	35.29	10.95	2.42	2.33
Oct-95	31.97	32.85	26.40	15.16	16.58	1.48
Nov-95	15.88	39.53	30.69	7.66	15.55	1.71
Dec-95	31.79	13.22	9.80	11.67	5.02	0.39
Jan-96	36.95	22.58	33.85	10.37	1.29	1.45
Regime III: [.97-1.06]						
Feb-96	0.00	0.04	0.19	0.00	0.00	0.01
Mar-96	1.93	0.48	0.80	0.67	0.07	0.08
Apr-96	0.00	0.00	8.79	0.00	0.00	2.59
May-96	0.04	0.05	28.56	0.00	0.00	7.96
Jun-96	0.02	5.16	0.33	0.00	1.95	0.02
Jul-96	0.00	0.00	0.47	0.00	0.00	0.02
Aug-96	0.00	0.00	.	0.00	0.00	.
Sep-96	0.19	1.11	2.21	0.01	0.09	0.12
Oct-96	0.00	.	0.01	0.00	.	0.00
Nov-96	0.00	0.46	3.27	0.00	0.01	0.09
Dec-96	3.61	1.12	3.45	0.20	0.03	0.05
Jan-97	4.92	25.59	38.11	0.12	2.46	1.50
Feb-97	0.00	50.37	97.96	0.00	1.55	2.75
Regime IV: [1.05-1.14]						
Mar-97	0.00	0.00	0.11	0.00	0.00	0.01
Apr-97	0.00	0.00	0.00	0.00	0.00	0.00
May-97	0.00	0.06	54.80	0.00	0.00	20.21
Jun-97	0.00	0.00	0.02	0.00	0.00	0.00
Jul-97	0.00	0.00	0.00	0.00	0.00	0.00

Table 7: Relationship between Realignment Intensities and Fundamentals, 8/95-7/97

This table reports the estimated coefficients from OLS regressions of the estimated monthly realignment intensities on a set of macroeconomic indicators. The indicators are: RER – real/US$ real exchange rate, INFL – Brazilian inflation rate, OUTPUT – index of industrial production, TRADE – trade balance, FRES – Brazilian foreign reserves, and MONEY – high-powered money. All variables except INFL are expressed as the log of the ratio of the value for Brazil of the corresponding measure to that for the U.S. Standard errors appear in italics below each reported coefficient.

	35-Day	60-Day	91-Day
RER	0.17	0.19	-0.29
	0.97	*1.33*	*-1.27*
INFL	0.00	0.00	0.00
	-1.05	*-0.42*	*0.19*
OUTPUT	0.02	0.07	-0.02
	0.68	*2.12**	*-0.42*
TRADE	0.00	-0.02	0.02
	0.14	*-2.04**	*0.93*
FRES	-0.01	-0.05	0.01
	-0.35	*-4.32**	*0.13*
MONEY	0.04	0.04	-0.01
	*2.05**	*1.86***	*-0.37*
Adj. R2	0.38	0.60	-0.14
N. Obs.	21	25	25

Figures 1a-1d: Real/Dollar Spot Rate and 30, 60 and 90 day Forward Rates

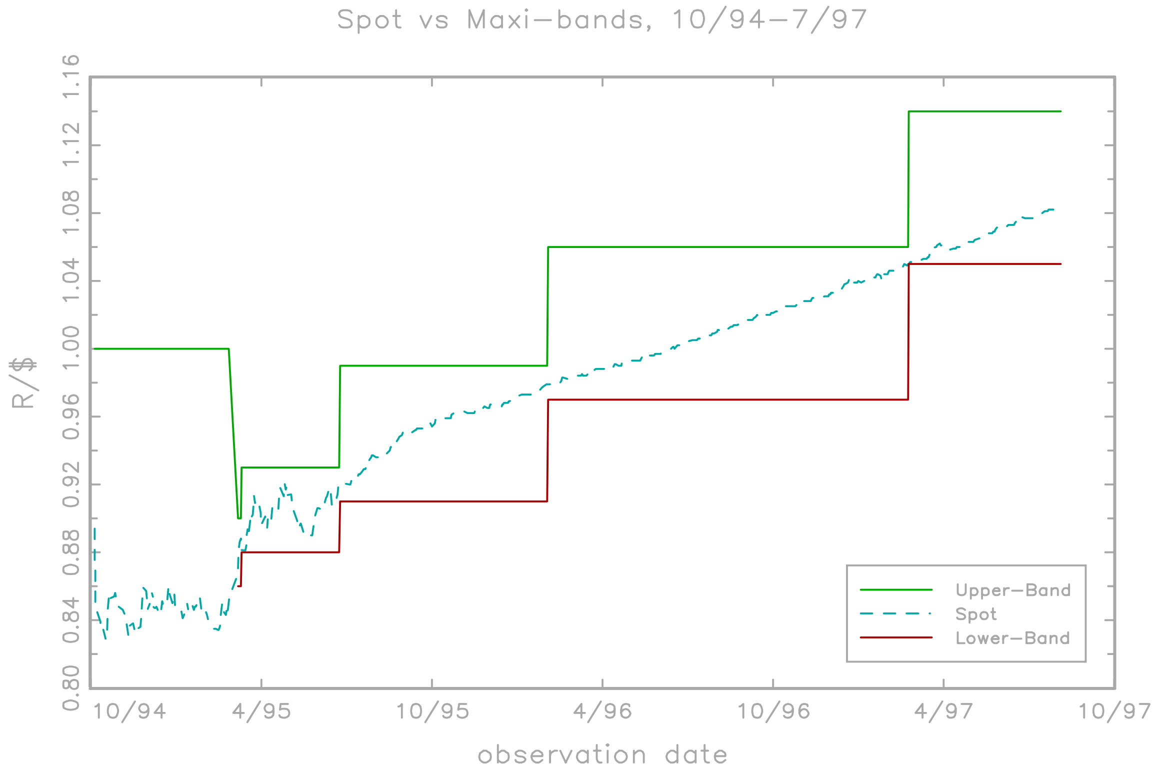

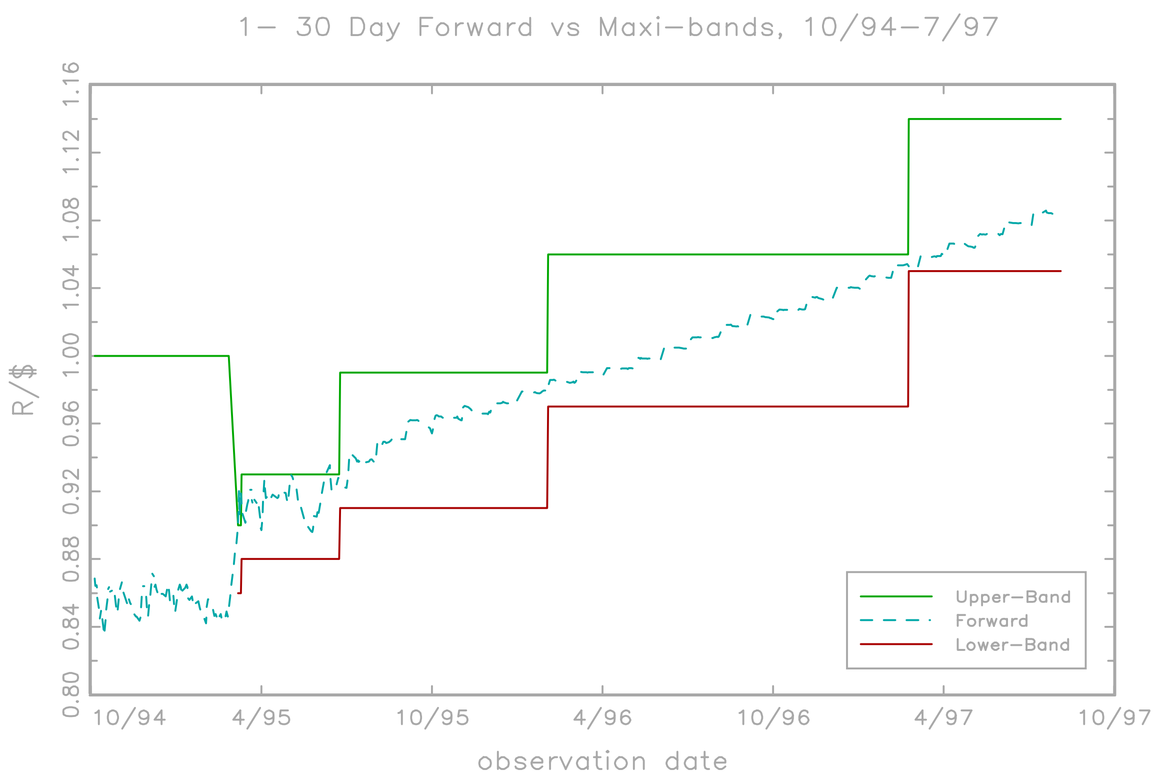

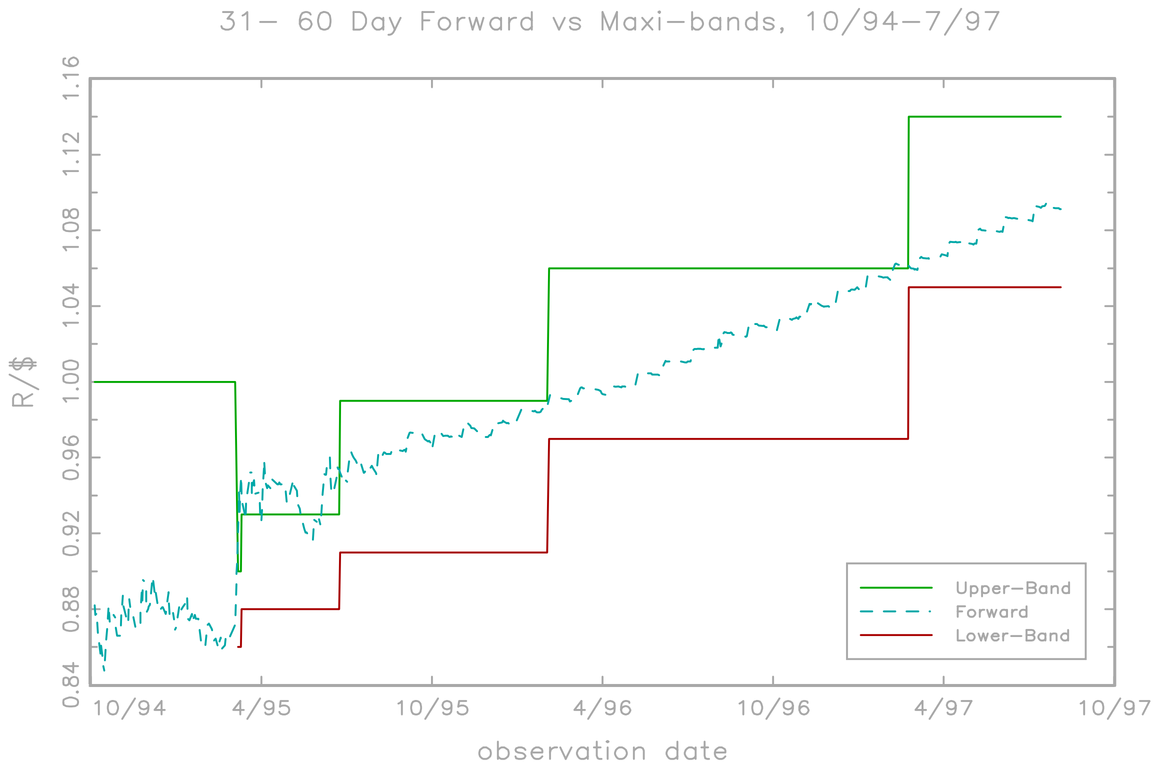
31– 60 Day Forward vs Maxi–bands, 10/94–7/97
R/$
0.84
0.88
0.92
0.96
1.00
1.04
1.08
1.12
1.16
10/94
4/95
10/95
4/96
10/96
4/97
10/97
observation date
Upper–Band
Forward
Lower–Band

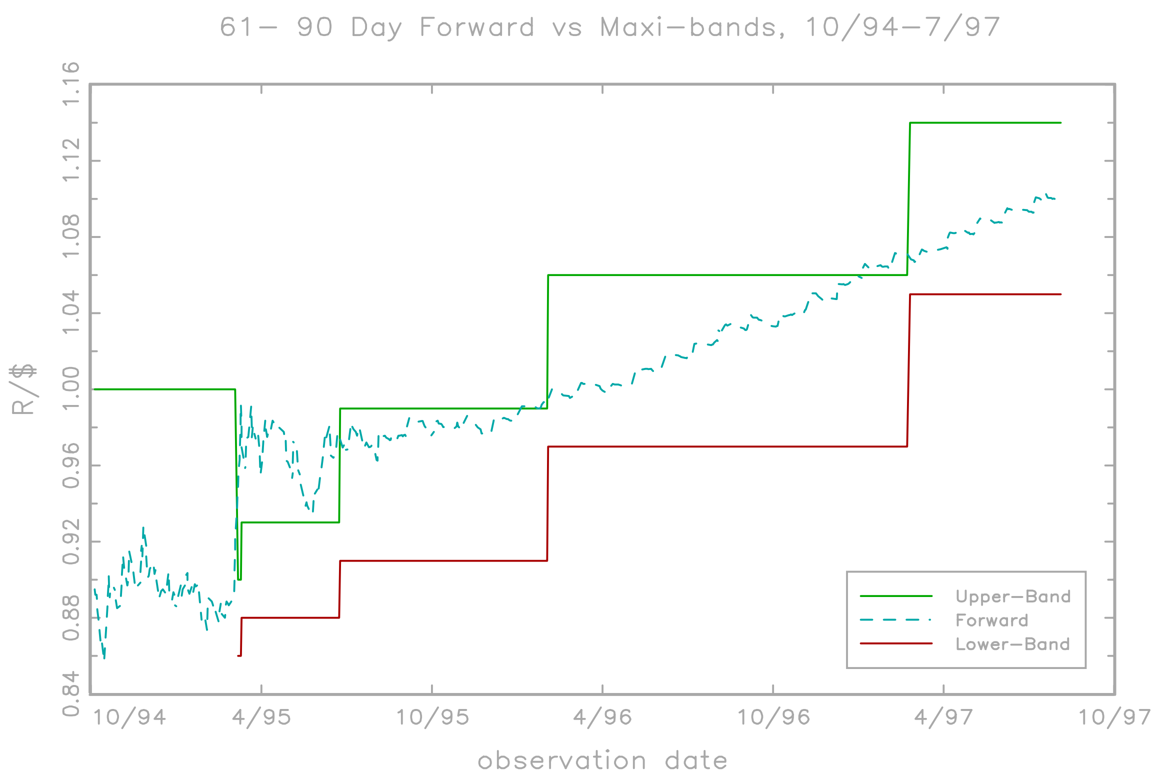
61– 90 Day Forward vs Maxi–bands, 10/94–7/97
R/$
0.84
0.88
0.92
0.96
1.00
1.04
1.08
1.12
1.16
10/94
4/95
10/95
4/96
10/96
4/97
10/97
observation date
Upper–Band
Forward
Lower–Band

Figures 2a-2c: Implied Exchange Rate Probability Distributions 10/94-7/97, 35, 60 and 91 days

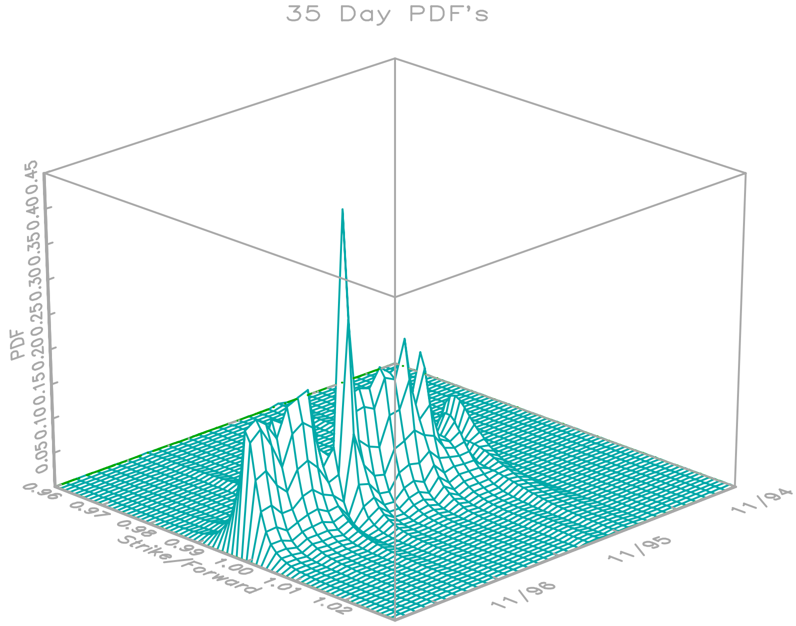

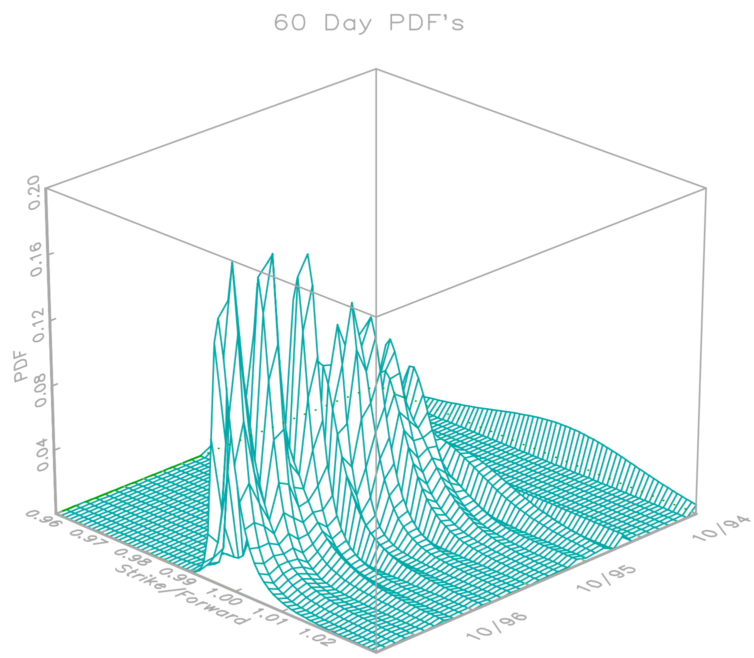

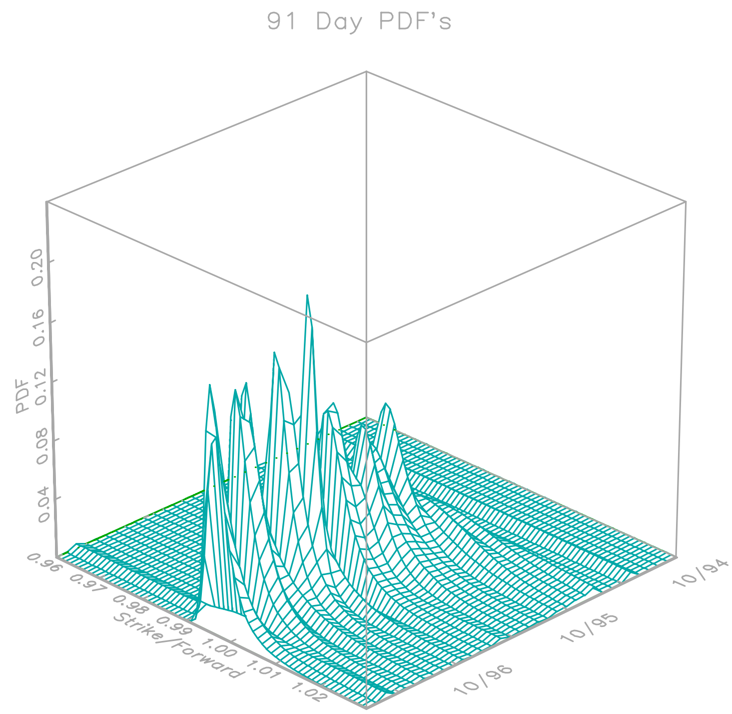
91 Day PDF's
PDF
0.04
0.08
0.12
0.16
0.20
0.96
0.97
0.98
0.99
1.00
1.01
1.02
Strike/Forward
10/94
10/95
10/96

Figures 3a-3l: Arbitrage-Based Tests of Exchange Rate Credibility

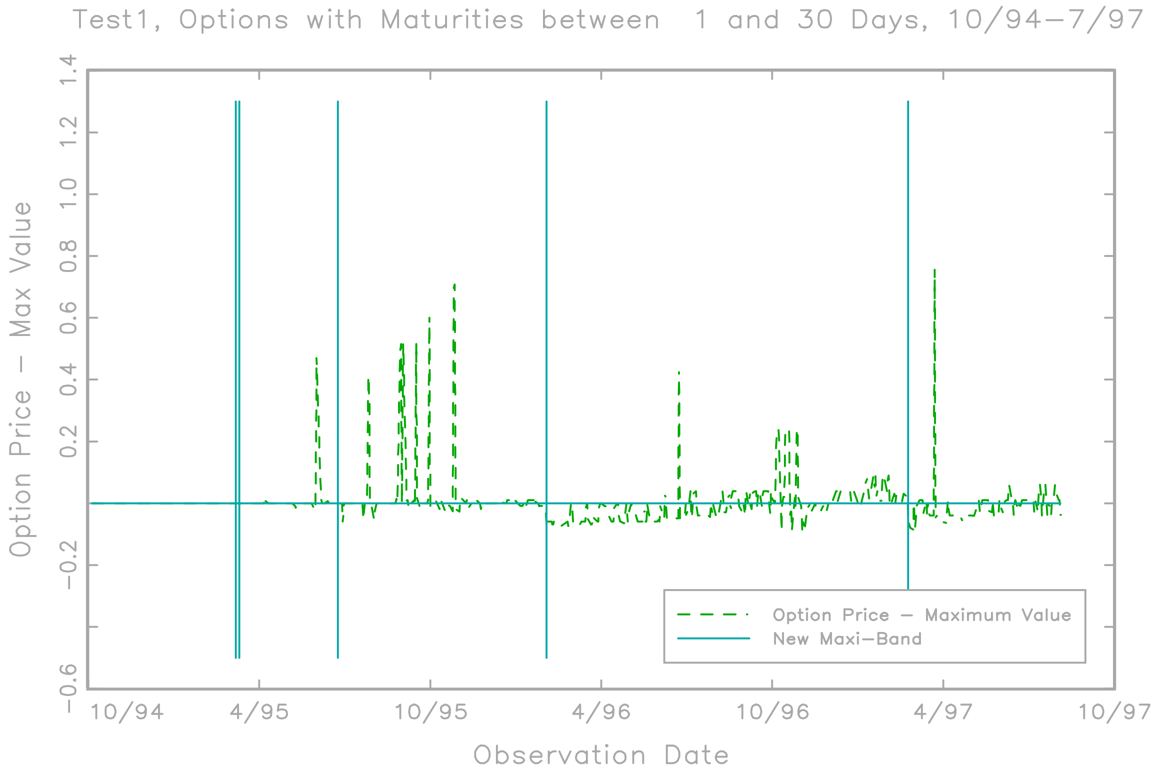

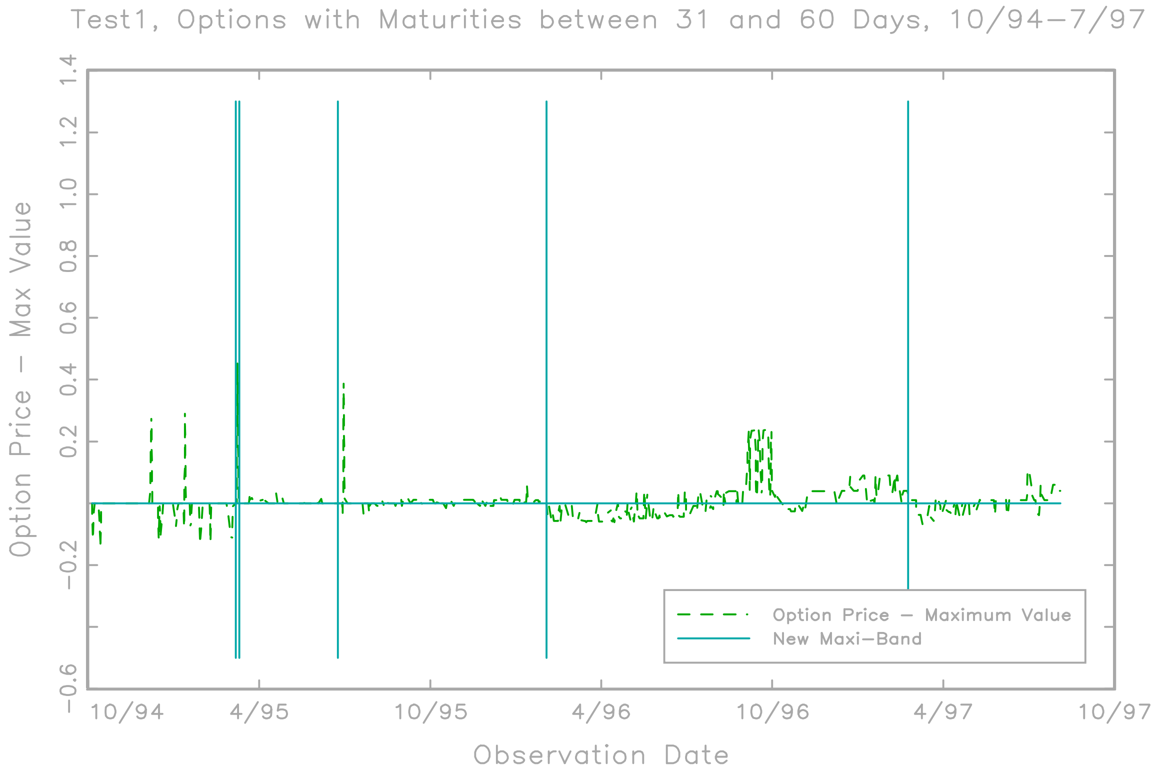
Test1, Options with Maturities between 31 and 60 Days, 10/94–7/97
Option Price – Max Value
Observation Date
Option Price – Maximum Value
New Maxi–Band
-0.6
-0.2
0.2
0.4
0.6
0.8
1.0
1.2
1.4
10/94
4/95
10/95
4/96
10/96
4/97
10/97


Fraction of Options Violating Test1
Fraction of Options Violating Test1
Observation Date
Fraction of Options Violating
New Maxi–Band
-0.6
-0.2
0.2
0.4
0.6
0.8
1.0
1.2
1.4
10/94
4/95
10/95
4/96
10/96
4/97
10/97

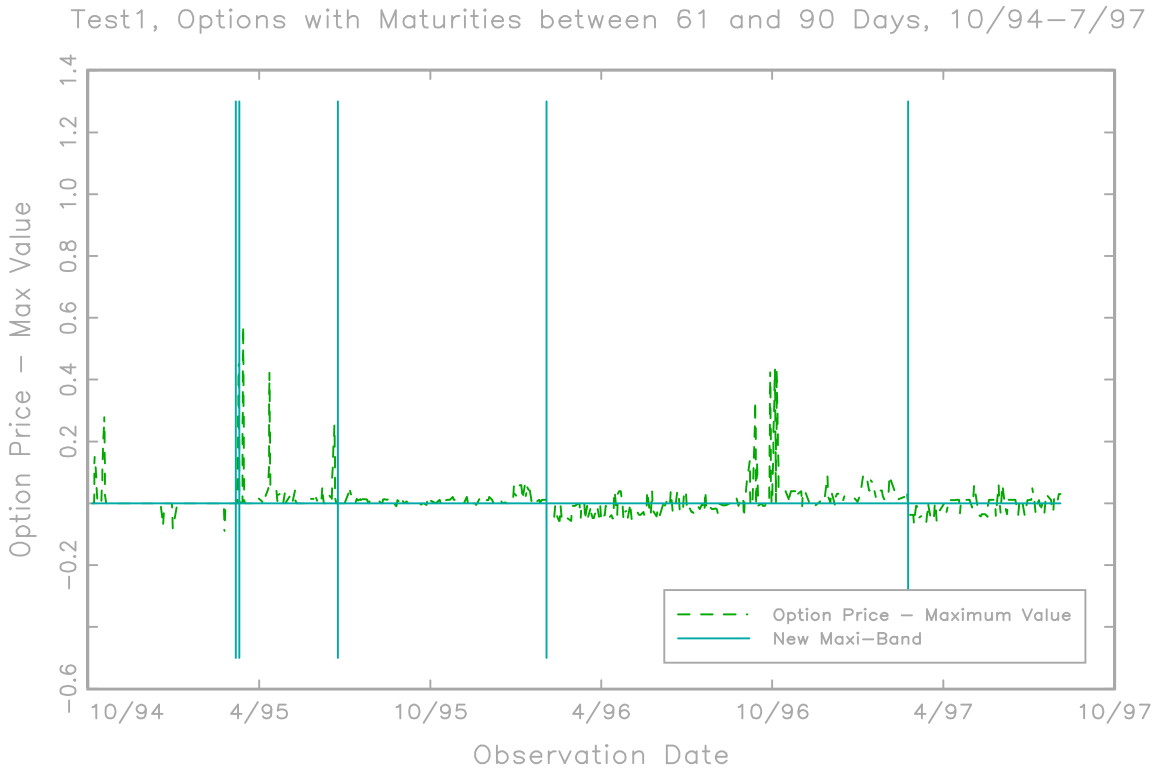
Test1, Options with Maturities between 61 and 90 Days, 10/94–7/97
Option Price – Max Value
1.4
1.2
1.0
0.8
0.6
0.4
0.2
–0.2
–0.6
10/94
4/95
10/95
4/96
10/96
4/97
10/97
Observation Date
Option Price – Maximum Value
New Maxi–Band


Fraction of Options Violating Test1
Fraction of Options Violating Test1
1.4
1.2
1.0
0.8
0.6
0.4
0.2
–0.2
–0.6
10/94
4/95
10/95
4/96
10/96
4/97
10/97
Observation Date
Fraction of Options Violating
New Maxi–Band

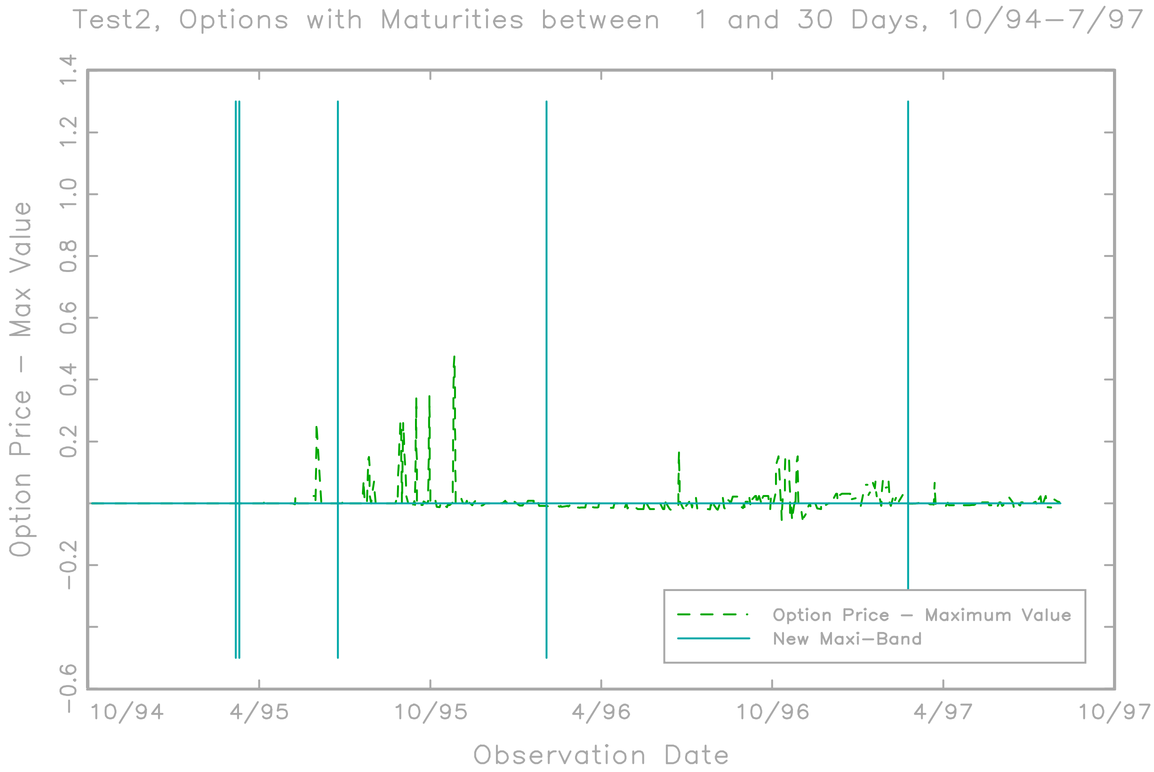
Test2, Options with Maturities between 1 and 30 Days, 10/94–7/97
Option Price – Max Value
1.4
1.2
1.0
0.8
0.6
0.4
0.2
–0.2
–0.6
10/94
4/95
10/95
4/96
10/96
4/97
10/97
Observation Date
Option Price – Maximum Value
New Maxi–Band


Fraction of Options Violating Test2
Option Price – Max Value
1.4
1.2
1.0
0.8
0.6
0.4
0.2
–0.2
–0.6
10/94
4/95
10/95
4/96
10/96
4/97
10/97
Observation Date
Fraction of Options Violating
New Maxi–Band

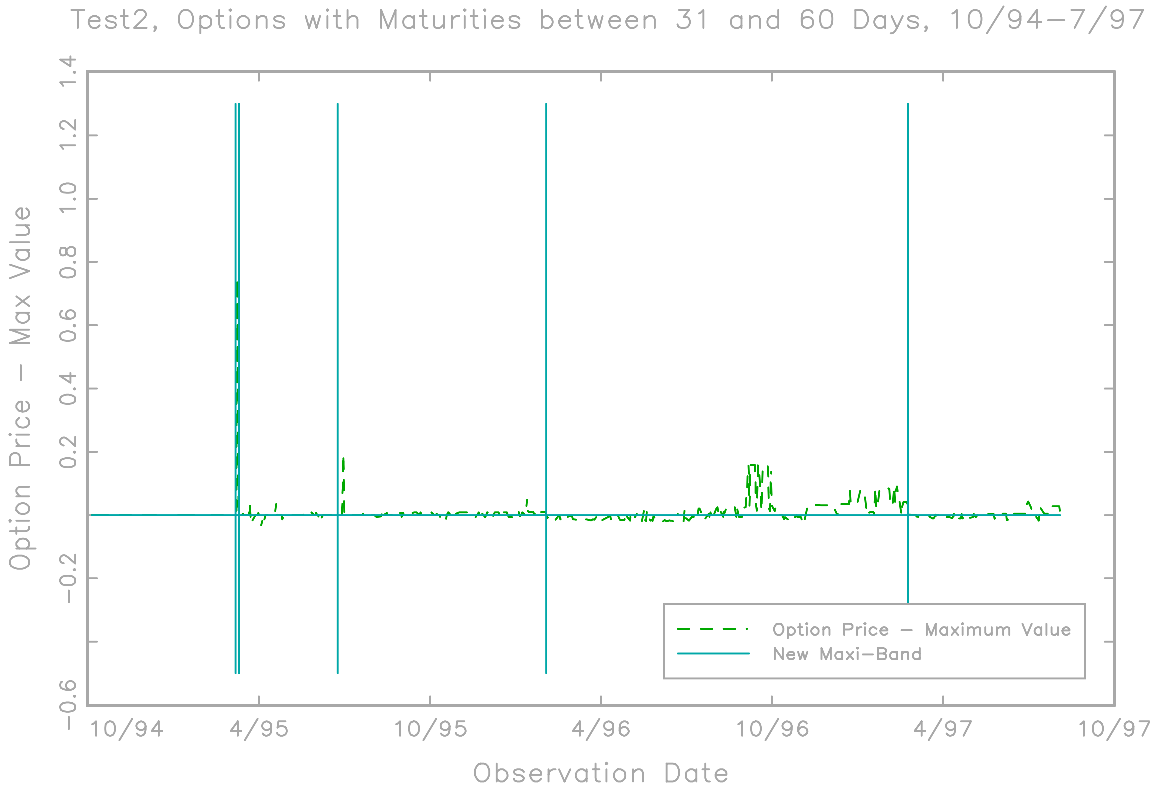
Test2, Options with Maturities between 31 and 60 Days, 10/94–7/97
Option Price – Max Value
1.4
1.2
1.0
0.8
0.6
0.4
0.2
–0.2
–0.6
10/94
4/95
10/95
4/96
10/96
4/97
10/97
Observation Date
Option Price – Maximum Value
New Maxi–Band

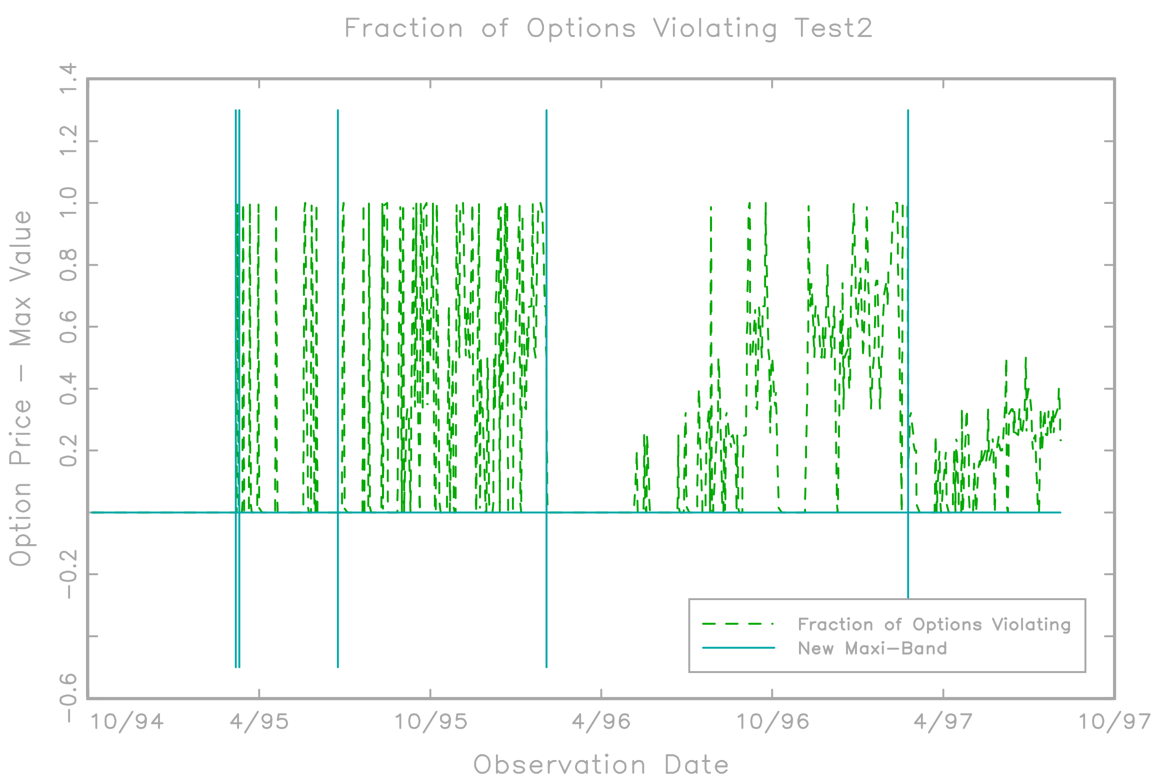
Fraction of Options Violating Test2
Option Price – Max Value
1.4
1.2
1.0
0.8
0.6
0.4
0.2
–0.2
–0.6
10/94
4/95
10/95
4/96
10/96
4/97
10/97
Observation Date
Fraction of Options Violating
New Maxi–Band

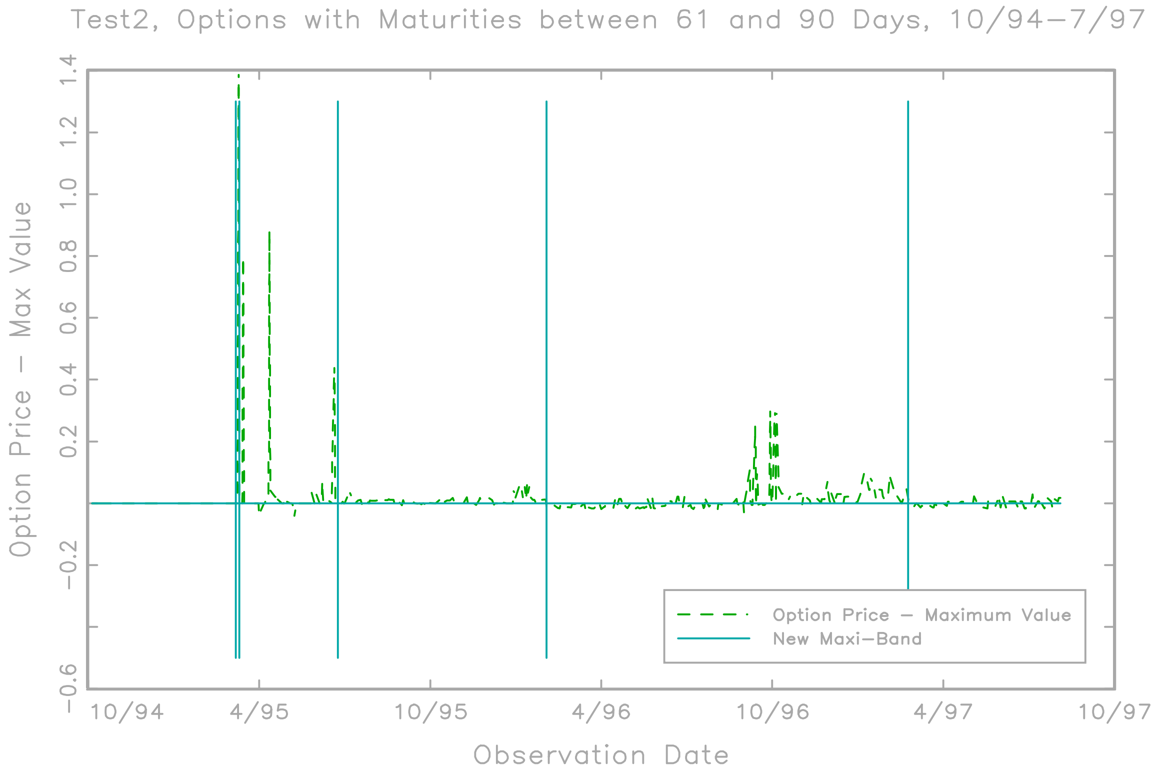
Test2, Options with Maturities between 61 and 90 Days, 10/94–7/97
Option Price – Max Value
1.4
1.2
1.0
0.8
0.6
0.4
0.2
–0.2
–0.6
10/94
4/95
10/95
4/96
10/96
4/97
10/97
Observation Date
Option Price – Maximum Value
New Maxi–Band


Fraction of Options Violating Test2
Option Price – Max Value
1.4
1.2
1.0
0.8
0.6
0.4
0.2
–0.2
–0.6
10/94
4/95
10/95
4/96
10/96
4/97
10/97
Observation Date
Fraction of Options Violating
New Maxi–Band

Figures 4a-4c: Probabilities and Intensities of Realignment, 5/95-8/97

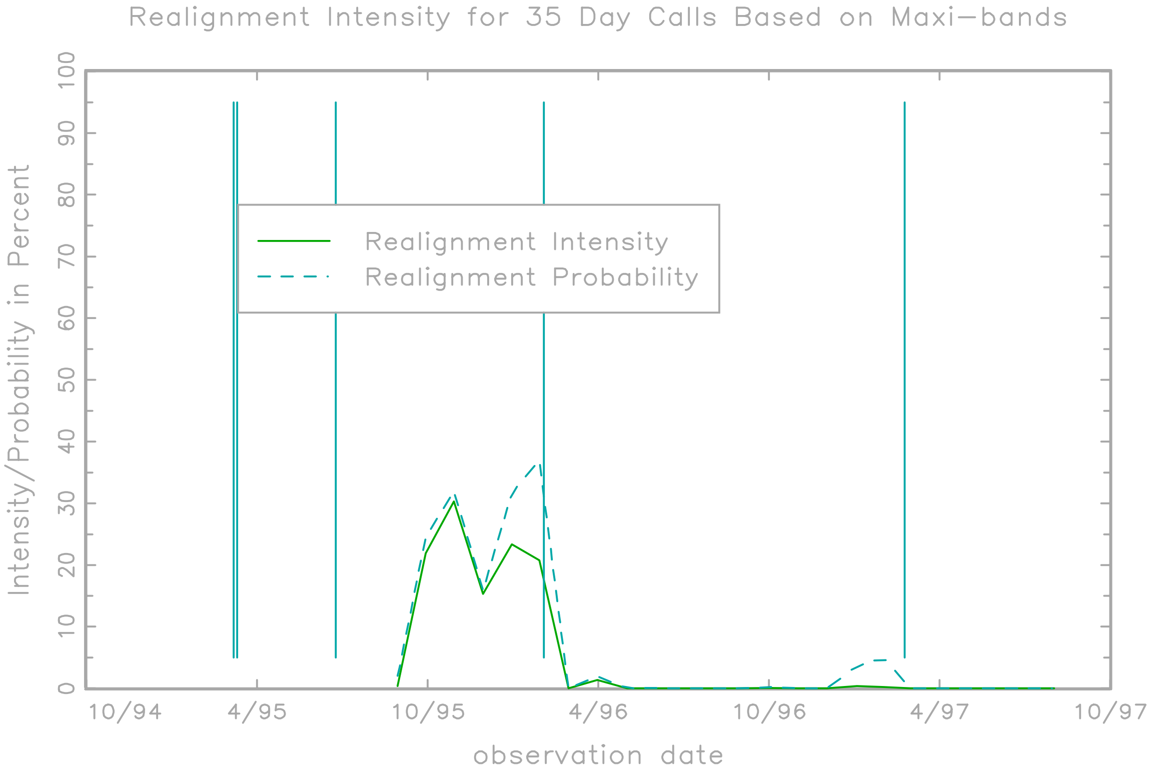

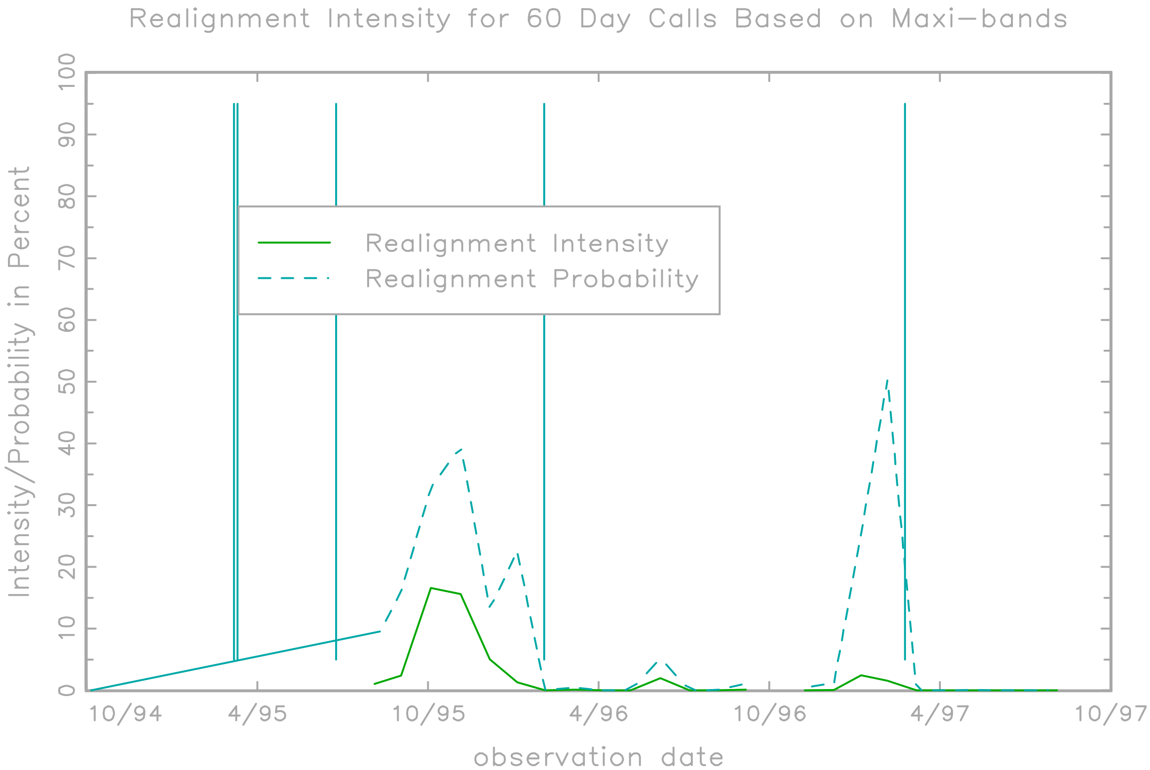

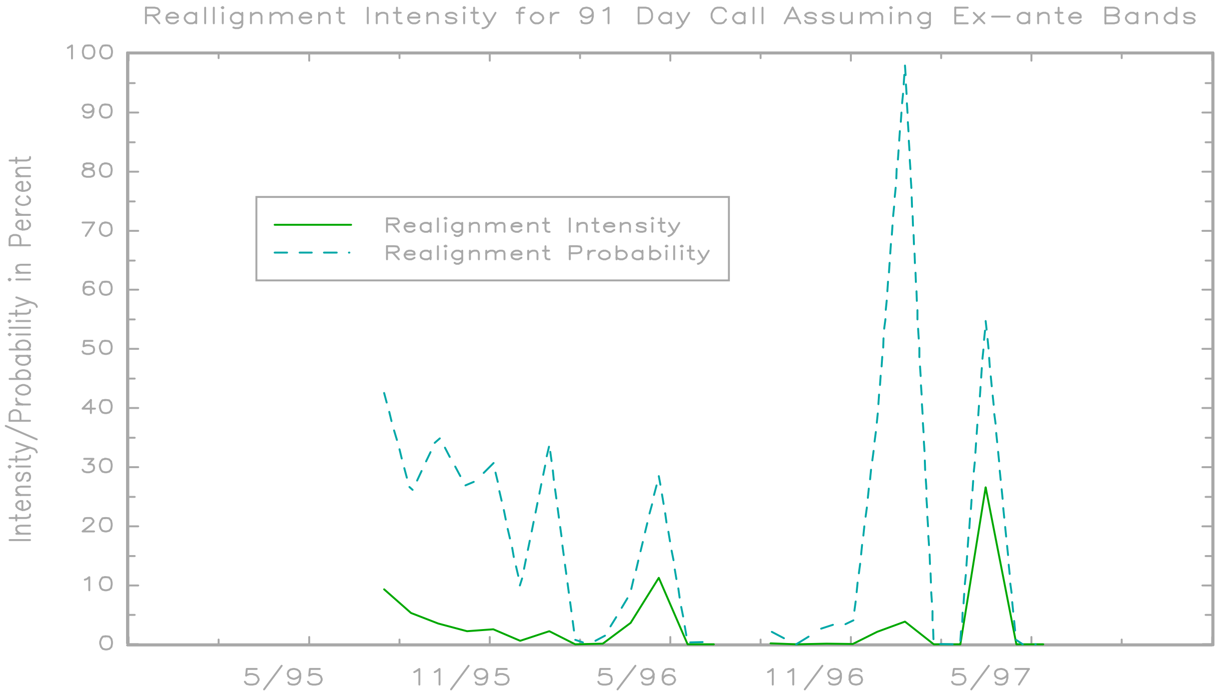
Realignment Intensity for 91 Day Call Assuming Ex-ante Bands
Intensity/Probability in Percent
100
90
80
70
60
50
40
30
20
10
0
Realignment Intensity
Realignment Probability
5/95
11/95
5/96
11/96
5/97

Discussion of the paper by Campa, Chang and Refalo

"An Options-Based Analysis of Emerging Markets Exchange Rate Expectations: Brazil's *Real* Plan, 1994 – 1997"

Paul Söderlind*

Data and the Results

This paper studies the credibility of the (moving) Brazilian exchange rate bands, the "maxiband" between early 1995 and mid 1997. The main idea is to explore the information in data on US dollar/Brazilian Real futures and options from the Commodities and Futures Exchange in Sao Paolo, Brazil. A series of biweekly PDFs of the future exchange rate (at the 30, 60 and 90 days horizon) is estimated. This gives information about how the probability of a realignment evolved over time. The data is also checked for if no-arbitrage conditions implied by complete credibility of the exchange rate band are satisfied.

The main results from the estimated PDFs are as follows:

- In general, the Real regime became increasingly more credible over time.
- However, there were temporary decreases in credibility in May 1996 (probably associated with the defeat of the social security reform in congress) and in May 1997 (probably associated with the revelation of the government's bribes to some congress members).
- When the spot rate approached the upper boundary, as in early 1997, the current band was no longer credible.

All these results are clearly visible from the estimated PDFs, in particular for the 90 days horizons. However, only the first is easily seen from the no-arbitrage test. It would be useful to add a discussion of why there are so few significant results for the shorter horizons (30 to 60 days). It would also be interesting to get some more information about the "power" of the arbitrage tests. For instance, under which circumstances can we expect these tests to perform well?

According to Figure 1 of the paper, the most dramatic movements of the spot and futures rates occurred around the establishment of the first maxiband in March 1995. This is discussed relatively little in the paper. It would be interesting to say some more about how market expectations changed around this time. It would also be very interesting to extend the sample to include more recent events (Asian crisis, collapse of the Real).

The lack of data lurks behind many of the methodological choices made in the paper. It would be interesting to see some more details on how much data there is. One possibility is to show a figure

* Stockholm School of Economics and CEPR.

with time on the horizontal axis and strike prices on the vertical and to mark all points of time and strikes where there was trade. Dates and ranges of strike prices with little trade is likely to give less precise estimates of the PDF, so such a figure would help the readers to get a grasp of the quality of the results. It might also be useful to include some more details on exactly which data is used in the estimation and tests (min, max, last traded?).

The Methods

My remaining comments are on the assumptions behind the different tests and measures for European call options used in the paper.

Denote the (future) exchange rate at the expiry of the option by S_T, the lower exchange rate boundary $\underline{S}$, and the upper $\overline{S}$. The idea of the *first arbitrage test* is that if the upper exchange rate boundary is perfectly credible, $\Pr(S_T > \overline{S}) = 0$, then a European call option with a strike prices, *K*, at or above $\overline{S}$ should be worth zero. (The distinction between the risk neutral and true distribution is not important here, since zero probability in the true distribution means zero probability in the risk neutral distribution). If no such option is available, then for $K < \overline{S}$ the highest possible value of the option price, C_K, is $(\overline{S} - K)/(1+i)$, where i is the interest rate times time to expiry. This would be the option price if all probability mass is located at the upper boundary. Any call option price above this value indicates that the band is not completely credible. This test requires a minimum of assumptions and data, and is therefore very attractive. It is unclear, however, if 100% credibility is a reasonable hypothesis. In any case, it would be very useful to illustrate the power of this test in various circumstances – perhaps with some examples.

The *second arbitrage test* is somewhat tighter, but comes at the cost of assuming that the lower boundary is completely credible, $\Pr(S_T \geq \underline{S}) = 1$. In that case an option with strike price $K = \underline{S}$ is always in the money and is therefore worth $C_{\underline{S}} = (F - \underline{S})/(1+i)$, where F is the forward price. This second test is that any option with strike price within the band must, by convexity, have a price below the straight line between $C_{\underline{S}}$ and $C_{\overline{S}}$, where the latter is zero if the band is 100% credible. The new assumption of perfect credibility of the lower boundary is probably reasonable in most cases.

The *devaluation intensity* is defined as the price of a European call option of $C_{\overline{S}}$, which can be thought of a weighted probability of realignment. If no data is available for this strike price, then a lower bound can be derived by once again assuming that the lower boundary is perfectly credible. By convexity, $C_{\overline{S}}$ must be above a straight line between $C_{\underline{S}}$ and C_K for K inside the band. This seems to be a good complement to the no-arbitrage tests in the sense that it tries to capture whether a lack of credibility is due to expectations of a small or large devaluation.

The PDFs are estimated by a modified version of Shimko's approach, where a smooth curve is fitted to the volatility smile, which is then inverted (via the Black-Scholes formula) to option prices. This option pricing function is differentiated (following Breeden-Litzenberger) twice with respect to the strike price to get the risk neutral PDF (divided by $1 + i$). This method seems to be flexible and there is evidence that it works well in many cases. However, the properties seem to be somewhat dependent on how the interpolation between the available implicit volatilities is done and also on the number/range of existing strike prices. It would be useful to discuss this in the text, and to highlight how much data has been available – and to give an assessment of the quality of the estimated PDFs (especially, the mass in the upper tail).

Discussion on "An Options-Based Analysis of Emerging Market Exchange Rate Expectations: Brazil's *Real* Plan, 1994 – 1997"

by Campa, Chang and Refalo (1999)

Disscussant: M^a Cruz Manzano

Firstly, I would like to thank the Bank for International Settlements and the organizers for the celebration of this interesting seminar.

The availability of indicators on agents expectations is a very relevant issue for monetary authorities and financial analysts, given the predominant role of those for the monetary transmission mechanism in developed financial markets. Agents assign subjective probabilities to each of the possible values of a variable in the future and, as a result, a specific probability distribution characterised and summarised the agents "feelings" about future realisations of variables. These distributions are not observed and an attempt is usually made to estimate some of the moments of the underlying distribution. Options markets, by definition of these assets, collect a very rich set of information on such distributions. But it has not been until very recently that such information has started to be analysed. The paper of Jose M^a Campa, Kevin Chang and James Refalo is an example of this growing body of analysis.

These authors exploit the information content of currency option data in terms of agents' expectations on future exchange rates and apply the analysis to the Brazilian case. The paper is an interesting contribution to the analysis of credibility of target zones regimes and, not less important, to the analysis of exchange rate developments in emerging countries.

There exists a large body of literature on credibility measures of target zone regimes both from a theoretical and an empirical point of view. But, to my knowledge, there are few studies on the case of emerging economies. It is not necessary to show the importance of these economies for world economic developments as the current world situation is stressing.

Campa, Chang and Refalo apply an analysis of credibility measures to the target zone regime governing the real –US $ exchange rate. As is explained in the paper, the Brazilian regime combine the imposition of a maxi-band- an upper and a lower band for the exchange rate- and a system of crawling peg in which movements within the band are controlled. The analysis of credibility is very close related to that of developments in expectations and in this regard, options are a privileged source of information as the authors show.

They use a PDF approach to examine the credibility of the crawling peg system and to build a realignment intensity indicator through the estimate of the exchange rate expectations distribution. In addition, they analyse the credibility of the upper band imposed on the exchange rate using some other credibility measures – based on arbitrage and convexity hypothesis on options pricing – which are explained in an article of Campa and Chang published in the American Economic Review in 1996.

Let me make some comments on only two points of the paper: one of them related to the data used and the other in relation to the conclusions reached by the authors.

The main contribution of the paper is the estimation of PDFs for expected exchange rates of the Brazilian real. Because of this, data used to estimate these functions are very relevant.

Focusing my attention on this point, I would like to stress that the analysis carried out uses daily data on options but, as is explained by the authors, because of data limitations – insufficient range of strikes, the need of comparing PDFs with the same time horizon,etc – only a semi-monthly series of probability distribution functions are obtained. Hence, options observations are aggregated over a period of 15 days making the hypothesis that the relationship between volatility and the strike forward ratio remains constant.

With the data aggregation mentioned, a volatility smile is fitted as a quadratic function of the strike and, finally, PDFs are smoothed using an exponential smoothing technique.

In my opinion, the manipulation of data carried out is a point to be carefully discussed because it could distort the analysis in a relevant way.

In fact, the estimation of *daily* PDFs is a crucial issue in markets in which news and economic and political events are rapidly incorporated into expectations and, hence, into prices. This is the case of exchange rates markets. In this regard, the aggregation carried out in the paper could be blurring relevant information because is taking into account events happened in a period of fifteen days, which could have drastically affected expectations distributions. In addition, smoothing techniques used could worsen the problem.

Aggregation could be particularly negative if one of the purposes of the credibility analysis is to analyze the content of options to anticipate realignments of the exchange rate bands, as is pointed out by the authors. In my opinion, if the role of options as leading indicators of realignments is to be stressed, a daily frequency in the analysis is required.

For this reason, it is highly advisable to calibrate the disadvantages and advantages of using daily data, particularly when aggregation does not solve some of the problems of having sparse data and it is doubtful that other more serious concerns do not emerge as a consequence of data manipulation.

On the other hand, aggregation of data is also particularly problematic when an attempt to explain realignment intensities – the product of the probability of devaluation and the expected size of it – by economic variables is made. The absence of significance in the regressions carried out, could be attributed to the fact that, in this type of markets, realignment probabilities can be subject to daily and sudden changes. External or domestic events and news can cause them. Then, it would be more convenient to analyse how certain events have affected, on a daily basis, the probabilities of realignment.

With respect to the estimation of the PDFs presented in the paper, some more details on the results would be desirable. In fact, because of the manipulation of option data carried out, it would be suitable to provide some measures of goodness to calibrate the quality of the exercise. In my view, some diagnosis -such as, for example, the comparison of first moments of distributions with forward exchange rates or the sum of probabilities estimated, is essential to evaluate to what extent the estimate of risk-neutral probability functions is a good, or at least a sensible, approximation of subjective probabilities.

My second comment is on the conclusions reached by the authors. They point out that their findings reinforce the superior ability of options relative to other indicators, to anticipate realignments of the exchange rate bands. In my opinion, there is no proof of such a statement. The paper highlights the relevance of the information content of options to characterise agents expectations on exchange rates and hence how they "feel" about the possibility of realignment. But the paper does not provide any evidence, in the Brazilian case, to state that options permit to anticipate realignments and no proof is given about the superior performance of options as leading indicators of realignments in relation to others.

To conclude, I would like to stress again the relevant contribution of the paper to the analysis of expectations in financial markets and, particularly, to the study of emerging economies.

Bernardino Adão and Jorge Barros Luís

"Interest rate spreads implicit in options: Spain and Italy against Germany"

Discussants:

Allan M Malz

Christian Upper

Interest Rate Spreads Implicit in Options: Spain and Italy against Germany*

Bernardino Adão
Banco de Portugal
Universidade Católica Portuguesa

and

Jorge Barros Luís
Banco de Portugal
University of York

Abstract

The options premiums are frequently used to obtain probability density functions (pdfs) for the prices of the underlying assets. When these assets are bank deposits or notional Government bonds it is possible to compute probability measures of future interest rates. Recently, in the literature there have been many papers presenting methods of how to estimate pdfs from options premiums. Nevertheless, the estimation of probabilities of forward interest rate functions is an issue that has never been analysed before. In this paper, we propose such a method, that can be used to study the evolution of the expectations about interest rate convergence. We look at the cases of Spain and Italy against Germany, before the adoption of a single currency, and conclude that the expectations on the short-term interest rates convergence of Spain and Italy vis-à-vis Germany have had a somewhat different trajectory, with higher expectations of convergence for Spain.

* Please address any correspondence to Bernardino Adão, Banco de Portugal, Av. Almirante Reis, n.71, 1150 Lisboa, Portugal.

I. Introduction

Derivative prices supply important information about market expectations. They can be used to obtain probability measures about future values of many relevant economic variables, such as interest rates, currency exchange rates and stock and commodity prices (see, for instance, Bahra (1996) and Söderlind and Svensson (1997)).

However, many times market practitioners and central bankers want to know the probability measure of a combination of economic variables, which is not directly associated with a traded financial instrument. This paper presents and illustrates a simple method of obtaining a density function of a combination of economic variables, $f(x,y)$, when there is an option for the variable x and another for the variable y but there is not an option for the variable $f(x,y)$.

To get the implied probability density function on the combination $f(x,y)$, when no options are traded on $f(x,y)$ itself, requires knowledge of the implied correlation between x and y. If variables x and y are exchange rates the implied correlation may be obtained from currency option implied volatilities. Campa and Chang (1997) show how correlations can be calculated via the triangular relationship that exists between options on different currency pairs. However, there is no known procedure of computing implied correlations when variables x and y are not exchange rates. In such cases, proxies (for instance, historical correlations or cross-section correlation between the futures prices for different maturities) are the only alternative. Given that the proxies may not coincide with the market implied correlations, it is necessary to carry out a sensitivity check of the results to the correlation assumption.

We restrict our attention to the case $f(x,y) = x\text{-}y$, being x a 3-month forward interest rate of the Spanish peseta or the Italian lira and y a 3-month forward interest rate of the German mark. Such method is particularly relevant for assessing convergence probabilities of monetary policies. This was a relevant issue in the assessment of the transition towards EMU, having been an important subject of empirical work by central banks and investment banks over the last years.

In this context, "EMU convergence calculators" were developed to compute the probability of a given country joining the EMU.[1] These convergence calculators have as domain the set of forward interest rate differentials and as counterdomain the interval [0,1]. However, this procedure of computing convergence indicators is controversial, to say the least, since the link between forward interest rate differentials and convergence probabilities is not clear. In fact, those differentials change due to other factors than the referred probabilities, namely those related with the business cycle and the economic convergence.

Moreover, the EMU calculators are based on the estimation of the interest rate spread that would happen if the country did not join the EMU. This spread has been estimated using the estimators of a regression between the observed spreads and term structure variables related to international risk-aversion and liquidity during a recent period in which the EMU probability was near zero. The identification of such sample period is difficult. Besides, the estimators' values are subject to the Lucas critique or other structural breaks. Therefore, the regression results are evidently conditioned by the sample period considered, as well as by the explanatory variables chosen.

The method presented in this paper goes further than previous methods that compute probabilities of future events, as it is based on risk-neutral probabilities that can be estimated directly from the prices of traded financial instruments. However, it has the disadvantage of the information about options premiums traded in exchanges being limited to one year horizons, while, for instance, indicators built from the current term structure of interest rates enable the computation of forward interest rates for longer horizons.

We focus on the probability of short-term interest rate convergence instead of on the probability of EMU participation for two reasons. The first has to do with the fact that a zero interest rate differential does not imply a unitary probability of monetary union. The second has to do with the fact that the monetary integration in January 1999 did not preclude positive interest rate differentials in earlier settlement dates.[2]

In fact, the European monetary unification, on 1/1/1999, implied a zero interest rate differential between Germany and the remaining participating countries only at that time. The announcement of the participating countries took place in the Brussels summit of 1-3 May 1998. The bilateral parities were defined as the ERM-EMS central exchange rates, which were different from the spot exchange rates. Thus, the associated spot interest rate differentials would likely be different from zero and would converge to zero, as EMU starting date would approach. Therefore, a small probability of a country short-term interest rates convergence for dates before 1/1/1999 did not necessarily mean that the market expected that country would stay out of the European Monetary Union.

The evidence presented in the paper suggests that the options market participants did not consider likely that there would be complete convergence of interest rates between Germany, Spain and Italy before June 1998. There is also evidence that the Italian short-term interest rate convergence was expected to be behind the Spanish one. Moreover, there is evidence that the option markets expected both spreads to be smaller in June of 1998 than in March 1998.

This paper has four more sections. The second section describes the model; the third section contains the estimation technique; the data and the empirical results obtained are presented in the fourth section and the fifth section concludes.

II. Model

Let S_1 and S_2 be two futures contracts. We assume that $(\ln S_{1t}, \ln S_{2t})$ are stochastic variables distributed as a mixture of two bivariate normal distribution. Thus, the bivariate probability density function (pdf) of (S_{1t}, S_{2t}) is a mixture of two bivariate lognormal, given by:

(1)

$$
\begin{aligned}
f(S_{1t},S_{2t}) = {} & \theta_t \left((2\pi)^2 |\mathbf{\Sigma}_{1t}|\right)^{-\frac{1}{2}} \exp\left(-\frac{1}{2}\left([\ln S_{1t}\ \ln S_{2t}] - \boldsymbol{\mu}'_{1t}\right)\mathbf{\Sigma}_{1t}^{-1}\left([\ln S_{1t}\ \ln S_{2t}]' - \boldsymbol{\mu}_{1t}\right)\right)\times \\
& \times \frac{1}{S_{1t}S_{2t}} + (1-\theta_t)\left((2\pi)^2 |\mathbf{\Sigma}_{2t}|\right)^{-\frac{1}{2}} \times \\
& \times \exp\left(-\frac{1}{2}\left([\ln S_{1t}\ \ln S_{2t}] - \boldsymbol{\mu}'_{2t}\right)\mathbf{\Sigma}_{2t}^{-1}\left([\ln S_{1t}\ \ln S_{2t}]' - \boldsymbol{\mu}_{2t}\right)\right)\frac{1}{S_{1t}S_{2t}}
\end{aligned}
$$

where $\boldsymbol{\mu}'_{kt} = [\mu_{k1t}\ \mu_{k2t}]$, $\mathbf{\Sigma}_{kt} \equiv \begin{bmatrix} \sigma^2_{k1t} & \rho_{kt}\sigma_{k1t}\sigma_{k2t} \\ \rho_{kt}\sigma_{k1t}\sigma_{k2t} & \sigma^2_{k2t} \end{bmatrix}$ and $\theta \in [0,1]$, being k = 1,2. The parameters $\boldsymbol{\mu}'_{kt}$ are the expected value vectors of the corresponding bivariate normal distributions, $\mathbf{\Sigma}_{kt}$ are the covariance matrices of the corresponding bivariate normal distributions and θ_t and $(1-\theta_t)$ are the weights of the distributions.

The marginal pdf for S_{it}, $f'(S_{it})$, is:

(2)
$$f'(S_{it}) = \theta_t \left((2\pi)^2 \sigma_{1it}^2\right)^{-1/2} \exp\left(-\frac{1}{2\sigma_{1it}^2}(\ln S_{it} - \mu_{1it})^2\right)\frac{1}{S_{it}} + $$
$$+ (1-\theta_t)\left((2\pi)^2 \sigma_{2it}^2\right)^{-1/2} \exp\left(-\frac{1}{2\sigma_{2it}^2}(\ln S_{it} - \mu_{2it})^2\right)\frac{1}{S_{it}}$$

To get an unconditional pdf for S_{1t} - S_{2t} we transform the variables.[3] Define Y_{1t}=S_{1t}-S_{2t} and Y_{2t} = S_{2t}. The bivariate pdf of Y_{1t} and Y_{2t} is:

(3)
$$f^*(Y_{1t}, Y_{2t}) = \theta_t \left((2\pi)^2 |\mathbf{\Sigma}_{1t}|\right)^{-1/2} \times$$
$$\times \exp\left(-\frac{1}{2}\left([\ln(Y_{1t}+Y_{2t})\ \ln Y_{2t}] - \boldsymbol{\mu}_{1t}'\right)\mathbf{\Sigma}_{1t}^{-1}\left([\ln(Y_{1t}+Y_{2t})\ \ln Y_{2t}]' - \boldsymbol{\mu}_{1t}\right)\right)\times$$
$$\times \frac{1}{(Y_{1t}+Y_{2t})Y_{2t}} + (1-\theta_t)\left((2\pi)^2 |\mathbf{\Sigma}_{2t}|\right)^{-1/2} \times$$
$$\times \exp\left(-\frac{1}{2}\left([\ln(Y_{1t}+Y_{2t})\ \ln Y_{2t}] - \boldsymbol{\mu}_{2t}'\right)\mathbf{\Sigma}_{2t}^{-1}\left([\ln(Y_{1t}+Y_{2t})\ \ln Y_{2t}]' - \boldsymbol{\mu}_{2t}\right)\right)\times$$
$$\times \frac{1}{(Y_{1t}+Y_{2t})Y_{2t}}$$

Thus, the marginal pdf for $Y_{1\mathrm{t}}$ corresponds to the pdf for the interest rate differential, being determined as:

(4) $$f^{**}(Y_{1t}) = \int_{-\infty}^{+\infty} f^*(Y_{1t}, Y_{2t})\, dY_{2t}$$

Since the algebraic integration of $f^*(Y_{1t}, Y_{2t})$ with respect to Y_{2t} revealed itself complex, we decided to get the marginal pdf of Y_{1t} through a computer code, which integrates numerically $f^*(Y_{1t}, Y_{2t})$ with respect to Y_{2t}.[4]

III. Estimation

To compute the conditional distribution of interest rates, as well as the marginal distribution of interest rate differentials, it is necessary to know the vectors of expected values $\boldsymbol{\mu}'_{kt}$ and the matrices of covariances $\boldsymbol{\Sigma}_{kt}$. The expected values and the variances in the matrices $\boldsymbol{\Sigma}_{kt}$ can be estimated from the options premiums on futures contracts. As it is well known, when investors are risk-neutral, an European call option on a futures contract at time t with strike price X and term to maturity τ, $C(X,\tau)$, obeys the expression:

$$
\begin{aligned}
(5)\qquad C(X,\tau) &= E_t e^{-r_{t,\tau}\tau} \max[S_{iT} - X, 0] \\
&= e^{-r_{t,\tau}\tau} \int_X^{\infty} (S_{iT} - X) q(S_{iT}) dS_{iT}
\end{aligned}
$$

where E_t is the conditional (on information known at date t) expected value operator, S_{iT} is the price of the underlying asset at maturity date T, r a riskless interest rate for maturity τ (being $\tau = T\text{-}t$), and $q(S_{iT}) \sim \theta LN(\mu_{1i},\sigma_{1i}) + (1-\theta) LN(\mu_{2i},\sigma_{2i})$ (i.e., a combination of two lognormal pdfs, being i = 1, 2). The solution to the problem below gives an estimate for the nine parameters of $q(S_{iT})$:

$$
(6)\qquad \underset{\substack{\theta,\mu_{11t},\mu_{12t},\mu_{21t},\mu_{22t},\\ \sigma_{11t},\sigma_{12t},\sigma_{21t},\sigma_{22t}}}{Min} \sum_{j=1}^{M} \left[\hat{C}_{1,j}(X_{1,j},\tau) - C^0_{1,j}\right]^2 + \sum_{j=1}^{N} \left[\hat{C}_{2,j}(X_{2,j},\tau) - C^0_{2,j}\right]^2,
$$

where M and N represent the number of strike prices available for each option, $X_{1,j}$ and $X_{2,j}$ the strike prices observed for each option, $\hat{C}_{1,j}$ and $\hat{C}_{2,j}$ are the estimated option premiums for each option and strike price for the maturity τ and $C^0_{1,j}$ and $C^0_{2,j}$ are the premiums observed for each option and strike price.[5]

There is a relationship between the elements off the diagonal of the matrices $\mathbf{\Sigma}_{1t}$ and $\mathbf{\Sigma}_{2t}$ and the correlation between $\ln S_{1t}$ and $\ln S_{2t}$. It can be shown, using the moment generating function technique, that the resulting correlation between $\ln S_{1t}$ and $\ln S_{2t}$ is:

$$(7)\ \rho=\frac{\theta\left(\mu_{11t}\mu_{12t}+\rho_{1t}\sigma_{11t}\sigma_{12t}\right)+(1-\theta)\left(\mu_{21t}\mu_{22t}+\rho_{2t}\sigma_{21t}\sigma_{22t}\right)-\left(\theta\mu_{11t}+(1-\theta)\mu_{21t}\right)\left(\theta\mu_{12t}+(1-\theta)\mu_{22t}\right)}{\prod_{j=1}^{2}\left[\theta\left(\mu_{1jt}^{2}+\sigma_{1jt}^{2}\right)+(1-\theta)\left(\mu_{2jt}^{2}+\sigma_{2jt}^{2}\right)-\left(\theta\mu_{1jt}+(1-\theta)\mu_{2jt}\right)^{2}\right]^{1/2}}$$

The assumptions about the elements off the diagonal matrices, $\mathbf{\Sigma}_{kt}$'s, will be discussed in the next section.

IV. Data and empirical results

We applied the model described to the conditional and unconditional differentials between 3-month German mark and Italian lira interest rates, on one hand, and German mark and Spanish peseta, on the other. For that purpose we used daily quotes between 18/3/1997 and 7/7/1997 of LIFFE's futures options on 3-month interest rates (for the mark and the lira) and MEFF Renta Fija's futures options on 3-month interest rates (for the peseta), with maturity on March 1998 and June 1998.[6]

Given that a fast nominal convergence process marked this period, we suspected that neither the historical correlation nor the futures prices correlation might estimate correctly the true correlation between the interest rates.[7] Thus, we opted for assuming several correlation figures and assessing the sensitivity of the results to those figures. We verified that the higher the values for $\bullet_1$ and $\bullet_2$ we consider, the more probability would be concentrated in higher differentials. That has to do with the fact that during the sample period the 3-month interest rate differentials were still substantial (see figure 1).

Consequently, the upper and the lower bounds of the interest rate convergence probability are obtained when the $\bullet_K$'s are near –1 and 1, respectively. Therefore, we initially considered $\bullet_1$, $\bullet_2$ = 0.99,[8] in order to characterise the less favourable scenario to lower interest rate differentials. The unconditional interest rate differential pdfs, estimated according to (6), evidence lower expected differentials for the Spanish peseta and for June 1998 (figure 2).

According to the shape of the estimated pdfs, the statistical measures of the distributions evidence a significant positive skewness in both differential pdfs for March 1998, as the mode is consistently lower than the median and the latter is consistently lower than the mean (figure 3). This difference is reduced only in the last sample days. The shape of the estimated pdfs is different for June 1998, as they consistently exhibit a negative skewness.

The results suggest that major improvements about the prospects on interest rate convergence of the Spanish peseta in March 1998 were achieved between March 1997 and July 1997, as the mean decreased from around 2.1 to below 1.5. Further convergence was expected to be done between March 1998 and June 1998, given that the mean of the distribution for June 1998 was between 0.75 and 1 in the sample considered. In spite of the large Italian lira differential, the evolution was not so remarkable, but the results obtained also show an expectation of lower spreads in June 1998 than in March 1998.

The distribution functions for the interest rate differentials can be used to get an indicator about the expectations of the lira and the peseta short-term interest rate convergence. This indicator could be the cumulative probability in zero, or any low enough differential that may be considered as corresponding to interest rate convergence.

The distribution functions for the interest rate differentials (peseta-mark and lira-mark) revealed small lower bounds to the probabilities for non-positive interest rate differentials (figure 4). However, the figures for the differentials in June 1998 (which were for Spain between 10% and 14% and for Italy up to 4%) were clearly higher than for March 1998. This may reflect that the markets were not completely sure of the peseta and lira integration in the European Monetary Union, or that they expected the continuation of convergence between June 1998 and the starting of the European Monetary Union in 1/1/1999.

As it was referred at the beginning of this section, until now we have assumed a correlation coefficient for both distributions of 0.99. Consequently, it is important to assess the sensitivity of the results to the correlation coefficient figures.

That analysis was performed for the Spanish differentials in March 1998. The results obtained confirm the conjecture that the higher the correlation coefficient, the smaller the probability of small interest rate differentials. Furthermore, our results also show that higher correlation coefficients imply lower dispersion and less smooth curves (figure 5).[9]

We also performed a sensitivity test to the estimation method. We compare the results obtained with a mixture of two lognormal distributions specification, previously presented, with those obtained with a lognormal distribution specification. In general, the pdfs obtained with the two specifications are rather different. The two-lognormal distribution is more asymmetric, sometimes has more than one mode and has fatter tails. As a result, the indicator we used to assess interest rate convergence, the probability of non-positive interest rate differential, assumes higher values in the two-lognormal specification.

In figure 6 we have the Spanish differential pdfs for March 1998 in 20/3/97 and 7/7/1997 (assuming 0.99 correlation coefficients). The day 7/7/1997 represents better the more frequent differences associated with the two estimation methods. In that sense, that day is a more standard day than the day 20/3/97.

V. Conclusion

The option premiums have been recently used to extract information about the expected future behaviour of many economic variables, in particular interest rates. Nevertheless, this literature does not have anything to say about the expected future behaviour of interest rate differentials. It is this paper objective to perform that task, with the estimation of the density function of short-term interest rate differentials.

The estimation of unconditional pdfs for the peseta-mark and lira-mark differentials shows that the conjecture that Spain was ahead of Italy in the convergence process was correct, at least with respect to short-term interest rates.[10] Moreover, the exercise shows that financial markets expected further convergence in the short-term interest rates after June 1998.

The interest rate differentials pdfs can also be used, for instance, to identify the convergence of long-term interest rates, supplying, in this case, useful information about the expectations of sovereign and liquidity risk and/or accomplishment of the long-term interest rates convergence criteria.[11]

ACKNOWLEDGEMENTS

The authors would like to thank to LIFFE and MEFF Renta Fija for making available the data used in this paper. A special acknowledgement is due to Nuno Cassola for encouragement and for the intensive discussions he had with us. The paper has also improved substantially from very useful suggestions and comments by an anonymous referee. The second author acknowledges support of the Lisbon Stock Exchange.

[1] See, for instance, De Grauwe (1996), Dillén and Edlund (1997), Favero *et al.* (1997) and JPMorgan (1997).

[2] At the time the exercise was done March and June 1998 were the maturity dates available nearest to January 1999.

[3] The conditional pdfs are obtained from the bivariate and marginal pdfs. The conditional pdf for S_{1t} given S_{2t} is the ratio $\frac{f(S_{1t}, S_{2t})}{f(S_{2t})}$. If we had chosen θ = 1, after substitution and algebraic manipulation, the following expression for the conditional pdf of S_{1t}, could have been obtained:

$$\exp\left\{-\frac{1}{2\sigma_{11t}^2\left(1-\rho_{1t}^2\right)}\left[\ln S_{1t} - \mu_{11t} - \frac{\rho_{1t}\sigma_{11t}}{\sigma_{12t}}\left(\ln S_{2t} - \mu_{12t}\right)\right]^2\right\} \times \left[2\pi\sigma_{11t}^2\left(1-\rho_{1t}^2\right)\right]^{-1/2}\frac{1}{S_{1t}}.$$

[4] The grid for Y_{2t} was chosen between with an interval of 10 basis point.

[5] About different alternative estimation techniques of the distribution for S_T see, e.g., Bahra (1996).

[6] Even though the options traded in LIFFE are American, since they are pure options, they can be treated as European options.

[7] Simple historical correlation coefficients, as well as exponentially weighted moving average of the correlation coefficients (between the futures contracts prices with term to maturity between September 1996 and June 1998, in the period between March 14, 1995 and the day for which the pdf is estimated), are in general positive and high. On the other hand, daily cross-section correlation coefficients between the interest rates implicit in futures prices for different settlement dates are in general non-positive.

[8] For the period under consideration this corresponds to a correlation coefficient between $\ln S_{1t}$ and $\ln S_{2t}$ in the interval [0.3,0.7]. These values are almost always below the corresponding historical correlation coefficients.

[9] Similar conclusions are also obtained for conditional density functions.

[10] As we are estimating risk-neutral density functions, any difference in risk-aversion patterns is not taken into account, which means that some of the probability differences between the Italian and the Spanish spread, as well as their time changes, may have been motivated by risk-aversion differences or variations.

[11] It is also possible to extract information about long-term interest rate differentials expectations from the prices of *DIFF* future contracts prices traded at MEFF Renta Fija. It would be interesting to evaluate the consistency of the information obtained by this method, confronting the expected values estimated with the differentials implicit in the quotes of the *DIFF* contracts.

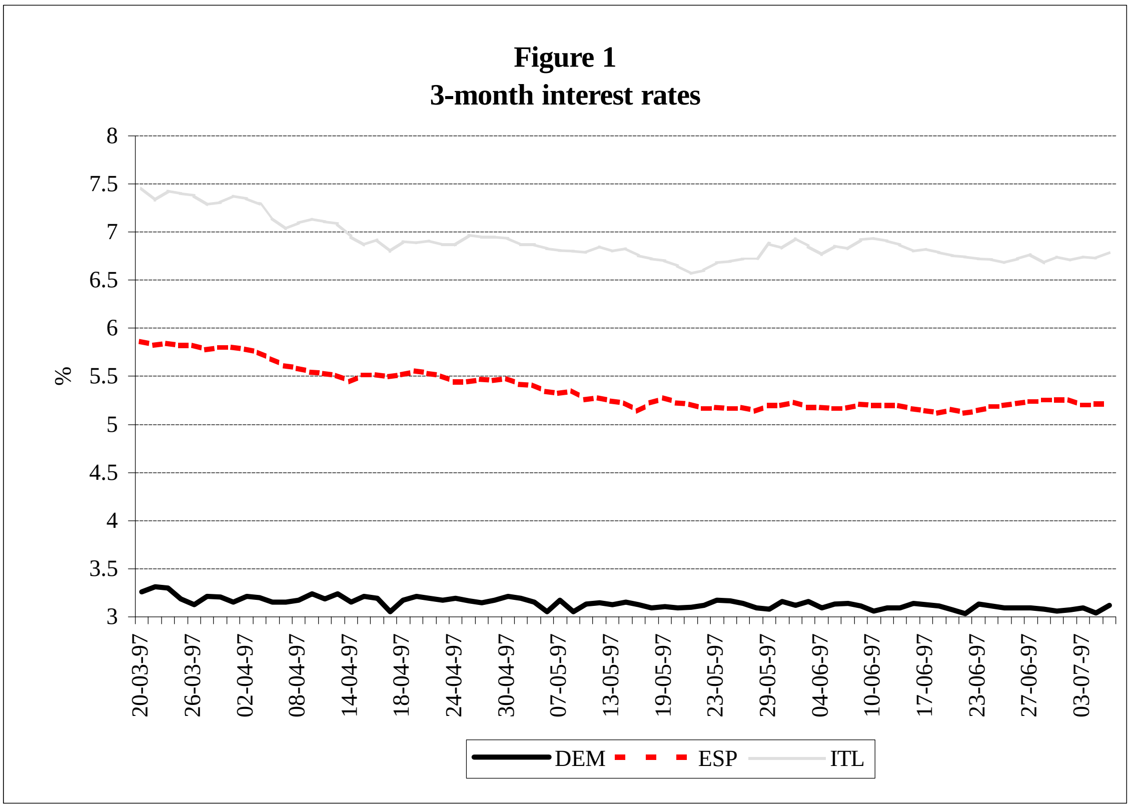
Figure 1
3-month interest rates
%
8
7.5
7
6.5
6
5.5
5
4.5
4
3.5
3
20-03-97
26-03-97
02-04-97
08-04-97
14-04-97
18-04-97
24-04-97
30-04-97
07-05-97
13-05-97
19-05-97
23-05-97
29-05-97
04-06-97
10-06-97
17-06-97
23-06-97
27-06-97
03-07-97
DEM
ESP
ITL

Figure 2
Unconditional density functions of 3-month interest rate differentials vis-à-vis German mark (ρ1, ρ2=0.99)

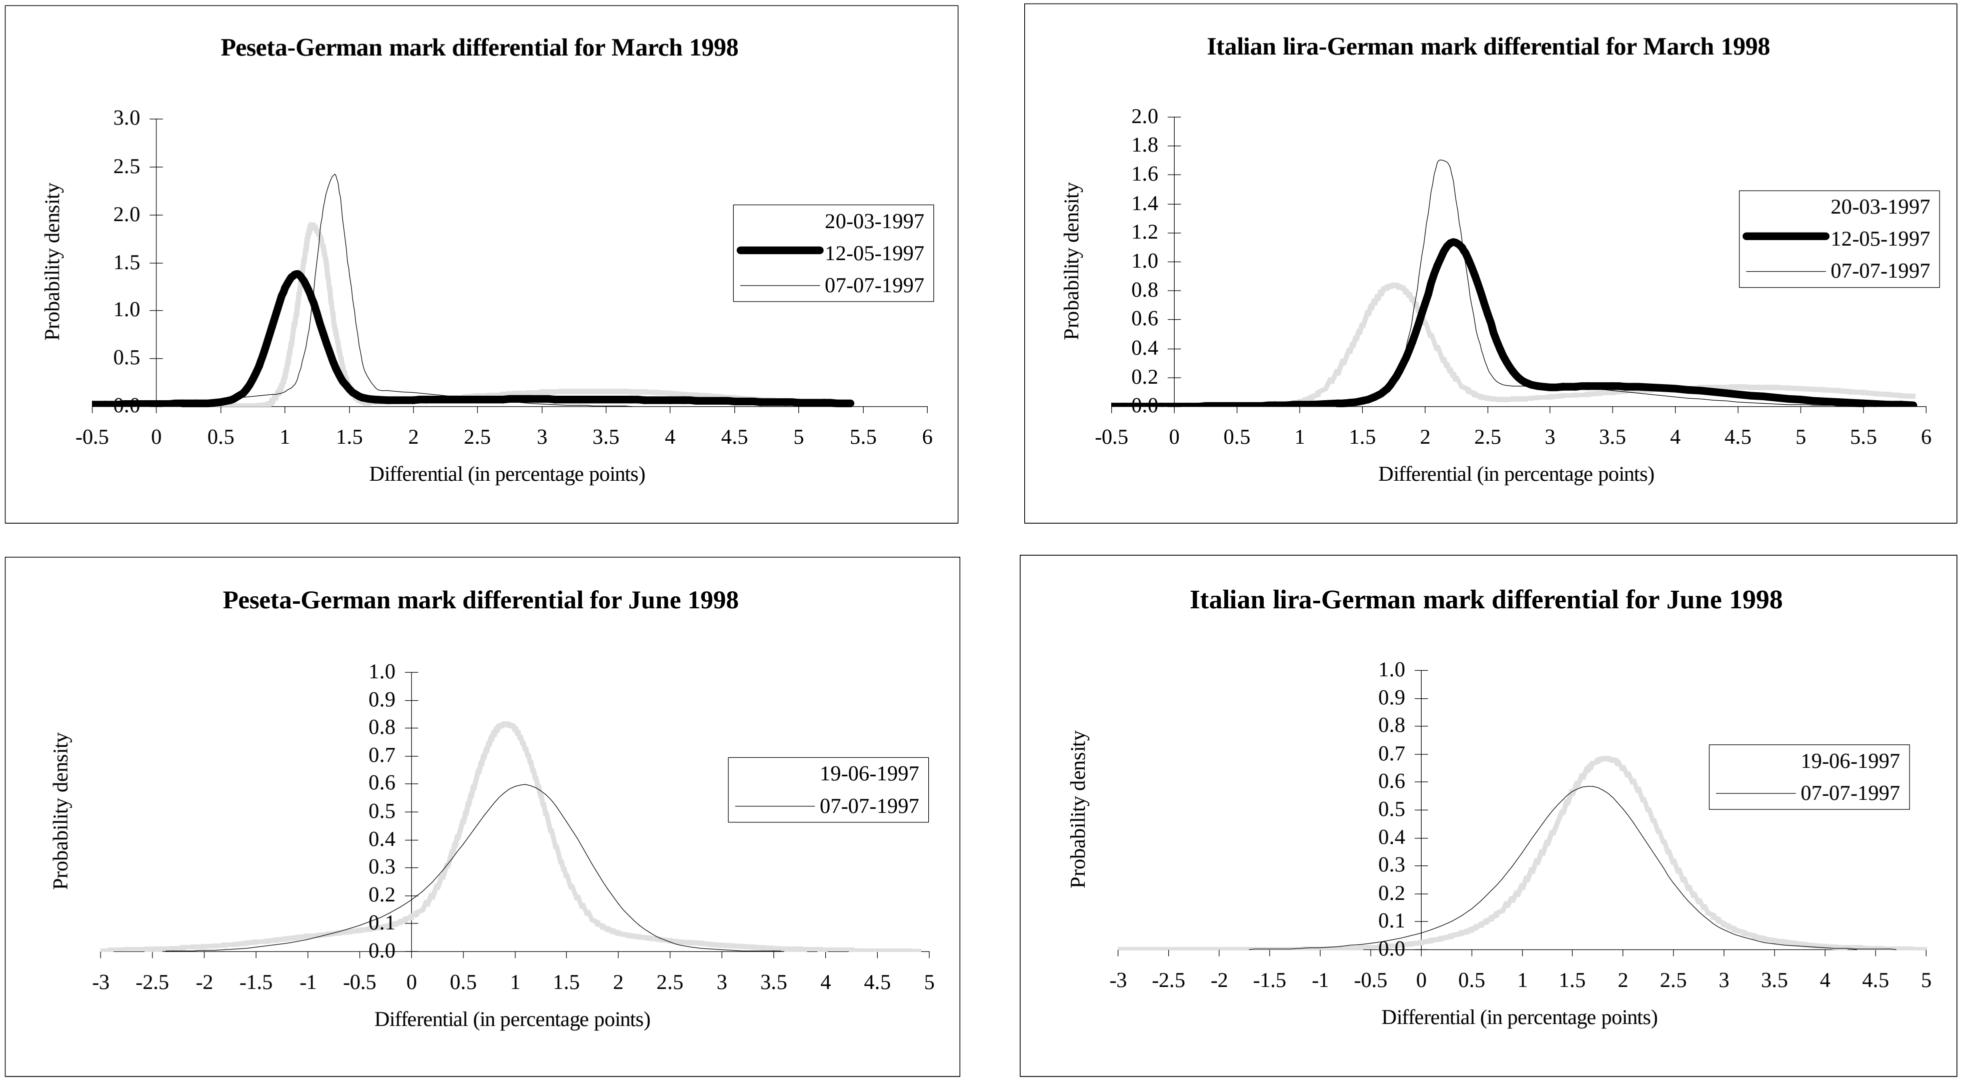

Figure 3
Statistical measures of 3-month interest rate differentials vis-à-vis German mark (ρ1, ρ2=0.99)

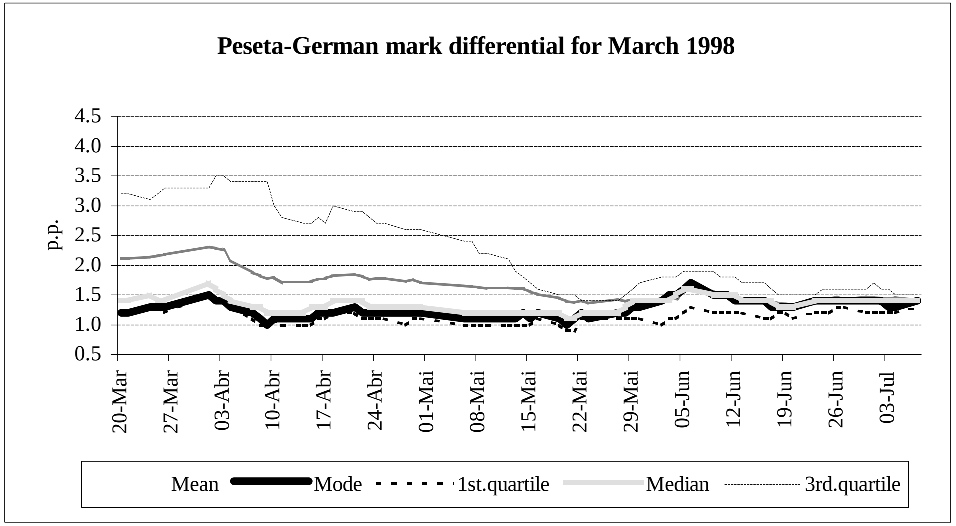

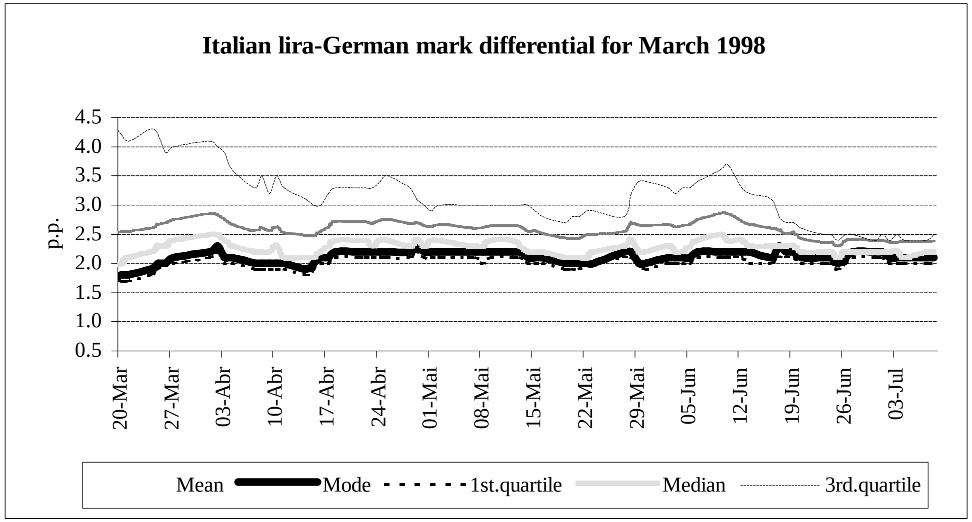

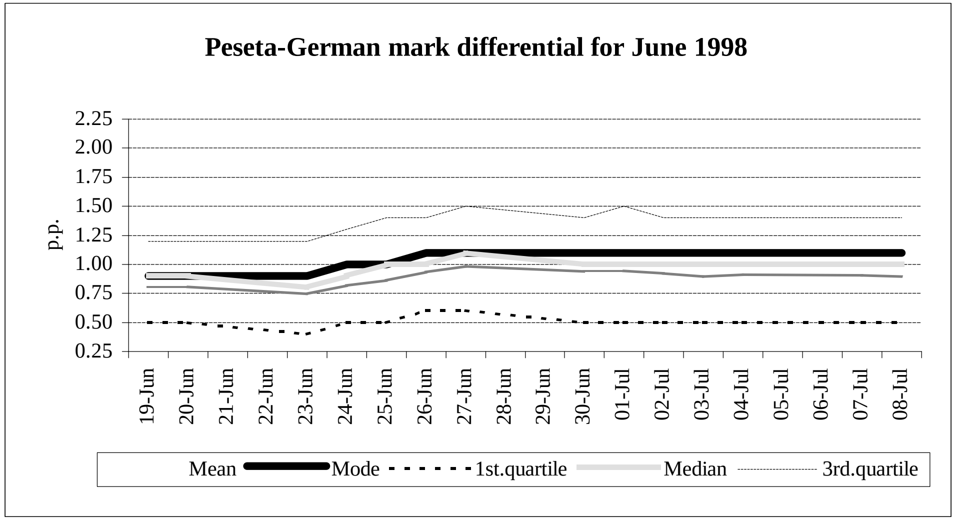

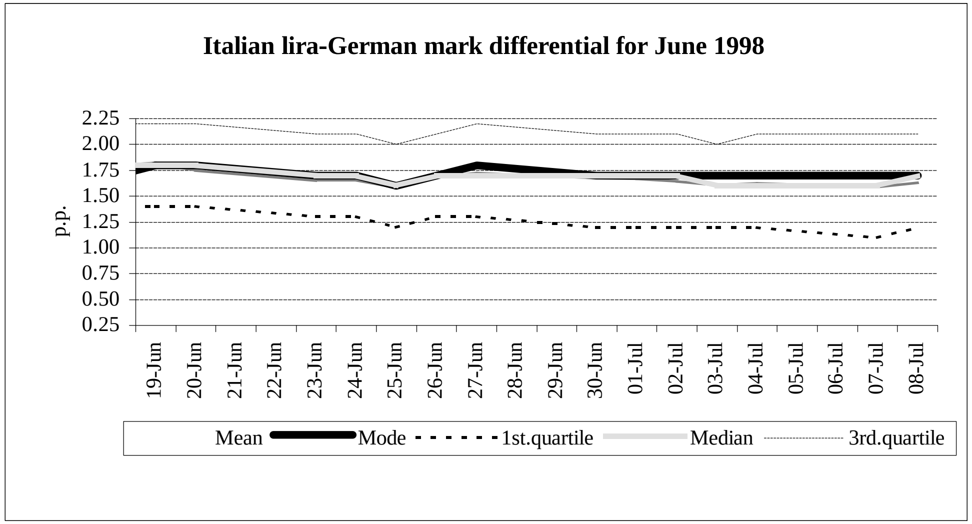

Figure 4
Probability of a non-positive interest rate spread vis-à-vis Germany (ρ1, ρ2=0.99)

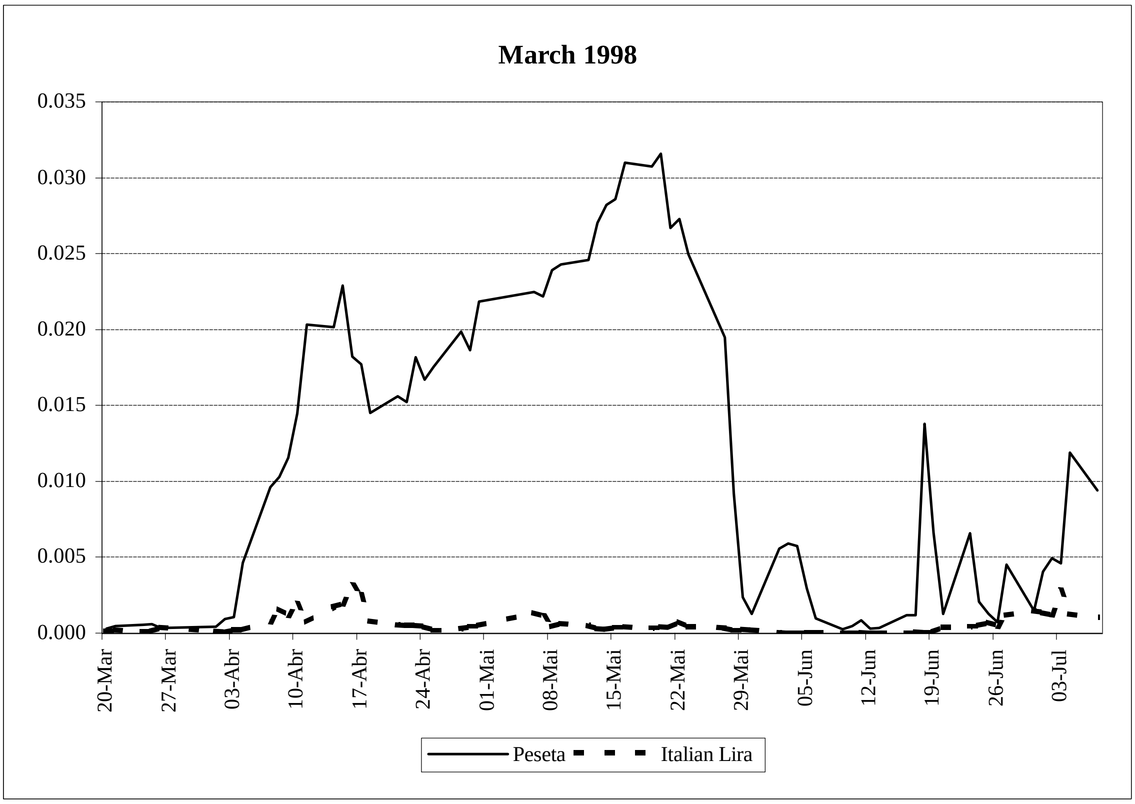

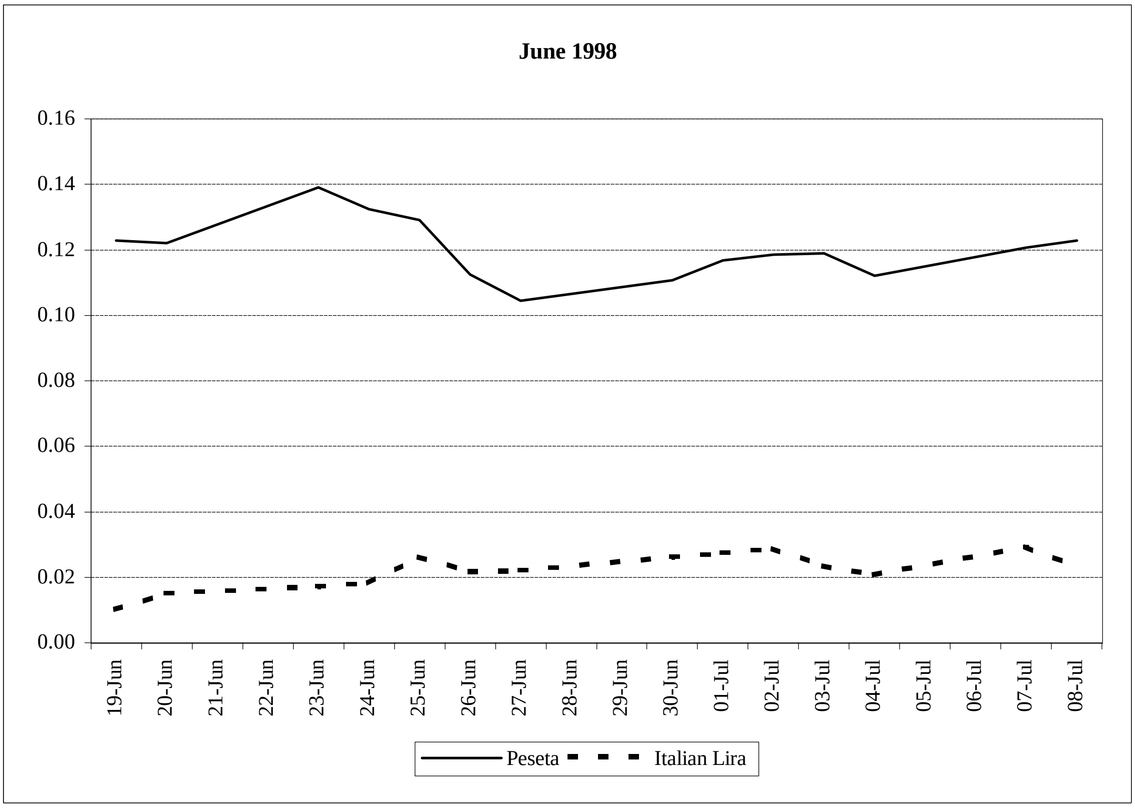

Figure 5
Sensitivity to the correlation coefficient of the unconditional density functions of 3-month interest rate differentials between the Spanish peseta and the German mark for March 1998

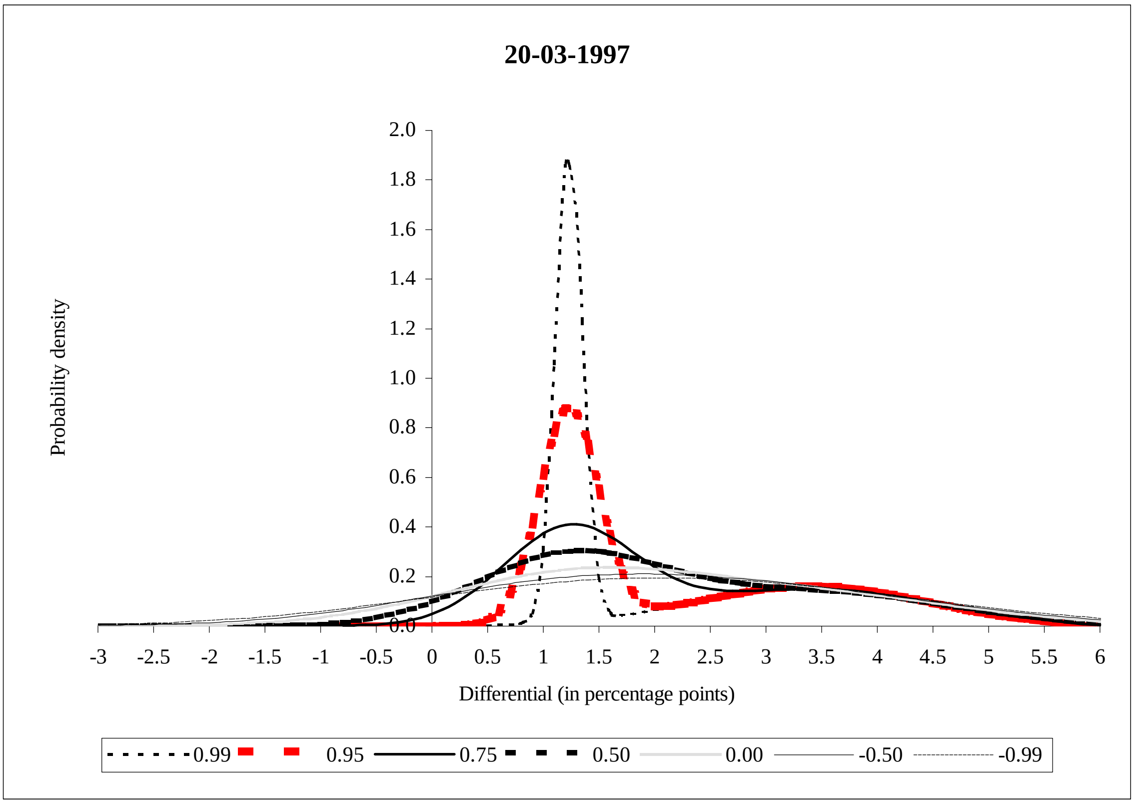

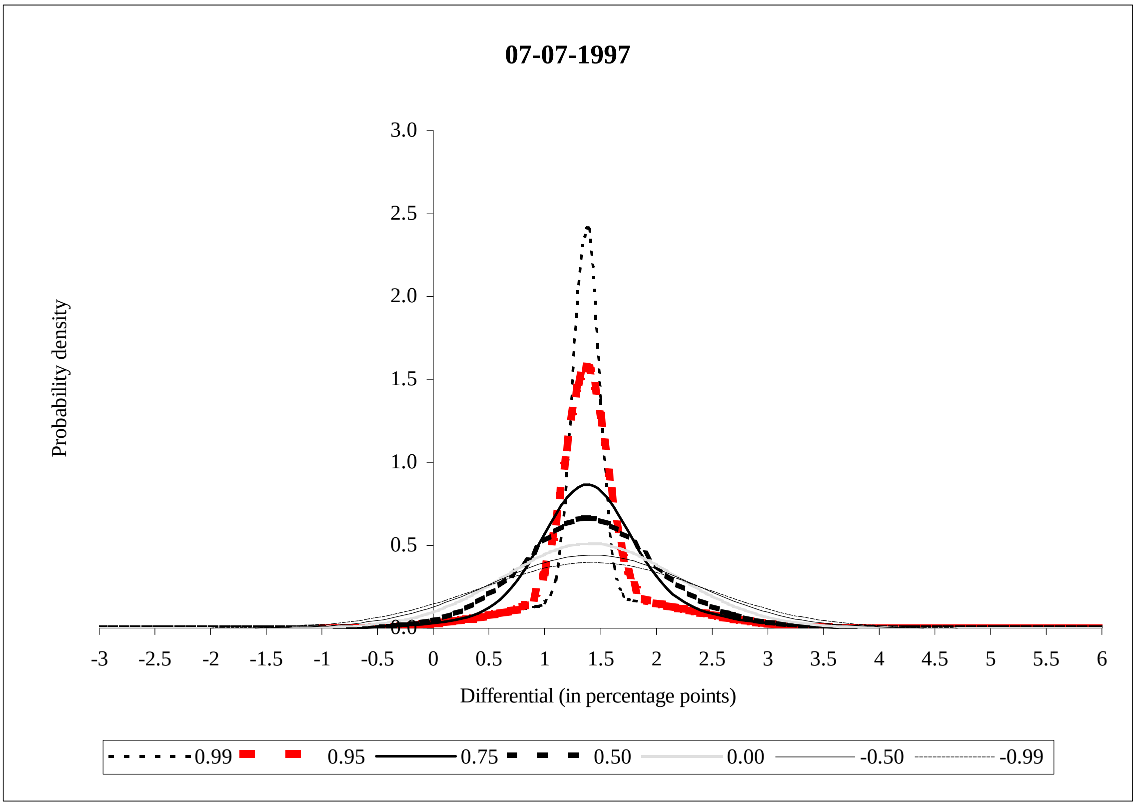

Figure 6

Sensitivity to the estimation method of the unconditional density functions of 3-month interest rate differentials between the Spanish peseta and the German mark for March 1998

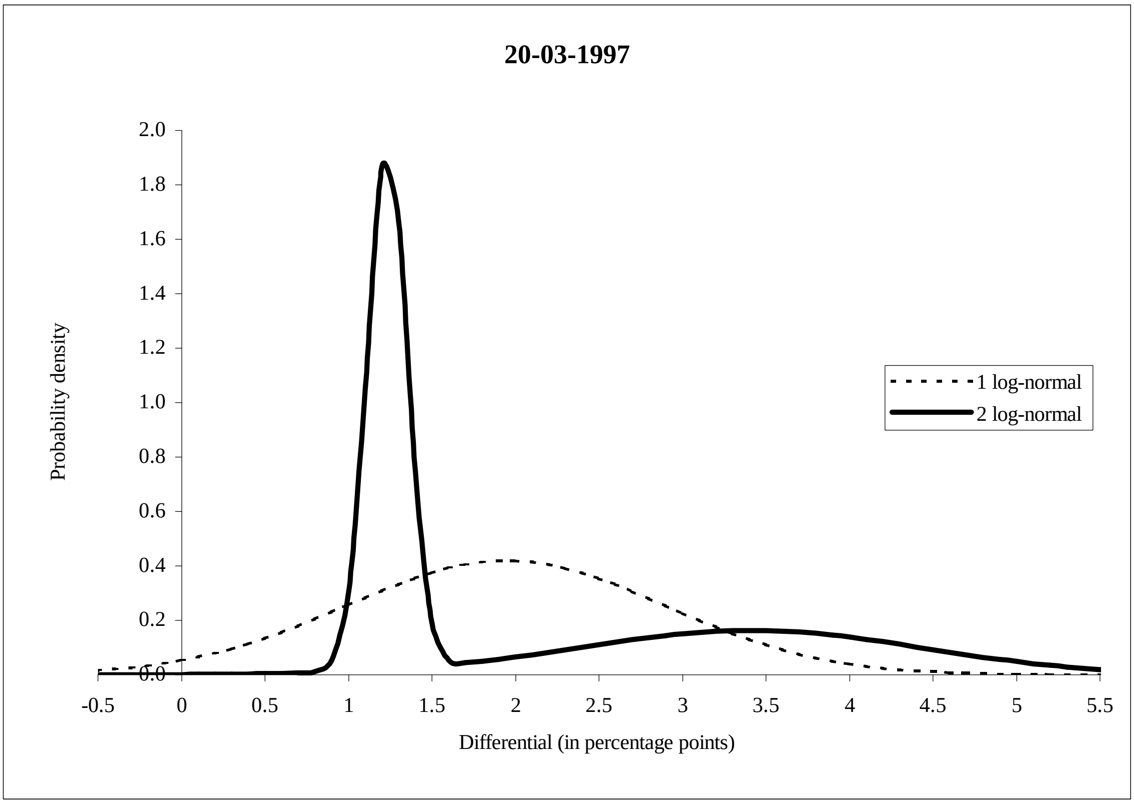

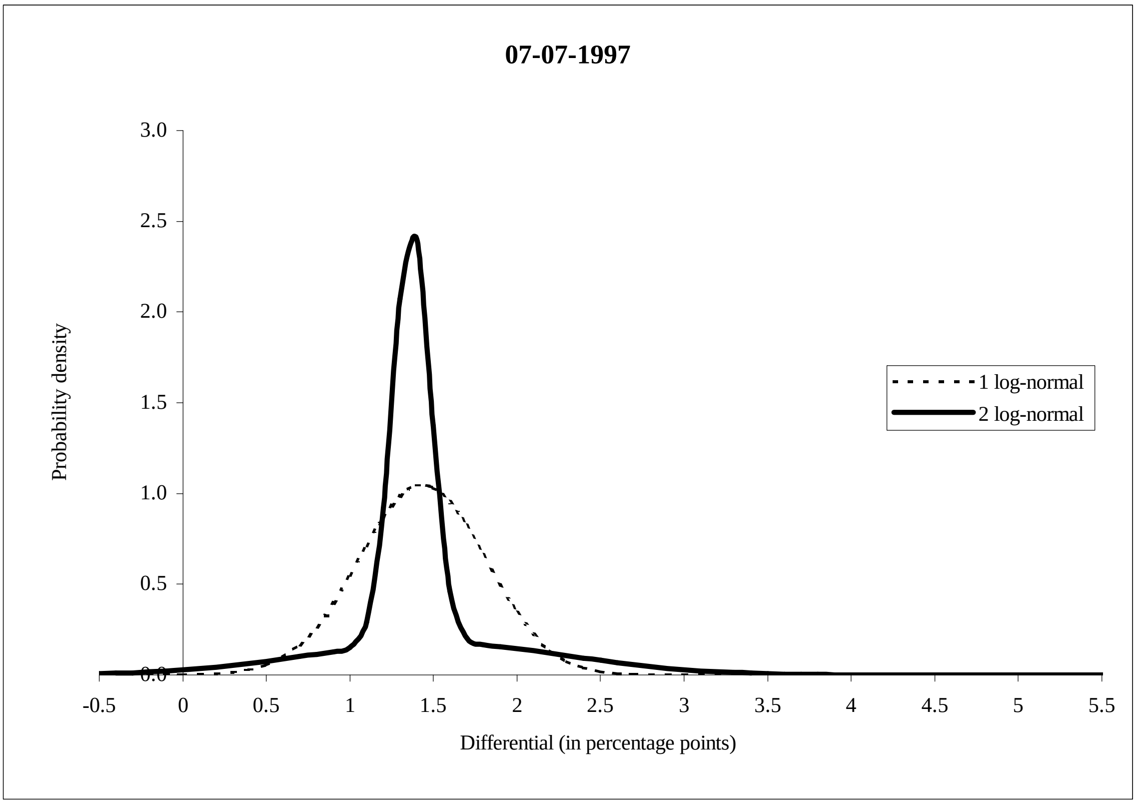

REFERENCES

Bahra, B. (1996), Implied Risk-Neutral Probability Density Functions from Option Prices: Theory and Application, *Bank of England Economic Bulletin*, August, 299-309.

Campa, José M. and P. H. Kevin Chang (1997), "The Forecasting Ability of Correlations Implied in Foreign Exchange Options", NBER Working Paper Series, Working Paper 5974.

De Grauwe, P. (1996), Forward Interest Rates as Predictors of EMU, *CEPR Discussion Paper Series* No. 1395.

Dillén, Hans and Martin Edlund (1997), "EMU expectations and interest rates", Bank of Sweden, Quarterly Review 2/97.

Favero, Carlo, Francesco Giavazzi, Fabrizio Iacone and Guido Tabellini (1997), "Extracting Information from Asset Prices: The methodology of EMU calculators".

JPMorgan (1997), "EMU calculator handbook", Foreign exchange research: technical series, January 15.

Söderlind, Paul and Svensson, L.E.O. (1997), "New Techniques to Extract Market Expectations from Financial Instruments"; Working Paper 5877, NBER Working Paper Series.

Comments on the Adão and Barros paper

Allan M. Malz
The RiskMetrics Group
allan.malz@riskmetrics.com

June 1999

This paper combines market data on a pair of interest rates with an estimate of the correlation of those rates to arrive at estimates of the probability distribution of the interest rate spread. It applies this method to interest rate spreads between German mark money markets on the one hand and Italian lira and Spanish peseta rates on the other, with a view to assessing expectations regarding convergence of interest rates in the approach to European Monetary Union (EMU) at the beginning of 1999. Since perfect interest rate convergence was presumed to be a consequence of credibly fixing exchange rates with no fluctuation limits, market views on interest rate convergence serve also as indicators of expectations regarding the credibility, timing and membership of EMU.

The estimation exercise requires an estimate of the expected correlation of interest rates. The authors base most of the results they report on a “benchmark" estimate in which the correlation is assumed to be close to 0.99. Given the estimated parameters, this assigns the lowest possible probabilities to narrow interest rate spread outcomes. It thus represents the most cautious assessment from the point of view of a policy maker aiming at convergence.

The authors estimate the probability density function by fitting a mixture of bivariate joint lognormals to prices of options on short-term interest rate futures. This is a natural extension of the by-now standard procedure of applying mixture distributions in PDF estimation. The correlation between the two interest rates is then determined by both the estimated parameters and the postulated values of the correlations of the lognormal distribution entering into the mixture. The highest possible correlation between the rates themselves is arrived at by postulating the correlations among the mixture components to be +0.99.

The paper draws several broad conclusions from the estimates:

- The market expected substantial interest rate convergence to occur between March and June 1998.
- Convergence was not, however, expected to be complete by June 1998.
- Interest-rate convergence was lower for the lira than for the peseta

One consequence of this procedure is that the outcome of a negative future interest-rate spread vis-à-vis the mark is assigned a positive, even substantial, probability. The probability of negative spreads is highest for the authors' benchmark case in which the mixture components have correlations of +0.99. However, a negative spread is difficult to interpret: the flip side of the “cautious" benchmark is this rather implausible scenario.

What were the correlations? They varied widely over 1996-1998, as seen in the accompanying chart. The red line represents lira and the blue line peseta. Correlations are based on one year of daily data.

Figure 1: Correlations of 3-month cash rates

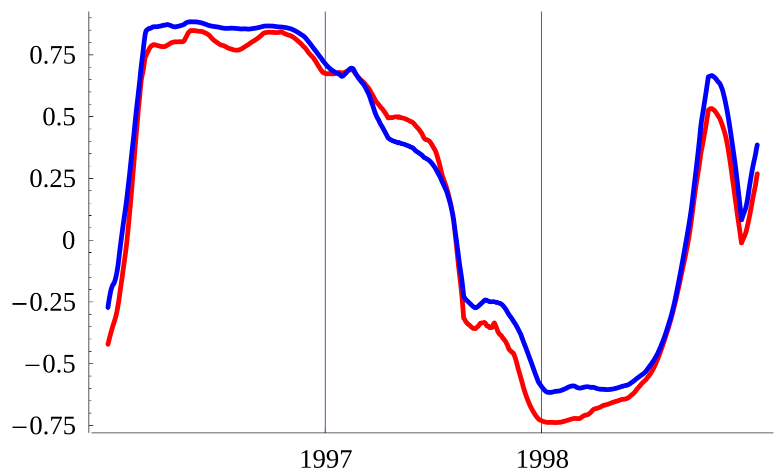

As can be seen, correlations were moving fast towards -1 during the period studied. Later, in 1998, with much convergence already attained and rates across Europe dropping, the correlations moved back towards +1. The chart suggests that market participants may have anticipated much smaller correlations that +0.99.

It would have been interesting to know whether the PDFs of the individual interest rates were also skewed, or if only the German or only the non-German PDFs were skewed.

The paper pushes the envelope on the uses of risk-neutral PDFs. It would be useful, however, to display additional information pertinent to convergence, since it is difficult to interpret PDFs in isolation from other markets and other dimensions of the money markets. In particular, it would be interesting to see

- The evolution of the term structure: How rapidly had longer term rates converged? How had the slopes of the yield curves changed over time? How do forward rate agreements behave?
- Exchange rates and deviations from central parity.
- Historical implied volatility for interest rates and currencies.

It would also be useful to discuss the liquidity of the lira and peseta contracts.

Comments on Adão & Barros Luís

by Christian Upper[1]

Outline of the paper

The paper by Adão and Barros Luís extends the literature of implied probability density functions (PDFs) to variables for which no options. More precisely, they derive an implied PDF for spread between Spanish and German, and Italian and German 3-month interest rates. Since there do not exist any options on the spread, this information has to be extracted from the prices of options on the underlying interest rates.

For this purpose, they set up a bivariate probability density function for the underlying interest rates, which is then transformed into the bivariate density of the spread and of one of the interest rates. Integrating over the latter yields the implied PDF for the spread. They then estimate the parameters of the implied *marginal* PDFs for each contract. This does not, however, yield any estimate for the *correlation* between the two contracts. Instead, they try out different values and check to which extend this affects their results. They find that high values for the correlation coefficient lead to more mass being concentrated in higher differentials.

Their results indicate that option market participants did not expect complete convergence of interest rates by either March or June 1998, although the mean differential implied in the March contracts declines over the sample period, more so for the peseta than for the lire. There did not find any equivalent decline for the June 1998 contract.

[1] The opinions expressed here are the author's own and should not be attributed to the Deutsche Bundesbank.

Comments

Let me first comment on the methodology. The main difficulty is the choice of ρ. Historical correlations cannot be used due to non-stationarity under the null hypothesis of convergence. Implied correlations a la Campa & Chang (1998) could in principle be used if exchange rate at maturity were known, but this was not the case here. Hence the approach of choosing a ρ and then undertaking a sensitivity analysis seems reasonable. The problem is that densities become very flat for low values of ρ, which makes it very difficult to say anything about convergence. A possible alternative would have been to wait a bit until the January 1999 contract became traded, assume that the currencies would enter monetary union at their EMS central rates, and use Chang & Campa's approach to compute implicit correlations.

Let me turn to data issues. The different results they get for the two contract maturities may be due to the low liquidity of the June contract. In fact, even the March contract had a large residual maturity and may thus not be very liquid. Unfortunately, I cannot offer any solution to this problem. Perhaps we should take seriously only implied PDFs that are computed for short residual maturities.

A more serious question is why we should have expected any convergence in short term interest rates by March or June 1998. For both Italy and Spain it seemed clear, although possibly for different reasons, that they would attempt to keep interest rates higher than those in Germany as long as possible. There was certainly no point in aligning monetary policy as early as half a year before EMU.

ISBN 92-9131-074-3