

Confidence Intervals and Constant-Maturity Series for Probability Measures Extracted from Options Prices

William Melick and Charles Thomas

Introduction

Market participants and researchers have always used information contained in financial prices to analyze economic and financial developments. Over the past three decades, larger and deeper financial markets have increased the amount and variety of information available, while declining computing costs have allowed more sophisticated techniques and models to be considered. As well, the rapid rise in derivatives trading has widened the set of information that can be extracted from the markets. For example, several techniques have recently been developed to extract probability density functions (PDFs) for the underlying asset from options prices—see Melick and Thomas (1997) and references therein.

This paper provides some initial findings on two issues arising from the extraction of PDFs. First, many heavily traded options are traded on listed exchanges with contracts expiring at fixed dates. This imparts a maturity dependence to summary statistics (e.g., moments or probabilities of being above or below a certain price) calculated from the PDFs implied by these options. That is, the summary statistics are limited in that there will only be as many observations as the number of days the option contract is traded (often a year at most), and the statistics will not be comparable because each applies to a slightly different maturity period. These limitations frustrate many attempts to make historical comparisons of summary statistics, or to use such statistics in time-series regression applications. Second, calculations from the PDFs are essentially point

estimates. To date, little work has been done to quantify the uncertainty around any point estimate generated from a PDF.

The paper is organized as follows. Section 1 presents an intuitive explanation of the extraction of PDFs from options prices, and presents examples of analysis from the Federal Reserve Board that makes use of such PDFs. Section 2 develops techniques to construct constant-maturity series for summary measures based on PDFs extracted from exchange-traded instruments. Foreign-currency options on futures will be used as an example, since the constant-maturity options contracts traded in the over-the-counter (OTC) foreign-exchange market can be used as a benchmark against which to compare the results from the exchange-traded market. Section 3 derives confidence intervals for the extracted PDFs using several methods. A summary and conclusions are contained in Section 4.

1 Extracting PDFs: Technique and Examples

Recovering market expectations from options markets is not a new exercise; most familiar is the calculation of implied volatility from options prices. More recently, information retrieval has shifted focus from a single parameter such as volatility to recovering the entire density (or alternatively the stochastic process) for the underlying asset. Recent examples of density recovery are Shimko (1993), Rubenstein (1994), Sherrick, Garcia, and Tirupattur (1996), Bahra (1996), Malz (1997), and Melick and Thomas (1997); examples of the recovery of the stochastic process are Bates (1991) and Malz (1996).¹

As shown in Cox and Ross (1976), a European option's price can be expressed as a discounted product of the probability that at expiration the option is in the money and the expected payoff of the option given that at expiration it is in the money. Therefore, the price of a European call option with exercise price X and an underlying asset price of f can be written as

$$c[X] = e^{-r \cdot t} \cdot \left(\int_X^\infty \gamma(f) df \cdot \left(\frac{\int_X^\infty f \cdot \gamma(f) df}{\int_X^\infty \gamma(f) df} - X \right) \right) \quad (1)$$

where $\gamma(f)$ is a density function for the value of the asset price, f , at the contract's expiration and $e^{-r \cdot t}$ is the discount factor for the period until the contract's expiration. The density function in equation (1) incorporates both

1. Söderlind and Svensson (1997) provide a nice review of the recovery techniques.

actuarial beliefs and preferences towards risk; $\gamma(f)$ is the Martingale-equivalent or risk-neutral PDF for the underlying asset. Armed with an assumption about the functional form of $\gamma(f)$, the price of an option at any strike can be calculated.² Naturally, this calculation can be reversed, and observed options prices can be used to infer the size and shape of the distribution of the underlying asset price at expiration. As shown by Breeden and Litzenberger (1978), options with a continuum of strike prices can be used to trace out the underlying asset price's entire PDF. Unfortunately, in practice the set of strike prices is limited, and some a priori structures or assumptions are needed to map the options prices into a PDF.

In addition, most exchange-traded options are American options on futures, necessitating some alterations to the formulas provided by Cox and Ross (1976). These alterations are found in Melick and Thomas (1997), where bounds on American options on futures are used to express an option's price in terms of the risk-neutral PDF. These expressions are then inverted, via an algorithm that minimizes the sum of squared deviations of actual options prices from predicted options prices, to estimate the parameters of the PDF.

Two caveats apply to the estimated PDFs. First, there are many densities that are observationally equivalent with respect to the information in a set of options prices. It is an a priori structure, such as the functional form for the estimated density, that allows us to choose one particular PDF. For example, it is always possible to construct a series of uniform (rectangular) densities that perfectly fit the observed options prices, although the resulting PDF is often implausible.³ The assumed functional form for the recovered density can be thought of as a smoothed version of these uniform densities. Second, the unknown extent to which attitudes towards risk are incorporated in options prices complicates the interpretation of any implied PDF. By way of analogy, one might attempt to extract perceived probabilities of a fire from the prices paid for fire insurance. Buyers of the insurance are willing to pay more than the price determined by the true or actuarial odds of a fire on their property. However, if there is competition among sellers, and each seller is able to distribute his/her risk so that the policy represents a small increment to the risk of the ultimate insurers' overall portfolio, then the insurance will be priced near its actuarial fair value. Thus, the implied probabilities of a fire recovered from insurance prices are a co-mingling of the true perceived probabilities and risk appetites.

2. Options prices calculated with the Black–Scholes pricing model assume that the price of the underlying asset at expiration will be drawn from a lognormal distribution.

3. See Neuhaus (1995) for examples.


For this reason, analysis at the Federal Reserve Board that makes use of implied PDFs tends to focus on “snapshot” comparisons over relatively short time periods, during which it is at least plausible that preferences towards risk will have remained constant. For example, Figure 1 plots six implied PDFs for three-month Eurodollar futures from the fall of 1997, beginning just prior to the 7 November release of strong October payroll data. The PDFs are separated by only a week or two, so the assumption that risk preferences are constant over the period might be reasonable. If the assumption is correct, then any changes in the PDFs are the result of changes in the market’s perception of actuarial probabilities, likely related to changing views about U.S. monetary policy. Before 7 November, there was relatively little mass above 5.925 per cent (denoted by the vertical line). After the payroll numbers were announced, the mass in the right tail began to increase. The odds of a Fed tightening increased over the period shown in Figure 1. Using simple back-of-the-envelope calculations involving (i) the spread between the Fed funds rate and three-month Eurodollars and (ii) term premiums in the Eurodollar market, a rate of 5.925 per cent would have been consistent with a 25-basis-point increase in the Fed funds rate. The mass of the PDF to the right of 5.925 per cent then gives the odds of a tightening of 25 basis points or more. Comparing the middle panels, the right-hand probability increased almost 4 percentage points between 10 November and 12 November, the date of the November Federal Open Markets Committee meeting. As of 14 November the probability above 5.925 per cent stood at 34.6. At that time, judging from current Eurodollar futures quotes, the right-hand hump evident from 7 November to 13 November appeared to be consistent with two scenarios: (i) a 25-basis-point tightening in December with no further moves, or (ii) no tightening in December and a 50-basis-point increase in February.

As a second example, the three panels of Figure 2 plot, for different dates in October 1997, the market’s implied PDF for the 3-month Euromark futures rate at the expiration of the March 1998 contract on 16 March. The densities are derived from options on Euromark futures that trade in London. The three panels provide an indication of the evolution of market expectations around the time of the 9 October 30-basis-point increase in the German repo rate. The initial effect of the rate increase, aside from shifting the density to the right, was to further delineate two clusters of probability that had been present before the rate hike. Given the distance between the current 3-month interest rate (the spot rate) and the peak of the left-most cluster in the top and middle panels, it seems reasonable to conclude that the left-most cluster likely corresponded to the view that there would be no further tightening by the Bundesbank between October and June 1998 (the end of the three-month period beginning in March). The right-most cluster in the top two panels can be associated with an alternative view that the


Figure 1

PDFs for Eurodollar Futures – December 1997 Contract


6 November 1997 Futures (F) = 5.780




7 November 1997 Futures (F) = 5.800




10 November 1997 Futures (F) = 5.820



12 November 1997 Futures (F) = 5.830



13 November 1997 Futures (F) = 5.850



14 November 1997 Futures (F) = 5.860






Figure 2

**Three-Month Euromark Futures Density Functions
March 1998 Contract**


8 October 1997



10 October 1997



17 October 1997



October repo rate increase was just an initial step, and that additional rate increases would come. The distance between the peaks of the two clusters in the middle panel, about 40 basis points, gives a rough indication of the amount of further tightening expected, as of 10 October, according to the alternative view. After 10 October, the market focused on hawkish comments by various Bundesbank council members, for example the statement by Ottmar Issing that European central banks must act in response to inflationary signs so as not to leave “a mess” for the European Central Bank. These comments all but eliminated market perceptions of an appreciable chance of no further rate hikes. Indeed, expectations of an even larger rate increase, of as much as 100 basis points, began to appear in the form of a far-right mass of probability. The far-right mass in the bottom panel is a bit above 4.5 per cent, the level at which some market commentators suggested at the time that EMU short rates would converge. Of course, since then market views have changed; the current conventional wisdom holds that German rates are likely to show little increase before the EMU.

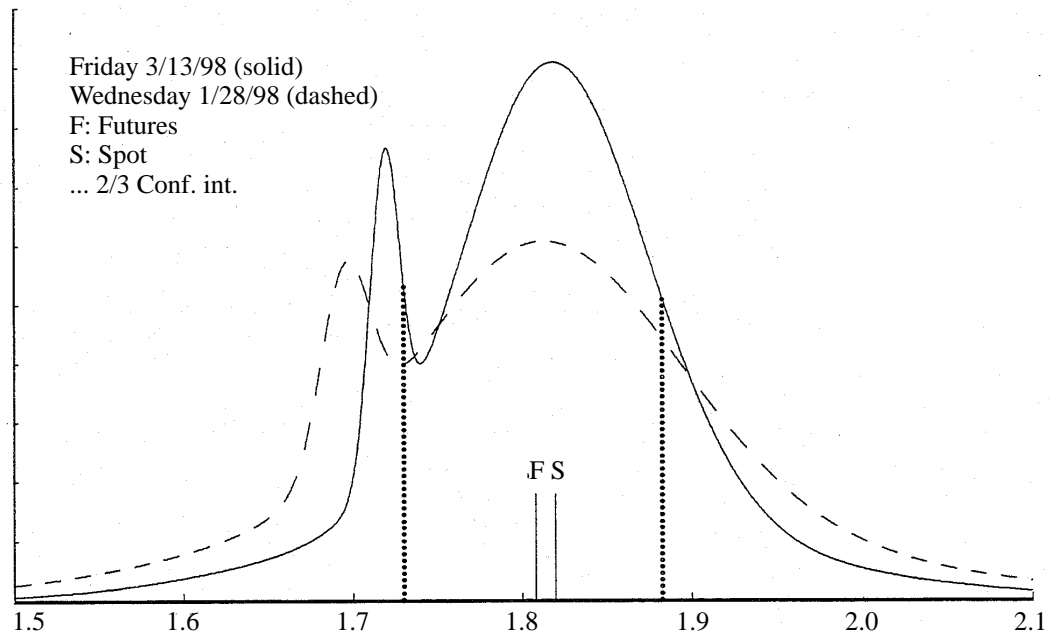
Implied PDFs can also be used as a check on forecasts developed by other measures. For example, the forecasting exercise at the Fed involves judgmental forecasts for exchange rates. Implied PDFs provide a market-based assessment of the reasonableness of these forecasts. Figure 3 shows the PDF and summary measures calculated for the DM/\$ rate from options traded on the Chicago Mercantile Exchange (CME) over the eight weeks or so between forecast rounds. These measures can then be used to gauge the judgmental forecast, providing a useful cross-check on staff perceptions.

Figure 4 provides summary measures calculated from PDFs implied by options, in this case skewness calculations for the Standard & Poor's 500 futures contract. The plot covers the period 10 February through 4 April. The chart plots the ratio of the probability that the futures price on 19 June 1997 will be 10 per cent below its current value to the probability that the futures price on 19 June 1997 will be 10 per cent above its current value. This ratio generally increased from the middle of February through the middle of April, indicating some combination of an increased relative likelihood of 10-per-cent declines and more willingness by market participants to hedge against such declines. Another pattern that is apparent is that sharp drops in the futures price tend to temporarily lower the ratio, although it usually increases in the next several days following the drop.

Finally, Figure 5 presents an analysis of the recent foreign exchange intervention by Japanese authorities that reportedly totalled in the neighbourhood of \$20 billion. The upper panel displays density functions before (Wednesday, 8 April 1998) and after (Monday, 13 April 1998) the

Figure 3

**Deutschemark Distribution and Forecast (DM/\$)
June 1998 Futures Contract**




Summary Statistics

Date	Mean	Median	Low band	High band	Relative 2/3 band	Dispersions 9/10 band	Coef. of skewness	Futures DM/\$
1/28/98	1.8063	1.8042	1.6971	1.9092	0.82	0.88	0.0177	1.7982
2/04/98	1.7951	1.7914	1.7013	1.8954	0.77	0.91	0.0307	1.7873
2/11/98	1.8116	1.8076	1.7149	1.9071	0.78	0.84	0.0375	1.8047
2/18/98	1.8155	1.8103	1.7208	1.9108	0.80	0.79	0.0535	1.8090
2/25/98	1.8130	1.8117	1.7209	1.8995	0.78	0.79	0.0130	1.8073
3/04/98	1.8125	1.8092	1.7316	1.8944	0.73	0.74	0.0381	1.8077
3/11/98	1.8253	1.8234	1.7479	1.8995	0.70	0.75	0.0226	1.8208
3/13/98	1.8121	1.8126	1.7297	1.8825	0.72	0.73	-0.0065	1.8080

Figure 4

**S&P 500 Futures and Relative Probability of a 10 Per Cent Decline to 10 Per Cent Increase
June 1997 Contract**




intervention on Thursday and Friday, 9–10 April 1998. The bottom panel presents summary statistics from the two plotted density functions, as well as a density estimated for Thursday, 9 April.⁴ The intervention removed a small right “shoulder” from the density, pushing probability mass to the left with a concentration near the supposed “line in the sand” of 130 yen per dollar. As is often the case, the dispersion of the density widened after the intervention; this is confirmed by the relative dispersions shown in the summary statistics. All in all, the intervention, from the perspective of the options market, seems to have had a modest effect.

4. The Chicago Mercantile Exchange was closed on Good Friday.

Figure 5

Yen Distribution (Yen/\$)
June 1998 Futures Contract



Summary Statistics

Date	Mean	Median	Low band	High band	Relative 2/3 band	Dispersions 9/10 band	Coef. of skewness	Futures Yen/\$
4/08/98	130.36	130.02	123.99	136.64	1.21	1.27	0.0491	129.94
4/09/98	129.55	129.27	123.12	135.70	1.22	1.30	0.0398	129.13
4/13/98	128.89	128.70	122.46	135.16	1.28	1.39	0.0265	128.47

When material like that in Figures 1 to 5 is presented, it is common for two questions to be raised. First, there is usually an interest in any time series that can be developed from the PDFs, allowing the current measures to be placed in historical perspective and perhaps providing an expectations variable that can be used in a traditional, regression-based macroeconomic estimation. Second, there is often an interest in any confidence intervals that can be placed on the PDFs and their associated calculations. The next two sections deal with these questions in turn.

2 Constant-Maturity Series

Any attempt to construct a time series of measures derived from PDFs implied by options prices is usually frustrated by the fact that exchange-traded contracts approach a given expiration date as time passes, imparting a maturity dependence to most measures that are calculated. In addition, on any given date, several contracts are trading, forcing some selection to be made—a selection that must allow for the replacement of a contract as it expires. These problems are not limited to PDF calculations; they are present even in the relatively simple matter of constructing a time series for a given futures price (Ma, Mercer, and Walker 1992). In the remainder of the section, the problems of time-to-maturity effects and contract-switch effects will be referred to under the general heading of maturity dependence.

With regard to PDFs, two approaches for correcting the problem are possible. First, the maturity dependence could be explicitly incorporated in the functional form assumption for the PDF; in the above call option valuation $\gamma(f)$ would become $\gamma(f, t)$. This is done in the Black–Scholes model, where the standard deviation of the total price change over the life of the option is assumed to vary with the square-root of time to maturity. Butler and Davies (1998) consider such a correction for PDFs implied for the three-month Eurosterling interest rate. Alternatively, the PDFs can be estimated freely, and any calculations based on the PDFs can be subsequently adjusted for maturity dependence.

No matter which method is chosen, researchers are usually left with no way of checking the results of the maturity correction. Usually, the final construct is some sort of time-independent or constant-maturity series that has no analogue in the market. However, this is not the case for foreign exchange options, where a constant-maturity contract trades on the OTC market. Therefore, the foreign exchange options market is a useful and probably unique laboratory for exploring various maturity dependence correction methods.

Table 1 describes the data sets used to compare the measures derived from the OTC and exchange-traded markets.

For the CME data, settlement options prices are used to estimate a PDF for each available contract on every trading day. The options are American, therefore the technique of Melick and Thomas (1997) is used to recover the risk-neutral PDF. The OTC data are one-month European options on the spot exchange rate with prices quoted in implied volatility terms using the Black–Scholes model to translate into currency units.⁵ The

5. See Malz (1997) for a discussion of the OTC market conventions.

Table 1
Data Sets

	Exchange traded			Over-the-counter (OTC)	
	Chicago Mercantile Exchange (CME)				
	Contracts	Trading days	Range	Trading days	Range
Deutschemark	52	7,881	2/24/84–8/30/96	743	9/13/93–8/30/96
Yen	44	6,421	3/17/86–8/30/96	743	9/13/93–8/30/96

implied volatilities are indicative, at-the-money (ATM) quotes⁶ taken from market-makers, they are not transaction prices. Each day on the OTC market, quotes are provided on a one-month contract, so any series based on the quotes is by definition a constant-maturity series. The OTC data will then be used to judge the effectiveness of several maturity corrections constructed for the CME data.

Naturally, the comparison between the CME and OTC data will be useful only to the extent that the two markets are tied to each other. Anecdotal evidence suggests that they are. For example, major trading houses have staffers who monitor the two markets for arbitrage opportunities. Figure 6 provides further evidence of the tight link between the two markets. Plotted here are the indicative ATM quotes from the OTC market (the solid line) and annualized ATM implied volatilities taken from the nearby contract on the CME (the dotted line).⁷ The series are very similar, with a simple correlation of 0.96. Deviations between the two appear to be related to contract-switch points for the nearby contract constructed from the CME data. After the switch to a new contract the CME nearby typically has 120 days to maturity, compared to the OTC's 30 days. Given the usual term structure of implied volatility (higher for longer-dated contracts), after the switch the CME volatility is a bit above that from the OTC dataset. In any event, the two markets are very closely integrated. This implies that the OTC data can be used to judge the effectiveness of any maturity-dependence correction developed for the CME data.


Developing a maturity correction naturally requires a measure calculated from a PDF that suffers from maturity dependence. As an

6. At-the-money options are those for which the strike price is at, or very near to, the current price for the underlying security or commodity.

7. These CME volatilities are calculated using the single call option that is closest to being ATM. The volatility is recovered using the Barone-Adesi and Whaley (1987) approximation. The nearby contract is defined as the contract that, among those with more than 30 days to expiration, is closest to expiration.

Figure 6

Deutschemark Chicago Mercantile Exchange (CME) and Over-the-Counter (OTC) Implied Volatilities: Correlation = .959



egregious offender, the scaled interquartile range (IQR) will be used as the measure derived from the CME PDFs that suffers from maturity dependence. Denoting Q_z as the value such that

$$\int_0^{Q_z} \gamma(f) df = z,$$

the IQR is then given by the ratio

$$\frac{Q_{.75} - Q_{.25}}{F},$$

where F is the futures price. Figure 7 demonstrates the strong maturity dependence of the IQR from the deutschemark CME data. The more distant the expiration date, the larger is the IQR; there is a strong suggestion of a logarithmic relationship. Figure 8 plots the nearby series for the deutschemark IQR, dramatically demonstrating the maturity dependence as the IQR jumps as contracts are rolled over. Three methods will be used to construct maturity corrections for the IQR. Each method will involve using

Figure 7
Scaled Interquartile Range (IQR) and Horizon


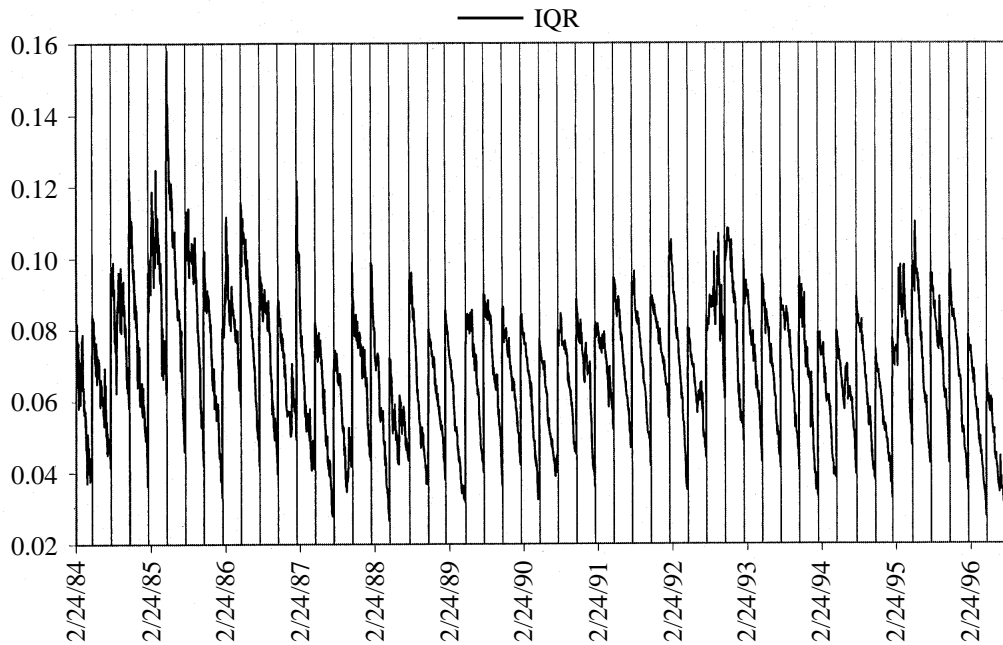


Figure 8
Deutschemark CME Nearby Scaled Interquartile Range (IQR)
Vertical lines denote contract switches



the residuals from a regression that attempts to purge the IQR of maturity dependence. The three regression equations are:

$$\ln(IQR_{i,t}) = \xi_i + \kappa_i \cdot \ln(d_{i,t}) + \mu_{i,t}; \quad (2)$$

$$\ln(IQR_{i,t}) = \alpha + \beta \cdot \ln(d_{i,t}) + \gamma_i + \varepsilon_{i,t}; \quad (3)$$

$$\ln(IQR_t) = \varpi + \rho \cdot \ln(d1_t) + \delta \cdot \ln(d2_t) + \eta \cdot s_t + \lambda_t; \quad (4)$$

where:

$IQR_{i,t}$ = scaled interquartile range from contract i on day t ;

$d_{i,t}$ = days to expiration for contract i on day t ;

$d1_t$ and $d2_t$ = days to expiration when (1) no contract switch
(2) switch;

s_t = dummy for contract switch.

The first two regressions estimate maturity adjustments using the IQR from all available contracts on each trading day, hence the double subscript. The first equation is more general in that it allows the effect of days to expiration to vary across contracts. The second equation imposes the restriction that the effect of days to expiration is the same across contracts, although it allows for a contract-specific constant term. The third equation has a nearby series as the dependent variable—that is, a series where only one contract has been included for each trading day, in this instance the contract with at least 30 days to expiration that was closest to expiration. The double log functional form is used in each equation, as a result of the pattern shown in Figure 7.

The coefficient estimates for the three equations are not really of interest. The real question is whether they have captured all of the maturity-dependence effects. To make that judgment, a nearby time series is created from the residuals from the three equations. For the first two equations, the residual for the contract with at least 30 days to expiration that was closest to expiration was chosen. The residuals for the third equation already form a nearby series. In Figures 9 through 11, residuals from the deutschemark equations are compared with the IQR calculated from PDFs implied by the OTC options prices to determine if the time dependence has been corrected. Table 2 presents correlations between the estimated residuals and the OTC measures.

Both the figures and Table 2 confirm that the approach embodied in equation (4) provides the best fit with the OTC data. That is, the best process seems to be (i) construct the nearby series for the measure of interest (in this case a scaled interquartile range), (ii) correct the series, via regression, for time-dependence and contract switches. The alternative approach, used in



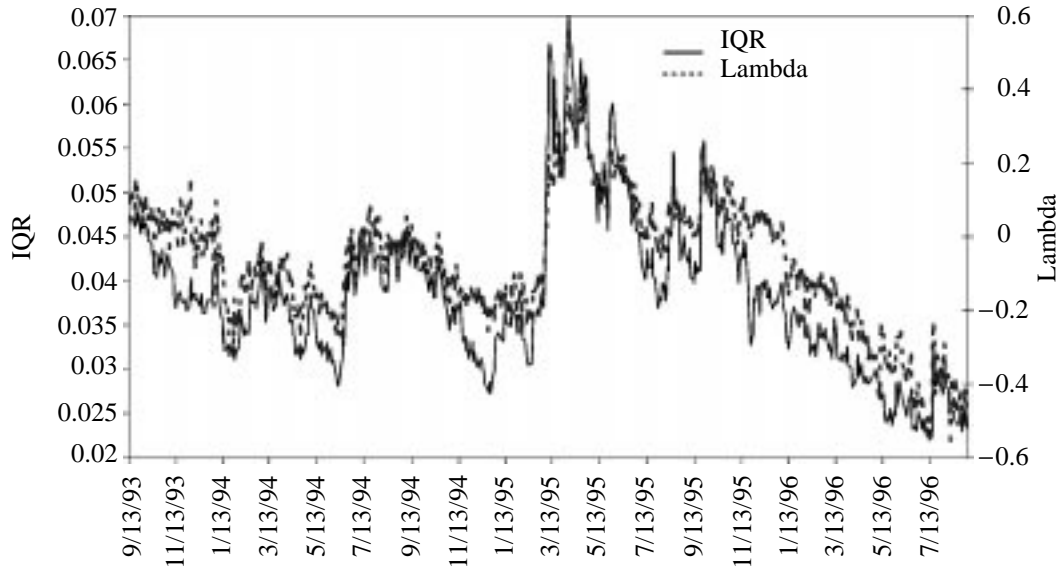
Figure 9**Over-the-Counter Scaled Interquartile Range (IQR) and Mu–Deutschemark****Figure 10****Over-the-Counter Scaled Interquartile Range (IQR) and Epsilon–Deutschemark**

Figure 11**Over-the-Counter Scaled Interquartile Range (IQR) and Lambda–Deutschemark**

equations (2) and (3), of first correcting for maturity dependence and then constructing the nearby series, does not work as well. This is confirmed in Table 3, which presents results from regressions of the change in the nearby residuals from equations (2) and (3) on the contract-switch dummy.

The second column of the table presents regression results with the absolute value of the change in the residual as the dependent variable. For the deutschemark, the contract-switch dummy is not significant in the regressions involving either the change in the residual or the absolute value of the change in the residual. Surprisingly, this is not the case for the yen, where the contract-switch dummy is significant in the regressions for the absolute value of the change in the residual. For the yen, the constructed nearby series jump when the contract is switched, although not up or down in a predictable fashion. This jump in the yen residuals puts them at a further disadvantage to the residuals from equation (4).

3 Uncertainty of the Estimated Distribution

In this section, we discuss issues associated with quantifying the uncertainty surrounding the estimated distributions and the inferences drawn from these distributions. We begin with a short review of the estimation

Table 2

**Correlation Between Chicago Mercantile Exchange Interquartile Range Residuals and Over-the-Counter Interquartile Range
13 September 1993 to 30 August 1996**

	Deutschemark	Yen
$\hat{\mu}$ — equation (2)	.679	.731
$\hat{\varepsilon}$ — equation (3)	.840	.865
$\hat{\lambda}$ — equation (4)	.928	.924

Table 3

***t*-statistic for Contract Switch Dummy in Regressions of
Nearby Interquartile Range (IQR) Residuals on Contract
Switch Dummy**

	Deutschemark 2/24/84 to 8/30/96		Yen 3/7/86 to 8/30/96	
	Dependent variable			
	IQR	Absolute (IQR)	IQR	Absolute (IQR)
$\hat{\mu}$ — Equation (2)	-0.249	0.126	-0.422	-1.609
$\hat{\varepsilon}$ — Equation (3)	-0.405	-0.579	-0.609	-2.163

method and the theory behind it. We then discuss several methods to obtain confidence bounds and provide examples to demonstrate how these methods can lead to very different results.

3.1 Review of the estimation procedure

Throughout this section we will focus on distributions derived from European options, where the theory and computations are relatively straightforward.

From Cox and Ross (1976) we know that the equilibrium price of a European call option can be written as follows:

$$\begin{aligned}
\tilde{c}_t[X] &= e^{-rt} \int_{-\infty}^{\infty} \gamma_t[f] \cdot \max[f - X, 0] df \\
&= e^{-rt} \int_X^{\infty} \gamma_t[f] \cdot [f - X] df,
\end{aligned} \tag{5}$$

where X is the strike price of the option, r is the risk-free interest rate, t is the time remaining until expiry of the option, and $\gamma[\bullet]$ is a risk-neutral, or Martingale-equivalent distribution function over all possible values of the futures price, f , at the date of the option's expiration. The tilde over the c denotes this as the theoretical, equilibrium price of the option.

Similarly, the equilibrium price of a put option can be written as follows:

$$\begin{aligned}
\tilde{p}_t[X] &= e^{-rt} \int_{-\infty}^{\infty} \gamma_t[f] \cdot \max[X - f, 0] df \\
&= e^{-rt} \int_0^X \gamma_t[f] \cdot [X - f] df.
\end{aligned} \tag{6}$$

We note that $\gamma[\bullet]$ is independent of the option's strike price and the option's type (put or call). That is, the same gamma is used to price all options on f at a given point in time.

Observed options prices can differ from these theoretical prices for several reasons. Trades based on liquidity considerations may temporarily move prices away from their equilibrium levels. Similarly, as new information becomes available, it may take some time until it is fully disseminated and incorporated in prices. Finally, prices are quoted at discrete ticks. Our estimations are based on end-of-day settlement prices, which mitigates the first two sources of error. The final source of error, the rounding to the nearest tick, remains a problem whose effect we try to quantify below.

We define the observational error, $\varepsilon^o[X]$, as the difference between the observed options prices, $c[X]$ or $p[X]$, and the theoretical price given above:

$$\varepsilon_{ct}^o[X] \equiv c_t[X] - e^{-rt} \int_X^{\infty} \gamma_t[f] \cdot [f - X] df \tag{7}$$

$$\varepsilon_{pt}^o[X] \equiv p_t[X] - e^{-rt} \int_0^X \gamma_t[f] \cdot [X - f] df \quad (8)$$

Theory places no restrictions on $\gamma[\bullet]$ other than that it is a distribution function, i.e., that it is always positive and integrates to 1 over the range of possible prices for f . To estimate $\gamma[\bullet]$ we approximate it with a flexible parametric distribution function $g[\bullet; \theta]$. Since this is an approximation, it introduces a second source of error. We define the model error, $\varepsilon^m[X]$, as the difference between the theoretical option value under $\gamma[\bullet]$ and the option value under the approximate function $g[\bullet; \theta]$:

$$\varepsilon_{ct}^m[X; \theta] \equiv \tilde{c}_t[X] - e^{-rt} \int_X^\infty g_t[f; \theta] \cdot [f - X] df \quad (9)$$

$$\varepsilon_{pt}^m[X; \theta] \equiv \tilde{p}_t[X] - e^{-rt} \int_0^X g_t[f; \theta] \cdot [X - f] df. \quad (10)$$

Thus, for any estimated set of parameters θ , the total pricing error $\varepsilon[X, \theta]$ can be expressed as the sum of an observational error and a modelling error:

$$\begin{aligned} \varepsilon_{ct}[X; \theta] &\equiv c_t[X] - e^{-rt} \int_X^\infty g_t[f; \theta] \cdot [f - X] df \\ &= \varepsilon_{ct}^o[X] + \varepsilon_{ct}^m[X; \theta] \end{aligned} \quad (11)$$

$$\begin{aligned} \varepsilon_{pt}[X; \theta] &\equiv p_t[X] - e^{-rt} \int_0^X g_t[f; \theta] \cdot [X - f] df \\ &= \varepsilon_{pt}^o[X] + \varepsilon_{pt}^m[X; \theta]. \end{aligned} \quad (12)$$

In practice, the functional form we use for $g[f]$ is a mixture of lognormals. In the examples used below, it is a mixture of two lognormals and $g[f]$ can be written as follows:

$$g[f; \theta] = \pi_1 \cdot \ln[f; \mu_1, \sigma_1] + \pi_2 \cdot \ln[f; \mu_2, \sigma_2],$$

where $\theta \equiv (\pi_1, \mu_1, \sigma_1, \pi_2, \mu_2, \sigma_2)$ and $\ln[\bullet; \mu, \sigma]$ is the lognormal distribution.

To ensure that $g[f]$ is a PDF, we impose that

$$1 \geq \pi_i \geq 0 \quad (i = 1, 2); \quad \pi_1 + \pi_2 = 1.$$

For reasons that are discussed below, we also impose an additional constraint on the dispersion parameters of the individual lognormals:

$$\sigma_i \geq \underline{\sigma} = .02 \quad (i = 1, 2).$$

The estimation, then, takes the following form:

$$\hat{\theta}_t = \underset{\theta}{\text{Argmin}} \left(\sum_{x \in X_{ct}} \epsilon_{ct}^2 [x; \theta] + \sum_{x \in X_{pt}} \epsilon_{pt}^2 [x; \theta] \right),$$

(where X_{ct} and X_{pt} are the available strike prices), subject to:

$$1 \geq \pi_i \geq 0 \quad (i = 1, 2);$$

$$\pi_1 + \pi_2 = 1;$$

$$\sigma_i \geq \underline{\sigma} = .02 \quad (i = 1, 2).$$

This can be solved as a constrained maximum-likelihood or non-linear least-squares problem.

3.2 An example

To illustrate the estimation technique and the issues associated with quantifying our uncertainty, we use an example using options on German short-term interest rates. These options are traded on the London International Financial Futures and Options Exchange (LIFFE) and, with a minor modification, can be treated as European.⁸ On 25 September 1997 there were 23 options trading on the December 1997 contract. The options had strike prices ranging from 94 to 98.5, which translate into interest rates ranging from 1.5 per cent to 6 per cent.⁹ Strike prices vary by 25 basis points and the options prices are quoted to the basis point.

8. Technically, options traded on the LIFFE are American in nature, in that they may be exercised prior to expiration. However, the margining scheme used on the LIFFE ensures that early exercise is never optimal. In addition, the margining scheme eliminates the need to discount the future value of the option when pricing it today. That is, for this contract, we omit the $\exp(-r \cdot t)$ in the pricing formulas.

9. Prior to estimation, the options' strike prices are translated into their interest rate equivalent by subtracting them from 100, and swapping the labels for puts and calls. With this translation, the mixture of lognormals, with its 0 lower support and unlimited upper support, remains a reasonable functional form. Without this translation, the estimated distribution would have an implicit upper support of 100 per cent and give some weight to negative interest rates.

The 25-basis-point spread between the strike prices led us to choose 0.02 as the lower bound for the dispersion parameters (σ). When the dispersion parameter for one of the component lognormals is at this lower bound, more than 95 per cent of the mass of this component distribution lies between adjacent strike prices. The information in options prices is such that it cannot distinguish between two distributions that have the same mass and mean between two adjacent strikes. Thus, this lower bound keeps the optimization routine from trying to distinguish between observationally equivalent distributions.¹⁰ For this day, one of the dispersion parameters was pinned at its lower bound.

The top panel of Figure 12 plots the estimated density on this day. We note the extra mass (relative to a single lognormal) in the range between 3.6 per cent and 4 per cent. Below, we try to quantify our confidence in our estimate of the mass above 3.6 per cent.

The bottom panel of Figure 13 plots the residuals in the option-pricing equation. We note that the maximum absolute error is on the order of 0.005, one-half of a single pricing tick, indicating that the estimated PDF does a good job of explaining the observed options prices. However, we also note that the errors are clearly not independent of strike price and option type. This pattern of the error terms is an issue when choosing a method to construct confidence bands.

3.3 Two methods for constructing confidence bands

The issue at hand is the uncertainty associated with the estimated density function and the uncertainty associated with inferences drawn from it. For example, we may like to make some statement about the probability of f falling above some point or the expectation of f , conditional on it being above some point. In general we have little interest in the individual parameter values themselves. Thus, the standard errors for the parameter estimates produced by most estimation packages are of little use.

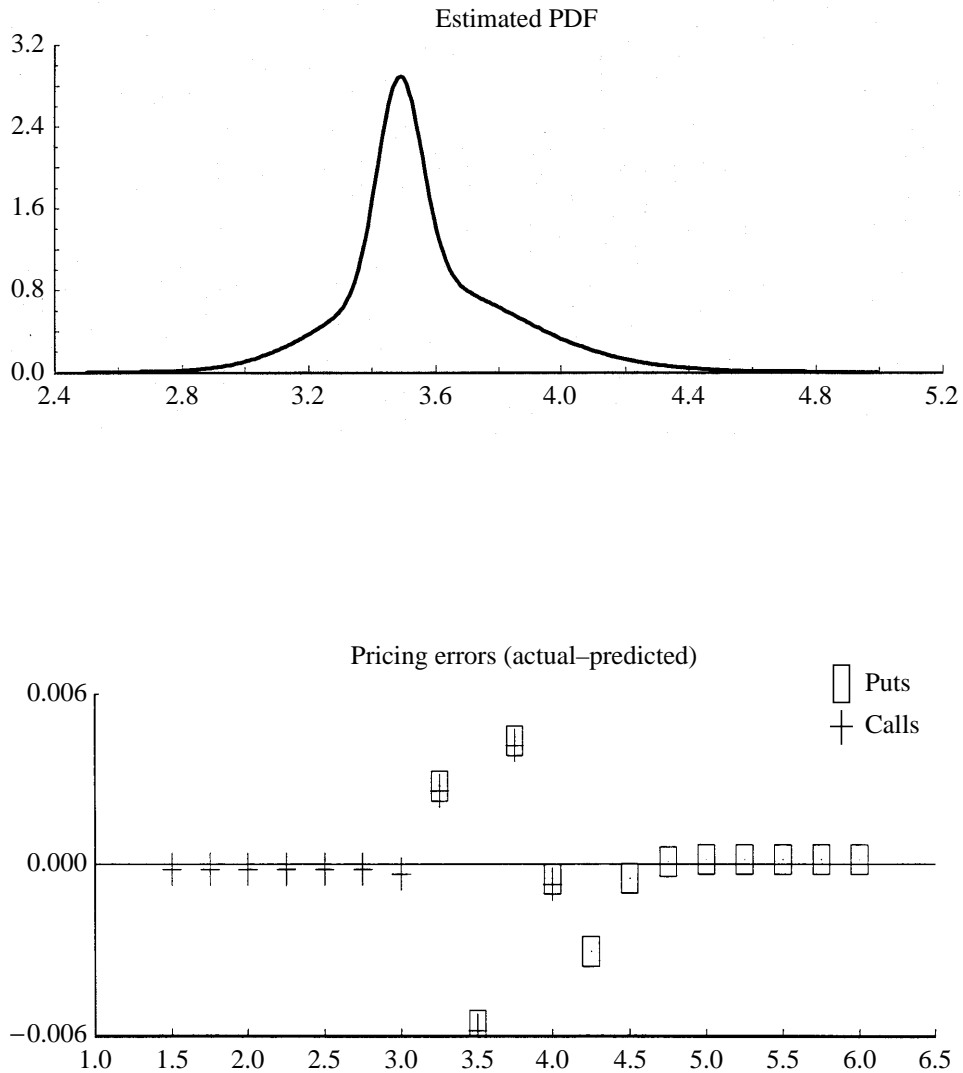
We applied the Monte Carlo and the bootstrap methods to gauge the uncertainty associated with the estimated distributions and the inferences drawn from them.¹¹ The two methods yield very different results, highlighting the special nature of this estimation problem. The constraints placed on the parameters during estimation add some complexity to the

10. See Melick and Thomas (1997, 98–99) for a discussion of observationally equivalent density functions.

11. The delta method gives similar results to the Monte Carlo method for confidence bands around the estimated PDF. However, the delta method requires derivatives for the function of interest, which are not always available. For this reason we focus on the Monte Carlo, which has wider applicability.

Figure 12


Euromark Futures–December 1997 Contract on 25 September 1997




computation of confidence bounds, but they are not the source of the difference between the results for the Monte Carlo and the bootstrap.

3.4 The Monte Carlo method

By constrained maximum likelihood we obtain a point estimate, $\hat{\theta}$, for the parameter vector and from the hessian a covariance matrix, Σ , for the

Figure 13**Monte Carlo Results**

Frequency histogram (500 trials) for probability (futures > 3.6 per cent)



estimated parameters. The Monte Carlo exploits the fact that under certain regularity conditions $(\hat{\theta} - \theta) \sim N(0, \Sigma)$, where θ is the “true” parameter vector. We make 500 draws from this distribution and add them back to the estimated parameters. This yields

$$\hat{\Theta} = \{\tilde{\theta}^1, \tilde{\theta}^2, \dots, \tilde{\theta}^{500}\}$$

a pseudo-distribution for the true parameter vector. From this pseudo-distribution we can construct confidence bands for the PDF and other functions of θ .

The top panel of Figure 13 provides an indication of the uncertainty associated with the estimated PDF. The solid line is the PDF from the parameter point estimates. The dashed lines (which lie almost on top of the solid line) correspond to the 95 per cent confidence bands. From the plot it is clear that the Monte Carlo methodology indicates that we have little uncertainty about where the PDF is.

As noted above, we are often interested in making inferences from the estimated PDFs and would like to know the uncertainty associated with these. For example, from the parameter point estimates we would say the market assigns a probability of 33 per cent to the interest rate being above 3.6 per cent on the options’ expiration date. That is,

$$Pr_{\hat{\theta}}[f \geq 3.6\%] = G[3.6; \hat{\theta}] \equiv \int_{3.6}^{\infty} g[f; \hat{\theta}] df = 33\%.$$

To gauge the robustness of this 33 percentage point estimate, we compute $G[3.6\%; \hat{\theta}^i]$, the probability associated with the realized interest rate being above 3.6 per cent, for each $\hat{\theta}^i \in \hat{\Theta}$. A histogram of these probabilities is given in the bottom panel of Figure 13. As we would expect from the tight bands around the estimated PDF, the Monte Carlo indicates there is little uncertainty around our 33 percentage point estimate.

There are several reasons why we may question the confidence bands coming from the Monte Carlo. The validity of the Monte Carlo method relies on the independence of the error terms and certain regularity conditions. As the bottom panel of Figure 12 shows, it is clear that the errors are not independent, and that the constraints on the estimated parameters invalidate some of the regularity conditions underlying the Monte Carlo.

3.5 The bootstrap method

The bootstrap method is designed to handle situations such as this, where we are reluctant to impose any structure on the error terms. The idea is to create a pseudo-sample by drawing (with replacement) from the

available observations and then estimate the model based on this pseudo-sample. Repeating this many times generates a set of parameter estimates. The distribution of the parameter estimates within this set will mimic the true distribution of parameter estimates, provided our original set of observations is a representative sample of reality.

The top panel of Figure 14 shows the PDF from the maximum-likelihood estimate of the parameters and the 90 per cent confidence bands obtained from the bootstrap method. We note that the confidence bands are much wider than those obtained from the Monte Carlo. The bottom panel plots a histogram, from the bootstrap estimates, of the probability that the interest rate will be above 3.6 per cent. As expected from the wide confidence bands on the PDF, this histogram indicates that, according to the bootstrap, we have little confidence in our point estimate of the probability that the interest rate will be above 3.6 per cent.

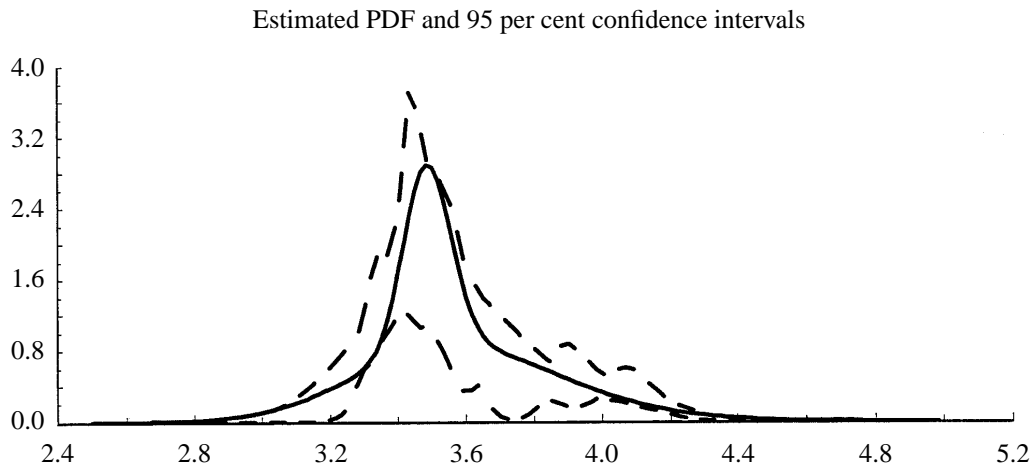
The wide variance of the bootstrap estimates has more to do with the special nature of the estimation problem than with the structure of the error terms. As noted earlier, a call option's price can be expressed as the product of the probability mass for f above the strike price and the expectation of f , conditional on it being above the strike price. There are many distributions for which this product is the same. What allows us to identify the underlying distribution, or choose among those with the same product in the tail above the strike, is the fact that for most of the support we have an observed option price for the next-higher strike. The price of this option at the next-higher strike embodies its own mass and conditional expectation. Taken together, the two options prices identify the mass and conditional expectation between the two strikes. Thus, the estimation routine has the flavor of an inductive construction, where each piece depends importantly on the part that went before.

When the bootstrap constructs its random pseudo-samples from the observations, it does not respect this inductive nature of the estimation. As a result, for many of the pseudo-samples there are relatively large gaps between the strikes, and often the highest and lowest strikes are under-represented. When the gaps between the strikes is large, and when the highest and lowest strikes are given little weight, the distribution is poorly identified, and it is understandable that a wide variety of estimates emerge from the pseudo-samples.

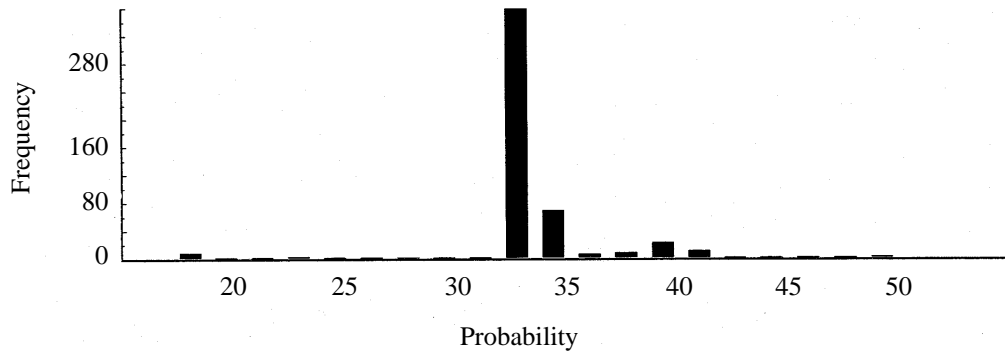
Conclusions

As we demonstrated in Section 1, the information contained in options prices can be used to address many issues of interest to policy-makers. However, many of the most interesting issues involve comparing

Figure 14
Bootstrap Results



Frequency histogram (500 trials) for probability(futures > 3.6 per cent)



market sentiment over relatively long time periods. Recent techniques for extracting information from options prices are only now being adapted to make such comparisons. The results of Section 2 show that regression techniques are promising in their ability to adjust summary measures to allow comparisons over long time spans.

As with any statistical exercise, the estimations based on options data are subject to questions of precision. Section 3 demonstrates the pitfalls of applying standard Monte Carlo and bootstrap methods. In short, the Monte Carlo methodology assumes certain regularity (normality and independence) conditions that are clearly violated in the options data. Thus, we question the extremely tight confidence intervals that the Monte Carlo methodology generates. The bootstrap methodology, which does not require these regularity conditions, still does not adequately quantify the uncertainty. Its problems arise from the particular interdependence of probability measures derived from options prices with adjacent strike prices. Future research in this area will require explicit modelling of this dependence.

References

- Bahra, B. 1996. "Probability Distributions of Future Asset Prices Implied by Option Prices." *Bank of England Quarterly Bulletin* 36 (August): 299–311.
- Barone-Adesi, G. and R. Whaley. 1987. "Efficient Analytic Approximation of American Option Values." *Journal of Finance* 42 (April): 301–20.
- Bates, D. 1991. "The Crash of '87: Was It Expected? The Evidence From Options Markets." *Journal of Finance* 46 (July): 1009–44.
- Breedon, D. and R. Litzenberger. 1978. "Prices of State-contingent Claims Implicit in Option Prices." *Journal of Business* 51 (4): 621–51.
- Butler, C. and H. Davies. 1998. "Assessing Market Views on Monetary Policy: The Use of Implied Risk Neutral Probability Distributions." Paper presented to Bank for International Settlements/Centre for Economic Policy Research conference "Asset Prices and Monetary Policy." January.
- Cox, J. and S. Ross. 1976. "The Valuation of Options for Alternative Stochastic Processes." *Journal of Financial Economics* 3 (1–2): 145–66.
- Ma, C., J. Mercer, and M. Walker. 1992. "Rolling over Futures Contracts: A Note." *Journal of Futures Markets* 12 (2): 203–17.
- Malz, A. 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 (October): 717–48.
- . 1997. "Estimating the Probability Distribution of the Future Exchange Rate from Option Prices." *The Journal of Derivatives* 5 (2): 18–36.
- Melick, W. and C. 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): 91–115.
- Neuhaus, H. 1995. "The Information Content of Derivatives for Monetary Policy: Implied Volatilities and Probabilities." Discussion Paper 3/95, Economic Research Group, Deutsche Bundesbank.
- Rubinstein, M. 1994. "Implied Binomial Trees." *Journal of Finance* 49 (July): 771–818.
- Sherrick, B., P. Garcia, and V. Tirupattur. 1996. "Recovering Probabilistic Information from Option Markets: Tests of Distributional Assumptions." *Journal of Futures Markets* 16 (5): 545–60.
- Shimko, D. 1993. "Bounds of Probability." *Risk* 6 (4): 33–7.
- Söderlind, P. and L. Svensson. 1997. "New Techniques to Extract Market Expectations from Financial Instruments." *Journal of Monetary Economics* 40 (October): 383–429.