Distributions Implied by Exchange Traded Options: A Ghost's Smile ?

Martin Cincibuch*

Abstract

A new and easily applicable method for estimating risk neutral distributions (RND) implied by American futures options is proposed. It amounts to inverting the Barone-Adesi and Whaley method (1987) (BAW method) to get the BAW implied volatility smile. Extensive empirical tests show that the BAW smile is equivalent to the volatility smile implied by corresponding European options. Therefore, the procedure leads to a legitimate RND estimation method. Further, the investigation of the currency options traded on the Chicago Mercantile Exchange and OTC markets in parallel provides us with insights on the structure and interaction of the two markets. Unequally distributed liquidity in the OTC market seems to lead to price distortions and an ensuing interesting 'ghost-like' shape of the RND density implied by CME options. Finally, using the empirical results, we propose a parsimonious generalisation of the existing methods for estimating volatility smiles from OTC options. A single free parameter significantly improves the fit.

*Czech National Bank and Center for Economic Research and Graduate Education in Prague. I would also like to acknowledge the support of the Czech Grant Agency under Grant 402/00/0814. The article is based on research pursued during an Academic Project for the Centre for Central Banking Studies (CCBS) at the Bank of England. I would like to thank everybody in the CCBS for their strong support and all of their kind help.
I would also like to thank Ron Anderson, Aris Bikos, Neil Cooper, Joe Ganley, Gordon Gemmill, Toni Gravelle, Allan Malz, Richhild Moessner, Jose Noguera and Nikolaos Panigirtzoglou for their valuable suggestions and worthwhile discussions. All errors and omissions are mine alone.

1 Introduction and summary

Risk neutral distributions[1] summarise much of the available information associated with market prices and therefore they are attractive for market, academic and central bank economists. As Bliss and Panigirtzoglou (2000) note, RNDs estimated from liquid assets may be used by market participants for pricing exotic derivatives. Further, from the point of view of the central bank, option markets provide information in addition to that provided by spot and futures markets, and implied risk neutral distributions represent a convenient tool for interpreting this additional information. Clews, Panigirtzoglou and Proudman (2000) describe the methods used at the Bank of England for estimating distributions implied by interest rate futures, which enter as a regular input at its Monetary Policy Committee briefings. At other central banks, RNDs are estimated from currency options and used for monitoring the foreign exchange market, for example, at the Bank of Canada or the Czech National Bank.

Next, due to the forward-looking nature of option prices, accurate estimates of implied distributions might arguably enhance VaR modelling. There is an increasing amount of literature pointing to the shortcomings of risk modelling based on the assumption that market price data follow a stochastic process which only depends on past observations. (e.g. Daníelsson (2000), Ahn et. al. (1999), Artzner et. al. (1999)). On the other hand, empirical evidence suggests that option-based measures of uncertainty are a better predictor of future volatility of the underlying asset than statistical time-series models. Christensen and Prabhala (1998) offer such evidence for S&P index options; Jorion (1995) for currency options for major currency pairs; and Bouc and Cincibuch (2001) for Czech koruna options. The question whether options also carry useful information about the fat-tailedness of the distribution of future assets' returns and about other deviations from lognormality is an important one from the risk management point of view. And indeed, the first step to answering such a question is to have a reliable estimate of the risk neutral distribution implied by the option prices.

The results presented in this article are threefold. First, we discuss a new method for estimating risk neutral distributions (RND) implied by American futures options. In contrast to other methods that utilise lower and upper bounds for the prices of American options, this method amounts to inverting the Barone-Adesi and Whaley method (1987) (BAW method)[2] to get the BAW implied volatility from the option prices and then approximating the BAW volatility smile with the weighted smoothing spline. Using the full history of yen[3] futures options traded on the Chicago Mercantile Exchange (CME) and comparing them with relevant option prices from the interbank over-the-counter (OTC) market[4], we found good support for the hypothesis that the BAW volatility implied by a American futures option does not differ significantly from the Black-Scholes volatility implied by the price of the European option with the same exercise price and maturity. Further and more importantly, we found that BAW volatilities derived from a pair of put and call CME options with the same exercise price are very close to each other. Indeed, the model independent and arbitrage based put-call parity stipulates that Black-Scholes volatilities implied by European puts and calls with the same exercise price are equal. Therefore, we argue that the BAW inversion leads to an appropriate 'European' volatility smile and that the approach is a legitimate RND estimation method. The method is numerically stable and easy to apply, and it circumvents convergence problems often encountered with parametric methods.

Second, by investigating the CME and OTC markets in parallel, we gained insight on the structure and interaction of the two markets. Intensive arbitrage seems to take place more for certain exercise prices than for others, which may be explained by the varying OTC liquidity over the price space. It leads to price distortions and an ensuing interesting 'ghost-like' shape of the RND density implied by CME options.

Third, using these empirical results, we show how it is possible to improve RND estimation from a low number of OTC option prices. In the literature, two quadratic extrapolation methods have been suggested, but they either break the non-arbitrage constraints imposed on the volatility function or they do not fit well the observed CME data. To overcome this problem, we suggest a parsimonious generalisation of these methods, which significantly improves the fit.

An improved way of estimating foreign exchange RNDs from the OTC market is useful, because data from this market have several convenient features from the practical point of view. In general, the OTC FX market is quite deep, and contrary to exchange traded options, it exists for most currencies. In addition, OTC quotes are usually available for fixed maturities. However, it is often the case that only a limited number of benchmark exercise prices are readily accessible[5] from the OTC market and therefore some extrapolation has to be made. In essence, we enhance this extrapolation by information from CME prices.

In the next section the equivalence between estimating the volatility smile and risk neutral distribution is established, and the methods of estimating RNDs are classified according to their generality. The third and fourth sections present the method for estimating RNDs from American futures options, and the fifth describes data and estimation results. In the sixth section, interaction between the OTC market and CME is discussed. Methodology improvements for estimating RNDs from OTC data are proposed in the seventh section. The eighth section provides conclusions.

2 Classification of methods for estimating RNDs from European options

Under the assumption of no arbitrage and frictionless markets, the price of a traded security can be expressed as an expected discounted security payoff; the expectation is taken with respect to an appropriate risk-neutral density f (Cox, Ingersoll and Ross (1976), Ross (1976)). For European call option price c with strike X and maturity T, this result is formalised as

$$c(S, X; r, T) = e^{-rT} \int_0^\infty \max(S_T - X, 0) f(S_T) dS_T, \tag{1}$$

where r and S denote the appropriate domestic interest rate and the current spot price and S_T the random spot price at the option's maturity T. Formula (1) may be then understood as a definition of risk neutral density. Breeden and Litzenberger (1978) showed that if the risk neutral distribution is continuous, its discounted density is equal to the second derivative of the European call option price with respect to the strike price:

$$f(X) = e^{rT} \frac{\partial^2}{\partial X^2} c(S, X; r, T). \tag{2}$$

Although it is well established that the assumptions of the benchmark Black-Scholes model (1973) are too restrictive and do not hold in the real world[6], their formula is still used in practice. The volatility parameter is the only unobservable variable in the Black-Scholes formula (This standard relationship appears in Appendix A.2 as Equation (26)), and therefore this formula can be used as a mapping that converts volatilities into prices and vice versa. Deviations of the real data from the benchmark model are reflected by the fact that the quoted so-called implied volatility, is not constant across strike prices and maturities[7]. Estimating RNDs from option prices can be seen as a relaxation of the distributional assumption of the Black-Scholes model. Indeed, under this model, the RND is lognormal and the volatility smile degenerates to a horizontal line.

Equations (1) and (2) and Black-Scholes pricing formula (26) establish the equivalence between risk neutral distributions with the mean equal to the forward rate and volatility smiles. If an RND density f is given, then prices of call options may by derived using (1) and consequently the implied volatility smile can be calculated from them with the inverted Black-Scholes formula (29). Conversely, the Black-Scholes formula (26) transforms a volatility smile to the call price function and the RND density is obtained from it with the Breeden and Litzenberger equation (2).

Equation (1) or (2) underlies any method of constructing RNDs from European option prices. Bahra (1997) and Chang and Melick (1999) sort the methods from the operational point of view. While Bahra (1997) recognises four types of techniques, Chang and Melick (1999) dichotomise the methods according to the extent to which they are based either on Equation (1) or (2). Another, and in our view, natural approach is to classify the methods according to their generality.

In the first and most restrictive category, we put those methods which assume a specific stochastic process driving the security price. A prominent example of this approach is obviously the Black-Scholes model (1973). More recently, Malz (1996) fitted a lognormal-jump diffusion process to FX options.

Secondly, a more general approach is to assume only a specific functional form of the terminal distribution. As Melick and Thomas (1997) note, it is more general, because a single terminal distribution might result from different stochastic processes. Melick and Thomas (1997), Bahra (1997) and Gemmill and Safeklos (1999) assume that the density of the terminal distribution is a linear combination of lognormal densities. Shimko (1993) does not directly assume a functional form for the distribution, but equivalently imposes a functional form of the volatility smile, assuming that implied volatility is a quadratic function of strike prices. And similarly, Malz (1997) assumes that implied volatility is a quadratic function, but he uses another measure of moneyness, an option's delta. He argues that this approach is superior to fitting functions in strike price space, because it avoids the violation of specific no-arbitrage conditions for the volatility function. In addition, as Bliss and Panigirtzoglou (2000) note, due to its character the delta space gives more weight to exercise prices close to the at-the-money (ATM) value. Thus, the delta space allows for a better approximation of these central options, which presumably have more informational value. This method is suitable for OTC options, for which only a small number of data points are often available. The shortage of observations makes virtually all other methods of little use, because for example three points are not enough to fit a mixture of two lognormals.

Finally, there are approaches that make no assumptions about the global nature of the density function, which are sometimes called non-parametric methods. Their advantages are discussed by Jackwerth and Rubinstein (1996), who construct a smooth density function, constrained by Equation (1), by minimising the norm that measures the density function's second derivative. Buchen and Kelly (1996) take an interesting step by deriving a functional form (dependent on the number of observed option prices) of the distribution that maximises entropy given an observed set of option prices. Bliss and Panigirtzoglou (2000) combine the Malz approach(1997) of using delta smiles with the Campa, Chang and Rieder (1997) method of smoothing spline interpolation, another nonparametric method.

Not surprisingly, comparative studies have shown that this last approach is more flexible than less general methods. Cooper (1999) compares the mixture-of-lognormals method with the volatility-smile-smoothing- method by running Monte-Carlo tests and finds that the latter outperforms the former. Also, Bliss and Panigirtzoglou (2000), using short sterling and FTSE 100 index contracts, find that the volatility smile method performs better than the double-lognormal one in terms of robustness.

Because of its flexibility and relative ease of implementation, we chose this last approach as a base method for our research. Deriving RNDs from prices of American futures options, we first control for the early exercise premium and for the mismatch in maturities between options and underlying futures. By doing so, we can construct approximate European equivalents to American futures options. Then, we apply the Bliss and Panigirtzoglou (2000) approach to these implied volatility approximations. In order to take into consideration possible price distortions due to low trading activity, we amend their method by weighting individual observations by a gently increasing function of their trading volume.

3 Estimating RND at maturity from American futures options

There are several approaches for estimating RND from European options, but not many researchers have dealt with American options. Melick and Thomas (1997) assumed that the terminal RND is a linear combination of three lognormal distributions and derived upper and lower bounds for American option prices in terms of this mixture distribution. They weighted these bounds differently for out-of-the-money (OTM) and in-the-money (ITM) strikes, and in the end, they fit a combination of the bounds to observed option prices to estimate ten parameters of the distribution. Flamouris and Giamouridis (2002) also used similar bounds for American futures options but assumed a different functional form of the terminal distribution. Instead of a linear combination of lognormals, the RND was estimated as the sum of a lognormal density and a combination of the third and fourth order terms of the Edgeworth series expansion around the lognormal distribution. As a single weight of the bounds was kept for all strikes, only four parameters remained for estimation.

Two problems might be associated with these methods. Notwithstanding the prominence of the lognormal distribution, which is in their centre, both methods are based on an ad-hoc specification and it might limit their ability to capture interesting economic phenomena. Second, because of the rather complicated functional forms involved, both methods lead to complex optimisation techniques.

A different approach that is proposed here to estimate RNDs from American futures options does not impose a priori any functional form of the distribution and is easy to implement using fast and stable numerical procedures. In essence, we attempt to reduce the American futures option problem to the European one and look for the implied volatility. For this purpose, we employ an inverse procedure to the analytical approximation for the prices of American futures options[8] derived by Barone-Adesi and Whaley (1987) (BAW). We empirically test the appropriateness of this approach.

Barone-Adesi and Whaley (1987) report extensive comparisons of American option prices calculated via their approximation with results obtained by precise methods like finite difference or compound option methods. The BAW method performs very well for various maturities, implied volatilities and levels of option moneyness. The largest reported mispricing amounts to three tenths of a percent of the dollar option price, which is negligible both in the context of volatility of option prices and market bid-ask spreads. The method is widely used in practice because of the ease of its implementation and its speed, which is also the greatest advantage from the point of view of this article. There are hundreds of thousands of strike price-option price pairs in the dataset and therefore speed, stability and good convergence of the numerical procedures are crucial[9]. In comparison with the finite difference method, Barone-Adesi and Whaley (1987) report that the method is about 2000 times faster. Also, Broadie and Detemple (1996), who conducted a large scale evaluation of many recent methods for computing American option prices, report method performance of this order.

Whaley (1986) employed the BAW method to investigate the validity of the underlying model. Similarly Melick and Thomas (1997) used the method on a single lognormal benchmark model to evaluate the mixture lognormal method. It is worth emphasising that we use the BAW in a different way, analogous to how the market uses the Black-Scholes formula in the European option context; i.e., as a mere mapping between option prices and volatility. However, since standard assumptions underlying the Black-Scholes (and also BAW) model are not valid in reality, it remains to be seen whether the method of inverting the BAW method is able to get rid of the actual early exercise premium. In other words, our proposed method for estimating RNDs hinges on the hypothesis that the BAW volatility[10] implied by the price of an American option equals the Black-Scholes volatility implied by the price of the corresponding European option with the same maturity and traded at the same time.

The intuition behind of why such a hypothesis might be reasonable stems from widespread quoting of options in volatility terms. It might well be the case that, in general, market participants suppose that violations of benchmark model assumptions are completely reflected in a (European) volatility smile and its term structure. Thus, they price American options by transforming this smile or term structure via BAW or a similar method. Indeed, this behavioural assumption is supported by the fact that the Chicago Mercantile Exchange, which is one significant market trading in American futures options, uses the BAW method as a standard pricing model[11].

We validate the legitimacy of the hypothesis by two empirical tests. Firstly, we test whether BAW volatility implied by an American futures option does not differ significantly from the Black-Scholes volatility implied by the price of the European option with the same exercise price and maturity. Secondly, we check whether the BAW volatilities derived from a pair of put and call CME options with the same exercise price do not differ. Indeed, the model independent and arbitrage based put-call parity stipulates that the Black-Scholes volatilities implied by European puts and calls with the same exercise price are equal.

4 Adjustment for different option and underlying futures maturity dates

Having corrected for the early exercise premium using the inversion of the BAW method, we have to make yet another adjustment due to possible maturity mismatches that might occur between currency options and their underlying futures.

For example, CME currency futures mature four times a year in a so-called March quarterly cycle (i.e., March, June, September and December). Currency futures options mature every month and the underlying futures contract of an option is the nearest futures contract in the March quarterly cycle whose termination of trading follows the option's last day of trading by more than two business days. Therefore, an option maturing in January, for example, is written on the future maturing in March. Even for an option maturing in March, there is some maturity mismatch of about two weeks, because options generally mature on the second Friday immediately preceding the third Wednesday of their contract month, but futures mature on the second business day immediately preceding the third Wednesday of their contract month[12].

Let t denote the time of the trade, τ the maturity date of the option and T the maturity date of its underlying futures. The price of a European option written on such a future with strike price X can be denoted as $c_t\left(F^T, X, \tau\right)$. We aim to transform it to the option on a future that matures also at τ. Note also that $c_t\left(F^\tau, X, \tau\right) = c_t\left(S, X, \tau\right)$, where $c_t\left(S, X, \tau\right)$ denotes a European option on a spot rate with maturity date τ. This transformation is allowed by a well known homogeneity property of European options (see the Appendix A.3). Denoting domestic and foreign interest rates by r and r^* and taking into account the option's homogeneity and that $F^T = F^\tau e^{(r-r^*)(T-\tau)}$, we may write

$$c_t\left(F^T, X, \tau\right) = e^{(r-r^*)(T-\tau)} c_t\left(F^\tau, \bar{X}, \tau\right) = e^{(r-r^*)(T-\tau)} c_t\left(S, \bar{X}, \tau\right), \tag{3}$$

where

$$\bar{X} = X e^{-(r-r^*)(T-\tau)}. \tag{4}$$

Equations (3) and (4) show that, in order to transform option prices, it is enough to discount both the strike price and option price by a factor of $e^{-(r-r^*)(T-\tau)}$.

5 Application: Currency options from the Chicago Mercantile Exchange

The method described above is applicable for American futures options in general, but it is designed primarily for the Chicago Mercantile Exchange - one of the largest organised markets for various futures and futures options. As was noted above, there are two ways for checking the empirical relevance of the proposed technique for eliminating the early exercise premium. The first one is to evaluate its results using the prices of actual European options. Therefore, we constructed BAW volatilities implied by CME prices of currency options and made comparisons with suitable volatilities derived from OTC currency option market quotes. If our hypothesis is right and the early exercise premium is priced so that market volatilities of the European options are input into the BAW (or into a similar model), then arbitrage between the two markets would take place and BAW and European volatilities would be close to each other up to the difference related to the transaction costs. The second check of the hypothesis involves the put-call parity, which holds for European, but not for American options. As a consequence of the parity relationship, if the early exercise premium is accounted for correctly (or in the same way as the market accounts for it), BAW volatilities implied by CME put and call options with a common strike would be equal.

For CME traded options, we have a full contract history of close-of-business data for dollar-yen currency futures and dollar-yen currency futures options. The total option turnover over the period 1992-2000 was about 12 million contracts, which represents approximately 1000 billion dollars in a notional amount. We calculated BAW volatility from daily settlement prices for all actual trades. Further, we adjusted strike prices for the difference in maturities between CME options and their underlying futures[13]. Next, we transformed price-BAW volatility space into delta-BAW volatility space. Then, following Bliss and Panigirtzoglou (2000), we approximated the CME delta volatility smile by a weighted smoothing[14] spline. The weights were calculated using logarithms of the trading volume[15].

For the OTC market, we worked with time series of dollar-yen option quotes since 1992 provided by two large market makers. The data consist of time series of at-the-money-forward (ATMF) volatilities, 25-delta risk reversals and 25-delta strangles for one-month options together with appropriate forward rates. From the OTC quotes, we backed out implied volatilities for three exercise prices (for a technical description see Appendix A.2).

The statistical results of comparing OTC volatilities derived from OTC contracts and BAW-implied volatilities for 25-delta, ATMF and 75-delta strikes for maximum maturity mismatch 6 calendar days are summarised in Table I. Although the statistical hypothesis that the OTC volatilities and BAW volatilities are equal may be comfortably rejected at any reasonable level of statistical significance, the differences are quite small on average. This result could be expected if arbitrage between the two markets takes place. With a very high degree of confidence, the mean difference is only one or two tenths of a percentage point, which is a value of little economic significance.

Figure 1 demonstrates the distribution of the difference for ATMF volatility over time. While the difference is quite low for most of the observations[16], from time to time the difference is greater. The most distinct are observations from October 1998 when the difference was more than 10 percentage points. They might be naturally explained by aberrant market conditions of that period and ensuing low liquidity hindering effective arbitrage. When this exceptional period is excluded from the sample, means and standard deviations of the differences become smaller. The resulting standard deviation of the difference of about 0.5% is consistent with the bid/ask spreads usually observed in the option market[17]. Yet, for some remaining observations the difference is still larger than 1 percentage point. One reason for which it might happen is that the arbitrage between the markets is complicated by the fact that the maturity of the CME contracts perfectly matches the maturity of benchmark OTC contracts only several times a year. In the analysis, however, we took into account contracts with only an approximate maturity match. We set the maximum maturity mismatch between CME contracts and the benchmark one-month OTC contract arbitrarily to 6 days (less than one week). Indeed, the liquidity of contracts with broken maturity dates is significantly lower, and consequently, wider bid/ask spreads make arbitrage less powerful.

Another reason for a wider difference between volatilities, which is sometimes observed, might stem from the time discrepancy of price quotes. While the OTC volatilities were attributed to actual trades that took place during the trading day in London, the CME data represent settlement prices that are determined after the close of business in Chicago. The volatility might jump in the meantime if some significant information hits the market.

As was discussed above, another check of the proposed method for analysing American currency futures options involves put-call parity. It turns out that the parity is satisfactorily fulfilled in most of the cases when the best fit could be found, not surprisingly, for most liquid strikes. A typical example of this result is shown in Figure 2, where crosses represent BAW volatility implied by CME calls and squares show BAW volatility of CME puts. Call and put BAW volatilities are close to each other for common strikes.

Statistically, the dataset contains 25,775 strike prices for which both put and call options were traded and prices were available. In the majority of cases, the distance between call and put BAW volatilities with a common strike was quite small. Table II shows the distribution of these distances. Summary statistics of this sample (with the 10 biggest outliers disregarded as errors) are presented in Table III.

6 Interaction between markets: A ghost's smile ?

Figures 2 and 3 illustrate interactions between the OTC and CME markets. Figure 2 shows delta space CME and OTC volatility smiles as of August 4, 1998[18] when the maturity date of options made them directly comparable with one-month OTC contracts. The CME smile represents a trading volume of 1247 call contracts and 2972 put contracts[19]. However, the volume was not evenly distributed over the exercise prices, with the bulk of the trades taking place for exercise prices around 25, 50 and 75 delta. The solid line is a weighted smoothing spline interpolating BAW volatilities derived from settlement prices of OTM options. The three shadowed circles represent OTC market quotes and the dashed line is the Malz (1997) quadratic volatility smile. Transformation of the smiles into risk neutral densities is shown in Figure 3, where the thick solid line represents the CME distribution and the dashed line is the OTC Malz density.

As was discussed above, a small difference between OTC quotes for 25-delta calls, ATMF options and 25-delta puts and the respective CME options suggests that arbitrage was taking place between the markets. On the other hand, it is obvious that the Malz approximation does not fit the farther-out CME strikes well. Reflecting the more pronounced smile, the CME distribution is more concentrated around the mean and has heavier tails than the Malz density[20].

The CME density in Figure 3 exhibits three modes, which give it a somewhat 'ghostly shape'. We might discuss whether these spikes are consequences of some real economic phenomenon or whether they are just artefacts stemming from the numerical method. Since these three spikes occur very often among the daily observations and since they persistently arise for any reasonable value of smoothing parameter of the natural spline, we argue that they are not a fluke. We think that the proximity of the modes to the position of the benchmark OTC strikes suggests that they are induced by arbitrage interaction between the OTC and CME markets. In the OTC FX option market, some strike prices (and maturities) play the role of benchmarks. The most important strikes are ATMF and 25-delta call and 25-delta put. Indeed, most of the banks operating on the OTC market are ready to price an option of any strike and maturity. However, liquidity for a nonstandardised strike and broken maturity, for example, 5-week 11-delta call option, is relatively low. Therefore, the bank which would sell such an option has to price it using its internal model. These models are calibrated using the volatilities of the benchmark contracts. In contrast, benchmark volatilities are discovered by supply and demand. Thus, higher liquidity associated with OTC benchmark strikes makes the arbitrage possible. The gravity of the deeper OTC market is detectable even in Figure 2, where some bending of the CME smile can be seen. In another sense, however, the nature of the spikes is also ghostly, because they do not reflect the shape of market expectations. They are rather a consequence of market imperfections.

Another natural question which Figure 2 evokes is whether the steep volatility smile for deltas farther out and the ensuing fatness of the distribution tails reflect the genuine shape of the RND or if it is a result of some other market imperfections. Indeed, the sudden change in the first derivative and the steepness of the delta smile is an artefact due to the nonlinear transformation from price space to delta space, because it puts more weight on strikes around the ATMF position and 'shrinks' the OTM regions of price space. Nevertheless, it does not affect the relative value ATMF and OTM options; and indeed, the function looks more natural in the dollar space.

Let us consider the hypothesis that relatively high prices of OTM options are caused by risk or liquidity premiums. In general, the price of an option might be very volatile. Therefore, risk-averse speculative buyers should require some additional compensation for risk and thus should be prepared to pay lower rather than higher prices. Conversely, speculative sellers of options might demand a higher price to compensate for the higher risk of their liabilities. Indeed, market participants also trade options to hedge their risks, but again there are presumably hedgers in both directions. Therefore, unless some asymmetry in hedging needs or market power exists between sellers and buyers, it is not clear how risk considerations might explain the volatility smile. Similarly, low liquidity of OTM options increases uncertainty about the future option price, but without a significantly different position between buyers and sellers this should affect only bid/ask spreads, not average realised prices.

Since we cannot think of any reasonable asymmetric factors that would affect demand and supply of currency products, we hypothesise that the pronounced volatility smile actually reflects the market perception of the heavy-tailed underlying risk-neutral distribution.

7 Improving the method for the estimation of OTC RNDs

Contrary to exchange traded options, the OTC market exists for most currencies and it is quite a deep market. In addition, OTC quotes are usually available for fixed maturities. However, it is often the case that only a limited number of benchmark exercise prices are readily accessible from the OTC market, and therefore, some extrapolation has to be made. In the literature, two quadratic extrapolation methods have been suggested, but they either break the non-arbitrage constraints imposed on the volatility function or they do not fit observed CME data well. Here we propose some improvement.

Both methods use quadratic functions. The first one, which was proposed by Shimko (1993), fits volatilities in the dollar space, and the second one, suggested by Malz (1997), uses Black-Scholes delta instead as a measure of the options' moneyness. Malz (1997) noted that for far OTM strikes the quadratic function in the dollar space breaks non-arbitrage constraints for the volatility function. However, as the above discussed comparison of CME and OTC smiles shows, the quadratic extrapolation in the delta smile underestimates volatilities of the OTM exchange traded options. Figure 4 shows the typical result that Shimko's (1993) function tends to fit data better, but somewhat overestimates far OTM volatilities.

Malz's underestimation of the smile is a consequence of the delta space's nature, which attributes too low weights to OTM strikes. Therefore, we seek a different moneyness space that would allow for a better fit of the quadratic function. Since the transformation between dollar and delta spaces is governed by the cumulative distribution function of the standard normal distribution, it is natural to generalise it by using normal distribution with variance as an additional free parameter. In Appendix A.2, the OTC quoting convention is summarised and the delta function is defined by Formula (30). Therefore, if the cumulative distribution function of the normal distribution with mean m and standard deviation G is denoted as $N(., m, G)$, the generalised delta function might be defined as

$$\Delta_c^G (X, \sigma; S, T, r, r^*, G) \equiv e^{-r^* T} N(d_1(X), 0, G). \tag{5}$$

The parameter G determines how much weight is attributed to far OTM strikes, and other symbols have the usual meaning. It can be determined using exchange traded options, or it might be estimated from the shape of the actual distribution of returns from the spot exchange rate.

Figure 4 illustrates how this approach improves the fit for $G = 6$, and Figure 5 summarises how the fit of the quadratic smile in the generalised delta space depends on the parameter G. In order to get this relationship, for a given G we summed distances of BAW volatilities from the G-quadratic smile which extrapolates 25-1-delta, ATMF and 75-1-delta CME BAW volatilities for a given range of maturities. From Figure 5, it is obvious that the fit monotonically improves with G, but for higher values of G the improvement is only marginal. However, the use of a generalised delta with G values of 4 or 5 makes the fit much better than is the case in Malz's original smile based on $G = 1$. When the same data are used as for Figure 3, but the OTC volatilities are extrapolated in generalised delta space, then all the statistics characterising CME and OTC densities are close to each other, including kurtosis.

Indeed, this generalised delta approach is not necessarily constrained to the CDF of the normal distribution. Another distribution might be alternatively used. It is easy to show that the generalised delta space approach ensures non-arbitrage constraints on the volatility function.

8 Conclusion

Three main results are presented in this article. First, good empirical evidence was found for the hypothesis that volatility smiles calculated from CME American currency futures prices by inversion of the Barone-Adesi and Whaley pricing model (1987) do not differ significantly from Black-Scholes volatilities implied by appropriate European currency options. Moreover, it was found that BAW volatilities derived from a pair of put and call CME options with the same exercise price are very close to each other. It follows that the inverted BAW procedure is able to get rid of the early exercise premium of American options. Therefore, standard procedures for estimating risk neutral distributions from European options might be easily adapted for CME American futures options. The natural smoothing spline method was then applied to estimate the European smile from the implied volatilities. All in all, the proposed approach is easy to apply and circumvents having to use American option price bounds as well as potential convergence problems stemming from nonlinear parameterisations used elsewhere in the literature. It is worth emphasising that the BAW model was chosen for the sake of computational ease, but inversion of other pricing procedures based on the constant volatility geometric Brownian process are likely to give similar results.

Second, by means of investigating the CME and OTC markets in parallel, we gained insight on the structure and interaction of the two markets. More intensive interaction seems to take place for benchmark exercise prices than for the others, which may be explained by the varying OTC liquidity over the price space. It seems to lead to some price distortions and an interesting 'ghost-like' shape of the RND density implied by CME options.

Third, using these empirical results, we attempt to improve the RND estimation from a low number of OTC option prices. Often only data for a limited number of benchmark OTC exercise prices are readily accessible, and therefore, some extrapolation is necessary. In the literature, two quadratic extrapolation methods have been suggested, but they either break the non-arbitrage constraints imposed on the volatility function or they do not fit the observed CME data well. To overcome this problem, we suggest a parsimonious generalisation of these methods which significantly improves the fit. Indeed, we enhance OTC smile extrapolation by information from CME prices.

A Mathematical Appendix

A.1 Barone-Adesi and Whaley approximation for American futures options and its inversion

Let F denote the futures price and T represent maturity, σ volatility and X the strike price of an American option on futures. Let c represent Black's price of the European call option. Barone-Adesi and Whaley(1987) gave the following approximation for the price of the American futures call option, which we denote by C^{BAW}:

$$C^{BAW} = c^B(F,T;X) + A_2(F^*)\left[\frac{F}{F^*}\right]^{q_2}, \text{ for } F < F^* \tag{6}$$

$$C^{BAW} = F - X, \text{ for } F \geq F^*, \tag{7}$$

$$c^B(F,T;X) = e^{-rT}\left[FN\left(d_1^B(F)\right) - XN\left(d_2^B(F)\right)\right], \tag{8}$$

$$A_2(F^*) = \left[\frac{F^*}{q_2}\right]\left\{1 - e^{-rT}N\left(d_1^B(F^*)\right)\right\}, \tag{9}$$

$$d_1^B(F) = \frac{\ln\left(\frac{F}{X}\right) + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}, \; d_2^B(F) = d_1^B(F) - \sigma\sqrt{T}, \tag{10}$$

$$q_2 = \frac{1}{2}\left[1 + \sqrt{1+4k}\right] \text{ and} \tag{11}$$

$$k = \frac{2r}{\sigma^2\left[1 - e^{-rT}\right]}. \tag{12}$$

The critical value F^*, above which American futures should be exercised immediately, is defined as a solution of

$$F^* - X = c^B(F^*,T;X) + A_2(F^*). \tag{13}$$

Formulae (6) - (13) were derived under the assumptions of the validity of the standard Black-Scholes model specifying that the futures follow a geometric Brownian motion stochastic process. Nevertheless, they also define a procedure that maps given σ, F, X, T and r to a single number C^{BAW}, i.e., the BAW American futures call option price. Let us summarise this procedure by the function

$$C^{BAW} = \Lambda_c(\sigma; F, X, T, r). \tag{14}$$

To define the BAW implied volatility, we construct an inverse function to Λ_c. The function $\Lambda_c^{-1}(p; F, X, T, r)$ is defined as

$$\sigma = \Lambda_c^{-1}\left(\Lambda_c(\sigma; F, X, T, r); F, X, T, r\right) \tag{15}$$

for admissible values of the variables σ, F, X, T, r. Then the BAW implied volatility for an American call quote C is defined as

$$\sigma_c^{BAW} = \Lambda_c^{-1}(C; F, X, T, r). \tag{16}$$

The approximation for the American futures put option, which we denote by P^{BAW}, is similar:

$$P^{BAW} = p^B(F,T;X) + A_1(F^{**})\left[\frac{F}{F^*}\right]^{q_1}, \text{ for } F < F^{**} \tag{17}$$

$$P^{BAW} = X - F, \text{ for } F \geq F^{**}, \tag{18}$$

$$p^B(F,T;X) = e^{-rT}\left[XN\left(-d_2^B(F)\right) - FN\left(-d_1^B(F)\right)\right], \tag{19}$$

$$A_1(F^{**}) = -\left[\frac{F^{**}}{q_1}\right]\left\{1 - e^{-rT}N\left(-d_1^B(F^{**})\right)\right\} \text{ and} \tag{20}$$

$$q_1 = \frac{1}{2}\left[1 - \sqrt{1+4k}\right]. \tag{21}$$

The critical value F^{**}, above which American futures options should be exercised immediately, is defined as a solution of

$$X - F^{**} = p^B(F^{**},T;X) + A_1(F^{**}). \tag{22}$$

Let us summarise formulae (17) - (22) by a function Λ_p:

$$P^{BAW} = \Lambda_p (\sigma; F, X, T, r). \quad (23)$$

Similarly to the case of call options, to get the BAW volatility σ_p^{BAW} implied by an American futures put option P we invert the method and construct the function $\Lambda_p^{-1}(p; F, X, T, r)$ defined as

$$\sigma = \Lambda_p^{-1} (\Lambda_p (\sigma; F, X, T, r); F, X, T, r) \quad (24)$$

for admissible values of σ, F, X, T and r. BAW implied volatility for the put is then defined as

$$\sigma_p^{BAW} = \Lambda_p^{-1} (P; F, X, T, r). \quad (25)$$

The functions Λ_c^{-1} and Λ_p^{-1}seem to be rather intricate, especially because each iterative step involves other iterative procedure evaluations. Nevertheless, this nested numerical method turns out to be very straightforward and fast. Since the shape of the option price behaves similarly as a function of the volatility parameter in both the American and European cases, the implied volatility function may be implemented by the same methods, i.e., the Newton method. As a first step it is suitable to use the European implied volatility as the initial value. Then, in each iterative step, the implementation of the functions Λ_c or Λ_p has to be called twice, because it is convenient to use numerical rather than analytical derivatives. Each call involves an iterative procedure to find the critical early exercise point, but with a suitable choice of initial values in this procedure as suggested by Barone-Adesi and Whaley (1987), it takes 4 iterations at most. Thus, only rarely is the total number of iterative steps in the implementation of Λ_c^{-1} or Λ_p^{-1} larger than 10.

A.2 OTC market conventions

This section details OTC market conventions for quoting currency options (see also Malz (1997)).

The price of a European currency option is determined by market forces and is mainly affected by the spot price of the underlying security S, the strike price X, the option's maturity T, the domestic and foreign interest rates r and r^*, and by the level of future uncertainty. To model the uncertainty, the Black-Scholes model introduces a volatility parameter σ. Although this model does not reflect reality completely, it is used for its very convenient properties as a tool for quoting prices. Thus, market participants in the OTC market quote prices in volatility terms. Deviations from the benchmark model are reflected in the fact that the quoted volatility is not constant across strike prices and maturities. Furthermore, to concentrate on the most option-like parts of the option price and abstract from erratic changes in the currency spot rate, the OTC market developed a convention for measuring the moneyness of options by the option's delta rather than by the variables S and X. And finally, prices of combinations like risk reversals and strangles are quoted, rather than prices of plain vanilla calls or puts.

`Transformation from option-strike space to volatility-delta space`

Let $c \equiv c(X; S, T, r, r^*)$ and $p \equiv p(X; S, T, r, r^*)$ denote prices of European call and put options with strike price X, respectively, and let other parameters in the brackets be defined as usual. To describe the transformation from price-strike space to volatility-delta space, it is convenient to define a European call implied volatility function, $\sigma^{impl}(c, X; S, T, r, r^*)$. The function $\sigma^{impl} \equiv \sigma^{impl}(c, X; ...)$ is implicitly defined as a solution to the Black-Scholes formula:

$$c = Se^{-r^*T} N\left(d_1\left(\sigma^{impl}\right)\right) - Xe^{-rT} N\left(d_2\left(\sigma^{impl}\right)\right), \quad (26)$$

where

$$d_1\left(\sigma^{impl}\right) = \frac{\ln\left(\frac{S}{X}\right) + \left[r - r^* + \frac{1}{2}\sigma^{impl2}\right]T}{\sigma^{impl}\sqrt{T}} \text{ and} \quad (27)$$

$$d_2\left(\sigma^{impl}\right) = d_1\left(\sigma^{impl}\right) - \sigma^{impl}\sqrt{T}. \quad (28)$$

Given S, T, r and r^*, the function σ^{impl} transforms the market price of a call option with strike price X to the corresponding implied volatility. Formally, let us write the mapping from price-strike space (c, X) to volatility-delta space (σ, δ) as

$$\sigma = \sigma^{impl}(c, X; S, T, r, r^*) \tag{29}$$

$$\delta = \Delta_c(X, \sigma; S, T, r, r^*) \equiv e^{-r^*T} N(d_1(X)), \text{ where} \tag{30}$$

$$d_1(X) = \frac{\ln\left(\frac{S}{X}\right) + \left[r - r^* + \frac{1}{2}\sigma^2\right] T}{\sigma\sqrt{T}}. \tag{31}$$

The function $\Delta_c(X, \sigma; S, T, r, r^*)$ defined in equation(30) is the well-known formula for the Black-Scholes delta of a European call.

A similar mapping is defined for put options. Note that the delta of a European put can be written in terms of the delta for a call $\Delta_p(X, \sigma; S, T, r, r^*) = \Delta_c(X, \sigma; S, T, r, r^*) - e^{-r^*T}$. Also, due to put-call parity, a put and a call with the same strike price X imply the same volatility. The mapping from (p, X) to (σ, δ) is then given by

$$\sigma = \sigma^{impl}(p + Se^{-r^*T} - Xe^{-rT}, X; S, T, r, r^*) \tag{32}$$

$$\delta = \Delta_p(X, \sigma; S, T, r, r^*) + e^{-r^*T}. \tag{33}$$

Thus both transformations from (c, X) to (σ, δ) and from (p, X) to (σ, δ) map the interval $(0, \infty) \times (0, \infty)$ to the interval $(0, \infty) \times \left(0, e^{-r^*T}\right)$.

`Transformation from volatility-delta space to option-strike space.`

To describe the inverse transformation from volatility-delta space to call-strike space, it is convenient to define the function $X^{impl} \equiv X^{impl}(\sigma, \delta_c; S, T, r, r^*)$, which generates the appropriate (implied) strike price of a call option from the values of σ and δ_c. It is defined as a solution of

$$\delta_c = \Delta_c\left(X^{impl}, \sigma; S, T, r, r^*\right). \tag{34}$$

Thus, the mapping from volatility-delta space to price-strike space, $(\sigma, \delta_c) \to (c, X)$, is given by

$$c = Se^{-r^*T} N(d_1(\sigma)) - Xe^{-rT} N(d_2(\sigma)) \tag{35}$$

$$X = X^{impl}(\sigma, \delta_c; S, T, r, r^*). \tag{36}$$

While the Newton method is very suitable for the numerical implementation of the function σ^{impl}, one has to be more cautious when numerically calculating the function X^{impl}. Note that the second derivative $\frac{\partial^2}{\partial x^2}\Delta_c(x, \sigma; S, T, r, r^*)$ changes sign once for $x \in (0, \infty)$ with inflexion point $X^* = Se^{\left[r - r^* - \frac{1}{2}\sigma^2\right]T}$. However, it is still possible to use this numerical approach. The function $\Delta_c(X; ...)$ is concave for $X \leq X^*$ and convex for $X \geq X^*$ Using this property, the Newton method might be amended. Then it is likely to be faster than a more general numerical method.

The following market conventions are used. A call option with delta δ_c is referred to as $100\delta_c$, i.e., a 25-delta call is a call option with $\delta_c = 0.25$. Similarly, a put option with delta δ_p is referred to as $-100\delta_p$, i.e., a 25-delta put is a put option with $\delta_p = -0.25$. Moreover, since for short maturity options the term e^{-r^*T}, which facilitates transformation between the delta of a put and the delta of a call, is close to one, a call counterpart to the 25-delta put is often referred to as a 75-delta call instead of a $100\left(e^{-r^*T} - 0.25\right)$delta-call. Another abbreviation is used for an ATMF call. Denote the market quotes of ATMF volatility as $(\sigma_{atmf}, \delta_{atmf})$. Sometimes, people refer to it as a 50-delta call, i.e.$(\sigma_{atmf}, 0.5)$. In fact, it is easy to show that, given a market quote of volatility σ_{atmf}, the delta of a call option with $X = F = Se^{(r-r^*)T}$ is $\delta_{atmf} = \Delta_c(F, \sigma_{atmf}; S, T, r, r^*) = e^{-r^*T} N\left(\frac{1}{2}\sigma_{atmf}\sqrt{T}\right)$. For example $\Delta_c(F, \sigma_{atmf}; S, 0.5, 6\%, 4\%)$ =0.52.

Actually, the market convention of mapping the (σ, δ_c) space to the (c, X) space seems to be somewhat odd. Note that for a given call delta (e.g. $\delta_c = 0.25$), a change in quoted volatility also represents a change in the implied strike, as the function $X^{impl}(\sigma, \delta_c; S, T, r, r^*)$ indicates. Therefore, it is theoretically possible to have two (or many) volatility-delta pairs that are transformed to one strike only. However, it apparently does not pose a problem for the market. The reason is that prices are not quoted close to each other and also that 'perverse' quotes would break no-arbitrage conditions for the volatility function.

`Usually quoted contracts`

In the OTC market, currency options are usually quoted in terms of ATMF implied volatilities, risk reversals and strangles. Both risk reversals and strangles are combinations of call and put options, which are equally OTM. Their moneyness is measured by their delta. While a buyer of a risk reversal acquires a long position in an OTM call option and a short position in an OTM put option, a buyer of a strangle buys both of them, i.e., gets long positions in OTM calls and puts. The price of a risk reversal is quoted in volatility terms as a difference between implied volatilities of an appropriate call and put. Similarly, strangle prices are quoted as an average volatility premium paid for the strangle components above the ATMF implied volatility. Let σ_{atmf}, $rr(.)$ and $str(.)$ denote market quotes of ATMF implied volatility, risk reversal and strangle. The parameter in the brackets refers to the delta of the call of the constituents. Malz (1997) shows how to back out pairs (σ, δ_c) from these market quotes. For convenience, denote these pairs as a function $\sigma(\delta_c)$. Then the following relationships hold:

$$rr(\delta_c) = \sigma(\delta_c) - \sigma\left(e^{-r^*T} - \delta_c\right) \tag{37}$$

$$str(\delta_c) = \frac{1}{2}\left[\sigma(\delta_c) + \sigma\left(e^{-r^*T} - \delta_c\right)\right] - \sigma_{atmf}. \tag{38}$$

So, for example, a 25-delta risk reversal $rr(0.25)$ is the difference between the implied volatility of a call option with a delta of 0.25 and the implied volatility of a put option with a delta of -0.25. A call option that has the same volatility as the 25-delta put is the one with a delta of $(e^{-r^*T} - 0.25)$. Note that sometimes the strangle is quoted without subtraction of σ_{atmf} in equation (38).

Equations (37) and (38) are easy to invert, and one obtains

$$\sigma(\delta_c) = rr(\delta_c) + \frac{1}{2}rr(\delta_c) + \sigma_{atmf} \tag{39}$$

$$\sigma\left(e^{-r^*T} - \delta_c\right) = str(\delta_c) + \frac{1}{2}rr(\delta_c) + \sigma_{atmf}. \tag{40}$$

Thus, equations (39) and (40) show how to get from a market quote of σ_{atmf}, $rr(\delta_c)$ and $str(\delta_c)$ to $\sigma(\delta_c)$ and $\sigma\left(e^{-r^*T} - \delta_c\right)$.

A.3 Homogeneity of European call options

Proposition 1 *Let $c_t(F, X, T, r)$ denote the price of a European call option at time t with underlying security F, strike price X and maturity T. Let r represent the risk free rate. Let $k > 0$. Then $c_t(kF, kX, T, r) = kc_t(F, X, T, r)$.*

Proof. *Further, let f denote the RND of F at maturity of the option, and similarly, let g denote the RND of kF at maturity of the option. Note that $g(kF) = \frac{1}{k}f(F)$. Then*

$$\begin{aligned} c_t(kF, kX) = e^{-rT}\int_0^\infty \max(kF_T - kX, 0)\, g(kF_T)d(kF_T) = \\ = e^{-rT}k\int_0^\infty \max(F_T - X, 0)\,\frac{1}{k}f(F_T)kdF_T = \\ = e^{-rT}k\int_0^\infty \max(F_T - X, 0)\, f(F_T)dF_T = \\ = kc_t(F, X) \end{aligned}$$

■

B Footnotes

1 It is a well established result that the price of a derivative asset which depends only on prices of traded securities may be expressed as its expected payoff discounted by the risk-free rate, where the expectation is taken over by the risk neutral distribution (1996). However, in general, true and risk neutral distributions may be quite distinct; the only necessary restriction is that the distributions share a common support. Grundy (1991) examines the relationship between option prices and the true distribution of the underlying asset and finds that imposing simple restrictions on the true distribution leads to useful bounds of its non-central moments. Rubinstein (1994) shows that for standard utility functions the true distribution tends to be slightly shifted with respect to the risk neutral one, but its shape remains very similar.

2 The BAW method is a good approximative analytical solution to the pricing problem for American futures options if the underlying futures follow a geometric Brownian motion with a constant volatility. We use this approach because of its computational speed and because it is widely used in practice. Indeed any other numerical pricing method could have been used.

3 We used the dollar-yen exchange rate because it has a long enough history and it is one of the most liquid CME currency pairs with the greatest number of active strike prices.

4 There are major structural differences between the two markets, which complicate the analysis. While OTC options are typically European ones having the spot exchange rate as the underlying variable, the CME contracts are American options on currency futures. Moreover, unlike OTC options which are in essence of fixed maturity, exchange traded options are of a fixed maturity date. Further, for a given CME currency-option contract, the underlying futures may have a significantly different maturity date. The markets differ also in liquidity. It is estimated that the OTC market is much deeper than the on-exchange one: Exchange-traded contracts are estimated to represent less than 3% of daily turnover in the foreign exchange market (Bank of International Settlements (2000)). On the other hand, the exchange market is more transparent. While exchanges provide information about prices, volumes and open interest in traded contracts, the OTC market is more opaque as trades are concluded on a bilateral basis.

5 If demanded, most of the banks that trade options would provide a price for an option with any strike. Such a price, however, would be most likely derived from the current market conditions characterised by the prices of a handful of standardised contracts.

6 The Black-Scholes model stipulates no arbitrage in perfect and frictionless markets, a constant short-term riskless rate and furthermore that the price of the underlying security follows geometric Brownian motion with drift and a constant volatility parameter. However, non-constant Black-Scholes volatility implied by option prices is observed in most of the markets.

7 The relation between strike price and implied volatility is called the volatility smile, and the relation between maturity and implied volatility is sometimes denoted as the term structure of volatility.

8 We express the price of an American option as a sum of the price of a European option and some early exercise premium. This method uses the fact that the premium has to satisfy the same pricing equation as option prices with appropriate boundary conditions. Under the assumptions of the benchmark Black-Scholes model, and when one term in the equation is neglected, it is possible to derive an analytical solution for the early exercise premium. The neglected term tends to be very small, especially for options with very short and very long maturities.

9 In fact, for each option the American option pricing method has to be invoked several times, because similar to the Black-Scholes model, searching for unknown volatility from the option price is in itself an iterative procedure. However, since the option price and volatility are in a well-behaved and close-to-linear relation, convergence is very fast. In the Appendix A.1, we summarise the BAW method and our approach to calculating its inversion.

10 The BAW method might be theoretically replaced by any method for pricing American futures options under the assumption of the geometric Brownian motion with constant volatility. Using the inverted BAW method for calculating the implied volatility is only a matter of computational ease and speed. Suitable methods might be selected using the results of Broadie and Detemple (1996).

11 In Chicago, there are two markets where American futures options are traded, RTH (Regular Trading Hours) and the GLOBEX2 electronic trading system, which ensures that in selected contracts trading continues virtually 24 hours a day. While during regular trading hours prices are quoted in dollars, on the GLOBEX2, volatility quotes are possible. The BAW method is used for transforming volatility quotes into dollar prices.

12 See CME Rulebook. It is available at www.cme.com.

13 Regarding interest rates, we used appropriate LIBOR deposit rates for the dollar and yen. As is typical in the literature, we assume that the short-term interest rates are non-stochastic and therefore, using the results of Cox, Ingersoll and Ross (1981), we treat futures contracts as if they were forward contracts. We believe that especially for currency futures options this assumption is quite innocuous. There is evidence in the literature that the difference between futures and forwards is rather small. Whaley (1986) for example uses this assumption even for options on long-term interest rate products.

14 If $\{y_i\}$ is a set of observations (implied volatilities), a smoothing cubic spline s minimizes

$$\lambda \sum_i w_i \left(y_i - s_i\right)^2 + (1-\lambda) \int \left(s''\right)^2 dx$$

for a smoothing parameter λ.

15 We want to attribute more weight to highly traded contracts, but in a less than proportional manner. We therefore choose a natural logarithm because this is a function that grows slower than any root function.

16 The concentration of the full sample around the mean and its heavy tails are illustrated by the high level of kurtosis shown in Table I. It also indicates that the differences are not normally distributed and that the actual confidence intervals for the means are arguably even tighter.

17 The arbitrage between markets is possible if the midprices differ for more than the sum of bid/ask spreads of the two markets. The bid/ask spread for the benchmark OTC currency option contracts for the developed market is about 0.2%.

18 It was also the last opportunity to compare exchange traded options with OTC options before the 1998 dollar/yen currency turmoil.

19 The face value of one underlying futures contract is 12.5 million yen, i.e., the face value of all options traded on that day was more than 350 million dollars.

20 It is indicated by the higher kurtosis of the CME distribution as shown in Figure 3. Other statistics, i.e., mean, modus, volatility, skewness and the Pearson statistic, are of similar value for both distributions. The means of the distributions do not differ significantly from the theoretical value of the forward price. The OTC quoted ATMF volatility was 15.65%

that day. The volatilities calculated from the CME distribution (defined as $\frac{\sqrt{variance}}{forward*\sqrt{T}}$) were fairly near this value, but the volatility of the OTC distribution was somewhat lower. The distributions on this day were both very symmetric. The asymmetry of the distribution, as measured by skewness or the Pearson statistic, is close to zero.

References

Ahn, D. H., Boudukh, J., Richardson, M., & Whitelaw, R.F., (1999). Optimal Risk Management Using Options. Journal of Finance 54 (1), pp. 359-375.

Artzner, P., Delbaen, F., Eber, J., & Heath, D., (1999). Coherent Measure of Risk. Mathematical Finance 9, pp. 203-228.

Bahra, B., (1997). Implied Risk Neutral Probability Density Functions from Option Prices: Theory and Application. Working Paper, Bank of England, ISSN 1368-5562.

Barone-Adesi, G., & Whaley, R. E., (1987). Efficient analytic approximation of American option values. Journal of Finance 32, pp. 301-320.

BIS Quarterly Review and BIS Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity, (2000). Bank of International Settlements, Basle, 2000.

Black, F., (1976). The Pricing of Commodity Contracts. Journal of Financial Economics 3, pp. 167-79.

Black, F., & Scholes, M., (1973). The Pricing of Options and Corporate Liabilities. Journal of Political Economy 81, pp. 637-59.

Bliss, R. R., & Panigirtzoglou, N., (2000). Testing the Stability of Implied Probability Density Functions. Bank of England Working Paper 114.

Buchen, P., & Kelly, M., (1996). The Maximum Entropy Distribution of an Asset Inferred from Option Prices. Journal of Finance and Quantitative Analysis 31(1), pp. 143-59.

Broadie, M., & Detemple, J., (1996). American Option Valuation: New Bounds, Approximations, and a Comparison of Existing Methods. The Review of Financial Studies 9, No. 4, pp 1211-1250.

Breeden, D. T., & Litzenberger, R. H., (1978). Price of State Contingent Claims Implicit in Option Prices. Journal of Business 51 (1978), 621-651.

Campa, J.M.,Chang, P.H.K., & Rieder, R.L., (1997). Implied Exchange Rate Distributions: Evidence from OTC Option Markets. NBER Working Paper 6179.

Bouc, P., & Cincibuch, M., (2001). Interpretation of Czech FX Options. Czech National Bank, Prague, WP 36.

Chang, P.H.K., & Melick, W.R., (1999). Background note. Proceedings of the BIS workshop on implied PDFs, Bank of International Settlements, Basle.

Christensen, B.J., & Prabhala, N.R., (1998). The Relation between Implied and Realized Volatility. Journal of Financial Economics 50, pp. 125-150.

Clews, R., Panigirtzoglou, N., & Proudman, J., (2000). Recent Developments in Extracting Information from Options Markets. Bank of England Quarterly Bulletin 40(1) February 2000, 50-57.

Cooper, N., (1999). Testing Techniques for Estimating Implied RNDs from the Prices of European-Style Options. Proceedings of the BIS workshop on implied PDFs, Bank of International Settlements, Basle.

Cox, J.C., Ingersoll, J.E., & Ross, S.A., (1981). The relation between forward prices and futures prices. Journal of Financial Economics 3, 145-166.

Cox, J., & Ross, S., (1976). The Valuation for Alternative Stochastic Processes. Journal of Financial Economics 3, pp. 145-66.

Daníelsson, J., (2000). The Emperor Has No Clothes: Limits to Risk Modelling. Financial Markets Group of London School of Economics, Mimeo.

Flamouris, D., & Giamouridis, D., (2002). Estimating Implied PDFs from American Options on Futures: A New Semiparametric Approach. The Journal of Futures Markets 22, No.1, pp 1-30.

Gemmill, G., & Safeklos, A., (1999). How Useful Are Implied Distributions? Evidence from Stock-Index Options. Proceedings of the BIS workshop on implied PDFs, Bank of International Settlements, Basle.

Grundy, B. D., (1991). Option Prices and the Underlying Asset's Return Distribution. The Journal of Finance 46 (3), pp. 1045-1069.

Jackwerth, C. J., & Rubinstein, M., (1996). Recovering Probability Distributions from Option Prices. Journal of Finance 51.

Jorion, P., (1995).Predicting Volatility in the Foreign Exchange Market. Journal of Finance 50 (2), pp. 507-528.

Hull, J., (1996). Options, Futures, and Other Derivative Securities. Prentice-Hall International, Inc. ISBN 0-13-015822-4.

MacMillan, L.W., (1986). Analytic Approximation for the American Put Option. Advances in Futures and Options Research 1, 1986, pp. 119-39.

Malz, A.M., (1996). Using Option Prices to Estimate Realignment Probabilities in the EMS: The Case of Sterling-Mark. Journal of International Money and Finance 15, pp. 717-48.

Malz, A.M., (1997). Estimating the Probability Distribution of the Future Exchange Rate from Option Prices. The Journal of Derivatives, Winter 1997, pp. 18-36.

Melick, W. R., & Thomas, Ch. P., (1997). Recovering an Asset's Implied PDF from Option Prices: An Application to Crude Oil during the Gulf Crisis. Journal of Financial and Quantitative Analysis (32) 1, March 1997, pp. 91-115.

Melick, W. R., (1999). Results of the Estimation of Implied PDFs from a Common Dataset. Proceedings of the BIS workshop on implied PDFs, Bank of International Settlements, Basle.

Moessner, R., (2000). Over-the-counter Interest Rate Options. Bank of England, Mimeo.

Ross, S., (1976). Options and Efficiency. Quarterly Journal of Economics 90, pp. 75-89.

Rubinstein, M., (1994). Implied binomial trees. Journal of Finance, 49

Shimko, D., (1993). Bounds of Probability. Risk 6 (4), pp. 33-37.

Soederlind, P., & Svensson, L., (1997). New Techniques to Extract Market Expectations from Financial Instruments. Journal of Monetary Economics 40. pp. 383-429.

Whaley, R. E., (1986).Valuation of American Futures Options: Theory and Empirical Tests. Journal of Finance. 31.

List of Tables

List of Figures

	Mean	Its 99% c.i.	Std. deviation	Kurtosis
Full sample, 882 obs.				
25-delta	0.08%	(0.01%,0.14%)	0.72%	93.50
50-delta	0.15%	(0.10%,0.21%)	0.60%	107.18
75-delta	0.25%	(0.19%,0.31%)	0.70%	102.96
Ex IV 1998, 852 obs.				
25-delta	0.06%	(0.01%,0.11%)	0.53%	7.46
50-delta	0.13%	(0.09%,0.17%)	0.44%	8.27
75-delta	0.21%	(0.17%,0.26%)	0.51%	5.84
Maximum maturity mismatch: 6 days				

Table I: Difference between BAW-implied and OTC volatilities

Distance	less than 0.5%	0.5% to 1%	1% to 5%	more than 5%
Number of observations	25458	296	112	9

Table II: Distances between BAW volatilities implied by calls and puts with common strikes

	Mean	Its 99% c.i.	Std. deviation	Kurtosis
Plain distances				
Call-put	-1.03e-003%	(-3.72e-003%,1.66e-003%)	1.68e-001%	154.81
OTM-ITM	2.89e-002%	(2.62e-002%,3.15e-002%)	1.65e-001%	162.19
Distances weighted by the trading volume				
Call-put	2.47e-007%	(1.11e-007%,3.83e-007%)	8.52e-006%	1739.33
OTM-ITM	4.40e-007%	(3.04e-007%,5.76e-007%)	8.51e-006%	1745.79

Table III: BAW volatility: Put-call parity property

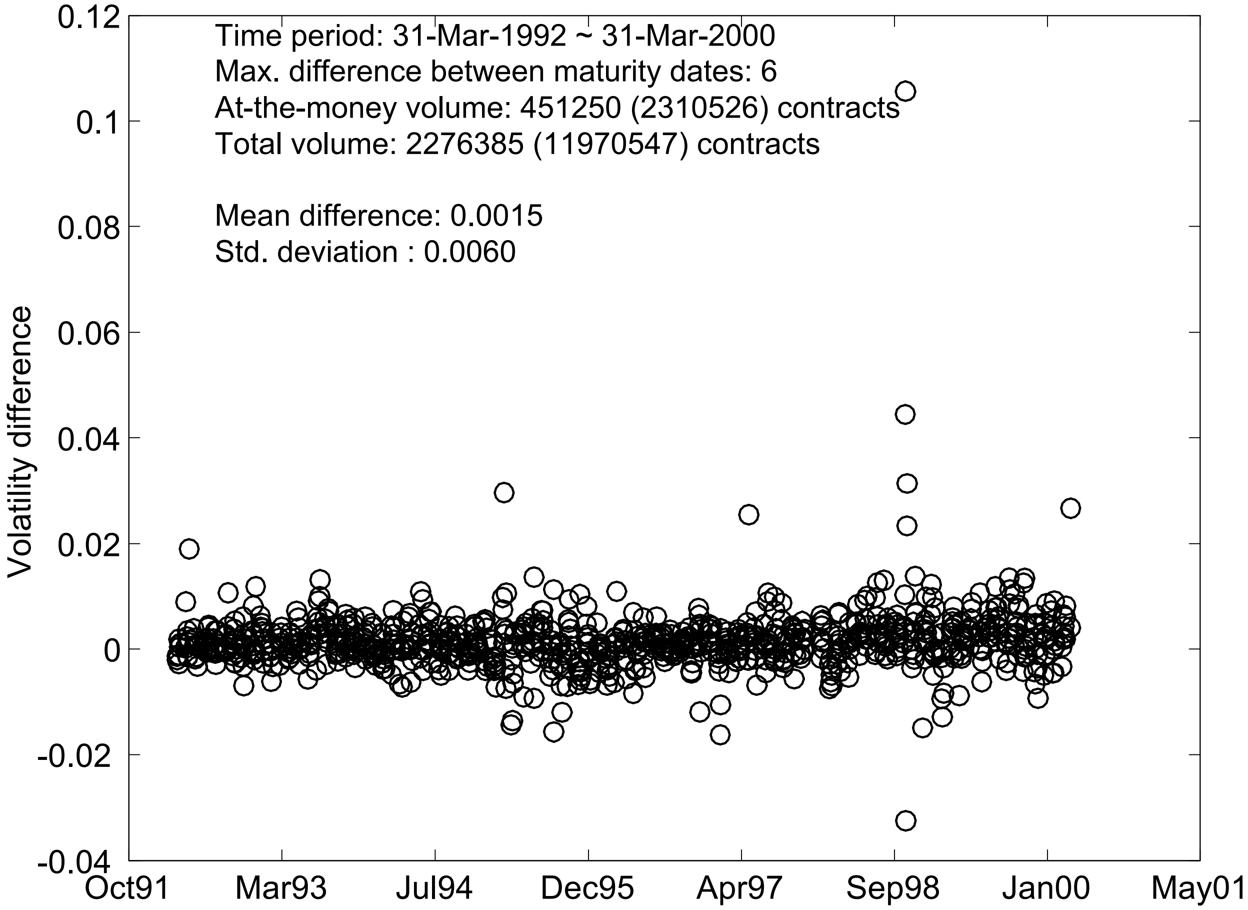

Figure 1: Difference between OTC and CME implied volatility (ATMF, BAW adjusted)

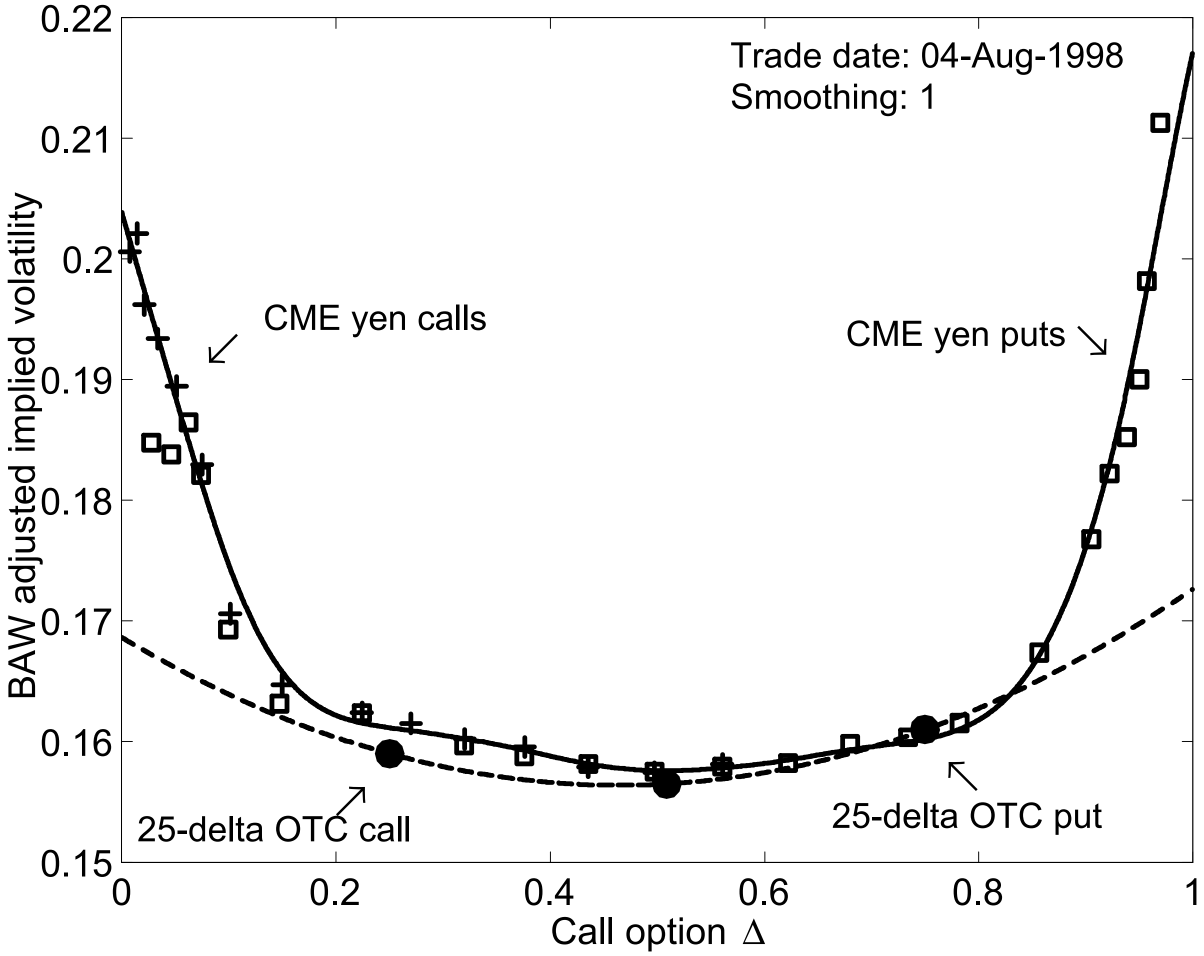

Figure 2: Δ-space smiles: Natural spline for CME and Malz quadratic extrapolation for OTC

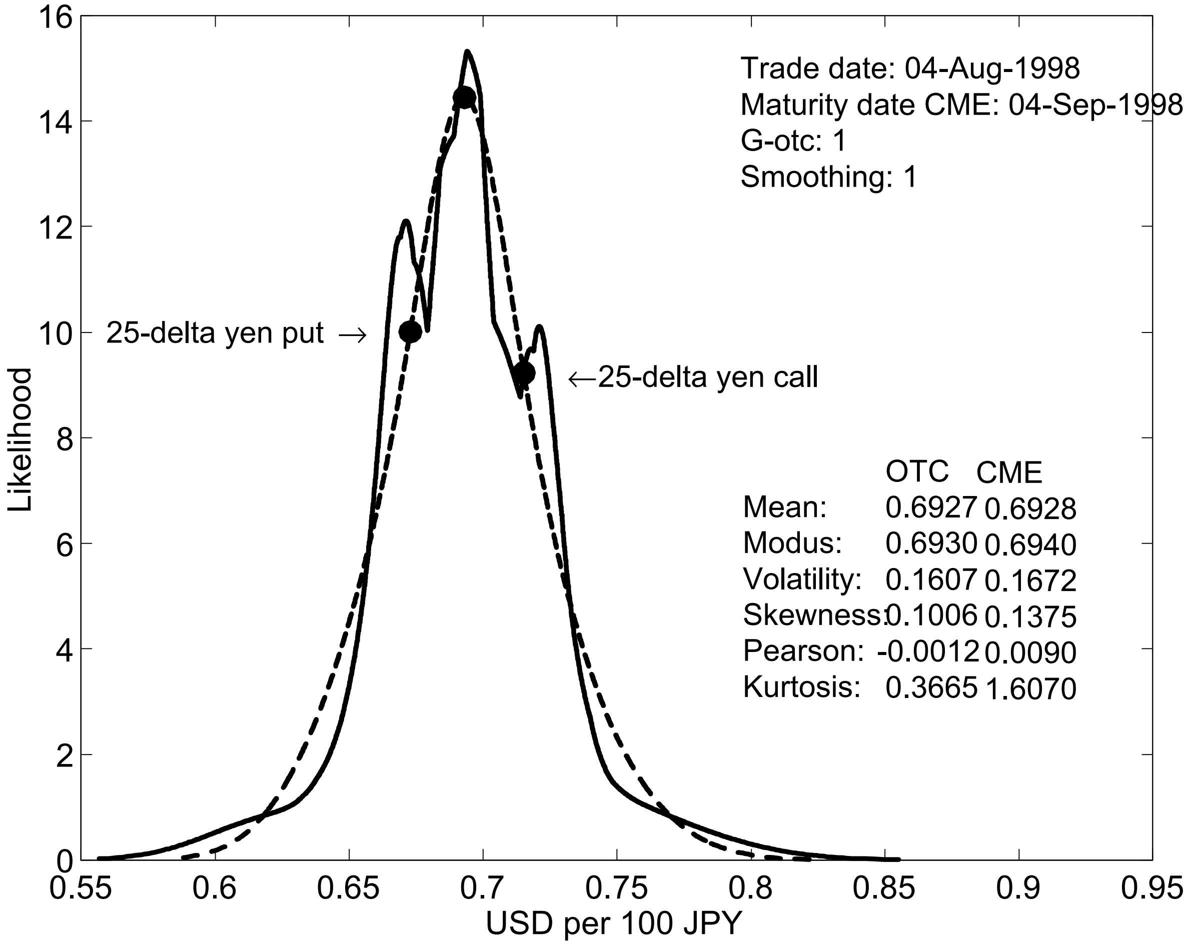

Figure 3: CME and OTC implied distributions

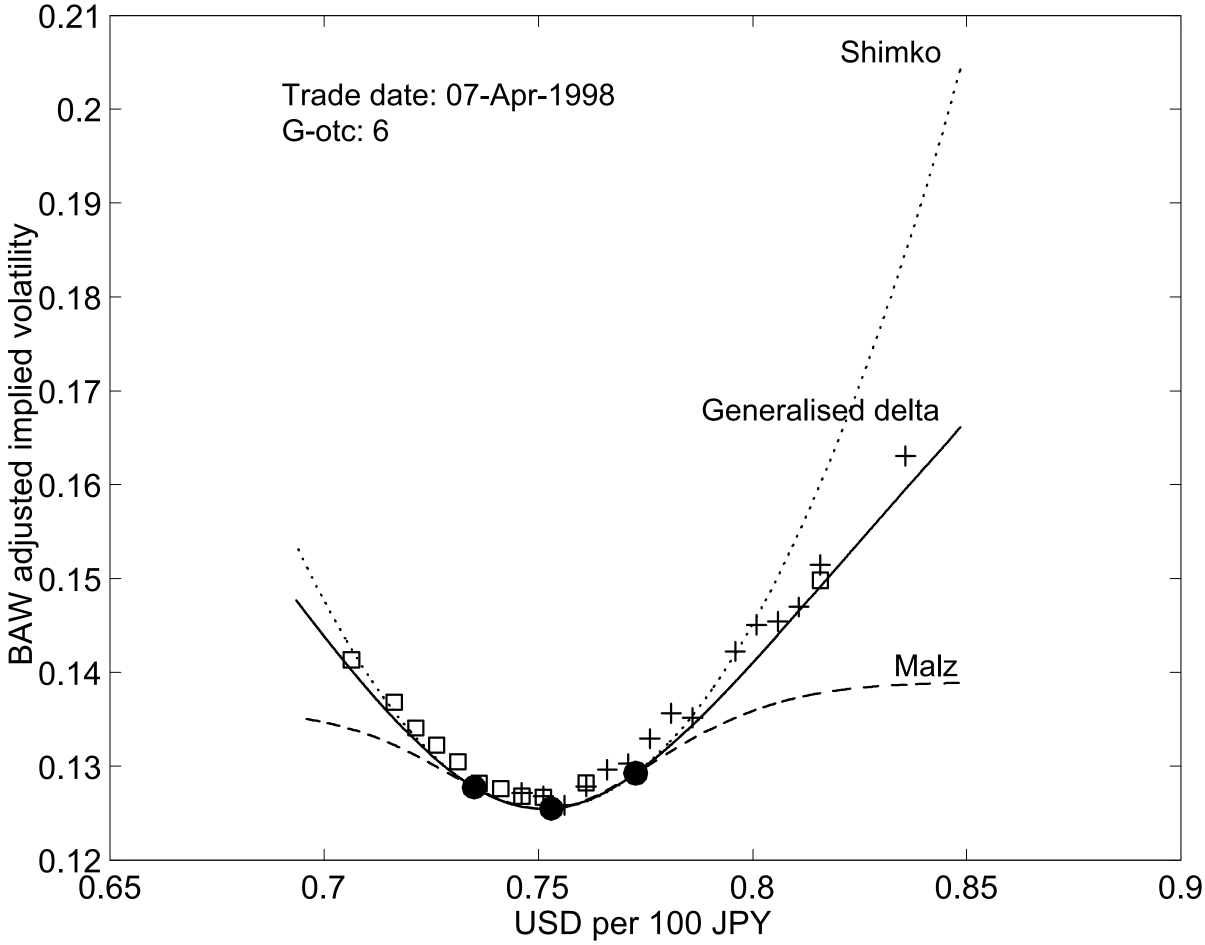

Figure 4: Comparison of methods for OTC smiles

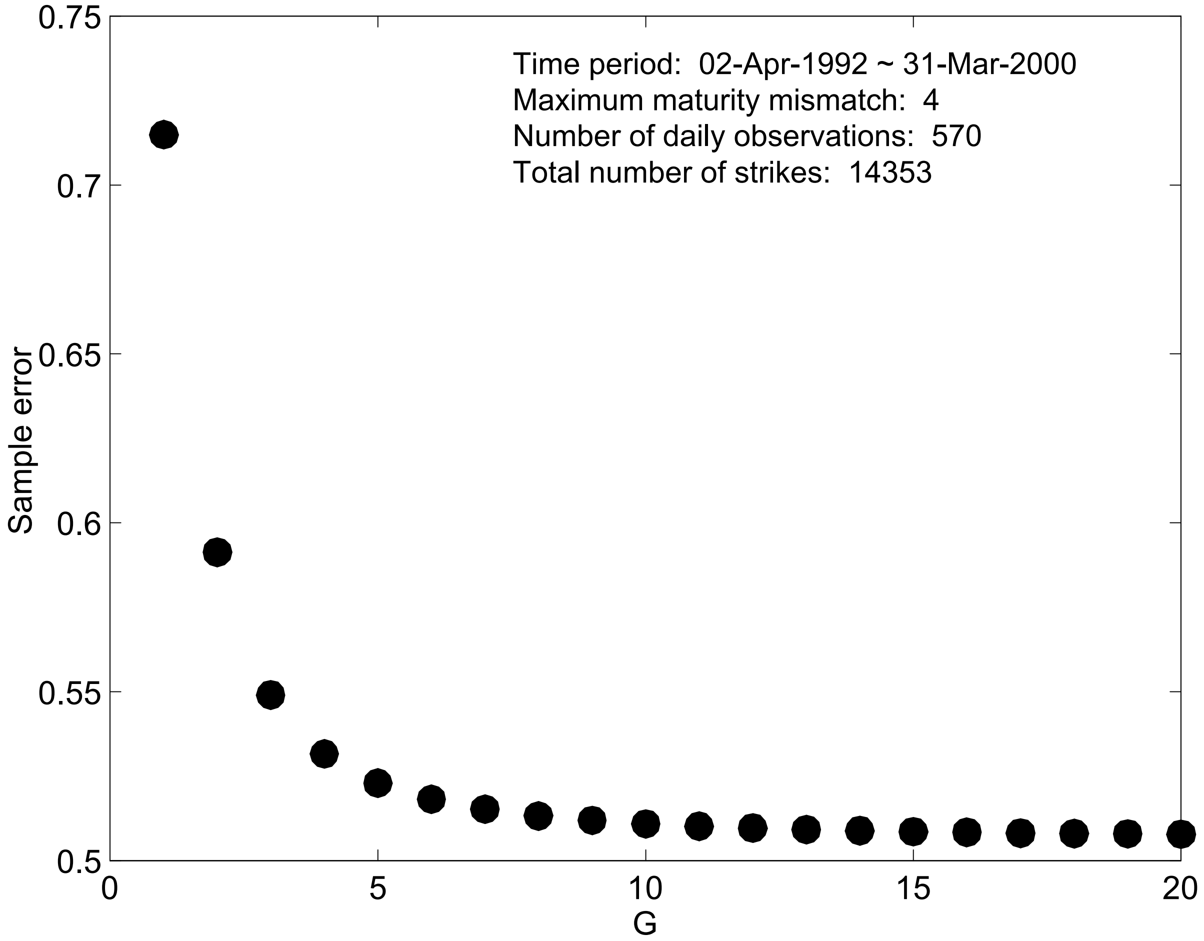

Figure 5: Fit of generalised delta smiles