The information content of currency options: evidence from emerging economies

by

Marian Micu*

ROSES, Université de Paris I Panthéon Sorbonne and Bank for International Settlements

September 2004

*The views expressed in this paper are those of the author and do not necessarily reflect those of the BIS.

Abstract

Currency option markets in emerging economies developed rapidly during the last couple of years. These developments were associated with an increased uncertainty about the exchange rate fluctuations in most of these countries. An important issue for central bankers and market analysts is to understand the information content of these derivative products.

Do option prices really provide good predictions about future exchange rate distributions? Are these options consistently priced along the term structure? Is there any evidence for systematic risk premia in option prices? All these questions have been more or less addressed in the case of foreign exchange options written on currencies of developed countries. However, so far, there is hardly any study on the information content of options written on emerging market currencies.

The objective of this doctoral thesis is twofold. First, it extends and provides a systematic analysis of the methods used to extract market expectations from currency options. Second, it applies this methodology to currency options written on emerging market currencies and offers some insights into the quality of the information which can be extracted from such options.

The information extracted from currency option prices can be of great help for central banks in setting up their monetary and exchange rate policies. For instance, in a country with a currency board arrangement, option markets could provide ex ante information about the confidence of market participants in the peg.[1] Option markets also allow to test the "signaling channel" of foreign exchange interventions, by looking at the impact of interventions on exchange rate expectations.[2] In some cases, the information extracted from option prices can provide a measure of contagion or be used in the calculation of foreign exchange risk exposures.[3] Some studies also derived measures of risk premia from option prices.[4]

This thesis is structured in four chapters. The first two chapters present the theoretical background of the methods used for pricing currency options and extracting risk-neutral probability density functions from their prices. In addition, we propose a new technique for recovering risk-neutral density functions based on the piecewise cubic Hermite interpolation of the implied volatility. The last two chapters provide two empirical applications aimed at testing the information content of foreign exchange options. First, we show how the information extracted from risk-neutral probability density functions can be used to provide an ex ante measure of the forward forecasting bias. Second, we test the expectations hypothesis on the term structure of implied volatilities and check whether implied volatilities can be successfully employed as predictors of future realised volatilities.

[1] See, for example the papers by Malz (1996), Campa and Chang (1996), etc.

[2] See, for instance, Galati, Melick and Micu (2004).

[3] Campa and Chang (2001) showed how to use the information extracted from multivariate exchange rate distributions to calculate implied correlations of exchange rate returns.

[4] See, inter alia, Ait-Sahalia and Lo (2002), or Bliss and Panigirtzoglou (2004).

To our knowledge, so far, there is no study which has systematically tested the information content of options written on emerging market currencies. One reason for this may be that, until very recently, data on these derivative instruments was not not readily available. Our data consists of market quotes of over-the-counter options on twelve emerging market currencies against the US dollar. These currencies are: the Brazilian real, Chilean peso, Czech koruna, Indonesian rupiah, Malaysian ringgit, Mexican peso, Polish zloty, South African rand and South Korean won. In particular, we use the information from straddles, strangles and risk-reversals corresponding to 50%, 25% and 10% levels of delta, as well as Eurocurrency interest rates recorded by currency option traders at a world major foreign exchange dealer bank. Observations are for options with maturities of one, three and six months, with a daily frequency from 10 November 1997 to 10 November 2002.

In the first chapter we present the theory on options valuation. The baseline model is the Garman-Kohlhagen model, which is the Black and Scholes model adapted for the pricing of currency options. In this chapter, we explain one of the basic concepts in option pricing, namely the *dynamic hedging (or delta hedging).* Dynamic hedging is an important assumption of the Black and Scholes model and it is the foundation of the risk-neutral approach to option pricing. The basic idea underlying the dynamic hedging assumption is that an option writer can continually change the exposure in the underlying asset, so that the change in the price of the option is fully compensated by the change in the price of the underlying asset. To understand this concept, suppose that a bank sold a call option. If the price of the underlying asset increases, it is more likely that the option will end up in-the-money. In this case, the bank would have to deliver the underlying asset at the expiration of the option contract to the buyer of the option for less than it actually costs in the spot market. The bank could decide to buy the asset at expiration in the spot market and deliver it to the option holder at a loss. Alternatively, the bank could have bought the asset when it entered into the option contract. However, if instead of going up, the price of the underlying asset decreases, the option will be out-of-the money and the bank may end up with an asset which is worth less than it was at the time of the purchase. In order to avoid either of these losses, the bank can cover its short position by replicating the payoff of the option it has sold. In the first chapter we show how this is possible by using option's sensitivity parameters.

In the foreign exchange over-the-counter market dealers quote option prices in terms of their implied volatilities rather than the actual price. Moreover, options are frequently sold in combinations. The most common combination traded in the interbank market is the *straddle*. The straddle price is often referred to as the at-the-money volatility. There exist also option combinations with which traders can take directional positions. The most representative are the strangles and the risk reversals, both consisting of out-of-the money put and call options. Chapter I provides a thorough review of these strategies and shows how they relate to the higher moments of risk-neutral exchange rate distributions.

The standard option pricing formulae are based on strong assumptions about the stochastic process underlying the exchange rate returns. The second part of the first chapter presents alternative stochastic processes of exchange rate distributions. One of these processes is the pure-jump (Poisson) process introduced by Cox and Ross (1976a). In contrast to the diffusion process, assumed in the Black and Scholes model, the jump process assumes that exchange rate returns follow deterministic movements upon which are superimposed discrete jumps.

A second class of models reviewed in this chapter is represented by the jump-diffusion models proposed by Merton (1976). These stochastic processes are formed by combining a continuous diffusion process and a discrete jump process. The prototype is obtained when the diffusion process is a Brownian motion and the jump process is an independent compound Poisson process with normally distributed jump amplitudes.

A third class of models presented are the stochastic volatility models. Various stochastic volatility models have been proposed by Hull and White (1987a), Johnson and Shanno (1987), Scott (1987) and Wiggins (1987). Estimation of stochastic volatility models has proved difficult because of the two-dimensional specification of the defining stochastic differential equations. Heston (1993) provided a closed-form solution for the price of a European call option when the spot asset is correlated with volatility and also showed how this model could be adapted to incorporate stochastic interest rates.

A natural extension of the Black and Scholes model is to incorporate the effect of stochastic interest rates. Merton (1973) was the first to propose such a model. However, in most cases, option prices obtained through the stochastic interest rate model do not differ substantially from those obtained with the Black and Scholes model.

Another class of option pricing models reviewed in this chapter is represented by the stable Paretian distribution models. These models were introduced by Mandelbrot (1963). The Paretian stable distribution is bell-shaped with fatter tails than the normal distribution. Negative prices can be ruled out by assuming that it is the change in the logarithm of the price, rather than the price itself, that follows a stable Paretian distribution. In practice, however, stable Paretian distributions are hardly used to price foreign exchange options.

A well-spread class of option pricing models is the constant elasticity of variance models. This type of models was proposed by Cox and Ross (1976a). The appeal of these models is that they are consistent with Black's (1976) observation that changes in volatility are negatively correlated with the price change in the underlying asset. The constant elasticity of variance models predicts that the Black and Scholes volatility should change deterministically over times as a function of the underlying asset price. However, given that asset prices have unit roots and typically non-zero drift, the constant elasticity of variance models imply that variance either approaches infinity or zero in the long run.

A final class of models presented in Chapter I is the implied binomial tree models. These models have been introduced by Rubinstein (1994) and Derman and Kani (1994) and are a generalisation of the Black and Scholes model in discrete time. Implied binomial tree models start with a given exchange rate distribution for a particular future date and then build a binomial tree whose final terminal distribution is equal to the given distribution. A prerequisite in the setting up of an implied binomial tree is the determination of the risk-neutral distribution at the expiry of the option contract. This is easily carried out by using the risk-neutral valuation approach. Under the risk-neutral valuation, the value of an option is equal to its expected payoff at maturity, discounted at the risk-free interest rate. In a market which rules out the presence of arbitrage opportunities, the expected value of the future spot rate should equal the forward rate.

All of these alternative option pricing models present advantages and drawbacks. Despite some of these models giving sometimes better approximations of currency option prices, the existence of a time-varying risk premium embedded in the distribution of actual exchange rate returns makes it difficult to compare different models.

In Chapter II, we present five common methods used to extract risk-neutral distributions from option prices and compare these methods using our currency options data. Risk-neutral probability density functions (PDFs) provide the probabilities attached by a risk-neutral agent to particular outcomes for future values of exchange rates. In this chapter, we present the methods introduced by Shimko (1993), Madan and Milne (1994), Malz (1996), Melick and Thomas (1997) and Bliss and Panigirtzoglou (2001). In addition, we propose a new method for recovering risk-neutral density functions based on the piecewise cubic Hermite interpolation of the implied volatility function. The comparisons of these techniques is made on the basis of summary statistics derived from the estimated PDFs. These summary statistics refer to measures of location, dispersion, asymmetry, fat-tailness and various quartiles.

Shimko's (1993) method consists of fitting the volatility smile with a parabola to obtain a continuous function of call prices within the range covered by the data. The probability density function of future exchange rate distributions is obtained by taking the second derivative of the option price with respect to the strike price (as shown in Breeden and Litzenberger (1978)). The main advantage of this method is given by the stability of the optimisation algorithm, which has an analytic solution. The major drawback is that irregular volatility functions are (sometimes) poorly approximated with a parabola.

Madan and Milne's (1994) approach consists of modelling the prices of contingent claims as elements of a separable Hilbert space that has a countable orthogonal basis.

A Hilbert space basis is in general difficult to construct because it requires the knowledge about the stochastic process of the underlying asset returns. However, Madan and Milne showed that, under fairly general conditions, one can specialise the Hilbert space basis to the family of Hermite polynomials. Under this assumption, one can infer the underlying risk-neutral density from traded security prices. This model has been applied to extract risk-neutral probability distributions from options written on stock index futures (Madan and Milne (1994), Coutant (1999)), interest rate futures (Abken, Madan and Ramamurtie (1996), McManus (1999), Coutant et al. (2001)) and exchange rates (Jondeau and Rockinger (2000)). A drawback of this method is that it occasionally provides unreliable estimates for the tails of the risk-neutral PDFs.

Malz (1996) questioned one of the basic assumptions of the Garman Kohlhagen model, namely the continuity and the normal distribution of the stochastic process that characterises the exchange rate returns. Alternatively, he proposed a jump-diffusion model to estimate the realignment probabilities for the sterling-mark exchange rate during the ERM crisis in 1992. In his model, option pricing formulae consist of a weighed sum of the lognormal solutions, where the weights are given by the probability of no jumps occurring and one jump occurring over the lifetime of the option. The advantage of this method resides in its good ability to characterise the evolution of exchange rate prices. Its main disadvantage is that the results may (sometimes) depend on the initial values used in the optimisation algorithm.

A well-known method used to extract risk-neutral distributions from option prices has been proposed by Melick and Thomas (1997). This method is based on a finite mixture of lognormal distributions. Its main advantage is that it imposes little structure on the process by which exchange rates evolve and permits the estimation of relatively flexible functional forms for their distributions. A reasonably flexible functional form, such as the mixture of a finite number of distributions, can easily accommodate a wide variety of shapes for the terminal distribution. As Melick and Thomas (1997) pointed out, starting with an assumption about the terminal risk-neutral distribution function, rather than stochastic process by which the underlying price evolves, has the advantage of being a more general approach. This is because a given stochastic process implies a unique terminal distribution, but the converse is not true, that is, any given risk-neutral distribution function may be consistent with many different stochastic processes. However, placing structure on the terminal distribution rather than on the stochastic process is not without costs. Thus, the recovered distribution does not give any guidance about the evolution of the asset price prior to expiration. This means that the resultant distribution cannot be directly used for constructing dynamic hedges or replication strategies for the option.

A more recent method has been proposed by Bliss and Panigirtzoglou (2002). In their study, Bliss and Panigirtzoglou (2002) used the Newton-Raphson algorithm to recover the implied volatilities from observed option prices. The implied volatility function is then smoothed with a natural spline. The spline function can also be weighted. A natural weight would be the trading volume, as the information content of option prices is directly linked to the liquidity of these instruments. However, detailed data on options trading volumes is quite scarce and unreliable for index options, and nonexistent for currency options. As an alternative the *vega* of the option could be used. The risk-neutral PDF is then easily obtained by taking the second derivative of fitted option prices with respect to the strike price.

Our proposed method for recovering risk-neutral probability density functions from option prices is based on a relatively recent interpolation technique, namely the piecewise cubic Hermite interpolation. Piecewise polynomial interpolation provides an alternative to the practical and theoretical difficulties associated with high-degree polynomial interpolation. In Hermite interpolation the derivatives as well as the values of the interpolating function are specified at the data points. Specifying derivative values simply adds more equations to the linear system that determines the parameters of the interpolating function. In order to have a well-defined solution, the number of equations and the number of parameters to be determined must be equal. To provide adequate flexibility while maintaining simplicity and computational efficiency, piecewise cubic polynomials are the most common choice for Hermite interpolation. A Hermite cubic interpolant is a piecewise cubic polynomial interpolant with a continuous first derivative. Although piecewise cubic Hermite interpolation eliminates the problems of excessive instability of the tails of implied risk-neutral density functions, it appears to sacrifice smoothness of the interpolating function.

Our empirical analysis in Chapter II focuses on three main issues. First, we test whether there are significant differences between the PDF measures (moments and quartiles) across methods. Second, we check the accuracy of our selected methods in replicating observed option prices. Third, we rank the methods based on their robustness to option pricing errors. In order to compare the implied PDFs across methods, we look at conventional measures of location, dispersion, asymmetry, fat-tailness and various tail percentiles of the estimated risk-neutral distributions. The goodness-of-fit of the estimated PDFs is tested by constructing some synthetic measures of option pricing errors. These measures are based on cross-sectional differences between theoretical and observed option prices. In particular, we use the mean squared error, the mean squared percentage error, the mean absolute error and the mean absolute percentage error. The pricing errors are calculated for each day across strike prices, and then averaged across time. Our tests for the stability of the implied risk-neutral distributions are based on Monte Carlo simulations. Thus, we perturb implied volatilities with random uniformly distributed amounts of maximum $\pm 10\%$ of their actual level, and compare the dispersion of the estimated PDF moments and quartiles. This approach is similar to that followed by Bliss and Panigirtzoglou (2002). The difference between our method and theirs is that we perturb implied volatilities rather than actual option prices. Our approach has the advantage of eliminating by construction possible violations of the option pricing arbitrage conditions.

We find that the average errors vary substantially from a method to another and tend to increase with the maturity of option contracts. Thus, the average mean absolute percentage errors for the piecewise cubic Hermite interpolation method is of 0.24% for options with 1-month time to maturity, and increases to 0.26% and 0.28% for 3- and 6-month option contracts. The mean squared percentage errors are much smaller, namely 0.002%, 0.002%, and 0.003% for 1-, 3- and 6-month time to maturity options. The spectrum of errors for the other methods is relatively wide compared with the piecewise cubic Hermite interpolation method, ranging from 1.86% to 20.98%, in the case of the mean absolute percentage errors, and from 0.41% to 9.88%, in the case of mean squared percentage errors. The quadratic interpolation method exhibits the largest pricing errors. We also find that there is a relative trade-off between the goodness-of-fit and the stability of the estimated methods. Thus, methods which have a better accuracy in fitting observed option prices appear to be more sensitive to option pricing errors, while the most stable methods have a fairly disappointing fitting. However, for the first two PDF moments as well as for the quartiles of the risk-neutral distributions we find that the estimates do not differ significantly across methods. This suggests that there is a large scope for selection between these methods without essentially sacrificing the accuracy of the analysis. Nonetheless, depending on the particular use of these PDFs, some methods may be more suitable than the others.[5]

In Chapter III we show how the information extracted from the risk-neutral probability density functions can be used to predict the forward forecasting bias. In this chapter, we also show that, due to the difference in hedging mechanisms between forward and option markets, it is more likely to find a good measure of the foreign exchange risk premium in the options market rather than in the forward market. This is explained by the fact that forward contracts can be perfectly hedged and do not require re-hedging. On the other hand, option contracts need to be continuously re-hedged until their expiry. If the exchange rates are volatile, the costs of re-hedging might be substantial. Moreover, if the foreign exchange markets are less liquid, the re-hedging might turn out to be very expensive for an option writer.

According to the hypothesis of the uncovered interest rate parity, expected changes in the nominal exchange rate should be positively related to the difference in the nominal interest rates across countries. In particular, this implies that the slope coefficient from the regression of the change in exchange rate on interest rate differential should

[5]However, we should bear in mind that all these techniques only provide *risk-neutral* measures of exchange rate distributions. In practice, traders may require high levels of risk premia for their exposures to emerging market currencies. In this case, our estimated risk-neutral distributions may be much less reliable than one would be inclined to believe.

be one. However, this hypothesis has been rejected in a large number of studies, using data for many different countries and time periods (see Hodrick (1987), Froot and Thaler (1990), Lewis (1995) and Engle (1996a) for some extensive surveys). This common empirical finding is called the forward premium (discount) puzzle.

Several explanations have been suggested in the literature to explain this puzzle. Some authors questioned the validity of the empirical models. Issues regarding the exchange rates stationarity or the ergodicity of the data samples analysed were the most popular subjects in this category. Other authors claimed that the hypothesis of the foreign exchange market inefficiency cannot be ruled out. These authors used either simple filter rules or proposed different ways to look at the "bubbles" formation in the FX market. However, the majority of empirical and theoretical papers on the forward premium puzzle focused on the risk premium. The proposed approaches were either based on pure statistical methods, general equilibrium or portfolio-balance models. Despite the wide variety of approaches, the explanations proposed for the forward premium puzzle are still inconclusive.

In Chapter III we propose an ex ante *level* measure of the forward term premium extracted from the risk-neutral probability density functions. Our results show that, regardless of method used to extract risk neutral distributions from option prices or the maturity of option contracts, market participants expected higher depreciations of the exchange rates of emerging economies than those implied by the interest rate parity. This result might indicate either that investors charged higher risk premia for their exposures to emerging market currencies, or that they indeed anticipated larger depreciations of these currencies than those predicted by the interest rate parity. Moreover, we find that the implied depreciation probability of these currencies increases with the maturity of option contracts. This may suggest that foreign exchange market participants are willing to pay higher risk premia to get longer-term protection against the depreciation of an emerging market currency. Our proposed ex ante measure of the forward term premium appears to explain relatively well the realised forward forecasting bias. Thus, we are able to explain up to 57% of the 1-month and up to 85% of the 3- and 6-month forward forecasting biases. We also find that for the currencies with an implicit or explicit peg to the US dollar the fit is substantially better. This result may suggest that the expectation errors might also play an important role in the explanation of the forward forecasting bias.

The last chapter of this thesis provides various tests of the information content of implied volatilities. We first evaluate how good are implied volatilities in predicting future realised volatilities and compare them with GARCH-based volatility forecasts. We find that implied volatilities do not always predict realised volatilities better than GARCH models. Furthermore, our estimates show that the fit decreases slightly with the maturity of option contracts. This indicates that short-term volatilities tend to be better predicted than long-term ones. A possible explanation of these results may be that options with maturities greater than one month are less liquid and therefore their information content may be less reliable.

This chapter also examines the expectations hypothesis of implied volatilities. In particular, it tests whether long-dated volatility quotes are consistent with expected future short-dated volatilities, assuming that traders in the foreign exchange market have rational expectations. Under the assumption that volatility follows a mean-reverting process with constant long-run mean and a constant coefficient of mean reversion, according to the expectations hypothesis, for a given change in the implied volatility of the short-maturity option, there should be a smaller change in the volatility of the long-maturity option. Our estimates show that the expectations hypothesis is rejected for all currencies in our sample. We also find relatively high values for the estimated risk premia. Moreover, the estimated risk premia included in at-the-money volatilities appear to increase with the maturity of option contracts, which suggests that traders require relatively high term premia for their exposures to emerging market currencies.

Lastly, our estimates show significant overreactions of long-term volatilities to changes in the short-term volatility quotes. These findings imply that the information content of at-the-money volatilities extracted from option prices written on these currencies is rather inconsistent.

The analysis provided in this doctoral thesis could be extended by testing various aspects of the informational efficiency of the foreign exchange market which are not covered here. One of these aspects is the reaction of option prices to macroeconomic releases. This will allow to ckeck the speed of adjustment of currency option prices to new information.[6] Anecdotal evidence suggests that the liquidity in option markets dries up substantially during a market stress period. Another possible extension would be to test whether option prices react consistently to market stress events.

Our general conclusion is that, currency option markets in emerging economies could provide useful information about future exchange rate distributions. However, this information is not always consistent and it may require some time until these markets become fully operational and have a better informational efficiency.

[6]For instance, Micu (2003) found that option prices react asymmetrically to macroeconomic announcements. Thus, it appeared that news which potentially lead to the depreciation of emerging market currencies matter more than those which have an expected positive impact on these currencies.