The information content of volatilities implied from currency options: empirical evidence from emerging market countries

Marian Micu*

Bank for International Settlements and

University of Paris 1 Panthéon Sorbonne

- First draft -

January 15, 2006

Abstract

In this paper we test whether volatilities implied from currency options provide better forecasts for future realised volatilities than those obtained with GARCH-type methods. Unlike previous studies that compare the GARCH forecasts with *at-the-money* implied volatilities, we focus on the standard deviation of the whole implied risk-neutral distribution. We find that, in most cases, for the 1-month horizon, implied volatilities from risk-neutral exchange rate distributions provide a better fit than the volatilities implied from at-the-money options or those estimated with a GARCH model. However, our results indicate that the fit decreases with the forecasting horizon. Thus, for 3- and 6-month horizons, there is no clear cut answer on which volatilities offer the best prediction of future realised volatilities. A possible explanation of these results may be that options with maturities higher than one month are less liquid and therefore their information content may be less reliable.

JEL Classifications: F31, G13, G15

Keywords: exchange rates, probability density functions, emerging economies, implied volatilities, GARCH volatilities

*Bank for International Settlements, Centrabahnplatz 2, CH-4002, Basel, Switzerland. Tel: +4161 280 92 79, Fax: +4161 280 91 00, E-mail: Marian.Micu@bis.org. The views expressed in this paper are those of the author and not necessarily the views of the BIS. I would like to thank José Campa, Willi Fritz, Des McManus and Will Melick for helpful discussions on extracting risk-neutral probability density functions from option prices. I am also grateful to John Hawkins for helpful comments on a previous draft of this paper.

1 Introduction

Volatility is the most important variable in the pricing of derivative securities. Volatility is very often associated with uncertainty. In the Black and Scholes (1973) model, the uncertainty associated with future price changes in the underlying asset is the only unknown variable in the pricing function. To price an option, one needs to be able to forecast the volatility of the underlying asset. In finance, volatility is often used to refer to standard deviation of assets' returns.

One of the relatively recently proposed measures for the volatility of asset returns is the *realised volatility*. The term *realised volatility* has been used in Andersen and Bollerslev (1998) to denote the sum of intra-day squared returns at short intervals such as 15 or 5 minutes. Andersen and Bollerslev (1998) showed that such a volatility estimator provides an accurate approximation of the latent process that defines volatility.

Given the apparent lack of any structural dynamic economic theory explaining the variation in higher order moments of asset returns distributions, particularly instrumental in the development of volatility models has been the autoregressive conditional heteroskedastic (ARCH) class of models introduced by Engle (1982). Parallel to the success of standard linear time-series models, arising from the use of the conditional versus the unconditional mean, the key insight offered by the GARCH models lies in the distinction between the conditional and the unconditional second order moments. While the unconditional covariance matrix for the variables of interest may be time-invariant, the conditional variances and covariances often depend non-trivially on the past changes in these variables.

The development of option markets and the general consensus over the Black and Scholes model as a baseline option pricing model, determined an extensive use of implied volatilities as market-based volatility forecasts. However, the existence of different volatilities for the same underlying asset and time to maturity casts doubts on the pricing accuracy of the Black and Scholes model. One possible explanation for this puzzling evidence is that the Black and Scholes model assumes a constant volatility over the lifetime of an option contract. A possible correction would be to use a stochastic volatility measure (see Hull and White (1987)). However, such a measure turned out to be insufficient to explain this puzzle. Other explanations based on market microstructure and measurement errors (eg liquidity, bid-ask spread and tick size) and investor risk preferences have also been proposed.[1] However, none of these models can fully account for the over-pricing of far out-of-the-money options. More importantly, the effect of skewness of the the risk-neutral implied exchange rate distribution on the estimated variance of these distributions has not been accounted for in the previous studies.

To our knowledge, so far, no study has been carried out on the information content of volatilities derived from options written on emerging market currencies. One reason for this may be that these markets were still underdeveloped until a couple of years ago and the data was not readily available until very recently.

In this study, we test whether volatilities implied from currency options provide better forecasts of future realised volatilities than GARCH-type models. We use options data for the Brazilian real, Indonesian rupiah, Mexican peso, South African rand, South Korean won and Thai bath. Observations are for options with maturities of one, three and six months, with a daily frequency, from 10 November 1997 to 10 November 2002.

Since options of different strike prices produce different implied volatilities, a decision has to be made as to which of these implied volatilities should be used, or which weighting scheme should be adopted. The most common solution is to choose the volatility derived from at-the-money options. There are at least two reasons for making this choice. First, in a stochastic volatility model, the volatility implied by at-the-money options appears to provide the closest approximation to the average volatility over the lifetime of an option contract, provided that volatility risk premium is either zero or constant. Second, at-the-money option contracts are the most liquid, and hence the at-the-money volatilities derived from options written on such contracts are least prone to measurement errors. However,

[1] Given that volatility is not a directly tradable asset, the hedging mechanism used in the Black and Scholes model may not apply and the risk neutral valuation principle has to be modified since volatility may include a risk premium. Different approaches to this problem have been adopted. Hull and White (1987) assumed that the volatility risk is not priced. Wiggins (1987) derived various specifications of volatility risk premium according to different assumptions for risk preferences. Heston (1993) provided a specification where volatility risk premium is proportional to variance and extracted the volatility risk premium from option prices in the same manner as implied volatility is extracted. He also found that, in the case where volatility is stochastic and uncorrelated with the change in the underlying asset, the Black and Scholes model overprices the at-the-money options and underprices both in- and out-of-the-money options. Moreover, he found that the degree of overpricing increases with the time to maturity.

the use of at-the-money volatilities alone would lead to an omission of potential information contained in options with strike prices different from the current forward rate.

Another possibility would be to incorporate the information provided by in- or out-of-the-money options by the estimating the standard deviation of the whole risk-neutral probability density function. The advantage of this measure is that it can theoretically account better for possible asymmetries of future exchange rate distributions. A possible drawback is that options with strike prices away from the current forward rate might be less liquid. Thus, this measure may be plagued by potential measurement errors.

We find that volatility is highly persistent and the exchange rate distributions exhibit leptokurtosis. Moreover, it appears that positive changes in exchange rates (meaning a depreciation of emerging market currencies) lead to higher volatilities than negative changes. Our results show that, in most cases, for the 1-month horizon, implied volatilities from risk-neutral exchange rate distributions provide a better fit than the volatilities forecasted with the GARCH-type model or those implied by the at-the-money option prices. However, our estimates indicate that the fit decreases with the forecasting horizon. Thus, short-term volatilities appear to be better predicted than the long-term ones. Moreover, for 3- and 6-month horizons, there is no clear cut answer on which volatilities provide the best prediction of future realised volatilities.

The reminder of this paper is organised as follows. In Section 2 we review the main findings in the literature on the predictive power of implied volatility. Section 3 compares the volatility forecasts derived from GARCH models with those obtained from the risk-neutral density forecasts. Section 4 describes the data. Our empirical estimates are presented in Section 5. Section 6 concludes.

2 Volatility prediction in the foreign exchange market

Given the existence of a multitude of volatility measures and forecasting models, an unavoidable question is which of these models provide the best volatility forecast. There is an extensive literature that compares different volatility forecasts. However, it is beyond the scope of this paper to review it. A recent good review is offered in Poon and Granger (2003).

In this paper, we will only review some of the papers which compared implied volatility with historical-based volatility forecasts in the context of foreign exchange markets.

One of the earliest studies that documented a strong predictive power for the implied volatility was proposed by Scott and Tucker (1988). They used transactions data for the British pound, the Canadian dollar, Deutsche mark, Japanese yen, and the Swiss franc traded on the Philadelphia Stock Exchange (PHLX) for the period March 14, 1983 through March 13, 1987. To assess the predictive accuracy of implied volatility they run regressions having the realised standard deviation over the lifetime of the option contract as a dependent variable and the implied volatility as an independent variable. Their estimates showed a strong relationship between these variables for 3-, 6-, and 9-month horizons, with adjusted R^2 in the range of 40-50%.[2]

Most of the empirical papers which compared historical and implied volatilities found that implied volatilities provide a better fit for the realised volatilities than historical methods. Xu and Taylor (1995) used daily closing option prices for the British pound, Deutsche mark, Japanese yen, and Swiss franc, quoted against the US dollar at the Philadelphia Stock Exchange (PHLX). Their specifications for historical volatilities were based on two GARCH-type models: GARCH(1,1) and Exponential GARCH. The estimates suggested that the Philadelphia currency options market is informationally efficient. Their conclusions were in contrast with those from earlier papers which identified a lack of informational efficiency for US stock options markets (eg Lamoureux and Lastrapes (1993) and Canina and Figlewski (1993)). They suggested that the existence of a lower cost for arbitrage trading in the foreign exchange compared with the stock market may be a possible explanation of this finding.

Jorion (1995) carried out a similar exercise by using data on currencies traded on the Chicago Mercantile Exchange (CME). For each option contract he matched the implied at-the-money volatility with the sequence of price movements on the underlying futures contract until option expiration. The *realised volatility* in his study is calculated as the variance of continuously compounded futures returns. He found that statistical time-series models, even when given the advantage of ex-post parameter estimates, are outperformed

[2] These results were also confirmed by more recent studies (see, for instance, Li (2002) for a survey).

by option-implied forecasts.

More recent models used long-memory specifications for the volatility forecast, and were built on volatility compiled from high frequency intra-day returns, while the implied volatility remained to be constructed from less frequent daily option prices. One example in this direction is the paper by Pong et al. (2003) which compared forecasts of the realised volatility of the pound, mark and yen exchange rates against the dollar with forecasts obtained from a short memory ARMA model, a long memory ARFIMA model, a GARCH model and option implied volatilities. They found that intra-day rates provided the most accurate forecasts for the one-day and one-week forecast horizons, while implied volatilities were at least as accurate as the historical forecasts for the one-month and three-month horizons. They argued that the superior accuracy of the historical forecasts relative to implied volatilities came from the use of high-frequency returns, rather than from a long-memory specification.

3 GARCH models versus risk-neutral density forecasts

The estimation and the forecasting of volatility is usually accomplished with time series methods. However, such models have a couple of drawbacks. One of these is that they are by nature backward-looking. This means that they assume the whole set of information about future exchange rates is fully included in their historical prices. Another inconvenience is that, in most of these models, the parameters do not change with the arrival of new information, ie the parameters are not time-varying. A third limitation of time series models is that they only offer some information about a (central) point estimate of future exchange rate prices or volatilities but not about the future distribution.

Option prices provide a more comprehensive description of the expected future distribution of exchange rates. More importantly, the forecasts based on implied risk-neutral distributions are forward-looking. However, these forecasts only provide the *risk-neutral* probability value of a set of future possible exchange rate prices. The risk-neutrality assumption implies that investors are assumed not to charge an additional premium for a change in the uncertainty about future exchange rates. It also means that all investors are assumed to have similar risk preferences. Nevertheless, in the real world market participants are not risk-neutral. The effect of risk on investors' decision-making process depends on their wealth, their utility function and the perceived level of risk in the market (see Pratt (1964)). Furthermore, investors have heterogeneous risk preferences and therefore risk-neutral distributions might not perfectly match the "real" distributions used by market participants to price foreign exchange options.

GARCH models have been widely used in the last two decades to estimate and forecast volatility in financial markets. Univariate GARCH models consist of two equations. The first, the mean equation, describes the observed data as a function of other variables plus an error term. The second, the variance equation, specifies the evolution of the conditional variance of the error term from the mean equation as a function of past conditional variances and lagged errors. The specification of the mean equation is not without interest. However, in our paper, we only focus on the variance equation.

The common denominator of existing GARCH models is that they postulate that a transformation of the conditional standard deviation is linearly related to (nonnegative) functions of past and present shocks, plus a moving average of transformed standard deviations. In the case of Bollerslev's (1986) GARCH(p, q) model, the transformations for the shocks and the standard deviations are simple quadratic forms, so that:

$$h_t = \sigma_t^2 = \omega + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2 \tag{1}$$

where $h_t = \sigma_t^2$ denotes the variance of exchange rate returns at time t and ε_t is the residual of an OLS regression which has the logarithmic changes in exchange rate prices as dependent variable and a constant as independent variable. The conditional variance equation in (1) can also be expressed as:

$$\varepsilon_t^2 = \omega + \sum_{i=1}^{q} (\alpha_i + \beta_i)\, \varepsilon_{t-i}^2 - \sum_{j=1}^{p} \beta_j h_{t-j} + \eta_t \tag{2}$$

where $\eta_t = \varepsilon_t^2 - h_t$ and $m = \max(p, q)$. Therefore, ε_t^2 will have the usual properties of an (autoregressive moving average) ARMA(m, p) process, so standard identification procedures for the orders of p and m can be carried out on the ε_t^2 series (see Bollerslev (1988)).

Although the use of the conditional normal distribution in (1) will generate a leptokurtic unconditional distribution for exchange rate returns, initial estimation work revealed that it still did not adequately account for the degree of fat-tailedness in the unconditional distribution. Thus, to produce a more adequate representation of this data, an alternative is to use a leptokurtic conditional distribution for the demeaned exchange rate returns along with the GARCH conditional variance model. Following Baillie and Bollerslev (1989), in our study we employ a Student-t distribution. For a standardised t distribution with ν degrees of freedom, the log-likelihood function is given by

$$\begin{aligned} \ln(L) &= T\left[\ln\Gamma\left(\frac{\nu+1}{2}\right)-\ln\Gamma\left(\frac{\nu}{2}\right)-\frac{1}{2}\ln(\nu-2)\right] \\ &\quad -\frac{1}{2}\sum_{t=1}^{T}\left[\ln h_t+(\nu+1)\ln\left(1+\frac{\varepsilon_t^2}{h_t(\nu-2)}\right)\right] \end{aligned} \quad (3)$$

where Γ denotes the usual gamma function (see Bollerslev (1987)).[3]

Despite the apparent success of GARCH(1,1) models, sometimes they cannot capture some important features of the data. One of these features is the asymmetric response of volatility to positive and negative news. Statistically, this effect is present in stock and bond prices when an unexpected drop in price (bad news) increases predictable volatility more than an unexpected increase in price (good news) of similar magnitude. However, for exchange rates it is not always clear how this impact should work because of the uncertainty about the portfolios composition of FX investors. Engle and Ng (1993) proposed a non-linear asymmetric GARCH model to deal with this issue. The conditional variance in a non-linear asymmetric GARCH model is given as follows:

$$h_t=\alpha+\beta h_{t-1}+\omega\left(\varepsilon_{t-1}+\gamma\sqrt{h_{t-1}}\right)^2 \quad (4)$$

Another well-known model proposed to capture asymmetric responses of volatility to

[3] A t-distribution is symmetric around 0, with the variance and the fourth moment equal to $Var(\varepsilon_t)=h_t$, $E\left(\varepsilon_t^4\right)=3(\nu-2)(\nu-4)^{-1}h_t^2$, and $\nu>4$. It is well known that for $1/\nu\to 0$ the t-distribution approaches a normal distribution with variance $h_{t|t-1}$, but for $1/\nu>0$ the $t-$distribution has "fatter tails" than the corresponding normal distribution.

shocks in exchange rate returns is Nelson's (1991) exponential GARCH (or EGARCH) model. The conditional variance in an exponential GARCH model is specified as:

$$\ln(h_t) = \alpha + \beta \ln(h_{t-1}) + \gamma \frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + \omega \left[\frac{|\varepsilon_{t-1}|}{\sqrt{h_{t-1}}} - \sqrt{2/\pi} \right] \quad (5)$$

where ω, β, γ, and α are constant parameters. The EGARCH model is asymmetric because the level of $\varepsilon_{t-1}/\sqrt{h_{t-1}}$ is included in the variance equation. Since γ, the coefficient attached to $\varepsilon_{t-1}/\sqrt{h_{t-1}}$, is usually negative, positive returns shocks generate less volatility than negative shocks, all else being equal.

For the estimation of the implied standard deviation of the whole risk-neutral exchange rate distribution we use the Hermite polynomials method proposed by Madan and Milne (1994).[4] Madan and Milne (1994) modelled the prices of contingent claims as elements of a separable Hilbert space that has a countable orthogonal basis. They noticed that one may think of the basis elements as analogous to factors in asset pricing. Thus, pricing in terms of a Hilbert space is analogous to the use of discount bonds as a basis for pricing fixed income securities or the construction of branches of a binomial tree in pricing options. However, a Hilbert space basis is in general difficult to construct because it requires a knowledge of the stochastic process of the underlying asset prices. Madan and Milne showed that, under fairly general conditions, one can specialise the Hilbert space basis to the family of Hermite polynomials. Using 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)). In the case of exchange rates, the Black and Scholes model is a parametric special case of the Madan and Milne (1994) model. Thus, the Hermite polynomials approximation is equivalent to performing a Fourier expansion to the baseline lognormal solution obtained from the Black and Scholes model. More precisely, the risk-neutral distribution is obtained through successive

[4] Various techniques have been proposed to extract risk-neutral PDFs from option prices. Most of these methods provide similar estimates for the first two moments of the risk-neutral distribution. However, for the higher moments of the distributions, the skewness and kurtosis, the estimates appear to be model-dependent (see, for instance, the papers by Melick and Thomas (1998) or McManus (1999)).

orthogonal perturbations to a normalised density function.[5] Thus, in the case of European currency options, the Hermite polynomial adjustments are constructed with respect to the normalised stochastic variable:

$$z = \frac{\ln\left(\frac{S_T}{F_t}\right) - \left(\mu - \frac{1}{2}\sigma^2\right)\tau}{\sigma\sqrt{\tau}}, \qquad z \sim N(0,1) \tag{6}$$

where S_t is the current spot exchange rate, τ is the time to maturity of the option contract and σ is the instantaneous standard deviation of the Wiener process that characterises exchange rate changes. $F_t = S_t e^{(r-r^*)\tau}$ is the forward price and $\mu = \ln(S_t) + \left(r - r^* - \sigma^2/2\right)\tau$ is the mean of the diffusion process from the Black and Scholes model. r and r^* are the risk-free interest rates in domestic and foreign currency. Thus, the risk-neutral normal distribution used as reference for the Hermite polynomials approximation is:

$$n(z) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{z^2}{2}\right) \tag{7}$$

The Hermite polynomials approximation of the risk-neutral density function, $p^{HP}(z)$, can be written as:

$$p^{HP}(z) = \lambda(z)\, n(z) \tag{8}$$

where $\lambda(z)$ denotes the departures from the reference distribution $n(z)$, which are captured by an infinite summation of Hermite polynomials, that is:

$$\lambda(z) = \sum_{k=0}^{\infty} b_k \phi_k(z) \tag{9}$$

where b_k are constants which have to be estimated and

$$\phi_k(z) = \frac{(-1)^k}{k!} \frac{1}{n(z)} \frac{\partial^k n(z)}{\partial z^k} = -\frac{1}{\sqrt{k}} \frac{\partial \phi_{k-1}(z)}{\partial z} + \frac{1}{\sqrt{k}} z \phi_{k-1}(z) \tag{10}$$

is an orthogonal system of standardised Hermite polynomials. The first four standardised

[5] In other words, rather than assuming specific expressions for *the change* in the risk-neutral probabilities, as one does under the martingale approach for option valuation, Madan and Milne (1994) assume a parametric structure for the risk-neutral density function itself.

Hermite polynomials are:[6]

$$
\begin{aligned}
\phi_0(z) &= 1 \\
\phi_1(z) &= z \\
\phi_2(z) &= \frac{1}{\sqrt{2}}\left(z^2-1\right) \\
\phi_3(z) &= \frac{1}{\sqrt{6}}\left(z^3-3z\right) \\
\phi_4(z) &= \frac{1}{\sqrt{24}}\left(z^4-6z^2+3\right)
\end{aligned} \tag{11}
$$

With the above notations and $d_2 = \frac{1}{\sigma\sqrt{\tau}}\left[\ln\left(\frac{S_t}{K}\right) + \left(r - r^* - \frac{\sigma^2}{2}\right)\tau\right]$, one can write the Black and Scholes formula for the European-style call currency option as:[7]

$$
\begin{aligned}
C_{BS}(z) &= e^{-r\tau}\int_{-d_2}^{\infty}\left[F_t e^{\left(\mu-\frac{\sigma^2}{2}\right)\tau+z\sigma\sqrt{\tau}} - K\right] n(z)\,dz \qquad (12) \\
&= e^{-r^*\tau} S_t \int_{-d_2}^{\infty} e^{\left(\mu-\frac{\sigma^2}{2}\right)\tau+z\sigma\sqrt{\tau}} n(z)\,dz - e^{-r\tau} K \int_{-d_2}^{\infty} n(z)\,dz
\end{aligned}
$$

By substituting the risk-neutral density function $p^{HP}(z)$ from equations (8) and (9), we obtain the Hermite polynomial approximation of the call option price:

$$
\begin{aligned}
C_{HP}(z) &= e^{-r^*\tau} S_t \int_{-d_2}^{\infty} e^{\left(\mu-\frac{\sigma^2}{2}\right)\tau+z\sigma\sqrt{\tau}} \sum_{k=0}^{\infty} b_k \phi_k(z)\, n(z)\,dz \\
&\quad - e^{-r\tau} K \int_{-d_2}^{\infty} \sum_{k=0}^{\infty} b_k \phi_k(z)\, n(z)\,dz
\end{aligned} \tag{13}
$$

To evaluate the price of this option, one needs to replicate its payoff and estimate the coefficients b_k. The expected payoff of the option can be expressed as:

[6] Higher-order Hermite polynomials can be easily calculated using the recurrence relationship: $\phi_k(z) = \frac{z}{\sqrt{k}}\phi_{k-1}(z) - \sqrt{\frac{k-1}{k}}\phi_{k-2}(z)$. The polynomials are orthogonal because $\int_{-\infty}^{\infty}\phi_k(z)\,\phi_j(z)\,n(z)\,dz$ equals one if $k = j$ and zero otherwise.

[7] To simplify the presentation we only give the derivation for the call option. The analysis for the put option valuation with a Hermite polynomial approximation is straightforward and follows the same reasoning as for the call option.

$$g_{C_{HP}}(z) = \int_{-d_2}^{\infty} \left[F_t e^{\left(\mu - \frac{\sigma^2}{2}\right)\tau + z\sigma\sqrt{\tau}} - K \right] \sum_{k=0}^{\infty} b_k \phi_k(z) \, n(z) \, dz \qquad (14)$$

and the call option price from equation (13) can be represented as follows:

$$C_{HP}(z) = e^{-r\tau} \sum_{k=0}^{\infty} \alpha_k b_k \qquad (15)$$

where

$$\alpha_k = \int_{-\infty}^{\infty} g_{C_{HP}}(z) \, \phi_k(z) \, n(z) \, dz \qquad (16)$$

Madan and Milne (1994) showed that, given the assumed probability model, the Hermite polynomial coefficients α_k are well defined and hence the b_k can be inferred from the observed option prices. The α_k coefficients are defined as:

$$\alpha_k = \frac{1}{\sqrt{k!}} \frac{\partial^k \Phi(u)}{\partial u^k} |_{u=0} \qquad (17)$$

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

$$\begin{aligned} \Phi(u) &= S_t e^{(r-r^*)\tau} e^{\mu\tau + \sigma\sqrt{\tau}u} N[d_1(u)] - KN[d_2(u)] \\ d_1(u) &= \frac{\ln(S_t/X) + (r - r^*)\tau}{\sigma\sqrt{\tau}} + \frac{1}{2}\sigma\sqrt{\tau} + u \\ d_2(u) &= d_1(u) - \sigma\sqrt{\tau} \end{aligned} \qquad (18)$$

For practical estimation purposes, the infinite sum of Hermite polynomials must be truncated at a finite order in z. In our study we truncate it at the fourth order.[8] In order to ensure that the risk-neutral PDF for the Hermite polynomial approximation behaves as a density function, the following restrictions are usually imposed (see Abken, Madan and Ramamurtie (1996)): $\beta_0 = 1$, $\beta_1 = 0$ and $\beta_2 = 0$. Under these restrictions, the risk-neutral probability density function for the fourth-order Hermite polynomials approximation is given

[8] Given the data restrictions, higher-order approximations are practically infeasable for foreign exchange over-the-counter options.

by:

$$p^{HP}(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}} \left(1 + \frac{3b_4}{\sqrt{24}} - \frac{3b_3}{\sqrt{6}} z - \frac{6b_4}{\sqrt{24}} z^2 + \frac{b_3}{\sqrt{6}} z^3 + \frac{b_4}{\sqrt{24}} z^4\right) \tag{19}$$

and the risk-neutral PDF for the exchange rate at the maturity of the call option is:

$$p^{HP}(S_T) = \frac{1}{S_T \sigma \sqrt{\tau}} p^{HP} \left[\frac{\ln(S_T/F_t) - \left(\mu - \frac{1}{2}\sigma^2\right)\tau}{\sigma\sqrt{\tau}}\right] \tag{20}$$

The mean μ, the standard deviation σ and the parameters b_3 and b_4 are estimated by minimising the sum of squared differences between theoretical and observed option and forward prices. The implied standard deviation of the PDF is $\sigma_{p^{HP}(S_T)} = \sigma\sqrt{\tau}$.

4 Data

While options on currencies and currency futures of developed countries are traded both on exchanges and over-the-counter, the options written on currencies of emerging countries are almost exclusively traded over-the-counter. Over-the-counter options are European-type. They are usually traded for standardised maturities. Price quotes are expressed as implied volatilities corresponding to defined levels of options' delta, which traders by agreement substitute into the Black and Scholes formula to determine the option premium. Since the volatility is the only unobservable parameter in the Black and Scholes model, these volatilities uniquely determine the options' prices. However, this does not necessarily mean traders believe that the Black and Scholes formula gives a fair evaluation of option prices. This market convention simply allows a direct mapping from implied volatility quotes into option prices.

Our data consists of market quotes of over-the-counter options on six emerging market currencies against the US dollar. These currencies are: the Brazilian real, Indonesian rupiah, Mexican peso, South African rand, South Korean won and Thai bath. 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 major global foreign exchange dealer bank.[9] Observations are for options with maturities of one, three and six months, with a daily frequency from 10 November 1997 to 10 November 2002.

Although the Black and Scholes model assumes constant volatilities across exercise prices, the implied volatility quoted by option traders for our selected currencies typically varies as a function of options' strike prices. This reflects a departure from the Black and Scholes formula's assumptions, implying that traders assume a non-lognormal distribution for the future exchange rates when they price these options. In general, our data shows that the implied volatility is lowest for at-the-money options, increasing for both in- and out-of-the money options. This pattern is referred to as the "volatility smile" and is consistent with the leptokurtosis of the distribution of future exchange rate returns. In general, the probability of future exchange rate realisations is not symmetrically distributed around the at-the-money strike price. For most of the selected currencies in our sample, we noticed a greater probability of a large depreciation. This is translated into positive risk-reversal quotes and positive skewness of the risk-neutral distributions.

A potential concern about the data is the liquidity of the options contracts. If the daily quotes on these options are illiquid, the information content of the implied PDFs may become noisy. A graphical inspection of our data series suggests that for some short periods stale prices are present, especially for risk-reversals and strangles. This may reduce the information content of estimated PDFs.

5 Empirical evidence

In order to test whether implied volatilities from risk-neutral probability density functions provide better forecasts for realised volatilities than GARCH-type methods, we define $\sigma_{t,T}$ to be the realised volatility over the lifetime of the option contract, measured from day t to day T. The future daily variance is defined as the arithmetic average of squared returns from the trading day of the option until the expiry date, without adjustment for the mean:

[9] Given that these interest rates are usually used by market participants for interbank loans, they should, theoretically, include a default risk premium. However, on this maket, the participants are top global financial and banking institutions which have short-term risks of default close to zero.

$$\sigma_{t,T}^2 = \frac{1}{T-t}\sum_{k=t}^{T-t} R_{t+k}^2 \tag{21}$$

where R_t^2 are the squared returns at time t, and T denotes the maturity of the option contract. The predictive power of a volatility forecast can be estimated by regressing the realised volatility on forecast volatility:

$$\sigma_{t,T} = \alpha + \beta\widehat{\sigma}_{t,T} + \varepsilon_{t,T} \tag{22}$$

where $\widehat{\sigma}_{t,T}$ is the volatility forecast measured on day t for the period between t and T. We use for $\widehat{\sigma}_{t,T}$ the standard deviation derived from the risk-neutral probability density functions or from at-the-money implied volatilities as well as the volatility forecast with our selected GARCH models.

The results of the tests are presented in Appendix A. We first examine which GARCH model fits better our exchange rate data. The results for the GARCH(1,1) model are presented in Table A.1. We find that the volatility of exchange rate returns is highly persistent, with γ coefficients between 0.63 and 0.90.

In order to test the specification of the GARCH$(1,1)$ model, we use two tests: the Lagrange Multiplier (LM) and the Kolmovorov-Smirnov (KS). The null hypothesis of the Lagrange Multiplier test is that the standardised residuals of the model are not correlated up to the fifth lag.[10] We find that for all countries but South Africa the autocorrelation of the residuals is removed by using a GARCH$(1,1)$ model. However, our Kolmogorov-Smirnov tests for the normality of the residuals reject in all cases the null hypothesis that the standardised residuals follow a normal distribution.

To correct for the non-normality of the residuals, we employ a GARCH$(1,1)$ model with t-distributed errors. This model captures some important features of the non-normality of the residuals (ie the leptokurtosis). Indeed, the hypothesis that the distribution of the residuals is close to the normal distribution ($1/\nu = 0$) is rejected in all cases. However, when we test whether the standardised residuals from this model follow a normal distribution, we

[10] We have tested various lag-lengths but the results did not change. We only report in the table the values of the test for five lags of the estimated residuals.

find that the model is unable to correct for the non-normality of the residuals (albeit the LM test shows that, in all cases, we cannot reject the null hypothesis of no autocorrelation).

The asymmetric response of volatility to positive and negative changes in exchange rates is tested with a nonlinear asymmetric GARCH and an exponential GARCH (EGARCH) model. The results of these tests are presented in Tables A.3. and A.4. For both models, the cases of misspecification were more frequent. Thus, the LM test rejects the null hypothesis of no autocorrelation of the residuals for South Korea and South Africa, in the case of the nonlinear asymmetric GARCH model, and for Mexico, South Korea, Thailand and South Africa, in the case of the exponential GARCH model. However, we do find significant asymmetric responses of volatilities to unexpected exchange rate changes. This is observed through positive and statistically significant ω coefficients. The difference between our findings and those observed in the stock market is that the values of the ω are positive. This suggests that positive changes in our selected exchange rates (meaning a depreciation of emerging market currencies) lead to higher volatilities than negative changes.

In order to select the most appropriate GARCH model, we use the likelihood ratio test. The likelihood ratio test is a statistical test of the goodness-of-fit between two models. A relatively more complex model is compared to a simpler model to see whether it fits a particular dataset significantly better. The likelihood ratio test is only valid if used to compare hierarchically nested models. That is, the more complex model must differ from the simple model only by the addition of one or more parameters. Adding additional parameters will always result in a higher likelihood score. However, there comes a point when adding additional parameters is no longer justified in terms of significant improvement in fit of a model to a particular data sample. The likelihood ratio test provides one objective criterion for selecting among possible models. To carry out this test we first need to obtain the likelihood scores from the maximisation of the likelihood function attached to each estimated model.[11] The likelihood ratio test is obtained as

$$LR = 2\left(\ln \ell_1 - \ln \ell_2\right) \tag{23}$$

[11] For the specification of these likelihood functions see Bollerslev (1986, 1987), Baillie and Bollerslev (1989), Nelson (1991) and Engle and Ng (1993).

where ℓ_1 and ℓ_2 are the likelihood scores derived from the first and the second model respectively. The LR approximately follows a chi-square distribution. To determine whether the difference in likelihood scores is statistically significant, we must consider the number of degrees of freedom of the χ^2 distribution, which is equal to the number of additional parameters in the most complex model. The results of our tests show that the simple GARCH$(1, 1)$ model and sometimes the GARCH$(1, 1)$ model with t-distributed errors outperformed other competing models.

We choose the simple GARCH$(1, 1)$ model and forecast volatility for 1-, 3- and 6-month horizons. We then test equation (22). We estimate the equation by GMM, using Hansen's (1982) method, and correct the standard errors for heteroskedasticity using White's (1980) model. The results of the tests are presented in Tables A.5. - A.7. We find two interesting results. First, in most cases, for the 1-month horizon, implied volatilities from the whole risk-neutral distribution provide a better fit than at-the-money implied volatilities or the volatilities forecasted with the GARCH$(1, 1)$ model. The only exception is Indonesia where GARCH forecasts seem to provide a better fit. The second interesting finding is that the fit decreases with the forecasting horizon. Thus, short-term volatilities appear to be better predicted than the long-term ones. Furthermore, for 3- and 6-month horizons, there is no clear cut answer on which volatilities provide the best prediction of future realised volatilities. A possible explanation of these results may be that options with maturities higher than one month are less liquid and therefore their information content may be less reliable. From these regressions we can conclude that, for short-term maturities, implied volatilities from the whole risk-neutral exchange rate distribution offer better predictions than historical models. However, the forecasting performance of these volatilities appears to decrease with the maturity of option contracts, making them less suitable forecasts for longer term horizons.

6 Conclusions

In this paper we test whether volatilities implied from the whole risk-neutral distribution of exchange rate prices provide better forecasts for the future realised volatility than those obtained from GARCH-type models. We use daily data for options with maturities of one, three and six months, for the US dollar exchange rates of Brazil, Indonesia, Mexico, South Africa, South Korea and Thailand, over the period 10 November 1997 to 10 November 2002. We find that volatilities implied from the risk-neutral density functions predict future realised volatilities better than GARCH-type models or at-the-money implied volatilities at the 1-month horizon. However, our estimates indicate that the fit decreases with the forecasting horizon. Thus, for 3- and 6-month horizons, there is no clear cut answer on which volatilities offer the best prediction of future realised volatilities.

References

[1] Abken, Peter A., Madan, Dilip B., and Ramamurtie, Sailesh, (1996), "Estimation of Risk-Neutral and Statistical Densities By Hermite Polynomial Approximation: With an Application to Eurodollar Futures Options", *Federal Reserve Bank of Atlanta Working Paper 96-5*.

[2] Andersen, Torben G., and Bollerslev, Tim, (1998), "Answering the skeptics: yes, standard volatility models do provide accurate forecasts", *International Economic Review*, Vol. 39, Issue 4, pp. 885-905.

[3] Baillie, Richard T., and Bollerslev, Tim, (1989), "Common stochastic trends in a system of exchange rates", *Journal of Finance*, Vol. 44, pp. 167-181.

[4] Black, Fischer, and Scholes, Myron, (1973), "The Pricing of Options and Corporate Liabilities", *Journal of Political Economy*, Vol. 81, pp. 637-59, May-June.

[5] Bollerslev, Tim, (1986), "Generalised autoregressive conditional heteroskedasticity", *Journal of Econometrics*, Vol. 31, pp. 307-327.

[6] Bollerslev, Tim, (1987), "A Conditional Heteroskedasticity Time Series Model for Speculative Prices and Rates of Return", *Review of Economics and Statistics*, Vol. 69, pp. 542-547.

[7] Bollerslev, Tim, (1988), "On the Correlation Structure for the Generalized Autoregressive Conditional Heteroskedasticity Process", *Journal of Time Series Analysis*, Vol. 9, pp. 121-131.

[8] Canina, Linda, and Figlewski, Stephen, (1993), "The Informational Content of Implied Volatility", *Review of Financial Studies*, Vol. 6, No. 3, pp. 659-681.

[9] Coutant, Sophie, (1999), "Implied risk aversion in option prices using Hermite polynomials", in "*Estimating and interpreting probability density functions*", Bank for International Settlements.

[10] Coutant, Sophie, Rockinger, Michael, and Jondeau, Eric, (2001), "Reading PIBOR Futures Options' Smile: The 1997 Snap Election", *Journal of Banking and Finance*, Vol 25, No. 11, pp 1957-1987.

[11] Engle Robert F., (1982), "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation", *Econometrica*, Vol. 50, Issue 4, pp. 987-1007.

[12] Engle, Robert F., and Ng, Victor K., (1993), "Measuring the Impact of News on Volatility", *Journal of Finance*, Vol. 48, No. 5, pp. 1749-1778.

[13] Hansen, Lars Peter, (1982), "Large sample properties of generalized method of moments estimators", *Econometrica*, Vol. 50, pp. 1029-1054.

[14] Heston, Steven L., (1993), "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options", *Review of Financial Studies*, Vol. 6, pp. 327-343.

[15] Hull, John, and White, Allan, (1987), "The Pricing of Options on Assets with Stochastic Volatilities", *Journal of Finance*, Vol. 42, No. 2, pp. 281-300.

[16] Jondeau, Eric, and Rockinger, Michael, (2000), "Reading the smile: the message conveyed by methods which infer risk neutral densities", *Journal of International Money and Finance*, Vol. 19, pp. 885-915.

[17] Jorion, Philippe, (1995) "Predicting volatility in the foreign exchange market", *Journal of Finance*, Vol. 50, No. 2, pp. 507-528.

[18] Lamoureux, Christopher G. and Lastrapes, William D., (1993), "Forecasting stock-return variance: toward an understanding of stochastic implied volatilities", *Review of Financial Studies*, Vol. 6, pp. 293-326.

[19] Li Kai, (2002), "Long-memory versus option-implied volatility prediction", *Journal of Derivatives*, Vol. 9, Issue 3, pp. 9-25.

[20] Madan, Dilip B., and Milne, Frank, (1994), "Contingent claims valued and hedged by pricing and investing in a basis", *Mathematical Finance*, Vol. 4, pp. 223-245.

[21] McManus, Des J., (1999), "The information content of interest rate futures options", *Working Paper, No. 99-15*, Bank of Canada.

[22] Melick, William R., and Thomas, Charles P., (1998), "Confidence Intervals and Constant-Maturity Series for Probability Measures Extracted from Options Prices", Bank of Canada, mimeo.

[23] Nelson, Daniel B., (1991), "Conditional Heteroskedasticity in Asset Returns: A New Approach", *Econometrica*, Vol. 59, No. 2, pp. 347-370.

[24] Pong, Shiu-Yan, Shackleton, Mark, Taylor, Stephen J., and Xu, Xinzhong, (2003), "Forecasting Currency Volatility: A Comparison of Implied Volatilities and AR(FI)MA Models ", University of Lancaster, mimeo.

[25] Poon, Ser-Huang, and Granger, Clive W., (2003), "Forecasting volatility in financial markets", *Journal of Economic Literature*, Vol. 41, pp. 478-539.

[26] Pratt, John W., (1964), "Risk Aversion in the Small and in the Large", *Econometrica*, Vol. 32, No. 1/2, pp. 122-136.

[27] Scott, Elton, and Tucker, Alan L., (1988), "Predicting Currency Return Variance", *Journal of Banking and Finance*, Vol. 13, No. 4, pp. 351-368.

[28] White, Halbert, (1980), "A Heteroskedasticity-Consistent Covariance Matrix and a Direct Test for Heteroskedasticity", *Econometrica*, Vol. 48, pp. 817-838, May.

[29] Wiggins, James B., (1987), "Option Values Under Stochastic Volatility: Theory and Empirical Estimates", *Journal of Financial Economics*, Vol. 19, pp. 351-372.

[30] Xu, Xinzhong, and Taylor, Stephen J., (1995), "Conditional volatility and the informational efficiency of the PHLX currency options market", *Journal of Banking and Finance*, Vol. 19, Issue 5, pp. 803-821.

Appendix A

Table A.1. GARCH(1,1) model

The conditional variance equation of exchange rate returns is specified as: $h_t = \alpha + \beta\varepsilon_{t-1}^2 + \gamma h_{t-1}$, where α, β and γ are constant parameters. The sample period is from 10 November 1997 to 10 November 2002. Regressions use daily observations. Robust standard errors are presented in parentheses. * indicates a coefficient significantly different from zero at the 10%, ** at the 5% and *** at the 1% confidence level. The null hypothesis of the Langrange Multiplier (LM) test is that the standardised residuals are not correlated up to the fifth lag. The null hypothesis of the Kolmogorov-Smirnov (KS) test is that the standardised residuals follow a normal distribution. P-values for the LM and KS tests are presented in parentheses.

Country	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\gamma}$	Log likelihood	LM test	KS test
Brazil	.037 (0.0357)	0.167*** (0.0441)	0.823*** (0.0452)	-1597.3	0.157 (1.00)	0.14 (0.00)
Mexico	0.0445*** (0.012)	0.205*** (0.0492)	0.634*** (0.0679)	-896.46	0.966 (0.965)	0.0728 (0.00)
Indonesia	0.131 (0.084)	0.227*** (0.0658)	0.773*** (0.0489)	-2430.8	0.654 (0.985)	0.134 (0.00)
South Korea	0.00333** (0.00149)	0.143*** (0.0267)	0.857*** (0.0243)	-1047.8	8.66 (0.123)	0.0867 (0.00)
Thailand	0.00231 (0.00182)	0.0945*** (0.0191)	0.903*** (0.0174)	-922.61	0.759 (0.98)	0.0762 (0.00)
South Africa	0.00466 (0.00291)	0.0991*** (0.0205)	0.901*** (0.021)	-1405.2	16.8 (0.00488)	0.0742 (0.00)

Table A.2. GARCH(1,1) model with t-distributed errors

The conditional variance equation of exchange rate returns is specified as: $h_t = \alpha + \beta\varepsilon_{t-1}^2 + \gamma h_{t-1}$, with the likelihood function $\ln L = T\left[\ln\Gamma\left(\frac{\upsilon+1}{2}\right) - \ln\Gamma\left(\frac{\upsilon}{2}\right) - \frac{1}{2}\ln(\upsilon-2) - \frac{1}{2}\sum_{t=1}^{T}\left[\ln h_t + (\upsilon+1)\ln\left(1+\varepsilon_t^2 h_t^{-1}(\upsilon-2)^{-1}\right)\right]\right]$, where α, β and γ are constant parameters, υ denotes the degrees of freedom and T the number of observations. The sample period is from 10 November 1997 to 10 November 2002. Regressions use daily observations. Robust standard errors are presented in parentheses. * indicates a coefficient significantly different from zero at the 10%, ** at the 5% and *** at the 1% confidence level. The null hypothesis of the Langrange Multiplier (LM) test is that the standardised residuals are not correlated up to the fifth lag. The null hypothesis of the Kolmogorov-Smirnov (KS) test is that the standardised residuals follow a normal distribution. P-values for the LM and KS tests are presented in parentheses.

Country	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\gamma}$	$1/\upsilon$	Log likelihood	LM test	KS test
Brazil	0.000728 (4.52)	0.167*** (0.00036)	0.833*** (0.0113)	0.832*** (0.00922)	-1055.6	0.0051 (1.00)	0.122 (0.00)
Mexico	0.0287 (4.17)	0.18*** (0.00636)	0.735*** (0.0386)	0.735*** (0.0217)	-835.81	1.73 (0.885)	0.0758 (0.00)
Indonesia	0.00528 (3.26)	0.221*** (0.00415)	0.779*** (0.0208)	0.778*** (0.0214)	-2056.7	1.66 (0.894)	0.0984 (0.00)
South Korea	0.00725 (3.76)	0.227*** (0.00287)	0.773*** (0.0276)	0.773*** (0.0167)	-944.91	6.5 (0.261)	0.0871 (0.00)
Thailand	0.005 (4.05)	0.163*** (0.00036)	0.837*** (0.00951)	0.836*** (0.00811)	-727.57	0.00712 (1.00)	0.0724 (0.00)
South Africa	0.00434 (3.9)	0.125*** (0.00394)	0.875*** (0.0145)	0.875*** (0.00527)	-1285.8	8.87 (0.114)	0.0669 (0.00)

Table A.3. Nonlinear asymmetric GARCH model

The conditional variance equation of exchange rate returns is specified as follows: $h_t = \alpha + \beta h_{t-1} + \omega\left(\varepsilon_{t-1} + \gamma\sqrt{h_{t-1}}\right)^2$, where α, β, γ and ω are constant parameters. The sample period is from 10 November 1997 to 10 November 2002. Regressions use daily observations. Robust standard errors are presented in parentheses. * indicates a coefficient significantly different from zero at the 10%, ** at the 5% and *** at the 1% confidence level. The null hypothesis of the Langrange Multiplier (LM) test is that the standardised residuals are not correlated up to the fifth lag. The null hypothesis of the Kolmogorov-Smirnov (KS) test is that the standardised residuals follow a normal distribution. P-values for the LM and KS tests are presented in parentheses.

Country	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\gamma}$	$\hat{\omega}$	Log likelihood	LM test	KS test
Brazil	0.0223 (-0.37)	0.147*** (0.0232)	0.833*** (0.0331)	0.833*** (0.0119)	-1586.7	0.166 (0.999)	0.14 (0.00)
Mexico	0.00891 (0.43)	0.125*** (0.0383)	0.756*** (0.134)	0.756*** (0.0954)	-867.02	1.19 (0.946)	0.0716 (0.00)
Indonesia	0.388*** (0.00)	0.388 (2.89)	0.612 (0.819)	0.6118 (1.09)	-2468.5	0.772 (0.979)	0.135 (0.00)
South Korea	0.00357 (0.0633)	0.152*** (0.00171)	0.848*** (0.0205)	0.848*** (0.0135)	-1044	9.79 (0.0814)	0.0875 (0.00)
Thailand	0.00198 (0.0758)	0.092*** 0.00178)	0.904*** (0.0137)	0.904*** (0.00335)	-919.32	0.926 (0.968)	0.0731 (0.00)
South Africa	0.00322 (0.0891)	0.083*** (0.0022)	0.917*** (0.0114)	0.917*** (0.00241)	-1404.5	21.9 (0.0005)	0.0773 (0.00)

Table A.4. Exponential GARCH (EGARCH) model

The conditional variance equation of exchange rate returns is specified as follows: $\ln(h_t) = \alpha + \beta \ln(h_{t-1}) + \gamma \frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + \omega \left[\frac{|\varepsilon_{t-1}|}{\sqrt{h_{t-1}}} - \sqrt{\frac{2}{\pi}} \right]$, where α, β, γ and ω are constant parameters. The sample period is from 10 November 1997 to 10 November 2002. Regressions use daily observations. Robust standard errors are in parentheses. * indicates a coefficient significantly different from zero at the 10%, ** at the 5% and *** at the 1% confidence level. The null hypothesis of the Langrange Multiplier (LM) test is that the standardised residuals are not correlated up to the fifth lag. The null hypothesis of the Kolmogorov-Smirnov (KS) test is that the standardised residuals follow a normal distribution. P-values for the LM and KS tests are presented in parentheses.

Country	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\gamma}$	$\hat{\omega}$	Log likelihood	LM test	KS test
Brazil	-0.14 (0.949)	0.00618 (0.0538)	0.238*** (0.0495)	0.238*** (0.0552)	-1711.1	0.708 (0.983)	0.151 (0.00)
Mexico	-0.183 (0.956)	0.07*** (0.0158)	0.078*** (0.011)	0.077*** (0.0134)	-881.78	12.6 (0.0269)	0.0716 (0.00)
Indonesia	-0.168 (0.983)	0.043 (0.0534)	0.54*** (0.107)	0.540*** (0.199)	-2507.8	5.16 (0.396)	0.138 (0.00)
South Korea	-0.151 (1.04)	0.015*** (0.00)	0.146*** (0.00367)	0.146*** (0.00544)	-1093.9	60.0 (0.00)	0.098 (0.00)
Thailand	-0.0227 (1.05)	0.016*** (0.0058)	0.023*** (0.0052)	0.023*** (0.00564)	-1053.6	29.5 (0.00)	0.0868 (0.00)
South Africa	-0.0628 (1.03)	0.033*** (0.00166)	0.077*** (0.0113)	0.077*** (0.00647)	-1416.6	60.2 (0.00)	0.08 (0.00)

Table A.5. Predictability regression for options with 1-month time to maturity

The table shows the estimates of the regression: $\sigma_{t,T} = \alpha + \beta\hat{\sigma}_t^T + \varepsilon_{t,T}$, where $\sigma_{t,T}$ is the realised volatility over the lifetime of the option contract and $\hat{\sigma}_t^T$ is the Madan and Milne's (1994) estimated volatility or the forecasted annualised conditional standard deviation obtained from the GARCH(1,1) model. Regressions use daily observations. The estimation period is: 10 November 1997 – 10 October 2002. Standard errors have been corrected for the induced overlap and heteroskedasticity using Hansen's (1982) and White's (1980) methods. Asymptotic t–statistics are in parentheses. * denotes a coefficient significantly different from zero at the 10%, ** at the 5% and *** at the 1% confidence level.

Country	Implied			GARCH		
	$\hat{\alpha}$	$\hat{\beta}$	$\overline{R}^2$	$\hat{\alpha}$	$\hat{\beta}$	$\overline{R}^2$
Brazil	-1.69*** (-3.4)	0.932*** (23.3)	0.47	-2.45*** (-2.84)	0.893*** (17.3)	0.40
Mexico	2.81*** (8.81)	0.417*** (14.6)	0.22	-3.43** (-2.18)	1.36*** (7.12)	0.16
Indonesia	5.99*** (6.68)	0.568*** (19.5)	0.32	-3.42** (-2.3)	0.878*** (18.4)	0.45
South Korea	0.298 (1.5)	0.695*** (34.6)	0.64	1.33*** (5.04)	0.741*** (22.7)	0.58
Thailand	-0.649*** (-2.96)	0.772*** (35.7)	0.75	0.398 1.61)	0.891*** (29.4)	0.66
South Africa	1.08** (2.36)	0.766*** (21.2)	0.35	4.19*** (9.01)	0.618*** (15.9)	0.30

Table A.6. Predictability regression for options with 3-month time to maturity

The table shows the estimates of the regression: $\sigma_{t,T} = \alpha + \beta\hat{\sigma}_t^T + \varepsilon_{t,T}$, where $\sigma_{t,T}$ is the realised volatility over the lifetime of the option contract and $\hat{\sigma}_t^T$ is the Madan and Milne's (1994) estimated volatility or the forecasted annualised conditional standard deviation obtained from the GARCH(1,1) model. Regressions use daily observations. The estimation period is: 10 November 1997 – 10 August 2002. Standard errors have been corrected for the induced overlap and heteroskedasticity using Hansen's (1982) and White's (1980) methods. Asymptotic t–statistics are in parentheses. * denotes a coefficient significantly different from zero at the 10%, ** at the 5% and *** at the 1% confidence level.

Country	Implied			GARCH		
	$\hat{\alpha}$	$\hat{\beta}$	$\bar{R}^2$	$\hat{\alpha}$	$\hat{\beta}$	$\bar{R}^2$
Brazil	2.91*** (3.05)	0.614*** (9.97)	0.17	-1.81* (-1.7)	0.762*** (15.3)	0.19
Mexico	4.39*** (15.2)	0.279*** (12.5)	0.17	-2.24 (-0.946)	1.27*** (4.44)	0.03
Indonesia	29.4*** (40.9)	-0.001*** (-23.7)	0.00	-8.1*** (-3.67)	0.825*** (15.7)	0.34
South Korea	1.23*** (7.1)	0.555*** (40.4)	0.66	2.06*** (8.93)	0.586*** (24.2)	0.55
Thailand	-0.153 (-0.616)	0.68*** (29.2)	0.64	1.32*** (5.3)	0.737*** (26.3)	0.59
South Africa	3.23*** (6.49)	0.69*** (21.6)	0.21	5.73*** (12.8)	0.514*** (17.1)	0.25

Table A.7. Predictability regression for options with 6-month time to maturity

The table shows the estimates of the regression: $\sigma_{t,T} = \alpha + \beta\hat{\sigma}_t^T + \varepsilon_{t,T}$, where $\sigma_{t,T}$ is the realised volatility over the lifetime of the option contract and $\hat{\sigma}_t^T$ is the Madan and Milne's (1994) estimated volatility or the forecasted annualised conditional standard deviation obtained from the GARCH(1,1) model. Regressions use daily observations. The estimation period is: 10 November 1997 – 10 May 2002. Standard errors have been corrected for the induced overlap and heteroskedasticity using Hansen's (1982) and White's (1980) methods. Asymptotic t–statistics are in parentheses. * denotes a coefficient significantly different from zero at the 10%, ** at the 5% and *** at the 1% confidence level.

Country	Implied			GARCH		
	$\hat{\alpha}$	$\hat{\beta}$	$\overline{R}^2$	$\hat{\alpha}$	$\hat{\beta}$	$\overline{R}^2$
Brazil	2.26*** (4.23)	0.645*** (26.6)	0.27	4.89*** (4.79)	0.405*** (11.5)	0.04
Mexico	4.94*** (22.9)	0.23*** (17.2)	0.16	-9.45** (-2.15)	2.15*** (4.07)	0.04
Indonesia	28.8*** (50.1)	0.009 (0.714)	0.00	-11.4*** (-4.88)	0.71*** (16.3)	0.34
South Korea	2.01*** (10.1)	0.446*** (32.4)	0.59	1.74*** (7.34)	0.528*** (25.0)	0.57
Thailand	0.516** (2.47)	0.57*** (31.7)	0.60	1.63*** (7.78)	0.619*** (29.5)	0.60
South Africa	8.47*** (14.9)	0.373*** (10.3)	0.06	7.97*** (17.9)	0.377*** (14.9)	0.14