Tracking Brazilian Exchange Rate Volatility

Sandro Canesso de Andrade
Haas School of Business, UC-Berkeley

Eui Jung Chang
Banco Central do Brasil

Benjamin Miranda Tabak$^\nabla$
Banco Central do Brasil

February, 2003

Abstract

This paper examines the relation between dollar-real exchange rate volatility implied in option prices and subsequent realized volatility. It investigates whether implied volatilities contain information about volatility over the remaining life of the option which is not present in past returns. Using GMM estimation consistent with telescoping observations evidence suggests that implied volatilities give superior forecasts of realized volatility if compared to GARCH(p,q), and Moving Average predictors, and that econometric models forecasts do not provide significant incremental information to that contained in implied volatilities.

$^\nabla$ Corresponding author. Benjamin.tabak@bcb.gov.br.

1.Introduction

The huge literature on modeling and forecasting volatility in the past decades poses no questions on the relevance of the theme for financial academics and practitioners.

This literature has witnessed the introduction of different models for forecasting volatility of financial assets and performance comparisons of these models. One of the main questions is whether volatilities extracted from option prices give superior forecasts to econometric models such as GARCH models.

Volatilities implied in option prices are considered to be "the market's forecast" of future volatility during the option's remaining life. Recent research provides abundant evidence that implied volatilities contain information about subsequent realized volatility which is not captured by econometric models built upon time series of past returns.

Jorion (1995) examines options on currency futures traded at the Chicago Mercantile Exchange, and finds that their implied volatilities are upward-biased estimators of future volatility, but outperform standard time-series models in terms of informational content. In fact, he shows that the statistical models he tested offered no incremental informational to implied volatilities. The author performs tests for the period of 1985 through 1992 for the Deustche Mark, Swiss Franc and for 1986 to 1992 for the Japanese Yen. Use daily observations and used the Hansen-White procedure to correct standard errors for overlap and heteroskedasticity.

Xu and Taylor (1995) achieved similarly strong results for options on spot currencies traded at the Philadelphia Stock Exchange. The authors analyze four exchange rates (British Pound, Deutsche Mark, Japanese Yen and Swiss Franc) from 1985 to 1991. Nonetheless, when using high-frequency (five minutes) Taylor and Xu (1997) documented that econometric models offered incremental information to implied volatilities, and vice versa.

Fleming (1998) studies options on the S&P 100 equity index traded at the Chicago Board Options Exchange. His conclusions are very similar to Jorion's, i.e., implied volatilities are upward-biased predictors, but subsume information of standard statistical models. The author uses the sample period from October 1985 through April 1992, and excludes all observations that overlap the October 1987 crash. Christensen and Prabhala (1998) study the same market with a much longer data set, and also find that implied volatility is upward-biased and more informative than daily returns when forecasting volatility. Their sample span the period beginning in November 1983 through May 1995. Still considering S&P 100 index options, Blair et alli (2000) use high-frequency data to build time-series models and to measure realized volatility, and find evidence that the incremental information provided by statistical models is insignificant.

Amin and Ng (1997) focus on the Chicago Mercantile Exchange market for options on short term forward interest rates, known as eurodollar options. They show that implied volatilities contain more information about

future volatility than statistical time series models, but the explanatory power of implied volatilities is enhanced by the use of historical information.

Malz (2000) examines, among others, the Chicago Board of Trade market for options on futures of the 30-year T-bond, and concludes that historical volatility contains much less information about future volatility than implied volatility.

For commodities, Kroner et alli (1993) find that volatility forecasts combining implied volatility and GARCH-based estimates tend do perform better than each method by itself.

Vasilellis and Meade (1996) find that volatility forecasts implied by the options market are better forecasts than those based on equity market data. Nonetheless, they found that GARCH forecasts have significant incremental information. They use the period starting in March 1986 through September 1991 and twelve companies quoted in the London Stock Exchange.

Gwilyn and Buckle (1999) analyze the period from 21 June 1993 to 19 May 1995 for FTSE100 index options on a daily basis and found evidence that implied volatilities contain more information than historical volatilities. However, their evidence suggest that implied volatilities are biased. Gwilyn (2001) investigated the information content of implied volatilities in the same context (using FTSE100 options) for the period of 21 June 1993 to 19 May 1995 and found evidence that implied volatility (although biased)) contains more information than forecasts based on simple historical volatility and GARCH models.

To the best of our knowledge the only published paper that compares correlations implied from options prices with subsequent realized correlations is Campa and Chang (1998). They work with over-the-counter options on spot currencies, and obtain results in line with the related research on implied volatilities, i.e., historically based forecasts contribute no incremental information to implied correlations. They evaluate Dollar-Mark, Dollar-Yen and Mark-Yen options from January 1989 to May 1995 using daily data.

Summarizing, recent literature offers clear evidence that option prices embed information about future asset returns volatility that cannot be extracted from past returns. In this paper we examine whether this conclusion also apply to calls on the dollar-real spot exchange rate traded at the Brazilian Bolsa de Mercadorias & Futuros (BM&F), in the period of February 1999 to June 2002[1].

We use as our option pricing model the standard Garman-Kohlhagen (1983) extension of the Black-Scholes (1973) model. As historically-based models, we use the moving average standard deviation with a moving window of 20 days, and a GARCH (1,1) model.

It is worth noting that the main objective of this paper is not to test whether the Garman-Kohlhagen pricing model is adequate for the dollar-real call market, but to examine the ability of implied volatilities computed with this simple model to provide information about subsequent realized

---

[1] There was a major change of regime in January 1999, when Brazil moved from a quasi-fixed to a floating exchange rate. Before February 1999, the dollar-real options market was very illiquid, and restricted to deep out-of-the-money calls.

volatility and to test whether time series forecasts contain additional information to implied volatilities.

The remainder of the paper is organized as follows. Section 2 describes in detail the data we use in this study. Section 3 outlines the empirical methodology and presents results. Section 4 concludes the paper, and suggests directions for further research.

2. Methodology and Data Sampling

The primary data of this study are daily dollar-real calls close prices from 01 February 1999 to 28 June 2002, provided by BM&F. This period covers 744 trading days. The average daily notional value traded at this market in the period was US$ 270 million, what places it among the most important call markets for emerging markets currencies.

Dollar-real calls at BM&F are of the European style, and mature on the first business day of the corresponding month of expiration. Thus, our data span 41 expiration cycles. The first cycle is made of calls maturing on the first business day of March 1999, and the last one of calls maturing on the first business day of July 2002.

Our analysis also uses daily dollar-real futures and interest rate futures (named DI futures) adjustment prices[2] provided by BM&F. These futures contracts also mature on the first business day of the corresponding month.

---

[2] BM&F futures adjustment price, used for settlement of daily margins, is the average price of transactions done in the last 30 minutes of the day, weighted by the volume of each transaction. They are more reliable than close prices, since they cannot be eventually distorted by a single manipulative transaction.

We also utilize daily dollar-real spot prices provided by the Central Bank of Brazil (average price) and by Bloomberg (high and low prices).

2.1 Sampling procedure

In the period considered, liquidity at the BM&F dollar-real call market was highly concentrated on contracts maturing on the two nearer expiration dates. In general, liquidity of calls maturing on the second expiration date was very thin until around 12 business days prior to the first expiration date. Then, liquidity began to shift gradually from calls of the first expiration date to calls of the second expiration date.

Using the Garman-Kohlhagen pricing model, it can be shown that the price-sensitivity of options to volatility approaches zero as the option reaches its maturity. To limit the effect of option expirations, in our sampling procedure we aim at picking options which are the nearest, but with at least 10 business days, to maturity[3]. Unfortunately, on some occasions liquidity on second expiration calls is still too reduced at 10 days prior to the maturity of first expiration calls. In such situations we have to select calls with less than 10 but never less than 6 business day to maturity. The average range of each of the 41 expiration cycles considered is from 28 until 9 business days to expiration.

In each cycle, on every trading day, we select the closest-to-the-money[4] call, considering the adjustment price in the dollar-real futures market on

---

[3] Xu and Taylor (1995) and Fleming (1998) use options with at least 10 and 15 calendar days to expiration, respectively. Jorion (1995) selects options maturing in more than 3 business days.
[4] The closest-to-the money call for each expiration date is the one whose strike price is nearer to the futures price maturing on the same date.

that day. There are two reasons in choosing the closest-to-the-money option over the others. First, using Garman-Kohlhagen's model it can be shown that under usual circumstances the closest-to-the-money option for each expiration date is the one whose price is more sensitive to the volatility of the underlying asset.

The second reason for selecting the closest-to-the money option relates to the apparent inconsistency of recovering a volatility forecast from an option pricing model of the Black-Scholes family, which assume that volatility is known and constant. The point is that Feinstein (1989) demonstrated that for short-term at-the-money options, the Black-Scholes formula is almost linear in its volatility argument. Under the assumption that volatility is uncorrelated to returns, Feinstein showed that linearity turns Black-Scholes implied volatility into a virtually unbiased estimator of future volatility for those options, considering the class of stochastic volatility option pricing models introduced by Hull and White (1987), which assume that either investors are indifferent towards volatility risk or volatility risk is nonsystematic. Finally, it is worth mentioning that in the period considered the closest-to-the-money call on each trading day was always one of the most liquid ones.

2.2 Calculating implied volatilities

For every trading day, implied volatility is calculated from the close price of the call selected by our sampling procedure, which is the closest to the money and the more liquid one.

Measurement errors could be caused by the nonsynchronicity of prices in both spot and option markets. Thus, instead of using directly the spot market price we have computed implied volatilities using the price of the dollar-real future contract expiring in the same day of the option contract. We used the Garman-Kohlhagen model, applying the cost-of–carry arbitrage formula that links future to spot prices. Therefore, implied volatility $\sigma_{i,t}$ is computed by solving the equation below

$$C_t = \frac{1}{(1+r_t)^{T_t}} \left[F_t \text{N}(d_t) - E_t \text{N}\left(d_t - \sigma_{i,t}\sqrt{T_t}\right) \right] \tag{1}$$

where

$$d_t = \frac{\ln\left(F_t / E_t\right)}{\sigma_{i,t}\sqrt{T_t}} + \frac{1}{2}\sigma_{i,t}\sqrt{T_t} \tag{2}$$

C_t is the call option price, E_t corresponds to the exercise price, T_t denotes the number of days to maturity, r_t is the daily interest rate, F_t is the adjustment price of the dollar-real future expiring in T_t days, and $N(\cdot)$ is the standard normal distribution function. The daily interest rate is the one implied in the adjustment price of the short term interest rate future contract (called DI future) that expires in T_t days.

2.3 Time series predictors of future volatility

We wish to test the informational content of implied volatilities in comparison to time series models built upon past returns. Returns are computed using the average daily prices of the dollar-real spot exchange rate, and we consider two time series models as benchmarks in our tests.

One is a fixed volatility model, in which the volatility estimate is the sample standard deviation $MA(20)_t$, computed with a moving window including the last 20 returns.

$$MA(20)_t = \sqrt{\frac{1}{20}\sum_{k=0}^{19}\left(r_{t-k} - \overline{r}_t\right)^2} \qquad (3a)$$

where $\overline{r}_t = \frac{1}{20}\sum_{k=0}^{19} r_{t-k}$ and $r_t = \ln\left(\frac{S_t}{S_{t-1}}\right)$, where S_t is the average price of the dollar-real exchange rate on day t.

We also computed a variation, which is an exponentially weighted moving average $EWMA(20)_t$ with a 20 returns moving window.

$$EWMA(20)_t = \sqrt{\frac{1-\lambda}{1-\lambda^{20}}\sum_{k=0}^{19}\lambda^i\left(r_{t-k} - \overline{r}_t\right)^2} \qquad (3b)$$

where $\overline{r}_t$ and r_t are the same as used in expression (3a) and $\lambda = 0.94$ [5].

The other time series benchmark is a model of the GARCH family, introduced by Bollerslev (1986). The model is estimated from a sample of daily returns covering February 1999 to June 2002. The GARCH(p,q) model is:

$$r_t = \mu + \varepsilon_t, \ \ \varepsilon_t \sim N(0; h_t) \qquad (4a)$$

$$h_t = \alpha_0 + \sum_{i=1}^{p}\alpha_i\varepsilon_{t-i}^2 + \sum_{i=1}^{q}\beta_i h_{t-i} \qquad (4b)$$

In line with Hsieh(1989), we consider the GARCH(1,1) model to be a parsimonious representation that fits data relatively well, since results not reported here show that higher orders have nothing extra to offer. The

GARCH(1,1) model also serves as a benchmark for assessing the informational content of implied volatility vis-à-vis time-series models in Lamoureux e Lastrapes (1993), Jorion (1995), Fleming (1998) and Campa and Chang (1998).

Results of the GARCH(1,1) estimation for the period of February 1999 until June 2002 are on Table 1.

<div align="center">

Table 1: GARCH estimation

μ	α_0	α_1	β_1
0.148e-3	1.65e-6**	0.136831*	0.838974*
(0.213e-3)	(8.11e-7)	(0.031254)	(0.027531)

</div>

* rejection of the null with 99% confidence
** rejection of the null with 95% confidence

Results are in line with previous research, showing that the GARCH(1,1) model is highly significant. Thus, volatility is time-varying and shocks are persistent. Note that $(\alpha_1 + \beta_1)$ equals 0.976, therefore the process is stationary.

We consider the in-sample forecast for the average conditional volatility over the remaining life of the option, generated by the GARCH(1,1) model estimated for the whole period[6]. This forecast is denoted here as *GARCH(1,1)$_t$*. Heynen et alli (1994) demonstrated that:

$$GARCH_t^2 = \frac{\hat{\alpha}_0}{1 - \hat{\alpha}_1 - \hat{\beta}_1} + \left(\hat{h}_{t+1} - \frac{\hat{\alpha}_0}{1 - \hat{\alpha}_1 - \hat{\beta}_1} \right) \frac{1 - \left(\hat{\alpha}_1 + \hat{\beta}_1 \right)^{T_t}}{T_t \left(1 - \hat{\alpha}_1 - \hat{\beta}_1 \right)}$$

---

[5] This expression is widely used in risk management for volatility forecasting purposes.

It is important to emphasize that the possibility of using in-sample forecasts, i.e., the possibility to use *ex post* parameter estimates, represents an "unfair" advantage we give to the GARCH model over implied volatility[7].

2.4 Measuring realized volatility in the spot market over the option's remaining life

The size of interval in which we measure realized volatility ranges from 38 business days, the call with the longest time to maturity picked in our sampling procedure, to 6 business days, the one with the shortest time to maturity. Because volatility cannot be directly observed, we measure realized volatility in two alternative ways. First, we compute the sample standard deviation of returns SD_t, using average daily prices in the dollar-real spot market.

$$SD_t = \sqrt{\frac{1}{T_t} \sum_{k=1}^{T_t} \left(r_{t+k} - \overline{r_t} \right)^2} \tag{5}$$

where $\overline{r_t} = \frac{1}{T_t} \sum_{k=1}^{T_t} r_{t+k}$ and $r_t = \ln\left(\frac{S_t}{S_{t-1}}\right)$ and S_t is the average price of the dollar-real exchange rate on day t.

We acknowledge the fact that when the interval size is small, the measurement error of realized volatility could be substantial. Taylor and Xu (1997) and Andersen and Bollerslev (1998) show that measurement errors

---

[6] We also tested the one-day-ahead conditional volatility $\sqrt{h_{t+1}}$, and qualitative results are the same.

[7] We could not test out-of-sample forecasts by GARCH models because estimations that mix in a sample data from two fundamentally different exchange rate regimes (refer to footnote number 1) are not

in the estimation of realized volatility might distort conclusions about the informational content of volatility forecasts. These authors suggest the use of high-frequency intra-day data. Due to its unavailability, we aim to improve the quality of our measures of realized volatility by using the Parkinson (1980) estimator, which improves the efficiency of realized volatility measures by using information embedded in daily high and low prices[8]. The Parkinson estimator is:

$$PK_t = \sqrt{\frac{1}{4\ln(2)}\frac{1}{T_t}\sum_{k=1}^{T_t}\left(H_{t+k} - L_{t+k}\right)^2}$$ (6)

where H_t and L_t are respectively the natural logarithm of the highest and the lowest price of the dollar-real spot exchange rate on day t.

Garman and Klass (1980) proved this is an unbiased estimator of volatility, which is around five times more efficient than the sample standard deviation[9].

3. Empirical Results

3.1 Implied volatility versus realized volatility

correctly specified, thus in the first months of 1999 there are not enough observations to allow estimation of GARCH models.

[8] The Parkinson (1980) estimator assumes that returns follow a continuous time Geometric Brownian motion with zero drift. Although this is certainly not true for the period as whole, as evidenced by the GARCH estimation, we assume that volatility in each of the intervals in which we measure realized volatility is constant.

[9] In fact, Garman and Klass (1980) point out that the Parkinson estimator would be downward biased in case of infrequent trading. We assume that the dollar-real spot rate market is not influenced by infrequent trading.

Following Fleming (1998), we evaluate the predictivity ability of implied volatilities by regressing realized volatility (SD_t or PK_t) on implied volatility ($\sigma_{i,t}$)[10]. We estimate α and β in the moment vector.

$$g_T\left(\alpha_1,\beta_1\right)=\frac{1}{NK}\sum_{t=1}^{NK}\left(SD_t-\alpha_1-\beta_1\sigma_{i,t}\right)Z_t \tag{7a}$$

$$g_T\left(\alpha_2,\beta_2\right)=\frac{1}{NK}\sum_{t=1}^{NK}\left(PK_t-\alpha_2-\beta_2\sigma_{i,t}\right)Z_t \tag{7b}$$

where NK is the number of observations, and Z_t represents a vector of instruments.

The series are specified in levels and each series has a high serial correlation. The main source of serial correlation is the fact that data overlap substantially. This is due to the fact that, in order to gain maximum efficiency within a limited sample period, we sample data daily (774 days), while forecasts intervals are determined by monthly option expiration cycles (41 cycles).

Table 2. Descriptive Statistics

	Time series $GARCH(1,1)_t$	Time series $MA(20)_t$	Time series $EWMA(20)_t$	Implied $\sigma_{i,t}$	Realized SD_t	Realized PK_t
Mean	0.1364	0.1338	0.1293	0.1514	0.1247	0.1160
Median	0.1186	0.1150	0.1094	0.1389	0.1112	0.1091
Max.	0.7030	0.6464	0.6791	0.7569	0.4527	0.3707
Min.	0.0684	0.0360	0.0367	0.0302	0.0205	0.0291
St.Dev.	0.0692	0.0914	0.0872	0.0835	0.0654	0.0608
Skewn.	3.5937	2.7503	2.7693	2.6049	1.1885	1.2779
Kurtosis	20.7834	12.7558	13.3955	14.5916	4.6307	4.9806

If volatility series possess a unit root, regressions specified as above are spurious. Therefore, we need to test the non-stationarity of the series before

---

[10] This approach is also taken by Canina and Figlewski (1993), Jorion (1995), Amin and Ng (1997),

performing regressions. Using both Dickey-Fuller (1979) and Phillips-Perron (1988) tests we reject the unit root hypothesis for all series, as evidenced by Table 3[11].

Table 3. Unit root tests

	ADF test Statistic	Phillips-Perron test statistic
$\sigma_{i,t}$	- 5.77 *	- 11.23 *
$MA(20)_t$	- 6.67 *	- 4.03 *
$GARCH(1,1)_t$	- 6.33 *	- 6.34 *
$EWMA(20)_t$	-6.93 *	-3.95 *
SD_t	- 4.97 *	- 5.13 *
PK_t	- 2.76 ***	- 3.07 **

* Reject the null of a unit root with 99% confidence.
** Reject the null of a unit root with 95% confidence.
*** Reject the null of a unit root with 90% confidence.

If a volatility forecast contains information about subsequent realized volatility, then the slope should be statistically distinguishable from zero. If the forecast is unbiased, then the intercept should be zero and the slope should be one. Due to the possibility of measurement errors in independent variables, Scott (1992) and Fleming (1998) use GMM estimation instead of GLS, in order to deal with the error-in-variables problem. We performed GMM estimation, using lagged independent variables as instruments. The informational content can be gauged by the coefficient of determination R^2 [12].

Christensen and Prabhala (1998), Campa and Chang (1998) and Blair et alli (2000).

[11] Scott (1992) and Fleming (1998) point out that even when non-stationarity is rejected, the spurious regression problem may still affect inference based on small samples. They tested the following alternative specification that is free from the spurious regression problem:

$$SD_t - \sigma_{i,t-1} = \alpha_1 + \beta_1 \left(\sigma_{i,t} - \sigma_{i,t-1} \right) + \varepsilon_t \text{ or } PK_t - \sigma_{i,t-1} = \alpha_2 + \beta_2 \left(\sigma_{i,t} - \sigma_{i,t-1} \right) + \varepsilon_t .$$

We also performed regressions, not reported in this study, with this specification, and verified that qualitative results are the same as those of the regression in levels reported here.

[12] The R^2 provides a direct assessment of the variability in realized volatility that is explained by the estimates. It is considered a simple gauge of the degree of predictability in the volatility process, and hence of the potential economic significance of the volatility forecasts.

Data overlap induces residual autocorrelation, as evidenced by low Durbin-Watson statistics in all regressions (below 0.5, not reported). This could yield inefficient slope estimates and spurious explanatory power. Following Jorion (1995), Amin and Ng (1997) and Campa and Chang (1998), we correct this using asymptotic standard errors computed from an heterokedasticity and autocorrelation consistent covariance matrix. In this paper we use Fleming's (1998) covariance matrix.

A consistent estimator for the covariance matrix is given by

$$\Omega_T = \frac{1}{NK} \sum_{t=1}^{NK} \left(\sum_{l=-K+1}^{K-1} \phi_{t,t-l} \left(\hat{\varepsilon}_t, \hat{\varepsilon}\sigma_{i,t} \right)^{'} \left(\hat{\varepsilon}_{t-l}, \hat{\varepsilon}\sigma_{i,t-l} \right) \right) \tag{8}$$

where $\hat{\varepsilon}$ is the GMM residual from equations (7a) and (7b) and $\phi_{t,t-l}$ is a dummy variable equal to one when contracts months represented by observations t and t-l overlap and zero otherwise.

Results for the regressions of realized volatility, as measured by the standard deviation (SD_t) or by the Parkinson estimator (PK_t), on implied volatlity are shown on Table 4. Wald tests for unbiasedness ($\alpha = 0$ and $\beta = 1$) are reported.

Table 4. Regressions of realized volatility on implied volatility

Dependent Variable	Intercept	Slope	Wald Test	Adjusted R²
SD_t	0.0340*	0.5998*	31.66*	42.01 %
	(0.0095)	(0.0651)	63.32	
PK_t	0.0251*	0.6019*	61.84*	44.54 %
	(0.0074)	(0.0515)	123.68	

* rejection of the null with 99% confidence.
Fleming (1998) corrected Standard errors in parenthesis

T-statistics on the coefficients of implied volatilities in both regressions are very high, 9 and 12, strongly rejecting the null hypothesis that implied volatilities carry no information about future volatility. Wald tests for unbiasedness also reject the null at the 99% level in both regressions, providing evidence that implied volatilities are biased predictors of future volatility.

Figure 1 provides enough evidence that the direction of the bias is upward, i.e., implied volatilities tend to overstate future volatility. This finding is consistent with Jorion (1995), Fleming (1998) and Bates (2000). Table 2 show that in the period considered implied volatility overstated realized volatility by an average of 5 percentage points on an annualized basis.

Slope coefficients less than one suggest that implied volatility is too volatile: on average a change in implied volatility does not fully translate into changes in realized volatility, but need to be scaled down.

In line with our expectation, and with Andersen and Bollerslev (1998), the R^2 of regressions suggest that the Parkinson estimator is more adequate in measuring realized volatility than the sample standard deviation of returns.

3.2 Implied volatility versus time series volatility forecasts

In the previous item we found that implied volatility is an upward-biased estimator that does carry information about future volatility. At this point

we want to compare the informational content of implied volatility vis-à-vis time series models.

To begin with, we perform regressions of realized volatility (SD_t or PK_t) on time-series volatility forecasts ($MA(20)_t$, $EWMA(20)_t$ and $GARCH(1,1)_t$)[13] and compare adjusted R^2's with the regressions using implied volatility.

$$g_T\left(\alpha_1, \beta_1\right) = \frac{1}{NK} \sum_{t=1}^{NK} \left(SD_t - \alpha_1 - \beta_1 time\_series\_forecasts_t\right) Z_t \qquad (9a)$$

$$g_T\left(\alpha_1, \beta_1\right) = \frac{1}{NK} \sum_{t=1}^{NK} \left(PK_t - \alpha_1 - \beta_1 time\_series\_forecasts_t\right) Z_t \qquad (9b)$$

To evaluate the incremental information implied volatility offers over historically-based forecasts, we also regress realized volatility on implied volatility and on time-series forecasts at the same time, again following Day and Lewis (1992) and Fleming (1998)[14].

$$g_T\left(\alpha_1, \beta_1\right) = \frac{1}{NK} \sum_{t=1}^{NK} \left(SD_t - \alpha_1 - \beta_1 time\_series\_forecasts_t - \beta_2 \sigma_{i,t}\right) Z_t \qquad (10a)$$

$$g_T\left(\alpha_1, \beta_1\right) = \frac{1}{NK} \sum_{t=1}^{NK} \left(PK_t - \alpha_1 - \beta_1 time\_series\_forecasts_t - \beta_2 \sigma_{i,t}\right) Z_t \qquad (10b)$$

In this kind of "encompassing regression", if an independent variable contains no useful information regarding the evolution of the dependent variable, we would expect the coefficient of that independent variable to be insignificantly different from zero.

Results of the regressions using the standard deviation as a measure of realized volatility are on Table 5a, and using the Parkinson estimator are on

---

[13] Table 3 shows that we can reject the null hypothesis of non-stationarity for time series forecasts.
[14] This approach of comparing multiple forecasts, often called "encompassing regression", is discussed in Fair and Shiller (1990), and is also used by Lamoureux and Lastrapes (1993), Jorion (1995), Christensen and Prabhala (1998) and Campa and Chang (1998).

Table 5b. Results of the regressions of Table 4 are repeated for expositional convenience.

Table 5a. Encompassing regressions using standard deviation realized volatility (SD_t)

Intercept	$\sigma_{i,t}$	$GARCH(1,1)_t$	$EWMA(20)_t$	$MA(20)_t$	Adjusted R^2
0.0340*	0.5998*				42.01 %
(0.0095)	(0.0651)				
0.0477*		0.5643*			38.11 %
(0.0120)		(0.0906)			
0.0676*			0.4407*		37.44 %
(0.0080)			(0.0614)		
0.0687*				0.4179*	36.13 %
0.0076				0.0558	
0.0356*	0.8920*	-0.3356			31.80 %
(0.0079)	(0.2231)	(0.2471)			
0.0223	0.9246*		-0.2880		31.52 %
(0.0121)	(0.2188)		(0.1852)		
0.0237	0.8842*			-0.2434	33.62 %
0.0116	0.1866			0.1484	

* Reject the null with 99% confidence
** Reject the null with 95% confidence
Fleming (1998) corrected Standard errors in parenthesis

Table 5b. Encompassing regressions using Parkinson realized volatility (PK_t).

Intercept	$\sigma_{i,t}$	$GARCH(1,1)_t$	$EWMA(20)_t$	$MA(20)_t$	Adjusted R^2
0.0251*	0.6019*				44.54 %
(0.0074)	(0.0515)				
0.0391*		0.5637*			40.39 %
(0.0096)		(0.0708)			
0.0594*			0.4368*		40.31 %
(0.0064)			(0.0443)		
0.0610*				0.4100*	39.00 %
0.0060				(0.0382)	
0.0268*	0.9138*	-0.3582			31.94 %
(0.0062)	(0.2046)	(0.2471)			
0.0112	0.9837*		-0.3385		29.02 %
(0.0124)	(0.2327)		(0.1852)		
0.0115	0.9742*			-0.3186**	29.83 %
(0.0120)	(0.2066)			(0.1484)	

* Reject the null with 99% confidence
** Reject the null with 95% confidence.
Fleming (1998) corrected Standard errors in parenthesis

The R^2 of the regressions using only one independent variable indicate that implied volatility contains more information about future volatility than historically-based forecasts, considering both measures of realized volatility. When realized volatility is measured by the standard deviation (SD_t), the R^2 of the regression on implied volatility is 42.01%, while on time series forecasts is only 37.23% on average. When the Parkinson estimator (PK_t) is used, implied volatility explains 44.54% of the variation of realized volatility, while time series forecasts explain on average only 39.90%.

When we regress realized volatility on more than one independent variable, results clearly show that implied volatility contains information about future volatility which is not captured by statistical models built upon past returns, since its coefficient is always significantly different from zero. As to incremental information offered by time series forecasts over implied volatility, the results are conclusive. If we use the standard variation (SD_t) to measure realized volatility, Table 5a shows that implied volatility is the only significant variable, subsuming historically-based forecasts. However, when the Parkinson estimator (PK_t) is used, Table 5b shows that the coefficients of historically-based forecasts are significantly different from zero only for the $MA(20)_t$, which could suggest that time series forecasts would offer some incremental information to implied volatility. Nonetheless, the additional explanatory power, measured by the increment in R^2 is negative. Thus we can conclude that implied volatility forecasts of realized volatility subsume other time-series forecasts.

3.3. The role of the forecasting horizon

An interesting issue would be to test whether implied volatilities perform better than $GARCH(1,1)_t$, $EWMA(20)_t$ and $MA(20)_t$ models for different forecasting horizons. We have divided our sample by grouping forecasts with a fixed time from expiration. We have built forecasts from 10 to 27 days from expiration and estimated (9a) and (9b) using time series benchmarks and implied volatilities. In this case, we use GMM with the Newey and West (1987) correction for autocorrelation and heteroskedasticity as we have a moving interval of fixed length.

In Figure 1 we present results for the Adjusted R^2 for regressions using implied, $GARCH(1,1)_t$ and $MA(20)_t$ and compare the information content of these forecasts for different forecasting horizons. As we can see implied volatilities seem to have a better performance independent of the forecasting horizon.

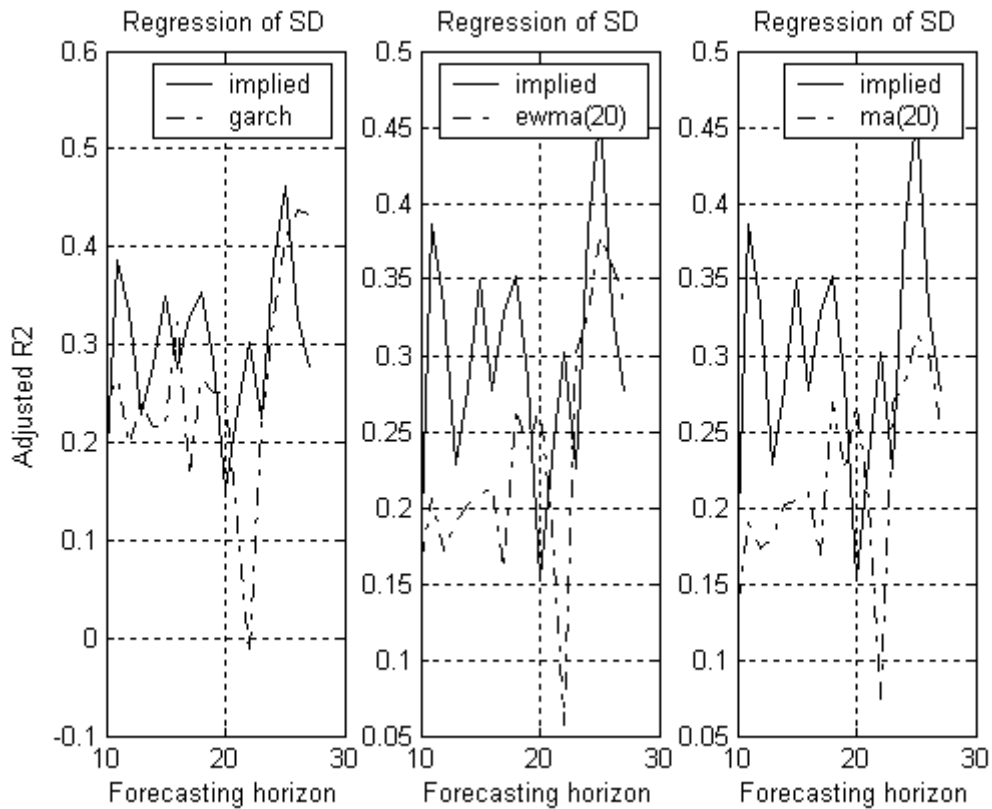Figure 1. R^2 for different forecasting horizons (standard deviation)

Figure 2 presents basically the same thing but uses the Parkinson estimator for realized volatility. Again, implied volatilities seem to possess greater information content than its econometric counterparts.

Figure 2. R² for different forecasting horizons (Parkinson)

Our results suggest that implied volatilities contain information that is not present in models built upon past returns and this is true for different forecasting horizons.

4. Conclusions

This article has presented evidence that implied volatilities contain information that is not present in past returns for the Brazilian exchange rate in the period after the devaluation of the Real in early 1999 through June 2002. Nonetheless, empirical results suggest that implied volatilities are upward biased as found in other studies. This bias may be due to systematic measurement errors, to market imperfections or to a model

misspecification among others. We have used closest to the money options to minimize adverse effects that could be caused by misspecification of the option pricing model[15]. We believe that a positive volatility risk premium is the main reason driving the upward-bias of implied volatilities. In emerging markets such as Brazil higher exchange rate volatility periods are almost always associated with worsened economic or political fundamentals, which also cause a decrease in total market wealth. Thus, agents who are short volatility are in fact selling insurance to the rest of the market, and therefore must be compensated in equilibrium. A few stylized facts corroborates this interpretation: in Brazil exchange rate volatility tends to be negatively correlated with stock market returns and with the level of exchange rate [16], and implied volatilities recovered from exchange-rate put options are also upward-biased. Nonetheless, a more thorough investigation into these issues is left for future research.

Time series forecasts as given by a $GARCH(1,1)_t$, an $EWMA(20)_t$ and a $MA(20)_t$ for subsequent realized volatilities do not add any significant information for forecasts based on implied volatilities. The adjusted R^2 statistics are the highest for implied volatility if compared to its econometric counterparts. Results are robust to the use of a correction of the covariance matrix that takes into account the telescoping nature of observations used in the study. Additionally, implied volatilities seem to perform better for most forecasting horizons in terms of information content. Empirical results are robust when we define realised volatility in terms of daily squared

---

[15] See Backus et al. (1997) and Feinstein (1989).

returns or to the use of the Parkinson volatility estimator, which uses the

high and low prices.

References

Amin, K. and V. Ng, 1997. Inferring future volatility from the information in implied volatility in eurodollar options: a new approach. Review of Financial Studies, 10, 333-367.

Andersen, T.G. and T. Bollerslev, 1998. Answering the skeptics: yes, standard volatility models do provide accurate forecasts. International Economic Review 39, 885-905.

Backus, D., S. Foresi, K. Li and L. Wu, 1997. Accounting for biases in Black-Scholes. Working Paper, New York University.

Bates, D.S., 2000. Post-´87 crash fears in the S&P 500 futures options market. Journal of Econometrics 94, 181-238.

Blair, B.V., S.H. Poon and S.J. Taylor, 2000. Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high frequency index returns. Lancaster University, mimeo.

Black, F. and M. Scholes, 1973. The valuation of options and corporate liabilities. Journal of Political Economy 81, 637-654.

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 31, 307-327.

Campa, J.M. and P.H.K. Chang, 1998. The forecasting ability of correlations implied in foreign exchange options. Journal of International Money and Finance 17, 855-880.

Canina, L. and S. Figkewski. 1993. The informational content of implied volatility. Review of Financial Studies 6, 659-681.

Christensen, B.J.and N.R. Prabhala, 1998. The relation between implied and realized volatility. Journal of Financial Economics 50, 125-150.

Day, T. and C. Lewis, 1992. Stock market volatility and the information content of stock index options. Journal of Econometrics 52, 267-287.

Dickey, D.A. and W.A. Fuller, 1979. Distribution of the estimators for autoregressive time series with a unit root. Journal of the American Statistical Association 74, 427–431.

Fair, R.C. and R.J. Shiller, 1990. Comparing information in forecasts from econometric models. American Economic Review 80, 375-389.

---

[16] Recall that the Brazilian market as whole is at least partially segmented from the rest of the world, and is on aggregate short in US dollars versus Brazilian reais.

Feinstein, S. 1989. The Black-Scholes formula is nearly linear in σ for at-the-money options; therefore implied volatilities from at-the-money options are virtually unbiased. Working paper, Federal Reserve Bank of Atlanta.

Fleming, J., 1998. The quality of market volatility forecasts implied by S&P 100 index option prices. Journal of Empirical Finance 5, 317-345.

Garman, M.B. and M.J. Klass, 1980. On the estimation of security price volatilities from historical price. Journal of Business, 53, 67-78.

Garman, M.B. and S. Kohlhagen, 1983. Foreign currency option values. Journal of International Money and Finance 2, 231-238.

Gwilym, O. A. and Buckle, M. 1999. Volatility Forecasting in the framework of the option expiration cycle. The European Journal of Finance 5, 73-94.

Gwilym, O. A. 2001. Forecasting volatility for options pricing for the U.K. stock market. Journal of Financial Management and Analysis 14, 55-62.

Heynen, R.C., A. Kemna and T. Vorst, 1994. Analysis of the term structure of implied volatilities. Journal of Financial and Quantitative Analysis 29, 31-56.

Hsieh, D., 1989. Modeling heteroskedasticity in daily foreign exchange rates. Journal of Business and Economic Statistics 7, 307-317.

Hull, J.C. and A.White, 1987. The pricing of options on assets with stochastic volatilities. Journal of Finance 42, 281-300.

Jorion, P., 1995. Predicting volatility in the foreign exchange market. Journal of Finance 50, 507-528.

Kroner, K.R., K.P. Kneafsey and S. Claessens, 1993. Forecasting volatility in commodity markets. World Bank Policy Research Working Paper No. 1226.

Lamoureux, C.G. and W.D. Lastrapes, 1993. Forecasting stock-return variance: towards an understanding of stochastic implied volatilities. Review of Financial Studies 6, 293-326.

Malz, A., 2000. Do implied volatilities provide early warning of market stress ? RiskMetrics Journal 1, 41-60.

Newey, W. and K. West, 1987. A simple positive semi-definite, heteroscedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703-708.

Parkinson, M. 1980. The extreme value method for estimating the variance of the rate of return. Journal of Business 53, 61-65.

Phillips, P.B. and P. Perron, 1988. Testing for a unit root in time series regression. Biometrika 75, 335–346.

Scott, L.O., 1992. The information content of prices is derivative security markets. IMF Staff Papers 39, 596-625.

Taylor, S.J. and X. Xu, 1997. The incremental volatility information in one million foreign exchange quotations. Journal of Empirical Finance 4, 317-340.

Vasilellis, G.A. and Meade, N., 1996. Forecasting volatility for portfolio selection. Journal of Business Finance & Accounting 23, 125-143.

Xu, X. and S.J. Taylor, 1995. Conditional volatility and the informational efficiency of the PHLX currency options markets. Journal of Banking and Finance 19, 803-821.