

Is it worth tracking dollar/real implied volatility?*

Sandro Canesso de Andrade[§]
Benjamin Miranda Tabak[□]

RESUMO

Este artigo examina a relação entre a volatilidade implícita em opções cambiais dólar/real e a volatilidade realizada, do período de fevereiro de 1999 a junho de 2000. Os resultados encontrados estão em linha com a literatura recente, sugerindo que a volatilidade implícita obtida por meio de um modelo simples de precificação de opções, embora seja um estimador viesado para cima da volatilidade futura, traz informação sobre a volatilidade realizada que não está presente nos retornos passados. Os resultados são robustos a dois modelos de séries temporais alternativos que exploram a informação implícita nos retornos, um modelo de volatilidade fixa e um GARCH(1,1), mesmo permitindo previsões in-sample do modelo GARCH(1,1). Os resultados também são robustos a duas maneiras de se calcular volatilidade realizada.

Palavras-chave: opções cambiais, volatilidade implícita, previsão, informação.

ABSTRACT

In this paper we examine the relation between dollar-real exchange rate volatility implied in option prices and subsequent realized volatility, in the period of February 1999 to February 2001. Our results are in line with recent literature, suggesting that the implied volatility obtained from a simple option pricing model, although an upward-biased estimator of future volatility, does provide information about volatility over the remaining life of the option which is not present in past returns. Results are robust to the choice of two alternative time series models to explore information embedded in returns, a fixed volatility and a GARCH(1,1) model, even allowing for in-sample forecasts by the GARCH(1,1) model. Results are also robust to the choice of measuring realized volatility in two alternative ways.

Key words: currency options, implied volatility, forecast, information.

JEL classification: G13, G14, C53.

* The authors thank participants at the XXII Meeting of the Brazilian Econometric Society, at the 28th meeting of the Academy of Economics and Finance and two anonymous referees for helpful comments and Verdi Monteiro for providing the data used in the paper.

The views expressed herein do not necessarily reflect the views of the Central Bank of Brazil.

§ Central Bank of Brazil.

□ Central Bank of Brazil and Catholic University of Brasilia. Email: benjamin.tabak@cbcb.gov.br

Recebido em março de 2001. Aceito em agosto de 2001.

1 Introduction

The ability to forecast second moments is useful in many applications, such as financial risk control, asset and liability management, and the pricing and hedging of derivative securities.

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, extracted by the use of relatively simple option pricing models, contain information about subsequent realized volatility which is not captured by statistical models built upon past returns. The conclusions are similar for many different markets, as well as various statistical techniques.

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 information to implied volatilities. Xu and Taylor (1995) achieved similarly strong results for options on spot currencies traded at the Philadelphia Stock Exchange. However, their conclusions change when they build statistical models and measure realized volatility using high-frequency (five minutes) returns. In this case, Taylor and Xu (1997) document that statistical 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: implied volatilities are upward-biased predictors, but subsume information of standard statistical models. 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. Still considering S&P 100 index options, Blair *et alii* (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.

In the case of commodities, Kroner *et alii* (1995) find that volatility forecasts combining implied volatility and GARCH-based estimates tend to perform better than each method by itself.

To 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: historically based forecasts contribute no incremental information to implied correlations.

In short: 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 February 2001.¹ Our option pricing model is 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.

At this point it is important to stress that the main objective of this article 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 volatility.

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 Data

The primary data of this study are daily dollar-real calls close prices from 02 February 1999 to 13 February 2001, provided by BM&F. This period covers 499 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.

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.

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 24 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 March 2001.

Our analysis also uses daily dollar-real futures and interest rate futures (named DI futures) adjustment prices² provided by BM&F. These futures contracts also mature on the first business day of the corresponding month. 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 24 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-money⁴ call, considering the adjustment price in the dollar-real futures market on that day. There are two reasons in

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 distorted by a single manipulative transaction.

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.

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 Computing implied volatilities

On every trading day, one implied volatility is calculated from the close price of the call selected by our sampling procedure.

In order to avoid measurement errors caused by the nonsynchronicity of prices in the spot and option markets, we compute implied volatilities using the price of the dollar-real future contract expiring in the same day of the option contract, instead of using directly the spot market price. Thus, we substitute the spot price for the future price in the Garman-Kohlhagen model, applying the cost-of-carry arbitrage formula that links future to spot prices. Therefore, from each observed call price C_p , implied volatility $\sigma_{i,t}$ is computed by numerically solving the equation

$$C_t = \frac{1}{(1+r_t)^{T_t}} \left[F_t N(d_t) - E_t N(d_t - \sigma_{i,t} \sqrt{T_t}) \right], \text{ where } d_t = \frac{\ln\left(\frac{F_t}{E_t}\right)}{\sigma_{i,t} \sqrt{T_t}} + \frac{1}{2} \sigma_{i,t} \sqrt{T_t}$$

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 benchmarks

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 $MA(20)_p$, computed with a moving window including the last 20 returns.

$MA(20)_t = \sqrt{\frac{1}{20} \sum_{k=0}^{19} (r_{t-k} - \bar{r}_t)^2}$, where $\bar{r}_t = \frac{1}{20} \sum_{k=0}^{19} r_{t-k}$, $r_t = \ln\left(\frac{S_t}{S_{t-1}}\right)$, and S_t is the average price of the dollar-real exchange rate on day t .

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 February 2001. The GARCH(p,q) model is:

$$r_{t+1} = \mu + \varepsilon_{t+1}, \quad \varepsilon_{t+1} \sim N(0; h_{t+1}), \quad h_{t+1} = \alpha_0 + \sum_{i=1}^p \alpha_i \varepsilon_{t+1-i}^2 + \sum_{i=1}^q \beta_i h_{t+1-i}$$

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 February 2001 are on Table I.

Table I
GARCH Estimation

μ	α_0	α_1	β_1	$h(0)$
175019E-03	123836E-05*	.099535*	.865716*	192169E-02**
.253975E-03	.366612E-06	.016217	.019153	.908642E-03

* 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 0.96, 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.⁵ This forecast is denoted here as $GARCH_t$. Heynen *et alii* (1994) demonstrated that:

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

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.⁶

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

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

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 35 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.

$SD_t = \sqrt{\frac{1}{T_t} \sum_{k=1}^{T_t} (r_{t+k} - \bar{r}_t)^2}$, where $\bar{r}_t = \frac{1}{T_t} \sum_{k=1}^{T_t} r_{t+k}$, $r_t = \ln\left(\frac{S_t}{S_{t-1}}\right)$, 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 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.⁷ The Parkinson estimator is:

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

Parkinson (1980) proved this is an unbiased estimator of volatility, which is around five times more efficient than the sample standard deviation.⁸

7 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.

8 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.

3 Empirical results


Descriptive statistics for time series volatilities, implied volatilities and realized volatilities are shown on Table II. All variables are in percent per annum, i.e., annualized by a factor of $\sqrt{252}$

Table II
Descriptive Statistics

	Time series $GARCH_t$	Time series $MA(20)_t$	Implied $\sigma_{i,t}$	Realized SD_t	Realized PK_t
Mean	0.1173	0.1254	0.1361	0.1045	0.0947
Median	0.0947	0.0912	0.1125	0.0892	0.0835
Max.	0.6520	0.6642	0.7569	0.4528	0.3814
Min.	0.0667	0.0263	0.0302	0.0205	0.0246
St. Dev.	0.0762	0.1055	0.0893	0.0609	0.0546
Skew.	3.7727	2.8557	3.2009	1.8211	2.2507
Kurtosis	19.605	11.830	17.244	7.3920	9.6386

Figure 1 displays, in percent per annum, the time variation of implied volatility and realized volatility as measured by the Parkinson estimator (PK_t). It is evident from both series that volatility is time varying. Figure 1 seems to suggest that implied volatility systematically overstate subsequent realized volatility.

Figure I
Implied and Realized Volatility (% annum)



3.1 Implied volatility versus realized volatility

Following Day and Lewis (1992), we evaluate the predictivity ability of implied volatilities by regressing realized volatility (SD_t or PK_t) on implied volatility ($\sigma_{i,t}$).⁹

$$realized_t = \alpha + \beta implied_t + \varepsilon_t$$

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 (321 days), while forecasts intervals are determined by monthly option expiration cycles (17 cycles).

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 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 III.¹⁰

Table III
Unit Root Tests

	ADF Test Statistic	Phillips-Perron Test Statistic
$\sigma_{i,t}$	-6.29*	-7.67*
$MA(20)_t$	-6.94*	-3.57*
$GARCH_t$	-7.16*	-9.91*
SD_t	-3.81*	-5.84*
PK_t	-4.55*	-5.70*

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

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

9 This approach is also taken by Canina and Figlewski (1993), Jorion (1995), Amin and Ng (1997), Christensen and Prabhala (1998), Campa and Chang (1998) and Blair *et alii* (2000).

10 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: $realized_t - implied_{t-1} = \alpha + \beta (implied_t - implied_{t-1}) + \varepsilon_t$.

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.

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. The informational content can be gauged by the coefficient of determination R^2 ¹¹

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 Newey and West (1987) covariance matrix.

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.

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

Table IV
Regressions of Realized Volatility on Implied Volatility

Dependent Variable	Intercept	Slope	Wald Test	Adjusted R ²
SD_t	0.0305* (0.0082)	0.5432* (0.0644)	50.23*	52.43 %
PK_t	0.0160* (0.0056)	0.5788* (0.0433)	94.37*	68.34 %

* rejection of the null with 99% confidence.

Asymptotic Newey-West (1987) standard errors in parenthesis.

11 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.

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 II shows that in the period considered implied volatility overstated realized volatility by an average of 3 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 needs to be scaled down.

In line with our expectation, and with Andersen and Bollerslev (1998), the R^2 of the 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 not 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 ($MA(20)_t$ and $GARCH_t$)¹² and compare adjusted R^2 's with the regressions using implied volatility.

$$realized_t = \alpha + \beta \text{ time_series_forecast}_t + \varepsilon_t$$

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).¹³

12 Table III shows that we can reject the null hypothesis of non-stationarity for time series forecasts.

13 This approach of comparing multiple forecasts, often called "encompassing regression", is discussed in Fair and Shiller (1990), and also used by Lamoureux and Lastrapes (1993), Jorion (1995), Christensen and Prabhala (1998) and Campa and Chang (1998).

$$realized_t = \alpha + \beta_1 implied_t + \beta_2 time_series_forecast_t + \varepsilon_t$$

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 Va, and using the Parkinson estimator are on Table Vb. Results of the regressions of Table IV are repeated for expositional convenience.

Table Va
Encompassing Regressions Using Standard Deviation Realized Volatility (SDt)

Intercept	$\sigma_{i,t}$	$GARCH_t$	$MA(20)_t$	Adjusted R ²
0.0305* (0.0082)	0.5432* (0.0644)			52.43 %
0.0424* (0.0084)		0.5277* (0.0705)		44.84 %
0.0483* (0.0070)			0.4450* (0.0591)	41.32 %
0.0285* (0.0083)	1.1879* (0.3942)	-0.7315*** (0.4156)		33.09 %
0.0263** (0.0115)	0.6706* (0.2076)		-0.1043 (0.1387)	50.42 %
0.0301** (0.0122)	1.1832* (0.4016)	-0.7863*** (0.4220)	0.0433 (0.1694)	33.02 %

* Reject the null with 99% confidence.

** Reject the null with 95% confidence.

*** Reject the null with 90% confidence.

Asymptotic Newey-West (1987) standard errors in parenthesis.

Table Vb
Encompassing Regressions Using Parkinson Realized Volatility (PK_t)

Intercept	$\sigma_{i,t}$	GARCH _t	MA(20) _t	Adjusted R ²
0.0160* (0.0056)	0.5788* (0.0433)			68.27 %
0.0243* (0.0062)		0.5999* (0.0547)		67.74 %
0.0430* (0.0048)			0.4094* (0.037)	64.09 %
0.0154** (0.0062)	0.7613* (0.2387)	-0.2071 (0.2452)		62.05 %
0.0149** (0.0070)	0.6118* (0.1254)		-0.0270 (0.0883)	67.22 %
0.0160** (0.0073)	0.7596* (0.2425)	-0.2267 (0.2409)	0.0155 (0.0985)	62.04 %

* Reject the null with 99% confidence.

** Reject the null with 95% confidence.

Asymptotic Newey-West (1987) standard errors are in parenthesis.

The R² 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 R² of the regression on implied volatility is 52.43%, while on GARCH and MA(20) forecasts is only 44.84% and 41.32%, respectively. When the Parkinson estimator (PK_t) is used, implied volatility explains 68.27% of the variation of realized volatility, while the GARCH forecast also does a good job, explaining 67.74%, while the MA(20) forecasts explains only 64.09%.

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 mixed. If we use the standard variation (SD_t) to measure realized volatility, Table Va shows that implied volatility is the only significant variable at the 99% significance level while the GARCH forecasts are significant at the 90% level. Nonetheless, using both implied

volatility and GARCH forecasts lowers the explanation power, which favors the implied volatility forecasts.

When the Parkinson estimator (PK_t) is used, Table Vb shows that the coefficients of historically-based forecasts are not significantly different from zero, suggesting that time series forecasts do not offer some incremental information to implied volatility.

4 Conclusions and directions for further research

Our results strongly suggest that the dollar-real volatility implied in prices of calls traded at BM&F, recovered by the use of the Garman-Kohlhagen option pricing model, contains information about subsequent realized volatility which is not present in past returns. Therefore, it is worth tracking dollar-real implied volatility in order to infer about future volatility, since forecasts that only use past returns are non-optimal, in the sense that they do not incorporate all public information available. This conclusion is in line with recent research, and is of interest to risk managers, asset and liability managers, players in the derivative markets, as well as financial regulators.

It is important to stress that results are robust to two alternative ways of measuring realized volatility, the standard deviation and the Parkinson estimator. It is also worth mentioning that the time series models were given the advantage of *ex post* parameter estimates.

Although dollar-real implied volatility is informative, our results show that it is an upward-biased estimator of future volatility. This finding is consistent with the results of Jorion (1995), Fleming (1998) and Bates (2000). Therefore, in order to build a superior forecast using implied volatility, one has to correct its bias.¹⁴

There are two possible sources for the upward-bias. First, it may be due to misspecification of the option pricing model, i.e., the market's forecast is not biased, but we rely on an inadequate pricing model to recover it. Second, the bias may be related to market imperfections, i.e., there are arbitrage opportunities, or transaction costs distort prices. A thorough investigation of the bias is beyond the scope of this study and left for further research.

14 By a simple linear model, for example.

However, we point out that model misspecification due to non-normality of returns, coupled with the fact that we never actually sample at-the-money but rather near-the-money calls, cannot be invoked to explain the bias. This is because the kind of non-normality occurred in the period of February 1999 to February 2001, positive skewness (3.20) and very high excess kurtosis (17.24), biases downwards, and not upwards, implied volatilities of near-the-money options priced by the Black-Scholes family of option pricing models, as proved by Backus *et alii* (1997).

Although we recognize that the problem of model misspecification exists, we believe that it is essentially related to the volatility risk premium. Especially in the beginning of the period considered, a few months after the change of the exchange rate regime, increases in the level of the dollar-real exchange rate were perceived to be associated with increases in the currency volatility.¹⁵ As the market as whole has been short dollars against Brazilian reais, increases in the level of the dollar-real exchange rate are associated with decreases in total market wealth. Thus, increases in volatility tend to be accompanied by decreases in market wealth, and because of that we regard volatility risk as systematic. Therefore, as investors are not indifferent to taking volatility risk, they demand a premium for being short volatility, as if they were "selling insurance" to the rest of the market.

We understand that there have been relatively few suppliers of this kind of insurance in the Brazilian market, thus, it is possible that the volatility risk premium does not account for the full magnitude of the bias we found. Then, in addition to model misspecification due to unpriced volatility risk, the bias may be also caused by market inefficiency to the strategy of systematically selling implied volatility and buying realized volatility. Therefore, it is worth testing if one can earn abnormal profits, after taking into consideration volatility risk, in the strategy of systematically shorting near-the-money dollar-real calls and delta-hedging currency exposure up to the maturity of the options. However, one has to take into consideration that there are high transaction costs involved in this strategy, and the bias may only signal an arbitrage opportunity after they are accounted for. Anyway, as mentioned before, a more elaborate and quantitative investigation into the sources of the bias, along the lines of Fleming (1999) and Bates (2000), is left for further research.

15 Using the GARCH(1,1) conditional volatility, we cannot reject the hypothesis of a positive correlation between volatility and returns on the dollar-real exchange rate at the 5% level. This phenomenon is also apparent in the positive skewness of the distribution of returns in the period.

Finally, we see as natural continuations of this research agenda the application of the same methodology outlined here to other Brazilian options markets, and the introduction of high frequency data to compute implied volatilities, build time series models and measure realized volatility.

References

- Amin, K., Ng, V. Keywords: currency options, implied volatility, forecast, information Inferring future volatility from the information in implied volatility in eurodollar options: a new approach. *Review of Financial Studies* 10, p. 333-367, 1997
- Andersen, T. G.; Bollerslev, T. Answering the skeptics: yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39, p. 885-905, 1998.
- Backus, D., Foresi, S., Li, K., Wu, L. Accounting for biases in Black-Scholes. *Working Paper*, New York University, 1997
- Bates, D. S. Post-'87 crash fears in the S&P 500 futures options market. *Journal of Econometrics* 94, p. 181-238, 2000.
- Blair, B. V.; Poon, S. H., Taylor, S. J. *Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high frequency index returns*. Lancaster University, 2000. Mimeografado.
- Black, F., Scholes, M. The valuation of options and corporate liabilities. *Journal of Political Economy* 81, p. 637-654, 1973.
- Bollerslev, T. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31, p. 307-327, 1986.
- Campa, J. M., Chang, P. H. K. The forecasting ability of correlations implied in foreign exchange options. *Journal of International Money and Finance* 17. p. 855-880, 1998.
- Canina, L., Figlewski, S. The informational content of implied volatility. *Review of Financial Studies* 6, p. 659-681, 1993.
- Christensen, B. J.; Prabhala, N. R. The relation between implied and realized volatility. *Journal of Financial Economics* 50, p. 125-150, 1998.

- Day, T., Lewis, C. Stock market volatility and the information content of stock index options. *Journal of Econometrics* 52, p. 267-287, 1992.
- Dickey, D. A., Fuller, W. A. Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association* 74, p. 427-431, 1979.
- Fair, R. C., Shiller, R. J. Comparing information in forecasts from econometric models. *American Economic Review* 80, p. 375-389, 1990.
- Feinstein, S. The Black-Scholes formula is nearly linear in s for at-the-money options; therefore implied volatilities from at-the-money options are virtually unbiased. *Working paper*, Federal Reserve Bank of Atlanta, 1989.
- Fleming, J.; Ostdiek, B., Whaley, R. E. Predicting stock market volatility: a new measure. *Journal of Futures Markets* 15, p. 265-302, 1995.
- Fleming, J. The quality of market volatility forecasts implied by S&P 100 index option prices. *Journal of Empirical Finance* 5, p. 317-345, 1998.
- _____ The economic significance of the forecast bias of S&P 100 index option implied volatility. *Advances in Futures and Options Research* 10, p. 219-251, 1999.
- Garman, M. B., Klass, M. J. On the estimation of security price volatilities from historical price. *Journal of Business*, 53, p. 67-78, 1980.
- Garman, M. B., Kohlhagen, S. Foreign currency option values. *Journal of International Money and Finance* 2, p. 231-238, 1983.
- Heynen, R. C., Kemna, A., Vorst, T. Analysis of the term structure of implied volatilities. *Journal of Financial and Quantitative Analysis* 29, p. 31-56, 1994.
- Hsieh, D. Modeling heteroskedasticity in daily foreign exchange rates. *Journal of Business and Economic Statistics* 7, p. 307-317, 1989.
- Hull, J. C., White, A. The pricing of options on assets with stochastic volatilities. *Journal of Finance* 42, p. 281-300, 1987
- Jorion, P. Predicting volatility in the foreign exchange market. *Journal of Finance* 50, p. 507-528, 1995.
- Lamoureux, C. G., Lastrapes, W. D. Forecasting stock-return variance: towards an understanding of stochastic implied volatilities. *Review of Financial Studies* 6, p. 293-326, 1993.

- Kroner, K. F., Kneafsey, K. P.; Claessens, S. Forecasting volatility in commodity markets. *Journal of Forecasting* 14, p. 77-95, 1995.
- Parkinson, M. The extreme value method for estimating the variance of the rate of return. *Journal of Business* 53, p. 61-65, 1980.
- Phillips, P. B., Perron, P. Testing for a unit root in time series regression. *Biometrika* 75, p. 335-346, 1988.
- Malz, A. Do implied volatilities provide early warning of market stress ? *RiskMetrics Journal* 1, p. 41-60, 2000.
- Newey, W.; West, K. A simple positive semi-definite, heteroscedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, p. 703-708, 1987.
- Scott, L. O. The information content of prices in derivative security markets. *IMF Staff Papers* 39, p. 596-625, 1992.
- Taylor, S. J., Xu, X. The incremental volatility information in one million foreign exchange quotations. *Journal of Empirical Finance* 4, p. 317-340, 1997
- Xu, X., Taylor, S. J. Conditional volatility and the informational efficiency of the PHLX currency options markets. *Journal of Banking and Finance* 19, p. 803-821, 1995.

