Modeling and Forecasting the Volatility of Brazilian Asset Returns: A Realized Variance Approach

Marcelo C. Carvalho*
Marco Aurélio S. Freire**
Marcelo C. Medeiros***
Leonardo R. Souza****

Abstract
The goal of this paper is twofold. First, using five of the most actively traded stocks in the Brazilian financial market, this paper shows that the normality assumption commonly used in the risk management area to describe the distributions of returns standardized by volatilities is not compatible with volatilities estimated by EWMA or GARCH models. In sharp contrast, when the information contained in high frequency data is used to construct the realized volatility measures, we attain the normality of the standardized returns, giving promise of improvements in Value-at-Risk statistics. We also describe the distributions of volatilities of the Brazilian stocks, showing that they are nearly lognormal. Second, we estimate a simple model to the log of realized volatilities that differs from the ones in other studies. The main difference is that we do not find evidence of long memory. The estimated model is compared with commonly used alternatives in an out-of-sample forecasting experiment.

***Keywords*:** Realized volatility; high frequency data; risk analysis; volatility forecasting; GARCH models.

JEL codes: C22; C51; C53.

Resumo
O artigo tem por objetivo modelar a volatilidade diária de cinco dos ativos mais negociados na bolsa de valores de São Paulo. A abordagem utilizada é baseada em dados intra-diários e o estimador conhecido como variância realizada é adotado. As principais conclusões são: em primeiro lugar, os retornos diários padronizados pela volatilidade realizada são aproximadamente normais. Além disso, as log-volatilidades também apresentam distribuicões bem próximas da normal. Finalmente, ao contrário da literature corrente, não há evidências

Submitted in November 2005. Revised in May 2006. This research has been partially supported by the CNPq. A previous version of this paper has been presented in seminars at IBMEC-RJ, PUC-Rio, and EPGE/FGV-RJ. We wish to thank the participants of these occasions for helpful comments. Our thanks also go to Álvaro Veiga, Cristiano Fernandes, and Marcelo Fernandes for insightful discussions. The suggestions and remarks from an anonymous referee are also gratefully acknowledged. Finally, the views expressed here do not represent those of the United Nations.

*Department Economic Statistics and Decision Support, Stockholm School of Economics, Stockholm, Sweden. E-mail: marcelo.carvalho@hhs.se

**Bank Boston, São Paulo, SP, Brazil. E-mail: msfreire@bkb.com.br

***Department of Economics, Pontifical Catholic University of Rio de Janeiro, Rio de Janeiro, RJ, Brazil. E-mail: mcm@econ.puc-rio.br

****Energy Statistics Section United Nations – Department of Economic and Social Affairs, United Nations. E-mail: souzal@un.org

©2004 Sociedade Brasileira de Finanças

de memória longa na série de volatilidade e um simples modelo de memória curta é suficiente para modelar e prever as séries diárias de volatilidade.

***Palavras-chave*:** volatilidade realizada; dados de alta freqüência; análise de risco; previsão de volatilidade; modelos GARCH.

1. Introduction

Given the rapid growth in financial markets and the continual development of new and more complex financial instruments, there is an ever-growing need for theoretical and empirical knowledge of the volatility in financial time series. It is widely known that the daily returns on financial assets, especially of stocks, are difficult to predict, although the volatility of the returns seems to be relatively easier to forecast. Therefore, it is hardly surprising that financial volatility has played such a central role in modern pricing and risk-management theories. There is, however, an inherent problem in using models where the volatility measure is necessary as the conditional variance is latent, and hence is not directly observable. The conditional variance can be estimated, among other approaches, by the (Generalized) Autoregressive Conditional Heteroskedasticity, or (G)ARCH, family of models proposed by Engle (1982) and Bollerslev (1986), stochastic volatility models (Taylor, 1986), or exponentially weighted moving averages (EWMA), as advocated by the Riskmetrics methodology (Morgan, 1996); see McAleer (2005) for a recent exposition. However, as observed by Bollerslev (1987), Malmsten and Teräsvirta (2004), and Carnero et al. (2004), among others, most of the latent volatility models fail to describe satisfactorily several stylized facts observed in financial time series. As a common example, standardized returns still have excess kurtosis. Another empirical fact that standard latent volatility models fail to describe in an adequate manner is the low, but slowly decreasing, autocorrelations in the squared returns associated with high excess kurtosis of returns. Correctly describing the dynamics of the returns is important in order to obtain accurate forecasts of the future volatility which, in turn, is important in risk analysis and management. In this sense, the assumption of Gaussian standardized returns has been refuted in many studies, such that heavy-tailed distributions have been used instead.

The search for an adequate framework for the estimation and prediction of the conditional variance of financial assets returns has led to the analysis of high frequency intraday data. Merton (1980) noted that the variance over a fixed interval can be estimated arbitrarily, although accurately, as the sum of squared realizations, provided the data are available at a sufficiently high sampling frequency. More recently, Andersen and Bollerslev (1998) showed that ex post daily foreign exchange volatility is best measured by aggregating 288 squared five-minute returns. The five-minute frequency is a trade-off between accuracy, which is theoretically optimized using the highest possible frequency, and microstructure noise that can arise through the bid-ask bounce, asynchronous trading, infrequent trading, and price discreteness, among other factors (see Madhavan (2000) or Biais

et al. (2005) for very good recent surveys). The sum of intraday squared returns is widely known as the realized variance, and its squared-root is the realized volatility.

Ignoring the remaining measurement error, the ex post volatility essentially becomes "observable". Andersen and Bollerslev (1998), Hansen and Lunde (2005), and Patton (2005) used the realized volatility to evaluate the out-of-sample forecasting performance of several latent volatility models. This same approach was adopted by Mota and Fernandes (2004) to compare different volatility models applied to the index of the São Paulo stock market (IBovespa).

On the other hand, as volatility becomes "observable", it can be modeled directly, rather than being treated as a latent variable. Based on the theoretical results of Andersen et al. (2003), Barndorff-Nielsen and Shephard (2002), and Meddahi (2002) several recent studies have documented the properties of realized volatilities constructed from high frequency data. Just to mention few recent examples, Andersen et al. (2001a) studied several exchange rates series. Andersen et al. (2001a) and Andersen et al. (2001b) examined the distribution of US stock returns. Other examples are Martens (2001a,b), Bollen and Inder (2002), and Martens et al. (2004).

Several important characteristics of the returns and realized volatilities came out from these studies. First, the unconditional distribution of daily returns exhibit excess kurtosis. Daily returns are not autocorrelated (except for the first order in some cases). Second, daily returns standardized by the realized variance measure are almost Gaussian. Third, the unconditional distributions of realized variance and volatility are distinctly non-normal and extremely right skewed. On the other hand, the natural logarithm of the volatility is close to normality. Fourth, the log of the realized volatility displays a high degree of (positive) autocorrelation which dies out very slowly. Finally, realized volatility does not seem to have a unit root, but there is clear evidence of fractional integration, roughly of order 0.40.

The main goal of this paper is twofold. First, using five of the most actively traded stocks in the São Paulo Stock Exchange (Bovespa), this paper shows that the normality assumption commonly used in the risk management area to describe the distributions of returns standardized by volatilities is not compatible with volatilities estimated by EWMA or GARCH models. In sharp contrast, when we use the information contained in high frequency data to construct the realized volatility measures, we attain the normality of the standardized returns, giving promise to improve on Value at Risk statistics. We also describe the distributions of volatilities of the Brazilian stocks, showing that they are nearly lognormal. Second, we estimate a simple model to the log of realized volatilities that differs from the ones in other studies. The main difference is that we do not find evidence of long memory. The estimated model is compared with commonly used alternatives in an out-of-sample experiment.

The study of Brazilian data is important as most of the stylized facts found in the recent literature concerns the US or European countries. However, results from

emerging markets can be significantly distinct as a consequence, for example, of differences in market microstructures. Furthermore, as the availability of high-frequency Brazilian data is much more limited, it is important to check how good intraday based models perform in practical situations, as for example, in Value-at-Risk analysis. This paper sheds some light on this issue, by showing when combined with latent volatility models, realized volatility based models improve the construction of confidence intervals in a forecasting exercise, yielding a more precise risk measure.

The paper proceeds as follows. In Section 2, we briefly describe the calculation of the realized volatility. Section 3 describes the data used in the paper and carefully analyze the distribution of the standardized returns and realized volatility. In Section 4 we estimate a simple linear model to the realized volatility and an out-of-sample experiment is conducted to evaluate the forecasting performance of the estimated models. Finally, Section 5 concludes.

2. Realized Variance and Realized Volatility

Suppose that, along day t, the logarithmic prices of a given asset follow a continuous time diffusion, as follows:

$$dp(t+\tau) = \mu(t+\tau) + \sigma(t+\tau)dW(t+\tau),\ 0 \leq \tau \leq 1,\ t = 1, 2, 3, \ldots \quad (1)$$

where $p(t+\tau)$ is the logarithmic price at time $(t+\tau)$, $\mu(t+\tau)$ is the drift component, $\sigma(t+\tau)$ is the instantaneous volatility (or standard deviation), and $dW(t+\tau)$ is the standard Brownian motion. Usually, the drift $\mu(t+\tau)$ is assumed to be constant.

Andersen et al. (2003) and Barndorff-Nielsen and Shephard (2002) showed that the daily compound returns, defined as $r_t = p(t) - p(t-1)$, are Gaussian conditionally on $\mathcal{F}_t = \mathcal{F}\{\mu(t+\tau-1), \sigma(t+\tau-1),\ 0 \leq \tau \geq 1\}$, the σ-algebra (information set) generated by the sample paths of $\mu(t+\tau-1)$ and $\sigma(t+\tau-1)$, $0 \leq \tau \leq 1$, such that

$$r_t|\mathcal{F}_t \sim \mathsf{N}\left(\int_0^1 \mu(t+\tau)d\tau, \int_0^1 \sigma^2(t+\tau)d\tau\right) \quad (2)$$

The term $IV_t = \int_0^1 \sigma^2(t+\tau)d\tau$ is known as the *integrated variance*, which is a measure of the day-t ex post volatility. In this sense, the integrated variance is the object of interest.

In practical applications, prices are observed at discrete and irregularly spaced intervals and there are many ways to sample the data. Suppose that in a given day t, we partition the interval $[0, 1]$ in subintervals and define the grid of observation times $\mathcal{G} = \{\tau_1, \ldots, \tau_n\}$, $0 = \tau_0 < \tau_1 < \cdots, \tau_n = 1$. The length of the ith subinterval is given by $\delta_i = \tau_i - \tau_{i-1}$. We shall assume that the length of each

subinterval shrinks to zero as the number of intraday observations (n) increases. The integrated variance over each of the subintervals is defined as

$$IV_{i,t} = \int_{\tau_{i-1}}^{\tau_i} \sigma^2(t+\tau)d\tau$$

The most widely used sampling scheme is calendar time sampling (CTS), where the intervals are equidistant in calendar time, that is $\delta_i = \frac{1}{n}$. This time of sampling will be adopted throughout the paper.

Set $p_{i,t}$, $t = 1, \ldots, n$, to be the ith price observation during day t, such that $r_{t,i} = p_{t,i} - p_{t,i-1}$ is the ith intra-period return of day t. The realized variance is defined as

$$RV_t = \sum_{i=1}^{n} r_{t,i} \tag{3}$$

The realized volatility is the square-root of RV_t.

Under some additional mild regularity conditions, which includes the assumption of uncorrelated intraday returns, Andersen et al. (2003) showed that the realized variance using all data available, as defined in equation (3), is a consistent estimator of the integrated variance, such that $RV_t \xrightarrow{p} IV_t$. Barndorff-Nielsen and Shephard (2002) and Bandi and Russell (2005b) derived the asymptotic distribution of the realized variance as

$$\sqrt{n}\frac{1}{\sqrt{2IQ_t}}\left(RV_t - IV_t\right) \xrightarrow{d} \mathsf{N}(0,1), \tag{4}$$

where the integrated quarticity, IQ_t, is defined as

$$IQ_t = \int_0^1 \sigma^4(t+\tau)d\tau. \tag{5}$$

Furthermore, under similarly mild assumptions, the integrated quarticity is consistently estimated by the realized quarticity, which is defined as

$$RQ_t = \frac{1}{3}\sum_{i=1}^{n} r_{t,i}^4 \tag{6}$$

However, when the returns are correlated, the realized volatility will be a biased estimator of the daily volatility. Although, in the context of efficient markets, the finding of correlated intraday returns may at first sight appear puzzling, it has a sensible explanation in the context of the market microstructure literature; see Campbell et al (1997, chapter 3). When the returns are sampled at higher frequencies, market microstructure may introduce some autocorrelation in the intraday returns, thus, driving the realized variance to be a biased estimator of the daily variance. On the other hand, lower frequencies may lead to an estimator with a

higher variance. The effects of microstructure and the optimal sampling of intraday returns have been deeply discussed in several papers, such as, for example, Bandi and Russell (2005a,b, 2006), Oomen (2005), Zhang et al. (2005), Hansen and Lunde (2006), among others.

3. The Data

In this paper we use data of five out of the ten most actively traded stocks on the São Paulo Stock Market (Bovespa), namely: Bradesco (BBDC4), Embratel (EBTP4), Petrobrás (PETR4), Telemar (TNLP4), and Vale do Rio Doce (VALE5). The data set consists of intraday prices observed every 15 minutes from 01 October 2001 to 30 November 2003 (539 daily observations). We use data until 11 April 2003 (379 daily observations) for estimation and in-sample evaluation and the remaining observations for out-of-sample analysis. Figure 1 shows the daily returns. The dashed lines represent the out-of-sample period.

One important point to mention is the choice of the sampling frequency. Due to the lack of more frequently observed prices, recent techniques to optimally estimate the realized variance are not possible to be used. Our choice is to heuristically test the bias-efficiency trade-off involved for three different frequencies: 15 minutes, 30 minutes, and 45 minutes. First, we estimate the realized volatility using three different frequencies as mentioned above and average them over the sample. Table 1 shows the average of the daily realized volatility. As pointed out by Andersen et al. (2003), if microstructure effects are present the average of the realized volatility may be differ according to the sampling frequency. As we can see by inspection of Table 1, the mean is rather stable.

Table 1
Average daily realized volatility

Asset	15-minute window	30-minute window	45-minute window
Bradesco	0.0215	0.0199	0.0200
Embratel	0.0434	0.0399	0.0404
Petrobrás	0.0200	0.0188	0.0190
Telemar	0.0228	0.0218	0.0221
Vale	0.0172	0.0159	0.0159

Notes: The table shows the average of the daily realized volatility estimated using different sampling frequencies. The estimation period is 01 October 2001 – 11 April 2003.

On the other hand, to estimate the precision of the estimator we make use of equation (4). Table 2 shows the average size of the 95% confidence interval for the realized volatility calculated from (4). As can be observed it seems that a 15-minute frequency yields the most tight confidence intervals. Given the stability of the mean across sampling frequencies, we proceed the analysis using the 15 minutes frequency as it entails the tightest confidence intervals.

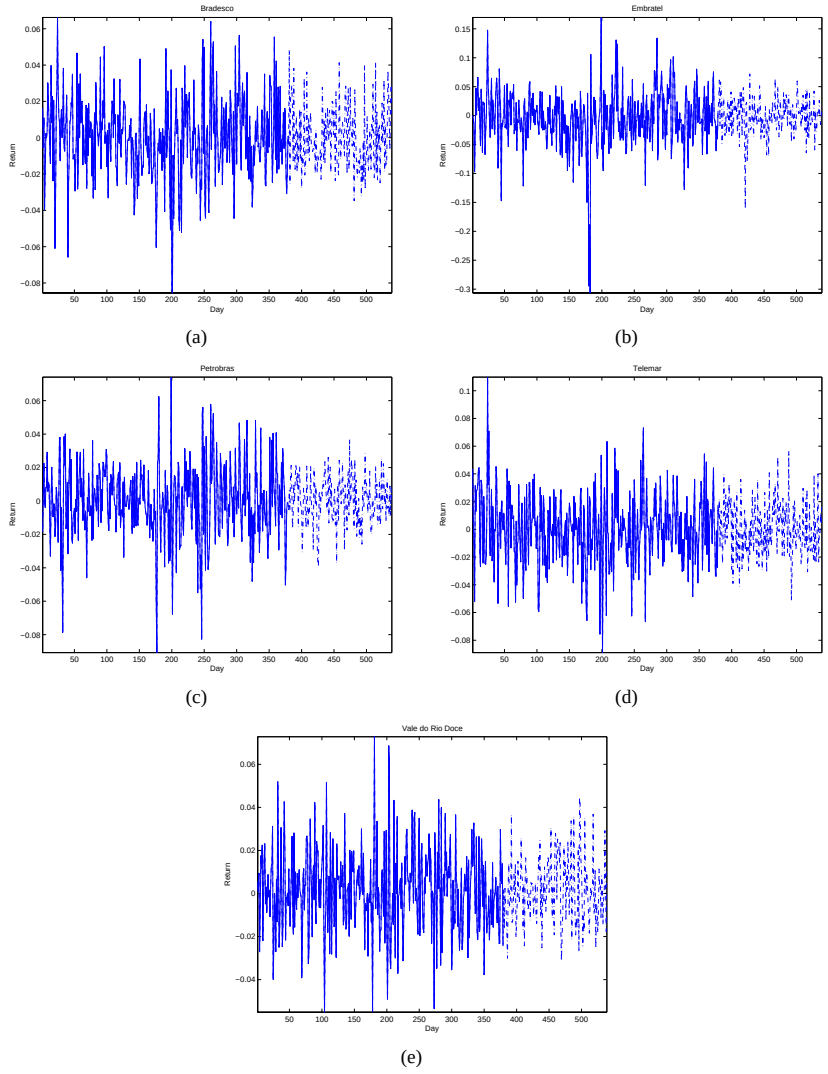

Figure 1
Daily returns. The dashed lines represent the out-of-sample period. Panel (a): Bradesco. Panel (b): Embratel. Panel (c): Petrobrás. Panel (d): Telemar. Panel (e): Vale do Rio Doce.

Table 2
Average of the confidence intervals of the daily realized volatility.

Asset	15-minute window	30-minute window	45-minute window
Bradesco	0.000871	0.001000	0.001100
Embratel	0.004800	0.005100	0.005200
Petrobrás	0.000867	0.000967	0.001100
Telemar	0.000930	0.001100	0.001300
Vale	0.000670	0.000724	0.000787

Notes: The table shows the average of the confidence interval of the daily realized volatility estimated using different sampling frequencies. The estimation period is 01 October 2001 – 11 April 2003.

3.1 The Distribution of Standardized Returns and Realized Volatility

Table 3 shows the mean, the standard deviation, the skewness, the kurtosis, and the p-value of the Jarque-Bera normality test for each of the daily returns on the five stocks considered in this paper, As expected, all the five series have excess of kurtosis, specially Embratel. Four of the series are negatively skewed, whereas Vale do Rio Doce is positive skewed. The Jarque-Bera test strongly rejects the null hypothesis of normality for all the five series.

Table 3
Daily returns: Descriptive statistics

Asset	Mean	Standard deviation	Skewness	Kurtosis	Jarque-Bera
Bradesco	1.64×10^{-4}	0.0225	-0.1669	4.0555	9.00×10^{-5}
Embratel	-4.90×10^{-3}	0.0511	-0.9743	8.9420	0
Petrobrás	-1.15×10^{-4}	0.0224	-0.2611	4.4836	5.80×10^{-9}
Telemar	3.91×10^{-4}	0.0256	-0.0107	4.0867	1.28×10^{-4}
Vale	1.60×10^{-3}	0.0192	0.1847	3.8186	2.20×10^{-3}

Notes: The table shows the mean, the standard deviation, the skewness, the kurtosis, and the p-value of the Jarque-Bera normality test for each one of the five series considered in this paper. The estimation period is 01 October 2001 – 11 April 2003.

Table 4 shows descriptive statistics for the standardized returns. To compare the realized volatility approach with other methods to compute the daily volatility, we estimate the following models: GARCH(1,1), EGARCH(1,1) Nelson (1991), and the asymmetric GJR-GARCH(1,1) (Glosten et al., 1993). In addition we also compute the volatility with the Riskmetrics methodology that is based on a exponentially weighted moving average of the squared returns (EWMA) with a decay factor $\lambda = 0.94$ as suggested in Morgan (1996).[1] For each of the daily standardized returns on the five stocks considered in this paper, Table 4 shows the mean, the standard deviation, the skewness, the kurtosis, and the p-value of the Jarque-Bera normality test. It seems that the realized volatility methodology produces (nearly) Gaussian standardized returns for all the five series. The same result does not hold for the other models. The only exceptions are the GARCH(1,1), EGARCH(1,1), and GJR-GARCH(1,1) models estimated for Bradesco and Vale do Rio Doce and the EGARCH(1,1) and the GJR-GARCH(1,1) for Petrobrás. Altogether, two main

[1] We also tested the EWMA with decay factor $\lambda = 0.89$ as suggested by Morgan (1996) for the case of emerging markets. However, the results were mostly identical.

conclusions emerge from the results in Table 4. Firstly, the use of realized volatility leads to standardized returns with distributions that are not statistically different from Gaussian. Secondly, asymmetrical models, such as the EGARCH and GJR-GARCH specifications, can improve on the GARCH(1,1) alternative by generating normal standardized returns.

Figure 2 shows the histograms of the returns and standardized returns when the daily variance is estimated by the realized volatility approach. Table 5 shows descriptive statistics for the realized volatility. It is clear that, for all the five series, the realized volatility is positively skewed and non-Gaussian, as the Jarque-Bera test strongly rejects the null hypothesis of normality. On the other hand, as shown in Table 6, the natural logarithm of the realized volatility is nearly Gaussian for Bradesco and Telemar. For Petrobrás and Vale, the p-value of the Jarque-Bera test is larger when the logarithms are considered, but normality s still rejected at any reasonable significance level. For Embratel, the log realized volatility is still strongly non Gaussian. The main reason for the rejection of the null hypothesis of normality is the high values for the skewness statistic. Figures 3 and 4 show the evolution and the histogram of the realized volatility and the log realized volatility.

Table 4
Daily standardized returns: descriptive statistics

Asset	Mean	Standard deviation	Skewness	Kurtosis	Jarque-Bera
		Panel I: Realized Volatility			
Bradesco	0.0152	0.9956	0.0831	2.7107	0.3890
Embratel	-0.0963	0.9976	0.1003	2.4461	0.0577
Petrobrás	-0.0076	1.0088	0.0122	2.4885	0.1133
Telemar	0.0236	1.0370	0.0672	2.5952	0.2177
Vale	0.1056	1.0748	0.0387	2.7161	0.4728
		Panel II: EWMA ($\lambda = 0.94$)			
Bradesco	$-7.37 \times 10-4$	1.0334	-0.1169	3.7898	0.0061
Embratel	-0.1089	1.0309	-0.4971	4.8002	$6.88 \times 10-15$
Petrobrás	-0.0171	1.0483	-0.4817	4.7595	$1.85 \times 10-14$
Telemar	0.0032	1.0357	-0.1474	3.7688	0.0060
Vale	0.0811	1.0312	-0.0796	4.2522	$8.95 \times 10-6$
		Panel II: GARCH(1,1)			
Bradesco	0.0030	1.0005	-0.0757	3.5296	0.1063
Embratel	0.0067	1.0010	-0.3006	4.0334	$1.80 \times 10-5$
Petrobrás	-0.0014	0.9976	-0.1800	3.5370	0.0434
Telemar	-0.0064	1.0068	-0.0598	3.9023	0.0019
Vale	0.0021	0.9995	0.0180	3.5813	0.0815
		Panel III: EGARCH(1,1)			
Bradesco	0.0088	1.0006	-0.0365	3.5296	0.1407
Embratel	-0.0008	1.0010	-0.3467	4.0334	$1.43 \times 10-7$
Petrobrás	0.0054	0.9989	-0.0989	3.5370	0.5153
Telemar	-0.0069	1.0089	-0.0971	3.9023	0.0164
Vale	0.0009	0.9996	0.0292	3.5813	0.0930
		Panel IV: GJR-GARCH(1,1)			
Bradesco	0.0098	1.0007	-0.0343	3.4954	0.1599
Embratel	0.0053	1.0010	-0.3212	4.0521	$8.98 \times 10-6$
Petrobrás	0.0045	0.9992	-0.1131	3.2742	0.3978
Telemar	-0.0059	1.0074	-0.0527	3.9142	0.0017
Vale	-0.0013	0.9994	-0.0250	3.4940	0.1644

Notes: The table shows the mean, the standard deviation, the skewness, the kurtosis, and the *p*-value of the Jarque-Bera test of the daily standardized returns. The estimation period is 01 October 2001 – 11 April 2003.

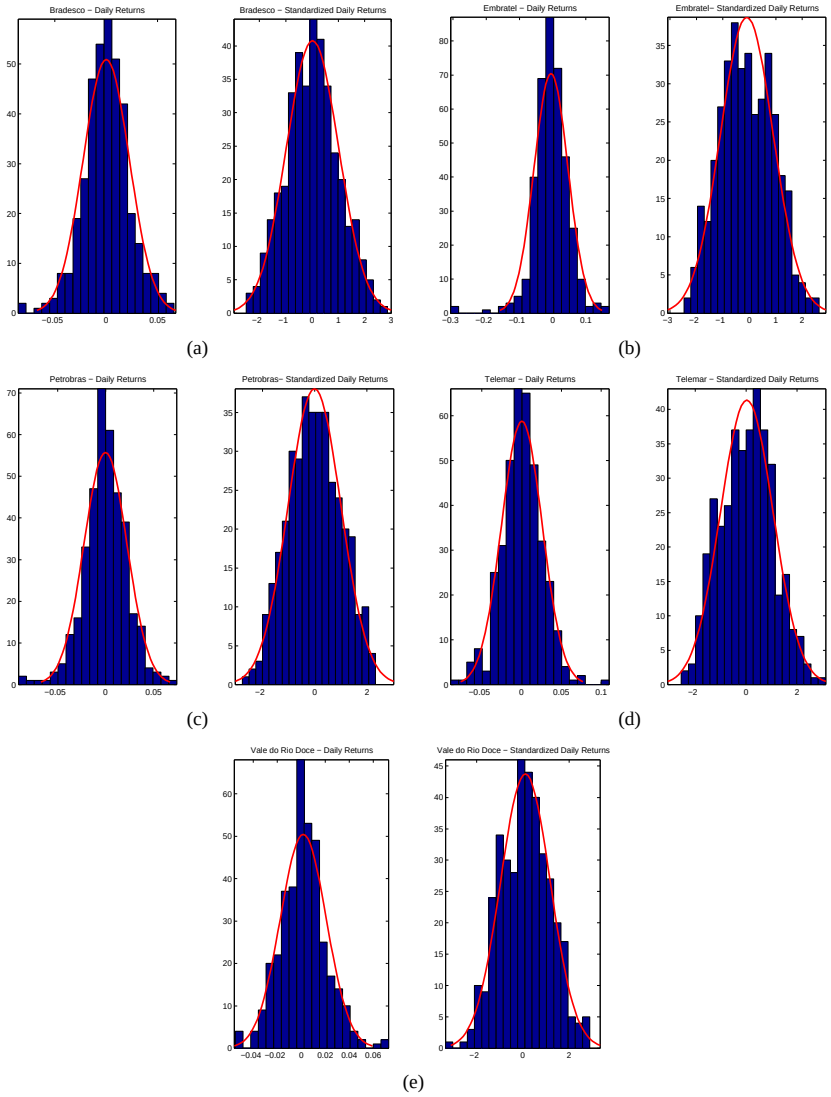

Figure 2
Histograms of the daily returns and standardized daily returns. Panel (a): Bradesco. Panel (b): Embratel. Panel (c): Petrobrás. Panel (d): Telemar. Panel (e): Vale do Rio Doce.

Table 5
Realized volatility: descriptive statistics

Asset	Mean	Standard deviation	Skewness	Kurtosis	Jarque-Bera
Bradesco	0.0215	0.0079	1.7447	9.6061	0
Embratel	0.0434	0.0225	5.2217	53.4585	0
Petrobrás	0.0200	0.0091	2.0659	9.5716	0
Telemar	0.0228	0.0079	0.7927	3.5341	3.44×10^{-10}
Vale	0.0172	0.0084	2.4204	12.0763	0

Notes: The table shows the mean, the standard deviation, the skewness, the kurtosis, and the p-value of the Jarque-Bera test of the daily realized volatilities. The estimation period is 01 October 2001 – 11 April 2003.

Table 6
Daily log realized volatilities: descriptive statistics

Asset	Mean	Standard deviation	Skewness	Kurtosis	Jarque-Bera
Bradesco	-3.9002	0.3451	0.0753	3.5299	0.1062
Embratel	-3.2200	0.3806	0.8128	5.1035	0
Petrobrás	-3.9939	0.3929	0.4693	3.4192	2.82×10^{-4}
Telemar	-3.8394	0.3473	-0.1016	2.7485	0.414
Vale	-4.1530	0.4081	0.5341	3.7153	2.87×10^{-6}

Notes: The table shows the mean, the standard deviation, the skewness, the kurtosis, and the p-value of the Jarque-Bera test of the daily log realized volatilities. The estimation period is 01 October 2001 – 11 April 2003.

4. Modeling and Forecasting Realized Volatility

4.1 In-sample analysis

In order to compare the performance of different methods and models to extract the daily volatility, we estimate 95% confidence intervals for the daily returns and check the number of observations that are outside the interval. Table 7 shows the number of exceptions of the different coverage probabilities. Observing Table 7 several facts emerge. Apart from the case of Vale, the 99% confidence intervals computed from realized volatilities are slightly overestimated. All the other models strongly underestimate the confidence intervals. This is mainly due to the fact that we consider the standard normal as the distribution of the standardized returns, which is certainly not the case when the latent volatility models are used. Similar results are found for the 95% confidence interval. When the 90% confidence interval is considered, than all the alternative models and methods seem to slightly underestimate the confidence intervals, apart from the cases of Embratel and Telemar.

In order to test if the coverage of different models and methods are statistically different from the nominal ones, Table 8 shows the p-values of the tests of unconditional coverage, independence, and conditional coverage (Christoffersen, 1998). According to tests, the realized volatility seems to produce "correct" intervals, with the GJR-GARCH model being the second best alternative to build confidence intervals. All the other models/methods fail in at least in one of the tests. For example, EWMA fails the unconditional and conditional coverage tests for Bradesco and Embratel when the 95% confidence interval is considered.

Table 7
In-sample analysis: frequency of observations of the returns that are outside a given confidence interval

Asset	Realized Volatility	EWMA ($\lambda = 0.94$)	GARCH(1,1)	EGARCH(1,1)	GJR-GARCH(1,1)
		Panel I: 99% Confidence Interval			
Bradesco	0.0026	0.0264	0.0185	0.0211	0.0211
Embratel	0.0026	0.0264	0.0211	0.0290	0.0211
Petrobras	0.0026	0.0158	0.0211	0.0158	0.0185
Telemar	0.0053	0.0237	0.0132	0.0185	0.0132
Vale	0.0158	0.0185	0.0185	0.0132	0.0158
		Panel II: 95% Confidence Interval			
Bradesco	0.0449	0.0686	0.0660	0.0765	0.0712
Embratel	0.0422	0.0660	0.0554	0.0528	0.0554
Petrobras	0.0422	0.0501	0.0554	0.0528	0.0501
Telemar	0.0528	0.0686	0.0607	0.0686	0.0660
Vale	0.0554	0.0554	0.0475	0.0501	0.0475
		Panel III: 90% Confidence Interval			
Bradesco	0.1029	0.0976	0.1055	0.1108	0.1029
Embratel	0.1003	0.1003	0.0923	0.0844	0.0923
Petrobras	0.1108	0.1082	0.1082	0.1055	0.1108
Telemar	0.1161	0.1108	0.0950	0.0976	0.1003
Vale	0.1266	0.1055	0.1029	0.1082	0.1029

Notes: The estimation period is 01 October 2001 – 11 April 2003.

Table 8
In-sample analysis: p-value of the test of the null hypothesis of correct unconditional coverage, independence, and correct conditional coverage, at nominal significance level 0.05

Asset	Realized Volatility	EWMA	GARCH(1,1)	EGARCH(1,1)	GJR-GARCH(1,1)
			Panel I: Unconditional Coverage		
			99% Confidence Interval		
Bradesco	0.0866	0.0078	0.1383	0.0585	0.0585
Embratel	0.0866	0.0078	0.0585	0.0025	0.0585
Petrobrás	0.0866	0.2930	0.0585	0.2930	0.1383
Telemar	0.3098	0.0223	0.5515	0.1383	0.5515
Vale	0.2930	0.1383	0.1383	0.5515	0.2930
			95% Confidence Interval		
Bradesco	0.6402	0.1149	0.1731	0.0275	0.0737
Embratel	0.4755	0.1731	0.6346	0.8062	0.6346
Petrobrás	0.4755	0.9906	0.6346	0.8062	0.9906
Telemar	0.8062	0.1149	0.3550	0.1149	0.1731
Vale	0.6346	0.6346	0.8214	0.9906	0.8214
			Panel II: Independence		
			99% Confidence Interval		
Bradesco	0.9419	0.2535	0.6073	0.5564	0.5564
Embratel	0.9419	0.4610	0.5564	0.4167	0.5564
Petrobrás	0.9419	0.0742	0.1497	0.0742	0.1084
Telemar	0.8840	0.5076	0.7143	0.6073	0.7143
Vale	0.6600	0.6073	0.6073	0.7143	0.6600
			95% Confidence Interval		
Bradesco	0.2057	0.8214	0.7790	0.8684	0.9561
Embratel	0.2343	0.0510	0.7852	0.1348	0.1159
Petrobrás	0.7005	0.8062	0.8673	0.9520	0.9616
Telemar	0.3832	0.6402	0.0490	0.1176	0.0897
Vale	0.1159	0.8214	0.9616	0.9616	0.8743
			Panel III: Conditional Coverage		
			99% Confidence Interval		
Bradesco	0.2298	0.0151	0.2921	0.1404	0.1404
Embratel	0.2298	0.0220	0.1404	0.0074	0.1404
Petrobrás	0.2298	0.1168	0.0591	0.1168	0.0919
Telemar	0.5907	0.0590	0.7832	0.2921	0.7832
Vale	0.5222	0.2921	0.2921	0.7832	0.5222
			95% Confidence Interval		
Bradesco	0.4025	0.1003	0.3801	0.0868	0.2017
Embratel	0.3821	0.0841	0.7533	0.3173	0.2595
Petrobrás	0.7200	0.1947	0.8808	0.9689	0.9988
Telemar	0.6634	0.6704	0.0940	0.0848	0.0937
Vale	0.2595	0.8514	0.9988	0.9988	0.9627

Notes: The estimation period is 01 October 2001 – 11 April 2003.

Figure 5 shows the autocorrelation and partial autocorrelation functions for the estimated log realized volatilities. By inspection of Figure 5 the natural logarithm of the realized volatilities, on the contrary of the international empirical evidence, is not very persistent. Table 9 presents the statistics and the respective p-values of the Augmented-Dickey-Fuller (ADF) and Philipps-Perron (PP) tests for unit

root and Lo's (1991) test for long-memory. The unit-root hypothesis is strongly rejected for all the five series. Furthermore, according to Lo's test there is no statistical evidence against the short-memory hypothesis.

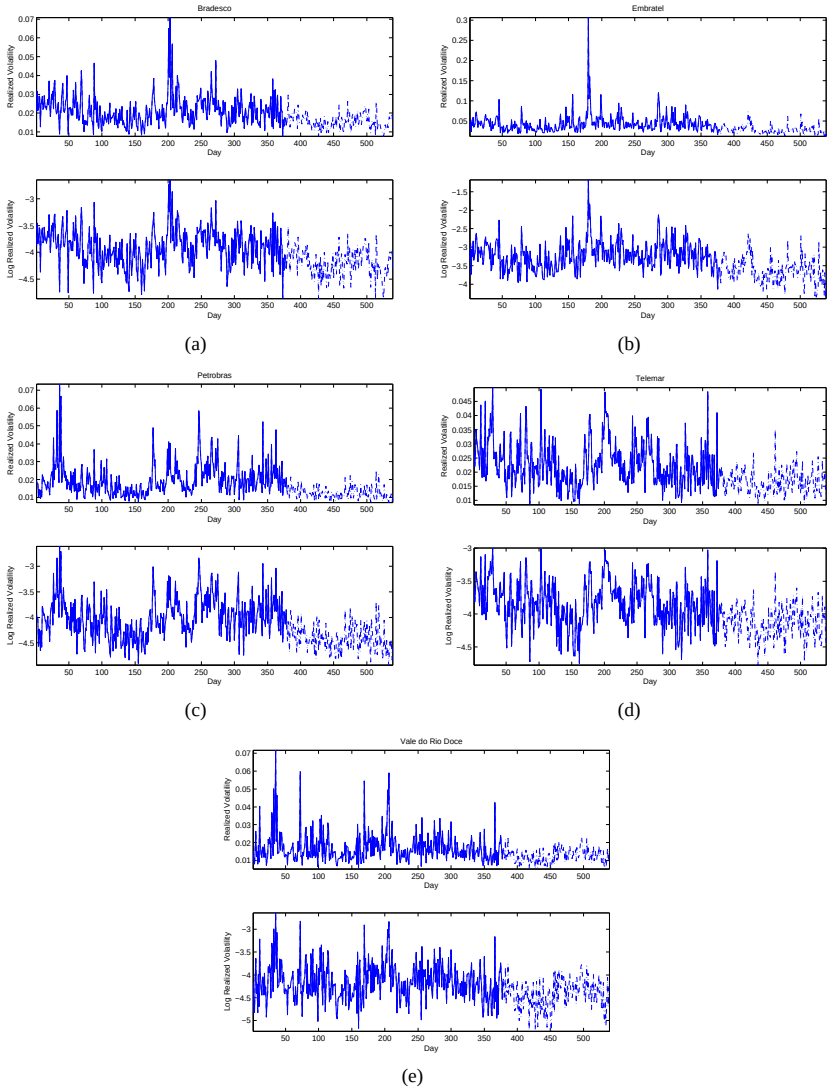

Figure 3
Daily realized volatilities. Panel (a): Bradesco. Panel (b): Embratel. Panel (c): Petrobrás. Panel (d): Telemar. Panel (e): Vale do Rio Doce.

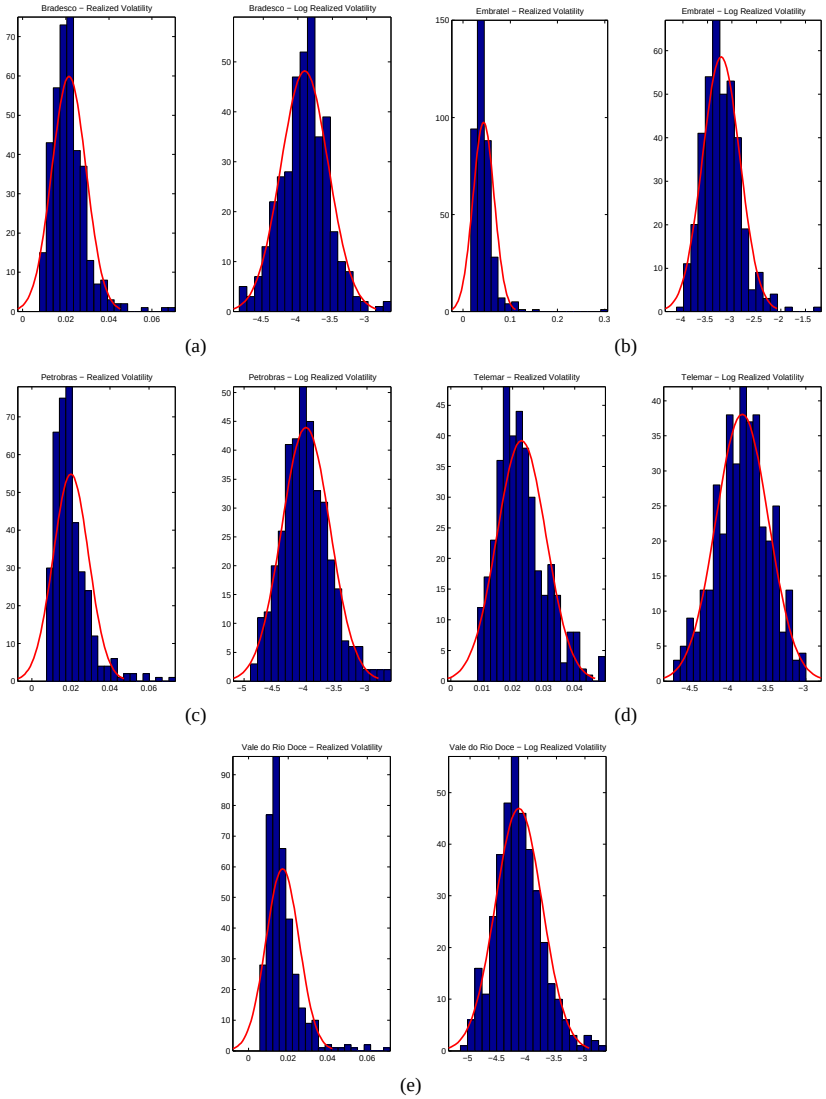

Figure 4
Histograms of the daily realized volatilities and log daily realized volatilities. Panel (a): Bradesco. Panel (b): Embratel. Panel (c): Petrobrás. Panel (d): Telemar. Panel (e): Vale do Rio Doce.

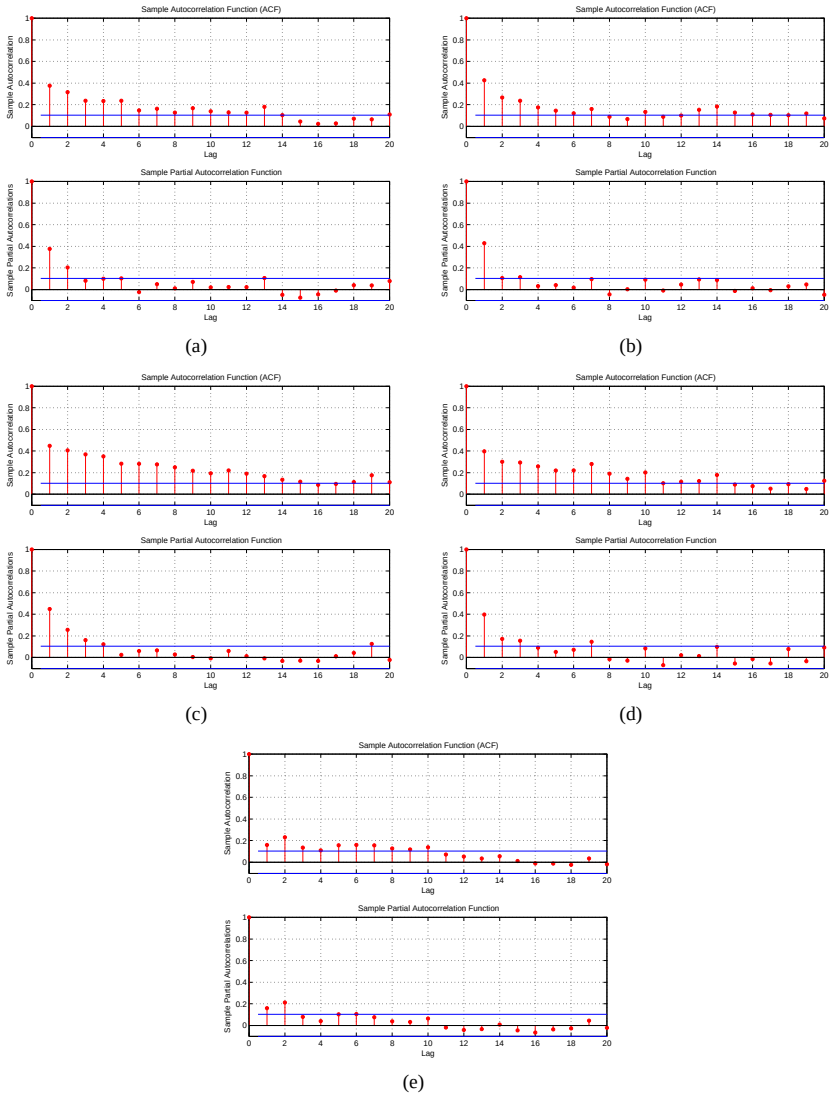

Figure 5
Autocorrelation and partial autocorrelation functions of the log realized volatility. Panel (a): Bradesco. Panel (b): Embratel. Panel (c): Petrobrás. Panel (d): Telemar. Panel (e): Vale do Rio Doce.

Table 9
In-sample analysis: Unit-root tests

Asset	Dickey-Fuller	Philipps-Perron	Lo's (1991) test
Bradesco	-8.77 (0)	-14.38 (0)	1.50 (>0.10)
Embratel	-12.22 (0)	-12.97 (0)	1.53 (>0.10)
Petrobras	-5.19 (0)	-13.97 (0)	1.57 (>0.10)
Telemar	-6.92 (0)	-14.29 (0)	1.08 (>0.10)
Vale	-10.16 (0)	-17.69 (0)	1.15 (>0.10)

Notes: The table shows the statistics of ADF and PP unit-root tests as well as Lo's (1991) test for short memory against long memory applied to the log of the realized volatilities. The numbers in parentheses are the p-values. The estimation period is 01 October 2001 – 11 April 2003.

Based on the evidence of no long memory in the log realized volatility series, we proceed by estimating a simple model for each series defined as

$$\log(h_t) = \alpha + \beta r_{t-1}^2 + \phi \log(h_{t-1}) + \delta \log(h_{t-1}) \times (r_{t-1} < 0) + \theta \varepsilon_{t-1} + \varepsilon_t, \quad (7)$$

where $\{\varepsilon_t\}_{t=1}^T$ is a sequence of independent and identically distributed random variables with zero mean and variance σ^2, $\varepsilon_t \sim \text{IID}\left(0, \sigma^2\right)$. The details of the estimated models are described in Table 10, which shows the estimated parameters and several diagnostic statistics. $R^2_{\text{adj.}}$ is the adjusted R^2, JB s the p-value of the Jarque-Bera normality test, $\text{LM}_{\text{SC}}(i)$ is the p-value of the Lagrange Multiplier (LM) test for ith order serial correlation in the estimated residuals, and $\text{LM}_{\text{ARCH}}(i)$ is the p-value of LM test for ith order ARCH effects in residuals.

It is important to stress some points with respect to the estimated model. Firstly, the moving average term is included in order to remove remaining autocorrelation in the residuals. Increasing the number of lags does not seem to yield uncorrelated residuals. Secondly, the leverage effect is only significant on the cases of Petrobrás and Telemar. In addition, apart from Bradesco and Telemar, all the residuals seem to be non Gaussian. Furthermore, there is no evidence of remaining serial correlation. However, there is some evidence of ARCH effects (volatility of volatility) for Bradesco and Vale, which may indicate the presence of time-varying conditional kurtosis.

Table 10
In-sample analysis: estimated models

$$\log(h_t) = \alpha + \beta r_{t-1}^2 + \phi \log(h_{t-1}) + \delta \log(h_{t-1}) \times (r_{t-1} < 0) + \theta \varepsilon_{t-1} + \varepsilon_t$$

Parameters	Bradesco	Embratel	Petrobrás	Telemar	Vale
α	-1.02 (0.20)	-2.69 (0.56)	-1.01 (0.25)	-1.35 (0.21)	-1.50 (0.49)
β	130.14 (25.31)	17.56 (5.58)	87.99 (26.10)	123.89 (18.04)	151.66 (49.71)
ϕ	0.88 (0.03)	0.59 (0.09)	0.89 (0.03)	0.84 (0.03)	0.83 (0.06)
δ	—	—	-0.01 (0.006)	-0.02 (0.005)	—
θ	-0.82 (0.04)	-0.33 (0.09)	-0.70 (0.06)	-0.77 (0.05)	-0.75 (0.07)
$R^2_{\text{adj.}}$	0.27	0.23	0.32	0.30	0.12
JB	0.18	0	0	0.33	0
$\text{LM}_{\text{SC}}(1)$	0.12	0.48	0.88	0.32	0.08
$\text{LM}_{\text{SC}}(4)$	0.45	0.87	0.84	0.85	0.06
$\text{LM}_{\text{ARCH}}(1)$	0.05	0.70	0.90	0.55	0.68
$\text{LM}_{\text{ARCH}}(4)$	0.42	0.70	0.98	0.93	0.007

Notes: $R^2_{\text{adj.}}$ is the adjusted R^2, JB s the p-value of the Jarque-Bera normality test, $\text{LM}_{\text{SC}}(i)$ is the p-value of the Lagrange Multiplier (LM) test for ith order serial correlation in the estimated residuals, and $\text{LM}_{\text{ARCH}}(i)$ is the p-value of LM test for ith order ARCH effects in residuals. The numbers between parentheses below the estimates are the White´s heteroskedasticity-robust standard errors. The estimation period is 01 October 2001 – 11 April 2003.

4.2 Out-of-sample analysis

To evaluate the forecasting performance of the models estimated before, we conduct an out-of-sample experiment. Figure 6 shows the daily returns and the 95% confidence interval computed with the forecasted volatilities. The dashed lines represent the out-of-sample period. Table 11 shows the frequency of observations of the absolute returns that are greater than the 99%, 95%, 90% confidence intervals over the out-of-sample period. We also consider the combination of the realized volatility with the latent volatility models. Some conclusions emerge from the table. The GARCH, EGARCH, and GJR-GARCH models underestimate the confidence intervals in the forecasting period for all the confidence levels considered and for all series. Apart from the case of Vale for the 95% confidence level, the EWMA method seems to correctly forecast the coverage. The realized volatility performs better than the GARCH family but slightly underestimate the confidence intervals specially for the 99% level. When forecast combination is considered the results are greatly improved, with almost no difference between distinct combinations.

Table 11
Out-of-sample analysis: Frequency of observations of the daily absolute returns are greater than a 95% confidence interval

Asset	RV	EWMA ($\lambda = 0.94$)	GARCH	EGARCH	GJR-GARCH	RV + GARCH	RV + EGARCH	RV + GJR-GARCH
				Panel I: 99% Confidence Interval				
Bradesco	0.0187	0.0125	0	0	0	0.0063	0.0063	0.0063
Embratel	0.0187	0.0063	0.0063	0.0063	0.0063	0.0063	0.0063	0.0063
Petrobrás	0.0250	0.0063	0	0	0	0.0063	0.0063	0.0063
Telemar	0.0187	0.0125	0.0063	0.0063	0.0063	0.0125	0.0125	0.0125
Vale	0.0437	0.0125	0	0	0	0.0125	0.0063	0.0125
				Panel II: 95% Confidence Interval				
Bradesco	0.0563	0.0437	0.0250	0.0125	0.0250	0.0437	0.0437	0.0375
Embratel	0.0750	0.0437	0.0187	0.0250	0.0187	0.0313	0.0313	0.0437
Petrobrás	0.0750	0.0375	0.0313	0.0313	0.0250	0.0437	0.0375	0.0375
Telemar	0.0813	0.0437	0.0187	0.0187	0.0187	0.0563	0.0563	0.0563
Vale	0.0938	0.0813	0.0313	0.0313	0.0313	0.0563	0.0563	0.0563
				Panel III: 90% Confidence Interval				
Bradesco	0.1125	0.1063	0.0750	0.0563	0.0688	0.0938	0.0875	0.0875
Embratel	0.0875	0.0813	0.0500	0.0688	0.0500	0.0750	0.0680	0.0750
Petrobrás	0.1250	0.0813	0.0563	0.0500	0.0500	0.0875	0.0813	0.0750
Telemar	0.1375	0.0938	0.0625	0.0563	0.0625	0.0875	0.0813	0.0875
Vale	0.1375	0.1187	0.0813	0.0688	0.0750	0.1250	0.1125	0.1313

Notes: The estimation period is 01 October 2001 – 11 April 2003.

5. Conclusions

This paper aims at verifying whether or not stylized facts depicted in the literature for the realized variance in the US equity market hold for intensely-traded equities in the São Paulo Stock Exchange Market (BOVESPA). For this purpose, we analyzed empirically the statistical properties related to these stylized facts for five of the most actively traded stocks on BOVESPA, namely, Bradesco (BBDC4); Embratel (EBTP4); Petrobrás (PETR4); Telemar (TNLP4); and Vale do Rio Doce (VALE5). Intra-day data (prices observed every 15 minutes) from 10/01/2001 to 11/30/2003 are utilized in this analysis.

Two main results can be drawn from this analysis. First, when the intraday returns were utilized for estimating the daily variance, the standardized log-returns display a (nearly) Gaussian distribution, as opposed to when EWMA or GARCH models are employed to estimate the daily variance. This result can be used to improve Value at Risk estimates, particularly those of a parametric kind. Second, we find no evidence of long memory in the log of the realized variance. This second result stands in sharp contrast with one of the above-mentioned stylized facts, as the log of the realized variance tends to display strong evidence of long memory in the US stock market.

Also, an out-of-sample assessment of prediction intervals is carried out using a simple model, as well as standard methods and models in the literature. While using our model produces slightly undersized and using GARCH-type models slightly oversized forecast intervals, the plain EWMA yields forecast intervals with coverage closer to the nominal value. However, on average the EWMA intervals are less precise (larger) than those yielded by our linear model. On the

other hand, combining the realized variance approach with GARCH-type models improves the coverage of the forecast intervals to as close to the nominal coverage as the EWMA intervals.

References

Andersen, T. & Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39:885–906.

Andersen, T., Bollerslev, T., Diebold, F. X., & Labys, P. (2001a). The distribution of realized exchange rate volatility. *Journal of the American Statistical Association*, 96:42–55.

Andersen, T., Bollerslev, T., Diebold, F. X., & Labys, P. (2001b). Exchange rate returns standardized by realized volatility are (nearly) Gaussian. *Multinational Finance Journal*.

Andersen, T., Bollerslev, T., Diebold, F. X., & Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*, 71:579–625.

Bandi, F. & Russell, J. (2005a). Market microstructure noise, integrated variance estimators, and the limitations of asymptotic approximations: A solution. Unpublished manuscript, Graduate School of Business, University of Chicago.

Bandi, F. & Russell, J. (2005b). Microstructure noise, realized volatility, and optimal sampling. Unpublished manuscript, Graduate School of Business, University of Chicago.

Bandi, F. & Russell, J. (2006). Separating market microstructure noise from volatility. *Journal of Financial Economics*, 79:655–692.

Barndorff-Nielsen, O. & Shephard, N. (2002). Econometric analysis of realised volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society B*, 64:253–280.

Biais, B., Glosten, L., & Spatt, C. (2005). Market microstructure: A survey of microfoundations, empirical results, and policy implications. *Journal of Financial Markets*, 8:217–264.

Bollen, B. & Inder, B. (2002). Estimating daily volatility in financial markets utilizing intraday data. *Journal of Empirical Finance*, 9:551–562.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 21:307–328.

Bollerslev, T. (1987). A conditionally heteroskedastic time series model for speculative prices and rates of return. *The Review of Economics and Statistics*, 69:542–547.

Carnero, M. A., Peña, D., & Ruiz, E. (2004). Persistence and kurtosis in GARCH and stochastic volatility models. *Journal of Financial Econometrics*, 2:319–342.

Christoffersen, P. F. (1998). Evaluating interval forecasts. *International Economic Review*, 39:841–862.

Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of UK inflations. *Econometrica*, 50:987–1007.

Glosten, L., Jagannanthan, R., & Runkle, R. (1993). On the relationship between the expected value and the volatility of the nominal excess returns on stocks. *Journal of Finance*, 48:1779–1801.

Hansen, P. R. & Lunde, A. (2005). A forecast comparison of volatility models: Does anything beat a GARCH(1,1) model? *Journal of Applied Econometrics*, 20:873–889.

Hansen, P. R. & Lunde, A. (2006). Realized variance and market microstructure noise (with discussion). *Journal of Business and Economic Statistics*, 24:127–218.

Madhavan, A. (2000). Market microstructure: A survey. *Journal of Financial Markets*, 3:205–258.

Malmsten, H. & Teräsvirta, T. (2004). Stylized facts of financial time series and three popular models of volatility. Working Paper Series in Economics and Finance 563, Stockholm School of Economics.

Martens, M. (2001a). Forecasting daily exchange rate volatility using intraday returns. *Journal of International Money and Finance*, 20:1–23.

Martens, M. (2001b). Measuring and forecasting S&P 500 index-futures volatility using high-frequency data. *Journal of Futures Markets*, 22(22):497–518.

Martens, M., Van Dijk, D., & De Pooter, M. (2004). Modeling and forecasting S&P 500 volatility: Long memory, structural breaks and nonlinearity. Discussion Paper 04-067/4, Tinbergen Institute.

McAleer, M. (2005). Automated inference and learning in modeling financial volatility. *Econometric Theory*, 21:232–261.

Meddahi, N. (2002). A theoretical comparison between integrated and realized volatility. *Journal of Applied Econometrics*, 17:479–508.

Merton, R. (1980). On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics*, 8:323–361.

Morgan, J. P. (1996). *J. P. Morgan/Reuters Riskmetrics – Technical Document*. J. P. Morgan, New York.

Mota, B. & Fernandes, M. (2004). Desempenho de estimadores de volatilidade na bolsa de valores de São Paulo. *Revista Brasileira de Economia*, 58:429–448.

Nelson, D. B. (1991). Conditinal heteroskedasticity in asset returns: A new approach. *Econometrica*, 59:347–370.

Oomen, R. (2005). Properties of bias-corrected realized variance under alternative sampling schemes. *Journal of Financial Econometrics*, 3:555–577.

Patton, A. (2005). Volatility forecast evaluation and comparison using imperfect volatility proxies. Unpublished manuscript, London School of Economics.

Taylor, S. J. (1986). *Modelling Financial Time Series*. John Wiley, Chichester, UK.

Zhang, L., Mykland, P., & Aït-Sahalia, Y. (2005). A tale of two time scales: Determining integrated volatility with noisy high frequency data. *Journal of the American Statistical Association*, 100:1394–1411.