

TEXTO PARA DISCUSSÃO

No. 532

Asymmetric effects and long memory in
the volatility of Dow Jones stocks

Marcel Scharth
Marcelo C. Medeiros



ASYMMETRIC EFFECTS AND LONG MEMORY IN THE VOLATILITY OF DOW JONES STOCKS

MARCEL SCHARTH AND MARCELO C. MEDEIROS

ABSTRACT. Does volatility reflect a continuous reaction to past shocks or changes in the markets induce shifts in the volatility dynamics? In this paper, we provide empirical evidence that cumulated price variations convey meaningful information about multiple regimes in the realized volatility of stocks, where large falls (rises) in prices are linked to persistent regimes of high (low) variance in stock returns. Incorporating past cumulated daily returns as an explanatory variable in a flexible and systematic nonlinear framework, we estimate that falls of different magnitudes over less than two months are associated with volatility levels 20% and 60% higher than the average of periods with stable or rising prices. We show that this effect accounts for large empirical values of long memory parameter estimates. Finally, we analyze that the proposed model significantly improves out of sample performance in relation to standard methods. This result is more pronounced in periods of high volatility.

KEYWORDS: Realized volatility, long memory, nonlinear models, asymmetric effects, regime switching, regression trees, smooth transition, value-at-risk, forecasting, empirical finance.

ACKNOWLEDGMENTS: The authors are grateful to Caio Almeida, Dick van Dijk, Eric Hillebrand, Michael McAleer, Timo Teräsvirta, and Álvaro Veiga for useful discussion and to Leonardo Souza for providing ARFIMA codes and helpful comments. Financial support from CAPES and CNPq is thankfully acknowledged. Material from this paper has been presented at seminars at the Econometric Institute (Erasmus University, Rotterdam), Stockholm School of Economics, University of São Paulo, Pontifical Catholic University of Rio de Janeiro (PUC-Rio), and at the XXVIII Meeting of the Brazilian Econometric Society.

1. INTRODUCTION

Does stock return volatility reflect a long-lived reaction to past shocks or structural breaks induce shifts in the volatility dynamics? Long range dependence (highly persistent autocorrelations) is a well documented stylized fact of the volatility of financial time series. This effect was first analyzed by Taylor (1986) for absolute values of stock returns. Ding, Granger, and Engle (1993) and de Lima and Crato (1993) considered powers of returns. More recently, Andersen, Bollerslev, Diebold, and Ebens (2001) studied the case of realized volatility¹. Even though the traditional GARCH (Generalized Autoregressive Conditional

¹Realized variance is defined as the sum of squared intraday returns sampled at a sufficiently high frequency, consistently approximating the integrated variance over the fixed interval where the observations are summed. Realized volatility is the squared-root of the realized variance. In practice, high frequency measures are contaminated by microstructure noise such as bid-ask bounce, asynchronous trading, infrequent trading, price discreteness, among others; see Biais, Glosten, and Spatt (2005). Ignoring the remaining measurement error, this *ex post* volatility measure can be modeled as an “observable” variable, in contrast to the latent variable models. See Andersen, Bollerslev, Diebold, and Labys (2003) and Barndorff-Nielsen and Shephard (2002) for the theoretical foundations of realized volatility. Several recent papers have proposed corrections to estimation of RV in order to take the microstructure noise into account; see McAleer and Medeiros (in press) for a review. In this paper we refer to realized volatility as a consistent estimator of the squared root of the integrated variance.

Heteroscedasticity) models of Engle (1982) and Bollerslev (1986) are able to describe the recurrent clusters in volatility, the short run dynamics of those models were shown to be an incomplete description of the data. Volatility breeds volatility; but then could volatility today reflect a particularly volatile week a year ago? How do markets keep the memory of past movements?

Modeling the long range dependence in the volatility of stocks and foreign exchange rates is among one the greatest empirical successes of fractionally integrated models; see Baillie (1996) for an exposition. Fractionally integrated processes ($I(d)$, where $0 < d < 1$) can be seen as a halfway paradigm between the short memory ($I(0)$) process and the infinite memory ($I(1)$) alternative. Long memory processes are able to engender hyperbolic patterns in autocorrelations, as verified in many empirical applications. Although no theoretical foundation has been developed to substantiate the long memory specification or elucidate the high persistence from past shocks, $I(d)$ processes emerged as a consonant description of the data generating process of volatility series, becoming the standard approach for modeling and forecasting realized volatility (Andersen, Bollerslev, Diebold, and Labys 2003). Early models that account for the long memory in volatility are the Fractionally Integrated GARCH (FIGARCH) model proposed by Baillie, Bollerslev, and Mikkelsen (1996) and the long memory stochastic volatility (LMSV) model discussed in Comte and Renault (1996) and Breidt, Crato, and de Lima (1998).

More recently, new theoretical results clarified how long memory properties are not distinctive of fractionally integrated models. Diebold and Inoue (2001) showed analytically that stochastic regime switching is easily confused with long memory, even in large samples, as long as only a small amount of regime switches occurs in a observed sample path. Granger and Hyung (2004) showed that occasional structural breaks generate slowly decaying autocorrelations and other properties of $I(d)$ processes. Simulation results in both papers underline the relevance of those results in empirical applications; see also Mikosch and Starica (2004) and Hillebrand (2005).

However, the empirical question revealed itself elusive. While the new literature kindled a debate around the possibility that the long memory observed in the volatility of stocks and exchange rates is spurious, empirical studies evaluated that structural breaks cannot fully account for the degree of persistence in the data. This suggests that both long memory and structural changes can describe the volatility of asset returns (Lobato and Savin 1998, Martens, van Dijk, and de Pooter 2004, Beltratti and Morana 2006, Morana and Beltratti 2004, Hyung and Franses 2002). Nevertheless, the estimation of structural breaks mirrors the original difficulty: Fractional integration also biases common structural breaks detection methodologies, such as the one derived by Bai (1997), towards the detection of spurious breaks. Moreover, no satisfactory answers emerged from statistical hypothesis tests, which requires unrealistically large samples; see Ohanissian, Russell, and Tsay (2004).

In this paper, we propose a new empirical approach related to the hypothesis of structural changes and regime switches. We inquire how *changes in the markets* affect volatility. We provide empirical evidence that long-term price variations convey meaningful information about multiple regimes in the realized volatility of stocks, where large falls (rises) in prices are linked to persistent regimes of high (low) variance in stock returns. What happens for instance if returns were slightly positive in previous months

and then prices plummet in the next? From the asymmetric effects literature, it is known that negative returns are related to subsequent increases in volatility. Econometric models such as Nelson (1991)'s Exponential GARCH (EGARCH) and the GJR-GARCH of Glosten, Jagannathan, and Runkle (1993) have been proposed to capture this effect. Nevertheless, the literature so far focused almost exclusively on the relation observed over one or few days. For example, Andersen, Bollerslev, Diebold, and Ebens (2001) ran a regression with a lagged negative return dummy and conclude that the economic impact of the leverage effect on the realized variance of stocks belonging to the Dow Jones Industrial Average Index (DJIA) is marginal. An exception is Bollerslev, Litvinova, and Tauchen (2005), who examined evidence on the negative correlation between stock market movements and stock market volatility over intraday sampling frequencies. The authors show that a sharp decline in the market over a five-minute interval is typically associated with a rise in market volatility that persists for up to several days after the initial shock.

Focusing on realized volatility (RV) series of sixteen Dow Jones Industrial Average (DJIA) stocks over the period from 1994 to 2003, we consider the following questions: Are volatility levels the same in periods of significant losses for investors like the end of 2002 (the DJIA reached a 4 year bottom) and periods like the year 2003 (the DJIA went up 25%)? Can negative returns over some horizon be associated with regimes of higher volatility? We pursue the argument by incorporating past cumulated daily returns in the modeling framework of volatility series. If price variations matter, what are the magnitudes that can be associated to regime switching behavior? What are the relevant horizons? To tackle these considerations, our econometric strategy is developed around a flexible and systematic modeling cycle based on the tree-structured smooth transition regression model (STR-Tree) of da Rosa, Veiga, and Medeiros (2003) and Medeiros, da Rosa, and Veiga (2005).

Our main result shows that the effect of falls and rises in prices on volatility is in fact highly significant and accounts for the high fractional differencing parameter estimates, even in samples spanning several years. For example, we show that the daily volatility series of the IBM stock can be described by a nonlinear model where falls of different magnitudes over less than two months are associated with volatility levels 20% and 60% higher than the average of periods with stable or rising prices. Based on those findings, we propose a new model to describe and forecast realized volatility. When compared with alternative specifications with short and long memory, the model proposed in this paper has a superior forecasting performance, which is even more pronounced in periods of high volatility. A model that allow for smoothly changing parameters across time (in order to capture possible structural breaks) is also estimated. However, the regime switching mechanism controlled by past cumulated returns turns out to be statistically superior. The results are uniform across 15 of the 16 series considered in this paper.

Other economic connections to long memory and regime switching in volatility have been proposed before. Beltratti and Morana (2006) found a close association between structural breaks in stock market volatility and structural breaks in the volatility of macroeconomic variables such as M1 growth and the Federal Funds rate, relating the observed evidence to monetary policy reaction to the state of the business cycle. Previously, Hamilton and Susmel (1994) analyzed that the conditional variance process of the US stock market can be described by a switching regime model with three persistent states, where the high

volatility state is prompted by general business downturn. Kim and Kim (1996) have suggested that the switch to the more turbulent state may be caused by higher variance in a fad component of the returns, instead of fundamentals. In the context of fractional integration, Andersen and Bollerslev (1997) demonstrate that by interpreting the volatility as a mixture of numerous heterogeneous short-run information arrivals, the observed volatility process may exhibit long-run dependence.

Our objective is therefore to bring the stylized fact in volatility into a more meaningful empirical framework. If we can relate structural changes to our candidate variable, the econometric issue of spurious structural change detection loses importance. We highlight the importance of this aspect by reporting evidence that long memory processes are at least an incomplete description of the volatility process of stocks, where weak in-sample performance seems to be closely related to the empirical issue of the excessive variation in estimates of the fractional differencing parameter through time, first documented by Granger and Ding (1996).

On the pragmatical side, the advantage of our approach is that an endogenous financial variable is potentially a much more useful bridge to risk management and option pricing. In contrast to ARFIMA (Autoregressive Fractionally Integrated Moving Average) or structural breaks models, our modeling makes it possible to use estimated relations to project future volatility scenarios. Ohanissian, Russell, and Tsay (2004) showed the relevance of this aspect by simulating different models with long memory properties as “true” data generating processes and breaking down the consequences for option pricing. They documented significant pricing errors from missteps in the long memory specification.

The rest of the paper is structured as follows. Section two briefly discusses the tree-structured smooth transition regression model describing the inference procedures, model building strategy and estimation. In Section three, we describe the data, the specification of our model and present the estimations for models with structural breaks and asymmetric effects. The relation between asymmetric effects and long memory is investigated in Section four. Section five contain an analysis of point and value at risk forecasting performances. Section six concludes.

2. MODELING FRAMEWORK

In this section, we present the non-linear econometric model used in the paper. The discussion of the tree-structured smooth transition regression (STR-Tree) model is based on da Rosa, Veiga, and Medeiros (2003) and Medeiros, da Rosa, and Veiga (2005), where details and proofs can be found.

2.1. A Brief Introduction to Regression Trees. Let $\mathbf{x}_t = (x_{1t}, \dots, x_{qt})' \in \mathbb{X} \subseteq \mathbb{R}^q$ be a vector which contains q explanatory variables (covariates or predictor variables) for a continuous univariate response $y_t \in \mathbb{R}$, $t = 1, \dots, T$. Suppose that the relationship between y_t and \mathbf{x}_t follows a regression model of the form

$$(1) \quad y_t = f(\mathbf{x}_t) + \varepsilon_t,$$

where the function $f(\cdot)$ is unknown and, in principle, there are no assumptions about the distribution of the random term ε_t . A regression tree is a nonparametric model based on the recursive partitioning of

the covariate space \mathbb{X} , which approximates the function $f(\cdot)$ as a sum of local models, each of which is determined in $K \in \mathbb{N}$ different regions (partitions) of \mathbb{X} . The model is usually displayed in a graph which has the format of a binary decision tree with $N \in \mathbb{N}$ parent (or split) nodes and $K \in \mathbb{N}$ terminal nodes (also called leaves), and which grows from the root node to the terminal nodes. Usually, the partitions are defined by a set of hyperplanes, each of which is orthogonal to the axis of a given predictor variable, called the *split variable*. The most important reference in regression tree models is the Classification and Regression Trees (CART) approach put forward by Breiman, Friedman, Olshen, and Stone (1984). In this context, the local models are just constants.

To mathematically represent a regression-tree model, we introduce the following notation. The root node is at position 0 and a parent node at position j generates left- and right-child nodes at positions $2j + 1$ and $2j + 2$, respectively. Every parent node has an associated split variable $x_{s_j t} \in \mathbf{x}_t$, where $s_j \in \mathbb{S} = \{1, 2, \dots, q\}$. Furthermore, let \mathbb{J} and \mathbb{T} be the sets of indexes of the parent and terminal nodes, respectively. Then, a tree architecture can be fully determined by \mathbb{J} and \mathbb{T} .

EXAMPLE 1. Consider a regime switching volatility model that allows for multiple regimes associated with asymmetric effects, where the influence of a negative return on volatility for the next day depends on the behavior of returns on the past week. Define $r_{5,t}$ as the cumulated return over a horizon of five days and r_t as the daily return. Suppose the daily volatility (σ_t) follows a piecewise constant process where the conditional mean depends on the sign of the return in the previous day. This effect itself is weaker on “good weeks” (or a positive return over the last five days) than on “bad weeks” (or a negative return over the last five days), such that $\sigma_t = \omega_1 + \varepsilon_t$ if $r_{t-1} \geq 0$, $\sigma_t = \omega_2 + \varepsilon_t$ if $r_{t-1} < 0$ and $r_{5,t-1} \geq 0$ and $\sigma_t = \omega_3 + \varepsilon_t$ if $r_{t-1} < 0$ and $r_{5,t-1} < 0$. ε_t is a white noise, and $\omega_3 > \omega_2 > \omega_1$ are constants. This model can be described in the regression tree with two parent nodes at positions 0 and 2 ($N = 2$, $\mathbb{J} = \{0, 2\}$) and three leaves or terminal nodes at positions 1, 5 and 6 ($K = 3$, $\mathbb{T} = \{1, 5, 6\}$). See Figure 1.

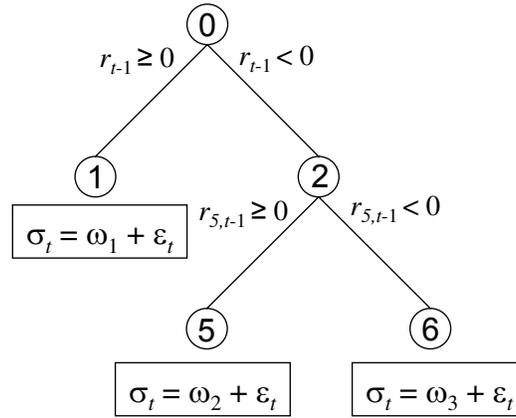


FIGURE 1. Graphical representation of the volatility model described in Example 1.

2.2. Tree-Structured Smooth Transition Regression. The STR-Tree model is an extension of the regression tree model, where the sharp splits are replaced by smooth splits given by a logistic function defined as

$$(2) \quad G(x; \gamma, c) = \frac{1}{1 + e^{-\gamma(x-c)}}.$$

The parameter γ , called the *slope parameter*, controls the smoothness of the logistic function. The regression tree model is nested in the smooth transition specification as a special case obtained when the slope parameter approaches infinity. The parameter c is called the *location parameter*.

Define $\log(RV_t)$ as the logarithm of the daily realized volatility. In this paper, $\log(RV_t)$ follows an augmented specification of the STR-Tree model defined as:

DEFINITION 1. Let $\mathbf{z}_t \subseteq \mathbf{x}_t$, such that \mathbf{x}_t is defined as in (1) and $\mathbf{z}_t \in \mathbb{R}^p$, $p \leq q$. The sequence of real-valued vectors $\{\mathbf{z}_t\}_{t=1}^T$ is stationary and ergodic. Set $\tilde{\mathbf{z}}_t = (1, \mathbf{z}_t)'$ and $\mathbf{w}_t \in \mathbb{R}^d$ is a vector of linear regressors, such that $\mathbf{w}_t \not\subseteq \mathbf{x}_t$. The time series $\{\log(RV_t)\}_{t=1}^T$ follows a Smooth Transition Regression Tree model, STR-Tree, if

$$(3) \quad \log(RV_t) = H_{\mathbb{J}\mathbb{T}}(\mathbf{x}_t, \mathbf{w}_t; \boldsymbol{\psi}) + \varepsilon_t = \boldsymbol{\alpha}'\mathbf{w}_t + \sum_{i \in \mathbb{T}} \boldsymbol{\beta}'_i \tilde{\mathbf{z}}_t B_{\mathbb{J}i}(\mathbf{x}_t; \boldsymbol{\theta}_i) + \varepsilon_t$$

where

$$(4) \quad B_{\mathbb{J}i}(\mathbf{x}_t; \boldsymbol{\theta}_i) = \prod_{j \in \mathbb{J}} G(x_{s_j, t}; \gamma_j, c_j)^{\frac{n_{i,j}(1+n_{i,j})}{2}} [1 - G(x_{s_j, t}; \gamma_j, c_j)]^{(1-n_{i,j})(1+n_{i,j})}$$

and

$$(5) \quad n_{i,j} = \begin{cases} -1 & \text{if the path to leaf } i \text{ does not include the parent node } j; \\ 0 & \text{if the path to leaf } i \text{ includes the right-child node of the parent node } j; \\ 1 & \text{if the path to leaf } i \text{ includes the left-child node of the parent node } j, \end{cases}$$

where $H_{\mathbb{J}\mathbb{T}}(\mathbf{x}_t, \mathbf{w}_t; \boldsymbol{\psi}) : \mathbb{R}^{q+1} \times \mathbb{R}^d \rightarrow \mathbb{R}$ is a nonlinear function indexed by the vector of parameters $\boldsymbol{\psi} \in \boldsymbol{\Psi}$ and $\{\varepsilon_t\}$ is a martingale difference sequence. Let \mathbb{J}_i be the subset of \mathbb{J} containing the indexes of the parent nodes that form the path to leaf i . Then, $\boldsymbol{\theta}_i$ is the vector containing all the parameters (γ_k, c_k) such that $k \in \mathbb{J}_i$, $i \in \mathbb{T}$.

The functions $B_{\mathbb{J}i}(\mathbf{x}_t; \boldsymbol{\theta}_i)$, $0 < B_{\mathbb{J}i}(\mathbf{x}_t; \boldsymbol{\theta}_i) < 1$, are known as *membership functions* and it is easy to show that $\sum_{i \in \mathbb{T}} B_{\mathbb{J}i}(\mathbf{x}_t; \boldsymbol{\theta}_i) = 1$, $\forall \mathbf{x}_t \in \mathbb{R}^{q+1}$.

The parameters of (3) are estimated by nonlinear least-squares (NLS) which is equivalent to quasi-maximum likelihood estimation. Let $\hat{\boldsymbol{\psi}}$ be the quasi-maximum likelihood estimator (QMLE) of $\boldsymbol{\psi}$ given by

$$(6) \quad \hat{\boldsymbol{\psi}} = \underset{\boldsymbol{\psi} \in \boldsymbol{\Psi}}{\operatorname{argmin}} Q_T(\boldsymbol{\psi}) = \underset{\boldsymbol{\psi} \in \boldsymbol{\Psi}}{\operatorname{argmin}} \frac{1}{T} \sum_{t=1}^T q_t(\boldsymbol{\psi}) = \underset{\boldsymbol{\psi} \in \boldsymbol{\Psi}}{\operatorname{argmin}} \left\{ \frac{1}{T} \sum_{t=1}^T [\log(RV_t) - H_{\mathbb{J}\mathbb{T}}(\mathbf{x}_t, \mathbf{w}_t; \boldsymbol{\psi})]^2 \right\}.$$

Under mild regularity conditions, Medeiros, da Rosa, and Veiga (2005) showed that $\widehat{\psi}$ is consistent and asymptotically normal.

2.3. Growing the Tree. In this section we briefly present the modeling cycle adopted in this paper. The choice of relevant variables, the selection of the node to be split (if this is the case), and the selection of the splitting (or transition) variable are carried out by sequence of Lagrange Multiplier (LM) tests following the ideas originally presented in Luukkonen, Saikkonen, and Terasvirta (1988) and vastly used in the literature.

Consider that $\log(RV_t)$ follows a STR-Tree model with K leaves and we want to test if the terminal node $i^* \in \mathbb{T}$ should be split or not. Write the model as

$$(7) \quad \begin{aligned} \log(RV_t) = & \alpha' \mathbf{w}_t + \sum_{i \in \mathbb{T} - \{i^*\}} \beta'_i \tilde{\mathbf{z}}_t B_{\mathbb{J}i}(\mathbf{x}_t; \boldsymbol{\theta}_i) \\ & + \beta'_{2i^*+1} \tilde{\mathbf{z}}_t B_{\mathbb{J}2i^*+1}(\mathbf{x}_t; \boldsymbol{\theta}_{2i^*+1}) + \beta'_{2i^*+2} \tilde{\mathbf{z}}_t B_{\mathbb{J}2i^*+2}(\mathbf{x}_t; \boldsymbol{\theta}_{2i^*+2}) + \varepsilon_t, \end{aligned}$$

where

$$\begin{aligned} B_{\mathbb{J}2i^*+1}(\mathbf{x}_t; \boldsymbol{\theta}_{2i^*+1}) &= B_{\mathbb{J}i^*}(\mathbf{x}_t; \boldsymbol{\theta}_{i^*}) G(x_{i^*t}; \gamma_{i^*}, c_{i^*}) \\ B_{\mathbb{J}2i^*+2}(\mathbf{x}_t; \boldsymbol{\theta}_{2i^*+2}) &= B_{\mathbb{J}i^*}(\mathbf{x}_t; \boldsymbol{\theta}_{i^*}) [1 - G(x_{i^*t}; \gamma_{i^*}, c_{i^*})]. \end{aligned}$$

In a more compact form, Equation (7) maybe written as

$$(8) \quad \begin{aligned} \log(RV_t) = & \alpha' \mathbf{w}_t + \sum_{i \in \mathbb{T} - \{i^*\}} \beta'_i \tilde{\mathbf{z}}_t B_{\mathbb{J}i}(\mathbf{x}_t; \boldsymbol{\theta}_i) \\ & + \phi'_1 \tilde{\mathbf{z}}_t B_{\mathbb{J}i^*}(\mathbf{x}_t; \boldsymbol{\theta}_{i^*}) + \phi'_2 \tilde{\mathbf{z}}_t B_{\mathbb{J}i^*}(\mathbf{x}_t; \boldsymbol{\theta}_{i^*}) G(x_{i^*t}; \gamma_{i^*}, c_{i^*}) + \varepsilon_t, \end{aligned}$$

where $\phi_1 = \beta_{2i^*+2}$ and $\phi_2 = \beta_{2i^*+1} - \beta_{2i^*+2}$.

In order to test the statistical significance of the split, a convenient null hypothesis is $\mathcal{H}_0 : \gamma_{i^*} = 0$ against the alternative $\mathcal{H}_a : \gamma_{i^*} > 0$. An alternative null hypothesis is $\mathcal{H}'_0 : \phi_2 = 0$. However, it is clear in (8) that under \mathcal{H}_0 , the nuisance parameters ϕ_2 and c_{i^*} can assume different values without changing the likelihood function, posing an identification problem; see Davies (1977, 1987).

A solution to this problem, proposed in Luukkonen, Saikkonen, and Terasvirta (1988), is to approximate the logistic function by a third-order Taylor expansion around $\gamma_{i^*} = 0$. After some algebra we get

$$(9) \quad \begin{aligned} \log(RV_t) = & \alpha' \mathbf{w}_t + \sum_{i \in \mathbb{T} - \{i^*\}} \beta'_i \tilde{\mathbf{z}}_t B_{\mathbb{J}i}(\mathbf{x}_t; \boldsymbol{\theta}_i) + \alpha'_0 \tilde{\mathbf{z}}_t B_{\mathbb{J}i^*}(\mathbf{x}_t; \boldsymbol{\theta}_{i^*}) \\ & + \alpha'_1 \tilde{\mathbf{z}}_t B_{\mathbb{J}i^*}(\mathbf{x}_t; \boldsymbol{\theta}_{i^*}) x_{i^*t} + \alpha'_2 \tilde{\mathbf{z}}_t B_{\mathbb{J}i^*}(\mathbf{x}_t; \boldsymbol{\theta}_{i^*}) x_{i^*t}^2 \\ & + \alpha'_3 \tilde{\mathbf{z}}_t B_{\mathbb{J}i^*}(\mathbf{x}_t; \boldsymbol{\theta}_{i^*}) x_{i^*t}^3 + e_t, \end{aligned}$$

where $e_t = \varepsilon_t + \phi_2 B_{\mathbb{J}i^*}(\mathbf{x}_t; \boldsymbol{\theta}_{i^*}) R(x_{i^*t}; \gamma_{i^*}, c_{i^*})$ and $R(x_{i^*t}; \gamma_{i^*}, c_{i^*})$ is the remainder. The parameters α_k , $k = 0, \dots, 3$ are functions of the original parameters of the model.

Thus the null hypothesis becomes

$$(10) \quad \mathcal{H}_0 : \alpha_1 = \alpha_2 = \alpha_3 = 0.$$

Under \mathcal{H}_0 , $R(x_{i^*t}; \gamma_{i^*}, c_{i^*}) = 0$ and $e_t = \varepsilon_t$, such that the properties of the error process remain unchanged under the null and thus asymptotic inference can be used. The test statistic is given by²:

$$(11) \quad LM = \frac{1}{\hat{\sigma}^2} \sum_{t=1}^T \hat{u}_t \hat{\nu}'_t \left\{ \sum_{t=1}^T \hat{\nu}_t \hat{\nu}'_t - \sum_{t=1}^T \hat{\nu}_t \hat{\mathbf{h}}'_t \left(\sum_{t=1}^T \hat{\mathbf{h}}_t \hat{\mathbf{h}}'_t \right)^{-1} \sum_{t=1}^T \hat{\mathbf{h}}_t \hat{\nu}'_t \right\}^{-1} \sum_{t=1}^T \hat{\nu}_t \hat{u}_t$$

where $\hat{u}_t = y_t - H_{\mathbb{J}\mathbb{T}}(\mathbf{x}_t, \mathbf{w}_t; \hat{\psi})$, $\hat{\sigma}^2 = \frac{1}{T} \sum_{t=1}^T \hat{u}_t^2$, $\hat{\mathbf{h}}_t = \frac{\partial H_{\mathbb{J}\mathbb{T}}(\mathbf{x}_t, \mathbf{w}_t; \psi)'}{\partial \psi} \Big|_{\gamma_{t_0}}$, and

$$\hat{\nu}_t = \left[\tilde{\mathbf{z}}_t B_{\mathbb{J}i^*}(\mathbf{x}_t; \hat{\boldsymbol{\theta}}_{i^*}) x_{i^*t}, \tilde{\mathbf{z}}_t B_{\mathbb{J}i^*}(\mathbf{x}_t; \hat{\boldsymbol{\theta}}_{i^*}) x_{i^*t}^2, \tilde{\mathbf{z}}_t B_{\mathbb{J}i^*}(\mathbf{x}_t; \hat{\boldsymbol{\theta}}_{i^*}) x_{i^*t}^3 \right]'$$

Under \mathcal{H}_0 , LM has an asymptotic χ^2 distribution with $m = 3(p+1)$ degrees of freedom.

As the assumption of normal and homoskedastic errors is usually violated in financial data, we carry out a robust version of the LM test, following the results in Wooldridge (1990). The test is implemented as follows:

- (1) Estimate the model with K regimes. If the sample size is small and the model is thus difficult to estimate, numerical problems in applying the maximum likelihood algorithm may lead to a solution such that the residual vector is not precisely orthogonal to the gradient matrix of $H_{\mathbb{J}\mathbb{T}}(\mathbf{x}_t, \mathbf{w}_t; \hat{\psi})$. This has an adverse effect on the empirical size of the test. To circumvent this problem, we regress the residuals \hat{u}_t on $\hat{\mathbf{h}}_t$ and compute the sum of squared residuals $SSR_0 = \sum_{t=1}^T \hat{u}_t^2$. The new residuals \tilde{u}_t are orthogonal to $\hat{\mathbf{h}}_t$.
- (2) Regress $\hat{\nu}_t$ on $\hat{\mathbf{h}}_t$ and compute the residuals \mathbf{r}_t .
- (3) Regress a vector of ones on $\tilde{\varepsilon}_t \mathbf{r}_t$ and calculate the sum of squared residuals SSR_1 .
- (4) The value of the test statistic is given by

$$(12) \quad LM_{\chi^2}^r = T - SSR_1.$$

Under H_0 , $LM_{\chi^2}^{hn}$ has an asymptotic χ^2 distribution with m degrees of freedom.

3. EMPIRICAL RESULTS

In this section we discuss how different specifications of the STR-Tree model actually describe the realized volatility series of DJIA stocks. Are there statistically significant structural breaks and/or regime shifts? What are the magnitudes and durations of those regimes? Are the level changes economically relevant? What do the estimation of structural breaks say about the stock market in the period? What are the in-sample fitting and out-of-sample forecasting properties of these models in relation to alternative models, such as the ARFIMA model?

The empirical analysis focuses on the realized volatility of sixteen Dow Jones Industrial Average index stocks: Alcoa (AA), American International Group (AIG), Boeing (BA), Caterpillar (CAT), General Electric (GE), General Motors (GM), Hewlett Packard (HP), IBM, Intel (INTC), Johnson and Johnson (JNJ),

²See Teräsvirta (1994) and Medeiros, da Rosa, and Veiga (2005) on the technical conditions for the validity of the test statistic.

Coca-Cola (KO), Merk (MRK), Microsoft (MSFT), Pfizer (PFE), Wal-Mart (WMT), and Exxon (XON). The raw intraday data are constituted of tick-by-tick quotes extracted from the NYSE Trade and Quote (TAQ) database. The period of analysis starts in January 3, 1994, and ends in December 31, 2003. Trading days with abnormally small trading volume and volatility caused by the proximity of holidays (for example, Good Friday) are excluded, leaving a total of 2541 daily observations.

We start by removing non-standard quotes, computing mid-quote prices, filtering possible errors, and obtaining one second returns for the 9:30 am to 16:05 p.m. period. Following the results of Hansen and Lunde (2006), we adopt the *previous tick* method for determining prices at precise time marks. Based on the results of Hasbrouck (1995), who reports a median 92.7% information share at the NYSE for Dow Jones stocks, and Blume and Goldstein (1997), who conclude that NYSE quotes match or determine the best displayed quote most of the time, we use NYSE quotes (or NASDAQ, for Microsoft and Intel) if they are close enough to the time marks in relation to other updates.

In order to estimate our measure of the daily realized volatility, we use the two time scales estimator of Zhang, Mykland, and Ait-Sahalia (2005) with five-minute grids, which is a consistent estimator of the daily realized volatility. The final dependent variable is the daily logarithm of the realized volatility. We also consider dummies for the days of the week as in Martens, van Dijk, and de Pooter (2004) and dummies for the following macroeconomic announcements: Federal Open Market Committee meetings (FOM), The Employment Situation Report from the Bureau of Labor Statistics (ESR), CPI and PPI.

In Section 3.1 we present the modeling cycle adopted in the empirical experiment. We carefully discuss variable selection and model specification. In order to evaluate the benefits of the STR-Tree model over standard models, we conduct a full sample study in Section 3.2, using data from 1994 to 2003. The goal of this analysis is to point out how the STR-Tree models may be useful to describe interesting stylized facts of financial time series, such as, long range dependence and asymmetries. We highlight our results to the particular case of the IBM stock. For all the others 15 stocks the results are rather similar and will be omitted for conciseness. Four versions of the STR-Tree model are estimated: A pure structural break model (STR-Tree/SB), where time is the single transition variable; an asymmetric effects model (STR-Tree/AE), where past cumulated returns of the stock over different horizons (reflecting “long-run” dynamics of the market) are the candidates for controlling the regime switches; an asymmetric effects model (STR-Tree/DJIA) where past cumulated returns of the DJIA index are used as transition variable; and finally, a combination of structural breaks and asymmetric effects model (STR-Tree/AE+SB), where both time and past cumulated returns are considered as split variables. We show that the asymmetric effects model successfully describe the long range dependence in the volatility of the stocks. Furthermore, using market returns (DJIA) or firm-specific returns causes no important difference in terms of in-sample performance. In-sample results are compared with the Heterogenous Autoregressive (HAR) model put forward by Corsi (2004) and the linear ARFIMA model.

In Section 3.3 we conduct an out-of-sample forecasting experiment, considering the last four years of the sample: From January 3, 2000 to December 31, 2003, covering 983 days. Each model is re-estimated daily using the full sample until that date and then used for point and value at risk forecasting for the

horizons of one, five, ten and 20 days ahead. The specification of the STR-Tree models is revised monthly. Point forecasts for the nonlinear models are calculated through conditional simulation, as well as interval forecasts for all models. For reference, we also include predictions from linear Autoregressive (AR), GARCH(1,1), and exponentially weighted moving average (EWMA) models. With respect to the latter, we take a different approach from the literature and compute an EWMA of the realized volatility itself. The STR-Tree/DJIA is not used to compute forecasts over one day ahead due the non-availability of the realized variance series for the index, which is essential in the conditional simulation.

3.1. Specification. Following the specific-to-general principle, we start the cycle from the root node (depth 0). Our general basic linear equation is given by:

$$(13) \quad \begin{aligned} \log(RV_t) = & \alpha_1 \log(RV_{t-1}) + \dots + \alpha_k \log(RV_{t-k}) + \delta_1 I[Mon]_t \\ & + \delta_2 I[Tue]_t + \delta_3 I[Wed]_t + \delta_4 I[Thu]_t + \delta_5 I[Fri]_t + \delta_6 I[FOMC]_t + \delta_7 I[EMP]_t \\ & + \delta_8 I[CPI]_t + \delta_9 I[PPI]_t + \varepsilon_t, \end{aligned}$$

where $I[Mon]_t$, $I[Tue]_t$, $I[Wed]_t$, $I[Thu]_t$, and $I[Fri]_t$ are days-of-the-week dummies; $I[FOMC]_t$, $I[EMP]_t$, $I[CPI]_t$, and $I[PPI]_t$ are dummies indicating dates for the following macroeconomic announcements: Federal Open Market Committee meetings, the Employment Situation report, CPI and PPI. Some authors discuss the relation between macroeconomic announcements and jumps; see, for example, Barndorff-Nielsen and Shephard (2006) and Huang (2006).

The first step in the modeling cycle is to use equation (13) to select the number of autoregressive lags and relevant days-of-the-week and announcement effects (variables that will be in \mathbf{w}_t), rendering the primary specification that will be contrasted against non-linearity. Autoregressive (AR) coefficients are tested up to the 15th order. Seeking a parsimonious specification, we base this selection on the Schwarz Information Criterion (SBIC), which initially selects autoregressive lags 1–3, 5, 7, 10 for all stocks, keeps the Monday dummy for some stocks and both the Monday and Friday dummies for others. The SBIC also selects the FOMC and EMP announcements. We verified that the inclusion of a moving average (MA) term could importantly cut down the number of AR terms, but we choose the less parsimonious AR specification since the computational burden for estimating an MA coefficient in a nonlinear framework is high and there are sufficient degrees of freedom. The presence of an MA coefficient could be justified by the existence of both persistent and non-persistent components in volatility, such as measurement noise or jump components³. We consider the importance of jump components in Section 3.3.3.

The next step is to select the set of variables in vectors \mathbf{x}_t and \mathbf{z}_t . Over the next sections, the candidate split variables \mathbf{z}_t falls in three cases: Structural breaks (time is the unique transition variable), asymmetric effects (lagged returns and lagged cumulated returns over past two to 120 days), and finally, the combination of structural breaks and asymmetric effects. A fourth possibility, explored by Martens, van Dijk, and de Pooter (2004), is the inclusion of lags of the realized volatility itself as split variables. However, this particular choice of asymmetry revealed not significant in all cases analyzed. At each node, the transition variable is selected as the one that minimizes the p -value of the robust version of the LM test.

³See Andersen, Bollerslev, and Diebold (2005) and Tauchen and Zhou (2005).

The elements of the vector \mathbf{z}_t are selected as a trade-off between parsimony/interpretability and fitting properties. In the structural break case we include the first two lags of the logarithm of the realized volatility, such that $\mathbf{z}_t = (\log(RV_{t-1}), \log(RV_{t-2}))'$. In the asymmetric effects model we set $\mathbf{z}_t = \emptyset$, such that $\tilde{\mathbf{z}}_t$ in Equation 3 is just a constant⁴.

3.2. Structural Breaks, Regime Switches and Long Memory: A Full Sample Evaluation. We start by following the recent literature and examining the effects of possible structural breaks on volatility levels (see, for example, Granger and Hyung (2004), Martens, van Dijk, and de Pooter (2004), Morana and Beltratti (2004)). The final estimated model for the case of IBM is given by

$$\begin{aligned} \log(RV_t) = & \underset{(0.164)}{0.261} \log(RV_{t-1}) + \underset{(0.078)}{0.224} \log(RV_{t-2}) + \underset{(0.021)}{0.084} \log(RV_{t-3}) \\ & + \underset{(0.020)}{0.074} \log(RV_{t-5}) + \underset{(0.019)}{0.044} \log(RV_{t-7}) + \underset{(0.018)}{0.047} \log(RV_{t-10}) \\ & - \underset{(0.013)}{0.064} I[Mon]_t - \underset{(0.014)}{0.063} I[Fri]_t + \underset{(0.032)}{0.067} I[FOMC]_t + \underset{(0.023)}{0.094} I[EMP]_t \\ & + \left\{ \underset{(0.048)}{0.005} + \underset{(0.164)}{0.261} \log(RV_{t-1}) + \underset{(0.078)}{0.224} \log(RV_{t-2}) \right\} \\ & \times G \left(t; \underset{(6.154)}{13.359}, \underset{(0.136)}{1.744} \right) G \left(t; \underset{(12.716)}{7.003}, \underset{(0.101)}{3.273} \right) \\ & + \left\{ \underset{(0.021)}{0.140} + \underset{(0.036)}{0.449} \log(RV_{t-1}) + \underset{(0.037)}{0.156} \log(RV_{t-2}) \right\} \\ & \times G \left(t; \underset{(6.154)}{13.359}, \underset{(0.136)}{1.744} \right) \left[1 - G \left(t; \underset{(12.716)}{7.003}, \underset{(0.101)}{3.273} \right) \right] \\ & + \left\{ \underset{(0.014)}{0.118} + \underset{(0.033)}{0.409} \log(RV_{t-1}) + \underset{(0.080)}{0.033} \log(RV_{t-2}) \right\} \\ & \times \left[1 - G \left(t; \underset{(6.154)}{13.359}, \underset{(0.136)}{1.744} \right) \right] \left[1 - G \left(t; \underset{(12.716)}{7.003}, \underset{(0.101)}{3.273} \right) \right] + \hat{\varepsilon}_t \end{aligned}$$

The final model has 23 estimated parameters. Although it may seem overparametrized, we stress the fact that we have a large number of observations. Two breaks are estimated: One in August 1998 (volatility and persistence go up; unconditional mean of the daily realized volatility goes from 1.50% to 2.10%, a 40% increase) and another one in April 2003 (volatility markedly falls; unconditional mean goes down from 2.10% to 1.15%, a 45% decrease). Note that the standard errors for the slope parameter estimates are quite high. Nevertheless, this is not an indication that the nonlinear effects are not significant. Due to the identification problem previously discussed in Section 2.3, the distribution for the usual t-statistic is not standard under $\mathcal{H}_0 : \gamma = 0$. The LM test is the adequate way to assess the statistical relevance of the structural changes; see Eitrheim and Terasvirta (1996) for a discussion.

⁴More general specifications of \mathbf{z}_t while statistically significant, brought no important out-of-sample gains, besides excessively increasing the number of estimated parameters and occasionally causing numerical problems in the estimation algorithm.

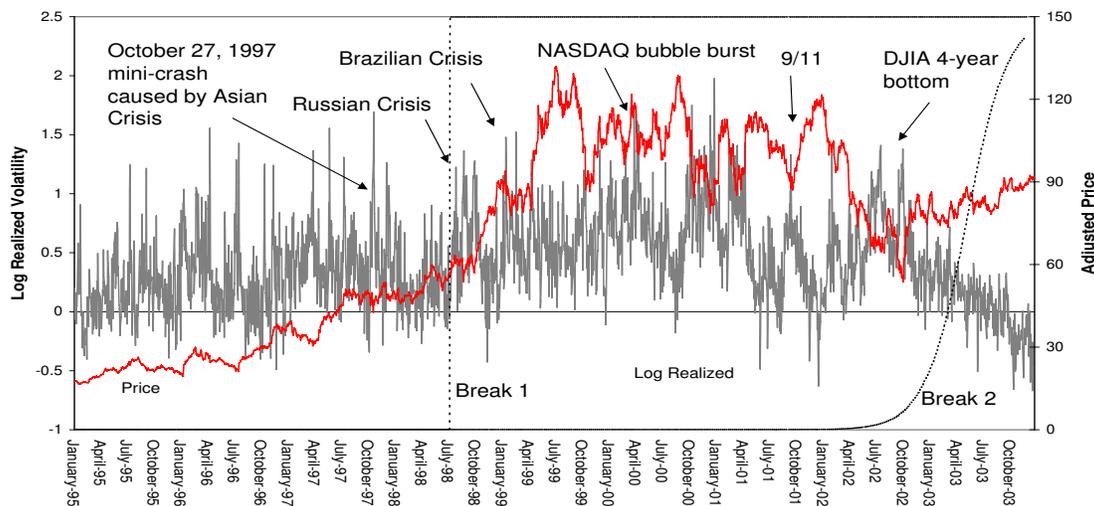
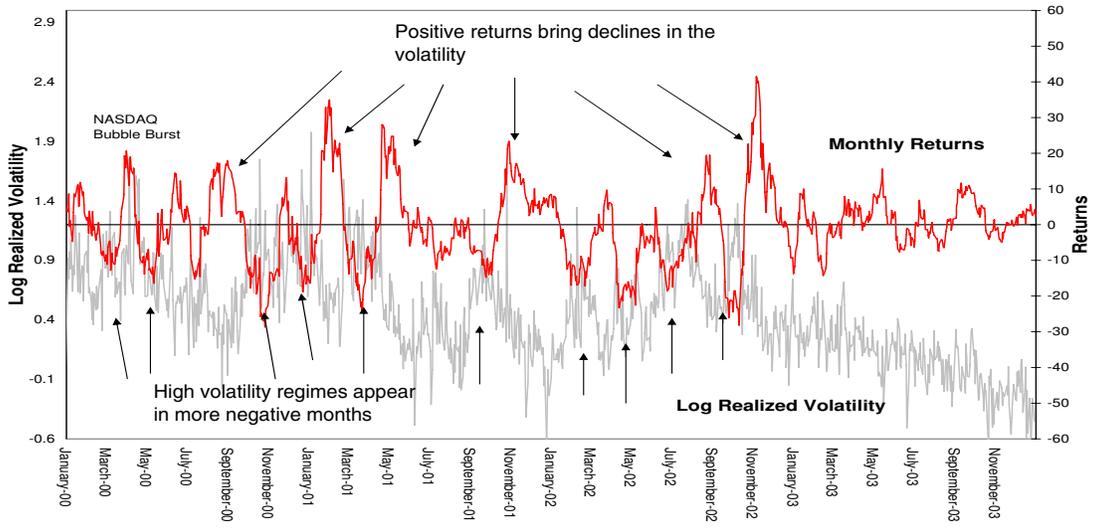


FIGURE 2. IBM daily log realized volatility (1995–2003) and the transition functions.

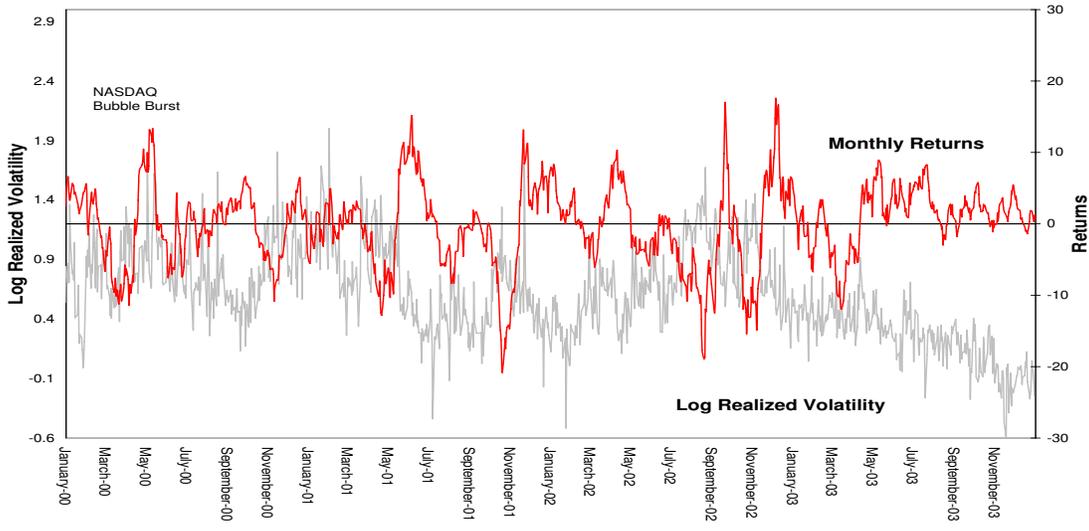
Figure 2 contextualize the timing of the breaks, depicting the two estimated transition functions, the log realized volatility for the period and the evolution of the stock price adjusted for dividends for the 1995–2003 period. The first break coincides exactly with the Russian Crisis in 1998, whilst the second one limits two distinct dynamics for the DJIA: While the index would reach a four year bottom by October 2002, the following year is a highly positive one for the index, which climbed 25% through the period. Figure 2 is suggestive of other similar relations: There are several clusters of high volatility associated with periods of large falls in the stock price, followed by sharp declines in volatility after the price jumps up again. Some examples are the periods of the October 1997 mini-crash, the Russian crisis, the NASDAQ bubble burst, the two clusters the end of 2000/beginning of 2001, the 9/11 period, and the bear market of 2002. The subsample between the first break and the second one (or the high volatility period) is marked by greater incidence of these price decreases. In the next section, we turn attention to this specific aspect.

3.2.1. *Asymmetric Effects.* The motivation for the estimation of lagged cumulated returns as a source of multiple regimes in volatility in the STR-Tree model is illustrated in Figure 3, which shows the realized volatility and monthly returns of IBM and the DJIA index for from 2000 to 2003. There seems to be a recurring pattern of shifts to higher volatility levels related to interludes of negative returns and reversals to low volatility levels in positive months. The single exception is the period just before the Nasdaq bubble burst.

As mentioned before, we estimate two asymmetric effects models: In the first one, past cumulated returns of the stock over different horizons are the candidates for controlling the regime switches (STR-Tree/AE) and the second one has past cumulated returns of the DJIA index as transition variables (STR-Tree/DJIA).



(a)



(b)

FIGURE 3. Panel (a): Realized volatility and monthly IBM returns. Panel (b): Realized volatility and monthly DJIA returns.

The estimated tree structure for the first model is shown in Figure 4 and is determined by the sets $\mathbb{T} = \{1, 6, 11, 23, 24\}$ and $\mathbb{J} = \{0, 2, 5, 12\}$.

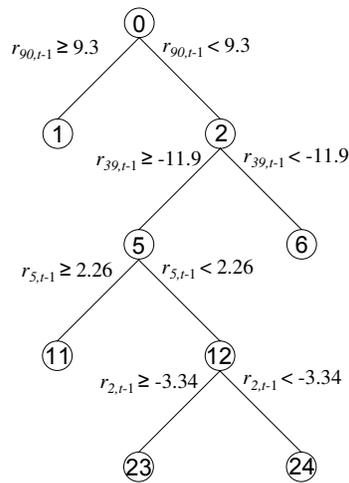


FIGURE 4. Estimated tree for IBM daily log realized volatility.

The final estimated STR-Tree/AE model is given by

$$\begin{aligned}
\log(RV_t) = & 0.386 \log(RV_{t-1}) + 0.118 \log(RV_{t-2}) + 0.107 \log(RV_{t-3}) \\
& (0.022) \quad (0.023) \quad (0.021) \\
& + 0.091 \log(RV_{t-5}) + 0.065 \log(RV_{t-7}) + 0.078 \log(RV_{t-10}) \\
& (0.020) \quad (0.019) \quad (0.018) \\
& - 0.068 I[Mon]_t - 0.064 I[Fri]_t + 0.068 I[FOMC]_t + 0.092 I[EMP]_t \\
& (0.012) \quad (0.014) \quad (0.032) \quad (0.023) \\
& + 0.081 \times G \left(r_{90,t-1}; 2.000, 0.541 \right) \\
& (0.013) \quad (1.082) \quad (0.344) \\
& + 0.184 \times \left[1 - G \left(r_{90,t-1}; 2.000, 0.541 \right) \right] \times \left[1 - G \left(r_{39,t-1}; 2.000, -0.955 \right) \right] \\
& (0.030) \quad (1.082) \quad (0.344) \quad (1.018) \quad (0.319) \\
& - 0.004 \times \left[1 - G \left(r_{90,t-1}; 2.000, 0.541 \right) \right] \times G \left(r_{39,t-1}; 2.000, -0.955 \right) \\
& (0.046) \quad (1.082) \quad (0.344) \quad (1.018) \quad (0.319) \\
& \quad \times G \left(r_{5,t-1}; 2.000, 0.479 \right) \\
& \quad (1.794) \quad (0.469) \\
& + 0.069 \times \left[1 - G \left(r_{90,t-1}; 2.000, 0.541 \right) \right] \times G \left(r_{39,t-1}; 2.000, -0.955 \right) \\
& (0.044) \quad (1.082) \quad (0.344) \quad (1.018) \quad (0.319) \\
& \quad \times \left[1 - G \left(r_{5,t-1}; 2.000, 0.479 \right) \right] \times G \left(r_{2,t-1}; 2.423, -1.091 \right) \\
& \quad (1.794) \quad (0.469) \quad (1.211) \quad (0.284) \\
& + 0.447 \times \left[1 - G \left(r_{90,t-1}; 2.000, 0.541 \right) \right] \times G \left(r_{39,t-1}; 2.000, -0.955 \right) \\
& (0.127) \quad (1.082) \quad (0.344) \quad (1.018) \quad (0.319) \\
& \quad \times \left[1 - G \left(r_{5,t-1}; 2.000, 0.479 \right) \right] \times \left[1 - G \left(r_{2,t-1}; 2.423, -1.091 \right) \right] + \widehat{\varepsilon}_t \\
& \quad (1.794) \quad (0.469) \quad (1.211) \quad (0.284)
\end{aligned}$$

Note that the transition variables are divided by their respective standard deviation.

The model is described by five highly statistically significant regimes determined by four levels of asymmetric effects. The first node indicates a low volatility regime linked to a rising market in the horizon of four months. On the other extreme, a decline of 12% or more over nearly two months introduce a regime of high variance, while superior returns over this same period bring intermediate volatility levels and short run leverage effects. Negative returns over two days also induce a regime of high variance. The estimated transition functions are illustrated in Figure 5.

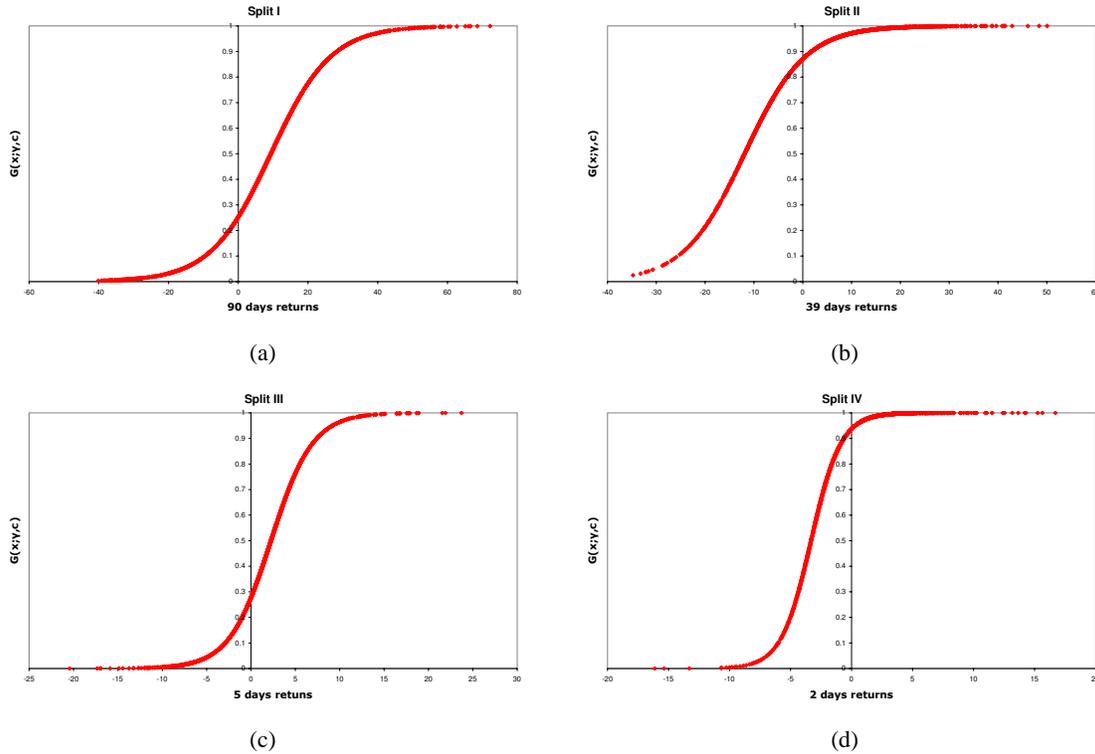


FIGURE 5. Estimated Transition Functions.

Based on the estimated regimes and the transition graphs displayed in Figure 5, we divide the observations in five different regimes. We split the observations according to the value of the transition functions (below or above 0.5). Table 1 reports the number of observations on each group and the respective mean and standard deviation of the realized volatility. Group one refers to the observations associated to the terminal node number one in Figure 4. Groups two and three include observations associated to the terminal node 11 and 23 (high returns, low volatility), respectively. Groups four and five relate to observations associated to nodes six and 24 (low returns, high volatility).

Concerning the STR-Tree/DJIA and the STR-Tree/SB+AE, the final estimated tree architectures are described in Figures 6 and 7.

TABLE 1. VOLATILITY REGIMES FOR IBM.

Mean and standard deviation of realized volatility for observations divided by a classification based on the STR-Tree/AE model with lagged cumulated returns as split variables.

Group	Mean	Standard Deviation	Number of Observations
1	1.57	0.54	1264
2	1.71	0.69	494
3	1.76	0.72	368
4	2.39	0.88	96
5	2.46	0.82	254
All	1.75	0.71	2476

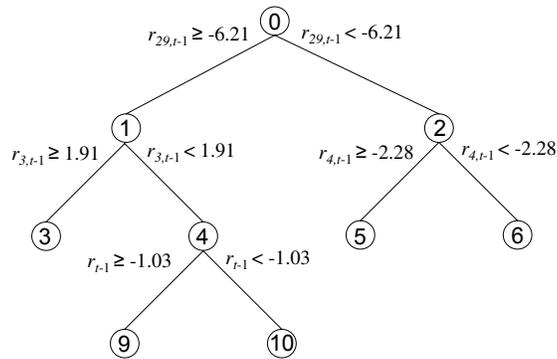


FIGURE 6. Estimated tree for IBM log realized volatility with cumulated returns of the DJIA index as transition variables.

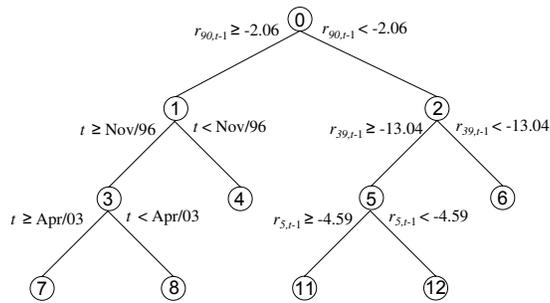


FIGURE 7. Estimated tree for IBM log realized volatility with cumulated returns of and time as transition variables.

3.2.2. *Autoregressive Fractionally Integrated Moving Average.* We now turn to the comparison of volatility models. We start with the standard ARFIMA(p, d, q) defined as

$$(14) \quad \phi_p(L)(1-L)^d(\log(RV_t) - \mu) = \theta_q(L)\varepsilon_t,$$

where d denotes the fractional differencing parameter, L the lag operator, ε_t is a white noise, $\phi_p(L)$ and $\theta_q(L)$ are polynomials of order p and q , having all roots lying outside the unit circle. For each series, we estimate several ARFIMA(p,d,q) specifications by maximum likelihood; see Baillie (1996). The best combination of p and q is selected by SBIC. The method leads to a choice of an ARFIMA(0, d ,0) for all series. Predictions for the ARFIMA(0, d ,0) model are computed through a truncation of the infinite autoregressive representation after the 150th lag. The final estimated model is given by:

$$(15) \quad (1 - L)^{0.516(0.057)} \left\{ \begin{aligned} &\log(RV_t) - \frac{0.502}{(0.009)} + \frac{0.059I[Mon]_t}{(0.019)} + \frac{0.032I[Fri]_t}{(0.020)} \\ &- \frac{0.100I[FOMC]_t}{(0.045)} - \frac{0.081I[EMP]_t}{(0.038)} \end{aligned} \right\} = \widehat{\varepsilon}_t$$

ARFIMA models have been estimated for realized volatility in Andersen, Bollerslev, Diebold, and Labys (2003), Areal and Taylor (2002), Beltratti and Morana (2005), Deo, Hurvich, and Lu (2006), Martens, van Dijk, and de Pooter (2004), Thomakos and Wang (2003), among others.

3.2.3. *Heterogenous Autoregressive.* The HAR (Heterogeneous Autoregressive) model proposed by Corsi (2004) is grounded on the Heterogeneous ARCH (HARCH) model developed by Müller, Dacorogna, Dav, Olsen, Pictet, and von Weizsacker (1997). It is specified as a multi-component volatility model with an additive hierarchical structure, leading to an additive time series model of the realized volatility which specifies the volatility as a sum of volatility components over different horizons. The model has been used for instance in Andersen, Bollerslev, and Diebold (2005) for its estimation simplicity and capacity to reproduce the autocorrelation patterns of long memory models over shorter horizons. Define the h -horizon normalized realized volatility by

$$(16) \quad \log(RV_t)_{t+h} = \frac{\log(RV_{t+1}) + \log(RV_{t+2}) + \dots + \log(RV_{t+h})}{h}$$

The estimated HAR model is given by:

$$(17) \quad \begin{aligned} \log(RV_t) = & - \frac{1.046}{(0.091)} + \frac{0.374 \log RV_{t-1}}{(0.023)} + \frac{0.068 \log RV_{t-2}}{(0.026)} + \frac{0.247 \log(RV_t)_{t-5}}{(0.046)} \\ & + \frac{0.225 \log(RV_t)_{t-22}}{(0.032)} - \frac{0.066I[Mon]_t}{(0.012)} - \frac{0.053I[Fri]_t}{(0.013)} \\ & + \frac{0.072I[FOMC]_t}{(0.029)} + \frac{0.093I[EMP]_t}{(0.025)} + \widehat{\varepsilon}_t \end{aligned}$$

We add a second order autoregressive term to the typical formulation of the model to account for remaining autocorrelation in small lags.

3.2.4. *Summary and Comparison of Results.* Table 2 shows summary statistics for the residuals of the four models, where JB is the p -value of the Jarque-Bera normality test, $Q(k)$ indicates the p -value of suitable tests of serial correlation up to the k^{th} lag (Ljung-Box portmanteau test for the ARFIMA and HAR models and a LM-type test for the nonlinear models; see a description of the latter in Medeiros and Veiga (2003)),

and $Q(k)^2$ gives the p -value of the same test for the squared residuals. The R^2 statistics are corrected according to Andersen, Bollerslev, and Meddahi (2005).

The table shows that the STR-Tree/AE model has superior in-sample fitting as measured by R^2 , while the STR-Tree/DJIA model is the best by the SBIC. The ARFIMA model has a remarkably inferior fitting performance than the others. All models generate highly skewed and leptokurtic residuals, which can be explained by forty outliers to the right of the distribution.

The $Q(k)$ statistics by their turn indicate all models with the exception of the ARFIMA model leave no significant remaining autocorrelation structure in the residuals up to the 20th lag at 5%. This could be due to ignored AR or MA terms in the ARFIMA, but less parsimonious models have been estimated and none of them was capable of reverting this result. Finally, there is strong evidence of dependence on squared residuals, but unlike the results of Beltratti and Morana (2005) for exchange rates, there is no indication whatsoever of long memory on the conditional variance of volatility.

TABLE 2. ESTIMATION DIAGNOSTICS.

The table shows summary statistics for the residuals of six different models estimated for the log realized volatility of IBM: The STR-Tree model with lagged cumulated returns as split variables (STR-Tree/AE), the STR-Tree model with time as the split variable (STR-Tree/SB), the STR-Tree model with time and cumulated returns as transition variables (STR-Tree/SB+AE), a STR-Tree model with cumulated returns of the DJIA index as transition variables (STR-Tree/DJIA), an ARFIMA(0, d ,0) model with exogenous variables and the HAR model. JB is the p -value of the Jarque-Bera normality test. $Q(k)$ indicates the p -value of adequate tests for serial correlation up to the k^{th} lag. $Q^2(k)$ gives the p -value of the same tests for the squared residuals. SBIC is the Schwarz Information Criterion. The R^2 is corrected as in Andersen, Bollerslev, and Meddahi (2005).

	STR-Tree/AE	STR-Tree/SB	STR-Tree/SB+AE	STR-Tree/DJIA	ARFIMA	HAR
R^2	0.631	0.619	0.624	0.621	0.505	0.610
SD	0.223	0.226	0.225	0.225	0.255	0.229
Skewness	0.697	0.725	0.707	0.736	0.336	0.707
Kurtosis	4.703	4.535	4.780	4.815	4.166	4.503
JB	0.000	0.000	0.000	0.000	0.000	0.000
$Q(5)$	0.367	0.432	0.382	0.189	0.000	0.637
$Q(10)$	0.115	0.308	0.157	0.079	0.000	0.275
$Q(20)$	0.399	0.422	0.432	0.101	0.000	0.530
$Q^2(10)$	0.012	0.001	0.006	0.032	0.000	0.001
$Q^2(20)$	0.032	0.008	0.041	0.086	0.000	0.008
SBIC	-2.905	-2.889	-2.918	-2.919	-2.699	-2.918

3.2.5. *Long Memory Analysis.* To assess the long memory characteristics of the estimated STR-Tree models for IBM, we run 1000 simulations of alternative models (with the same length as the sample) and evaluate estimates of the fractional differencing parameter (d). We also include AR simulations using the linear parameters of the STR-Tree/AE estimation to emphasize how the non-linear effects do generate hyperbolic patterns in autocorrelations beyond the possibly misleading effect of persistent autoregressive structures.

We apply two methods for the estimation of the long memory parameter: The widely used log periodogram estimator (GPH) of Geweke and Porter-Hudak (1983) and the bias reduced estimator of Andrews and Guggenberger (2003). We employ two values for the number of ordinates ℓ used in each regression: $T^{1/2}$, the usual rule of thumb value suggested by Geweke and Porter-Hudak (1983) (simulation-based), and the value selected by the plug-in method of Hurvich and Deo (1999), which points to $T^{0.65}$ for all series. T is the sample size.

For each set of simulations, we also evaluate the power of the Ohanissian, Russell, and Tsay (2004) test of true long memory process, which is based on the invariance property of the long memory parameter over temporal aggregation under the null. Andersen, Bollerslev, Diebold, and Ebens (2001), for example, examine this property for DJIA stocks as evidence of long memory.

Table 3 reports the mean and standard deviation (in parenthesis) of the fractional differencing parameter (d) estimates for the log realized volatility of IBM (entire sample) and over the simulations. The first line of the table reveals that the model with regime switching accounts for a large degree of long memory, even in large samples. In line with the literature, the same is also true for the model with structural breaks. The table also shows that the Ohanissian, Russell, and Tsay (2004) test has little power against these alternatives. For the log realized volatility series the test does not reject the null hypothesis, albeit sensibly to the specification (ℓ and the number of aggregations) and the sample itself. For instance, if the first week is removed from the sample, the test rejects the null of long memory at the 5% level. Unfortunately, the test can almost always be tailored to favor one of the alternatives.

Initially documented by Granger and Ding (1996), an important issue with the ARFIMA approach is the excessive variance of the fractional differencing parameter estimates over time, possibly involving extensive periods in non-stationary regions. This problem is illustrated in Figure 8, which shows the evolution of GPH estimates ($\ell = T^{0.65}$) in a rolling window of three years over our sample. The estimates range from around 0.3 to 0.8.

An interesting feature of the STR-Tree/AE model is that it can possibly account for this fact. We illustrate this through a partial simulation of the model using the actual return series as transition variables, dividing the sample by the first estimated break in model STR-Tree/SB. Even though this simulation is *ad hoc* and tends to underestimate the capacity of the model of generate persistent autocorrelations, it can provide an useful indication of this ability. Table 4 shows the results, including the estimate for the log realized volatility series. As suggested by Figure 8, all estimates for the log realized volatility point to a significantly lower estimative for the first part of sample. In fact, this is the source of the weak in-sample performance of the ARFIMA model analyzed in section 3.2.4 – the high d estimate for the overall series

TABLE 3. LOG-PERIDOGRAM ESTIMATES - SIMULATIONS AND LOG REALIZED VOLATILITY.

The table reports the mean and standard deviation (in parenthesis) of the fractional differencing parameter (d) estimates for IBM daily log realized volatility and over 1000 simulations of three models: The STR-Tree model with lagged cumulated returns as split variables (STR-Tree/AE), the STR-Tree model with time as the split variable (STR-Tree/SB), and the AR model. GPH and AG stand for the Geweke and Porter-Hudak (1983) and Andrews and Guggenberger (2003) estimators, respectively. The number of ordinates used in each regression is indicated in the first row. Two values for this parameter are employed: 0.5, the usual rule of thumb for the GPH, and 0.65, selected by the plug-in method of Hurvich and Deo (1999). The last column gives the results for the Ohanissian, Russell, and Tsay (2004) test of the null of a true long memory process: The first three numbers indicate the percentage of simulations where the null is rejected at the 5% level, while the last line indicates the p -value of the test for the Log Realized Volatility of IBM.

Model	$\ell = T^{0.5}$		$\ell = T^{0.65}$		$\ell = T^{0.7}$
	GPH	AG	GPH	AG	LMT
STR-Tree/AE	0.48 (0.15)	0.30 (0.25)	0.60 (0.08)	0.44 (0.17)	33.8%
STR-Tree/SB	0.42 (0.08)	0.51 (0.12)	0.50 (0.04)	0.42 (0.09)	25.5%
AR	0.14 (0.11)	0.02 (0.16)	0.38 (0.05)	0.11 (0.11)	94.5%
Log Realized Vol	0.60 (0.10)	0.35 (0.17)	0.46 (0.05)	0.59 (0.10)	0.556

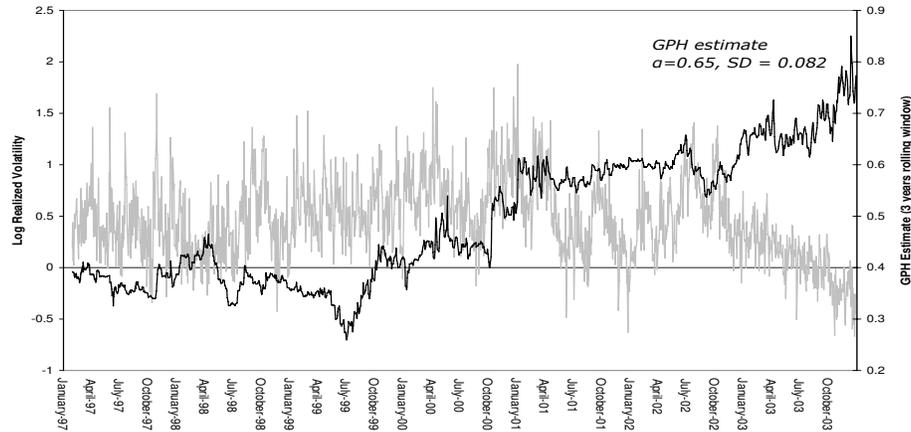


FIGURE 8. GPH Estimates in a Rolling Window.

produce large errors in the first subsample as well as dependence in the residuals (which are also induced by the period of the second break). Back to the table, although average estimates for the partial simulations are lower than the ones in the nonstationary region for the realized volatility in the second subsample, the model in fact seems to be able to reproduce this behavior.

3.3. Forecasting Analysis. We base the out-of-sample analysis on the four last years of the sample, ranging from January 3, 2000 to December 31, 2003, covering 983 days. Each model is re-estimated daily

TABLE 4. LOG-PERIODOGRAM ESTIMATES - PARTIAL SIMULATIONS AND LOG REALIZED VOLATILITY.

The table reports the mean and standard deviation (in parenthesis) of the fractional differencing parameter (d) estimates of two subsamples of the daily log realized volatility of IBM and over 1000 (partial) simulations of two models: the STR-Tree model with lagged cumulated returns as split variables (STR-Tree/AE) and the STR-Tree model with time as the split variable (STR-Tree/SB). GPH and AG stand for the Geweke and Porter-Hudak (1983) and Andrews and Guggenberger (2003) estimators, respectively.

Jan/1994 to Aug/1998	GPH ($\ell = T^{0.5}$)	GPH ($\ell = T^{0.65}$)	AG ($\ell = T^{0.65}$)
STR-Tree Partial Simulation	0.33 (0.13)	0.52 (0.07)	0.29 (0.14)
Log Realized Vol	0.34 (0.13)	0.29 (0.07)	0.36 (0.14)
Sep/1998 to Dec/2003	GPH ($\ell = T^{0.5}$)	GPH ($\ell = T^{0.65}$)	AG ($\ell = T^{0.65}$)
STR-Tree Partial Simulation	0.46 (0.11)	0.60 (0.06)	0.43 (0.12)
Log Realized Vol	0.65 (0.12)	0.66 (0.07)	0.74 (0.14)

using the full sample until that date and then used for point and value at risk forecasting for the horizons of one, five, ten and 20 days ahead. The specification of the STR-Tree models is revised monthly. Point forecasts for the nonlinear models are calculated through conditional simulation, as well as interval forecasts for all models. For reference, we also include predictions generated by a GARCH(1,1) model and an exponentially weighted moving average (EWMA). With respect to the latter, we take a different approach from the literature and compute an EWMA of the realized volatility with decay parameter set to 0.8.

3.3.1. *Point Forecasts.* The point forecasts results are reported in Tables 5. The evaluation of forecasts is based on the mean absolute error (MAE) criterion and the estimation of the Mincer-Zarnowitz regression

$$RV_t = \alpha + \beta \widehat{RV}_{t|t-1,i} + \varepsilon_{t,i}$$

where RV_t is the observed realized volatility on day t and $\widehat{RV}_{t|t-1,i}$ is the one-step-ahead forecast of model i for the volatility on day t . If the model i is correctly specified then $\alpha = 0$ and $\beta = 1$. We compute the (robust) p -value of the F test for this joint hypothesis and report the (corrected) R^2 of the regression as a measure of the ability of the model to track variance over time. However, the presence of heteroskedasticity hinders the computation of appropriate statistics for five, ten, and 20 days.

We also report two tests for superior predictive ability. The first one is the Harvey, Leybourne, and Newbold (1997) modification of the Diebold and Mariano (1995) test of equal predictive accuracy. Each concurrent model is compared against the ARFIMA model. Let $g(e_{1t})$ and $g(e_{2t})$ denote the loss function for the prediction errors e_{1t} and e_{2t} of models 1 and 2 on day t . For the MAE, $g(e_{it}) = |RV_t - \widehat{RV}_{t|t-j,i}|$ and for the R^2 , $g(e_{it}) = [RV_t - \widehat{RV}_{t|t-j,i}]^2$. The null hypothesis is $\mathbb{E}[g(e_{1t}) - g(e_{2t})] = 0$.

The second test is the Superior Predictive Ability (SPA) test developed by Hansen (2005). The null hypothesis is that a given model is not inferior to any other competing models in terms of a given loss function.

For one day-ahead forecasts, the STR-Tree/AE models are superlative both in terms of MAE and R^2 , significantly outperforming the ARFIMA model (with errors 5% smaller on average) and being the only ones not rejected by the SPA tests. In the sequence, there is little distinction between the ARFIMA, AR and HAR models in terms of R^2 , while the last two are slightly better in terms of MAE (the differences are significant at 10% and 5% respectively). The model with structural breaks is markedly inferior to those alternatives. The superiority of the model with asymmetric effects in terms of the MAE is reproduced in all stocks. When the R^2 is considered the STR-Tree/AE model is superior in 12 series (80%). The ARFIMA, HAR and EWMA models alternate as the second best in terms of R^2 , while the HAR specification has an edge in terms of the MAE; see also Table 8.

The advantage of the STR-Tree/AE model in terms of the MAE is preserved when the five days horizon is considered. The EWMA model significantly outperforms the ARFIMA model. The performance of the ARFIMA, HAR and AR models are relatively similar with respect to the MAE, with an advantage for fractional integration in R^2 . The results for ten and 20 days are similar: The STR-Tree/AE model is still the best in terms of the MAE, significantly exceeding the ARFIMA model, and being almost identical to the EWMA when R^2 is considered. However, the model with asymmetric effects and structural breaks become greatly superior in R^2 for 20 days forecasts. The null hypothesis of the SPA test is no longer rejected at 5% for ARFIMA, AR and EWMA specifications; HAR predictions come moderately behind.

Back to the other stocks, for ten days forecasts the STR-Tree/AE model is the best in MAE for twelve stocks, the EWMA model for two and the HAR for only one. As in one day forecast, neither ARFIMA, HAR or EWMA forecasts consistently appear as the second best, even though the latter achieves some advantage. On the other hand, a different pattern emerge for the R^2 : The EWMA model is the best in ten stocks, the STR-Tree/AE model in three, the HAR model in one and the STR-Tree/SB+AE model also in one.

TABLE 5. FORECASTING RESULTS.

The table reports the out-of-sample forecasting results of for the IBM daily realized volatility for the period between 2000 and 2003 (983 trading days, excluding days affected by holidays), where each model is re-estimated daily and used for predictions one, five, ten and 20 days ahead. MAE is the mean absolute error. R^2 is the (corrected) R-squared of $RV_t = \alpha + \beta \widehat{RV}_{t|t-j,i} + \varepsilon_{t,i}$, where $\widehat{RV}_{t|t-j,i}$ is the prediction of model i for the realized volatility on day t and RV_t is the observed realized volatility on that day. F is the p -value of the (heteroskedasticity robust) F test of the joint hypothesis that $\alpha = 0$ and $\beta = 1$. HLN is the p -value of the Harvey, Leybourne, and Newbold (1997) test of equality of the mean of loss functions, where the models are compared with the ARFIMA. SPA is the p -value of the Superior Predictive Ability test developed by Hansen (2005). The null hypothesis is that a given model is not inferior to any other competing models in terms of a given loss function. EWMA is the exponential weighted moving average of realized volatility itself.

	1 day							5 days						
	MAE	HLN	SPA	$\overline{R^2}$	HLN	SPA	F	MAE	HLN	SPA	$\overline{R^2}$	HLN	SPA	F
STR-Tree/AE	0.322	0.000	0.960	0.641	0.004	0.275	0.009	0.397	0.000	0.975	0.499	0.012	0.947	–
STR-Tree/SB	0.365	0.000	0.000	0.592	0.018	0.000	0.000	0.474	0.000	0.000	0.424	0.000	0.002	–
STR-Tree/DJIA	0.324	0.000	0.456	0.644	0.002	0.921	0.049	–	–	–	–	–	–	–
STR-Tree/SB+AE	0.340	0.485	0.004	0.610	0.304	0.011	0.938	0.409	0.185	0.285	0.495	0.071	0.841	–
HAR	0.332	0.027	0.026	0.618	0.418	0.003	0.000	0.412	0.338	0.038	0.468	0.068	0.026	–
ARFIMA	0.339	–	0.001	0.617	–	0.009	0.169	0.414	–	0.032	0.478	–	0.228	–
AR	0.334	0.092	0.001	0.616	0.497	0.004	0.000	0.410	0.215	0.020	0.467	0.066	0.021	–
EWMA	0.348	0.031	0.001	0.598	0.015	0.006	0.412	0.407	0.098	0.517	0.492	0.032	0.733	–
GARCH	0.490	0.000	0.000	0.368	0.000	0.000	0.002	0.527	0.000	0.000	0.289	0.000	0.000	–
	10 days							20 days						
	MAE	HLN	SPA	$\overline{R^2}$	HLN	SPA	F	MAE	HLN	SPA	$\overline{R^2}$	HLN	SPA	F
STR-Tree/AE	0.447	0.003	0.969	0.388	0.048	0.878	–	0.507	0.012	0.982	0.251	0.150	0.399	–
STR-Tree/SB	0.532	0.000	0.000	0.314	0.000	0.002	–	0.604	0.000	0.000	0.172	0.000	0.002	–
STR-Tree/DJIA	–	–	–	–	–	–	–	–	–	–	–	–	–	–
STR-Tree/SB+AE	0.460	0.321	0.446	0.392	0.072	0.826	–	0.510	0.025	0.890	0.288	0.004	0.777	–
HAR	0.466	0.311	0.039	0.353	0.025	0.058	–	0.535	0.160	0.001	0.227	0.149	0.122	–
ARFIMA	0.463	–	0.287	0.370	–	0.565	–	0.524	–	0.489	0.237	–	0.269	–
AR	0.458	0.249	0.131	0.359	0.092	0.067	–	0.518	0.253	0.269	0.230	0.230	0.122	–
EWMA	0.463	0.473	0.390	0.390	0.028	0.907	–	0.536	0.090	0.233	0.252	0.107	0.370	–
GARCH	0.555	0.000	0.000	0.230	0.000	0.000	–	0.591	0.000	0.000	0.149	0.000	0.008	–

We also examine the forecasting performance of the different models by year. After 2003 the volatility consistently and sharply declined through that period, inducing autocorrelations in the residuals of all models. The results for 2000–2002 are presented in Table 6, where we concentrate on the ARFIMA and STR-Tree/AE models only. In the table, one, two or three asterisks next to MAE and/or R^2 indicate that the model has statistically significantly lower MAE/sum of squared residuals by the Harvey, Leybourne, and Newbold (1997) test at the 10%, 5% and 1% levels, respectively.

In 2000, the STR-Tree/AE is superior for one and five days ahead forecasts (significant at 5%), while the criteria diverge for ten and 20 days: The ARFIMA outperforms the STR-Tree/AE in terms of the MAE and the reverse happens with the R^2 . The contradiction suggests a volatility level unaccounted for by the STR-Tree/AE estimations, which otherwise demonstrated superior capacity to track variations in volatility. In 2001 and 2002, however, the STR-Tree/AE consistently and strongly outperforms the ARFIMA model in all horizons and criteria.

The statistics for 2003 are given in Table 7. For one day forecasts, the performances of the AR, EWMA, STR-Tree/AE and HAR models are very similar and superior to ARFIMA, while the EWMA and HAR models have better MAE and the ARFIMA model higher R^2 for 20 days. MAEs are considerably smaller than in previous years, suggesting a lower variance of the log realized volatility in the period. In fact, 20 days forecasts for the ARFIMA model have lower MAE than one day forecasts in all the previous years. The table also shows that the STR-Tree/SB model is strongly outperformed by ARFIMA and EWMA in the period. The apparent contradiction posed by the weak performance of the break model can be seen in light of the analysis of Granger and Hyung (2004), who show that the prediction with structural breaks models tend to be weaker even if the true process is a break process: Since there is a lag in the detection of the break, moving average models perform better, a quality that is also shared by spurious ARFIMA estimations.

TABLE 6. FORECASTING RESULTS BY YEAR: 2000–2002.

The table reports the out-of-sample forecasting results of the STR-Tree/AE, STR-Tree/DJIA, and ARFIMA models for each year between 2000 and 2002, where each model is re-estimated daily and used for predictions one, five, ten and 20 days ahead. MAE is the mean absolute error. R^2 is the corrected R-squared of the following regression: $RV_t = \alpha + \beta \widehat{RV}_{t,i} + \varepsilon_{t,i}$, where $\widehat{RV}_{t,i}$ is the prediction of model i for the realized volatility on day t and RV_t is the “observed” realized volatility on that day. One, two or three asterisks next to the MAE and/or the R^2 indicate that the model has statistically significantly lower MAE/sum of squared residuals by the Harvey, Leybourne, and Newbold (1997) test at the 10%, 5% and 1% levels respectively.

	1 day					
	2000		2001		2002	
	MAE	R^2	MAE	R^2	MAE	R^2
ARFIMA	0.459	0.309	0.373	0.504	0.352	0.618
STR-Tree/AE	0.451	0.336**	0.350***	0.550***	0.328***	0.644***
STR-Tree/DJIA	0.451	0.335**	0.353***	0.556***	0.326***	0.652***
	5 days					
	2000		2001		2002	
	MAE	R^2	MAE	R^2	MAE	R^2
ARFIMA	0.536	0.129	0.465	0.390	0.454	0.357
STR-Tree/AE	0.547	0.153*	0.420***	0.405	0.428***	0.432***
STR-Tree/DJIA	–	–	–	–	–	–
	10 days					
	2000		2001		2002	
	MAE	R^2	MAE	R^2	MAE	R^2
ARFIMA	0.567***	0.082	0.537	0.233	0.525	0.190
STR-Tree/AE	0.608	0.095	0.485***	0.250	0.479***	0.288***
STR-Tree/DJIA	–	–	–	–	–	–
	20 days					
	2000		2001		2002	
	MAE	R^2	MAE	R^2	MAE	R^2
ARFIMA	0.605***	0.016	0.634	0.097	0.583	0.062
STR-Tree/AE	0.633	0.024	0.567***	0.114	0.529***	0.148***
STR-Tree/DJIA	–	–	–	–	–	–

TABLE 7. FORECASTING RESULTS BY YEAR: 2003.

The table reports the out-of-sample forecasting results for the IBM daily realized volatility for the year 2003, where each model is re-estimated daily and used for predictions one, five, ten and 20 days ahead. MAE is the mean absolute error. R^2 is the (corrected) R-squared of $RV_t = \alpha + \beta \widehat{RV}_{t|t-j,i} + \varepsilon_{t,i}$, where $\widehat{RV}_{t|t-j,i}$ is the prediction of model i for the realized volatility on day t and RV_t is the observed realized volatility on that day. F is the p -value of the (heteroskedasticity robust) F test of the joint hypothesis that $\alpha = 0$ and $\beta = 1$. HLN is the p -value of the Harvey, Leybourne, and Newbold (1997) test of equality of the mean of loss functions, where the models are compared with the ARFIMA. SPA is the p -value of the Superior Predictive Ability test developed by Hansen (2005). The null hypothesis is that a given model is not inferior to any other competing models in terms of a given loss function. EWMA is the exponential weighted moving average of realized volatility itself.

	1 day							20 days						
	MAE	HLN	SPA	$\overline{R^2}$	HLN	SPA	F	MAE	HLN	SPA	$\overline{R^2}$	HLN	SPA	F
STR-Tree/AE	0.157	0.002	0.907	0.598	0.067	0.923	0.000	0.236	0.000	0.000	0.482	0.008	0.055	–
STR-Tree/SB	0.201	0.000	0.000	0.573	0.418	0.339	0.000	0.524	0.000	0.000	0.456	0.002	0.004	–
STR-Tree/DJIA	0.161	0.028	0.001	0.599	0.066	0.949	0.000							
STR-Tree/SB+AE	0.165	0.169	0.010	0.573	0.441	0.069	0.000	0.297	0.005	0.000	0.457	0.000	0.020	–
HAR	0.156	0.000	0.951	0.593	0.032	0.900	0.010	0.191	0.000	0.848	0.478	0.000	0.027	–
ARFIMA	0.170	–	0.000	0.569	–	0.166	0.000	0.274	–	0.000	0.546	–	0.880	–
AR	0.159	0.000	0.011	0.589	0.076	0.450	0.003	0.207	0.000	0.000	0.478	0.000	0.021	–
EWMA	0.158	0.005	0.695	0.586	0.158	0.606	0.010	0.200	0.000	0.539	0.479	0.000	0.001	–
GARCH	0.322	0.000	0.000	0.413	0.001	0.000	0.000	0.527	0.000	0.000	0.276	0.000	0.004	–

TABLE 8. ONE-DAY-AHEAD FORECASTING RESULTS FOR ALL SERIES.

The table reports the out-of-sample forecasting results for the daily realized volatility of 15 Dow Jones stocks, where each model is re-estimated daily and used for predictions one, five, ten and 20 days ahead. MAE is the mean absolute error. R^2 is the (corrected) R-squared of $RV_t = \alpha + \beta \widehat{RV}_{t|t-j,i} + \varepsilon_{t,i}$, where $\widehat{RV}_{t|t-j,i}$ is the prediction of model i for the realized volatility on day t and RV_t is the observed realized volatility on that day. The figures between parenthesis are the p -value of the Superior Predictive Ability test developed by Hansen (2005). The null hypothesis is that a given model is not inferior to any other competing models in terms of a given loss function. EWMA is the exponential weighted moving average of realized volatility itself.

Series	STR-Tree/AE	STR-Tree/SB	STR-Tree/AE+SB	ARFIMA	HAR	EWMA
AA	0.450 (0.933)	0.476 (0.001)	0.560 (0.067)	0.456 (0.301)	0.474 (0.000)	0.465 (0.034)
AIG	0.359 (0.913)	0.371 (0.003)	0.372 (0.005)	0.364 (0.260)	0.369 (0.041)	0.371 (0.039)
BA	0.393 (0.837)	0.414 (0.002)	0.404 (0.057)	0.397 (0.469)	0.405 (0.099)	0.409 (0.063)
CAT	0.398 (0.904)	0.423 (0.000)	0.423 (0.000)	0.404 (0.152)	0.412 (0.044)	0.411 (0.063)
GE	0.340 (0.873)	0.369 (0.000)	0.363 (0.000)	0.349 (0.118)	0.355 (0.008)	0.361 (0.004)
GM	0.374 (0.920)	0.409 (0.000)	0.388 (0.003)	0.380 (0.181)	0.388 (0.003)	0.389 (0.007)
HP	0.574 (0.903)	0.604 (0.000)	0.585 (0.084)	0.579 (0.372)	0.599 (0.002)	0.583 (0.314)
INTC	0.436 (0.821)	0.490 (0.000)	0.443 (0.331)	0.448 (0.055)	0.459 (0.004)	0.466 (0.001)
JNJ	0.368 (0.806)	0.380 (0.015)	0.385 (0.008)	0.372 (0.579)	0.379 (0.139)	0.380 (0.113)
KO	0.335 (0.904)	0.360 (0.000)	0.346 (0.014)	0.341 (0.164)	0.348 (0.006)	0.339 (0.473)
MRK	0.367 (0.886)	0.389 (0.001)	0.377 (0.034)	0.370 (0.634)	0.381 (0.010)	0.378 (0.064)
MSFT	0.347 (0.827)	0.380 (0.000)	0.364 (0.000)	0.357 (0.133)	0.363 (0.013)	0.369 (0.013)
PFE	0.426 (0.036)	0.433 (0.001)	0.433 (0.002)	0.419 (0.893)	0.423 (0.531)	0.421 (0.669)
WMT	0.397 (0.882)	0.408 (0.026)	0.413 (0.008)	0.400 (0.690)	0.409 (0.036)	0.408 (0.125)
XON	0.306 (0.882)	0.322 (0.001)	0.323 (0.065)	0.312 (0.110)	0.323 (0.000)	0.321 (0.005)

3.3.2. *Value at Risk.* The evaluation of value-at-risk forecasts is based on the likelihood ratio tests for unconditional coverage and independence of Christoffersen (1998). Our analysis is similar to Beltratti and Morana (2005), who study the benefits of value-at-risk with long memory.

Initially, consider only one day forecasts. Let $\widehat{q}_{t|t-1}^i(\alpha)$ be the $(1 - \alpha)$ interval forecast of model i for day t conditional on information on day $t - 1$. In our application, we consider 95% and 99% value-at-risk measures, i.e., $\alpha = 0.05$ and $\alpha = 0.01$, respectively. We construct the sequence of coverage failures for the lower α tail as:

$$F_{t|t-1} = \begin{cases} 1 & \text{if } r_{t+1} < \widehat{q}_{t+1|t}^i(\alpha) \\ 0 & \text{if } r_{t+1} > \widehat{q}_{t+1|t}^i(\alpha) \end{cases}$$

where r_t is the return observed on day t . The unconditional coverage (UC) is a test of the null $\mathbb{E}(F_{t+1|t}) = \alpha$ against $\mathbb{E}(F_{t+1|t}) \neq \alpha$. The test of independence is constructed against a first-order Markov alternative.

TABLE 9. TEN-DAYS-AHEAD FORECASTING RESULTS FOR ALL SERIES.

The table reports the out-of-sample forecasting results for the daily realized volatility of 15 Dow Jones stocks, where each model is re-estimated daily and used for predictions one, five, ten and 20 days ahead. MAE is the mean absolute error. R^2 is the (corrected) R-squared of $RV_t = \alpha + \beta \widehat{RV}_{t|t-j,i} + \varepsilon_{t,i}$, where $\widehat{RV}_{t|t-j,i}$ is the prediction of model i for the realized volatility on day t and RV_t is the observed realized volatility on that day. The figures between parenthesis are the p -value of the Superior Predictive Ability test developed by Hansen (2005). The null hypothesis is that a given model is not inferior to any other competing models in terms of a given loss function. EWMA is the exponential weighted moving average of realized volatility itself.

Series	STR-Tree/AE	STR-Tree/SB	STR-Tree/AE+SB	ARFIMA	HAR	EWMA
AA	0.583 (0.927)	0.644 (0.003)	0.585 (0.865)	0.603 (0.064)	0.605 (0.175)	0.597 (0.368)
AIG	0.460 (0.856)	0.472 (0.249)	0.472 (0.238)	0.460 (0.874)	0.470 (0.469)	0.481 (0.122)
BA	0.512 (0.902)	0.551 (0.000)	0.530 (0.188)	0.523 (0.203)	0.515 (0.845)	0.533 (0.225)
CAT	0.499 (0.855)	0.541 (0.011)	0.520 (0.135)	0.507 (0.368)	0.517 (0.321)	0.513 (0.366)
GE	0.452 (0.884)	0.489 (0.000)	0.459 (0.616)	0.467 (0.015)	0.457 (0.776)	0.464 (0.415)
GM	0.455 (0.921)	0.526 (0.000)	0.469 (0.088)	0.481 (0.000)	0.486 (0.008)	0.483 (0.054)
HP	0.744 (0.659)	0.759 (0.327)	0.756 (0.347)	0.746 (0.560)	0.745 (0.486)	0.731 (0.727)
INTC	0.607 (0.903)	0.754 (0.000)	0.640 (0.030)	0.629 (0.014)	0.640 (0.127)	0.625 (0.408)
JNJ	0.456 (0.899)	0.477 (0.020)	0.485 (0.003)	0.460 (0.592)	0.468 (0.345)	0.487 (0.048)
KO	0.411 (0.892)	0.448 (0.001)	0.433 (0.006)	0.414 (0.574)	0.429 (0.134)	0.430 (0.121)
MRK	0.436 (0.891)	0.466 (0.000)	0.446 (0.189)	0.437 (0.753)	0.441 (0.625)	0.440 (0.541)
MSFT	0.505 (0.862)	0.551 (0.000)	0.525 (0.081)	0.510 (0.250)	0.512 (0.727)	0.520 (0.354)
PFE	0.500 (0.540)	0.536 (0.000)	0.495 (0.938)	0.508 (0.180)	0.506 (0.243)	0.510 (0.172)
WMT	0.524 (0.478)	0.536 (0.001)	0.527 (0.296)	0.518 (0.734)	0.519 (0.535)	0.511 (0.689)
XON	0.395 (0.899)	0.410 (0.099)	0.400 (0.516)	0.396 (0.797)	0.427 (0.001)	0.432 (0.000)

For five, ten, and 20 days value-at-risk forecasts, we still compute the value-at-risk on a daily basis. Since the overlapping returns cause the events to be correlated, we follow Beltratti and Morana (2005) and implement a test based on Bonferroni bounds suggested by Diebold, Gunther, and Tay (1998). For k -step forecasts, a test of size bounded by θ can be implemented by performing individual tests of size $\frac{\theta}{k}$ on each of k subseries

$$\{F_{1+k|1}, F_{1+2k|1+k}, F_{1+3k|1+2k}, \dots\}, \{F_{2+k|2}, F_{2+2k|2+k}, F_{2+3k|2+2k}, \dots\}, \dots, \\ \{F_{k-1+k|k-1}, F_{k-1+2k|k-1+k}, F_{k-1+3k|k-1+2k}, \dots\}$$

and rejecting the null if there is a rejection on any of the subseries. We also report the number of subseries that are rejected.

The value-at-risk comparison of the ARFIMA, structural breaks and asymmetric effects models is organized in Table 10, showing that all models adequately forecast the coverage intervals at all horizons.

TABLE 10. VALUE AT RISK ANALYSIS.

The table reports the out-of-sample value-at-risk results of the ARFIMA, STR-Tree/SB and STR-Tree/AE models for the IBM volatility in the 2000-2003 period (983 trading days, excluding days affected by holidays), where each model is re-estimated daily and used for calculating 1% and 5% value-at-risk thresholds by conditional simulation. Failures is the percentage of days when returns over the next 1, 5, 10 and 20 fell in the α lower tail of the predicted distribution. Note that 5, 10 and 20 days % failures are affected by overlapping return sequences. UC and IND are the p -values of the likelihood ratio tests for unconditional coverage and independence (against a first order Markov alternative) developed by Christoffersen (1998). For 5, 10 and 20 days, we use a test based on Bonferroni Bounds as suggested by Diebold, Gunther, and Tay (1998). R is the number of subseries (out of 1, 5, 10 or 20 accordingly) where UC is rejected at the 5% level.

	1 day							
	1%				5%			
	% Failures	UC	R	IND	% Failures	UC	R	IND
ARFIMA	0.006	0.186	0	0.786	0.043	0.284	0	0.134
STR-Tree/AE	0.010	0.957	0	0.650	0.054	0.578	0	0.931
STR-Tree/SB	0.009	0.787	0	0.683	0.043	0.284	0	0.388
	5 days							
	1%				5%			
	% Failures	UC	R	IND	% Failures	UC	R	IND
ARFIMA	0.011	1.000	0	1.000	0.058	1.000	0	1.000
STR-Tree/AE	0.010	1.000	0	1.000	0.059	1.000	0	1.000
STR-Tree/SB	0.010	1.000	0	1.000	0.048	1.000	0	0.423
	10 days							
	1%				5%			
	% Failures	UC	R	IND	% Failures	UC	R	IND
ARFIMA	0.021	0.994	0	0.576	0.082	0.154	2	0.630
STR-Tree/AE	0.015	0.994	0	1.000	0.079	0.395	1	0.697
STR-Tree/SB	0.017	1.000	0	0.576	0.070	0.370	1	0.327
	20 days							
	1%				5%			
	% Failures	UC	R	IND	% Failures	UC	R	IND
ARFIMA	0.032	0.289	1	1.000	0.102	0.277	8	1.000
STR-Tree/AE	0.027	1.000	0	1.000	0.111	0.134	8	1.000
STR-Tree/SB	0.028	0.289	2	1.000	0.082	0.311	2	1.000

3.3.3. *The Effect of Jumps.* Our analysis so far has not explicitly considered the presence of less persistent elements in the volatility of stocks, in contrast with the smooth and very slowly mean-reverting part associated with long memory properties. Jump components have been receiving growing attention in the realized volatility literature. Building on theoretical results for bi-power variation measures, articles such as Andersen, Bollerslev, and Diebold (2005), Tauchen and Zhou (2005), and Barndorff-Nielsen and Shephard (2006) established related frameworks for non-parametric estimation of the jump component in asset return volatility. Empirically, Andersen, Bollerslev, and Diebold (2005) incorporates the distinction between jump and non-jump components into a forecasting model for the DM/USD exchange rate, the S&P500 market

index, and the 30-year U.S. Treasury bond yield realized volatility series and find substantial performance improvements in daily weekly, and monthly predictions.

To verify the direct impact of the jump component for our conclusions, we closely follow Andersen, Bollerslev, and Diebold (2005) and recalculate the previous forecasts using the lagged jump series as an explanatory variable for the STR-Tree/AE and HAR models. The new results are displayed in Table 11. In sharp contrast with the results of in Andersen, Bollerslev, and Diebold (2005), the outcome of additionally considering jumps in the realized volatility of IBM is marginal; for instance, the R^2 of daily forecasts raise from 0.641 to 0.644 and from 0.618 to 0.621 for the STR-Tree/AE and HAR models respectively.

TABLE 11. FORECASTING RESULTS: JUMPS.

The table reports the out-of-sample forecasting results of for the IBM volatility in the 2000–2003 (983 trading days, excluding days affected by holidays), where each model explicitly incorporate jump components, is re-estimated daily and used for predictions 1, 5, 10 and 20 days ahead. MAE is the mean absolute error. \widehat{RV}_t is the corrected r-squared of the following regression: $RV_t = \alpha + \beta \widehat{RV}_{t,i} + \varepsilon_{t,i}$, where $\widehat{RV}_{t,i}$ is the prediction of model i for the realized volatility on day t and RV_t is the “observed” realized volatility on that day. HLN is the p -value of the Harvey, Leybourne and Newbold (1997) test of equality of the mean of loss functions (in the table, the absolute deviation and the residuals of the regression above), where the models are compared with the ARFIMA model.

	1 day					
	MAE	HLN	SPA	R^2	HLN	SPA
STR-Tree/AE	0.324	0.000	0.340	0.644	0.001	0.785
HAR	0.334	0.079	0.001	0.621	0.259	0.004
	5 days					
	MAE	HLN	SPA	R^2	HLN	SPA
STR-Tree/AE	0.398	0.000	0.793	0.500	0.005	0.968
HAR	0.410	0.198	0.004	0.472	0.194	0.007
	10 days					
	MAE	HLN	SPA	R^2	HLN	SPA
STR-Tree/AE	0.450	0.008	0.504	0.386	0.068	0.742
HAR	0.463	0.480	0.014	0.355	0.033	0.041
	20 days					
	MAE	HLN	SPA	R^2	HLN	SPA
STR-Tree/AE	0.450	0.008	0.504	0.386	0.068	0.742
HAR	0.463	0.480	0.014	0.355	0.033	0.041

4. CONCLUSION

In this paper, we considered the hypothesis that cumulated price variations convey essential information concerning shifts in the level of stock volatility series and can be related to multiple regimes that induce highly persistent autocorrelations that are hard to distinguish from the patterns generated by fractionally

integrated processes – even in sample sizes spanning several years. We showed, using realized volatilities computed from intraday returns, that volatility levels in periods of losses for investors like the end of 2002 (when the DJIA index reached a 4 year bottom) are significantly higher than periods like 2003, when the index went up 25%; there is strong evidence of multiple regimes linked to return patterns in all series considered. For the particular case of IBM, we show that falls of different magnitudes over less than two months are associated with volatility levels approximately 20% and 60% higher when compared to periods of stable or rising prices. Cumulated past returns over different horizons provide relevant information concerning regime switches in volatility dynamics. The result was robust to the choice of firm-specific or market returns as transition variables.

We underline the importance of this analysis by presenting further evidence that fractionally integrated processes are an incomplete description of the volatility process of stocks, arguing that weak in-sample performances are closely related to the empirical issue of excessive variation in estimates of the fractional differencing parameter over time.

Empirical results, by their turn, indicate that the multiple regime model proposed in the paper is superior in terms of forecasting performance, specially in periods of high volatility. In 15 of the 16 series considered in the paper, the STR-Tree model with past cumulated returns as transition variables significantly outperforms several concurrent models, such as the AR, ARFIMA, HAR, GARCH and EWMA models. Surprisingly, the EWMA model seems to be very competitive, specially in when volatility is low, such as in 2003.

REFERENCES

- ANDERSEN, T., AND T. BOLLERSLEV (1997): “Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns,” *Journal of Finance*, 52, 975–1005.
- ANDERSEN, T., T. BOLLERSLEV, AND F. X. DIEBOLD (2005): “Roughing it Up: Including Jump Components in the Measurement, Modeling and Forecasting of Return Volatility,” NBER Working Papers 11775, National Bureau of Economic Research.
- ANDERSEN, T., T. BOLLERSLEV, F. X. DIEBOLD, AND H. EBENS (2001): “The Distribution of Realized Stock Return Volatility,” *Journal of Financial Economics*, 61, 43–76.
- ANDERSEN, T., T. BOLLERSLEV, F. X. DIEBOLD, AND P. LABYS (2003): “Modeling and Forecasting Realized Volatility,” *Econometrica*, 71, 579–625.
- ANDERSEN, T., T. BOLLERSLEV, AND N. MEDDAHI (2005): “Correcting the Errors: A Note on Volatility Forecast Evaluation Based on High-Frequency Data and Realized Volatilities,” *Econometrica*, 73(1), 279–296.
- ANDREWS, D., AND P. GUGGENBERGER (2003): “A Bias-reduced Log-periodogram Regression Estimator of the Long Memory Parameter,” *Econometrica*, 71, 675–712.
- AREAL, N., AND S. R. TAYLOR (2002): “The Realized Volatility of FTSE-100 Futures Prices,” *Journal of Futures Markets*, 22, 627–648.
- BAI, J. (1997): “Estimating Multiple Breaks One at a time,” *Econometric Theory*, 13, 315–352.
- BAILLIE, R., T. BOLLERSLEV, AND H. O. MIKKELSEN (1996): “Fractionally integrated generalized autoregressive conditional heteroskedasticity,” *Journal of Econometrics*, 74, 3–30.
- BAILLIE, R. T. (1996): “Long Memory Processes and Fractional Integration in Econometrics,” *Journal of Econometrics*, 73(1), 5–59.
- BARNDORFF-NIELSEN, O., AND N. SHEPHARD (2002): “Econometric analysis of realized volatility and its use in estimating stochastic volatility models,” *Journal of the Royal Statistical Society B*, 64, 253–280.

- (2006): “Econometrics of Testing for Jumps in Financial Economics Using Bipower Variation,” *Journal of Financial Econometrics*, 4, 1–30.
- BELTRATTI, A., AND C. MORANA (2005): “Statistical Benefits of Value-At-Risk with Long Memory,” *Journal of Risk*, 7(4).
- (2006): “Breaks and Persistency: Macroeconomic Causes of Stock Market Volatility,” *Journal of Econometrics*, 131(1-2), 151–177.
- BIAIS, B., L. GLOSTEN, AND C. SPATT (2005): “Market microstructure: a survey of microfoundations, empirical results, and policy implications,” *Journal of Financial Markets*, 8, 217–264.
- BLUME, M., AND M. GOLDSTEIN (1997): “Quotes, Order Flow, and Price Discovery,” *Journal of Finance*, 52, 221–244.
- BOLLERSLEV, T. (1986): “Generalized Autoregressive Conditional Heteroskedasticity,” *Journal of Econometrics*, 21, 307–328.
- BOLLERSLEV, T., J. LITVINOVA, AND G. TAUCHEN (2005): “Leverage and Volatility Feedback Effects in High-Frequency Data,” Working paper, Duke University.
- BREIDT, F., N. CRATO, AND P. DE LIMA (1998): “The detection and estimation of long memory in stochastic volatility,” *Journal of Econometrics*, 83, 325–348.
- BREIMAN, L., J. H. FRIEDMAN, R. A. OLSHEN, AND C. STONE (1984): “Classification and Regression Trees,” Discussion paper, Wadsworth: Belmont, CA.
- CHRISTOFFERSEN, P. (1998): “Evaluating Interval Forecasts,” *International Economic Review*, 39, 841–62.
- COMTE, F., AND E. RENAULT (1996): “Long memory continuous time models,” *Journal of Econometrics*, 73, 101–149.
- CORSI, F. (2004): “A Simple Long Memory Model of Realized Volatility,” Manuscript, University of Southern Switzerland.
- DA ROSA, J., A. VEIGA, AND M. C. MEDEIROS (2003): “Tree-structured smooth transition regression models based on CAR-Talgorithm,” Discussion paper 469, Pontifical Catholic University of Rio de Janeiro, Department of Economics.
- DAVIES, R. (1977): “Hypothesis testing when a nuisance parameter is present only under the alternative,” *Biometrika*, 64, 247–254.
- (1987): “Hypothesis testing when a parameter is present only under the alternative,” *Biometrika*, 74, 33–43.
- DE LIMA, P., AND N. CRATO (1993): “Long range dependence in the conditional variance of stock returns,” Discussion paper, Proceedings of the Business and Economics Statistics Section, Joint Statistical Meetings.
- DEO, R., C. HURVICH, AND Y. LU (2006): “Forecasting Realized Volatility Using a Long-Memory Stochastic Volatility,” *Journal of Econometrics*, 131, 29–58.
- DIEBOLD, F., AND A. INOUE (2001): “Long Memory and Regime Switching,” *Journal of Econometrics*, 105, 131–159.
- DIEBOLD, F. X., T. GUNTHER, AND A. TAY (1998): “Evaluating density forecasts,” *International Economic Review*, 39, 863–83.
- DIEBOLD, F. X., AND R. S. MARIANO (1995): “Comparing Predictive Accuracy,” *Journal of Business and Economic Statistics*, 13, 253–263.
- DING, Z., C. GRANGER, AND R. ENGLE (1993): “A long memory property of stock market returns and a new model,” *Journal of Empirical Finance*, 1, 83–106.
- EITRHEIM, O., AND T. TERASVIRTA (1996): “Testing the adequacy of smooth transition autoregressive models,” *Journal of Econometrics*, 74, 59–75.
- ENGLE, R. F. (1982): “Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation,” *Econometrica*, 50, 987–1007.
- GEWEKE, J., AND S. PORTER-HUDAK (1983): “The estimation and application of long memory time series models,” *Journal of Time Series Analysis*, 4, 221–238.
- GLOSTEN, L. R., R. JAGANNATHAN, AND D. E. RUNKLE (1993): “On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks,” *Journal of Finance*, 48(5), 1779–1801.
- GRANGER, C., AND Z. DING (1996): “Varieties of Long Memory Models,” *Journal of Econometrics*, 73, 61–77.
- GRANGER, C. W. J., AND N. HYUNG (2004): “Occasional Structural Breaks and Long Memory with an Application to the S P 500 Absolute Stock Returns,” *Journal of Empirical Finance*, 11(3), 399–421.

- HAMILTON, J., AND R. SUSMEL (1994): "Autoregressive Conditional Heteroskedasticity and Changes in Regime," *Journal of Econometrics*, 64, 307–33.
- HANSEN, P. (2005): "A Test for Superior Predictive Ability," *Journal of Business and Economic Statistics*, 23(4), 365–380.
- HANSEN, P. R., AND A. LUNDE (2006): "Realized variance and market microstructure noise (with discussion)," *Journal of Business and Economic Statistics*, 24, 127–218.
- HARVEY, D., S. LEYBOURNE, AND P. NEWBOLD (1997): "Testing the Equality of Mean Squared Errors," *International Journal of Forecasting*, 13, 281–291.
- HASBROUCK, J. (1995): "One security, many markets: determining the contributions to price discovery," *Journal of Finance*, 50, 1175–1198.
- HILLEBRAND, E. (2005): "Neglecting Parameter Changes in GARCH Models," *Journal of Econometrics*, 129, 121–138.
- HUANG, X. (2006): "Macroeconomic News Announcements, Financial Market Volatility and Jumps," Working paper, Duke University.
- HURVICH, C., AND R. DEO (1999): "Plug-in Selection of the Number of Frequencies in Regression Estimates of the Memory Parameter of a Long-Memory Time Series," *Journal of Time Series Analysis*, 20, 331–341.
- HYUNG, N., AND P. FRANSES (2002): "Inflation rates: long-memory, level shifts, or both?," Report 2002-08, Econometric Institute, Erasmus University.
- KIM, C., AND M. KIM (1996): "Transient Fads and the Crash of '87," *Journal of Applied Econometrics*, 11, 41–58.
- LOBATO, I. N., AND N. E. SAVIN (1998): "Real and Spurious Long-Memory Properties of Stock-Market Data," *Journal of Business and Economic Statistics*, 16, 261–268.
- LUUKKONEN, R., P. SAIKKONEN, AND T. TERASVIRTA (1988): "Testing linearity against smooth transition autoregressive models," *Biometrika*, 75, 491–499.
- MARTENS, M., D. VAN DIJK, AND M. DE POOTER (2004): "Modeling and Forecasting S&P 500 Volatility: Long Memory, Structural Breaks and Nonlinearity," Discussion Paper 04-067/4, Tinbergen Institute.
- MCALÉER, M., AND M. MEDEIROS (in press): "Realized Volatility: A Review," *Econometric Reviews*.
- MEDEIROS, M., J. DA ROSA, AND A. VEIGA (2005): "Tree-Structured smooth transition regression models," Discussion paper 510, Pontifical Catholic University of Rio de Janeiro, Department of Economics.
- MEDEIROS, M., AND A. VEIGA (2003): "Diagnostic Checking in a Flexible Nonlinear Time Series Model," *Journal of Time Series Analysis*, 24(4), 461–482.
- MIKOSCH, T., AND C. STARICA (2004): "Nonstationarities in Financial Time Series, the Long-Range Dependence, and the IGARCH Effects," *Review of Economics and Statistics*, 86, 378–390.
- MÜLLER, U., M. DACOROGNA, R. DAV, R. OLSEN, O. PICTET, AND J. VON WEIZSACKER (1997): "Volatilities of Different Time Resolutions – Analysing the Dynamics of Market Components," *Journal of Empirical Finance*, 4, 213–239.
- MORANA, C., AND A. BELTRATTI (2004): "Structural Change and Long Range Dependence in Volatility of Exchange Rates: Either, Neither or Both?," *Journal of Empirical Finance*, 11(4), 629–658.
- NELSON, D. (1991): "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica*, 59(2), 347–70.
- OHANISSIAN, A., J. RUSSELL, AND R. TSAY (2004): "True or Spurious Long Memory in Volatility: Does it Matter for Pricing Options?," Working paper, Graduate School of Business, University of Chicago.
- TAUCHEN, G., AND H. ZHOU (2005): "Identifying Realized Jumps on Financial Markets," Working paper, Duke University, Department of Economics.
- TAYLOR, S. (1986): *Modelling Financial Time Series*. John Wiley, Chichester, UK.
- TERÄSVIRTA, T. (1994): "Specification, Estimation, and Evaluation of Smooth Transition Autoregressive Models," *Journal of the American Statistical Association*, 89, 208–218.
- THOMAKOS, D., AND T. WANG (2003): "Realized Volatility in the Futures Market," *Journal of Empirical Finance*, 10, 321–353.
- WOOLDRIDGE, J. (1990): "A Unified Approach to Robust, Regression-Based Specification Tests," *Econometric Theory*, 6(1), 17–43.

Departamento de Economia PUC-Rio
Pontifícia Universidade Católica do Rio de Janeiro
Rua Marques de São Vicente 225 - Rio de Janeiro 22453-900, RJ
Tel.(21) 31141078 Fax (21) 31141084
www.econ.puc-rio.br
flavia@econ.puc-rio.br