



Universität St.Gallen

Modeling Tick-by-Tick Realized Correlations

Francesco Audrino and Fulvio Corsi

January 2008 Discussion Paper no. 2008-05

Editor:

Prof. Jörg Baumberger
University of St. Gallen
Department of Economics
Bodanstr. 1
CH-9000 St. Gallen
Phone +41 71 224 22 41
Fax +41 71 224 28 85
Email joerg.baumberger@unisg.ch

Publisher:

Department of Economics
University of St. Gallen
Varnbühlstrasse 19
CH-9010 St. Gallen
Phone +41 71 224 23 25
Fax +41 71 224 22 98

Electronic Publication:

<http://www.vwa.unisg.ch>

Modeling Tick-by-Tick Realized Correlations

Francesco Audrino and Fulvio Corsi

Author's address:

Prof. Dr. Francesco Audrino
Institute of Mathematics and Statistics
Bodanstrasse 6
9000 St. Gallen
Tel. +41 71 224 2431
Fax +41 71 224 2894
Email francesco.audrino@unisg.ch
Website www.istituti.usilu.net/audrinof/web/

Author's address:

Dr. Fulvio Corsi
Swiss Finance Institute
Via Buffi 13
CH-6904 Lugano
Tel. +41 58 666 4243
Fax +41 58 666 46 47
Email fulvio.corsi@lu.unisi.ch

Abstract

We propose a tree-structured heterogeneous autoregressive (tree-HAR) process as a simple and parsimonious model for the estimation and prediction of tick-by-tick realized correlations. The model can account for different time and other relevant predictors' dependent regime shifts in the conditional mean dynamics of the realized correlation series. Testing the model on S&P 500 and 30-year treasury bond futures realized correlations, we provide empirical evidence that the tree-HAR model reaches a good compromise between simplicity and flexibility, and yields accurate single- and multi-step out-of-sample forecasts. Such forecasts are also better than those obtained from other standard approaches.

Keywords

High frequency data; Realized correlation; Stock-bond correlation; Tree-structured models; HAR; Regimes.

JEL Classification

C13; C22; C51; C53

1 Introduction

Asset returns cross-correlation is pivotal to many prominent financial problems, such as asset allocation, risk management and option pricing. Recently, the use of high frequency data has been advocated to improve the precision of asset volatility estimation, yielding to the so-called realized volatility (RV) approach proposed in a series of breakthrough papers by Andersen et al. (2001b, 2003), Barndorff-Nielsen and Shephard (2001, 2002a, 2002b, 2005), and Comte and Renault (2001). Regarding the realized volatility approach, the idea of employing high frequency data in the computation of covariances between two assets leads to the analogous concept of *realized covariance* (or covariation); for more details see Martens (2004), Hayashi and Yoshida (2005), Griffin and Oomen (2006), Palandri (2006), Sheppard (2006), and Voev and Lunde (2007). Recently, Corsi and Audrino (2007) proposed a modified tick-by-tick realized covariance estimator in cases where a rounding in the price stamps is needed, a typical situation for many practical financial data sets. *Realized correlations* are then constructed as quotients between realized covariances and products of realized standard deviations.¹

Similar to the smooth transition heterogeneous autoregressive technique proposed by McAleer and Medeiros (2007) which modeled the realized volatility of sixteen US stocks with good forecasting results, we propose a regime-dependent, tree-structured heterogeneous autoregressive (tree-HAR) model for the estimation and prediction of the tick-by-tick realized correlation series. The conditional mean dynamics of the realized correlation series follow local linear HAR processes and are subjected to regime shifts depending on past values of certain relevant predictor variables, such as, for example, past returns, past realized volatilities or time. The local HAR processes are standard linear models where the explanatory variables are past realized correlations at three different horizons: daily, weekly and monthly. For more details, see Corsi (2004) and Corsi et al. (2006). This structure allows the model to take into account two important features exhibited by most real data realized correlation series: long memory and structural changes. Another feature

¹This paper investigates only the fraction of daily realized correlations corresponding to the time when the markets are open. Refer to Andersen et al. (2007) for a more general discussion about how the night correlation (from the overnight returns) can be incorporated into our framework.

of the tree-HAR model is that it belongs to the class of tree-structured threshold regime models, and can therefore be easily estimated and regimes can be confidently interpreted in terms of relevant predictor variables; see for example, Audrino and Bühlmann (2001) and Audrino and Trojani (2006).

We test the accuracy of the tree-HAR model on the series of daily tick-by-tick realized correlations between S&P 500 and 30-year treasury bond futures, and show that the proposed model is able to mimic well the dynamic properties of the daily realized correlation process and can provide accurate daily and multi-step (that is weekly and monthly) out-of-sample forecasts.

This study contributes to the literature in two ways. First, we collect empirical evidence that realized correlations show drastic regime shifts, supporting the evidence already found in other studies. The presence of structural breaks in the stock–bond correlations is already established in several works in the recent literature. Among these, Cappiello et al. (2006) provided statistical evidence of a structural break in stock–bond correlations due to the introduction of the euro. Li (2002), Ilmanen (2003) and Christiansen and Rinaldo (2006) report that US stock–bond correlations went from positive to negative after 1997. Guidolin and Timmermann (2004) found empirical evidence that a four-state Markov regime-switching model is needed to capture the joint dynamics of US stock and bond returns. In their empirical study, Pastor and Stambaugh (2003) found that changes in stock–bond correlations are related to different levels of liquidity. Thus, allowing for structural breaks depending on a large set of predictor variables (for example, liquidity) may be relevant for improving forecasting accuracy.

This paper differs from the above-mentioned studies mainly because we analyze tick-by-tick realized correlations. We contribute to the literature on US stock–bond correlations by estimating local dynamics and incorporating structural breaks in a threshold-type model. The estimated tree-HAR model for daily US stock–bond realized correlations has three optimal regimes (endogenously estimated from the data) which can be interpreted as follows. The first regime is in reaction of US market crashes: in particular, the first regime is characterized by large negative past S&P 500 daily returns, the conditional mean dynamics of the realized correlations are highly persistent, and the volatility of

realized correlations is large. The second and third regimes are both characterized by relatively positive² past S&P 500 daily returns, but for two different time periods. We identify a structural break in time corresponding to the period February to March 1992. This structural break may be a consequence of the western European monetary crisis of 1992–1993. After March 1992 the persistence of the conditional mean dynamics and the volatility of the realized correlations significantly increase. Moreover, we also find that past individual realized volatilities are relevant predictors for future realized correlations.

Second, we perform a series of out-of-sample tests for the superior predictive ability (SPA; see Hansen, 2005) of our model against a number of competitors using different goodness-of-fit statistics, to verify whether the greater flexibility allowed by the tree-HAR model (with a corresponding higher number of parameters to be estimated) has any value for forecasting. We empirically show that the tree-HAR model systematically outperforms the competitors, particularly when multi-period forecasts are considered.

The remainder of the paper is organized as follows: Section 2 proposes the tree-HAR process as a model for the estimation and forecast of realized correlations. Section 3 presents the empirical application to a bivariate series of S&P 500 and 30-year US treasury bond futures tick-by-tick data. Section 4 summarizes and concludes.

2 Modeling Realized Correlations

2.1 The Model

Empirical evidence on strong temporal dependence of realized correlations has been already showed in Andersen et al. (2001a, 2003), and Ferland and Lalancette (2006), among others. This evidence, together with our empirical results reported in Section 3, suggests that realized correlation series are best described by long-memory type of models.

Corsi (2004) and Corsi et al. (2006) recently proposed a class of time series models called heterogeneous autoregressive (HAR) models, which successfully model the long-memory behavior of financial variables in a simple and parsimonious way.

²As is usual for threshold regimes, “positive” here means above the optimal threshold value.

The basic idea was introduced in a study by Müller et al. (1997), where the long memory observed in the volatility was explained as the superimposition of only a few processes operating on different time scales. Corsi (2004) proposed a stochastic additive cascade of three different realized volatility components corresponding to the three main different time horizons present in the market: daily, weekly, and monthly. This stochastic volatility cascade leads to simple AR-type models in the realized volatility which also feature a consideration of realized volatilities defined over different time horizons (the HAR-RV models). Although the HAR models do not formally belong to the class of long-memory models, they are able to reproduce a memory decay which is almost indistinguishable from that observed in the empirical data.

The above-mentioned empirical evidence on the high degree of persistence of correlations suggests that the parsimonious HAR models could also be successfully applied to model the time series of realized correlations. Figure 1 shows the autocorrelogram of the empirical stock–bond correlations of our real data application of Section 3 superimposed on the one for correlations obtained simulating data from the HAR model with parameters similar to those estimated from real data.

Insert Figure 1 about here

Figure 1 shows that the HAR model is very good at mimicking the persistence present in real data correlations.

A second important stylized fact which must be taken into account when building up a model for the realized correlations’ dynamics is the (possible) presence of structural breaks. Various studies in the recent literature on stock–bond correlations already report that stock–bond correlations went from positive to negative after 1997; see for example, Ilmanen (2003). The reasons given to explain this pattern vary. One relates to market uncertainty and risk, introducing the “flight-to-quality” effect, which suggests the phenomenon of fleeing from stock to bond markets in times of worsening economic conditions (see for example, Ilmanen, 2003; or Connolly et al., 2005). Another explanation for the change of sign in stock–bond correlations relates to differences in inflation expectations or in other macroeconomic announcements (see for example, Li, 2002; or Christiansen and

Ranaldo, 2006). Pastor and Stambaugh (2003) found that a kind of “flight-to-quality” effect appears in months with exceptionally low liquidity, that is months in which liquidity drops severely tend to be months in which stocks and fixed-income assets move in opposite directions. In two recent studies, Guidolin and Timmermann (2004) and Audrino and Trojani (2007) incorporated the possible regime shifts in the conditional dynamics of (realized) correlations using regime-dependent models. Guidolin and Timmermann (2004) analyzed the joint dynamics of US stock and bond returns using a Markov regime-switching model, and found empirical evidence of the presence of four different regimes.

Along these lines, we propose a tree-structured local HAR model for the dynamics of tick-by-tick realized correlations which is able to take into account the above-discussed stylized facts of realized correlation series: long-memory and structural breaks. Tree-structured models belong to the class of threshold regimes models, where regimes are characterized by some threshold for the relevant predictor variables. This class of model was introduced by Audrino and Bühlmann (2001) in the financial volatility literature, and was generalized recently to capture simultaneous regime shifts in the first and second conditional dynamics of returns series, with good results for different forecasting applications (see for example, Audrino and Trojani, 2006).

Let $\{\widetilde{RC}\}_{t \geq 1}$ be the daily Fisher-transformed (FT) series of the tick-by-tick realized correlations $\{RC\}_{t \geq 1}$,³ that is:

$$\widetilde{RC}_t = \frac{1}{2} \log \left(\frac{1 + RC_t}{1 - RC_t} \right), \quad RC_t \in [-1, 1]$$

We then model the series $\{\widetilde{RC}\}_{t \geq 1}$ as:

$$\widetilde{RC}_{t+1} = \mathbb{E}_t[\widetilde{RC}_{t+1}] + \sigma_{t+1}U_{t+1} \quad (1)$$

where $\{U_t\}_{t \geq 1}$ is a sequence of i.i.d. innovations following the distribution p_U with expected value 0 and variance 1, and $\mathbb{E}_t[\cdot]$ denotes (as usual) the conditional expectation given the

³When considering Fisher-transformed correlations we do not have to impose any restriction on the parameters in the model to ensure the final estimates and forecasts to lie in the $[-1, 1]$ interval. By performing the tree-HAR analysis on the original correlations, we found qualitatively exactly the same results; for example, the optimal regime structure was found to be the same.

information up to time t . The conditional dynamics of the FT correlations are given by:

$$\mathbb{E}_t[\widetilde{RC}_{t+1}] = \sum_{j=1}^k (a_j + b_j^{(d)} \widetilde{RC}_t + b_j^{(w)} \widetilde{RC}_t^{(w)} + b_j^{(m)} \widetilde{RC}_t^{(m)}) I_{[\mathbf{X}_t^{\text{pred}} \in \mathcal{R}_j]} \quad (2)$$

$$\sigma_{t+1}^2 = \sum_{j=1}^k \sigma_j^2 I_{[\mathbf{X}_t^{\text{pred}} \in \mathcal{R}_j]}, \quad \sigma_j^2 > 0, j = 1, \dots, k \quad (3)$$

where $\theta = (a_j, b_j^{(d)}, b_j^{(w)}, b_j^{(m)}, \sigma_j^2 : j = 1, \dots, k)$ is a parameter vector which parameterizes the local HAR dynamics in the different regimes, k is the number of regimes (endogenously estimated from the data), and $\widetilde{RC}_t^{(w)}$ and $\widetilde{RC}_t^{(m)}$ are respectively the weekly and monthly FT-realized correlations, obtained as simple averages of 5 respectively 22 daily FT-realized correlations. The regimes are characterized by partition cells \mathcal{R}_j of the relevant predictor space G of $\mathbf{X}_t^{\text{pred}}$:

$$G = \bigcup_{j=1}^k \mathcal{R}_j, \quad \mathcal{R}_i \cap \mathcal{R}_j = \emptyset \quad (i \neq j)$$

In our study, the relevant predictor variables in $\mathbf{X}_t^{\text{pred}}$ are past-lagged FT-realized correlations, and past-lagged realized volatilities and returns of the two instruments under investigation. All such predictor variables are considered at three different time horizons: daily, weekly, and monthly. We also consider as an additional predictor variable time.

To completely specify the conditional dynamics given in equations (2) and (3) of the FT-realized correlations, we determined the shape of the partition cells \mathcal{R}_j , which are admissible in the tree-HAR model. Similar to the standard classification and regression trees (CART) procedure (see Breiman et al., 1994), the only restriction we impose is that regimes must be characterized by rectangular partition cells with edges determined by thresholds on the predictor variables. Such partition cells are practically constructed using a binary tree. Introducing this restriction has two major advantages: it allows a clear interpretation of the regimes in terms of relevant predictor variables, and also allows us to estimate the model using large-dimensional predictor spaces G .

As an illustration, in our empirical application on US stock–bond realized correlations

presented in Section 3, the optimal partition cells are of the form:

$$\begin{aligned}\mathcal{R}_1 &= \{\mathbf{X}^{\text{pred}} : R_{\text{S\&P } 500} \leq -15.21\} \\ \mathcal{R}_2 &= \{\mathbf{X}^{\text{pred}} : R_{\text{S\&P } 500} > -15.21 \text{ and } t \leq \text{March } 1991\} \\ \mathcal{R}_3 &= \{\mathbf{X}^{\text{pred}} : R_{\text{S\&P } 500} > -15.21 \text{ and } t > \text{March } 1991\}\end{aligned}$$

where $R_{\text{S\&P } 500}$ denotes the (annualized) daily returns of the US S&P 500 Index, and t denotes time. We find a first regime characterized by large losses of the US market index, and second and third regimes in reaction of positive and moderate losses of the US market, with an important structural break in time corresponding to March 1992. Section 3.2 contains a more structured discussion and interpretation of these results.

2.2 Estimation

The tree-HAR model introduced in equations (1) to (3) can be estimated using quasi-maximum likelihood (QML). Conditional on some reasonable starting values, the negative quasi-log-likelihood for model (1)-(3) is given by:

$$-l(\theta; (\widetilde{RC}, \mathbf{X}^{\text{pred}})_1^n) = \frac{n}{2} \log(2\pi) + \frac{1}{2} \sum_{t=1}^n \log(\sigma_t^2(\theta)) + \frac{1}{2} \sum_{t=1}^n \frac{(\widetilde{RC}_t - \mathbb{E}_{t;\theta}[\widetilde{RC}_{t-1}])^2}{\sigma_t^2(\theta)} \quad (4)$$

Therefore, for any fixed sequence of partition cells, the tree-HAR model can be estimated by QML. The choice of the optimal partition cells (that is, splitting variables and threshold values) involves a model choice procedure for non-nested hypotheses. Similar to CART, and as already discussed in Audrino and Bühlmann (2001) and Audrino and Trojani (2006), for general tree-structured models, the model selection of the optimal splitting variables and threshold values can be performed via a tree-structured partial search. Within any data-determined tree structure the optimal model is selected using the Bayesian-Schwartz information criterion (BIC). As pointed out by Hansen (1996), the use of model selection criteria to decide if the inclusion of another regime is relevant in threshold regression models such as the tree-HAR model has the following drawback: the location of the split cannot be estimated consistently when an irrelevant regime is added to the model. To overcome the problem, and also to ensure computational feasibility, we

searched for threshold values over fixed grid points that are empirical quantiles of the different predictor variables.⁴ For more details about the flexible procedure used to estimate the model, refer to Section 2.3 and Appendix A in Audrino and Trojani (2006). Proof of the consistency of the conditional mean and volatility estimates in the tree-HAR model under a possible model misspecification can be derived from Theorem 1 in Audrino and Bühlmann (2001).

3 Empirical application

3.1 Data and stylized facts

We considered a tick-by-tick bivariate returns series of the S&P 500 Index and 30-year US treasury bond futures for the period from January 1990 to October 2003, for a total of 3,391 daily observations. The data come from the Price-data.com database with time stamps rounded at the one minute frequency. Combining the first-last (FL) realized covariance estimator introduced by Corsi and Audrino (2007) as a generalization of the Hayashi and Yoshida (2005) realized covariance estimator to overcome the problem of working with rounded price time stamps together with a tick-by-tick realized volatility estimator, we are now able to construct a realized correlation measure where both the volatilities and the covariances are computed from tick-by-tick data. As is usual, correlations are computed as quotients between covariances and products of standard deviations. Our study employ the multi-scales discrete sine transform (DST) realized volatility estimator proposed by Curci and Corsi (2003), which consists in a multi-frequency regression-based approach made robust by a discrete sine transform filter which optimally decorrelates the

⁴McAleer and Medeiros (2007) and Medeiros and Veiga (2008) recently proposed a sequence of tests to determine the optimal number of regimes for a class of smooth transition models for the dynamics of financial (realized) volatility which circumvents the problem of identification in a way that controls the significance level of the tests in the sequence and computes an upper bound to the overall significance level. Such a strategy can be easily adapted to the case of fitting tree-HAR models. When choosing optimal splitting variables and threshold values using the above-mentioned sequence of tests, we obtained qualitatively the same regimes and, consequently, forecasting accuracy.

price signal from microstructure noise.⁵

The upper part of Figure 2 shows the time series of 3,391 daily tick-by-tick correlations we obtained for the time period between 1990 and 2003,⁶ together with their autocorrelogram. In the bottom part of Figure 2 Fisher-transformed (FT) realized correlations are plotted, again together with their autocorrelogram.

Insert Figure 2 about here

The two empirical stylized facts of realized correlations discussed in Section 2.1 are clearly pointed out in Figure 2: long memory and structural breaks. The same features are also apparent in the FT correlation series. When restricting the discussion only on structural breaks in time, from a preliminary visual inspection of the (Fisher-transformed) correlation dynamics in Figure 2, we can recognize two important changes of regime: the first occurring around the end of 1993 to the beginning of 1994, and the second one around the end of 1997. The correlation between the two series oscillates around a positive stable value of about 20 percent until 1994, around 40 percent from 1994 to 1997, while after the end of 1997, the correlation starts to exhibit a stronger dynamics and becomes predominantly negative.⁷ These structural changes in the dynamics of the correlation between S&P 500 and US bonds is also apparent from the different behavior of the autocorrelation function computed for the three periods January 1990 to the end of 1993 (henceforth called the 90-94 period), from the beginning of 1994 to the end of 1997 (94-98 period) and from the beginning of 1998 to the end of the sample (98-03 period), illustrated in Figure 3.

Insert Figure 3 about here

⁵We also constructed tick-by-tick realized volatilities using the two scales estimator of Zhang et al. (2005), obtaining similar results.

⁶In two days out of 3,391 we obtained an estimated tick-by-tick realized correlation slightly outside the $[-1, 1]$ interval. In those two cases, we arbitrarily set the correlation absolute value at 0.99.

⁷In the first case we report evidence of a regime shift in the positive direction; whereas in the second case evidence of a regime shift in the negative direction (with a consequent change of sign of the correlations).

In the first 90-94 period the level of the autocorrelation is very low and quickly decaying. In the second period 94-97 the autocorrelation level and its persistence significantly increase. After the end of 1997 the memory of the process, particularly at small and moderate lags, rises further. This points to a consistent increase in the memory persistence of the stock–bond correlation in the most recent years. This new stylized fact of the stock–bond correlation would not be so easily identifiable without the employment of high frequency data and a precise realized correlation measure.

Another interesting effect shown in Figure 3 is how time structural changes affect the global autocorrelation function computed on the full sample inducing an artificially high level in the autocorrelation coefficients. This phenomenon is discussed in detail by Diebold and Inoue (2001). Nonetheless, even without this structural break effect, the autocorrelation function of the realized correlation remains highly persistent, as shown by the separated sub-sample autocorrelation functions.

Various studies in the recent literature on stock–bond correlations empirically show that stock–bond correlation dynamics are subjected to structural breaks in time. Most of these studies report that stock–bond correlations went from positive to negative after 1997; see, among others, Ilmanen (2003), Connolly et al. (2005), Li (2002), and Christiansen and Rinaldo (2006). The reasons given to explain this pattern vary. One relates to market uncertainty and risk, introducing the “flight-to-quality” effect, which suggests the phenomenon of fleeing from stock to bond markets in times of worsening economic conditions. Another explanation for the change of sign in bond–stock correlations relates to differences in inflation expectations or in other macroeconomic announcements.

In this study, we wanted to generalize this analysis to allow for structural breaks depending on a larger set of predictor variables, directly incorporated in a tree-HAR model specification. In addition to time, we considered as predictors past-lagged US S&P 500 and bond returns and realized volatilities, and past-lagged realized correlations at three different time horizons (daily, weekly, and monthly). The results of this in-sample analysis are shown in the next section.

3.2 Estimation results

The analysis starts by estimating the tree-HAR model (1)-(3) on the whole data sample (from January 1990 to October 2003). Optimal coefficients, as well as the optimal regimes, are reported in Table 1. Model-based bootstrap standard errors are given in parentheses.

Insert Table 1 about here

Table 1 shows that almost all coefficients in the local dynamics of the conditional mean and variance of the Fisher-transformed (FT) realized correlations (with only one exception) are highly significant. The first regime is characterized by large losses of the daily S&P 500 Index, and therefore of the US stock market. In this regime volatility is high and the conditional mean dynamics are highly persistent. Moreover, the long-term mean is negative. This result supports and can be interpreted through the “flight-to-quality” effect already cited several times in the literature; see for example, Ilmanen (2003). In this regime characterized by bad stock market conditions, stocks and fixed-income assets tend to move in opposite directions, suggesting the phenomenon of fleeing from stock to bond markets.

The second and third regimes are characterized by “relatively good” stock market conditions. Under such conditions, we identified a structural break in time around March 1992 that may be caused by the western European monetary crisis of 1992–1993. Before that date, FT-realized correlations are moderately persistent, the long-run mean is positive and small, and volatility is also small. Such dynamics can be reasonably well associated with stable economic and market conditions. On the contrary, after March 1992 volatility increases (of approximately a factor 2), conditional mean dynamics of FT-realized correlations become highly persistent, and the long-run mean drastically increases.

To end this section, we report the results of different goodness-of-fit statistics: the AIC and BIC criteria, and the mean absolute errors (MAE), the mean squared errors (MSE), and Ljung-Box (LB) tests at three different lags (5, 20, 50) for the residuals U_t in (1). For the purposes of comparison, we report the same statistics for different competitors introduced in the literature: the standard AR(1) model, the ARMA(1,1)

model, the ARIMA(1,1,1) model introduced for non-stationary time series, and the global HAR model. Results are summarized in Table 2

Insert Table 2 about here

The superior ability of the tree-HAR model in estimating the dynamics of the FT-realized correlations is clearly shown by the values of the different goodness-of-fit statistics. For all performance measures considered, the tree-HAR model yields the best results. When neglecting to incorporate in the model long-memory or the possible presence of structural breaks, the estimates obtained for the conditional dynamics of the FT-realized correlations are highly inaccurate.

3.3 Forecasting results

To better validate the usefulness of the tree-HAR model for the real data under investigation, we examine its forecasting ability, and compare this with the alternative methods introduced in the last section. We performed a series of out-of-sample tests to assess the forecasting power of the tree-HAR model for single-period and multi-period predictions. Regarding goodness-of-fit statistics, we again considered the out-of-sample MAE and MSE of the residuals. In addition to these performance measures, we also report results for the out-of-sample log-likelihood in equation (4) in the single-period out-of-sample test, and the R^2 obtained when regressing realizations against forecasts at the same time t (Mincer-Zarnowitz regression).

3.3.1 Single-period forecasts

To derive the daily forecasts we used a rolling strategy. The models are re-estimated every month (22 trading days) using all past data available in the sample. The initial in-sample period is from January 1990 to December 1999. Consequently, we obtained 926 out-of-sample daily forecasts (until October 2003). Results are summarized in Table 3. p -values of superior predictive ability (SPA) tests introduced by Hansen (2005) for the null-hypothesis that the chosen model is not inferior to any of the alternatives are given in parentheses.

Insert Table 3 about here

The tree-HAR model yields the best results for three out of the four goodness-of-fit statistics considered; with respect to the MSE statistic, the ARIMA(1,1,1) model is slightly better. Differences measured by the MAE and MSE statistics are in most cases very small and not statistically significant. Only the simple AR(1) and ARMA(1,1) models are clearly outperformed by the competitors. However, with respect to the out-of-sample likelihood, the tree-HAR model yields significant improvements in the accuracy of the FT-realized correlation forecasts over the competitors. Summarizing, Table 3 suggests that when the foci are daily forecasts, the tree-HAR model is the best. Improvements will be marginal, depending on the performance measure.

3.3.2 Multi-period forecasts

Practical asset allocation applications typically require correlation forecasts at longer time horizons. Therefore, we performed two out-of-sample tests at weekly (that is, five days) and monthly (that is, 22 days) horizons to assess the accuracy of the multi-period forecasts obtained using the different approaches. Such multi-period predictions are constructed using filtered historical simulation (FHS); see Barone-Adesi et al. (1998, 1999). Briefly, FHS works as follows. We generate 10,000 future scenarios at 5 (22) days horizons, bootstrapping the residuals estimated from the different models.⁸ The forecast for the 5 (22) days ahead FT-realized correlation is then given by the median of the empirical distribution of the simulated future scenarios.

As in the previous out-of-sample experiment, we used the same rolling strategy and initial in-sample period. Results are summarized in Table 4. Again, p -values of SPA tests are reported in parentheses.

Insert Table 4 about here

⁸We considered the stationary bootstrap of Politis and Romano (1994) to account for the possible remaining autocorrelation in the residuals (see especially the Ljung-Box test results given in Table 2). Modification of the implementation to the block bootstrap of Künsch (1989) is straightforward.

The better forecasting power of the tree-HAR model with respect to all competitors for multi-period predictions is borne out by the results of the SPA tests. Especially for longer-time ahead forecasts (that is, 1 month), the predictions obtained using the tree-HAR model outperform those gleaned from the alternative approaches. Gains are, in most cases, statistically significant.

4 Conclusions

Combining realized covariances with realized volatilities, we obtained a realized correlation measure where both the volatilities and the covariances are computed from tick-by-tick data. We then propose a tree-HAR model as a simple and parsimonious representation for the conditional dynamics of the (Fisher-transformed) realized correlations. The tree-HAR model is able to take into account two important stylized facts of realized correlations: strong temporal dependence (that is, long-memory) and structural breaks.

We estimated the tree-HAR model on the realized correlation series of S&P 500 and US bond returns, finding empirical evidence that the conditional dynamics are subjected to regime changes depending on different values of past S&P 500 daily returns. We also identified a structural break in time, around March 1992.

We then tested the prediction accuracy of the model using SPA tests for different goodness-of-fit statistics finding empirical evidence of its strong predictive power. The tree-HAR model significantly outperforms the competitors, particularly when the final goal is multi-period forecasting.

References

- Andersen, T., Bollerslev, T., Diebold, F., and Ebens, H. (2001a). The distribution of realized stock return volatility. *Journal of Financial Economics* **61** (1), 43–76.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P. (2001b). The distribution of exchange rate volatility. *Journal of the American Statistical Association* **96**, 42–55.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica* **71**(2), 579–625.
- Andersen, T. G., Bollerslev, T., and Huang, X. (2007). A reduced form framework for modeling volatility of speculative prices based on realized variation measures. CREATES Research Paper Series 14.
- Audrino, F. and Bühlmann, P. (2001). Tree-structured GARCH models. *Journal of the Royal Statistical Society, Series B*, **63**, 727–744.
- Audrino, F. and Trojani, F. (2006). Estimating and predicting multivariate volatility regimes in global stock markets. *Journal of Applied Econometrics* **21**, No. 3, 345–369.
- Audrino, F. and Trojani, F. (2007). A general multivariate GARCH model with dynamic conditional correlations. VWA Discussion Paper Series 25, University of St. Gallen.
- Barndorff-Nielsen, O. E. and Shephard, N. (2001). Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society, Series B*, **63**, 167–241.
- Barndorff-Nielsen, O. E. and Shephard, N. (2002a). Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society, Series B*, **64**, 253–280.
- Barndorff-Nielsen, O. E. and Shephard, N. (2002b). Estimating quadratic variation using realized variance. *Journal of Applied Econometrics* **17**(5), 457–477.

- Barndorff-Nielsen, O. E. and Shephard, N. (2005). How accurate is the asymptotic approximation to the distribution of realized volatility?. In D. W. F. Andrews and J. H. Stock (eds.), *Identification and Inference for Econometric Models. A Festschrift in Honour of T.J. Rothenberg*. Cambridge University Press, pp. 306–331.
- Barone-Adesi, G., Bourgoin, F., and Giannopoulos, K. (1998). Don’t Look Back. *Risk* **11**, August, 100–104.
- Barone-Adesi, G., Giannopoulos, K., and Vosper, L. (1999). VaR Without Correlations for Portfolio of Derivative Securities. *Journal of Futures Markets* **19** (April), 583–602.
- Breiman, L., Friedman, J. H., Olshen, R. A., and Stone, C. J. (1984). *Classification and Regression Trees*. Wadsworth, Belmont (CA).
- Cappiello, L., Engle, R. F., and Sheppard, K. (2006). Asymmetric dynamics in the correlations of global equity and bond returns. *Journal of Financial Econometrics* **4**, 537–572.
- Christiansen, C. and Rinaldo, A. (2006). Realized bond-stock correlation: macroeconomic announcement effects. Forthcoming in the *Journal of Futures Markets*.
- Comte, F. and Renault, E. (2001). Long memory in continuous time stochastic volatility models. *Mathematical Finance* **8**, 291–323.
- Connolly, R., Stivers, C., and Sun, L. (2005). Stock market uncertainty and the stock-bond return relation. *Journal of Financial and Quantitative Analysis* **40**, 161–194.
- Corsi, F. (2004). Simple long memory models of realized volatility. Unpublished manuscript, University of Lugano.
- Corsi, F. and Audrino, F. (2007). Realized covariance tick-by-tick in presence of rounded time stamps and general microstructure effects. Working paper, University of Lugano.
- Corsi, F., Mitnik, S., Pigorsch, C., and Pigorsch, U. (2006). The volatility of realized

- volatility. Forthcoming in the *Special Issue of Econometric Reviews on Realized Volatility and Long Memory*.
- Curci, G. and Corsi, F. (2003). A discrete sine transform approach for realized volatility measurement. NCCR FINRISK Working Paper No. 44.
- Diebold, F. X. and Inoue, A. (2001). Long memory and regime switching. *Journal of Econometrics* **105**, 131-159.
- Ferland, R. and Lalancette, S. (2006). Dynamics of realized volatilities and correlations: An empirical study. *Journal of Banking and Finance* **30** (7), 2109–2130.
- Griffin, J. E. and Oomen, R. C. A. (2006). Covariance measurement in the presence of non-synchronous trading and market microstructure noise. Unpublished manuscript.
- Guidolin, M. and Timmermann, A. (2004). An econometric model of nonlinear dynamics in the joint distribution of stock and bond returns. Forthcoming in the *Journal of Applied Econometrics*.
- Hansen, B. E. (2006). Inference when a nuisance parameter is not identified under the null hypothesis. *Econometrica* **64**, 413–430.
- Hansen, P. R. (2005). A test for superior predictive ability. *Journal of Business & Economic Statistics* **23**, 365–380.
- Hayashi, T. and Yoshida, N. (2005). On covariance estimation of non-synchronously observed diffusion processes. *Bernoulli* **11**(2), 359–379.
- Ilmanen, A. (2003). Stock-Bond Correlations. *Journal of Fixed Income* **13**, 55–66.
- Künsch, H. R. (1989). The jackknife and the bootstrap for general stationary observations. *Annals of Statistics* **17**, 1217–1241.
- Li, L. (2002). Macroeconomic factors and the correlation of stock and bond returns. Working paper, Yale International Center for Finance.
- Martens, M. (2004). Estimating unbiased and precise realized covariances. Social Science Research Network Electronic Library.

- McAleer, M. and Medeiros, M. C. (2007). A multiple regime smooth transition heterogeneous autoregressive model for long memory and asymmetries. Forthcoming in the *Journal of Econometrics*.
- Medeiros, M. C. and Veiga, A. (2008). Modeling multiple regimes in financial volatility with a flexible coefficient GARCH(1,1) model. Forthcoming in *Econometric Theory*.
- Müller, U., Dacorogna, M., Dave, R., Olsen, R., Pictet, O., and Weizsacker, J. (1997). Volatilities of different time resolutions – Analyzing the dynamics of market components. *Journal of Empirical Finance* **4**, 213–239.
- Palandri, A. (2006). Consistent realized covariance for asynchronous observations contaminated by market microstructure noise. Unpublished manuscript.
- Pastor, L. and Stambaugh, R. F. (2003). Liquidity risk and expected stock returns. *The Journal of Political Economy* **111**, 642–685.
- Politis, D. N. and Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association* **89**, 1303–1313.
- Sheppard, K. (2006). Realized covariance and scrambling. Unpublished manuscript.
- Voev, V. and Lunde, A. (2007). Integrated covariance estimation using high-frequency data in the presence of noise. *Journal of Financial Econometrics* **5**, 68–104.
- Zhang, L., Aït-Sahalia, Y., and Mykland, P.A. (2005). A tale of two time scales: determining integrated volatility with noisy high frequency data. *Journal of the American Statistical Association* **100**, 1394–1411.

US S&P 500 bond FT correlation optimal parameters

Regime structure	Local parameters				
\mathcal{R}_j	a_j	$b_j^{(d)}$	$b_j^{(w)}$	$b_j^{(m)}$	σ_j^2
$R_{\text{S\&P 500}} \leq -15.21$	-0.0510 (0.0135)	0.3517 (0.0546)	0.3123 (0.0914)	0.3432 (0.0934)	0.0615 (0.0110)
$R_{\text{S\&P 500}} > -15.21,$ $t \leq \text{March 1991}$	0.0968 (0.0254)	0.1385 (0.0361)	0.2698 (0.07982)	0.1913 (0.1231)	0.0186 (0.0033)
$R_{\text{S\&P 500}} > -15.21,$ $t > \text{March 1991}$	0.0092 (0.0040)	0.1927 (0.0201)	0.3829 (0.0337)	0.3799 (0.0367)	0.0403 (0.0031)

Table 1: Tree-HAR in-sample optimal parameters and regimes for the tick-by-tick Fisher-transformed (FT) realized correlation series of S&P 500 and US bond returns. The sample period is from January 1990 to October 2003, for a total of 3,391 daily observations. $R_{\text{S\&P 500}}$ and t denote the past-lagged daily S&P 500 return and time, respectively. Model-based bootstrap standard errors computed using 1,000 replications are given in parentheses.

In-sample performance results					
	AR(1)	ARMA(1,1)	ARIMA(1,1,1)	HAR	Tree-HAR
AIC	-189.39	-1174.8	-1222.3	-1223.1	-1427.6
BIC	-171.03	-1150.3	-1197.8	-1192.5	-1335.8
MAE	0.1681	0.1449	0.1440	0.1436	0.1425
MSE	0.0553	0.0412	0.0406	0.0406	0.0398
LB(5)	0	≈ 0	0.0975	0.0312	0.1556
LB(20)	0	≈ 0	0.0149	0.0161	0.0637
LB(50)	0	≈ 0	0.0382	0.0507	0.2715

Table 2: In-sample goodness-of-fit results for the tree-HAR model, in comparison with the classical AR(1), ARMA(1,1), ARIMA(1,1,1) models, and the global HAR model. Data are FT-realized correlations between January 1990 and October 2003, for a total of 3,391 daily observations. The performance measure considered are the Akaike and Bayesian-Schwartz information criteria (AIC, BIC), and the mean absolute errors (MAE), the mean squared errors (MSE), and p-values of standard Ljung-Box tests at three different lags of the residuals.

Single-period forecasting results				
Model	Loglik.	MAE	MSE	R ²
AR(1)	155.7 (0)	0.2055 (0)	0.0753 (0)	0.3911
ARMA(1,1)	−82.10 (0.036)	0.1615 (0.222)	0.0480 (0.049)	0.4795
ARIMA(1,1,1)	−93.69 (0.088)	0.1601 (0.777)	0.0469 (0.785)	0.4810
HAR	−88.22 (0.042)	0.1602 (0.842)	0.0474 (0.148)	0.4806
Tree-HAR	−109.1 (0.617)	0.1601 (0.407)	0.0471 (0.527)	0.5077

Table 3: Comparative results of one day forecasts of S&P 500 US bond FT-realized correlations obtained using the classical AR(1), ARMA(1,1), ARIMA(1,1,1) models, the global HAR model, and the tree-HAR model. The forecasting time period is between January 2000 and October 2003, for a total of 926 daily observations. Out-of-sample forecasts are computed using a rolling strategy: the models are re-estimated every month (22 days) using the whole past information in the data sample. Performance is measured according to the out-of-sample log-likelihood (Loglik.), mean absolute errors (MAE) and mean squared errors (MSE) of the residuals, and R² of regressing realizations against forecasts at the same time t . Values in parentheses are reported p -values of superior predictive ability (SPA) tests for the null-hypothesis that the a given model is not inferior to any of the alternatives.

Multi-period forecasting results: one week horizon			
Model	MAE	MSE	R ²
AR(1)	0.4243 (0)	0.2431 (0)	0.2810
ARMA(1,1)	0.1777 (0.085)	0.0600 (0.002)	0.3853
ARIMA(1,1,1)	0.1763 (0.274)	0.0565 (0.074)	0.3859
HAR	0.1756 (0.346)	0.0579 (0.005)	0.3861
Tree-HAR	0.1742 (0.776)	0.0557 (0.497)	0.4298

Multi-period forecasting results: one month horizon			
Model	MAE	MSE	R ²
AR(1)	0.5939 (0)	0.4346 (0)	0.0853
ARMA(1,1)	0.2329 (0)	0.0974 (0)	0.2261
ARIMA(1,1,1)	0.2118 (0.061)	0.0766 (0.051)	0.2234
HAR	0.2135 (0.006)	0.0813 (0)	0.2381
Tree-HAR	0.2066 (0.533)	0.0751 (0.501)	0.2837

Table 4: Comparative results of one week and one month forecasts of S&P 500 US bond FT-realized correlations obtained using the classical AR(1), ARMA(1,1), ARIMA(1,1,1) models, the global HAR model, and the tree-HAR model. The forecasting time period is between January 2000 and October 2003. Out-of-sample forecasts are computed using a rolling strategy: the models are re-estimated every month (22 days) using the whole past information in the data sample. Performance is measured according to the mean absolute errors (MAE) and mean squared errors (MSE) of the residuals, and R² of regressing realizations against forecasts at the same time t . Items in parentheses are reported p -values of superior predictive ability (SPA) tests for the null-hypothesis that the a given model is not inferior to any of the alternatives.

Autocorrelograms comparison of empirical and simulated data

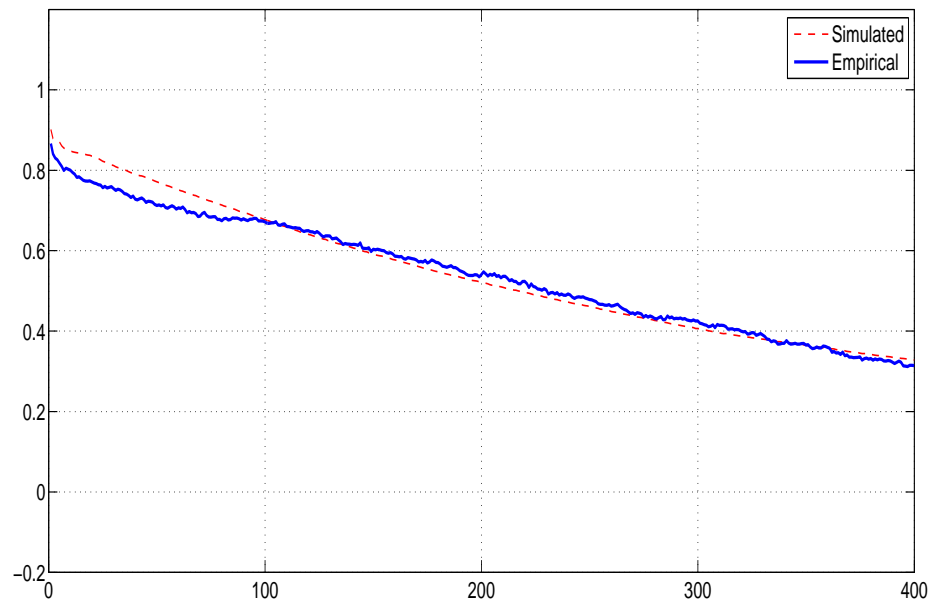


Figure 1: Comparison of autocorrelograms of correlations for the empirical US stock–bond data under investigation (solid) and for simulated data (dotted) from an HAR model with parameters estimated on the full US stock–bond real data sample.

S&P 500 US bond daily realized correlation from 1990 to 2003

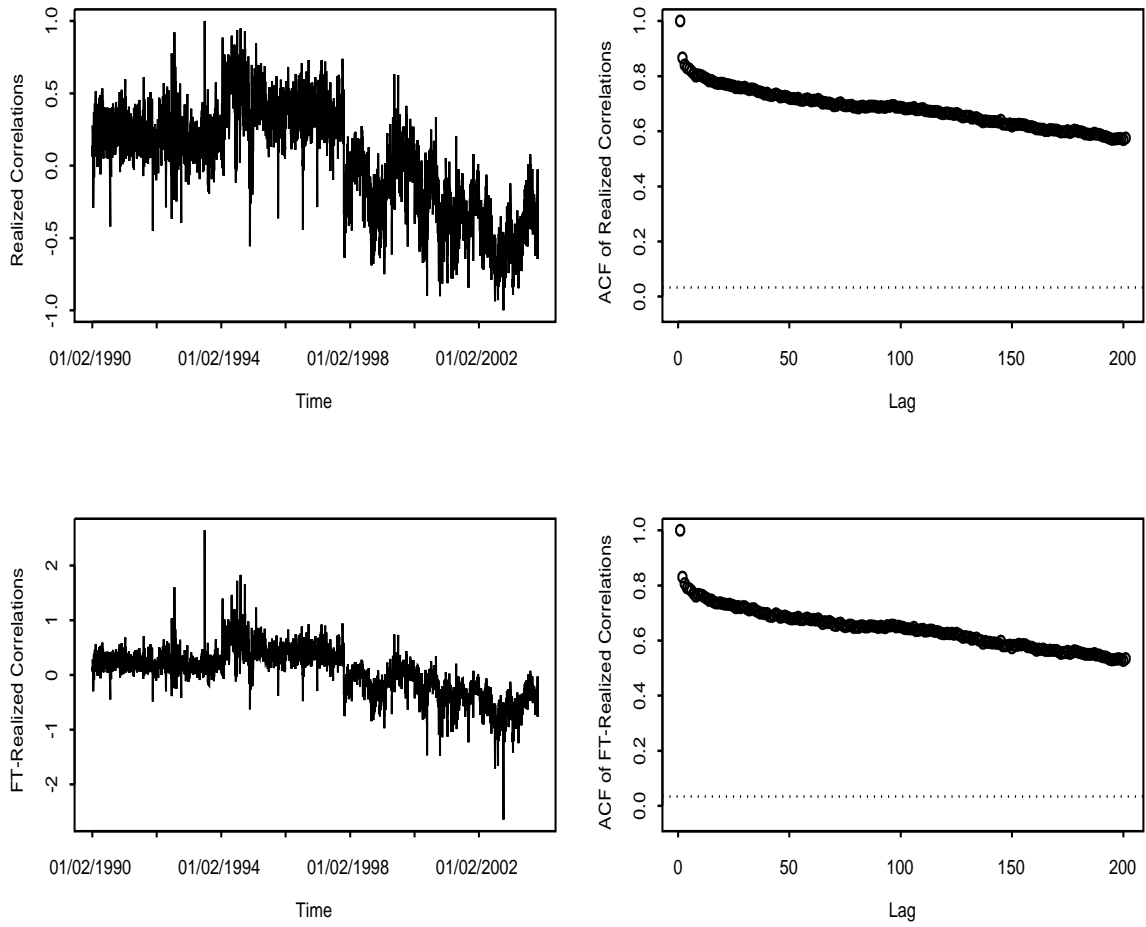


Figure 2: Time series of daily S&P 500 and 30-year US treasury bond realized correlations (upper-left panel) and Fisher-transformed (FT) realized correlations (lower-left panel) constructed using tick-by-tick data, together with their autocorrelograms (right panels). The time period under investigation is from January 1990 to October 2003.

Autocorrelogram of S&P 500 US bond correlation

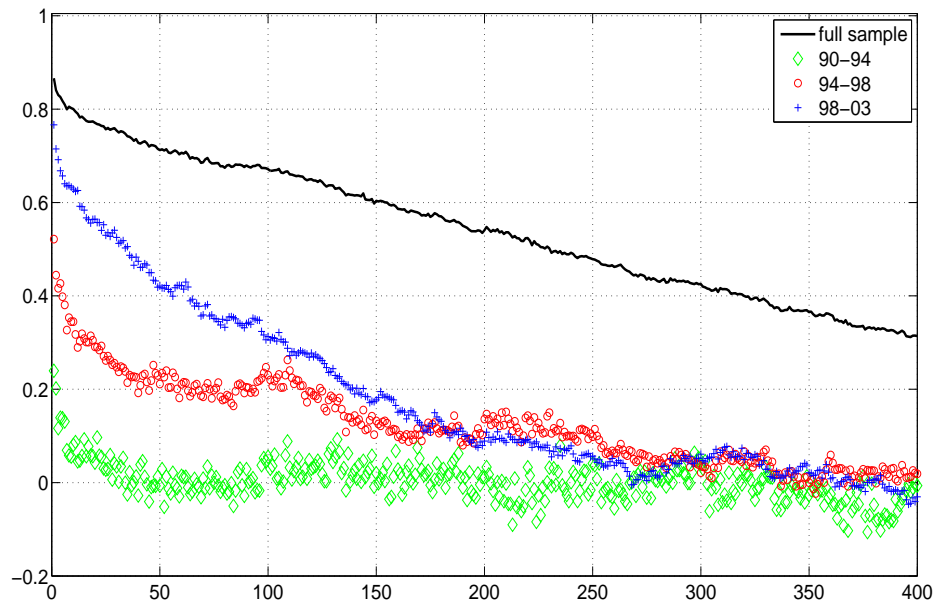


Figure 3: Autocorrelation functions of the S&P 500 and 30-year US treasury bond realized correlation for the full sample 1990–2003 and the three sub-samples (time regimes) 90-94, 94-98, and 98-03.