Economics Letters 125 (2014) 43–46

Contents lists available at ScienceDirect

Economics Letters

journal homepage: www.elsevier.com/locate/ecolet

Can we reject linearity in an HAR-RV model for the S&P 500? Insights from a nonparametric HAR-RV

Jerome Lahaye *, Philip Shaw

Fordham University, United States

HIGHLIGHTS

- We tests the linearity assumption underlying the heterogeneous autoregressive model for realized volatility (HAR-RV).
- We implement a consistent model specification test that is robust to both distributional and model misspecification.
- We find that, using a nonparametric HAR-RV, we are unable to reject the null of linearity.

ARTICLE INFO

Article history:
Received 22 April 2014
Received in revised form
26 June 2014
Accepted 5 July 2014
Available online 2 August 2014

JEL classification:
C12
C13
C14
C22
C52
C58
G1

Keywords:
Nonparametric
Realized volatility
HAR
Heterogeneous
Autoregressive

ABSTRACT

This article tests the linearity assumption underlying the popular heterogeneous autoregressive model for realized volatility (HAR-RV). We implement a consistent model specification test that is robust to both distributional and model misspecification. We find that, using a nonparametric HAR-RV (NPHAR-RV), we are unable to reject the null of linearity.

1. Introduction

The heterogeneous autoregressive model for realized volatility (HAR-RV) of Corsi (2009) and its variants are probably among the most widely used volatility models for financial returns.[1] The HAR-RV is a parsimonious stochastic volatility model that captures long memory through inclusion of lagged realized volatilities aggregated at different frequencies. Therefore, it provides an appealing alternative to ARFIMA models for realized volatility (e.g. Andersen et al., 2001). More precisely, the HAR-RV(1, 5, 21) model specifies day-t volatility as a linear function of lagged daily, weekly and monthly volatility, plus a disturbance.[2] The purpose of this paper is to test the commonly made linearity assumption of the HAR-RV(1, 5, 21) model for the S&P 500.

To do this, we use a *consistent* nonparametric model specification test presented in Li (1994) and Zheng (1996). Our paper complements those of Andrada-Felix et al. (2013) and Clements

* Correspondence to: Fordham University, 113 west 60th street, #924G NY, NY 10023, United States. Tel.: +1 212 636 6343.
E-mail addresses: jlahaye@fordham.edu (J. Lahaye), pshaw5@fordham.edu (P. Shaw).

1 As of February 2014, Corsi (2009) counts several hundreds of citations.

2 The numbers (1, 5, 21) specify the retained lagged regressor aggregation: 1, 5 and 21 days for the lagged daily, weekly and monthly volatility, respectively. The rationale for this aggregation choice is studied in Craioveanu and Hillebrand (2012).

http://dx.doi.org/10.1016/j.econlet.2014.07.003
0165-1765/

and Becker (2009). Our approach differs as we are the first to formally test the linear model first developed in Corsi (2009). In addition to this, Clements and Becker (2009) and Andrada-Felix et al. (2013) use a nearest neighbor estimator whereas we use a local constant approach. Both Andrada-Felix et al. (2013) and Clements and Becker (2009) focus on the difference in gains from forecasting while we focus on the validity of the linearity assumption made in Corsi (2009). Our results show that the linear HAR-RV(1, 5, 21) is in-fact justified by the data as we fail to reject the HAR-RV(1, 5, 21) as the true population model.

2. The HAR-RV model

An asset price daily realized variance is defined as the sum of squared intraday log-returns over the day. For M intraday returns over one day, day-t realized variance is:

$$RV_t = \sum_{j=1}^{M} r_{t,j}^2, \tag{1}$$

where $r_{t,j}$ is an intraday scaled log-price difference $100 \times (\log(p_{t,j}) - \log(p_{t,j-1}))$. It consistently estimates daily integrated variance under a Brownian semimartingale when the sampling frequency shrinks to zero. Therefore, the availability of high frequency data allows one to model volatility as virtually observed quantity, rather than an unobserved latent variable (such as ARCH/GARCH Engle, 1982 and Bollerslev, 1986 or latent stochastic volatility models (e.g. Taylor, 1994) or Harvey, 2013). One such model for observed realized volatility is the heterogeneous autoregressive model.

The HAR-RV(1, 5, 21) includes lagged daily, weekly and monthly volatility in a model for daily realized volatility. The model for RV_t writes then as:

$$RV_t = \beta_0 + \beta_1 RV_{t-1} + \beta_2 RV_{t-1}^{w} + \beta_3 RV_{t-1}^{m} + \epsilon_t, \tag{2}$$

where RV_{t-1}^{w} denotes the lagged weekly volatility ($\frac{1}{5}\sum_{i=1}^{5} RV_{t-i}$), RV_{t-1}^{m} is the lagged monthly volatility ($\frac{1}{21}\sum_{i=1}^{21} RV_{t-i}$), and ϵ_t is a disturbance.

In recent years, many studies have used the HAR as a benchmark model for different purposes. For example, and most recently to name just a few, Souček and Todorova (2014) study the role of jumps in volatility transmission, notably on the S&P 500 within an HAR model. The HAR-RV lag specification is tested in Craioveanu and Hillebrand (2012) and provides support for the HAR-RV(1, 5, 21) model. More generally, Liu et al. (2013) assess the quality of the 5-min RV estimator. This important contribution (Liu et al., 2013) shows that among a set of 400 volatility estimators and using broad-asset classes, no-one significantly outperforms 5-min RV (which we use here). Liu et al. (2013) reach this conclusion notably through out-of-sample forecast evaluation in an HAR model. Therefore, Liu, Patton, and Sheppard's (Liu et al., 2013) results argue in favor of using 5-min RV, which is our choice in the present study.

3. Data

We construct daily volatility on the S&P 500 using intraday returns sampled at 5-min intervals.[3] The data is provided by

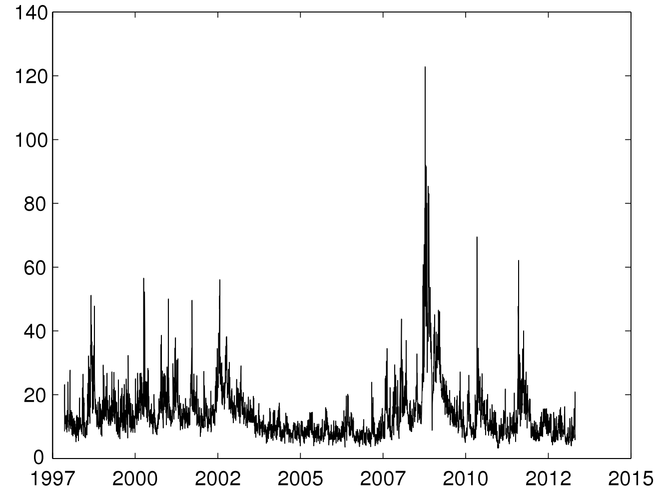

Fig. 1. S&P 500 daily realized volatility. Note: Daily realized volatility for 3836 trading days from November 11 1997 to April 26 2013. Daily volatilities are annualized: $\sqrt{260} \times \sqrt{RV_t}$, and computed with $M = 78$ 5-min intraday returns per day.

disktrading[4] and covers $T = 3836$ trading days from November 11 1997 to April 26 2013.

Returns at intraday period j and day t in percentage points are defined as:

$$r_{t,j} = 100 \times (\log(p_{t,j}) - \log(p_{t,j-1})), \quad j = 1 \ldots M, t = 1 \ldots T, \tag{3}$$

where $\log(p_{t,j})$ denotes the average log-price, computed with the last price of interval j and the first price of interval $j+1$. The S&P 500 market trading hours range from 9.30 am to 4 pm. Therefore, in line with the literature (see Andrada-Felix et al., 2013 and the reference therein), we ignore the overnight return. The first and last prices of the day ($p_{t,0}$ and $p_{t,M}$) are thus obtained as the first price of the first interval and the last price of the last interval, respectively. For the S&P 500 market sampled at 5-min, there are $M = 78$ intra-day intervals.

Fig. 1 describes the time series of annualized daily realized volatility for the S&P 500 3826 sample days, using 5-min returns. One can observe the relatively quiet period of low volatility between 2002 and 2007 before the subsequent unprecedented extreme levels during the financial crisis.

4. Specification test

In order to model realized volatility we take a flexible approach that utilizes methods that are robust to distributional and functional form misspecification. Suppose we are interested in modeling realized volatility RV_t using a set of covariates $X_{t-1} = [RV_{t-1}, RV_{t-1}^{w}, RV_{t-1}^{m}]$. In general we can think of RV_t as being some function of X_{t-1}:

$$RV_t = \phi(X_{t-1}) + \epsilon_t \tag{4}$$

where ϕ gives us the functional relationship between our explanatory variables and our measure of realized volatility. The only assumption we make is additive separability of the error term.[5] Typically, researchers assume ϕ is a linear function of the covariates such that $\phi(X_{t-1}) = X_{t-1}\beta$. In this paper we take a different approach to modeling volatility by allowing the function $\phi(X_{t-1})$ to be fully unknown. There are a number of advantages to taking

3 This sampling frequency is conventional in the literature and considered to provide a good trade-off between the need to sample at a frequency as high as possible, and the impossibility to do so because the noise starts to dominate the signal when the frequency is too high. Furthermore, this choice is motivated by Liu et al. (2013) who show that 5-min RV is not significantly outperformed by any other measure in a very broad set of 400 estimators.

4 This data (http://disktrading.is99.com/disktrading/) is used notably in Andrada-Felix et al. (2013) and Bos et al. (2012).

5 For nonadditive models see Matzkin (2003).