Erasmus University Rotterdam

Erasmus School of Economics

Bachelor Thesis Financial Econometrics

Introducing non-linearity to the HAR-RV model in order to control the market

Author:

Jelmar Faber (447466)

Supervisor:

Karel de Wit

Second assessor:

Xiao X Xiao

July 5, 2020

Abstract

Volatility is a well-known concept in financial markets for a considerable time. The HAR-RV volatility model is a renowned model that is proficient in forecasting volatility. Though, its ability to reproduce stylized facts is ambiguous. This paper will introduce a combination between the HAR-RV and a non-linear model to create a volatility model that aims to reproduce stylized facts to a higher degree. The resulting HAR-TAR model can reproduce the main stylized facts of financial returns relatively adequate. Moreover, the HAR-TAR models can outperform the HAR-RV model in the context of forecasting, but only in specific settings. Overall, this research shows that further investigation of the HAR-RV model in the area of non-linearity and other stylized facts can result in a deeper understanding of market behaviour.

Contents

1 Introduction

Over the past decades people have been more captivated with economics and markets. As interest increases people searched for ways to control or even exploit the market. In order to manage the market one must learn about its behaviour and characteristics. For instance, when a lot of uncertainty is present in the market, less transactions are occurring. Some of these key features are vital in order to create a beneficial market environment for the participants. In order to enable market participants to understand the behaviour of the market, researchers introduced the concept volatility according to Müller et al. (1993). The fluctuations of financial returns are captured by volatility. These fluctuations indicate how much uncertainty exists in the market, which directly affects market participants decisions. So, volatility is an extremely important tool to understand the behaviour of people in the market and thus the behaviour of the market itself. Overall, volatility is a popular measure to describe state of the market over time. As volatility is not something one observes it is rather difficult to obtain an exact and precise measure for the volatility. Therefore, models are required to create measures for the fluctuations in the market and support decision making. Over time a lot of research has been invested to capture volatility, however a model that is able to adequately and consistently forecast it, has not been found.

One of the first generation of volatility models is the autoregressive model according to Subbotin et al. (2009). These models are currently still used in practise even though they were introduced decades ago. Since, the autoregressive models seem to capture the most important characteristics of financial returns. As uncertainty existed regarding the proficiency of a volatility model researchers called for requirements that need to be met, when creating a volatility model. These requirements are paraphrased into stylized facts, which generally occur in financial returns. If a researcher creates a new model, it is mainly tested based on the stylized facts. There are numerous stylized facts as shown in Cont (2001). In the context of this paper the focus lies with the three most important stylized facts. The first main stylized fact is determined by the autocorrelations of the squared and absolute returns. The autocorrelations tend to remain significant for a long period of time, which show that the returns exhibit a long-memory property. The second main stylized fact is determined by heavy tails of the return distributions, which indicate that the returns are not normally distributed. Lastly, the third main stylized fact shows that returns exhibit a self-similar property according to Xu and Gençay (2003). This implies that the probability density functions of financial returns are similar at different scales. Meaning that there exists a certain relationship between the returns over time linked together by a factor.

The autoregressive models are generally capable of reproducing the main stylized facts, however only to a mediocre extent. Therefore, further research was required, this led to the HAR-RV model as shown in Corsi (2009). The volatility model is a restricted AR(22) model that is relatively effortless to understand. Though, the HAR-RV model can be classified as a type of autoregressive model, it provides superior characteristics of financial returns compared to other volatility models. Furthermore, it is capable of creating relatively adequate forecasts compared to standard autoregressive and ARFIMA models, as shown in Corsi (2009). However, the HAR-RV model is created in a setting where the three main stylized facts are the center of attention regarding characteristics of financial returns. Perhaps other stylized facts are overlooked and redirecting the focus to these stylized facts can turn out beneficial. In this paper, the center of attention is occupied by non-linearity of financial returns. In some markets the returns display a strong non-linear relation, therefore this can be a key characteristic according to Blasques et al. (2020). So, it can be beneficial to adjust the HAR-RV model in such a way that it can recreate the non-linearity property of financial returns.

Currently, a lot of research has been done regarding forecasting and estimation of volatility. A result of this research is the HAR-RV model, that will be examined in this paper. The model has been extended thoroughly in many successful ways where it has been able to forecast volatility relatively well. A venture that has not been examined yet is whether or not the HAR-RV model can be extended in order to reproduce the non-linearity of financial returns. This paper aims to develop the model further by means of combining it with the non-linear TAR model as shown in Tong (2012). In some circumstances the extended HAR-TAR model is able to reproduce the stylized facts more accurately compared to the HAR-RV model. Incidentally, this results in superior volatility forecasts in specific instances. The HAR-TAR volatility model will provide a better understanding of the movements of financial returns in markets. Ultimately, this model can be used by market participants, such as governments, to gain an enhanced perspective of the behaviour of the market. This will enable them to control or exploit the market and may create a less variable market that can lead to more collective prosperity.

2 Literature review

In order to understand where the current literature is lacking one must first derive how the present is founded. Volatility is a highly influential statistic that is thoroughly used, as shown in Subbotin et al. (2009). By means of volatility researchers and investors may gain a better understanding of the market according to Müller et al. (1993). For instance, if a lot of fluctuations are present in the market, this is reflected by means of a high volatility. In order to determine if a volatility model can be used in practise it is tested on various measures. The model is used to forecast volatility and it can be verified, to which extent the model is able to reproduce the stylized facts of financial returns according to Cont (2001). In this paper the main focus of stylized facts lies with long-memory, fat tails, self-similarity and non-linearity. The long-memory property proves challenging to reproduce according to Ding et al. (1993). The squared and absolute transformations of financial returns cause the difficulty. These tend to exhibit significant autocorrelations, which become gradually insignificant when the number of lags increase. So, in order to create a volatility model that can reproduce the stylized facts the autocorrelations of absolute and squared returns must be examined thoroughly. A volatility model should aim to reproduce the stylized facts adequately, therefore one ought to examine the possibilities to do so. One of the first volatility models that were introduced to reproduce the stylized facts are the autoregressive models. In order to obtain superior forecasts these models were altered and models, such as the ARFIMA model were created as shown in Granger and Joyeux (1980). These models showed promising results, however they had a flaw. The ARFIMA model and other comparable models are constructed with a fractional difference operator, such as $(1 - L)^d$. This operator causes a type of mixing between short-term and long-term parameters, therefore it becomes difficult to differentiate between short-memory and long-memory parameters according to Comte and Renault (1998). Due to the fact that it is challenging to differentiate between long-term and short-term parameters, the long-memory property becomes difficult to reproduce. To counteract this problem a continuous-time definition of these parameters is required.

At this point in time, it seems an impasse is reached, because autoregressive models cannot reproduce the long-memory property of financial returns according to Comte and Renault (1998). However, it turns out that this differs in practise according to LeBaron et al. (2001). Autoregressive models are able to reproduce the long-memory property to a certain extent, though this is not supported by theory. Therefore, the reproduction of this stylized fact by autoregressive models is viewed as apparent behaviour, according to LeBaron et al. (2001). This idea offers more opportunities to create relatively straightforward volatility models, because the amount of lags of an autoregressive model can remain relatively low in this context.

Although the reproduction of the long-memory property is not as challenging as before there is still major criticism regarding volatility models. In most volatility models it is assumed that there exists homogeneity among investors regarding their trading behaviour according to Müller et al. (1993). As one can imagine, this is not the case in practise, therefore a volatility model should incorporate Heterogenous Market Hypothesis according to Müller et al. (1993). This theory states there exists heterogeneity among investors. This concept was first implemented in Müller et al. (1997) to create the HARCH volatility model. The HAR-RV volatility model is based on the HARCH model and therefore implements the Heterogenous Market Hyptothesis, according to Corsi (2009). Incidentally, the HAR-RV model is a volatility cascade model, where long-term volatility has larger influence on short-term volatility than conversely. This implies the difference between long-term investors and short-term investors, which results in a new perspective regarding volatility models.

Currently, the HAR-RV model has been extended in various ways throughout the literature. For instance, Wen et al. (2016) introduce sixteen variations of the HAR-RV model, where they mainly include a jump element to ensure improved forecasts. They find that the jump element in the model indeed provides superior forecasts when implementing these various models. Furthermore, Bollerslev et al. (2016) exploit the lacking accuracy of the forecasts and aim to improve it by allowing the parameters of the model to change based on measurement error. Moreover, in Audrino and Knaus (2016) a lasso technique is performed in combination with the HAR-RV model. They question the amount of lags which are currently in the model, because the lasso technique led to other conclusions. Besides this, they show that the lasso technique performs as well as the HAR-RV model based on the stylized facts. Another interesting alteration of the HAR-RV volatility model is shown in Huang et al. (2016). In order to reproduce the stylized facts as good as possible the HAR-RV model is combined with the GARCH model. This should improve the capability of the model to reproduce the stylized facts, which holds especially for the long-memory property. Overall, various extensions of the model have been researched, however other stylized facts as shown in Cont (2001) are not taken into account.

All in all, various directions of the HAR-RV model have been investigated and exploited. Though the literature offers numerous extension of the HAR-RV model, the stylized facts have largely been neglected in this research. As mentioned earlier, there are other stylized facts that have not been thoroughly addressed yet, therefore I propose to investigate this venue. The non-linearity characteristic of some financial returns is relatively strong according to Blasques et al. (2020). So, a promising model may result if one is able to incorporate the non-linearity aspect of financial returns in the HAR-RV model. As the HAR-RV model is relatively straightforward, it would be preferred if the resulting model is as well, because it would be easy to implement in practise. There are various non-linear autoregressive models, the TAR model is one of these, as shown in Tong (2012). The TAR model is a regime switching model, where the regime is determined by a specific threshold. The threshold can be chosen in a way that it switches between states to capture the non-linearity of volatility adequately. I propose to combine the TAR model with the HAR-RV model to improve the reproduction of non-linearity in financial returns in some areas. Ultimately, this results in the HAR-TAR volatility model that may provide a better understanding of market behaviour.

3 Data

The data used to evaluate the volatility models consist of three different data sets. The data consists of tick-by-tick daily realized volatility series. The data sets are obtained from the Oxford-Man Institute website. They retrieve the prices from Thomson Reuters to compute the daily realized volatility. The prices used to compute the daily realized volatility are updated with a frequency of five minutes. In equation (1) the expression for the daily realized volatility is shown. The first data set is the AEX daily realized volatility series, ranging from 3rd of January 2000 to the 19th of May 2020 consisting of 5192 observations. The second data set is the FTSE daily realized volatility series ranging from 4th of January 2000 to 19th of May 2020. The third and last data set contains S&P500 daily realized volatility dating from 3rd of January 2000 to 19th of May 2020. No data cleaning has been done, doing so may erode the characteristics that exist in the data. As the HAR-TAR model is designed to reproduce these characteristics, cleaning would negate the purpose of this paper. Furthermore, the data sets originate from distinct economic markets, therefore the various models are tested in multiple contexts. This will provide a wide range of results and may therefore offer a robust conclusion.

Table 1: Summary statistics of the data sets

	AEX	FTSE	S&P500
Mean	1,867E-03	1,935E-03	1,772E-03
Std. Deviation	3.551E-03	4.754E-03	4,245E-03

Note. The mean and standard deviation are shown for the annualized daily realized volatility series of the AEX, FTSE and S&P500 data sets.

In order to create an evident view of the data some summary statistics are shown in Table 1. One may notice that the mean of realized volatilities are relatively equivalent and not significantly different from zero. The main difference between the data sets is exhibited through the standard deviation. The FTSE data set has the highest standard deviation, therefore the highest variation in data. Since a high amount of variation is challenging to forecast, one can assume that the FTSE forecasts are generally the worst. Furthermore, the HAR-TAR volatility model may provide superior forecasts for the data sets that have high variation. This may be due to non-linear characteristics that the data sets exhibit, though this may also be caused by other stylized facts.

4 Methodology

4.1 Realized volatility

In order to implement the aforementioned volatility cascade models, one should first specify the context in which these are created. It is assumed that the logarithmic prices result from a standard continuous time process according to Corsi (2009). As mentioned before, the stochastic volatility process ensures one can differentiate between long-term and short-term parameters according to Comte and Renault (1998). To implement stochastic volatility models in practise, an estimator for the instantaneous variance is required, which is also referred to as the integrated variance. The direct modelling of realized volatility as the integrated volatility results in a generally superior model, according to Andersen et al. (2001). Therefore, the realized volatility is used as an estimate for the daily volatility to create the HAR-RV model, as shown in Corsi (2009). In equation (1) an expression for the daily realized volatility is shown.

$$RV_t^{(d)} = \sqrt{\sum_{j=0}^{M-1} r^2_{t-j\cdot\Delta}} \tag{1}$$

The intraday returns $r^2_{t-j\cdot\Delta}$ used in equation (1) are computed according to equation (2).

$$r^2_{t-j\cdot\Delta} = p(t - j \cdot \Delta) - p(t - (j+1) \cdot \Delta), \qquad j = 1, \ldots, M \tag{2}$$

Where $p(\cdot)$ denotes the logarithmic price and $\Delta = 1d/M$ is the subinterval length on which the intraday returns are computed. The parameter M denotes the amount of intraday prices to model the realized volatility.

4.2 HAR-RV model

One of the foundations the HAR-RV volatility model is built on, is that volatility over longer time intervals has a stronger impact on volatility over a shorter time interval than conversely. To implement this, the HAR-RV model incorporates the Heterogenous Market Hypothesis as shown in Müller et al. (1993). This theory implies that there does not exist homogeneity between investors, meaning that investors differ from one another. In the context of the HAR-RV model, the difference between investors is paraphrased such that they differ regarding time horizons in which investors trade, as shown in Corsi (2009). Investors are split up in three main groups. Short-term investors, medium-term investors and long-term investors. The short-term investors represent daily rebalance frequency, medium-term investors rebalance on a weekly frequency and long-term investors rebalance on a monthly frequency. In the HAR-RV model, the realized volatilities drive the future daily realized volatility according to Corsi (2009). The expression for the HAR-RV model is shown in equation (3).

$$RV_{t+1d}^{(d)} = c + \beta^{(d)} RV_t^{(d)} + \beta^{(w)} RV_t^{(w)} + \beta^{(m)} RV_t^{(m)} + \omega_{t+1d} \tag{3}$$

Where $RV_t^{(d)}$, $RV_t^{(w)}$ and $RV_t^{(m)}$ are the daily, weekly and monthly realized volatilities respectively. The future daily realized volatility depends on the three lagged realized volatility components, therefore the long-term investors exert more influence on short-term investors than conversely in the model. As mentioned before, the HAR-RV model is created based on the apparent behaviour assumption according to Corsi (2009). Though, the HAR-RV model is a restricted AR(22) model, it should still be able to reproduce the main stylized facts relatively well. The HAR-RV model is the main foundation on which the HAR-TAR model is created, so one can now introduce this model.

4.3 HAR-TAR models

Though the HAR-RV model performs relatively well, it is still not able to fully capture all of the stylized facts. There are numerous stylized facts of financial returns as shown in Cont (2001), though a venue that is yet to be thoroughly examined is the non-linear characteristics of financial returns. Non-linearity is an important characteristic of financial returns according to Blasques et al. (2020), so extending the HAR-RV model to take non-linearity into account may prove beneficial. In order to reproduce non-linearity of financial returns, I propose to combine the HAR-RV model and a non-linear volatility model. The TAR volatility model is a relatively simple non-linear volatility model, as shown in Tong (2012). The resulting HAR-TAR volatility model is shown in equation (4).

$$RV_{t+1d}^{(d)} = c + \rho_t RV_t^{(d)} + \beta^{(w)} RV_t^{(w)} + \beta^{(m)} RV_t^{(m)} + \omega_{t+1d} \tag{4}$$

Where $\rho_t = \beta^{(d)} + \beta^{(n)} I(RV_{t-1}^{(d)} < k)$ and $RV_t^{(d)}$, $RV_t^{(w)}$ and $RV_t^{(m)}$ are the daily, weekly and monthly realized volatilities respectively. The variable ρ_t is the nonlinear element in the expression. The variable is mainly determined by the value of the lagged daily volatility compared to a threshold k. As stated before, this variable should improve the reproduction of non-linearity of financial returns regarding the volatility model. Assuming that the error term has zero-mean and is *i.i.d.* the standard least squares regression provides consistent estimates for $\beta^{(\cdot)}$ in the HAR-RV model according to Corsi (2009). As the HAR-TAR model is a linear transformation of the HAR-RV model, the estimates of the least squares regression provide consistent estimates as well. In order to create a robust view of the HAR-TAR model, multiple variations of the HAR-TAR model are introduced. The threshold value k differs for the varying HAR-TAR models. This paper introduces two variations of the HAR-TAR model, being the HAR-TAR(T) and the HAR-TAR(M) model. The HAR-TAR(T) model can be obtained through determining the threshold by means of the Threshold Grid Search Process (TGSP) algorithm as shown in Järås and Mohammadipour Gishani (2010). This algorithm determines the threshold k through splitting the data and performing multiple least squares regressions. The residuals are saved and the optimal threshold value is chosen such that the sum of squared residuals is minimised. When the threshold k is equal to the mean of the data the HAR-TAR(M) model can be obtained. The mean is an educated guess for the threshold, therefore one would expect the performance of the HAR-TAR(M) model to be weaker compared to the HAR-TAR(T) model. This is due to the fact that the threshold value is probably not optimal and will therefore likely result in larger residuals and incidentally cause weaker forecasts.

4.4 Simulation

In order to verify whether or not the HAR-RV and the HAR-TAR model are able to reproduce the stylized facts, a simulation is performed. The simulation that will be performed is done solely for the AEX data set, because this data set has the least variation in the data as shown in Table 1. The simulation for $M = 12$ intra-day returns is based on the simulation shown in Corsi (2009). To enact the simulation two expressions are required for the simulated returns and the simulated volatility, these are given in equation (5) and (6).

$$r_t^{(2h)} = \sigma_t^{(d)} \epsilon_t \tag{5}$$

$$\sigma_{t+2h}^{(d)} = c + \beta^{(d)} RV_t^{(d)} + \beta^{(w)} RV_t^{(w)} + \beta^{(m)} RV_t^{(m)} + \omega_{t+2h} \tag{6}$$

Where $RV_t^{(\cdot)}$ represents the respective daily, weekly and monthly realized volatility, $\sigma_{t+2h}^{(d)}$ the intra-day volatility and $r_t^{(2h)}$ the intra-day returns. To obtain realistic results, the parameters for the realized volatilities are calibrated based on the calibration in Corsi (2009). The values are assigned as follows, $\beta^{(d)} = 0.1$, $\beta^{(w)} = 0.08$ and $\beta^{(m)} = 0.08$. A least squares regression of equation (3) is performed using the AEX data set. The value for $\hat{c}$ is plugged into equation (6) and ω_t is assumed to be normally distributed with zero-mean and variance equal to the standard error of $\hat{c}$. Furthermore, it is assumed that $\epsilon_t \sim N(0, 1)$, as shown in Corsi (2009). To perform the simulation the intra-day returns are initialised with 22×12 zeros to compute the daily realized volatility as shown in equation (1). The daily realized volatilities are then used to compute the weekly and monthly realized volatility, as shown in Corsi (2009). After that, the intra-day volatility is simulated to implement in equation (5) and the simulated intra-day return is computed. This intra-day return is used to recalculate the realized volatilities and the intra-day volatility is simulated again to compute the new simulated intra-day return. This procedure is repeated until 150.000 days worth of daily realized volatilities can be computed. In the context of the HAR-TAR(T) regression, the parameters are calibrated equivalent to the simulation of the HAR-RV model, however in this setting the least squares regression of equation (4) is performed. The estimated value for $\beta^{(n)}$ is also obtained and implemented, along with the estimated constant and its standard error. Ultimately, this simulation is compared to the actual realized volatility series of the AEX data set to verify to what extent the stylized facts are reproduced.

5 Results

A description for the programming codes that are used to obtain the results is given in appendix A. The HAR-TAR(T) model should provide superior simulations to the HAR-TAR(M) model, so only the HAR-TAR(T) and HAR-RV model are simulated. In the context of forecasting both HAR-TAR models are evaluated, because the HAR-TAR(M) model can provide another meaningful benchmark.

5.1 Simulation

The results of the simulations are shown in this section. The intra-day returns are simulated corresponding to 150.000 daily realized volatilities to enable the simulations to converge to a distinct pattern. This pattern represents characteristics for the returns and daily realized volatility, according to Corsi (2009). One can view the daily actual return of the AEX data set and the simulated daily return series in Figure 1. A distinct characteristic of the actual daily returns is the variation of the series. There are two major occurrences around 2000 and 5000 daily returns, where the returns are relatively high compared to the other returns in the series. This is due to the fact that the daily realized volatility of the actual series is also relatively high, as shown in Figure 2. In both of the simulated return series there are instances where some returns are relatively higher compared to the other returns in the series. In absolute terms, this difference approximates the difference shown in the actual return series. However, the key discrepancy between the simulated daily returns and the actual daily returns exists in the relatively low returns. In the context of the HAR-RV simulated

Figure 1: Actual and simulated daily return series

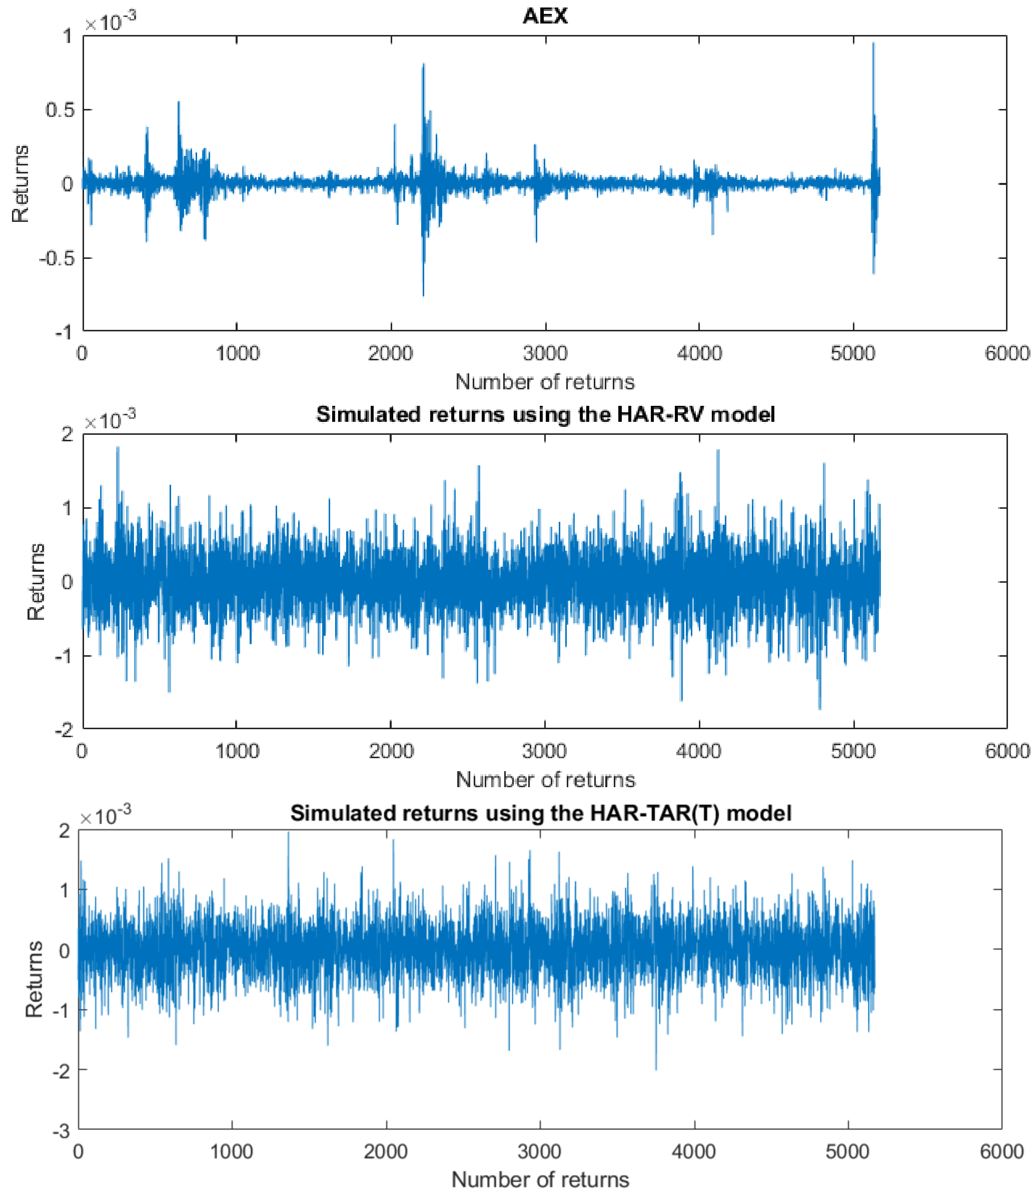

Note. Simulated and actual daily return series. The Figure shows plots from the actual AEX daily return series, simulated daily return series using the HAR-RV model and simulated daily return series using the HAR-TAR(T) model from top to bottom respectively.

returns, these are as large as the largest returns in the actual return series. This causes overall less contrast between returns, which causes a return series where there generally exists less variation in the context of the simulated daily returns. This implies that the simulated daily returns only seem to capture the variation in the actual returns series to a relatively mediocre extent. In the context of the HAR-TAR(T) simulated daily returns, the variation in the simulated series is richer compared to the HAR-RV simulated daily returns. This is due to the fact that the HAR-TAR(T) simulated daily returns exhibit generally more distinct returns. However, the large amount of variation that is present in the actual daily returns is not fully captured.

In Figure 2 the actual and simulated daily realized volatility series is shown for the HAR-RV and the HAR-TAR(T) model. As is the case with the daily return series, the actual AEX daily realized volatilities show a relatively large amount of variation in the series. The simulated daily realized volatility series show this amount of variation to a certain extent. However, the variation that is captured in the simulated daily volatility series is relatively small. There are patterns that the simulated series capture rather well. An example of this is the clustering of volatility at certain points. This is especially visible in the simulated daily realized volatility series using the HAR-TAR(T) model. When the realized volatility increases, it occurs gradually. This increase of volatility continues until the peak is reached and then decreases. The actual daily realized volatility series exhibits this characteristic as well, however the build up to the peak is faster and the peak itself is relatively higher. This major difference between the peaks of the actual and simulated realized volatilities may originate from the fact that the simulated series are heavily influenced by lagged realized volatilities. This causes the built up to the peak to be longer. Furthermore, the averaged lagged realized volatilities cause the peaks to be flattened and incidentally relatively smaller. Although, the simulations do not precisely replicate the actual daily realized volatility series, they grasp the essential fluctuations in the series. Furthermore, there exists relatively more variation in the HAR-TAR(T) simulated realized volatilities compared to the HAR-RV simulated realized volatilities. This is due to the fact that the built up to realized volatility peaks occur relatively faster, resulting in a realized volatility series with richer variation.

In Figure 3 the partial autocorrelations are shown for actual and simulated daily realized volatility series. The partial autocorrelations of the actual AEX daily realized volatility series remain significant for approximately six lags and tend to converge to insignificance after. This characteristic is portrayed in the simulated partial autocorrelations as well. This means that the models can replicate the long-memory property of financial returns relatively well.

Figure 2: Actual and simulated daily realized volatility series

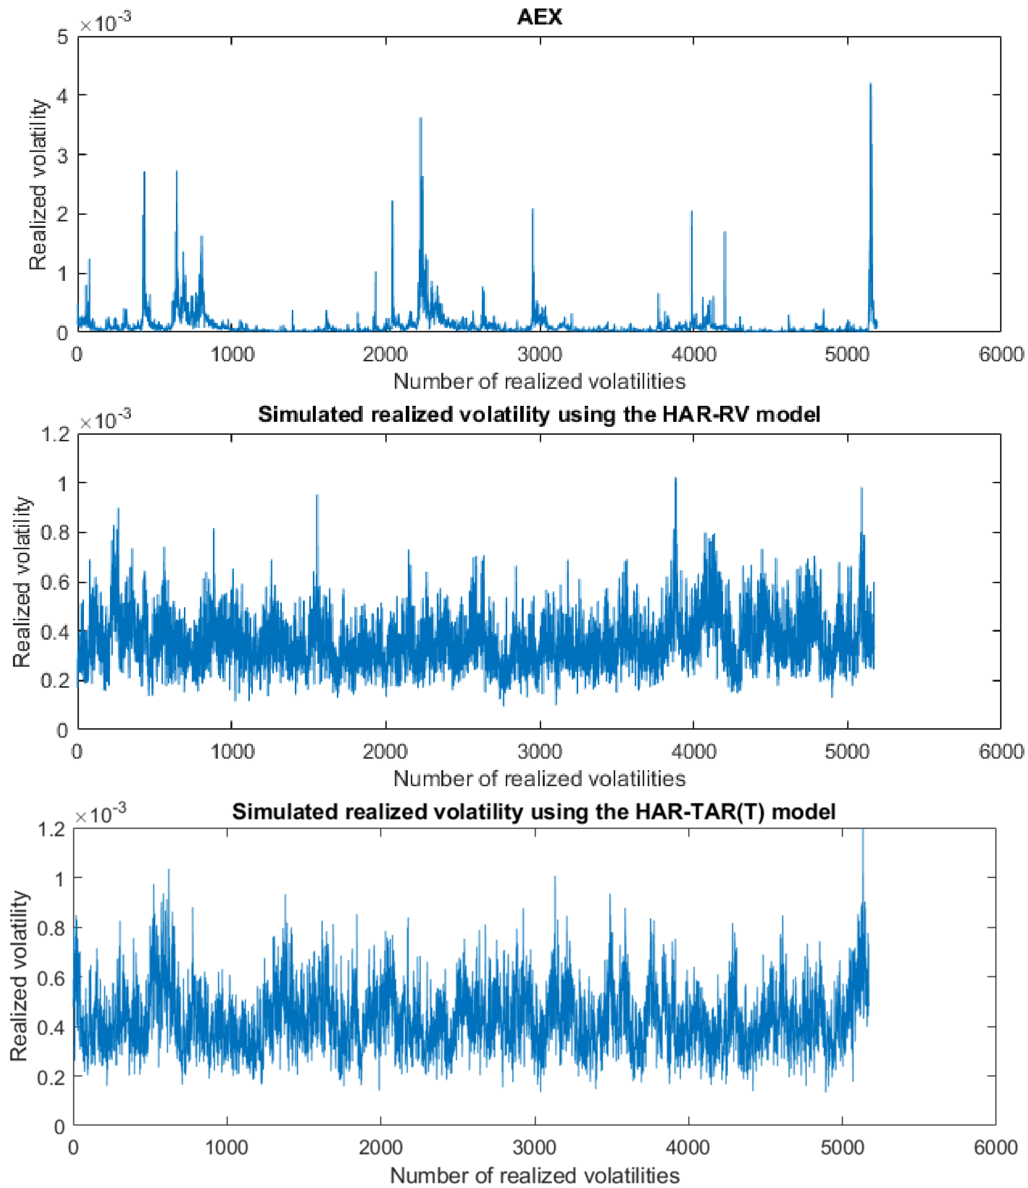

Note. Simulated and actual daily realized volatility series. The Figure shows plots from the actual AEX daily realized volatility series, simulated daily realized volatility series using the HAR-RV model and simulated daily realized volatility series using the HAR-TAR(T) model from top to bottom respectively.

There is a slight difference between the actual partial autocorrelations and the simulated partial autocorrelations. This occurs at the third lag, in the context of the actual series the autocorrelation is insignificant, however the simulated model shows significance at that lag. Furthermore, the HAR-TAR(T) partial autocorrelations tend to converge to insignificance slightly later than the HAR-RV partial autocorrelations.

Figure 3: Actual and simulated partial autocorrelations of daily realized volatility

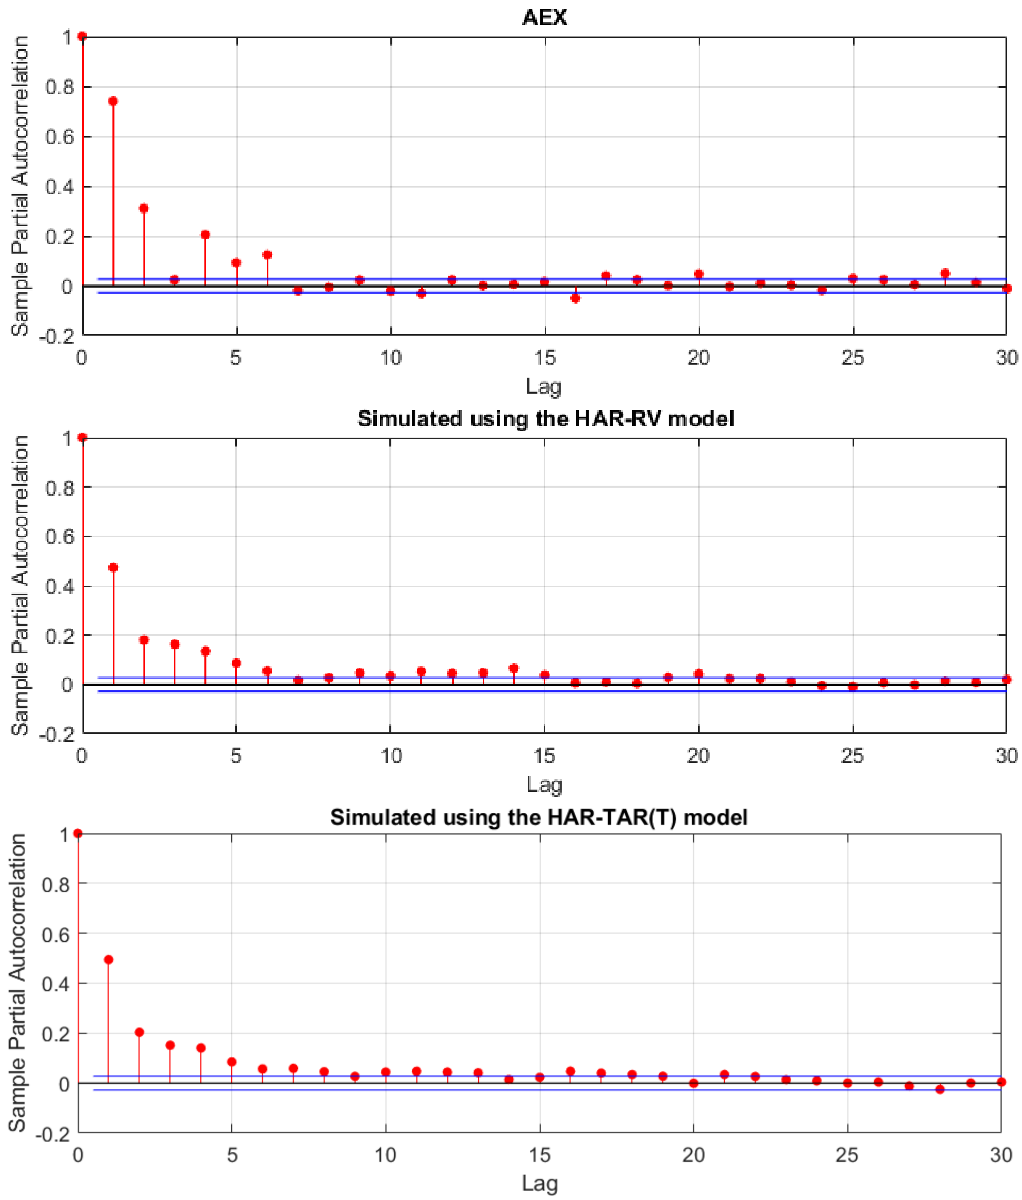

Note. Simulated and actual partial autocorrelations of daily realized volatility series. The Figure shows plots from the actual AEX partial autocorrelations, simulated partial autocorrelations using the HAR-RV model and simulated partial autocorrelations using the HAR-TAR(T) model from top to bottom respectively.

5.2 Evaluation of models

While it is important to examine the capabilities of the volatility models to reproduce the stylized facts, other aspects regarding the model are relevant as well. Such as, the fit of the volatility model to the data and its ability to adequately forecast volatility.

5.2.1 HAR-RV

The HAR-RV model is the main foundation of the HAR-TAR models, therefore it serves as a proper benchmark to the HAR-TAR models. In Table 2 the least squares estimates of the HAR-RV model are shown for the three data sets.

Table 2: HAR-RV least squares estimates

	AEX	FTSE	S&P500
c	1,274E-05	2,036E-05	1,220E-05
	(4,306)	(3,911)	(2,558)
β^d	0,405	0,112	0,271
	(5,260)	(1,760)	(2,571)
β^w	0,415	0,580	0,529
	(5,727)	(4,083)	(3,746)
β^m	0,072	0,143	0,091
	(1,385)	(1,807)	(1,161)

Note. The least square estimates are shown of the HAR-RV model for the AEX, FTSE and S&P500 data sets. Where t-statistics are shown in parentheses based on Newey-West standard errors with serial correlation order 5.

The estimates in Table 2 are obtained performing the least squares regression of equation (3). The least squares estimates are consistent, though the Newey-West standard errors are implemented to take possible serial correlation into account. From the regression results one can state that the constant and the weekly realized volatility are statistically significant across all data sets. Furthermore, the daily realized volatility is significant for the AEX and S&P500 data set, however this is not the case for the FTSE data set. This may be due to the volatile nature of the data set as it can lead to noisier estimates, therefore the coefficient is not significant at a 5% significance level. Moreover, the monthly realized volatility is not significant at a 5% significance level across all data sets. This is an interesting observation, as the monthly realized volatility is an average of daily realized volatilities. This should provide less noisier estimates for the monthly realized volatility, however it seems the opposite is the case.

5.2.2 HAR-TAR(T)

In order to verify whether or not the HAR-TAR volatility model is capable of reproducing the non-linear property of some financial returns a least squares regression is performed. If the coefficient $\beta^{(n)}$ is significant there exists non-linearity in financial returns and the HAR-TAR model is able to capture a substantial part of it. In Table 3 the least squares estimates of the HAR-TAR(T) model are shown across the data sets.

Table 3: HAR-TAR(T) least squares estimates

	AEX	FTSE	S&P500
c	1,667E-05	2,180E-05	1,718E-05
	(4,501)	(3,603)	(3,172)
β^d	0,408	0,113	0,275
	(5,234)	(1,743)	(2,564)
β^n	-0,084	-0,028	-0,122
	(-1,615)	(-0,440)	(-2,048)
β^w	0,410	0,578	0,522
	(5,568)	(4,053)	(3,650)
β^m	0,070	0,142	0,090
	(1,343)	(1,799)	(1,152)

Note. The least square estimates are shown of the HAR-TAR(T) model for the AEX, FTSE and S&P500 data sets. Where t-statistics are shown in parentheses based on Newey-West standard errors with serial correlation order 5.

When one examines Table 3 a majority of the results are complementary to the results in Table 2. The weekly realized volatility and the constant are significant across all data sets at a 5% significance level. Furthermore, the daily realized volatility is only significant for the AEX and S&P500 data set and the monthly realized volatility is not significant across all data sets. The major contrast lies with the added non-linear explanatory variable. It appears that the variable is only significant for the S&P500 data set at a 5% significance level. As mentioned before, not all financial returns exhibit non-linear characteristics, so it is not surprising that the non-linear element is only significant in the context of the S&P500 data set. So, it is expected that the HAR-TAR(T) model can outperform the HAR-RV model in the context of the S&P500 data set and may provide weaker forecasts regarding the AEX and FTSE data sets.

5.2.3 Fit of models

The fit of an econometric model is highly important, because it indicates whether or not this model can describe the variation in the data adequately. The AIC and SIC measures are calculated according to Hurvich and Tsai (1991). In Table 4 a multiple hypothesis test is shown for the HAR-RV model and the HAR-TAR(T) model, along with the respective AIC and BIC values across the data sets.

Table 4: Goodness-of-fit tests for the HAR-RV and HAR-TAR(T) model

	AEX	FTSE	S&P500
F-test HAR-RV	17,444	9,332	23,687
F-test HAR-TAR(T)	17,247	9,320	23,357
	(1,908)	(1,908)	(1,908)
	AIC		
AR(22)	-7,752E+04	-7,118E+04	-7,416E+04
HAR-RV	-6,655E+04	-6,049E+04	-6,326E+04
HAR-TAR(T)	-6,656E+04	-6,049E+04	-6,326E+04
	BIC		
AR(22)	-7,736E+04	-7,102E+04	-7,400E+04
HAR-RV	-3,786E+04	-3,216E+04	-3,509E+04
HAR-TAR(T)	-3,787E+04	-3,216E+04	-3,510E+04

Note. The F-tests represent the multiple hypothesis test between the unrestricted AR(22) model and the restricted HAR-RV or HAR-TAR(T) model with the 1% critical value in parentheses and the respective AIC and BIC values for the models.

To determine whether or not the fit of the HAR-TAR(T) and the HAR-RV model is adequate for the data sets, multiple hypothesis tests are performed as shown in Table 4. The HAR-RV and the HAR-TAR(T) models are tested against an unrestricted AR(22) model. The multiple hypothesis tests across all data sets indicate that the null hypothesis is rejected for both the HAR-RV and the HAR-TAR(T) model. This implies that the fit of the models is relatively adequate and reflects the key amount of variation in the data more compared to the unrestricted AR(22) model. The values for the multiple hypothesis test imply that the fit of the HAR-RV model is slightly better compared to the fit of the HAR-TAR(T) model to the data. Furthermore, the AIC values given in Table 4 are relatively equivalent between the AR(22) model and the volatility cascade models. The major difference occurs in the context of the BIC values, where the AR(22) model shows relatively superior performance across all data sets compared to the volatility cascade models. This difference may be due to the fact that the BIC measure penalises more compared to the AIC measure. The lower BIC value implies that the AR(22) model has a better fit to the data compared to the volatility cascade models.

5.2.4 Forecast evaluation

Another popular measure to test the performance of a volatility model is by forecasting, besides the reproduction of stylized facts. The volatility cascade models are compared to the AR(1) and AR(3) model in order to examine whether or not the volatility cascade models are able to forecast relatively adequately. In Table 5 the performance measures for the one-day-ahead in-sample forecasts are given for the data sets.

Table 5: One day ahead in-sample forecasts

	AEX			FTSE			S&P500		
	RMSE	MAE	R^2	RMSE	MAE	R^2	RMSE	MAE	R^2
AR(1)	1,510E-04	5,620E-05	0,548	2,616E-04	7,816E-05	0,243	1,935E-04	6,301E-05	0,481
AR(3)	1,434E-04	5,186E-05	0,592	2,467E-04	6,898E-05	0,327	1,771E-04	5,674E-05	0,565
HAR-RV	1,426E-04	4,967E-05	0,598	2,410E-04	6,493E-05	0,360	1,782E-04	5,558E-05	0,561
HAR-TAR(T)	1,426E-04	4,990E-05	0,598	2,410E-04	6,497E-05	0,360	1,781E-04	5,590E-05	0,562
HAR-TAR(M)	1,426E-04	4,990E-05	0,598	2,410E-04	6,497E-05	0,360	1,781E-04	5,590E-05	0,562

Note. One-day-ahead in-sample forecast evaluations of AR(1), AR(3), HAR-RV, HAR-TAR(T) and HAR-TAR(M) models for the various data sets. Performance measures are the Root Mean Squared Error (RMSE), Mean Absolute Error (MAE) and R^2 of the Mincer-Zarnowitz regression.

To evaluate the in-sample forecasts three different performance measures are used, as shown in Table 5. The Root Mean Squared Error (RMSE), Mean Absolute Error (MAE) and the R^2 from the Mincer-Zarnowitz regression. In the context of the AEX data set one can observe that all volatility models provide relatively adequate in-sample forecasts, as the values for the performance measures are relatively equivalent. The in-sample forecasts evaluated with the FTSE data set are relatively equivalent as well except for the AR(1) model which performs relatively weak. However, in the context of the S&P500 data set, the AR(3) model is surprisingly able to provide superior in-sample forecasts compared to the other models. Furthermore, the difference between the performance of the HAR-TAR(T) and the HAR-TAR(M) model is negligible. Overall, the volatility cascade models are generally able to produce relatively adequate in-sample forecasts.

When one intends to implement a volatility model in practise it is vital to learn about its out-of-sample performance, since this approximates the performance of the model in practise. Therefore, out-of-sample forecasts are evaluated with three different forecast horizons.

Table 6: Out-of-sample forecasts for different forecast horizons

	1 day			1 week			2 weeks		
	RMSE	MAE	R^2	RMSE	MAE	R^2	RMSE	MAE	R^2
					AEX				
AR(1)	1,525E-04	5,205E-05	0,488	1,901E-04	8,246E-05	0,358	2,099E-04	9,461E-05	0,029
AR(3)	1,459E-04	4,795E-05	0,530	1,909E-04	6,748E-05	0,595	2,257E-04	8,506E-05	0,069
HAR-RV	1,503E-04	4,554E-05	0,539	1,958E-04	5,596E-05	0,619	1,826E-04	6,595E-05	0,656
HAR-TAR(T)	1,509E-04	4,600E-05	0,536	1,677E-04	5,506E-05	0,628	2,436E-04	7,748E-05	0,077
HAR-TAR(M)	1,508E-04	4,640E-05	0,537	1,691E-04	5,600E-05	0,620	2,528E-04	7,928E-05	0,043
					FTSE				
AR(1)	2,880E-04	8,016E-05	0,186	3,122E-04	1,092E-04	0,031	3,162E-04	1,121E-04	0,006
AR(3)	2,745E-04	7,139E-05	0,244	3,003E-04	8,307E-05	0,433	3,181E-04	9,904E-05	0,359
HAR-RV	2,692E-04	6,621E-05	0,308	3,012E-04	7,480E-05	0,535	2,957E-04	8,257E-05	0,424
HAR-TAR(T)	2,697E-04	6,660E-05	0,308	2,798E-04	6,954E-05	0,491	3,046E-04	8,222E-05	0,335
HAR-TAR(M)	2,702E-04	6,715E-05	0,305	2,846E-04	7,346E-05	0,470	3,140E-04	9,356E-05	0,325
					S&P500				
AR(1)	2,043E-04	6,340E-05	0,511	2,697E-04	9,943E-05	0,221	2,889E-04	1,093E-04	0,030
AR(3)	2,023E-04	5,916E-05	0,516	2,500E-04	7,223E-05	0,552	2,729E-04	8,452E-05	0,423
HAR-RV	2,104E-04	5,814E-05	0,526	2,539E-04	6,745E-05	0,607	3,171E-04	8,072E-05	0,534
HAR-TAR(T)	2,108E-04	5,909E-05	0,525	2,399E-04	6,595E-05	0,531	3,741E-04	9,033E-05	0,031
HAR-TAR(M)	2,107E-04	5,883E-05	0,525	2,393E-04	6,697E-05	0,540	3,156E-04	8,875E-05	0,096

Note. Out-of-sample forecast evaluations of 1-day, 1-week and 2-week ahead forecasts for AR(1), AR(3), HAR-RV, HAR-TAR(T) and HAR-TAR(M) models. All models are re-estimated using a moving window of 1000 observations. Performance measures are the Root Mean Squared Error (RMSE), Mean Absolute Error (MAE) and R^2 of the Mincer-Zarnowitz regression.

The evaluation of the one-day-ahead, one-week-ahead and two-weeks-ahead forecasts are shown in Table 6. The performance measures shown in Table 6 are computed with forecasts that are computed using a moving window of 1000 days to provide a credible view of the forecasts. When one examines the performance measures, it becomes evident that the HAR-RV model is generally superior, however not in all instances. Only in the context of the AEX data set for the one-week-ahead forecasts, the HAR-TAR models are outperforming the HAR-RV model. In the context of one-day-ahead forecasts the differences between the volatility cascade models are relatively small. As the forecast horizon increases, the differences between the performance measures do as well. The AR(1) and the AR(3) model are generally outperformed by the HAR-TAR models, however in some instances the AR(3) model outperforms the HAR-TAR models. This only holds with a relatively high forecast horizon, though it is surprising as the difference between performance measures is in some cases relatively large. In the context of the AEX and FTSE data set this is expected to a certain extent, as these data sets showed no significant signs of non-linearity. However, the difference is also relatively large in the context of the S&P500 data set. Perhaps a model that has a relatively low amount of lags is a better fit to the S&P500 data set. Furthermore, Table 6 shows there is a distinct difference between the out-of-sample performance of the HAR-TAR(T) and the HAR-TAR(M) model. There are a few instances where the HAR-TAR(M) model is superior, however in general the HAR-TAR(T) model outperforms the HAR-TAR(M) model.

6 Conclusion

Volatility is an immensely important concept, because it can help people to understand the behaviour of the market. As one manages to understand how the market will react to situations, opportunities occur to control or exploit the market. Gaining a more profound understanding of volatility can especially be beneficial in creating more control while economies take a downfall. In order to better understand the market, volatility models were introduced and implemented to take advantage of uncertain times paired with economic downfall.

There are two major ways to test whether a volatility model is proficient, through reproduction of stylized facts and forecasting volatility. In this paper the main focus lies with the reproduction of stylized facts. They prove to be a main guideline and are vital to understand to which extent the volatility models reproduce these special characteristics. The renowned HAR-RV volatility model was created as a tool to forecast volatility and there was only focus on reproduction of long-memory, heavy tails and self-similarity of financial returns. Though, the HAR-RV model seems relatively capable in reproducing these characteristics there are several others. In this paper the HAR-TAR model is introduced, this is essentially a combination between the HAR-RV model and the non-linear TAR model. The HAR-TAR volatility model is created to improve the reproduction of stylized facts of financial returns compared to the HAR-RV model.

While not all financial returns exhibit non-linear characteristics, the HAR-TAR model could provide a more adequate reproduction of the stylized facts. Furthermore, a least squares regression shows that in some contexts there exists non-linearity in financial returns, so the HAR-TAR models could expand knowledge on fluctuations in the market. Through a simulation it becomes clear that the HAR-TAR model is able to reproduce characteristics of financial returns to a higher degree compared to the HAR-RV model. Perhaps, this is due to the non-linear element, nevertheless the simulation using the HAR-TAR model shows promising results. Especially, in the area of creating variation in the simulated volatility series and creating clustered volatility with richer dynamics compared to the simulation with the HAR-RV model.

On the other hand, forecasting volatility remains an important function a volatility model can have. Overall, the HAR-TAR model is not able to provide forecasts that are as adequate as the forecasts computed with the HAR-RV model. This holds especially when the forecast horizon increases, however in the context of one-week-ahead forecasts there is an instance where the HAR-TAR models are able to outperform the HAR-RV model.

The results presented in this paper seem promising, however there are a few drawbacks. The main reproduction of some stylized facts seem to have improved when the HAR-TAR model is implemented. However, it is not clear to what extent the non-linear characteristic of financial returns is reproduced. Furthermore, the HAR-TAR model seemed to falter compared to the HAR-RV model in the context of reproducing significant autocorrelation patterns. Although, the discrepancy is minor, it can harm the further interpretation of volatility in general, instead of contributing to its understanding. Overall, the HAR-TAR model shows that combinations between the HAR-RV model and a model based on stylized facts can prove beneficial in the reproduction of stylized facts. Extended research can investigate whether the stylized facts are consistently reproduced to an adequate extent. Moreover, other models can be created that aim to reproduce other stylized facts in combination with the HAR-RV model to obtain an improved volatility model.

Though, one cannot conclude that the HAR-TAR models are overall superior to the HAR-RV model, there is still a light at the end of the tunnel. This research showed that the combination of the HAR-RV model with a non-linear model shows promising results in the context of reproducing stylized facts. This may expand the existing knowledge on fluctuations in the market and may offer entities such as governments a tool that can help them control the market. This may in turn lead to a more stable market that can contribute to collective prosperity.

References

Torben G Andersen, Tim Bollerslev, Francis X Diebold, and Paul Labys. The distribution of realized exchange rate volatility. *Journal of the American statistical association*, 96(453):42–55, 2001.

Francesco Audrino and Simon D Knaus. Lassoing the har model: A model selection perspective on realized volatility dynamics. *Econometric Reviews*, 35(8-10):1485–1521, 2016.

Francisco Blasques, Siem Jan Koopman, and André Lucas. Nonlinear autoregressive models with optimality properties. *Econometric Reviews*, 39(6):559–578, 2020.

Tim Bollerslev, Andrew J Patton, and Rogier Quaedvlieg. Exploiting the errors: A simple approach for improved volatility forecasting. *Journal of Econometrics*, 192(1):1–18, 2016.

Fabienne Comte and Eric Renault. Long memory in continuous-time stochastic volatility models. *Mathematical finance*, 8(4):291–323, 1998.

Rama Cont. Empirical properties of asset returns: stylized facts and statistical issues. 2001.

Fulvio Corsi. A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2):174–196, 2009.

Zhuanxin Ding, Clive WJ Granger, and Robert F Engle. A long memory property of stock market returns and a new model. *Journal of empirical finance*, 1(1):83–106, 1993.

Clive WJ Granger and Roselyne Joyeux. An introduction to long-memory time series models and fractional differencing. *Journal of time series analysis*, 1(1):15–29, 1980.

Zhuo Huang, Hao Liu, and Tianyi Wang. Modeling long memory volatility using realized measures of volatility: A realized har garch model. *Economic Modelling*, 52:812–821, 2016.

Clifford M Hurvich and Chih-Ling Tsai. Bias of the corrected aic criterion for underfitted regression and time series models. *Biometrika*, 78(3):499–509, 1991.

Jacob Järås and Azadeh Mohammadipour Gishani. Threshold detection in autoregressive non-linear models, 2010. Student Paper.

Blake LeBaron et al. Stochastic volatility as a simple generator of apparent financial power laws and long memory. *Quantitative Finance*, 1(6):621–631, 2001.

Ulrich A Müller, Michel M Dacorogna, Rakhal D Davé, Olivier V Pictet, Richard B Olsen, and J Robert Ward. Fractals and intrinsic time: A challenge to econometricians. *Unpublished manuscript, Olsen & Associates, Zürich*, page 130, 1993.

Ulrich A Müller, Michel M Dacorogna, Rakhal D Davé, Richard B Olsen, Olivier V Pictet, and Jacob E Von Weizsäcker. Volatilities of different time resolutions—analyzing the dynamics of market components. *Journal of Empirical Finance*, 4(2-3):213–239, 1997.

Alexander Subbotin, Thierry Chauveau, and Kateryna Shapovalova. Volatility models: from garch to multi-horizon cascades. *Nova Science Publishers, Forthcoming*, 2009.

Howell Tong. *Threshold models in non-linear time series analysis*, volume 21. Springer Science & Business Media, 2012.

Fenghua Wen, Xu Gong, and Shenghua Cai. Forecasting the volatility of crude oil futures using har-type models with structural breaks. *Energy Economics*, 59:400–413, 2016.

Zhaoxia Xu and Ramazan Gençay. Scaling, self-similarity and multifractality in fx markets. *Physica A: Statistical Mechanics and its Applications*, 323:578–590, 2003.

Appendices

A Description of programming codes

Name of program	Description
AR1	Program to compute one-day-ahead in-sample forecasts for the AR(1) model.
AR1out	Program to compute k-days-ahead out-of-sample forecasts for the AR(1) model with $k = 1, 5, 10$.
AR3	Program to compute one-day-ahead in-sample forecasts for the AR(3) model.
AR3out	Program to compute k-days-ahead out-of-sample forecasts for the AR(3) model with $k = 1, 5, 10$.
AR22	Program to compute the AIC and BIC values and obtain residuals from a least squares regression for the AR(22) model.
dailyret	Program to transform the intra-day returns in the simulation to daily returns.
fit	Program to perform the multiple hypothesis test between the unrestricted AR(22) model and the restricted HAR-RV or HAR-TAR(T) model.
forharrv	Program to compute in-sample one-day-ahead forecasts for the HAR-RV model.
forhartar	Program to compute in-sample one-day-ahead forecasts for the HAR-TAR models.
harrv	Program to perform a least squares regression for the HAR-RV model. The AIC and BIC measures are computed as well.
harrvout	Program to compute k-days-ahead out-of-sample forecasts for the HAR-RV model with $k = 1, 5, 10$.
hartar	Program to perform a least squares regression for the HAR-TAR(T) model. The AIC and BIC measures are computed as well.
hartarout	Program to compute k-days-ahead out-of-sample forecasts for the HAR-TAR models with $k = 1, 5, 10$.
meanabserror	Program to compute the MAE for the in-sample AR(1) forecasts.
meanabserror3	Program to compute the MAE for the in-sample AR(3) forecasts.
meanabserrorhar	Program to compute the MAE for the in-sample HAR-RV and HAR-TAR forecasts.

MZreg1	Program to compute the R^2 of the Mincer-Zarnowtiz regression for the in-sample AR(1) forecasts.
MZreg3	Program to compute the R^2 of the Mincer-Zarnowtiz regression for the in-sample AR(3) forecasts.
MZreghar	Program to compute the R^2 of the Mincer-Zarnowtiz regression for the in-sample HAR-RV and HAR-TAR forecasts.
NeweyWest	Program to compute the Newey-West standard errors of regression estimates with serial correlation order 5.
plotten	Program to plot the figures for the simulation with axes.
PMout	Program to compute performance measures of out-of-sample forecast for all models. With the MAE, RMSE and R^2 of the Mincer-Zarnowtiz regression as performance measures.
realvol	Program to compute the simulated daily realized volatiliy based on the simulated intra-day returns.
rootmeansqrerror	Program to compute RMSE of the in-sample AR(1) forecasts.
rootmeansqrerror3	Program to compute RMSE of the in-sample AR(3) forecasts.
rootmeansqrerrorhar	Program to compute RMSE of the in-sample HAR-RV and HAR-TAR forecasts.
simulatieharrv	Program to perform the simulation for the HAR-RV model and to compute the returns for the actual data.
simulatieharrv	Program to perform the simulation for the HAR-TAR(T) model.
threshold	Program to compute the optimal threshold for the HAR-TAR(T) model based on the TGSP algorithm.
ttest	Program to compute the t statistic to evaluate the least squares estimates.