Modeling the volatility of BFCIUS and creating a Financial Conditions Index

Yan Kit Lee 331522

June 28, 2013

Abstract

As the economy is still recovering from one of the greatest financial downturns in the last decades, we have experienced that this period of recession is paired with financial stress. This financial stress might even be the reason that we have suffered from events like the credit crunch, and thus should be prevented from happening as much as possible. Therefore, this research looks into two different aspects of financial conditions. By forecasting several indicators of financial stress, I make an effort to convert the financial conditions into numerical values, resulting in a financial conditions index. Moreover, the index from Bloomberg is analyzed extensively in this research, in order to model the volatility of the financial conditions. This paper shows, that by correct modeling, the volatility of the financial conditions can be predicted.

1 Introduction

Stress in the financial sector is a phenomenon that is not desirable, as it can lead to financial disasters, like defaults of major financial institutions. However, sometimes it can be hard to tell whether or not there is financial stress, because there is not a real definition for this term. Earlier this year, Lee et al. [2013] have tried to create an index to measure this stress. This Financial Stress Index (FSI) includes information from variables that should effect the financial sector. Of course such an index is not unique, it has been created by many people and institutions. One of these institutions is Bloomberg. Bloomberg is an enormous company with a massive financial database to supply the world with financial data. In fact, up to one third of the worlds financial data is supplied by Bloomberg. Their index to measure the financial conditions, the Bloomberg Financial Conditions Index for the United States (BFCIUS), is generally considered accurate.

One of the properties of financial data is that the volatility is not constant, but time varying. This is called volatility clustering. Especially in data that measures the financial conditions over time, this property must be present. Therefore, if one would like to model and forecast financial data, it is important to take the volatility clustering into account. This is something Lee et al. [2013] have not done in their research. For the modeling of their index, the FSI, it was assumed that the volatility is constant over the whole period, which could be incorrect. If the volatility is somewhat predictable one could even consider modeling it. Therefore the goal of this research is to model the volatility of this BFCIUS and I eventually hope to be able to predict the volatility of the financial stress in the future.

Another unfortunate finding from Lee et al. [2013] was the poor forecast ability of one of the indexes. This particular FSI scored well on predicting the financial downturns, but it was too difficult to model, which led to inaccurate forecasts. Therefore, in this paper I want to see if I can improve on that part. I want to take a slightly different approach by modeling and forecasting the individual variables in stead of the index. Using these forecasted variables I will build another index, a Financial Conditions Index (FCI), which hopefully can closely match the BFCIUS. Note that this index is exactly the opposite of the FSI. The FCI goes down in periods of stress, whereas the FSI goes up.

This paper is structured as follows. Section 2 describes what data has been used for this research. Section 3 shows what methods are used to model the volatility of the BFCIUS and to create a FCI. In section 4 the results of this research are presented. Section 5 concludes.

2 Data

The variables used in this research are listed in Table 1. This list of variables and spreads is almost the same as the one in Lee et al. [2013], the only variable that is new is the BFCIUS, which is plotted in Figure 1. This index tracks the overall stress in the U.S. money market, bond market, and equity market and provides a useful gauge to assess the availability and cost of credit in the U.S. financial market. The credit crunch in 2008 is clearly visible as the BFCIUS steeply drops around observation 2600.

All these variables should be somehow related to the financial conditions in the United States of America. The last two variables on the list require some further explanation. The Banking β is constructed by taking $\beta = \frac{\text{cov(XLF,SPY)}}{\text{var(SPY)}}$ and the Financial σ is constructed by regressing VIX on SPY, where the residuals of this regression represent the Financial σ.

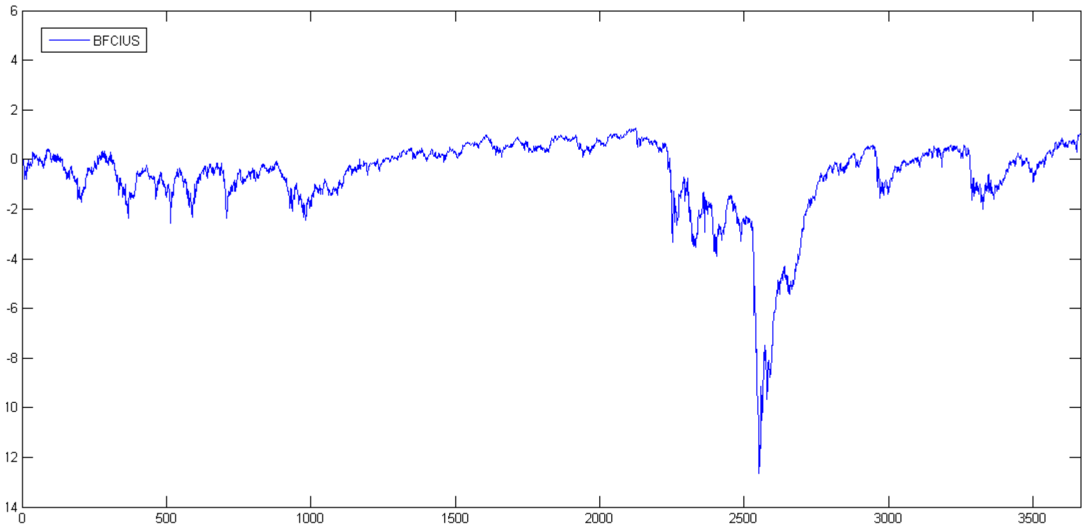

Figure 1: BFCIUS from 4 January 1999 until 16 January 2013

2.1 Interpolation

The set of data consists of daily observations from 4 January 1999 until 16 January 2013. Many of these financial variables do not change in the weekends, so leaving these out will result in a dataset of 3663 observations. Unfortunately, some variables have missing data points in the sample. However, this is not a major problem, because the data gaps are not larger than two observations, and most of the times not even larger than one observation. Therefore, this problem is fixed by using the linear interpolation method. This method works as follows. If there are one or two missing data points in a variable, a line is constructed between the observations preceding and following this gap. The missing data points are assigned a value corresponding to this line. The observations which missed out on one or more variables would have gone to waste if the data is not interpolated, but now I can use the full dataset.

2.2 Taking differences

The next step is to transform the variables to their first differences. After this transformation each variable has a positive sign if the value is higher than the previous observation, and a negative sign if it is lower than the previous observation. In this research all work is done with the differences

Indicators	Explanation in case of increase
BFCIUS	Increase indicates better financial conditions
High Yield ("risky bonds")	Compensation for risk increases, which indicates less secure return
LIBOR	Compensation for risk of interbank loans
DTB3 (Daily Treasury Bill rate	Compensation for risk of treasury bills
DED3 (Daily 3-month €/$ deposit rate)	Compensation for exchange rate risk
DFF (Daily Federal Funds rate)	Compensation for risk for overnight loans for depository institutions
VIX: implied volatility (of options)	Indicates the market's expectation of stock market. Increase shows more stress
XLF: Financial Select Sector SPDR	Increase indicates confidence in the financial sector
SPY: Exchange traded fund of S&P 500	Increase indicates confidence in the S&P 500
EUDex (€/$ exchange rate)	The higher the rate, the less valuable the dollar, indicates a decrease of the U.S. economy
2Y, 10Y, 30Y (2, 10, 30 year government bonds)	Increase indicates compensation for increased risk on government bonds
Yield curve: 10y - 3M treasury bill	Indicates that the short term risk decreases
AAA- 10Y	Indicates that the return on AAA (rated by Standard and Poor's) bonds increases
Baa - 10Y	Indicates that the return on Baa (rated by Moody's) increases
HY - 10Y	Indicates that the return on high yield bond increases
TED spread (LIBOR - yield on treasury bills)	Indicates that the interbank rate increases, compensates for higher risk
Banking β	Returns of the banking sector are more volatile, than the return on the overall market
Financial σ	Indicates the volatility of the financial market

Table 1: List of variables. Note that the spreads starting from Yield Curve were created by myself.

of the original variables. In the rest of this paper these differences will just be referred to as 'variables'.

Since the variables are already daily changes, the chances are high that the series are stationary. However, there is still a small probability of the presence of an unit root. Therefore, I apply the Augmented Dicky-Fuller (ADF) test to see if θ is indeed smaller than 1.

$$y_t = \alpha + \theta y_{t-1} + \delta_1 \Delta y_{t-1} + \ldots + \delta_k \Delta y_{t-k} + \epsilon_t \tag{1}$$

Here k is the number of lags in the ADF regression. If the null hypothesis of the ADF test is rejected, the data is already stationary and it is not necessary to take the second differences. Luckily, this is the case for all the variables.

For the first part of the research, the modeling of the volatility, only the BFCIUS is needed, but for the second part, all variables are initially used to construct a FCI. As mentioned before, the variables are now the first differences of the original variables, in other words the change from day $t-1$ to t. All variables are forecasted using econometric methods. If a one-period-ahead forecast is made today, we have an estimated value for tomorrow. And as we go further in the future, the accuracy of the forecast also gets worse, but this information is very valuable. If financial conditions can be predicted a long time in advance, this gives the opportunity to anticipate on it. Therefore, the idea arises to change the daily observations into weekly observations. For each variable I now have the change from one week to next week. Now, if a one-period-ahead forecast is made today, this is the estimate for the next week. This way the horizon has actually become 5 times as large.

2.3 Standardizing

The set of variables are all of different magnitudes, so before modeling them I have to make one more adjustment: standardizing. There are several standardizing methods, but here I choose to use only method: standardizing recursively. By standardizing recursively I only use the information available until period t and not the entire data set. This way, only historical data is of importance to the observation on period t and not the future observations. This means that for the standardization of variable x on period t, x_t is subtracted by the mean of x in time period $[1, t]$ and divided by the standard deviation of x in time period $[1, t]$.

3 Methodology

In this section I present the different methods I used in the research. As mentioned before, this research consists of two parts: the modeling of the volatility of the BFCIUS and the constructing of a FCI by forecasting the individual variables. Consequently, this section is also divided into two parts. The methods for modeling the volatility will be discussed first.

3.1 Modeling the volatility of BFCIUS

In this section, the differences of BFCIUS is the only variable used, so I define y_t as the value of the first difference of BFCIUS in period t. Before capturing the volatility of the variable, one should first model the variable itself (as far as possible). This can be done by regressing the variable on its lags. To see if this is appropriate I have to look at the correlogram of this variable. The correlogram shows whether or not there is autocorrelation in the series and until which order. If there is significant autocorrelation, I model y_t using autoregressive terms.

3.1.1 Test for conditional heteroskedasticity

To be able to model the volatility, there should be presence of conditional heteroskedasticity. If this is the case, I can use the Autoregressive Conditional Heteroskedasticity model (ARCH), introduced by Engle [1983]. The possibility of conditional heteroskedasticity can be tested using the ARCH Lagrange Multiplier(ARCH-LM) test. Actually, this test is also a general test for possible non-linearities in time series. The null hypothesis is that the model is correctly specified, so a rejection of the this null hypothesis would mean a misspecification, which often continues in including ARCH terms, hence the name ARCH-LM test. The test is done as follows:

- Estimate the model for y_t and let e_t be the series of residuals.
- Estimate the ARCH model by regressing the squared residuals e_t^2 on a constant and on p squared lagged residuals $e_{t-1}^2, \ldots, e_{t-p}^2$.
- Now calculate $LM = nR^2$, where R^2 is from the regression of the second step. The LM-test is asymptotically χ_p^2 distributed. The null hypothesis is rejected for large values of LM.

If the null hypothesis of the ARCH-LM test cannot be rejected, I assume the series are already modeled correctly. But if the null hypothesis is rejected, I can advance into building an ARCH model.

3.1.2 ARCH and GARCH

If the variance of a time series depends on the past, it is appropriate to use ARCH models. This can be combined with the AR model for y_t. The general case of this combination model is discussed in Baillie and Bollerslev [1992].

$$y_t = \mu + \theta_1 y_{t-1} + \ldots + \theta_r y_{t-r} + \epsilon_t \tag{2}$$

with $\epsilon_t = z_t \sigma_t$ and z_t is assumed to have a normal distribution,

$$\sigma_t^2 = \omega + \alpha_1 \epsilon_{t-1}^2 + \ldots + \alpha_p \epsilon_{t-p}^2 \tag{3}$$

Here r is the order of the AR model, and p is the order of the ARCH model. As variances are non-negative, the conditions $\omega > 0$ and $\alpha_1, \ldots, \alpha_p \geq 0$, must hold. In case $\alpha_1 = \ldots = \alpha_p = 0$, the conditional variance is a constant, and the series ϵ_t is conditionally homoskedastic.

In most cases, p needs to be taken quite large to capture the dynamic patterns in the conditional volatility; and more parameters also mean more restrictions. Bollerslev [1986] introduced the Generalized ARCH model (GARCH), which solves this problem. If I transform it into a GARCH(1,1) model, Equation (3) now looks like this:

$$\sigma_t^2 = \omega + \alpha_1 \epsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \tag{4}$$

So now the conditional variance is explained by the preceding error term and conditional variance in the previous period. Here again, all parameters must be non-negative and α_1 must be strictly positive for β_1 to be identified. To see why this GARCH model does not need many squared residual terms, I can rewrite Equation (4) as follows:

$$\sigma_t^2 = \omega + \alpha_1 \epsilon_{t-1}^2 + \beta_1(\omega + \alpha_1 \epsilon_{t-2}^2 + \beta_1 \sigma_{t-2}^2) \tag{5}$$

By continuously replacing the σ^2 on the right hand side, this equation can be expressed like this:

$$\sigma_t^2 = \sum_{i=0}^{\infty} \beta_1^i \omega + \alpha_1 \sum_{i=1}^{\infty} \beta_1^{i-1} \epsilon_{t-i}^2 \tag{6}$$

It is clear now that the GARCH(1,1) model corresponds to an ARCH(∞) model. The parameters of the GARCH model and the AR model for y_t are estimated simultaneously using Maximum Likelihood.

3.1.3 Tests for misspecification

Now the volatility is modeled in a way that previous shocks have the same effect on the conditional volatility at time t no matter what the sign of these shocks is. However, one could argue that positive and negative shocks may have different impacts on the volatility in the following periods. In time series data it is often found that periods of great volatility are initiated by large negative shocks. This suggests that one should allow for a difference in effects of positive and negative shocks in the GARCH model.

Before actually building such an asymmetric GARCH model, one should test whether this is really necessary. Therefore, I check the probability of a misspecification in the GARCH(1,1) model by looking at the standardized residuals:

$$\hat{z}_t = \frac{\hat{\epsilon}_t}{\hat{\sigma}_t} \tag{7}$$

If the model is correctly specified, the following properties must hold:

- The standardized residuals have mean 0 and variance 1.
- No autocorrelation in the standardized residuals
- No autocorrelation in the squared standardized residuals
- Kurtosis is equal to 3

If these properties do not hold, there could be a misspecification and one should consider using a different distribution for z_t. For example, the Student t distribution could be better, as it allows for excess Kurtosis in z_t, which often occurs in financial data. The degrees of freedom of the Student t distribution can be estimated as an extra parameter together with the other parameters with Maximum Likelihood.

Another test one can do is the Sign Bias (SB) test, introduced by Engle and Ng [1993]. This test shows whether positive and negative shocks affect future volatility differently. The null hypothesis of the SB test is that there is no difference in effect. The test can be performed as follows:

- Create a dummy variable S_t^-, which takes the value 1 if ϵ_{t-1} is negative and 0 otherwise.
- Regress the squared standardized residuals on a constant and S_t^-.
- The SB test statistic is the t-statistic of S_t^-. The null hypothesis is rejected for large values of the t-statistic.

A rejection of the null hypothesis indicates a difference in effects of the positive and negative shocks, which suggests an asymmetric GARCH model would fit the volatility better.

3.1.4 Threshold GARCH model

The Threshold GARCH model (TGARCH), introduced by Glosten et al. [1993], is such an asymmetric model that allows for a different effect between positive and negative innovations. The threshold in our case is 0, and we add a variable to Equation (4) to capture the effects of observations that exceed this threshold.

$$\sigma_t^2 = \omega + \alpha_1 \epsilon_{t-1}^2 + \gamma_1 \epsilon_{t-1}^2 I[\epsilon_{t-1}^2 < 0] + \beta_1 \sigma_{t-1}^2 \tag{8}$$

Here $I[A]$ is an indicator function that takes the value 1 if A happens and zero otherwise. Also the following conditions must hold: $\omega > 0$, $\alpha > 0$, $\gamma > 0$ and $\beta \geq 0$. To be covariance stationary there is even an extra condition: $\alpha + \gamma/2 + \beta < 1$. If this condition does not hold the variance has no unconditional mean. This becomes clear if I take the expectation of Equation (8):

$$\mathbb{E}(\sigma_t^2) = \omega + \alpha_1 \mathbb{E}[\epsilon_{t-1}^2] + \gamma_1 \mathbb{E}[\epsilon_{t-1}^2 I[\epsilon_{t-1}^2 < 0]] + \beta_1 \mathbb{E}[\sigma_{t-1}^2] \tag{9}$$

$$\bar{\sigma}^2 = \omega + \alpha_1 \bar{\sigma}^2 + \gamma_1 \bar{\sigma}^2 \frac{1}{2} + \beta_1 \bar{\sigma}^2 \tag{10}$$

$$\bar{\sigma}^2 = \frac{\omega}{1 - \alpha_1 - \gamma_1/2 - \beta_1} \tag{11}$$

I come to this expression by assuming that the median of z_t is 0, so $\mathbb{E}[I[\epsilon_{t-1}^2 < 0]] = P(\epsilon_{t-1} < 0) = 0.5$. If the covariance stationarity condition does not hold, it is clear that the unconditional mean in Equation (11) cannot be calculated.

3.1.5 Forecasting volatility

Using the models in the previous section I can now forecast the volatility of the BFCIUS. I assume that the final model, TGARCH(1,1) model, fits the data the best, then the 1-step-ahead forecast for the conditional variance looks like this:

$$\hat{\sigma}_{t+1|t}^2 = \mathbb{E}[\omega + \alpha_1 \epsilon_t^2 + \gamma_1 \epsilon_t^2 I[\epsilon_t^2 < 0] + \beta_1 \sigma_t^2] \tag{12}$$

$$\hat{\sigma}_{t+1|t}^2 = \hat{\omega} + (\hat{\alpha_1} + \hat{\gamma_1}/2 + \hat{\beta_1})\sigma_t^2 \tag{13}$$

Here $\hat{\omega}, \hat{\alpha_1}, \hat{\gamma_1}, \hat{\beta_1}$ are the estimated parameters from the TGARCH(1,1) model.

Using the expression in Equation (13) I can recursively compute the h-step-ahead forecast for $h \geq 2$.

$$\hat{\sigma}_{t+h|t}^2 = \hat{\omega} + (\hat{\alpha_1} + \hat{\gamma_1}/2 + \hat{\beta_1})\hat{\sigma}_{t+h-1|t}^2 \tag{14}$$

I can also make a direct expression of the h-step-ahead forecast:

$$\hat{\sigma}_{t+h|t}^2 = \hat{\omega} \sum_{i=0}^{h-2} (\hat{\alpha_1} + \hat{\gamma_1}/2 + \hat{\beta_1})^i + (\hat{\alpha_1} + \hat{\gamma_1}/2 + \hat{\beta_1})^{h-1} \hat{\sigma}_{t+1}^2 \tag{15}$$

where $\hat{\sigma}_{t+1}^2$ can be obtained using Equation (13). Once these forecasts are made, the logical next step would be to evaluate these forecasts. The common way to do this is by calculating the Mean Squared Prediction Error (MSPE).

$$\text{MSPE} = \frac{1}{P} \sum_{j=0}^{P-1} (\hat{\sigma}_{t+h+j|t+j}^2 - \sigma_{t+h+j}^2)^2, \tag{16}$$

Here P is the number of forecasts made.

Unfortunately, there is a problem in this approach. To evaluate the volatility forecast, I need the true volatility σ^2_{t+h+j}, which is unobserved. To be able to calculate the MSPE, I need an estimate of $\sigma^2_{t+h+j|t}$. Most of the times, the squared shock ϵ^2_{t+h+j} is used as an estimate, as $\mathbb{E}[\epsilon^2_{t+h+j|t+h+j-1}] = \sigma^2_{t+h+j}$, which makes it an unbiased estimate. On the other hand, if I express this estimate as $\epsilon^2_{t+h+j} = z^2_{t+h+j}\sigma^2_{t+h+j}$, it is clear that the estimate is very noisy, because of the square of z_{t+h+j}. As a solution for this problem, it is suggested by Franses and van Dijk [To be published] to use data which is sampled more frequently to get a more accurate measure of volatility. This is usually referred to as realized volatility. Unfortunately, intraday data of the BFCIUS is not available over such a great time period, thus the realized volatility cannot be used. So despite the fact that the estimate ϵ^2_{t+h+j} is noisy, I will still use it to compare the forecasts with.

3.2 Construction of a Financial Conditions Index

In this section the methods to construct an own Financial Conditions Index (FCI) is presented. As mentioned earlier, the steps taken to eventually produce a FCI are in a different order than in Lee et al. [2013]. Here I first start with analyzing and modeling all individual variables.

3.2.1 Autoregressive models

Let m be the number of variables in and T the number of observations the dataset, so that I have m different series of $y_{t,j}$, with $j = 1, \dots, m$ and $t = 1, \dots, T$. For each series $y_{t,j}$ I check its correlogram to see its structure of autocorrelation. If there is indeed autocorrelation present in the series, I model them using an AR model.

$$y_{t,j} = \alpha + \beta_1 y_{t-1,j} + \dots + \beta_p y_{t-p,j} + \epsilon_t \qquad \text{for } j = 1, \dots, m \tag{17}$$

Here p is the order of the AR(p) model, which I determine for each series $y_{t,j}$ separately by looking at the correlograms. The values of the parameters $\alpha, \beta_1, \dots, \beta_p$ can be estimated using OLS.

3.2.2 Heterogeneous Autoregressive models

As shown in Lee et al. [2013] the Heterogeneous Autoregressive model (HAR), introduced by Corsi [2009], seems to be a good method to model the FSI with. So lets see if this HAR model fits the individual variables as well. The HAR model is actually a simple model that can create a long memory by using only a few explanatory variables. I want the model to have the information of the past month captured, but I do not want to include 20 autoregressive terms in the model. What the HAR model does is it creates a variable that represents the information of the past month by taking the mean of these 20 lagged values.

$$y^{(m)}_{t,j} = \frac{1}{20}(y_{t-1,j} + \dots + y_{t-20,j}) \qquad \text{for } j = 1, \dots, m \tag{18}$$

In the same way I also create a variable that captures the information of the past week.

$$y^{(w)}_{t,j} = \frac{1}{5}(y_{t-1,j} + \dots + y_{t-5,j}) \qquad \text{for } j = 1, \dots, m \tag{19}$$

Using these two variables and a simple AR(1) term I now have HAR model with the lagged daily, weekly and monthly values of each series as explanatory variables.

$$y_{t,j} = \alpha + \beta_1 y^{(d)}_{t,j} + \beta_2 y^{(w)}_{t,j} + \beta_3 y^{(m)}_{t,j} + \epsilon_t \qquad \text{for } j = 1, \dots, m \tag{20}$$

The values of the parameters α, β_1, β_2 and β_3 can be estimated using OLS.

3.2.3 Forecasting the individual series

Using the two models I have just shown, I can make out-of-sample forecasts for the individual series. As none of the AR models have an order higher than 2, there is a high probability that the forecasts with a large horizon are inaccurate. Therefore, I limit the forecasts of the AR models to the horizon of 1.

$$\hat{y}_{t+1|t,j} = \hat{\alpha} + \hat{\beta}_1 y_{t,j} + \ldots + \hat{\beta}_p y_{t-p+1,j} \qquad \text{for } j = 1, \ldots, m \tag{21}$$

In the Data section I already mentioned the problem of the inaccurate forecasts with a large horizon. That is why I also transformed the daily data into weekly data. So with the other data set of weekly observations, I can follow the same steps as with the daily steps. The difference is that the one-period-ahead forecast now actually is a horizon of 1 week. Unlike the AR model, the HAR model does have a long memory, so in theory this model should be able to produce more accurate forecasts for large horizons. The HAR forecasts with horizon h can be computed in the following way:

$$\hat{y}_{t+h|t,j} = \hat{\alpha} + \hat{\beta}_1 y_{t,j}^{(d)} + \hat{\beta}_2 y_{t,j}^{(w)} + \hat{\beta}_3 y_{t,j}^{(m)} \qquad \text{for } j = 1, \ldots, m \tag{22}$$

The forecasts of the two models can be evaluated by calculating the MSPE. Here P is the number of forecasts.

$$\text{MSPE} = \frac{1}{P} \sum_{j=0}^{P-1} (\hat{y}_{t+h+j|t+j} - y_{t+h+j})^2 \tag{23}$$

3.2.4 Constructing the out-of-sample index

Having the out-of-sample forecasts of the variables, I finally arrive to the part where I construct the FCI. The first step in this is selecting the right variables to build the index. I select by looking at the correlation of the variables with the BFCIUS. The BFCIUS is an index that is considered to describe the financial conditions accurately, so the variables that form my own FCI should have enough correlation with the BFCIUS. The variables with an absolute correlation value higher than 0.50 with the BFCIUS will be selected, and the negatively correlated variables are multiplied by -1 to capture the right effect.

With this selection of variables I use two different techniques to construct a FCI: taking the unweighted sum of the variables, and taking the weighted sum of the first k principal components. The first technique is quite straight forward. Say, I have a selection of s variables which are correlated with the BFCIUS. Then the unweighted sum FCI (USFCI) is constructed as follows:

$$USFCI_t = \sum_{j=1}^{s} y_{t,j} \tag{24}$$

The second technique requires some more work. Principal Component Analysis (PCA) evaluates a number of eigenvectors, representing independent linear combinations, equal to the s number of variables in the selection, and also a same number of eigenvalues evaluated from the covariance matrix of the data. The eigenvectors are evaluated as follows.

$$\Sigma z_j = \lambda_j z_j \tag{25}$$

Here Σ represents the $s \times s$ covariance matrix of the dataset, z_j the jth $s \times 1$ eigenvector and λ_j the jth 1×1 eigenvalue. If I multiply these eigenvectors with the selected variables, this results in the s principal components.

$$X_{tj} Z = PC_{tj} \tag{26}$$

Here X_{tj} is the $T \times s$ matrix of the variables and Z is the $s \times s$ matrix of all the eigenvectors.

The next step is to determine k, the amount of principal components I want to include. There are a several criteria to choose this number k. One could look at the scree plot, find the 'elbow' point and include the eigenvalues up until that point. However, the 'elbow' point can be hard to determine, as it can be seen in different locations. One could also include all principal components which eigenvalues exceed the value 1, because 1 is the average eigenvalue, and if it exceeds this value its principal component should be important enough. The problem with this criterion is that it is difficult to say whether the eigenvalue is significantly higher than 1 or not. Eventually I choose the criterion that the total proportion of the variance that the principal components explain should be at least 60%. I keep on adding principal components until this level of explained variance is reached. Now I know how to create the principal components and how to determine the number k, I can finally create Principal Components FCI (PCFCI):

$$PCFCI_t = \frac{1}{s} \sum_{j=1}^{k} \lambda_j PC_{tj} \tag{27}$$

The eigenvalues tell how much explanatory power each principal component has, so in this way the principal components are weighted according to their importance in relation to the data.

Note that in period t, I do not know the observations of the variables in $t+1, t+2 \ldots$, so this is where the forecasted series plays a role. I eventually want to know the values of the PCFCI in the out-sample. In each period t in the out-sample, I have all the observed values of the variables until period t, and then I add the forecasted values of a chosen forecast horizon h. With these observations and forecasts I apply the steps of PCA as mentioned earlier. In other words, PCA is done recursively for each period t, so in each period t I forecast the $\widehat{PCFCI}_{t+h}$.

4 Results

4.1 Modeling the volatility of BFCIUS

In this section I present the results of modeling the volatility of BFCIUS. As mentioned in the Methodology section, the first step is to check for the presence of autocorrelation in y_t, which is the first difference of BFCIUS. The correlogram, shown in the Appendix, displays no significant autocorrelation in the series y_t. This means that it cannot be modeled using autoregressive terms, and as I do not include other explanatory variables, y_t only consists of its mean μ and the innovations ϵ_t.

Next I have to check the disturbance terms ϵ_t for conditional heteroskedasticity by doing the ARCH-LM test described in Section 3.1.1. By regressing the squared residuals on a constant and its lag, the test yields the following result: $LM = 127.88$. This LM test is asymptotically $\chi^2_{(1)}$ distributed, so with a confidence level of 95%, a value for the LM test of 3.84 or higher would reject the null hypothesis. As $127.88 > 3.84$ the null hypothesis is rejected, and I assume that the disturbance terms ϵ_t are conditional heteroskedastic.

Knowing that ϵ_t is conditional heteroskedastic, I can build conditional heteroskedasticity models. As shown in Section 3.1.2, the GARCH model is an extended version of the ARCH model, so I will estimate the GARCH(1,1) model. The estimated coefficients are shown in Equation (28), the standard errors are between parentheses.

$$\sigma_t^2 = 0.000(0.000) + 0.201(0.009)\epsilon_{t-1}^2 + 0.807(0.008)\sigma_{t-1}^2 \tag{28}$$

To see if the volatility is modeled correctly I do a couple of tests. These tests require the standardized residuals $\hat{z}_t$, so I compute these according to formula (7). In Figure 2 the histogram of $\hat{z}_t$ are shown. The red line represents the normal distribution. It is clear that the standardized

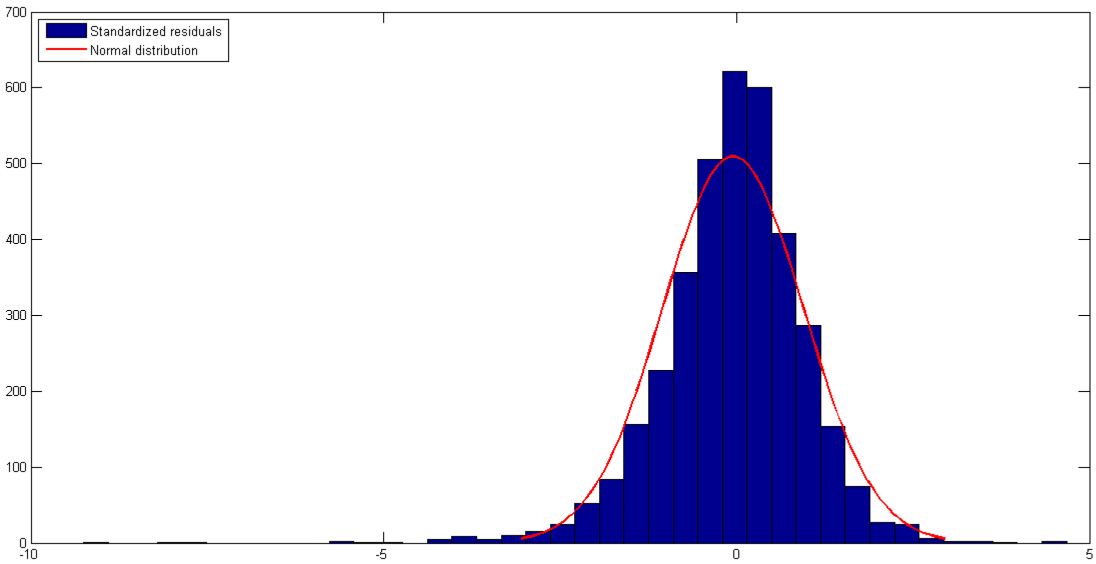

Figure 2: Standardized residuals of GARCH(1,1) fitted with a normal distribution. The standardized residuals are computed as follows: $\hat{z}_t = \frac{\hat{\epsilon}_t}{\hat{\sigma}_t}$

residuals are not normally distributed. First I have to check if the conditions listed in Section 3.1.3 hold. The mean and variance of $\hat{z}_t$ are respectively -0.049 and 0.998, so the mean differs a little from 0, but the variance is almost exactly equal to 1. If I look at the correlograms of $\hat{z}_t$ and $\hat{z}_t^2$, there is no significant autocorrelation. The kurtosis is 8.611, which can also be observed from the histogram, is much higher than 3, so this condition does not hold. This suggests that perhaps a different distribution for z_t is more appropriate.

The next test I do is the SB test, to see if the the sign of a shock affects the volatility. The regression of $\hat{z}_t^2$ on a constant and the dummy S_t^- gives a t-statistic of 1.969 for S_t^-. With a p-value of 0.049, the null hypothesis of no difference in effect is rejected on a confidence level of 95%.

The result of the SB test indicates that the volatility model should include a variable to capture the difference in effect. Therefore, the switch from GARCH to TGARCH is made. Also, the properties of $\hat{z}_t$ suggest that a Student t distribution might fit better, as this distribution allows for excess kurtosis. Thus, I estimate a TGARCH(1,1) model with a Student t distribution for z_t. The estimated coefficients are shown here, the standard errors are between parentheses.

$$\sigma_t^2 = 0.000(0.000) + 0.044(0.014)\epsilon_{t-1}^2 + 0.187(0.022)\epsilon_{t-1}^2 I[\epsilon_{t-1}^2 < 0] + 0.862(0.011)\sigma_{t-1}^2 \qquad (29)$$

The total effect of a negative shock is a calculated by adding up the two coefficients: $0.044 + 0.187 = 0.231$. This means that the effect of a negative innovation is about 5 times as large as the effect of a positive innovation on the conditional volatility.

If I check these coefficients for the final condition $\alpha + \gamma/2 + \beta < 1$, I see that the coefficients add up to 0.999. Looking at the standard errors, this number is probably too close to 1, so I perform a Wald test to check this hypothesis. With a p-value of 0.845 the sum of the coefficients is not significantly smaller than 1, which means that it is not covariance stationary. Although it does not hold the condition of covariance stationarity, it is still strict stationary. And as Nelson [1990] has shown, if it is strict stationary, it is still a well defined model.

With these estimates of the TGARCH(1,1) model, I forecast the conditional volatility. The forecast sample is defined as the period 1 January 2010 until 16 January 2013. Filling in the estimated coefficients from the previous section into Equation (15) gives the h-step-ahead forecasts.

Figure 3 shows the 1-step-ahead volatility forecast plotted with squared residuals. Despite the fact that the squared shocks are a noisy estimate for the real volatility, it matches the forecasts quite well. In periods of high volatility, around observation 100 and between 400 and 500, the forecasts seem to follow the 'real' volatility upwards and in periods of low volatility it stays down.

I also made 5-steps and 20-steps ahead forecasts, which are also quite accurate. These plots can be found in the Appendix. The MSPEs for respectively the 1-step, 5-steps and 20-steps ahead forcasts are 0.0025, 0.0031 and 0.0035.

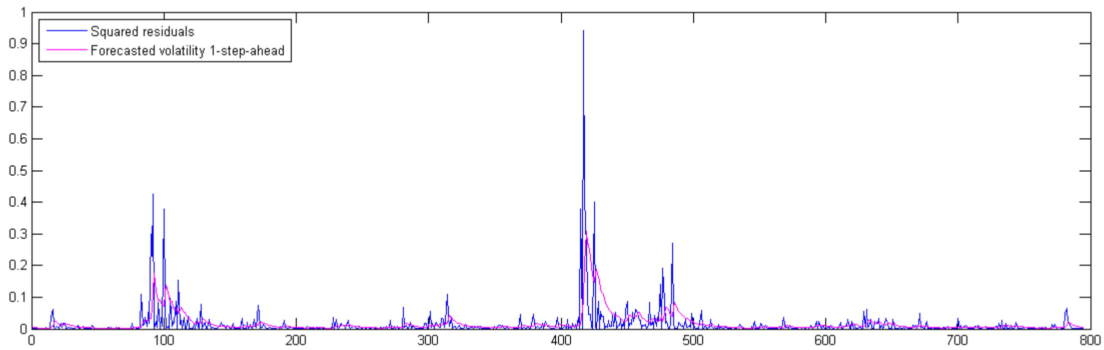

Figure 3: Forecasted volatility and the squared residuals. The forecasts are computed using this formula: $\hat{\sigma}_{t+h|t}^2 = \hat{\omega}\sum_{i=0}^{h-2}(\hat{\alpha_1} + \hat{\gamma_1}/2 + \hat{\beta_1})^i + (\hat{\alpha_1} + \hat{\gamma_1}/2 + \hat{\beta_1})^{h-1}\hat{\sigma}_{t+1}^2$

4.2 Construction of a Financial Conditions Index

In this section I present the results of the second part of the research: the construction of a FCI. As explained earlier, in this part of the research I initially work with a total of 20 variables which are all modeled and forecasted individually. However, not every variable plays a significant role in constructing the FCI. Therefore, I choose to only present the results of the variables which eventually were significant in the construction of the FCI.

4.2.1 Modeling the variables

As explained in the Methodology section, the series are modeled in two different ways: AR and HAR. To make AR models I first have to check the autocorrelation patterns in the variables. It turns out that in some series there was no significant autocorrelation present, so these were not possible to model in this way. However, these series also turn out to have very little correlation with the BFCIUS, so they are not used in the construction of the index. Table 2 shows the autocorrelation structure of the variables.

Variable	AR terms	Correlation with BFCIUS
VIX	2	-0.940
XLF	1	-0.809
SPY	1	0.867
Yield curve	1	0.580
HY-10Y	1	-0.624
Financial σ	1	-0.697

Table 2: Variables which have the highest correlation with BFCIUS in the forecast sample period 1 January 2010 - 16 January 2013. The corresponding AR terms are determined by looking at the autocorrelation pattern in the correlograms of the variables.

The variables listed in this table are modeled using AR and HAR. The variables that are listed in Table 1 in the Data section but not in this table, are not used in the rest of the research.

4.2.2 Forecasting the individual series

With the estimated AR and HAR models I now make out-of-sample forecasts of the 6 variables. The MSPEs of the 1-step-ahead forecasts, are presented in Table 3. Unfortunately, the MSPEs vary a lot between the variables, which could mean that some variables are not modeled correctly. Especially the variables VIX and SPY are forecasted very badly, as their MSPEs are above 1 for both the AR and the HAR models. This could negatively affect the accuracy of the FCI, because these two variables are very highly correlated with the BFCIUS in the out-sample. In general the HAR forecasts seem to be slightly better than the AR forecasts, except for the variables Yield curve and HY-10Y. It seems that these two variables can be forecasted better by using only the information from the last observation.

Variable	MSPE - AR	MSPE - HAR
VIX	1.545	1.256
XLF	0.4075	0.3403
SPY	1.185	1.063
Yield curve	0.5122	0.684
HY-10Y	0.4351	0.8936
Financial σ	0.2409	0.1953

Table 3: MSPEs of 1-step-ahead forecasts, which are calculated using MSPE $= \frac{1}{P}\sum_{j=0}^{P-1}(\hat{y}_{t+h+j|t+j} - y_{t+h+j})^2$, the higher the number, the more inaccurate the forecast.

As mentioned in the Data section, I also created a weekly data set to be able to make larger forecast horizons. Sadly, the patterns of the weekly differences are even more difficult to model and to forecast than the daily differences. With MSPEs of the important variables VIX and SPY that are larger than 4, the forecasts are considered very inaccurate and therefore will not be used.

4.2.3 Constructing the out-of-sample index

With the forecasts made in the previous section, I now arrive to the final part of the research, the construction of the index. Unfortunately, the forecasts of the variables are not so good, so the expectation is that the index is not of great quality. Figure 4 shows the USFCIs constructed using the AR and HAR forecasts, plotted together with the BFCIUS. Note that the figure does not show the entire forecast period, but it is zoomed into a smaller period to show the differences between the indexes more clearly. The figure with complete forecast sample can be found in the Appendix. Surprisingly, the USFCI constructed with the AR forecasts seems to be better than the other one, as the green line matches the blue line more than the red line does. Besides, if I check the correlation between the two USFCIs and the BFCIUS, the USFCI-AR wins with 0.760 over 0.092. The large difference is rather unexpected, because the MSPEs of the forecasted variables do not differ a lot between the two models and are even in slight advantage of the HAR model. An explanation could be that the HAR model overfits all the variables, and taking the sum of these variables results in a very inaccurate index.

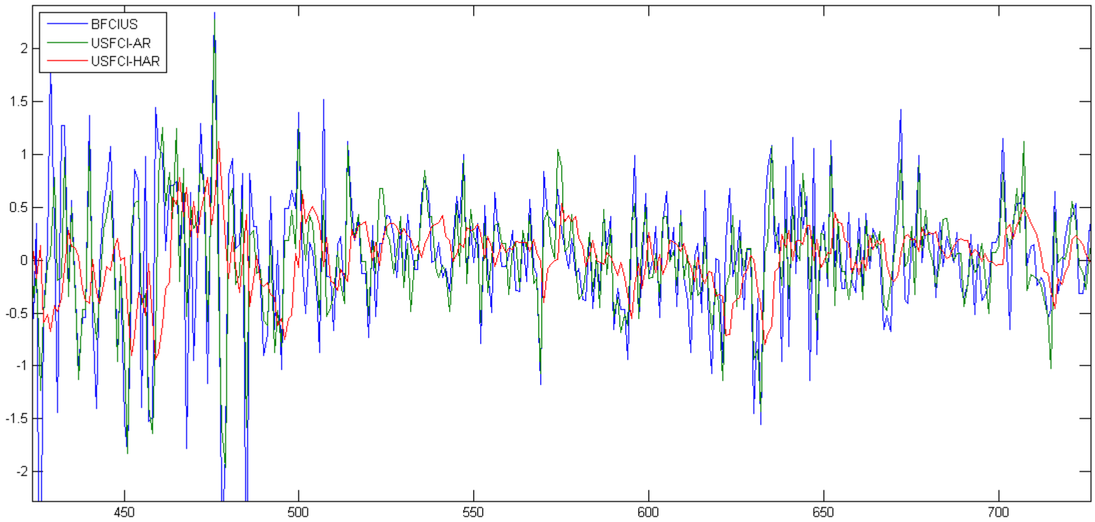

Figure 4: USFCIs constructed using AR and HAR forecasts

Where one of the USFCIs can match the BFCIUS a little bit, the PCFCIs turn out to be no success at all. Somehow this method of constructing the index does not seem to work at all, as Figure 9 in the Appendix shows. Unfortunately, I cannot really figure out why it does not follow the pattern of the BFCIUS at all. Thus the results of the PCFCIs are of no use.

5 Conclusion

In this research I have dealt with different aspects of financial stress. On one hand I have tried to model the volatility of financial stress, and on the other hand I tried to convert the financial conditions into numerical values by creating a financial conditions index. These two tasks were not really related to each other and it has turned out that they also have very different results.

The purpose of volatility modeling is to capture the clusters of volatility. In financial data periods of high and low volatility come in clusters, and therefore the assumption of homoskedasticity is often incorrect. The conditional heteroskedasticty test has shown that this is indeed the case for the BFCIUS. The outcome of several tests has led to the use of ARCH, GARCH and eventually TGARCH models to capture the patterns of volatility. This TGARCH turned out to be the most suitable as I discovered that positive and negative shocks have different effects on the volatility, and this models allows for that. Therefore, this model is used to make out-of-sample forecasts of the volatility. Luckily, the forecasts errors were quite small and so the use of this model seems to be appropriate.

The second part of this research turned out to be rather disappointing. The findings did not help to reach the goal I had in mind and where actually that poor that I suggest no further research to be done on this matter. The idea was to model and forecast the individual series, because Lee et al. [2013] has shown that modeling the volatile FSIs was somewhat difficult to do. But this research has shown that modeling the variables is even a greater problem. The patterns of the important variables such as the VIX, could not be captured by only using autoregressive terms. Probably more complex models are needed to describe the changes of these variables, and therefore my simple AR and HAR models gave poor forecasts, which consequently lead to the creation of an inaccurate index.

A Additional graphs

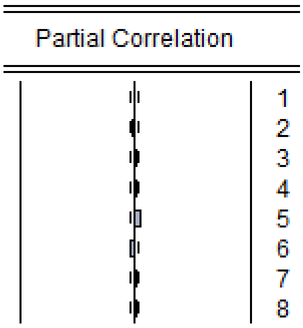

Figure 5: Correlogram of the first difference of BFCIUS

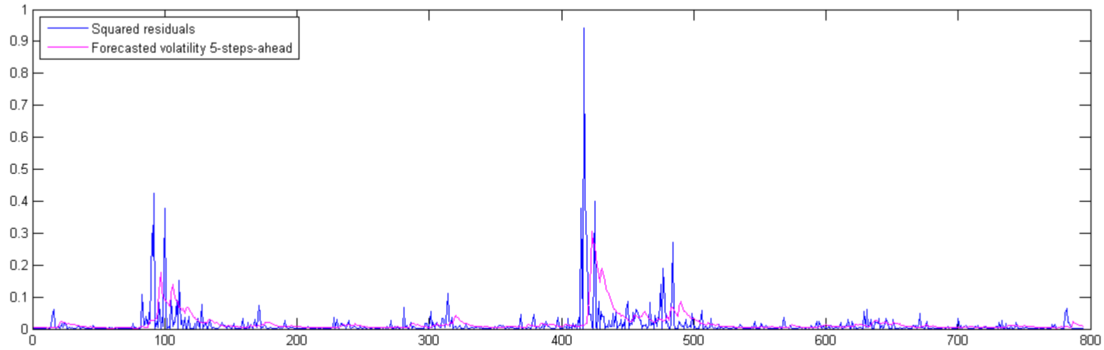

Figure 6: Forecasted volatility and the squared residuals. The forecasts are computed using this formula: $\hat{\sigma}^2_{t+h|t} = \hat{\omega}\sum_{i=0}^{h-2}(\hat{\alpha_1} + \hat{\gamma_1}/2 + \hat{\beta_1})^i + (\hat{\alpha_1} + \hat{\gamma_1}/2 + \hat{\beta_1})^{h-1}\hat{\sigma}^2_{t+1}$, here h is 5.

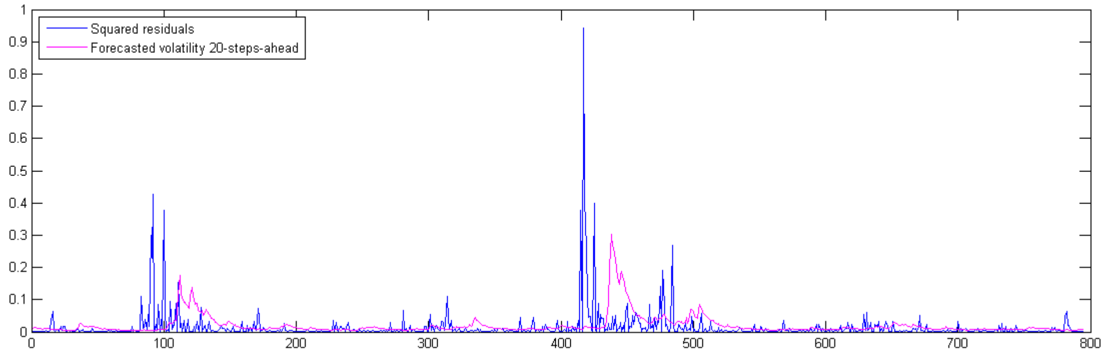

Figure 7: Forecasted volatility and the squared residuals. The forecasts are computed using this formula: $\hat{\sigma}^2_{t+h|t} = \hat{\omega}\sum_{i=0}^{h-2}(\hat{\alpha_1} + \hat{\gamma_1}/2 + \hat{\beta}_1)^i + (\hat{\alpha_1} + \hat{\gamma_1}/2 + \hat{\beta}_1)^{h-1}\hat{\sigma}^2_{t+1}$, here h is 20.

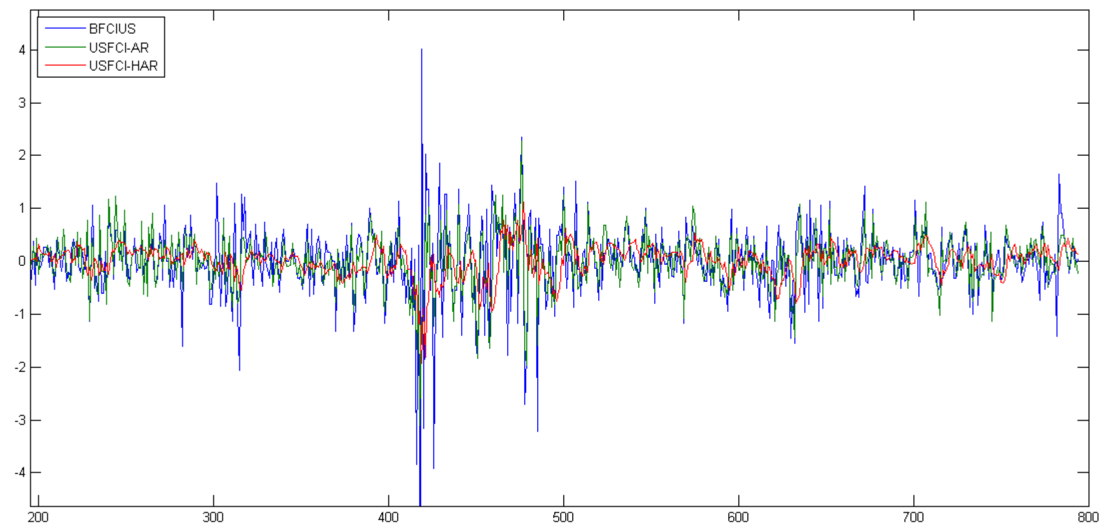

Figure 8: USFCIs constructed using AR and HAR forecasts. USFCI is constructed as follows: $USFCI_t = \sum_{j=1}^{s} y_{t,j}$.

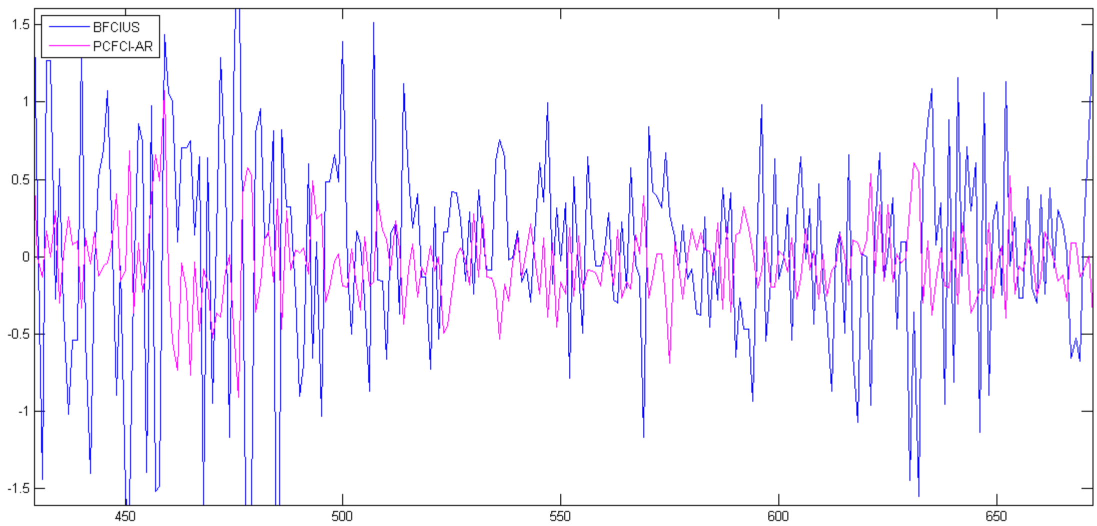

Figure 9: PCFCI constructed using AR forecasts. PCFCI is constructed as follows: $PCFCI_t = \frac{1}{s}\sum_{j=1}^{k}\lambda_j PC_{tj}$.

References

Richard T Baillie and Tim Bollerslev. Prediction in dynamic models with time-dependent conditional variances. *Journal of Econometrics*, 52(1):91–113, 1992.

Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics*, 31(3):307–327, 1986.

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

Robert F Engle. Estimates of the variance of us inflation based upon the arch model. *Journal of Money, Credit and Banking*, 15(3):286–301, 1983.

Robert F Engle and Victor K Ng. Measuring and testing the impact of news on volatility. *The journal of finance*, 48(5):1749–1778, 1993.

Philip Hans Franses and Dick van Dijk. *Time Series Models for Business and Economic Forecasting*. Cambridge University Press, To be published.

Lawrence R Glosten, Ravi Jagannathan, and David E Runkle. On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance*, 48(5):1779–1801, 1993.

Yan Kit Lee, Josephine Boerman, Philo Meerman, and Rik de Wilde. Modeling the financial stress index. *Seminar Finance in Econometrics, Erasmus School of Economics*, 2013.

Daniel B Nelson. Stationarity and persistence in the garch (1, 1) model. *Econometric theory*, 6 (03):318–334, 1990.