Bachelor Thesis Double Degree Econometrics & Economics

A Regime-Switching GARCH-MIDAS Approach to Modelling Stock Market Volatility

Author
K.W. Naumann (432211)

Supervisor
B. van Os
Second Assessor
dr. M.D. Zaharieva

Final Version – 6^{th} July 2019

Abstract

This paper investigates macroeconomic effects on the volatility of the S&P500 stock market index in the GARCH-MIDAS setup of Engle et al. (2013). This setup is extended to a two-regime Markov-switching model. The results indicate a significant effect of macroeconomic variables on volatility. There are four variables leading volatility: real consumption, consumer sentiment, housing starts and the term spread. Implementation of the Markov-switching extensions shows that a second regime primarily corresponds to "outliers" in volatility. Descriptively it is on par with the single-regime GARCH-MIDAS in the 2000-2010 period, but in terms of forecasting the regime-switching model is inferior, with decreasing accuracy as the forecast horizon increases. This suggests that regime information is sufficiently contained in the macroeconomic variables, such that explicitly accounting for regimes becomes unnecessary when considering a diversified stock index for a developed country.

Contents

1 Introduction

The vast majority of individuals are risk-averse and would thus most likely prefer to see a steady increase in their wealth as opposed to large upward or downward swings. Spikes in volatility, the degree of variability of returns to an investment (commonly measured as the variance of returns), suggest stable returns are unlikely and large swings in wealth are more probable. Consequently, volatility is a term often associated with anguish by individuals who do not want to face the risk of large losses. On the other hand, there are *antifragile* (Taleb, 2012) elements whose potential is unlocked in the face of volatility. One must only think of the derivatives market, where deeply out-of-the-money options can see significant gains in value as volatility increases, due to an increased probability of maturing in-the-money. Similarly, active strategy hedge funds would consider low volatility periods dangerous to their profits. This is corroborated by the rise in passive exchange-traded funds (ETFs), and decreasing assets under management by the hedge fund industry in the post-2008 global financial crisis volatility lull. The implication of these various complicated economic relationships to volatility is that forecasting volatility levels is of crucial importance, not just in the short-run but also in the long-term. Asset managers will be keen on rebalancing portfolios to optimise their risk-return tradeoff in preparation for changes in volatility levels. Failure to do so could result not only in decreased profits but might lead to investors pulling their money from the funds. Consequently, volatility prediction and being able to prepare in advance is a competitive edge for any investor and fund manager alike. Therefore, the implementation of strong predictive models for volatility is crucial in light of the fact that asset managers are predicted to be in charge of more than USD100 trillion in global assets under management by 2020 (PWC, 2017).

The analysis and forecasting of volatility has come a long way since the introduction of the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models by Engle (1982) and Bollerslev (1986). However, the link between stock market volatility and macroeconomic factors had already been postulated by Officer in 1973 and subsequently by Schwert in his article "Why does stock market volatility change over time?" (1989). Only in 2013 did Engle, Ghysels, and Sohn address this parsimoniously by incorporating macroeconomic factors into the GARCH-class models through a MIDAS (MIxed DAta Sampling) term. They showed that the inclusion of macroeconomic factors significantly improves volatility estimates, and improves the accuracy of volatility predictions at long horizons. However, the model developed by Engle et al. is still susceptible to structural breaks in the volatility process, decreasing the accuracy of forecasts if the break goes undetected, an issue crucial to risk-management (see Andreou & Ghysels (2003) or Ang & Timmermann (2012)).

In this paper, I develop the Markov-switching GARCH-MIDAS (MS-GM) model as an extension of Engle, Ghysels and Sohn's approach. This model incorporates different volatility regimes in order to improve structural stability and potentially account for structural breaks. In fact, regime-switching models are one type of stochastic break process in which there are multiple different data generating processes that "mix" at each point in time depending on a Markov chain. The motivation for specifically using a two-regime Markov-switching model is the relation of volatility to the business cycle. Volatility is countercyclical, implying that it is higher during recession periods as opposed to periods of economic growth. In fact, it is hypothesised that the underlying volatility process differs during booms and busts, in the sense that it would be expected that during a downturn there is a GARCH process with higher unconditional variance, while during an upswing the underlying GARCH process has a lower unconditional variance. Indeed, Hamilton & Lin (1996) find volatility to be higher during recessions. The effect of modelling this way is that forecasts are improved, especially in the short-term (e.g. see Hamilton & Susmel (1994) or Klaassen (2002)).

First, the results of Conrad & Loch are replicated to show that the GARCH-MIDAS implementation is correct. Their paper is chosen, as opposed to the original paper of Engle et al. (2013), as they apply a wider selection of macroeconomic variables, which give more flexibility in model choice in case one macroeconomic factor is rejected due to identification issues (explained in section 4.2). In fact, the original paper by Engle et al. (2013) is a subset of this paper, as Conrad and Loch also consider the industrial production and inflation time series, akin to Engle et Al.

This extension is important to the scientific community as a regime-switching model has only been partially implemented by Pan et Al. (2017) in an application to the oil markets. Most importantly, the MS-GM model nests both the simple GARCH process and the single-regime GARCH-MIDAS, consequently it could serve as a basis for further developing a more generalised model.

This paper proceeds as follows: looking at the volatility of the S&P500 index, it is determined what kind of macroeconomic effects on volatility exist. Thereafter the MS-GM model is estimated over the same period and macroeconomic variables. The in-sample fit of the MS-GM model is assessed and contrasted with the single-regime model. Thereafter, the forecasting performance of both variants is investigated in a short-term and long-term scenario. Summa sumarum, an answer to the following research question is sought:

Does a Markov-switching GARCH-MIDAS volatility model improve upon the descriptive and predictive performance of the GARCH-MIDAS of Engle et al. (2013)?

Given the results of the respective models, it is found that, as with Engle et al. (2013) and Conrad & Loch (2015), macroeconomic variables have a significant influence on volatility. In fact, the consumer sentiment, real consumption, housing starts and term spread actually contain leading information for volatility. Implementation of the Markov-switching extensions shows that descriptively it is on par with the fit of the single-regime GARCH-MIDAS. However, in terms of forecasting the regime-switching model is inferior. This suggests that regime information is sufficiently contained in the macroeconomic variables, such that explicitly accounting for regimes becomes unnecessary when considering a diversified stock index for a developed country.

The remainder of the paper is structured as follows: section 2 outlines the state of current research and embeds the MS-GM. Section 3 gives an overview of the data under investigation, while section 4 outlines the methodology for the MS-GM model and how the model will be evaluated. In section 5 the results of the model implementation are presented and analysed. Lastly, section 6 concludes and suggests some further avenues for research.

2 Literature Review

The use of Autoregressive Conditional Heteroskedasticity (ARCH) models to capture the persistence and clustering in volatility dates back to the seminal work of Engle (1982). Since this time, the Generalized ARCH (GARCH) class models developed by Bollerslev (1986) have been a staple in the analysis of volatility, together with stochastic volatility models. Over the course of time, they have been applied and extended numerously, for instance through the exponential-GARCH variation of Nelson (1991), and the asymmetric GARCH of Engle & Ng (1993). However, as Schwert (1989) points out, stock market volatility is linked to macroeconomic activity. In fact, it is empirically suggested that volatility is better described by component models that use a form of long-term volatility. Consequently, several models have been proposed to achieve exactly this goal. Engle & Lee (1999) develop an additive mean-reverting volatility model, with a long-term mean component and short-term transitory component modeled by different GARCH(1,1) processes. Several others have also proposed a variety of two-factor models: Ding & Granger (1996) propose a class of long-memory models, Alizadeh et al. (2002) use price-range based stochastic volatility models, Chernov et al. (2003) consider multiple stochastic volatility factors, and Adrian & Rosenberg (2008) consider an additive model of short and long-run volatility. Despite these innovations, the major innovations for GARCH-type processes came in 2008 and 2013. In 2008, Engle and Rangel propose the multiplicative two-component spline GARCH, which uses a gradually changing deterministic component, and a short-term GARCH. Subsequently, in 2013, Engle, Ghysels, and Sohn develop the GARCH-MIDAS (MIxed DAta Sampling) model, that facilitates a direct inclusion of low-frequency macroeconomic data into the long-term volatility component. Specifically, it makes use of the standard GARCH(1,1) process in modelling short-term volatility, and a MIDAS term developed by Ghysels et al. (2006) for the long-term process. Engle et al. find that the macroeconomic series they include have an approximately 30% contribution to the volatility in the most recent period, while it accounts for roughly half of the predicted volatility across their entire sample. Most prominent is the improved forecasting ability when considering longer-term horizons. When comparing the mean-squared forecasting error of the models including a macroeconomic variable with those of the spline-GARCH they find it to be consistently lower.

Since 2013, the GARCH-MIDAS model has been applied in multiple contexts. Asgharian et al. (2013) extend the analysis of macroeconomic effects by looking at the principal components of macroeconomic as a dimension-reduction technique. They suggest this improves the predictive quality of the model, however, this is contested by Conrad & Loch (2015) who don't find significant improvements. In fact, Conrad & Loch (2015) apply the GARCH-MIDAS to a wide set of macroeconomic variables related to the United States. They find these variables to be significant influences on volatility, and are able to categorise them into *leading* (e.g. housing starts) and *coincidental* variables for volatility based on their weighting schemes, which reflect the countercyclical nature of volatility. There are also several applications to the Chinese market, including inflation and production (Girardin & Joyeux, 2013), and the effect of "hot money" (Wei, Yu, et al., 2017). In turn, Conrad et al. (2018) consider cryptocurrencies and their long-term components.

Furthermore, a variety of applications focus on the analysis of commodities, such as agricultural commodity price volatility investigated by Dönmez & Magrini (2013), oil price volatility in terms of supply and demand characteristics (Pan et al., 2017), U.S. monetary policy (Amendola et al., 2017), as well as economic policy uncertainty (Wei, Liu, et al., 2017). The majority of these papers find a significant effect of the macroeconomic variables on the volatility series under consideration. More importantly, these papers find that the volatility forecasts at longer horizons (for example in Wei, Liu, et al. (2017)) are improved substantially over the basic GARCH models. There remains debate as to parameter reduction in macroeconomic variables however, as the current literature primarily finds including a single macroeconomic variable or realised variance as the most accurate and parsimonious forecasting specification. Principal components have not significantly improved the forecasting performance, indeed it appears that the model confidence set approach of Wei, Liu, et al. could be the ideal solution for reducing parameters.

Despite the inclusion of a long-term macroeconomic component, Engle, Ghysels and Sohn (2013) still find that "the full sample models are not immune to breaks" when using industrial production and inflation as factors. Thus, they are forced to split their sample into sub-samples in order to improve fit. Structural breaks in the volatility processes of asset prices have been found by multiple researchers. A concise overview is provided in Andreou & Ghysels (2002), who find structural breaks in the parameters aligned with the Russian and Asian crises. The fact that structural breaks are not accounted for is a significant limitation in the GARCH-MIDAS model, as the majority of breaks can only be determined ex-post, with one of the many testing methodologies provided (see for example Bai & Perron (1998)). However, in order to forecast accurately, structural breaks should ideally be accounted for in the model itself. Using a regime-switching model is one way to include some of these potential breaks into the model specification and improve forecasting accuracy. Regimes are a particular kind of structural break in which the overall data generating process depends on several underlying volatility processes that are activated at each period based on a probability. This allows for the discrete shifts in volatility that characterise a structural break. The reason for applying regimes to model volatility, particularly the two-regime model, stems from the countercyclical nature of volatility. Volatility tends to increase during economic downturns, and remain at low levels during upswings. Thus two GARCH processes that differ in unconditional variance could characterise the volatility process. Indeed, researchers such as Hamilton & Lin (1996) use a regime approach and find volatility to be higher during recessions.

Regime-switching models determined by a Markov process have been introduced to economics by Hamilton (1988; 1989). There are several key results worth mentioning in connection to volatility modelling. Initially suggested by Diebold (1986), research into the existence of regimes in volatility was conducted by Lamoreux & Lastrapes (1990) as well as Kim & Kon (1999), who find that the introduction of different regimes into the GARCH process reduces the persistence parameters in the GARCH process. Consequently, not accounting for such changes can lead to model misspecification. On a theoretical level, Mikosch & Stărică (2004) as well as Hillebrand (2005) suggest that the process driving the persistence in a GARCH(1,1) model to unity is due to non-stationarity brought about by "changes in the unconditional variance" of the underlying volatility process. This suggests that the application of multiple regimes of different GARCH(1,1) processes in the short-term volatility component could capture these changes in unconditional variance. Implementing such a regime-changing model, Hamilton & Susmel (1994) find that it improves fit and forecasting, and distinguish between high- and low-volatility regimes, where high volatility regimes are loosely associated with recessions. Klaassen (2002) extends the generalised regime switching model developed by Gray (1996) to dollar exchange rates, and thus resolves an issue wherein single-regime forecasts overestimate the level of volatility. Most recently, Haas et al. (2004) develop a new approach Markov-switching GARCH by using an ARCH(∞) specification. They also find significantly improved forecasts in exchange rate volatility as the regimes are able to account for the changes in unconditional variance observed by Hillebrand and Mikosch and Stărică. On the other hand, Haas et al. (2004) does find that in some cases the regime estimates are not significant, such that it is not necessarily the case that all volatility series have a regime structure.

In this paper, I propose a GARCH-MIDAS specification that includes a regime-switching unconditional variance component to improve the structural stability of the GARCH-MIDAS model and account for some of the structural breaks in a parsimonious manner. In the single-regime model of Engle et al., the macroeconomic variables create a time-varying unconditional variance. This opens up the possibility that the information about different regimes, with a high and low volatility regime, is already contained in the macroeconomic information and thus explicitly modelling regimes may not be necessary. In this regard, the following question is assessed:

(S_1) *Does the information in macroeconomic variables incorporate changes in regime sufficiently to render a second regime insignificant?*

If the regimes account for the majority of changes in the variance, then it could be that the significance of the macroeconomic variables decreases or their weighting schemes differ because they contain little additional information. In this respect, it is also interesting to investigate whether the *lead* and *coincidental* weighting structures that were found in Conrad & Loch (2015) are retained in the presence of regimes.

Based on the literature reviewed, as well as work by Stărică et al. (2005), it is likely that if regimes are present then the inclusion of two regimes would lead to an improvement in the forecasting ability of the GARCH-MIDAS model. In essence, two sub-questions will be investigated:

(S_2) *Does a Markov-switching GARCH-MIDAS specification improve **short-term** forecasting in comparison to the single-regime GARCH-MIDAS specification?*

(S_3) *Does a Markov-switching GARCH-MIDAS specification improve **long-term** forecasting in comparison to the single-regime GARCH-MIDAS specification?*

3 Data

This section gives a comprehensive overview of the stock return and macroeconomic data used. The period under consideration dates from 1^{st} January 1969 until 31^{st} December of 2010[1]. The data has been sourced from Conrad & Loch (2015) and was originally collected from the St. Louis Federal Reserve Database (FRED), the University of Michigan, the Chicago Federal Reserve, the Oxford MAN Institute, and the Bureau of Economic Analysis.

Stock Market Returns: The daily percentage log returns[2], $r_{i,t}$, based on close prices of the United States' S&P500 index are considered. Monthly realised volatility is then calculated as $RV_t = \sum_{i=1}^{N_t} r_{i,t}^2$, where N_t is the number of days in month t. Lastly, the *true* daily realised variance based on 5-minute intra-day sub-sampling is collected across the period for the 1^{st} January 2000 until the 31^{st} December 2010 in order to assess forecasting performance. Realised variance is used for forecast evaluation as volatility is not directly observable, and while the squared daily returns are an unbiased proxy they are *noisy*. Realised variance on the other hand is unbiased and non-noisy.

Macroeconomic Variables: The eleven macroeconomic variables considered in this paper are observed at a quarterly frequency, and include: the real consumption (CONS), the University of Michigan's Consumer Sentiment Index (CSI), real GDP, housing starts (HOUSE), inflation[3] (INFL), the industrial production index (IPI), the Chicago Federal Reserve's National Activity Index (NAI), the Institute for Supply Management's new orders index (NOI), corporate profit, the term spread as measured by the difference in 10-year Treasury bond and three-month Treasury bill, and lastly the unemployment rate (UNEMP). In order to apply the variables in the context of the GARCH-MIDAS models, the following transformations are taken to guarantee stationarity based on an augmented Dickey-Fuller test at the 10% significance level. The NAI, NOI, term spread and RV are included without manipulation. The first difference of the CSI and unemployment rate are taken. For all other variables the annualised percentage growth[4] are considered. Table 1 contains summary statistics and figure 1 shows time-series plots of the transformed variables.

The descriptive statistics lead to several immediate observations. Primarily, it appears that all series exhibit excess kurtosis and skew in either direction. This suggests that none of these series is normally distributed, and that observations around the mean are of extremely high frequency. Additionally, for housing, profit, realised variance and term spread the maximum is of much larger magnitude than the minimum, suggesting that outliers in these series could be present. If these extreme observations correlate with spikes in volatility, then such outliers in macroeconomic series (for instance the spikes in housing starts in the early 1980s) could be determinant for volatility in the sense that they contain a lot of information.

[1] This sample includes the 1973 oil crisis, 1987 Black Monday, the 1997 Asian financial crisis, the 1998 Russian financial crisis, the 1999 Argentine crisis, the dot-com bubble burst in 2001, and the 2008 Global Financial crisis

[2] $100 \times \log(P_t/P_{t-1})$, where P_t is the price level at time t

[3] measured by the change in GDP deflator

[4] $100 \times \left[(X_t/X_{t-1})^4 - 1\right]$, where X_t is the level of the macroeconomic variable at time t

Table 1 – Descriptive statistics of the transformed data (1^{st} Jan. 1969 - 31^{st} Dec. 2010)

	Obs.	Max	Mean	Median	Min	SD	Skew	Kurtosis	AC(1)
S&P500 Return Data									
Returns	10852	10.957	0.023	0.041	-22.900	1.088	-1.019	28.611	0.015
Macroeconomic Variables									
Δ CONS	172	10.185	2.947	3.351	-11.926	2.956	-1.249	7.427	0.078
CSI	172	16.267	-0.156	-0.333	-14.700	5.350	0.107	3.608	-0.078
Δ GDP	172	11.158	2.442	2.586	-10.367	3.190	-0.969	5.820	0.493
Δ HOUSE	172	236.050	5.890	-0.805	-69.026	43.567	1.792	10.105	0.124
INF	172	13.691	3.729	3.027	-0.327	2.628	1.206	4.208	0.826
Δ IPI	172	21.156	2.150	3.194	-29.033	6.641	-1.056	6.586	0.543
NAI	172	1.917	-0.017	0.117	-3.407	0.884	-1.409	6.243	0.734
NOI	172	71.900	54.744	55.850	27.267	7.726	-0.747	4.031	0.738
Δ PROFIT	172	180.271	12.328	11.409	-70.805	29.607	1.514	10.047	0.134
RV	172	1143.404	74.747	43.961	11.607	121.562	6.399	51.112	0.377
SPREAD	172	3.800	1.660	1.772	-1.430	1.281	-0.425	2.291	0.878
Δ UNEMP	172	1.767	0.031	-0.033	-0.9667	0.382	1.319	6.620	0.501

[1] The daily realised variance series starts on 1^{st} Jan. 2000

The presented statistics include the number of observations, maximum, mean, median, minimum, standard deviation, skewness, kurtosis, and first-order autocorrelations. The delta, Δ, denote series that have been transformed, as described in the data section (section 3).

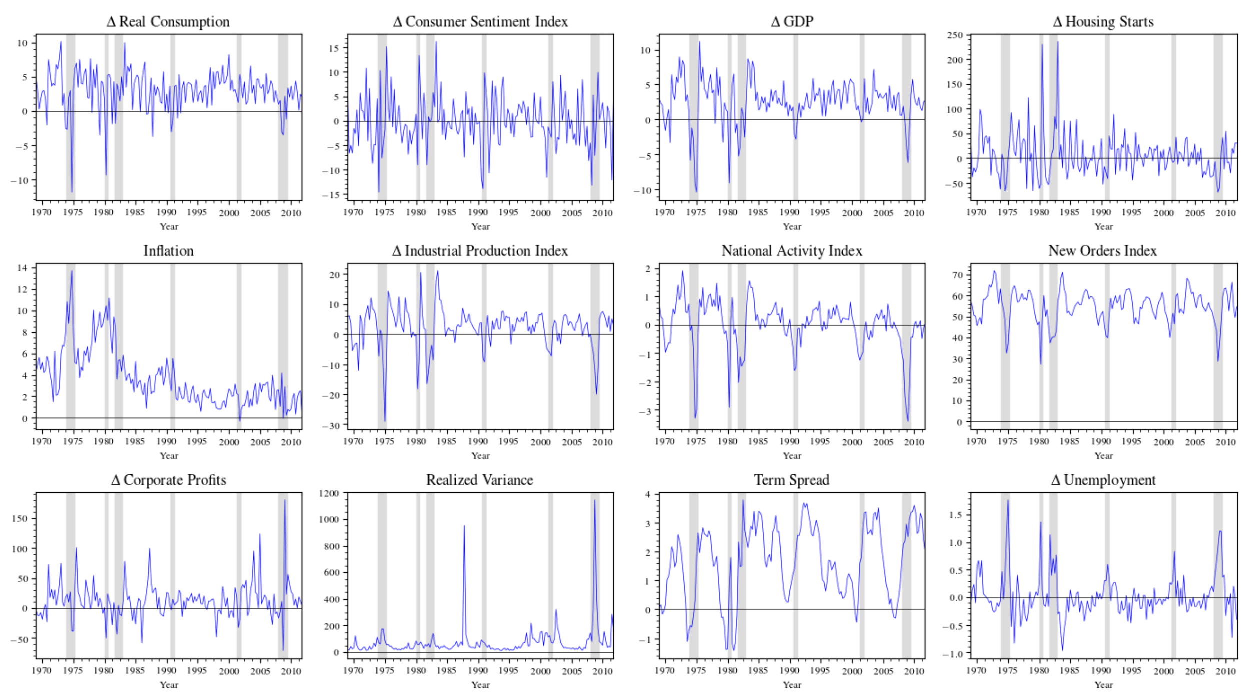

Figure 1 – Time-series graphs of the transformed data. The shaded areas represent recessions as determined by NBER

The time-series plots of figure 1 give a strong indication of the variation in each of the variables. For instance, inflation appears to have a consistently declining shape whereas the consumer sentiment index and term spread have a stronger series of peaks and troughs. Looking at the recession shadings, in particular, the cyclical nature of the macroeconomic variables can be approximated. For instance, housing starts, industrial production, and national activity are pro-cyclical in nature, experiencing troughs during recessions. Meanwhile realised variance and unemployment a strongly counter-cyclical, as would be expected. This suggests that the coefficients of the macroeconomic variables could also correspond to the cyclical nature of the macroeconomic variables in contrast to the counter-cyclical nature of volatility.

4 Methodology

This section outlines the econometric techniques applied within this paper. Firstly, the GARCH-MIDAS (GM) model, as developed by Engle et al. (2013), is introduced (section 4.1). Thereafter, the Markov-switching GARCH-MIDAS (MS-GM) model is developed (section 4.3), before presenting the forecasting and evaluation techniques applied (sections 4.4 & 4.5).

4.1 The GARCH-MIDAS Model

To begin with, the return, $r_{i,t}$, on the S&P500 on day i of month t is defined by

$$r_{i,t} = \mu + \sqrt{\tau_t \times g_{i,t}}\epsilon_{i,t}, \tag{1}$$

where μ is the time-invariant mean return, and $\epsilon_{i,t}$ is a exogenous shock scaled by the product of long-term volatility τ_t and short-term volatility $g_{i,t}$. It should be noted that $i \in \{1, \ldots, N_t\}$, where N_t is the number of days in quarter $t \in \{1, \ldots, T\}$. In this context it is assumed that the distribution of the shocks, $\epsilon_{i,t}$, conditional on all available information is standard normal, i.e. $\epsilon_{i,t}|R_{i-1,t} \sim N(0,1)$ with $R_{i-1,t}$ denoting all information up to day $i-1$ of month t.

Initially, the results of Conrad & Loch (2015) are replicated, thus their asymmetric threshold GARCH model for the short-term volatility is applied. As noted in the introduction, Conrad & Loch's results include the variables of the original paper by Engle et al., thus making their results a quasi-subset to those of Conrad & Loch. Hence, short-term volatility can be written as

$$g_{i,t} = (1-\alpha-\beta-\gamma/2) + (\alpha+\gamma\mathbb{I}_{[r_{i-1,t}-\mu<0]})\frac{(r_{i-1,t}-\mu)^2}{\tau_t} + \beta g_{i-1,t} \quad s.t.\ \alpha > 0, \beta \geq 0,\ \alpha+\beta+\gamma/2 < 1, \tag{2}$$

where $\mathbb{I}_A$ is an indicator function equal to 1 if condition A is satisfied, and 0 otherwise. The asymmetric GARCH(1,1) model implies that current volatility is the sum of a baseline level, $1-\alpha-\beta-\gamma/2$, the returns shock of the prior period scaled by α with additional weight γ if negative, and the prior period volatility scaled by β. In this model, $\alpha+\beta+\gamma/2$ measures the persistence in volatility and encapsulates the "long memory" characteristic. When using only a short-term GARCH model to approximate the volatility process, the literature consistently finds $\alpha+\beta+\gamma/2 \approx 1$, suggesting a very high degree of persistence is to be expected. With the assumptions given in equation 2, $\frac{r_{i,t}-\mu}{\sqrt{\tau_t}} = \sqrt{g_{i,t}}\epsilon_{i,t}$ is a covariance-stationary asymmetric GARCH(1,1) process.

The innovation of Engle, Ghysels and Sohn (2013) lies in the definition of the long-term volatility component, τ_t, via a MIDAS (MIxed DAta Sampling) term. Specifically, τ_t can be summarised as

$$\log(\tau_t) = m + \theta \sum_{k=1}^{K} \varphi_k(\omega_1, \omega_2) X_{t-k}, \tag{3}$$

where the log of the long-term volatility is the sum of a baseline constant m and $K \geq 0$ lags of the macroeconomic variable X weighted by the MIDAS weighting function $\varphi_k(\cdot)$ which is presented in equation 4. An important implication is that the use of the multiplicative components $g_{i,t}$ and τ_t imply that the GARCH-MIDAS in effect follows an asymmetric GARCH process with a time-varying unconditional variance τ_t. This follows from the fact that the unconditional variance of the short-term asymmetric GARCH (equation 2) is equal to 1. The use of $\log(\tau_t)$ ensures that the long-term volatility component is non-negative. Furthermore, it is assumed that the component τ_t is covariance stationary.

The variable X in the definition of τ_t (equation 3) refers to one of the macroeconomic variables introduced in the data section (section 3). The addition of more than one macroeconomic variable is not considered, on the basis that Guérin and Marcellino (2013) find a severe non-convergence problem as the number of parameters in Markov-switching MIDAS models increases.

The MIDAS weighting function is the beta function of Ghysels, Sinko and Valkanov (2006) and can be written as

$$\varphi_k(\omega_1, \omega_2) = \frac{(k/K)^{\omega_1 - 1}(1 - k/K)^{\omega_2 - 1}}{\sum_{j=1}^{K} (j/K)^{\omega_1 - 1}(1 - j/K)^{\omega_2 - 1}}. \tag{4}$$

This formulation of the weighting function requires only two parameters and guarantees that all weights are non-negative and sum up to 1. Additionally, it can generate a variety of forms including equal weights ($\omega_1 = \omega_2 = 1$), a slow or fast decline, and hump shapes.

The restriction of $\omega_2 \geq 1$ is enforced in order to guarantee that the weights are declining overall. In the literature, it is also common to fix $\omega_1 = 1$, in order to reduce the number of parameters and guarantee declining weights. However, Conrad & Loch (2015) discover hump-shaped weighting structures for some variables[5]. Consequently, both a restricted and unrestricted estimation is undertaken for each macroeconomic variable and compared via a likelihood ratio test.

The estimation of the GARCH-MIDAS model is done via quasi-maximum likelihood estimation (QMLE) following the determination of exogenous parameters. Firstly, the number of lags is set at $K = 12$ to replicate the work of Conrad & Loch (2015). Given $K = 12$, the remaining parameters are gathered in the vector $\Theta = \{\alpha, \beta, \gamma, m, \theta, \omega_1, \omega_2\}$. The log-likelihood function given by

$$\text{LLF}(\Theta) = -\frac{1}{2} \sum_{t=1}^{T} \sum_{i=1}^{N_t} \log(2\pi) + \log(\tau_t(R_{i-1,t})\, g_{i,t}(R_{i-1,t})) + \frac{r_{i,t} - \mu}{\tau_t(R_{i-1,t})\, g_{i,t}(R_{i-1,t})} \tag{5}$$

is then maximised in these parameters. The estimation effectively begins on 01-01-1973, as this is the first available observation of τ_t based on $K = 12$ lags, i.e. a three-year lag period.

[5] housing starts, consumer sentiment, and term spread

4.2 Identification of the GARCH-MIDAS

It has been noted that identification issues can arise in the GARCH-MIDAS model. If $\theta \to 0$, the parameters ω_1 and ω_2 are likely not identified, as they may take any value without altering τ_t. To fully resolve the issue of non-identification is beyond the scope of this paper. However, two measures are taken in order to reduce the likelihood of suffering from non-identification. In the first place, a global optimisation algorithm[6] is used. This should increase the likelihood of having the correct ω_j in case of non-identification as a larger variety of starting values are accounted for such that the QMLE optimum is less likely to be local. This means that even with $\theta \to 0$ it is more probable to find the correct ω_i. Furthermore, based on the work of Andrews & Cheng (2012), an *identification-category selection* (ICS) procedure is applied to determine whether θ is finite, and thus weakly-identified or unidentified. This requires the the calculation of A_n, which is defined by

$$A_n = \left(n\hat{\theta}'_n \hat{\Sigma}^{-1}_{'\theta\theta,n} \hat{\theta}_n / d_\theta\right)^{\frac{1}{2}}, \tag{6}$$

where $\hat{\Sigma}^{-1}_{'\theta\theta,n}$ is the upper-left d_θ quadrant of the variance of θ. In the case of the GARCH-MIDAS model only θ is considered as the key parameter, such that the formulation of A_n reduces to the t-statistic of θ. A_n is then compared to critical value $\kappa_n = (\ln n)^{-\frac{1}{2}} = (\ln 9592)^{-\frac{1}{2}} = 3.028$, where $n = 9592$ is the number of dates for which volatility is estimated. If $A_n > \kappa_n$ the model is considered identified, and normal analysis ensues. Else, the model with this specific macroeconomic variable will be considered unidentified, and not further treated.

4.3 The Markov-Switching GARCH-MIDAS Model

Time-series of volatilities often show evidence of multiple structural breaks. Regime switches are one type of structural break that can occur, and are useful in volatility modelling due to the countercyclical nature of volatility, as explained in the theoretical framework (section 2). In this paper the GARCH-MIDAS model is extended to a two-regime Markov-switching GARCH-MIDAS (MS-GM) model, where regimes have different unconditional variances.

In order to allow for regimes in the short-term volatility component, it is assumed that $g_{i,t}$ depends on the latent and unobservable process $s_{i,t}$, which can take values $j \in \{0, 1\}$ at each point in time. This implies that $g_{i,t}$ now takes the form of

$$g^{(j)}_{i,t} = \alpha_{0j} + (\alpha_1 + \gamma \mathbb{1}_{[r_{i-1,t}-\mu<0]}) \frac{(r_{i-1,t} - \mu)^2}{\tau_t} + \beta \overline{g_{i-1,t}} \quad s.t.\ \alpha_{1j}, \beta_j \geq 0, \tag{7}$$

where the Asymmetric GARCH(1,1) process includes a constant for the unconditional variance, $\frac{\alpha_{0j}}{1-\alpha_1-\beta-\gamma/2}$, that depends on the realisation $s_{i,t} = j$. The latent process $s_{i,t}$ is governed by a first-order ergodic[7] homogeneous Markov Chain with probability transition matrix

$$P = \begin{bmatrix} p_{00} & 1 - p_{00} \\ 1 - p_{11} & p_{11} \end{bmatrix}, \tag{8}$$

[6] The *basinhopping* approach of the *Scipy* package in Python, which is based on the work of Wales & Doye (1997)

[7] The assumption of ergodicity is required for the Hamilton filter applied to calculate the recursive smooth inference probabilities in the Expectation Maximisation Algorithm

where $p_{ij} = \mathbb{P}(s_{i,t} = j | s_{i-1,t} = i)$, implying that the current state of $s_{i,t}$ depends solely on the prior state of the latent variable, $s_{i-1,t}$.

This new GARCH(1,1) model suffers from the issue of path-dependence, due to the presence of $\overline{g_{i-1,t}}$, which depends on the previous state of $s_{i,t}$. This problem has been approached by Gray (1996) and Klaassen (2002) but their solutions suffer from analytical intractability according to Haas et al. (2004), whose solution of re-writing of the model as an ARCH(∞) is thus followed.

In order to re-write the GARCH process the restriction $\max\{\beta_0, \beta_1\} < 1$ is imposed. This should have no tangible effect on the resulting parameter estimates as it is commonly observed that $\alpha + \beta \approx 1$, but persistence remains below one. Additionally, Lamoreux and Lastrapes (1990) find that the estimates of the persistence parameter, $\alpha + \beta$, decline as regimes are introduced. Given the restriction, the GARCH(1,1) model can be reformulated into an ARCH(∞) formulation where $g_{i,t}^{(j)}$ only depends on the parameters of regime j that can be written as

$$g_{i,t}^{(j)} = \alpha_{0j}(1-\beta_j)^{-1} + \alpha_1 \sum_{q=1}^{\infty} \beta^{q-1} \frac{(r_{i-q,t} - \mu)^2}{\tau_t} + \gamma \sum_{q=1}^{\infty} \beta^{q-1} \mathbb{1}_{[r_{i-q,t} - \mu < 0]} \frac{(r_{i-q,t} - \mu)^2}{\tau_t} \quad s.t. \;\; \alpha_{1j}, \beta_j \geq 0. \tag{9}$$

In equation 9, the elements of the summation term are considered 0 up to the point $i = 1, t = 1$ where the sample data starts. The implementation of the MS-GM model requires that in the long-term volatility component, τ_t, the restriction $m = 0$ is applied. This is done in order for α_{0j} to be identified, as the regime-switch includes the constants in $g_{i,t}^{(j)}$, thus the additional constant m in the τ_t specification would interfere.

Based on the work of Pan et al. (2017), the parameters are estimated by maximising the log-likelihood of equation 10. In order to avoid starting value dependency, this is once again optimised across several starting values through the basinhopping algorithm. This also aims to avoid the numerous local-minima that can occur in the n^2 different possible paths for the latent process s. The likelihood function can be written as

$$\begin{aligned} \log(L(R_{N_t,T}; \Gamma)) = \sum_{t=1}^{T} \sum_{i=1}^{N_t} \log[&\mathbb{P}(s_{i,t} = 0 \mid R_{i-1,t}) f(r_{i,t} \mid s_{i,t} = 0, R_{i-1,t}; \Gamma) \\ &+ \mathbb{P}(s_{i,t} = 1 \mid R_{i-1,t}) f(r_{i,t} \mid s_{i,t} = 1, R_{i-1,t}; \Gamma)], \end{aligned} \tag{10}$$

where $f(\cdot)$ represents the density function of the normal distribution.

The probability $\mathbb{P}(s_{i,t} = j \mid R_{i-1,t})$ is determined recursively by applying a Hamilton filter. Specifically, using the formulation of equation 11, the inference and one-step forecast probabilities are calculated by the recursions of equations 12 and 13.

$$\Omega_{i,t|i,t} = \begin{bmatrix} \mathbb{P}(s_{i,t} = 0 \mid R_{i,t}) \\ \mathbb{P}(s_{i,t} = 1 \mid R_{i,t}) \end{bmatrix} \tag{11}$$

$$\Omega_{i,t|i,t} = \left[\Omega'_{i,t|i-1,t} \mathbf{f}_{i,t}\right]^{-1} \Omega_{i,t|i-1,t} \odot \mathbf{f}_{i,t} \tag{12}$$

$$\Omega_{i+1,t|i,t} = P' \Omega_{i,t|i,t} \tag{13}$$

Where $\odot$ denotes the Hadamard (elementwise) product of two matrices of the same dimension.

4.4 Forecasting

In order to compare the single-regime and two-regime GARCH-MIDAS models, short-term and long-term forecasting accuracy is considered. For the case of short-term forecasting, the 15-step ahead forecast is considered. Whereas for the case of long-term forecasting the 75- and 125-step forecasts are evaluated. These values correspond roughly to three-weeks, three months and five months in calendar time, assuming that there are five trading days per week.

Forecasting with the single-regime GARCH-MIDAS follows the straightforward point-forecast methods of a GARCH model. The asymmetric h-step ahead forecast can be formulated as

$$\hat{g}_{i+h,t|R_{i,t}} = 1 + \left(\alpha + \beta + \frac{1}{2}\gamma\right)^{h-1} (g_{i+1,t} - 1). \tag{14}$$

With the GARCH-MIDAS, forecasts of volatility also depend on τ. As with the GARCH models, the one-step-ahead τ_{t+1} is predetermined at point t. This means that for the 15-, 75- and 125-day horizons τ_t is determined a priori. Forecasting τ_t would be beyond the scope of this paper, as this requires direct forecasts of macroeconomic variables. It is common practice in the literature to thus take $\tau_{t+q} = \tau_{t+1}\ \forall q > 1$ (for instance in Conrad & Loch (2015)).

In the two-regime Markov-switching model, the h-step ahead forecast takes a slightly different from, namely

$$\begin{aligned}\mathbb{E}\left[\tau_t g_{i+h,t} \mid R_{i,t}\right] &= \tau_t \mathbb{E}\left[g_{i+h,t} \mid R_{i,t}\right] \\ &= \tau_t \left[\mathbb{P}(s_{i+h,t} = 0|R_{i,t})\mathbb{E}(g^{(0)}_{i+h,t} \mid R_{i,t}) + \mathbb{P}(s_{i+h,t} = 1|R_{i,t})\mathbb{E}(g^{(1)}_{i+h,t} \mid R_{i,t})\right]. \end{aligned} \tag{15}$$

In short, the h-step forecast is calculated for each regime separately, and then the weighted by the probability of being in regime $j \in 0, 1$ in h periods.

The h-step forecast for each regime can thus be written as

$$\mathbb{E}(g^{(j)}_{i+h,t} \mid R_{i,t}) = \alpha_{0j}(1-\beta)^{-1} + \alpha_1 \sum_{q=1+h}^{\infty} \beta_j^{q-1} \frac{(r_{i-q,t} - \mu)^2}{\tau_t} + \gamma \sum_{q=1+h}^{\infty} \beta^{q-1} \mathbb{1}_{[r_{i-q,t}-\mu<0]} \frac{(r_{i-q,t} - \mu)^2}{\tau_t}, \tag{16}$$

which follows from the ARCH(∞) representation of section 4.3. This representation guarantees the information is restricted to historical information available at day i of quarter t and is calculated for both regimes individually. it should be noted that this expectation converges to the unconditional volatility of each regime as $h \to \infty$. The probability of being in each regime h days ahead can be calculated via

$$\Omega_{i+h,t|R_{i,t}} = (P')^h \times \Omega_{i,t|i,t}. \tag{17}$$

The forecasts will proceed by estimating the parameters of the model on the sample leading up to 1^{st} January 1999. Thereafter, forecasts starting with the 1^{st} January 2000 will be generated for each of the models. This guarantees that no future information is included in the model estimates, and should thus give a more powerful assessment of the out-of-sample performance.

The forecasts will be evaluated against the intra-day subsampled realised variance using the mean squared prediction error (MSPE, equation 18) as well as the QLIKE (equation 19) criterion from Conrad & Kleen (2018), who find that the QLIKE is the preferred criterion for evaluating volatility forecasts on a theoretical level. This is because the QLIKE criterion is robust to noise in the volatility proxy. This is in line with the earlier findings of Patton (2011). For both criteria, lower statistics are preferable and perfect forecasts would result in both being equal to zero.

$$MSPE = \frac{1}{N}\sum_{n=1}^{N}(\sigma_n^2 - \hat{\sigma}_n^2)^2 \tag{18}$$

$$QLIKE = \frac{1}{N}\sum_{n=1}^{N}\left[\frac{\sigma_n^2}{\hat{\sigma}_n^2} - \ln\left(\frac{\sigma_n^2}{\hat{\sigma}_n^2}\right) - 1\right] \tag{19}$$

For completeness the Mincer-Zarnowitz regression, $\sigma^2 = \beta_0 + \beta_1\hat{\sigma}^2$ where $\hat{\sigma}^2$ is the volatility forecast and σ^2 is the true volatility proxied by the intra-day subsampled realised variance, will also be considered. However, it will only be considered whether the parameter estimates of the linear regression are in line with expectations (i.e. intercept $\beta_0 = 0$ and slope $\beta_1 = 1$). The R^2 measure is not considered as Conrad & Kleen (2018) find it to be unreliable due to the R^2 being high when the squared error is also high.

4.5 Evaluation & Interpretation

There are several ways in which the results of the estimated models will be evaluated. Given that a model is identified, parameter significance is assessed through t-statistics. Robust standard errors are computed by making use of the asymptotic theory of direct Maximum Likelihood Estimation. Specifically, the numerical approximation[8] of the Hessian matrix is used.

For an economic interpretation, a variance decomposition approach is considered. Specifically, based on Engle et al. (2013), the variance ratio of equation 20 is calculated for each model. A higher variance ratio implies that the long-term component explains a larger share of the expected volatility. However, it should be noted that a low variance ratio does not imply this is not the case, as it could be that the underlying variable moves very smoothly but still determines volatility.

$$\text{VAR(X)} = \frac{Var(\log(\tau_t))}{Var(\log(\tau_t g_{i,t}))} \tag{20}$$

Lastly, the in-sample fit of the models will be contrasted. Using the realised variance across the period from the 1^{st} of January 2000 onward, all identified models will be evaluated against the realised variance for this period. The assessment will take the same form as the forecasting evaluation. Specifically, the mean squared error (analogous to the MSPE, equation 18) and the QLIKE criterion (equation 19) are evaluated. Additionally, the Mincer-Zarnowitz regression will be performed. This investigation should lead to an answer as to whether the Markov-switching model can better describe the underlying volatility process.

[8]This makes use of the numdifftools package in Python

5 Results

This section outlines the empirical results obtained, and will proceed as follows: first, the results of the single-regime GARCH-MIDAS are presented in section 5.1. This is a replication and slight improvement of the results already obtained by Conrad & Loch (2015), and serves to show that the model was implemented correctly. Thereafter, section 5.2 outlines the in-sample performance of the Markov-switching GARCH-MIDAS. Finally, sections 5.3 and 5.4 compare the in-sample and forecasting performance of these models.

5.1 The single-regime GARCH-MIDAS

The results of the single-regime GARCH-MIDAS are presented in three ways: table 2 presents the parameters of the model as estimated by quasi-maximum likelihood, while figure 2 visualises the implied weighting structures in the MIDAS term and figure 3 shows the time series graphs of the annualised volatility as determined by the chosen models. In order to analyse the results, first, the short-term GARCH is analysed, before considering the long-term component in terms of the countercyclical nature of volatility, identification issues as well as the lead-lag structures in τ_t.

Beginning with the short-term volatility component, all models exhibit similar coefficients in the short-term asymmetric threshold GARCH(1,1) component (table 2). Namely, $\mu \approx 0.026, \alpha \approx 0.019, \beta \approx 0.91$, and $\gamma \approx 0.095$, all of which are significant at the 1% level. These results also mirror those of Conrad & Loch (2015) exactly. In line with the general findings of single-regime GARCH models, persistence is approaching unity, i.e. $\alpha + \beta + \frac{1}{2}\gamma \approx 0.985$. Ostensibly, this is due to the fact that regimes are not accounted for, as the addition of the long-term MIDAS does not appear to affect the finding of near-unity persistence. As is common to threshold GARCH models, negative return shocks have a larger effect on volatility than positive ones due to the leverage effect.

The true heart of the GARCH-MIDAS models lies in the long-term component, τ_t. Economically speaking, the sign of the θ coefficient for the MIDAS term including macroeconomic volatility reveals the counter-cyclical nature of volatility. For indicators that are generally pro-cyclical, such as real consumption, housing starts, national activity or the term spread, the sign of θ is negative. This implies that volatility increases whenever there is a decrease in these indicators which is associated with an economic contraction. Meanwhile, generally counter-cyclical macroeconomic indicators, such as realised variance, inflation, and unemployment, have positive θ coefficients. In fact, the signs of the θ parameters correspond to the type and cyclicality of the macroeconomic variables in lists of business cycle indicators such as those of Stock & Watson (1989) or Hertzberg & Beckman (1989). Visually, this is captured loosely in the volatility spikes that appear around periods of recession (grey shaded areas in figure 3).

Almost all coefficients θ are significant at the 1% level. Based on the identification test of Andrews & Cheng (2012) mentioned in the methodology (section 4.2) it is found that the model with inflation and the restricted case of the real consumption models are not identified. The effect of this can be seen in the undynamic flat shape of the annualised τ_t for inflation in figure 3, as compared to those of the term spread or consumer sentiment index. Consequently, the inflation

Table 2 – Quasi Maximum Likelihood Estimation results of the regular GARCH-MIDAS for all macro-economic variables on the full sample period (1^{st} Jan. 1973 - 31^{st} Dec. 2010)

	μ	α_1	β	γ	m	θ	ω_1	ω_2	Log-Likelihood	BIC	VAR(X)	Identified
Consumption	0.0263***	0.0190***	0.9183***	0.0930***	0.2904**	-0.0959**	1.0	3.2186*	-12,789.962	2.6747	4.00	
	(0.0084)	(0.0040)	(0.0058)	(0.0083)	(0.1548)	(0.0435)	(0.9328)	(2.1966)	[0.0829]			
	0.0269***	0.0194***	0.9191***	0.0910***	0.3361***	-0.1127***	2.5952***	7.2962**	-12,788.459	2.6744	6.16	✓
	(0.0084)	(0.0040)	(0.0057)	(0.0083)	(0.1295)	(0.0318)	(1.0994)	(3.3734)				
CSI	0.0269***	0.0193***	0.9160***	0.0932***	-0.0176	-0.1282***	1.0	1.3025***	-12,785.897	2.6739	6.73	✓
	(0.0084)	(0.0041)	(0.0060)	(0.0084)	(0.081)	(0.0341)	(0.3326)	(0.4497)	[0.0058]			
	0.0273***	0.0193***	0.9173***	0.0918***	-0.0158	-0.1328***	2.2456***	3.1048***	-12,782.088	2.6731	10.33	✓
	(0.0084)	(0.0040)	(0.0059)	(0.0082)	(0.0823)	(0.0288)	(0.6685)	(1.3285)				
GDP	0.0265***	0.0182***	0.9191***	0.0923***	0.1885**	-0.0784***	1.0	4.0021**	-12,788.041	2.6743	5.23	✓
	(0.0084)	(0.0040)	(0.0057)	(0.0082)	(0.1017)	(0.0233)	(1.0345)	(2.4273)	[0.6490]			
	0.0265***	0.0181***	0.9191***	0.0922***	0.1934**	-0.0808***	1.2991	4.5077**	-12,787.938	2.6743	5.49	✓
	(0.0084)	(0.0040)	(0.0057)	(0.0082)	(0.1002)	(0.0220)	(1.0636)	(2.5543)				
Housing Starts	0.0266***	0.0205***	0.9143***	0.0931***	0.0763	-0.0157***	1.0	1.5912***	-12,781.366	2.6729	12.17	✓
	(0.0084)	(0.0041)	(0.0061)	(0.0084)	(0.0822)	(0.0037)	(0.3827)	(0.56)	[0.0011]			
	0.0268***	0.0183***	0.9166***	0.0931***	0.0665	-0.0170***	2.7260***	4.1504***	-12,776.072	2.6718	18.22	✓
	(0.0084)	(0.0041)	(0.006)	(0.0082)	(0.0801)	(0.003)	(0.8174)	(1.3334)				
Inflation	0.0271***	0.0195***	0.9214***	0.0887***	-0.0957	0.0267	1.0	3.0681	-12,793.840	2.6755	0.83	
	(0.0084)	(0.004)	(0.0055)	(0.008)	(0.1589)	(0.0326)	(4.0058)	(6.5469)	[0.6272]			
	0.0271***	0.0195***	0.9214***	0.0887***	-0.1036	0.0287	1.3175	3.4707	-12,793.722	2.6755	0.96	
	(0.0084)	(0.004)	(0.0055)	(0.008)	(0.1393)	(0.0264)	(3.1321)	(5.3806)				
IPI	0.0264***	0.0174***	0.9184***	0.0930***	0.0787	-0.0424***	1.0	4.0000**	-12,787.496	2.6742	6.11	✓
	(0.0084)	(0.0041)	(0.0058)	(0.0082)	(0.0833)	(0.0118)	(0.8676)	(2.3242)	[0.47090]			
	0.0264***	0.0173***	0.9186***	0.0928***	0.0789	-0.0429***	1.7154*	5.7751**	-12,787.237	2.6742	6.55	✓
	(0.0084)	(0.0041)	(0.0058)	(0.0082)	(0.083)	(0.0109)	(1.0821)	(3.2163)				
NAI	0.0262***	0.0168***	0.9174***	0.0944***	-0.0266	-0.3030***	1.0	6.7115***	-12,783.232	2.6736	10.45	✓
	(0.0084)	(0.0041)	(0.0058)	(0.0083)	(0.0778)	(0.0601)	(1.3852)	(5.5538)	[0.9993]			
	0.0262***	0.0167***	0.9174***	0.0944***	-0.0272***	-0.3042***	1.0210	6.6982***	-12,783.129	2.6733	10.49	✓
	(0.0084)	(0.0041)	(0.0058)	(0.0083)	(0.0778)	(0.0601)	(1.3851)	(5.5524)				
NOI	0.0255***	0.0147***	0.9175***	0.0974***	2.6482***	-0.0490***	1.0	3.7683**	-12,776.102	2.6718	14.77	✓
	(0.0084)	(0.0041)	(0.0058)	(0.0083)	(0.4695)	(0.0085)	(0.6860)	(2.1174)	[0.8818]			
	0.0255***	0.0147***	0.9175***	0.0975***	2.6581***	-0.0492***	0.9297*	3.5941**	-12,776.091	2.6718	14.77	✓
	(0.0084)	(0.0041)	(0.0058)	(0.0083)	(0.4672)	(0.0085)	(0.6499)	(1.9694)				
Profit	0.0261***	0.0171***	0.9192***	0.0934***	0.2311***	-0.0190***	1.0	2.0655***	-12,782.080	2.6703	10.65	✓
	(0.0084)	(0.0042)	(0.0058)	(0.0081)	(0.0973)	(0.0042)	(0.3891)	(0.7619)	[0.7931]			
	0.0262***	0.0170***	0.9194***	0.0932***	0.2315***	-0.0191***	1.1018***	2.2236***	-12,782.040	2.6703	10.88	✓
	(0.0084)	(0.0042)	(0.0057)	(0.0081)	(0.0969)	(0.0041)	(0.404)	(0.8202)				
RV	0.0270***	0.0202***	0.9068***	0.1001***	-0.2741**	0.0033***	1.0	3.0342**	-12,784.027	2.6735	10.55	✓
	(0.0084)	(0.0042)	(0.0093)	(0.0105)	(0.0851)	(0.0006)	(0.3999)	(1.7381)	[0.2603]			
	0.0270***	0.0205***	0.9044***	0.1019***	-0.2985***	0.0036***	0.4982	1.8220**	-12,783.394	2.6734	11.68	✓
	(0.0084)	(0.0043)	(0.0079)	(0.0097)	(0.0854)	(0.0007)	(0.3935)	(0.9439)				
Term Spread	0.0274***	0.0191***	0.9147***	0.0918***	0.4060***	-0.2644***	1.0	1.4367*	-12,778.401	2.6723	12.05	✓
	(0.0084)	(0.0042)	(0.0062)	(0.0084)	(0.0998)	(0.0431)	(0.7683)	(0.974)	[0.0525]			
	0.0269***	0.0181***	0.9147***	0.0928***	0.3660***	-0.2425***	3.8882*	5.0164*	-12,776.520	2.6719	13.44	✓
	(0.0084)	(0.0042)	(0.0063)	(0.0085)	(0.0943)	(0.0388)	(2.4747)	(3.2187)				
Unemployment	0.0270***	0.0184***	0.9181***	0.0917***	-0.0285	0.5597***	1.0	5.5832*	-12,788.808	2.6745	4.89	✓
	(0.0084)	(0.0041)	(0.0058)	(0.0082)	(0.0809)	(0.167)	(1.2217)	(4.1778)	[0.6717]			
	0.0270***	0.0184***	0.9182***	0.0916***	-0.0285	0.5755***	1.3071	6.3019*	-12,788.718	2.6745	5.09	✓
	(0.0084)	(0.0041)	(0.0058)	(0.0082)	(0.0811)	(0.1608)	(1.2291)	(4.377)				

The BIC refers to the Bayesian Information Criterion (Schwarz, 1978). The variance ratio is taken as in equation 20, and can be interpreted in % terms. The numbers in brackets represent the robust standard errors of the respective coefficients. Meanwhile, the value underneath the log-likelihood represent the p-value of a likelihood ratio test for the null hypothesis of $\omega_1 = 1$. The annotations of *,**,*** represent parameter significance at the 10%, 5% and 1% levels based on a standard t-test. Additionally, demarcation with a ✓implies that the model is identified based on the identification category selection procedure of Andrews & Cheng (2012).

model is dropped from further analysis.

Overall, based on the Bayesian Information Criterion it appears as though the restricted new orders index, as well as the unrestricted term spread, offer the best model fit. Unsurprisingly, the unidentified inflation model performs worst based on the BIC. It should be noted that the absolute difference in BIC values is not very large for these different models, suggesting a very similar fit. The variance ratios, on the other hand, offer a different picture. Here the housing starts, new orders index, and term spread dominate the other models. Unemployment, consumption, and GDP explain the lowest fraction of their respective models. This is in line with the idea of leading and coincident indicators, as it could be hypothesised that leading indicators are more key in explaining volatility than are coincident indicators such as unemployment.

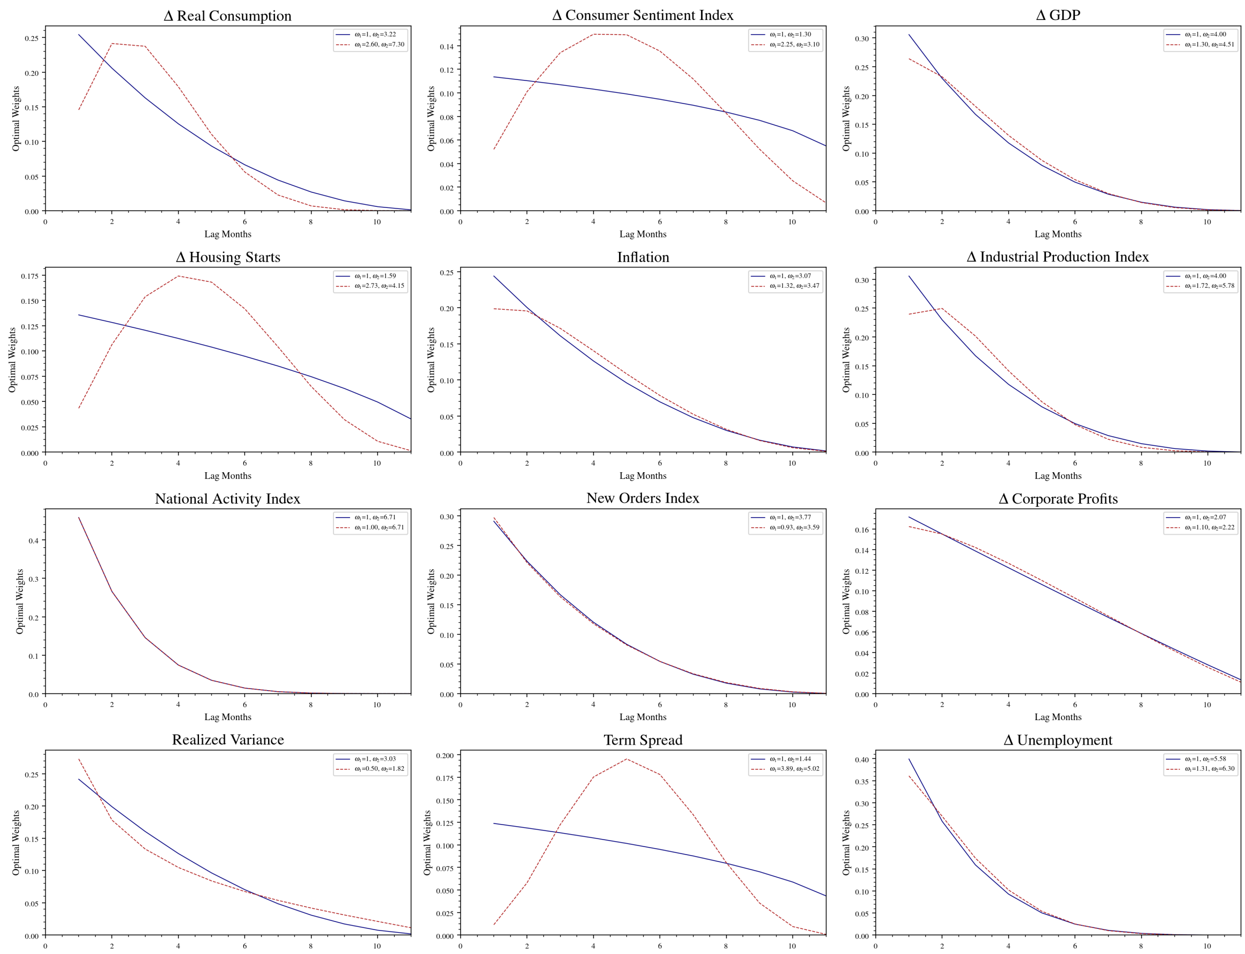

Figure 2 – Restricted and unrestricted weighting schemes of the asymmetric simple GARCH-MIDAS of Conrad & Loch (2015). Note that there is the restriction $\omega_2 > 1$. Orange lines represent the unrestricted case, while blue lines represent the restricted case of $\omega_1 = 1$

The lead-lag structure of the macroeconomic models is akin to that found by Conrad & Loch (2015). Specifically, the models using real consumption, the consumer sentiment index, housing starts, and term spread have a *hump* shaped weighting structure (see figure 2), as determined by a likelihood ratio test with a 10% significance level. The weighting shape suggests that these variables contain *leading* information for future volatility. For the remaining variables, the unrestricted weighting scheme very closely mirrors that of the restricted case (see figure 2). Consequently, these variables appear to be *coincident* indicators of the business cycles. This is confirmed by the inclusion of consumer confidence, term spread, and housing starts in leading economic indicators such as that of the OECD (2012). Interestingly, new manufacturing orders are also in the leading indicators of Hertzberg & Beckman (1989), but the likelihood ratio test does not reject the restriction of monotonically decreasing weights. In this sense, also the real consumption index is surprising, as it is not in the index of variables. It should be noted that this is only a weak rejection of $\omega_1 = 1$, whereas CSI, housing starts and term spread more strongly reject the restricted model.

Looking at the term spread first, the coefficient θ is negative, implying that an increase in the term spread has a dampening effect on volatility. With a maximum weight of 0.195 on the 5^{th} lag, this means that a 1% increase in the term spread 5 quarters ago leads to a $e^{-0.2425\times0.1953} - 1 = -0.046 \approx 4.6\%$ reduction in volatility. This effect makes economic sense, volatility is countercyclical while the term spread is pro-cyclical. Researchers, including Estrella & Hardouvelis (1991), have found that changes in the term spread are a predictor of future real economic activity. This makes the term spread a predictive variable in the business cycle analysis, and thus also for countercyclical volatility. An interpretation of this could be taken along the following lines: a negative or small term spread implies high short-term interest rates and low long-term interest rates. This represents an excess demand for the long-term "safe haven" treasury bonds, which guarantees interest payments at a fixed rate for a longer period of time. These portfolio allocation choices are symptomatic of a lack of confidence in the future short-term state of the economy, which correlates with higher variance in the returns of the stock-market, and thus volatility. Investigating the term-spread volatility time-series of figure 3, shows that the long-term τ component mirrors the volatility process very closely. Indeed, unlike some of the other variables it also moves with the 2001 dotcom bubble, which was concentrated in the technology sectors and thus might not have such pronounced macroeconomic effect as other prior crises.

Considering housing starts, Leamer (2007) makes the case that these are an important leading indicator for U.S. recessions. Kydland et al. (2016) suggests that housing starts are an important indicator due to the high dependence on mortgage financing. From an intuitive standpoint, the purchase or construction of a house is a large investment for the majority of individual households in the United States. Assuming that the majority of individuals are risk-averse, a positive outlook on the future potential of the economy and consequently increases in personal income and value in the new home are important factors promoting a purchase decision. Thus increases in housing starts reflect positive attitudes toward the future of the economy. Additionally, many housing projects will be partially financed by mortgages in order to spread the payments for the real estate across time. Hence, the available rate of financing is an important consideration in starting a house. Mortgage rates, in turn, are negatively correlated with future GDP, such that in a low-rate environment it is likely that an upturn will follow as more people start houses based on cheap

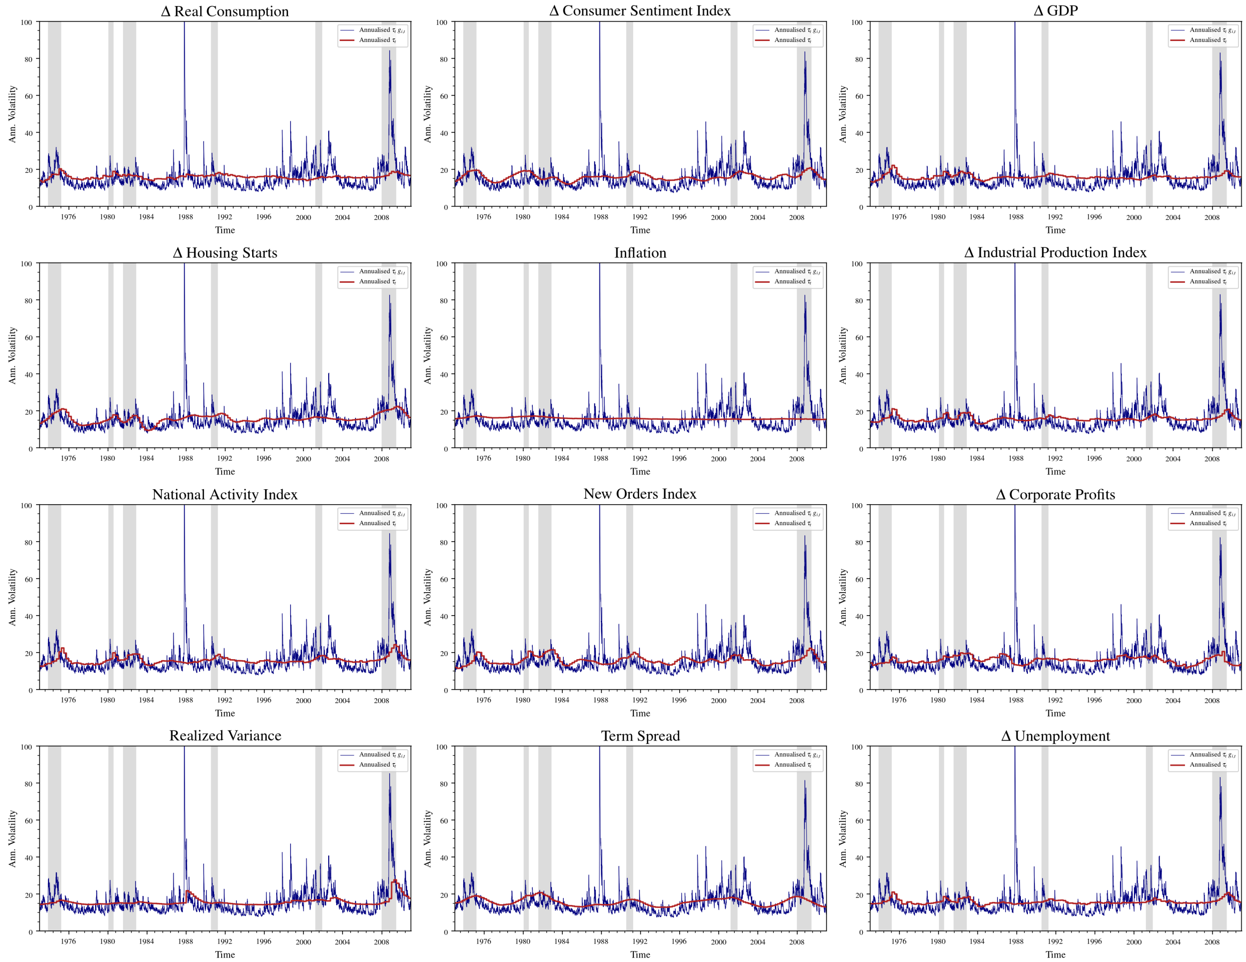

Figure 3 – Time Series plots of both the annualised total volatility $\tau_t g_{i,t}$ (blue) and the long-term component τ_t (orange). For real consumption, CSI, housing starts and term spread the unrestricted model is used, for all other variables the restriction $\omega_1 = 1$ applies. Grey shadings indicate the periods of recession, as determined by the National Bureau of Economic Research

financing. The relationship of the interest rates and future GDP has been outlined for the term spread already and is analogous here. The effects of this leading information are reflected in the estimated volatility time series shown in figure 3. Here it can be seen that the annualised τ mirrors the level of volatility extremely closely for the period until 1991. Thereafter, the dotcom bubble of the late 1990s and early 2000s cause high volatility that was limited to the digital sector, and thus isolated from the real economy. Meanwhile, there again is a very visible spike in volatility around the 2008/9 financial crisis which depended heavily on mortgage-related derivatives. However, it should be noted that the spike presented by housing is not as extreme as the actual volatility. Concluding this analysis, it appears that the leading information of housing starts is thus reflected in the GARCH-MIDAS through unrestricted weights and negative θ.

Lastly, the likelihood ratio test finds that the Consumer Sentiment Index, as well as real consumption, are significantly better described by an unrestricted weighting scheme. As with the term spread and housing start, it appears as though this variable contains information that leads the volatility process. From an economic perspective, this is not surprising. An increase in the consumer sentiment index indicates a positive outlook on the economy from a consumer perspective. This translates into a higher willingness to purchase new products. Most importantly, these indicators imply that the savings rate of individuals has decreased and spending on goods and services has increased. This, in turn, translates into increased production and profits, and potentially wages. Improved economic conditions spell an upturn in the business cycle, and consequently dampening of stock market volatility. The leading information may come from two factors. On the one hand, an increased willingness to spend would also include larger items or services that may be linked to larger periods of deliberation. Additionally, consumer sentiment might have a lagged response to new information concerning the state of the economy. Combining these factors suggests that an increase in consumer sentiment and real consumption will mark an extended period of willingness to spend and improvement in the economy. A look at figure 3 suggests that the τ related to the consumer sentiment index mirrors the shape of the volatility extremely well. Meanwhile, the real consumption time series suggests that there is little variation in the long-term component, but the changes that do occur mirror the changes in the volatility level. However, this might be a spurious finding of leading information as generally real consumption is considered a coincidental or lagging indicator of the business cycle.

Overall, the single-regime GARCH-MIDAS model presents economically reasonable results, with leading business cycle indicators also leading volatility. Identification issues were not pronounced, except for the case of inflation. Indeed, the inclusion of the macroeconomic variables explain between five and twenty percent of the variance in volatility.

5.2 The Markov-Switching GARCH-MIDAS

In this section, the in-sample results of the two-regime Markov-switching GARCH-MIDAS are analysed. In first place, the short-term component is considered. Thereafter, the countercyclical nature of volatility, parameter significance and identification issues are examined. Lastly, an investigation into the effects of regime-switching is conducted. The parameter estimation results are presented in table 3, with a visualisation of the volatility time-series for the identified models in figure 5 and corresponding weighting structures in figure 4.

In terms of the short-term GARCH component, $g_{i,t}$, observations similar to the single-regime GARCH-MIDAS (section 5.1) can be made. Specifically, the parameters α_1, β and γ are all significant at the 1% level. These parameters are extremely similar across the different models, with $\alpha_1 \approx 0.015$, $\beta \approx 0.93$, and $\gamma \approx 0.07$ (see table 3). These results also show that persistence remains high, at $\alpha_1+\beta+\frac{1}{2}\gamma \approx 0.98$, which is almost identical to the single-regime GARCH–MIDAS. This is to be expected[9] as the only difference in the two short-term components is the unconditional variance, modeled through the intercept α_{0j}. However, it should be noted that the AR component

[9] The literature finds that the persistence decreases when regimes are incorporated into all the coefficients. This is not done here due to convergence issues mentioned in section 4.3. Thus the same α_1 and β is in both regimes, and persistence does not necessarily need to decrease

β has increased in comparison to the single-regime case. Additionally, in the regime-switching case, the coefficients α_1 and γ are smaller in magnitude, probably due to the fact that some of the unexpected returns are accounted for by the different high and low volatility regimes. For instance, when in a high volatility regime, "unexpected" returns would be less unexpected by virtue of the high volatility regime, which implies that large swings in security prices, reflected in large and often negative returns, are expected. The ratio of the two coefficients remains the same however, with $\gamma > 0$ reflecting the leverage effect often observed with volatility.

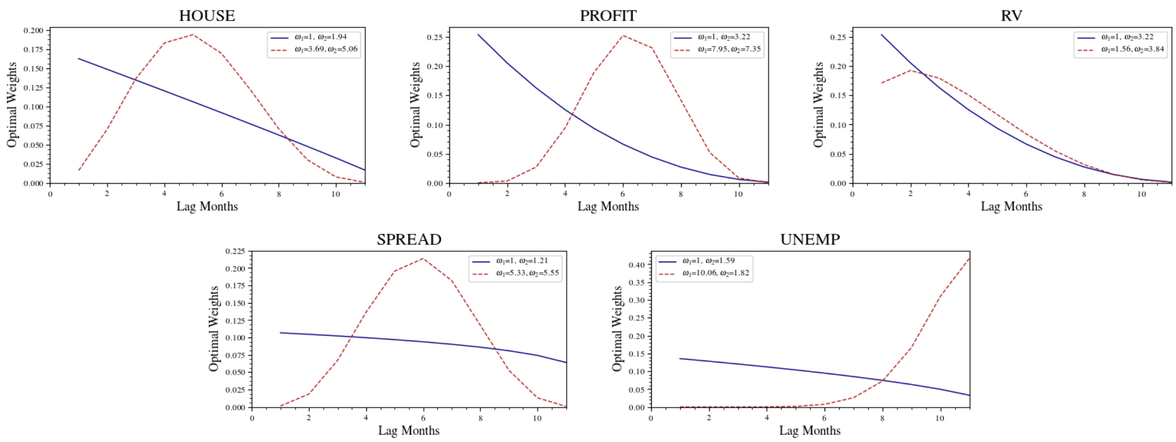

Figure 4 – Time Series plots of both the annualised total volatility $\tau_t g_{i,t}$ (blue) and the long-term component τ_t (red) for the identified models in the two-regime Markov-switching setup. Additionally, for housing starts, profit, term spread and unemployment, the unrestricted models are taken based on a likelihood ratio test. Grey shadings indicate the periods of recession, as determined by the National Bureau of Economic Research

The counter-cyclical nature of volatility is confirmed again by the signs of the θ coefficients in the two-regime GARCH-MIDAS. As discussed in the context of the results of the single-regime GARCH-MIDAS, the macroeconomic variables that are pro-cyclical have a negative coefficient. The same applies here, pro-cyclical variables such as the CSI, or NAU show negative θ, meanwhile countercyclical variables such as inflation or RV have positive coefficients. The notable exception to this is the model with unemployment. In the restricted model, the expected positive sign of θ is confirmed, as unemployment is counter-cyclical, however in the unrestricted model the sign is negative. At first glance, this seems economically counterintuitive. Considering the unrestricted weighting scheme for unemployment (figure 4) it appears that the model is misspecified. Unlike the steadily decreasing weighting scheme of the restricted model, the unrestricted model puts the largest weight onto the furthest observation. This implies that the most important observation stems from three years (12 quarters) prior to the date of the volatility that is estimated. When considering the time-series plots of the change in the unemployment rate (figure 1 of section 3) this suggests that the lag is so far back that the unemployment considered is actually in the "prior state" of the business cycle. For instance, during the 2008-09 financial crisis, volatility was at an extreme level and the economy in a recession (see figures 3 or 5), however, the unemployment change that has most weights in the MIDAS term stem from the 2005-06 period when unemployment changes were, in fact, negative (figure 1). Consequently, the resulting long-term implication is that there is a large positive τ, or in other words, a high unconditional variance. This suggests that while unemployment is countercyclical, information from the prior business cycle might have predictive power. In this respect, it is possible that the number of MIDAS lags, $K = 12$, is too small and increasing this could result in a hump shape with a peak between 3-5 years as this is the common length of a business cycle. An alternative explanation for this far lag could be that unemployment lags the business cycle, as it often takes time before it is possible to release workers from a company.

Identification issues in the two-regime Markov-switching GARCH-MIDAS are much more pronounced than in the single-regime GARCH-MIDAS. Based on the results in table 3 the majority of the models are not identified. The models that remain identified are those which in the single-regime setup rejected the restriction of $\omega_1 = 1$ and contain leading information for the volatility. Additionally, for the spread, RV and unemployment both the restricted and unrestricted models are identified. Likewise, the restricted profit model is identified. Inspecting the θ coefficients actually reveals that all of the coefficients are of the same sign but of smaller magnitude in the two-regime model as compared to the single-regime model (with the exception of unemployment, as detailed above). In effect, this suggests that the information contained in these models was in part related to explaining the changes in regimes, i.e. the sudden shifts in unconditional variance, such that explicitly accounting for a regime-structure takes the place of this information as it is contained in the macroeconomic variables. The result is that a very reduced set of models is available, and these models contain additional information on either the state of the economy or of the stock market that is relevant to the development of volatility across time. For instance, corporate profits could be relevant in giving information about specific firm "health" in addition to the sudden changes in unconditional variance that generally mark a regime switching model.

Using the handful of models that are identified (restricted profits and RV, as well as unrestricted CSI, housing, unemployment, and spread), the structure of the regimes is now assessed. Overall, the low volatility regime is represented by $\alpha_{00} \approx 0.015$, with a range of 0.0079 at the minimum and 0.0723 at the maximum. Meanwhile, the high volatility regime is not consistent among models. There is a wide range of values represented, ranging from 0.45 to 4.17 in the case of the NOI. On average, the low-volatility regime has $\alpha_{00} \approx 0.015$ and the high-volatility regime has $\alpha_{01} \approx 0.6$ (excluding the large NOI observations). This translates to annualised unconditional volatility of $\sqrt{252 \times 0.75} \approx 13.75\%$, and $\sqrt{252 \times 30} \approx 86.95\%$ for the low- and high-volatility regimes respectively. This shows the stark difference in baseline values of the two regimes, and considering the small standard errors suggests that they are significantly different from one another (i.e. a 95% confidence interval for both coefficients would not overlap).

The probability of remaining in the regime is high, at $p_{00} \approx 99\%$, for the low-volatility regime while it is extremely small for the extremely high-volatility regime at $p_{11} \approx 20\%$. A t-test reveals that p_{00} is in fact significantly different from 1 for all models, except for NOI, and that p_{11} is significantly different from 0 for all models. Based on the probabilities and stark differences in intercept, this suggests that regimes are indeed present. However, a view of figure 5, suggests a

Table 3 – Estimation results of the two-regime Markov-switching GARCH-MIDAS for all macro-economic variables on the full sample period (1^{st} Jan. 1973 - 31^{st} Dec. 2010)

	μ	α_{00}	α_{01}	α_1	β	γ	θ	ω_1	ω_2	p_{00}	p_{11}	Log-Likelihood	BIC	VAR(X	Identified
Cons.	0.0326*** (0.0082)	0.0126*** (0.0027)	0.7908*** (0.3084)	0.0155*** (0.0038)	0.9334*** (0.0052)	0.0710*** (0.0076)	-0.0829** (0.0398)	1.0	3.2537* (2.3585)	0.9912*** (0.0034)	0.2080** (0.1229)	12,641.093 [0.0578]	2.6465	2.95	
	0.0328*** (0.0082)	0.0079*** (0.0013)	0.4689*** (0.1579)	0.0157*** (0.0040)	0.9336*** (0.0055)	0.0695*** (0.0076)	0.0789*** (0.0314)	7.5501 (5.9686)	1.6106 (1.3017)	0.9912*** (0.0032)	0.2415** (0.1398)	12,639.293	2.6462	3.25	
CSI	0.0323*** (0.0082)	0.0110*** (0.0017)	0.6365*** (0.2224)	0.0153*** (0.0039)	0.9318*** (0.0055)	0.0715*** (0.0077)	-0.1147*** (0.0468)	1.0	1.4231** (0.7192)	0.9913*** (0.0033)	0.2126** (0.1269)	12,639.504 [0.1379]	2.6462	4.83	
	0.0325*** (0.0082)	0.0110*** (0.0017)	0.6961*** (0.2474)	0.0155*** (0.0038)	0.9323*** (0.0055)	0.0702*** (0.0076)	-0.1270*** (0.0374)	1.6304*** (0.5445)	1.947**3 (0.8636)	0.9920*** (0.0030)	0.2177** (0.1235)	12,638.403	2.6460	6.12	✓
GDP	0.0328*** (0.0082)	0.0113*** (0.0036)	0.6616** (0.2947)	0.0158*** (0.0039)	0.9329*** (0.0056)	0.0708*** (0.0078)	-0.0497 (0.0730)	1.0	3.9818 (6.0232)	0.9904*** (0.0037)	0.1912*** (0.1185)	12,642.014 [1.0]	2.6467	2.12	
	0.0327*** (0.0082)	0.0113*** (0.0041)	0.6387** (0.3032)	0.0155*** (0.0038)	0.9332*** (0.0055)	0.0708*** (0.0079)	-0.0505 (0.0949)	1.3010 (10.6732)	4.5069 (11.9495)	0.9903*** (0.0037)	0.2107** (0.1220)	12,642.061	2.6467	2.15	
Housing	0.0327*** (0.0081)	0.0109*** (0.0018)	0.5677*** (0.1124)	0.0159*** (0.0039)	0.9320*** (0.0056)	0.0712*** (0.0078)	-0.0084** (0.0038)	1.0	1.9439*** (0.1333)	0.9900*** (0.0026)	0.2169*** (0.0917)	12,641.039 [0.0028]	2.6465	3.41	
	0.0324*** (0.0082)	0.0119*** (0.0018)	0.5872*** (0.1990)	0.0149*** (0.0038)	0.9317*** (0.0056)	0.0717*** (0.0077)	-0.0128*** (0.0035)	3.6862*** (1.4534)	5.058***9 (2.0056)	0.9899*** (0.0038)	0.2009** (0.1091)	12,636.583	2.6456	8.44	✓
Inflation	0.0333*** (0.0082)	0.0081*** (0.0014)	0.5362*** (0.1948)	0.0165*** (0.0038)	0.9345*** (0.0052)	0.0685*** (0.0075)	0.0326 (0.0305)	1.0	3.2185 (5.6122)	0.9909*** (0.0036)	0.1928* (0.1286)	12,642.986 [1.000]	2.6469	1.39	
	0.0334*** (0.0082)	0.0081*** (0.0014)	0.5489*** (0.2058)	0.0166*** (0.0038)	0.9345*** (0.0052)	0.0684*** (0.0075)	0.0317 (0.0321)	1.3154 (3.752)	3.4717 (6.172)	0.9910*** (0.0037)	0.1874* (0.1285)	12,643.004	2.6469	1.30	
IPI	0.0328*** (0.0082)	0.0114*** (0.0025)	0.6386*** (0.2600)	0.0154*** (0.0039)	0.9323*** (0.0057)	0.0712*** (0.0078)	-0.0355 (0.0302)	1.0	3.1925 (2.5492)	0.9902*** (0.0039)	0.1853* (0.1130)	12,641.194 [0.1578]	2.6466	3.40	
	0.0328*** (0.0082)	0.0112*** (0.0021)	0.6551*** (0.2352)	0.0156*** (0.0038)	0.9330*** (0.0056)	0.0702*** (0.0078)	-0.0379** (0.0179)	3.1956 (3.3958)	6.7150 (6.0428)	0.9905*** (0.0036)	0.1874** (0.1096)	12,640.196	2.6464	4.37	
NAI	0.0329*** (0.0082)	0.0105*** (0.0020)	0.5777*** (0.2289)	0.0153*** (0.0039)	0.9320*** (0.0058)	0.0718*** (0.0080)	-0.1682 (0.1503)	1.0	6.8274 (10.9804)	0.9901*** (0.0041)	0.1878* (0.1162)	12,641.843 [1.0]	2.6467	3.02	
	0.0329*** (0.0082)	0.0106*** (0.0018)	0.6221*** (0.2484)	0.0154*** (0.0039)	0.9320*** (0.0057)	0.0717*** (0.0079)	-0.1738*** (0.0668)	1.0202 (0.8069)	6.6984 (–)	0.9906*** (0.0038)	0.1847* (0.1157)	12,641.848	2.6467	3.18	
NOI	0.0325*** (0.0082)	0.0345 (0.0403)	1.6980 (1.8316)	0.0148*** (0.0041)	0.9319*** (0.0059)	0.0730*** (0.0084)	-0.0217 (0.0191)	1.0	3.7262 (4.4378)	0.9892*** (0.0045)	0.1832* (0.1157)	12,641.142 [0.0703]	2.6466	2.78	
	0.0323*** (0.0082)	0.0723* (0.4094)	4.1737 (31.9092)	0.0147*** (0.0053)	0.9315*** (0.0091)	0.0728*** (0.0128)	-0.0341*** (0.0780)	2.3057* (12.3006)	4.9250* (22.6068)	0.9909*** (0.0073)	0.1825** (0.1091)	12,639.504	2.6462	6.00	✓
Profit	0.0334*** (0.0082)	0.0087*** (0.0013)	0.5122*** (0.1735)	0.0168*** (0.0038)	0.934***7 (0.0052)	0.0683*** (0.0075)	0.0 (0.0038)	1.0	3.2182* (0.0167)	0.9894*** (0.0029)	0.1887* (0.1384)	12,643.501 [0.0002]	2.6471	0.00	
	0.0330*** (0.0082)	0.0125*** (0.0023)	0.5016*** (0.2189)	0.0153*** (0.0039)	0.9312*** (0.0060)	0.0725*** (0.0081)	-0.0143*** (0.0050)	7.9481** (2.3280)	7.3531** (2.2550)	0.9853*** (0.0057)	0.1396* (0.1115)	12,636.435	2.6456	7.25	✓
RV	0.0332*** (0.0082)	0.0106*** (0.0017)	0.4471*** (0.1591)	0.0166*** (0.0040)	0.9246*** (0.0066)	0.0768*** (0.0087)	0.0029*** (0.0008)	1.0	3.2152* (2.1479)	0.9889*** (0.0046)	0.1964* (0.1247)	12,639.718 [0.3915]	2.6463	6.50	✓
	0.0333*** (0.0082)	0.0099*** (0.0016)	0.4675*** (0.1652)	0.0163*** (0.0039)	0.9274*** (0.0068)	0.0746*** (0.008611)	-0.2379*** (0.0457)	1.5598** (0.9055)	3.8354** (2.1545)	0.9898*** (0.0040)	0.2088** (0.1257)	12,639.351	2.6457	6.06	✓
Spread	0.0331*** (0.0082)	0.0182*** (0.0040)	0.8714*** (0.3583)	0.0155*** (0.0039)	0.9295*** (0.0062)	0.0718*** (0.0082)	-0.2379*** (0.0584)	1.0	1.2127 (0.9942)	0.9898*** (0.0044)	0.1852* (0.1147)	12636.973 [0.0102]	2.6457	7.34	✓
	0.0330*** (0.0082)	0.0190*** (0.0035)	0.8468*** (0.3283)	0.0148*** (0.0039)	0.9290*** (0.0061)	0.0727*** (0.0082)	-0.2350*** (0.0457)	5.3309* (3.4098)	5.5463* (3.4188)	0.9893*** (0.0045)	0.1694** (0.1076)	12,633.673	2.6450	9.26	✓
Unemp.	0.0335*** (0.0082)	0.0088*** (0.0012)	0.5334*** (0.1838)	0.0169*** (0.0038)	0.9344*** (0.0052)	0.0684*** (0.0075)	0.1000*** (0.0002)	1.0	1.5897*** (0.1539)	0.9896*** (0.0039)	0.1794* (0.1229)	12643.363 [0.0005]	2.6470	0.08	✓
	0.0311*** (0.0082)	0.0102*** (0.0015)	0.5894*** (0.1993)	0.0142*** (0.0039)	0.9339*** (0.0053)	0.0718*** (0.0075)	-0.6265*** (0.1954)	10.0575 (12.1232)	1.8206 (2.3425)	0.9915*** (0.0036)	0.2599** (0.1385)	12,637.371	2.6458	4.94	✓

The BIC refers to the Bayesian Information Criterion (Schwarz, 1978). The variance ratio is taken as in equation 20, and can be interpreted in % terms. The numbers in brackets represent the robust standard errors of the respective coefficients. Meanwhile, the value underneath the log-likelihood represent the p-value of a likelihood ratio test for the null hypothesis of $\omega_1 = 1$. The annotations of *,**,*** represent parameter significance at the 10%, 5% and 1% levels based on a standard t-test. Additionally, demarcation with a ✓implies that the model is identified based on the identification category selection procedure of Andrews & Cheng (2012).

different picture. In figure 5, the upper time series plots show the volatility series based on the two-regime estimates, while the lower plots show the probability of being in regime 1 at any point in time, i.e. $\mathbb{P}(s_{i,t} = 1 \mid R_{i,t})$.

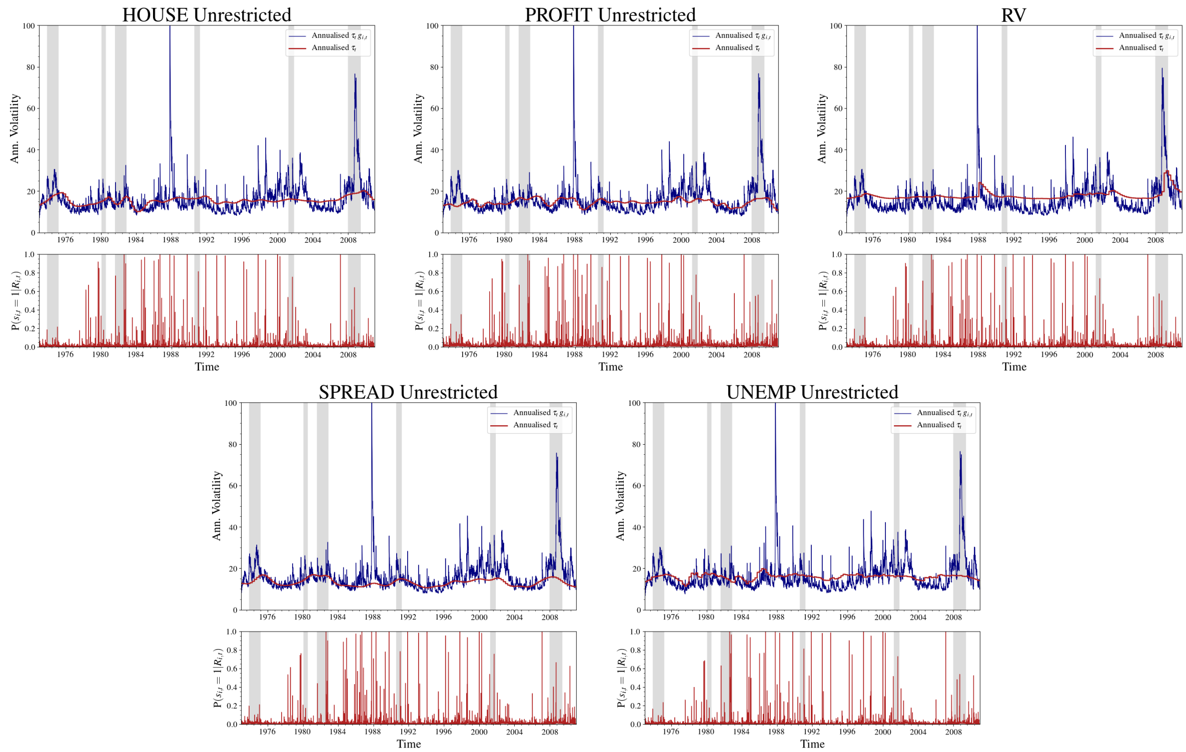

Figure 5 – Time Series plots of both the annualised total volatility $\tau_t g_{i,t}$ (blue) and the long-term component τ_t (red) for the identified models in the two-regime Markov-switching setup. Additionally, for housing starts, profit, term spread and unemployment, the unrestricted models are taken based on a likelihood ratio test. Grey shadings indicate the periods of recession, as determined by the National Bureau of Economic Research

Hamilton (1990) refers to changes in the regime as "occasional discrete shifts". These changes in the regime are governed by probabilities, in this case through the workings of a Markov chain. The workings of which are reflected in the results of the two-regime Markov-switching model. The main finding related to the incorporation of a second regime is now presented.

Considering the annualised total volatility in comparison to the annualised long term component in figure 5, it is clear that the total volatility closely resembles the shapes of the long-term volatility, in effect, there is a GARCH process that moves around the long-term τ component. This closely mirrors the results of the single-regime model (see figure 3), and is due to the same explanations already given there. Incorporating a second regime, it appears that the probability of the second high-volatility regime jumps upward whenever there is a short and extreme "shock" in the annualised volatility that is not accompanied by a corresponding spike in the long-term component. These shocks are characterised by a distinct lack of volatility buildup going into the spike. To illustrate this phenomenon, consider the housing series in figure 5. In the lead up to 1980, there are several spikes in the annualised volatility. However, these do not cause corresponding increases in the probability of high-volatility regime as there is a distinct increase in the volatility up until that point. Meanwhile, events such as the 1987 black Monday stock crash which do not have a build up of volatility across several weeks prior receive jumps in the probability of the high regime. In effect, it appears that for all identified Markov-switching series the second regime is “active” for the sudden jumps in volatility that are not fully explained or cannot be explained by the macroeconomic information that governs the predominant volatility regime. Consequently, it appears that looking only at the observations characterised in part by the high-volatility regimes, these are the “outliers” to the macroeconomic explanation of volatility. In this sense, naming the regimes as *continuous regime*, for the macroeconomically dominated regime, and *shock regime* for the high-volatility outlier regime could be more apt. It is also likely that the model is overfitted because the outliers in the sample are accounted for by the regime with increased unconditional variance. Generally, overfitting results in improved in-sample fit with worse forecasting performance. This is explored in the next sections. Overall, this suggests again that the macroeconomic variables already contain sufficient information for the relation of the business cycle to volatility, such that the modelling of regimes that generally correspond with the business cycle is not necessary.

The results found for the S&P500 differ from those obtained by Pan et al. (2017) who find much higher persistence for the second regime when considering global oil markets. This is not entirely surprising, however, as it is possible that through the aggregation of multiple stocks into one index, the S&P might be less prone to speculative attacks as compared to the crude oil market. Thus if a single stock experiences extreme swings, this does not imply the entire index will. This aggregation is similar to the diversification that is often undertaken in portfolios, and suggests that macroeconomic variables, which affect the majority of the index constituents, would have a larger explanatory power. Meanwhile, the oil market might be less driven by macroeconomic fundamentals (Pan et Al. only use production and a demand index) than the S&P is, as it is only a single commodity. Hence, individual unexpected jumps in volatility might be more common, and not explained by supply and demand but instead by regimes. As a consequence, incorporating regimes into the oil market model of Pan et Al. might add more information, as compared to the case of the S&P500 analysed here.

5.3 Comparison of In-Sample Fit

The in-sample fit of the single-regime and two-regime Markov-switching GARCH-MIDAS models is compared analogously to the forecasts. Specifically, the intra-day sampled realised variance is used as the baseline *true* volatility, and MSPE, QLIKE and the Mincer-Zarnowitz regression are performed on the available data (1^{st} January 2000 until the 31^{st} December 2010). The results are displayed in table 4.

The results show that both models fit the original data extremely well. In terms of the MSPE and QLIKE statistics, the performance of all macroeconomic models in both regime cases is very similar. The MSPE and QLIKE are slightly lower in the single-regime case, with an average of 4.1575 and 0.1980 in comparison to the 4.2417 and 0.2075 of the two-regime case.

Table 4 – In-sample fit comparison of the single-regime GARCH-MIDAS and two-regime Markov-switching GARCH-MIDAS models using intra-day Realised Variance as baseline (1^{st} Jan. 2000 - 31^{st} Dec. 2010)

			Mincer–Zarnowitz				
	MSPE	QLIKE	β_0		β_1		R^2
Single-Regime Results							
Consumption (U)	4.2554	0.1983	0.0057	(0.12)	0.8023***	(60.91)	57.57%
CSI (U)	4.1880	0.2015	-0.0183	(-0.41)	0.8139***	(61.30)	57.87%
GDP	4.2059	0.1971	**-0.0051**	(-0.11)	0.8142***	(60.85)	57.52%
Housing (U)	4.1376	0.2050	-0.0493	(-1.10)	0.8314***	(61.18)	57.78%
IPI	4.1631	0.1989	-0.0205	(-0.46)	0.8217***	(61.12)	57.73%
NAI	4.1786	0.2007	-0.0131	(-0.30)	0.8077***	(61.86)	58.32%
NOI	4.1171	**0.1936**	-0.0140	(-0.32)	0.8229***	(61.63)	58.14%
Profit	4.1918	0.1890	0.0066	(0.15)	0.8217***	(60.39)	57.15%
RV	4.1394	0.2006	-0.0342	(-0.78)	0.8142***	(62.21)	**58.60%**
Term Spread (U)	**3.9907**	0.1945	-0.0455	(-1.02)	**0.8671*****	(61.25)	57.83%
Unemployment	4.1651	0.1983	-0.0164	(-0.37)	0.8200***	(61.16)	57.76%
Two-Regime Results							
Housing (U)	4.2501	0.2153	-0.0978	(-2.09)	0.8715***	(57.97)	55.13%
Profit	4.2212	**0.1991**	**-0.0563**	(-1.21)	0.8689***	(58.08)	55.22%
RV	4.3552	0.2114	-0.0572	(-1.23)	0.8302***	(58.28)	55.40%
Term Spread (U)	**4.1188**	0.2042	-0.1014	(-2.17)	**0.9079*****	(58.43)	**55.52%**
Unemployment (U)	4.2765	0.2084	-0.0607	(-1.30)	0.8631***	(57.63)	54.84%

The table shows in-sample fit for the period 01-01-2000 until 31-12-2010 based on the model estimation in comparison to the intra-day realised variance. MSPE refers to Mean Squared Prediction error. The Mincer-Zarnowitz regression is of the form $\sigma^2 = \beta_0 + \beta_1 \hat{\sigma}^2$, where $\hat{\sigma}^2$ is the estimated volatility. The values in brackets show the t-statistics of the respective coefficients. For the single- and two-regime models, the best-in-category variable has been bolded. The (U) denotes those models where the unrestricted weighting scheme is used. Note: non-identified models have been omitted from this analysis.

Differences between the two models arise when considering the Mincer-Zarnowitz results. Overall the results appear similar, with intercepts that are not significantly different from 0 and slopes that are significant and tend to 1. When looking more closely, it appears that the slope coefficients of the single-regime model are below those of the two-regime model, while the intercepts are slightly higher. This implies that the two-regime model is less likely going to over-estimate the magnitude of the volatility as compared to the single-regime model. The interpretation of the second regime as *shock* regime is an explanation of this. Accounting for the shocks with discrete jumps in unconditional variance means that the magnitude estimates in each regime are less affected by aberrant observations. A visual representation of this is shown in figure 6, where single- and two-regime term spread GARCH-MIDAS is overlaid on the realised variance. The figure shows that the realised variance often is below the estimated volatility, which remains at the upper level of the realised variance observations. Additionally, the peaks of the two-regime model (gold line in figure 6) match the sudden jumps in volatility more closely than the single-regime model. However, neither of the models fully match the discrete jumps of the 2000 to 2003 period or of the early 2008 period. The results of the in-sample assessment lead to a provisional answer to the question of whether the information in macroeconomic variables incorporates regime changes (question S_1

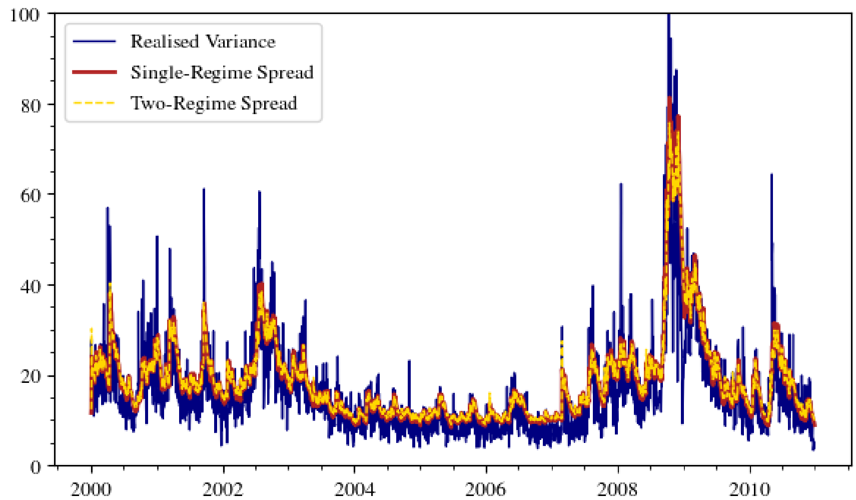

Figure 6 – Time series plot of the Term Spread models in both the single and two-regime cases, overlaid on the intra-day subsampled realised variance

of section 2). Namely that because the model fit between both cases is strikingly similar, this suggests that there is not much additional information in modelling discrete regime switches as this information appears to already be contained in the changes in macroeconomic variables.

5.4 Forecasting Performance Evaluation

In this section, the forecasting performance of the single- and two-regime GARCH-MIDAS models are compared at three horizons, 15-days, 75-days, and 125-days. The dominance of the GARCH-MIDAS over the simple GARCH(1,1) model in terms of long-horizon forecasting has been shown multiple times in the literature (e.g. see Conrad & Loch (2015)) and is thus not repeated here. The results of the forecasting exercise are shown in table 5 for the single-regime model, and in table 6 for the two-regime model. In both cases, the non-identified models have been omitted.

Beginning with the short-term forecasts, it can be stated that while the performance of the single- and two-regime models is similar, the two-regime model performs slightly worse. This becomes evident when considering the MSPE and QLIKE values. For the single-regime case, these average 6.9 and 0.35 respectively, while for the two-regime case they are 7.5 and 0.5. The same effect occurs with the long-term forecasts, where for 75-days the MSPE and QLIKE of the two-regime are 15% and 380% higher than the single-regime forecasts. Similar observations are made for the 125-day forecasts where differences are 12% and 260%.

Before evaluating the performance of the two-regime model, an investigation into the single-regime performance is undertaken. Considering the forecasting framework (see equations 14 and 15 of section 4.4), it is observed that as the forecast horizon h increases, the forecast itself converges to the unconditional variance, i.e. as $h \to \infty \Rightarrow g_{i+h,t} = \mathbb{E}(g_{i+h,t} \mid R_{i,t}) \to 1$ by the nature of the Asymmetric GARCH(1,1) specification of Conrad & Loch (2015) (see equation 2 of section 4.1). Since the forecast of σ^2 also depends on τ_t this implies that in the very long horizons,

Table 5 – Forecasting evaluation of the single-regime GARCH-MIDAS using intra-day Realised Variance as baseline (1^{st} Jan. 2000 - 31^{st} Dec. 2010)

			Mincer–Zarnowitz				
	MSPE	QLIKE	β_0		β_1		R^2
15-Day Forecasts							
Consumption (U)	0.834	0.948	0.298	(4.897)	0.742	(32.219)	27.5 %
CSI (U)	0.820	0.976	0.249	(4.081)	0.758	(33.000)	28.5 %
GDP	0.846	0.952	0.310	(5.038)	0.737	(31.329)	26.4 %
Housing (U)	0.825	0.946	0.184	(2.893)	0.792	(31.745)	26.9 %
IPI	0.839	0.957	0.284	(4.588)	0.751	(31.481)	26.6 %
NAI	0.831	0.943	0.262	(4.309)	0.734	(32.920)	28.4 %
NOI	0.842	**0.937**	0.311	(5.083)	0.741	(31.545)	26.7 %
Profit	0.891	0.951	0.416	(6.805)	0.680	(29.629)	24.3 %
RV	1.0	1.0	0.445	(7.034)	0.578	(27.210)	21.3 %
Term Spread (U)	**0.817**	0.968	**0.177**	(2.742)	**0.927**	(30.854)	25.8 %
Unemployment	0.835	0.946	0.277	(4.493)	0.752	(31.800)	27.0 %
75-Day Forecasts							
Consumption (U)	0.882	0.967	0.540	(5.504)	0.811	(11.502)	4.6 %
CSI (U)	**0.855**	0.866	0.270	(2.675)	**0.983**	(14.145)	6.8 %
GDP	0.889	0.991	0.608	(6.233)	0.755	(10.762)	4.1 %
Housing (U)	0.856	**0.851**	0.075	(0.654)	1.083	(13.769)	6.5 %
IPI	0.885	0.969	0.539	(5.300)	0.801	(10.916)	4.2 %
NAI	0.872	0.877	0.530	(5.681)	0.740	(12.504)	5.4 %
NOI	0.883	0.935	0.618	(6.776)	0.748	(11.768)	4.8 %
Profit	0.891	0.898	0.757	(9.423)	0.640	(12.261)	5.2 %
RV	1.0	1.0	0.985	(12.509)	0.319	(8.772)	2.7 %
Term Spread (U)	0.905	1.061	**-0.205**	(-1.266)	1.734	(11.017)	4.2 %
Unemployment	0.881	0.941	0.505	(4.979)	0.821	(11.332)	4.5 %
125-Day Forecasts							
Consumption (U)	0.906	0.951	0.699	(4.181)	0.767	(4.876)	0.9 %
CSI (U)	0.875	0.797	**0.051**	(0.328)	1.277	(9.856)	3.4 %
GDP	0.910	0.966	0.824	(4.933)	0.647	(4.086)	0.6 %
Housing (U)	**0.862**	**0.751**	-0.317	(-1.923)	1.493	(11.503)	4.6 %
IPI	0.905	0.928	0.706	(4.028)	0.740	(4.573)	0.8 %
NAI	0.892	0.831	0.675	(4.917)	0.696	(6.322)	1.4 %
NOI	0.901	0.884	0.638	(4.508)	0.832	(6.385)	1.5 %
Profit	0.885	0.778	0.499	(4.475)	**1.008**	(10.040)	3.6 %
RV	1.0	1.0	1.452	(15.368)	0.009	(0.169)	0.0 %
Term Spread (U)	0.906	0.946	-0.350	(-1.571)	2.018	(8.422)	2.5 %
Unemployment	0.901	0.900	0.605	(3.513)	0.827	(5.289)	1.0 %

The table shows forecast performance at three horizons. MSPE and QLIKE refer to the ratio as compared to the RV model. The Mincer-Zarnowitz regression is of the form $\sigma^2 = \beta_0 + \beta_1 \hat{\sigma}^2$, where $\hat{\sigma}^2$ is the estimated volatility. The values in brackets show the t-statistics of the respective coefficients. For each forecast horizon, the best-in-category variable has been bolded. The (U) denotes those models where the unrestricted weighting scheme is used. Note: non-identified models have been omitted from this analysis.

given the assumption of $\tau_{t+h} = \tau_{t+1}$ used here, the forecast simply reduces to τ_{t+1}. In effect this means that the long-term volatility forecast is converging to the one-step macroeconomic effect. Considering again the plots of the fit (figure 3), the majority of long-term components mirror the lower bounds of volatility (with the exception of RV, which tends to be above the volatility level,

Table 6 – Forecasting evaluation of the two-regime Markov-switching GARCH-MIDAS using intra-day Realised Variance as baseline (1^{st} Jan. 2000 - 31^{st} Dec. 2010)

			Mincer–Zarnowitz				
	MSPE	QLIKE	β_0		β_1		R^2
15-Day Forecasts							
Housing (U)	**0.997**	1.057	0.073	(1.095)	1.697	(30.939)	25.9%
Profit	0.99	1.034	0.17	(2.564)	1.503	(29.605)	24.3%
RV	1.0	**1.0**	0.312	(4.69)	**1.257**	(26.874)	20.9%
Term Spread (U)	1.044	1.052	-0.306	(-3.989)	2.172	(30.213)	25.0%
Unemployment (U)	1.002	1.043	**0.048**	(0.706)	1.675	(30.035)	24.8%
75-Day Forecasts							
Housing (U)	1.016	1.167	-2.089	(-8.094)	11.467	(14.098)	6.8%
Profit	1.02	1.237	-2.091	(-9.24)	12.457	(16.196)	8.8%
RV	1.0	1.0	**0.432**	(2.9)	**2.726**	(7.508)	2.0%
Term Spread (U)	**0.984**	**0.734**	-0.78	(-4.119)	5.236	(12.408)	5.3%
Unemployment (U)	1.019	1.175	-4.465	(-10.014)	18.721	(13.404)	6.2%
125-Day Forecasts							
Housing (U)	1.005	1.065	-1.97	(-7.196)	11.376	(12.815)	5.7%
Profit	1.011	1.173	-2.718	(-9.601)	15.382	(15.064)	7.7%
RV	1.0	1.0	1.415	(8.599)	**0.136**	(0.324)	0.0%
Term Spread (U)	**0.966**	**0.601**	**-1.37**	(-7.224)	6.609	(15.619)	8.2%
Unemployment (U)	1.009	1.1	-4.606	(-7.931)	19.862	(10.502)	3.9%

The table shows forecast performance at three horizons. MSPE and QLIKE refer to the ratio as compared to the RV model. The Mincer-Zarnowitz regression is of the form $\sigma^2 = \beta_0 + \beta_1 \hat{\sigma}^2$, where $\hat{\sigma}^2$ is the estimated volatility. The values in brackets show the t-statistics of the respective coefficients. For each forecast horizon, the best-in-category variable has been bolded. The (U) denotes those models where the unrestricted weighting scheme is used. Note: non-identified models have been omitted from this analysis.

perhaps due to effects of shocks). This explains the positive intercept coefficient of the Mincer-Zarnowitz regressions. Figure 6 of the in-sample spreads shows that with intra-day realised variance as baseline the estimated volatility tends to be in the upper range of intra-day variance fluctuations. This in turn explains the $\beta_1 < 1$ phenomenon of the mincer-Zarnowitz regressions in the single-regime case. Overall, the single-regime GARCH-MIDAS models appear to give forecasts that are quite accurate. Within this class of models, the short-term 15-day forecasts are most accurate when using the term spread. Ostensibly, this is due to changes in market expectations which are almost instantaneously reflected in the term spread, while other variables are lagged and thus might not reflect changes in market expectations as promptly. Intermediate-term forecasts are dominated by consumer sentiment, with housing also showing high accuracy (unbiased and $\beta_1 \approx 1$). This is also true of the long-term 125-day forecasts. Here, housing starts are again dominant. This is due to the large investment that starting a house requires, these investments will mostly be made when the individuals that are investing are positive in their outlook on the economy, else there would most likely not be a sufficient rate of return on investment to warrant the initial outlay. The paper of Conrad & Loch (2015), as well as the remaining literature, have generally observed that at such horizons the forecasting performance of the GARCH-MIDAS class of models is definitively superior to the use of only a GARCH component. When considering the forecasts overall, the models that have an unrestricted weighting scheme appear to perform best in terms of forecasting performance. For instance, term spread, housing, and consumer sentiment are often among the best forecasting models in terms of MSPE and QLIKE. This is in line with the interpretation that these variables contain leading information about the underlying volatility process. It should be noted that at the longer-term horizon (125-days), the horizon is almost at exactly the weighting peak of the unrestricted models. For example, the highest weight of the housing starts occurs at a 4-month lag, which is almost exactly 125-days prior. At this horizon, the unrestricted models, with exception of consumption, are all unbiased. However, they appear to underestimate the magnitude of volatility, as indicated by $\beta_1 > 1$ for CSI, housing and term spread, whereas all other models are downward biased and overestimate the volatility ($\beta_0 > 0$ and $\beta_1 < 1$).

Turning to the forecasts using the two-regime model, the results are not promising. As described above, the QLIKE characteristics, as well as the MSPE, are significantly higher when using this two-regime approach. The reason for the under-performance of the two-regime model is again due to the nature of long-term forecasting and can be seen when considering the coefficients of the Mincer-Zarnowitz regression. In the case of the Markov-switching model, the same convergence characteristics as for the single-regime model are in play. In this case, the volatility for each regime converges to the regime specific unconditional variance, $\frac{\alpha_{0j}}{1-\alpha_1-\beta-\gamma/2}$, which is approximately $\bar{g}^{(0)} \approx 0.75$ and $\bar{g}^{(1)} \approx 30$. Likewise, the probability of being in the respective regimes converges to the steady-state probabilities. Taking $p_{00} \approx 99\%$ and $p_{11} \approx 20\%$ these converge to steady state probabilities $\vartheta_0 \approx 98.77\%$ and $\vartheta_1 \approx 1.23\%$. Summa summarum, on average the unconditional volatility that the two-regime models converge to is $\tau_{t+1}\left(\vartheta_0\bar{g}^{(0)} + \vartheta_1\bar{g}^{(1)}\right) = 1.1 \times \tau_{t+1}$. At first glance, this does not appear significantly different from the single-regime models, which converge to τ_{t+1}. However, the forecasting performance of these models is considerably worse. In fact, for every category analysed here (MSPE, QLIKE, β_0 and β_1) the Markov-switching model is inferior to the single-regime model, and performance drops much more rapidly in the horizon. In the short-term 15-day scenario, the Markov-models show a significantly higher β_1 coefficient as compared to the standard models, suggesting that they significantly underestimate the magnitude of future volatility. In fact, this is true across forecast horizons. The reason for this is the predominant time spent in regime 0, which has a much lower unconditional variance than the single-regime GARCH-MIDAS model. This leads to forecasts that are consistently below the intra-day volatility, which becomes more extreme as the forecast horizon increases.

In terms of ranking the models by forecasting ability, the term-spread appears to consistently outperform the other models regardless of the horizon. This is in contrast to the optimal long-term models when considering the single-regime models, where housing starts and consumer sentiment outperform the term spread at long horizons.

Overall, the answer to questions S_2 and S_3 of section 2 is that a Markov-switching model does not improve forecasting ability in comparison to the original GARCH-MIDAS models. In fact, the forecasting ability of the single-regime GARCH-MIDAS is superior in every metric considered here.

6 Conclusion

In conclusion, the answer to the research question posed in the introduction (section 1) is that the two-regime Markov-switching GARCH-MIDAS model does not yield significant improvements in descriptive ability, and is inferior in terms of forecasting. This follows from the observation that the second regime of the Markov-switching model is primarily activated when there are sudden spikes in volatility, and thus acts somewhat as an outlier detection mechanism (section 5.2). This implies that the macroeconomic variables considered in this GARCH-MIDAS setup already incorporate information related to regime changes. Thus the time-varying unconditional variance represented by τ_t sufficiently accounts for periods of high and low volatility. In terms of in-sample fit the two-regime model is on par with the single-regime model (section 5.3), while under-performing severely in terms of forecasting (section 5.4). These observations could be in line with a diagnosis of overfitting in the case of the two-regime model.

There are several limitations inherent in the results presented in this paper. In the first place, the index considered is the S&P500. The United States is a developed country, with a functioning economy and free capital market. With a developed economy it is possible, as shown in this and other papers, to distill the effects on volatility of the macroeconomy. However, this might not be the case for less developed capital markets or those subject to capital controls. The clear impacts of housing or consumer sentiment, which contain regime information, might not be as clear for a country such as Vietnam or Egypt, where regimes in volatility might be present. Additionally, it should be noted that a limited period of time has been considered, which stops shortly after the financial crisis of 2008. Consequently, it is possible that the current effects of the macroeconomy onto the volatility process might have changed. For instance, the prolonged period of depressed long-term interest rates that have arisen from central banks' quantitative easing programs might have changed the relation of the term spread to the underlying volatility process.

While the results of this paper hold for the particular case of the S&P500, further research is warranted into the macroeconomic effects on volatility and the presence of regimes. For instance, it should be investigated whether regimes still exist for individual assets or smaller groups of assets, and in which cases the macroeconomic variables are sufficient to describe regimes. In considering the S&P500 index, the aggregation of 500 individual stocks reduce the impact of individual shocks and thus there are more clear macroeconomic effects than there could be on an individual stock. Different approaches to changes in the volatility process should also be explored, for instance, it could be plausible that the significance of the macroeconomic variables is dependent on the state of the economy, such that in a downturn effects might be larger. This could be explored through a Markov-switching MIDAS component, or an activation function for the MIDAS component when there are *significant* changes in the macroeconomic fundamentals.

References

Adrian, T., & Rosenberg, J. (2008). Stock Returns and Volatility: Pricing the Short-Run and Long-Run Components of Market Risk. *Journal of Finance*, *63*(6), 2997–3030.

Alizadeh, A., Brandt, M., & Diebold, F. (2002). Range-Based Estimation of Stochastic Volatility Models. *Journal of Finance*, *57*(3), 1047–1091.

Amendola, A., Candila, V., & Scognamillo, A. (2017). On the influence of US monetary policy on crude oil price volatility. *Empirical Economics*, *52*(1), 155–178.

Andreou, E., & Ghysels, E. (2002). Detecting Multiple Breaks in Financial Market Volatility Dynamics. *Journal of Applied Econometrics*, *17*(5), 579–600.

Andreou, E., & Ghysels, E. (2003). *Quality Control for Financial Risk Management: Monitoring Disruptions in the Distribution of Risk Exposure.* (Discussion Paper, UNC)

Andrews, D., & Cheng, X. (2012). Estimation and inference with weak, semi-strong, and strong identification. *Econometrica*, *80*(5), 2153–2211.

Ang, A., & Timmermann, A. (2012). Regime Changes and Financial Markets. *Annual Review of Financial Economics*, *4*, 313–337.

Asgharian, H., Hou, J., & Javed, F. (2013). The Importance of the Macroeconomic Variables in Forecasting Stock Return Variance: A GARCH-MIDAS Approach. *Journal of Forecasting*, *32*(7), 600–612.

Bai, J., & Perron, P. (1998). Estimating and Testing Linear Models with Multiple Structural Changes. *Econometrica*, *66*(1), 47–78.

Bollerslev, T. (1986). Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, *31*(3), 307–327.

Chernov, M., Gallant, A., Ghysels, E., & Tauchen, G. (2003). Alternative models for stock price dynamics. *Journal of Econometrics*, *116*(1), 225–257.

Conrad, C., Custovic, A., & Ghysels, E. (2018). Long- and Short-Term Cryptocurrency Volatility Components: A GARCH-MIDAS Analysis. *Journal of Risk and Financial Management*, *11*(2), 23–35.

Conrad, C., & Kleen, O. (2018). *Two are better than one: volatility forecasting using multiplicative component GARCH models.* (Unpublished Working Paper. A preliminary version of this paper circulated under the title “On the Statistical Properties of Multiplicative GARCH Models” (2016))

Conrad, C., & Loch, K. (2015). Anticipating Long-Term Stock Market Volatility. *Journal of Applied Econometrics*, *30*(7), 1090–1114.

Diebold, F. (1986). Modeling The persistence Of Conditional Variances: A Comment. *Econometric Reviews*, *5*(1), 51–56.

Ding, Z., & Granger, C. (1996). Modeling Volatility Persistence of Speculative Returns: A New Approach. *Journal of Econometrics*, *73*(1), 185–215.

Dönmez, A., & Magrini, E. (2013). *Agricultural Commodity Price Volatility and it's Macroeconomic Determinants: A GARCH-MIDAS Approach.* (ULYSSES Working Paper no. 5)

Engle, R. (1982). Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of the United Kingdom Inflation. *Econometrica*, *50*(4), 987–1007.

Engle, R., Ghysels, E., & Sohn, B. (2013). Stock Market Volatility and Macroeconomic Fundamentals. *The Review of Economics and Statistics*, *95*(3), 776–797.

Engle, R., & Lee, G. (1999). A Permanent and Transitory Component Model of Stock Return Volatility. In R. Engle & H. White (Eds.), *Cointegration, Causality, and Forecasting: A Festschrift in Honor of Clive W.J. Granger* (p. 475—497). New York: Oxford University Press.

Engle, R., & Ng, V. (1993). Measuring and Testing the Impact of News on Volatility. *The Journal of Finance*, *48*(5), 1749–1778.

Estrella, A., & Hardouvelis, G. A. (1991). The term structure as a predictor of real economic activity. *The journal of Finance*, *46*(2), 555–576.

Ghysels, E., Sinko, A., & Valkanov, R. (2006). MIDAS Regressions: Further Results and New Directions. *Econometric Reviews*, *26*(1), 53–90.

Girardin, E., & Joyeux, R. (2013). Macro fundamentals as a source of stock market volatility in China: A GARCH-MIDAS approach. *Economic Modelling*, *34*, 59–68.

Gray, S. (1996). Modeling the Conditional Distribution of Interest Rates as a Regime-Switching Process. *Jounral of Financial Economics*, *42*, 27–62.

Guérin, P., & Marcellino, M. (2013). Markov-Switching MIDAS Models. *Journal of Business & Economic Statistics*, *31*(1), 45–56.

Haas, M., Mittnik, S., & Paolella, M. (2004). A New Approach to Markov-Switching GARCH Models. *Journal of Financial Econometrics*, *2*(4), 493–530.

Hamilton, J. (1988). Rational-Expectations Econometric Analysis of Changes in Regime: An Investigation of the Term Structure of Interest Rates. *Journal of Economic Dynamics and Control*, *12*(2), 385–423.

Hamilton, J. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*, *57*(2), 357–384.

Hamilton, J. (1990). Analysis of Time Series Subject to Changes in Regime. *Journal of Econometrics*, *45*(1), 39–70.

Hamilton, J., & Lin, G. (1996). Stock market volatility and the business cycle. *Journal of Applied Econometrics*, *11*(5), 573–593.

Hamilton, J., & Susmel, R. (1994). Autoregressive Conditional Heteroskedasticity and Changes in Regime. *Journal of Econometrics*, *64*(1), 307–333.

Hertzberg, M. P., & Beckman, B. A. (1989). Business cycle indicators: Revised composite indexes. *Business Conditions Digest*, *29*(1).

Hillebrand, E. (2005). Neglecting parameter changes in GARCH models. *Journal of Econometrics*, *129*(1), 121–138.

Kim, D., & Kon, S. (1999). Structural change and time dependence in models of stock returns. *Journal of Empirical Finance*, *6*(3), 283–308.

Klaassen, F. (2002). Improving GARCH Volatility Forecasts with Regime-Switching GARCH. *Empirical Economics*, *27*, 363–394.

Kydland, F. E., Rupert, P., & Šustek, R. (2016). Housing dynamics over the business cycle. *International Economic Review*, *57*(4), 1149–1177.

Lamoreux, C., & Lastrapes, W. (1990). Persistence in Variance, Structural Change, and the GARCH Model. *Journal of Business and Economic Statistics*, *8*(2), 225–234.

Leamer, E. E. (2007). *Housing is the business cycle* (Tech. Rep.). National Bureau of Economic Research.

Mikosch, T., & Stărică, C. (2004). Nonstationarities in Financial Time Series, the Long-Range Dependence, and the IGARCH Effects. *Review of Economics and Statistics*, *86*(1), 378–390.

Nelson, D. (1991). Conditional Heteroskedasticity in Asset Returns: A New Approach. *Econometrica*, *59*(2), 347–370.

OECD. (2012). *OECD System of Composite Leading Indicators.* Retrieved from `http://www.oecd.org/sdd/leading-indicators/41629509.pdf`

Officer, R. (1973). The Variability of the Market Factor of the New York Stock Exchange. *The Journal of Business*, *46*(3), 434–453.

Pan, Z., Wang, Y., Wu, C., & Yin, L. (2017). Oil price volatility and macroeconomic fundamentals: A regime switching GARCH-MIDAS model. *The Journal of Empirical Finance*, *43*, 130–142.

Patton, A. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, *160*(1), 246–256.

PWC. (2017). *Asset & Wealth Management Insights - Asset Management 2020: Taking stock.* Retrieved from `https://www.pwc.com/gx/en/asset-management/asset-management-insights/assets/am-insights-june-2017.pdf`

Schwarz, G. (1978). Estimating the Dimension of a Model. *The Annals of Statistics*, *6*(2), 461–464.

Schwert, G. (1989). Why Does Stock Market Volatility Change Over Time? *The Journal of Finance*, *44*(5), 1115–1153.

Stock, J. H., & Watson, M. W. (1989). New indexes of coincident and leading economic indicators. *NBER macroeconomics annual*, *4*, 351–394.

Stărică, C., Herzel, S., & Nord, T. (2005). *Why does the GARCH(1,1) model fail to provide sensible longer-horizon volatility forecasts? .* (Working Paper, Chalmers University of Technology, Göteburg)

Taleb, N. (2012). *Antifragile.* Random House,US.

Wales, D., & Doye, J. (1997). Global Optimization by Basin-Hopping and the Lowest Energy Structures of Lennard-Jones Clusters Containing up to 110 Atoms. *The Journal of Physical Chemistry A*, *101*(28), 5111–5116.

Wei, Y., Liu, J., Lai, X., & Hu, Y. (2017). Which determinant is the most informative in forecasting crude oil market volatility: Fundamental, speculation, or uncertainty? *Energy Economics*, *68*, 141–150.

Wei, Y., Yu, Q., Liu, J., & Cao, Y. (2017). Hot money and China's stock market volatility: Further evidence using the GARCH–MIDAS model. *Physica A: Statistical Mechanics and it's Applications*, *492*, 923–930.

A Code Appendix

The code used to derive the results of this thesis has been attached as a zip file. It can be broken down into the following files:

1. `basin_settings` includes the settings used in the basinhopping optimisation algorithm
2. `thesis_settings` includes listings of the output and input files used for the code
3. `latexExport` includes functions to export tables to .tex format, it is used in the data section
4. `likelihood_functions` contains the log likelihood function of the single-regime models
5. `data` includes functions to process the raw data of the returns, macroeconomic variables, and realised variance
6. `result_analysis` functions used in the analysis of the in-sample results, including testing and identification
7. `forecasting` functions used to both forecast and analyse forecasts for the single-regime and two-regime models
8. `GM_estimation` implementation of the basinhopping optimisation for the single-regime case
9. `GM_analysis` implementation of the analysis for the single-regime case (i.e. testing)
10. `GM_forecasting` implementation of the forecasting process of the single-regime model
11. `ema_actual` implements the basinhopping optimisation approach for the switching model
12. `ema_analysis` implements the in-sample analysis and testing of the swtiching model
13. `ema_forecasting` implements the forecasting and forecast analysis of the switching model
14. `RV_fit_comparison` implements the in-sample comparison of the single regime and the switching model