Erasmus University Rotterdam

Erasmus School of Ecnonomics

Bachelor Thesis Econometrics and Operations Research

Rutger Buijs

433258

Conditional Volatility Time-Series Efficient Factors

Supervisor: dr. X. Xiao

Second assessor: dr. M.D. Zaharieva

July 5, 2020

Abstract

Factors showcase positive serial correlation which can exploited by timing past returns resulting in 'time-series efficient factors' according to Ehsani and Linnainmaa (2020). However, factors showcase heteroskedasticity as well, this research investigates whether conditioning on volatility next to past returns can benefit the Sharpe Ratio for a mean-variance investor. It presents great, statistically significant, improvements in terms of Sharpe Ratios for two of the five considered factors and shows promising results for the other factors. The framework of 'conditional volatility time-series efficient factors' can decrease factor risk considerably if modeled correctly.

Contents

1 Introduction **3**

2 Data **4**

2.1 Autocorrelation 5

2.2 Heteroskedasticity 7

3 Methodology **9**

3.1 Time-Series Efficient Factors 9

3.1.1 Efficient Factors with one month lagged signal 11

3.1.2 Efficient Factor with one year lagged signal 11

3.2 Conditional Volatility Time-Series Efficient Factors 13

3.2.1 Conditional Volatility Weights 14

3.3 Model Performance 15

4 Results **16**

4.1 Results of Time-Series Efficient Factors 17

4.2 Results of Conditional Volatility Time-Series Efficient Factors 18

4.2.1 What drives our results? 21

5 Conclusion **22**

A Figures **23**

A.1 Autocorrelation Functions 23

A.2 Autocorrelation Functions to assess heteroskedasticity 28

A.3 Delta Weight Figures 33

1 Introduction

This paper focuses on the 5 factor (5FF) Model from Fama and French (2015), whose factors are made time-series efficient following the approach of Ehsani and Linnainmaa (2020). The 5FF model explains stock returns by means of factors such as the original CAPM (Capital Asset Pricing Model) market premium variable from Sharpe (1964) and four on anomalies based factors. The factors considered in this research are; RMRF(value weighted Market portfolio returns minus Risk Free portfolio returns), SMB (Small Minus Big stock size), HML (High Minus Low book to market ratio), RMW (Robust Minus Weak profitability) and CMA (Conservative Minus Aggressive investment strategy). The 5FF model from Fama and French (2015) originated from the original three factor (3FF) model by by Fama and French (1992), which presented compelling evidence that the unconditional CAPM model did not account for portfolio returns sorted by size and book to market ratio. Ever since the publication of the 3FF model, various papers researched how to efficiently compose cross-sectional asset pricing models. Fama and French added the factors RMW and CMA to the 3FF model resulting in the 5FF model, but other researches have different views on which factors to incorporate in a cross-sectional asset pricing model to obtain a more efficient asset pricing model, as also reflected in Fama and French (2016).

Around the same time the 3FF model was introduced Fama and French (1992), Jegadeesh and Titman (1993) presented convincing results on the effectiveness of momentum in asset pricing, where they substantiated that the positive autocorrelation in asset returns results from market under-reaction on available information. The results from Jegadeesh and Titman (1993) initiated a wider research into the momentum and autocorrelation of asset returns, with Carhart (1997) introducing an extension on the 3FF model with the factor UMD (Up Minus Down momentum), accounting for persistence in asset returns. Years of research into cross-sectional factor models resulted in a great understanding of the effectiveness of various factors and anomalies to explain asset returns. Novy-Marx and Velikov (2016) considers twenty-three of the most popular asset pricing anomalies, among which different sorts of momentum factors, with a focus on creating a taxonomy of anomalies in the cross-section of expected asset returns. Alternative to the cross-sectional approach on momentum (e.g. the UMD factor), Gupta and Kelly (2019) investigates time-series factor momentum based on the factors itself instead of returns. Gupta and Kelly (2019) substantiated this approach by showing that for 49 out of 65 widely used factors the autocorrelations are significantly positive. In line with the findings from Gupta and Kelly (2019), Ehsani and Linnainmaa (2019) find as well that factor momentum explains the standard momentum as in Jegadeesh and Titman (1993), but as well as specific momentum factors such as: industry (adjusted) momentum, intermediate momentum and Sharpe momentum. They indicate that stock momentum builds upon factor momentum and with this insight created a time-series efficient 5FF model in Ehsani and Linnainmaa (2020) following the framework of Ferson and Siegel (2001) to condition on lagged factor returns. A few year earlier Moreira and Muir (2017) investigated 'time-series efficient factors' with a different on conditioning information. Where Ehsani and Linnainmaa (2020) condition on past factor returns, Moreira and Muir (2017) condition on past factor volatility. This paper tries to replicate the results obtained in Ehsani and Linnainmaa (2020) and further tries to improve the ability to model portfolio returns by researching whether extending the framework from Ehsani and Linnainmaa (2020) with an adapted framework based on Moreira and Muir (2017). With as central question within this research, can we improve the time-series efficient factor framework from Ehsani and Linnainmaa (2020) by incorporating conditional volatility, creating conditionall volatility time-series efficient factors'? We structure this paper by first covering the data properties in section 2, in which we assess presence of serial correaltion and heteroskedasticity. Section 3 provides the framework of the (conditional volatility) time-series efficient factors and elaborates on assessing model performance. Section 4 summarizes and tries to explain the obtained results for the considered factor models. Section 5 concludes on the findings of this research and states shortcomings that might initiate further research.

2 Data

The data used in this research is similar to that of Ehsani and Linnainmaa (2020) and is obtained from the French data library[1]. The factor data consists of the market returns R_M, the risk free returns R_F and the 6 value weighted portfolios; $SMB_{B/M}$, SMB_{RMW}, SMB_{CMA}, HML, CMW and RMW, where the SMB factor is constructed taking the average of the three SMB_x value weighted portfolios. The data set includes all stock value data from the NYSE, AMEX and NASDAQ listed firms from July 1963 to December 2018, with portfolio-restructuring on a yearly base at the end of June. To create the factors we take the 30th and 70th percentile as breakpoints of the data, except for SMB where the NYSE median is considered as breakpoint. For example, the factor HML is constructed by taking the average returns of the companies that are: with equal weights; Small with High-B/M ratio and Large with high-B/M ratio, subtracting the average returns of companies that are, again with equal weights: Small with Low-B/M ratio and Large with Low-B/M ratio. The same concept holds for the other factors, except for SMB

[1] The data library from French can be accessed with the following link: https://mba.tuck.dartmouth.edu/pages/faculty/ken.french

which is created by taking the average returns of 9 portfolios containing small firms minus that of 9 portfolios containing big firms. Basic statistical properties of the factors including the Sharpe Ratio, which is calculated by the average return divided by the standard deviation, are stated in table 1.

	Factor				
	MKTRF	SMB	HML	RMW	CMA
Mean	6.15	2.86	3.88	3.09	3.40
Standard Deviation	15.21	10.47	9.70	7.52	6.92
Sharpe Ratio	0.40	0.27	0.40	0.41	0.49
t-value	3.01	2.03	2.98	3.06	3.66

Table 1: General factor statistics
This table contains the annualized mean, standard deviation, Sharpe Ratio and corresponding t-value for each of the five factors considered. Factor data considered is from July 1963 to December 2018.

2.1 Autocorrelation

In order to improve the individual factors from the 5 factor model by means of timing them, the factors need to be significantly autocorrelated. Significance of the autocorrelation is assessed by means of the Ljung-Box Q-statistics from Ljung and Box (1978). In appendix section A.1 figures of the Autocorrelation Functions (ACF) for the different factors up to lag 12 are provided, including significance bounds set by the Ljung-Box Q statistic. As can be observed, serial correlation is present for all five factors, substantiating the approach of timing the factors.

Next to the more formal Ljung-Box Q-statistics, to assess the autocorrelation property of the factor returns, we also assess this property with a more empirical approach. In table 2 we separate the factor returns for each factor into two groups, high and low, based on the median. We do this for the monthly returns, left-side of the table, and for the average one year returns, right side of the table. We determine the potential of forecasting distinct factor returns when we allocate the prior (average) return into the high or low group. In table 2 we observe that for all the factors except for the excess market return factor, MKTRF, the difference in expected returns (H-L) based on the prior month return is significantly different from zero at a significance level of 0.01, which corresponds to a critical t-statistic of approximately 2.7. For the difference in expected returns based on the allocated prior one-year returns, we see much less convincing evidence of the potential of explaining factor returns. By conditioning on the prior one-year return we obtain for RMW an expected factor return of 0.42 and 0.11 for respectively the high and low group, the difference of 0.31 corresponds to a t-statistic of 1.82, which is the largest among the factors, however not significant at the 0.05 level (critical t-statistic is approximately 2.0). The findings in differences in expected factor returns are in line with the AR(1) coefficients ρ and their t-statistics stated in table 2, except for the difference in expected factor returns compared to the relatively low ρ coefficient for SMB conditioning on the prior month. This finding for SMB conditioning on the prior month could be explained by the significant lag at twelve months, see appendix A.1, which is not well covered by the AR(1) model conditioning on the past month. However, as could be observed in the Autocorrelation Functions (ACF) in appendix A.1, we conclude that an AR(1) model conditioning on the one month lag in general copes superior to the AR(1) model conditioning on the one year average lag, considering of how well it copes with the autocorrelation present in the data. In regard of the estimated slope coefficients ρ we see that for HML, RMW and CMA a relatively high AR(1) coefficient is estimated when we condition on the past month return, compared to all the other coefficients. This implies that the performance could benefit greatly from timing HML, RMW and CMA by signalling on their past month return, but in order to assess this we need to consider the co-movement in volatility as well.

Factor	Prior one month return				Prior one-year return			
	Low	High	H-L	$\rho_{(1,1)}$	Low	High	H-L	$\rho_{(12,1)}$
MKTRF	0.33	0.69	0.36	0.07	0.48	0.52	0.05	0.02
	(1.19)	(3.60)	(1.05)	(1.85)	(1.79)	(2.36)	(0.14)	(0.55)
SMB	-0.14	0.62	0.75	0.06	-0.08	0.58	0.66	0.08
	(-0.83)	(3.72)	(3.22)	(1.51)	(-0.47)	(3.33)	(2.77)	(2.26)
HML	-0.06	0.71	0.77	0.17	0.11	0.51	0.40	0.07
	(-0.42)	(4.57)	(3.60)	(4.43)	(0.68)	(3.45)	(1.81)	(2.16)
RMW	-0.11	0.62	0.73	0.17	0.11	0.42	0.31	0.07
	(-0.89)	(5.54)	(4.39)	(4.33)	(0.82)	(3.85)	(1.82)	(2.33)
CMA	-0.01	0.58	0.59	0.12	0.14	0.43	0.28	0.07
	(-0.07)	(5.02)	(3.82)	(3.21)	(1.34)	(3.68)	(1.82)	(2.09)
Average	-0.02	0.64	0.66		0.10	0.51	0.41	
	(-0.23)	(8.29)	(5.51)		(1.06)	(6.12)	(3.33)	

Table 2: Predictability of factor returns by allocating prior return into high or low return group This table reports expected factor returns (with t-values between parentheses), conditioning on the allocation of the; prior month return (left part) and the prior one-year average return (right part), into a high or low group. The division of the groups per factor is based on the median of the full sample. The lower row, denoted as Average, is formed by means of taking the average return of the factors allocated in the high/low group for each month and consecutively averaging the high/low based returns per factor over the full sample. For the Average row, we exclude the months in which none of the factors get allocated to the particular group. Stated in the table next to the expected factor returns and their differences, are the autocorrelation coefficients ρ of the factor returns is stated, with factor return correlations between respectively month t and $t-1$ and between t and $t-1$ to $t-12$. The data considered consists of factor returns from July 1963 through December 2018.

2.2 Heteroskedasticity

In order to be able to increase portfolio performance by timing factors we require the factor volatility not to move proportionally to the movements of expected factor returns as in 2. We require this since it is essential for the SR to improve, as it is a risk corrected performance measure. The assessment of this property of factor returns is closely related to the approach of assessing autocorrelation in the previous subsection 2.1. In table 3 we again separate the returns for each factor by the median into a high and low group. On the left side of the table we base this classification on monthly returns and on the right table side we base this on average past year returns. We determine the ability of the high and low classified previous factor returns to forecast distinct variances. In table 3 we observe that for the factor returns only CMA obtains a significant F-statistic (at a 5% significance level) that could harm the potential of timing factors. The F-statistic of MKTRF is also significant, however, this is significant in our advantage, since higher returns imply lower future variance, from which the SR could benefit. The same can be observed for RMW, however RMW is only significant in the case where we signal on the average one year lagged returns. For SMB and HML signalling on one the month lagged return we see high/low variance ratios of respectively 1.01 and 1.10, implying that lagged returns are not able to forecast distinct variances as they are insignificant and close to 1. This poses interesting opportunities when timing these factors considering their positive autocorrelations and return forecasting abilities as assessed in 2.1. Combining the insights from table 3 and table 2 implies that for factors such as MKTRF, RMW when the lagged factor return has been high, the next return is not only expected to have high returns, but as well as a relatively low risk. From this insight, and the relatively strong return forecasting ability of RMW we expect this factor to benefit greatly from timing it.

In order to benefit from timing factors we need the factor's variances not to move in line with the factor's returns, but Ehsani and Linnainmaa (2020) also make the assumption of homoskedasticity in factor returns. The empirical results in table 3 already foresees that for some of the factors homoskedasticity is not likely to be present. We formally assess whether the homoskedasticity assumption is justified by considering the Ljung-Box Q statistic of the squared residuals, where we obtain residuals by subtracting the factor mean of the factor returns. Plots of the Autocorrelation Functions (ACF) of the squared residuals can be found in appendix section A.2. Incorporated in the plots are as well the ACF's of the squared residuals when estimating an AR(1) and an AR(1)-GARCH(1,1) model. We can observe that for each factor the assumption of homoskedasticity does not seem to hold, nor does it for the factors estimated with an AR(1) model. However, by extending the model to an AR(1)-GARCH(1,1) model, we observe that the GARCH component copes with the present heteroskedasticity very well. One might argue, that the AR(1)-GARCH(1,1) model still does not sufficiently tackle the present heteroskedasticity for each factors as some lags are still slightly outside of the 95% confidence interval. For example the seventh and twelfth lag for SMB, the fifth lag of HML and second lag of RMW are still slightly significant. However, throughout our further research we assume the AR(1)-GARCH(1,1) model to produce homoskedastic error terms. The ability of the GARCH component to cope with heteroskedasticity also provides us with motivation to incorporate GARCH(1,1) obtained conditional volatility into the framework of time-series efficient factors.

Factor	Prior one month return			Prior one-year return		
	Low	High	F-stat $\frac{high}{low}$	Low	High	F-stat $(\frac{high}{low})$
MKTRF	26.2	12.37	0.47 (0.00)	23.09	16.1	0.7 (0.00)
SMB	8.98	9.06	1.01 (0.94)	8.44	9.92	1.18 (0.14)
HML	7.34	8.08	1.1 (0.39)	8.69	7.15	0.82 (0.08)
RMW	5.01	4.19	0.84 (0.10)	5.65	3.88	0.69 (0.00)
CMA	3.41	4.42	1.3 (0.02)	3.63	4.4	1.21 (0.08)
Average	5.52	3.89	0.71 (0.00)	5.4	4.43	0.82 (0.01)

Table 3: Predictability of factor variance by allocating prior return into high- or low return group
This table reports expected factor variance, conditioning on the allocation of the prior month return (left part) and the prior one-year average return (right part), into a high or low group. The division of the groups per factor is based on the median of the full sample. The lower row, denoted as Average, is formed by means of taking the average returns of the factors allocated in the high/low group for each month and consecutively take the variance of the high and low based returns over the full sample. For the Average row, we exclude the months in which none of the factors get allocated to the particular group. Stated in the table next to the expected factor variances for the high and low group is the variance ratio, high divided by low, with corresponding F-statistic values within parentheses. The data considered consists of factor returns from July 1963 through December 2018.

3 Methodology

3.1 Time-Series Efficient Factors

In this paper we consider the factors used in the 5FF framework from Fama and French (2015) as stated in section 1. We consider these factors independently and investigate the possibility to improve their Sharpe Ratios for a Markowitz (1959) mean-variance investor, which means, we try to maximize the ratio of expected returns to their volatility/risk. In order to obtain time series efficient factors we make use of the framework from Ferson and Siegel (2001) who use return forecasting signals to find the unique mean-variance efficient portfolio. As presented in Gupta and Kelly (2019) most factors are significantly serially correlated, which is in line with our findings for the five factors we consider in this paper, in section 2.1. This motivates to model time-series efficient factors, we do this according to the approach of Ehsani and Linnainmaa

(2020) with a simple AR(1) model.

$$R_t = \mu + \rho R_{t-1} + \epsilon_t \tag{1}$$

Where R_t denotes a particular factors' return at time t with corresponding unconditional mean and variance of $\frac{\mu}{1-\rho}$ and $\frac{\sigma_\epsilon^2}{1-\rho^2}$, ρ denotes the estimated AR(1) coefficient. To create time series efficient factors we thus condition on one lag of the factor return, following Ferson and Siegel (2001). The time-series efficient factor respective to the original factor is denoted below.

$$\hat{R}_t = X(\tilde{S}_t)R_t \tag{2}$$

Here $\hat{R}_t$ denotes the time-series Efficient factor at time t and R_t the original factor, $X(\tilde{S}_t)$ is the weight that transforms the factor conditioning on the return signal S_t. Obtained from Ehsani and Linnainmaa (2020) based on Ferson and Siegel (2001), the weight function $X(\tilde{S}_t)$ is denoted as:

$$X(\tilde{S}_t) = \frac{\mu_p}{\zeta} \cdot \frac{\mu(\tilde{S}_t)}{\mu^2(\tilde{S}_t) + \sigma_\epsilon^2(\tilde{S}_t)} \tag{3}$$

Here μ_p denotes the unconditional expected factor returns obtained from the original factor. The conditional expected portfolio returns $\mu(S_t)$, assuming an AR(1) model is used to condition the time-series Efficient factor on, and the constant ζ are defined below.

$$\begin{aligned} \zeta &= \mathbb{E}\left(\frac{\mu^2(\tilde{S_t})}{\mu^2(\tilde{S_t}) * \sigma_\epsilon^2(\tilde{S_t})}\right) \approx \frac{SR^2 + \rho^2}{SR^2 + 1} \\ \mu(\tilde{S}_t) &= \mu_p(1-\rho) + \rho R_{t-1} \end{aligned} \tag{4}$$

For a derivation of the constant ζ we refer to Ferson and Siegel (2001) who state that the constant could be interpreted by its relation to the slopes of both the conditional and unconditional mean-variance frontiers. The conditional mean conditions on the signal and used the slope coefficient, ρ, from the regression of the original factor returns on the signal.

By using the constant ζ and the conditional mean $\mu(\tilde{S}_t)$ from equation (4) for the weight function (3), we can obtain the weights $X(\tilde{S}_t)$. By weighting the original factor by $X(\tilde{S}_t)$ as in equation (2), the effect of the signal (for example one lagged return) is incorporated in the time-series efficient factor. In order to stay close to the original factor we bound the weights between $[0, 1]$, as so we do not allow to leverage or short the original factor. The next two subsections provide specific denotations of the statistics used for different signals.

3.1.1 Efficient Factors with one month lagged signal

In the case time-series Efficient factors are created using the one month lagged return R_{t-1} as signal, we obtain a straightforward specification of the model. In order to obtain the model specification we need to specify the slope coefficient of the factor returns on the signal, like an AR specification and their correlation coefficient.

$$\begin{aligned} \beta &= \frac{cov(R_t, \tilde{S}_t)}{var(\tilde{S}_t)} \\ \rho &= \frac{cov(R_t, \tilde{S}_t)}{\sqrt{var(\tilde{S}_t)} * \sqrt{var(R_t)}} \end{aligned} \tag{5}$$

As the signal's variance equals the variance of the factor's returns, by assuming homoskedasticity, the slope coefficient β equals the correlation coefficient, ρ. This results in the following specification of ζ and $\mu(\tilde{S}_t)$ obtained from Ehsani and Linnainmaa (2020), where only the AR(1) slope coefficient, which we now denote as ρ, is required.

$$\begin{aligned} \zeta &= \frac{SR^2 + \rho^2}{SR^2 + 1} \\ \mu(\tilde{S}_t) &= \mu_p(1-\rho) + \rho R_{t-1} \\ X(\tilde{S}_t) &= \frac{\mu_p}{\zeta} \cdot \frac{\mu(\tilde{S}_t)}{\mu^2(\tilde{S}_t) + \sigma_\epsilon^2} \end{aligned} \tag{6}$$

The constant ζ is calculated by means of unconditional Sharpe Ratio of the original factor, which is calculated by $SR = \frac{\mu_p}{\sigma_p}$ with σ_p being the unconditional standard deviation of the original factor. The obtained ζ and $\mu(\tilde{S}_t)$ can be plugged into the weight function (3), for which we take the unconditional variance σ_ϵ^2 as $\sigma_\epsilon^2(\tilde{S}_t)$ by assuming homoskedasticity. With the weights $X(\tilde{S}_t)$ one could then form the time-series Efficient factors following equation (2) with as signal the one month lagged factor return. The resulting factor optimally times the factor given its previous month returns and forms an unique frontier according to Ferson and Siegel (2001).

3.1.2 Efficient Factor with one year lagged signal

Next to time-series Efficient factors conditioning on their one month lagged return, we also condition on their one year lagged returns, in order to maintain an AR(1) model we take the average of these monthly returns and use that as the signal.

$$\tilde{S_{t,n}} = \frac{1}{n}\sum_{i=1}^{n} R_{t-i} = R_{t-1,t-n} \tag{7}$$

In equation (7) n denotes the amount of lagged months incorporated in the signal, which equals 12 as we use the average of one year lagged factor returns. Changing the signal from a one month lagged return to the average of the one year lagged returns changes some specifications in the model. This is due to the moments of the signal, which are different for the two signal specifications. By estimating an AR(1) model of the signal on the original factor returns, the slope coefficient β and correlation coefficient ρ from equation (5), equal each other in case the variance of the signal equals the variance of the factor returns. This is the case for a one month lag as signal, but not for the average one year lagged returns as can be seen below.

$$
\begin{aligned}
var(\tilde{S_{t,n}}) &= \frac{\sigma_\epsilon^2}{n(1-\rho^2)}\left(1+\frac{2}{n}\sum_{k=1}^{n}(n-k)\rho^k\right) \\
var(R_t) &= \frac{\sigma_\epsilon^2}{1-\rho^2}
\end{aligned}
\tag{8}
$$

Here ρ denotes the original correlation coefficient from the one month lagged factor returns on the current factor returns. Equation 8 results in the following specification of equation (5).

$$
\begin{aligned}
\beta_n &= \frac{\sum_{k=1}^{n}\rho^k}{1+\frac{2}{n}\sum_{k=1}^{n}(n-k)\rho^k} \\
\rho_n &= \frac{\frac{1}{\sqrt{n}}\sum_{k=1}^{n}\rho^k}{\sqrt{1+\frac{2}{n}\sum_{k=1}^{n}(n-k)\rho^k}}
\end{aligned}
\tag{9}
$$

Here β_n and ρ_n respectively denote the slope coefficient of an AR(1) regression and the correlation coefficient between the average of n lagged months of factor returns on the current factor return. These specification changes subsequently result in the following generalisation of the specification in (6) which exclusively signals on one month lagged return.

$$
\begin{aligned}
\zeta &= \frac{SR^2+\rho_n^2}{SR^2+1} \\
\mu(\tilde{S_{t,n}}) &= \mu_p(1-\rho)+\beta_n R_{t-n,t-1} \\
X(\tilde{S_{t,n}}) &= \frac{\mu_p}{\zeta}\cdot\frac{\mu(\tilde{S_{t,n}})}{\mu^2(\tilde{S_{t,n}})+\sigma_\epsilon^2}
\end{aligned}
\tag{10}
$$

Where the resulting weights $X(\tilde{S_{t,n}})$ can be obtained by plugging the constant ζ and conditional mean $\mu(\tilde{S_{t,n}})$ in (3), where again we take the unconditional variance σ_ϵ^2 as $\sigma_\epsilon^2(\tilde{S_{t,n}})$ in equation (3) by assuming homoskedasticity. The obtained weights can be used to obtain the time-series Efficient factors as in equation (2), with the average of the one year lagged monthly factor returns as signal. For a more detailed derivation of the model specifications covered in this section we refer to appendix A.3 of Ehsani and Linnainmaa (2020).

3.2 Conditional Volatility Time-Series Efficient Factors

The time-series efficient factors we obtain following the approach described in the previous subsections are able to benefit from the present autocorrelation as can be observed in A.1. An important assumption in our approach to obtain time-series efficient factors is that we assume homoskedasticity, for example equation (5) from subsection 3.1.1 equates the correlation coefficient ρ and slope coefficient β. By assuming homoskedasticity we also use the unconditional variance σ_ϵ^2 for $\sigma_\epsilon^2(\tilde{S_{t,n}})$ in the weight function (3). However, as we conclude in subsection 2.2 there is strong evidence of presence of heteroskedasticity in the factor returns. By timing the factors we can benefit from the present autocorrelation, we investigate whether we can benefit from the present heteroskedasticity as well, by introducing conditional variance into the weighting function.

Well acclaimed models to account for time varying volatility are the ARCH and GARCH models from Engle (1982) and Bollerslev (1986). We implement a GARCH(1,1) model in our approach resulting in the AR(1)-GARCH(1,1) specification as formulated below.

$$\begin{aligned} R_t &= \mu + \rho R_{t-1} + \epsilon_t \\ \sigma_\epsilon^2(\tilde{S}_t) &= \omega + \alpha\epsilon_{t-1}^2 + \gamma\sigma_\epsilon^2(\tilde{S}_{t-1}) \\ \textit{where } \epsilon_t \mid \tilde{S}_t &\sim \mathcal{N}(0, \sigma_\epsilon^2(\tilde{S}_t)) \end{aligned} \tag{11}$$

Here we make the assumption that the resulting errors ϵ_t are normally i.i.d. distributed. To assess whether the assumption of normality of the residuals seems plausible we will perform the Jarque-Bera test from Jarque and Bera (1980). We estimate the GARCH(1,1) models using Matlab's 'Econometric Modeler' app. We use the conditional variance $\sigma_\epsilon^2(\tilde{S}_t)$, which we obtain from our AR-GARCH specification, in two different models to determine weights. The first approach uses the conditional variance obtained from the AR-GARCH model in weight function (3), which differs from the previous specifications where the unconditional variance is used. By introducing conditional variance into the weight function we drop the assumption of homoskedasticity, as contradicted in our findings stated in section 2.2.

3.2.1 Conditional Volatility Weights

Motivated by our findings of the present heteroskedasticity in the factor returns and the ability of a GARCH model to cope with this, we extend the weighting function of Ehsani and Linnainmaa (2020). We create a new weighting function based on Moreira and Muir (2017), who pose an alternative approach of timing factors, by timing their variance. The weighting function times variance such that risk exposure will be decreased when recent volatility has been high. Despite that their main focus lies within modelling weights conditioning on variance, they are able to increase factor's SR's by timing the factor's volatility as formulated in the equation below:

$$\hat{R}_t = \frac{c}{\sigma_\epsilon(\tilde{S}_t)} R_t \tag{12}$$

Here $\frac{c}{\sigma^2(\tilde{S}_t)}$ could be interpreted as a weight function conditioning on the conditional volatility, with c denoting a particular constant accounting for the average exposure.
We introduce conditional volatility in the framework of Ehsani and Linnainmaa (2020) by multiplying the original weight function (3) by the weight function of Moreira and Muir (2017). We implement the Moreira-Muir weights in two versions of the Ehsani-Linnainmaa weights, one with conditional variance already included and one still with the unconditional variance, resulting in two different updated weight functions.

$$X^*(\tilde{S}_t) = \frac{c}{\sigma_\epsilon(\tilde{S}_t)} \cdot \frac{\mu_p}{\zeta} \cdot \frac{\mu(\tilde{S_{t,n}})}{\mu^2(\tilde{S_{t,n}}) + \sigma_\epsilon^2} \tag{13}$$

$$X^{**}(\tilde{S}_t) = \frac{c}{\sigma_\epsilon(\tilde{S}_t)} \cdot \frac{\mu_p}{\zeta} \cdot \frac{\mu(\tilde{S_{t,n}})}{\mu^2(\tilde{S_{t,n}}) + \sigma_\epsilon^2(\tilde{S}_t)} \tag{14}$$

The above weight functions combine the approaches of Ehsani and Linnainmaa (2020) and Moreira and Muir (2017). As we implement the conditional volatility weight into our established framework we deviate slightly from the specification of Moreira and Muir (2017). We base the weights on volatility's, rather than variances and we calculate the constant c with the median of the conditional volatility e.g. we let the constant c equal the median of the obtained volatility estimations and divide that by the estimated volatility. Compared to Moreira and Muir who take the constant c such that the resulting efficient factor has the same unconditional variance as the original factor. With our specification of the constant c we ensure that the for half of the observations the Moreira-Muir weights has a positive effect on the resulting weight and for the other half a negative effect. Moreira and Muir (2017) state that the constant c has no impact on the estimated SR's, however this assumes no leverage constraints are in place which is not the case in our framework where we restrict the weights to be between zero and one, for this reason the specification of c is important for the models' performances. With our given specification for the constant c and use of volatility in stead of variance we ensure that the weights do not react heavily on the conditional volatility which would result in a rather binary weight function. The incorporation of the conditional volatility weight could be interpreted as scaling the Ehsani-Linnainmaa weight which exploits autocorrelation, with the Moreira-Muir weight which exploits heteroskedasticity. We will condition on volatility in two different ways. Motivated by Moreira and Muir (2017) we condition on the realized volatility, RV, for which we use daily factor data and is obtained as defined below.

$$\sigma(\tilde{S}_t) = RV^2 = \sum_{d=1/n}^{1} (R_{t+d} - \frac{\sum_{d=1/n}^{1} R_{t+d}}{n})^2 \tag{15}$$

Here n denotes the amount of datapoints in one month. Next to conditioning on Realized volatility we also use conditional volatility from the GARCH model.

3.3 Model Performance

As a Markowitz (1959) mean-variance investor, we want to have as minimal of a risk as possible for a certain given expected return. Derived by Ferson and Siegel (2001), the variance of a time-series efficient factor with a given expected return can be predicted by using moments of the original factor, as stated in the following equation.

$$\sigma_p^{2*} = \mu_p^2(\frac{1}{\zeta} - 1) \tag{16}$$

Since ζ is a constant and we require a certain expected return on our portfolio, μ_p, we can easily obtain the corresponding variance. Subsequently, this can be used to obtain the predicted time-series efficient factor's Sharpe ratio following the equation below.

$$SR^* = \frac{\mu_p}{\sigma_p^*} = \frac{1}{\sqrt{\frac{1}{\zeta} - 1}} \tag{17}$$

Here SR^* is the predicted time-series efficient factor's Sharpe Ratio, which should intuitively be at least equal to that of the original factor, since the weight $X(S_t)$ equalling 1 would imply the time-series efficient factor equalling the original factor. This is also derived by Ferson and Siegel

(2001) of which the result is stated below.

$$\frac{SR^*}{SR} = \sqrt{\frac{1 + (\frac{\rho}{SR})^2}{1 - \rho^2}} \geq 1 \qquad |\, \rho \,|< 1 \tag{18}$$

The ratio of the Sharpe factors equals at least 1 where the ratio can converge to 1 when the autocorrelation coefficient, ρ, approaches zero. This emphasizes the importance of the autocorrelation of the factor for expected improvements on the Sharpe ratio compared to the original factor. We see as well that for a portfolio that had originally a relatively low SR, we can expect more improvement than for a portfolio that already had a significantly large SR. Making use of equation (16) we can predict efficient factors.

To assess model performance we calculate z-statistic values exhibiting the significance of the achieved, or predicted, SR improvements. We examine the SR improvement's significance with the test from Jobson and Korkie (1981), corrected by Memmel (2003).

$$\begin{aligned} z &= \frac{\sigma_o\mu_e - \sigma_e\mu_o}{\sqrt{\theta}} \\ \theta &= \frac{1}{T}\left(2\sigma_e^2\sigma_o^2 - \sigma_e\sigma_o\sigma_{e,o} + \frac{1}{2}\sigma_o^2\mu_e^2 + \frac{1}{2}\sigma_e^2\mu_o^2 - \frac{\mu_e\mu_o}{\sigma_e\sigma_o}\sigma_{e,o}^2\right) \end{aligned} \tag{19}$$

Here θ denotes the variance of the SR where the subscripts 'e' and 'o' respectively denote time-series efficient and original, resulting in z denoting the z-statistic from which we can assess significance. We calculate the z-statistic for the realized efficient factors with the observed covariance, however, we can also estimate the significance of the predicted time-series efficient factor's SR. We then estimate the z-statistic with the theoretical covariance $\sigma_{e,o}^{2*}$ obtained from Ehsani and Linnainmaa (2020).

$$\sigma_{e,o}^{2*} = \frac{\mu_p(1 - \rho^2)}{SR^2 + \rho^2} \tag{20}$$

4 Results

This section evaluates the performance of the time-series efficient factors for the different signals and weight functions where we compare them to the original, not timed, factors. We do this by means of comparing them by their obtained Sharpe Ratios, which serves as a widely used criterion for a mean-variance investor who aims to minimize their risk for a given expected return. We focus first on the results of the time efficient factors for which we follow the established framework by Ehsani and Linnainmaa (2020). In the second subsection we will investigate the results obtained by the time-series efficient factors incorporating conditional volatility, where the frameworks of Ehsani and Linnainmaa (2020) and Moreira and Muir (2017) are combined.

4.1 Results of Time-Series Efficient Factors

In section 3 methodology, we cover the established framework by Ehsani and Linnainmaa (2020) where factors are timed on their past returns. We elaborate on the process of timing factors signalling on their one month lagged return in subsection 3.1.1 and do likewise for timing factors on their lagged one year average return in subsection 3.1.2. In this section we will assess the performance of the time-series efficient factors based on their SR's. Next to the actual time-series efficient factors, we also estimated predictions of improvements in terms of SR that would be obtained by timing the factors (signalling on one month lagged return) in subsection 3.3. In table 4 we can see, as expected and substantiated in equation (18), that for all the factors the realized efficient version improves the SR compared to the original factor. However we also incorporate z- and p-values on the significance of the the SR improvements according to (19), this poses a statistical affirmation on the significance of the improvements.

As we describe in section 2, we especially see for RMW that it has great potential of performance improvement by timing the factor, due to it's ability to forecast high returns and relatively low volatility when the past month's return has been high and that a low previous return forecasts a low return with higher volatility. As shown in subsection A.1 we obtain the highest AR(1) coefficients for HML, RMW and CMA equalling respectively 0.17, 0.17 and 0.12 providing us with the intuition that timing those factors on their past month return could greatly benefit the SR, this intuition is validated by looking at table 4. We observe that the improvements for the two factors HML and RMW with one month lag as signal are significant at a 5% significance level and for CMA at the 10% level. HML, RMW and CMA make up for our largest achieved improvements in respect of their SR's with respective improvements of 0.22, 0.24, 0.13. Notable as well is that for MKTRF and SMB the time-series efficient factors with a one-year lag have less improvement in their SR compared to the one month lagged version, which is not apparent in their p-values. However, comparing the achieved improvements of the factors signalling on the past month return to the factors signalling on the past year return we see that SR's in the former case benefit more from timing. The only factor that statistically benefits more significantly, however less in absolute terms, is SMB. Overall, we come to to the conclusion that past month returns perform superior to the past average year returns as signal to time the factors. This is in line with our preliminary data research, from which we see that the past average year as signal performs inferior compared to the one month lagged signal. It lacks in general, in terms of being able to benefit of the present autocorrelation shown by; having less significant AR(1) coefficients and by inferior estimates of returns 2. Hence, we decide to further assess the ability of improving the time-efficient factor framework signalling on the past month return.

Factor type	Statistic	Factor				
		MKTRF	SMB	HML	RMW	CMA
Original Factor	SR	0.40	0.27	0.40	0.41	0.49
Predicted Efficient Factor (Signal r_{t-1})	SR	0.48	0.34	0.73	0.72	0.66
	ΔSR	0.07	0.07	0.33	0.31	0.17
	z-value	0.96	0.79	2.52	2.46	1.69
	p-value	0.34	0.43	0.01	0.01	0.09
Realized Efficient Factor (Signal r_{t-1})	SR	0.53	0.38	0.62	0.65	0.62
	ΔSR	0.12	0.10	0.22	0.24	0.13
	z-value	1.61	1.38	2.28	2.24	1.68
	p-value	0.11	0.17	0.02	0.03	0.09
Realized Efficient Factor (Signal r_{t-12})	SR	0.44	0.33	0.48	0.54	0.51
	ΔSR	0.03	0.06	0.08	0.12	0.02
	z-value	1.55	1.77	1.61	1.46	0.96
	p-value	0.12	0.08	0.11	0.14	0.34

Table 4: Time-series efficient factor results based on the Ehsani-Linnainmaa framework
This table reports the Sharpe Ratios of the original factors, predicted time-efficient factors and realized efficient factors signalling on the one month lagged return and the one year average lagged return. ΔSR denotes the improvement of the factor's SR compared to the original factor's SR. Denoted below the ΔSR's are the corresponding z-value and p-value denoting the statistical significance of the improvements. The data considered consists of factor returns from July 1963 through December 2018.

4.2 Results of Conditional Volatility Time-Series Efficient Factors

In the data section we show that heteroskedasticity is present the factor data and show that future volatility is predictable by the GARCH(1,1) model. We show this predictability in the figures in appendix section A.2 and discuss the possibility of exploiting heteroskedasticity in section 2.2. Motivated by these findings and the research of Moreira and Muir (2017), who times factors on volatility, we estimate conditional volatility time-series efficient factors. Next to conditioning on GARCH(1,1) estimated volatility we also condition on realized volatility and variance, motivated by Moreira and Muir (2017) who primarily focus on timing factors with realized variance. The presented framework in section 3.2 combines the established frameworks of Ehsani and Linnainmaa (2020) and Moreira and Muir (2017) for which the results will be covered in this section. In table 5 we denoted the SR's for four different conditional volatility time-series efficient factors and also added the original factor and Ehsani-Linnainmaa AR(1) factor (which we will denote as EL(AR) from now on) for comparison reasons. The four conditional volatility time-series efficient models[2] we consider are the; EL(AR)MM(GARCH), EL(AR)MM(RV), EL(AR)MM(RV^2) and the EL(AR-GARCH) models all conditioning on past month data. In table 5 we see that the improvements vary significantly between the different factors compared to the EL(AR) factors. We observe that MKTRF and SMB do not benefit from timing volatility in our frameworks, however, for HML, RMW and CMA factors we are able to improve the SR's. The EL(AR)MM(GARCH) model, which we motivated primarily in our previous section, is able to improve the SR's in comparison to the EL(AR) model for HML and CMA, very slightly with respective SR improvements of 0.02 and 0.01. However, for RMW we report a SR of 0.78, corresponding to an improvement compared to the EL(AR)MM(RV) framework of 0.13 SR points which is statistically significant at the 5% level. The results on the EL(AR)MM(GARCH) model imply that timing volatility could benefit the time-series efficient factor greatly. By altering the conditional volatility model in the MM-weight from the GARCH specification to the Realized Volatility (RV) we obtain even better results. For all the considered models, we obtain the largest improvements following the EL(AR)MM(RV) framework with SR improvements (compared to the EL(AR) portfolios) of 0.16, 0.18 and 0.06 respectively for for HML, RMW and CMA. The improvements for HML and RMW are also statistically significant with corresponding p_2 values of respectively 0.02 and 0.03, indicating significant statistical improvement compared to the EL(AR) factors, following the test statistic of Jobson and Korkie (1981) in equation. Notably is the large difference in SR's between the EL(AR)MM(GARCH) and EL(AR)MM(RV) portfolios for HML of respectively 0.64 and 0.78, which seems not be in line with the relation between these frameworks for the other factor's SR's. We also assessed the EL(AR)MM(RV^2) portfolios conditioning on variance, which by being more sensitive to volatility movements, imposes more impact of MM-weight on the resulting weight. As we can see in table 5, this does not benefit the resulting SR's compared to it's counterpart EL(AR)MM(RV) conditioning on volatility. By implementing conditional volatility directly into the Ehsani-Linnainmaa framework as we do in the EL(AR-GARCH) model as specified in subsection 3.2 we obtain an improvement for HML in terms of SR compared to the EL(AR) portfolio. Nonetheless, implementing conditional volatility directly into the framework from Ehsani and Linnainmaa (2020), provides us with considerably worse results compared to other models where the framework is extended with the framework of Moreira and Muir (2017). Overall, we conclude that the EL(AR)MM(RV) model performs best among the conditional volatility time-series efficient factors.

[2]We denote the conditional volatility time-series efficient factors as EL(..)MM(..), where EL denotes the Ehsani Linnainmaa weight and MM denotes the Moreira-Muir weight. Between parentheses is stated which model(s) are used to obtain the signal of the particular weight.

Time-series efficient factor framework			Factor's Sharpe Ratio				
Ehsani-Linnainmaa Signal	Moreira-Muir Signal	Statistic	MKTRF	SMB	HML	RMW	CMA
		SR	0.40	0.27	0.40	0.41	0.49
AR(1)		SR	0.53	0.38	0.62	0.65	0.62
		μ_p	1.48	0.8	0.96	0.76	0.87
		σ_p	9.74	7.32	5.36	4.06	4.86
		p_1 value	0.11	0.17	0.02	0.03	0.09
AR(1)	GARCH(1,1) (volatility)	SR	0.52	0.35	0.64	0.78	0.63
		μ_p	1.39	0.70	0.74	0.57	0.78
		σ_p	9.27	7.02	3.97	2.53	4.26
		p_1 value	0.14	0.34	0.01	0.00	0.08
		p_2 value	1.22	1.87	0.55	0.05	0.76
AR(1)	RV	SR	0.52	0.35	0.78	0.82	0.68
		μ_p	1.34	0.70	0.91	0.6	0.76
		σ_p	8.84	6.98	4.04	2.53	3.86
		p_1 value	0.15	0.36	0.00	0.00	0.03
		p_2 value	1.04	1.60	0.02	0.03	0.27
AR(1)	RV^2	SR	0.47	0.28	0.77	0.79	0.63
		μ_p	1.13	0.54	0.88	0.56	0.63
		σ_p	8.33	6.78	3.98	2.45	3.45
		p_1 value	0.47	0.98	0.00	0.00	0.17
		p_2 value	1.59	1.90	0.17	0.23	0.91
AR(1)-GARCH(1,1)		SR	0.44	0.37	0.62	0.76	0.62
		μ_p	1.39	0.65	0.85	0.80	0.77
		σ_p	11.07	6.04	4.73	3.66	4.27
		p_1 value	0.61	0.25	0.02	0.00	0.12
		p_2 value	1.90	1.06	0.97	0.26	0.94

Table 5: Time-series efficient factor performance, with conditional volatility
This table contains the results of the time-series efficient factors in terms of SR and its components; portfolio mean and standard deviation. Below the SR and its components are the p_1 and p_2 values, respectively denoting the p-value on improving the original factor and the EL(AR1) factor. On the left side of the table we state the used framework of the times-series efficient factors. Every time-series efficient factors uses the weights from the Ehsani-Linnainmaa framework with distinct conditional variance/volatility structures in the Moreira-Muir framework, respectively: GARCH(1,1) volatility, Realized Volatility (RV) and Realized Variance (RV^2), but also the EL(AR-GARCH) model without MM weight. If no specification is given, the associated weight function is not used to obtain the factor. The first factor, as special case, denotes the original factor, as no signal specification is given. All signals are based on one month lagged factor data. The data considered consists of factor returns from July 1963 through December 2018.

4.2.1 What drives our results?

As stated in the previous subsection, we are able to improve the SR's of the HML, RMW, and CMA factors by conditioning on volatility. However, we can't obtain better results for MKTRF and SMB by conditioning on volatility. This section tries to point out what causes the improvements and for what reasons we do not obtain improvements. The reason for the EL(AR)MM(RV) model to outperform the EL(AR)MM(GARCH) model is unclear, however it might be due to the fact that Realized Volatility conditions directly on (on average) 21 datapoints while the GARCH only conditions two direct datapoints. Another plausible explanation for this result is that we are not able to confirm normality in the error terms in equation 13 for the GARCH(1,1) model according to the Jarque Bera test statistic. The only factor for which normality in the error terms could not be rejected at the 5% level is CMA with a corresponding p-value of 0.01. For this reason we focus primarily on the EL(AR) and the conditional volatility time-series efficient factor following EL(AR)MM(RV) as this framework provided us with the best results. In table 5 we stated the average returns and unconditional standard deviations for the factors. In this table we see that the EL(AR) model reports the highest average return for all factors, this indicates that improvements in SR are caused by relatively low standard deviations. For the factors HML, RMW and CMA in the EL(AR)MM(RV) framework we report a standard deviation of 4.04, 2.53 and 4.26 which corresponds to decreases of 25%, 38% and 21% while the returns only decrease with 5%, 21% and 13% compared to the EL(AR) obtained factors. In conclusion, despite the lower average returns, we are able to decrease the standard deviation such that for those factors the SR's increase considerably. For MKTRF and SMB the decrease in standard deviation does not weigh up against the decrease in returns, as we can not increase the SR. In appendix A.3 are graphs of the effect of extending the EL(AR) model to the EL(AR)MM(RV) model with ΔWeights of the two frameworks and above that we plot the conditional volatility according to the Realized Volatility. In these graphs we can see that around half of the time, the updated weights are higher and lower in the other half . However, the weights are somewhat higher in times stable factor returns but lower considerably in times of risk, this is due to our choice of the constant c in equation 13 for which we took the median (instead of for example the mean). Letting the constant c equal the mean would result in less sensitivity on high volatility spikes, since the mean is generally much higher due to the rather large upward volatility spikes compared to downward volatility 'spikes'. By taking the median we ensured that the model's weights would be much lower around times of high risk as can be seen around the years of 2001 and 2008 for example. We can not answer the question why MKTRF and SMB are not able to benefit from conditioning on volatility with certainty. However, in figure 11 we see much more volatility for MKTRF compared to the other factors as also stated in 1. Despite not formally tested, we think for SMB less clustering of volatility is present observable in figure 12, this would negatively affect the ability to model conditional volatility well. In the heteroskedasticity figures in appendix A.2 we see that the amount of heteroskedasticity that is present is significant for MKRTF and SMB, nonetheless it is less extreme as for the other factors, which could imply less possible benefit from conditioning on volatility. Another observation we make for both MKTRF and SMB is that they have smaller negative ΔWeights during times of stressed factor returns, this could substantiate the lower decrease in standard deviation reported in 5. Finally we, conclude that for the factors HML, RMW and CMA, one can attain a much lower unconditional volatility without harming the average returns too much when conditioning on volatility, resulting in great improvements of SR. To assess whether this can be achieved for the MKRTF and SMB factor too, more research into the modelling of conditional volatility and research into the optimal approach of extending Ehsani and Linnainmaa (2020) framework with the Moreira and Muir (2017) framework is required.

5 Conclusion

In this research we estimated time series efficient factors following the framework established by Ehsani and Linnainmaa (2020). Motivated by the presence of heteroskedasticity and the results from Moreira and Muir (2017), the main goal of this research is to investigate whether the time-series efficient factors can benefit from timing their volatility next to their past returns. We have managed to extend the framework from Ehsani and Linnainmaa with the adapted framework from Moreira and Muir, resulting in conditional volatility time-series efficient factors. In general, the EL(AR)MM(RV) framework performs best, which is able to improve SR's for the factors HML, RMW and CMA of which HML and RMW are statistically significant. The time-series efficient factors conditioning on volatility that we estimated have slightly lower average returns, however the significant decrease in volatility resulted in great SR improvements. This research has posed an innovative timing model by combining the insights from Ehsani Linnainmaa and of Moreira Muir. However, it does not thoroughly investigate the modelling of conditional volatility in time efficient framework, nor assessed the optimal calibration of the relation between the two weights. Hence, we believe that timing volatility could benefit the MKTRF and SMB factor, even though this paper is not able to present improvements in performance. We conclude that incorporating conditional volatility in the time-series efficient factor framework poses a promising opportunity for mean-variance investors.

A Figures

In this appendix section we provide various figures substantiating the research. In subsection A.1 we provide figures containing ACF's for the five considered factors to assess autocorrelation. In subsection A.2 we present figures of the ACF's on the squared residuals to assess heteroskedasticity. Subsection A.3 contains figures of the ΔWeights, EL(AR) weights minus EL(AR)MM(RV) weights, for the different factors. The figures also incorporate monthly realized volatility.

A.1 Autocorrelation Functions

In this subsection we provide figures containing ACF's for the five considered factors where we distinguish between the original present autocorrelation and the autocorrelation left after estimating an AR(1) model. From the resulting figures we can assess whether autocorrelation is present and whether an AR(1) model seems to cope with this data aspect well.

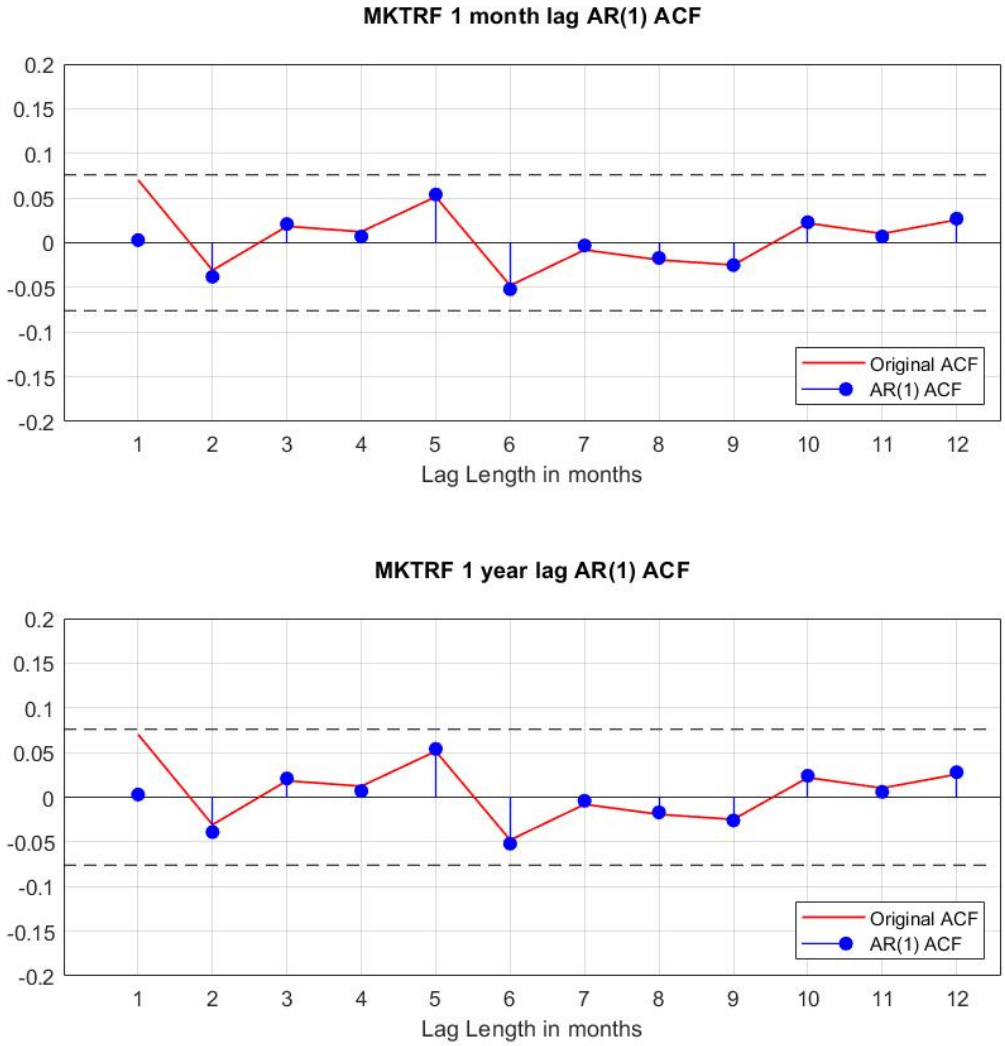

Figure 1: Autocorrelation Function for the original MKTRF factor and the AR(1) models with one month and one year average lag. The dotted horizontal line represents the Ljung-Box Q-statistic at a 5% significance level of no autocorrelation. The X-axis represents the amount of lags in months and the Y-axis represents the autocorrelation coefficient.

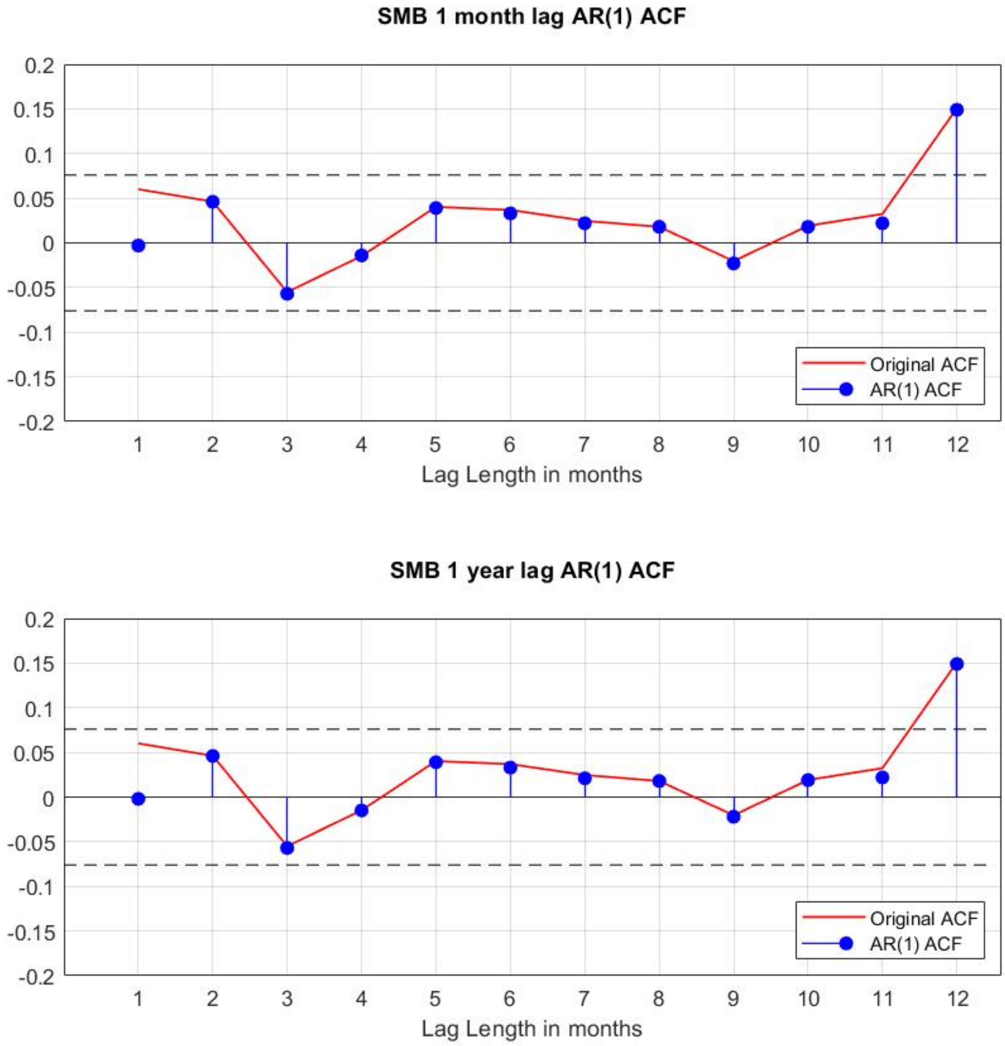

Figure 2: Autocorrelation Function for the original SMB factor and the AR(1) models with one month and one year average lag. The dotted horizontal line represents the Ljung-Box Q-statistic at a 5% significance level of no autocorrelation. The X-axis represents the amount of lags in months and the Y-axis represents the autocorrelation coefficient.

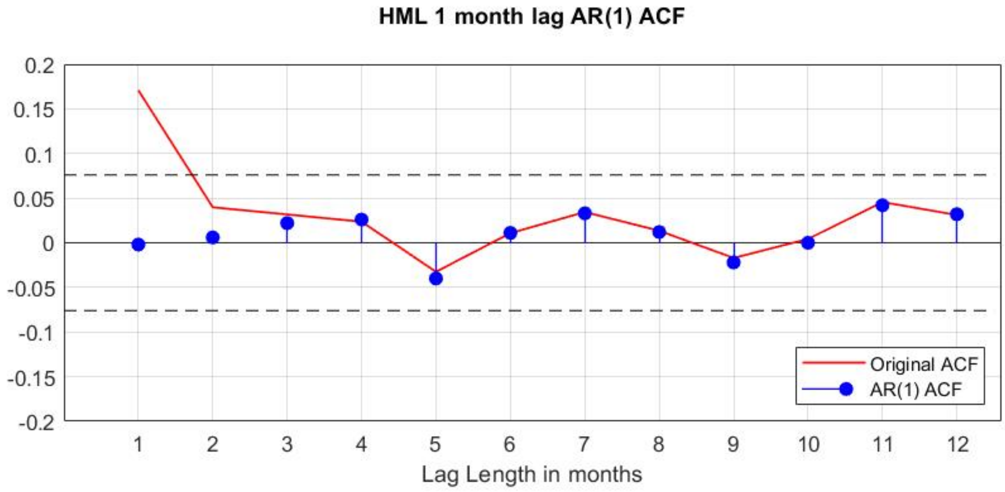

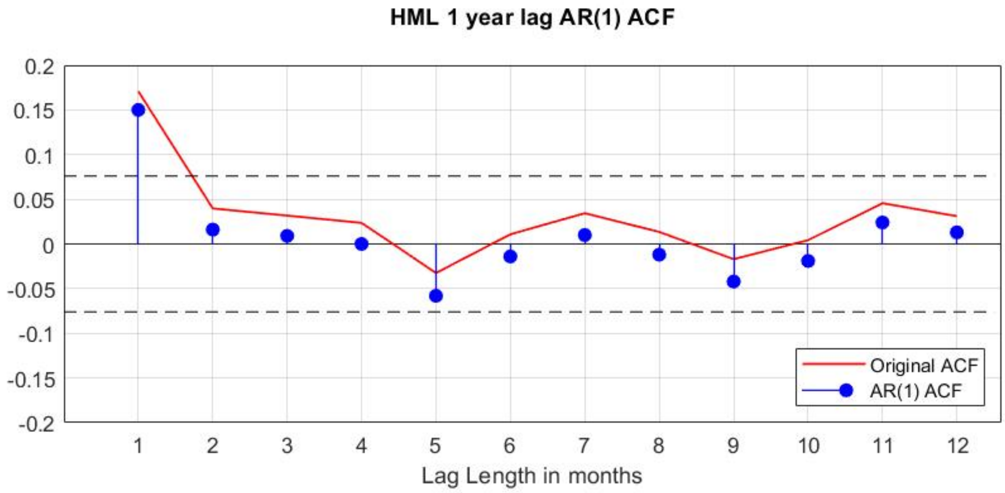

Figure 3: Autocorrelation Function for the original HML factor and the AR(1) models with one month and one year average lag. The dotted horizontal line represents the Ljung-Box Q-statistic at a 5% significance level of no autocorrelation. The X-axis represents the amount of lags in months and the Y-axis represents the autocorrelation coefficient.

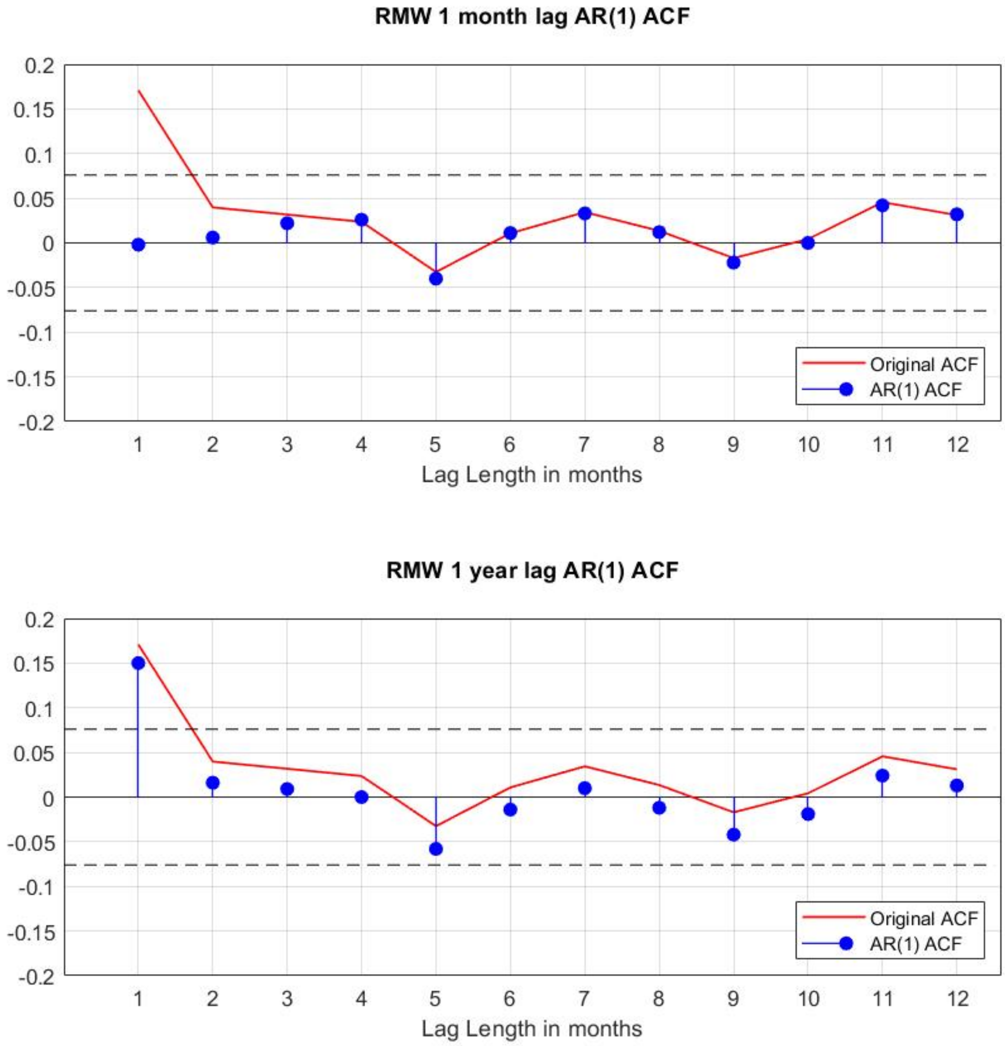

Figure 4: Autocorrelation Function for the original RMW factor and the AR(1) models with one month and one year average lag. The dotted horizontal line represents the Ljung-Box Q-statistic at a 5% significance level of no autocorrelation. The X-axis represents the amount of lags in months and the Y-axis represents the autocorrelation coefficient.

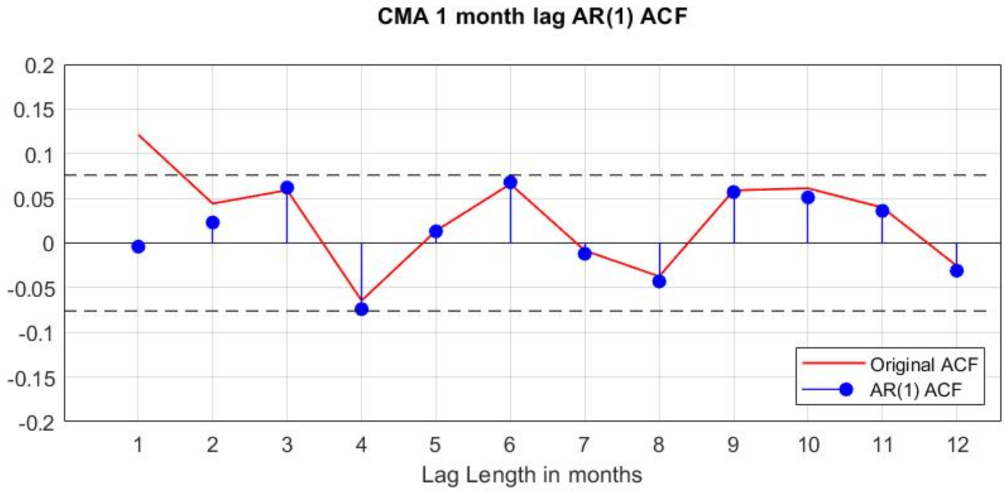

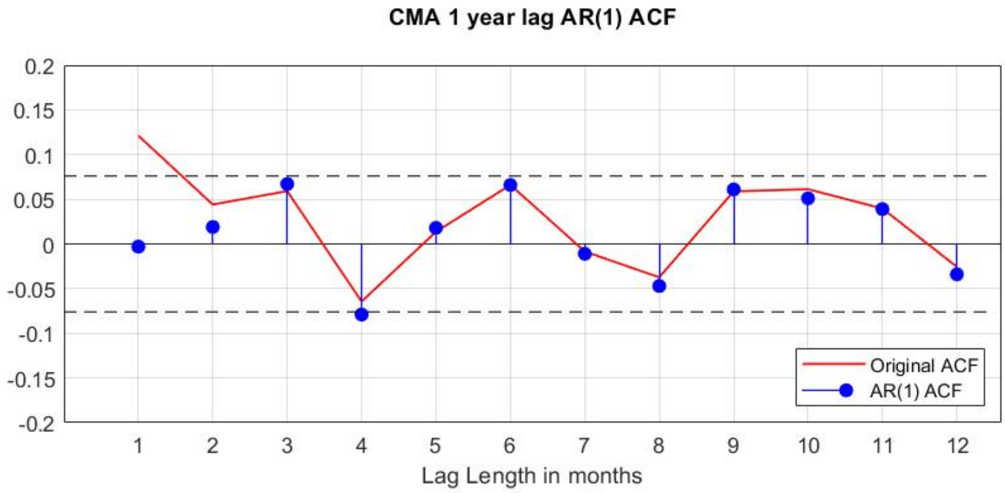

Figure 5: Autocorrelation Function for the original CMA factor and the AR(1) models with one month and one year average lag. The dotted horizontal line represents the Ljung-Box Q-statistic at a 5% significance level of no autocorrelation. The X-axis represents the amount of lags in months and the Y-axis represents the autocorrelation coefficient.

A.2 Autocorrelation Functions to assess heteroskedasticity

This Subsection presents ACF's of the squared residuals of the factors, showing presence of heteroskedasticity. We do this for the squared residuals obtained from subtracting the mean of the original factor and for the squared residuals obtained by estimating a GARCH(1,1) model. These figures are used to assess the ability of the GARCH model to cope well with the present heteroskedasticity.

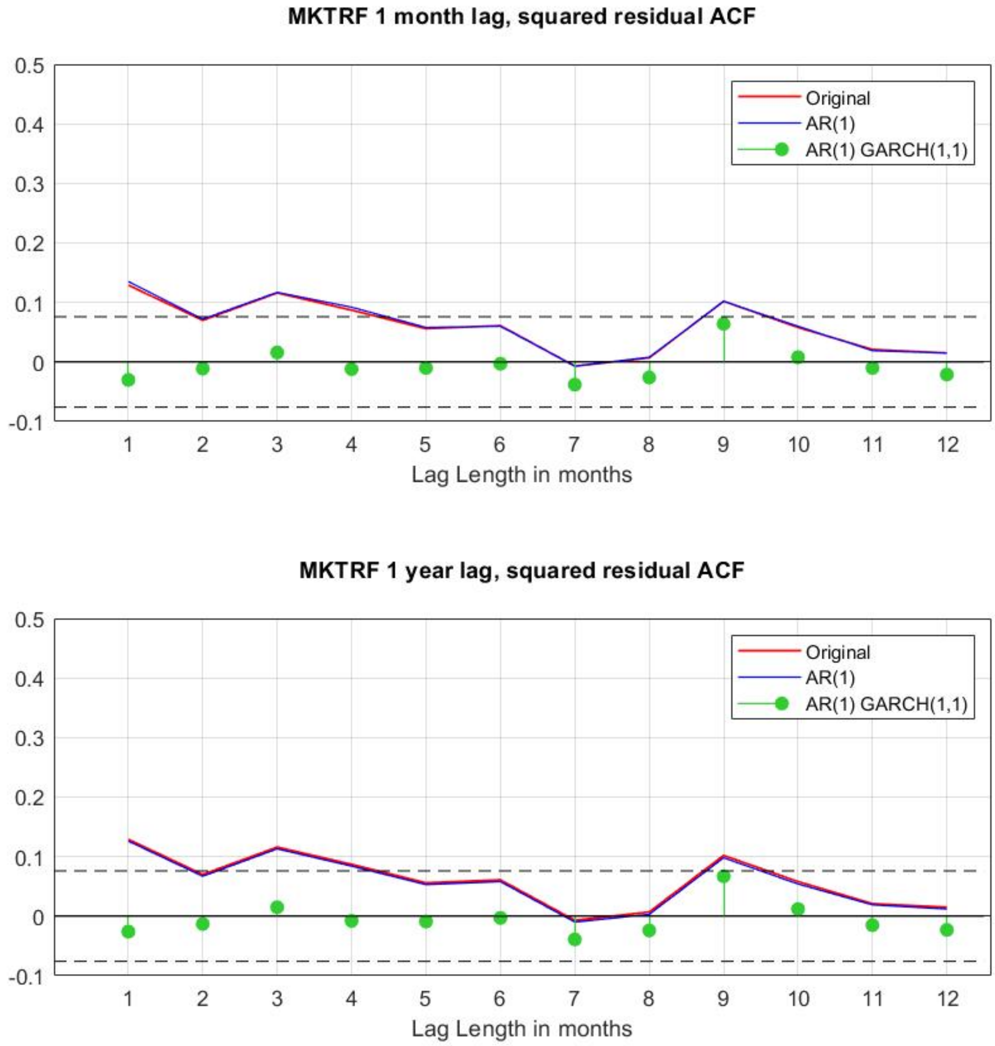

Figure 6: Autocorrelation Function of the squared residuals of the original MKTRF factor, for which we regress a constant on the factor, and the AR(1) and AR(1)-GARCH(1,1) models with one month and one year average lag. The dotted horizontal line represents the Ljung-Box Q-statistic at a 5% significance level of no autocorrelation. The X-axis represents the amount of lags in months and the Y-axis represents the autocorrelation coefficient.

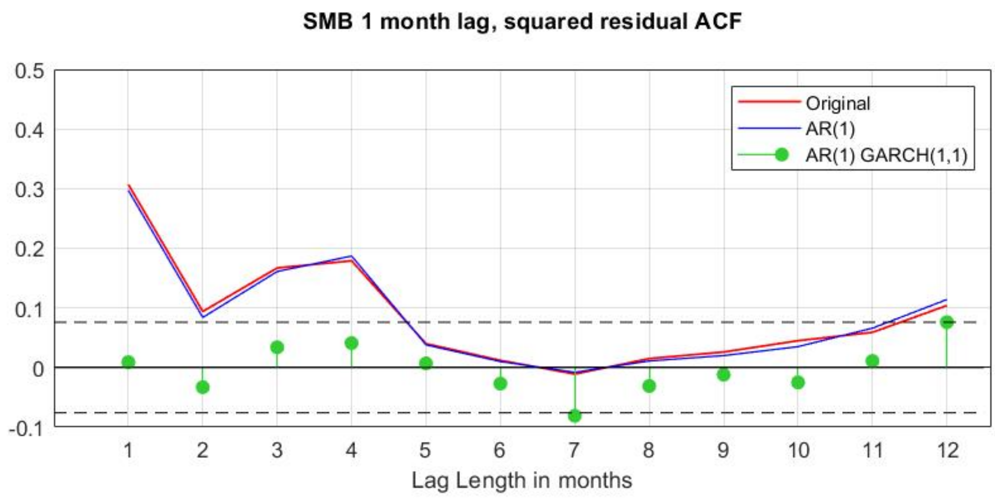

Figure 7: Autocorrelation Function of the squared residuals of the original SMB factor, for which we regress a constant on the factor, and the AR(1) and AR(1)-GARCH(1,1) models with one month and one year average lag. The dotted horizontal line represents the Ljung-Box Q-statistic at a 5% significance level of no autocorrelation. The X-axis represents the amount of lags in months and the Y-axis represents the autocorrelation coefficient.

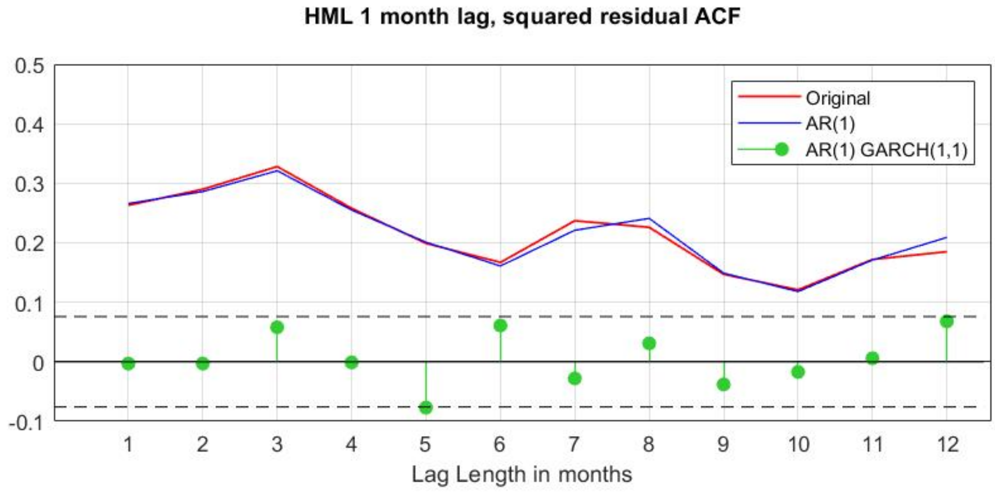

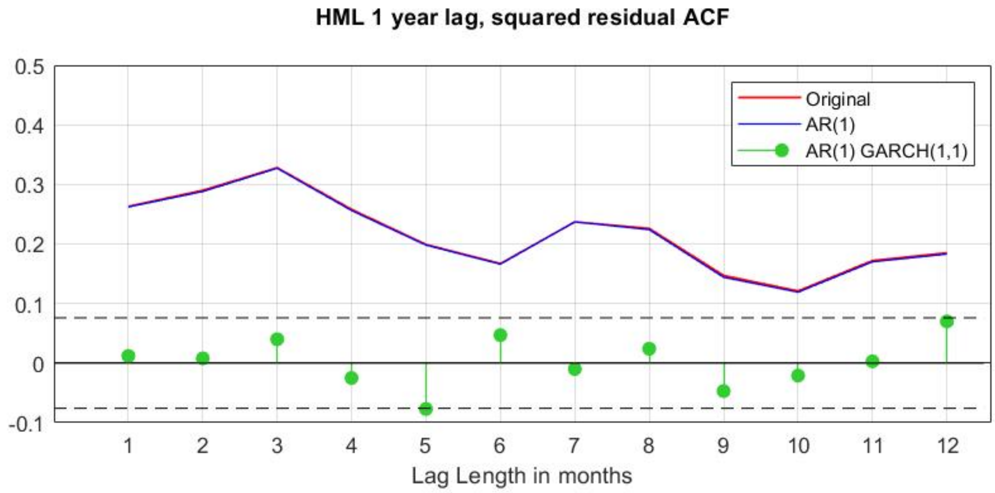

Figure 8: Autocorrelation Function of the squared residuals of the original HML factor, for which we regress a constant on the factor, and the AR(1) and AR(1)-GARCH(1,1) models with one month and one year average lag. The dotted horizontal line represents the Ljung-Box Q-statistic at a 5% significance level of no autocorrelation. The X-axis represents the amount of lags in months and the Y-axis represents the autocorrelation coefficient.

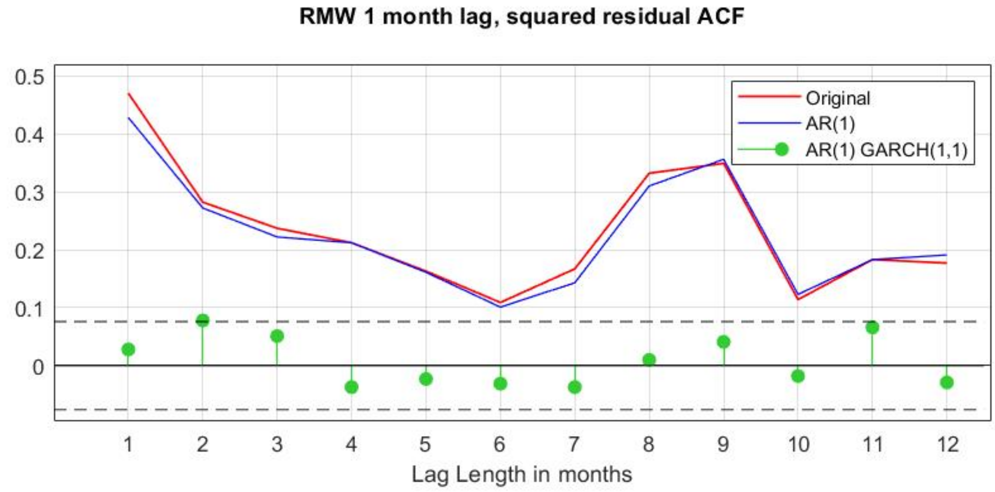

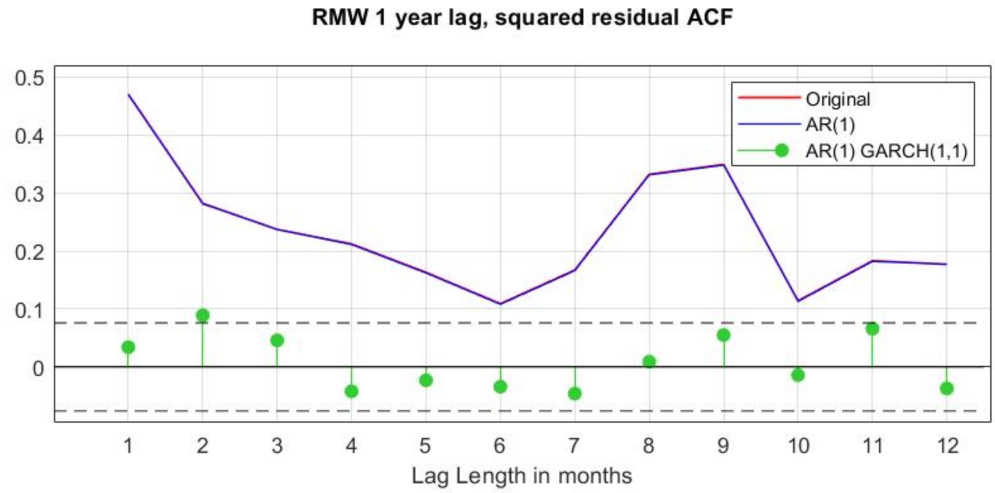

Figure 9: Autocorrelation Function of the squared residuals of the original RMW factor, for which we regress a constant on the factor, and the AR(1) and AR(1)-GARCH(1,1) models with one month and one year average lag. The dotted horizontal line represents the Ljung-Box Q-statistic at a 5% significance level of no autocorrelation. The X-axis represents the amount of lags in months and the Y-axis represents the autocorrelation coefficient.

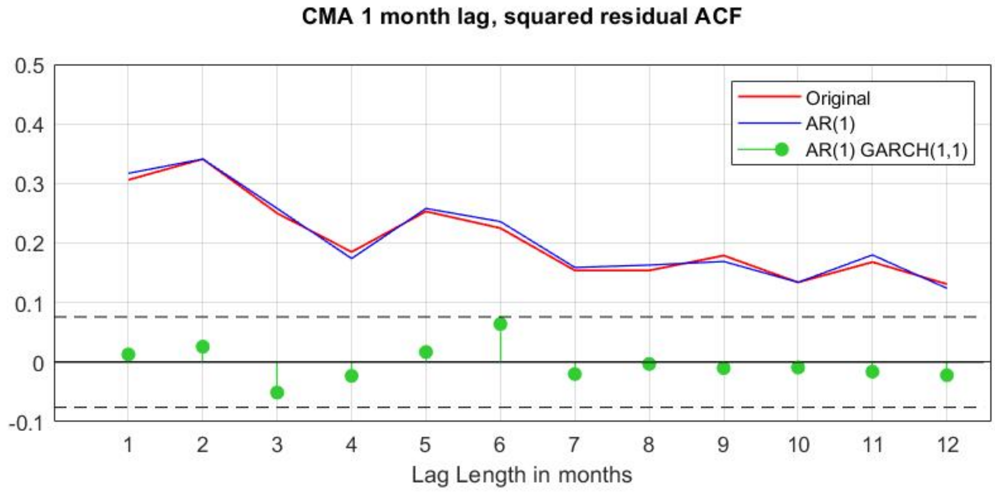

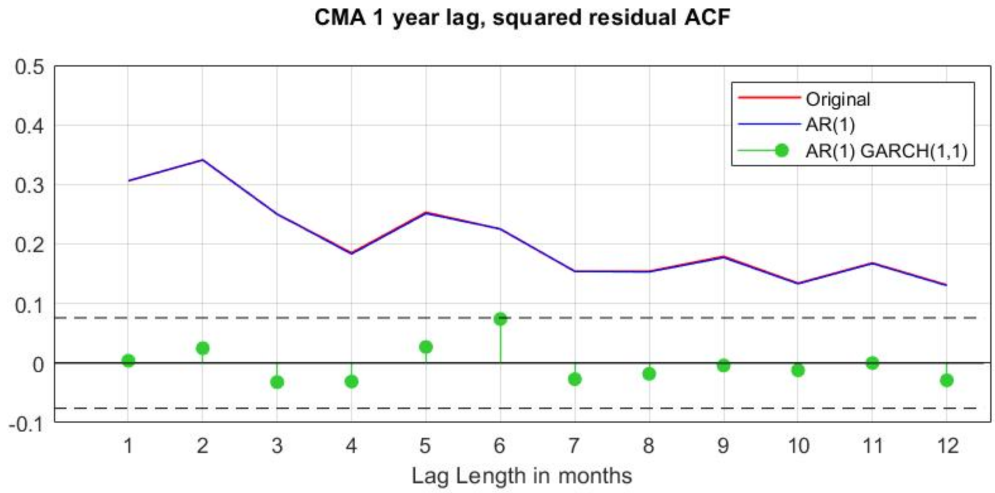

Figure 10: Autocorrelation Function of the squared residuals of the original CMA factor, for which we regress a constant on the factor, and the AR(1) and AR(1)-GARCH(1,1) models with one month and one year average lag. The dotted horizontal line represents the Ljung-Box Q-statistic at a 5% significance level of no autocorrelation. The X-axis represents the amount of lags in months and the Y-axis represents the autocorrelation coefficient.

A.3 Delta Weight Figures

This subsection presents graphs of the ΔWeights, EL(AR) weights minus EL(AR)MM(RV) weights, for the different factors. The figures also contain the monthly Realized Volatility, such that we can assess the relation between the updated weights and the estimated conditional volatility. We consider the figures in order to explain the results of the conditional volatility time-series efficient factors compared to the orginal, unconditional volatility, time-series efficient factor counterpart.

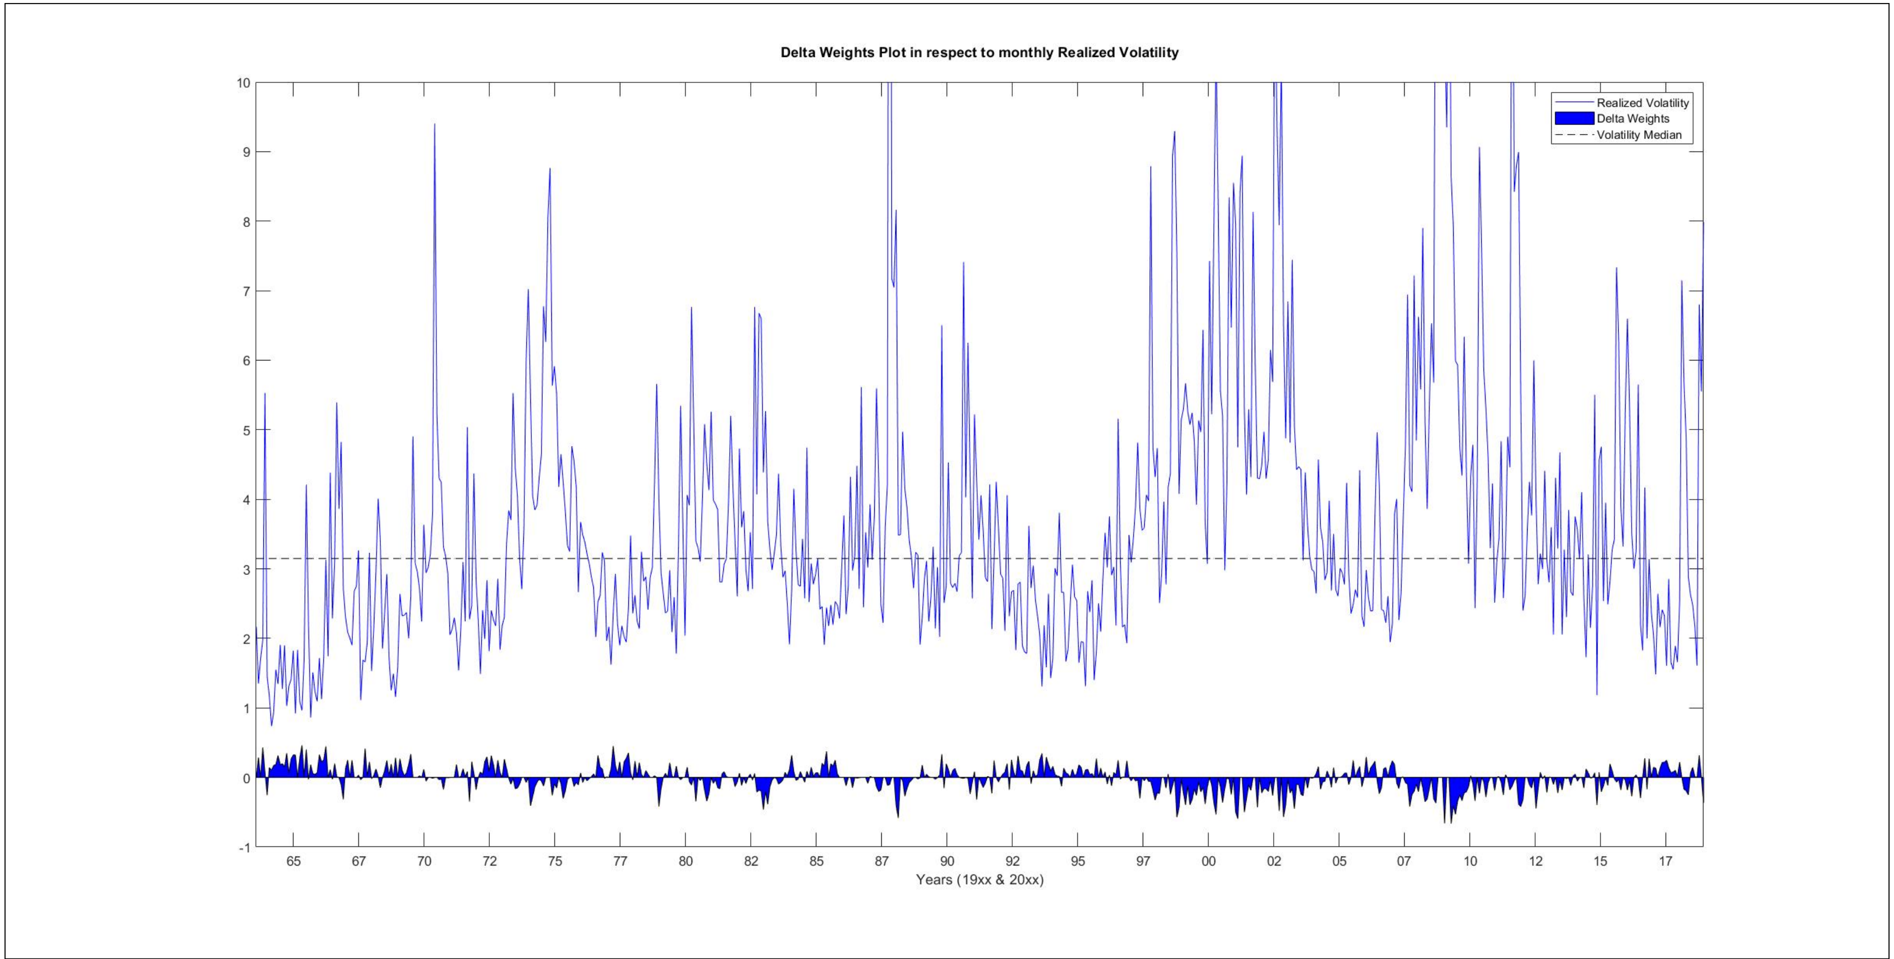

Figure 11: Graph of the ΔWeights
the ΔWeights, EL(AR)MM(RV) minus the EL(AR) weights, for MKRTF are plotted to assess the effect on the weights by conditioning on volatility. The effect on the weights can be compared to the realized volatility plotted above. The blue line corresponds with the estimated realized volatility, the dark blue area corresponds to the ΔWeights and the ragged black line is the realized volatility median.

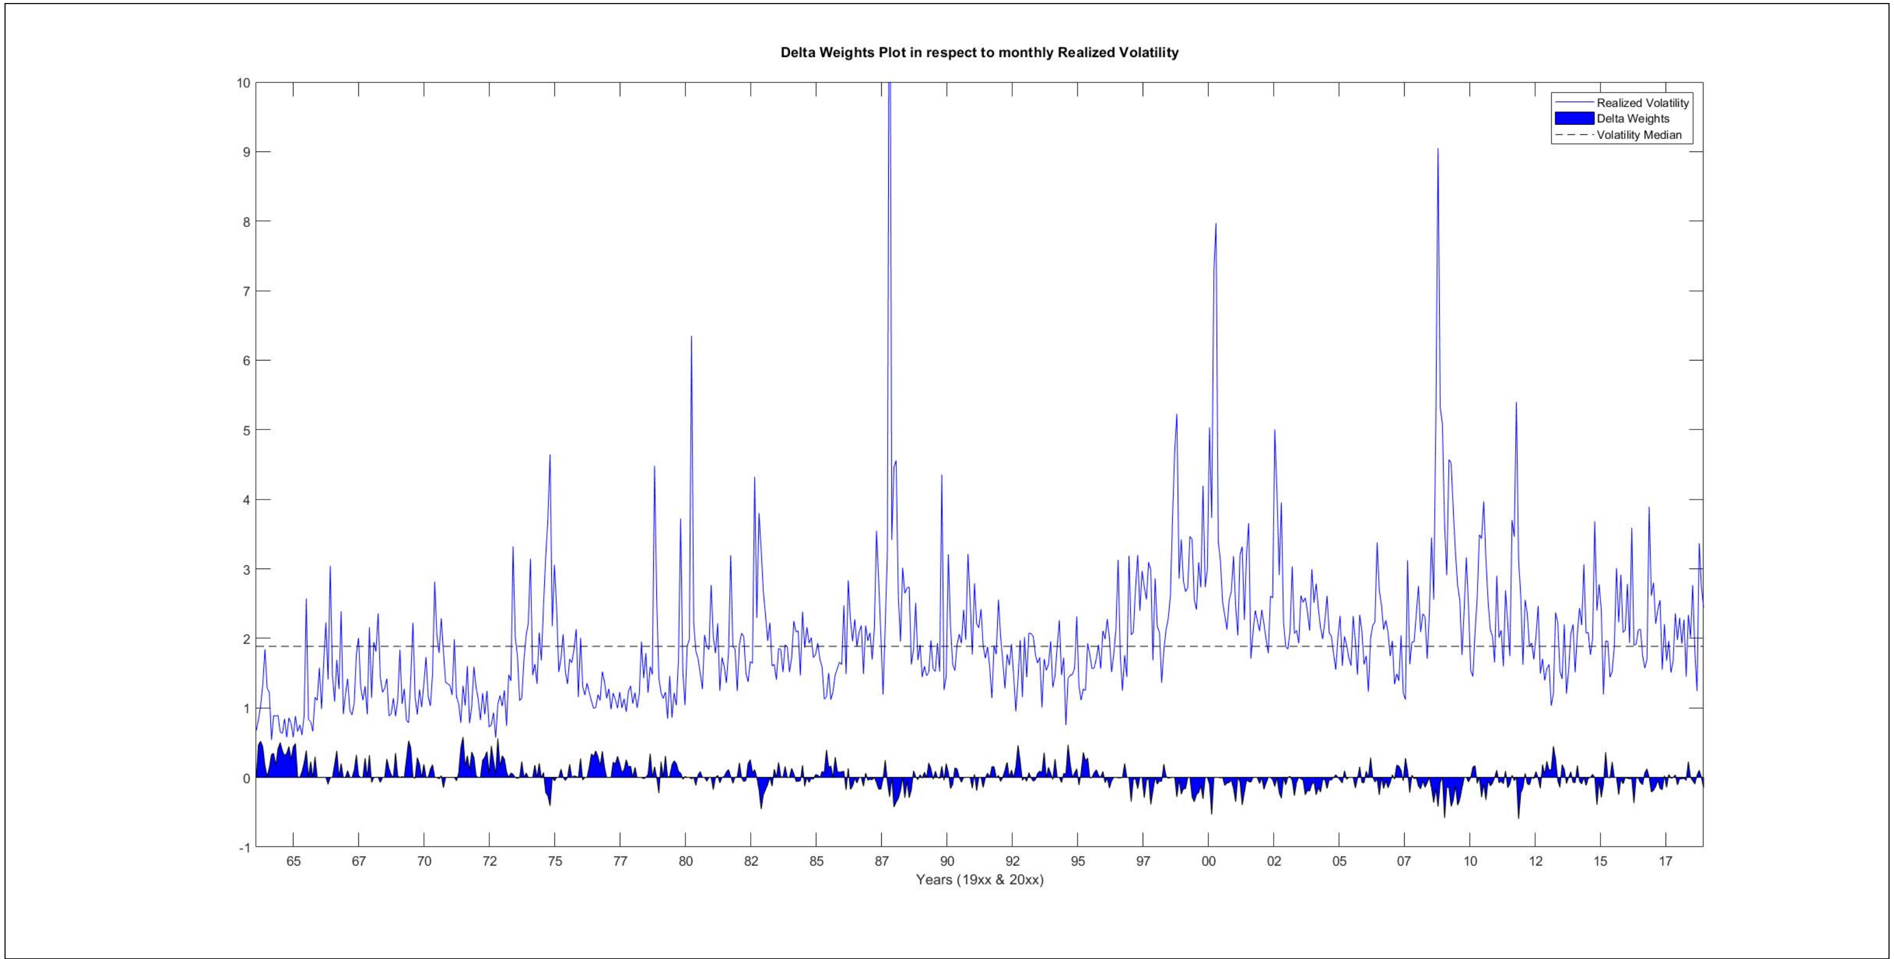

Figure 12: Graph of the ΔWeights
the ΔWeights, EL(AR)MM(RV) minus the EL(AR) weights, for SMB are plotted to assess the effect on the weights by conditioning on volatility. The effect on the weights can be compared to the realized volatility plotted above. The blue line corresponds with the estimated realized volatility, the dark blue area corresponds to the ΔWeights and the ragged black line is the realized volatility median.

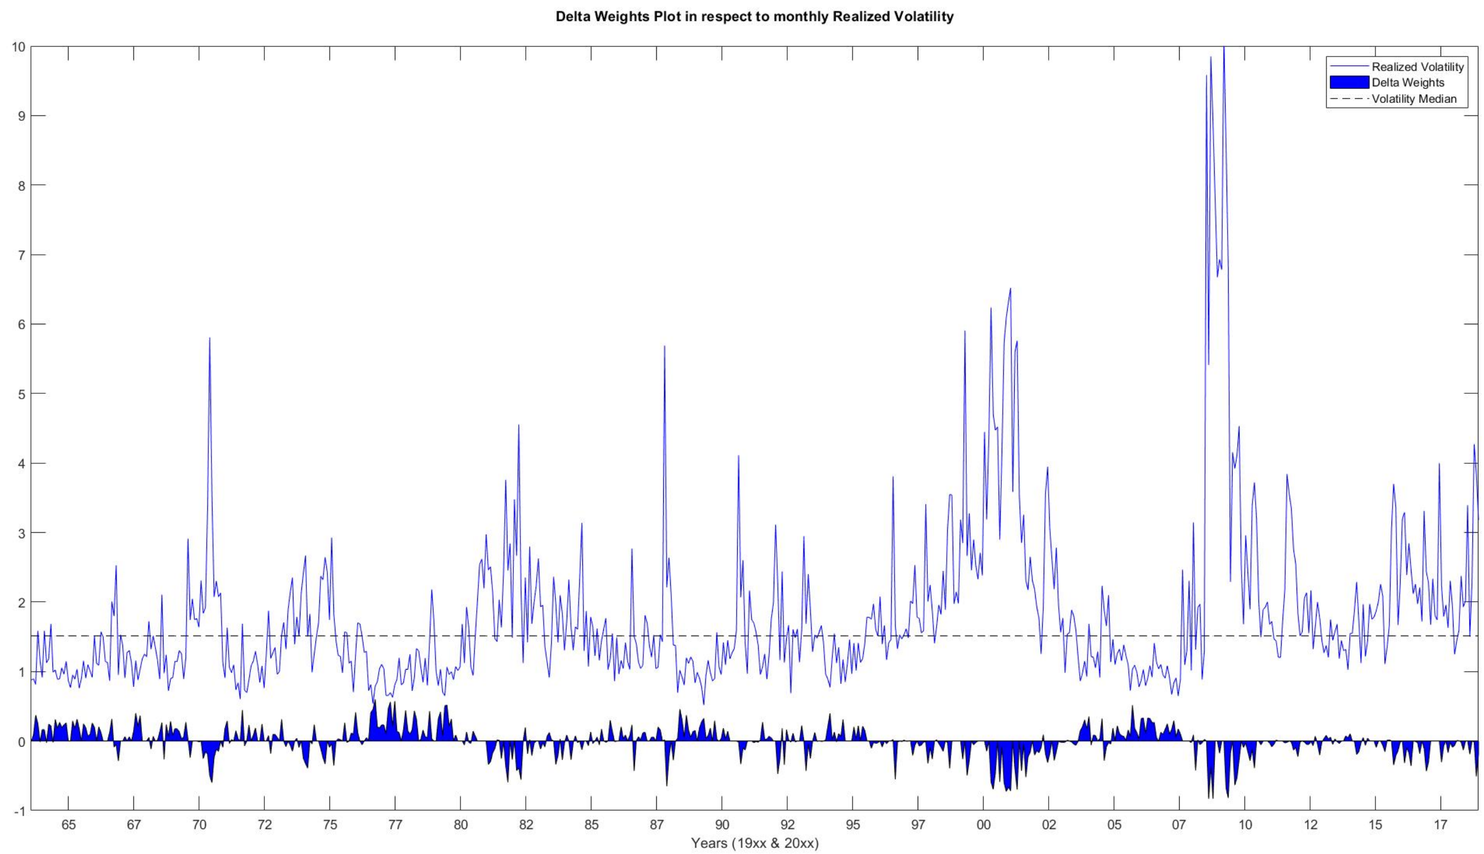

Figure 13: Graph of the ΔWeights
the ΔWeights, EL(AR)MM(RV) minus the EL(AR) weights, for HML are plotted to assess the effect on the weights by conditioning on volatility. The effect on the weights can be compared to the realized volatility plotted above. The blue line corresponds with the estimated realized volatility, the dark blue area corresponds to the ΔWeights and the ragged black line is the realized volatility median.

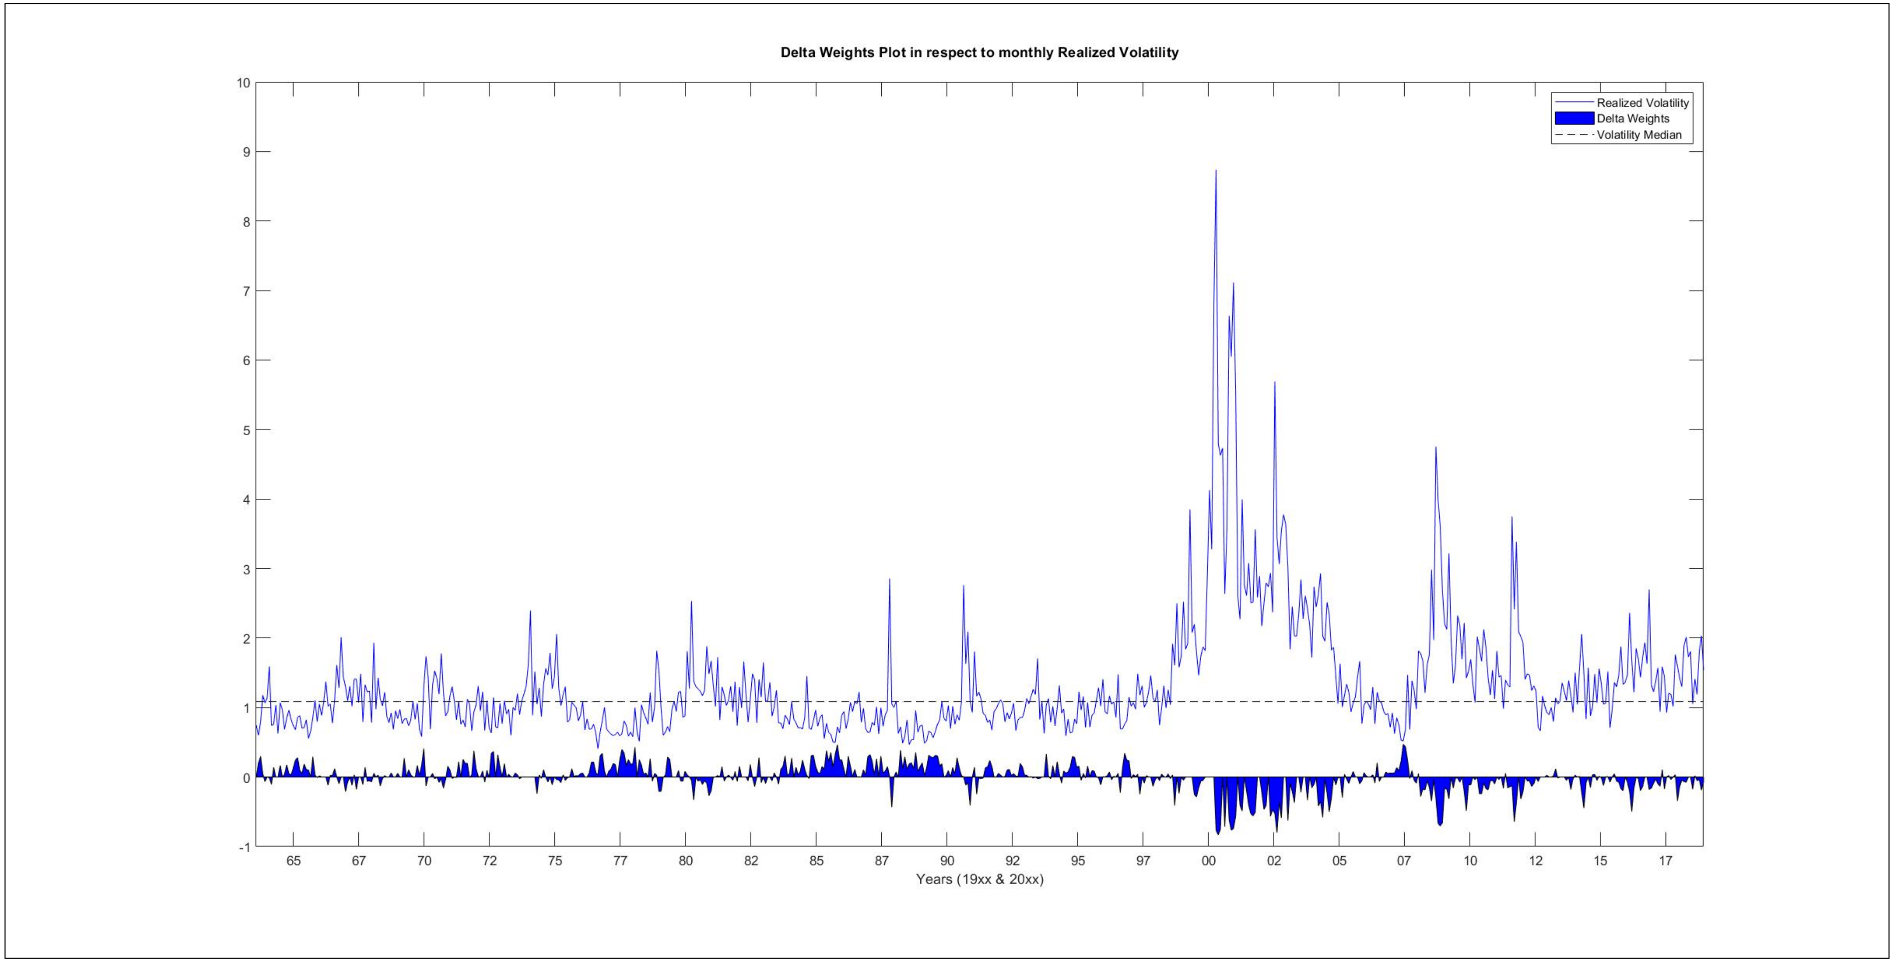

Figure 14: Graph of the ΔWeights
the ΔWeights, EL(AR)MM(RV) minus the EL(AR) weights, for RMW are plotted to assess the effect on the weights by conditioning on volatility. The effect on the weights can be compared to the realized volatility plotted above. The blue line corresponds with the estimated realized volatility, the dark blue area corresponds to the ΔWeights and the ragged black line is the realized volatility median.

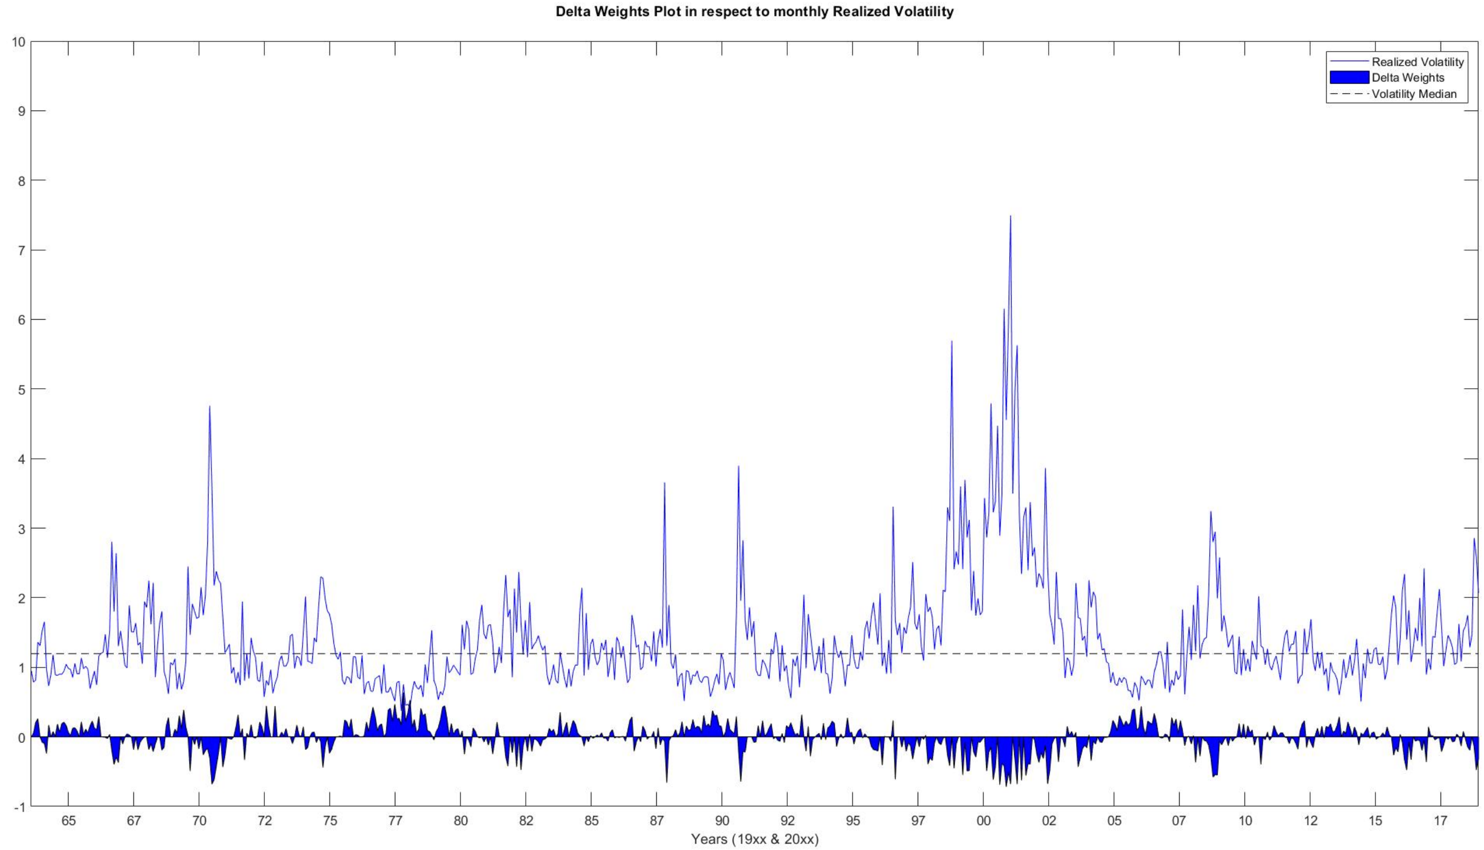

Figure 15: Graph of the ΔWeights
the ΔWeights, EL(AR)MM(RV) minus the EL(AR) weights, for CMA are plotted to assess the effect on the weights by conditioning on volatility. The effect on the weights can be compared to the realized volatility plotted above. The blue line corresponds with the estimated realized volatility, the dark blue area corresponds to the ΔWeights and the ragged black line is the realized volatility median.

References

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

Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of finance*, 52(1):57–82.

Ehsani, S. and Linnainmaa, J. T. (2019). Factor momentum and the momentum factor. Technical report, National Bureau of Economic Research.

Ehsani, S. and Linnainmaa, J. T. (2020). Time-series efficient factors. *Available at SSRN 3555473*.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica: Journal of the Econometric Society*, pages 987–1007.

Fama, E. F. and French, K. R. (1992). The cross-section of expected stock returns. *the Journal of Finance*, 47(2):427–465.

Fama, E. F. and French, K. R. (2015). A five-factor asset pricing model. *Journal of financial economics*, 116(1):1–22.

Fama, E. F. and French, K. R. (2016). Dissecting anomalies with a five-factor model. *The Review of Financial Studies*, 29(1):69–103.

Ferson, W. E. and Siegel, A. F. (2001). The efficient use of conditioning information in portfolios. *The Journal of Finance*, 56(3):967–982.

Gupta, T. and Kelly, B. (2019). Factor momentum everywhere. *The Journal of Portfolio Management*, 45(3):13–36.

Jarque, C. M. and Bera, A. K. (1980). Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Economics letters*, 6(3):255–259.

Jegadeesh, N. and Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of finance*, 48(1):65–91.

Jobson, J. D. and Korkie, B. M. (1981). Performance hypothesis testing with the sharpe and treynor measures. *The Journal of Finance*, 36(4):889–908.

Ljung, G. M. and Box, G. E. (1978). On a measure of lack of fit in time series models. *Biometrika*, 65(2):297–303.

Markowitz, H. (1959). *Portfolio selection: Efficient diversification of investments*, volume 16. John Wiley New York.

Memmel, C. (2003). Performance hypothesis testing with the sharpe ratio. *Finance Letters*, 1(1).

Moreira, A. and Muir, T. (2017). Volatility-managed portfolios. *The Journal of Finance*, 72(4):1611–1644.

Novy-Marx, R. and Velikov, M. (2016). A taxonomy of anomalies and their trading costs. *The Review of Financial Studies*, 29(1):104–147.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The journal of finance*, 19(3):425–442.