WISEflow
Europe/Oslo(CEST)
30 Jun 2022

Handelshøyskolen BI

GRA 19703 Master Thesis

Thesis Master of Science 100% - W

Predefinert informasjon

Startdato:	16-01-2022 09:00	**Termin:**	202210
Sluttdato:	01-07-2022 12:00	**Vurderingsform:**	Norsk 6-trinns skala (A-F)
Eksamensform:	T		
Flowkode:	202210||10936||IN00||W||T		
Intern sensor:	(Anonymisert)		

Deltaker

Navn:	Jørgen Horvei og Maria Bakken

Informasjon fra deltaker

Tittel *:	Does Volatility Timing Enhance Portfolio Performance?
Navn på veileder *:	Geir Høidal Bjønnes

Inneholder besvarelsen konfidensielt materiale?:	Nei	**Kan besvarelsen offentliggjøres?:**	Ja

Gruppe

Gruppenavn:	(Anonymisert)
Gruppenummer:	20
Andre medlemmer i gruppen:	

Does Volatility Timing Enhance Portfolio Performance?

Master Thesis

MSc in Business - Major in Finance

Maria Bakken Jørgen Horvei

Oslo, July 1 - 2022

Abstract

In this paper, we replicate the methodology of Moreira and Muir's "Volatility-Managed Portfolios" (2017). We investigate whether it is possible to benefit from volatility timing in smaller equity markets by testing the strategy on systematic risk factors in Norway. The strategy is constructed by scaling monthly returns by the inverse of their previous month's realized variance. We find that the strategy only performs well in unrealistic trading environments where all costs and restraints associated with transactions are non-existent.

Supervisor

Geir Høidal Bjønnes

Acknowledgements

This work is the culmination of a lengthy and winding road that is higher education. Its ups and downs endured by the endless support of our families, friends, colleagues, and professors. We would like to sincerely thank our supervisor, Geir Høidal Bjønnes, for the excellent supervision, clarification, and comments throughout this process. His insights and advice have been invaluable to us.

Table of Contents

1. Introduction

In finance, the risk-return tradeoff is considered one of the most fundamental and robust theories there is (Markowitz, 1952). The simple principle that potential returns rise with a corresponding increase in risk, entails that high levels of uncertainty should be associated with a greater probability of higher returns. However, this relation does not always seem to hold empirically, as theory suggests. In Moreira and Muir's paper "Volatility-Managed Portfolios" (2017) published in The Journal of Finance in 2017, they propose a new trading strategy based on volatility timing, which by construction has the same volatility as the market but generates higher returns. They document evidence on the effectiveness of the strategy by pointing at its ability to obtain positive and significant alphas, increased Sharpe ratios and large utility gains for mean-variance investors. These findings are exceptionally interesting as they contradict well-established theory and challenge the basic notion of the risk-return tradeoff.

The strategy is based upon volatility in the market and is carried out by constructing portfolios that scale monthly returns by the inverse of their previous month's realized variance. By doing this, one decreases risk exposure when variance was recently high and increase risk exposure when variance was recently low. Each month, the strategy adjusts risk exposure to the portfolio according to the variation in realized variance. The outcome of performing this procedure is the creation of what is referred to as volatility-managed portfolios. The analysis is instituted from the vantage point of a mean-variance investor, who adjusts his or her investment allocation according to the attractiveness of the mean-variance tradeoff, $\frac{\mu_t}{\sigma_t^2}$. The strategy is tested on several systematic factors that summarize pricing information for a wide cross section of assets; the market, value, momentum, profitability, return on equity, investment, and betting-against-beta factors, as well as the currency carry trade.

In this research paper, we examine the time-series relation between risk and return by replicating the methodology of Moreira and Muir (2017) on the Norwegian stock market. Our aim is to investigate whether it is possible to benefit from volatility timing in smaller equity markets, which has been neglected by other researchers. Like the original paper, we test the strategy on numerous systematic risk factors. The strategy is executed on the Fama-French three-factor model (1993), which implies the overall excess market risk (MKT), small minus big (SMB) and the high minus low (HML) factors. Additionally, we test the strategy on the liquidity factor (LIQ) from Næs, Skjeltorp and Ødegaard (2011) and on both the Carhart's momentum factor (1997) and the Fama-French's momentum factor (2015). Each factor is scaled based on their scaling factor, which is plotted in Figure 1.

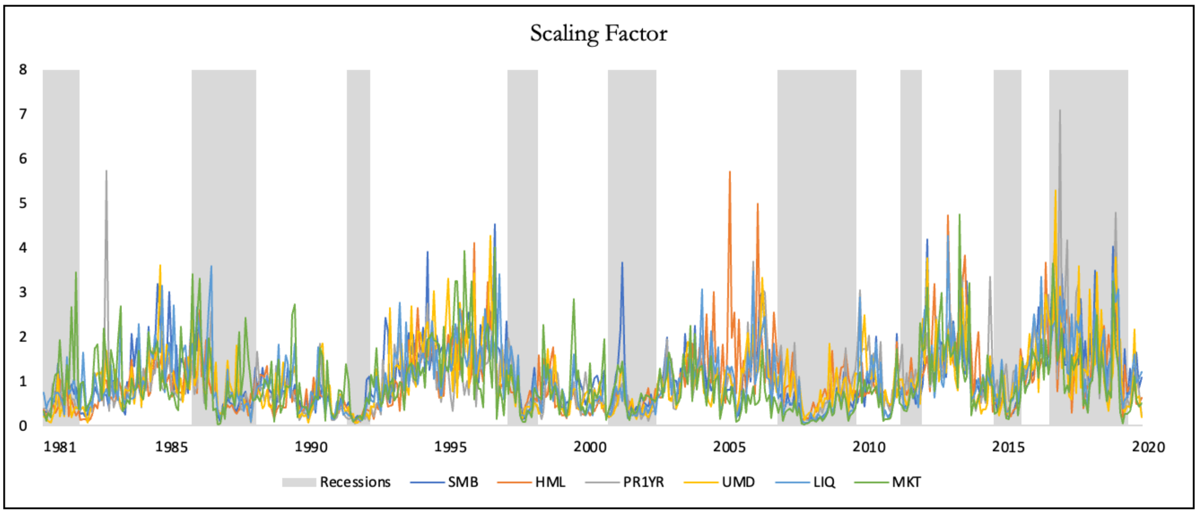

Figure 1. Time series of average risk exposure by factor. This figure plots the weights of the monthly volatility-managed factor portfolio in Norway, computed using lagged realized variance. The volatility-managed strategy takes more risk when volatility is low and less risk when volatility is high. Light shaded bars indicate NORREC recessions and visualize the strategy's empirical pattern of risk exposure.

To evaluate the strategy's performance, we run univariate regressions of the volatility-managed factors on the unmanaged factors. All alphas turn out positive and statistically significant at the 10% level, except HML. Our best performing factor is the momentum factor by Carhart, holding an alpha of 9.42%. The market factor has an alpha of 3.95%, an appraisal ratio of 0.30 and reports an increase in utility of 14.83%. A positive and significant alpha implies that the volatility-managed strategy expands the mean-variance efficient frontier. We extend our analysis to a multifactor environment by constructing mean-variance efficient portfolios. The MVE portfolios endure the same procedure as the single factors but are constructed using various combinations of factors. We consider four different factor combinations; one Fama-French three-factor model (1993), one that contains all factors, and two separate portfolios where UMD and LIQ are put together with the Fama-French three-factor model. We observe positive and significant alphas for all factor combinations, ranging from 2.77% to 3.63%. Additionally, we examine how the strategy works during shorter and more realistic investment horizons, and construct 10-year subsamples for both single-factor and multifactor portfolios. Not surprisingly, the subsamples generate the best results during recession periods. Otherwise, we find that the strategy is not sufficient at a 10-year horizon.

In the last and most extensive part of this paper, we carry out a comprehensive validation process to evaluate the robustness of our results. The robustness checks investigate the strategy from several different perspectives, to better understand the profitability of volatility timing. We conduct similar tests as to the original paper, but also perform additional robustness checks beyond those conducted by Moreira and Muir (2017). The analysis establishes that business cycle risk is an unlikely explanation for our results. This is done by regressing all volatility-managed factors on the original factors while also including an additional interaction term consisting of a NORREC recession dummy. To validate the approach of using constant weights for the construction of MVE portfolios, we normalize each factor by a common volatility factor. We compute the first principal component of realized variance across all factors and normalize each factor by $\frac{1}{RV_t^{PC}}$. The relatively strong comovement in factor volatilities validates the approach. Moreira and Muir (2017) concentrate their analysis around well-established systematic factors. As a contribution to their research, we consider the possibility of implementing the strategy on less diversified, industry portfolios. We find that almost all industries in the Norwegian market produce positive and significant alphas. Another important contribution to the research is our evaluation of different rebalancing frequencies. We rebalance weekly, biweekly, monthly, quarterly, and annually to see if the strategy performs better at other frequencies. Our results suggest that alphas decline with horizon, to the point where the rebalancing frequency is so low that it is unable to capture the effect of a volatility shock.

Moreover, we implement leverage constraints to explore how the strategy performs with market frictions. As it turns out, the strategy is unobtainable in the Norwegian stock market without the use of leverage. Even with 50% leverage, the strategy is unable to produce significant alphas. By looking into the effect of transaction costs, we observe a remarkably lower break-even point in our analysis compared to Moreira and Muir (2017). In fact, the strategy is unable to produce significant alphas and our results suggest that the strategy fails to survive transaction costs. This results in a blurred line between a profitable strategy and an unprofitable strategy as transaction costs are hard to estimate properly. Finally, we consider inabilities of easily exiting positions in times of market distress by looking into liquidity risk. By reducing profits during indications of market distress, we find that after accounting for liquidity costs, the strategy underperforms relative to the buy-and-hold strategy.

The paper proceeds as follows; In section 2, we review the relevant literature. Section 3 gives a brief description of the data. Section 4 shows the methodological approach. In section 5 we present our main results. In section 6 we conduct robustness checks and discuss implications regarding these and our main results. Section 7 presents the conclusion.

2. Literature Review

The idea of volatility-based trading strategies accelerated when Fleming, Kirby, and Ostdiek (2001:2003) thoroughly discussed the matter in their papers from 2001 and 2003. They highlighted the economic significance of time-varying, predictable volatility and found that strategies that time volatility outperform portfolios that are not, based on the same expected return and volatility. In a systematic approach, Fleming et. al. (2001) considers a mean-variance investor acting on optimization rules to allocate funds, across four asset classes; stocks, bonds, gold, and cash. As variance and covariance can be predicted with greater precision than return, expected return is treated as constant while the portfolio weights change according to the covariance matrix. The weights are determined by minimizing the mean squared error of the estimator. To measure the value of volatility timing, a fee is estimated based on the likelihood that an investor would be willing to pay to switch from the static non-managed portfolio to the dynamic volatility-managed portfolio. Fleming et. al. (2001) finds that 92% of the simulated cases are resulting in a positive fee, indicating that a risk averse investor would switch from the optimal static to the dynamic portfolio. On average, the estimated fee exceeded 170 bps per year. Their findings functioned as a building block for further research on the effect of volatility timing.

One of the most prominent repercussions was by Barroso and Santa-Clara (2015) which found that managing the risk of momentum strategies give rise to large economic gains. The methodology used in this research coincide with the one later adopted by Moreira and Muir (2017). Historically, we find that momentum strategies have performed very well with an average monthly excess return of 1.75% from 1927 to 2011. However, occasional large crashes are observed, which paints the strategy as unappealing to many investors. Barroso and Santa-Clara (2015) sought to mitigate this risk and proposed a different method to manage risk. The long-short portfolio is scaled by its realized variance in the previous six months, while targeting a constant volatility. By securing a constant target volatility, we observe similar risk for volatility-managed portfolios and its non-managed counterparts. They find an increase in Sharpe ratio from 0.53 for the non-managed portfolio to 0.97 for the volatility-managed portfolio. However, the greatest benefit from this approach is the reduction in crash risk. Barroso and Santa-Clara (2015) report that excess kurtosis decrease from 18.24 to 2.68 and that the left skewness increase from -2.47 to -0.42. The scaled momentum strategy is reported as robust across subsamples, and on international data, indicating value for markets outside the U.S.

Research related to timing volatility accelerated during the financial crisis in 2008, based on investors behaviors. Research has found that top investors along with older investors are more likely to deleverage their positions during market distress (Nagel et. al., 2016). This contradicts conventional wisdom to increase risk-taking or hold risk-taking constant. A widely held view during the aftermath of the financial crisis in 2008, was that those who reduced their positions missed out on a once in a decade opportunity (Cochrane, 2008). However, we evidently find that strategies that deleveraged its positions during the financial crisis obtained better results looking back, and the ability to return to the market whenever volatility receded, was present (Moreira & Muir, 2017). These tendencies can be explained by higher attention to their portfolios, which results in the ability to time the market better than others (Nagel et. al., 2016).

Our thesis is directly related to the paper of Moreira and Muir (2017), "Volatility-Managed Portfolios" from 2017. In their continuation of the research, a new paper follows in 2019 where they build on the idea that long-term investors should time volatility (Moreira & Muir, 2019). This paper is the first to study a long-term investor's portfolio in a framework that is adaptable enough to fit the most dominant facts about the aggregate stock market. It embodies the portfolio problem of a long-term investor that allocates wealth between a risky asset and a riskless asset under circumstances where both volatility and expected returns are time varying. Additionally, the research incorporates the importance of parameter uncertainty related to the dynamics of volatility and expected returns. The main findings suggest substantially decreased risk exposure when there is an increase in volatility, and that ignoring volatility variations generates large utility losses. Furthermore, Moreira and Muir (2019) report an 2.4% increase in wealth per year resulting from the volatility timing, both with an investment horizon of 5- and 20 years. These benefits outperform expected return timing, especially when parameter uncertainty is accounted for. Nevertheless, the paper provides empirical evidence that the benefits of volatility timing are independent of the investor's horizon (Moreira & Muir, 2019).

Not surprisingly, multiple research papers have made the attempt to disproof the controversial and highly unorthodox study carried out by Moreira and Muir (2017). In the following, we will discuss two such papers. Cederburg, O'Doherty, Wang, and Yan (2019) contributes to the research by studying a comprehensive set of 103 equity strategies and examine the value of volatility-managed portfolios for real-time investors. The research finds that there is no statistical or economical evidence that indicates systematically higher Sharpe ratios for volatility-managed portfolios compared to non-managed portfolios in real time. Findings show that the volatility-managed portfolio outperforms in 53 out of the 103 cases, while the non-managed portfolio outperforms in 50 out of the 103 cases, which indicates no significant difference between the two strategies. Moreover, Cederburg et. al. (2019) points to the difficulties of implementing the strategy by Moreira and Muir (2017) in real time, as a consequence to investors being required to combine the volatility-scaled as unscaled versions of a given portfolio using ex post optimal weights. When trading strategies are formed based upon the information available at the time being, their performance may differ from the ex-post results. The main reasons behind this are first; that the conditional risk-return trade-off could be unstable over time, for a given factor. It is likely that past results are undesirable estimations of future potential, in view of their poor informative character. Secondly, estimation risk raises concerns as weights are often unstable and one finds that the corresponding optimal portfolios tend to perform inadequately out of sample.

A direct response to Moreira and Muir's research is the research paper called; "*Volatility-Managed Portfolios: Does It Really Work?*" written by Liu, Tang and Zhou (2018). Their findings point to several things that are worth mentioning as it is closely related to our ongoing research, the first one being the identification of look-ahead bias. The paper reports the presence of look-ahead bias in the procedures of Moreira and Muir (2017) and evaluates the correction of it in the data. Incorporating look-ahead bias correction, results in maximum drawdowns of 68-93% compared to the maximum drawdowns of 56% without the correction, which makes the strategy highly difficult to implement. Secondly, the paper uncovers that the Sharpe ratio of the volatility-managed portfolio does not differ much compared to the buy and hold portfolio. When performing out of sample tests, Liu et. al. (2018) find that the Sharpe ratio is only marginally higher and that the differences are never statistically significant. Further investigation indicates that the strategy outperforms the market only during recessions or times of financial crises. Additionally, the research paper moves to break down the sample period into subplots of approximately 20-year time periods. Reported results find that in half of the subsamples, the strategy of timing volatility underperforms relative to the market. Exceptionally, the subsample including the financial crisis in 2008 reveals that the volatility managed portfolio outperforms the market. This enhances the beforementioned note that timing volatility only produce significant results in times of financial distress. Furthermore, these indications weaken the likelihood of implementing the strategy in practice.

Equally common are research papers investigating alternative explanations to the significant abnormal returns obtained by Moreira and Muir (2017) in their "Volatility-Managed Portfolios" paper. Barroso and Detzel (2021) explore the likelihood of limits to arbitrage explaining the benefits of volatility-managed portfolios. They argue that volatility-managed portfolios contradict with traditional investment behavior and that the strategy's outperformance cannot be explained by rational asset-pricing models. Therefore, they suggest that it would be reasonable to believe that the phenomenon persisted by reason of arbitrage frictions. The research paper finds that managed market factors' performance is concentrated in high-sentiment times, which is consistent with prior theory in which sentiment traders buy more aggressively than they sell short and underreact to volatility shocks. Accordingly, Barroso and Detzel's (2021) results find that investors are better off with not timing volatility when sentiment is low.

3. Data Description

The sample period ranges from July 1981 through November 2020 and is consistent for all data used in our empirical analyses. We use Norwegian factor data and Norwegian industry data. Our main sample consists of daily- and monthly returns on Norwegian factors collected from Bernt Arne Ødegaard's (2021) website[1]. The source of the raw data stem from daily observations of stock market data from Oslo Stock Exchange data service. The Norwegian factors include the size factor (SMB), value factor (HML), Carhart momentum factor (PR1YR), Fama-French momentum factor (UMD), liquidity factor (LIQ)[2], market return (MKT) and the risk-free rate (RF). The data sample composes to 9894 daily observations, which converts into 473 months. All portfolios are value weighted, and the least liquid stocks are filtered out. More comprehensive statistics on the Norwegian factors are assembled in appendix 1, 2 and 3. The Norwegian industry data is composed of stocks traded on the Oslo Stock Exchange that are sorted into industry portfolios based on their Global Industry Classification Standard.

The risk-free rate needs some extra attention since the estimations differ over our sample period. From 1986 and onwards, the Norwegian Interbank Offered Rate, NIBOR has been used as an estimation of the risk-free rate. Since monthly NIBOR rates only are available from 1986, some imperfect proxies are used for the previous periods. From 1982-1986, the overnight NIBOR rates are used as an approximation, while the shortest possible bond yield for treasuries in Eitrheim et al. (2007) is used from 1981-1982. It is, however, highly unlikely that these imperfect proxies will have any substantial impact on our results.

[1] https://ba-odegaard.no/financial_data/ose_asset_pricing_data/index.html
[2] See (Naes et al., 2008) for more information on the liquidity factor.

4. Methodology

Motivated by the portfolio problem of a mean-variance investor who adjusts his or her allocation according to the attractiveness of the mean-variance trade-off $\frac{\mu_t}{\sigma_t^2}$, we evaluate the volatility-managed strategy's performance in the Norwegian equity market. We investigate whether volatility timing is able to produce positive and significant alphas, and formally test if the strategy can increase utility and expand the mean-variance frontier.

4.1 Portfolio Construction

We replicate the methodology of Moreira and Muir (2017) and construct volatility-managed portfolios by scaling monthly factor returns by the inverse of their conditional variance. Our volatility-managed portfolios are made up by well-established risk factors that summarize pricing information contained in a wide set of assets and thus we focus our attention on the behavior of these factors. Each month, the strategy adjusts risk exposure to the portfolio according to the variation in our measure of conditional variance. We use the previous month's realized variance as a proxy for conditional variance, such that the managed portfolios scale by the inverse of their previous month's realized variance. The volatility-managed portfolio of a given factor in month $t + 1$ is then

$$f_{t+1}^{VM} = \frac{c}{\hat{\sigma}_t^2(f)} f_{t+1} \tag{1}$$

where f_{t+1} is the return of the non-managed portfolio, $\hat{\sigma}_t^2(f)$ is the portfolio's realized variance, and the constant c controls the average exposure of the strategy. We choose c such that the unconditional standard deviation of the managed portfolio equals that of the non-managed portfolio[3]. Monthly realized variance is calculated based on daily returns

$$\hat{\sigma}_t^2(f) = RV_t^2(f) = \sum_{d=1/td}^{1} \left(f_{t+d} - \frac{\sum_{d=1/td}^{1} f_{t+d}}{td} \right)^2 \tag{2}$$

where td is the number of trading days in each month. Future monthly variance is thus effectively modelled as a random walk without drift. An appealing aspect of this approach is

[3] The choice of c has no effect on the Sharpe ratio of the volatility-managed portfolio, and thus the fact that we use the full sample to compute c does not bias the results.

that it is easily implementable by an investor in real time, as it does not require any parameter estimation.

4.2 Empirical Methodology

The portfolio construction leaves us with two sets of portfolios for each factor; the original monthly time-series of unmanaged factor returns f_{t+1} and a monthly time-series of volatility-managed factor returns f_{t+1}^{VM}. To examine whether the strategy expands the efficient frontier, we run univariate time-series regressions of the volatility-managed portfolios on the original unmanaged portfolios factor by factor. The model specification is given by;

$$f_{t+1}^{VM} = \alpha + \beta f_{t+1} + \epsilon_{t+1} \tag{3}$$

For systematic factors that summarize pricing information for a wide cross-section of assets and strategies, a positive α implies that the volatility-managed strategy expands the mean-variance efficient frontier. These single-factor alphas have economic interpretation when the individual factors accurately describe the opportunity set of investors or if each factor captures a different dimension of risk (Moreira & Muir, 2017). Additionally, a significant and positive α can indicate that volatility timing is able to increase Sharpe ratios relative to the original unmanaged factors. To quantify this increase, we report annualized Appraisal ratios AR;

$$AR = \sqrt{12}\frac{\alpha}{RMSE} \tag{4}$$

where α is the unconditional alpha and RMSE is the Root Mean Square Error of the regression (3). Formally, RMSE is the standard deviation of the residuals;

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(\hat{y}_n - y_n)^2} \tag{5}$$

where y_n is the observed values and $\hat{y}_n$ is the predicted values from our model.

Utility gains are calculated to further evaluate the economic relevance of our results. Here we assess the benefits of volatility timing by comparing the increase in Sharpe ratio for the volatility managed portfolio relative to the buy and hold portfolio. More specifically, we compute the percentage utility gain as;

$$\Delta U_{MV}(\%) = \frac{SR_{new}^2 - SR_{old}^2}{SR_{old}^2} \tag{6}$$

where $SR_{new}^2 = \sqrt{SR_{old}^2 + \left(\frac{\alpha}{\sigma_\varepsilon}\right)^2}$ and SR_{old} is the Sharpe ratio given by the original unscaled factor.

4.3 Mean-Variance Efficient Portfolios

We extend our analysis to a multifactor environment by constructing new volatility-managed portfolios that combine the multiple factors from before. We let F_{t+1} be a vector of factor returns and b the static weights that generate the maximum in-sample Sharpe ratio. We define the Mean-Variance Efficient (MVE) portfolio as $f_{t+1}^{MVE} = b'F_{t+1}$. We then construct the volatility-managed MVE portfolio as

$$f_{t+1}^{MVE,\sigma} = \frac{c}{\hat{\sigma}_t^2(f_{t+1}^{MVE})} f_{t+1}^{MVE} \tag{7}$$

where f_{t+1}^{MVE} is the return of the non-managed MVE portfolio, $\hat{\sigma}_t^2(f_{t+1}^{MVE})$ is the MVE portfolio's realized variance, and c controls the average exposure of the strategy. As before, we choose c such that the unconditional standard deviation of the managed portfolio equals that of the non-managed portfolio. What is interesting about this is that the volatility-managed portfolio shifts the conditional beta on the MVE portfolio without altering the relative weights across the individual factors. The result is a strategy that focuses entirely on the time-series aspect of volatility timing.

5. Main Results

Before reporting our main results, we refer to appendix 5 which contains a replication of Moreira and Muir's study made for the U.S factors; MKT, SMB, HML, RMW and CMA. We do this to verify that our method is equivalent to theirs. The following sections disclose our main results, which resembles those presented by Moreira & Muir (2017).

5.1 Single-Factor Portfolios

We first conduct our analysis factor by factor by regressing our volatility-managed factor portfolios on the original unmanaged portfolios. As most of the factors are lowly correlated as reported in appendix 3, they capture different dimensions of risk. In our single-factor analysis, we find evidence of the same empirical pattern documented by Moreira and Muir (2017). Our results show pervasiveness across factors and imply that the results are uniquely driven by the time-series relationship between risk and return.

Table 1, Panel A reports the results from running the univariate time-series regression from (3). We document positive and statistically significant alphas for PR1YR, UMD and LIQ of 9.41%, 9.21% and 3.78%, respectively at the 5% significance level. Our HML factor is significant at the 5% level but reports a negative alpha of -3.83%. The SMB factor and the EMKT factor report positive alphas but are both just marginally significant at the 10% level. Notably, we find that SMB is positive and HML is negative, while Moreira and Muir find the opposite. The largest alpha obtained is for the Carhart momentum factor, consistent with previous research on volatility management and equivalent to Moreira and Muir's (2017) results. Barroso and Santa-Clara (2015) find that the risk of momentum is both predictable and highly variable over time and discover that managing momentum eliminates crashes and nearly doubles the Sharpe ratio. Overall, our single-factor results look a lot like those presented by Moreira and Muir (2017). The positive and significant alphas imply that the volatility-managed portfolios expand the mean-variance efficient frontier, that is our risk-adjusted performance is better for the volatility-managed factors than it is for the original unmanaged factors. To quantify the relative increase in Sharpe ratios, we report annualized Appraisal ratios. Our Appraisal ratios signify large utility gains for all factors except HML (-0.32). The largest Appraisal ratio belongs to the Carhart momentum factor (0.80), while all the others range between (0.28 - 0.71).

Table 1
Volatility-Managed Factor Alphas

In **Panel A**, we run time-series regressions of each volatility-managed factor on the non-volatility-managed factor in the following manner; $f_t^{VM} = \alpha + \beta f_t + \epsilon_t$. The volatility-managed factor f_t^{VM} is scaled by the factor's inverse realized variance in the preceding month $f_t^{VM} = \frac{c}{RV_{t-1}^2} f_t$. In **Panel B**, we use SMB, HML, UMD, LIQ and excess MKT as controlling factors in the regression. The data are monthly and the sample period is 1981 to 2020. Standard errors are reported in parentheses and are controlled for heteroskedasticity. The accompanying heteroskedasticity adjusted t-statistic indicates whether the factor is statistically significant. Alpha values are annualized in percent per year by multiplying monthly factors by 12. RMSE is reported in annualized percent per year to calculate Appraisal ratio, which is further multiplied by √12.

Panel A: Univariate Regressions						
	SMB	**HML**	**PR1YR**	**UMD**	**LIQ**	**EMKT**
SMB	0.76					
	(0.07)					
	11.65					
HML		0.75				
		(0.06)				
		12.52				
PR1YR			0.71			
			(0.07)			
			10.69			
UMD				0.75		
				(0.06)		
				13.33		
LIQ					0.78	
					(0.06)	
					13.46	
EMKT						0.77
						(0.06)
						12.64
Alpha	2.83%	-3.83%	9.42%	9.21%	3.78%	3.95%
	(1.68)	(1.92)	(1.92)	(2.04)	(1.56)	(2.4)
	1.69	-2.00	4.90	4.52	2.42	1.65
N	472	472	472	472	472	472
R^2	0.58	0.56	0.50	0.56	0.62	0.59
RMSE	34.44	41.64	40.68	45.24	33.84	45.84
AR	0.28	-0.32	0.80	0.71	0.39	0.30
Panel B: Alphas Controlling for Factors						
Alpha	3.27%	-2.60%	7.85%	7.09%	5.48%	2.80%
	(1.68)	(1.68)	(2.04)	(2.04)	(1.80)	(2.52)
	1.95	-1.55	3.85	3.48	3.05	1.11

In Panel B we test the robustness of our results by checking whether our alphas remain significant when we control for exposure to well-known risk factors. We use SMB, HML, UMD, LIQ and EMKT as controlling factors in the regression, and have purposely left out PR1YR in favor of UMD to handle multicollinearity. From appendix 3 we observe high correlation between PR1YR and UMD and including both may undermine the statistical significance of our results. Our results encounter marginal changes when we are controlling for other factors. Naturally, both the significance level and the magnitude of our alphas are reduced, but the strategy still produce positive alphas for all significant factors. The results thereby suggest that our volatility-managed portfolios increase utility for mean-variance investors operating in the Norwegian market.

Our market factor deserves special attention as it evidently is the most easily available factor to implement for average real time investors. Indeed, since this is purely a long-only portfolio it makes it more accessible compared to more sophisticated portfolios that either require access to shorting or for investors to create the factors themselves. Figure 2 exhibits our strategy's performance for the market factor relative to the original buy-and-hold strategy. The top panel plots the cumulative nominal returns of the volatility-managed market factor compared to the original buy-and-hold strategy from 1981 to 2020. We invest 1 NOK in 1981 and plot each strategy's cumulative returns on a log scale. The volatility-managed market factor realizes relatively steady gains, which accumulates to 254 NOK at the end of our sample versus 220 NOK for the buy-and-hold strategy.

The lower panels display the percentage drawdowns and the one-year rolling average returns of each strategy. We also report the maximum drawdown of each strategy explicitly. The drawdowns measure downside volatility and the maximum drawdown represents the largest loss from a peak of the cumulative return sustained by our portfolio. The maximum drawdown is an informative figure on the effectiveness of forming mean-variance portfolios in terms of mitigating the most extreme losses. The rolling returns are smoothed over several one-year periods to paint a more accurate picture of the performance of each strategy. These are particularly useful when we examine return behavior for specific holding periods throughout the history of our sample. It helps us determine when the volatility-managed strategy outperforms the buy-and-hold portfolio. The panels in Figure 2 illustrates the volatility-managed portfolio's outperformance of the market both in terms of cumulative returns and its response to downside volatility. The volatility-managed strategy is constructed to take more risk when volatility is low and less risk when volatility is high, which results in its largest losses being concentrated around times of low volatility. Noteworthy, this stands in contrast to the fact that large market losses tend to happen when volatility is high. When we compare our results from Figure 2 with our realized volatility plot in Figure 3, we see that the managed strategy avoids the most substantial losses during high volatility episodes like the global financial crisis.

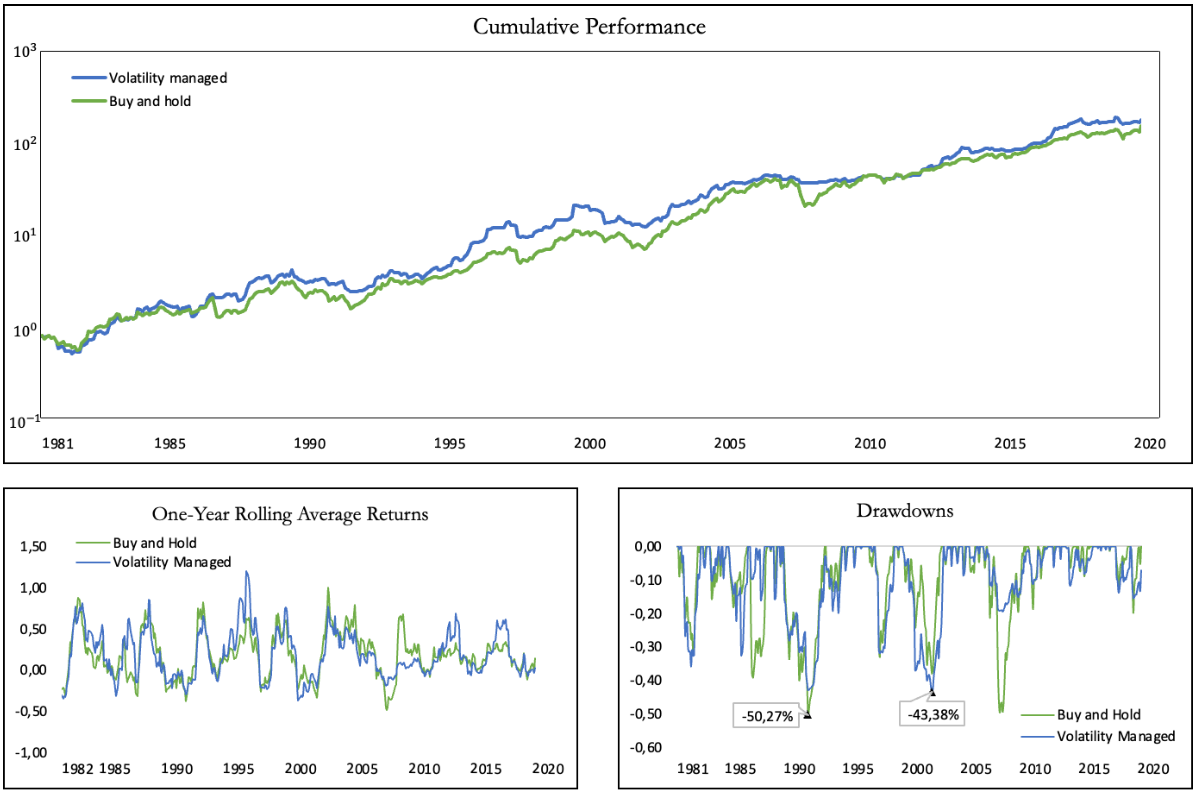

Figure 2. Cumulative returns to the volatility-managed market return. The top panel plots the cumulative returns to the buy-and-hold strategy versus the volatility-managed strategy for the market from 1981-2020. The y-axis is on a log scale and both strategies have the same unconditional monthly standard deviation. The lower left panel plots one-year rolling average returns from each strategy and the lower right panel shows the drawdown of each strategy. The maximum drawdown for each strategy is depicted explicitly in the lower right panel.

We evaluate the economic relevance of our results further by computing the percentage utility gains $\Delta U_{MV}(\%)$. Our results imply large utility gains for mean-variance investors. In fact, we report an increase in utility of 14.83% for the market portfolio stemming from volatility timing. Although this is considerably less than the 65% increase in lifetime utility reported by Moreira and Muir (2017), it still substantiates the economic relevance of our results.

To further assess the strategy's performance, we split the data into smaller subsamples. Moreira and Muir (2017) merely perform subsample tests for mean-variance efficient portfolios, but we find it appropriate to do this for the single factors as well. Testing the results across shorter time intervals makes our analysis more rigorous and helps identify specific periods where the managed strategy performs exceptionally well. We divide our initial sample into four 10-year subsamples and run the same univariate regressions as before (3). Since our sample period begins in July 1981 and ends in November 2020, the last subsample is 8 months shorter than the first three. We considered analyzing even shorter samples but noticed that the statistical power already was weak enough for the 10-year subsamples.

Table 2 unveils the managed strategy's performance across the four subsamples. The results are generally weaker, and the strategy do not perform as well for shorter time periods. None of the managed factor portfolios produce significant alphas in all subperiods. This could be a sign of weakness to the strategy, especially since longer investment horizons can be unrealistic for many investors. It is also particularly interesting that the managed market portfolio does not produce any significant alphas across all four subperiods. Additionally, for the first period from 1981-1991, none of our single factor alphas are statistically significant. Moreira and Muir (2017) achieve great results during recessions, which is consistent with the proclaimed features of the strategy; that it performs exceptionally well during times of high volatility. Our subsample analysis does not provide any clear indications of this outperformance, especially if we look at the 2001-2011 period. Notwithstanding, Figure 2 visualizes that the volatility-managed market portfolio outperforms the market during the global financial crisis.

Table 2
Single Factor Subsample Analysis

In this table, we run equivalent univariate time-series regressions of each volatility-managed factor on the non-volatility managed factor, but instead of using the full sample, we run the regressions across four 10-year subsamples (August 1981 - July 1991, August 1991 - July 2001, August 2001 - July 2011 & August 2011 - November 2020). The data are monthly and alpha values are annualized in percent per year by multiplying monthly factors by 12. The accompanying heteroskedasticity adjusted t-statistic indicates whether the factor is statistically significant.

Subsample Analysis - Single Factor Alphas						
	SMB	**HML**	**PR1YR**	**UMD**	**LIQ**	**EMKT**
1981 - 1991						
Alpha	0.80%	0.13%	3,21%	4.24%	5.00%	5.21%
T-statistic	0.28	0.05	1.07	1.38	1.40	1.01
N	120	120	120	120	120	120
1991 - 2001						
Alpha	5.96%	-10.31%	7.28%	11.63%	0.25%	5.42%
T-statistic	1.47	2.35	2.57	2.40	0.09	1.22
N	120	120	120	120	120	120
2001 - 2011						
Alpha	4.48%	-5.51%	9.83%	7.97%	3.68%	0.36%
T-statistic	2.02	1.70	3.17	2.63	1.39	0.11
N	120	120	120	120	120	120
2011 - 2020						
Alpha	-2.57%	7.37%	10.51%	13.69%	7.55%	5.38%
T-statistic	0.77	1.70	2.18	2.81	2.63	1.47
N	112	112	112	112	112	112

5.2 Multifactor Portfolios

In this section we report our results for the mean-variance efficient portfolios as we extend our analysis to a multifactor environment. The MVE portfolios are constructed using various combinations of the multiple factors. We consider four different factor combinations; one Fama-French three-factor model (1993), one that contains all factors, and two separate portfolios where UMD and LIQ are put together with the Fama-French three-factor model.

Table 3 reports our findings for the multifactor analysis. The table displays positive and significant alphas for all volatility-managed MVE portfolios, consistent with those of Moreira and Muir (2017). Alphas range from 2.77% to 3.63% and are generally much more significant than the single-factor portfolios. In fact, all alphas are significant at the 1% level. The largest alpha comes from the portfolio made up by all six factors, while the lowest belongs to the Fama-French three-factor model. Including UMD to the Fama-French three-factor model slightly increases the alpha to 2.80%, while including LIQ produces an alpha of 3.15%. The managed portfolios generate higher Sharpe ratios than their unmanaged counterparts and Appraisal ratios are high for all factor combinations. Appraisal ratios range from 0.30 to 0.91 and imply large utility gains for mean-variance investors. Additionally, since the weights are constructed in-sample, the original MVE Sharpe ratios are likely to be overstated relative to the truth. This implies that the increase in Sharpe ratios documented by the strategy is likely to be understated.

Table 3
Mean-Variance Efficient Factors

In this table, we create mean-variance efficient (MVE) portfolios by combining various factors. The sample can be thought of as the relevant information set for a given investor in the Norwegian stock market. We volatility time each of these MVE portfolios and report alphas obtained from regressing the volatility-managed MVE portfolio on the original MVE portfolio. Just like the volatility-managed single-factor portfolio, the volatility-managed multi-factor portfolios scale the portfolios by the inverse of its previous month's realized variance. Additionally, we report the annualized Sharpe ratio of the original MVE portfolios and the volatility-managed MVE portfolio, together with annualized Appraisal ratios. The factor combinations considered are MVE (All): SMB, HML, MKT, UMD, and LIQ, then the Fama-French three-factor model, and lastly the Fama-French three-factor model in combination with UMD and in combination with LIQ. All output is annualized and standard errors (reported in parentheses) are controlled for heteroscedasticity. The accompanying heteroskedasticity adjusted t-statistic indicates whether the factor is statistically significant.

Mean-Variance Efficient Portfolios - Full Sample Analysis					
	EMKT	**MVE (ALL)**	**FF3**	**FF3 + UMD**	**FF3 + LIQ**
Alpha	3.95%	3.63%	2.77%	2.80%	3.15%
	(2.40)	(0.60)	(0.84)	(0.72)	(0.72)
	1.65	6.06	3.29	3.89	4.38
N	472	472	472	472	472
R^2	0.59	0.63	0.64	0.64	0.65
RMSE	45.84	13.80	17.04	15.84	15.72
Original SR	0.78	1.54	1.30	1.38	1.32
Vol-Managed SR	0.83	1.79	1.42	1.51	1.50
Appraisal Ratio	0.30	0.91	0.56	0.61	0.69

It is important to make note of the feasibility of investing in these rather complex portfolios. The construction of them makes it unlikely for investors to be able to invest in real time unless investors are already engaged in the multifactor scheme. Accordingly, as pointed out by Moreira and Muir (2017); the notion behind this research is to see if volatility timing generates positive alphas and expands the mean-variance efficient frontier for investors already invested in multifactor portfolios.

As for the single factors, we analyze the MVE portfolios' performance across shorter time intervals. We construct equivalent 10-year subsamples with the same timespan as before. Table 4 presents our results, which vary across the four subsamples. The earlier and later periods document weaker and less significant alphas, while the two mid sectional periods obtain higher and more significant alphas. MVE (All) is the only portfolio that produces significant alphas in the first and the last subsample. All other MVE portfolios report insignificant alphas in these periods. From 1991-2001 and from 2001-2011 our subsample analysis produces positive and significant alphas for all MVE portfolios.

Table 4
Multifactor Subsample Analysis

In this table, we volatility time each of the MVE portfolios, but instead of using the full sample, we run the regressions across four 10-year subsamples (August 1981 - July 1991, August 1991 - July 2001, August 2001 - July 2011 & August 2011 - November 2020). We report alphas obtained from regressing the volatility-managed MVE portfolio on the original MVE portfolio across each of these subperiods. The factor combinations considered are MVE (All): SMB, HML, MKT, UMD, and LIQ, then the Fama-French three factor model, and lastly the Fama-French three factor model in combination with UMD and in combination with LIQ. All output is annualized and standard errors (reported in parentheses) are controlled for heteroscedasticity. The accompanying heteroskedasticity adjusted t-statistic indicates whether the factor is statistically significant.

Mean-Variance Efficient Portfolios - Subsample Analysis					
	EMKT	**MVE (ALL)**	**FF3**	**FF3 + UMD**	**FF3 + LIQ**
1981 - 1991					
Alpha	5.21%	2.95%	1.28%	1.52%	1.92%
T-statistic	1.01	1.96	0.73	0.89	1.17
N	120	120	120	120	120
1991 - 2001					
Alpha	5.42%	3.33%	5.32%	4.27%	5.46%
T-statistic	1.22	2.24	2.80	2.23	3.07
N	120	120	120	120	120
2001 - 2011					
Alpha	0.36%	3.06%	2.61%	2.50%	2.78%
T-statistic	0.11	3.01	2.32	2.28	2.69
N	120	120	120	120	120
2011 - 2020					
Alpha	5.38%	2.44%	1.13%	1.34%	1.66%
T-statistic	1.47	1.68	0.63	0.98	1.00
N	112	112	112	112	112

6. Discussion

Our results have so far been consistent with the findings of Moreira and Muir (2017). The strategy produces positive and significant alphas in most cases and appears to expand the mean-variance efficient frontier. In the forthcoming sections, we carry out a comprehensive validation process where we conduct an extensive battery of tests to evaluate the robustness of our results. These are complementary to our main analysis as they provide additional evidence on the performance of the volatility-managed strategy.

6.1 Business Cycle Risk

The scaling factor controls the average exposure to the strategy and diminishes in periods of high volatility as shown in Figure 1. This means that our strategy reduces the weight in risky assets and verifies that the volatility-managed strategy takes more risk when volatility is low and less risk when volatility is high. Figure 2 illustrates this further, as we see that the volatility-managed market factor has a lower standard deviation through recession periods like the global financial crisis where volatility was high. It is therefore already unlikely that business cycle risk is the explanation for our positive and significant results.

Table 5 makes this point more evident. We regress all managed factors on the original factors, while also including an additional interaction term consisting of a NORREC recession dummy. The coefficient on this term represents the conditional beta of our strategy on the original factor during recession periods relative to non-recession periods. Our results show lower beta coefficients for all factors when they interact with the recession dummy. This indicates that the volatility-managed strategy takes less risk during recessions. For instance, we find for the volatility-managed market factor a non-recession beta of 0.91 while the recession beta is -0.36, making the beta of our volatility-managed portfolio conditional on a recession equal to 0.55. The recession dummies' effect on our betas correspond to those of Moreira and Muir (2017), and it seems implausible that business cycle risk is the explanation for our results. This becomes even clearer, when we plot the time-series of realized volatility in Figure 3 as it reveals that all factor volatilities tend to rise in recessions.

Table 5
Recession Betas by Factor

In this table, we regress each volatility-managed factor on the non-volatility managed factor and include recession dummies $1_{rec,t}$ using NORREC recessions, which we interact with the original factors; $f_t^{VM} = \alpha_0 + \alpha_1 1_{rec,t} + \beta_0 f_t + \beta_1 1_{rec,t_t} x f_t + \epsilon_t$. This gives the relative beta of the volatility-managed factor conditional on recessions compared to the unconditional estimate. Standard errors are reported in parentheses and are controlled for heteroskedasticity. The accompanying heteroskedasticity adjusted t-statistic indicates whether the factor is statistically significant. We find that $\beta_1 < 0$, so that betas for each factor are relatively lower in recessions.

Business Cycle Risk - Recession Betas by Factor						
	SMB	**HML**	**PR1YR**	**UMD**	**LIQ**	**EMKT**
SMB	0.91					
	(0.09)					
	10.31					
SMBx1rec	-0.36					
	(0.12)					
	-3.14					
HML		0.82				
		(0.09)				
		9.34				
HMLx1rec		-0.20				
		(0.11)				
		-1.88				
PR1YR			0.81			
			(0.08)			
			10.70			
PR1YRx1rec			-0.22			
			(0.13)			
			-1.72			
UMD				0.88		
				(0.08)		
				10.66		
UMDx1rec				-0.32		
				(0.11)		
				-2.90		
LIQ					0.92	
					(0.07)	
					13.94	
LIQx1rec					-0.27	
					(0.10)	
					-2.63	
EMKT						0.92
						(0.07)
						13.94
EMKTx1rec						-0.30
						(0.11)
						-2.81
N	472	472	472	472	472	472
R^2	0.61	0.57	0.51	0.58	0.63	0.61

6.2 Volatility Comovement

Earlier on, when we extended our analysis to a multifactor environment, we imposed constant weights across factors to construct the MVE portfolios. To validate this approach, we normalize each factor by a common volatility factor, to illustrate the comovement among factors. Moreira and Muir (2017) argue that when realized volatility is highly correlated across factors, normalizing by a common volatility factor should not drastically change the results. To see this, we plot the time-series of monthly realized volatility for each factor, in Figure 3.

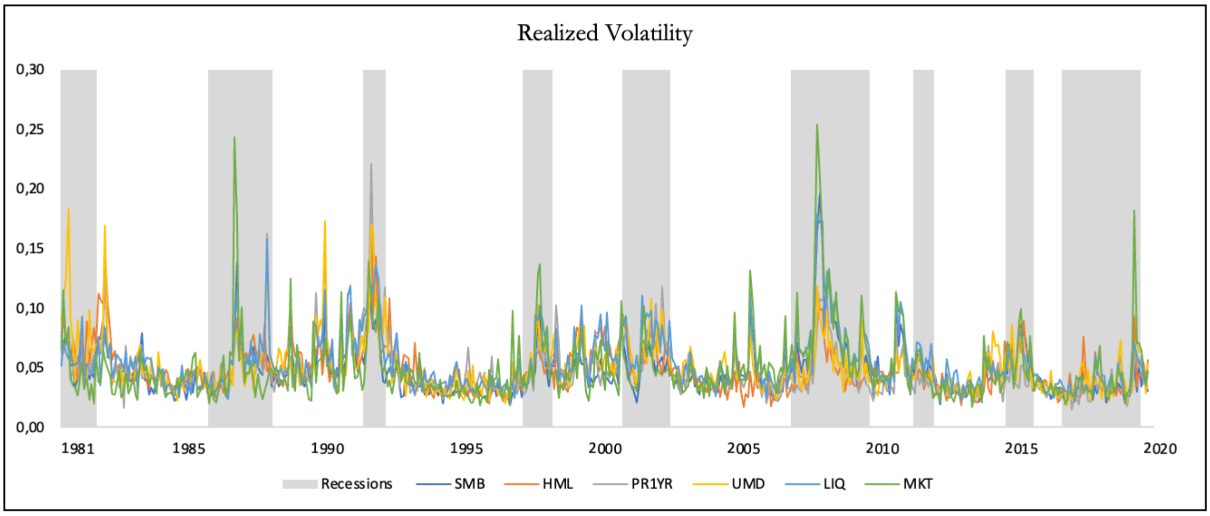

Figure 3. Time series of realized volatility by factors. This figure plots the time series of monthly realized volatility of each individual factor in Norway. We emphasize the common comovement in volatility across factors and that volatility generally increase for all factors in recessions. Light shaded bars indicate NORREC recessions and exhibits a clear business cycle pattern in volatility.

An interesting observation from Figure 3 is the existence of volatility clusters. We see that the realized volatility follows the same pattern across factors, which indicates a common comovement in volatility. The inclusion of NORREC recessions reveals a clear business cycle pattern in volatility, which generally rise for all factors in recessions. To quantify the visual representation of comovement between factor volatilities, we present the correlation matrix of realized volatility in Appendix 4. It shows high correlation between all factor volatilities, except the market factor, which is just moderately correlated with HML, PR1YR and UMD.

To confirm the visual observation of volatility comovement across factors, we compute the first principal component of realized variance across all factors and normalize each factor by $\frac{1}{RV_t^{PC}}$. We then run univariate time-series regressions of each managed factor on the unmanaged factors as usual. This stands in contrast to normalizing by each factor's own realized variance. The principal component analysis merges variables such that the fusion of variables contains a major part of the global variance, i.e., it maximizes the variance. Simultaneously, it merges the variables such that the fusion of variables minimizes the squared distance with the raw data. In this case, the raw data is the time series of realized variance and thus, the first principal component captures most of the variability in realized variance across factors. It explains how much of the variance contained in the data that is obtained from the first principal component.

Table 6
Normalizing by Common Volatility

In this table, we construct volatility-managed portfolios using a common volatility factor. Specifically, we construct volatility-managed strategies for each factor using the first principal component of realized variance across all factors. Each factor is thus normalized by the same variable, in contrast to normalizing by each factor's own past realized variance. We run time-series regressions of each volatility-managed factor on the unmanaged factor portfolio. Standard errors are reported in parentheses and are controlled for heteroskedasticity. The accompanying heteroskedasticity adjusted t-statistic indicates whether the factor is statistically significant. Alpha values are annualized in percent per year by multiplying monthly factors by 12. RMSE is reported in annualized percent to calculate Appraisal ratio, which is further multiplied by $\sqrt{12}$.

	SMB	HML	PR1YR	UMD	LIQ	EMKT
Alpha	1.38%	-1.51%	4.48%	3.60%	1.52%	1.85%
	(0.72)	(0.72)	(0.84)	(0.84)	(0.60)	(1.20)
	1.91	-2.10	5.33	4.29	2.53	1.54
N	472	472	472	472	472	472
R^2	0.94	0.94	0.92	0.93	0.94	0.92
RMSE	13.44	15.96	16.32	17.64	13.80	19.92
AR	0.36	-0.33	0.95	0.71	0.38	0.32

Table 6 reports the results from normalizing with the common volatility factor. The results are weaker than the main results and our alphas decrease more than what is documented by Moreira and Muir (2017). However, all positive alphas from the main results are still positive, and the excess market factor is the only factor that has become insignificant. All other factors are significant, and the common volatility timing seems to work about the same for most factors. The relatively strong comovement in factor volatilities validates the approach of imposing a constant weight across factors to construct the MVE portfolios.

6.3 Industry Portfolios

Moreira and Muir (2017) concentrate their analysis around well-established systematic factors. As a contribution to their research, we consider the possibility of implementing the strategy on other, less diversified portfolios. We conduct a supplementary analysis where we implement the strategy on disparate Norwegian industry portfolios. The portfolio construction procedure is once again applied by running new univariate time-series regressions of volatility-managed industry portfolios on the unmanaged industry portfolios.

Table 7, Panel A reports the results of managing volatility for the Norwegian industry portfolios. Almost all industries produce positive and statistically significant alphas. Material, Consumer Discretionary (ConsD), Consumer Staples (ConsS) and Finance all reveal large positive alphas, statistically significant at the 5% level. Energy, Health, and IT report positive alphas, statistically significant at the 10% level, while Industry is insignificant at all three levels. Material produces the largest alpha of 14.02% and Finance obtains the highest Appraisal ratio. These findings imply that managing volatility for Norwegian industry portfolios leads to higher risk-adjusted returns and utility gains for mean-variance investors. In Panel B we test the robustness by controlling for exposure to the systematic risk factors; SMB, HML, UMD, LIQ and EMKT. We get a disperse set of results, where Energy, Material and IT become insignificant alongside Industry. This implies that the excess return of these factors can be explained by exposure to some of our risk factors. All other industries remain significant, and Consumer Staples almost doubles its alpha from the exposure.

Table 7
Volatility-Managed Factor Alphas – Industry Portfolios

In **Panel A**, we run time-series regressions of each volatility-managed industry portfolio on the non-volatility-managed industry portfolio in the following manner; $I_t^{VM} = \alpha + \beta I_t + \epsilon_t$. The volatility-managed industry I_t^{VM} is scaled by the industry portfolio's inverse realized variance in the preceding month $I_t^{VM} = \frac{c}{RV_{t-1}^2} I_t$. In **Panel B**, we use SMB, HML, UMD, LIQ and excess MKT as controlling factors in the regression. The data are monthly and the sample period is 1981 to 2020. Standard errors are reported in parentheses and are controlled for heteroskedasticity. The accompanying heteroskedasticity adjusted t-statistic indicates whether the factor is statistically significant. Alpha values are annualized in percent per year by multiplying monthly factors by 12. RMSE is reported in annualized percent to calculate Appraisal ratio, which is further multiplied by √12.

Panel A: Univariate Regressions								
	Energy	**Material**	**Industry**	**ConsD**	**ConsS**	**Health**	**Finance**	**IT**
Energy	0.81							
	(0.05)							
	15.72							
Material		0.50						
		(0.19)						
		2.67						
Industry			0.78					
			(0.06)					
			12.24					
ConsD				0.67				
				(0.08)				
				8.33				
ConsS					0.74			
					(0.06)			
					11.76			
Health						0.40		
						(0.05)		
						7.52		
Finance							0.65	
							(0.07)	
							9.64	
IT								
								0.66
								(0.08)
								8.43
Alpha	4.74%	14.02%	1.40%	9.82%	7.48%	5.63%	8.30%	8.49%
	(2.52)	(5.52)	(2.76)	(4.08)	(2.64)	(3.24)	(3.24)	(4.92)
	1.88	2.54	0.51	2.41	2.83	1.74	2.56	1.73
N	472	472	472	472	472	472	473	474
$\mathbf{R^2}$	0.64	0.26	0.60	0.45	0.54	0.16	0.42	0.43
RMSE	56.04	122.40	53.04	90.24	57.84	92.88	63.36	102.36
AR	0.29	0.40	0.09	0.38	0.45	0.21	0.45	0.29
Panel B: Alphas Controlling for Factors								
Alpha	6.10%	4.09%	3.81%	9.86%	14.28%	11.82%	6.05%	6.83%
	(3.84)	(5.76)	(3.00)	(4.68)	(3.36)	(4.92)	(3.60)	(5.76)
	1.59	0.71	1.27	2.11	4.25	2.40	1.68	1.19

A noteworthy detail about our industry analysis is that the industries are less diversified than the systematic factors. By using alpha as our performance measure, we generally assume that the portfolios are well diversified. Since there are differences in diversification between the factor portfolios and the industry portfolios, one should be careful comparing alphas and Appraisal ratios directly.

6.4 Horizon effect

We have been implementing the strategy at a monthly rebalancing frequency, equal to the original paper (Moreira & Muir, 2017). The rebalancing frequency is of paramount importance because of the way our portfolios are constructed. When we scale each period's portfolio return by the inverse of the previous periods realized variance, we rely on the relationship between lagged volatility and current volatility. This relationship might vary across different rebalancing frequencies and brings up the question of how the strategy would perform at other frequencies. Executing the strategy on less frequent rebalancing periods can be useful to better understand the full dynamic between volatility shocks, expected returns and the price of risk. It allows us to conform our findings with the empirical theory that movement in variance and expected returns are countercyclical. However, lower rebalancing frequencies indicates repositioning based upon realized variance which is no longer representative to current variance. This can produce less accurate results. Nonetheless, more frequent rebalancing results in a more precise rebalancing.

To study horizon effects, we first look at the dynamics of risk and return through a vector autoregressive (VAR) model, which is suited to document how volatility and expected return respond to a volatility shock over a certain period. Like Moreira and Muir (2017), we run a VAR model at a monthly frequency with one lag of (log) realized variance and realized returns. We then plot the impulse response functions of the expected variance and expected return of the market portfolio for a shock to the realized variance, in Figure 4.

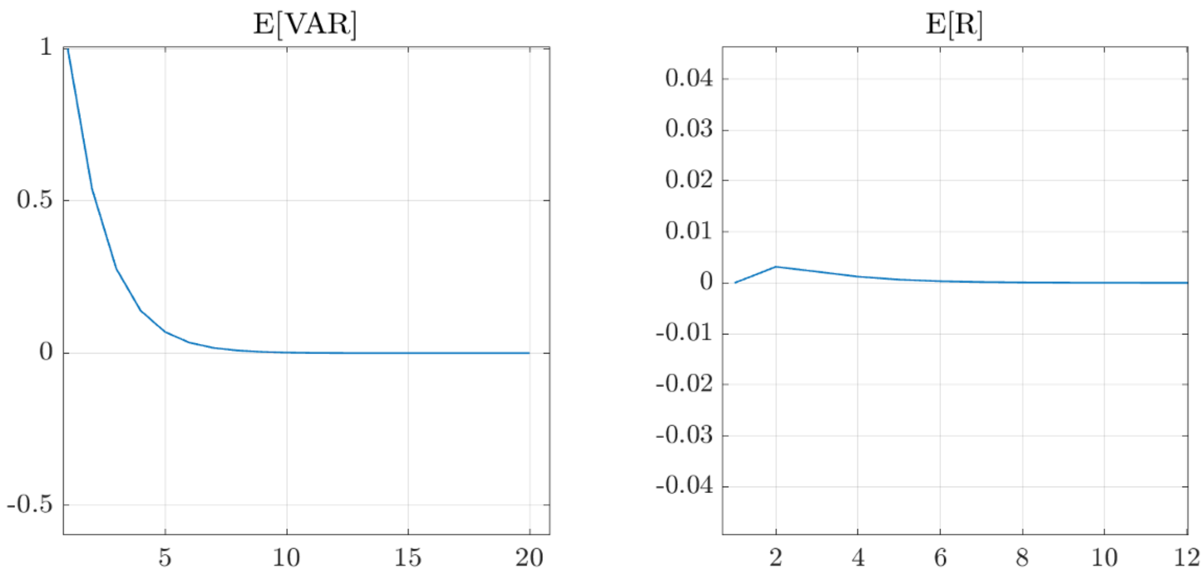

Figure 4. Impulse response of expected variance and expected return. This figure plots the impulse response of the expected variance and expected return of the market portfolios for a shock to the realized variance. The x-axis is in months.

Our results indicate that variance spikes on impact but fades away rather quickly at the rate of approximately 10 months, consistent with variance being strongly mean reverting. The expected return does not rise as much on impact, which is consistent with the findings of Moreira and Muir (2017). Although, our results show that the rise on impact is moderately lower, and that the elevation retrieves faster. These results indicate that the risk-return trade-off deteriorates on impact, but gradually ameliorates as volatility declines through a recession. However, as Moreira and Muir (2017) points out, the estimated response of expected returns to a volatility shock should be interpreted carefully, as return predictability regressions are poorly estimated.

To further evaluate the horizon effects, we repeat our analysis on a weekly, biweekly, quarterly, and annual basis. We form portfolios as before, but instead of scaling monthly returns by the inverse of their previous month's realized variance, we now use the respective rebalancing frequencies. To illustrate, a weekly rebalanced portfolio would be constructed by scaling weekly returns by the inverse of its previous weeks realized variance. When this is done for all abovementioned frequencies, we run similar time-series regressions and report the results in Figure 5.

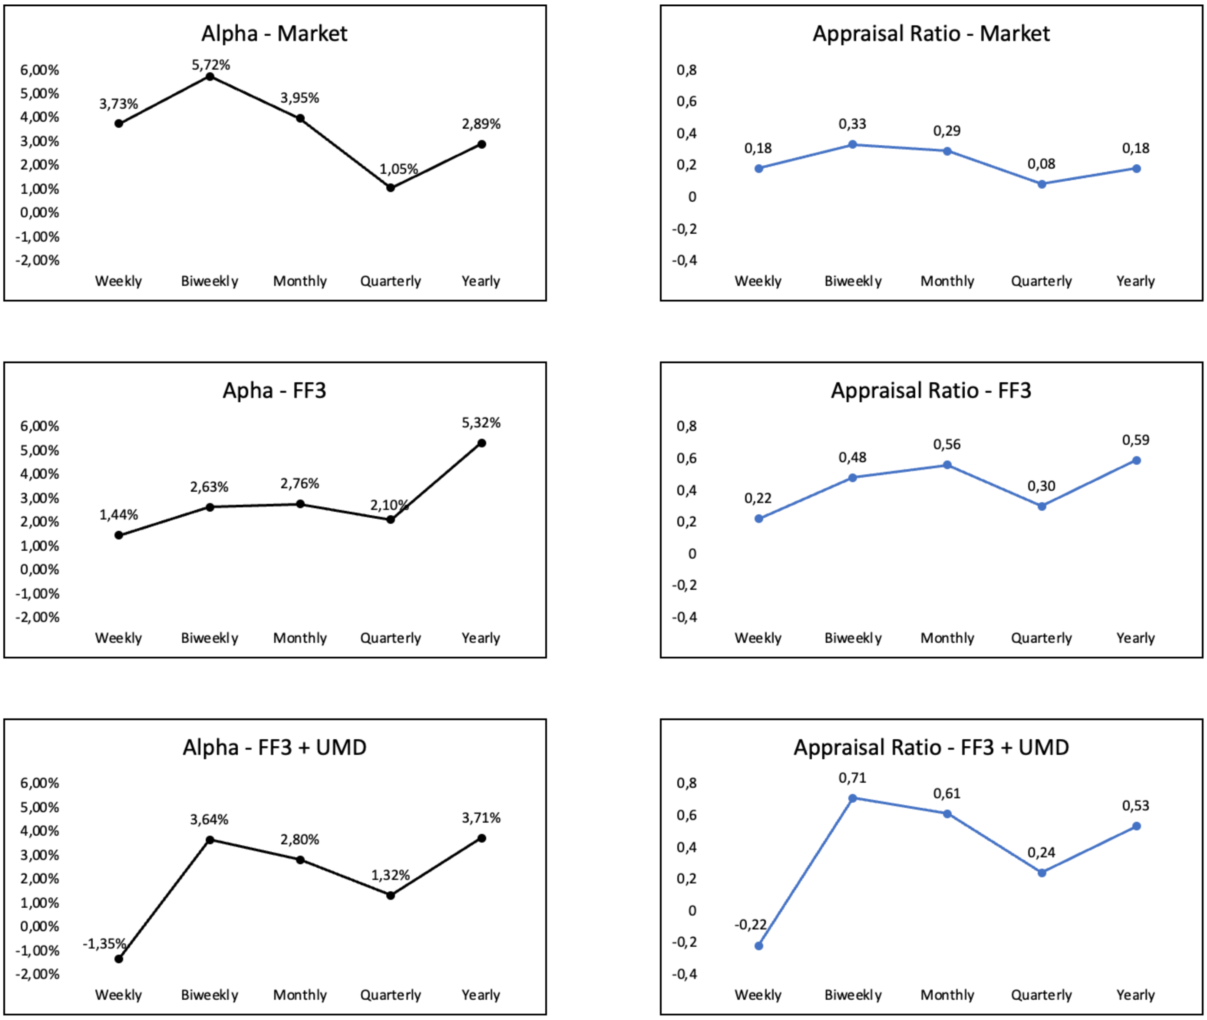

Figure 5. Results by rebalancing period. This figure plots the alphas and Appraisal ratios by rebalancing period. The scaled portfolios have been constructed using the inverse of the previous period's realized variance, customized to each frequency. The top two panels do this for the market, the middle panels for the MVE portfolio formed from the Fama-French three-factors, and the bottom panels add the UMD factor. All numbers are annualized for ease of interpretation.

Figure 5 exhibits the alphas and Appraisal ratios for the market, the Fama-French three-factor portfolio, and the Fama-French three-factor + UMD portfolio. We observe incoherent results for the weekly rebalancing, due to lack of data points. When dividing the data into weekly quantum's, we perceive three to five observations per week, which can lead to misrepresentation of the data. Given the data's poor portrayal at a weekly basis, no formal economic analysis can be performed. Biweekly rebalancing reports positive and significant alphas for all three portfolios. Contrarily, quarterly rebalancing produces lower and insignificant alphas in all three cases. Lastly, annual rebalancing returns high and significant alphas for the two MVE portfolios, but an insignificant alpha for the market factor. The biweekly rebalancing seems superior for the market portfolio, however this is before accounting for transaction costs. More frequent rebalancing implies higher transaction costs and might therefore affect the outcome of our results. Our results suggest that alphas decline with horizon, to the point where the rebalancing frequency is so low that it is unable to capture the effect of a volatility shock. The alphas thereby encounter a rise at an annual rebalancing frequency. These results are consistent with our VAR analysis.

6.5 Leverage Constraints

Up until this point, the strategy has been evaluated while deeming the market frictionless. Real life implications involve costs or restraints which we will look into in the following. In this section, we explore the importance of leverage and investigate how the volatility-managed strategy performs under certain leverage constraints. We consider various alternative strategies that capture volatility timing while simultaneously reducing trading activity. This involves using realized variance as before, realized volatility, and two leverage constrained strategies limited to no leverage and 50% leverage. Note that a 50% leverage constraint is consistent with standard margin requirements.

Table 8 documents the upper distribution of the weights for the volatility-managed market portfolio. The baseline strategy reports weights of 0.78, 1.31, 1.94 and 4.40 for the 75^{th}, 90^{th} and 99^{th} percentiles respectively. It uses modest leverage most of the time but employs ample leverage in the upper part of the distribution, when realized variance is low. Realized volatility makes the upper weights less extreme, with the 99^{th} percentile at 2.39 instead of 4.40. The results are corresponding to those of Moreira and Muir (2017). However, this is not the case when we impose leverage limitations on our strategy. The volatility-managed strategy turns out to be unobtainable without having access to leverage. While constructing the portfolios, we choose c so that the volatility-managed portfolio has the same unconditional standard deviation as the buy-and-hold portfolio, which is not possible without using leverage. A condition to the strategy is violated, and the strategy is no longer feasible. This constitutes limitations, as the average investor are unable to benefit from the volatility timing approach and can thus be deemed a weakness to the strategy.

Table 8
Volatility Timing and Leverage

In this table, we evaluate alternative volatility-managed strategies for the market portfolio and report alphas, Sharpe ratios, Appraisal ratios and distribution of weights used in each strategy. Specifically, we consider using inverse realized volatility instead of variance, using the original inverse realized variance but restricting risk exposure to be below 1 (i.e., with a no leverage constraint) or 1.5 (i.e., with a maximum 50% leverage). We focus on upper percentiles of weights to determine how much leverage is typically used in each strategy.

					Distribution of Weights w			
w_t	**Description**	α	**Sharpe**	**Appraisal**	**P50**	**P75**	**P90**	**P99**
$\frac{1}{RV_t^2}$	Realized Variance	3.95%	0.83	0.30	0.78	1.31	1.94	4.40
$\frac{1}{RV_t}$	Realized Volatility	2.58%	0.84	0.31	1.00	1.30	1.59	2.39
$\min\left(\frac{c}{RV_t^2}, 1\right)$	No Leverage	0.00%	-	-	1.00	1.00	1.00	1.00
$\min\left(\frac{c}{RV_t^2}, 1.5\right)$	50 % Leverage	2.66%	0.82	0.26	1.08	1.50	1.50	1.50

These findings contradict the ones made by Moreira and Muir (2017), who report that most investors can benefit from volatility timing even under tight leverage constraints. Our results for the Norwegian market imply that one can only benefit from the strategy by having access to leverage. Nonetheless, even with a 50% leverage constraint our alpha is not statistically significant. Moreira and Muir (2017) argue that the strategy can be realistically implemented in real time. Our findings indicate that the strategy requires leverage and real time implementation is not necessarily feasible.

6.6 Transaction Costs

Managing volatility is a complex process, often exacerbated by high transaction costs inherent in implementing the strategy. Transaction costs are of high relevance as they diminish returns and reduce alphas. The performance of the strategy therefore depends on the relative magnitude of these costs. To evaluate whether the strategy survives transaction costs, we consider various transaction cost assumptions and test the same alternative strategies as we did for the leverage constraints. Though, we exclude the no leverage strategy as it is not implementable. We use the same empirical transaction cost approximations as Moreira and Muir (2017) and evaluate the impact on our managed market factor. First, we look at transaction costs of 1 bp stemming from Fleming, Kirby, and Ostdiek (2003). Then, we consider 10 bps, as described by Frazzini, Israel, and Moskowitz (2014). Finally, to account for rising transaction costs during recessions, we examine the effect of 14 bps.

We report the average absolute change in monthly weights, expected return and alpha of each strategy before transaction costs. We then report the alphas after accounting for transaction costs and derive the implied trading costs in basis points needed to break even. Unsurprisingly, we only observe minor changes to our alphas when incorporating 1 bp of transaction costs, while higher levels of transaction costs evidently lead to larger reductions. All alphas remain positive, but none are statistically significant for transaction costs of 10 bps or higher. What is also particularly interesting is the amount of basis points needed to break-even, which are far less than what we observe in Moreira and Muir's (2017) evaluation. This indicates a sensitivity, which should be considered carefully. In general, the break-even costs are considered low and alphas insignificant at relatively moderate levels of transaction costs. These results suggest that the strategy fails to survive transaction costs and are thus inconsistent with those of Moreira and Muir (2017).

Table 9
Transaction Costs of Volatility Timing

In this table, we evaluate the volatility timing strategy for the market portfolio when including transaction costs. We also consider alternative strategies that capture the idea of volatility timing but significantly reduce trading activity implied by the strategy. Specifically, we consider using inverse realized volatility instead of variance, using the original inverse realized variance but restricting risk exposure to be below 1 (i.e., with a no leverage constraint) or 1.5 (i.e., with a maximum 50% leverage constraint). We report the average absolute change in monthly weights ($|\Delta w|$), expected return and alpha of each of these strategies. Additionally, we report the alpha of each strategy when including various trading costs (1 bps, 10 bps and 14 bps). Finally, the last column reports the implied trading costs in basis points needed to drive our alphas to zero in each of the cases. Alpha values and expected returns are annualized in percent by multiplying monthly factors by 12.

					α After Trading Costs			
w	**Description**	$\lvert\Delta w\rvert$	$E[R]$	α	**1bps**	**10bps**	**14bps**	**Break Even**
$\frac{1}{RV_t^2}$	Realized Variance	0.54	16.25%	3.95%	3.83%	2.75%	2.27%	32bps
$\frac{1}{RV_t}$	Realized Volatility	0.31	17.27%	2.58%	2.46%	1.38%	0.90%	21bps
$\min\left(\frac{c}{RV_t^2}, 1.5\right)$	50 % Leverage	0.33	16.55%	2.66%	2.54%	2.54%	1.46%	23bps

As to how realistic these measures are, there is certainly questions rising. One being the use of static transaction costs during the whole sample period, as literature generally suggests larger transaction costs in the past. It is also questionable whether these are realistic measures of transaction cost during times of market distress and high liquidity costs. We investigate this further in the next subsection related to liquidity concerns.

6.7 Liquidity Costs

The volatility-managed strategy is designed to make monthly adjustments based on movements in volatility. Low volatility means high risk exposure, while high volatility implies low risk exposure and makes the strategy downscale its weights. Since the strategy relies on predefined rules of execution, it presupposes that the market is invariably liquid. This is not always the case, especially when volatility spikes. Not being able to liquidate your positions in a timely manner and at reasonable prices, is what we often refer to as liquidity risk. In times when market volume dries up typically during times of distress, the lack of buy and sell orders can cause the market to fluctuate much more rapidly than usual. Volatility puts an extra strain on the liquidation process and could jeopardize the market value of the asset.

Inability to easily exit positions or selling securities well below market prices may affect the performance of the volatility-managed strategy. Moreira and Muir (2017) neglect these issues and do not consider liquidity costs when implementing the strategy. Traditionally, we find that academia has had little focus on this area, and many prefer to think of the market as perfect and frictionless. Unfortunately, this is not the case, and we find that liquidity risk is highly relevant to account for when analyzing volatility timing strategies. Excess transaction costs arising from liquidity concerns can become very high and reduce large amounts of profits. According to Jean-Philippe Bouchaud from Capital Fund Management, nearly two-thirds of trading gains evaporates due to liquidity problems (Bouchaud, et. al., 2018). This results in a blurred line between a profitable and unprofitable trading strategy when such losses are taken into consideration.

Whilst explicit transaction costs can be accounted for, implicit costs such as market impact, are difficult to estimate. In the previous subsection, our alphas became insignificant after accounting for moderate levels of transaction costs. Since Moreira and Muir (2017) already operate with fixed cost assumptions, we allow ourselves to do the same thing. To illustrate the possible impact of illiquidity we use the two-thirds assumption and test the volatility-managed strategy's performance after accounting for illiquidity effects. We use NORREC recessions as an indicator of market distress and remove two-thirds of profits during those times. Thereafter, in Figure 6, we plot the cumulative returns of the volatility-managed strategy after accounting for these losses versus the buy-and-hold strategy for the market portfolio. We report a positive but insignificant alpha of 2.07 %. However, Figure 6 indicates an underperformance of the volatility-managed strategy relative to the original buy-and-hold strategy after accounting for liquidity costs.

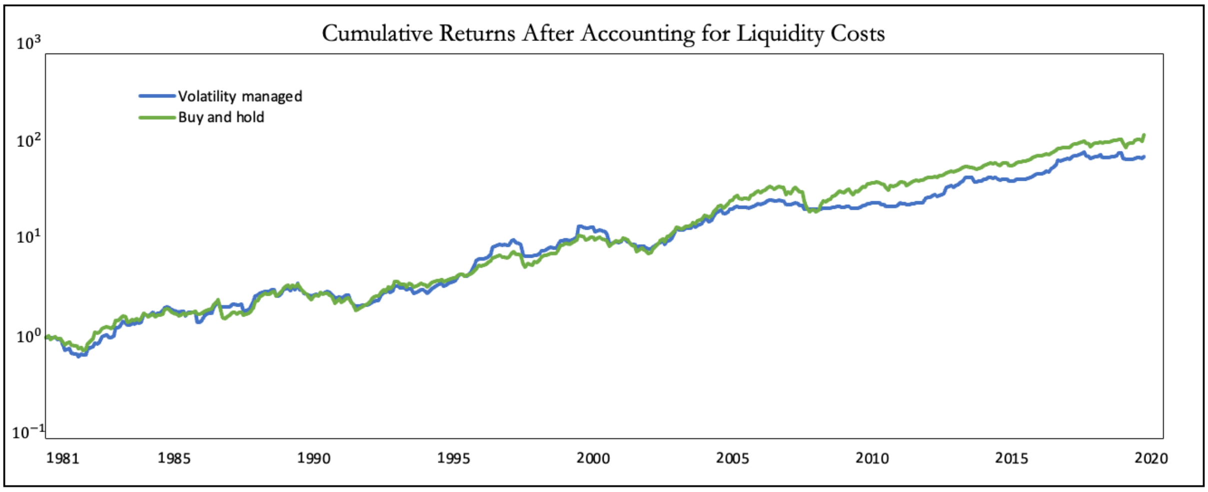

Figure 6. Cumulative returns to the volatility-managed market return after liquidity costs. This figure plots the cumulative returns to the buy-and-hold strategy versus the volatility-managed strategy for the market portfolio from 1981-2020 after accounting for liquidity costs. The y-axis is on a log scale and both strategies have the same unconditional monthly standard deviation.

We would like to emphasize the uncertainties related to the two-thirds assumption and that assuming static liquidity costs during all recessions might be unreasonable. Estimating the exact costs of illiquidity would be a complex task and since it is not the main purpose of our study, we leave that for future research.

7. Conclusion

In this paper, we replicate the methodology of Moreira and Muir's "Volatility-Managed Portfolios" (2017) on the Norwegian stock market. We investigate whether it is possible to benefit from volatility timing in smaller equity markets by constructing portfolios that scale monthly returns by the inverse of their previous month's realized variance. The strategy is tested on several systematic factors in Norway and intends to take advantage of variations in the risk-return tradeoff.

We have found positive and significant alphas that substantiate the strategy's eminence in a frictionless market. Our single factor portfolios produce large and significant alphas for all factors except HML, while our mean-variance efficient portfolios report positive and significant alphas for all factor combinations. The strategy does not seem to work well across shorter 10-year subsamples but outperforms the market during high volatility episodes. Business cycle risk seems unlikely to be the explanation behind our results and the strong comovement in factor volatilities validates the use of constant weights to construct MVE portfolios. We find evidence that supports the strategy's implementation on less diversified industry portfolios, and horizon effects suggest that alphas decline with time. However, market frictions exist, and implementing them has an enormous effect on our results. The strategy is unobtainable in the Norwegian stock market without the use of leverage, and even with 50% leverage, the strategy is unable to produce significant results. This also applies to transaction costs, as our results indicate that the strategy fails to survive even moderate levels of transaction costs. Finally, we find that after accounting for liquidity costs, the strategy underperforms relative to the buy-and-hold strategy.

After a meticulous study of the volatility-managed strategy's performance in the Norwegian stock market, it can be concluded that the strategy only performs well in unrealistic trading environments where all costs and restraints associated with transactions are non-existent.

8. Bibliography

Barroso, P., & Detzel, A. (2021). *Do Limits to arbitrage explain the benefits of volatility-managed portfolios?* Journal of Financial Economics, Vol. 140, No. 3, 744-767. https://www-sciencedirect-com.ezproxy.library.bi.no/science/article/pii/S0304405X21000775

Barroso, & Santa-Clara, P. (2015). Momentum has its moments. Journal of Financial Economics, 116(1), 111–120. https://doi.org/10.1016/j.jfineco.2014.11.010

Bouchaud, J., Bonart, J., Donjer, J., & Gould, M. (2018). *Trades, Quotes and Prices.* Cambridge University Press.

Carhart, M. (1997). On persistence in mutual fund performance. The Journal of Finance, 52(1), 57-82.
On Persistence in Mutual Fund Performance - Carhart - 1997 - The Journal of Finance - Wiley Online Library

Cederburg, S., O'Doherty, M., Wang, F., & Yan, X. (2019). *On the performance of Volatility-Managed Portfolios.* Journal of Financial Economics. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3357038

Cochrane, J. (2008). *Is Now the Time to Buy Stocks?* Wall Street Journal. https://www.wsj.com/articles/SB122645226692719401

Eitrheim, Ø., Grytten, & O., Klovland, J. (2007). *Historical Monetary Statistics for Norway – Part III*. Norges Bank https://norges-bank.brage.unit.no/norges-bank-xmlui/handle/11250/2506317

Fama. E., & French, K. (2015). A five-factor asset pricing model. Journal of Finaancial Economics, 116(1), 1-22.
A five-factor asset pricing model - ScienceDirect

Fama, F., & French, K. (1993). Common risk factors in the returns on stocks and bonds. Journal of Financial Economics, 33(1), 3-56.
PII: 0304-405X(93)90023-5 (ucsd.edu)

Fleming, Kirby, C., & Ostdiek, B. (2001). The Economic Value of Volatility Timing. Journal of Finance, 56(1), 329–352.
https://doi.org/10.1111/0022-1082.00327

Fleming, Kirby, C., & Ostdiek, B. (2003). The economic value of volatility timing using "realized" volatility. Journal of Financial Economics, 67(3), 473–509.
https://doi.org/10.1016/S0304-405X(02)00259-3

Frazzini, A., Israel, R., & Moskowitz, T. (2014) *Trading costs of asset pricing anomalies.* AQR Capital Management. Trading Costs of Asset Pricing Anomalies by Andrea Frazzini, Ronen Israel, Tobias J. Moskowitz :: SSRN

Liu, F., Tang, X., and Zhou, G. (2018*). Volatility-Managed Portfolios: Does It Really Work?* Journal of Portfolio Management, 46.
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3283395

Markowitz, H. (1952). *Portfolio Selection.* Journal of Finance, 7(1), 77-91.
https://www.jstor.org/stable/2975974

Moreira, A. & Muir, T. (2017). *Volatility-managed portfolios*. The Journal of Finance, 72(4), 1611-1643. https://www-jstor org.ezproxy.library.bi.no/stable/pdf/26652549.pdf?refreqid=excelsior%3A46bca2844abcfbd5f4fe139b3f8db14a

Moreira, A & Muir, T. (2019). *Should long-term investors time volatility?* Journal of Finance Economics, 131(3), 507-527.
https://www-sciencedirect-com.ezproxy.library.bi.no/science/article/pii/S0304405X18302745

Nagel, S., Reck, D., Stuart, B., & Slemrod, J. (2016). Who sold during the crash of 2008-09? Evidence from tax-return data on daily sales of stock. Who Sold During the Crash of 2008-9? Evidence from Tax-Return Data on Daily Sales of Stock – Frank Hawkins Kenan Institute of Private Enterprise (unc.edu)

Næs, R., Skjeltorp, J. A., & Ødegaard, B. (2011). *Stock market liquidity and the business cycle.* The Journal of Finance, 66(1), 139-176.
Stock Market Liquidity and the Business Cycle - NÆS - 2011 - The Journal of Finance - Wiley Online Library

Ødegaard, B. (2021). *Empirics of the Oslo stock exchange: Asset pricing results.* 1980-2015.empirics_ose_basics_2021_03.pdf (ba-odegaard.no)

9. Appendices

Appendix 1 - Table 1
Monthly Summary Statistics of Norwegian Factor Returns

Summary Statistics - Daily Observations						
	SMB	**HML**	**PR1YR**	**UMD**	**LIQ**	**MKT**
Mean	0.049 %	0.023 %	0.024 %	0.035 %	0.085 %	0.101 %
St.dev	0.011	0.012	0.012	0.012	0.013	0.013
Median	0.051 %	0.008 %	0.060 %	0.055 %	0.040 %	0.123 %
Max	8.61 %	7.57 %	10.26 %	8.09 %	13.64 %	11.37 %
Min	-10.81 %	-8.69 %	-17.25 %	-17.16 %	-8.83 %	-17.81 %
N	9894	9894	9894	9894	9894	9894

Appendix 2 - Table 2
Daily Summary Statistics of Norwegian Factor Returns

Summary Statistics - Monthly Observations						
	SMB	**HML**	**PR1YR**	**UMD**	**LIQ**	**MKT**
Mean	0.73 %	0.36 %	0.96 %	0.75 %	0.15 %	1.88 %
St.dev	0.044	0.053	0.048	0.057	0.045	0.059
Median	0.67 %	0.05 %	1.03 %	0.86 %	-0.009 %	2.35 %
Max	21.08 %	22.16 %	15.43 %	25.48 %	16.42 %	19.72 %
Min	-16.63 %	-19.64 %	-16.78 %	-24.27 %	-17.66 %	-23.79 %
N	473	473	473	473	473	473

Appendix 3 - Table 3
Correlation Matrix of Norwegian Factor Returns

Correlation Matrix – Monthly Observations						
	SMB	**HML**	**PR1YR**	**UMD**	**LIQ**	**MKT**
SMB	1	-0.190	0.120	0.110	0.560	-0.400
HML	-0.190	1	-0.001	-0.035	0.022	0.067
PR1YR	0.120	-0.001	1	0.780	-0.034	-0.140
UMD	0.110	-0.035	0.780	1	-0.048	-0.096
LIQ	0.560	0.022	-0.034	-0.048	1	-0.057
MKT	-0.40	0.067	-0.140	-0.096	-0.570	1

Appendix 4 - Table 4
Correlation Matrix of Realized Volatility

Correlation Matrix – Monthly Observations						
	SMB	**HML**	**PR1YR**	**UMD**	**LIQ**	**MKT**
SMB	1	0.65	0.62	0.60	0.84	0.79
HML	0.65	1	0.64	0.65	0.62	0.48
PR1YR	0.62	0.64	1	0.82	0.72	0.47
UMD	0.60	0.65	0.82	1	0.65	0.47
LIQ	0.84	0.62	0.72	0.65	1	0.75
MKT	0.79	0.48	0.47	0.47	0.75	1

Appendix 5 - Table 5
Volatility-Managed Factor Alphas - US

We run time-series regressions of each volatility-managed factor on the non-volatility-managed factor in the following manner; $f_t^{VM} = \alpha + \beta f_t + \epsilon_t$. The volatility-managed factor f_t^{VM} is scaled by the factor's inverse realized variance in the preceding month $f_t^{VM} = \frac{c}{RV_{t-1}^2} f_t$. The data are monthly and the sample period is from 1926 to 2016. Standard errors are reported in parentheses and are controlled for heteroskedasticity. Alpha values are annualized in percent per year by multiplying monthly factors by 12.

Univariate Regressions					
	EMKT	**SMB**	**HML**	**RMW**	**CMA**
EMKT	0.60				
	(0.02)				
SMB		0.60			
		(0.02)			
HML			0.54		
			(0.02)		
RMW				0.58	
				(0.03)	
CMA					0.67
					(0.02)
Alpha	4.89%	-0.52%	1.68%	2.70%	0.32%
	(0.13)	(0.07)	(0.09)	(0.07)	(0.05)
N	1073	1073	1073	629	629
R-squared	0.36	0.36	0.29	0.33	0.46