

Becker, Janis; Leschinski, Christian

Working Paper

Estimating the volatility of asset pricing factors

Hannover Economic Papers (HEP), No. 631

Provided in Cooperation with:

School of Economics and Management, University of Hannover

Suggested Citation: Becker, Janis; Leschinski, Christian (2018) : Estimating the volatility of asset pricing factors, Hannover Economic Papers (HEP), No. 631, Leibniz Universität Hannover, Wirtschaftswissenschaftliche Fakultät, Hannover

This Version is available at:

<https://hdl.handle.net/10419/200643>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Estimating the Volatility of Asset Pricing Factors

Janis Becker, Christian Leschinski¹

Leibniz University Hannover

This version: May 20, 2018

Abstract

Models based on factors such as size, value, or momentum are ubiquitous in asset pricing. Therefore, portfolio allocation and risk management require estimates of the volatility of these factors. While realized volatility has become a standard tool for liquid individual assets, this measure is not available for factor models, due to their construction from the CRSP data base that does not provide high frequency data and contains a large number of less liquid stocks.

Here, we provide a statistical approach to estimate the volatility of these factors. The efficacy of this approach relative to the use of models based on squared returns is demonstrated for forecasts of the market volatility and a portfolio allocation strategy that is based on volatility timing.

JEL-Numbers: C58, G11, G12, G17, G32

Keywords: Asset Pricing · Realized Volatility · Factor Models · Volatility Forecasting

¹Corresponding author:
Phone: +49-511-762-5383
Fax: +49-511-762-3923
E-Mail: leschinski@statistik.uni-hannover.de

1 Introduction

Volatility permeates finance, since it is central for everything from risk management to asset allocation. The fact that volatility is unobserved therefore poses a special challenge to practitioners that has been alleviated by the increased availability of high frequency data and the advent of realized volatility, which led to major improvements of volatility estimates relative to GARCH models. For economy wide risk factors, such as the size and value factors used in asset pricing, however, high frequency data is not available, so that one still has to rely on GARCH-type models.

In this paper we propose a methodology to overcome this issue and estimate factor volatility with a precision comparable to that of realized volatility estimates. This is achieved by constructing approximate high frequency returns of the respective risk factors.

There is a wide consensus that the cross section of asset returns is best described by factor models that proxy for economy wide risk factors. In addition to the established market, size, and value factors of Fama and French (1993), and the momentum factor of Carhart (1997), a plethora of anomalies has been uncovered in the literature that largely failed to attain the status of additional factors (cf. Stambaugh and Yuan (2016)). Recently, Fama and French (2015), Hou, Xue, and Zhang (2015), Stambaugh and Yuan (2016), and Fama and French (2017) suggest investment, profitability, and mispricing factors that subsume a large proportion of these anomalies.

In a simplified form these factors are typically constructed as follows. First, all stocks in the asset universe are sorted according to some firm characteristic. Then, two value weighted portfolios are formed from those stocks whose firm characteristics fall into the highest and lowest $x\%$ -quantile. The factor return is then obtained as the return from buying one of these portfolios and selling the other.

For risk management and portfolio formation purposes it is, however, not only the return but also the volatility of these factors that is of interest. Return volatility is a key variable for the pricing of options, speaks directly to the risk-return trade-off central to portfolio allocation, and even finds its way into government regulations.

For liquid individual assets the unobservability of volatility has been alleviated through the increased availability of high frequency data and the advent of realized volatility. Given that returns of the asset can be observed frictionless in arbitrarily small time intervals, realized volatility provides a consistent estimate of the quadratic variation of the stock return. For a review of these concepts cf. Andersen and Benzoni (2009).

While this approach is straightforward for individual assets, the calculation of realized volatilities for empirical factors is challenging. This is because the COMPUSTAT and CRSP data bases that are typically used to construct the factor returns do not provide

high frequency data. To calculate realized factor volatilities, it would therefore be necessary to match the stocks in these data bases with those from a high frequency data provider.

This, however, is not straightforward. High frequency data is typically only available for the most liquid stocks that are traded regularly in short time intervals. The CRSP portfolios that are used to construct empirical factor models, on the other hand, contain much more illiquid stocks that are simply not traded often enough to calculate realized volatilities. Furthermore, high frequency data bases are not necessarily free of survivorship bias, and finally – even if these hindrances would not exist – the matching of data bases typically constitutes a large effort and there tend to be non-negligible matching errors.

Practitioners or researchers that need to estimate factor volatilities are therefore restricted to one of two choices: either use squared returns as a volatility measure as for example in Moreira and Muir (2017), or estimate the underlying volatility process through a GARCH model. Both approaches have major drawbacks. Squared returns provide an unbiased but inconsistent estimate of the true variance and were the standard measure considered in the GARCH literature prior to the emergence of realized volatility. It is, however, well known that squared returns are extremely noisy. Andersen and Bollerslev (1998) show that, despite the high degree of persistence in stock return volatility, even the true model is only able to explain five to ten percent of the daily fluctuation in squared returns. Estimates based on GARCH models, on the other hand, will be less volatile but are biased and inconsistent if the model is misspecified.

The main contribution of this paper is therefore to propose an estimation method for factor volatility that is close in precision to realized volatility. Our approach is applicable whenever the researcher has access to daily factor return series and some high frequency data base. The idea is to approximate the factor return using a linear combination of the returns in the data base. In the first step, an appropriate linear combination is estimated using ridge regression. In the second step, the realized volatility of this approximate factor is calculated and used as an estimate for the volatility of the actual factor.

The details of this procedure are discussed in Section 2. Subsequently, Section 3 provides simulation results that demonstrate its favorable performance. The empirical validity and usefulness of this approach is demonstrated in a number of ways in Section 4. First, we analyze the relationship between our estimate and the squared returns for the factors considered by Fama and French (2015) and show that both are estimates of the same underlying volatility process. Second, we consider the example of the market factor where we can use the realized volatility of the S&P 500 to evaluate the accuracy of volatility forecasts. Here, we find that using our measure improves forecasts of the

factor volatility considerably compared to squared returns and GARCH-type models. Finally, we extend the analysis of Moreira and Muir (2017), who show that trading strategies based on timing factor volatility yield substantial alphas. It is shown that using our volatility estimate instead of the original approach based on squared returns improves the performance of this trading strategy. Conclusions are then discussed in Section 5.

2 Estimating Factor Volatility: Methodology

If asset returns are driven by a given factor model, then it holds true that the return of each asset is a linear combination of the returns of these factors and an idiosyncratic error term. Let there be K factors and denote the return of factor $k = 1, \dots, K$ at time t by f_{kt} . Then the return of asset i at time t according to this model is given by

$$r_{it} = \sum_{k=1}^K \lambda_{ik} f_{kt} + \varepsilon_{it}, \quad (1)$$

where $\varepsilon_{it} \sim (0, \sigma_\varepsilon^2)$, λ_{ik} is the loading of the i th asset on the k th factor, and $i = 1, \dots, N$. It is assumed that the ε_{it} have limited cross-sectional and serial dependence and that they are independent of all the λ_{ik} and f_{kt} .

Conversely, it follows that the return of each factor can be approximated by a linear combination of the asset returns. For suitable β_{ik} , we therefore have

$$f_{kt} = \sum_{i=1}^N \beta_{ik} r_{it} + \nu_{kt}, \quad (2)$$

where ν_{kt} represents the approximation error which can be expected to be small for large N , since the idiosyncratic errors ε_{it} in (1) average out.

The rationale behind this approach becomes clear if we rewrite model (1) for a vector of N assets. With $R_t = (r_{1t}, \dots, r_{Nt})'$, $F_t = (f_{1t}, \dots, f_{Kt})'$, $\lambda_i = (\lambda_{i1}, \dots, \lambda_{iK})'$, $\varepsilon_t = (\varepsilon_{1t}, \dots, \varepsilon_{Nt})'$ and $\Lambda = (\lambda_1, \dots, \lambda_N)'$, we obtain

$$R_t = \Lambda F_t + \varepsilon_t.$$

If Λ was known (and $\Lambda' \Lambda$ invertible), we could estimate F_t , by

$$(\Lambda' \Lambda)^{-1} \Lambda' R_t = F_t + (\Lambda' \Lambda)^{-1} \Lambda' \varepsilon_t = F_t + \varepsilon_t^*.$$

Since Λ is $N \times K$, and ε_t is $N \times 1$, ε_t^* is $K \times 1$. Therefore, every element of ε_t^* is a weighted

average of the innovation terms $\varepsilon_{1t}, \dots, \varepsilon_{Nt}$ and the vector ε_t^* converges to zero by a suitable law of large numbers (cf. Stock and Watson (2011) for a related discussion of cross-sectional averaging and statistical factor models).

The coefficient vector $\beta_k = (\beta_{1k}, \dots, \beta_{Nk})'$ in (2) corresponds to the k th row of the matrix $(\Lambda' \Lambda)^{-1} \Lambda'$. Since the returns f_{kt} of the observed factors are readily available, the problem in estimating β_k is that it is N dimensional and therefore potentially very variable if the time dimension T is not large enough. In fact, it is likely that $N > T$ in empirical applications, so that standard estimation methods cannot be applied.

Our objective is not to recover which stocks are part of the portfolios that are used to derive the factor returns. Instead, we want to obtain a good approximation of the factor returns in terms of mean squared error (MSE). We therefore resort to regularization and estimate β_k using ridge regression. The estimator is given by

$$\hat{\beta}_k = \arg \min_{\beta_{1k}, \dots, \beta_{Nk}} \left\{ \sum_{t=1}^T \left(f_{kt} - \sum_{i=1}^N \beta_{ik} r_{it} \right)^2 + \gamma \sum_{i=1}^N \beta_{ik}^2 \right\}, \quad (3)$$

with $\gamma > 0$. This is a least squares estimator with an additional penalty term that shrinks the coefficients towards zero. The size of the penalty term depends on the parameter γ that can be selected using cross validation. While the introduction of the penalty term introduces some bias, the rationale behind ridge regression is that for suitable γ , the reduction in variance outweighs the size of the bias, so that $\hat{\beta}_k$ is more accurate than the OLS estimator in terms of the mean squared error. Moreover, γ lowers the effective degrees of freedom, so that $N > T$ is permitted if γ is sufficiently large.

To obtain an estimate of the volatility V_{kt} , denote the m th of M intraday returns of asset i on day t by $r_{it}^{(m)}$. Then, from (1), we can approximate the m th intraday return of factor k on day t by

$$\hat{f}_{kt}^{(m)} = \sum_{i=1}^N \hat{\beta}_{ik} r_{it}^{(m)}. \quad (4)$$

Consequently, an estimator analogous to realized volatility is given by

$$\hat{V}_{kt} = \sum_{m=1}^M \left(\hat{f}_{kt}^{(m)} \right)^2. \quad (5)$$

We refer to $\hat{V}_{kt}$ as the Ridge-RV estimator. Obviously, $\hat{V}_{kt}$ is not consistent, since $\hat{\beta}_k$ is not a consistent estimate of β_k . However, if $\hat{f}_{kt}^{(m)}$ is a good approximation of the true unobserved high frequency return $f_{kt}^{(m)}$ of the k th factor, then $\hat{V}_{kt}$ is an approximation of the realized volatility of factor k .

To summarize, our method proceeds as follows

- 1) Regress the daily factor return f_{kt} on the daily returns of the N stocks in the data base to obtain the coefficient vector $\hat{\beta}_k$ from (3).
- 2) Obtain estimates $\hat{f}_{kt}^{(m)}$ of the intraday returns of the factors using (4).
- 3) Estimate the volatility of the factor from the estimated intraday returns $\hat{f}_{kt}^{(m)}$ using the Ridge-RV estimator in (5).

3 Monte Carlo Simulation

To demonstrate the usefulness of the Ridge-RV estimator, we conduct a simulation study that is tailored to resemble the setup in the empirical applications in Section 4.

It is well known that stock volatilities tend to have long memory and are well described by fractionally integrated processes (cf. Andersen et al. (2001)). A fractionally integrated process X_t is given by

$$(1 - B)^d X_t = v_t, \quad (6)$$

where B defined by $BX_t = X_{t-1}$ is the lag operator, v_t is a short memory process, and $-1/2 < d \leq 1$. The fractional difference operator $(1 - B)^d$ is defined in terms of generalized binomial coefficients. For details confer the original contributions of Granger and Joyeux (1980) or Hosking (1981). A process that fulfills (6) — such as the well known ARFIMA model — is referred to as $I(d)$. Standard short memory processes are included for $d = 0$ and unit root processes are obtained for $d = 1$.

To generate long memory in the daily volatilities V_{kt} of the K factors, we use the long memory stochastic volatility framework of Breidt, Crato, and De Lima (1998) and simulate T daily observations for each factor using

$$V_{kt} = \exp(X_{kt}), \text{ with } X_{kt} \sim ARFIMA(0, d, 0).$$

The log-volatilities therefore follow a fractionally integrated model. Applying the exponential function guarantees that all volatilities are positive. The V_{kt} obtained this way are used as the true daily volatilities.

Based on these, we subsequently draw M intraday factor returns $f_{kt}^{(m)} \sim N(0, V_{kt}/M)$ for each day and factor. The daily factor returns are obtained as $\sum_{m=1}^M f_{kt}^{(m)}$, so that they have variance V_{kt} . Using these intraday factor returns, we can simulate intraday returns

of N stocks. In analogy to Equation (1), the m th return of stock i at day t evolves as

$$r_{it}^{(m)} = \sum_{k=1}^K \lambda_{ik} f_{kt}^{(m)} + \varepsilon_{it}^{(m)},$$

with $\varepsilon_{it}^{(m)} \sim N(0, \sigma_\varepsilon^2/M)$ being a noise component. As for the daily factor returns, daily stock returns are obtained as the sum over the M intraday returns so that $r_{it} = \sum_{m=1}^M r_{it}^{(m)}$. All parameters are chosen such that the situation in our empirical application in Section 4 is replicated as closely as possible. This means we consider $K = 6$ factors whose correlation matrix matches the correlation matrix of the six factors considered there, we chose the memory parameter d to be 0.6 for all factors as the literature suggests the memory parameter of volatility to be in this region (cf. Wenger, Leschinski, and Sibbertsen (2018)), we simulate $M = 78$ intraday returns which corresponds to five minute stock data, the factor loadings λ_{ik} used for the simulation of stock returns are given by regression estimates of the factor loadings of $N = 500$ randomly chosen stocks that were in the S&P 500 at some point in the last 20 years, and σ_ε^2 evolves as the residual variance of this regression. Moreover, we set $T = 750$.

Based on this simulated data we then apply the procedure described in Section 2 based on Equations (3) to (5). Using the intraday factor returns $f_{kt}^{(m)}$, we can also compute the actual realized volatility. As a comparison, we further fit a GARCH(1,1) and a FIGARCH(1,d,1) model as proposed by Baillie, Bollerslev, and Mikkelsen (1996), and we consider the squared daily factor returns as an estimate of V_{kt} , too.

The results from 1,000 Monte Carlo repetitions can be found in Table 1 that shows the bias compared to the true volatility V_{kt} and the RMSE of all the procedures considered. As expected, the results are qualitatively similar for all factors and indicate the RV to be the best estimator. Our Ridge-RV estimator delivers only slightly worse results which are caused by a small downward bias due to the penalty term in (3).

This bias is, however, negligible compared to the variance of the estimates. The squared returns are unbiased, but their large variance leads to an RMSE that is several times larger than that of the Ridge-RV estimator. The GARCH model cannot remedy the noise problem, and is biased since it does not allow for long memory, but the data generating process is $I(d)$. The FIGARCH enhances on this problem but it is noisier than the GARCH model.

To summarize, in line with the literature the simulation study indicates that the RV is the best estimator for the volatility. In situations where the intraday returns of a portfolio cannot be observed, however, the Ridge-RV estimator is the best choice.

		F1	F2	F3	F4	F5	F6
RV	RMSE	1.042	1.195	1.088	1.278	1.612	1.238
	Bias	0.002	-0.005	0.007	0.000	-0.011	-0.006
Ridge-RV	RMSE	1.073	1.376	1.156	1.581	1.742	1.432
	Bias	-0.112	-0.194	-0.161	-0.233	-0.203	-0.183
Squared Return	RMSE	8.383	9.787	8.708	12.188	12.486	10.163
	Bias	0.000	0.014	-0.041	0.073	0.103	-0.010
GARCH(1,1)	RMSE	8.337	10.401	8.717	13.266	14.237	10.738
	Bias	0.998	1.234	1.042	1.440	1.708	1.234
FIGARCH(1,d,1)	RMSE	13.504	14.603	13.747	14.773	15.980	15.098
	Bias	0.520	0.315	0.421	0.235	0.074	0.209

Table 1: Simulation results: $\text{RMSE} \times 1000$ and $\text{Bias} \times 1000$ for different volatility estimation approaches. The true volatility processes of the six factors (F1, F2,...) evolve as $V_{kt} = \exp(X_{kt})$, with $X_{kt} \sim \text{ARFIMA}(0, d, 0)$. Moreover, the correlation matrix of the simulated processes matches the correlation matrix of the six factors considered in the empirical application.

4 Empirical Analysis

As an example for the application of the procedure described in Section 2, we consider the market (MKT), size (SMB), and value (HML) factors included in the 3-factor model of Fama and French (1993), the profitability (RMW) and investment (CMA) factors added in the 5-factor model of Fama and French (2015), and the momentum factor (MOM) included by Carhart (1997). These factors are commonly used in the asset pricing literature and their validity is widely accepted. Daily returns of these factors are freely available on the [homepage of Kenneth R. French](#).

In addition to the daily factor returns we require daily returns r_{it} and high-frequency returns $r_{it}^{(m)}$ for the estimation of (2) and the calculation of approximate 5-minute factor returns $f_{kt}^{(m)}$ from (4).

Since it is common to calculate realized volatilities from 5-minute returns, we extract five-minute prices of all stocks that were part of the S&P 500 at some point between 1996 and 2017 from the *Thomson Reuters Tick History* data base. This results in a total amount of 1,367 stocks that are considered. Since high frequency data is often subject to minor recording mistakes, it is common practice to apply some form of data cleaning. Here, we adopt the approach of Barndorff-Nielsen et al. (2009), which comprises, among other things, the removal of observations with negative stock prices and abnormal high or low entries in comparison to other observations on the same day.

Due to the long time span, it cannot be expected that the coefficients β_{ik} stay constant over time. The loading of individual stocks on factors can change as competitors are

	MKT	SMB	HML	MOM	RMW	CWA
R^2	97.72	76.60	84.51	88.00	79.99	87.17

Table 2: Ridge regression results: coefficient of determination R^2 in percent for the ridge regression as considered in (2).

acquired that have a different exposure to market risk, small firms grow into large firms, and growth stocks turn into value stocks as companies mature. We therefore conduct the estimation of the coefficient vector $\hat{\beta}_k$ according to (3) in a rolling window of size W . For the factors MKT, SMB, and HML which are based on firm characteristics that are relatively stable over time we set $W = 750$. The factors MOM and CMA that are based on more dynamic features are estimated in a window of size $W = 125$.

To demonstrate the empirical validity of our factor volatility estimates, the next section shows a number of model diagnostics. Afterwards, Section 4.2 demonstrates that volatility forecasts can be improved by using our measure and Section 4.3 presents an application to portfolio management that highlights the importance of factor volatility forecasting.

4.1 In-Sample Model Diagnostics

When trying to evaluate the performance of the Ridge-RV estimator, we face the problem that the true volatility process is unobserved and realized volatilities are not available for the factors. Only squared returns can be observed. We therefore consider a number of model diagnostics that demonstrate the satisfactory performance of our procedure, before turning to the applications in Sections 4.2 and 4.3.

The Ridge-RV estimate is based on the approximation of the factor of interest by a linear combination of stock returns. If this approximation in (2) is sufficiently accurate, so are those in (4) and (5). Consequently, we should obtain large coefficients of determination R^2 in the regression shown in (2). Table 2 shows that the measure ranges from 97.72 percent for the market factor to 76.60 percent for the size factor. Since the R^2 is above 75% for all factors the ridge regression seems to approximate the factor returns with sufficient accuracy.

Squared returns and Ridge-RV are both estimates of the same unobserved volatility process. They can therefore both be understood as differently perturbed versions of it. A first approach to test the validity of the Ridge-RV estimator in this empirical setup is therefore to test for fractional cointegration between the squared returns and $\hat{V}_{kt}$. Fractional cointegration is a natural generalization of cointegration to fractionally integrated series. Two time series X_t and Y_t are said to be fractionally cointegrated, if both are $I(d)$ and there exists a linear combination $X_t - \alpha - \beta Y_t = u_t$, so that u_t is $I(d-b)$

	MKT	SMB	HML	MOM	RMW	CMA	
CH	5.731	2.972	3.843	4.036	0.986	5.205	(1.697)
SRF	4.640	1.499	2.910	2.596	2.517	4.194	(1.960)

Table 3: Fractional cointegration test results: test statistics and critical values for the tests by Chen and Hurvich (2006) (CH) and Souza et al. (2017) (SRF). Here, the null of no fractional cointegration between log squared returns and log Ridge-RVs is tested against the alternative of fractional cointegration. The values in brackets are critical values at the five percent level.

for some $0 < b \leq d$. As in standard cointegration, both series must be highly persistent and they are (fractionally) cointegrated if a linear combination of them has reduced persistence. The extension lies in the fact that the reduction of persistence does not have to be from $I(1)$ to $I(0)$, but can be from $I(d)$ to $I(d-b)$.

When modeling volatility time series it is common practice to work with the log of the volatility series since it is better approximated by the normal distribution (cf. Andersen et al. (2001)). If $\ln \sigma_t^2$ denotes the true volatility process, then $\ln f_{kt}^2 = \ln \sigma_t^2 + \omega_{kt}$ and $\ln \hat{V}_{kt} = \ln \sigma_t^2 + \eta_{kt}$, where ω_{kt} and η_{kt} are the respective estimation errors. Therefore, if $\ln \sigma_t^2$ is $I(d)$, then $\hat{V}_{kt}$ can only be a reasonable estimator of $\ln \sigma_t^2$, if it is fractionally cointegrated with $\ln f_{kt}^2$, so that $\ln \hat{V}_{kt} - \ln f_{kt}^2 = \eta_{kt} - \omega_{kt}$ is $I(d-b)$.

Figure 1 plots the logarithms of the squared returns and our volatility estimate over time. Two main observations can be made. First, our measure is comoving with the squared factor returns, which is a first indication for the existence of a fractional cointegrating relationship. Larger values of the squared factor returns are associated with larger values of the Ridge-RV and vice versa. This holds for all factors and all time periods, except for the size factor where a short time period between 2004 and 2006 exists for which the two time series seem to diverge. Second, the Ridge-RV appears to be far less perturbed than the squared returns.

To formally test the hypothesis of fractional cointegration between both volatility measures, we apply the tests of Chen and Hurvich (2006) and Souza et al. (2017) for the null hypothesis of no fractional cointegration. Under the alternative a fractional cointegration relationship exists.

Table 3 reports the results of the tests. In line with Figure 1 the test by Chen and Hurvich (2006) rejects the null of no fractional cointegration for all factors, except for the RMW factor, and the test by Souza et al. (2017) rejects the null for all factors, except for the size factor. Therefore, we can conclude that squared returns and Ridge-RV are fractionally cointegrated.

All of the statistics presented so far show that our Ridge-RV estimator works well. However, as discussed above, all of the evidence provided is indirect, since the actual volatility process is unobserved. For the market factor, we can, however, conduct one




Figure 1: Time series plots of the logarithms of Ridge-RV and squared returns for the six factors.




Figure 2: Both plots display the Ridge-RV estimate of market factor volatility and the true volatility of the market factor approximated by the realized volatility of the S&P 500. While the left plot shows the two measures over time, the right plot displays a scatter plot.

experiment that provides insight into the actual accuracy of the Ridge-RV estimate. Even though we do not have realized volatilities for the market factor, it is well known that the value weighted CRSP return which is generally regarded as the best available market proxy is highly correlated with that of the S&P 500. The correlation coefficient is about 99 percent, meaning that the direction of the variation and its scaling over time is essentially the same. For the S&P 500 it is possible to obtain intraday prices, meaning that we can calculate realized volatilities. Consequently, we can compare our estimate of the market volatility with the realized volatility of the S&P 500. As Andersen and Benzoni (2009) stress, the realized volatility is the natural ex post measure of the underlying volatility process to consider. Figure 2 shows that the two measures are close to identical. In fact, they have a correlation of 93 percent, are fractionally cointegrated, and regressing our volatility estimate on the realized volatility yields an insignificant intercept and a slope that is almost one (0.98).

We therefore conclude that our estimate is appropriate for describing the volatility of the market factor. Even though the results in Tables 2 and 3 indicate that the procedure works slightly better for the market factor than for the other factors, the degree of precision obtained for the market implies that the Ridge-RV should still be a good estimate for the volatility of the other factors.

It should be noted, however, that the procedure is based on the assumption that the factors under consideration are actually relevant for the cross section of stock returns. This may be an issue if one wishes to apply the procedure to any of the many weak factors discussed in the literature.

4.2 Out-of-Sample Forecasts of Market Volatility

For portfolio allocation and risk management purposes, accurate forecasts are needed in addition to ex post and on-line estimates of the factor volatility. In this section we therefore compare the performance of forecasts using squared returns and GARCH-type models with those using Ridge-RV.

When trying to evaluate these forecasts, we again face the problem that the true factor volatility is unobserved. As shown by Andersen and Bollerslev (1998), considering squared returns as a proxy for the true factor volatility when evaluating volatility forecasts is not suitable since the tremendous amount of noise in the return generating process inevitably causes a poor performances of the forecasting models. On the other hand, it is seems tautological to show superior performance of our Ridge-RV measure when considering it as the true factor volatility. We therefore proceed as in the previous section and conduct a forecast comparison for the volatility of the market factor, where we can use realized volatilities of the S&P 500 to proxy for the true factor volatility. This makes for a fair comparison, since both types of models (Ridge-RV and models based on squared returns) do not use the realized volatilities of the S&P 500 in any way. The Ridge-RV is predicted using the HAR model of Corsi (2009). We refer to this forecast as the HAR-Ridge-RV model. As a benchmark, we also consider the standard HAR-RV model, which is possible for the market but not for the other factors. It can thus be interpreted as the "infeasible" model that we try to approximate when predicting factors such as SMB, HML, or others. As feasible benchmark models we include a GARCH(1,1) and due to the long range dependence in factor volatility we also use a FIGARCH(1,d,1) model fitted to the squared returns. All estimations are carried out in a rolling window of 750 observations.

For the evaluation of the forecasts we consider the RMSE and the QLIKE loss function, since Patton (2011) shows that these are the only commonly used measures that preserve the true ordering of the forecasts if they are evaluated on a perturbed volatility proxy. Furthermore, we report the R^2 from Mincer-Zarnowitz (Mincer and Zarnowitz (1969)) regressions given by $RV_{kt} = b_{0k} + b_{1k} \hat{V}_{kt}^{(h)} + u_{kt}$. Here, RV_{kt} is the observed realized volatility, $\hat{V}_{kt}^{(h)}$ the predicted volatility based on all information available in $t-h$ with h being the forecast horizon, and u_{kt} is an error term. Consequently, larger values of the coefficient of determination R^2 in this regression imply that the forecasts are performing better in predicting the true volatility.

Table 4 shows the results of this forecasting exercise for 1-step, 5-step, and 22-step forecasts. It can be seen that for all forecasting horizons and for all evaluation measures the HAR-Ridge-RV model performs better than all of the models based on squared daily returns. For 1-step forecasts, for example, the RMSE of the HAR-Ridge-RV model is

	1-step			5-steps			22-steps		
	RMSE	QLIKE	R^2	RMSE	QLIKE	R^2	RMSE	QLIKE	R^2
GARCH(1,1)	0.208	0.290	0.491	0.232	0.383	0.382	0.280	0.536	0.164
FIGARCH(1,d,1)	0.219	0.293	0.443	0.229	0.373	0.397	0.275	0.523	0.178
HAR-Ridge-RV	0.174	0.252	0.624	0.199	0.374	0.482	0.237	0.545	0.257
HAR-RV	0.172	0.208	0.604	0.201	0.324	0.469	0.238	0.486	0.265

Table 4: Forecast results: $\text{RMSE} \times 10^3$, QLIKE, and R^2 from Mincer-Zarnowitz regressions for the competing models and different forecast horizons. GARCH and FIGARCH use squared returns to forecast the market factor volatility, HAR-Ridge-RV uses the Ridge-RV estimate, and HAR-RV uses the true volatility given by the realized volatility of the S&P 500.

0.174, QLIKE is 0.252 and the R^2 is 0.624, while for the GARCH model, which is the best model using squared returns, the RMSE is 0.208, QLIKE is 0.29 and the R^2 is 0.491. As can be expected, the forecasting performance of the models becomes worse on longer horizons. The ranking of the models, however, stays the same.

When comparing the forecasts based on our volatility estimate with the HAR-RV forecasts based on the realized volatility of the S&P 500, it can be seen that the two models deliver qualitatively similar results.

Consequently, forecasts based on Ridge-RV achieve their objective to approximate those that are obtained if realized volatilities are available and they strongly outperform forecasts of the market volatility compared to models using squared returns. For factors other than the market, where realized volatilities are not available, they can therefore be expected to provide results that are far better than standard approaches. Whether the performance actually carries over to other factors is analyzed in an indirect way in the next section that also provides an example for the application of our method for practical purposes.

4.3 Volatility Timing Using Ridge-RV

As an illustration of the potential applications of Ridge-RV, we reconsider a volatility timing strategy recently proposed by Moreira and Muir (2017), who show that timing the volatility of the risk factors considered here can lead to substantial alphas. This is because the risk premia associated with the factors appear to be relatively stable over time, whereas their volatility exhibits considerable time variation.

Moreira and Muir (2017) consider the returns of a strategy that entails to scale monthly factor returns by the inverse of their previous month's volatility. Consequently, the strategy invests more heavily if volatility is low and stays out of the market if volatility is high. Since a regression of the unweighted factor returns on the volatility weighted

	MKT	SMB	HML	MOM	RMW	CWA
Monthly Re-Balancing						
MM17	6.49**	-0.12	1.31	8.70***	-0.25	-0.44
Ridge-RV	6.67**	1.53*	3.01*	9.92***	0.29	0.86
Daily Re-Balancing						
GARCH(1,1)	1.90	-0.10	0.15	9.34***	0.05	-0.99
FIGARCH(1,d,1)	2.10	0.00	0.84	9.28***	0.09	-1.07
HAR-Ridge-RV	1.25	1.93*	3.25*	9.81***	0.07	-0.08

Table 5: Volatility timing results: Annualized intercepts (alphas) of a time series regression with the volatility weighted factor return as endogenous variable and the ordinary factor return as exogenous variable. The first row reports the values of the original strategy with monthly re-balancing based on average squared returns. This is the benchmark. The second row contains the results of the same analysis based on our Ridge-RV measure. The last three rows correspond to daily factor re-balancing using volatility forecasts made by GARCH, FIGARCH and HAR-Ridge-RV. Moreover, *** (**) [*] indicates alpha to be significantly larger zero at the 1% (5%) [10%] level.

factor returns yields a significant intercept, this strategy expands the mean-variance frontier.

We extend the strategy in the following two ways using our estimate of factor volatility.¹ First, Moreira and Muir (2017) use a monthly volatility estimate obtained by averaging squared daily returns. As squared returns are extremely noisy, replicating their procedure using our less noisy Ridge-RV estimate should therefore be beneficial. Comparing the first two rows of Table 5 shows that this is the case for all six factors.

It can further be seen that alpha is significant for four out of six factors when using our measure while it is only significant for two out of six factors when using the measure based on daily squared returns.

Second, since volatilities exhibit considerable variation on a day-to-day basis, there is reason to assume that the volatility trading strategy might be even more successful if the strategy is executed with daily portfolio re-balancing. Therefore, we replicated the procedure using daily volatility forecasts from the HAR-Ridge-RV model and, as a comparison, using GARCH(1,1) and FIGARCH(1,d,1) forecasts.

The lower panel of Table 5 show that for all factors, except for the market factor, the alphas obtained when using our volatility measure are larger than those of the GARCH models. Moreover, the table reports that daily volatility timing leads to alphas that are significant for three out of six factors. Similar to the results obtained by

¹It should be noted that, due to the requirement of intraday data, all results are based on the period from 2002 until 2017. Therefore, the results presented here differ from those in Moreira and Muir (2017) as they use a longer time period.

the monthly strategy, alpha is particularly large for the momentum factor and not significant for the two more recently introduced factors RMW and CMA. Daily volatility timing using the Ridge-RV is further beneficial compared to the original strategy for five out of six factors. However, when comparing among the daily and monthly Ridge-RV based strategies, we find that less frequent portfolio re-balancing leads to better results. Furthermore, frequent trading would likely cause considerable transaction costs that negate the abnormal returns generated by the strategy.

The finding that the strategy improves when using our measure instead of squared returns for both the daily and the monthly horizon underlines that our measure is capable of forecasting factor volatility better than squared returns. In particular the improvements in the volatility timing strategy for the SMB and HML factors is evidence that the good performance of Ridge-RV for the market volatility documented in Sections 4.1 and 4.2 carries over to other factors.

5 Conclusion

Although the volatilities of economy wide risk factors such as the size and value factors of Fama and French (1993) are of importance for risk management and portfolio allocation purposes, the development of methods for their estimation has lagged behind that for liquid individual assets or indices, where intraday returns are available.

The Ridge-RV approach suggested in this paper circumvents the lack of high frequency data for factor returns and provides a volatility measure that is closely related to realized volatility. This is achieved by approximating the daily factor returns by a linear combination of the returns of assets for which intraday returns are available. Holding the weights in the linear combination constant then allows to obtain approximate high frequency factor returns that are the basis for the estimation of the factor volatility. Due to the large number of parameters in the linear combination that have to be estimated, it is necessary to apply a regularized estimation method such as ridge regression. Even though this introduces some bias, our simulations show that the bias is negligible in comparison to the reduction in variance relative to existing methods.

The subsequent applications to the market, size, value, momentum, investment, and profitability factors demonstrate that the proposed measure performs well in practice and outperforms competing approaches such as GARCH-type models. We therefore find that adopting the proposed approach has the potential for significant improvements in asset allocation decisions and risk management.

Acknowledgements

We want to thank Philipp Sibbertsen and Fabian Hollstein for their helpful comments. Furthermore, we are grateful for financial support from the University of Hannover through the program "Wege in die Forschung".

References

- Andersen, T. and L. Benzoni (2009). “Realized volatility”. In: *Handbook of Financial Time Series*. Springer, pp. 555–575.
- Andersen, T. and T. Bollerslev (1998). “Answering the skeptics: Yes, standard volatility models do provide accurate forecasts”. In: *International Economic Review*, pp. 885–905.
- Andersen, T., T. Bollerslev, F. Diebold, and P. Labys (2001). “The distribution of realized exchange rate volatility”. In: *Journal of the American Statistical Association* 96(453), pp. 42–55.
- Baillie, R., T. Bollerslev, and H. Mikkelsen (1996). “Fractionally integrated generalized autoregressive conditional heteroskedasticity”. In: *Journal of Econometrics* 74(1), pp. 3–30.
- Barndorff-Nielsen, O., P. Hansen, A. Lunde, and N. Shephard (2009). “Realized kernels in practice: Trades and quotes”. In: *The Econometrics Journal* 12(3), pp. C1–C32.
- Breidt, J., N. Crato, and P. De Lima (1998). “The detection and estimation of long memory in stochastic volatility”. In: *Journal of Econometrics* 83(1-2), pp. 325–348.
- Carhart, M. (1997). “On persistence in mutual fund performance”. In: *The Journal of Finance* 52(1), pp. 57–82.
- Chen, W. and C. Hurvich (2006). “Semiparametric estimation of fractional cointegrating subspaces”. In: *The Annals of Statistics* 34(6), pp. 2939–2979.
- Corsi, F. (2009). “A simple approximate long-memory model of realized volatility”. In: *Journal of Financial Econometrics* 7(2), pp. 174–196.
- Fama, E. and K. French (1993). “Common risk factors in the returns on stocks and bonds”. In: *Journal of Financial Economics* 33(1), pp. 3–56.
- Fama, E. and K. French (2015). “A five-factor asset pricing model”. In: *Journal of Financial Economics* 116(1), pp. 1–22.
- Fama, E. and K. French (2017). “International tests of a five-factor asset pricing model”. In: *Journal of Financial Economics* 123(3), pp. 441–463.
- Granger, C. and R. Joyeux (1980). “An introduction to long-memory time series models and fractional differencing”. In: *Journal of Time Series Analysis* 1(1), pp. 15–29.
- Hosking, J. (1981). “Fractional differencing”. In: *Biometrika* 68(1), pp. 165–176.
- Hou, K., C. Xue, and L. Zhang (2015). “Digesting anomalies: An investment approach”. In: *The Review of Financial Studies* 28(3), pp. 650–705.

- Mincer, J. and V. Zarnowitz (1969). “The evaluation of economic forecasts”. In: *Economic forecasts and expectations: Analysis of forecasting behavior and performance*. NBER, pp. 3–46.
- Moreira, A. and T. Muir (2017). “Volatility-Managed Portfolios”. In: *The Journal of Finance* 72(4), pp. 1611–1644.
- Patton, A. (2011). “Volatility forecast comparison using imperfect volatility proxies”. In: *Journal of Econometrics* 160(1), pp. 246 –256.
- Souza, M., V. Reisen, G. Franco, and P. Bondon (2017). “The estimation and testing of the cointegration order based on the frequency domain”. In: *Journal of Business & Economic Statistics*. DOI: 10.1080/07350015.2016.1251442.
- Stambaugh, R. and Y. Yuan (2016). “Mispricing factors”. In: *The Review of Financial Studies* 30(4), pp. 1270–1315.
- Stock, J. and M. Watson (2011). “Dynamic factor models”. In: *Oxford Handbook on Economic Forecasting*.
- Wenger, K, C. Leschinski, and P. Sibbertsen (2018). “The memory of volatility”. In: *Quantitative Finance and Economics* 2(1), pp. 137 –159.