

Forecasting Multidimensional Tail Risk at Short and Long Horizons

April 2017

Abstract

We define Multidimensional Value at Risk (MVaR) as a natural generalization of VaR. This generalization makes possible a number of important applications. For example, many techniques developed for VaR can be applied directly to MVaR. As an illustration, we employ VaR forecasting and evaluation techniques. One of our forecasting models builds on the progress made in the volatility literature and decomposes MVaR into long-term trend and short-term cycle components. We compute short- and long-term MVaR forecasts for several multidimensional time series and discuss their (un)conditional accuracy.

Keywords: Multidimensional Risk, Multidimensional Value at Risk, Two-Factor Decomposition, Long Horizon Forecasting

JEL classification: C52; C53.

1. Introduction

The interest in multidimensional tail (MT) events is driven by its importance in economics, finance, insurance and in many other areas of applied probability, statistics and decision theory. In economics and finance, modeling and forecasting MT events is paramount for many important applications such as portfolio decisions (e.g., Ang and Bekaert, 2002), risk management (e.g., Embrechts et al. 2002; Meine, et al. 2016), multidimensional options (e.g., Cherubini and Luciano, 2002), credit derivatives, collateralised debt obligations and insurance (e.g., Hull and White 2006; Kalemanova et al. 2007; Su and Spindler, 2013), contagion, spillovers and economic crises (Bae et al. 2003; Zheng, et al. 2012; Hautsch, Schaumburg and Schienle, 2015), systemic risk and financial stability (Adrian and Brunnermeier, 2016; Gonzáles-Rivera, 2014) and market integration (e.g., Bartram et al. 2006; Lehkonen, 2015).

Tail events are closely related to extreme risk that is generally defined as the potential for significant adverse deviation from expected results. In the univariate context, a measure of extreme risk widely used in practice is the Value at Risk (VaR). VaR is defined as the maximum loss on a portfolio over a certain period of time that can be expected with a nominal probability. However, modern risk management generally involves more than one risk factor and is particularly concerned with the evaluation and balancing of their impacts. For example, multifactor models (e.g., Chen et al., 1986; Ferson and Harvey, 1998) are used to measure and manage exposure to each of the multiple economy-wide risk factors.

This paper discusses a new angle on modeling and forecasting multidimensional tail events. Building on related recent literature (e.g., Prékopa, 2012; Polanski and Stoja, 2012; Torres, et al., 2015), we apply a generalized version of VaR, Multidimensional Value at Risk (MVaR), that is defined as a value that delimits a multidimensional tail with a nominal probability mass under a given density function. MVaR can be seen as an illustration of the multiple sources of risk: If VaR is a univariate risk measure, which instead of the variance takes into account the entire tail density, then MVaR is a measure of multidimensional risk that instead of the covariances takes into account the entire multidimensional tail.

Why should we care about MVaR when in typical portfolio applications it is the portfolio VaR that matters and not the multidimensional tail risk of the components of the portfolio? Although VaR might be the appropriate risk measure in portfolio applications, MVaR is useful in other circumstances where risk sources cannot be aggregated to form an informative risk measure or the portfolio interpretation of a collection of variables is not natural, useful or possible.

A prominent example of the importance of properly accounting for the distributional characteristics of the multiple sources of risk comes from stress testing of portfolios or financial systems. Typically, stress testing frameworks begin by developing individual scenarios with a negative outlook (tail events) for the evolution of certain economic drivers (e.g. GDP growth, interest rates, unemployment, stock market performance, investor sentiment) and then proceed to evaluate the impact of these on portfolios or

systemically important institutions (e.g., Bank of England, 2015). Treating these drivers individually presents a problem as they are obviously interdependent. Moreover, it would be difficult, if not nonsensical, to construct a portfolio of these factors and use its VaR as a tail risk measure. For example, what are the appropriate weights and their interpretations for each source of risk in such a portfolio? A sensible alternative in this case is to consider the sources of risk jointly. In this case, MVaR can considerably simplify the task.

Another example, related to stress testing that highlights the importance of MVaR is systemic risk. This is the risk of collapse faced by the financial system as a whole when one of its constituent parts gets into financial distress. Due to the interconnectivity of the financial institutions, a shock faced by one institution in the form of a tail event, increases the probability other financial institutions experiencing similar tail events, leading to a domino effect (e.g., Gai and Kapadia, 2010; Rogers and Veraart, 2013; Hautsch, Schaumburg and Schienle, 2014). In this case, it would be inappropriate and uninformative to treat the financial system as portfolio of banks and compute its VaR.

Therefore, while it is important to have a measure of the aggregate tail risk, often it is also important to know the direct dependence on, interrelationships among as well as the co-dynamics of the specific sources of tail risk. By focusing on the joint distribution of the individual sources of tail risks, we provide a framework to characterize the co-dependence of these risks.

An important advantage of MVaR is that techniques and applications developed for VaR can, in principle, be applied directly to MVaR. In this paper, we illustrate this with short- and long-term MVaR forecasting and evaluation. First, to obtain one-step ahead MVaR forecasts we employ the Conditional Autoregressive Value at Risk (CAViaR) of Engle and Manganelli (2004). However, CAViaR is a purely statistical model and does not distinguish between long-term and persistent movements in the tails, driven perhaps by the macroeconomic and company fundamentals, and transitory movements due to investor sentiment or other short-lived effects. With this in mind, we investigate a new two-factor forecasting model that we apply to MVaR. The model has several advantages. It is simple to estimate and it can easily produce multi-step ahead forecasts. Our Two-Factor Model (2FM) decomposes MVaR into a long-term trend and short-term cycle which can then be examined for relationships with economic and other variables. Finally, we use the scaling property of financial and economic time series to forecast MVaR at different frequencies and horizons. We evaluate the MVaR forecasts by employing adapted conditional and unconditional evaluation techniques of VaR forecasts. This paper is, to the best of our knowledge, the first to raise these issues in relation to (multidimensional) tail events.

2. Multidimensional Value at Risk

For the continuous and strictly increasing CDF F (PDF f) of a unidimensional random variable Y on the real line, the VaR at the nominal level a is usually identified with the quantile q_a for which $F(q_a) = a$. More generally, VaR can be defined as the cutoff q_a such that the probability mass under f of the interval $\{y \in R: y/d \geq q_a\}$ for a non-zero

number d is equal to a . Depending on the values of d , this definition can apply to the left ($d < 0$) or to the right ($d > 0$) tail of a distribution and it also allows for normalization.

In analogy to VaR, for a joint CDF F (PDF f) of a vector $\mathbf{Y} = (Y_1, \dots, Y_N)$ of N random variables on R^N with continuous and strictly increasing marginal CDFs, the Multidimensional Value at Risk (MVaR) in direction $\mathbf{d} \in R^N$ at the nominal probability level a is the unique cut-off value $q_a^{\mathbf{d}} \in R$ such that the set,

$$\mathcal{M}_a^{\mathbf{d}} = \{\mathbf{y} \in R^N: y_i/d_i \geq q_a^{\mathbf{d}}, \forall d_i \neq 0\}, \quad (1)$$

has probability mass a under f . We refer to the set $\mathcal{M}_a^{\mathbf{d}}$ as MVaR-region or multidimensional tail. In Figure 1, we illustrate the construction of the multidimensional tail $\mathcal{M}_a^{\mathbf{d}}$ as a Cartesian product of univariate tails (VaR-intervals),

$$\mathcal{M}_a^{\mathbf{d}} = \{y_1 \in R: y_1/d_1 \geq q_a^{\mathbf{d}}\} \times \dots \times \{y_N \in R: y_N/d_N \geq q_a^{\mathbf{d}}\},$$

where the probability mass for each VaR-interval $\{y_i \in R: y_i/d_i \geq q_a^{\mathbf{d}}\}$ can be computed from the corresponding marginal CDF.

[Figure 1]

We also say that $\mathbf{x} \in R^N$ is an extreme observation when \mathbf{x} lies in the MVaR-region. The directional vector \mathbf{d} (together with the significance level a) defines the region of interest

and it has also a distinct financial interpretation. For example, in the case of systemic risk the choice of the directional vector hinges on the particular economic metric of interest to the regulator. This could be, for example, how much the regulator may have to ‘pour into’ an institution in distress to prevent it from ‘infecting’ its counterparties, where Core Equity Tier 1 (CET1) capital as one of the most important macroprudential policy ratios for financial stability, is an obvious candidate. If a bank gets into distress and ‘eats up’ its CET1 ratio, then the regulator may be forced to bail it out by providing funding equal to CET1 to take the bank’s capital to its pre-distress level. Suppose a financial system is made up of three banks with CET1 of 2, 1 and 4. Then, a directional vector of particular interest for the regulator of this financial system is $d = (d_1, d_2, d_3)' = -(2, 1, 4)'$ as it succinctly represents the exposure of the economy to the systemic risk originating from these three banks and thus, the relative level of capital which the regulator may need to pour in to bail out these banks in case of their failure.

In spite of their conceptual simplicity, working directly with MVaR can prove challenging in higher dimensions. However, the relevant MVaR inference can be easily obtained by transforming points in the domain of f into scalars. Specifically, for each $\mathbf{x} \in R^N$ we define the point $\mathbf{x}^d \in R^N$ on the line along the directional vector $\mathbf{d} \in R^N$ as follows,

$$\mathbf{x}^d = v^d(\mathbf{x}) \cdot \mathbf{d}, \text{ where } v^d(\mathbf{x}) = \min_{d_i \neq 0} \{x_i/d_i\}. \quad (2)$$

We illustrate in Figure 2 and show in the Appendix the following property of the projection $v^d(\mathbf{x}) \in R$,

$$v^{\mathbf{d}}(\mathbf{x}) \geq q_a^{\mathbf{d}} \Leftrightarrow \mathbf{x} \in \mathcal{M}_a^{\mathbf{d}}. \quad (3)$$

Intuitively, observation \mathbf{x} lying in the MVaR-region implies that its projection $v^{\mathbf{d}}(\mathbf{x})$ exceeds the MVaR $q_a^{\mathbf{d}}$ and vice-versa.

[Figure 2]

3. Forecasting MVaR

For $N = 1$ and $d_1 = -1$, the multidimensional tail $\mathcal{M}_a^{\mathbf{d}}$ takes the form $\{y \in R: y \leq -q_a^{\mathbf{d}}\}$, i.e., $-q_a^{\mathbf{d}}$ is the a -quantile under the PDF f . Then, the a -quantile computed from a series of *i.i.d.* observations drawn from f is the natural estimator of (M)VaR $q_a^{\mathbf{d}}$. In higher dimensions, the MVaR $q_a^{\mathbf{d}}$ for $\mathbf{d} = -\mathbf{1}$ can be estimated in a similar manner as the a -quantile of the projections $v^{\mathbf{d}}(\mathbf{x}_t)$ of multidimensional observations \mathbf{x}_t . When estimating MVaR from projections, we omit the reference to the directional vector \mathbf{d} and write simply q_a .

In the remainder of this section, we apply three different MVaR forecasting methods to obtain forecasts over a horizon of k -steps ahead.¹ The methods presented in Subsections 3.1 and 3.2 are useful for forecasting daily MVaR one-step ahead and k -step ahead, where $k = 1$ and $k \geq 5$ refer to short- and long-term horizon forecasts, respectively. The method

¹ We also apply these techniques to VaR and find that the models do a similarly good job at forecasting VaR. As MVaR encompasses VaR, in order to preserve space we do not report these results. They are available upon request.

presented in Subsection 3.3 allows for forecasting low frequency (e.g., monthly) MVaR, which due to the limited number of such observations in practice would be difficult otherwise.

3.1. Conditional Autoregressive Value at Risk

Several approaches to short-term VaR forecasting have been proposed (e.g., Kuester et al., 2006; Nieto and Ruiz, 2016 for surveys of the VaR forecasting techniques). Some estimate the volatility of the time series first (e.g., by a GARCH model) and then compute VaR, often under the assumption of normality. Others use rolling historical quantiles (e.g., Boudoukh et al., 1998) or rely on extreme value theory (e.g., Danielsson and de Vries, 2000). Engle and Manganelli (2004) on the other hand, propose a different approach to VaR estimation and forecasting. Instead of modeling the whole distribution from heteroscedasticity-adjusted returns, they model the quantile directly from raw returns. As VaR is closely linked to volatility which is clustered in financial data, a natural way to model VaR is to use an autoregressive process. Engle and Manganelli (2004) specify the evolution of the quantile over time by the Conditional Autoregressive Value at Risk (CAViaR) model and estimate its parameters by quantile regression. CAViaR allows for many specifications of the autoregressive process which can be used for MVaR forecasting. In our empirical exercise in Section 4, we use their asymmetric slope function,

$$q_{a,t+1} = \beta_1 + \beta_2 q_{a,t} + \beta_3 \max(v_t^d, 0) - \beta_4 \min(v_t^d, 0), \quad (4)$$

where the next period quantile $q_{a,t+1}$ is a function of the current period quantile $q_{a,t}$ and projection v_t^d .

The quantile regression estimation of the parameter vector $\boldsymbol{\beta} = (\beta_1, \beta_2, \beta_3, \beta_4)$ in (4) boils down to the solution of the minimization problem,

$$\min_{\boldsymbol{\beta}} \frac{1}{T} \sum_{t=1}^T [a - I(v_t^d < q_{a,t})][v_t^d - q_{a,t}], \quad (5)$$

where $q_{a,t}$ is computed by (4), $I(\cdot)$ is the indicator function and a is the nominal probability. In our empirical study, we use CAViaR not as a competing, but as a complementary short-term MVaR forecasting model and obtain the k -step ahead forecasts of the MVaR $q_{a,t+k}$ with a technique that we present next.

3.2. Two-Factor Model

Similar to GARCH, CAViaR is a purely statistical model which cannot be easily related to macroeconomic or company fundamentals. However, tail events – similar to volatility – must be connected to fundamentals (see, for example, Bloom, 2009; Massacci, 2016). Moreover, evidence increasingly suggests that volatility is characterised by a multi-factor structure, with different dynamic processes governing its long-term and short-term dynamics. Engle and Lee (1999) introduce a component GARCH model which decomposes volatility into a permanent long-run trend component and a transitory short-run component that is mean-reverting towards the long-run trend. They find that a two-factor model provides a better fit to the data than an equivalent one-factor model (see also

Alizadeh et al. 2002; Brandt and Jones, 2006). Importantly, the two-factor specification makes possible linking the long-term trend of volatility to macroeconomic variables (e.g., Engle and Rangel, 2008). There is a significant number of VaR forecasting models in the literature but models that link VaR to macroeconomic fundamentals are as yet elusive. While perhaps the spline-GARCH model of Engle and Rangel (2008) may be extended to MVaR, it would be computationally demanding. The Two-Factor Model that we present here offers a simple and efficient way to decompose MVaR into a long-term trend and a short-term cycle. This decomposition would then allow for the linking of the long-term trend to macroeconomic and company fundamentals while the short-term cyclical component may be related to transient investor sentiment or other short-lived effects. For brevity, we do not pursue this idea in this paper but are investigating it in a separate project.

The finding that volatility has both a highly persistent factor and a strongly stationary factor has important implications for modeling and forecasting VaR. As VaR is closely related to volatility (e.g., Takahashi, Watanabe and Omori, 2016), any improvements in volatility forecasts are inherited by VaR forecasts. Motivated by the interpretation of two-factor volatility models, we explore an alternative, simple approach to modeling and forecasting MVaR over both short and long horizons. Specifically, we hypothesize that MVaRs follow a two-factor process given by

$$q_{a,t} = \varrho_{a,t} + \varphi(q_{a,t-1} - \varrho_{a,t-1}) + \varepsilon_t, \quad (6)$$

where $q_{a,t}$ is the long-term trend component of MVaR, $q_{a,t} - \varrho_{a,t}$ is the short-term cyclical deviation from the long-term trend and ε_t is a random error term with zero mean and constant variance. We assume that the long-term trend $\varrho_{a,t}$ is a stationary but highly persistent process but leave its precise dynamics unspecified. The parameter φ measures the speed of reversion of the cyclical component of MVaR to the long-term trend.

We implement the Two-Factor Model given by (6) in two steps. In the first step, we extract the long-run component $\varrho_{a,t}$ non-parametrically from the historical estimate of the a -quantile $\tilde{q}_{a,t}$. There are several techniques to extract the long-run component from a time series (see, for example, Durbin and Koopman, 2012). Here, we use the low-pass filter of Hodrick and Prescott (1997) which extracts a low frequency non-linear trend from a time-series and is often employed in applied macroeconomics. We also experimented with other filters such as the Christiano and Fitzgerald (2003) band pass filter and for some values of the oscillation parameters we obtained similar results.²

To implement the Two-Factor MVaR model with the Hodrick-Prescott filter, we set the smoothing parameter to the commonly used value of 100 multiplied by the squared frequency of the data, which for daily data (assuming 240 trading days per year) is 5,760,000 (see, for example, Baxter and King, 1999). In the second step, we estimate an autoregressive model for the cyclical component:

² To preserve space, we do not present the results of the MVaR forecasts with the Christiano-Fitzgerald (2003) filter. They are available upon request.

$$\tilde{q}_{a,t} - \tilde{q}_{a,t-1} = \varphi(\tilde{q}_{a,t-1} - \tilde{q}_{a,t-2}) + e_t, \quad (7)$$

where e_t is a zero mean random error. In order to forecast MVaR using the 2FM, we assume that the long-term trend follows a random walk over the forecast horizon, so that the k -steps ahead forecast $\hat{q}_{a,t+k} = \tilde{q}_{a,t}$ for all $k > 0$, and use the estimated autoregressive parameter from (7) to forecast the cyclical component. The k -step ahead MVaR forecast is therefore given by

$$\hat{q}_{a,t+k} = (1 - \tilde{\varphi}^k)\tilde{q}_{a,t} + \tilde{\varphi}^k\tilde{q}_{a,t} \quad (8)$$

This is a weighted average of the current estimate of the long-term trend $\tilde{q}_{a,t}$ and the current estimate of $\tilde{q}_{a,t}$. For the very long-term horizon, i.e., as $k \rightarrow \infty$, $\hat{q}_{a,t+k} \rightarrow \tilde{q}_{a,t}$, with a speed that is determined by the estimated coefficient $\tilde{\varphi}$.

3.3. Scaled MVaR

The fact that financial returns at lower frequencies can be computed as the sum of returns at higher frequencies suggests that we can use the latter to estimate the MVaR of the former. So far we have focused on the highest frequency which in our empirical section is one day. However, often risk forecasts at lower frequencies are needed. For example, Basel Accords require financial institutions to model risk using a 10-day holding period. It has become the industry standard to estimate daily VaR and then scale it up by $10^{1/2}$ in order to get the 10-day VaR. This is known as the square-root-of-time rule (SQRT-rule). The SQRT-rule originates in the scaling property of *i.i.d.* Gaussian variables X_1, \dots, X_k ,

$$X_1 + X_2 + \dots + X_k \stackrel{d}{=} k^{1/2} \cdot X_1$$

As the financial asset returns strongly violate the assumption of normality, neither moments of distributions (such as volatility) nor their quantiles should be scaled according to the SQRT-rule.³

Generally, the distribution of the random variables X_1, \dots, X_k displays a scaling behavior if it holds that,

$$X_1 + X_2 + \dots + X_k \stackrel{d}{=} k^\delta \cdot X_1,$$

where δ is the scaling exponent. Then, the a -quantile satisfies,

$$q_a(\sum_{i=1}^k X_i) = k^\delta \cdot q_a(X_1) \tag{9}$$

For many empirical distributions, the scaling property (9) is a good approximation only for nominal probability a sufficiently close to zero. For these distributions, one can estimate an extreme event at high frequencies for which there is an abundance of data (e.g., daily) and then use the scaling laws to estimate the extreme event at the lower

³ Indeed, the Basel Committee in its technical guidance paper (Basel Committee on Banking Supervision, 2002) no longer suggests that the SQRT-rule be used, but that “in constructing VaR models estimating potential quarterly losses, institutions may use quarterly data or convert shorter period data to a quarterly equivalent using an analytically appropriate method supported by empirical evidence”.

frequency of interest (e.g., monthly; see Mandelbrot, 1997; McNeil and Frey, 2000; Gabaix, 2009). Taking the logarithm of (9),

$$\ln\left(q_a\left(\sum_{i=1}^k X_i\right)\right) = \ln(q_a(X_1)) + \delta \cdot \ln(k), \quad (10)$$

makes it obvious why a straight line on the log-log plot is called a signature of scaling law.

4. Empirical Evaluation of MVaR Forecasts

4.1. Statistical Evaluation of MVaR forecasts

There is a vast number of alternative methods for evaluating VaR forecasts (see, for example, Nieto and Ruiz, 2016 for a recent review). Due to their intuitive appeal and popularity among practitioners, we focus in what follows on three simple and mutually complementary tests. Although these tests have been designed for testing VaR accuracy, they clearly also apply to the univariate projection series $(v^{\mathbf{d}}(\mathbf{x}_t))_{t=1}^T$.

Under the correct forecasting model, the proportion of MVaR violations, i.e., the proportion of projections $v^{\mathbf{d}}(\mathbf{x}_t)$ of observation \mathbf{x}_t that verify (3) should approach the nominal probability a for a sufficiently large sample. We refer to this procedure as unconditional accuracy. On the other hand, the conditional accuracy requires that the number of projections exceeding MVaR should be unpredictable when conditioned on past violations. In other words, MVaR violations should be serially uncorrelated. To assess both types of accuracy, we resort to the original unconditional accuracy test of Kupiec (1995)

and the test of independence by Christoffersen (1998), for consistency labeled here conditional accuracy test.⁴

The test statistic of the unconditional accuracy test of Kupiec (1995) is given by,

$$t_u = (\hat{a} - a) / \sqrt{\hat{a}(\hat{a} - a) / T} \quad (11)$$

where \hat{a} is the percentage of actual MVaR exceptions (violations), a is the nominal probability of exceptions and T is the number of observations. Intuitively, an unconditionally accurate model has an exception rate \hat{a} that is close to a .

The second, more stringent criterion regards the conditional accuracy. The likelihood ratio test of Christoffersen (1998) examines the serial independence of MVaR violations and is given by

$$LR_C = 2(\ln L_A - \ln L_0) \quad (12)$$

where,

$$L_A = (1 - \Pi_{01})^{T_{00}} \Pi_{01}^{T_{01}} (1 - \Pi_{11})^{T_{10}} \Pi_{11}^{T_{11}},$$

$$L_0 = (1 - \Pi)^{T_{00} + T_{10}} \Pi^{T_{01} + T_{11}} (1 - \Pi_{11})^{T_{10}} \Pi_{11}^{T_{11}},$$

⁴ Christoffersen (1998) proposes also a test of conditional coverage that simultaneously tests for unconditional and conditional accuracy. As we are interested in testing these hypotheses separately, we omit it here.

and

$$\Pi_{ij} = \frac{T_{ij}}{T_{i0} + T_{i1}},$$

$$\Pi = \frac{T_{01} + T_{11}}{T_{00} + T_{01} + T_{10} + T_{11}}.$$

T_{ij} is the number of times that state j follows state i . Here, state 0 obtains if no exceedence of MVaR forecast occurs and state 1 if such exceedence occurs. This statistic has an asymptotic χ^2 distribution with one degree of freedom, $LR_c \rightarrow \chi^2(1)$.

Engle and Manganelli (2004) remark that unconditional and conditional accuracy are necessary but not sufficient conditions to assess the performance of a quantile forecasting model. They construct an example where unconditional exceedances are correct and serially uncorrelated but the conditional probability of violation, given the quantile forecast, differs dramatically from the nominal level. Their dynamic quantile (DQ) test aims at avoiding such errors. Complementary to Kupiec (1995) and Christoffersen (1998) tests, we use a version of the DQ statistic to test the null that the conditional coverage, given the MVaR forecast, is equal to the nominal level a ,

$$DQ = \frac{\mathbf{hit}' \mathbf{q}_a \mathbf{q}'_a \mathbf{hit}}{a(1-a) \mathbf{q}_a' \mathbf{q}_a} \quad (13)$$

where \mathbf{hit} and \mathbf{q}_a are $T \times 1$ column vectors containing $hit_t = I(v_t^d < q_{a,t}) - a$ and the MVaR forecasts $q_{a,t}$, respectively. This statistic has an asymptotic χ^2 distribution with one degree of freedom, $DQ \rightarrow \chi^2(1)$.

4.2. Data

We use three different datasets to evaluate the performance of the MVaR forecasting models: the main US and European stock indices as well as EU bond indices. The US stock index dataset contains daily closing prices for S&P 500, Dow Jones and Nasdaq considered here as proxies for the performance of the underlying general sectors; the European stock index dataset contains daily closing prices of FTSE100 (UK), DAX (Germany), CAC40 (France) and MIB30 (Italy) used here as proxies for the health of respective economies. Finally, the European bond index dataset contains daily closing prices of 10 year government bonds considered here as proxies for country risk. From the raw prices, we compute the continuously compounded daily returns covering the period from 1 September 1996 to 31 October 2015, 5000 daily observations for each return series. We use the first 2000 observations for the initial estimation and the remaining 3000 observations for evaluating the out-of-sample forecasts in which the estimation window is rolled forward daily.

For each set of returns, we compute the corresponding vector of standard deviations \mathbf{SD} . The projection $v^d(\mathbf{x}_t)$ for each observation \mathbf{x}_t in this set is then computed by (2) for the directional vector $\mathbf{d} = -\mathbf{SD}$. Note that the projections would be identical if we computed them from the standardized returns using the directional vector $\mathbf{d} = -\mathbf{1}$. For consistency

with the VaR literature, we multiply each projection $v^d(\mathbf{x})$ by -1 so that more extreme negative returns correspond to lower values of $-v^d(\mathbf{x})$.

In what follows, we refer to the daily (k -days period) returns and MVaRs as frequency- l (frequency- k). For example, frequency-5 and frequency-20 MVaRs are computed from weekly and monthly returns respectively.

Table 1 reports summary statistics for the daily log return series for the sample. Panel A reports the mean, standard deviation, skewness, excess kurtosis and the Bera-Jarque statistic for the log returns and their projections. Panel B reports the first six autocorrelation coefficients, the Ljung-Box Q statistic for autocorrelation up to six lags for the projections and the p-values. All series are highly non-normal with negative skewness and positive excess kurtosis. The excess kurtosis for bond returns is almost half that of the stock returns. The projected series are highly autocorrelated and have, by construction, different empirical properties from the returns from which they originate.

As discussed above, asset returns are conditionally heteroscedastic. Therefore, to account for this feature of the data, we have also performed the analysis for returns standardized by the square root of the volatility obtained from a GARCH(1,1) model estimated over a rolling window. To preserve space, we present them in Tables 1A – 3A in the Online Appendix.

[Table 1]

Panel A of Figure 3 plots the projected US stock index returns (US Projections) and their “realized” daily MVaR over the period 2 January 2012 to 31 October 2015. The “realized” MVaR is estimated as the historical fifth quantile in the estimation window rolled forward daily. It is clear that the “realized” MVaR is slowly evolving. Panel B plots the same “realized” MVaR (note the different scale from Panel A) together with its long-term trend estimated using the Hodrick-Prescott filter over the sample. The trend is a smoothed version of the “realized” MVaR and closely tracks it although there are periods, for example during 2013, when the deviation is evident. Panel C of Figure 3 plots the resulting cyclical component of the “realized” MVaR using the Hodrick-Prescott filter. It is clear that the long-term trend in MVaR is time-varying and highly persistent, while the cyclical component is strongly mean-reverting, lending support to the two-factor representation of MVaR.

[Figure 3]

Figure 4 shows log-log plots of the frequencies $\{2^i\}_{i=0}^7$ days (x-axis) vs. the empirical frequency- 2^i MVaR estimates for US projections (y-axis) and the corresponding fitted straight lines. Estimates of frequency- 2^i MVaRs have been computed from non-overlapping intervals of length 2^i , $i = \{0, \dots, 7\}$ (i.e. one day to 6.4 months) spanning the whole sample of 5,000 observations. We find a good linear fit for all our datasets which indicates scaling in the tails of the projected return distributions. For the US (EU) projections and for $\alpha = 1\%$, 2.5% and 5% the scaling exponents δ are 0.52 (0.53), 0.56 (0.57) and 0.59 (0.55) respectively, implying that the underlying distributions have fat

tails. These estimates differ markedly from the estimates of around 0.42 in Hauksson et al. (2001) for the univariate VaRs.

[Figure 4]

5. Results

The out-of-sample MVaR estimation is performed using the last 3000 observations. For the out-of-sample forecasts, we moved a window of $T = 2000$ observations along the time axis. For each window $\mathbf{v}_t^d = (v_{t-T+1}^d, \dots, v_t^d)$ where $t = T, \dots, 3000 + T$, we first estimate the parameter vector $\boldsymbol{\beta}$ in (4) by solving the minimization problem (5) numerically and φ in (7) by a simple regression of the deviations $\tilde{q}_{a,t} - \tilde{Q}_{a,t}$ on their one-lagged values. For each window, we compute also frequency- k returns in non-overlapping intervals of length $k = 2^i$, $i = \{0, \dots, 4\}$, within this window. From these returns, we estimate the frequency- k MVaR by the relevant quantiles and the scaling exponent by regressing the frequency- k log-MVaR on the log-frequencies $\log(k)$.

Subsequently, we use the estimated parameters to obtain MVaR forecasts as follows. For the CAViaR and 2FM, the k -day ahead forecast $\hat{q}_{a,t+k}$ of the daily MVaR is given directly by (4) and (8) respectively, where in the case of CAViaR $k = 1$. Finally, for the Scaling Model the formula (9) delivers at date t a forecast of the frequency- k MVaR for the period $(t + 1, \dots, t + k)$ (i.e., weekly, monthly and quarterly MVaR for $k = 5, 20$ and 60).

The performance statistics for the MVaR forecasting models are presented in Tables 2 – 4. These tables report the actual exception rates (\hat{a}) as well as the p-values of the t_u , LR_c and DQ statistics to test the null hypotheses of unconditional and conditional accuracy for different MVaR specifications and nominal probability levels across the three datasets.

In line with previous evidence, CAViaR performs well for stock indices for one-day ahead forecasts both, conditionally and unconditionally. Indeed, the p-values of the t_u statistics indicate that the null of unconditional accuracy cannot be rejected for all three nominal probabilities. Further, the p-values of the LR_c and DQ statistics suggest that the conditional accuracy performance is satisfactory. The results for the bond return projections are the exception. In all three cases, CAViaR generates exceptions that are considerably below the required nominal probability a . Perhaps, this should be expected as CAViaR is a model for forecasting the quantiles of series that are more prone to tail events. Focusing on the stock indices datasets (Tables 2 and 3), there are differences in performance for different levels of a : it appears that CAViaR is more accurate for higher a . For example, in the case of US indices for $a = 5\%$ the actual exceedance rate \hat{a} is 5.2%, whereas for $a = 1\%$ the actual exceedance rate is 1.3%. This finding is similar to findings in the VaR literature (e.g., Kuester et al. 2006).

[Table 2]

[Table 3]

[Table 4]

The Two-Factor Model on the other hand appears to perform well for all three portfolios and at all nominal levels. At the longer end of the forecast horizon (60-day, i.e. approximately three months ahead), the forecast errors start to become considerable and the p-values of the LR_c statistics suggest that the conditional accuracy performance of the model is inadequate. However, for the shorter horizons, the performance on balance, seems acceptable. Interestingly, the performance of the 2FM appears more balanced with regard to α relative to CAViaR. For example, in the case of one-day ahead MVaR forecasts for US indices and $\alpha = 5\%$, the actual exceedance rate is $\hat{\alpha} = 5.6\%$, whereas for $\alpha = 1\%$ it is 1.7. However, in the case of European indices and for $\alpha = 5\%$ and 1%, these statistics are 5.4% and 1.3% respectively. This pattern can be observed for the longer horizon forecasts, although the relative errors of forecasts increase with horizon. For example, in the case of 60-day ahead MVaR forecasts for US indices and $\alpha = 5\%$ and 1%, the actual exceedance rates $\hat{\alpha}$ are 6% and 2.2%, while for the European indices these statistics are 5.8% and 1.7% respectively.

Importantly, the Two-Factor Model performs remarkably well unconditionally for the bond indices and it would appear that the forecasts are more accurate than in the case of stock indices. Moreover, the accuracy does not deteriorate substantially with horizon (Table 4). For example, in the case of one-day ahead forecasts for $\alpha = 5\%$ the actual exceedance rate $\hat{\alpha}$ is 5%, whereas for $\alpha = 1\%$ the exceedance rate is 1.2%. In the case of a 60-day ahead forecasts these statistics are 5.1% and 1.4% respectively. The errors are smaller for the shorter horizons. However, the conditional accuracy tests suggest that violations are serially correlated for the one-day and 60-day ahead forecasts for $\alpha = 1\%$

but they improve for the intermediate horizons. For $\alpha = 2.5\%$ the conditional accuracy does not appear to change much with horizon and for $\alpha = 5\%$ it improves slightly with horizon. Thus, on balance the Two-Factor Model produces unconditionally accurate MVaR forecasts for all datasets.

The Scaling Model delivers frequency- k MVaR forecasts of reasonable unconditional accuracy, especially for shorter periods, except perhaps for the bond return projections. However, the p-values of the Christoffersen (1998) test indicate that MVaR violations are highly serially correlated. This is not surprising given that we move a relatively long window of 2000 observations one day at each step. As a result, the resulting scaling forecasts change very slowly and cannot anticipate clusters of turbulence.

There is also an interesting performance discrepancy between bonds and stocks. For bonds, the Scaling Model consistently generates pessimistic forecasts with actual exception rates below the nominal ones. For stocks, on the other hand, the Scaling Model generates optimistic forecasts that are violated more often than they should. Somewhat surprisingly, the actual exception rate for $\alpha = 5\%$ tends to increase for longer periods. For example, for US indices the actual exception rates are 0.40, 0.45, 0.48, 0.53 and 0.52 for horizons of 1, 5, 10, 20 and 60 days ahead respectively. However, the Scaling Model forecasts in this empirical exercise should be treated with caution as the scaling exponents (slopes of the regression lines in the log-log plots) have been estimated in each window from five 2^i -MVaRs ($i = 0, \dots, 4$) only.

For all three models, we observe that the p-values of the DQ and the t_u statistics are well aligned (except in a few instances as e.g. for the 1% Scaling forecast 10-days ahead). However, there is no obvious relationship between the p-values of the DQ and LR_c statistics. The intuition for this regularity is exemplified by a constant forecast. If this forecast generates a correct unconditional coverage, then the DQ statistic (13) takes on the value of zero and thus, a p-value of one even if violations are serially correlated. On the other hand, an unconditional actual coverage that deviates significantly from the nominal level will lead to a large value of the DQ statistic (13) and thus, a low p-value.

An interesting question is how does the performance of the MVaR forecasts compare relative to that of the VaR forecasts of a portfolio made up of the same underlying series. We investigated this issue for an equally-weighted portfolio⁵ and found that, on balance, equally-weighted portfolio VaR forecasts are comparable with MVaR forecasts. Moreover, as we discuss in the Introduction the advantage of MVaR relative to VaR is in situations where a portfolio cannot be constructed and thus, a portfolio VaR cannot be obtained.

Our results above rely on quantile estimates that are computed from samples of projections of multidimensional observations on the directional vector. It is well-known that sample quantiles are convergent and biased estimators, whose asymptotic variance can be derived by the delta-method. A confidence interval (CI) for their true value can be

⁵ We would like to thank the reviewer for suggesting this analysis. To preserve space, we present these results in Tables 4A – 6A in the Online Appendix.

constructed by exploiting the binomial property of quantiles (see, for example, Serfling, 1980). By this property, the exact confidence coefficient for the quantile q_a in an ordered sample (x_1, \dots, x_n) is calculated from the binomial distribution with parameters n (sample size) and a ,

$$\Pr(x_i < q_a < x_j) = \sum_{k=i}^{j-1} \binom{n}{k} a^k (1-a)^{n-k} \quad (14)$$

In Table 5, we report the CIs and their lengths, computed by (14) from the projections of observations, where x_i and x_j are chosen such that the probability on the r.h.s. of (14) approximates the nominal confidence coefficient that we set at 95%. We report further the CIs computed by Monte Carlo simulations from samples that were bootstrapped from the relevant data set or generated from the multivariate Student-t distribution with parameters estimated in the same set.⁶ We find, in particular, that CIs computed from Student-t distribution are significantly shorter than the ones computed by (14) and by bootstrapping. This observation may cast doubts on the suitability of the Student-t distribution as a modelling tool for MVaR estimation. We observe further that the length of CIs decreases in a , which suggests stronger confidence in (forecasting) results for higher values of a .

[Table 5]

In line with the 2FM (cf., equation (6)) we argue that an MVaR forecast has two components. We conjecture that the first component is slowly evolving and captures the

⁶ We would like to thank the Associate Editor for suggesting this analysis.

evolution of macroeconomic or other (e.g., company) fundamentals. The second component captures the fast and occasionally violent but transitory movements perhaps reflecting investor sentiment or other short-lived effects. Changes in sentiment can trigger strong liquidity shocks with a significant impact on volatility (Campbell, Grossman and Wang, 1993). In the short run, a change in one set of prices may influence investor sentiment triggering changes in a seemingly unrelated set of prices (Eichengreen and Mody, 1998), thus leading to multidimensional tail risk.

In this context, unconditional Kupiec (1995) and conditional Christoffersen (1998) tests can be intuitively linked to these two components of forecasts. The unconditional accuracy test effectively examines whether a model is consistent with the fundamentals and generates, over the long term, the correct exception rates. The conditional accuracy test, on the other hand, examines how well a forecasting model responds to the twists and turns of the market “animal spirits” which, by definition, are of a behavioral nature with little or no relationship to the long-term fundamentals.

This decomposition highlights the difficulty of long-term (M)VaR forecasting. A comprehensive forecasting model should not only capture the long-term general movements in fundamentals but also anticipate short-lived bursts of turbulence. As it is almost impossible to accurately forecast, well in advance, the latter component, it is too demanding to expect any long-term (M)VaR forecasting model to be conditionally accurate. Therefore, we argue that the adequacy of long-term (M)VaR forecasts should be judged primarily on the basis of the unconditional accuracy test. The conditional accuracy

test, on the other hand, is relevant mainly for short-term (M)VaR forecasts. The practical implication of these observations is that institutions can only get an indication of average long-term exposures from these models but need to monitor their short-term exposures with short-term, conditionally accurate forecasting models such as CAViaR.

6. Conclusions

Aggregation of multiple sources of risk sidelines questions which are paramount for hedging, risk management and financial stability. Interesting answers can be obtained by considering the individual sources of risks jointly. We propose a simple and flexible framework to capture multidimensional tail risk. This framework allows for adapting the techniques and applications developed for unidimensional tail risk which is relatively straightforward even in higher dimensions.

We apply this framework to forecast multidimensional tail events out-of-sample at different horizons and evaluate them statistically. While short horizon forecasts are both conditionally and unconditionally accurate, we find that long horizon forecasts are unconditionally accurate but fail the conditional accuracy tests. However, we argue that this is to be expected. Conditional accuracy is too demanding a criterion for any long horizon (multidimensional) tail event forecasting model. Given our understanding of, and ability to model (multidimensional) tail events, only short horizon forecasts should be subjected to conditional accuracy tests. Long horizon forecasting models of (multidimensional) tail events should be judged primarily on their ability to generate unconditionally accurate forecasts. In this context, it would be interesting to understand

the relationship of the long-term trend and short-term cycle of MVaR to macroeconomic and other fundamentals and investor sentiment, respectively.

Appendix

Proof of (3):

$$\begin{aligned}
 v^d(\mathbf{x}) \geq q_a^d &\Leftrightarrow \mathbf{x} \in \mathcal{M}_a^d. \\
 \Rightarrow: v^d(\mathbf{x}) = \min_{i:d_i \neq 0} \frac{x_i}{d_i} \geq q_a^d &\Rightarrow \frac{x_i}{d_i} \geq q_a^d, \quad \forall i: d_i \neq 0 \Rightarrow \mathbf{x} \in \mathcal{M}_a^d. \\
 \Leftarrow: \mathbf{x} \in \mathcal{M}_a^d &\Rightarrow \frac{x_i}{d_i} \geq q_a^d, \quad \forall i: d_i \neq 0 \Rightarrow \min_{i:d_i \neq 0} \frac{x_i}{d_i} = v^d(\mathbf{x}) \geq q_a^d.
 \end{aligned}$$

References


- Adrian, T., Brunnermeier, M., 2016, CoVaR. *American Economic Review*, 106, 1705-1741.
- Alizadeh, S., Brandt, M. and Diebold, F. 2002, Range-Based Estimation of Stochastic Volatility Models, *Journal of Finance* 57, 1047–92.
- Ang, A. and Bekaert, G. 2002, International Asset Allocation with Regime Shifts, *Review of Financial Studies*, 15, 1137-1187.
- Bae, K-H., Karolyi, G.A. and Stulz, R.M. 2003, A New Approach to Measuring Financial Contagion, *Review of Financial Studies*, 16, 717-763.
- Bartram, S.M., Taylor, S.J. and Wang, Y.-H. 2006, The Euro and European Financial Market Dependence, *Journal of Banking and Finance*, 31, 1461-1481.
- Bank of England, 2015, The Bank of England’s approach to stress testing the UK banking system, Technical Document available at <http://www.bankofengland.co.uk/financialstability/Documents/stresstesting/2015/approach.pdf>
- Basel Committee on Banking Supervision, 2002, Overview of the Amendment to the Capital Accord to Incorporate Market Risk. Technical Report.
- Baxter, M., King, R. 1999, Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series, *Review of Economics and Statistics*, 81, 575-593.
- Bloom, N. 2009, The Impact of Uncertainty Shocks, *Econometrica*, 77, 623-685.
- Boudoukh, J., Richardson, M., and Whitelaw, R.F. 1998, The Best of Both Worlds, *Risk*, 11, 64-67.

- Brandt, M. and Jones, C. 2006, Volatility Forecasting with Range-Based EGARCH Models, *Journal of Business and Economic Statistics* 79, 61-74.
- Campbell, J. Y., Grossman, S. J., and Wang, J. 1993, Trading Volume and Serial Correlation in Stock Returns, *Quarterly Journal of Economics*, 108, 905-939.
- Chen, N-F., Roll, R. and Ross, S.A. 1986, Economic Forces and the Stock Market, *Journal of Business*, 59, 383-403.
- Cherubini, U. and Luciano, E. 2002, Bivariate Option Pricing with Copulas, *Applied Mathematical Finance*, 9, 69-85.
- Christiano, L. and Fitzgerald, T. 2003. The Band Pass Filter, *International Economic Review* 44, 435-465.
- Christoffersen, P. F. 1998, Evaluating Interval Forecasts, *International Economic Review*, 39, 841-862.
- Danielsson, J., and de Vries, C. G., 2000, Value-at-Risk and Extreme Returns, *Annales d'Economie et de Statistique*, 60, 239–270.
- Durbin, J. and Koopman, S. J. 2012, *Time Series Analysis by State Space Methods*, Oxford Statistical Science Series.
- Eichengreen, B. and Mody, A. 1998, What Explains Changing Spreads on Emerging-Market Debt: Fundamentals or Market Sentiment?, NBER Working Paper 6408.
- Embrechts, P., McNeil, A., and Straumann, D. 2002, Correlation and Dependence Properties in Risk Management: Properties and Pitfalls, in *Risk Management: Value at Risk and Beyond*, ed. M. Dempster, Cambridge University Press.
- Engle, R., and Lee, G. 1999, A Permanent and Transitory Model of Stock Return Volatility, in *Cointegration, Causality, and Forecasting: A Festschrift in Honor of Clive W. J. Granger*, eds. R. F. Engle and H. White, New York: Oxford University Press, 475-497.
- Engle, R. and Manganelli, S. 2004, CAViaR: Conditional Autoregressive Value at Risk by Regression Quantiles, *Journal of Business and Economic Statistics*, 22, 367-381.
- Engle, R. and Rangel, J.G. 2008, The Spline GARCH Model for Low Frequency Volatility and its Global Macroeconomic Causes, *Review of Financial Studies*, 21, 1187-1222.
- Ferson, W. and Harvey, C.R. 1998, Fundamental Determinants of National Equity Market Returns: A perspective on Conditional asset pricing, *Journal of Banking and Finance* 21, 1625-1665.
- Gabaix, X. 2009, Power Laws in Economics and Finance, *Annual Review of Economics*, 1, 255-293.
- Gai, P., and Kapadia, S. 2010, Contagion in financial networks, *Proceedings of the Royal Society Series A*, 466, 2401-2423.
- González-Rivera, G. 2014. Predicting Rare Events: Evaluating Systemic and Idiosyncratic Risk, *International Journal of Forecasting*, 30, 688–690.

- Hauksson, H.A., Dacorogna, M., Domenig, T., Muller, U., and Samorodnitsky, G. 2001, Multivariate extremes, aggregation and risk estimation, *Quantitative Finance*, 1, 79-95.
- Hautsch, N., Schaumburg, J. and Schienle, M. 2014. Forecasting systemic impact in financial networks, *International Journal of Forecasting*, 30, 781–794.
- Hautsch, N., Schaumburg, J. and Schienle, M. 2015, Financial Network Systemic Risk Contributions, *Review of Finance*, 19, 685-738.
- Hodrick, R., Prescott, E. 1997, Post-War US Business Cycles: an Empirical Investigation, *Journal of Money, Credit and Banking*, 29, 1-16.
- Hull, J.C. and White, A.D. 2006, Valuing Credit Derivatives Using an Implied Copula approach, *Journal of Derivatives*, Winter, 8, 28-39.
- Kalemanova, A., Schmid, B. and Werner, R. 2007. The Normal Inverse Gaussian Distribution for Synthetic CDO Pricing, *Journal of Derivatives*, 14, 80-93.
- Kuester, K., Mittnik, S., Paoletta, M.S. 2006, Value-at-Risk Prediction: A Comparison of Alternative Strategies, *Journal of Financial Econometrics*, 4, 53-89.
- Kupiec, P. H. 1995, Techniques for Verifying the Accuracy of Risk Measurement Models, *Journal of Derivatives*, 3, 73-84.
- Lehkonen, H. 2015, Stock Market Integration and the Global Financial Crisis, *Review of Finance*, 19, 2039-2094.
- Massacci, D. 2016, Tail Risk Dynamics in Stock Returns: Links to the Macroeconomy and Global Markets, *Management Science*, forthcoming.
- Mandelbrot, B. 1997, *Fractals and Scaling in Finance*. New York, Springer.
- McNeil, A.J. and Frey, R. 2000, Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance*, 7, 271-300.
- Meine, C., Supper, H. and Weiß, G. 2016, Is Tail Risk Priced in Credit Default Swap Premia? *Review of Finance*, 20, 287-336.
- Nieto, M.R. and Ruiz, E. 2016, Frontiers in VaR forecasting and backtesting, *International Journal of Forecasting* 32, 475–501.
- Polanski, A. and Stoja, E. 2012, Efficient Evaluation of Multidimensional Time-Varying Density Forecasts with Applications to Risk Management, *International Journal of Forecasting*, 28, 343-352.
- Prékopa, A. 2012. Multivariate Value at Risk and Related Topics. *Annals of Operations Research* 193, 49–69.
- Rogers, L. and Veraart, L. 2013, Failure and Rescue in an Interbank Network. *Management Science*, 59, 882-898.
- Serfling, R. J. (1980). *Approximation Theorems of Mathematical Statistics*. John Wiley & Sons, New York.
- Su, L. and Spindler, M. 2013, Nonparametric Testing for Asymmetric Information, *Journal of Business and Economic Statistics*, 31, 208-225.


- Takahashi, M., Watanabe, T. and Omori, Yasuhiro. 2016, Volatility and Quantile Forecasts by Realized Stochastic Volatility Models with Generalized Hyperbolic Distribution, *International Journal of Forecasting*, 32, 437-457.
- Torres, R., Lillo, R. E. & Laniado, H. 2015, A Directional Multivariate Value at Risk, *Insurance: Mathematics and Economics*, 65, 111–123.
- Zheng, S., Shi, N., and Zhang, Z. 2012, Generalized Measures of Correlation for Asymmetry, Nonlinearity, and Beyond, *Journal of the American Statistical Association*, 107, 1239-1252.

Figure 1: MVaR-region as a Cartesian product VaR-intervals for $N=2$



Notes: MVaR-region \mathcal{M}_a^d (dark shaded area) in the direction of the vector \mathbf{d} . Note that the upper left corner of \mathcal{M}_a^d corresponds to the point $q_a^d \cdot \mathbf{d}$.


Figure 2: Projections (2) of Observations Inside and Outside of MVaR-region



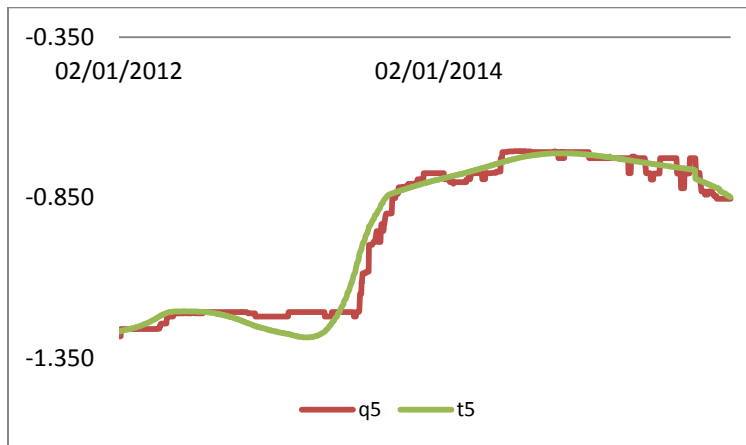
Notes: All points inside (outside) MVaR-region \mathcal{M}_a^d (shaded area) are projected inside (outside) \mathcal{M}_a^d .

Figure 3 Decomposition of US Return Projections MVaR into Trend and Cycle Components

Panel A: Projected Returns and their “Realized” Fifth Quantile




Panel B: Trend of the “Realized” Fifth Quantile Estimated from HP Filter



Notes: Panel A shows the “realized” MVaR estimator (q5) of the US stock indices return projections computed by equation (2). The sample period in the figure is 02/01/2012 to 31/10/2015 (1000 observations). Panel B shows the “realized” MVaR estimator (q5) and its long-run trend (t5) estimated with a Hodrick-Prescott filter with a smoothing parameter of 5,760,000. Panel C shows the cyclical component of the MVaR (c5) defined as the difference between the original series and the trend.


Figure 3 Decomposition of US Return Projections MVaR into Trend and Cycle Components

Panel C: The Cyclical Component of the “Realized” Fifth Quantile



Notes: Panel A shows the “realized” MVaR estimator of the US stock indices return projections computed by equation (2). The sample period in the figure is 02/01/2012 to 31/10/2015 (1000 observations). Panel B shows the “realized” MVaR estimator (q5) and its long-run trend (t5) estimated with a Hodrick-Prescott filter with a smoothing parameter of 5,760,000. Panel C shows the cyclical component of the MVaR (c5) defined as the difference between the original series and the trend.

Figure 4: MVaR Scaling for US Stock Indices



Notes: A log-log plot of empirical frequency- k MVaR (y-axis) at 1% (top left), 2.5% (top right) and 5% (bottom) computed from the returns of US stock indices at different frequencies (x-axis, k days). The respective scaling parameters (slopes) are 0.52, 0.56 and 0.59. The sample period is 1/09/1996 to 31/10/2015 (5000 observations).

Table 1: Summary Statistics and Autocorrelations**Panel A: Summary Statistics**

	Mean	Standard Deviation	Skewness	Excess Kurtosis	Bera-Jarque
DJ30	0.023%	1.155%	-0.143	7.871	12924.742
SP500	0.023%	1.227%	-0.231	7.994	13356.985
NASDAQ	0.030%	1.615%	-0.050	5.391	6056.037
US Projections	-24.84%	97.88%	-0.634	7.698	12681.444
FTSE100	0.027%	1.122%	-0.215	6.071	7716.688
DAX	0.031%	1.265%	0.078	8.855	16339.224
CAC40	0.035%	1.294%	-0.086	4.524	4269.501
MIB30	0.026%	1.371%	-0.162	4.243	3772.200
E-S Projections	-40.59%	98.06%	-0.813	9.600	19752.610
UK Bonds	0.027%	0.381%	-0.006	2.146	959.166
German Bonds	0.025%	0.338%	-0.260	2.411	1267.072
French Bonds	0.025%	0.345%	-0.224	2.993	1908.178
Italian Bonds	0.031%	0.427%	0.537	3.414	1977.693
E-B Projections	-61.01%	98.67%	-1.634	13.544	40437.863

**Panel B: Autocorrelations
Projected Returns**

	1	2	3	4	5	6	Q	P-value
US Projections	-0.046	-0.031	0.009	0.004	-0.008	-0.005	16.356	0.012
E-S Projections	0.056	-0.019	-0.032	0.056	-0.017	0.009	40.056	0.000
E-B Projections	0.215	0.127	0.106	0.132	0.103	0.098	557.405	0.000

Notes: Panel A reports the mean, standard deviation, skewness, excess kurtosis and the Bera-Jarque statistic for daily log close-to-close returns for US stock indices DJ30, SP500 and Nasdaq, European stock indices FTSE100, DAX, CAC40 and MIB30 and 10 year bond prices for UK, Germany, France and Italy. The corresponding projections are computed for the directional vector of standard deviations of the relevant variables. The sample period is 1/09/1996 to 31/10/2015 (5000 observations). Panel B reports the first six autocorrelation coefficients and the Ljung-Box Q statistic for autocorrelation up to six lags, for projected US stock, EU stock and EU bond returns. P-values are also reported.

Table 2: MVaR Out-of-Sample Forecasting Results for US Stock Indices

k	Model	$\alpha = 1\%$				$\alpha = 2.5\%$				$\alpha = 5\%$			
		$\hat{\alpha}$	p-val. (t_u)	p-val. (LR_c)	p-val. (DQ)	$\hat{\alpha}$	p-val. (t_u)	p-val. (LR_c)	p-val. (DQ)	$\hat{\alpha}$	p-val. (t_u)	p-val. (LR_c)	p-val. (DQ)
1	CAViaR	0.013	0.111	0.298	0.290	0.027	0.571	0.924	0.295	0.052	0.622	0.217	0.183
	2FM	0.017	0.002	0.074	0.009	0.030	0.109	0.023	0.874	0.056	0.132	0.256	0.695
	Scaling	0.009	0.852	0.000	0.533	0.023	0.465	0.005	0.301	0.040	0.005	0.011	0.077
5	2FM	0.018	0.001	0.101	0.004	0.032	0.036	0.004	0.138	0.057	0.111	0.164	0.760
	Scaling	0.006	0.012	0.000	0.010	0.016	0.000	0.000	0.000	0.454	0.226	0.000	0.171
10	2FM	0.018	0.001	0.018	0.004	0.033	0.010	0.002	0.055	0.057	0.091	0.183	0.713
	Scaling	0.009	0.441	0.000	0.125	0.015	0.000	0.000	0.000	0.048	0.636	0.000	0.318
20	2FM	0.018	0.001	0.019	0.011	0.034	0.007	0.003	0.192	0.057	0.098	0.181	0.885
	Scaling	0.011	0.577	0.000	0.257	0.014	0.000	0.000	0.000	0.053	0.417	0.000	0.271
60	2FM	0.022	0.000	0.010	0.001	0.038	0.000	0.023	0.055	0.060	0.024	0.088	0.156
	Scaling	0.019	0.000	0.000	0.018	0.019	0.027	0.000	0.005	0.052	0.621	0.000	0.167

Notes: The table reports the actual exception rate ($\hat{\alpha}$) for each MVaR forecasting model out of 3000 observations, (i.e. the proportion of times the forecasted MVaR is exceeded), the p-value of the t-statistic to test the null hypothesis of unconditional accuracy (formula (11)) and the p-values of the LR and DQ statistics (formulas (12) and (13), respectively) to test the null hypothesis of conditional accuracy for different confidence levels. The out-of-sample period of 3000 observations is 14 August 2007 to 31 October 2015. For CAViaR and 2FM models the daily MVaR forecasts are k -day ahead, while for Scaling the forecasts are for frequency- k MVaR.

Table 3: MVaR Out-of-Sample Forecasting Results for European Stock Indices

k	Model	$\alpha = 1\%$				$\alpha = 2.5\%$				$\alpha = 5\%$			
		$\hat{\alpha}$	p-val. (t_u)	p-val. (LR_c)	p-val. (DQ)	$\hat{\alpha}$	p-val. (t_u)	p-val. (LR_c)	p-val. (DQ)	$\hat{\alpha}$	p-val. (t_u)	p-val. (LR_c)	p-val. (DQ)
1	CAViaR	0.012	0.247	0.562	0.735	0.026	0.731	0.413	0.319	0.048	0.609	0.324	0.211
	2FM	0.013	0.147	0.106	0.679	0.027	0.433	0.275	0.575	0.054	0.333	0.000	0.903
	Scaling	0.009	0.431	0.500	0.183	0.023	0.465	0.098	0.159	0.046	0.296	0.005	0.290
5	2FM	0.013	0.189	0.096	0.768	0.027	0.427	0.278	0.492	0.055	0.224	0.000	0.818
	Scaling	0.011	0.595	0.000	0.755	0.024	0.729	0.000	0.198	0.049	0.747	0.000	0.316
10	2FM	0.013	0.108	0.119	0.523	0.028	0.261	0.135	0.580	0.056	0.189	0.000	0.613
	Scaling	0.015	0.017	0.000	0.007	0.025	0.887	0.000	0.466	0.051	0.653	0.000	0.348
20	2FM	0.014	0.043	0.146	0.576	0.031	0.052	0.094	0.924	0.056	0.176	0.000	0.571
	Scaling	0.021	0.000	0.000	0.000	0.034	0.007	0.000	0.170	0.057	0.114	0.000	0.267
60	2FM	0.017	0.005	0.050	0.172	0.032	0.040	0.000	0.130	0.058	0.059	0.000	0.325
	Scaling	0.026	0.000	0.000	0.002	0.036	0.002	0.000	0.084	0.058	0.059	0.000	0.217

Notes: The table reports the actual exception rate ($\hat{\alpha}$) for each MVaR forecasting model out of 3000 observations, (i.e. the proportion of times the forecasted MVaR is exceeded), the p-value of the t-statistic to test the null hypothesis of unconditional accuracy (formula (11)) and the p-values of the LR and DQ statistics (formulas (12) and (13), respectively) to test the null hypothesis of conditional accuracy for different confidence levels. The out-of-sample period of 3000 observations is 14 August 2007 to 31 October 2015. For CAViaR and 2FM models the daily MVaR forecasts are k -day ahead, while for Scaling the forecasts are for frequency- k MVaR.

Table 4: MVaR Out-of-Sample Forecasting Results for European Bond Indices

k	Model	$\alpha = 1\%$				$\alpha = 2.5\%$				$\alpha = 5\%$			
		$\hat{\alpha}$	p-val. (t_u)	p-val. (LR_c)	p-val. (DQ)	$\hat{\alpha}$	p-val. (t_u)	p-val. (LR_c)	p-val. (DQ)	$\hat{\alpha}$	p-val. (t_u)	p-val. (LR_c)	p-val. (DQ)
1	CAViaR	0.006	0.011	0.623	0.096	0.019	0.010	0.110	0.166	0.039	0.003	0.060	0.025
	2FM	0.012	0.314	0.009	0.944	0.025	0.908	0.103	0.161	0.050	1.000	0.053	0.301
	Scaling	0.007	0.049	0.586	0.346	0.019	0.024	0.445	0.051	0.042	0.037	0.134	0.058
5	2FM	0.012	0.311	0.077	0.924	0.026	0.809	0.063	0.205	0.050	0.920	0.025	0.168
	Scaling	0.007	0.089	0.000	0.127	0.013	0.000	0.000	0.000	0.041	0.042	0.000	0.003
10	2FM	0.012	0.307	0.078	0.790	0.026	0.711	0.064	0.255	0.051	0.838	0.029	0.173
	Scaling	0.005	0.000	0.000	0.013	0.013	0.000	0.000	0.000	0.041	0.011	0.000	0.007
20	2FM	0.012	0.378	0.071	0.340	0.026	0.775	0.067	0.278	0.050	0.997	0.113	0.122
	Scaling	0.007	0.054	0.000	0.048	0.011	0.000	0.000	0.000	0.046	0.336	0.000	0.157
60	2FM	0.014	0.068	0.001	0.262	0.025	0.955	0.098	0.289	0.051	0.870	0.201	0.238
	Scaling	0.004	0.000	0.000	0.008	0.009	0.000	0.000	0.000	0.053	0.370	0.000	0.442

Notes: The table reports the actual exception rate ($\hat{\alpha}$) for each MVaR forecasting model out of 3000 observations, (i.e. the proportion of times the forecasted MVaR is exceeded), the p-value of the t-statistic to test the null hypothesis of unconditional accuracy (formula (11)) and the p-values of the LR and DQ statistics (formulas (12) and (13), respectively) to test the null hypothesis of conditional accuracy for different confidence levels. The out-of-sample period of 3000 observations is 14 August 2007 to 31 October 2015. For CAViaR and 2FM models the daily MVaR forecasts are k -day ahead, while for Scaling the forecasts are for frequency- k MVaR.

Table 5: 95% Confidence Intervals (CI) and their Lengths for MVaR Estimates

Dataset	CI	$\alpha = 1\%$	$\alpha = 2.5\%$	$\alpha = 5\%$
EU equities	Formula (14)	(-3.62, -2.82), 0.80	(-2.48, -2.19), 0.29	(-1.99, -1.73), 0.26
	Bootstrap	(-3.49, -2.70), 0.79	(-2.46, -2.16), 0.30	(-1.97, -1.73), 0.24
	Student-t	(-3.29, -2.85), 0.44	(-2.57, -2.31), 0.26	(-2.08, -1.90), 0.18
EU bonds	Formula (14)	(-4.27, -3.26), 1.02,	(-2.99, -2.55}, 0.43	(-2.33, -2.11), 0.23
	Bootstrap	(-4.22, -3.21), 1.01	(-2.93, -2.52}, 0.40	(-2.32, -2.11), 0.21
	Student-t	(-3.91, -3.57), 0.34	(-2.92, -2.72}, 0.20	(-2.29, -2.16), 0.12
US indices	Formula (14)	(-3.43, -2.63), 0.79	(-2.32, -1.94), 0.38	(-1.69, -1.48), 0.21
	Bootstrap	(-3.29, -2.58), 0.70	(-2.31, -1.91), 0.40	(-1.67, -1.47), 0.20
	Student-t	(-3.13, -2.90), 0.23	(-2.43, -2.30), 0.13	(-1.99, -1.90), 0.09

Notes: Formula (14) computes the (approximate) 95%-CI from the projections of the last 3,000 observations (between 1/05/2004 and 31/10/2015) in the relevant data set. Bootstrap (Student-t) computes the 95%-CI from 1,000 samples of size 3,000 each, drawn from projections of the last 3,000 observations in the relevant data set (from multivariate Student-t with parameters estimated from the last 3,000 observations in the relevant data set).