Multifractality and Value at Risk Forecasting of Exchange Rates*

Jonathan A. Batten†
Harald Kinateder‡
Niklas Wagner§

This version: April 2, 2012

*The authors would like to thank Chee-Jin Yap for collecting data. All omissions and errors remain with the authors.

†Department of Finance, Hong Kong University and Technology HKUST

‡Department of Business and Economics, Passau University

§Corresponding author. Department of Business and Economics, Passau University, 94030 Passau, Germany. Phone: +49 851 509 3241, Fax: +49 851 509 3242, E-mail: nwagner@uni-passau.de, nwagner@alum.calberkeley.org.

Multifractality and Value at Risk Forecasting of Exchange Rates

Abstract

The empirical properties of risky asset returns typically reveal autocorrelation functions that decline rapidly and are statistically insignificant beyond a few lags. However, the autocorrelation functions of squares and absolute values of returns, typically decline slowly, and in fact may persist for many years. Assuming the return series are also stationary, then this behavior indicates that the series is not independent, and is instead either short-range dependent, or in some cases also long range dependent. These scaling properties are well known by financial market practitioners, although they are commonly ignored for modeling convenience.

This paper contributes to the empirical literature that investigates the risk scaling properties of financial asset returns. We do so by investigating the presence of complex, nonlinear behavior in the EUR/USD spot foreign exchange rate, which is the most important currency pair traded in over-the-counter (OTC) markets. Our findings demonstrate a strong case against the use of models that accommodate multifractality and instead favor those that address fat-tailed, or leptokurtic, features present in these distributions.

Our findings also provide additional insights into the scaling properties of the spot returns. For example, the scaling functions are initially linear, although ultimately concave, which is consistent with returns' distributions of infinite low order moments. This result would support the need for modeling that addresses multifractality, although modifying commonly used models cannot readily meet these requirements. We then present evidence of a class of models that can incorporate the empirically observed features.

Keywords: Exchange Rates, Forecasting, High Frequency Exchange Rates, Modeling Asset Returns, Multifractality, Value at Risk, VaR.

JEL Classification: C22

The multifractal model of asset returns (MMAR), as introduced by Mandelbrot et al. (1997), covers fat-tailed return distributions, long-range dependence in different powers of returns and scale inconsistency. Amongst the standard empirical properties of risky asset returns much more can be inferred from the scaling properties of the returns. We find the empirical scaling functions of spot EUR/USD exchange rates are initially linear and ultimately concave, which is strongly suggestive of returns distributions with infinite low order moments, or alternatively that multifractal behavior is a necessary modeling requirement. In this paper we derive a modified MMAR approach, which models trading time by a series of volume ticks and allows the MMAR to be used for out-of-sample Value at Risk (VaR) forecasting.

1 Introduction

A widely used approach to risk management in financial intermediaries is a statistical procedure described as Value at Risk (VaR). In general, this approach estimates the worst expected loss over a target horizon that will not be exceeded with a given level of confidence. Importantly, VaR is used by regulators to determine the appropriate amount of risk capital needed by the intermediary to avoid default arising from adverse movements in asset prices. In addition, it can be used by management to support decisions involving the setting of position limits and performance evaluation. Typically VaR is calculated based on daily return data, although doing so ignores the risk –and potential losses– associated with the liquidation of positions due to adverse *intraday* price movements.

Nonetheless, an important caveat to VaR estimation is the well-known fact

that financial assets returns (typically estimated as the interperiod change of the logarithm of the price) do not display ideal statistical properties. Instead, they exhibit multifractal or multiscaling features that are characterized by a special form of time-invariance, which combines fat-tailed returns (leptokurtosis) with long memory.[1] The presence of these features in financial time series demonstrates the criticality of adopting more accurate risk measures. Importantly, failure to do so may result in standard risk measurement and management techniques misrepresenting the appropriate level of risk capital required to support an intermediary during a crisis event, which in the event of failure could lead to broader system wide consequences.

The main contribution of this paper is that we demonstrate a more accurate intra-day risk measure that better forecasts these potential losses, thereby providing financial market practitioners and regulators with a better determination of the required level of risk capital. To do so, we employ a modified version of the multifractal model of asset returns (MMAR) that was originally proposed by Mandelbrot et al. (1997). This model has the benefit of addressing the complex properties of financial price returns by estimating various degrees of long memory at different powers, and importantly is able to accommodate the presence of fat-

[1]The empirical evidence is extensive. For example, Fisher and Calvet (2002) find multifractality in a (Deutsche Mark) DMK/USD high frequency series and Xu and Gençay (2003) also prove 5-minute USD/DMK returns are multifractal. Eisler and Kertész (2004) report multiscaling behavior for a high frequency Deutscher Aktien Index (DAX) stock series and high frequency observations of the 200 most liquid stocks at New York Stock Exchange (NYSE). A stationary process exhibits long memory, when its autocorrelation function behaves asymptotically as a power law. For alternative high frequency VaR approaches, see Sun et al. (2009).

tails (or leptokurtosis).[2]

Alternate approaches, such as fractionally integrated GARCH (Generalized autoregressive conditional heteroskedasticity) models (or FIGARCH) have the same decay rate for all moments and are not scale-consistent. Consequently, Fisher and Calvet (2004), as well as Lux (2008) find MMAR outperform both GARCH and FIGARCH models applied to foreign exchange rate series, although Fillol (2003) shows that MMAR is a better model for replicating the scaling properties observed in the French CAC-40 (Cotation Assisté en Continu) stock series.

In a recent study Fleming and Kirby (2011) find a strong correlation between the innovations to volume and volatility and suggest that volume could be utilised to improve risk prediction. Given this conclusion, and following Clark (1973), we model trading time by a series of trading volume ticks and provide a modified MMAR approach for out-of-sample VaR forecasting. These modifications overcome many of the limitations of the previous multifractal models such as their combinatorial nature and their restriction to a bounded interval. Note that Fisher and Calvet (2001) overcome these shortcomings by introducing a Markov-switching multifractal model, while Lux (2008) provides an alternative generalized method.

In the present paper we apply MMAR to the VaR forecasting of the EUR/USD

[2]See for example, Liu and Lux (2005). Different degrees of long memory in various powers of returns are a stylized fact of financial returns. For example, see Ding et al. (1993), Lobato and Savin (1998) and Cont (2001).

spot exchange rate. The EUR/USD spot exchange rate is the major currency pair traded in the over-the-counter spot foreign exchange market, with daily turnover now estimated to be excess of US$1 trillion (BIS (2010)). In order to test the forecasting ability of our novel VaR approach, we study the out-of-sample accuracy of VaR predictions for both 12-hour and daily (24-hour) forecast horizons. While these forecast periods are somewhat arbitrary they are consistent with the trading activities expected of global financial intermediaries with a subsidiary, or branch, always open during the 24-hour trading day.

One additional novelty to our approach is that we use high frequency observations. Our data consists of round-the-clock EUR/USD spot exchange rate prices quoted by market participants on the Reuters trading platform during the period January 5, 2006 to December 31, 2007. These prices are bundled into 5-minute time stamped intervals with the spot price and the trading ticks recorded. We find that the EUR/USD returns are multifractal, with the moments showing different scaling exponents. Our novel MMAR approach is then compared with forecasts based on historical simulation and a benchmark GARCH approach. These comparisons show that the MMAR approach is the only model, which produces admissible VaR forecasts for the 12-hour forecast horizon. However, for the daily horizon, we find the MMAR outperforms all alternatives.

The paper is organized as follows. Section 2 contains a brief review of the multifractality literature and introduces our modified version of the multifractal model of asset returns. In Section 3 we present the multiple-period VaR forecast-

ing concept for our MMAR model and outline competing standard VaR models. Section 4 contains an empirical study of the various models' forecasting ability using the EUR/USD foreign exchange rate series. Finally, Section 5 concludes.

2 Multifractal Modeling of FX Returns

We model foreign exchange (FX) returns using a modified version of the multifractal model of asset returns (MMAR). Before our MMAR approach is introduced, we provide a brief overreview of the multifractality concept.

2.1 Multifractality

Mandelbrot et al. (1997) define a multifractal process as one that possesses a nonlinear scaling function. By definition, a linear scaling function is therefore a feature of unifractal series. However, uniscaling processes do not model returns appropriately, as the majority of financial data sets do not exhibit linear scaling functions (see for example, Cont (2001), Fisher and Calvet (2002) and Xu and Gençay (2003)).

In the following, we formally define multifractal processes and introduce the scaling and partition functions. The partition function is necessary to calculate the scaling function.

Definition 2.1 *A time series process* $\{Y_t\}_{1\leq t\leq T}$ *is called multifractal, if it has*

stationary increments and satisfies

$$\mathbb{E}\big(|Y_{t,h} - Y_t|^q\big) = c(q)h^{\tau(q)+1}, \tag{1}$$

for all $1 \leq h \leq T$*, and* $0 \leq q \leq Q$*. The scaling function* $\tau(q)$ *and the prefactor* $c(q)$ *are both deterministic functions of* q*.*

The scaling function $\tau(q)$ takes into account the influence of the time h on the moments q. For $q = 0$ all scaling functions have the same intercept $\tau(0) = -1$. According to Definition 2.1, unifractal processes like Brownian motion (BM) and fractional Brownian motion (FBM) have linear scaling functions. The scaling law of BM is

$$\tau_{BM}(q) = \frac{q}{2} - 1, \tag{2}$$

and for FBM

$$\tau_{FBM}(q) = qH - 1, \tag{3}$$

for $0 < H < 1$.

Fisher and Calvet (2002) propose the following regression model to calculate the scaling function $\tau(q)$:

$$\log S(q,h) = \tau(q) \log h, \tag{4}$$

where $S(q,h)$ denotes the partition function. The partition function $S(q,h)$ of a time series process $\{Y_t\}_{1 \leq t \leq T}$ is defined for each moment q and obtained by

partitioning the series into $N = \lfloor T/h \rfloor$ non overlapping subintervals of length $h \in [1, T]$:

$$S(q, h) = \sum_{i=0}^{N-1} |Y_{i \cdot h+h} - Y_{i \cdot h}|^q. \tag{5}$$

The partition function is an indicator of the stability of the moments depending on the time horizon.

2.2 Multifractal Model of Asset Returns

We now formalize the construction of the MMAR, which models FX returns $\{R_t\}_{1 \leq t \leq T}$ by compounding a FBM with a multifractal trading time. The FX return is defined as $R_t \equiv \log X_t - \log X_{t-1}$, where X_t denotes a FX rate.

Assumption 2.1 *The MMAR assumes that the FX return process $\{R_t\}_{1 \leq t \leq T}$ follows*

$$R_t = B_t^{(H)}[\theta_t], \tag{6}$$

where $B_t^{(H)}$ is a FBM with self-affinity index $0 < H < 1$ and θ_t is a multifractal trading time. $B_t^{(H)}[\theta_t]$ is a compound process.

The multifractal trading time controls the tails of the process $\{B_t^{(H)}[\theta_t]\}_{1 \leq t \leq T}$. In different periods of time, the trading can be slower or faster. As the FX return process is assumed to follow $B_t^{(H)}[\theta_t]$, smaller or larger values of θ_t correspond to smaller and larger volatility, respectively (Eisler and Kertész (2004)).

The most important feature of FBM as discussed in Mandelbrot and van Ness (1968) is that FBM is able to model different degrees of persistence. For $H = 0.5$, FBM is ordinary BM with independent and identically distributed (i.i.d.) increments. Yet, $B_t^{(H)}$ is antipersistent when $0 < H < 0.5$ and displays long memory for $0.5 < H < 1$.

Additionally, we impose the following assumptions:

Assumption 2.2 *The processes* $\{B_t^{(H)}\}_{1 \leq t \leq T}$ *and* $\{\theta_t\}_{1 \leq t \leq T}$ *are independent.*

Assumption 2.3 *The multifractal trading time* θ_t *is given by the normalized volume ticks* $\overline{V}_t$*,*

$$\overline{V}_t = \frac{V_t}{\sum_{t=1}^{T}(V_t)}, \tag{7}$$

where $V_t \geq 0$ *denotes the number of volume ticks at time* t*. It is assumed that* $\overline{V}_t$ *is multifractal as given by the scaling law in equation (1).*

In our setting θ_t is not a virtual trading time, but reflects the normalized number of trades that occur. The motivation for this method is twofold. First, Ma (2005) studies stock trading time and finds that stock trading time confirms the multifractal trading time hypothesis suggested by Mandelbrot et al. (1997). Second, the actual volume ticks reflect the behavior of the market participants more accurately than a virtual trading time. The use of volume ticks is an innovation compared with previous MMAR approaches. For example, former MMAR approaches model θ_t by log-normal or Poisson cascade processes (Fisher and Calvet (2001)).

Given Assumptions 2.1, 2.2, and 2.3, we use a method based on the factorization of the autocovariance matrix, K, to simulate sample paths of $B_t^{(H)}[\theta_t]$ at times $1 \leq t \leq T$:

$$B_t^{(H)}[\theta_t] = \big[g^T \times A\big]\sigma \tag{8}$$

with

$$g \sim \mathcal{N}(0,1) \tag{9}$$

and

$$K = AA^T, \tag{10}$$

where $K \in \mathbb{R}^{T\times T}$ is a positive-definite matrix and $g \in \mathbb{R}^{T\times 1}$ is a vector of standard Gaussian numbers, while σ is the average of the unconditional standard deviations of R_t and $\overline{V}_t$. The symmetric matrix $A \in \mathbb{R}^{T\times T}$ is the square root of the matrix K. The elements of the matrix K are

$$K_{i,j} = \frac{1}{2}\big(\Theta_t^{2H} + \Theta_s^{2H} - |\Theta_t - \Theta_s|^{2H}\big), \tag{11}$$

for each given time $1 \leq s \leq t \leq T$, where $\{\Theta_t\}_{1\leq t\leq T}$ is the cumulative trading time that increases with real time t,

$$\Theta_t = \sum_{i=1}^{t} \theta_i. \tag{12}$$

3 MMAR Market Risk Estimation

Given some probability $0 < \alpha < 1$, the $\mathcal{F}_{t-1}$-measurable VaR for the h-period-ahead time interval $(t-1, t+h-1]$, $VaR^{\alpha}_{t,h}$, is defined as

$$VaR^{\alpha}_{t,h} = -\inf\left\{r : F_{t,h}(r) \geq \alpha\right\}, \tag{13}$$

which yields $\mathbb{P}(R_{t,h} \leq -VaR^{\alpha}_{t,h}|\mathcal{F}_{t-1}) = \alpha$. The low-frequency FX return (or h-period FX return), $R_{t,h} = \log X_{t+h-1} - \log X_{t-1}$, is given by $R_{t,h} = \sum_{\tau=1}^{h} R_{t+\tau-1}$. $F_{t,h}(r)$ denotes the time-$(t-1)$ conditional distribution of the low frequency FX return, i.e. $F_{t,h}(r) = \mathbb{P}(R_{t,h} \leq r|\mathcal{F}_{t-1})$. We now present the VaR representation for our modified MMAR approach.

To calculate the VaR, we simulate a total of D MMAR sample paths $R^{(d)}_t, ..., R^{(d)}_t$ for $1 \leq t \leq T$ and $1 \leq d \leq D$, where $R^{(d)}_t$ is the simulated one-period MMAR return, R_t, at time t in the dth sample. Since the MMAR creates a self-affine return sample path, we derive the corresponding h-period MMAR return, $\{R^{(d)}_{t,h}\}_{1 \leq t \leq T}$, with

$$R^{(d)}_{t,h} = h^H R^{(d)}_t, \tag{14}$$

for $h > 0$ and $0 < H < 1$. The conditional distribution of $R^{(d)}_{t,h}$ is given by

$$F^{(d)}_{t,h}(r^{(d)}) = \mathbb{P}\left(R^{(d)}_{t,h} \leq r^{(d)}|\mathcal{F}_{t-1}\right). \tag{15}$$

From equation (15), we obtain D quantile estimates $F^{(d)\leftarrow}_{t,h}(\alpha)$. Given this, the

multiple-period VaR is derived by the mean of the D quantile estimates:

$$VaR_{t,h}^{\alpha} = -\left(\frac{1}{D}\sum_{d=1}^{D} F_{t,h}^{(d)\leftarrow}(\alpha)\right). \tag{16}$$

4 Empirical Analysis

This section investigates the predictive performance of our MMAR VaR approach and two alternative VaR models. We study the models' performance using high frequency foreign exchange data.

4.1 Dataset and Descriptive Statistics

The 5-minute round-the-clock EUR/USD series is obtained from the Reuters' trading platform, and covers January 5, 2006, through to December 31, 2007, for a total of 138,418 high frequency observations. The data is bundled into 5-minute sequential time stamped intervals with the spot price and the ticks recorded. A tick is a trade through the platform for the standard minimum amount of about 3-5 million euro. Since the highest frequency is 5-minutes ($h = 1$), one hour corresponds to $h = 12$ and one day is equal to $h = 288$.[3] Table 1 shows that with an increasing sampling frequency unconditional volatility rises. Additionally, one can clearly observe the declining kurtosis leading to approximately Gaussian data. 5-minute EUR/USD returns are fat-tailed (compare also Figure 2) and

[3]For $h = 1$ the return sampled at the highest frequency is $R_{t,1} = R_t$.

highly short-term dependent since one can observe significant autocorrelation of absolute EUR/USD returns. Yet, this effect declines, when the sampling frequency increases.

[Insert Table 1 about here]

The EUR/USD returns at 5-minute frequency are illustrated in Figure 1. In Figure 2 Quantile-Quantile (Q-Q) plots of the sample EUR/USD returns for 5-minute and daily sampling are provided. In both panels the empirical quantiles (points) are plotted against the theoretical quantiles of a normal distribution (straight line). If the two distributions are similar, the points of the empirical distribution should lie on the line. The Q-Q plot shows that 5-minute EUR/USD returns are fat-tailed. With an increasing sampling frequency (daily) the EUR/USD returns are less tailed, but have still fatter tails than a normal distribution.

[Insert Figure 1 about here]

[Insert Figure 2 about here]

Figure 3 illustrates 5-minute EUR/USD trading ticks V_t and the cumulative trading time Θ_t.

[Insert Figure 3 about here]

4.2 Evidence of Multifractality

Figure 4 presents EUR/USD return partition functions for moments $1 \leq q \leq 5$ in the scaling region up to one week. The partition functions are plotted against h using logarithmic scales.[4] For $q \in [1, 3]$ the lines are approximately linear. Generally higher moments are more sensitive to deviations from scaling as they capture information in the tails of $R_{t,h}$. The partition function for moments $q = 4$ and $q = 5$ shows linear behavior up to $\log h = 3$ which corresponds to a time horizon of about 100 minutes.

[Insert Figure 4 about here]

Figure 5 illustrates estimated scaling functions of EUR/USD returns and Brownian motion. The scaling function of ordinary Brownian motion (dashed line) as defined in equation (2) is linear, which is typical for unifractal processes. The scaling function of EUR/USD appears completely different; it is non-linear which is a characteristic of a multifractal process (see Definition 2.1).

[Insert Figure 5 about here]

4.3 Out-of-Sample VaR Forecasts

Model Estimation

In order to examine out-of-sample performance, we calculate the 1% VaR from January 6, 2006 to December 31, 2007 for 12-hour ($h = 144$) and daily ($h =$

[4] One week corresponds to 1.440 observations: $\log h = 7.27$.

288) forecast horizons. The matrix A which is the square root of the matrix K is computed by the Cholesky decomposition and the self-affinity index H is estimated using the variance of residuals approach of Peng et al. (1994). We use $D = 5,000$ MMAR sample paths to calculate multiple-period VaR. The VaR results of our MMAR approach are compared to two alternatives.

First, instead of simulated multifractal return series, we use historical simulation of EUR/USD returns. The multiple-period VaR based on this historical simulation is:

$$VaR_{t,h}^{\alpha} = -\big(\sqrt{h}F_t^{\leftarrow}(\alpha)\big), \tag{17}$$

where $F_t^{\leftarrow}(\alpha)$ denotes the empirical α-quantile of R_t.

Second, we also compare the VaR results of our MMAR approach to a VaR model in the location-scale class. The multiple-period VaR based on GARCH is:

$$VaR_{t,h}^{\alpha} = -\big(\mu_{t,h} + \sigma_{t,h}\ F^{\leftarrow}(\alpha)\big), \tag{18}$$

where $F^{\leftarrow}$ is the inverse of the innovations distribution. The multiple-period mean, $\mu_{t,h}$, as well as the multiple-period volatility, $\sigma_{t,h}$, are determined by the GARCH(1,1) model of Bollerslev (1986). Estimation of the GARCH(1,1) model with skewed student-t innovations is carried out by a maximum likelihood method.

All model parameters of our VaR models are calculated using a moving window, which consists of 288 5-minute high frequency EUR/USD returns. The

time-t VaR forecasts are updated every time h, which results in $\lfloor \frac{T-288}{h} \rfloor$ successive VaR predictions.

VaR Forecast Results

The out-of-sample VaR forecast results are summarized in Table 2. We provide the average value of the VaR estimates, $\overline{VaR}$, and the empirical coverage rate, $\widehat{\alpha}$. We use the Christoffersen (1998) LR statistics for VaR performance evaluation. Appendix A contains a detailed explanation.[5]

[Insert Table 2 about here]

We consider a violation of a backtesting statistic to have occured when a corresponding p-value of the respective test statistic is below 5 percent. The LR_{uc} statistic tests for the correct violation level, which is $\alpha = 1\%$ in our setting. VaR forecasts may be correct on average, but produce violation clusters, a phenomenon ignored by unconditional coverage. The LR_{ind} statistic tests for independence of the VaR violations. The conditional coverage statistic, LR_{cc}, combines both concepts. We therefore consider it as our most important backtesting statistic. The VaR predictions of the three models for the 12-hour horizon from January 6, 2007 to December 31, 2008 are plotted in Figure 6. The dash dotted line refers to the MMAR approach, the dashed line is historical

[5] All models are validated based on non-overlapping h-period EUR/USD returns in order to provide independent out-of-sample tests.

simulation and the solid line is the GARCH based VaR model. In the following, we analyze the models' performance in detail.

Unconditional Coverage versus Conditional Coverage

The unconditional coverage test indicates that for the 12-hour horizon of GARCH and for all horizons of historical simulation, the predicted VaR exceeds the actual loss levels. This leads to insufficient coverage of potential losses since the VaR level is too low. Our MMAR approach passes all backtesting statistics for p-values above 0.05. For both horizons, we find that our MMAR model achieves better conditional coverage than historical simulation and the GARCH based VaR. All models pass the independence test indicating that there is no severe clustering of VaR violations over time. As a result, backtesting violations result due to VaR predictions exceeding actual levels.

12-hour versus daily Forecast Horizon

Comparing the two forecast horizons, we find that historical simulation and GARCH produce better VaR results (and therefore higher p-values of all backtesting statistics) for the daily horizon. This is because returns sampled at a higher frequency often obtain more extreme returns as lower frequency returns. Our MMAR model has less problems with such an environment. First, the MMAR is characterized by a special form of time-invariance, which combines extreme returns with long memory (Fisher and Calvet (2002)). Second, our

MMAR approach uses trading volume. This is important, as increased trading activities can produce extreme returns.

The predicted VaR of GARCH does not violate any backtesting statistic for the daily horizon. The MMAR performs best for the 12-hour horizon: We find that the p-value of LR_{cc} (which is 0.900) and the empirical violation level $\widehat{\alpha} = 1.04\%$ both report excellent VaR forecasting. Concerning the daily horizon, we report good backtesting results for the MMAR, but inferior to the 12-hour horizon. To conclude, for both horizons, our MMAR model dominates the VaR results of historical simulation and GARCH and never violates a backtesting statistic.

[Insert Figure 6 about here]

4.4 Economic Implications

In the last section we discussed the VaR forecasting results from a statistical point of view. Now, we present some economic implications of our results. Let's assume a bank takes 1 million euro in a long (bought) position against the US dollar and intends to close it after 12 hours. On average, we expect the MMAR to determine 9.800 euro as risk capital. The expected average risk capital of historical simulation and GARCH based VaR would be 7.000 and 8.000 euro, respectively. Given this and the empirical violation rate $\widehat{\alpha}$ reported in Table 2, we find that the bank would be considerably undercapitalised with its long

position when using historical simulation.[6]

A lack of regulatory capital can also affect stock price movements of banks. Authorities like the the Federal Reserve (Fed) perform stress tests to check whether banks are able to maintain a sufficient level of capital in the event of extreme price movements. Banks using inefficient VaR models may lack sufficient risk capital and therefore may face problems when these stress tests are undertaken. It is not uncommon for a banks' stock price to respond negatively when the results are publicly disclosed. On the other hand, stock prices of banks performing well in a stress test are likely to respond positively.

5 Conclusion

In this paper we address the need by financial intermediaries to provide more accurate forecasts of intra-day VaR. Doing so is critical not just for regulatory purposes but also for better internal risk management. One approach to solving this difficult issue is to use a MMAR based VaR technique for risk prediction. The MMAR as introduced by Mandelbrot et al. (1997) covers fat-tailed return distributions, long-range dependence in different powers of returns and scale inconsistency. We model trading time using a series of trading volume ticks and provide a modified MMAR approach for out-of-sample VaR forecasting. The

[6] As the empirical violation rate $\widehat{\alpha}$ of MMAR is roughly 1% which is the target rate, the risk capital of MMAR is very close to the required risk capital. The other models have significantly higher empirical violation rates implying a lack of capital.

novel MMAR approach is a quantile based VaR approach using FX return and volume data to enhance risk forecasts. We obtain superior forecast results as compared with other alternate VaR models (historical simulation and GARCH). While our results confirm the importance of considering multifractality when attempting to predict risk, further refinements of standard approaches are important for efficient market risk prediction.

Another important fact is that our MMAR model is a consistent approach which does not need any further model specifications. Other standard approaches for VaR including extreme value theory, or GARCH, need additional specifications. For example, GARCH based VaR needs an assumption for the innovations distribution. This is also a relevant aspect for estimating VaR via extreme value theory because one has to prefilter raw returns in order to obtain independent returns. When raw returns are prefiltered by a GARCH filter than the choice of the innovations distribution influences the quality of VaR results, as well.

References

Bank for International Settlements (2010): Triennial Central Bank Survey Report on global foreign exchange market activity in 2010, Monetary and Economic Department BIS, Basel, Switzerland. December: 1-95. Source Table B.6, page 15.

Bollerslev, Tim (1986): Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 31: 307–327.

Christoffersen, Peter (1998): Evaluating interval forecasts, *International Economic Review* 39: 841–862.

Clark, Peter K. (1973): A subordinated stochastic process model with finite variance for speculative prices, *Econometrica* 41: 135–156.

Cont, Rama (2001): Empirical properties of asset returns: Stylized facts and statistical issues, *Quantitative Finance* 1: 223–236.

Ding, Zhuanxin, Robert F. Engle, and Clive W. J. Granger (1993): A long memory property of stock market returns and a new model, *Journal of Empirical Finance* 1: 83–106.

Eisler, Zoltan and Kertész Janos (2004): Multifractal model of asset returns with leverage effects, *Physica A* 343: 603–622.

Fillol, Jérôme (2003): Multifractality: Theory and evidence an application of the French stock market, *Economics Bulletin* 3: 1–12.

Fisher, Adlai and Calvet Laurent (2001): Forecasting multifractal volatility, *Journal of Econometrics* 105: 27–58.

Fisher, Adlai and Calvet Laurent (2002): Multifractality in asset returns: Theory and evidence, *The Review of Economics and Statistics* 84: 381–406.

Fisher, Adlai and Calvet Laurent (2004): How to forecast long-run volatility: regime-switching and the estimation of multifractal processes, *Journal of Financial Econometrics* 2: 49–83.

Fleming, Jeff and Kirby Chris (2011): Long memory in volatility and trading volume, *The Journal of Banking and Finance* 35: 1714–1726.

Liu, Ruipeng and Thomas Lux (2005): Long memory in financial time series: Estimation of the bivariate multi-fractal model and its application for value-at-risk, Working Paper, University of Kiel.

Lobato, Ignacio N. and Savin Nathan E. (2002): Real and spurious long-memory properties of stock market data, *Journal of Business and Economic Statistics* 16: 261–283.

Lux, Thomas (2008): The Markov-Switching multi-fractal model of asset returns: GMM estimation and linear forecasting of volatility, *Journal of Business and Economic Statistics* 26: 194–210.

Ma, Gang (2005): Multiscaling trading time, Working Paper, Brandeis University.

Mandelbrot, Benoit B., Adlai Fisher, and Laurent Calvet (1997): A multifractal model of asset returns, Working Paper, Cowles Foundation.

Mandelbrot, Benoit B. and J. W. Van Ness (1968): Fractional Brownian motion, fractional noises and applications, *SIAM Review* 10: 422–437.

Peng, Chuang-Kang, Sergey V. Buldyrev, Shlomo Havlin, Michael Simons, H. Eugene Stanley, and Ary L. Goldberger (1994): Mosaic organization of DNA nucleotides, *Physical Review E* 49: 1685–1689.

Sun, Wei, Svetlozar Rachev, Ye Chen, and Frank J. Fabozzi (2009): A new approach for using Lévy processes for determining high frequency value-at-risk predictions, *European Financial Management* 15: 340–361.

Xu, Zhaoxia and Gençay Ramazan (2003): Scaling, self-similarity and multifractality in FX markets, *Physica A* 323: 578–590.

A Backtesting VaR

Market risk models predict VaR with error. The validity of a VaR prediction model is measured based on predicted versus actual loss levels. To evaluate the out-of-sample performance of the proposed models we follow the concept of Christoffersen (1998). The indicator (or hit) function $I_t = \mathbf{1}_{\{R_{t,h} < -VaR^{\alpha}_{t,h}\}}$ represents the history of observations, $t = 1, ..., T$, for which losses in excess of the predicted VaR occur.

A.1 Unconditional Coverage

When a VaR model is designed perfectly, the number of observations that fall outside the predicted VaR should be exactly in line with the given VaR level, such that $\mathbb{E}(I_t|\mathcal{F}_t) = \alpha$ holds. Hence, the test of unconditional coverage is

$$H_0 : \mathbb{E}(I_t|\mathcal{F}_t) = \alpha \quad \text{vs.} \quad H_1 : \mathbb{E}(I_t|\mathcal{F}_t) \neq \alpha.$$

Under the null hypothesis, the likelihood-ratio (LR) test statistic follows as

$$LR_{uc} = -2\ln[\mathcal{L}(\alpha)/\mathcal{L}(\widehat{\alpha})] \sim \chi^2(1), \tag{19}$$

where $\mathcal{L}(\alpha)$ is the binomial likelihood with parameter α and $\widehat{\alpha} = \frac{1}{T}\sum_{t=1}^{T} I_t$ is the maximum likelihood estimator of α.

A.2 Independence

Besides the above requirement VaR violations should be independent, which requires an additional test. Let n_{ij} denote the number of observations for which $I_t = j$ occurred following $I_{t-1} = i$ and assume that $\{I_t\}$ is a first-order Markov chain with transition probabilities $\pi_{ij} = \mathbb{P}(I_t = j | I_{t-1} = i)$. This yields the likelihood

$$\mathcal{L}(\Pi) = (1 - \pi_{01})^{n_{00}} \pi_{01}^{n_{01}} (1 - \pi_{11})^{n_{10}} \pi_{11}^{n_{11}}.$$

Maximum likelihood estimators for the transition probabilities are:

$$\widehat{\pi}_{01} = \frac{n_{01}}{n_{00} + n_{01}}, \quad \text{and} \quad \widehat{\pi}_{11} = \frac{n_{11}}{n_{10} + n_{11}}.$$

Under the null hypothesis of independence, $\mathbb{P}(I_t = 0) = \pi_0 = \pi_{01} = \pi_{11}$, which implies

$$\mathcal{L}(\pi_0) = (1 - \pi_0)^{n_{00} + n_{10}} \pi_0^{n_{01} + n_{11}}.$$

The maximum likelihood estimate for $\widehat{\pi}_0$ is

$$\widehat{\pi}_0 = \frac{n_{01} + n_{11}}{n_{00} + n_{10} + n_{01} + n_{11}}.$$

Based on $\widehat{\pi}_0$ and $\widehat{\Pi}$, the independence LR test statistic is

$$LR_{ind} = -2 \ln[\mathcal{L}(\widehat{\pi}_0) / \mathcal{L}(\widehat{\Pi})] \sim \chi^2(1). \tag{20}$$

A.3 Conditional Coverage

The LR_{ind} statistic (20) tests for independence, but it does not take coverage into account. Christoffersen (1998) therefore proposes a combined test statistic:

$$\begin{aligned} LR_{cc} &= LR_{uc} + LR_{ind} \\ &= -2\ln[\mathcal{L}(\alpha)/\mathcal{L}(\widehat{\Pi})] \sim \chi^2(2). \end{aligned} \tag{21}$$

B Tables

Table 1: EUR/USD Return Statistics

Summary statistics of EUR/USD returns for various aggregation levels h. The Phillips-Perron unit root test statistics indicate that the null has to be rejected in favor of the stationarity alternative. Sample period: January 5, 2006 to December 31, 2007.

Frequency	5 minutes	1 hour	1 day	1 week
Mean	1.4×10^{-6}	0.00002	0.00040	0.00205
Std. Dev.	0.0003	0.0009	0.0045	0.0109
Skewness	0.85	0.42	0.05	-0.48
Kurtosis	124.73	16.78	4.08	3.23
AC of $R_t(1)$	0.068	-0.001	-0.045	-0.179
AC of $R_t(2)$	-0.023	0.005	-0.001	0.083
AC of $R_t(3)$	-0.015	0.016	0.023	-0.118
AC of $\|R_t\|(1)$	0.176	0.133	-0.017	0.047
AC of $\|R_t\|(2)$	0.160	0.075	0.024	0.136
AC of $\|R_t\|(3)$	0.135	0.062	0.031	0.049

Table 2: VaR Forecasts

Results of 1%-VaR predictions for 12-hour ($h = 144$) and daily ($h = 288$) forecast horizons. $\overline{VaR}$ denotes average sample VaR. LR-statistics are as defined in Appendix A, where * and ** denote rejection of the null at the 5 and 1 percent significance levels, respectively (corresponding p-values in parenthesis). All VaR forecasts are based on a moving window which consists of 288 5-minute high frequency EUR/USD returns. MMAR denotes the multifractal model of asset returns, HS is the historical simulation and GARCH refers to the generalized autoregressive conditional heteroskedasticity model.

Horizon	Size		MMAR	HS	GARCH
12-hour	959	$\overline{VaR}$	0.0098	0.0070	0.0080
		$\widehat{\alpha}$	1.04	2.60	2.19
		LR_{uc}	0.02	17.34**	10.24**
			[0.895]	[0.000]	[0.001]
		LR_{ind}	0.21	1.34	0.94
			[0.646]	[0.247]	[0.332]
		LR_{cc}	0.23	18.68**	11.18**
			[0.892]	[0.000]	[0.004]
1 day	479	$\overline{VaR}$	0.0137	0.0098	0.0109
		$\widehat{\alpha}$	0.84	1.46	1.25
		LR_{uc}	0.14	10.24**	0.29
			[0.709]	[0.001]	[0.593]
		LR_{ind}	0.07	0.07	0.15
			[0.795]	[0.795]	[0.696]
		LR_{cc}	0.21	10.31**	0.44
			[0.902]	[0.006]	[0.803]

C Figures

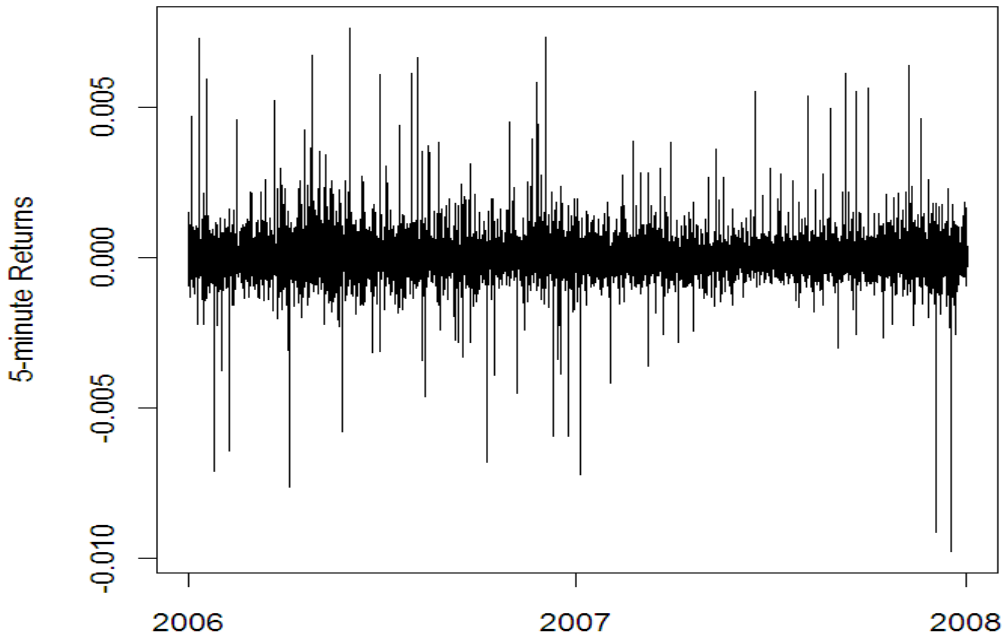

Figure 1: Continuously compounded 5-minute EUR/USD returns. Sample period: January 5, 2006 to December 31, 2007.

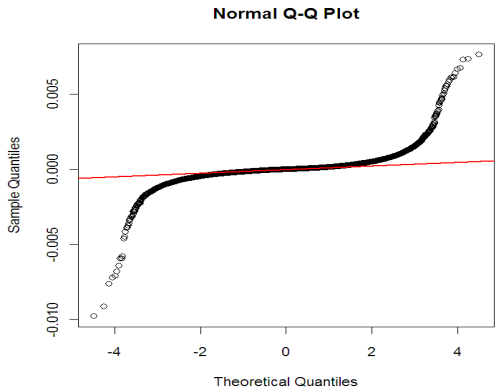

(a) 5-minute EUR/USD Returns

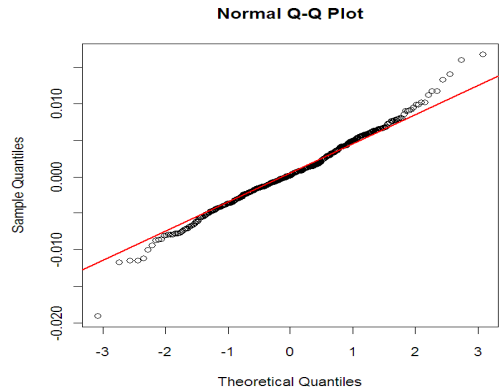

(b) Daily EUR/USD Returns

Figure 2: Quantile-Quantile (Q-Q) plots of EUR/USD returns for 5-minute in panel (a) and daily sampling in panel (b) are given. In both panels the empirical quantiles (points) are plotted against the theoretical quantiles of a normal distribution (straight line). If the two distributions are similar, the point of the empirical distribution should lie on the line. With an increasing sampling frequency the returns are less tailed, but have still fater tails as the normal distribution. Sample period: January 5, 2006 to December 31, 2007.

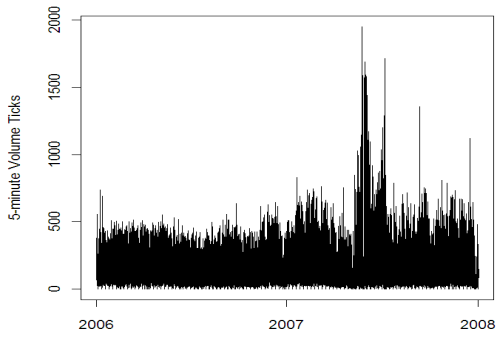

(a) 5-minute Volume Ticks

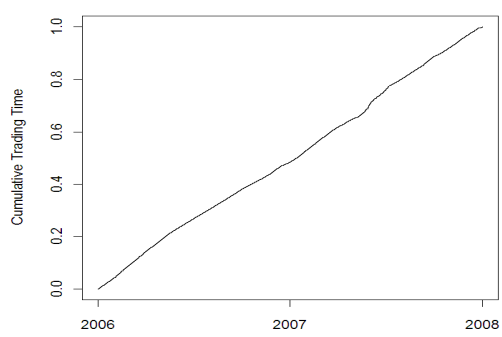

(b) Cumulative Trading Time

Figure 3: Panel (a) illustrates 5-minute EUR/USD trading ticks V_t. A tick is a trade through the platform for the standard minimum amount of about 3-5 million euro. Panel (b) contains the cumulative trading time Θ_t. Sample period: January 5, 2006 to December 31, 2007.

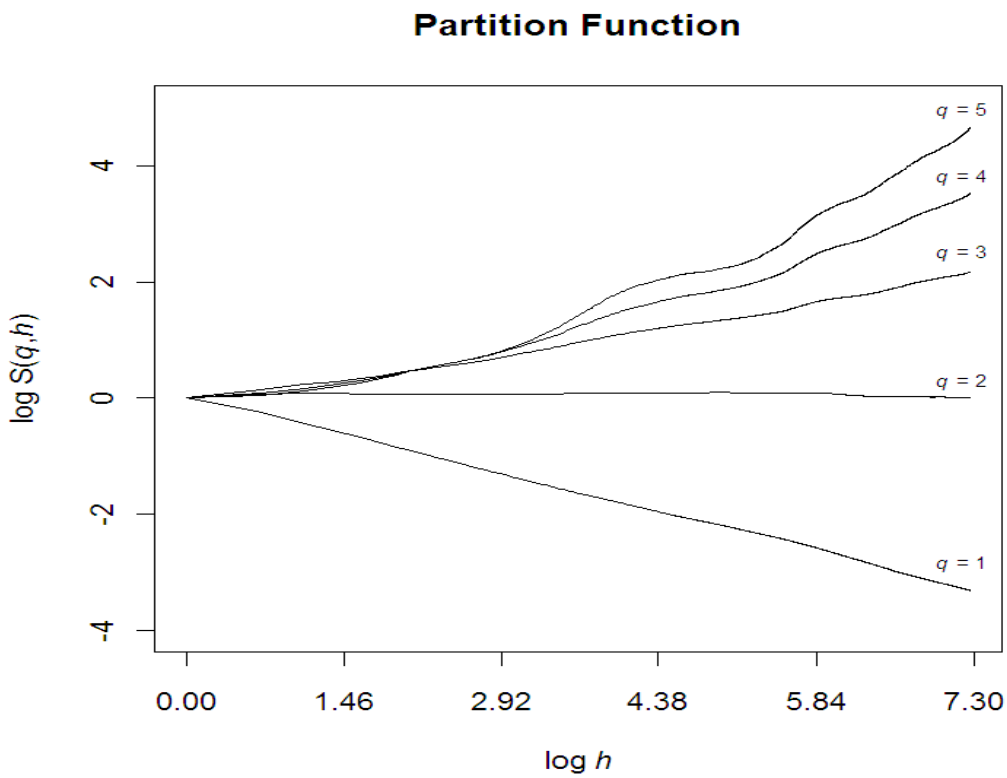

Figure 4: EUR/USD return partition functions for moments $1 \leq q \leq 5$. For each moment q the curves represent the ranges $\log h$ from 5-minutes up to one week. 5-minutes correspond to $\log 1 = 0$ and one week is $\log 1440 = 7.27$. The renormalization for all partition functions is $\log S(q, \text{5-minutes}) = 0$. This procedure allows us to plot all curves on the same graph. Sample period: January 5, 2006 to December 31, 2007.

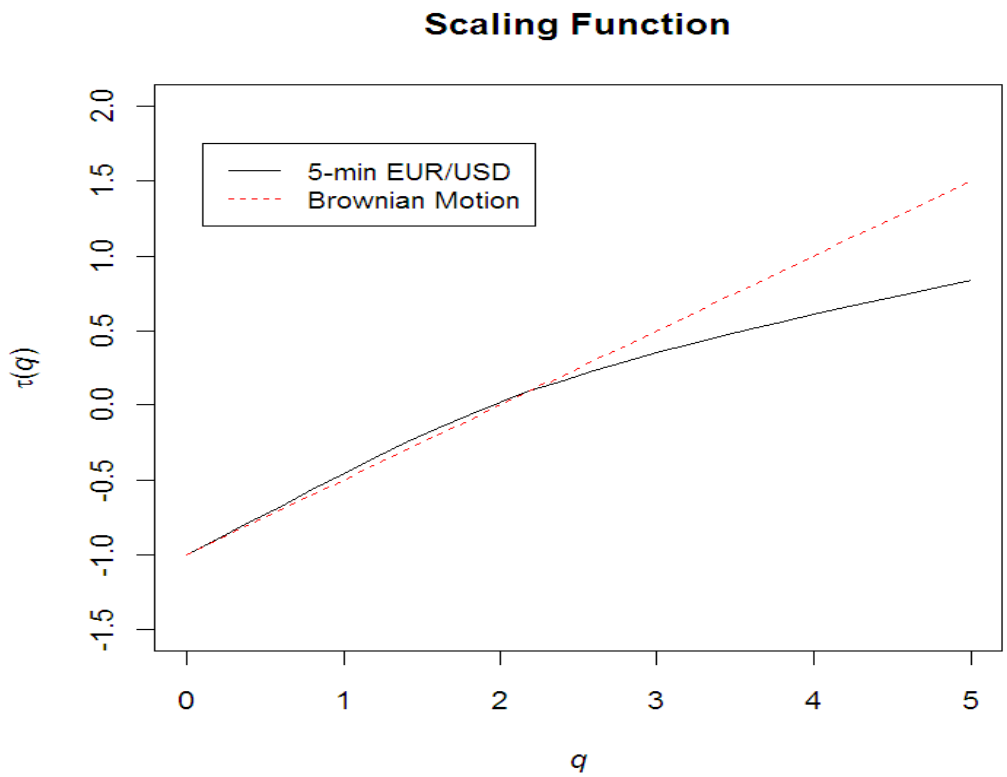

Figure 5: Estimated scaling functions of EUR/USD returns and Brownian motion. For each partition function $S(q, h)$, we estimate the slope by OLS regression (5) to obtain $\tau(q)$. The scaling function of ordinary Brownian motion (dashed line) is linear which is typical for unifractal processes. Multifractal processes have non-linear scaling functions (see Definition 2.1). The estimated $\tau(q)$ (solid line) of EUR/USD returns is non-linear and concave. Sample period: January 5, 2006 to December 31, 2007.

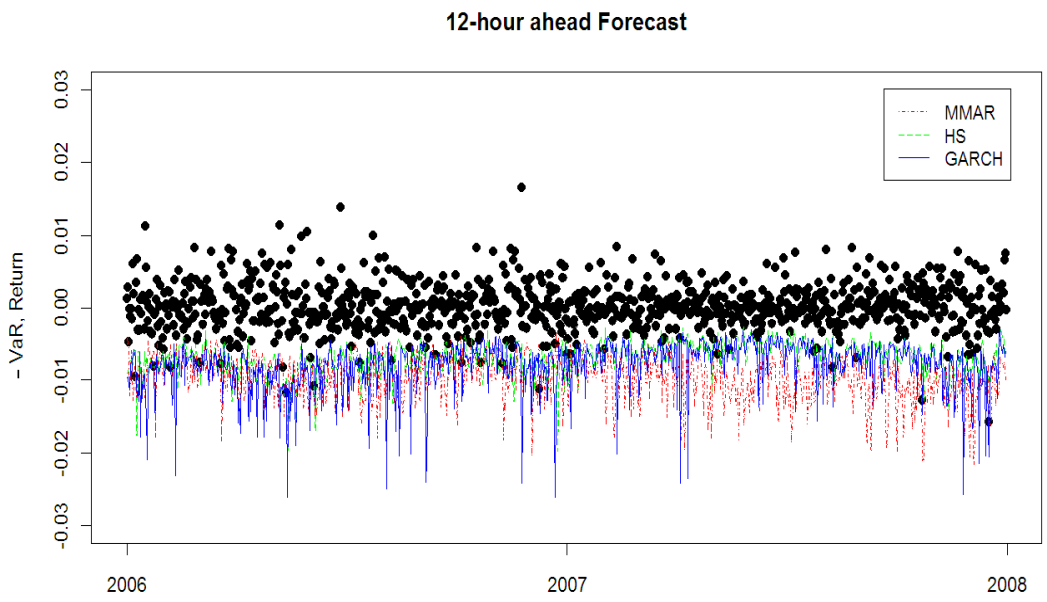

Figure 6: 12-hour ahead VaR forecasts for our MMAR approach (dash dotted line), historical simulation (dashed line) and GARCH (solid line). The forecasting period is January 6, 2006 to December 31, 2007.