Forecasting Value-at-Risk Using High Frequency Information*

Huiyu Huang†
Grantham, Mayo, Van Otterloo
and Company LLC

Tae-Hwy Lee‡
University of California, Riverside

February 2013

Abstract

In prediction of quantiles of *daily* S&P 500 returns we consider how we use high-frequency 5-minute data. We examine methods that incorporate the high frequency information either indirectly through combining forecasts (using forecasts generated from returns sampled at different intra-day interval) or directly through combining high frequency information into one model. We consider subsample averaging, bootstrap averaging, forecast averaging methods for the indirect case, and factor models with principal component approach for both cases. We show, in forecasting daily S&P 500 index return quantile (VaR is simply the negative of it), using high-frequency information is beneficial, often substantially and particularly so in forecasting downside risk. Our empirical results show that the averaging methods (subsample averaging, bootstrap averaging, forecast averaging), which serve as different ways of forming the ensemble average from using high frequency intraday information, provide excellent forecasting performance compared to using just low frequency daily information.

Key Words: VaR, Quantiles, Subsample averaging, Bootstrap averaging, Forecast combination, High-frequency data.
JEL Classifications: C53, G32, C22

---

*The opinions expressed in this paper are those of the authors and do not necessarily reflect those of Grantham, Mayo, Van Otterloo and Company LLC.

†GMO Emerging Markets, 2150 Shattuck Ave, Suite 900, Berkeley, CA 94704, U.S.A. Phone: +1 (510) 809-4594. Email: evelyn.huang@gmo.com.

‡Corresponding Author. Department of Economics, University of California, Riverside, CA 92521-0427, U.S.A. Phone: +1 (951) 827-1509. E-mail: taelee@ucr.edu.

1 Introduction

Due to increasing fragility in financial markets and the extensive use of derivative products, effective management of financial risks has become ever more important. A risk measurement, namely "Value-at-Risk" (VaR), has received great attention from both regulatory and academic world due to its simplicity despite the defective property that it lacks the coherency in the sense of Artzner *et al* (1999). Numerous papers have studied various aspects of VaR methodology. Typically the VaR is computed in *daily* frequency, such as for the 1 to 10-day ahead forecasts of the tail quantiles. In this paper, we discuss how we can improve the accuracy of daily out-of-sample VaR forecasts by using high frequency intra-day information.

Consider a financial return series $\{r_t\}_{t=1}^T$, generated by the probability law $\Pr(r_t \le r|\mathcal{F}_{t-1}) \equiv F_t(r)$ conditional on the information set $\mathcal{F}_{t-1}$ (σ-field) at time $t-1$. Suppose $\{r_t\}$ admits the stochastic process

$$r_t = \mu_t + \varepsilon_t = \mu_t + \sigma_t z_t, \tag{1}$$

where $\mu_t = E(r_t|\mathcal{F}_{t-1})$, $\sigma_t^2 = E(\varepsilon_t^2|\mathcal{F}_{t-1})$, and $\{z_t\} \equiv \{\varepsilon_t/\sigma_t\}$ has the conditional distribution function $G_t(z) \equiv \Pr(z_t \le z|\mathcal{F}_{t-1})$. The VaR with a given tail probability $\alpha \in (0,1)$ is defined as the negative of conditional quantile (denoted $q_t(\alpha)$)

$$\Pr(r_t \le q_t(\alpha)|\mathcal{F}_{t-1}) = F_t(q_t(\alpha)) = \alpha. \tag{2}$$

The quantile can be estimated either from inverting the distribution function

$$q_t(\alpha) = F_t^{-1}(\alpha) = \mu_t + \sigma_t G_t^{-1}(\alpha), \tag{3}$$

where the second equality holds if $F_t(r)$ belongs to location-scale family, or from the quantile regression. A quantile model involves the specification of the conditional distribution $F_t(\cdot)$ or the quantile regression function $q_t(\alpha)$. The former can also be estimated by its components, namely, μ_t, σ_t^2, $G_t(\cdot)$. For instance, Clements, Galvao and Kim (2008) make distributional assumptions for future daily returns coupled with forecasts of realized volatility to get quantile forecasts. In this paper we focus on the latter, the quantile regression with incorporating the high frequency information.

The rich dynamics in ultra-high-frequency financial data has extensively been used in estimation of quadratic variation (integrated variance) in so called realized volatility literature, and also used in high frequency duration models. There is considerable amount of literature addressing the use of high frequency data for realized volatility and duration models, whereas little work has been done for quantiles which is practically more relevant and central to financial risk management. This paper contributes to the literature by examining whether it pays to incorporate the intraday data and, more importantly, how to incorporate them to achieve better performance for forecasting daily quantiles (VaR is simply the negative of it).

The existing literature on high-frequency data is mainly devoted to volatility issue of estimating quadratic variation, mixed data sampling (MIDAS) approach of Ghysels et al (2006), and the autoregressive conditional duration models of Engle and Russell (1998). In this paper we use the high-frequency data for a different purpose. Motivated by the bagging (bootstrap aggregating) approach of Breiman (1996), which replicates the true distribution by bootstrap distribution in computing the ensemble average, we consider returns computed from high frequency intraday observations as multiple replications of the true distribution of daily returns. Suppose there are hourly time series available while we are interested in generating daily forecasts (e.g., 1 day ahead or 10 day ahead). Assuming that each of hourly information are generated from identical distribution (without intraday pattern or diurnal cycles), we may consider the high frequency (hourly) series as multiple replications of lower frequency (daily) series. In that context, we regard the high frequency data as *subsamples* of daily series in different time within a day.

Viewing the high frequency information in this manner leads us to consider the subsample averaging idea from the realized volatility (RV) literature and use it for forecasting quantiles instead. The original subsampling idea can be traced back to Zhou (1996) and further studied by Zhang, Mykland and Aït-Sahalia (ZMA, 2005) and Barndorff-Nielsen, Hansen, Lunde and Shephard (BNHLS, 2011). Andersen, Bollerslev, Diebold, and Labys (ABDL, 2001) and Barndorff-Nielsen and Shephard (2002) establish that RV, defined as the sum of squared intraday returns of small intervals, is an asymptotically unbiased estimator of the unobserved quadratic variation as the interval length approaches zero. However, in the presence of market microstructure noise, such nice property of RV is contaminated. Recent innovative works investigating this issue include Aït-Sahalia, Mykland and Zhang (AMZ, 2005), Bandi and Russell (2005), Hansen and Lunde (2006), ZMA (2005), and BNHLS (2008, 2011). When the observed price process is the true underlying price process plus microstructure noise, it is shown that RV will be overwhelmed by the noise and explodes when the sampling frequency approaches infinity. Therefore, it may be optimal to sample less frequently than the case in the absence of noise. Zhou (1996) proposes for the first time an unbiased data-driven estimator of volatility and a subsample averaging volatility estimator. ZMA (2005) and BNHLS (2011) establish improved estimators for quadratic variation through subsampling. The bias-adjusted estimator of ZMA (2005) based on the subsample averaging method is able to eventually push the estimation bias to zero. BNHLS (2011) show that subsampling is highly advantageous for RV estimators based on discontinuous kernels.

Motivated by this subsampling approach that is shown to outperform ones using directly the highest frequency series and to avoid arbitrariness in choosing sampling frequencies, we propose a subsample averaging quantile forecast in construction similar to the bagging approach and to the simple average combination of forecasts. We also compare this approach with other forecasts utilizing the highest frequency information directly. There are vast amount of literature on predicting daily market returns using daily close data only. We deem that proper use of high-frequency intraday data should not only be helpful for achieving more accurate estimation of volatility, but also be beneficial for forecasting daily VaR/quantiles. The question is how to incorporate the vast amount of intraday high frequency information for daily low frequency modeling or forecasting of VaR/quantiles.

In Huang and Lee (2010), the relative advantages of combination of forecasts (CF) over combination of information (CI) are discussed. Under circumstances of highly correlated predictors, important predictors omitted, and/or low signal-to-noise ratio, CF is more likely to win over CI. In this paper, we proceed by considering these two alternate approaches of using intraday information and compare the CF approaches (combination of individual forecasts obtained from using one 5-min intraday information at a time) with the CI approaches (combination of all 5-min intraday information into one model) for their forecasting ability for daily market return VaR/quantiles prediction. It is well-known that the mean of market daily return is very hard to predict whereas its quantiles may be predictable, particularly in tails. See Lee and Yang (2006) for some evidence from using bagging. Practically, quantile forecasts are very important for risk management purpose. See for example Chernozhukov and Umantsev (2001). Therefore, we implement the quantile forecasts using high-frequency data through various CF and CI methods. To avoid being overly parameterized, we consider subsample averaging, some CF methods, and factor models with principal component approach. We compare their prediction errors and find that in daily S&P 500 return quantile forecasts, generally CF with simple weighting schemes, subsample averaging in particular, are found to be superior to others.

The rest of the paper is organized as follows. Section 2 describes the data and its organizations. Section 3 discusses the quantile forecasting methods we consider, and Section 4 presents out-of-sample VaR/quantile forecasting results. Section 5 concludes.

2 Data

The data we use consists of S&P 500 index values at 5-minute intervals recorded in between 9:35 a.m. and 4:00 p.m. from June 9, 1997 to May 30, 2003, a total of 1,501 days and 117,078 observations.[1] In cleaning the data, those periods of market closings are treated as no variation in index values, thus there exists 78 ticks each trading day. From this pool of 5-minute index data, we construct 78 "daily" returns, each having a time span of twenty-four hours but pointing at different times in a day.

Specifically, let us denote the daily close return by $r_t^{(0)} = p_t^{4:00\text{pm}} - p_{t-1}^{4:00\text{pm}}$, where $p_t^{4:00\text{pm}}$ denotes the logarithm of the S&P 500 index value at 4:00 p.m. on day t. $\{p_t^{4:00\text{pm}}\}_{t=1}^T$ is one sub-sample of the entire S&P 500 5-minute index data, obtained by systematic sampling at time 4:00 p.m. on each day hence it represents one 78th of the entire 5-minute high-frequency information. Therefore, $r_t^{(0)}$ is the daily return based on this particular sub-sample, which is also the conventional daily close return.

We define a sub-sample daily return, $r_t^{(1)} = p_t^{9:35\text{am}} - p_{t-1}^{9:35\text{am}}$, and similarly other sub-sample daily returns $r_t^{(j)}$ for $j = 2, \ldots, m = 77$, and we have $r_t^{(m=77)} = p_t^{3:55\text{pm}} - p_{t-1}^{3:55\text{pm}}$.

3 Forecasting Quantiles Using High Frequency Information

In this section we describe various methods of using high frequency information in quantile forecasting. These includes Daily-Close quantile forecasts, Bagging-Daily-Close quantile forecast, Subsample-Averaging (SA) quantile forecast, Combining-Forecast (CF) quantile forecast, Combining-Information (CI) quantile forecast, and some variations of SA, CF, CI methods, as presented in Table 1.

3.1 Daily Close Quantile Forecasts

Suppose our objective is to predict the α-quantile of daily close return $r_{T+h}^{(0)}$ conditioning on the information up to time T: $Q_\alpha(r_{T+h}^{(0)}|X_T)$, where h denotes the forecast horizon. Typically, $X_T \subset \mathcal{F}_T$

[1] We are grateful to George Jiang who generously shared this high-frequency intraday data with us. The data are extracted from the contemporaneous index levels recorded with the quotes of SPX options from the CBOE.

is the vector of past values of daily returns in the same daily frequency as $\{r_t^{(0)}\}_{t=1}^T$. As in Lee and Yang (2006) we use a quantile regression model

$$Q_\alpha\left(r_{t+h}^{(0)}|X_t\right) = X_t^{(0)\prime}\beta^{(0)}(\alpha) + \varepsilon_{t+h}, \tag{4}$$

to predict the quantiles of daily S&P 500 return with $h = 1$, using Chernozhukov and Umantsev's (2001) linear polynomial model (with the quadratic term capturing the volatility clustering effect on quantiles) such that

$$X_t^{(0)} = \left(1 \ r_t^{(0)} \ r_t^{(0)^2}\right)'. \tag{5}$$

The quantile forecast $\hat{Q}_\alpha\left(r_{T+h}^{(0)}|X_T\right) = X_T^{(0)\prime}\hat{\beta}^{(0)}(\alpha)$ from this quantile regression model is denoted "Daily-Close" in Table 1, which is estimated by minimizing the "tick" (or check) function of Koenker and Basset (1978)

$$\rho_\alpha(\varepsilon) = [\alpha - 1(\varepsilon < 0)]\,\varepsilon. \tag{6}$$

To examine whether it pays to incorporate the intraday data, we use this Daily-Close as the benchmark model to be compared with other models incorporating intraday information in our empirical analysis in Section 4.

3.2 CI Quantile Forecasts

Besides the past realization of daily close return $r_t^{(0)}$, now we have other sub-samples $r_t^{(j)}$, $j = 1, \ldots, m$, constructed from the high-frequency intraday data. For the purpose of accommodating more information, it is natural to expand X_t to include the rich dynamics in all other intraday sub-samples $\left(r_t^{(0)} \ r_t^{(1)} \ \ldots \ r_t^{(m)}\right)$. As an example, for the linear polynomial model, to utilize high frequency information, the new regressors will be

$$X_t^m = \left(1 \ r_t^{(0)} \ r_t^{(0)^2} \ r_t^{(1)} \ r_t^{(1)^2} \ \ldots \ r_t^{(m)} \ r_t^{(m)^2}\right)'. \tag{7}$$

By doing this, one follows the CI scheme. That is, combining all the high-frequency information into one model altogether and generate an ultimate forecast:

$$\hat{Q}_\alpha^m\left(r_{T+h}^{(0)}|X_T^m\right) = X_T^{m\prime}\hat{\beta}_m(\alpha), \tag{8}$$

where the coefficient $\hat{\beta}_m(\alpha)$ is estimated by running the quantile regression of $r_t^{(0)}$ on X_{t-h}^m. Denote this quantile forecast "CI-Unrestricted" in Table 1.

Now re-define $X_t^m = \left(r_t^{(0)} \ r_t^{(0)^2} \ r_t^{(1)} \ r_t^{(1)^2} \ \ldots \ r_t^{(m)} \ r_t^{(m)^2}\right)'$ excluding the constant term. We note that this is a fairly large vector with dimension 156 ($m = 77$ in our empirical study). Moreover, consecutive subsample returns, $r_t^{(j)}$ and $r_t^{(j+1)}$, may be highly correlated. Therefore methods that can mitigate high dimension and multicollinearity problems are desirable in order to improve out-of-sample forecast performance. We consider factor model with principal component approach (see Stock and Watson 2002, 2011). The procedure here starts by assuming that $\left(X_t^m, r_{t+h}^{(0)}\right)$ admits a factor model representation with k common latent factors f_t

$$X_t^m = \Lambda f_t + e_t, \tag{9}$$

$$Q_\alpha(r_{t+h}^{(0)}|X_t^m) = \left(1\ f_t'\right) \hat{\beta}_f(\alpha) + \varepsilon_{t+h}, \tag{10}$$

and $\hat{\beta}_f(\alpha)$ is obtained by running quantile regression of $r_t^{(0)}$ on $\left(1\ \hat{f}_{t-h}'\right)'$. This quantile forecast is denoted as "CI-PC" in Table 1.[2]

If k is unknown, it can be estimated by some conventional information criteria (IC) or ones tailored to quantile forecasting models. Machado (1993) provides variants of the Schwarz Information Criterion for model selection in M-estimation which include quantile regression as a special case. According to Koenker (2005), for median regression, the criterion $SIC = \ln\left(\hat{\zeta}_T\right) + \frac{1}{2} p \ln T$, where $\hat{\zeta}_T \equiv \frac{1}{T-h}\sum_{t=1}^{T-h} \rho_\alpha\left(\varepsilon_{t+h}\right)$, $\alpha = 0.5$, $\rho_\alpha(\cdot)$ is the tick loss function in (6), and p is the dimension of the model, yields a consistent model selection procedure under certain assumptions. To our knowledge, however, a consistent model selection procedure and corresponding information criterion for *factor* quantile regressions have not been established yet. We leave the theoretical study of this procedure for our future research. In our empirical study, we consider conventional information criteria such as AIC and BIC where the estimated number of factors k is selected by $\min_{1\le k\le k_{max}} \mathrm{IC}_k = \ln(\mathrm{SSR}(k)/T) + g(T)k$, where k_{max} is the hypothesized upper limit for the true number of factors and SSR is the sum of squared residuals. Here we simply extend the formulas for AIC (with $g(T) = 2/T$) and BIC (with $g(T) = \ln T/T$) straightly into the quantile case, but substituting "residual" in the above-mentioned term SSR by tick loss $\left[\alpha - \mathbf{1}\left(r_{t+h}^{(0)} < \hat{Q}_\alpha\left(r_{t+h}^{(0)}|X_t^m\right)\right)\right]\left(r_{t+h}^{(0)} - \hat{Q}_\alpha\left(r_{t+h}^{(0)}|X_t^m\right)\right)$, for $t = 1, \ldots, T-h$. Alternatively, we can fix k ex-ante at some small values such as 1, 2, or 3.

[2] To better comprehend this factor representation, if we think stock market as a pool of vast amount of information and price is determined by aggregate behavior of market participants, there may exist a few fundamental shocks to the market in the day that ultimately determine daily return at close. This small set of shocks (main forces) may be captured by the latent factors in the factor model (9), or technically, by a few principal components of the big explanatory variable set X_t^m, which contains as large as 156 variables capturing levels and volatilities of the market throughout the trading day.

3.3 CF Quantile Forecasts

The conventional CF methodology with one fixed forecast target, i.e., the α-quantile of $r_{T+h}^{(0)}$, is implemented as follows.

1. CF Step 1: Compute quantile forecasts from regressing $r_t^{(0)}$ on each individual sub-sample,

$$\hat{Q}_\alpha^{(j)}\left(r_{T+h}^{(0)}|X_T^{(j)}\right) = X_T^{(j)\prime}\hat{\beta}^{(j)}(\alpha), \tag{11}$$

   where $X_t^{(j)} = \left(1\ r_t^{(j)}\ r_t^{(j)2}\right)'$, and $\hat{\beta}^{(j)}(\alpha)$ is obtained by estimating the quantile regression of $r_t^{(0)}$ on $X_{t-h}^{(j)}$, for $j = 0, 1, \ldots, m$.

2. CF Step 2: Combine the quantile forecasts from Step 1 by some weighting methods for forecast combination. A simplest example of the combination methodology is the simple average:

$$\bar{Q}_\alpha^{\text{CF-Mean}}\left(r_{T+h}^{(0)}|X_T^m\right) = \frac{1}{m+1}\sum_{j=0}^{m}\hat{Q}_\alpha^{(j)}\left(r_{T+h}^{(0)}|X_T^{(j)}\right), \tag{12}$$

   i.e., taking the mean of all the individual quantile forecasts (denoted "CF-Mean" in Table 1). One can also use the median of the individual quantile forecasts as a combined forecast (denoted "CF-Median" in Table 1).

In Step 1, given the high-frequency 5-minute data we have, there are in total 78 individual quantile forecasts to be combined. Besides the simple methods such as combining these 78 forecasts by simple average or median without estimating any weighting parameters, we also use the regression approach forecast combination (Granger and Ramanathan, 1984) to explore more of the cross-sectional information contained in these 78 individual forecasts. Obviously, this regression approach is not working well because of the high dimensionality. Therefore, we consider the principal component methodology for combining the quantile forecasts generated in Step 1 and form the VaR forecast (denoted "CF-PC" in Table 1) by the following factor model of the quantile forecasts

$$\hat{\mathbf{Q}}_\alpha^m\left(r_t^{(0)}|X_{t-h}^m\right) = \Gamma\hat{F}_t + v_t, \tag{13}$$

where

$$\hat{\mathbf{Q}}_\alpha^m\left(r_t^{(0)}|X_{t-h}^m\right) = \left[\hat{Q}_\alpha^{(0)}\left(r_t^{(0)}|X_{t-h}^{(0)}\right), \hat{Q}_\alpha^{(1)}\left(r_t^{(0)}|X_{t-h}^{(1)}\right), \ldots, \hat{Q}_\alpha^{(m)}\left(r_t^{(0)}|X_{t-h}^{(m)}\right)\right]'. \tag{14}$$

Then we estimate $\hat{\beta}_F(\alpha)$ from

$$Q_{\alpha}^{\text{CF-PC}}\left(r_{t+h}^{(0)}|X_t^m\right) = \left(1\ \hat{F}_{t+h}'\right)\beta_F(\alpha) + \eta_{t+h}, \tag{15}$$

to obtain the final CF-PC quantile forecast

$$\hat{Q}_{\alpha}^{\text{CF-PC}}\left(r_{T+h}^{(0)}|X_T^m\right) = \left(1\ \hat{F}_{T+h}'\right)\hat{\beta}_F(\alpha). \tag{16}$$

The size of the common factor set $\hat{F}_t$, i.e., the number of principal components of $\hat{\mathbf{Q}}_{\alpha}^m\left(r_t^{(0)}|X_{t-h}^m\right)$, can be determined again by information criteria AIC or BIC as discussed in Section 3.2, or fixed ex-ante at some small values such as 1, 2, or 3.

3.4 Subsample-Averaging Quantile Forecasts

Inspired from the RV literature and as discussed in Section 1, here we propose a new methodology that extends the subsample idea for RV estimation (using intraday high-frequency information) to daily quantile forecasting. We call it *subsample average* (SA), which is unique to forecasting using high-frequency data. This approach is similar to the CF-Mean procedure discussed in 3.3 but it differs from CF-Mean in Step 1: instead of fixing left-hand-side (LHS) variable in (11) as the 4 p.m. daily close return $r_t^{(0)}$, we modify (11) by setting LHS variable as the corresponding subsample return $r_t^{(j)}$ for each $j = 0, 1, \ldots, m$.

1. SA Step 1:
$$\tilde{Q}_{\alpha}^{(j)}\left(r_{t+h}^{(j)}|X_t^{(j)}\right) = X_t^{(j)\prime}\tilde{\beta}^{(j)}(\alpha), \tag{17}$$
which give the SA quantile forecasts: $\tilde{Q}_{\alpha}^{(j)}\left(r_{T+h}^{(j)}|R_T^{(j)}\right) = X_T^{(j)\prime}\tilde{\beta}^{(j)}(\alpha)$ for $j = 0, 1, \ldots, m$.

2. SA Step 2: Taking simple average, we get the quantile forecast
$$\bar{Q}_{\alpha,T+h}^{\text{SA-Mean}}\left(r_{T+h}^{(j)}|X_T^m\right) = \frac{1}{m+1}\sum_{j=0}^{m}\tilde{Q}_{\alpha}^{(j)}\left(r_{T+h}^{(j)}|X_T^{(j)}\right), \tag{18}$$
which is denoted "SA-Mean" in Table 1. Similarly, "SA-Median" is to take median instead of average.

We compare the out-of-sample tick loss of SA-Mean and SA-Median against $r_{T+h}^{(0)}$, i.e., $\varepsilon = r_{T+h}^{(0)} - \bar{Q}_{\alpha,T+h}^{\text{SA}}\left(r_{T+h}^{(j)}|X_T^m\right)$ in Equation (6). Although the benchmark Daily-Close, CF, and CI quantile forecasts are all computed for $r_{T+h}^{(0)}$ by design, we consider returns computed from high frequency intraday subsample observations as multiple replications of the true distribution of daily returns (even if it is at a different time of a day). This is motivated by the bagging approach which replicates the true distribution by bootstrap distribution in computing the ensemble average. In this sense, we consider the high frequency subsamples just like the bootstrap samples, which both permit us to estimate the ensemble average in daily frequency. Thus we treat the subsample-averaging VaR/quantile forecasts as daily-close VaR/quantile forecasts.

3.5 Bagging Daily Close Quantile Forecasts

Bootstrap-averaging (bagging) replicates the true distribution by bootstrap distribution in computing the ensemble average of a stochastic process. This paper is also motivated from the idea of bootstrap-averaging, and we extend it to subsample-averaging. Subsample-averaging may replicate the true distribution by using returns computed from high frequency data in computing the ensemble average of a stochastic process.

We assume that each of 5-minute intraday information are generated from identical stationary distribution (ignoring intraday pattern or diurnal cycles). Then we may consider the high frequency series as multiple replications of lower frequency (daily) series. Under the assumption of the strict stationarity we use the high frequency data to generate subsamples of daily series in different time within a day. In Section 3.4, we consider high frequency observations as multiple replications of the true distribution of daily returns. The subsample returns, $r_t^{(j)}$ $(j = 0, 1, \ldots, m)$, are considered as $(m + 1)$ replications (draws) of the daily return series.

In this section, we consider bagging. Bagging predictor is a combined predictor formed over a set of training sets to smooth out the "instability" caused by parameter estimation uncertainty and model uncertainty. A predictor is said to be "unstable" if a small change in the training set will lead to a significant change in the predictor (Breiman, 1996). The mechanism of bagging has been explained in various ways. Breiman (1996) uses the Cauchy-Schwarz inequality for a squared error loss. Lee and Yang (2006) extend it to a convex loss (e.g., a tick function for quantiles) using the Jensen's inequality. Bühlmann and Yu (2002) show that, for a nonsmooth unstable predictor, bagging reduces variance of the first order term. In particular, they show that bagging can reduce the mean squared error of forecasts by averaging over the randomness of variable selection. Buja and Stuetzle (2006) and Friedman and Hall (2007) expand a smooth unstable function into linear and higher order terms, and show bagging reduces the variance of the higher order terms. Bagging also stabilizes prediction by equalizing the influence of outlying training observations. Stock and Watson (2011) show that bagging is asymptotically Bayesian shrinkage. Applications of bagging include inflation (Inoue and Kilian 2008), financial volatility (Hillebrand and Medeiros 2010), equity premium (Huang and Lee 2010), short-term interest rates (Audrino and Medeiros 2011), and employment data (Rapach and Strauss 2010).

Bagging for the α-quantile of daily close return $r_t^{(0)}$ is implemented as follows.

1. Bagging Step 1: Construct a bootstrap sample $\{r_t^{(0)*}\}$ according to the empirical distribution of daily close returns. Run the quantile regression in (4) by regressing $r_t^{(0)*}$ on $X_{t-h}^{(0)*} = \left(1 \; r_{t-h}^{(0)*} \; r_{t-h}^{(0)*2}\right)'$, obtain $\hat{\beta}^{(0)*}(\alpha)$, and compute the bootstrap Daily Close quantile forecast

$$\hat{Q}_{\alpha}^{(0)*}\left(r_{T+h}^{(0)*}|X_T^{(0)*}\right) = X_T^{(0)*\prime}\hat{\beta}^{(0)*}(\alpha). \tag{19}$$

   Repeat this B times.

2. Bagging Step 2: Combine the quantile forecasts from Step 1 by some weighting methods for forecast combination. A simplest example of the combination is the simple average:

$$\bar{Q}_{\alpha}^{\text{Bagging}}\left(r_{T+h}^{(0)}|X_T^{(0)}\right) = \frac{1}{B}\sum_{b=1}^{B}\hat{Q}_{\alpha}^{(0)*(b)}\left(r_{T+h}^{(0)*(b)}|X_T^{(0)*(b)}\right), \tag{20}$$

   i.e., taking the simple mean of all the bootstrap Daily Close quantile forecasts (denoted "Bagging-Mean"). One can also use the median of the individual quantile forecasts as a combined forecast (denoted "Bagging-Median"). As they are very similar, we only report the former in Table 1, denoted as "Bagging Daily-Close".

A concern of applying bagging to time series is whether a bootstrap can provide a sound simulation sample for dependent data, for which the bootstrap is required to be consistent. It has been shown that some bootstrap procedure (such as moving block bootstrap) can provide consistent densities for moment estimators and quantile estimators. See Hall, Horowitz, and Jing (1995) and Fitzenberger (1997). Therefore we use block bootstrapping in our empirical study as stock market returns are likely to exhibit time dependence. In the next section, we use the number of bootstrap samples $B = 50$ and the bootstrap block size fixed at 4.

4 Out-of-sample Quantile Forecasting Results

Table 1 presents the performance of each forecasting method for predicting one-day ahead ($h = 1$) daily close S&P 500 return quantiles. The 78 daily returns introduced in Section 2 are calculated by logarithm difference of corresponding index values and multiplied by 100. The size of out-of-sample period is $P = 500$ days from May 29th, 2001 to May 30th, 2003. The in-sample size is $R = 1000$ days, from June 10th, 1997 to May 25th, 2001. We use the rolling window scheme to estimate the parameters and set the size of estimation window at $R = 1000$. Quantile regressions are estimated using the interior-point algorithm by Portnoy and Koenker (1997). We report the out-of-sample mean tick loss ratios of each chosen forecasting scheme over Daily Close benchmark for the left-tail probability parameter $\alpha = 0.01, 0.05, \ldots, 0.99$, as defined in (2).

The results in Table 1 show that, except for CI-Unrestricted which obviously does not work well due to its dimensionality problem, in general we indeed gain by incorporating high-frequency intraday information into predicting daily close return quantiles. Evidently, we observe that generally the "averaging" methods (SA-Mean, SA-Median, CF-Mean, CF-Median, and Bagging) work better than various CI methods, for most values of α. All SA and CF methods with simple weighting schemes improve upon Daily-Close benchmark (less than 1 in loss ratios), sometimes quite substantial. This indicates the stableness of the averaging methods with simple weighting schemes (mean or median).

For those with factor model approaches, the maximum hypothesized number of factors, k_{max}, is set at 15, so number of factors k is chosen within interval $[1, 15]$ for AIC and BIC or fixed at 1, 2, or 3. We see that generally BIC performs better than AIC but usually worse than fixing factor numbers (at 1, 2 or 3). Therefore, the simple rule-of-thumb of fixing k at a small number, that is, using a few fixed number of factors to summarize the entire high-frequency information, works considerably better than estimating k by AIC or BIC.

We also find that CF is more robust than CI. This highlights the merits of CF that are illustrated in Huang and Lee (2010), and consistent with what is often found in the literature about CF.

The percentage loss reductions at tails, especially at the left tails, are much larger than those of middle quantiles. It shows that there are more rooms for the various forecasting schemes to improve upon Daily-Close model at tails by incorporating the high frequency information to forecast daily quantiles. Zoom in on the results for the left tail with $\alpha = 0.01$ which the financial risk management usually concerns with. Relative to the benchmark quantile forecast using Daily-Close returns, Bagging Daily-Close quantile forecast reduces the relative out-of-sample mean forecast tick loss ratio to 91.39% (improving by 8.61%). CF-Mean quantile forecast and CF-Median quantile forecast reduce the loss ratio to 92.26% and 93.39% respectively, which are slightly worse than Bagging

Daily-Close quantile Forecast. The subsample averaging methods (SA-Mean and SA-Median) are generating the best daily quantile forecasts with the out-of-sample mean tick loss ratios at 88% and 89%. That is an astounding 11-12% reduction of tick loss for daily VaR forecasting with $\alpha = 0.01$, which is achieved by incorporating the intraday 5-minute high-frequency information into modeling the lower-frequency daily series.

5 Conclusions

In this paper, we present and propose several methods of incorporating high frequency information in forecasting daily quantiles, which are shown to be particularly useful for lower tail daily Value-at-Risk forecasting. These methods are based on "averaging". Unlike a typical method of model-averaging whose average is taken over different models, our averaging involves only one model but is taken over multiple forecasts that are generated from using multiple lower (daily) frequency datasets constructed from higher (intraday) frequency data. Hence, it is not model averaging, but high-frequency data averaging or combining. We suggest three such data averaging/combining methods – combining forecast (CF), subsample averaging (SA), and bootstrap averaging (Bagging) – as explained in Section 3. All of these three averaging methods are designed to incorporate high frequency (intraday) information into lower frequency (daily) modelling/forecasting.

As demonstrated in Section 4 in forecasting VaR/quantiles of the S&P 500 index return, using high-frequency information is beneficial, often substantially and particularly so in forecasting downside risk. For daily S&P 500 return lower tail out-of-sample VaR forecasts, our empirical results show that the averaging methods via SA, Bagging, and CF (which serve as different ways of forming the ensemble average) from using high frequency intraday information have excellent forecasting performance when compared to just using low frequency daily-close information.

References

Aït-Sahalia, Y., Mykland, P.A. and Zhang, L. (2005), "How Often to Sample a Continuous-time Process in the Presence of Market Microstructure Noise," *Review of Financial Studies* 18: 351-416.

Andersen, T.G., Bollerslev, T., Diebold, F.X. and Labys, P. (2001), "The Distribution of Realized Exchange Rate Volatility," *Journal of the American Statistical Association* 96: 42-55.

Andersen, T.G., Bollerslev, T., Diebold, F.X. and Labys, P. (2003), "Modeling and Forecasting Realized Volatility," *Econometrica* 71: 579-625.

Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). "Coherent Measures of Risk," *Mathematical Finance* 9(3): 203–228.

Audrino, F. and Medeiros, M.C. (2011): "Modeling and forecasting short-term interest rates: The benefits of smooth regimes, macroeconomic variables, and bagging", *Journal of Applied Econometrics* 26(6): 999-1022.

Bandi, F.M. and Russell, J.R. (2008), "Microstructure Noise, Realized Variance, and Optimal Sampling," *Review of Economic Studies* 75: 339-369

Barndorff-Nielsen, O.E. and Shephard, N. (2002), "Econometric Analysis of Realised Volatility and Its Use in Estimating Stochastic Volatility Models," *Journal of the Royal Statistical Society Series B* 64: 853-223.

Barndorff-Nielsen, O.E., Hansen, P.R., Lunde, A., and Shephard, N. (2008), "Designing Realised Kernels to Measure the Ex-post Variation of Equity Prices in the Presence of Noise," *Econometrica* 76(6): 1481–1536

Barndorff-Nielsen, O.E., Hansen, P.R., Lunde, A., and Shephard, N. (2011), "Subsampling Realised Kernels," *Journal of Econometrics*, 160: 204-219

Breiman, L. (1996), "Bagging Predictors," *Machine Learning* 36, 105–139.

Bühlmann, P., and Yu, B. (2002), "Analyzing Bagging," *Annals of Statistics* 30, 927–961.

Buja, A. and Stuetzle, W. (2002), "Observations on Bagging", University of Pennsylvania and University of Washington, Seattle.

Chernozhukov, V. and Umantsev, L. (2001), "Conditional Value-at-Risk: Aspects of Modeling and Estimation," *Empirical Economics* 36, 221-292.

Clements, M.P., Galvao, A.B., and Kim, J.H. (2008), "Quantile Forecasts of Daily Exchange Rate Returns from Forecasts of Realized Volatility," *Journal of Empirical Finance* 15: 729-750

Engle, R.F. and Russell, J.R. (1998), "Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data," *Econometrica* 66(5): 1127-1162.

Fitzenberger, B. (1997), "The Moving Blocks Bootstrap and Robust Inference for Linear Least Squares and Quantile Regressions", *Journal of Econometrics* 82, 235-287.

Friedman, J.H. and P. Hall (2007), "On Bagging and Nonlinear Estimation", *Journal of Statistical Planning and Inference* 137, 669-683.

Ghysels, E., Sinko, A. and Valkanov, R. (2006), "MIDAS Regressions: Further Results and New Directions", *Econometric Reviews* 26: 53-90

Granger, C.W.J. and Ramanathan, R. (1984), "Improved Methods of Combining Forecasts," *Journal of Forecasting* 3: 197-204.

Hall, P., Horowitz, J.L. and Jing, B.-Y. (1995), "On Block Rules for the Bootstrap with Dependent Data," *Biometrika* 82(3): 561-574.

Hansen, P.R. and Lunde, A. (2006), "Realized Variance and Market Microstructure Noise," *Journal of Business and Economic Statistics* 24: 127-218.

Hillebrand, E. and Medeiros, M.C. (2010): "The Benefits of Bagging for Forecast Models of Realized Volatility," *Econometric Reviews* 29(5): 571-593.

Huang, H. and Lee, T.-H. (2010), "To Combine Forecasts or to Combine Information?" *Econometric Reviews* 29(5): 534-570.

Inoue, A., and Kilian, L. (2008), "How Useful is Bagging in Forecasting Economic Time Series? A Case Study of U.S. CPI Inflation," *Journal of the American Statistical Association* 103(482), 511-522.

Koenker, R. (2005), *Quantile Regression*, Cambridge University Press.

Koenker, R and G. Basset (1978), "Asymptotic Theory of Least Absolute Error Regression," *The Journal of the American Statistical Association* 73, 618-622.

Lee, T.-H., and Y. Yang (2006), "Bagging Binary and Quantile Predictors for Time Series," *Journal of Econometrics* 135, 465–497.

Machado, J.A.F. (1993), "Robust Model Selection and M-Estimaiion," *Econometric Theory* 9: 479-990.

Portnoy, S. and Koenker, R. (1997), "The Gaussian Hare And the Laplacean Tortoise: Computability of l_1 vs l_2 Regression Estimators," *Statistical Science* 12: 279-300.

Rapach, D., and Strauss, J. (2010): "Bagging or Combining (or Both)? An Analysis Based on Forecasting U.S. Employment Growth," *Econometric Reviews* 29(5): 511-533.

Stock, J.H. and Watson, M.W. (2002), "Forecasting Using Principal Components from a Large Number of Predictors," *Journal of the American Statistical Association* 97: 1167-1179.

Stock, J.H. and Watson, M.W. (2011), "Generalized Shrinkage Methods for Forecasting Using Many Predictors," Harvard University and Princeton University.

Zhang, L., Mykland, P.A. and Aït-Sahalia, Y. (2005), "A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data," *Journal of the American Statistical Association* 100: 1394-1411.

Zhou, B. (1996), "High-Frequency Data and Volatility in Foreign-Exchange Rates," *Journal of Business and Economic Statistics* 14: 45-52.

Table 1. Forecasting Quantiles Using High Frequency Information

This table presents the relative performance of each forecasting scheme for one-day ahead daily close S&P 500 return quantile prediction. We use 5-minute high frequency S&P 500 index data from June 9th, 1997 to May 30th, 2003, a total of 117,078 observations on 1,501 days. There are 78 5-minute price index observations during a day. We generate the 78 subsample "daily" return series from a time in a day to the same time in the following trading day. Each of the daily returns is with length 1,500 days. The daily returns are calculated by log difference of corresponding index values and multiplied by 100. Out-of-sample size P is 500 days, from May 29th, 2001 to May 30th, 2003, in-sample size R is 1,000 days, from June 10th, 1997 to May 25th, 2001. We use rolling window scheme to estimate the parameters and set the size of estimation window at R=1000. Here we report the out-of-sample mean forecast tick loss ratios of each chosen forecasting scheme relative to the Daily Close benchmark model. We report for different values of the left-tail probability parameter α, including α=0.5 for the median. A less-than-one number indicates that the chosen model outperforms the benchmark Daily Close, and **bolded** number indicates the smallest average tick loss ratio of a model among all methods for each α in each column. For CF-PC and CI-PC, the information criteria AIC and BIC in selecting number of factors are modified for the mean tick forecast errors instead of the mean squared errors. In all factor model approaches, the maximum hypothesized number of factors, k_{max}, is set at 15, and so the number of factors 'k' is chosen within interval [1,15]. In Bagging the number of bootstrap samples is 50 and bootstrap block size is fixed at 4.

	α=0.01	α=0.05	α=0.1	α=0.3	α=0.5	α=0.7	α=0.9	α=0.95	α=0.99
Daily-Close (mean tick loss×100)	4.5816	15.1202	24.7583	48.3697	54.5713	48.6378	26.7932	16.5165	5.0059
Bagging Daily-Close	0.9139	0.9826	0.9976	0.9935	**0.9971**	0.9980	0.9934	0.9938	0.9131
SA-Mean	**0.8803**	0.9723	**0.9877**	**0.9873**	0.9990	0.9967	**0.9882**	0.9764	0.9513
SA-Median	0.8903	0.9731	0.9916	0.9885	0.9994	0.9954	0.9905	0.9818	0.9699
CF-Mean	0.9226	0.9718	0.9928	0.9938	0.9992	0.9943	0.9912	**0.9750**	0.9403
CF-Median	0.9339	0.9746	0.9952	0.9939	0.9997	0.9952	0.9940	0.9805	0.9445
CF-PC (AIC)	1.1006	1.0221	1.0273	1.0025	1.0097	1.0047	1.0224	0.9772	2.2536
CF-PC (BIC)	1.0746	1.0103	0.9982	0.9948	0.9985	0.9944	1.0064	0.9938	2.0761
CF-PC (k=1)	0.9253	0.9726	0.9892	0.9962	0.9985	0.9982	0.9993	0.9832	0.9562
CF-PC (k=2)	0.9356	0.9792	0.9894	0.9973	0.9985	0.9983	0.9916	1.0034	0.9896
CF-PC (k=3)	0.9853	0.9823	1.0012	0.9983	0.9992	1.0008	0.9884	0.9985	0.9607
CI-Unrestricted	5.8885	2.0387	1.5619	1.1643	1.1971	1.2420	1.4362	1.8454	4.8916
CI-PC (AIC)	1.0956	1.0340	1.0337	0.9914	1.0006	0.9980	0.9888	1.0278	0.9905
CI-PC (BIC)	1.0456	1.0291	1.0284	0.9951	1.0004	**0.9871**	0.9891	0.9945	0.9967
CI-PC (k=1)	0.9471	**0.9675**	1.0121	0.9983	1.0004	1.0018	1.0030	0.9952	0.9167
CI-PC (k=2)	0.9589	0.9817	1.0132	0.9977	1.0033	1.0040	1.0095	0.9966	**0.8947**
CI-PC (k=3)	0.9974	0.9752	1.0014	0.9984	1.0025	0.9955	1.0018	0.9879	1.0224