QUANTITATIVE METHODS IN ECONOMICS
Vol. XV, No. 1, 2014, pp. 7 – 17

BOOSTING UNDER QUANTILE REGRESSION – CAN WE USE IT FOR MARKET RISK EVALUATION?

Katarzyna Bień-Barkowska
Institute of Econometrics, Warsaw School of Economics
e-mail: katarzyna.bien@sgh.waw.pl

Abstract: We consider boosting, i.e. one of popular statistical machine-learning meta-algorithms, as a possible tool for combining individual volatility estimates under a quantile regression (QR) framework. Short empirical exercise is carried out for the S&P500 daily return series in the period of 2004-2009. Our initial findings show that this novel approach is very promising and the in-sample goodness-of-fit of the QR model is very good. However much further research should be conducted as far as the out-of-sample quality of conditional quantile predictions is concerned.

Keywords: boosting, quantile regression, GARCH models, value-at-risk

INTRODUCTION

Boosting refers to an iterative statistical machine learning meta-algorithm which aims to enhance the predictive accuracy of different weak classification algorithms (weak learners), i.e. classifiers that evidence a substantial error rate. In brief, the method is recognized as very complex and efficient when making a new prediction rule by combining different and often inaccurate individual classification rules. Different examples of specific boosting algorithms have been proposed in the literature so far, and perhaps the most renown one is the Adaptive Boosting algorithm (i.e. AdaBoost) (see [Freund and Schapire 1997]). In short, the AdaBoost algorithm iteratively evokes a new weak classification rule which assigns more weights to these data points that eluded correct classification by former classifiers. In this manner the algorithm keeps reinforcing the focus of additional weak learners on inappropriately classified data, thus improving the final accuracy of prediction. Final classification is obtained by appropriate weighting votes of single classifiers. A thorough discussion of the boosting mechanism from the statistical perspective has been presented by [Friedman et al. 2000] or [Bühlmann and Hothorn 2007].

From an econometric viewpoint, boosting might be used as an optimization algorithm for choosing the best combination of explanatory variables (predictors) with respect to an economic question at hand. To this end, based upon the nature of the economic phenomenon under study as well as specific statistical features of dependent variable to be considered, many different cost functions can be easily accommodated in the boosting algorithm. These might be, for example, negative binominal log-likelihood for a binary classification problem, L1-norm loss function for median regression, L2-norm loss function for standard (mean) regression or a check function for quantile regression (see [Bühlmann and Yu 2003]; [Bühlmann 2006]; [Fenske at al. 2011]). Boosting methods have also been applied to density estimation by [Ridgeway 2002] or [Di Marzio and Taylor 2005] or to survival analysis by [Hothorn et al. 2006], [Lu and Li 2008] or [Chen et al. 2013]. In short, once the loss criterion is set, boosting algorithm performs sequential updates of an (parameter) estimator according to the steepest gradient descent of the loss function evaluated at the empirical data. At each iteration step, separate regression models (weak learners) are used to explain the negative of gradient of the evaluated cost function with the penalized ordinary least squares method (see [Fenske at al. 2011]).

The aim of this analysis is to provide a short pilot empirical study on possible application of boosting algorithm when combining different volatility estimates under a quantile regression (QR) framework (see [Koenker 2005]). We are inspired by the recent contribution of [Fenske at al. 2011], where the functional gradient boosting algorithm for quantile regression has been thoroughly discussed. For an empirical analysis we applied the software package (application 'mboost') developed under the R environment by [Hothorn et al. 2010] and [Hothorn et al. 2013] (see also [R Development Core Team 2008]). In this pilot study we intend to consider a boosting-based model for a conditional quantile of return distribution. The quantile regression model might be simply treated as a (percentage) value-at-risk model where the optimal combination of linear predictors has been selected (and accordingly weighted) from the set of volatility estimates based upon different specifications of GARCH models. In such a setup, individual parametric GARCH-based conditional quantile predictions might be even severely biased, whereas the boosting algorithm is awaited to combine them in an optimal way, hence enforcing the quality of emerging value-at-risk measures.

TEORETICAL FOUNDATIONS

The concept of value-at-risk is fundamentally related to the notion of a quantile function. If r_t denotes a return on portfolio between times $t-1$ and t, the corresponding (percentage) $VaR_{t,\alpha}$ would be equal to $q_\alpha(r_t)$ i.e. the conditional α-quantile of return distribution:

$$\Pr(r_t < VaR_{t,\alpha} \mid F_{t-1}) = q_\alpha(r_t), \tag{1}$$

where F_{t-1} denotes an information set available at $t-1$. In financial risk management, VaR constitutes a popular risk measure. From equation (1) it becomes clear, that VaR is a threshold value for (percentage) loss. Thus, the probability that marked-to-market return on portfolio value (over given time horizon) will be lower than VaR will be equal to the chosen probability level α. There are plenty of value-at-risk models proposed in the literature (see [Jorion 2000]). The most popular VaR models are based upon: the RiskMetrics approach, parametric GARCH models, semiparametric methods which combine parametric GARCH models with nonparametric distribution estimates (i.e. filtered historical simulation) or CAViaR models that directly depict conditional quantiles as observation-driven autoregressive processes (see [Engle and Manganelli 2004]).

There is a strong trend in the recent literature to improve the prediction accuracy of different forecasts by combining them. For a standard regression problem, simple averages or weighted averages of individual forecasts (i.e. averages weighted by inverses of prediction errors) are usually used. For example, [Aiolfi et al. 2010] show that the equally-weighted average of survey forecasts and forecasts from various time-series models leads to smaller out-of-sample prediction errors. Quite recently, combining the individual volatility forecasts (see [Amendola and Storti 2008], [Jing-Rong et al. 2011]) or even density forecasts attracted much attention. For example, [Hall and Mitchell 2007] set the weights of individual density forecasts as the weights that minimize the ‘distance’ (measured by the Kullback–Leibler information criterion) between the forecasted and the true (unknown) density. The most modern approach is to combine forecasts under the quantile regression framework. [Chiriac and Pohlmeier 2012] propose new methods for combining individual value-at-risk forecasts directly. They show how to mix information from different VaR specifications in an optimal way using a method-of-moments estimator. Alternatively, they combine individual VaR forecasts under the QR framework. [Jeon and Tylor 2013] enrich the CAViaR model of [Engle and Manganelli 2004] with the implied volatility measure that reflects the market’s expectation of risk and carries new information in comparison to historical volatility estimates.

In this pilot study we consider seven different volatility estimates $\hat{\sigma}_t$. Each of these is derived from a different GARCH specification:

1. Standard 'plain vanilla' GARCH(1,1) model of [Bollerslev 1986]:

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \tag{2}$$

where ε_t^2 denotes the residuals from the mean filtration process. (For the sake of parsimony, ARMA(1,1) model has been used in the conditional mean equation.)

2. Integrated GARCH(1,1) model of [Engle and Bollerslev 1986]:

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + (1-\alpha)\sigma_{t-1}^2 \tag{3}$$

3. Exponential GARCH(1,1) model of Nelson (1991):

$$\ln(\sigma_t^2) = \omega + \alpha\frac{\varepsilon_{t-1}^2}{\sigma_{t-1}^2} + \gamma\left(\left|\frac{\varepsilon_{t-1}}{\sigma_{t-1}}\right| - \sqrt{\frac{2}{\pi}}\right) + \beta\ln(\sigma_{t-1}^2) \tag{4}$$

4. GJR GARCH model of [Glosten et al. 1993]:

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \gamma I_{t-1}\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \tag{5}$$

where I_{t-1} denotes the indicator function ($I_{t-1} = 1$ if $\varepsilon_{t-1} \le 0$ and $I_{t-1} = 0$ otherwise).

5. The asymmetric power ARCH(1,1) (APARCH) model of [Ding et al. 1993]:

$$\sigma_t^\delta = \omega + \alpha\left(\left|\varepsilon_{t-1}\right| - \gamma\varepsilon_{t-1}\right)^\delta + \beta\sigma_{t-1}^\delta \tag{6}$$

where $\delta > 0$ denotes a parameter of the Box-Cox transformation of σ_t^2.

6. The absolute value GARCH (AVGARCH) model of [Taylor 1986] and [Schwert 1990]:

$$\sigma_t = \omega + \alpha\left|\varepsilon_{t-1}\right| + \beta\sigma_{t-1} \tag{7}$$

7. The Nonlinear Asymmetric GARCH model of [Engle and Ng 1993]:

$$\sigma_t^2 = \omega + \alpha\sigma_{t-1}^2\left(\left|\frac{\varepsilon_{t-1}}{\sigma_{t-1}} - \eta_2\right| - \eta_1\left(\frac{\varepsilon_{t-1}}{\sigma_{t-1}} - \eta_2\right)\right)\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \tag{8}$$

where η_1 denotes a "rotation" parameter and η_2 denotes a "shifts" parameter, respectively.

All the 'sigma' estimates obtained from the aforementioned models will be treated as explanatory variables in a boosting-based quantile regression analysis. Accordingly, we aim to search for the optimal weighting algorithm of these volatility estimates under the QR framework.

Under the QR setup the conditional quantile of return distribution is given as:

$$q_{\alpha}(r_t \mid \mathbf{x}_t) = \eta_{\alpha,t} = \mathbf{x}_t' \boldsymbol{\beta}_{\alpha} \tag{9}$$

where $\mathbf{x}_t$ denotes an $[L\text{x}1]$ vector of VaR predictors at time t (individual explanatory variables, including the obtained sigma estimates) and $\boldsymbol{\beta}_{\alpha}$ denotes a corresponding $[L\text{x}1]$ parameter vector. The parameter vector $\boldsymbol{\beta}_{\alpha}$ can be estimated by finding a minimum of the following QR optimization problem:

$$\arg\min \sum_{t=1}^{T} \rho_{\alpha}\left(r_t - \mathbf{x}_t' \boldsymbol{\beta}_{\alpha}\right) \text{ where } \rho_{\alpha}(u) = \begin{cases} u\alpha & u \geq 0 \\ u(\alpha - 1) & u < 0 \end{cases}.$$

The functional gradient boosting algorithm looks for the minimum of the empirical risk function: $\frac{1}{T}\sum_{t=1}^{T} L_t$, where L_t denotes its t-th contribution, which, in the case of a quantile regression problem, is given as: $L_t = \rho_{\alpha}(r_t - \eta_{\alpha,t})$ (where $\eta_{\alpha,t}$ denotes a theoretical value of a conditional α- quantile or, in other words, it is a linear combination of individual predictors of a given conditional α-quantile).

In the following, we present the outline of boosting strategy after [Fenske at al. 2011] with slight modifications (and changes in notation) in order to adjust the algorithm to the setup of our study.

1. Choose an appropriate starting value for parameter vector $\boldsymbol{\beta}_{\alpha} = \boldsymbol{\beta}_{\alpha}^{0}$. Define a maximum number of boosting iterations m_{stop} and set the iteration index at $m = 0$.

2. Compute $[T\text{x}1]$ vector of negative gradients of the empirical risk function (evaluated at each t):

$$u_{\alpha,t} = -\left.\frac{\partial L_t}{\partial \eta_{\alpha,t}}\right|_{\eta_{\alpha,t} = \hat{\eta}_{\alpha,t}^{[m-1]}}, \qquad t = 1,2,...,T \tag{10}$$

In the case of quantile regression, the first derivative of L_t with respect to $\eta_{\alpha,t}$ is:

$$u_{\alpha,t} = \rho'_\tau\left(r_t - \hat{\eta}_{\alpha,t}^{[m-1]}\right) = \begin{cases} \tau & if \quad r_t - \hat{\eta}_{\alpha,t}^{[m-1]} > 0 \\ 0 & if \quad r_t - \hat{\eta}_{\alpha,t}^{[m-1]} = 0 \\ \tau - 1 & if \quad r_t - \hat{\eta}_{\alpha,t}^{[m-1]} < 0 \end{cases} \tag{11}$$

3. With the OLS, fit possible explanatory variables to the obtained negative gradients in order to obtain the *m*-step estimates: $\hat{b}_{\alpha,l}^{[m]}$ (for $l = 1,2,...,\mathrm{L}$). These regressions are the base learners assigned to individual parameters $\beta_{\alpha,l}$. Estimation of $\hat{b}_{\alpha,l}^{[m]}$ is done by minimizing the standard L_2 loss: $\min(\mathbf{u}_\alpha - \hat{\mathbf{u}}_\alpha)'(\mathbf{u}_\alpha - \hat{\mathbf{u}}_\alpha)$ where $\hat{\mathbf{u}}_\alpha = \mathbf{x}_l b_{\alpha,l}$ (optimization is performed for each x_l variable separately).

4. If the best-fitting variable has an indicator $l*$ $(1 \le l* \le \mathrm{L})$, then the coefficient that corresponds to this variable is updated accordingly as: $\hat{\beta}_{\alpha,l*}^{[m]} = \hat{\beta}_{\alpha,l*}^{[m-1]} + \nu\,\hat{b}_{\alpha,l*}^{[m]}$ where $\nu \in (0,1]$ is a given step size, i.e. shrinkage parameter. All other parameters are kept constant: $\hat{\beta}_{\alpha,l}^{[m]} = \hat{\beta}_{\alpha,l}^{[m-1]}, \quad l \neq l*$

5. Increase m by one until $m = m_{stop}$ or go back to [2].

Functional gradient boosting has a very intuitive interpretation. In step [3] of the algorithm, L separate linear regression models are estimated, but only the best one (according to mean square criterion) is selected to update the m-step parameter vector $\boldsymbol{\beta}_\alpha^{[m]}$. Accordingly, at each iteration, boosting algorithm chooses only one variable that explains in the best way the negative gradient of the empirical loss function. In step [4] the parameter corresponding to this variable is changed proportionally to the value of achieved $\hat{b}_{\alpha,l}^{[m]}$, whereas some shrinkage should be made according to the chosen size of ν.

EMPIRICAL EXERCISE

Time series of daily log returns on S&P500 close prices between January 2004 and December 2009 has been selected as the dataset for the exercise. The huge heterogeneity of the time span under study allows to cover a 'calm' period of 2004-2006 and the very turbulent period of a recent financial turmoil of 2007-

2009. In Table 1 we present some standard backtesting measures of individual GARCH-based quantile estimates of return distribution. These are the results of popular unconditional coverage test of [Kupiec 1995] (i.e. test statistics LL_{UC}) and of conditional coverage test of [Christoffersen 1998] (i.e. test statistics LL_{CC}). Large values of the obtained test statistics evidence that the in-sample fit of all GARCH-based conditional quantile forecast is very poor. This can be well understood if we take into account the structural break of July 2007 (first signals of the upcoming turmoil) or the crash of September 2008 (the fall of Lehman Brothers) that should have been taken into consideration while constructing GARCH models. Moreover, all GARCH specifications have been estimated with the assumption of Gaussian distribution for the error terms, which significantly underestimates the true thickness of the lower distribution tail.

Table 1. Quality of (in-sample) VaR estimates under different GARCH specifications. LL_{UC} denotes the unconditional coverage statistics and LL_{CC} denotes the conditional coverage statistics. Bolded values denote statistically significant (at 5%) outcomes.

model	LL_{UC} $VaR_{0,05}$	LL_{CC} $VaR_{0,05}$	LL_{UC} $VaR_{0,01}$	LL_{CC} $VaR_{0,01}$
GARCH	**4.972**	**713.883**	**22.902**	**370.251**
IGARCH	1.237	**655.023**	**11.579**	**305.938**
EGARCH	2.809	**684.453**	**8.927**	**287.563**
GJR GARCH	3.841	**696.225**	**8.926**	**287.563**
APARCH	2.108	**672.681**	**8.927**	**287.563**
AVGARCH	2.447	**678.567**	**7.708**	**278.376**
NAGARCH	1.237	**655.023**	**11.579**	**305.938**

Source: own calculations.

Volatility estimates resulting from the seven different, but in fact incorrect GARCH specifications have been used as potential predictors in a boosting mechanism together with a one-day lagged “High-Low” price range measure for S&P500. As suggested by [Bühlmann, Hothorn 2007] we centered all individual predictors (by subtracting their mean value) before initializing boosting algorithm. The initial value for the intercept in the QR model has been selected as the unconditional 0.05-quantile or the unconditional 0.01-quantile, respectively. The shrinkage parameter has been set as $\nu = 0.05$, thus we allow for more shrinkage than [Fenske et al 2011], in order to account for a considerable multicollinearity between predictors. The optimal number of boosting iterations (*m*) has been selected with the application of a standard 25-fold bootstrap procedure in order to avoid overfitting of the learning mechanism.

Figure 1. Value of the empirical loss function for increasing number of boosting iterations. Results from 25 individual subsamples corresponding to the 25-fold bootstrap procedure (grey lines) and their average (black line) with respect to the 0.05-quantile (left panel) and the 0.01-quantile (right panel).

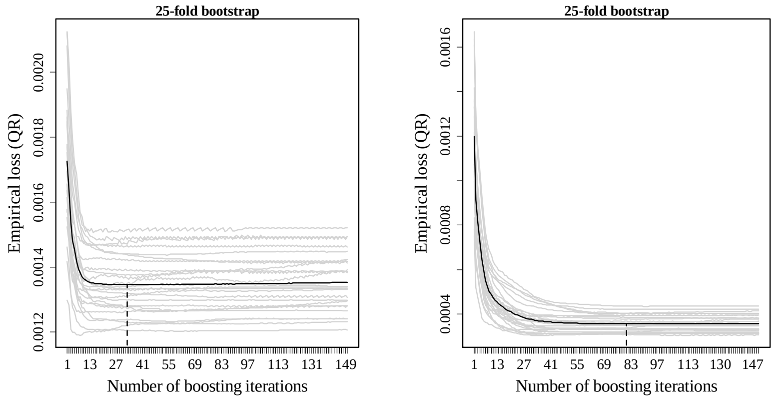

Source: own calculations with the application of the 'mboost' library.

The results of the boosting-based model for the 0.05-quantile are the following. Out of eight potential individual predictors, five were selected by the algorithm:

- standard GARCH-based volatility ($\hat{\beta}_{1,0.05} = -0.499$),
- IGARCH-based volatility ($\hat{\beta}_{2,0.05} = 0.160$),
- EGARCH-based volatility ($\hat{\beta}_{3,0.05} = -0.573$),
- GJR GARCH-based volatility ($\hat{\beta}_{4,0.05} = -0.911$), and the
- H-L price range ($\hat{\beta}_{8,0.05} = 0.060$).

In the case of 0.01-quantile regression, once again

- standard GARCH-type volatility is selected $\hat{\beta}_{1,0.01} = -0.888$, and then
- IGARCH-based volatility ($\hat{\beta}_{2,0.01} = 0.095$),
- EGARCH-type volatility ($\hat{\beta}_{3,0.01} = -0.774$),
- GJR GARCH-type volatility ($\hat{\beta}_{4,0.01} = -0.650$) and
- H-L price range ($\hat{\beta}_{8,0.01} = -0.047$).

Therefore, we can formulate the following conclusions. First, weights of volatility estimates selected by a boosting algorithm differ for 0.05-quantile and the 0.01-quantile, although the types of selected models stay the same. Second, majority of selected sigma-predictors have, as expected, negative impact for the conditional quantile. Third, leverage or non-linearity effects play an important role as suggested by a large parameter values for the asymmetric GARCH specifications, both for the 0.05-quantile and 0.01-quantile.

Figure 3. S&P returns between January 2004 and December 2009 (grey line) and the corresponding boosted Value at Risk at α=0.05 and α=0.01 (black lines).

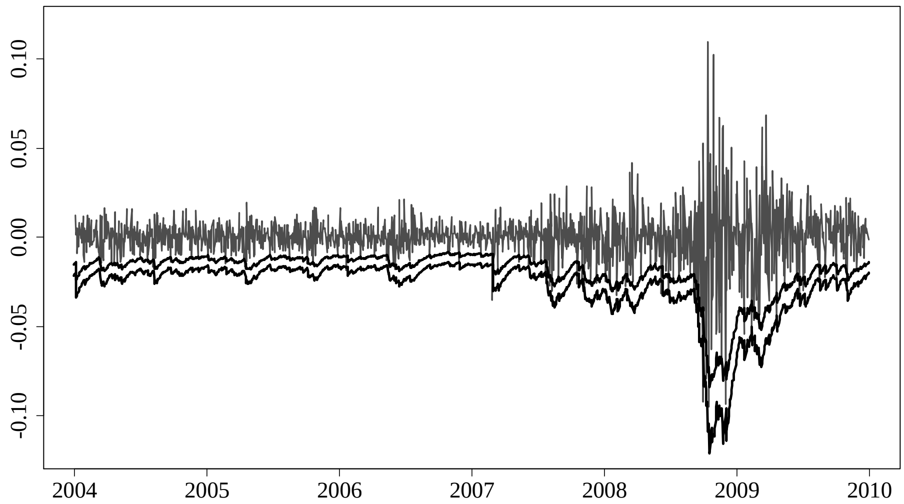

Source: own calculations with the application of the ‘mboost’ application.

In Figure 2 we present the return series under study together with the obtained boosting-based time-varying VaR estimates. We can observe that the estimated conditional quantiles seem to suitably react to the down- and upswings in the return series and to capture volatility clustering effects in a satisfactory manner. Moreover, boosting mechanism allows for a very good fit of the (in sample) QR model. According to the results of both, the unconditional coverage and the conditional coverage tests, the observed fraction of VaR exceedances does not significantly differ from the probability level set in the model. We are also not able to reject the null, that the exceedances are independent in time (LL_{UC} is equal to 0.002 and LL_{CC} is equal to 0.17 for $VaR_{0.05}$ and LL_{UC} is equal to 0.055 and LL_{CC} is equal to 0.399 for $VaR_{0.01}$).

This new approach seems to set forth a promising research direction in VaR modelling. Its merits lie in a properly chosen loss function, which, contrary to majority of GARCH specifications, does not impose any parametric assumptions on the distribution of financial returns. As opposed to GARCH models, it estimates the conditional quantile directly and in semiparametric fashion, which stays in line with the CAViaR approach. The dynamics of the model can be easily driven by different forms of volatility estimates or other variables as lagged transaction volumes or implied volatility estimates. The drawback of this approach is its sensitivity to selection of a shrinkage factor or maximum numbers of performed boosting iterations. The approach can be also ‘fragile’ to possible structural breaks in the series, which may pose a further need for a time-varying weights. Moreover, much more effort should be put on a proper evaluation of the out-of-sample properties of the model, which is of utmost importance as far as the model application in the risk management is concerned.

REFERENCES:

Aiolfi M., Capistran C., Timmermann A. (2010) Forecast Combinations, forthcoming in Forecast Handbook (Oxford) ed. Clements, M. and Hendry, D.

Amendola A., Storti G. (2008) A GMM procedure for combining volatility forecasts, Computational Statistics and Data Analysis, 52(6), 3047–3060.

Bollerslev T. (1986) Generalized autoregressive conditional heteroskedasticity, Journal of Econometrics, 31, 307-327.

Bühlmann P. (2006) Boosting for high-dimensional linear models, Annals of Statistics 34, 559–583.

Bühlmann P., Hothorn T. (2007) Boosting Algorithms: Regularization, Prediction and Model Fitting (with Discussion), Statistical Science, 22(4), 477–505.

Bühlmann P., Yu B. (2003) Boosting with the L2 Loss: Regression and Classification, Journal of the American Statistical Association, 98, 324–339.

Chen Y., Zhenyu J., Mercola D., Xiaohui X. (2013) A Gradient Boosting Algorithm for Survival Analysis via Direct Optimization of Concordance Index, Computational and Mathematical Methods in Medicine 2013, 1-8.

Chiriac R., Rohlmeier W. (2012) Improving the Value at Risk Forecasts: Theory and Evidence from the Financial Crisis, Journal of Economic Dynamics and Control 36, 1212–1228.

Christoffersen P. F. (1998) Evaluating Interval Forecasts, International Economic Review 39, 841–862.

Engle R.F., Bollerslev T. (1986) Modelling the persistence of conditional variances, Econometric Reviews, 5(1), 1-50.

Engle R.F., Manganelli S. (2004) CAViaR: Conditional Autoregressive Value at Risk by Regression Quantiles, Journal of Business & Economic Statistics 22(4), p. 367-381.

Engle R.F., Ng V.K. (1993) Measuring and testing the impact of news on volatility, Journal of Finance, 48(5) 1749 - 1778.

Fenske N., Kneib T., Hothorn T. (2011) Identifying Risk Factors for Severe Childhood Malnutrition by Boosting Additive Quantile Regression, Journal of the American Statistical Association, 106, 494-510.

Freund Y., Schapire R.E. (1997) A decision-theoretic generalization of on-line learning and an application to boosting, Journal of Computer and System Sciences 55, p. 119-139.

Friedman J., Hastie T., Tibshirani R. (2000) Additive Logistic Regression: A statistical view of boosting, The Annals of Statistics 28, 337-407.

Glosten L.R., Jagannathan R., Runkle D.E. (1993) On the relation between the expected value and the volatility of the nominal excess return on stocks, Journal of Finance, 48(5), 1779-1801.

Hall S.G., Mitchell J. (2007) Combining density forecasts, International Journal of Forecasting 23, 1-13.

Hothorn T., Bühlmann P., Dudoit S., Molinaro A. and Van Der Laan M. (2006) Survival ensembles, Biostatistics 7, 355–373.

Hothorn T., Bühlmann P., Kneib T., Schmid M., and Hofner B. (2010) Model-Based Boosting 2.0, Journal of Machine Learning Research, 11, p. 1851–1855.

Hothorn T., Bühlmann P., Kneib T., Schmid M., Hofner B., Sobotka F., Scheipl F. (2013) Mboost: Model-Based Boosting, R package version 2.2-3, available at http://CRAN.R-project.org/package=mboost.

Jeon J., Taylor J.W. (2013) Using CAViaR Models with Implied Volatility for Value at Risk Estimation, Journal of Forecasting 32, p. 62-74.

Jing-Rong D., Yu-Ke C., Yan Z. (2011) Combining Stock Market Volatility Forecasts with analysis of Stock Materials under Regime Switching, Advances in Computer Science, Intelligent System and Environment, Springer, 393-402.

Jorion P. (2000) Value at Risk. The new benchmark for managing financial risk, McGraw-Hill.

Koenker R. (2005) Quantile Regression. Economic Society Monographs, Cambridge University Press, New York.

Kupiec P. H. (1995) Techniques for verifying the accuracy of risk measurement models, The Journal of Derivatives, 3(2), 73-84.

Li Y., and Zhu J. (2008) L1-Norm Quantile Regression, Journal of Computational and Graphical Statistics 17, 163–185.

Lu W., Li L. (2008) Boosting method for nonlinear transformation models with censored survival data, Biostatistics 9, 658-667.

Menardi G., Tedeschi F., Torelli N. (2011) On the Use of Boosting Procedures to Predict the Risk of Default, [in:] Classification and Multivariate Analysis for Complex Data Structures, Springer, Berlin, 211-218.

R Development Core Team (2008) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, URL http://www.R-project.org.

Ridgeway G. (2002) Looking for lumps: Boosting and bagging for density estimation, Computational Statistics and Data Analysis 38, 379–392.

Schwert G. W. (1990) Stock volatility and the crash of '87, Review of Financial Studies, 3(1), 77-102.

Taylor S. J. (1986) Modelling financial time series, Wiley.