



869 | 25

Views | CrossRef citations to date | Altmetric

Original Articles

Robustify Financial Time Series Forecasting with Bagging

Sainan Jin, Liangjun Su & Aman Ullah

Pages 575-605 | Accepted author version posted online: 09 Aug 2013, Published online: 27 Nov 2013

Cite this article <https://doi.org/10.1080/07474938.2013.825142>



Sample our
Economics, Finance,
Business & Industry Journals
>> **Sign in here** to start your access
to the latest two volumes for 14 days

Full Article

Figures & data

References

Citations

Metrics

Reprints & Permissions

Read this article

Share

Abstract

In this paper we propose a revised version of (bagging) bootstrap aggregating as a forecast combination method for the out-of-sample forecasts in time series models. The revised version explicitly takes into account the dependence in time series data and can be used to justify the validity of bagging in the reduction of mean squared forecast error when compared with the unbagged forecasts. Monte Carlo simulations show that the new method works quite well and outperforms the traditional one-step-ahead linear forecast as well as the nonparametric forecast in general, especially when the in-sample estimation period is small. We also find that the bagging forecasts based on misspecified linear models may work as effectively as those based on nonparametric models, suggesting the robustification property of bagging method in terms of out-of-sample forecasts. We then reexamine forecasting powers of predictive variables suggested in the literature to forecast the excess returns or equity premium. We find

that, consistent with Goyal and Welch (2008), the historical average excess stock return forecasts may beat other predictor variables in the literature when we apply traditional one-step linear forecast and the nonparametric forecasting methods. However, when using the bagging method or its revised version, which help to improve the mean squared forecast error for “unstable” predictors, the predictive variables have a better forecasting power than the historical mean.

Keywords:

- Bagging
- Combined forecasts
- Nonparametric models
- Predictability
- Time series

JEL Classification:

- C14
- C53

ACKNOWLEDGMENTS

The authors are grateful to a referee for useful comments and suggestions. They are also thankful to Tae-Hwy Lee and Yundong Tu for discussions on the subject matter of this paper. The first author gratefully acknowledges financial support from the SMU research grant (Grant number: C244/MSS10E006).

Notes

In practice, one could draw the multiple training sets ($b = 1, \dots, B$) from P and employ a weighted version of the ensemble aggregating predictor, where $w_{b,t}$ is the weight function with $\sum_{b=1}^B w_{b,t} = 1$.

We also consider the cases of negative autocorrelation. We set $\rho = -0.5$ and -0.95 . The results seem to be similar to the results when $\rho = 0.5$ and 0.95 , respectively. To save space for the tables, the results are not reported here.

We try different values of R by setting $R = 50, 100$, and 200 . We also try different out-of-sample periods by setting P to be $100, 200$, and 500 . The results are similar and not reported here for brevity.

As suggested by one referee, The choice of equal weights is optimal in the iid case but not necessarily with dependent series like the ones considered in the paper. We also try to compute the weights by using Bayesian model averaging (BMA) technique as introduced in Lee and Yang ([2006](#)). The BMA gives a large weight to the b th bootstrap predictor at each period t when it has forecasted well over the past k periods and a small weight to the predictor at period t when it forecasted poorly over the past k periods. We have set $k = 1, 5$ and R . The results are similar to those based on equal weights and not reported here for brevity.

As one referee remarks, the choice of h affects the results and may not be optimal as chosen. We are dealing with dependent series and cross validation methods require blocking here too. We follow Hart and Vieu ([1990](#)) and set different leave-out sequences to take care of the dependence structure of the time series. We set $l_n = 0, 1, 2, 3, 4, 5$ as in Hart and Vieu ([1990](#)), where $l_n = 0$ corresponds to the ordinary leave-one-out cross validation, $l_n > 0$ corresponds to leave $2l_n + 1$ observations out, and the leave-out sequence is $\{X_j\}$ with $|j - t| \leq l_n$. The results for $l_n > 0$ are similar to those based on the usual leave-one-out least squares cross validation and thus not reported here.

We also try to choose the bandwidth by the “rule of thumb”: $h_l = c_0 s_l n^{-1/(4+q)}$, where s_l stands for the sample standard deviations of $X_{it,l}$, the l th regressor in X_{it} . We set $c_0 = 0.5, 1$, and 2 to examine the sensitivity of our test to the choice of bandwidth. It turns out that the results of our proposed methods and bagging methods are robust to different bandwidth choice. On the other hand, the usual one-step-ahead local constant and local linear predictors are sensitive to the bandwidth choice.

Note: In Tables 1–6, the results are based on the out-of-sample forecast MSE averaged over 200 repetitions. The in-sample period is $R = 20, 50, 100$, and 200 , respectively; the out-of-sample period is $P = 50$. For the bagging methods, the number of bootstrap resamples is $B = 100$.

Note: Sample begins: 1950 M1, Forecast begin: 1954 M1, Forecast end: 2005 M12. The benchmark predictor: historical mean.

Note: Sample begin: 1950 M1, Forecast begin: 1957 M1, Forecast end: 2005 M12. The benchmark predictor: historical mean.

Note: Sample begin: 1950 M1, Forecast begin: 1962 M1, Forecast end: 2005 M12. The benchmark predictor: historical mean.

As we discussed in the previous section, the one-step-ahead local constant and local linear predictors are sensitive to the bandwidth choice. If we try to choose the bandwidth by the “rule of thumb”: $h_1 = c_0 s_1 n^{-1/(4+q)}$, and set different values of $c_0 = 0.5, 1$, and 2 , the results are quite different.

The calculation is based on Eq. (13) in Campbell and Thompson ([2008](#)). For illustration, the paper considered an investor with a single-period horizon and mean-variance preference and calculated the expected excess return when the investor observes the predicting variable and the expected excess return when the investor does not observe the predicting variable. The difference between these two expected excess returns is , where γ is the coefficient of relative risk aversion, and S is the unconditional Sharpe ratio of the risky asset.

Related research

People also read

Recommended articles

Cited by
25

Information for

Authors
R&D professionals
Editors
Librarians
Societies

Opportunities

Reprints and e-prints
Advertising solutions
Accelerated publication
Corporate access solutions

Open access

Overview
Open journals
Open Select
Dove Medical Press
F1000Research

Help and information

Help and contact
Newsroom
All journals
Books

Keep up to date

Register to receive personalised research and resources by email

 Sign me up

