Econometric Reviews ›
Volume 29, 2010 - Issue 5-6: The Link Between Statistical Learning Theory and Econometrics: Applications in Econometrics, Finance, and Marketing

Original Articles

Bagging or Combining (or Both)? An Analysis Based on Forecasting U.S. Employment Growth

David E. Rapach & Jack K. Strauss

Pages 511-533 | Published online: 15 Sep 2010

Cite this article https://doi.org/10.1080/07474938.2010.481550

Abstract

Forecast ... h many
potentia ... e two
approac ... ent. The
first app ... ific
procedu ... ential
pred ... forecasts
from ... h individual
ARDL m ... ation
forecast ... g a mean
square f ... find that
bagging ... that

incorporating information from both bagging and combination forecasts based on principal components often leads to further gains in forecast accuracy.

Keywords: Bagging Combination forecasts Employment Forecast encompassing Principal components

JEL Classification: C22 C52 C53 E24

ACKNOWLEDGMENTS

The authors thank two anonymous referees for very helpful comments. The authors also gratefully acknowledge financial support from the Simon Center for Regional Forecasting at Saint Louis University. The results reported in this article were generated using GAUSS 6.1. The GAUSS program files are available at http://pages.slu.edu/faculty/rapachde/ Research.htm.

Notes

Lee and Yang (2006) use bagging techniques to develop binary and quantile forecasts of financial variables.

See Timmermann (2006) for a comprehensive review of forecast combining methods.

For example, employment growth—in the context of the so-called “jobless” recovery
from the ... 004
presiden ... ployment
growth i ... in Federal
Reserve ... ment growth
are impo

Note ... flation to
combina ... not consider
methods ... clusters, or
principa ... he best with
respect ... 8).

The t-statistics for the OLS estimates of δ_i in (1) are computed using Newey and West (1987) heteroskedasticity and autocorrelation consistent (HAC) standard errors based on a lag truncation of $h - 1$.

Inoue and Kilian (2008) consider a range of critical values. We obtain similar results using other conventional critical values such as 1.96.

Following Inoue and Kilian (2008), we use $m = h$ and $B = 100$.

“Recursive” indicates that the forecasts are generated using an expanding estimation window. The out-of-sample forecasts are “simulated”—as opposed to “real-time”—because they are based on revised data and not on the data actually available at the time of forecast formation. Real-time forecasting exercises are not feasible in the present article, as data vintages are not readily available for all of the variables we consider in our forecasting exercise. We follow much of the macroeconomic forecasting literature, including Stock and Watson (1999, 2003, 2004), in analyzing simulated out-of-sample forecasts.

Our results are not very sensitive to the maximum lag lengths.

Observe that the number of clusters serves to define the size of the first cluster, as none of the other clusters are used in generating the forecast. The greater the number of clusters, the smaller the size of the first cluster.

Using the taxonomy in Huang and Lee (2007), all of the combining methods we consider are classified as “combination of forecasts,” while the bagging forecasts—which ar —are classifie ion procedu x, which involves nselves (instead asts). We exper ng method forec to consider the appr f Bai and Ng (2006) t e as proxies for estim ng model.

The noti d Newbold (1973) and Chong and Hendry (1986). See Clements and Hendry (1998) for a textbook treatment of forecast encompassing.

A word of caution is in order with respect to the use of the MHLN $_h$ statistic in making inferences on the relative information content across forecasting models. Recent research demonstrates that a number of issues—such as the size of the in-sample period relative to the out-of-sample period, type of estimation window (for example, fixed, rolling, or recursive), and whether the models are nested or non-nested—can affect the asymptotic distribution of the test statistics; see Corradi and Swanson (2006) for an informative review of these issues. We recognize that, strictly speaking, all of the conditions required for the validity of the asymptotic distribution may not be met in our applications, so that inferences based on the MHLN $_h$ statistic serve as a rough guide to statistical significance.

Vendor performance is an index that measures how quickly companies receive deliveries from their suppliers. An increase in the index means that it is taking longer for companies to receive deliveries.

Notes: The entries for the AR benchmark model report the MSFE; the other entries report the MSFE for the ARDL model forecasts indicated on the left to the MSFE for the AR benchmark model forecasts.

Notes to Tables 2–5: The MSFE ratio reports the ratio of the MSFE for the BA model or combination forecasts indicated on the left to the MSFE for the AR benchmark model forecasts. H_0: BA encompasses CB (H_0: CB encompasses BA) corresponds a test of the null hypothesis that the BA model (combination) forecasts encompass the combination (BA mod ... is that the BA mode ... nodel) forecast ... el) forecast in the op ... statistic correspo ... indicate signific...

Never ... unity, indicatin ... oting that the MSF ... ne individual ARDL m ... eriod at the 3- and 6...

Note: The entries in the BA model and PC combination rows report the ratio of the MSFE for these forecasts to the MSFE for the AR benchmark model forecasts. The entries in the BA encompasses PC? (PC encompasses BA?) rows indicate whether the BA model (PC combination) forecasts encompass the PC combination (BA model) forecasts according to the results in Tables 2–5 using a 10% significance level. The entries for the average rows report the ratio of the MSFE for a forecast formed as a simple average of the BA model and PC combination forecasts to the MSFE for the AR benchmark model forecasts.

For the cases where both of the $MHLN_h$ statistics are significant, the weights on the BA and PC forecasts in Tables 2-5 are close to 0.50, so taking the mean of the two forecasts is reasonable. This procedure also avoids having to estimate the weights, making it easier to implement in practice.

We examined the robustness of our results along a number of dimensions and obtained similar results. For example, the results are very similar when we use the AIC instead of the SIC to select the lag lengths in (1) and (2). We also computed combination forecasts for a set of potential predictors that excludes manufacturing capital orders and manufacturing and trade sales, two variables that are available with a one-month lag relative to the other potential predictors (and so are not "coincident" with the other predictors). We again obtain very similar results. The complete results for these robustness checks are available at http://pages.slu.edu/faculty/rapachde/Research.htm.