International Journal of Forecasting 31 (2015) 1096–1103

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Optimal combination of survey forecasts

Cristina Conflitti [a,b], Christine De Mol [a,c], Domenico Giannone [a,d,e,f,*]

[a] ECARES, Université libre de Bruxelles, Belgium
[b] Bank of Italy, Italy
[c] Department of Mathematics, Université libre de Bruxelles, Belgium
[d] LUISS University of Rome, Italy
[e] CEPR, United Kingdom
[f] Federal Reserve Bank of New York, United States

ARTICLE INFO

Keywords:
Forecast combination
Forecast evaluation
Survey of Professional Forecasters
Real-time data
Shrinkage
High-dimensional data

ABSTRACT

We consider the problem of combining individual forecasts of real gross domestic product (GDP) growth and Harmonized Index of Consumer Prices (HICP) inflation from the Survey of Professional Forecasters (SPF) for the Euro area. Contrary to the common practice of using equal combination weights, we compute weights which are optimal in the sense that they minimize the mean square forecast error (MSFE) in the case of point forecasts and maximize a logarithmic score in the case of density forecasts. We show that this is a viable strategy even when the number of forecasts to be combined gets large, provided that we constrain these weights to be positive and to sum to one. Indeed, this enforces a form of shrinkage on the weights which ensures a reasonable out-of-sample performance of the combined forecasts.

1. Introduction

The idea of combining individual forecasts provided by different sources in order to achieve an improved accuracy and reliability is quite an old one. There is a vast body of literature on the subject, advocating the usefulness of forecast combination methods both from a theoretical point of view and on the basis of the results of empirical studies. The forecasts being combined can be either judgemental, provided for example by individual forecasters participating in surveys, or else provided by different quantitative models. In the present paper, we focus on survey data, and in particular on the ECB Survey of Professional Forecasters (SPF). The European Central Bank (ECB) has been conducting this survey at a quarterly frequency since the inception of the European Monetary Union (EMU). The survey participants are experts affiliated with financial and non-financial European institutions. They are asked to provide point and density forecasts for GDP growth, HICP inflation and unemployment at different horizons. For a detailed description of the survey, see the papers by Bowles et al. (2007, 2010) and Garcia (2003).

A simple and widely used combination method consists of simply averaging all available forecasts of a given variable, attributing equal weights to the individual predictions. However, the idea of determining optimal combination weights that minimize some objective criterion or cost function is more appealing. When combining point forecasts, a natural target to minimize is the mean square forecast error (MSFE), i.e. the variance of the combination

* Corresponding author at: ECARES, Université libre de Bruxelles, Belgium.

E-mail addresses: cristina.conflitti@bancaditalia.it (C. Conflitti), demol@ulb.ac.be (C. De Mol), Domenico.Giannone@ny.frb.org (D. Giannone).

http://dx.doi.org/10.1016/j.ijforecast.2015.03.009
0169-2070/

around the variable to be predicted. In practice, this minimization can be performed over some available historical periods, so that the optimal weights minimize an empirical least-squares criterion. In economics, this idea dates back to the work of Bates and Granger (1969) and Granger and Ramanathan (1984), and has been the subject of a great variety of developments, including the use of different optimality criteria, time-varying weights, nonlinear combination schemes, etc. For a review of the literature, we refer to the survey papers by Clemen (1989) and Timmermann (2006). More recently, a similar approach has been advocated for density forecasts using combination weights that maximize the so-called logarithmic score (Geweke & Amisano, 2011; Hall & Mitchell, 2007).

A closer look at this body of literature shows that, when dealing with applications, only rather small numbers of individual forecasts are considered for optimal combination, with optimality appearing to be given up as soon as these numbers become large. For example, the recent papers by Geweke and Amisano (2011, 2012) and Sloughter, Gneiting, and Raftery (2010) deal with combinations of just a handful of individual forecasts. In other papers that consider larger numbers of forecasts, either dimension-reduction techniques such as principal components are applied first (as per Chan, Stock, & Watson, 1999; Poncela, Rodrígues, Sánchez-Mangas, & Senra, 2011; Stock & Watson, 2004) or the weights are either taken to be equal or assigned on the sole basis of the previous performance of each forecaster, ignoring mutual dependence (two important recent examples are the papers by Clark & McCracken, 2010, for point forecasts and Jore, Mitchell, & Vahey, 2010, for densities). This strategy appears to be justified empirically by the fact that the resulting simple averaging schemes tend to outperform more sophisticated ones. Such a phenomenon is usually referred to as the "forecast combination puzzle", and has been documented recently for our dataset by Genre, Kenny, Meyer, and Timmermann (2013), who show that the simple equal-weight averages constitute a benchmark which is very hard to improve upon. This explains why this practice is still prevalent today among institutions such as the ECB. Interestingly, a similar phenomenon has been observed by De Miguel, Garlappi, and Uppal (2009) in portfolio optimization, a problem which shares with forecast combination the idea, due to Markowitz, of exploiting diversification in order to decrease the risk/variance.

In the present paper, we show that there is no need to give up optimality when going to a high-dimensional setting, i.e., when combining a large number of forecasts. The reason why previous works either stick to small combinations or, for large datasets, rely on simplified covariance modelling is probably related to two fundamental difficulties: (i) the presence of finite-sample errors and numerical instabilities in the estimation of the weights (see e.g. Smith & Wallis, 2009); and (ii) the need to solve the resulting high-dimensional optimization problem in a way that is efficient computationally. For the cases of both point and density forecast combinations, we argue that the determination of the optimal weights is stabilized by the constraint that they should be positive and add up to one, and we show that the computation of these optimal weights can be implemented easily using iterative algorithms that can handle a large number of forecasts efficiently.

In Section 2, we deal with the combination of point forecasts, defining the optimal weights as those that minimize the MSFE over some historical period, imposing the constraints that these weights must be positive (or, more precisely, nonnegative) and sum to one. Hence, the optimal combination problem reduces to a (possibly high-dimensional) constrained least-squares regression problem, where the complete covariance structure between forecasters is taken into account. We show that the combined use of these two – rather natural – constraints on the weights allows for a proper formulation of the problem, in the sense that it enforces an implicit shrinkage of the weights, rendering their computation stable with respect to errors in the data, even for large panels of forecasters, which is not generally the case for ordinary least-squares estimates in high-dimensional situations. Moreover, the problem turns out to be analogous to the determination of no-short minimum variance Markowitz portfolios, i.e., portfolios for which the weights are constrained to be nonnegative. As was established by Brodie, Daubechies, De Mol, Giannone, and Loris (2009), such portfolios are a special case of a larger family of sparse and stable portfolios that are derived through a constrained "lasso" regression problem. This implies that the weight vector that solves our optimization problem is sparse, i.e., that many of the weights are exactly zero.

The idea of optimal combination can be extended to the case of density forecasts, using an appropriate similarity measure between probability densities, such as a Kullback–Leibler divergence, instead of a least-squares distance. However, in the case of survey data, we miss a target density, since only the realized value of the variable to be forecast, say, GDP growth or HICP inflation, is available. Then, as was proposed by Hall and Mitchell (2007), we show in Section 3 that the optimal weights can be obtained by maximizing a logarithmic score function, under the constraints that the weights are nonnegative and sum to one, which ensures that the combination of densities is still a proper density. To compute such weights, we derive a simple iterative algorithm which scales well with the dimension of the panel, i.e., allows us to handle large datasets.

Section 4 contains our empirical analysis. The SPF point and density forecasts for GDP growth and HICP inflation are combined optimally, as described above, and compared with equal-weight combinations, such as those used by the ECB to summarize the results of each round of the survey. The evaluation is performed by means of a real-time out-of-sample forecasting exercise.

Finally, Section 5 contains the conclusions of our work, as well as some pointers to other potential applications of our combination framework.

2. Optimal combination of point forecasts

As in the paper by Granger and Ramanathan (1984), we address the problem of determining the optimal combination weights for point forecasts as a least-squares regression problem and hence, we use the full covariance structure between forecasters. However, in addition, we