

Statistical post-processing of dual-resolution ensemble forecasts

Sándor Baran¹  | Martin Leutbecher² | Marianna Szabó¹ | Zied Ben Bouallègue²

¹Department of Applied Mathematics and Probability Theory, Faculty of Informatics, University of Debrecen, Debrecen, Hungary

²European Centre for Medium-Range Weather Forecasts Shinfield Park, Reading, UK

Correspondence

Sándor Baran, Department of Applied Mathematics and Probability Theory, Faculty of Informatics, University of Debrecen, Kassai út 26, H-4028 Debrecen, Hungary.
Email: baran.sandor@inf.unideb.hu

Funding information

EFOP-3.6.3-VEKOP-16-2017-00002 project, co-financed by the Hungarian Government and the European Social Fund and the Hungarian National Research, Development and Innovation Office: Grant/Award Number: NN125679.

The computational cost as well as the probabilistic skill of ensemble forecasts depends on the spatial resolution of the numerical weather prediction model and the ensemble size. Periodically, e.g. when more computational resources become available, it is appropriate to reassess the balance between resolution and ensemble size. Recently, it has been proposed to investigate this balance in the context of dual-resolution ensembles, which use members with two different resolutions to make probabilistic forecasts. This study investigates whether statistical post-processing of such dual-resolution ensemble forecasts changes the conclusions regarding the optimal dual-resolution configuration.

Medium-range dual-resolution ensemble forecasts of 2 m temperature have been calibrated using ensemble model output statistics. The forecasts are produced with ECMWF's Integrated Forecast System and have horizontal resolutions between 18 and 45 km. The ensemble sizes range from 8 to 254 members. The forecasts are verified with SYNOP station data. Results show that score differences between various single- and dual-resolution configurations are strongly reduced by statistical post-processing. Therefore, the benefit of some dual-resolution configurations over single-resolution configurations appears to be less pronounced than for raw forecasts. Moreover, the ranking of the ensemble configurations can be affected by the statistical post-processing.

KEYWORDS

dual-resolution, ensemble model output statistics, ensemble post-processing, probabilistic forecasting

1 | INTRODUCTION

Ensemble forecast systems evolve as more computational resources become available. Defining how to upgrade a forecast system requires, among other things, decisions regarding the spatial resolution of the forecast model as well as the ensemble size. The required computational resources strongly depend on these decisions. Thus, it is necessary to find an appropriate compromise between increasing spatial resolution and increasing ensemble size. Recently, Leutbecher and Ben Bouallègue (2018, personal communication) looked at the question of the most skilful ensemble configuration for given computational resources in a *dual-resolution* setting, where k lower-resolution members and m higher-resolution members were combined. For 2 m temperature forecasts in the medium-range, they found that

dual-resolution ensembles with about equal number of lower- and higher-resolution members provided more skilful predictions than alternative single-resolution configurations with either only lower-resolution or only higher-resolution members. Their study looked at forecasts obtained from raw model output by simply pooling together all members.

However, raw ensemble forecasts tend to be underdispersive and can be subject to systematic bias. These deficiencies result in a lack of calibration and they have been documented for several different EPSs (e.g. Buizza *et al.*, 2005). Any lack of calibration calls for some form of statistical post-processing (Buizza, 2018). In the last decade, various methods of statistical calibration have been developed (e.g. Schmeits and Kok, 2010; Ruiz and Saulo, 2012; Williams *et al.*, 2014). Methods like Bayesian model averaging (BMA; Raftery *et al.*, 2005) and ensemble model output

statistics (EMOS; Gneiting *et al.*, 2005) provide full predictive distributions. Once the predictive distribution is given, its functionals (e.g. median or mean) can easily be calculated and considered as point forecasts.

The BMA predictive distribution of a future weather quantity is a weighted mixture of probability distributions corresponding to the individual ensemble members with weights determined by the predictive performance of the members during the training period. BMA models for various weather quantities differ in the distribution of the mixture components: for temperature or pressure a normal mixture is suggested (Raftery *et al.*, 2005), wind speed can be modelled using gamma (Sloughter *et al.*, 2010) or truncated normal (Baran, 2014) distributions, whereas for precipitation accumulation a discrete-continuous gamma model was developed (Sloughter *et al.*, 2007).

Here we concentrate on the essentially simpler EMOS, or non-homogeneous regression, approach where the predictive distribution is given by a single parametric probability density function (PDF) with parameters depending on the ensemble. Similar to the BMA, different weather quantities require different distributions. Temperature and pressure can again be modelled by normal distributions (Gneiting *et al.*, 2005), wind speed requires non-negative and skewed distributions such as truncated normal (Thorarinsdottir and Gneiting, 2010), generalized extreme value (Lerch and Thorarinsdottir, 2013), log-normal (Baran and Lerch, 2015) or their mixture (Baran and Lerch, 2016), whereas censored generalized extreme value (Scheuerer, 2014) and censored shifted gamma (Scheuerer and Hamill, 2015; Baran and Nemoda, 2016) distributions and their various combinations (Baran and Lerch, 2018) provide good models for precipitation accumulation.

In this paper, we examine whether dual-resolution ensembles are still preferable to single-resolution ensembles after statistical post-processing. Dual-resolution ensembles fall in the category of multi-model ensembles, where each contributing model has specific error characteristics. Previous studies on post-processing of multi-model forecasts have focused on the optimization of the weights attributed to each contributing model (Doblas-Reyes *et al.*, 2005; Casanova and Ahrens, 2009; DelSole *et al.*, 2013; Raynaud *et al.*, 2015). Here, following an EMOS approach, the target is to provide as a forecast a full probability distribution optimized in terms of probabilistic skill.

The forecast and verification data used in this study are described in section 2. Then, the EMOS calibration and verification methodologies are presented in section 3. The results and conclusions follow in sections 4 and 5, respectively.

2 | FORECAST AND OBSERVATION DATA

This study applies EMOS calibration to global medium-range ensemble forecasts with ECMWF's Integrated Forecast System (IFS) with three horizontal resolutions. The

lower-resolution ensembles have four to five times more members than the highest resolution ensemble

- 50 members at TCo639 (grid resolution ~ 18 km),
- 200 members at TCo399 (grid resolution ~ 29 km),
- 254 members at TCo255 (grid resolution ~ 50 km).

The 50-member TCo639 ensemble is the operational ECMWF medium-range ensemble while the lower-resolution ensembles are generated with the same model version as the operational ensemble at the time (cycle 41r2). Leutbecher (2018) analysed the data from the 200-member TCo399 experiment and describes the perturbation methodology for the initial conditions and the stochastic representation of model uncertainties, which is identical in the three ensembles. The period of investigation is boreal summer 2016. The ensemble forecasts are initialized once daily between 1 June and 31 August 2016.

The cost ratio between a TCo639 forecast and a TCo399 forecast is about 4:1 and the cost ratio between a TCo639 forecast and a TCo255 forecast is around 16:1. When constructing different dual-resolution configurations, one TCo639 member can be traded against four TCo399 members or against sixteen TCo255 members. Various configurations of higher- and lower-resolution ensemble forecasts are studied. We will consider TCo399–TCo639 dual-resolution ensembles and TCo255–TCo639 ensembles as well as two scenarios corresponding to different assumptions on available HPC resources. The *large supercomputer* (LHPC) scenario assumes the availability of current HPC resources of the ECMWF, whereas the *small supercomputer* (SHPC) scenario assumes one sixth of these resources; Table 1 gives an overview. Note that TCo399–TCo639 and TCo255–TCo639 combinations are based on different ensemble sizes of TCo639 members, so in fact one has two different LHPC scenarios. The reason is completely technical, as the cost of 50 TCo639 members is equivalent to the cost of 800 TCo255 members; however, with the current GRIB settings, the largest possible ensemble size to be handled is 255.

The focus is on forecasts of 2 m temperature. Post-processing and verification of the ensemble forecasts rely on measurements from surface synoptic observation (SYNOP) stations. Measurements are reported from various locations around the globe with low observation densities in tropical and sub-tropical regions and high densities in Europe and the North American continent. The number of observations available varies from day to day. A subset of about 4,500 stations with full availability over the verification period is used in this study.

Observations are compared to forecasts at the nearest grid point of the native forecast grid. To account for systematic representativity errors related to the coarse description of the orography in the model, an orographic correction is applied to the raw temperature forecasts. For each forecast, the correction corresponds to an adjustment linear with the

TABLE 1 Investigated dual-resolution mixtures

TCo399–TCo639				TCo255–TCo639			
Large HPC		Small HPC		Large HPC		Small HPC	
M_L	M_H	M_L	M_H	M_L	M_H	M_L	M_H
0	50	0	8	0	16	0	8
40	40	8	6	16	15	16	7
120	20	16	4	32	14	32	6
160	10	24	2	64	12	64	4
200	0	28	1	128	8	128	0
		32	0	254	0		




FIGURE 1 Stationwise differences between high-resolution (TCo639) and low-resolution (a–c, TCo399) and (d–f, TCo255) ensemble characteristics at day 5 over Europe (~1,300 stations). (a, d) show mean differences (K), (b, e) mean ensemble variance differences (K^2), and (c, f) mean root mean squared error differences (K). Statistics are computed for ensembles with 50 members each

height difference between station and model representation ($\Delta T/\Delta z = -0.0065$ K/m).

A basic comparison of the statistics of lower- and higher-resolution forecasts is shown in Figure 1. The distributions of mean differences at the station level are plotted: differences in terms of mean ensemble forecast, mean ensemble variance, and mean forecast error, as measured by the root mean squared error of the ensemble mean. The comparison is performed considering ensembles of the same size (50) in order to focus on the impact of the model resolution and not on the impact of the ensemble size. The statistics shown in Figure 1 are valid for forecasts at day 5, for stations located in Europe.

Ensembles with different model resolutions have on average very small mean differences, but they can exhibit large mean differences (up to 2 K) at certain locations (Figure 1a,d). The ensemble forecasts at high resolution (TCo639) have on average a larger mean variance (Figure 1b,e) and a smaller mean error (Figure 1c,f) at the station level than

ensemble forecasts at a lower resolution. Here again, the station-to-station variability can be large. Comparing (a, b, c) with (d, e, f) shows that the differences in terms of variance and error are slightly larger when the difference of the underlying model resolutions increases.

3 | ENSEMBLE MODEL OUTPUT STATISTICS

As mentioned earlier, the normal distribution provides a good model for temperature, however one has to appropriately link the predictive PDF with the ensemble members.

3.1 | Model and parameter estimation

The EMOS approach suggested by Gneiting *et al.* (2005) considers the mean and variance of the predictive PDF to be affine functions of the ensemble forecasts and the ensemble

variance, respectively. If $f_1, f_2, \dots, f_K$ denote the ensemble forecasts for temperature for a given location, time and lead time, the associated predictive distribution of temperature will be

$$\mathcal{N}(a + b_1 f_1 + \dots + b_K f_K, c + dS^2), \quad (1)$$

with

$$S^2 := \frac{1}{K-1} \sum_{k=1}^K (f_k - \bar{f})^2,$$

where $\bar{f}$ denotes the ensemble mean.

In a multi-model context with exchangeable ensemble groups, Gneiting (2014) suggests using the same coefficients within a given group (also Fraley *et al.*, 2010). Hence, if we have M ensemble members divided into K exchangeable groups, where the k th group contains $M_k \geq 1$ ensemble members ($\sum_{k=1}^K M_k = M$) and has mean $\bar{f}_k$, the EMOS predictive distribution will be

$$\mathcal{N}(a + b_1 \bar{f}_1 + \dots + b_K \bar{f}_K, c + dS^2), \quad (2)$$

with S^2 denoting again the ensemble variance.

According to the optimal score estimation approach of Gneiting and Raftery (2007), model parameters $a, b_1, \dots, b_K$ and c, d are determined by optimizing the mean value of a proper scoring rule as a function of the parameters over suitably chosen training data. Scoring rules measure the forecast skill by numerical values assigned to pairs of forecasts and observations, and for predictive distributions the most popular ones are the logarithmic score, that is, the negative logarithm of the predictive PDF evaluated at the verifying observation (Gneiting and Raftery, 2007) and the continuous ranked probability score (CRPS; Gneiting and Raftery, 2007; Wilks, 2011). Given a (predictive) cumulative distribution function (CDF) $F(y)$ and real value (observation) x the CRPS is defined as

$$\begin{aligned} \text{CRPS}(F, x) &:= \int_{-\infty}^{\infty} (F(y) - \mathbb{I}_{\{y \geq x\}})^2 dy \\ &= \mathbb{E}|X - x| - \frac{1}{2} \mathbb{E}|X - X'|, \end{aligned} \quad (3)$$

where $\mathbb{I}_H$ denotes the indicator of a set H , whereas X and X' are independent random variables with CDF F and finite first moment. The right-hand side of Equation 3 shows that the CRPS has the same unit as the observation, moreover for a normal distribution it has a simple closed form (Gneiting *et al.*, 2005). One should also mention that both CRPS and logarithmic score are negatively oriented, that is the smaller the better, and optimization with respect to the latter provides the maximum-likelihood estimates of the model parameters.

The choice of the training data is important for statistical post-processing. For estimating the EMOS model parameters, a rolling training period is applied and the estimates are obtained using ensemble forecasts and corresponding validating observations for the preceding n calendar days. Given a training period length, there are two traditional approaches for selecting the training data (Thorarinsdottir and Gneiting, 2010). In the global (regional) approach, parameters are

estimated using all available forecast cases from the training period resulting in a single universal set of parameters across the entire ensemble domain. It requires quite short training periods, but usually it is unsuitable for large and heterogeneous observation domains. For local parameter estimation, one has distinct parameter estimates for the different stations obtained only using training data of the given station. To avoid numerical stability problems, local EMOS requires much longer training periods (e.g. Hemri *et al.*, 2014 address optimal training period length for different weather quantities), but if the training data is large enough, it will usually outperform the regional approach. To combine the advantages of local and regional EMOS, Lerch and Baran (2017) introduced two semi-local methods where the training data for a given station is augmented with data from stations with similar characteristics. The choice of similar stations is based either on suitably defined distance functions or on clustering. Here we focus on the clustering-based semi-local estimation, where the observation sites are grouped into clusters using k -means clustering of feature vectors depending both on the station climatology and the forecast errors of the raw ensemble during the training period. A regional parameter estimation is then performed within each cluster. With the help of this method, one can get reliable parameter estimates even for short training periods and the obtained models may outperform the local EMOS approach (Lerch and Baran, 2017). Hence, section 4 focuses mainly on results corresponding to semi-local EMOS post-processing.

3.2 | Verification scores

The fundamental aim of probabilistic forecasting is to access the maximal sharpness of the forecast distribution subject to calibration (Gneiting *et al.*, 2007), where the former refers to the concentration of the predictive distribution and the latter to the statistical consistency between the predictive distributions and the validating observations. A standard tool of quantifying the predictive performance of probabilistic forecasts both in terms of calibration and sharpness is the mean CRPS over all forecast cases.

Besides the CRPS, one can also consider Brier scores (BS; Wilks, 2011, section 8.4.2) for the dichotomous event that the observed temperature x exceeds a given threshold y . For a predictive CDF $F(y)$, the Brier score is defined as

$$\text{BS}(F, x; y) := (F(y) - \mathbb{I}_{\{y \geq x\}})^2,$$

(e.g. Gneiting and Ranjan, 2011), and note that the CRPS is the integral of the BS over all possible thresholds. In the results provided in section 4, we consider as thresholds the 5, 10, ..., 90, and 95 percentiles of the corresponding station climatology for the verification period.

Further, let $q_\tau(F)$ denote the τ -quantile ($0 \leq \tau \leq 1$) of a CDF $F(y)$, that is

$$q_\tau(F) := F^{-1}(\tau) := \inf\{y : F(y) \geq \tau\},$$

and consider the loss function

$$\rho_\tau(x) := \begin{cases} \tau|x|, & \text{if } x \geq 0, \\ (1-\tau)|x|, & \text{if } x < 0. \end{cases}$$

Then for a given value x , the quantile score (QS; e.g. Bentzien and Friederichs, 2014) is defined as

$$QS_\tau(F, x) := \rho_\tau(x - q_\tau(F)).$$

In the present study QS values corresponding to the 2, 5, 10, 20, 50, 80, 90, 95, and 98 percentiles of the predictive distribution and the raw ensemble are considered. Note that, to evaluate CRPS, BS and QS of the raw ensemble, one has to replace the predictive CDF by the empirical one.

The improvement in BS and QS with respect to a reference predictive distribution F_{ref} can be measured with the help of the Brier skill score (BSS) and the quantile skill score (QSS) defined as

$$BSS(F, F_{\text{ref}}, x; y) := 1 - \frac{BS(F, x; y)}{BS(F_{\text{ref}}, x; y)}$$

and

$$QSS_\tau(F, F_{\text{ref}}, x) := 1 - \frac{QS_\tau(F, x)}{QS_\tau(F_{\text{ref}}, x)},$$

respectively (Friederichs and Thorarinsdottir, 2012). Obviously, in contrast to the BS and QS, the corresponding skill scores are positively oriented, that is the larger the better.

Finally, point forecasts such as ensemble/EMOS medians and means are evaluated using mean absolute errors (MAEs) and root mean squared errors (RMSEs), where the former is optimal for the median and the latter for the mean (Gneiting, 2011), although for the normal EMOS model these quantities coincide. As mean of QS_{50} , the quantile score of the median forecast, over all forecast cases is exactly half of the MAE, in what follows we will report only the RMSE values.

As suggested by Gneiting and Ranjan (2011), to assess the statistical significance of the differences between the verification scores, we make use of the Diebold–Mariano (DM) test of equal predictive performance, as it allows us to account for the temporal dependencies in the forecast errors (Diebold and Mariano, 1995). Baran and Lerch (2018) give more details about the DM test. Further, confidence intervals for mean score values and mean score differences are obtained with the help of 2,000 block bootstrap samples using the stationary bootstrap scheme with mean block length according to Politis and Romano (1994).

4 | RESULTS

4.1 | Preliminaries

As mentioned in the Introduction, EMOS calibration is applied to different dual-resolution ensemble forecasts for 2 m temperature. In what follows, let $f_{H,1}, f_{H,2}, \dots, f_{H,M_H}$ and $f_{L,1}, f_{L,2}, \dots, f_{L,M_L}$ denote the higher- and lower-resolution

ensemble members, respectively, and denote by $\bar{f}_H$ and $\bar{f}_L$ the corresponding ensemble means. As ensemble members of a given resolution can be considered exchangeable, model (2) reduces to

$$\mathcal{N}(a + b_H \bar{f}_H + b_L \bar{f}_L, c + dS^2), \quad (4)$$

with $b_L = 0$ for pure higher-resolution ($M_L = 0$) and $b_H = 0$ for pure lower resolution ($M_H = 0$) configurations and S^2 being the variance of the pooled ensemble defined by Equation 1. Following the ideas of Gneiting *et al.* (2005), model parameters are estimated by minimizing the mean CRPS of the predictive distributions and validating observations corresponding to the forecast cases of the training data.

For the case of TCo399–TCo639 combinations, stationwise statistical tests of equality of the variances and distributions of the two full-size ensembles ($M_H = 50, M_L = 200$) suggest treating the variances of the component ensembles separately by considering the predictive distribution

$$\mathcal{N}(a + b_H \bar{f}_H + b_L \bar{f}_L, c + d_H S_H^2 + d_L S_L^2), \quad (5)$$

with S_H^2 and S_L^2 denoting the variances of the high- and low-resolution components, respectively. However, at least for the TCo639–TCo399 combination, the verification scores of the more complex model (5) do not differ significantly from those of model (4), so all our calibration results are obtained using the latter.

As the dataset at hand covers only the boreal summer 2016 (92 calendar days), one has to consider a relatively short training period length that still allow reliable parameter estimation to leave enough data for model verification. As a trade-off between the two requirements, calibration is performed using a rolling 30-day training period. This means that verification scores are calculated for ensemble forecasts initialized between 1 July and 31 August 2016 and the corresponding validating observations. Obviously, the forecast periods are shifted by 1–15 days according to the lead times of the ensemble predictions.

As global EMOS is not justifiable for global data, local and semi-local approaches to parameter estimation are considered. Similar to Lerch and Baran (2017), the k -means clustering of stations is based on 24-dimensional feature vectors consisting of 12 equidistant quantiles of the climatological CDF and 12 equidistant quantiles of the empirical CDF of forecast errors of the ensemble mean over the training period. Two hundred clusters are considered, which yields a similar mean station number per cluster as in Lerch and Baran (2017). Local EMOS estimates 4 to 5 parameters from 30 forecast–observation pairs while semi-local EMOS estimates the same parameters from about 600 forecast–observation pairs. Therefore, the latter should be able to constrain the parameters much better.

In order to highlight the differences between local and semi-local approaches, a very short 10-day training period is also investigated.




FIGURE 2 (a, b) Mean CRPS values (the lower the better) of global dual-resolution ensemble forecasts for 2 m temperature and (c, d) the difference in mean CRPS (the lower the better) from the reference pure high-resolution ensemble with 95% confidence intervals, LHPC scenario. Resolutions are (a, c) TCo399-TCo639 and (b, d) TCo255-TCo639 [Colour figure can be viewed at wileyonlinelibrary.com]

To be fully consistent with the results for the raw ensemble, EMOS predictive distributions are obtained using the orographically corrected ensemble forecasts. Local EMOS does not require such a preliminary bias correction but, for semi-local EMOS, the use of this correction yields improved skill in terms of verification scores because local variability within a cluster can still be represented through the orographic correction.

4.2 | Calibration of mixtures for large supercomputer

For raw ensemble forecasts, both resolution combinations prefer balanced mixtures practically for all lead times, i.e. (40,40) for TCo399-TCo639 and (16,15) for TCo255-TCo639. This can be clearly observed in Figure 2 showing the mean CRPS values of dual-resolution ensemble forecasts for 2 m temperature and the difference in mean CRPS with respect to the pure high-resolution case as a function of lead time. It is expected that the optimal reduction of biases that are not shared among the different resolutions helps to explain why the balanced mixtures turn out to be best. Work by Leutbecher and Ben Bouall  gue (2018, personal communication) is investigating this aspect quantitatively.

Results in Figure 2 are consistent with other studies that examined the sensitivity of the skill of ensemble forecasts to horizontal resolution and ensemble size. For instance, Raynaud and Bout  tier (2017) show qualitatively consistent results

at shorter lead times for experimentation with a limited-area ensemble at resolutions of 2.5 and 1.3 km and ensemble sizes up to 34 members. Horizontal resolution is more important at shorter lead times while ensemble size matters more at longer lead times.

Statistical post-processing substantially changes this picture. As depicted in Figure 3, semi-local EMOS significantly decreases the mean CRPS for all lead times, and the differences between the predictive performance of the various mixture combinations is strongly reduced: more so for TCo399-TCo639 than for TCo255-TCo639. Local EMOS yields very similar results (not shown).

Differences in mean CRPS with respect to the pure high-resolution case are shown in Figure 4. After day 8, the pure low-resolution ensemble outperforms the pure high-resolution one both for TCo399-TCo639 and TCo255-TCo639 resolutions, whereas for short lead times, semi-local EMOS prefers balanced combinations, which are also optimal in the raw ensemble. Note that for longer lead times, score differences based on local EMOS are more variable leading to wider confidence intervals. Hence, in the remaining part of this section we present mainly the results of the semi-local EMOS calibration.

To investigate the statistical significance of differences between the mean CRPS values corresponding to various combinations, stationwise DM tests were performed. Figure 5 shows the proportion of stations with significantly different mean CRPS at a 5% level for different lead times for local




FIGURE 3 Mean CRPS values (the lower the better) of semi-local EMOS post-processed global dual-resolution ensemble forecasts for 2 m temperature, LHPC scenario. Resolutions are (a) TCo399–TCo639 and (b) TCo255–TCo639 [Colour figure can be viewed at wileyonlinelibrary.com]




FIGURE 4 Difference in mean CRPS (the lower the better) from the reference pure high-resolution case with 95% confidence intervals for (a, b) local and (c, d) semi-local EMOS post-processed global dual-resolution ensemble forecasts of 2 m temperature, LHPC scenario. Resolutions are (a, c) TCo399–TCo639 and (b, d) TCo255–TCo639 [Colour figure can be viewed at wileyonlinelibrary.com]

(lower triangle) and semi-local (upper triangle) parameter estimation approaches. In general, longer lead times result in a smaller proportion of stations with significant difference both for local and semi-local EMOS post-processing and the values in the first column and row of each matrix are consistent with Figure 4, respectively.

A natural question to investigate is whether stations with significant difference in mean CRPS between the optimal combination and the reference pure high-resolution case exhibit some clear spatial pattern. However, visualization of stations with significant difference on a 5% level on maps does not reveal any connection to their location.

The analysis of Brier skill scores with thresholds equal to the 5, 10, ..., 95 percentiles of the corresponding station sample climatology for the verification period leads us to

similar conclusions. BSS values with respect to the pure high-resolution case of semi-local EMOS post-processed forecasts displayed in Figure 6 are fully consistent with the graphs in the bottom row of Figure 4. At days 1 and 5, the balanced combinations have the best forecast skill for all thresholds, whereas at day 10 all other combinations outperform the pure high-resolution one.

Figure 7 shows differences of quantile skill scores for ensemble configurations post-processed with semi-local EMOS. All differences are with respect to the pure higher-resolution configuration. For the median (50% quantile; c, d), results are fully consistent with the CRPS differences. For the more extreme quantiles (2% and 98%; a, b and e, f respectively), score differences are comparable to those for the median but the confidence intervals tend to be larger.




FIGURE 5 Proportion of stations with significantly different mean CRPS at a 5% level for lead times (a, b) 1, (c, d) 5, (e, f) 10 days for local (lower triangle) and semi-local (upper triangle) parameter estimation approaches, LHPC scenario. Resolutions are (a, c, e) TC0399–TC0639 and (b, d, f) TC255–TC0639 [Colour figure can be viewed at wileyonlinelibrary.com]

In any case, quantile score differences between configurations are statistically not strongly significant for the longer lead times.

Finally, consider the root mean squared errors of EMOS median forecasts. Figure 8 displays difference in RMSE values of semi-local EMOS post-processed combinations from the reference pure high-resolution case. Note that

for both mixtures the graphs are very similar to the ones in Figure 4c,d and the same similarity of differences in mean CRPS (Figure 4a, b) and RMSE (not reported) can be observed for local EMOS as well.

One can conclude that for the LHPC scenario, all examined verification scores show a consistent picture. In case of the TC0399–TC0639 mixture up to day 9 the balanced




FIGURE 6 Brier skill scores (the higher the better) with respect to the reference pure high-resolution case with 95% confidence intervals of semi-local EMOS post-processed global dual-resolution ensemble forecasts for lead times (a, b) 1, (c, d) 5, (e, f) 10 days for 2 m temperature, LHPC scenario. Resolutions are (a, c, e) TCo399–TCo639 and (b, d, f) TCo255–TCo639 [Colour figure can be viewed at wileyonlinelibrary.com]

combination (40,40) has the best predictive performance, whereas from day 10 all combinations outperform the reference pure high-resolution one, however the difference between the scores of combinations involving TCo399 forecasts are often very small. This situation differs from the case of the raw ensemble where the balanced combination exhibits the best forecast skills for all lead times (Figure 2a, c). For the TCo255–TCo639 mixture the balanced combination (16,15) is clearly preferred up to day 5; from day 10 again all combinations outperform the reference pure high-resolution one, but with more pronounced differences than for the other mixture. For long lead times, the pure low-resolution ensemble performs very well and the ordering of combinations again differs from the one based on the raw ensemble forecasts, which identify a balanced combination as best for all lead times (Figure 2b, d).

4.3 | Calibration of mixtures for small supercomputer

Now, we examine dual-resolution configurations for the small HPC scenario. In terms of the raw ensemble, the balanced combinations ((8,6) for TCo399–TCo639 and (16,7) for TCo255–TCo639) are the most skilful for short lead times, whereas for longer lead times, the larger the ensemble size, the better, i.e. a pure low-resolution ensemble has the best forecast skill (Figure 9). Results differ from the large supercomputer scenario because the maximum ensemble size for TCo639 is much smaller here. The skill of the eight-member TCo639 ensemble is considerably worse than the skill of all other dual-resolution combinations from a lead time of 5 days (8 days) for the combinations with TCo399 (TCo255). This is consistent with the sensitivity of the CRPS to ensemble size as documented for instance by Leutbecher (2018).




FIGURE 7 Quantile skill scores (the higher the better) for percentiles (a, b) 2, (c, d) 50 and (e, f) 98 with respect to the reference pure high-resolution case with 95% confidence intervals of semi-local EMOS post-processed global dual-resolution ensemble forecasts for 2 m temperature, LHPC scenario. Resolutions are (a, c, e) TC0399–TC0639 and (b, d, f) TC0255–TC0639 [Colour figure can be viewed at wileyonlinelibrary.com]

Statistical post-processing with semi-local EMOS significantly improves the calibration in terms of the mean CRPS and reduces the differences between the various combinations consistent with LHPC results (Figure 10). The score differences are shown in Figure 11 for local and semi-local EMOS. Both parameter estimation approaches provide an overall consistent ranking: at smaller lead times, the more balanced combinations are best, whereas for longer lead times the larger the ensemble size, the better the performance.

Statistical significance of score differences between configurations has also been computed stationwise (not shown). The longer the lead time, the smaller is the proportion of stations where differences are significant. Compared with the LHPC scenario, at day 1 the pattern is similar to Figure 5a,b. At day 5, the proportions are smaller than for LHPC both for TC0399–TC0639 and TC0255–TC0639 mixtures, whereas at day 10, there are more stations where the difference in mean CRPS is significant in SHPC than in LHPC, which can be

explained by the fact that, at longer lead times, the ensemble size is particularly important in the SHPC scenario.

For TC0399–TC0639 mixtures, the Brier skill scores with respect to the reference pure high-resolution case (Figure 12a,c,e) at day 1 indicate the advantage of the most balanced combination (8,6); at day 5 all combinations involving low-resolution members outperform the post-processed pure high-resolution ensemble, whereas at day 10 the ordering clearly reflects the ensemble size, that is the larger the better. For the TC0255–TC0639 mixtures (Figure 12b,d,f), combinations with larger ensemble sizes are considered and both at days 1 and 5 the post-processed 128-member pure low-resolution ensemble underperforms the eight-member pure high-resolution one and the effect of the ensemble size on the predictive performance dominates only at day 10. Note that the results for the Brier scores are fully consistent with those for the CRPS (Figures 12 and 11, respectively).




FIGURE 8 Difference in RMSE values (the lower the better) from the reference pure high-resolution case with 95% confidence intervals of semi-local EMOS post-processed global dual-resolution ensemble forecasts for 2 m temperature, LHPC scenario. Resolutions are (a) TCo399–TCo639 and (b) TCo255–TCo639 [Colour figure can be viewed at wileyonlinelibrary.com]




FIGURE 9 (a, b) Mean CRPS values (the lower the better) of global dual-resolution ensemble forecasts for 2 m temperature and (c, d) the difference in mean CRPS (the lower the better) from the reference pure high resolution ensemble with 95% confidence intervals, SHPC scenario. Resolutions are (a, c) TCo399–TCo639 and (b, d) TCo255–TCo639 [Colour figure can be viewed at wileyonlinelibrary.com]




FIGURE 10 Mean CRPS values (the lower the better) of semi-local EMOS post-processed global dual-resolution ensemble forecasts for 2 m temperature, SHPC scenario. Resolutions are (a) TCo399–TCo639 and (b) TCo255–TCo639 [Colour figure can be viewed at wileyonlinelibrary.com]




FIGURE 11 Difference in mean CRPS (the lower the better) from the reference pure high-resolution case with 95% confidence intervals of (a, b) local and (c, d) semi-local EMOS post-processed global dual-resolution ensemble forecasts for 2 m temperature, SHPC scenario. Resolutions are (a, c) TCo399–TCo639 and (b, d) TCo255–TCo639 [Colour figure can be viewed at wileyonlinelibrary.com]

Further, similar to the LHPC scenario, the 50% quantile skill score differences with respect to the reference pure high-resolution case (not shown) are in line with the results for the CRPS. For the tails, the benefit of the various combinations with respect to the pure high-resolution case are of similar amplitude as that for the LHPC scenario, and especially for the TCo399–TCo639 mixture the differences are often non-significant.

Finally, evaluation of the accuracy of the mean of different post-processed dual-resolution ensemble forecasts in terms of the root mean squared error shows the same picture as the verification scores for probabilistic forecasts. The differences in RMSE values from the reference pure high-resolution case for semi-local EMOS (not shown) are very similar to those for the CRPS.

Similar to the LHPC scenario, all investigated verification scores indicate similar conclusions. The balanced combination (8,6) and (16,7) are preferred up to 4 days, although the differences compared to combinations (16,4) and (32,6), respectively, are very small. For TCo399–TCo639 after day 7, whereas for TCo255–TCo639 after day 9, the post-processed pure low-resolution ensemble shows the best predictive performance, and for long lead times the forecast skill depends on the ensemble size. This behaviour is rather similar to that seen in the raw ensemble (compare Figures 9 and 11).

4.4 | Calibration using a very short training period

Additional results have been obtained using a 10-day training period in order to compare improvements from local and

semi-local EMOS post-processing. The TCo399–TCo639 ensemble configurations of the LHPC scenario have been calibrated for lead times up to 5 days. For local EMOS, this implies ten forecast–observation pairs are used to estimate 4–5 parameters, whereas semi-local EMOS uses on average around 225 forecast–observation pairs for the same purpose for the price of having a universal set of parameters within each cluster. For verification, we use data of the same 62 calendar days as in sections 4.2 and 4.3, i.e. ensemble forecasts initialized between 1 July and 31 August 2016 and the corresponding validating observations.

Table 2 lists verification scores for lead times 1 and 5 days and the corresponding 95% confidence intervals. Note that all post-processing approaches for all combinations substantially outperform the raw ensemble. For the 30-day training period, there is no clear preference for local EMOS or semi-local EMOS. However, one should mention that the uncertainty in score differences of local EMOS post-processed forecasts is much higher, especially for longer lead times. In contrast, for the ten-day training period, semi-local EMOS is clearly superior to local EMOS. Thus, clustering-based semi-local estimation of EMOS parameters provides a reasonable alternative to the local approach, especially in situations where ensemble data cover only a rather short time period.

5 | CONCLUSIONS

The EMOS calibration leads to substantial improvements in skill for all examined single- and dual-resolution ensemble




FIGURE 12 Brier skill scores (the higher the better) with respect to the reference pure high-resolution case with 95% confidence intervals of semi-local EMOS post-processed global dual-resolution ensemble forecasts of 2 m temperature for lead times (a, b) 1, (c, d) 5, (e, f) 10 days, SHPC scenario. Resolutions are (a, c, e) TC0399–TC0639 and (b, d, f) TC0255–TC0639 [Colour figure can be viewed at wileyonlinelibrary.com]

forecasts. For instance, the CRPS decreases from values around 1.3 K to values just under 1.0 K at day 3 using the semi-local EMOS. The improvements are large, although not as large as those reported by Hemri *et al.* (2014). Their raw ensemble forecasts have significantly larger CRPS values than ours. We attribute this difference to having applied an orographic correction to the forecasts. Our 'raw forecasts' have been adjusted with a simple correction that is proportional to the altitude difference between the model orography and the station height while Hemri *et al.* (2014) showed results for uncorrected forecasts.

In terms of calibration methodology, the clustering-based semi-local estimation of EMOS parameters provides a reasonable alternative to the local approach, especially in situations where ensemble data cover only a rather short time period. This is fully in line with results reported by Lerch and Baran

(2017), where multi-model ensemble forecasts of wind speed over Europe and North Africa were calibrated.

The EMOS calibration parameters were obtained by optimizing skill in terms of the CRPS. EMOS improves not only the probabilistic forecasts but also the accuracy of point forecasts like the ensemble mean or the median. Moreover, when comparing ensemble configurations after calibration, score differences for the Brier score and the quantile score show fairly consistent rankings of the ensemble configurations across different event thresholds and probability levels, respectively.

What emerges as the best single- or dual-resolution configuration can change due to EMOS calibration. For instance, the TC0399–TC0639 dual-resolution configuration (40, 40) is best at all lead times for raw forecasts in the large super-computer scenario. After calibration, this configuration is still

TABLE 2 Verification scores of 2 m temperature (K) with 95% confidence intervals for the raw ensemble, local and semi-local EMOS post-processed forecasts using 10-day and 30-day training periods, TCo399–TCo639 mixture, LHPC scenario. Lower values imply more skilful forecasts for all metrics

Combination	Model	CRPS	RMSE	QS2	QS50	QS98
Day 1						
(0,50)	Raw ensemble	1.258 (1.237,1.279)	2.068 (2.037,2.103)	0.391 (0.364,0.419)	0.754 (0.744,0.765)	0.340 (0.330,0.350)
	Local EMOS, 10d	0.925 (0.910,0.939)	1.720 (1.692,1.747)	0.134 (0.130,0.139)	0.626 (0.616,0.636)	0.111 (0.108,0.113)
	Semi-local EMOS, 10d	0.894 (0.878,0.909)	1.683 (1.654,1.710)	0.105 (0.102,0.109)	0.619 (0.609,0.629)	0.086 (0.085,0.088)
	Local EMOS, 30d	0.851 (0.833,0.868)	1.617 (1.586,1.646)	0.097 (0.094,0.101)	0.591 (0.579,0.602)	0.081 (0.079,0.084)
	Semi-local EMOS, 30d	0.869 (0.850,0.887)	1.648 (1.613,1.679)	0.093 (0.090,0.097)	0.606 (0.593,0.618)	0.081 (0.079,0.083)
(40,40)	Raw ensemble	1.226 (1.205,1.249)	2.053 (2.020,2.088)	0.351 (0.327,0.377)	0.752 (0.741,0.764)	0.292 (0.283,0.301)
	Local EMOS, 10d	0.935 (0.920,0.950)	1.733 (1.707,1.757)	0.141 (0.137,0.146)	0.630 (0.620,0.640)	0.117 (0.114,0.120)
	Semi-local EMOS, 10d	0.889 (0.873,0.903)	1.671 (1.645,1.696)	0.106 (0.102,0.110)	0.615 (0.604,0.625)	0.085 (0.084,0.087)
	Local EMOS, 30d	0.846 (0.828,0.864)	1.605 (1.576,1.633)	0.099 (0.096,0.102)	0.587 (0.575,0.599)	0.082 (0.080,0.085)
	Semi-local EMOS, 30d	0.864 (0.845,0.881)	1.637 (1.606,1.667)	0.093 (0.090,0.097)	0.602 (0.589,0.615)	0.080 (0.078,0.082)
(120,20)	Raw ensemble	1.270 (1.250,1.292)	2.093 (2.063,2.127)	0.362 (0.335,0.388)	0.771 (0.761,0.781)	0.300 (0.291,0.309)
	Local EMOS, 10d	0.935 (0.920,0.950)	1.732 (1.706,1.757)	0.141 (0.138,0.146)	0.630 (0.621,0.639)	0.117 (0.114,0.120)
	Semi-local EMOS, 10d	0.894 (0.878,0.908)	1.679 (1.651,1.704)	0.106 (0.102,0.110)	0.619 (0.608,0.628)	0.086 (0.084,0.088)
	Local EMOS, 30d	0.847 (0.829,0.864)	1.606 (1.575,1.634)	0.099 (0.096,0.102)	0.588 (0.576,0.598)	0.082 (0.079,0.085)
	Semi-local EMOS, 30d	0.869 (0.851,0.886)	1.645 (1.611,1.674)	0.093 (0.090,0.097)	0.606 (0.593,0.613)	0.081 (0.079,0.083)
(160,10)	Raw ensemble	1.289 (1.267,1.311)	2.109 (2.078,2.139)	0.373 (0.346,0.402)	0.774 (0.764,0.785)	0.311 (0.302,0.320)
	Local EMOS, 10d	0.937 (0.922,0.953)	1.736 (1.710,1.760)	0.142 (0.138,0.147)	0.631 (0.622,0.641)	0.118 (0.115,0.120)
	Semi-local EMOS, 10d	0.896 (0.881,0.911)	1.685 (1.657,1.710)	0.106 (0.102,0.110)	0.621 (0.610,0.632)	0.087 (0.085,0.089)
	Local EMOS, 30d	0.848 (0.829,0.865)	1.608 (1.575,1.635)	0.099 (0.096,0.102)	0.588 (0.576,0.599)	0.082 (0.079,0.085)
	Semi-local EMOS, 30d	0.871 (0.852,0.889)	1.648 (1.613,1.678)	0.093 (0.090,0.097)	0.607 (0.594,0.619)	0.081 (0.078,0.083)
(200,0)	Raw ensemble	1.304 (1.282,1.326)	2.121 (2.092,2.152)	0.406 (0.378,0.433)	0.776 (0.766,0.787)	0.345 (0.334,0.356)
	Local EMOS, 10d	0.943 (0.928,0.959)	1.754 (1.727,1.779)	0.136 (0.132,0.141)	0.639 (0.628,0.649)	0.113 (0.111,0.115)
	Semi-local EMOS, 10d	0.913 (0.897,0.928)	1.718 (1.690,1.742)	0.108 (0.104,0.112)	0.632 (0.620,0.643)	0.089 (0.086,0.091)
	Local EMOS, 30d	0.869 (0.849,0.887)	1.650 (1.618,1.678)	0.099 (0.096,0.102)	0.604 (0.591,0.615)	0.083 (0.080,0.086)
	Semi-local EMOS, 30d	0.888 (0.867,0.908)	1.681 (1.646,1.713)	0.095 (0.092,0.098)	0.619 (0.606,0.632)	0.083 (0.080,0.085)
Day 5						
(0,50)	Raw ensemble	1.391 (1.369,1.417)	2.425 (2.382,2.471)	0.277 (0.260,0.296)	0.909 (0.896,0.924)	0.243 (0.234,0.253)
	Local EMOS, 10d	1.205 (1.175,1.235)	2.226 (2.173,2.279)	0.166 (0.159,0.176)	0.819 (0.800,0.839)	0.140 (0.135,0.144)
	Semi-local EMOS, 10d	1.142 (1.112,1.171)	2.135 (2.081,2.185)	0.129 (0.123,0.137)	0.794 (0.774,0.813)	0.107 (0.104,0.111)
	Local EMOS, 30d	1.099 (1.070,1.124)	2.070 (2.014,2.121)	0.118 (0.114,0.123)	0.767 (0.746,0.786)	0.102 (0.098,0.105)
	Semi-local EMOS, 30d	1.102 (1.073,1.128)	2.073 (2.018,2.122)	0.110 (0.106,0.116)	0.773 (0.751,0.791)	0.098 (0.095,0.101)
(40,40)	Raw ensemble	1.377 (1.354,1.402)	2.418 (2.373,2.468)	0.253 (0.237,0.270)	0.908 (0.893,0.925)	0.216 (0.209,0.224)
	Local EMOS, 10d	1.226 (1.195,1.256)	2.266 (2.204,2.322)	0.174 (0.167,0.182)	0.831 (0.809,0.852)	0.147 (0.142,0.152)
	Semi-local EMOS, 10d	1.142 (1.111,1.172)	2.138 (2.077,2.191)	0.129 (0.123,0.137)	0.794 (0.771,0.812)	0.108 (0.104,0.112)
	Local EMOS, 30d	1.101 (1.072,1.128)	2.073 (2.019,2.129)	0.119 (0.114,0.124)	0.768 (0.747,0.788)	0.103 (0.099,0.106)
	Semi-local EMOS, 30d	1.100 (1.071,1.127)	2.070 (2.018,2.122)	0.110 (0.106,0.115)	0.771 (0.750,0.790)	0.098 (0.095,0.101)
(120,20)	Raw ensemble	1.403 (1.380,1.429)	2.447 (2.401,2.497)	0.255 (0.238,0.272)	0.921 (0.906,0.937)	0.220 (0.213,0.227)
	Local EMOS, 10d	1.225 (1.195,1.256)	2.263 (2.204,2.321)	0.174 (0.166,0.183)	0.830 (0.809,0.851)	0.147 (0.142,0.152)
	Semi-local EMOS, 10d	1.144 (1.113,1.173)	2.140 (2.080,2.196)	0.129 (0.123,0.137)	0.795 (0.774,0.815)	0.108 (0.104,0.112)
	Local EMOS, 30d	1.101 (1.075,1.128)	2.073 (2.020,2.125)	0.119 (0.114,0.124)	0.768 (0.749,0.787)	0.103 (0.099,0.106)
	Semi-local EMOS, 30d	1.103 (1.075,1.130)	2.073 (2.022,2.123)	0.111 (0.107,0.116)	0.774 (0.754,0.792)	0.098 (0.095,0.101)
(160,10)	Raw ensemble	1.413 (1.391,1.439)	2.459 (2.416,2.509)	0.260 (0.242,0.277)	0.924 (0.910,0.940)	0.225 (0.218,0.233)
	Local EMOS, 10d	1.228 (1.199,1.258)	2.269 (2.208,2.328)	0.176 (0.168,0.185)	0.832 (0.811,0.852)	0.149 (0.144,0.153)
	Semi-local EMOS, 10d	1.147 (1.117,1.176)	2.148 (2.085,2.205)	0.130 (0.123,0.137)	0.797 (0.775,0.816)	0.108 (0.105,0.112)
	Local EMOS, 30d	1.105 (1.077,1.131)	2.079 (2.026,2.132)	0.120 (0.115,0.125)	0.770 (0.750,0.789)	0.103 (0.100,0.107)
	Semi-local EMOS, 30d	1.106 (1.076,1.131)	2.080 (2.027,2.130)	0.111 (0.107,0.117)	0.776 (0.755,0.793)	0.098 (0.095,0.102)
(200,0)	Raw ensemble	1.422 (1.397,1.446)	2.467 (2.421,2.514)	0.270 (0.253,0.287)	0.926 (0.912,0.942)	0.238 (0.231,0.246)
	Local EMOS, 10d	1.211 (1.180,1.242)	2.241 (2.182,2.296)	0.166 (0.158,0.175)	0.825 (0.805,0.845)	0.141 (0.136,0.144)
	Semi-local EMOS, 10d	1.152 (1.122,1.180)	2.152 (2.092,2.205)	0.129 (0.122,0.137)	0.801 (0.780,0.820)	0.109 (0.105,0.113)
	Local EMOS, 30d	1.108 (1.081,1.135)	2.087 (2.028,2.138)	0.118 (0.113,0.123)	0.773 (0.754,0.792)	0.103 (0.100,0.106)
	Semi-local EMOS, 30d	1.112 (1.084,1.139)	2.090 (2.030,2.140)	0.111 (0.107,0.117)	0.780 (0.759,0.798)	0.099 (0.096,0.102)

best until about day 7; however, at longer lead times the other configurations with at least 140 members are equally skilful. Likewise, 50 members at TCo639 are as skilful as 200 members at TCo399 resolution for the raw forecasts while, after calibration, the 200 lower-resolution members are slightly more skilful than the 50 higher-resolution members.

For the small supercomputer scenario, the overall ranking is similar before and after EMOS calibration. At lead times beyond day 7, the skill is mainly determined by the ensemble size and the pure low-resolution ensemble shows the best predictive performance.

The calibration of the ensemble forecasts with EMOS strongly reduces the differences in skill among the equal-cost configurations of single- and dual-resolution ensembles. This implies that making the correct selection of the 'best resolution/ensemble size configuration' is less important for those users who will use EMOS calibrated forecasts instead of the raw forecasts. For direct model output, non-trivial dual-resolution ensemble forecasts of 2 m temperature are considerably more skilful than a single-resolution configuration with the same computational cost. When EMOS calibration is applied, the benefit of a dual-resolution configurations is more marginal than without calibration.

Results were discussed for spatial resolutions up to 18 km that are affordable with computational resources available at ECMWF in 2018. Further work is planned to assess the benefit of dual-resolution ensembles at resolutions up to 9 km that will be affordable in the next 5 years. Moreover, we remark that there is ongoing research at ECMWF on calibration of dual-resolution precipitation forecasts using quantile mapping (Gascón *et al.*, 2019, personal communication).

Finally, a natural question is whether more sophisticated post-processing approaches provide the same answer. This is a possible direction of further research, although some preliminary tests for the large HPC scenario with the computationally far more challenging local BMA approach using 5 days lead time have already been made. In this particular case for the dataset at hand, BMA calibration significantly underperforms EMOS in terms of the mean CRPS, however the CRPS values of the different combinations induce a similar ranking as the corresponding EMOS scores.

ACKNOWLEDGEMENT

Sándor Baran and Marianna Szabó were supported by the project EFOP-3.6.3-VEKOP-16-2017-00002, co-financed by the Hungarian Government and the European Social Fund. Sándor Baran also acknowledges the support of the Hungarian National Research, Development and Innovation Office under grant no. NN125679. He is also grateful to the ECMWF for supporting his research stays in Reading and to the Heidelberg Institute for Theoretical Studies for hosting him as visiting scientist. Last, but not least, the authors are indebted

to the Associate Editor and two anonymous reviewers as their valuable comments greatly improved the quality of this paper.

ORCID

Sándor Baran  <https://orcid.org/0000-0003-1035-004X>

REFERENCES

- Baran, S. (2014) Probabilistic wind speed forecasting using Bayesian model averaging with truncated normal components. *Computational Statistics and Data Analysis*, 75, 227–238.
- Baran, S. and Lerch, S. (2015) Log-normal distribution based EMOS models for probabilistic wind speed forecasting. *Quarterly Journal of the Royal Meteorological Society*, 141, 2289–2299.
- Baran, S. and Lerch, S. (2016) Mixture EMOS model for calibrating ensemble forecasts of wind speed. *Environmetrics*, 27, 116–130.
- Baran, S. and Lerch, S. (2018) Combining predictive distributions for statistical post-processing of ensemble forecasts. *International Journal of Forecasting*, 34, 477–496.
- Baran, S. and Nemoda, D. (2016) Censored and shifted gamma distribution-based EMOS model for probabilistic quantitative precipitation forecasting. *Environmetrics*, 27, 280–292.
- Bentzen, S. and Friederichs, P. (2014) Decomposition and graphical portrayal of the quantile score. *Quarterly Journal of the Royal Meteorological Society*, 140, 1924–1934.
- Buizza, R. (2018) Ensemble forecasting and the need for calibration, Chapter 2 in *Statistical Postprocessing of Ensemble Forecasts*, Vannitsem, S., Wilks, D.S., Messner, J.W. (eds). Elsevier, Amsterdam, Netherlands.
- Buizza, R., Houtekamer, P.L., Toth, Z., Pellerin, G., Wei, M. and Zhu, Y. (2005) A comparison of the ECMWF, MSC, and NCEP global ensemble prediction systems. *Monthly Weather Review*, 133, 1076–1097.
- Casanova, S. and Ahrens, B. (2009) On the weighting of multimodel ensembles in seasonal and short-range weather forecasting. *Monthly Weather Review*, 137, 3811–3822.
- DelSole, T., Yang, X. and Tippett, M. (2013) Is unequal weighting significantly better than equal weighting for multi-model forecasting? *Quarterly Journal of the Royal Meteorological Society*, 139, 176–183.
- Diebold, F.X. and Mariano, R.S. (1995) Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13, 253–263.
- Doblas-Reyes, F., Hagedorn, R. and Palmer, T. (2005) The rationale behind the success of multi-model ensembles in seasonal forecasting II. Calibration and combination. *Tellus A*, 57, 234–252.
- Fraley, C., Raftery, A.E. and Gneiting, T. (2010) Calibrating multimodel forecast ensembles with exchangeable and missing members using Bayesian model averaging. *Monthly Weather Review*, 138, 190–202.
- Friederichs, P. and Thorarindottir, T.L. (2012) Forecast verification for extreme value distributions with an application to probabilistic peak wind prediction. *Environmetrics*, 23, 579–594.
- Gneiting, T. (2011) Making and evaluating point forecasts. *Journal of the American Statistical Association*, 106, 746–762.
- Gneiting, T. (2014) *Calibration of medium-range weather forecasts*. Technical Memorandum No. 719. ECMWF, Reading, UK.
- Gneiting, T., Balabdaoui, F. and Raftery, A.E. (2007) Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society. Series B*, 69, 243–268.
- Gneiting, T. and Raftery, A.E. (2007) Strictly proper scoring rules, prediction and estimation. *Journal of the American Statistical Association*, 102, 359–378.
- Gneiting, T., Raftery, A.E., Westveld, A.H. and Goldman, T. (2005) Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Monthly Weather Review*, 133, 1098–1118.
- Gneiting, T. and Ranjan, R. (2011) Comparing density forecasts using threshold- and quantile-weighted scoring rules. *Journal of Business & Economic Statistics*, 29, 411–422.
- Hemri, S., Scheuerer, M., Pappenberger, F., Bogner, K. and Haiden, T. (2014) Trends in the predictive performance of raw ensemble weather forecasts. *Geophysical Research Letters*, 41, 9197–9205.

- Iversen, T., Deckmyn, A., Santos, C., Sattler, K., Bremnes, J.B., Feddersen, H. and Frogner, I.L. (2011) Evaluation of 'GLAMEPS' – a proposed multimodel EPS for short-range forecasting. *Tellus A*, 63, 513–530.
- Lerch, S. and Baran, S. (2017) Similarity-based semi-local estimation of EMOS models. *Journal of the Royal Statistical Society. Series C*, 66, 29–51.
- Lerch, S. and Thorarinsdottir, T.L. (2013) Comparison of non-homogeneous regression models for probabilistic wind speed forecasting. *Tellus A*, 65, 21206.
- Leutbecher, M. (2018) Ensemble size: How suboptimal is less than infinity?. *Quarterly Journal of the Royal Meteorological Society*. <https://doi.org/10.1002/qj.3387>.
- Politis, D.N. and Romano, J.P. (1994) The stationary bootstrap. *Journal of the American Statistical Association*, 89, 1303–1313.
- Raftery, A.E., Gneiting, T., Balabdaoui, F. and Polakowski, M. (2005) Using Bayesian model averaging to calibrate forecast ensembles. *Monthly Weather Review*, 133, 1155–1174.
- Raynaud, L. and Bouttier, F. (2017) The impact of horizontal resolution and ensemble size for convective-scale probabilistic forecasts. *Quarterly Journal of the Royal Meteorological Society*, 143, 3037–3047.
- Raynaud, L., Pannekoucke, O., Arbogast, P. and Bouttier, F. (2015) Application of a Bayesian weighting for short-range lagged ensemble forecasting at the convective scale. *Quarterly Journal of the Royal Meteorological Society*, 141, 459–468.
- Ruiz, J.J. and Saulo, C. (2012) How sensitive are probabilistic precipitation forecasts to the choice of calibration algorithms and the ensemble generation method? Part I: Sensitivity to calibration methods. *Meteorological Applications*, 19, 302–313.
- Scheuerer, M. (2014) Probabilistic quantitative precipitation forecasting using ensemble model output statistics. *Quarterly Journal of the Royal Meteorological Society*, 140, 1086–1096.
- Scheuerer, M. and Hamill, T.M. (2015) Statistical post-processing of ensemble precipitation forecasts by fitting censored, shifted gamma distributions. *Monthly Weather Review*, 143, 4578–4596.
- Schmeits, M.J. and Kok, K.J. (2010) A comparison between raw ensemble output, (modified) Bayesian model averaging and extended logistic regression using ECMWF ensemble precipitation reforecasts. *Monthly Weather Review*, 138, 4199–4211.
- Sloughter, J.M., Gneiting, T. and Raftery, A.E. (2010) Probabilistic wind speed forecasting using ensembles and Bayesian model averaging. *Journal of the American Statistical Association*, 105, 25–37.
- Sloughter, J.M., Raftery, A.E., Gneiting, T. and Fraley, C. (2007) Probabilistic quantitative precipitation forecasting using Bayesian model averaging. *Monthly Weather Review*, 135, 3209–3220.
- Swinbank, R., Kyouda, M., Buchanan, P., Froude, L., Hamill, T.M., Hewson, T.D., Keller, J.H., Matsueda, M., Methven, J., Pappenberger, F., Scheuerer, M., Tittley, H.A., Wilson, L. and Yamaguchi, M. (2016) The TIGGE project and its achievements. *Bulletin of the American Meteorological Society*, 97, 49–67.
- Thorarinsdottir, T.L. and Gneiting, T. (2010) Probabilistic forecasts of wind speed: ensemble model output statistics by using heteroscedastic censored regression. *Journal of the Royal Statistical Society. Series A*, 173, 371–388.
- Wilks, D.S. (2011) *Statistical Methods in the Atmospheric Sciences*, 3rd edition. Elsevier, Amsterdam, Netherlands.
- Williams, R.M., Ferro, C.A.T. and Kwasniok, F. (2014) A comparison of ensemble post-processing methods for extreme events. *Quarterly Journal of the Royal Meteorological Society*, 140, 1112–1120.

How to cite this article: Baran S, Leutbecher M, Szabó M, Ben Bouallègue Z. Statistical post-processing of dual-resolution ensemble forecasts. *Q J R Meteorol Soc.* 2019;1–16. <https://doi.org/10.1002/qj.3521>