

# Combining predictive distributions for statistical post-processing of ensemble forecasts

Sándor Baran<sup>a,\*</sup>, Sebastian Lerch<sup>b,c</sup>

<sup>a</sup>*Faculty of Informatics, University of Debrecen, Kassai út 26, H-4028 Debrecen, Hungary*

<sup>b</sup>*Heidelberg Institute for Theoretical Studies, Schloss-Wolfsbrunnengasse 35, D-69118 Heidelberg, Germany*

<sup>c</sup>*Institute for Stochastics, Karlsruhe Institute of Technology, Englerstraße 2, D-76128 Karlsruhe, Germany*

---

## Abstract

Statistical post-processing techniques are now widely used to correct systematic biases and errors in calibration of ensemble forecasts obtained from multiple runs of numerical weather prediction models. A standard approach is the ensemble model output statistics (EMOS) method resulting in a predictive distribution given by a single parametric law with parameters depending on the ensemble members. In this article we assess the merits of combining multiple EMOS models based on different parametric families. In four case studies with wind speed and precipitation forecasts from two ensemble prediction systems we investigate the performance of state of the art forecast combination methods and propose a computationally efficient approach to determining linear pool combination weights. We study how forecast combination performs compared to the theoretically superior but cumbersome estimation of a full mixture model and assess which degree of flexibility of the forecast combination approach yields the best practical results for post-processing applications.

*Keywords:* Combining forecasts, comparative studies, density forecasts, ensemble model output statistics, precipitation, weather forecasting, wind speed.

---

## 1. Introduction

Nowadays, weather forecasts are typically based on the output of numerical weather prediction (NWP) models which describe the physical behavior of the atmosphere through nonlinear partial differential equations. Single deterministic predictions produced by single runs of such models fail to account for uncertainties in the initial conditions and the numerical model. Therefore, NWP models are nowadays typically run several times with varying initial conditions and model physics, resulting in an ensemble of forecasts, see Palmer (2002) and Gneiting & Raftery (2005) for reviews. Examples of ensemble prediction systems (EPSs) are the 51-member European Centre for Medium-Range Weather Forecasts (ECMWF) ensemble (Molteni et al.,

---

\*Corresponding author. Tel: +36 52 512900 75212, fax: +36 52 512996.

*Email address:* `baran.sandor@inf.unideb.hu` (Sándor Baran)

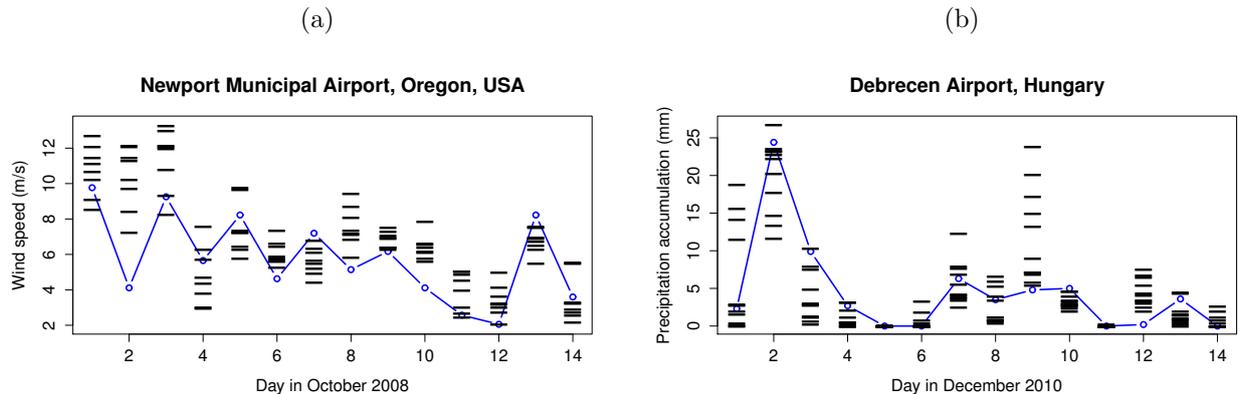


Figure 1: (a) Wind speed observations (blue line) and corresponding UWME forecasts (bars) for Newport Municipal Airport, Oregon, USA, for the first two weeks of October 2008; (b) observed precipitation accumulation (blue line) and the corresponding ALADIN-HUNEPS ensemble forecasts (bars) for Debrecen Airport, Hungary, for the first two weeks of December 2010.

1996), the eight-member University of Washington Mesoscale ensemble (UWME; Eckel & Mass, 2005), and the 11-member Aire Limitée Adaptation dynamique Développement International-Hungary Ensemble Prediction System (ALADIN-HUNEPS; Horányi et al., 2006) of the Hungarian Meteorological Service (HMS). The transition from single deterministic forecasts to ensemble predictions can be seen as an important step towards probabilistic forecasting, however, ensemble forecasts are often underdispersive, that is, the spread of the ensemble is too small to account for the full uncertainty, and subject to systematic bias. They thus require some form of statistical post-processing. To illustrate the systematic errors of ensemble forecasts, Figure 1(a) shows UWME wind speed forecasts for Newport Municipal Airport (OR) and the corresponding observations for the first two weeks of October 2008, and Figure 1(b) shows ALADIN-HUNEPS forecasts of precipitation accumulation at Debrecen Airport and the corresponding observations for the first two weeks of December 2010. Both time series illustrate the lack of an appropriate representation of the forecast uncertainty as the verifying observations frequently fall outside the range of the ensemble forecasts.

Over the past decade, various statistical post-processing methods have been proposed in the meteorological literature. In the Bayesian model averaging (BMA; Raftery et al., 2005) approach the forecast distribution is given by a weighted mixture of parametric densities, each of which depends on a single ensemble member with mixture weights being determined by the performance of the ensemble members in the training period. Within this article we build on the conceptually simpler ensemble model output statistics (EMOS) approach proposed by Gneiting et al. (2005), where the conditional distribution of the weather variable of interest given the ensemble predictions is modeled by a single parametric family. The parameters of the forecast distribution are connected to the ensemble forecast through suitable link functions. For example, the original EMOS approach models temperature with a Gaussian predictive distribution where the mean is an affine function of the ensemble member forecasts and the variance is an affine function of the

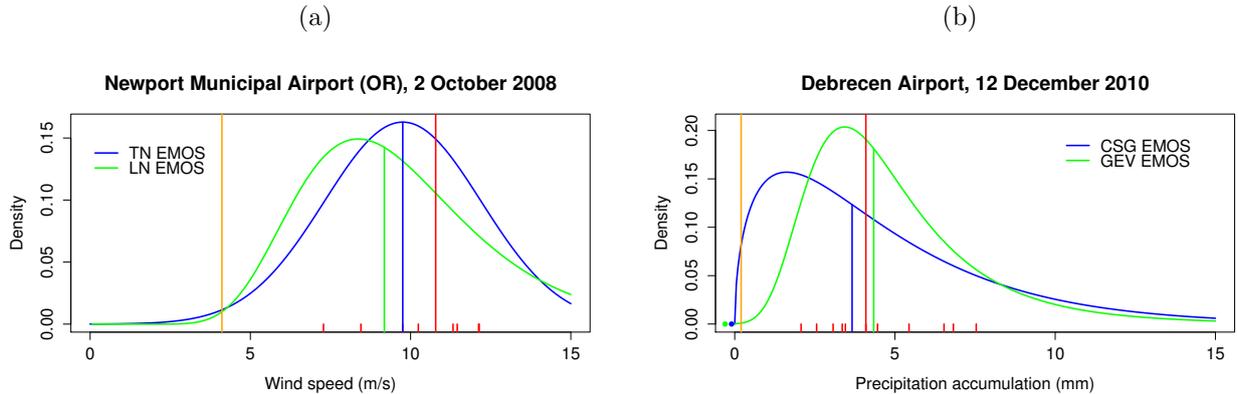


Figure 2: (a) Wind speed observations, the corresponding UWME forecasts and TN and LN EMOS predictive distributions for Newport Municipal Airport (OR) for 2 October 2008; (b) observed precipitation accumulation, the corresponding ALADIN-HUNEPS ensemble forecasts and CSG and GEV EMOS predictive distributions for Debrecen Airport for 12 December 2010. Ensemble members: red bars; ensemble median: vertical red line; observation: vertical orange line; predictive PDFs: blue/green lines; EMOS medians: vertical blue/green lines.

ensemble variance.

Over the last years the EMOS approach has been extended to other weather variables such as wind speed (Thorarinsdottir & Gneiting, 2010; Lerch & Thorarinsdottir, 2013; Baran & Lerch, 2015; Scheuerer & Möller, 2015), precipitation (Scheuerer, 2014; Scheuerer & Hamill, 2015; Baran & Nemoda, 2016), and total cloud cover (Hemri et al., 2016). To illustrate the EMOS approach to post-processing Figure 2(a) shows the observed wind speed, the corresponding UWME forecasts and truncated normal (TN) and log-normal (LN) EMOS predictive distributions (for details see Section 3.1) for Newport Municipal Airport for 2 October 2008. A different situation is shown in Figure 2(b) where the observed precipitation accumulation, the corresponding ALADIN-HUNEPS ensemble forecasts and estimated censored and shifted gamma (CSG) and censored generalized extreme value (GEV) EMOS predictive distributions (see Section 3.2) for Debrecen Airport for 12 December 2010 are plotted. In both examples, the spread of the ensemble forecasts is notably smaller than the spread of the post-processed forecast distribution.

The success of statistical post-processing relies on finding appropriate parametric families for the weather variable of interest. However, the choice of a suitable parametric model is a non-trivial task and often a multitude of competing models is available. The relative performance of these models usually varies for different data sets and applications.

Regime-switching combination models proposed by Lerch & Thorarinsdottir (2013) partly alleviate the limited flexibility of single parametric family models by selecting one of several candidate models based on covariate information. However, the applicability of this approach is subject to the availability of suitable covariates. For some weather variables, full mixture EMOS models can be formulated where the parameters

and weights of a mixture of two forecast distributions are jointly estimated (Baran & Lerch, 2016). However, such approaches are limited to specific weather variables, and the estimation is computationally demanding.

In this article we investigate the feasibility of an alternative, more generally applicable route towards improving the forecast performance that has recently received some interest, and has for example been called for in Yang et al. (2017). Motivated by recent promising results of Möller & Groß (2016) and Bassetti et al. (2017), we study whether combining predictive distributions from individual post-processing models is able to significantly improve the forecast performance. In a first step, individual EMOS models based on single parametric distributions are estimated. In a second step the forecast distributions are combined utilizing state of the art forecast combination techniques such as the (spread-adjusted) linear pool, the beta-transformed linear pool (Gneiting & Ranjan, 2013), and a recently proposed Bayesian, essentially non-parametric calibration approach (Bassetti et al., 2017). Further, we propose a computationally efficient 'plug-in' approach to determining combination weights in the linear pool that is specific to post-processing applications.

The main contribution of this article is an empirical assessment of the merits of combining forecast distributions from post-processing models. Specifically, we investigate how forecast combination performs compared to the theoretically superior but cumbersome estimation of a full mixture model. Secondly, we assess which degree of flexibility of the forecast combination approach yields the best practical results for post-processing applications.

An alternative approach to post-processing that circumvents the problem of choosing suitable parametric forecast distributions is the use of non-parametric methods, see for example Hamill & Whitaker (2006); Flowerdew (2014), and Taillardat et al. (2016). However, these approaches suffer from the limitation that the support of the forecast distribution is restricted to the range of observed values in the training sets. Further, these methods require sufficiently long training periods, and generally lead to high computational costs.

The remainder of this article is organized as follows. Section 2 contains a description of the ensemble systems and the observation data. In Section 3, the EMOS method is reviewed, and the individual EMOS models for wind speed and precipitation are introduced. Thereafter, Section 4 provides a description of the forecast combination approaches and the application to post-processing. The various EMOS models and forecast combination approaches are compared in four case studies in Section 5. The article concludes with a discussion in Section 6.

## 2. Data

We consider two different weather variables, wind speed and precipitation accumulation, and two distinct data sets of ensemble forecasts and corresponding validating observations for each weather quantity. The

wind speed data sets are identical to data used in Baran & Lerch (2015, 2016), whereas the precipitation data coincide with those studied in Baran & Nemoda (2016). For detailed descriptions of the ensemble forecasts and corresponding observations we refer to these articles and references therein.

Ensemble members that are generated with the help of random perturbations of initial conditions are statistically indistinguishable, and are referred to as exchangeable. The notion of exchangeability of ensemble members is important for the formulation of post-processing models, see Section 3 for details.

### *2.1. University of Washington mesoscale ensemble*

The UWME covers the Pacific Northwest region of North America with a horizontal resolution of 12 km and consists of eight members generated from different runs of the fifth generation Pennsylvania State University–National Center for Atmospheric Research mesoscale model (Grell et al., 1995). The initial and boundary conditions of the model runs are provided by different sources, the individual ensemble members can therefore be clearly distinguished and are considered to be non-exchangeable. The data set at hand contains 48 h ahead forecasts and corresponding validating observations for 10 m maximal wind speed (given in m/s) and 24 h precipitation accumulation (given in mm) for 152 stations in the Automated Surface Observing Network (National Weather Service, 1998) in the U.S. states of Washington, Oregon, Idaho, California and Nevada.

We focus on calendar year 2008 with additional forecasts and observations from the last months of 2007 used to allow for training periods of equal length for the model estimation. After removing days and locations with missing predictions and/or observations, stations where data are available only on very few days are also removed resulting in 101 stations with 27 481 individual forecast cases for wind speed and 75 stations with 20 448 individual forecast cases for precipitation.

### *2.2. ALADIN-HUNEPS ensemble*

The ALADIN-HUNEPS system covers large parts of continental Europe on an 8 km grid. It is obtained with dynamical downscaling of the global ARPEGE based PEARP system of Météo France (Horányi et al., 2011; Descamps et al, 2015) and consists of 11 ensemble members, 10 of which are exchangeable forecasts from perturbed initial conditions, and one of which is a control member from the unperturbed analysis.

We use ensembles of 42 h ahead forecasts of 10 m instantaneous wind speed (in m/s) and 24 h precipitation accumulation (in mm) issued for 10 major cities in Hungary together with the corresponding validation observations. Wind speed data are available for a one-year period from 1 April 2012 to 31 March 2013, and precipitation data are available between 1 October 2010 and 25 March 2011. Days with missing forecasts and/or observations are excluded from the analysis for both wind speed (6 days) and precipitation (2 days).

### 3. Ensemble model output statistics

Successful statistical post-processing of ensemble forecasts relies on finding and estimating appropriate parametric models for the conditional distribution of the weather variable of interest given the ensemble predictions. In case of the EMOS approach, the forecast distribution is given by a single parametric law with parameters depending on the ensemble forecast. While temperature can be modeled by a normal distribution (Gneiting et al., 2005), the choice of a suitable parametric family is much less straightforward for weather variables such as wind speed or precipitation. A multitude of post-processing approaches and modeling strategies has been proposed over the last years. In the following short review, we focus on EMOS models for wind speed and precipitation, and subsequently investigate methods to combine forecast distributions from different models.

#### 3.1. EMOS models for wind speed

Non-negative weather variables such as wind speed require skewed predictive distributions with non-negative support like Weibull (Justus et al., 1978) or gamma distributions (Garcia et al., 1988). Recently developed EMOS approaches utilize truncated normal (TN; Thorarinsdottir & Gneiting, 2010), gamma (Scheuerer & Möller, 2015), generalized extreme value (GEV; Lerch & Thorarinsdottir, 2013) and log-normal (LN; Baran & Lerch, 2015) distributions to model the conditional distribution of wind speed given the ensemble predictions. Here, we focus on the truncated normal and log-normal models.

##### 3.1.1. Truncated normal EMOS model

We denote by  $\mathcal{N}_0(\mu, \sigma^2)$  the TN distribution with location  $\mu$ , scale  $\sigma > 0$ , and cut-off at zero with probability density function (PDF)

$$g(x|\mu, \sigma) := \frac{\frac{1}{\sigma}\varphi((x-\mu)/\sigma)}{\Phi(\mu/\sigma)}, \quad \text{if } x \geq 0, \quad \text{and} \quad g(x|\mu, \sigma) := 0, \quad \text{otherwise,}$$

where  $\varphi$  and  $\Phi$  are the PDF and the cumulative distribution function (CDF) of the standard normal distribution, respectively. The predictive distribution of the EMOS model proposed by Thorarinsdottir & Gneiting (2010) is

$$\mathcal{N}_0(a_0 + a_1 f_1 + \dots + a_K f_K, b_0 + b_1 S^2) \quad \text{with} \quad S^2 := \frac{1}{K-1} \sum_{k=1}^K (f_k - \bar{f})^2, \quad (3.1)$$

where  $f_1, f_2, \dots, f_K$  denote the ensemble of distinguishable forecasts of wind speed for a given location and time, and  $\bar{f}$  is the ensemble mean. Location parameters  $a_0, a_1, \dots, a_K \in \mathbb{R}$  and scale parameters  $b_0 \in \mathbb{R}$ ,  $b_1 \geq 0$  of model (3.1) can be estimated from the training data, consisting of ensemble members and verifying observations from the preceding  $n$  days, by optimizing an appropriate verification score, see Section 3.3.

However, most operational EPSs generate forecasts using random perturbations of the initial conditions resulting in statistically indistinguishable ensemble members which are referred to as exchangeable. Examples include the 51-member ECMWF ensemble, as well as sub-ensembles of forecasts from single models that form groups of exchangeable members within multi-model EPSs such as the THORPEX Interactive Grand Global Ensemble (Swinbank et al., 2016) or the GLAMEPS ensemble (Iversen et al., 2011). To account for the generation of the forecasts, ensemble members within a given group of exchangeable members should share the same coefficients in the post-processing model (Fraley et al., 2010; Gneiting, 2014).

To formalize this notion, a generalized formulation of model (3.1) for the case of  $M$  ensemble members divided into  $K$  groups, where the  $k$ th group contains  $M_k \geq 1$  exchangeable ensemble members ( $\sum_{k=1}^K M_k = M$ ) introduced in Baran & Lerch (2015) is given by

$$\mathcal{N}_0 \left( a_0 + a_1 \sum_{\ell_1=1}^{M_1} f_{1,\ell_1} + \cdots + a_K \sum_{\ell_K=1}^{M_K} f_{K,\ell_K}, b_0 + b_1 S^2 \right).$$

Analogous concepts apply to all EMOS models discussed in the subsequent sections.

### 3.1.2. Log-normal EMOS model

As an alternative to the TN EMOS model, Baran & Lerch (2015) introduce an EMOS approach based on log-normal forecast distributions where the mean  $m$  and variance  $v$  of the predictive distribution are linked to the ensemble members as

$$m = \alpha_0 + \alpha_1 f_1 + \cdots + \alpha_K f_K \quad \text{and} \quad v = \beta_0 + \beta_1 S^2. \quad (3.2)$$

These quantities uniquely determine the location  $\mu$  and shape  $\sigma > 0$  of the underlying LN distribution  $\mathcal{LN}(\mu, \sigma)$  with PDF

$$h(x|\mu, \sigma) := \frac{1}{x\sigma} \varphi((\log x - \mu)/\sigma), \quad \text{if } x \geq 0, \quad \text{and} \quad h(x|\mu, \sigma) := 0, \quad \text{otherwise,}$$

via transformations

$$\mu = \log \left( \frac{m^2}{\sqrt{v + m^2}} \right) \quad \text{and} \quad \sigma = \sqrt{\log \left( 1 + \frac{v}{m^2} \right)}.$$

Similar to the TN EMOS model, estimates of parameters  $\alpha_0, \alpha_1, \dots, \alpha_K \in \mathbb{R}$  and  $\beta_0 \in \mathbb{R}, \beta_1 \geq 0$  are obtained by optimizing the mean of an appropriate verification score over all forecast cases in the training data.

### 3.1.3. Combination and mixture models

The TN and LN models described above model the conditional distribution of wind speed given the ensemble predictions with a single parametric forecast distribution. This approach relies on the choice of a suitable parametric family, and limits the flexibility of the model. For instance, it can be demonstrated that

the heavier tails of the LN model are more appropriate for modeling higher wind speeds in the right tail of the distribution, whereas the TN model is more appropriate for the bulk of the distribution, see Baran & Lerch (2016) for details.

Therefore, different combination and mixture models have been proposed in the literature. In the regime-switching combination approach (Lerch & Thorarinsdottir, 2013; Baran & Lerch, 2015) one of the candidate models is selected based on covariate information with suitably adapted parameter estimation procedures. For example, a TN model can be used if the median ensemble forecast is below a threshold  $\eta$ , and an LN model is used in case of median ensemble forecasts exceeding this threshold. Such combination models have been demonstrated to improve the predictive performance compared to the individual models, however, they require the choice of a suitable covariate, and the threshold parameter  $\eta$  has to be determined by repeating the model estimation over a grid of potential values, thereby limiting the flexibility and increasing the computational cost of such approaches.

In order to flexibly combine the advantages of lighter and heavier-tailed distributions and to avoid these problems in the process, Baran & Lerch (2016) propose a mixture model of the form

$$\psi(x|\mu_{TN}, \sigma_{TN}; \mu_{LN}, \sigma_{LN}; \omega) := \omega g(x|\mu_{TN}, \sigma_{TN}) + (1 - \omega)h(x|\mu_{LN}, \sigma_{LN}), \quad (3.3)$$

where the parameters of the component distributions  $g$  and  $h$  depend on the ensemble forecasts as specified in (3.1) and (3.2). The EMOS coefficients and the weight  $\omega \in [0, 1]$  of the mixture model (3.3) are estimated jointly using optimum score approaches. This mixture model approach results in significantly improved calibration (Baran & Lerch, 2016), however, it is computationally demanding and hinders using standard optimum score estimation based on the continuous ranked probability score due to the lack of an analytic expression of the objective function, see Section 3.3 for details. Similar mixture models where the different component distributions focus on specific regions of interest such as the bulk and the tail above a threshold value have been tested, but models based on truncated normal and generalized Pareto distributions result in worse predictive performance, see Baran & Lerch (2016).

In contrast to the joint estimation of all parameters in (3.3), the forecast combination approaches introduced in Section 4 are two-step procedures where in a first step, EMOS models based on a single parametric family are estimated, and in a second step, these models are combined as a weighted mixture by estimating an appropriate weight. In Section 5 the full mixture model (3.3) is used as a benchmark, whereas the regime-switching combination approach will not be considered any further.

### 3.2. EMOS models for precipitation

The discrete-continuous nature of precipitation accumulation requires a non-negative predictive distribution assigning positive mass to the event of zero precipitation. A popular choice is to consider a continuous distribution that can take both positive and negative values and left-censor it at zero (Scheuerer, 2014;

Scheuerer & Hamill, 2015; Baran & Nemoda, 2016), which thereby assigns the mass of negative values to zero precipitation accumulation.

### 3.2.1. Censored and shifted gamma EMOS model

Let  $G(\cdot|\kappa, \theta)$  denote the CDF of the  $\Gamma(\kappa, \theta)$  distribution with shape  $\kappa > 0$  and scale  $\theta > 0$  and let  $\delta > 0$ . Then the CDF of the shifted gamma distribution left censored at zero (CSG)  $\Gamma_0(\kappa, \theta, \delta)$  with shape  $\kappa$ , scale  $\theta$  and shift  $\delta$  is given by

$$G_0(x|\kappa, \theta, \delta) := G(x + \delta|\kappa, \theta), \quad \text{if } x \geq 0, \quad \text{and} \quad G_0(x|\kappa, \theta, \delta) := 0, \quad \text{otherwise,} \quad (3.4)$$

that is, mass  $G(\delta|\kappa, \theta)$  is assigned to the origin. In the CSG EMOS approach of Baran & Nemoda (2016) the mean  $m = \kappa\theta$  and variance  $\sigma^2 = \kappa\theta^2$  of the uncensored gamma distribution  $\Gamma(\kappa, \theta)$  are affine functions of the ensemble and ensemble mean, respectively, that is

$$m = a_0 + a_1 f_1 + \dots + a_K f_K \quad \text{and} \quad \sigma^2 = b_0 + b_1 \bar{f}. \quad (3.5)$$

### 3.2.2. Censored generalized extreme value EMOS model

The CDF of a GEV distribution  $\mathcal{G}\mathcal{E}\mathcal{V}(\mu, \sigma, \xi)$  with location  $\mu$ , scale  $\sigma > 0$  and shape  $\xi$  equals

$$H(x|\mu, \sigma, \xi) := \begin{cases} \exp\left(-\left[1 + \xi\left(\frac{x-\mu}{\sigma}\right)\right]^{-1/\xi}\right), & \xi \neq 0; \\ \exp\left(-\exp\left(-\frac{x-\mu}{\sigma}\right)\right), & \xi = 0, \end{cases} \quad \text{if } 1 + \xi(x - \mu)/\sigma > 0,$$

and zero otherwise, which for  $-0.278 < \xi < 1$  has a positive skewness and an existing mean

$$m = \begin{cases} \mu + \sigma \frac{\Gamma(1-\xi)-1}{\xi}, & \xi \neq 0; \\ \mu + \sigma\gamma, & \xi = 0, \end{cases}$$

where  $\gamma$  denotes the Euler-Mascheroni constant.

The EMOS model for precipitation accumulation proposed by Scheuerer (2014) is based on a censored GEV distribution  $\mathcal{G}\mathcal{E}\mathcal{V}_0(\mu, \sigma, \xi)$  with CDF

$$H_0(x|\mu, \sigma, \xi) = H(x|\mu, \sigma, \xi), \quad \text{if } x \geq 0, \quad \text{and} \quad H_0(x|\mu, \sigma, \xi) := 0, \quad \text{otherwise,} \quad (3.6)$$

where

$$m = \alpha_0 + \alpha_1 f_1 + \dots + \alpha_K f_K + \nu p_0 \quad \text{and} \quad \sigma = \beta_0 + \beta_1 \text{MD}(f), \quad (3.7)$$

with

$$p_0 := \frac{1}{K} \sum_{k=1}^K \mathbb{1}_{\{f_k=0\}} \quad \text{and} \quad \text{MD}(f) := \frac{1}{K^2} \sum_{k,\ell=1}^K |f_k - f_\ell|,$$

where  $\mathbb{1}_A$  denotes the indicator function of the set  $A$ .

### 3.2.3. Mixture models

Similar to wind speed, general mixture models with CSG and GEV component distributions of the form

$$\varrho(x|\kappa, \theta, \delta; \mu, \sigma, \xi; \omega) := \omega g_0(x|\kappa, \theta, \delta) + (1 - \omega)h_0(x|\mu, \sigma, \xi), \quad (3.8)$$

can be formulated, where  $g_0(\cdot|\kappa, \theta, \delta)$  and  $h_0(\cdot|\mu, \sigma, \xi)$  denote the generalized PDFs of the CSG and censored GEV distributions, respectively, and the dependence of parameters  $\kappa, \theta$  and  $\mu, \sigma$  on the ensemble is given by (3.5) and (3.7).

However, joint optimum score estimation of the parameters is more involved than in the case of the TN-LN mixture model (3.3) for wind speed due to the larger number of parameters and the discrete-continuous nature of the forecast distribution  $\varrho$ . As initial tests with the precipitation data sets introduced in Section 2 indicated problematic behavior of the numerical optimization algorithms potentially caused by the non-smooth dependence of the objective functions on the parameters, we do not pursue this approach any further and only consider the forecast combination approaches introduced in Section 4. Compared to jointly estimating all parameters, these methods separate the estimation into two steps and thereby result in more stable optimization problems.

### 3.3. Forecast evaluation and parameter estimation

In probabilistic forecasting the general aim is to maximize the sharpness of the predictive distribution subject to calibration (Gneiting et al., 2007). Calibration refers to the statistical consistency between the forecast and the observation, and given that the predictive distribution is calibrated, it should be as concentrated (or sharp) as possible. Calibration and sharpness can be assessed simultaneously with the help of proper scoring rules.

Proper scoring rules are loss functions that assign a numerical value to pairs of forecasts and observations. In the atmospheric sciences the most popular scoring rules are the continuous ranked probability score (CRPS; Matheson & Winkler, 1976; Gneiting & Raftery, 2007) and the logarithmic score (LogS; Good, 1952). Given a predictive CDF  $F(y)$  and an observation  $x$ , the CRPS is defined as

$$\begin{aligned} \text{CRPS}(F, x) &:= \int_{-\infty}^{\infty} (F(y) - \mathbb{1}_{\{y \geq x\}})^2 dy \\ &= \int_{-\infty}^x F^2(y) dy + \int_x^{\infty} (1 - F(y))^2 dy \\ &= \mathbf{E}|X - x| - \frac{1}{2} \mathbf{E}|X - X'|, \end{aligned} \quad (3.9)$$

where  $X$  and  $X'$  are independent random variables with CDF  $F$  and finite first moment. The last representation in (3.9) implies that the CRPS can be expressed in the same unit as the observation. The logarithmic score is the negative logarithm of the predictive density  $f(y)$  evaluated at the verifying observation, i.e.,

$$\text{LogS}(F, x) := -\log(f(x)).$$

Both CRPS and LogS are proper scoring rules (Gneiting & Raftery, 2007) which are negatively oriented, that is, smaller scores indicate better forecasts.

Proper scoring rules provide valuable tools for the estimation of model parameters. Following the general optimum score estimation approach of Gneiting & Raftery (2007), the parameters of a predictive distribution can be determined by optimizing the average value of a proper scoring rule as a function of the parameters over a suitably chosen training set. Optimum score estimation based on minimizing the LogS then corresponds to classical maximum likelihood (ML) estimation. If closed form expressions of the integral in (3.9) are available, minimum CRPS estimation, i.e. optimum score estimation based on minimizing the mean CRPS, often provides a valuable, more robust alternative to ML estimation (see e.g. Gneiting et al., 2005).

Analytic expressions of the CRPS are available for all individual EMOS models for wind speed and precipitation introduced in Sections 3.1 and 3.2, thereby allowing for efficient parameter estimation procedures by minimizing the mean CRPS over the forecast cases in the training periods. The closed form solutions are provided in the corresponding articles (Thorarinsdottir & Gneiting, 2010; Baran & Lerch, 2015; Scheuerer, 2014; Scheuerer & Hamill, 2015). Implementations for the statistical programming language R (R Core Team, 2017) are for example available in the `scoringRules` package (Jordan et al., 2017). The parameter estimation for the EMOS models is performed using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm (Press et al., 2007, Section 10.9) implemented in the `optim` function in R. In the case of precipitation we use a constrained version of the BFGS algorithm to ensure positivity of the EMOS coefficients. See Section 5 for details on the selection of the training sets over which the parameters are estimated.

By contrast, the CRPS is not available in closed form for the mixture models (3.3) and (3.8), or any of the forecast combination models introduced in Section 4. Therefore, each step of the optimization procedure requires numerical integration resulting in high computational costs. In case of the mixture model (3.3) for wind speed, we instead use ML estimation of the parameters. The forecast combination approaches introduced below partly alleviate this issue by separating the parameter estimation into two steps rather than estimating all parameters jointly.

#### 4. Forecast combination methods and application to statistical post-processing

We now describe state of the art methods for combining predictive distributions, which we will employ in a post-processing context. The combination approaches constitute two-step methods. The first step is given by the estimation of component models in the form of EMOS models based on suitable single parametric families. In a second step, the component models are combined by estimating the mixture weight and possibly more combination parameters. Compared to the previously discussed mixture model (3.3), the two-step approaches reduce the dimensionality of the optimization problem. Further, the combination approaches can be flexibly applied to any weather variable of interest given that suitable component models

are available, the model formulation is thus given in a general form.

Let  $G(x|\mathbf{f};\nu)$  and  $H(x|\mathbf{f};\theta)$  be predictive CDFs belonging to two different families of distributions depending on the ensemble  $\mathbf{f}$  via parameter vectors  $\nu$  and  $\theta$ , respectively. The EMOS models introduced in Sections 3.1 and 3.2 are later used as component distributions for wind speed and precipitation, respectively.

#### 4.1. Linear pool and spread-adjusted linear pool

We start by introducing the linear and spread-adjusted linear pool of forecast distributions which have been applied to post-processing ensemble forecasts by Möller & Groß (2016). The classical linear pool (LP) employs a mixture model with a predictive CDF of the general form

$$F^{\text{LP}}(x|\mathbf{f};\nu,\theta,\omega) := \omega G(x|\mathbf{f};\nu) + (1-\omega)H(x|\mathbf{f};\theta), \quad \omega \in [0,1]. \quad (4.1)$$

As a special case we also consider the equally weighted linear pool (EWLP) for which we simply set  $\omega = 1/2$ . Since this combination approach is virtually free of computational costs, it can serve as a convenient benchmark.

As demonstrated by Gneiting & Ranjan (2013), linear pooling of predictive distributions increases the dispersion of the forecasts. They propose spread-adjusted and beta-transformed linear pooling approaches that allow to correct for this deficiency. The spread-adjusted linear pool (SLP) results in a predictive distribution

$$F^{\text{SLP}}(x|\mathbf{f};\nu,\theta,\omega) := \omega G\left(\frac{x}{c}|\mathbf{f};\nu\right) + (1-\omega)H\left(\frac{x}{c}|\mathbf{f};\theta\right), \quad \omega \in [0,1] \quad (4.2)$$

with spread adjustment parameter  $c > 0$ . The linear pool is obtained for  $c = 1$ .

As noted by Gneiting & Ranjan (2013), the forecasts of the Bayesian model averaging approach of Raftery et al. (2005) take a similar functional form, but differ in that the combination parameters and the parameters of the individual component distributions are estimated simultaneously. Further, the forecast distributions of the mixture components in the BMA approach depend on a single ensemble member only, whereas the EMOS predictive distributions used here depend on the entire ensemble through suitable link functions with coefficients  $\nu, \theta$ .

The weight  $\omega \in [0,1]$  and the spread adjustment parameter  $c > 0$  have to be estimated from past forecast cases. Note that these need to be training samples where post-processed forecast distributions are available. Möller & Groß (2016) suggest to choose sets of candidate parameter values, e.g.  $\omega \in \{0, 0.05, \dots, 0.95, 1\}$  and  $c \in \{0.7, 0.75, \dots, 1.25, 1.3\}$ , and to apply the combination formulas (4.1) and (4.2) for all possible parameter combinations in order to select those parameter values corresponding to the lowest mean CRPS in the training sample. The CRPS of the forecast distributions is thereby computed using numerical integration. However, as tests indicated improvements in the predictive performance and lower computational costs, we instead determine the optimal parameter values by numerical optimization

with the CRPS as a target functional. To allow for a direct comparison with the component models, we use forecast-observation pairs from the same training sets that were used to estimate the EMOS coefficients. Various alternative approaches to estimate the combination weight in case of the linear pool have been proposed in the literature, including approaches based on other scoring rules such as the LogS (Hall & Mitchell, 2007) or weighted scoring rules (Opschoor et al., 2017), as well as Bayesian approaches (Billio et al., 2013; Del Negro et al, 2016).

#### 4.1.1. A plug-in variant of the linear pool for post-processing applications

In the following we propose a simple plug-in variant to determine the weight parameter in the linear pool that only requires a single numerical integration step rather than repeated numerical integration during the optimization procedure. In tables and figures of Section 5, this approach is abbreviated by LP-PI.

Let  $G(x|\mathbf{f};\nu)$  and  $H(x|\mathbf{f};\theta)$  be predictive CDFs as above, and let  $(x_k, \mathbf{f}_k)$ ,  $k = 1, 2, \dots, n$ , denote the pairs of verifying observations and ensemble forecasts in the training data. The basic idea of the proposed plug-in variant of the linear pool is to utilize the current EMOS parameters estimated for day  $n + \tau$  where  $\tau$  is the forecast horizon to compute the parameters of the corresponding component distributions over the entire training sample. In contrast to utilizing the respective EMOS coefficient vectors  $(\nu_k)$  and  $(\theta_k)$  for  $k = 1, 2, \dots, n$ , this reduces the number of required numerical integrations at the cost of not using the specific parameter values estimated for those days in the training period. Since no spread adjustment is applied, this approach shares the deficiencies of the standard linear pool described above.

As before, consider a linear pool mixture model with predictive CDF

$$F^{\text{LP}}(x|\mathbf{f};\nu, \theta, \omega) := \omega G(x|\mathbf{f};\nu) + (1 - \omega)H(x|\mathbf{f};\theta), \quad \omega \in [0, 1]. \quad (4.3)$$

Short calculation based on the integral representation in the second line of (3.9) shows

$$\begin{aligned} \text{CRPS}(F^{\text{LP}}(\cdot|\mathbf{f};\nu, \theta, \omega), x) &= \omega^2 \text{CRPS}(G(\cdot|\mathbf{f};\nu), x) + (1 - \omega)^2 \text{CRPS}(H(\cdot|\mathbf{f};\theta), x) \\ &+ 2\omega(1 - \omega) \left[ \int_{-\infty}^x G(y|\mathbf{f};\nu)H(y|\mathbf{f};\theta)dy + \int_x^{\infty} (1 - G(y|\mathbf{f};\nu))(1 - H(y|\mathbf{f};\theta))dy \right]. \end{aligned}$$

Let  $\nu_{\circ}$  and  $\theta_{\circ}$  denote the optimal parameters of the individual models in the training set estimated in the first step for day  $n + \tau$ , that is

$$\nu_{\circ} := \arg \min_{\nu} \overline{\text{CRPS}(G, \nu)} \quad \text{and} \quad \theta_{\circ} := \arg \min_{\theta} \overline{\text{CRPS}(H, \theta)},$$

where

$$\overline{\text{CRPS}(G, \nu)} := \frac{1}{n} \sum_{k=1}^n \text{CRPS}(G(\cdot|\mathbf{f}_k; \nu), x_k), \quad \overline{\text{CRPS}(H, \theta)} := \frac{1}{n} \sum_{k=1}^n \text{CRPS}(H(\cdot|\mathbf{f}_k; \theta), x_k).$$

We propose to use  $\nu_o$  and  $\theta_o$  as parameters of the mixture model (4.3) and then, in the modified second step, to optimize

$$\overline{\text{CRPS}(F^{\text{LP}}, \omega)} := \frac{1}{n} \sum_{i=1}^n \text{CRPS}(F^{\text{LP}}(\cdot | \mathbf{f}_k; \nu_o, \theta_o, \omega), x_k)$$

as a function of  $\omega$ . The minimum point of  $\overline{\text{CRPS}(F^{\text{LP}}, \omega)}$  is

$$\omega_o^* = \frac{\overline{\text{CRPS}(H, \theta_o)} - \overline{\mathcal{M}(G, H, \nu_o, \theta_o)}}{\overline{\text{CRPS}(G, \nu_o)} + \overline{\text{CRPS}(H, \theta_o)} - 2\overline{\mathcal{M}(G, H, \nu_o, \theta_o)}},$$

where

$$\overline{\mathcal{M}(G, H, \nu, \theta)} := \frac{1}{n} \sum_{k=1}^n \left[ \int_{-\infty}^{x_k} G(y | \mathbf{f}_k; \nu) H(y | \mathbf{f}_k; \theta) dy + \int_{x_k}^{\infty} (1 - G(y | \mathbf{f}_k; \nu)) (1 - H(y | \mathbf{f}_k; \theta)) dy \right],$$

and short calculation shows

$$\begin{aligned} \omega_o^* &= \frac{\sum_{k=1}^n \int_{-\infty}^{\infty} H(y | \mathbf{f}_k; \theta_o) (H(y | \mathbf{f}_k; \theta_o) - G(y | \mathbf{f}_k; \nu_o)) dy}{\sum_{k=1}^n \int_{-\infty}^{\infty} (H(y | \mathbf{f}_k; \theta_o) - G(y | \mathbf{f}_k; \nu_o))^2 dy} \\ &\quad - \frac{\sum_{k=1}^n \int_{x_k}^{\infty} (H(y | \mathbf{f}_k; \theta_o) - G(y | \mathbf{f}_k; \nu_o)) dy}{\sum_{k=1}^n \int_{-\infty}^{\infty} (H(y | \mathbf{f}_k; \theta_o) - G(y | \mathbf{f}_k; \nu_o))^2 dy}. \end{aligned}$$

Now, as  $\omega_o^*$  might fall outside the unit interval  $[0, 1]$ , we use

$$\omega_o := \min \{ \max \{ 0, \omega_o^* \}, 1 \}$$

as our final estimate of the weight. Finally, one can easily show that within the training sample

$$\overline{\text{CRPS}(F^{\text{LP}}, \omega_o)} \leq \min \{ \overline{\text{CRPS}(G, \nu_o)}, \overline{\text{CRPS}(H, \theta_o)} \}, \quad (4.4)$$

so the mean CRPS of the mixture model (4.3) with parameters  $(\nu_o, \theta_o, \omega_o)$  cannot exceed the optimal mean CRPS values of the components. However, (4.4) of course gives no guarantee that for a new out-of-sample pair  $(\tilde{x}, \tilde{\mathbf{f}})$  the CRPS of the mixture  $\text{CRPS}(F^{\text{LP}}(\cdot | \tilde{\mathbf{f}}; \nu_o, \theta_o, \omega_o), \tilde{x})$  does not exceed any of the corresponding individual CRPS values.

The above method can be generalized to a convex combination of  $r$  different parametric families. However, in this case the optimal weight vector is a coordinate-wise non-negative solution of a quadratic optimization problem with a single linear constraint, where the main diagonal of the corresponding  $r \times r$  symmetric matrix consists of the mean CRPS values of the component models, whereas the other entries, which are similar to  $\overline{\mathcal{M}(G, H, \nu, \theta)}$ , can be expressed via integrals.

#### 4.2. Beta-transformed linear pool

As an alternative to the spread-adjusted linear pool that allows for correcting for the lack of dispersion of the linear pool, Gneiting & Ranjan (2013) propose the beta-transformed linear pool (BLP) with predictive CDF

$$F^{\text{BLP}}(x | \mathbf{f}; \nu, \theta, \omega) := B_{\alpha, \beta}(\omega G(x | \mathbf{f}; \nu) + (1 - \omega)H(x | \mathbf{f}; \theta)), \quad \omega \in [0, 1]. \quad (4.5)$$

Here,  $B_{\alpha,\beta}$  denotes the CDF of the beta distribution with parameters  $\alpha > 0$  and  $\beta > 0$ .

Similar to the linear and spread-adjusted linear pool, the combination parameters  $\alpha, \beta, \omega$  have to be estimated from suitably chosen training data. We proceed as before and estimate the parameters by numerically optimizing the mean CRPS over the forecast cases that coincide with the training period used to determine the coefficients of the EMOS models. Note that the representation in the second line of equation (3.9) with lower bound 0 is beneficial to avoid numerical issues when computing the CRPS using numerical integration, particularly for precipitation forecasts.

#### 4.3. Bayesian non-parametric combination approach

Bassetti et al. (2017) recently proposed an extension of the BLP approach. Motivated by results on mixture distributions from theoretical statistics, they propose a forecast aggregation method based on a mixture of beta distributions. In case of a finite mixture with  $L$  components, the resulting predictive CDF is

$$F^{\text{BM}L}(x|\mathbf{f};\nu,\theta,\omega) := \sum_{\ell=1}^L w_{\ell} B_{\alpha_{\ell},\beta_{\ell}}(\omega G(x|\mathbf{f};\nu) + (1-\omega)H(x|\mathbf{f};\theta)), \quad \omega \in [0,1], \quad (4.6)$$

where  $\alpha_{\ell} > 0$ ,  $\beta_{\ell} > 0$ ,  $w_{\ell} \geq 0$ ,  $\ell = 1, 2, \dots, L$ , are the parameters of the beta mixture components. The BLP approach in (4.5) arises as a special case for  $L = 1$ .

As the number  $L$  of components is usually unknown, Bassetti et al. (2017) propose a Bayesian inference approach that allows to treat  $L$  as unbounded and random. This infinite beta mixture approach, referred to as BMC approach in the following, has CDF

$$F^{\text{BM}\infty}(x|\mathbf{f};\nu,\theta,\omega) := \sum_{\ell=1}^{\infty} w_{\ell} B_{\alpha_{\ell},\beta_{\ell}}(\omega G(x|\mathbf{f};\nu) + (1-\omega)H(x|\mathbf{f};\theta)), \quad \omega \in [0,1]. \quad (4.7)$$

Based on the slice sampling algorithm of Walker (2007) and Kalli et al. (2011) for infinite mixtures, Bassetti et al. (2017) give an algorithm that results in samples from mixture parameters  $\alpha_{\ell}, \beta_{\ell}, w_{\ell}$  and  $\omega$ , allowing to generate draws from the predictive distribution (4.7). Note that the algorithm in fact deals with finite mixtures, however, the number of components may differ from draw to draw. In order to obtain an estimate of a verification score for a given location and time, we average over the predictive CDFs obtained through iterations of the algorithm and compute the score for the mean CDF by discrete approximation of the first integral of (3.9) over a dense grid. Note that this approach is in line with theoretical considerations on forecast evaluation based on simulation output discussed in Krüger et al. (2016).

In a case study based on wind speed data from a single observation station in Bassetti et al. (2017), the BMC method shows very promising results and substantially outperforms the linear pool. Here, we apply the BMC method to larger data sets of ensemble forecasts of wind speed and precipitation at multiple stations. Due to the point mass at zero precipitation in the forecast distributions, some minor adjustments of the sampling algorithm described in Bassetti et al. (2017) are required. Specifically, in steps 4 and 5 of the

Gibbs sampling algorithm described in Section S1.2 of the supplementary materials of Bassetti et al. (2017), in case of zero observed precipitation, random values from the intervals between zero and the corresponding probabilities of no precipitation are chosen as values of the component CSG and GEV predictive CDFs. We found that without this adjustment, the marginal predictive CDF generally underestimates precipitation accumulation substantially. The value of the resulting CDF at 0, i.e.,  $F^{\text{BM}\infty}(0|\mathbf{f};\nu,\theta,\omega)$ , is approximated by linear interpolation of the values at the first two grid points.

## 5. Case studies

Here, we report the results of four case studies for the wind speed and precipitation data sets introduced in Section 2. Note that BMC results for each forecast case are based on the forecast distribution given by the mean of 50 predictive CDFs obtained from the post burn-in iterations of the sampling algorithm described in Bassetti et al. (2017).

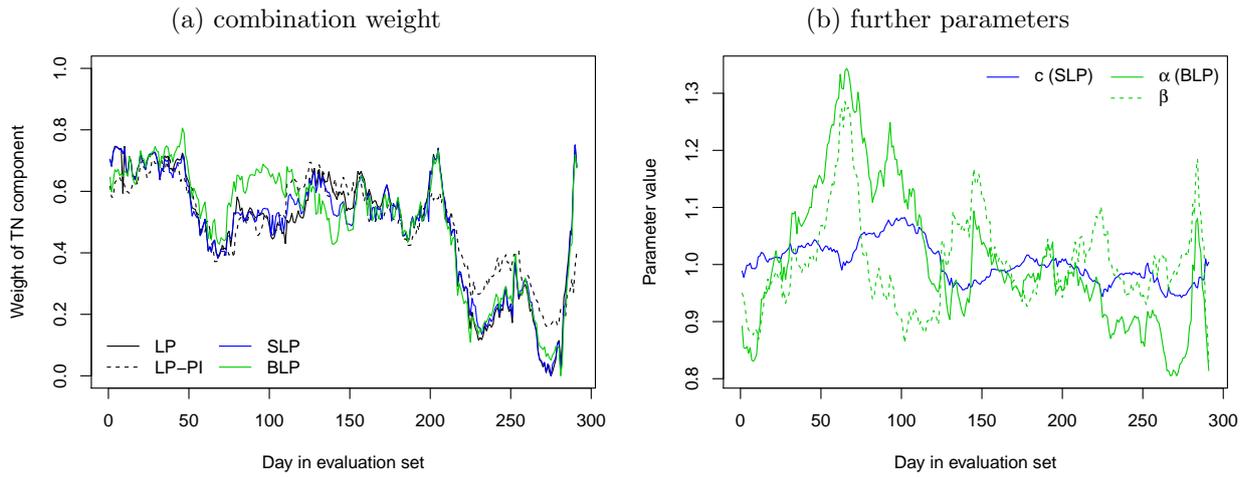
### 5.1. Wind speed

The post-processing models introduced in Section 3 are estimated using the optimum score estimation approach described in Section 3.3. The TN and LN component models are estimated by minimizing the mean CRPS over the training sets, whereas ML estimation is employed for the full mixture model (3.3). We remark that ML estimation has also been tested for the TN and LN EMOS models, however this approach results in slightly worse predictive performance. Following previous work (Baran & Lerch, 2015, 2016), we use rolling training periods of length 30 days (UWME data) and 43 days (ALADIN-HUNEPS data), and estimate the parameters regionally by combining forecast cases from all available observation stations to form a single training set for all stations. Note that alternative similarity-based semilocal approaches to selecting the training sets have been investigated in Lerch & Baran (2017).

Given the estimated coefficients of the component models, the combination parameters in the two-step combination approaches are estimated over the corresponding rolling training periods as described in Section 4. For the UWME data, forecast cases from calendar year 2007 were also used to obtain training periods of equal length for all models which are validated on the data of calendar year 2008. For the ALADIN-HUNEPS data, the first 43 days are not included in the evaluation period (that ranges from 27 June 2012 to 31 March 2013) in order to compare all models over equal training periods.

Figure 3 graphically illustrates the resulting combination parameters for the LP, LP-PI, SLP and BLP methods. Here, the BMC method is excluded as the parameters vary over the random draws of the algorithm and do not allow for a straightforward summary. For both data sets, the estimated weight parameters are generally very similar for all methods, with minor deviations for the LP-PI and BLP approaches. The spread-adjustment parameter  $c$  in the SLP method does not vary much over time, whereas the  $\alpha, \beta$  parameters in the BLP approach fluctuate much more rapidly.

UWME



ALADIN-HUNEPS

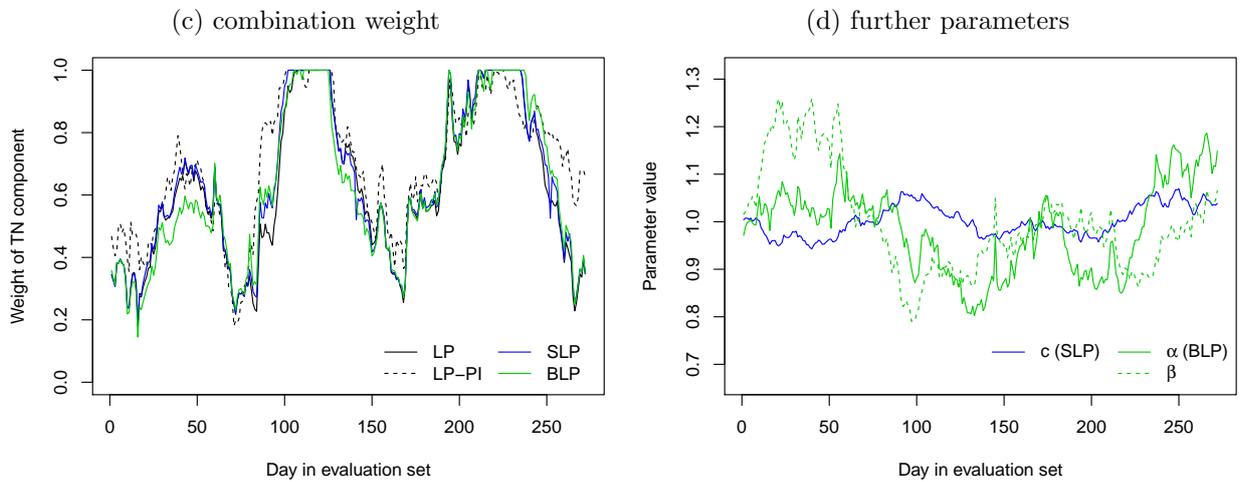


Figure 3: Illustration of mixture weights and other combination parameters for the LP, LP-PI, SLP and BLP combination methods over the corresponding verification periods for both wind speed data sets.

Table 1: Mean CRPS for probabilistic wind speed forecasts of the raw ensemble, the TN, LN and TN-LN (ML) EMOS models, and the forecast combination approaches.

Forecast	UWME	ALADIN-HUNEPS
Ensemble	1.353	0.804
TN	1.114	0.735
LN	1.113	0.740
TN-LN (ML)	1.100	0.731
EWLP	1.111	0.734
LP	1.111	0.735
LP-PI	1.111	0.735
SLP	1.111	0.737
BLP	1.110	0.738
BMC	1.106	0.738

Table 1 shows the mean CRPS values for all post-processing models and combination approaches for both data sets. All post-processing and combination methods substantially improve the raw ensemble forecasts. Among the post-processing models, the full TN-LN mixture model performs best, and the ranking of the TN and LN model depends on the data set at hand. For the UWME data, forecast combination methods outperform the individual TN and LN component models, but are unable to compete with the TN-LN mixture model. The relative differences between the combination approaches are small, with the BLP and BMC approaches showing slightly better results. By contrast, only the EWLP combination method is able to perform a bit better than the TN EMOS model for the ALADIN-HUNEPS data, and the SLP, BLP and BMC approaches result in slightly worse forecasts. Note that the BLP and SLP methods result in worse forecasts compared to the LP approach even though the latter arises as a special case for  $\alpha = \beta = 1$  and  $c = 1$ . A potential explanation for these observations is the danger of over-fitting in choosing the optimal combination parameter values in the training sample that might not be optimal for the corresponding out of sample evaluation set. Further, the ALADIN-HUNEPS data set is comprised of only 10 observation stations. The training sets thus contain fewer forecast cases compared to the UWME data which might favor combination methods with a lower number of parameters.

To assess the statistical significance of difference between the verification scores we utilize the Diebold-Mariano (DM; Diebold & Mariano, 1995) test which allows to account for the temporal dependencies in the forecast errors. For a pair of forecast methods  $F_1, F_2$ , denote the vector of CRPS differences by  $(d_1, d_2, \dots, d_n)$ , with

$$d_i(F_1, F_2) = \text{CRPS}(F_1^{(i)}, x_i) - \text{CRPS}(F_2^{(i)}, x_i),$$

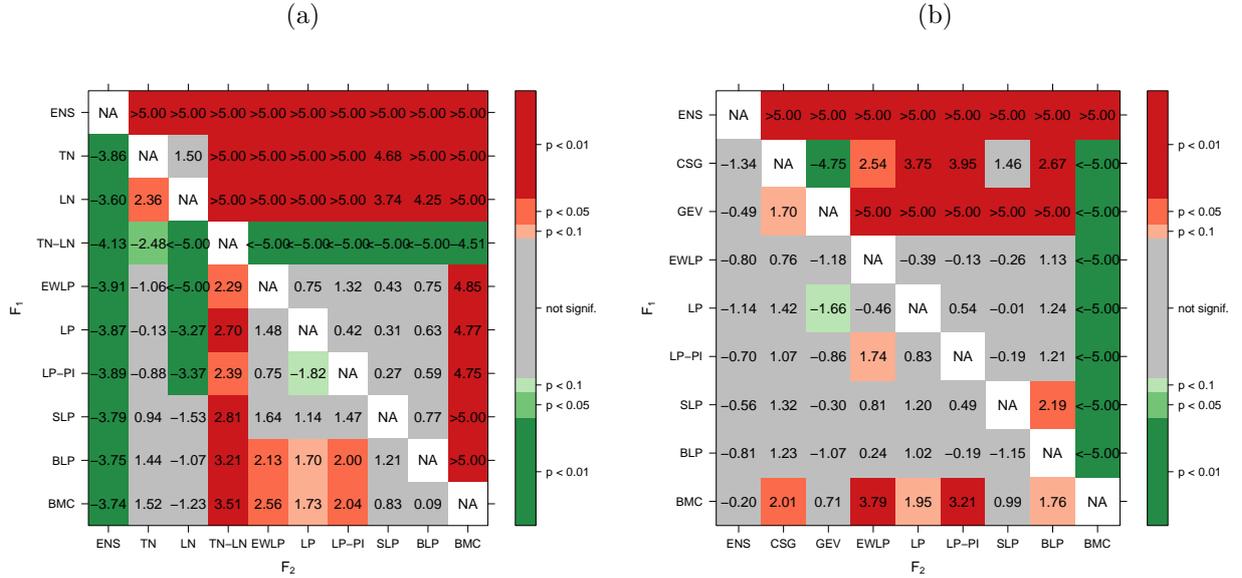


Figure 4: Summary DM test results for both (a) wind speed and (b) precipitation data sets and all pair-wise comparisons of forecasts. In both plots, the upper triangle contains results for the UWME data, and the lower triangle contains results for the ALADIN-HUNEPS data. The entry in row  $i$  and column  $j$  contains the value of the DM test statistic  $t_n$  for comparisons of  $F_1$  and  $F_2$ , where  $F_1$  is the forecast of the model in the  $i$ -th row, and  $F_2$  is the forecast of the model in the  $j$ -th column, color-coded so that green (red) entries indicate superior performance of the model in the corresponding row (column).

where  $F_j^{(i)}$  denotes the forecast distribution  $F_j$ ,  $j = 1, 2$ , for forecast case  $i = 1, 2, \dots, n$  in the evaluation set, and  $x_i$  denotes the corresponding observation.

DM tests are formal statistical tests of equal predictive performance based on the test statistic

$$t_n = \sqrt{n} \frac{\bar{d}}{\hat{\sigma}_d},$$

where  $\bar{d} = \frac{1}{n} \sum_{i=1}^n d_i$  and  $\hat{\sigma}_d$  is an estimator of the asymptotic variance of the score difference. Under standard regularity conditions,  $t_n$  is asymptotically standard normal under the null hypothesis of equal predictive performance of  $F_1$  and  $F_2$ . Negative values of  $t_n$  indicate superior predictive performance of  $F_1$ , and  $F_2$  is preferred if  $t_n$  is positive. As an estimator  $\hat{\sigma}_d^2$ , we use the sample autocovariance up to lag  $\tau - 1$  in case of  $\tau$  step ahead forecasts, see Baran & Lerch (2016) for details.

Figure 4(a) graphically summarizes the results of the pairwise DM tests for wind speed. For the UWME data illustrated in the upper triangle the tests reveal a high level of significance of the score differences between the component models and the forecast combination approaches. Similar results can be observed for the ALADIN-HUNEPS data (lower triangle), where, however, the differences between the component and combined models are generally smaller and less significant. Significant score differences can be observed between BLP and the approaches based on the linear pool. Evaluation of the different post-processing approaches using the mean logarithmic score results in a similar picture, we thus do not report the cor-

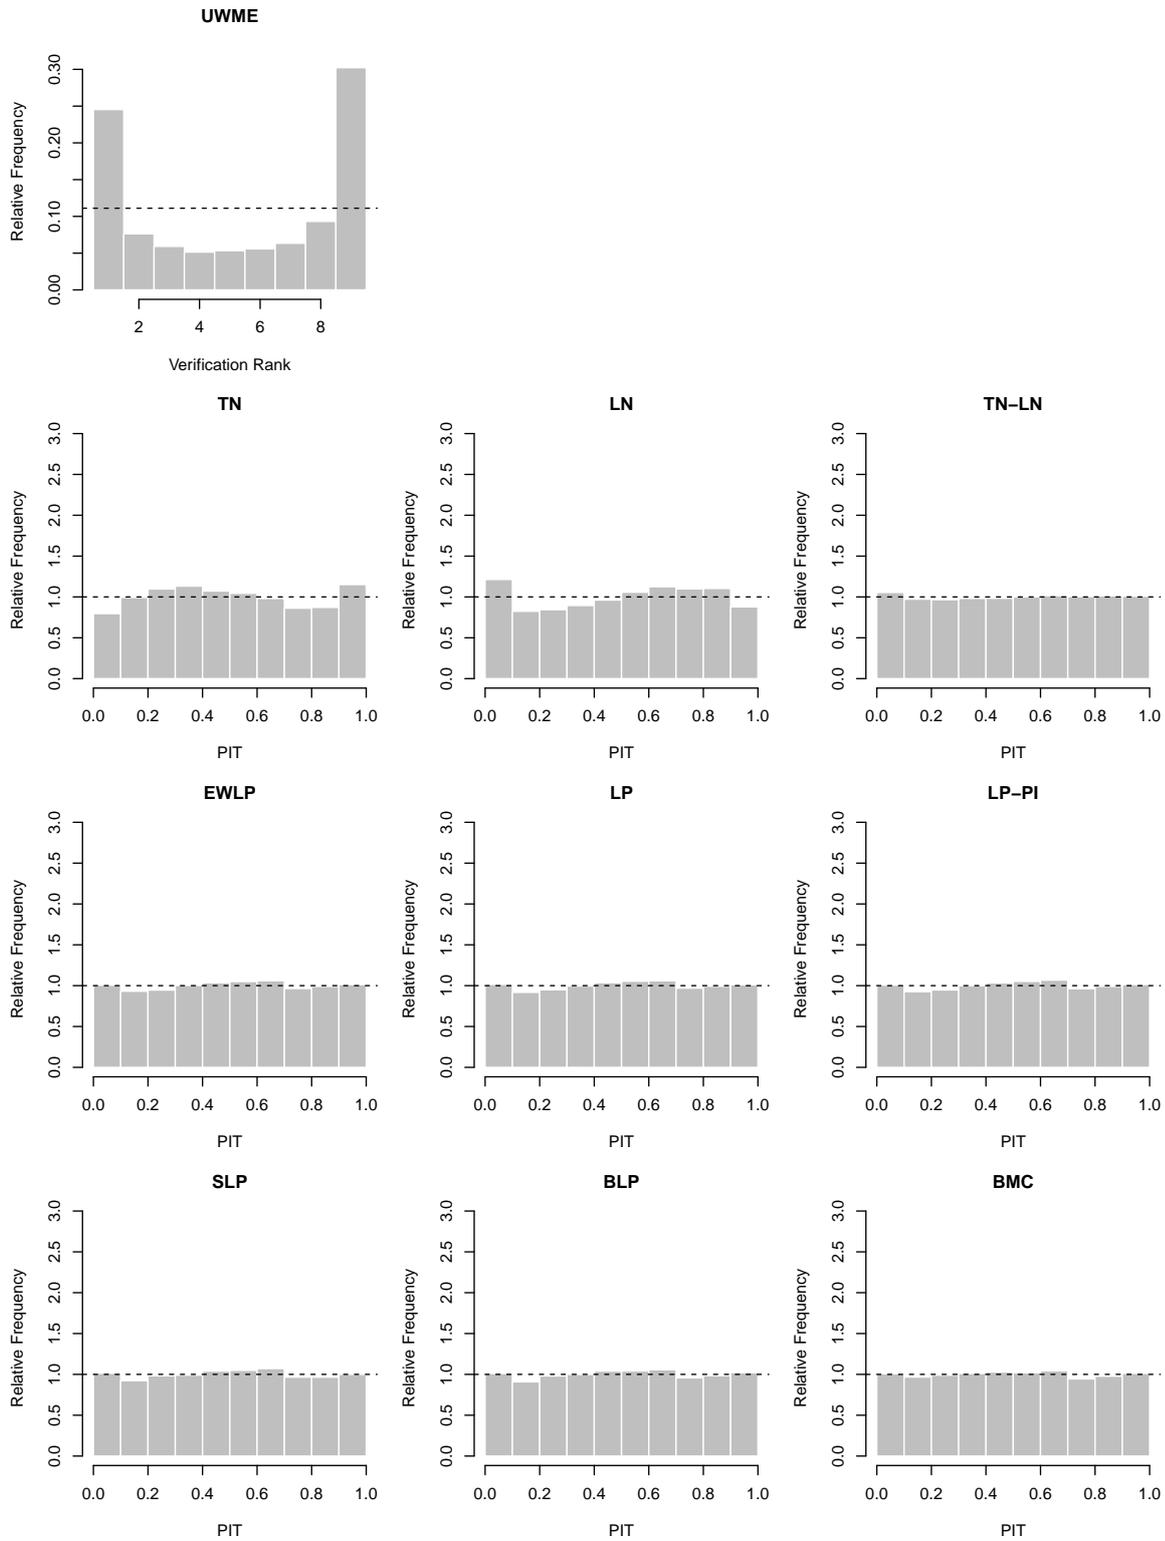


Figure 5: Verification rank histogram of raw ensemble forecasts and PIT histograms for post-processed and combined forecast distributions for the UWME data.

responding values. A notable difference is that the TN-LN mixture model shows the worst predictive performance for the ALADIN-HUNEPS data. However, note that using the logarithmic score generally results in less significant score differences compared to the CRPS.

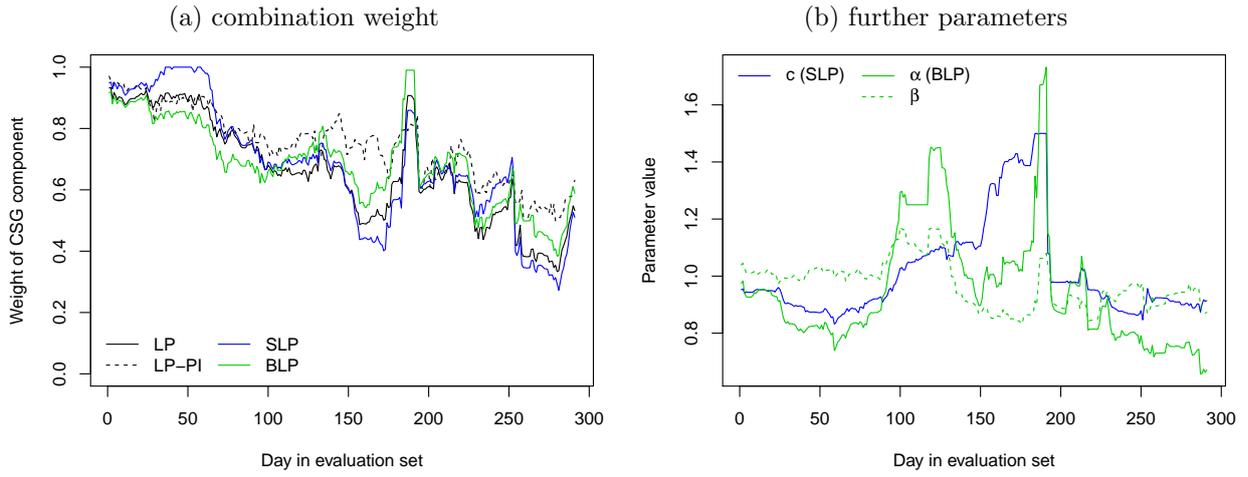
The improved predictive performance of the forecast combination approaches compared to the individual EMOS models based on single parametric distributions can be partially explained by the improved calibration of the predictive distributions that will be demonstrated in the following. Calibration of the raw ensemble and the post-processed forecasts can be assessed graphically with the help of verification rank and probability integral transform (PIT) histograms, respectively. The former is the histogram of ranks of the validating observations with respect to the corresponding ensemble predictions computed for all forecast cases (see e.g. Wilks, 2011, Section 7.7.2). For a calibrated ensemble, the observations and the ensemble forecasts should be exchangeable, resulting in a uniform verification rank histogram. The PIT is the value of the predictive CDF evaluated at the verifying observation (Raftery et al., 2005), PIT histograms can therefore be seen as continuous counterparts of verification rank histograms. The visual inspection of deviations from the desired uniform distribution of the verification ranks and PIT values allows to further detect possible reasons of miscalibration (Gneiting et al., 2007).

A verification rank histogram of the raw ensemble forecast and PIT histograms of the post-processed and combined forecast distributions for the UWME data are shown in Figure 5. Note that for the BMC forecasts PIT values are calculated for all 50 predictive CDFs of a given forecast case. Compared to the U-shaped verification rank histogram of the underdispersive raw ensemble forecasts where the observation takes too many high and low ranks, all post-processing approaches are better calibrated which is indicated by smaller deviations from the desired uniform distribution of the PIT values. The calibration of the individual TN and LN component models is not perfect, with the TN model showing systematic over-predictions of high wind speeds, and the LN model over-predicting low wind speed values. By contrast, all forecast combination approaches are able to correct for these deficiencies and are well calibrated, similar to the full TN-LN mixture model. The differences in calibration among the combination methods are small, as the PIT histograms are virtually indistinguishable. The results for the ALADIN-HUNEPS data are qualitatively similar, the corresponding Figure A.8 is shown in Appendix A.

## 5.2. Precipitation

Similar to wind speed, the post-processing models introduced in Section 3 are estimated using optimum score estimation approaches. We again report the results corresponding to CSG and GEV EMOS parameters minimizing the mean CRPS over the training data, as ML estimation of parameters in general yields models with slightly poorer forecast skills. The coefficients of the EMOS models (3.4) and (3.6) based on single CSG and GEV distributions are obtained using rolling training periods of lengths 70 (UWME) and 55 days (ALADIN-HUNEPS), which ensures comparability with Baran & Nemoda (2016).

UWME



ALADIN-HUNEPS

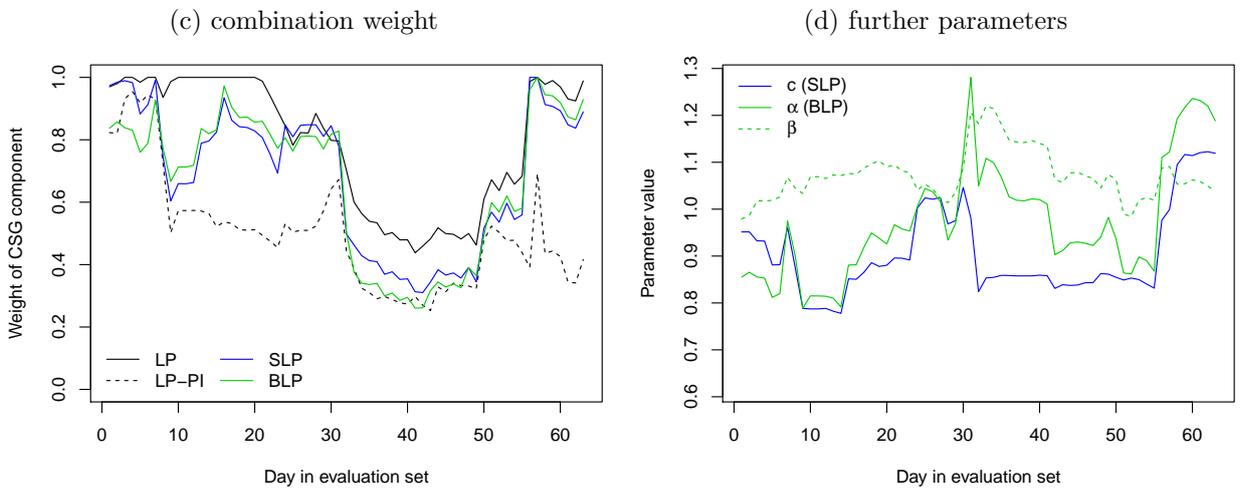


Figure 6: Illustration of mixture weights and other combination parameters for the LP, LP-PI, SLP and BLP combination methods over the corresponding verification periods for both precipitation data sets.

Table 2: Mean CRPS for probabilistic precipitation accumulation forecasts of the raw ensemble, the CSG and GEV EMOS models, and the forecast combination approaches.

Forecast	UWME	ALADIN-HUNEPS
Ensemble	2.884	0.269
CSG	2.198	0.258
GEV	2.227	0.264
EWLP	2.188	0.261
LP	2.189	0.259
LP-PI	2.189	0.262
SLP	2.190	0.263
BLP	2.184	0.261
BMC	2.288	0.267

Given the estimated coefficients of the CSG and GEV component models, the parameters of the two-step combination approaches are estimated as described for wind speed. For the UWME data, forecast cases from calendar year 2007 are again used to obtain training periods of equal length for all models, whereas the first 55 days are excluded from the evaluation period of the ALADIN-HUNEPS data. In this way UWME forecasts are again validated on data from calendar year 2008, whereas the verification period for ALADIN-HUNEPS precipitation forecasts is 21 January – 25 March 2011.

The estimates of the combination parameters over the evaluation period are shown in Figure 6. The mixture weights of the LP, LP-PI, SLP and BLP approaches exhibit relatively similar developments over time, and the spread-adjustment parameter  $c$  shows slightly higher variability compared to the wind speed forecasts.

Mean CRPS values for all post-processing models and forecast combination methods for both data sets are shown in Table 2. Compared to wind speed, the relative improvements of both EMOS models over the raw ensemble forecasts are smaller, particularly for the ALADIN-HUNEPS data. This observation is in line with various comparative studies of post-processing models for different variables (see for example Hemri et al., 2014). The CSG model outperforms the GEV model, and in case of the UWME data, the predictive performance is further improved by combining the forecasts via the EWLP, LP, LP-PI, SLP and BLP approaches which show small relative differences. In light of the larger relative score differences in favor of the CSG method it is worth noting that the estimated mixture weights of the CSG component in these combination approaches are between 0.3 and 0.7 for a large number of forecast cases, see Figure 6. As observed for wind speed, none of the combination methods is able to outperform the best component model for the ALADIN-HUNEPS data. Forecasts produced by the BMC method are worse than both

component models for both data sets and only marginally better than the raw ensemble forecasts in case of the ALADIN-HUNEPS data. These results call for an extension of the methodology towards forecast distributions with point masses beyond the provisional adaptation described in Section 4.3.

The statistical significance of the observed score differences is again assessed DM tests with the setup described above. The results are summarized in Figure 4(b). For the ALADIN-HUNEPS data none of the score differences show a high level of significance. For the UWME data, the differences between the individual post-processing models and the forecast combination approaches are more pronounced. The score differences between the EWLP, LP, LP-PI, SLP and BLP approaches are not significant, however, as expected the BMC forecasts perform substantially worse compared to all alternatives for both data sets.

For the UWME forecasts evaluation of the competing post-processing approaches using the mean logarithmic scores (not reported) results in a similar ranking as the order with respect to the mean CRPS. In case of ALADIN-HUNEPS data the relative predictive performances are somewhat different, which might be explained by the differences how CRPS and logarithmic score penalize forecast errors and the smaller number of forecast cases in the verification period.

In contrast to the wind speed forecast discussed above, the EMOS models for precipitation accumulation provide relatively well calibrated forecasts for both data sets, and the forecast combination methods only result in slightly improved calibration. The corresponding verification rank and PIT histograms are provided in Figures A.9 and A.10 in Appendix A.

### 5.3. Computational aspects

An important aspect of the empirical comparison of forecast combination approaches are the computational costs associated with increasingly complex combination methods. Clearly, any estimation of additional combination parameters will necessarily add computational cost compared to the component models.

Figure 7 graphically summarizes the distribution of computation times over days in the evaluation periods of all forecast combination approaches. The computation times of the EWLP approach correspond to the sum of the corresponding computation times required for estimating the parameters of the two component models. The computation times of all other combination approaches are the sum of these and the additional time required for determining the combination parameters. All computation times are obtained using a standard laptop with an Intel Core i7-4700 MQ ( $4 \times 2.40$  GHz) CPU and 20 GB RAM. Importantly, note that the results for the BMC are not directly comparable as they were obtained in Matlab (version 2016a) under Windows 10, whereas all other computation times were recorded in R (version 3.4.1) under Linux Fedora 26.

Generally, the BLP approach results in the highest computational costs, followed by the BMC, SLP and LP methods. For wind speed, all these approaches show higher computation times than the TN-LN mixture model, whereas the proposed plug-in variant of the linear pool only has marginally higher computation times

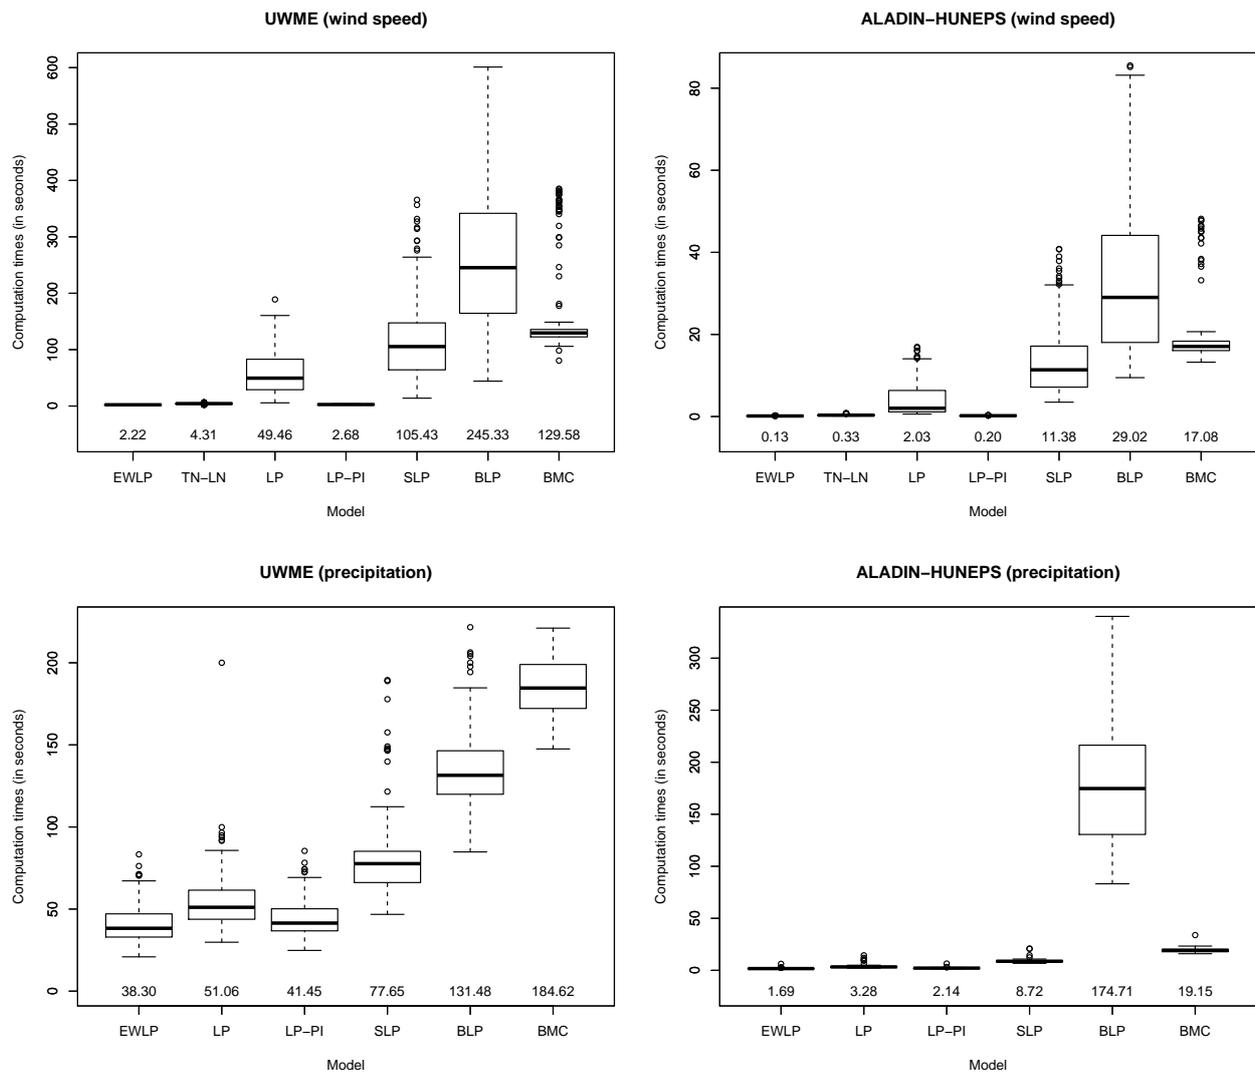


Figure 7: Distribution of computation times over days in the verification period for the different forecast combination approaches and the four data sets. Median computation times for the different models are shown above the labels on the horizontal axis.

than the sum of the component models. In light of the results presented in the previous subsections, the higher computational costs of the more complex methods are not necessarily justified by better predictive performance.

However, note that all computation times reported here are negligible compared to the computational costs of generating the ensemble forecasts (UWME, 2011).

## 6. Conclusions

We have investigated the feasibility of using forecast combination approaches to improve the predictive performance of statistical post-processing models based on single parametric families. In general, the results strongly depend on the data set at hand, and forecast combination may either provide slight improvements, or even result in worse forecasts compared to the best mixture component. The relative improvements obtained through forecast combination are larger in the case of wind speed where the PIT histograms reveal complementary systematic errors in calibration. However, none of the forecast combination methods is able to outperform a jointly estimated full TN-LN mixture EMOS models.

Compared to previous work of Möller & Groß (2016) and Bassetti et al. (2017), we generally do not find as substantial differences in predictive performance between the individual EMOS models and the combination approaches. Further, compared to the case study of wind speed data in Bassetti et al. (2017) the LP and BMC method show much less significant differences. Generally, we find that more complex combination approaches do not necessarily perform better than the simple LP approach. In fact, except for the BMC approach for the UWME wind speed data, no competing combination method is able to significantly outperform the linear pool or the plug-in variant. It might thus be interesting to further investigate which features of the data and mixture component models would be beneficial for more complex combination approaches.

The larger relative improvements of forecast combination for the UWME data sets may indicate that longer training periods are generally better suited to determine the combination parameters. Therefore, training sets expanding with time might improve the predictive performance. Further, local estimation of the combination parameters using only data from the single observation station of interest, or alternative similarity-based semilocal approaches (Lerch & Baran, 2017) might be better able to account for locally varying station-dependent features of the forecast errors of the individual EMOS models. Here we used equal regional training periods for all models to ensure direct comparability, and leave such extensions for future work.

We have proposed a plug-in variant of the linear pool that utilizes the most recent EMOS coefficient estimates to replace those in the training period used to determine the mixture weight, and thereby reduces the number of required numerical integrations to a single one. Despite discarding potentially useful

information for forecast combination, the LP-PI approach results in very similar mixture weights and predictive performance compared to the traditional linear pool, and may thus offer an alternative option if the computational costs of estimating the mixture weight are high.

Forecast combination methods may also offer a new approach to post-processing multi-model ensemble predictions such as the TIGGE forecasts (Swinbank et al., 2016). Instead of utilizing forecasts from all models as input of an EMOS model based on a single parametric distribution it might be helpful to post-process the ensemble predictions of the different models independently, and then subsequently combine the forecast distributions with the approaches discussed above. A further interesting starting point for future research might be the application of the combination methods to other weather variables such as total cloud cover (Hemri et al., 2016).

**Acknowledgments.** The work leading to this paper was partially done during visits of Sándor Baran at the Heidelberg Institute for Theoretical Studies in the framework of the visiting scientist program and DAAD program “Research Stays for University Academics and Scientists, 2017”. Sándor Baran was also supported by the János Bolyai Research Scholarship of the Hungarian Academy of Sciences and by the EFOP-3.6.1-16-2016-00022 project. The project is co-financed by the European Union and the European Social Fund. Sebastian Lerch gratefully acknowledges support by Deutsche Forschungsgemeinschaft (DFG) through the project “C7 – Statistical post-processing and stochastic physics for ensemble predictions” within Transregional Collaborative Research Center 165 “Waves to Weather”, and thanks the Klaus Tschira Foundation for infrastructural support at the Heidelberg Institute for Theoretical Studies. Helpful comments by Tilmann Gneiting are gratefully acknowledged. The authors thank Federico Bassetti, Roberto Casarin, and Francesco Ravazzolo for kindly providing a Matlab implementation of the BMC method, and further thank the University of Washington MURI group for providing the UWME data and Mihály Szűcs from the HMS for providing the ALADIN-HUNEPS data. The comments of the editor and two anonymous reviewers greatly improved the quality of this paper.

## References

- Baran, S. & 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. & Lerch, S. (2016). Mixture EMOS model for calibrating ensemble forecasts of wind speed. *Environmetrics*, 27, 116–130.
- Baran, S. & Nemoda, D. (2016). Censored and shifted gamma distribution based EMOS model for probabilistic quantitative precipitation forecasting. *Environmetrics*, 27, 280–292.
- Bassetti, F., Casarin, R. & Ravazzolo, F. (2017). Bayesian nonparametric calibration and combination of predictive distributions. *Journal of the American Statistical Association*, in press, doi:10.1080/01621459.2016.1273117.
- Billio, M., Casarin, R., Ravazzolo, F. & van Dijk, H.K. (2013). Time-varying combinations of predictive densities using nonlinear filtering. *Journal of Econometrics*, 177, 213–232.

- Del Negro, M., Hasegawa, R.B. & Schorfheide, F. (2016). Dynamic prediction pools: an investigation of financial frictions and forecasting performance. *Journal of Econometrics*, 192, 391–405.
- Descamps, L., Labadie, C., Joly, A., Bazile, E., Arbogast, P. & Cébron, P. (2015). PEARP, the Météo-France short-range ensemble prediction system. *Quarterly Journal of the Royal Meteorological Society*, 141, 1671–1685.
- Diebold F.X. & Mariano, R.S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13, 253–263.
- Eckel, F.A. & Mass, C.F. (2005). Effective mesoscale, short-range ensemble forecasting. *Weather and Forecasting*, 20, 328–350.
- Flowerdew, J. (2014). Calibrating ensemble reliability whilst preserving spatial structure. *Tellus A*, 66, 22662.
- Fraley, C., Raftery, A.E. & Gneiting, T. (2010). Calibrating multimodel forecast ensembles with exchangeable and missing members using Bayesian model averaging. *Monthly Weather Review*, 138, 190–202.
- Garcia, A., Torres, J.L., Prieto, E. & De Francisco, A. (1998). Fitting wind speed distributions: A case study. *Solar Energy*, 62, 139–144.
- Gneiting, T. (2014). Calibration of medium-range weather forecasts. ECMWF Technical Memorandum No. 719. Available at: <http://www.ecmwf.int/sites/default/files/elibrary/2014/9607-calibration-medium-range-weather-forecasts.pdf> [Accessed on 12 January 2018]
- Gneiting, T., Balabdaoui, F. & 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. (2005). Weather forecasting with ensemble methods. *Science*, 310, 248–249.
- Gneiting, T. & 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. & Goldman, T. (2005). Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Monthly Weather Review*, 133, 1098–1118.
- Gneiting, T. & Ranjan, R. (2013). Combining predictive distributions. *Electronic Journal of Statistics*, 7, 1747–1782.
- Good, I.J. (1952). Rational decisions. *Journal of the Royal Statistical Society Series B*, 14, 107–114.
- Grell, G.A., Dudhia, J. & Stauffer, D.R. (1995). A description of the fifth-generation Penn state/NCAR mesoscale model (MM5). Technical Note NCAR/TN-398+STR. National Center for Atmospheric Research, Boulder. Available at: <http://www2.mmm.ucar.edu/mm5/documents/mm5-desc-doc.html> [Accessed on 12 January 2018]
- Hall, S.G. & Mitchell, J. (2007). Combining density forecasts. *International Journal of Forecasting*, 23, 1–13.
- Hamill, T.M. & Whitaker, J.S. (2006). Probabilistic quantitative precipitation forecasts based on reforecast analogs: Theory and application. *Monthly Weather Review*, 134, 3209–3229.
- Hemri, S., Haiden, T. & Pappenberger, F. (2016). Discrete post-processing of total cloud cover ensemble forecasts. *Monthly Weather Review*, 2565–2577.
- Hemri, S., Scheuerer, M., Pappenberger, F., Bogner, K. & Haiden, T. (2014). Trends in the predictive performance of raw ensemble weather forecasts. *Geophysical Research Letters*, 41, 9197–9205.
- Horányi, A., Kertész, S., Kullmann, L. & Radnóti, G. (2006). The ARPEGE/ALADIN mesoscale numerical modeling system and its application at the Hungarian Meteorological Service. *Időjárás*, 110, 203–227.
- Horányi, A., Mile, M. & Szűcs, M. (2011). Latest developments around the ALADIN operational short-range ensemble prediction system in Hungary. *Tellus A*, 63, 642–651.
- Iversen, T., Deckmin, A., Santos, C., Sattler, K., Bremnes, J.B., Feddersen, H. & Frogner, I.-L. (2011). Evaluation of 'GLAMEPS' – a proposed multimodel EPS for short range forecasting. *Tellus A*, 63, 513–530.
- Jordan, A., Krüger, F., & Lerch, S. (2017). Evaluating probabilistic forecasts with the R package scoringRules. Working paper, preprint available at <https://arxiv.org/abs/1709.04743> [Accessed on 12 January 2018]
- Justus, C.G., Hargraves, W.R., Mikhail, A. & Graber, D. (1978). Methods for estimating wind speed frequency distributions. *Journal of Applied Meteorology and Climatology*, 17, 350–353.

- Kalli, M., Griffin, J.E., & Walker, S.G. (2011). Slice sampling mixture models. *Statistics and Computing*, 21, 93–105.
- Krüger, F., Lerch, S., Thorarinsdottir, T.L., & Gneiting, T. (2016). Probabilistic Forecasting and Comparative Model Assessment Based on Markov Chain Monte Carlo Output. Working paper, preprint available at <https://arxiv.org/abs/1608.06802> [Accessed on 12 January 2018]
- Lerch, S. & Baran, S. (2017). Similarity-based semilocal estimation of post-processing models. *Journal of the Royal Statistical Society Series C*, 66, 29–51.
- Lerch, S. & Thorarinsdottir, T.L. (2013). Comparison of non-homogeneous regression models for probabilistic wind speed forecasting. *Tellus A*, 65, 21206.
- Matheson, J.E. & Winkler, R.L. (1976). Scoring rules for continuous probability distributions. *Management Science*, 22, 1087–1096.
- Molteni, F., Buizza, R., Palmer, T.N. & Petroliagis, T. (1996). The ECMWF ensemble prediction system: Methodology and validation. *Quarterly Journal of the Royal Meteorological Society*, 122, 73–119.
- Möller, A. & Groß, J. (2016). Probabilistic temperature forecasting based on an ensemble AR modification *Quarterly Journal of the Royal Meteorological Society*, 142, 1385–1394.
- National Weather Service (1998). Automated Surface Observing System (ASOS) User’s Guide. Available at: <http://www.weather.gov/asos/aum-toc.pdf> [Accessed on 12 January 2018]
- Opschoor, A., van Dijk, D. & van der Wel, M. (2017). Combining density forecasts using focused scoring rules. *Journal of Applied Econometrics*, 32, in press.
- Palmer, T.N. (2002). The economic value of ensemble forecasts as a tool for risk assessment: From days to decades. *Quarterly Journal of the Royal Meteorological Society*, 128, 747–774.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T. & Flannery, B.T. (2007). *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, Cambridge.
- R Core Team (2017). R: A Language and Environment for Statistical Computing. Available at: <https://www.R-project.org/> [Accessed on 12 January 2018]
- Raftery, A.E., Gneiting, T., Balabdaoui, F. & Polakowski, M. (2005). Using Bayesian model averaging to calibrate forecast ensembles. *Monthly Weather Review*, 133, 1155–1174.
- Scheuerer, M. (2014). Probabilistic quantitative precipitation forecasting using ensemble model output statistics. *Quarterly Journal of the Royal Meteorological Society*, 140, 1086–1096.
- Scheuerer, M. & Hamill, T.M. (2015). Statistical post-processing of ensemble precipitation forecasts by fitting censored, shifted gamma distributions. *Monthly Weather Review*, 143, 4578–4596.
- Scheuerer, M. & Möller, D. (2015). Probabilistic wind speed forecasting on a grid based on ensemble model output statistics. *Annals of Applied Statistics*, 9, 1328–1349.
- Slughter, J.M., Raftery, A E., Gneiting, T. & 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. & Yamaguchi, M. (2016). The TIGGE project and its achievements. *Bulletin of the American Meteorological Society*, 97, 49–67.
- Taillardat, M., Mestre, O., Zamo, M. & Naveau, P. (2016). Calibrated ensemble forecasts using quantile regression forests and ensemble model output statistics. *Monthly Weather Review*, 144, 2375–2393.
- Thorarinsdottir, T.L. & 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.
- UWME (2011). University of Washington Mesoscale Ensemble (UWME) Information. Available at: [https://atmos.washington.edu/~ens/uwme\\_info.html](https://atmos.washington.edu/~ens/uwme_info.html) [Accessed on 12 January 2018]

- Walker, S.G. (2007). Sampling the Dirichlet mixture model with slices. *Communications in Statistics - Simulation and Computation*, 36, 45–54.
- Wilks, D.S. (2011). *Statistical Methods in the Atmospheric Sciences*. 3rd ed. Elsevier, Amsterdam.
- Yang, X., Sharma, S., Siddique, R., Greybush, S.J. and Mejia, A., 2017. Postprocessing of GEFS precipitation ensemble reforecasts over the US Mid-Atlantic region. *Mon. Wea. Rev.*, 145, 1641–1658.

## Appendix A. Additional figures

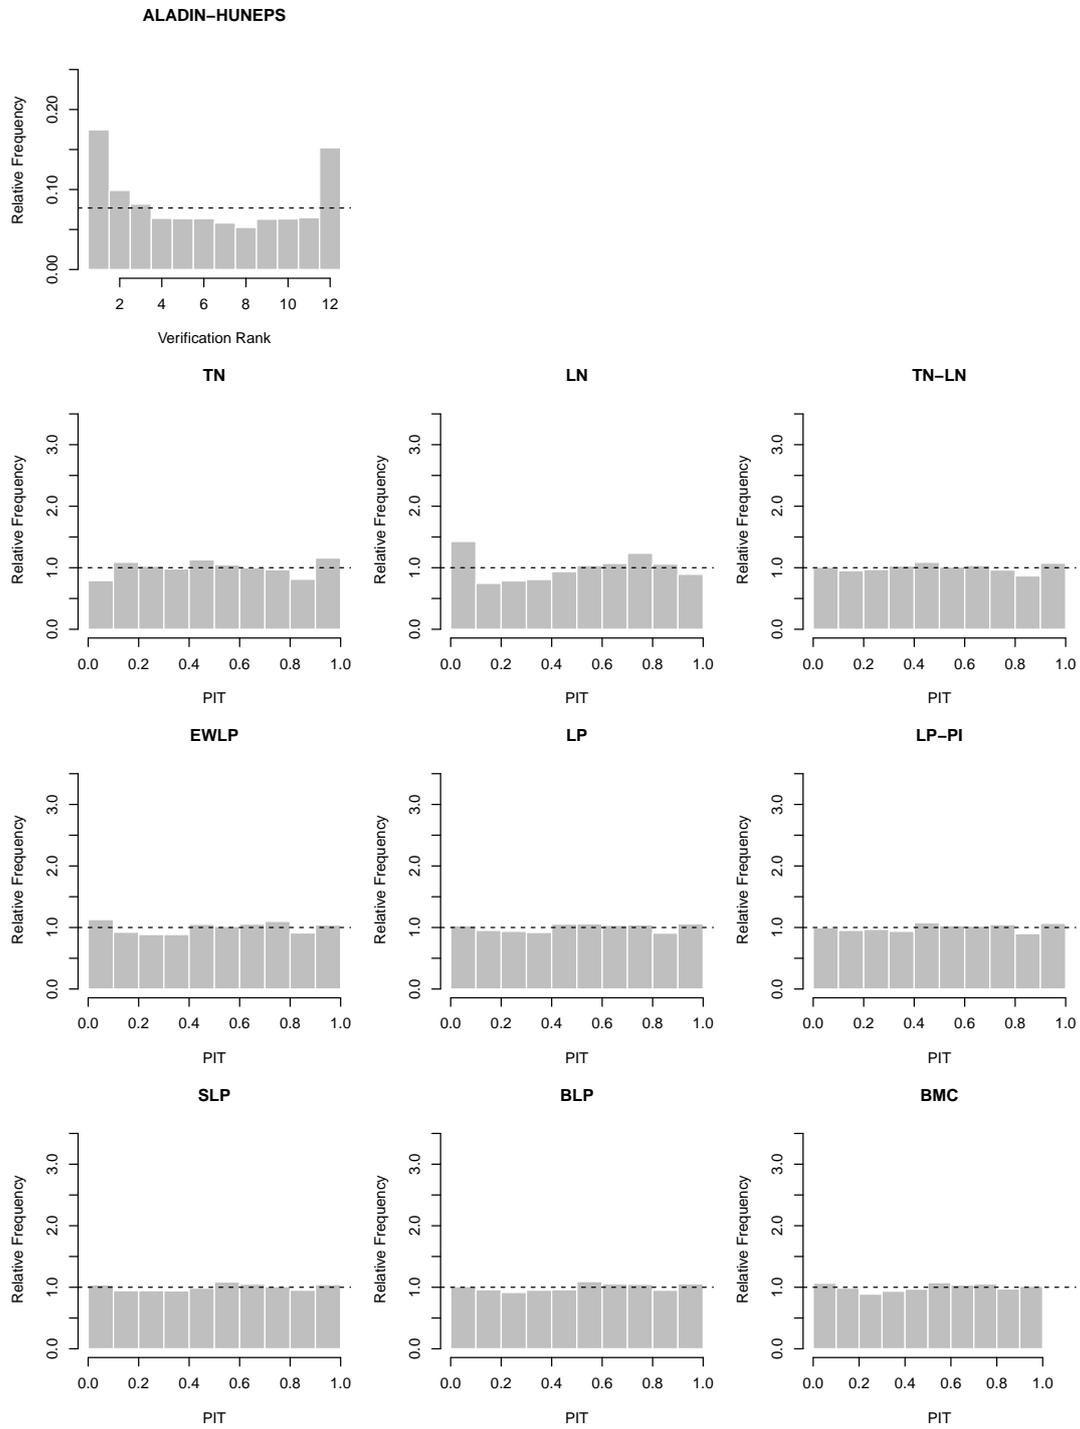


Figure A.8: Verification rank histogram of raw ensemble forecasts and PIT histograms for post-processed and combined forecast distributions for the ALADIN-HUNEPS wind speed data.

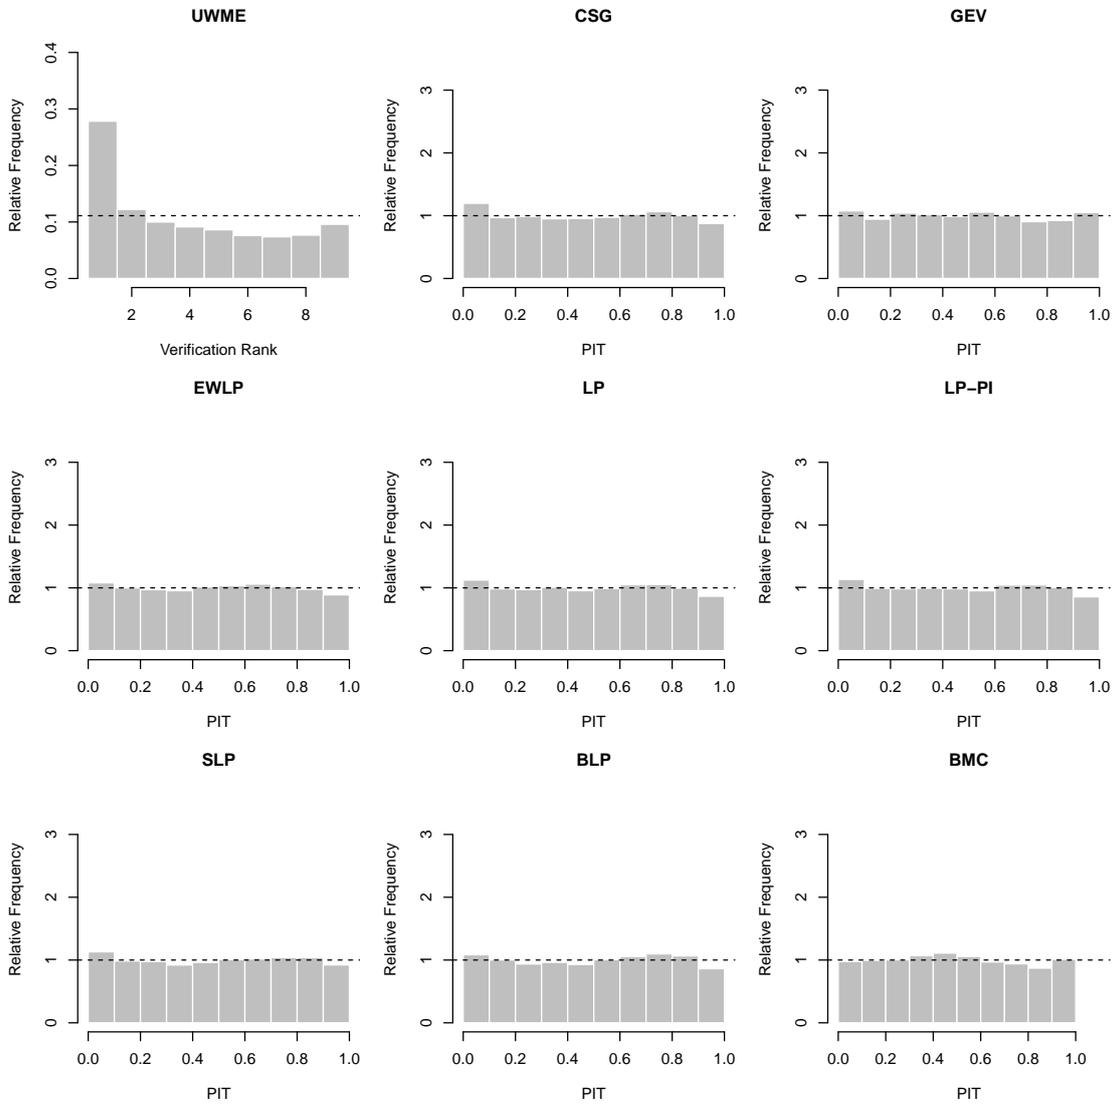


Figure A.9: Verification rank histogram of raw ensemble forecasts and randomized PIT histograms for post-processed and combined forecast distributions for the UWME precipitation data. Note that to account for the discrete-continuous nature of the models for precipitation accumulation, in case of zero observed precipitation the PITs are randomized in that a random value is chosen uniformly from the interval between zero and the probability of no precipitation (Sloughter et al., 2007). Similarly, in case of precipitation accumulation zero observations are randomized among all zero forecasts to compute the verification rank histograms for the raw ensemble forecasts.

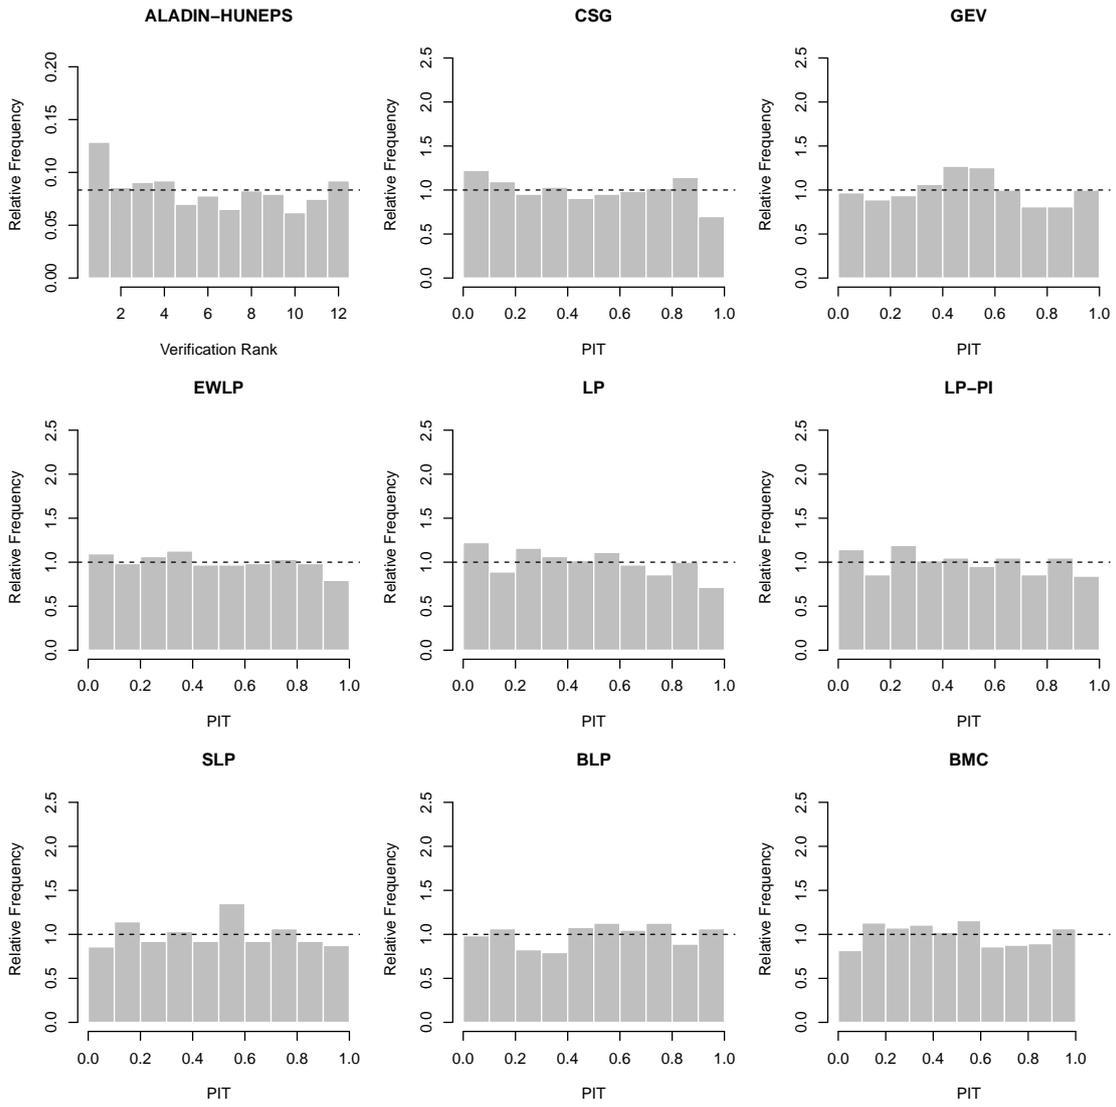


Figure A.10: Verification rank histogram of raw ensemble forecasts and randomized PIT histograms for post-processed and combined forecast distributions for the ALADIN-HUNEPS precipitation data.