

RESEARCH ARTICLE

Statistical post-processing of heat index ensemble forecasts: Is there a royal road?

Sándor Baran^{1,2}  | Ágnes Baran¹ | Florian Pappenberger² | Zied Ben Bouallègue² 

¹Faculty of Informatics, University of Debrecen, Hungary

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

Correspondence

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

Funding information

Hungarian National Research, Development and Innovation Office, Grant/Award Number: NN125679; the Hungarian Government and the European Social Fund, Grant/Award Number: EFOP-3.6.2-16-2017-00015

Abstract

We investigate the effect of statistical post-processing on the probabilistic skill of discomfort index (DI) and indoor wet-bulb globe temperature (WBGT_{id}) ensemble forecasts, both calculated from the corresponding forecasts of temperature and dew point temperature. Two different methodological approaches to calibration are compared. In the first case, we start with joint post-processing of the temperature and dew point forecasts and then create calibrated samples of DI and WBGT_{id} using samples from the obtained bivariate predictive distributions. This approach is compared with direct post-processing of the heat index ensemble forecasts. For this purpose, a novel ensemble model output statistics model based on a generalized extreme value distribution is proposed. The predictive performance of both methods is tested on the operational temperature and dew point ensemble forecasts of the European Centre for Medium-Range Weather Forecasts and the corresponding forecasts of DI and WBGT_{id}. For short lead times (up to day 6), both approaches significantly improve the forecast skill. Among the competing post-processing methods, direct calibration of heat indices exhibits the best predictive performance, very closely followed by the more general approach based on joint calibration of temperature and dew point temperature. Additionally, a machine learning approach is tested and shows comparable performance for the case when one is interested only in forecasting heat index warning level categories.

KEYWORDS

discomfort index, ensemble model output statistics, multivariate method, probabilistic forecasting, statistical post-processing, wet-bulb globe temperature

1 | INTRODUCTION

In this century, extreme temperatures have impacted 97 million people causing over 40 billion Euro economic damage (EM-DAT, 2020). In particular, heat stress has

detrimental effect on human health and human activities. In order to mitigate the effect of heat stress, dedicated warning systems have been developed (McGregor *et al.*, 2015; Morabito *et al.*, 2019). These systems require meteorological forecasts from which relevant indicators

expressing the impact of heat stress on humans are determined (Di Napoli *et al.*, 2019). A large number of different indices of various complexities exist and even a comparison of 165 different ones demonstrated that there is no universal acceptance of a single index (de Freitas and Grigorieva, 2017). In this paper, the focus is on two indices which can be easily derived from standard outputs of weather forecast models and are commonly used: the discomfort index (DI) and the indoor version of the wet-bulb globe temperature (WBGTid). Both are functions of temperature and dew point temperature, but have been chosen since they differ in their formulation and complexity which may have an impact on the final findings.

In our study, forecasts of DI and WBGTid are derived from the ensemble prediction system run at the European Centre for Medium-Range Weather Forecasts (ECMWF). Ensemble predictions provide different scenarios of the future, recognising the uncertainties inherited in weather forecasting (Leutbecher and Palmer, 2008). The derived probabilistic forecasts are an essential tool to support forecast-based decision making (Fundel *et al.*, 2019). It has been already established that heat index forecasts based on ensembles are skilful with respect to climatology up to 10 days (Pappenberger *et al.*, 2015). However, the value of these forecasts can be increased by correcting the forecasts through so-called statistical post-processing, taking account of inefficiencies in representing the aforementioned uncertainties and systematic biases. This also adjusts for the fact that the produced forecasts represent an average over a spatial area (a grid cell) whilst being used and applied locally on individual locations.

Statistical post-processing methods include parametric models such as Bayesian model averaging (BMA; Raftery *et al.*, 2005) and ensemble model output statistics (EMOS; Gneiting *et al.*, 2005), providing full predictive distributions of the weather variables at hand. We focus on the EMOS approach, which specifies the predictive distribution by a single parametric law with parameters connected to the ensemble members via appropriate link functions. Post-processing methods are often applied to individual weather quantities (Wilks, 2018). However, the dependence structure between different weather variables is of crucial importance in the calculation of heat stress indicators. This is acknowledged and handled if bivariate EMOS models (Schuhen *et al.*, 2012; Baran and Möller, 2017) or other approaches to joint calibrations such as the ensemble copula coupling (ECC) method of Schefzik *et al.* (2013) are deployed. More advanced methods (e.g., Ben Bouallegue *et al.*, 2016; Barnes *et al.*, 2019) exist, but are not investigated here.

This paper explores the fundamental question whether it is more efficient to base the calibration on the two input forecasts (temperature and dew point temperature)

or rather to calibrate directly the end-product (the heat index). We apply these two different methodological approaches to the calibration of DI and WBGTid ensemble forecasts. On the one hand, we first calibrate ensemble forecast vectors of temperature and dew point temperature, either with a bivariate EMOS model or with ECC combined with independent univariate EMOS models. The calibrated DI and WBGTid forecasts are derived from these calibrated input forecasts. On the other hand, we develop a new EMOS model for post-processing both DI and WBGTid ensemble forecasts, where the predictive PDFs follow a generalized extreme value (GEV) distribution. The benchmark method consists of a simple adjustment of the weather quantities to account for the scale mismatch between model output and point observations. A general overview of the applied procedures is given in Figure 1. The different approaches to post-processing are ranked in terms of performance; their implementation and limitations are also discussed.

Finally, we envisage the situation where one is interested only in the prediction of the warning level category of a heat index. In this case, statistical calibration can be considered as a classification problem resulting in forecast probabilities for the different categories. In addition to the use of the predictive distributions calculated from the above-mentioned calibration methods, we also test the forecast skill of a general machine learning-based approach.

The paper is organized as follows. Section 2 provides the fundamental formulae for the calculation of DI and WBGTid, followed by a description of the temperature and dew point temperature datasets. The applied calibration methods (including the novel GEV EMOS approach) are given in Section 3, and the verification procedure is discussed in Section 4. Section 5 presents the results of our study devoting separate sections to bivariate and univariate post-processing, and to calibration of discrete forecasts. Some figures are placed in the Appendix in order to improve readability. Finally, lessons learned and recommendations can be found in Section 6.

2 | DATA

2.1 | Heat index definitions

We start with the formal definition of the two heat indices of interest, which are both functions of two weather quantities: temperature and dew point (DP) temperature. Discomfort index I_d is calculated from temperature T (°C) and relative humidity (RH) H_r (%) as

$$I_d := T - 0.0055(100 - H_r)(T - 14.5). \quad (1)$$

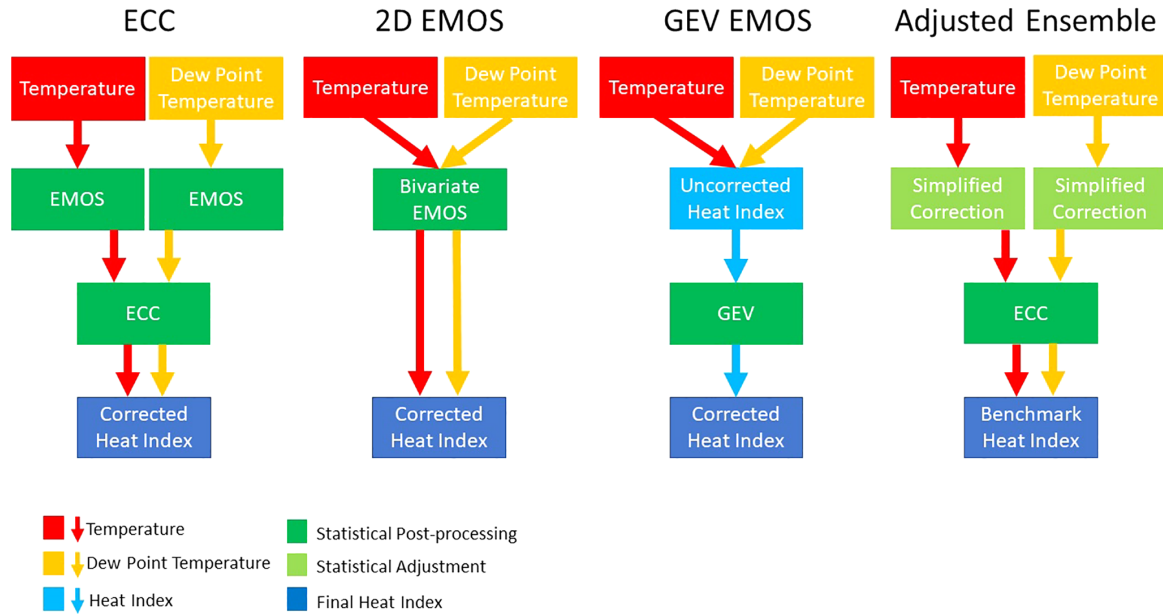


FIGURE 1 An overview of the different post-processing approaches considered in this paper

RH can easily be expressed using temperature and DP temperature T_{dp} (°C) using the Magnus formula with constants suggested by Sonntag (1990):

$$H_r := 100 \exp \left(\frac{17.62 T_{dp}}{243.12 + T_{dp}} - \frac{17.62 T}{243.12 + T} \right). \quad (2)$$

Indoor wet-bulb globe temperature T_{WBgid} is calculated as

$$T_{WBgid} := 0.67 T_{pwb} + 0.33 T, \quad (3)$$

where T_{pwb} denotes the psychrometric wet-bulb (PWB) temperature (°C). One can obtain this quantity using Bernard's formula (Bernard and Pourmoghani, 1999) providing it as a solution of equation

$$1556 P_d - 1.484 P_d T_{pwb} - 1556 P_w + 1.484 P_w T_{pwb} + 1010 (T - T_{pwb}) = 0,$$

where

$$P_d = 6.106 \exp \left(\frac{17.27 T_{dp}}{237.3 + T_{dp}} \right) \quad \text{and} \\ P_w = 6.106 \exp \left(\frac{17.27 T_{pwb}}{237.3 + T_{pwb}} \right)$$

are the saturation vapour pressures (hPa) at DP temperature T_{dp} and PWB temperature T_{pwb} , respectively (Lemke and Kjellstrom, 2012). It results that WBGTid is also a function of temperature and dew point and can be computed using, for example, the *HeatStress* package of R (Casanueva, 2019).

2.2 | Observations

Temperature and dew point observations are available at 1,459 surface synoptic observation (SYNOP) stations in the midlatitudes of Europe (35° – 65°N; 12.5°W – 42.5°E). The corresponding DI is derived by applying Equations (1) and (2), whereas WBGTid is obtained from Equation (3), where PWB is determined with the help of the *HeatStress* package. The selected dataset covers a period between 1 May and 30 September 2017.

Both DI and WBGTid observations are negatively skewed (skewnesses are –0.415 and –0.286, respectively). This should be taken into account in the choice of the parametric distribution family for EMOS post-processing. The left-skewed character of both indices is visible in Figure 2 showing the corresponding climatological histograms together with the PDFs of generalized extreme value (GEV), skew normal and split normal distributions (with matching means, variances and skewnesses). These three parametric distributions are potential candidates for EMOS modelling (Section 3.3).

2.3 | Ensemble forecasts

The operational ensemble run at ECMWF comprises 50 perturbed members. We focus on forecasts initialized at 1200 UTC and we consider 1–15 day lead times. For each observation, we associate the forecast at the nearest grid point. In order to account for the difference between model and station elevations, orographic correction is applied to both raw temperature and dew point forecasts. Practically,

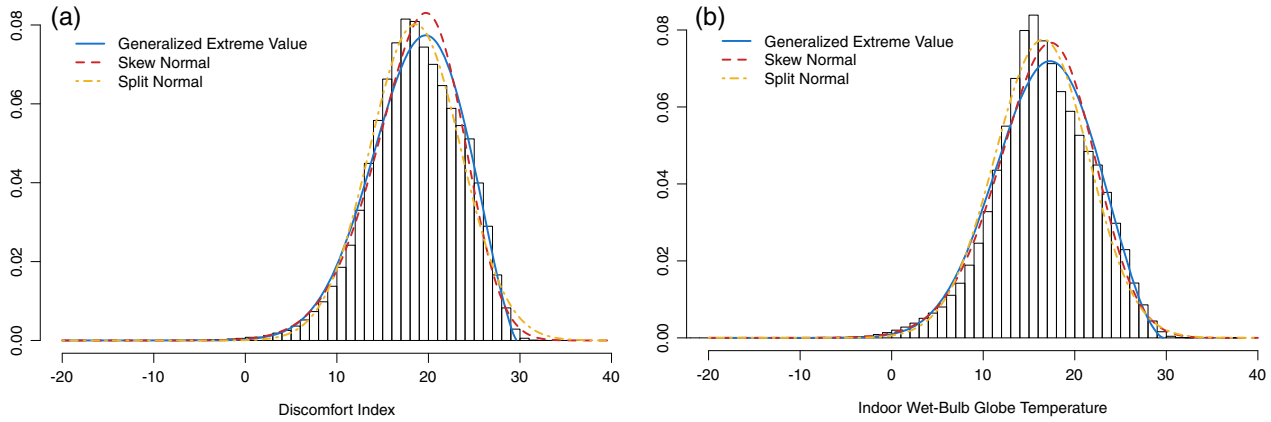


FIGURE 2 Climatological histograms of (a) DI and (b) WBGTid for the period 1 May–30 September 2017, and the PDFs of generalized extreme value, skew normal and split normal distributions with matching means, variances and skewnesses

we add $0.0065\Delta_e$ ($^{\circ}\text{C}$) to the forecasts, with Δ_e being the altitude difference between station and model representation. Temperature and dew point forecasts are used to derive the corresponding raw ensemble forecasts of DI and WBGTid.

Raw forecasts provide information on a grid (here the model grid resolution is ~ 18 km), whereas observations are point measurements. This scale mismatch leads to representativeness error in the forecast. The raw forecast uncertainty, valid at the model-resolution scale, can be increased in order to capture the temperature variability at smaller spatial scale. The method followed here is a simple down-scaling step, inspired by the pioneer work of Saeetra *et al.* (2004): it consists of adding to each member a draw from a centred Gaussian distribution with standard deviation

$$\sigma_{\text{pert}} := 0.75 + 0.18 |\Delta_e|^{1/4}. \quad (4)$$

This empirical formula is derived from a preliminary analysis of 2m temperature measurements of a high-density observation network (data description, methodology, and final derived model are presented in Ben Bouallègue 2020). The spread adjustment is applied independently to raw temperature and dew point ensemble forecasts, followed by a reordering of the adjusted forecasts using the rank of the raw forecasts before the calculation of DI and WBGTid forecasts (Figure 1). These corrected forecasts are referred to as the adjusted ensemble and are used as a reference for the computation of skill scores (Section 4.2).

3 | CALIBRATION METHODS

Our main approach to calibration is the computationally efficient EMOS method that has demonstrated excellent

predictive performance for a wide range of weather quantities. In what follows, let f_1, f_2, \dots, f_K denote the ensemble forecast of a given weather quantity for a given location, time and lead time, and \bar{f} is the ensemble mean. Further, let S^2 and MD denote the ensemble variance and the ensemble mean absolute difference, respectively, defined as

$$S^2 := \frac{1}{K-1} \sum_{k=1}^K (f_k - \bar{f})^2 \quad \text{and} \quad \text{MD} := \frac{1}{K^2} \sum_{k,\ell=1}^K |f_k - f_\ell|.$$

Multivariate ensemble forecasts are distinguished by bold notation $\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_K$, and, in this case, the ensemble variance is replaced by the ensemble covariance matrix

$$\mathbf{S}^2 := \frac{1}{K-1} \sum_{k=1}^K (\mathbf{f}_k - \bar{\mathbf{f}})(\mathbf{f}_k - \bar{\mathbf{f}})^T.$$

The 50 members of the operational ECMWF ensemble system are considered as statistically indistinguishable and, in this sense, exchangeable. In the following, if we have M ensemble members divided into K exchangeable groups, where the k th group contains $M_k \geq 1$ ensemble members ($\sum_{k=1}^K M_k = M$), notation \bar{f}_k ($\bar{\mathbf{f}}_k$) will be used for the corresponding group mean (vector).

3.1 | Normal EMOS model

Separate calibration of temperature and dew point ensemble forecasts is based on the EMOS model suggested by Gneiting *et al.* (2005). The predictive distribution is a Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$ with mean μ and variance σ^2 linked to the ensemble members via equations

$$\mu = a + b_1 f_1 + \dots + b_K f_K \quad \text{and} \quad \sigma^2 := c^2 + d^2 S^2, \quad (5)$$

where $a, b_1, \dots, b_K, c, d \in \mathbb{R}$. When an ensemble contains groups of statistically indistinguishable ensemble members, members within a given group should share the same parameters (Gneiting, 2014; Wilks, 2018) and the equation for the mean in (5) is replaced by

$$\mu = a + b_1 \bar{f}_1 + \dots + b_K \bar{f}_K. \quad (6)$$

3.2 | Bivariate normal EMOS model

The separate univariate EMOS calibration of temperature and dew point ensemble forecasts does not take into account the obvious dependence between these two weather quantities. In order to overcome this limitation, one can consider the joint post-processing of these two variables using an EMOS model with a bivariate normal predictive distribution $\mathcal{N}_2(\mu, \Sigma)$, where the mean vector μ and covariance matrix Σ are expressed as

$$\mu = \mathcal{A} + B_1 \bar{f}_1 + \dots + B_K \bar{f}_K \quad \text{and} \quad \Sigma = CC^T + DS^2D^T, \quad (7)$$

respectively. Here $\mathcal{A} \in \mathbb{R}^2$, whereas B_1, \dots, B_K and C, D are two-by-two real parameter matrices, where C is lower triangular. This way, CC^T is symmetric and non-negative definite. In the case of existence of groups of exchangeable ensemble members, the bivariate generalization of (6) is

$$\mu = \mathcal{A} + B_1 \bar{f}_1 + \dots + B_K \bar{f}_K. \quad (8)$$

Note that a bivariate EMOS model with a different parametrization of the covariance matrix has already been successfully applied for the calibration of wind vector ensemble forecasts (Schuhen *et al.*, 2012).

3.3 | Generalized extreme value EMOS model

In contrast to temperature and dew point (where the symmetric Gaussian distribution provides a reasonable fit), both DI and WBGTid observation distributions result in left-skewed histograms, calling for predictive PDFs with negative skewness. Potential candidates are the following: a skew normal, a split normal or a GEV distribution (Figure 2). In a preliminary analysis, we performed Kolmogorov–Smirnov tests for goodness of fit for each of the proposed distributions based on random samples of the size of the total number of forecast cases. The fit was rejected in all cases; however, based on the mean values of the test statistic, one can obtain a ranking of the distributions. For DI this ranking was *split normal – GEV – skew normal*, whereas for WBGTid it was *skew*

normal – GEV – split normal. Further, despite the upper limit in support caused by the negative skewness, the GEV distribution seems to best capture the tail behaviour of the observations (Figure 2). Hence, if the aim is to use the same distribution in EMOS modelling of both indices, in our case GEV will make a reasonable choice.

A GEV distribution $\mathcal{GEV}(\mu, \sigma, \xi)$ with location μ , scale $\sigma > 0$ and shape ξ can be characterized by its cumulative distribution function (CDF)

$$H(x | \mu, \sigma, \xi) := \begin{cases} \exp \left(- \left[1 + \xi \left(\frac{x - \mu}{\sigma} \right) \right]^{-1/\xi} \right), & \xi \neq 0; \\ \exp \left(- \exp \left(- \frac{x - \mu}{\sigma} \right) \right), & \xi = 0, \end{cases} \quad (9)$$

if $1 + \xi(x - \mu)/\sigma > 0$ and zero otherwise, which for $\xi < 1$ has a finite mean. For calibrating ensemble forecasts of DI and WBGTid, we suggest modelling location and scale parameters by

$$\mu = \alpha + \beta_1 \bar{f}_1 + \dots + \beta_K \bar{f}_K \quad \text{and} \quad \sigma := \gamma^2 + \delta^2 \text{MD} \quad (10)$$

with $\alpha, \beta_1, \dots, \beta_K, \gamma, \delta \in \mathbb{R}$, whereas the shape parameter ξ is considered to be independent of the ensemble members. Similar to (6) and (8), the exchangeable version of (10) considers the location μ as an affine function of the group means:

$$\mu = \alpha + \beta_1 \bar{f}_1 + \dots + \beta_K \bar{f}_K. \quad (11)$$

The expression of the mean (or location) as an affine function of the ensemble is a general feature in EMOS post-processing regardless of the weather quantity at hand (e.g., Gneiting *et al.*, 2005; Thorarinsdottir and Gneiting, 2010; Baran and Nemoda, 2016), whereas the dependence of the scale on the ensemble mean difference is similar to the censored GEV model of Scheuerer (2014) for calibrating precipitation ensemble forecasts. Exploratory tests with the dataset at hand show that the proposed model significantly outperforms the GEV EMOS model with scale depending on the ensemble mean, that is when $\sigma = \gamma^2 + \delta^2 \bar{f}$ (e.g., Lerch and Thorarinsdottir, 2013), and results in slightly (and not always significantly) better predictive performance than its natural modifications involving the ensemble variance S^2 , such as

$$\sigma = \sqrt{\gamma^2 + \delta^2 S^2} \quad \text{or} \quad \sigma = \gamma^2 + \delta^2 \sqrt{S^2}.$$

3.4 | Parameter estimation

Parameters of the normal and GEV EMOS models described in Sections 3.1 – 3.3 are estimated according to

the optimum score estimation principle of Gneiting and Raftery (2007). The estimates of the unknown parameters are obtained by optimizing the mean value of a proper scoring rule over an appropriate training dataset. The most popular proper scoring rules are the logarithmic (or ignorance) score (LogS; Good, 1952), that is the negative logarithm of the predictive PDF at the verifying observation, and the continuous ranked probability score (CRPS; Gneiting and Raftery, 2007; Wilks, 2011). The CRPS corresponding to a (predictive) CDF F and a real value (observation) x is defined as

$$\begin{aligned} \text{CRPS}(F, x) &:= \int_{-\infty}^{\infty} (F(y) - \mathbb{I}_{\{y \geq x\}})^2 dy \\ &= E|X - x| - \frac{1}{2}E|X - X'|, \end{aligned} \quad (12)$$

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 last representation in (12) implies that the CRPS can be expressed in the same unit as the observation. Both CRPS and LogS are negatively oriented scoring rules (the smaller the better), and the optimization with respect to LogS gives back the maximum likelihood (ML) estimation of the parameters. Furthermore, the CRPS for both normal and GEV distributions can be expressed in closed form (for the exact formulae, see Thorarinsdottir and Gneiting (2010) and Friederichs and Thorarinsdottir (2012), respectively), thereby allowing for efficient parameter estimation procedures often being more robust than the ML approach.

The logarithmic score is obviously well defined also in the case of multivariate predictive distributions. Further, the direct multivariate extension of the CRPS is the energy score (ES) introduced by Gneiting and Raftery (2007). Given a CDF F on \mathbb{R}^d and a d -dimensional vector \mathbf{x} , the energy score is defined as

$$\text{ES}(F, \mathbf{x}) := E\|\mathbf{X} - \mathbf{x}\| - \frac{1}{2}E\|\mathbf{X} - \mathbf{X}'\|, \quad (13)$$

where $\|\cdot\|$ denotes the Euclidean distance and, similar to the univariate case, \mathbf{X} and \mathbf{X}' are independent random vectors having distribution F . In most cases (including the case of the Gaussian model), ES cannot be given in a closed form, so it should be replaced by a Monte Carlo approximation based on a large sample drawn from F (Gneiting *et al.*, 2008). In our study, we apply ML estimation for bivariate models, whereas parameters of all univariate EMOS models are obtained by optimizing the CRPS. In the latter case, tests using ML estimation have not demonstrated significantly better results.

The appropriate choice of training data is also an important issue in statistical post-processing. In EMOS

modelling, the standard approach is the use of rolling training periods where the parameter estimates are obtained using ensemble forecasts and corresponding validating observations for the preceding n calendar days. Once the appropriate training period length is chosen, there are two general approaches to spatial selection of training data (Thorarinsdottir and Gneiting, 2010). In the local approach, model parameters are estimated using only training data of the given station, resulting in distinct parameter estimates for the different stations. The global (regional) approach uses training data of the whole ensemble domain and all observation stations share the same set of EMOS parameters. Global modelling requires a far shorter training period but is usually unsuitable for large and heterogeneous datasets like the one at hand. Alternatively, one can combine the advantages of local and global parameter estimation using a semi-local approach (Lerch and Baran, 2017): the training data for a given station is enhanced with data from stations with similar characteristics. Following Baran *et al.* (2019), clustering-based semi-local modelling where regional parameter estimation is also considered. The clusters are derived using k -means clustering of feature vectors depending both on the station climatology and the forecast errors of the raw ensemble mean over the training period.

3.5 | Ensemble copula coupling

The ensemble copula coupling (ECC) approach of Schefzik *et al.* (2013) is a general multivariate post-processing method, where after univariate calibration, the rank order information in the raw ensemble is used to restore correlations between the calibrated forecasts. In our case, we first perform univariate EMOS calibration of temperature and dew point forecasts using the normal EMOS model of Section 3.1. Then we draw N random samples of the same size as the raw ensemble from both individual EMOS predictive distributions (ECC-R; e.g., Schefzik, 2016) and reorder them to have the same rankings as the corresponding raw temperature and dew point forecasts. The obtained bivariate sample is considered as a sample from a calibrated bivariate predictive distribution of temperature and dew point.

In practice, when applying the ECC approach, 20 samples of the size of the raw ensemble (50) are drawn from the EMOS predictive distributions of temperature and dew point. Each sample is then reordered according to the rank structure of the corresponding raw ensemble forecasts resulting in 1,000 draws from the ECC predictive distribution, which can be used for forecast evaluation. In a similar way, verification scores of the bivariate EMOS

models are also estimated using 1,000 samples from the bivariate predictive distributions.

4 | VERIFICATION

4.1 | Proper scores

For a given lead time, competing forecasts in terms of probability distribution are compared with the help of the mean CRPS in the univariate case and mean ES value in the bivariate case over all forecast cases in the verification data. The skill of univariate forecasts for dichotomous events (observation x exceeding a given threshold y) is assessed with the mean Brier score (BS). For a predictive CDF $F(y)$, the Brier score is defined as

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

The integral of BS over all possible thresholds results in the CRPS. The goodness of fit of probabilistic forecasts to extreme values of the univariate weather quantity is evaluated using the threshold-weighted continuous ranked probability score (twCRPS)

$$\text{twCRPS}(F, x) := \int_{-\infty}^{\infty} (F(y) - \mathbb{I}_{\{y \geq x\}})^2 \omega(y) dy,$$

where $\omega(y) \geq 0$ is a weight function (Gneiting and Ranjan, 2011). The case where $\omega(y) \equiv 1$ corresponds to the CRPS defined in (12). In the following, we set $\omega(y) = \mathbb{I}_{\{y \geq r\}}$ in order to focus exclusively on values above a given threshold r .

4.2 | Skill scores

The improvement in terms of a given score S_F for a forecast F with respect to a reference forecast F_{ref} can be quantified using the corresponding skill score

$$S_F^{\text{skill}} := \frac{\overline{S}_{F_{\text{ref}}} - \overline{S}_F}{\overline{S}_{F_{\text{ref}}} - \overline{S}_{F_{\text{perf}}}}, \quad (14)$$

where \overline{S}_F and $\overline{S}_{F_{\text{ref}}}$ denote the mean score values corresponding to forecasts F and F_{ref} , respectively, and $\overline{S}_{F_{\text{perf}}}$ corresponds to the mean score value of a perfect deterministic forecast (with here $\overline{S}_{F_{\text{perf}}} = 0$). Based on (14), we compute the continuous ranked probability skill score (CRPSS), the energy skill score (ESS), the Brier skill score (BSS) and the threshold-weighted continuous ranked probability skill score (twCRPSS). Skill scores are positively oriented (the

larger the better), and in Section 5 the adjusted ensemble is used as a reference.

4.3 | Calibration assessment

A simple graphical tool for assessing calibration of univariate ensemble forecasts is the verification rank histogram, which is the histogram of ranks of verifying observations with respect to the corresponding ensemble forecasts (e.g., Wilks, 2011, section 8.7.2). In the case of a perfectly calibrated K -member ensemble, the ranks follow a uniform distribution on $\{1, 2, \dots, K+1\}$, and the deviation from uniformity can be quantified by the reliability index Δ defined by

$$\Delta := \sum_{r=1}^{K+1} \left| \rho_r - \frac{1}{K+1} \right|, \quad (15)$$

where ρ_r is the relative frequency of rank r (Delle Monache *et al.*, 2006). One can also generalize the verification rank histogram to multivariate ensemble forecasts. Thorarinsdottir *et al.* (2016) suggested several options in order to define ranks in a multivariate context. We apply the average ranking given by the average of the univariate ranks and the multivariate ordering proposed by Gneiting *et al.* (2008). In the case of a probabilistic forecast, one can plot the verification rank histogram and calculate a corresponding reliability index on the basis of a large number of ensembles sampled from the predictive PDF.

Calibration assessment of univariate predictive distributions is examined with the help of probability integral transform (PIT) histograms (the continuous counterparts of the verification rank histograms). By definition, PIT is the value of the predictive CDF at the validating observation, with a possible randomization at points of discontinuity (Gneiting and Ranjan, 2013). In case of perfect calibration, PIT follows a uniform law on the $[0, 1]$ interval.

4.4 | Forecast value

The forecast value compares mean expenses of users having access to different levels of information; it has the form of a skill score (Section 4.2). The framework is the one of a simple binary decision setting: facing a potential threatening event (i.e., a heat index exceeding a warning level), a user can take (or not) preventive action. This protective action has a cost C , whereas the user can endure a loss L if the event actually occurs but no protective action was taken anticipatively. The ratio C/L is called cost-loss ratio and is considered as a characteristic of the user in the decision-making process. The user takes action based on the forecast: if the forecast probability is greater than the

user's cost–loss ratio then preventive action with cost C is taken, otherwise the risk is taken to face a loss L in the case of the event. This rational behaviour is sometimes referred to as the Bayes action (Rodwell *et al.*, 2020 give a discussion on the concept and applications). When only climatological information is available, the user's cost–loss ratio is simply compared to the event base rate. In this framework, the forecast value is defined as the ratio of two mean expense differences involving the following elements: the mean expense of the user taking the Bayes action with respect to the forecast (\bar{E}_F), the mean expense of a user making a decision based on the observation climatology ($\bar{E}_{F_{\text{clim}}}$), and the mean expense of an omniscient user who is taking action only when the event actually occurs ($\bar{E}_{F_{\text{perf}}}$). Formally, the forecast value is defined as

$$V = \frac{\bar{E}_F - \bar{E}_{F_{\text{clim}}}}{\bar{E}_{F_{\text{perf}}} - \bar{E}_{F_{\text{clim}}}},$$

which can be expressed in a simple form using the elements of a contingency table (e.g., Wilks, 2001). In our implementation, the mean expenses of the different type of user are computed independently at each station. This means that we consider a climatology (estimated as the sample base rate) varying from one station to another.

5 | RESULTS

We consider two different methodological approaches to the calibration of DI and WBG Tid. In the first case, we start with joint post-processing of temperature and dew point ensemble forecasts, either using directly a bivariate normal EMOS model or applying first univariate normal EMOS models to both components and then ECC to restore the dependence structure. Samples from the bivariate predictive distributions are then used to obtain samples from the predictive distributions of DI and WBG Tid. The second approach is the direct post-processing of DI and WBG Tid ensemble forecasts with the help of the novel GEV model introduced in Section 3.3 (Figure 1).

5.1 | Bivariate calibration of temperature and dew point forecasts

We start with some considerations about the implementation of the EMOS approaches. The 50 ensemble members are regarded as exchangeable; their separate univariate and joint bivariate calibration is performed using normal EMOS models with distribution means linked to the ensemble means via (6) and (8), respectively, with

$K = 1$. Thus, the EMOS models require the estimation of four free parameters in the univariate case and 13 in the bivariate case. A large number of free parameters calls either for long training periods in local estimation or for a semi-local approach. Both local and clustering-based semi-local parameter estimations with 5, 10, 15, 25, 50 and 75 clusters are tested. Clustering of stations is based on 40 dimensional features equally representing the properties of both temperature and dew point forecasts and observations. In accordance with the results of Lerch and Baran (2017), the forecast skill does not really depend on the number of features applied, provided that one works with a sufficiently large number of features. Moreover, the longer the lead time of the post-processed forecasts, the more training data are needed to outperform the raw ensemble forecasts (see discussion in Feldmann *et al.*, 2019). As a trade-off, the length of the training period is fixed to 60 days. This choice leaves 79 calendar days (period 14 July–30 September 2017) as a verification period when taking into account the maximal lead time of 15 days.

First, calibration of temperature and dew point forecast is performed separately. Figure 3 shows the CRPSS of different semi-local EMOS models with respect to local EMOS. Local post-processing is optimal up to day 5, then semi-local approaches with various number of clusters start to dominate, exhibiting significantly better predictive performance. Hence, we propose to apply different types of EMOS configurations for different forecast lead times: local EMOS at the beginning of the forecast range and semi-local configuration(s) for longer lead times. Practically, local EMOS is considered for days 1–5 both for temperature and dew point, while for days 6–10 we use semi-local EMOS with 75 clusters. For dew point, this latter model is also kept for days 11–15, whereas for temperature we reduce the number of clusters to 5. The different mixtures of EMOS configurations are reported in Table 1. In the following, the method referred to as ECC relies on the reported mixed approach for the univariate post-processing step.

In the same spirit as Figure 3, the optimal training configuration is investigated by inspecting the multivariate performance in terms of ESS for the semi-local bivariate EMOS models with respect to the local one. As reported in Table 1, optimal performance in terms of ES can be obtained by applying local modelling for days 1–5, semi-local with 75 clusters for days 6–10, and semi-local with five clusters for longer lead times (referred to as 2D EMOS optimal). However, this mixture of EMOS configurations is sub-optimal when applied to derive predictive distributions of DI and WBG Tid (Section 5.2). Better forecast skill can be obtained using local EMOS for days 1–4, semi-local with 50 clusters for days 5–7, and

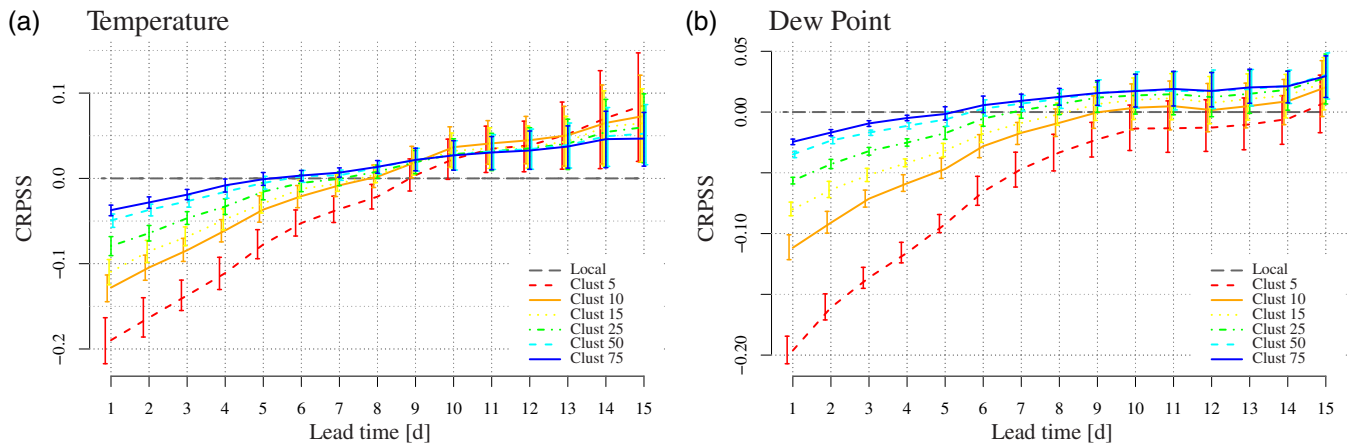


FIGURE 3 CRPSS of semi-local EMOS models with respect to the local EMOS for (a) temperature and (b) dew point together with 95% confidence intervals

TABLE 1 Training configurations of EMOS models for different forecast lead times

Post-processing method	Lead time (days)														
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Normal EMOS, temperature	⇐		Local		⇒	⇐		75 Clusters		⇒	⇐		5 Clusters		⇒
Normal EMOS, dew point	⇐		Local		⇒	⇐				75 Clusters					⇒
2D EMOS	⇐		Local		⇒	⇐		50 Clusters		⇒	⇐		5 Clusters		⇒
2D EMOS optimal	⇐		Local		⇒	⇐		75 Clusters		⇒	⇐		5 Clusters		⇒
GEV EMOS, DI and WBG Tid	⇐		Local		⇒	⇐		75 Clusters		⇒	⇐		5 Clusters		⇒

semi-local with five clusters until the end of the forecast range (referred to as 2D EMOS). For days 6–10 this means for example, lower mean CRPS, whereas at day 5 slightly (<1 %) higher mean CRPS but much lower reliability indices both for DI and WBG Tid.

In Figure 4, the different multivariate post-processing approaches are compared in terms of a multivariate skill score, ESS. Bivariate EMOS and ECC post-processing approaches significantly outperform both the raw ensemble and the adjusted ensemble up to day 9, while the advantage of the adjusted ensemble with respect to the raw ensemble disappears only after day 10. All performance differences shrink with forecast lead time. Bivariate EMOS generally outperforms ECC but both methods have similar performance between day 5 and day 10. When significant difference in mean ES between these two approaches is tested, it appears that the proportion of stations with significant difference (positive or negative) is less than 7% at day 1 but increases to more than 30% at day 10 (not shown).

Reliability of the bivariate forecast is now investigated. Figure 5 presents reliability indices (corresponding to average ranking and multivariate ordering) as a function of the forecast lead time. The reliability indices are




FIGURE 4 ESS with respect to the adjusted ensemble of bivariate forecasts together with 95% confidence intervals

rather consistent with one another: similar patterns are found for the two types of ranking (average and multivariate). One exception is the difference between raw and adjusted ensemble results which is much smaller when average ranks are considered. The raw ensemble results in strongly U-shaped histograms at short lead times, while the ensemble adjustment allows us to improve

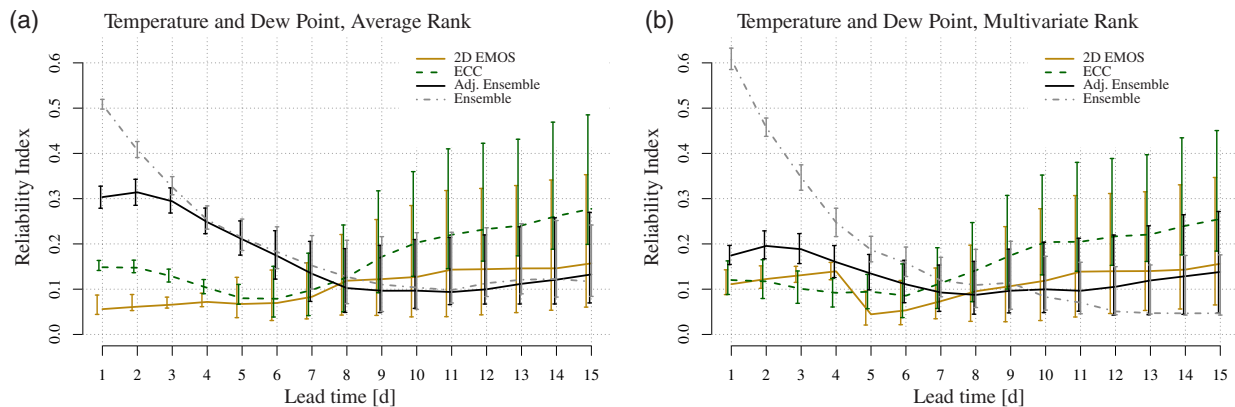


FIGURE 5 Reliability indices of bivariate forecasts corresponding to (a) average ranks and (b) multivariate ordering together with 95% confidence intervals

reliability in particular with respect to the multivariate ranking (Figures A1 and A2). 2D EMOS exhibits the better reliability (lowest reliability indices) up to day 8 and fairly good results for longer lead times. However, EMOS post-processed forecasts are slightly biased towards the end of the forecast range, producing low multivariate ranks of the verifying observations (Figures A1 and A2). ECC forecasts are highly hump-shaped at short lead times, rather well calibrated at medium lead times, and U-shaped for long lead times. Post-processing approaches applied here fail to improve reliability of the ensemble beyond day 10, which is in line with their predictive performance in terms of the ES.

In a nutshell, for short and medium lead times, bivariate post-processing can improve multivariate forecasts of temperature and dew point, and in particular reliability, in a multivariate sense, is increased. However, the post-processing methods applied here do not explicitly account for any relationship between temperature and dew point respecting the laws of physics. This limitation results sometimes in situations where a dew point forecast is greater than a temperature forecast. In that case, the former is set equal to the latter in order to avoid physical inconsistencies and allow the computation of the heat indices. This situation occurs less than 3% of the cases when applying 2D EMOS and can be more frequent when simply adjusting the ensemble forecasts.

5.2 | Post-processing of discomfort index and indoor wet-bulb globe temperature forecasts

Calibrated DI and WBGT_{id} forecasts can be derived from the jointly calibrated temperature and dew point forecasts, but DI and WBGT_{id} ensemble forecasts can also be directly post-processed using the GEV EMOS model introduced in

Section 3.3. The GEV model has only five free parameters in the case of a single group of exchangeable ensemble members. Considering a 60-day rolling training period and the same verification period 14 July–30 September 2017 as before, we inspect the CRPSS of various semi-local EMOS configurations with respect to local EMOS for DI and WBGT_{id} (not shown). Based on this exploratory analysis, we propose to adapt the EMOS configuration as a function of the forecast lead time as follows: for both indices local EMOS is applied between days 1 and 4, semi-local EMOS with 75 and 5 clusters for days 5–8 and 9–15, respectively. This model is referred to as GEV EMOS (Table 1).

In Figure 6, results in terms of CRPSS indicate that post-processing significantly improves both the raw and adjusted ensembles up to day 7 for DI and day 6 for WBGT_{id}. For both indices, 2D EMOS post-processed forecasts have the best predictive skill, followed closely by GEV EMOS forecasts, whereas ECC shows the worst predictive performance among the calibration approaches. However, the differences between models are minor and often non-significant. For lead times up to day 7, the differences in CRPS between the competing EMOS forecasts are significant at a 5% level for less than 6% of the stations. This number increases to more than 15% when differences are computed against the adjusted ensemble forecasts (not shown).

Verification rank and PIT histograms help diagnosing reliability issue in the forecast. Results for DI are shown in Figure A3, while histograms for WBGT_{id}, which are very similar, are not shown. For both indices, the raw ensemble forecasts are highly underdispersive for day 1; however, the dispersion improves with the forecast lead time. Moreover, raw forecasts exhibit a bimodal character, which is in line with the bivariate histograms of raw temperature and dew point forecasts. Only local GEV EMOS calibration removes this bimodality and in general, similar to bivariate forecasts, the longer the lead time, the smaller the effect of post-processing.

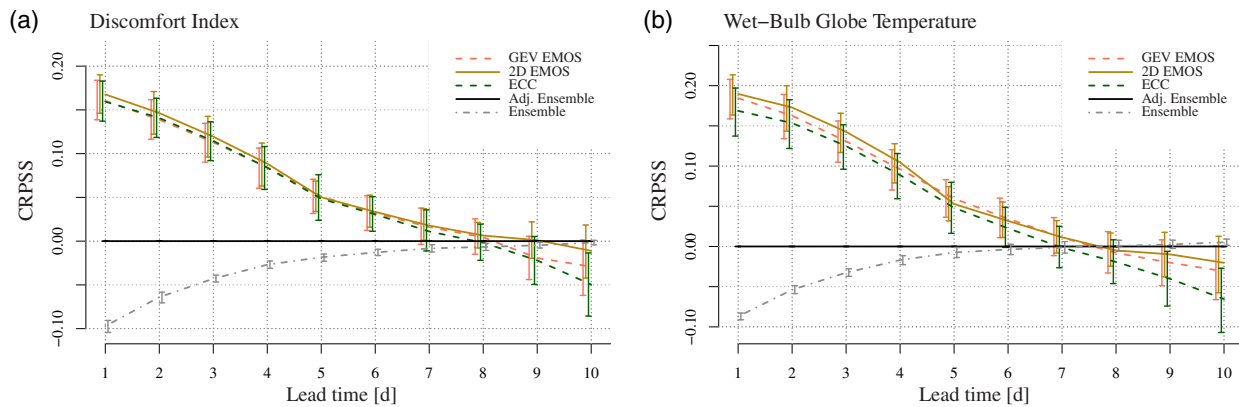


FIGURE 6 CRPSS with respect to the adjusted ensemble of (a) DI and (b) WBGTid forecasts together with 95% confidence intervals

Another perspective is now taken focusing on the forecast skill at the right tail of the observation distributions. Figure 7a,b depicts BSS and twCRPSS considering the event DI exceeding $r = 27^\circ\text{C}$. This threshold corresponds to the level at which most of the population suffers from heat-induced discomfort (Stathopoulou *et al.*, 2005). In terms of BSS, all EMOS approaches outperform the raw and adjusted ensemble up to day 10 with comparable skill. However, in terms of twCRPS, the GEV EMOS model clearly outperforms the other post-processing methods. For lower thresholds, the differences between the EMOS approaches is less pronounced, while for higher thresholds the results are noisier and the significance bars are larger (not shown).

For WBGTid, a threshold $r = 27.8^\circ\text{C}$ is considered. It corresponds to the threshold associated with the first warning level (green flag) (Patel *et al.*, 2013). Figure 7c,d shows BSS and twCRPSS: post-processed forecasts significantly outperform the raw and adjusted ensemble forecasts up to at least day 8, whereas ECC exhibits positive skill score up to day 10. Considering larger thresholds, small BS and twCRPS values combined with large uncertainty in the scores leads to quite noisy results which are difficult to interpret (not shown). The change in EMOS configuration (from local to semi-local) is clearly visible in Figure 7: jumps appear at days 5 and 8 for 2D EMOS, at days 5 and 9 for GEV EMOS, while the curves corresponding to ECC remain smooth.

One more step towards refined performance analysis is taken by looking at the forecast value. For the two events considered above, forecast value is plotted as a function of the user cost-loss ratio in Figure 8. As a general feature, not only the forecast value, but also the number of users that can benefit from the forecast, decreases with the forecast lead time. At day 10, the forecast of WBGTid has no value for a user who needs high certainty about the event occurrence (high cost-loss ratios), but it is still valuable for a user who is risk adverse (low cost-loss ratios). The

post-processed forecasts have value for some users even beyond day 10 (not shown). As a purely statistical characteristic of the score, the forecast value reaches a maximum for a user with cost-loss ratio equal to the climatological base rate. This is true locally, but the maximum value varies from one location to another and the corresponding pick is therefore smoothed out when aggregated over different locations. In any case, the event for DI is more probable than the event for WBGTid, which explains that the maximum is reached for different C/L values in the two cases. The raw ensemble has negative value for extreme C/L for both DI and WBGTid forecasts at 4 and 10 days, but the ensemble adjustment allows us to make the forecasts valuable for (almost) all types of users. Forecast post-processing increases the forecast value further in most of the cases. There are also slight differences in terms of forecast value between the different calibration approaches.

5.3 | Predictive distributions of categories of heat indices

Heat indices are often grouped into categories according to the severity of the effect of heat stress on human activity. Table 2 provides a classification of DI and heat categories of WBGTid based on Stathopoulou *et al.* (2005) and Patel *et al.* (2013), respectively. Instead of continuous heat indices, we focus now on the corresponding categories taking values $\{1, 2, 3, 4, 5, 6\}$ for DI and $\{1, 2, 3, 4, 5\}$ for WBGTid. In this case, a predictive distribution is specified by a probability mass function (PMF) and post-processing reduces to a classification problem, which is a typical application area of machine learning methods. Here, predictive distributions are obtained with the help of multilayer perceptron neural networks (MLP NN; Hastie *et al.*, 2009). We skip the intermediate step of obtaining DI and WBGTid ensemble forecasts from the

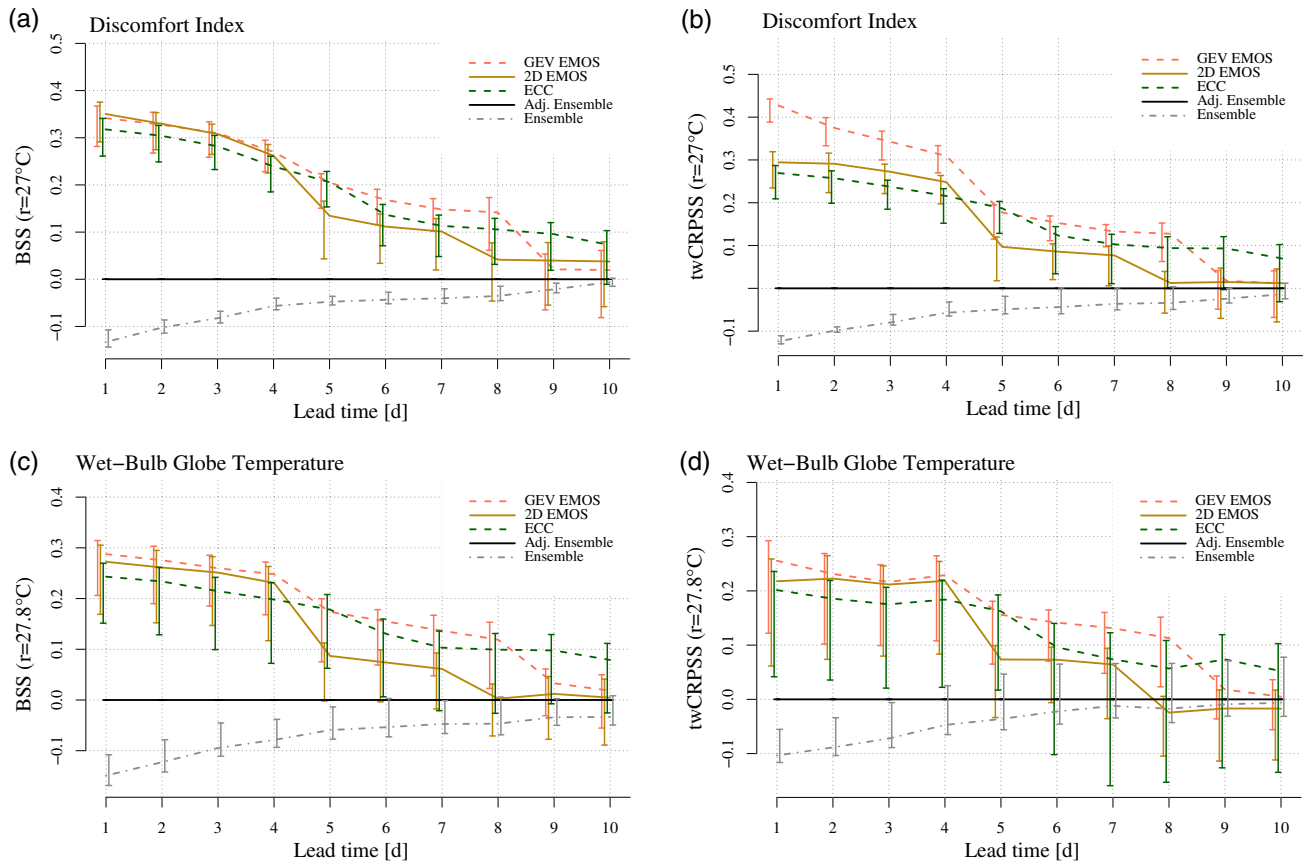


FIGURE 7 (a, c) BSS and (b, d) twCRPSS with respect to the adjusted ensemble corresponding to thresholds (a, b) $r = 27^\circ\text{C}$ for the DI and (c, d) $r = 27.8^\circ\text{C}$ for the WBGTid forecasts together with 95% confidence intervals

corresponding forecasts of temperature and dew point, and connect the latter directly to the DI and WBGTid categories. As primary input variables (also called features or covariates), we use the mean, variance and covariance of temperature and dew point ensemble forecasts, as well as the geographical coordinates (latitude, longitude), and the elevations of the SYNOP stations and of the corresponding model grid point. Instead of applying local or semi-local training, all forecast cases in the training periods are regarded as a single group. This nine-dimensional feature vector can be extended further with forecast errors of temperature and dew point. Ensemble mean errors of the 1 to 5 closest days to the end of the training period can be added to the input variables, increasing the number of covariates by 2 to 10 more features. In this case, the length of the training period is naturally reduced by the number of days where the forecasts errors are considered.

Predictive PMFs of the DI and WBGTid categories are obtained applying MLP NNs with a single hidden layer and cross-entropy as loss function. In our case-study, the different categories are highly imbalanced: the proportions of observations in the first category in Table 2 is $\sim 67\%$ for DI and $\sim 99\%$ for WBGTid, whereas the last category in Table 2

represents less than 1% in both cases. For this reason, we also consider a two-step procedure (referred to as dual net). In the first step a binary classification is performed, which for each feature vector results in the distribution $(p_1, 1 - p_1)$ of the binary outcome of belonging to the first (richest) category or not. Then a second net is trained using only feature vectors coming from the remaining κ classes, where $\kappa = 5$ for DI and $\kappa = 4$ for WBGTid. For each feature set, this step yields a distribution (q_1, \dots, q_κ) of the corresponding classes and the final distribution of index categories is obtained as $(p_1, (1 - p_1)q_1, \dots, (1 - p_1)q_\kappa)$.

Similar to the EMOS approaches, different training configurations are envisaged for different lead times. We tested the performance of single nets with 15, 20, 25, 30, 35, 40 and 50 neurons in the hidden layer, both with the nine-dimensional primary features and with extended feature vectors containing forecast errors of either 1 or 5 days. In the case of dual nets, we used 15, 25 and 40 neurons in the hidden layer of both components and considered the primary features plus the 1-day forecast error. In general, a large number of neurons leads to overfitting, and according to our observations, for longer lead times the use of forecast errors does not improve the

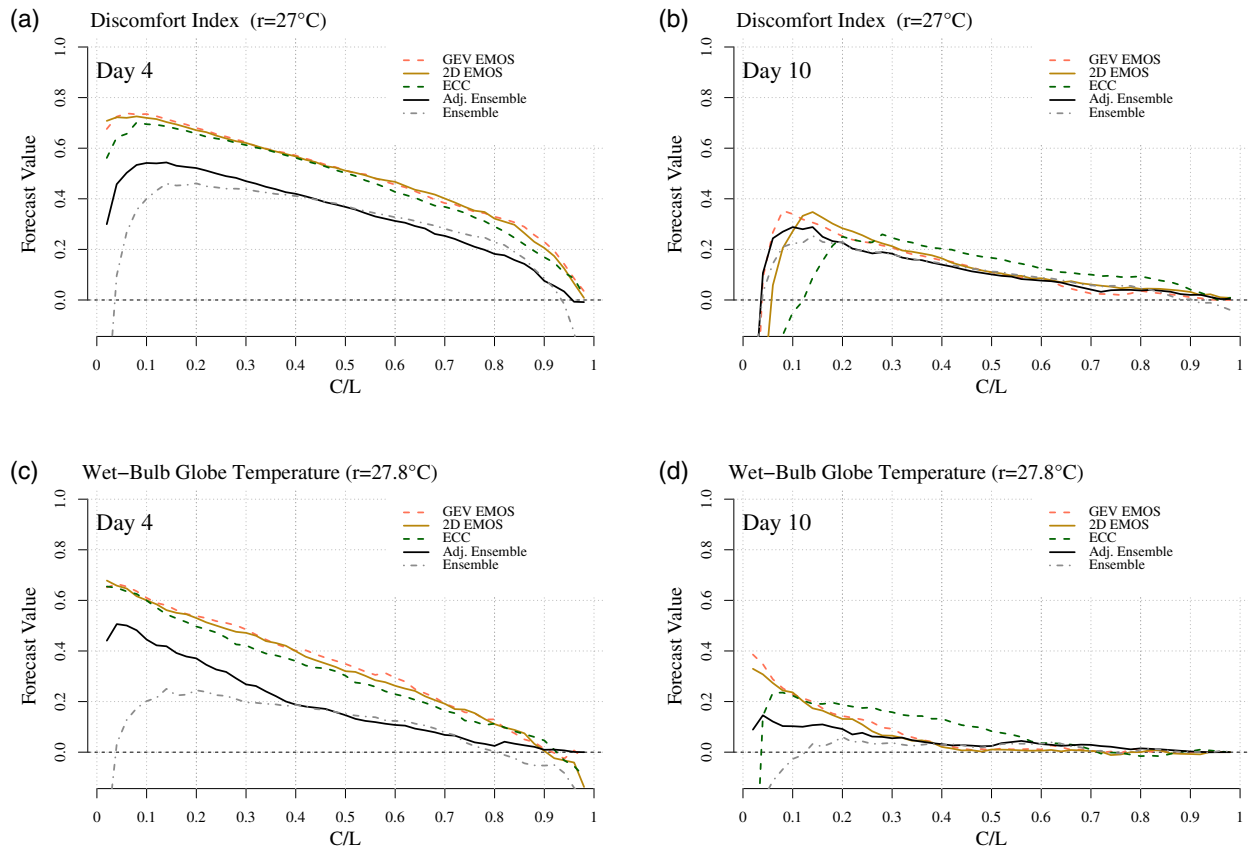


FIGURE 8 Aggregated forecast value for (a, c) day 4 and (b, d) day 10 corresponding to thresholds (a, b) $r = 27^\circ\text{C}$ for the DI and (c, d) $r = 27.8^\circ\text{C}$ for the WBGTid forecasts

TABLE 2 Classification of DI and heat categories of WBGTid based on Stathopoulou *et al.*, (2005) and Patel *et al.*, (2013), respectively

Discomfort Index (I_d)			Wet-Bulb Globe Temperature (T_{WBGTid})		
Category	Range ($^\circ\text{C}$)	Classification	Category	Range ($^\circ\text{C}$)	Flag colour
1	$I_d < 21$	No discomfort	1	$T_{\text{WBGTid}} < 27.8$	No flag
2	$21 \leq I_d < 24$	Under 50% population feels discomfort	2	$27.8 \leq T_{\text{WBGTid}} < 29.5$	Green
3	$24 \leq I_d < 27$	Over 50% population feels discomfort	3	$29.5 \leq T_{\text{WBGTid}} < 31.1$	Yellow
4	$27 \leq I_d < 29$	Most of population feels discomfort	4	$31.1 \leq T_{\text{WBGTid}} < 32.2$	Red
5	$29 \leq I_d < 32$	Everyone feels severe stress	5	$T_{\text{WBGTid}} \geq 32.2$	Black
6	$I_d \geq 32$	State of medical emergency			

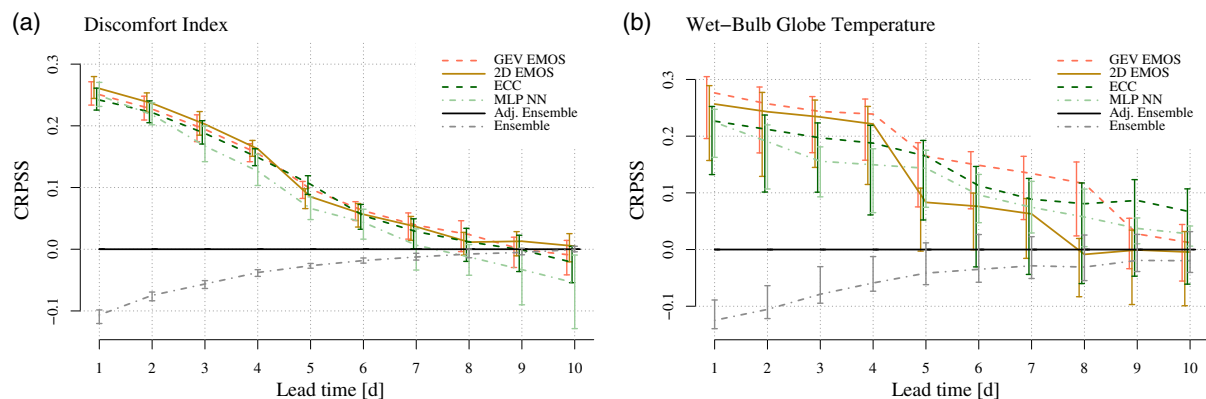
predictive performance, while the dual net can handle better the larger uncertainty in the forecasts. For DI categories, the following configuration is selected: a single net with 40 neurons making use of the forecast errors for 5 days as additional predictors for days 1–4, single net with 15 neurons for days 5–8, and dual net with 15–15 neurons for days 9–15 (Table 3). For WBGTid categories, single networks with medium numbers of neurons are preferred, and moreover, the use of forecast errors has a longer-lasting effect. So, the selected configuration

is the following: for days 1–8 we consider 30 neurons and 5-day forecast errors, for days 9–12 only the primary features and 35 neurons, whereas for the last 3 days we just reduce the number of neurons to 30 (Table 3). In the following discussion these two blends are referred to as MLP NN.

In Figure 9, CRPS is applied to forecast and observation categories. All calibration approaches significantly outperform both the raw and the adjusted forecasts of DI categories till day 6. The results for the different

TABLE 3 MLP NN configurations for different lead times

MLP NN configurations		Lead time (days)															
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
DI	Net	⇐	Single		⇒	⇐	Single		⇒	⇐			Dual			⇒	
	Neurons in hidden layer	⇐	40		⇒	⇐	15		⇒	⇐			15-15			⇒	
	Forecast errors	⇐	5		⇒	⇐	0		⇒	⇐			0			⇒	
WBG Tid	Net	⇐			Single				⇒	⇐	Single		⇒	⇐	Single		⇒
	Neurons in hidden layer	⇐			30				⇒	⇐	35		⇒	⇐	30		⇒
	Forecast errors	⇐			5				⇒	⇐	0		⇒	⇐	0		⇒

**FIGURE 9** CRPS with respect to the adjusted ensemble of (a) DI and (b) WBG Tid categories together with 95% confidence intervals

EMOS models are in line with those of the continuous forecasts (Figure 6a). The robust and rather generic MLP NN method is outperformed by the parametric models; however, the differences are often non-significant. A different ranking of the calibration methods can be observed for the WBG Tid categories in Figure 9b, which shows a high level of similarities with BSS and twCRPS plots in Figure 7c,d. The GEV and 2D EMOS forecast skill jumps at the lead times where the training configurations are changed. After day 7, 2D EMOS loses its skill against the adjusted ensemble while the machine learning approach shows a reasonable predictive performance up to day 10. Among the different post-processing approaches, MLP NN provides the sharpest predictive PMF in terms of the average width of the 90% central predictive intervals and in terms of variance for longer lead times too (not shown). Due to the extremely imbalanced nature of WBG Tid observation categories, during the training of a neural net one can easily face the situation where some categories are completely missing from the training data. This problem is less pronounced for the other investigated post-processing methods as the EMOS predictive PMFs are calculated using continuous forecasts and observations. Thus, the WBG Tid results of the different classification approaches should be interpreted cautiously.

6 | CONCLUSION

This study focuses on improving forecast heat indices such as discomfort index (DI) and indoor wet-bulb globe temperature (WBG Tid) driven by temperature and dew point forecasts as inputs through statistical calibration. Post-processing based on ensemble model output statistics (EMOS) approaches enables us to significantly improve continuous forecasts of heat indices, at least up to day 6. In particular, DI and WBG Tid forecast reliability is clearly increased in that case. Among different flavours of EMOS approaches, the generalized extreme value (GEV) model has the best overall performance, followed by the bivariate EMOS and ensemble copula coupling. The GEV EMOS is tailored to the end-products (DI and WBG Tid), whereas calibration based on joint post-processing of temperature and dew point forecasts can be applied to any heat index depending on these two weather quantities. For longer forecast ranges (typically beyond the first week in the forecast), post-processing does not improve the forecast and can even lead to a deterioration of the forecast skill which is explained by the non-stationarity of the forecast errors.

Moving from continuous forecasts to predicted categories of heat indices, we further compare the traditional

approach of post-processing to multilayer perceptron neural networks, demonstrating that both methods have comparable predictive power. Furthermore, we evidenced that local EMOS post-processing is outperformed by cluster-based EMOS approaches after day 4 if a limited training period is used. Decreasing the number of clusters with the forecast lead time helps to extend the benefit of EMOS post-processing for longer lead times. As a limit, a single-cluster approach is equivalent to a global training approach which is not suitable for forecasting over a large domain.

The present work highlights several avenues of potential future research. For example, as an alternative training dataset to the one above, one could consider the use of reforecasts. This is not investigated here, but should be considered to address efficiently the calibration of medium- to extended-range forecasts. When accessing large training datasets, appropriate predictors could be defined to account for seasonality and/or weather regime-dependent forecast errors. In this situation, machine learning approaches seem more flexible to deal with an increased number of predictors.

The assessment of forecast performance spans from the multivariate ES to the forecast value for a range of users defined by their cost–loss ratio, from the generic univariate CRPS to a focus on extreme events with BS and twCRPS. The cascade of verification results indicate that the ranking of the post-processing methods might differ from one score to another: the best method for one aspect of the forecast performance is not necessary the best solution for all applications. While an improvement with EMOS post-processing is expected to be effective up to day 6, the ensemble demonstrates potential skill beyond day 10 in terms of forecast value for some users.



The comparison of bivariate calibration approaches with direct post-processing leads us to formulate the following recommendations. Direct post-processing is generally more efficient; targeting one specific event/user could be even more appropriate. However, post-processing of the input variables appears to be a competitive approach. In that case, calibration of the dependence structure gives the best results. The simplest approach, which consist in the calibration of each component separately combined using the empirical copula of the ensemble, has reasonable performance as well. For more complex heat indices, calibration of the input variables can be envisaged if no EMOS model exists for the end-product but, as the above discussion tends to indicate, there is no royal road to post-processing of heat indices.

ACKNOWLEDGMENTS

Sándor Baran was supported by the Hungarian National Research, Development and Innovation Office under grant

no. NN125679. Ágnes Baran and Sándor Baran also acknowledge the support of the EFOP-3.6.2-16-2017-00015 project. The project was co-financed by the Hungarian Government and the European Social Fund. The authors are indebted to Claudia Di Napoli and David S. Richardson for their valuable suggestions and remarks. Last but not least, the authors are grateful to the reviewers and the editor for their valuable comments.

ORCID

Sándor Baran  <https://orcid.org/0000-0003-1035-004X>
Zied Ben Bouallègue  <https://orcid.org/0000-0003-2914-4203>

REFERENCES

- Baran, S. and Möller, A. (2017) Bivariate ensemble model output statistics approach for joint forecasting of wind speed and temperature. *Meteorology and Atmospheric Physics*, 129, 99–112.
- Baran, S. and Nemoda, D. (2016) Censored and shifted gamma distribution based EMOS model for probabilistic quantitative precipitation forecasting. *Environmetrics*, 27, 280–292.
- Baran, S., Leutbecher, M., Szabó, M. and Ben Bouallègue, Z. (2019) Statistical post-processing of dual-resolution ensemble forecasts. *Quarterly Journal of the Royal Meteorological Society*, 145, 1705–1720.
- Barnes, C., Brierley, C.M. and Chandler, R.E. (2019) New approaches to postprocessing of multi-model ensemble forecasts. *Quarterly Journal of the Royal Meteorological Society*, 145, 3479–3498.
- Ben Bouallègue, Z. (2020) Accounting for representativeness in the verification of ensemble forecasts. Technical Memorandum 865, ECMWF, Reading, UK.
- Ben Bouallègue, Z., Heppelmann, T., Theis, S.E. and Pinson, P. (2016) Generation of scenarios from calibrated ensemble forecasts with a dual ensemble copula coupling approach. *Monthly Weather Review*, 144, 4737–4750.
- Bernard, T.E. and Pourmoghani, M. (1999) Prediction of workplace wet bulb global temperature. *Applied Occupational and Environmental Hygiene*, 14, 126–134.
- Casanueva, A. (2019) HeatStress. <https://doi.org/10.5281/zenodo.3264929>; accessed 3 April 2020.
- de Freitas, C.R. and Grigorieva, E.A. (2017) A comparison and appraisal of a comprehensive range of human thermal climate indices. *International Journal of Biometeorology*, 61, 487–512.
- Delle Monache, L., Hacker, J.P., Zhou, Y., Deng, X. and Stull, R.B. (2006) Probabilistic aspects of meteorological and ozone regional ensemble forecasts. *Journal of Geophysical Research; Atmospheres*, 111(D24). <https://doi.org/10.1029/2005JD006917>
- Di Napoli, C., Pappenberger, F. and Cloke, H.L. (2019) Verification of heat stress thresholds for a health-based heat-wave definition. *Journal of Applied Meteorology and Climatology*, 58, 1177–1194.
- EM-DAT (2020) *The International Disaster Database*. Available at: <http://www.emdat.be>; accessed 3 April 2020. Centre for Research on the Epidemiology of Disasters (CRED), Université Catholique de Louvain, Brussels.
- Feldmann, K., Richardson, D.S. and Gneiting, T. (2019) Grid- versus station-based postprocessing of ensemble temperature forecasts. *Geophysical Research Letters*, 46, 7744–7751.

- Friederichs, P. and Thorarinsdottir, T.L. (2012) Forecast verification for extreme value distributions with an application to probabilistic peak wind prediction. *Environmetrics*, 23, 579–594.
- Fundel, V.J., Fleischhut, N., Herzog, S.M., Göber, M. and Hagedorn, R. (2019) Promoting the use of probabilistic weather forecasts through a dialogue between scientists, developers and end-users. *Quarterly Journal of the Royal Meteorological Society*, 145, 210–231.
- Gneiting, T. (2014) *Calibration of medium-range weather forecasts*. Technical Memorandum 719, ECMWF, Reading, UK.
- 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. and Ranjan, R. (2011) Comparing density forecasts using threshold- and quantile-weighted scoring rules. *Journal of Business and Economic Statistics*, 29, 411–422.
- Gneiting, T. and Ranjan, R. (2013) Combining predictive distributions. *Electronic Journal of Statistics*, 7, 1747–1782.
- 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., Stanberry, L.I., Grit, E.P., Held, L. and Johnson, N.A. (2008) Assessing probabilistic forecasts of multivariate quantities, with applications to ensemble predictions of surface winds (with discussion and rejoinder). *Test*, 17, 211–264.
- Good, I.J. (1952) Rational decisions. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 14, 107–114.
- Hastie, T., Tibshirani, R. and Friedman, J. (2009) *The Elements of Statistical Learning. Data Mining, Inference, and Prediction* (2nd ed.) Springer, New York, NY.
- Lemke, B. and Kjellstrom, T. (2012) Calculating workplace WBGT from meteorological data: a tool for climate change assessment. *Industrial Health*, 50, 267–278.
- Lerch, S. and Baran, S. (2017) Similarity-based semi-local estimation of EMOS models. *Journal of the Royal Statistical Society series C – Applied Statistics*, 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. and Palmer, T.N. (2008) Ensemble forecasting. *Journal of Computational Physics*, 227, 3515–3539.
- McGregor, G., Bessemoulin, P., Ebi, K. and Menne, B. (Eds.) (2015) *Heatwaves and Health: Guidance on Warning-system Development*. World Meteorological Organisation and World Health Organisation, Geneva, Switzerland.
- Morabito, M., Messeri, A., Noti, P., Casanueva, A., Crisci, A., Kotlarski, S., Orlandini, S., Schwierz, C., Spirig, C., Kingma, B.R., Flouris, A.D. and Nybo, L. (2019) An occupational heat–health warning system for Europe: the heat-shield platform. *International Journal of Environmental Research and Public Health*, 16(2890). <https://doi.org/10.3390/ijerph16162890>
- Pappenberger, F., Jendritzky, G., Staiger, H., Dutra, E., Di Giuseppe, F., Richardson, D.S. and Cloke, H.L. (2015) Global forecasting of thermal health hazards: the skill of probabilistic predictions of the universal thermal climate index (UTCI). *International Journal of Biometeorology*, 59, 311–323.
- Patel, T., Mullen, S.P. and Santee, W.R. (2013) Comparison of methods for estimating wet-bulb globe temperature index from standard meteorological measurements. *Military Medicine*, 178, 926–932.
- 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.
- Rodwell, M.J., Hammond, J., Thornton, S. and Richardson, D.S. (2020) User decisions, and how these could guide developments in probabilistic forecasting. *Quarterly Journal of the Royal Meteorological Society*, <https://doi.org/10.1002/qj.3845>
- Saetra, Ø., Hersbach, H., Bidlot, J.R. and Richardson, D.S. (2004) Effects of observation errors on the statistics for ensemble spread and reliability. *Monthly Weather Review*, 132, 1487–1501.
- Schefzik, R. (2016) Combining parametric low-dimensional ensemble postprocessing with reordering methods. *Quarterly Journal of the Royal Meteorological Society*, 142, 2463–2477.
- Schefzik, R., Thorarinsdottir, T.L. and Gneiting, T. (2013) Uncertainty quantification in complex simulation models using ensemble copula coupling. *Statistical Science*, 28, 616–640.
- Scheuerer, M. (2014) Probabilistic quantitative precipitation forecasting using ensemble model output statistics. *Quarterly Journal of the Royal Meteorological Society*, 140, 1086–1096.
- Schuhen, N., Thorarinsdottir, T.L. and Gneiting, T. (2012) Ensemble model output statistics for wind vectors. *Monthly Weather Review*, 140, 3204–3219.
- Sonntag, D. (1990) Important new values of the physical constants of 1986, vapour pressure formulations based on the IST-90 and psychrometer formulae. *Meteorologische Zeitschrift*, 70, 340–344.
- Stathopoulou, M.I., Cartalis, C., Keramitsoglou, I. and Santamouris, M. (2005). Thermal remote sensing of Thom's discomfort index (DI): comparison with *insitu* measurements, in Proceedings of SPIE 5983, Remote Sensing for Environmental Monitoring, GIS Applications, and Geology V, DOI 10.1117/12.627541, (to appear in print).
- 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. Statistical Society*, 173, 371–388.
- Thorarinsdottir, T.L., Scheuerer, M. and Heinz, C. (2016) Assessing the calibration of high-dimensional ensemble forecasts using rank histograms. *Journal of Computational and Graphical Statistics*, 25, 105–122.
- Wilks, D.S. (2001) A skill score based on economic value for probability forecasts. *Meteorological Applications*, 8, 209–219.
- Wilks, D.S. (2011) *Statistical Methods in the Atmospheric Sciences* (3rd ed.) Elsevier, Amsterdam, Netherlands.
- Wilks, D.S. (2018) Univariate ensemble postprocessing. Chapter 3 in Statistical Postprocessing of Ensemble Forecasts. Vannitsem, S., Wilks, D.S., Messner, J.W. (eds). Elsevier, Amsterdam, Netherlands.

How to cite this article: Baran S, Baran Á, Pappenberger F, Ben Bouallègue Z. Statistical post-processing of heat index ensemble forecasts: Is there a royal road?. *Q.J.R. Meteorol. Soc.* 2020;1–19. <https://doi.org/10.1002/qj.3853>

APPENDIX

Histograms

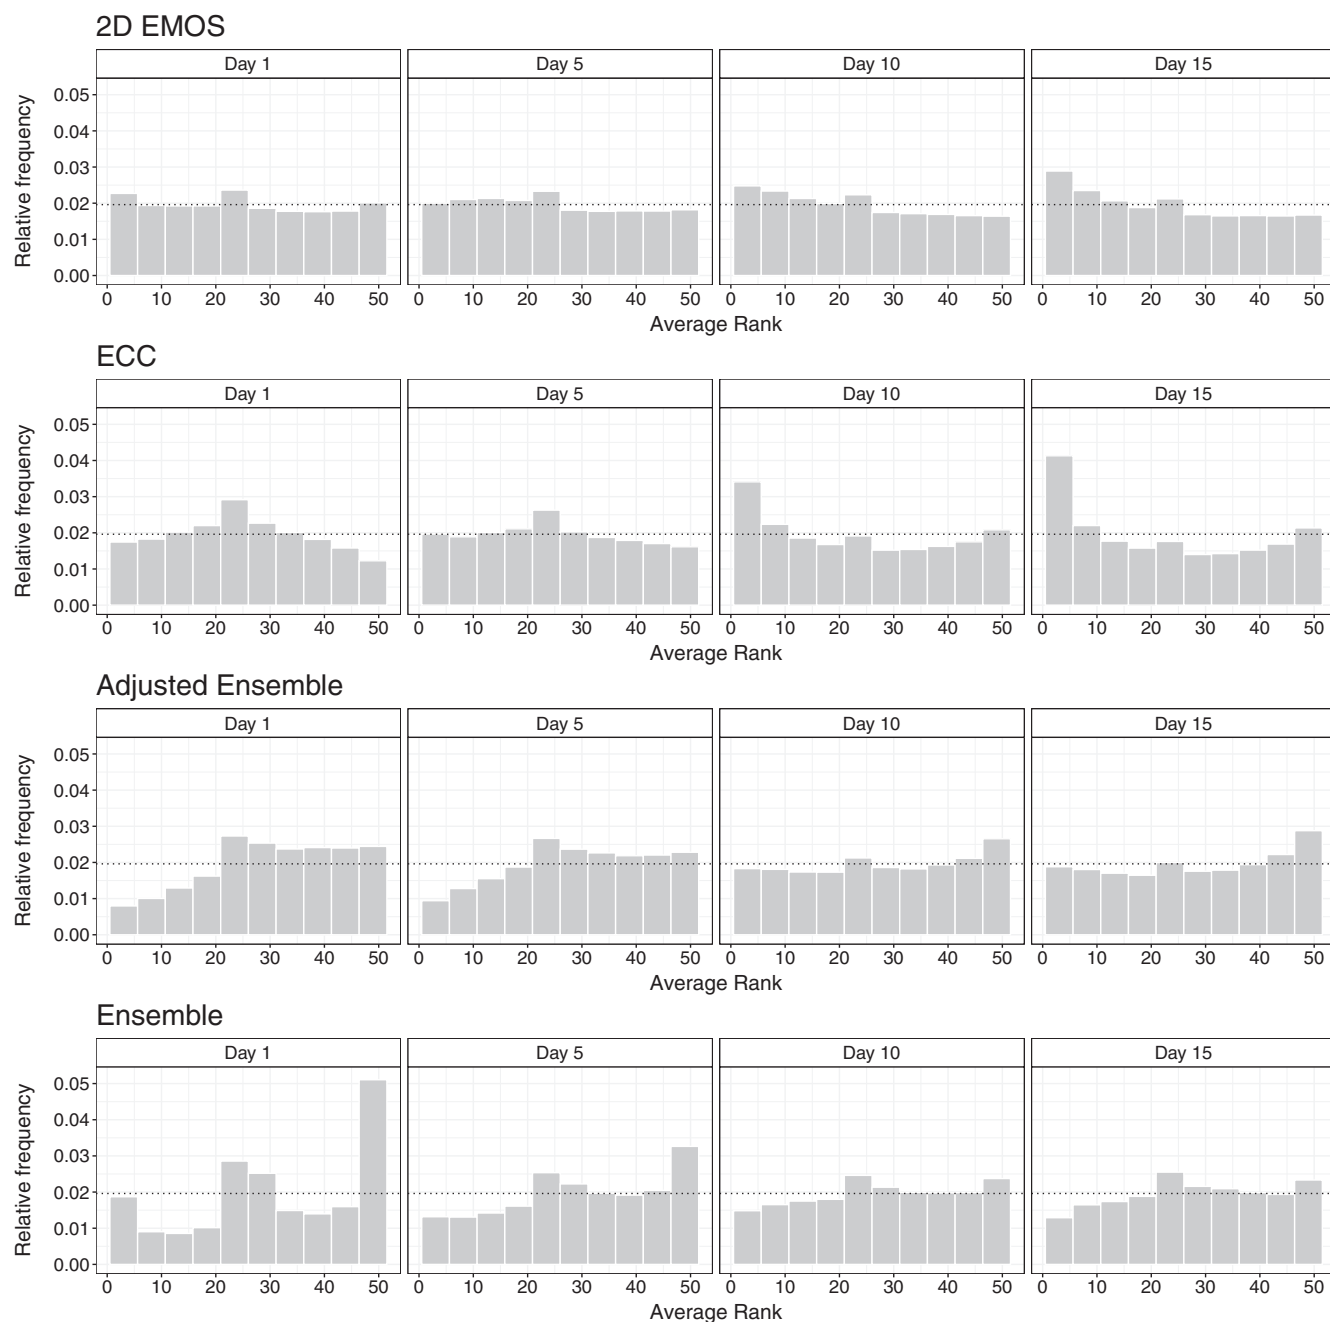




FIGURE A1 Rank histograms of bivariate forecasts with respect to average ranks for days 1, 5, 10 and 15. Figure 5a shows the corresponding reliability indices


2D EMOS



ECC



Adjusted Ensemble



Ensemble




FIGURE A2 Rank histograms of bivariate forecasts with respect to multivariate ordering for days 1, 5, 10 and 15. Figure 5b shows the corresponding reliability indices

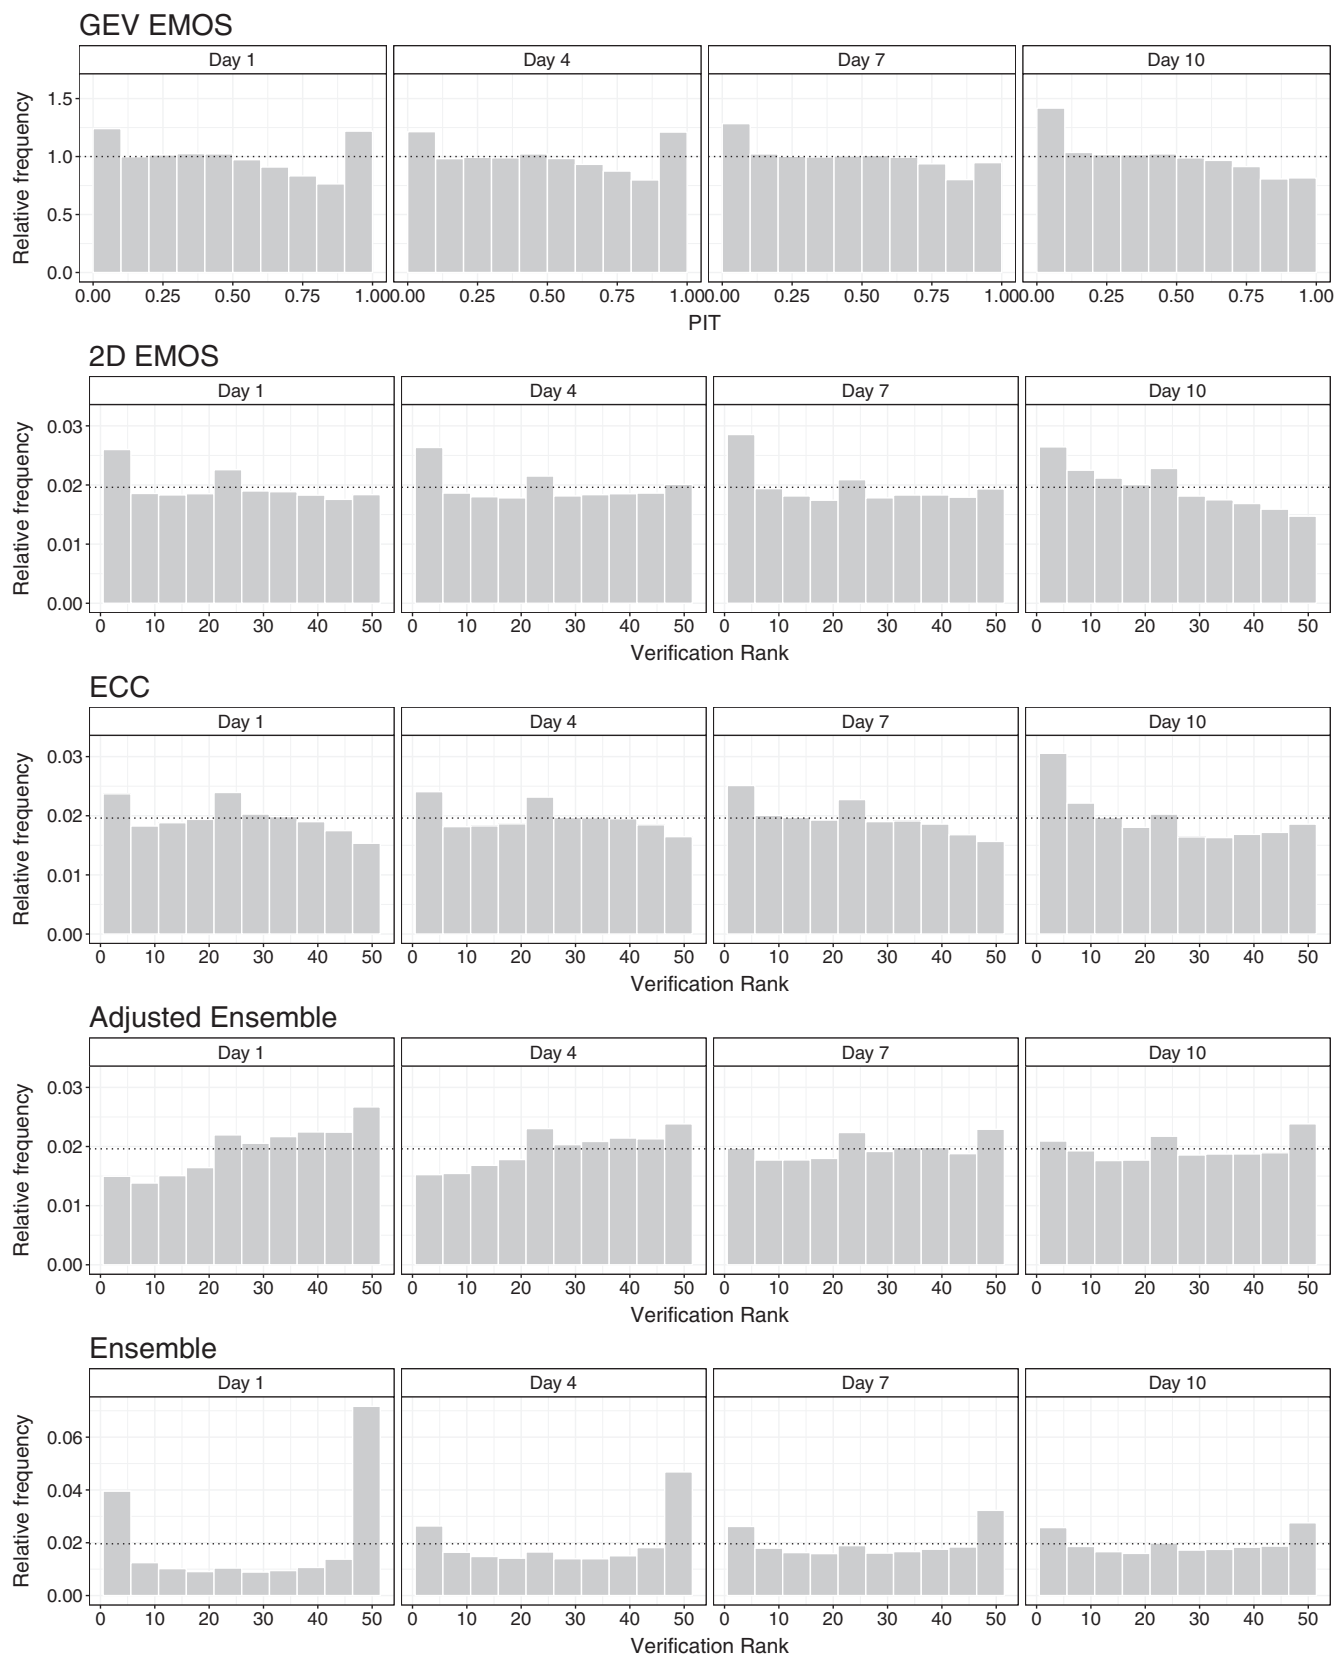


FIGURE A3 PIT histograms of the GEV EMOS post-processed, simulated verification rank histograms of the 2D EMOS post-processed, and verification rank histograms of the adjusted and raw DI ensemble forecasts for days 1, 4, 7 and 10