

# Regression diagnostics meets forecast evaluation: conditional calibration, reliability diagrams, and coefficient of determination

Tilmann Gneiting<sup>1,2</sup> and Johannes Resin<sup>3,2,1</sup>

<sup>1</sup>*Computational Statistics, Heidelberg Institute for Theoretical Studies*  
e-mail: [tilmann.gneiting@h-its.org](mailto:tilmann.gneiting@h-its.org)

<sup>2</sup>*Institute for Stochastics, Karlsruhe Institute of Technology (KIT)*

<sup>3</sup>*Alfred Weber Institute for Economics, Heidelberg University*  
e-mail: [johannes.resin@awi.uni-heidelberg.de](mailto:johannes.resin@awi.uni-heidelberg.de)

**Abstract:** A common principle in model diagnostics and forecast evaluation is that fitted or predicted distributions ought to be reliable, ideally in the sense of auto-calibration, where the outcome is a random draw from the posited distribution. For binary responses, auto-calibration is the universal concept of reliability. For real-valued outcomes, a general theory of calibration has been elusive, despite a recent surge of interest in distributional regression and machine learning. We develop a framework rooted in probability theory, which gives rise to hierarchies of calibration, and applies to both predictive distributions and stand-alone point forecasts. In a nutshell, a prediction is conditionally T-calibrated if it can be taken at face value in terms of an identifiable functional T. We introduce population versions of T-reliability diagrams and revisit a score decomposition into measures of miscalibration, discrimination, and uncertainty. In empirical settings, stable and efficient estimators of T-reliability diagrams and score components arise via nonparametric isotonic regression and the pool-adjacent-violators algorithm. For in-sample model diagnostics, we propose a universal coefficient of determination that nests and reinterprets the classical  $R^2$  in least squares regression and its natural analog  $R^1$  in quantile regression, yet applies to T-regression in general.

**MSC2020 subject classifications:** Primary 62G99, 62J20.

**Keywords and phrases:** Calibration test, canonical loss, consistent scoring function, model diagnostics, nonparametric isotonic regression, prequential principle, score decomposition, skill score.

Received October 2022.

## 1. Introduction

Predictive distributions ought to be calibrated or reliable (Dawid, 1984; Gneiting and Katzfuss, 2014). More generally, statistical models ought to provide

plausible probabilistic explanations of observations, be it in-sample or out-of-sample, ideally in the sense of auto-calibration, meaning that the outcomes are indistinguishable from random draws from the posited distributions. For binary outcomes, auto-calibration is the universal standard of reliability. In the general case of linearly ordered, real-valued outcomes, weaker, typically unconditional facets of calibration have been studied, with probabilistic calibration, which corresponds to the uniformity of the probability integral transform (PIT; Dawid, 1984; Diebold, Gunther and Tay, 1998), being the most popular notion. Recently, conditional notions have been proposed (Mason et al., 2007; Bentzien and Friederichs, 2014; Strähl and Ziegel, 2017), and there has been a surge of attention to calibration in the machine learning community, where the full conditional distribution of a response, given a feature vector, is of increasing interest, as exemplified by the work of Guo et al. (2017), Kuleshov, Fenner and Ermon (2018), Song et al. (2019), Gupta, Podkopaev and Ramdas (2020), Zhao, Ma and Ermon (2020), Sahoo et al. (2021) and Roelofs et al. (2022). While in the literature on forecast evaluation predictive performance is judged out-of-sample, calibration is relevant in regression diagnostics as well, where in-sample goodness-of-fit is assessed via test statistics or criteria of  $R^2$ -type. In many ways, the complementary perspectives of model diagnostics and forecast evaluation are two sides of the same coin.

In this paper, we strive to develop a theory of calibration for real-valued outcomes that complements the aforementioned strands of literature. Starting from measure theoretic and probabilistic foundations, we develop practical tools for visualizing, diagnosing and testing calibration, for both in-sample and out-of-sample settings, and applying both to full distributions and functionals thereof. Section 2 develops an overarching, rigorous theoretical framework in a general population setting, where we establish hierarchical relations between notions of unconditional and conditional calibration, with Figure 1 summarizing key results. We reduce a posited distribution to a typically single-valued statistical functional,  $T$ , and define conditional calibration in terms of said functional. While in general auto-calibration fails to imply calibration in terms of a functional, we prove this implication for functionals defined via an identification function, such as event probabilities, means, quantiles, and generalized quantiles. We plot recalibrated values of the functional against posited values to obtain  $T$ -reliability diagrams and revisit extant score decompositions to define nonnegative measures of miscalibration (MCB), discrimination (DSC) and uncertainty (UNC), for which the mean score satisfies  $\bar{S} = \text{MCB} - \text{DSC} + \text{UNC}$ . These considerations continue to apply when  $T$ -regression is studied as an end in itself, such as in mean (least squares) and quantile regression. In this setting, Theorem 2.26 establishes a general link between unconditional calibration and canonical score optimization, which nests classical results in least squares regression and the partitioning inequalities of quantile regression (Koenker and Bassett, 1978, Theorem 3.4).

In Section 3, we turn to empirical settings and statistical inference. We adopt and generalize the approach of Dimitriadis, Gneiting and Jordan (2021) that uses isotonic regression and the pool-adjacent-violators (PAV) algorithm

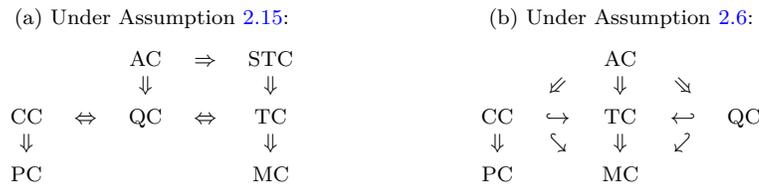


FIG 1. Preview of key findings in Section 2.3: Hierarchies of calibration (a) for continuous, strictly increasing cumulative distribution functions (CDFs) with common support and (b) under minimal conditions, with auto-calibration (AC) being the strongest notion. Conditional exceedance probability calibration (CC) is a conditional version of probabilistic calibration (PC), whereas threshold calibration (TC) is a conditional version of marginal calibration (MC). Quantile calibration (QC) differs from CC and TC in subtle ways. Strong threshold calibration (STC) is a stronger notion of threshold calibration introduced by Sahoo et al. (2021) for continuous CDFs. Hook arrows show conjectured implications.

(Ayer et al., 1955) to obtain consistent, optimally binned, reproducible, and PAV based (CORP) estimates of T-reliability diagrams and score components, along with uncertainty quantification via resampling. As opposed to extant estimators, the CORP approach yields non-decreasing reliability diagrams and guarantees the nonnegativity of the estimated MCB and DSC components. The regularizing constraint of isotonicity avoids artifacts and overfitting. For in-sample model diagnostics, we introduce a generalized coefficient of determination  $R^*$  that links to skill scores, and nests both the classical variance explained or  $R^2$  in least squares regression (Kvålseth, 1985), and its natural analogue  $R^1$  in quantile regression (Koenker and Machado, 1999). Subject to modest conditions  $R^* \in [0, 1]$ , with values of 0 and 1 indicating uninformative and immaculate fits, respectively.

In forecast evaluation, reliability diagrams and score components serve to diagnose and quantify performance on test samples. The most prominent case arises when T is the mean functional and performance is assessed by the mean squared error (MSE). As a preview of the diagnostic tools developed in this paper, we assess point forecasts by Tredennick et al. (2021) of (log-transformed) butterfly population size from a ridge regression and a null model. The CORP mean reliability diagrams and MSE decompositions in Figure 2 show that, while both models are reliable, ridge regression enjoys considerably higher discrimination ability.

The paper closes in Section 4, where we discuss our findings and provide a roadmap for follow-up research. While Dimitriadis, Gneiting and Jordan (2021) introduced the CORP approach in the nested case of probability forecasts for binary outcomes, the setting of real-valued outcomes treated in this paper is far more complex as it necessitates the consideration of statistical functionals in general. Throughout, we link the traditional case of regression diagnostics and (stand-alone) point forecast evaluation, where functionals such as conditional means, moments, quantiles, or expectiles are modeled and predicted, to model diagnostics and forecast evaluation in the fully distributional setting (Gneiting and Katzfuss, 2014; Hothorn, Kneib and Bühlmann, 2014). Appendices A–C include material of more specialized or predominantly technical character.

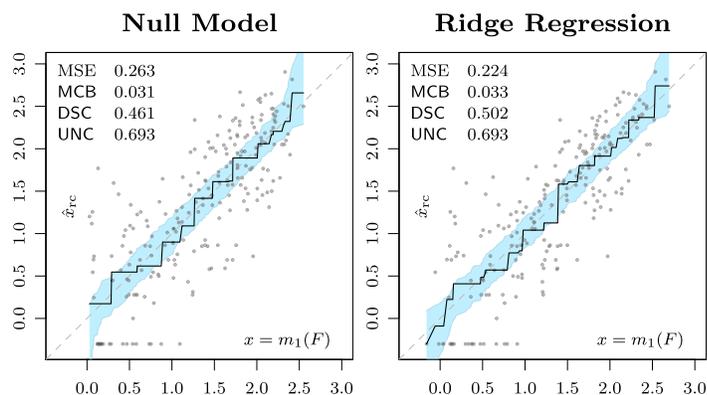


FIG 2. CORP mean reliability diagrams for point forecasts of (log-transformed) butterfly population size from the null model (left) and ridge regression (right) of [Tredennick et al. \(2021\)](#), along with 90% consistency bands and miscalibration (MCB), discrimination (DSC) and uncertainty (UNC) components of the mean squared error (MSE).

## 2. Notions of calibration, reliability diagrams, and score decompositions

Generally, we use the symbol  $\mathcal{L}$  to denote a generic conditional or unconditional law or distribution, and we identify distributions with their cumulative distribution functions (CDFs). We write  $\mathcal{N}(m, c^2)$  to denote a normal distribution with mean  $m$  and variance  $c^2$ , and we let  $\varphi$  and  $\Phi$  denote the density and the CDF, respectively, of a standard normal variable.

### 2.1. Prediction spaces and prequential principle

We consider the joint law of a posited distribution and the respective outcome in the technical setting of [Gneiting and Ranjan \(2013\)](#). Specifically, let  $(\Omega, \mathcal{A}, \mathbb{Q})$  be a *prediction space*, i.e., a probability space where the elementary elements  $\omega \in \Omega$  correspond to realizations of the random triple

$$(F, Y, U),$$

where  $Y$  is the real-valued outcome,  $F$  is a posited distribution for  $Y$  in the form of a CDF, and  $U$  is uniformly distributed on the unit interval.<sup>1</sup> Statements involving conditional or unconditional distributions, expectations, or probabilities, generally refer to the probability measure  $\mathbb{Q}$ , which specifies the joint distribution of the forecast  $F$  and the outcome  $Y$ . The uniform random variable  $U$

<sup>1</sup>Following the extant literature, we view both the predictive distributions, which might depend on stochastic parameters, and the outcomes as random elements ([Murphy and Winkler, 1987](#); [Gneiting, Balabdaoui and Raftery, 2007](#)). In parts of our paper, it suffices to consider a point prediction space, where the elements of the probability space correspond to realizations of the random tuple  $(X, Y)$  where  $X$  is a point forecast ([Ehm et al., 2016](#), eq. (20)).

allows for randomization. Throughout, we assume that  $U$  is independent of the  $\sigma$ -algebra generated by the random variable  $Y$  and the random function  $F$  in the technical sense detailed prior to Definition 2.6 in [Strähl and Ziegel \(2017\)](#).

Let  $\mathcal{A}_0 \subseteq \mathcal{A}$  denote the forecaster's *information basis*, i.e., a sub- $\sigma$ -algebra such that  $F$  is measurable with respect to  $\mathcal{A}_0$ . Then  $F$  is *ideal* relative to  $\mathcal{A}_0$  ([Gneiting and Ranjan, 2013](#)) if

$$F(y) = \mathbb{Q}(Y \leq y \mid \mathcal{A}_0) \quad \text{almost surely, for all } y \in \mathbb{R}.$$

If  $F$  is ideal relative to some sub  $\sigma$ -algebra  $\mathcal{A}_0$ , then it is *auto-calibrated* ([Tsyplakov, 2013](#)) in the sense that

$$F(y) = \mathbb{Q}(Y \leq y \mid F) \quad \text{almost surely, for all } y \in \mathbb{R},$$

which is equivalent to being ideal relative to the information basis  $\sigma(F) \subseteq \mathcal{A}_0$ . Extensions to prediction spaces with tuples  $(Y, F_1, \dots, F_k, U)$  that allow for multiple CDF-valued forecasts  $F_1, \dots, F_k$  with associated information bases  $\mathcal{A}_1, \dots, \mathcal{A}_k \subset \mathcal{A}$  are straightforward.

**Example 2.1** ([Gneiting and Ranjan, 2013](#); [Pohle, 2020](#)). Conditionally on a standard normal variate  $\mu$ , let the outcome  $Y$  be normal with mean  $\mu$  and variance 1. Then the *perfect* forecast  $F_1 = \mathcal{N}(\mu, 1)$  is ideal relative to the information basis  $\mathcal{A}_1 = \sigma(\mu)$  generated by  $\mu$ . The *unconditional* forecast  $F_2 = \mathcal{N}(0, 2)$  agrees with the marginal distribution of the outcome  $Y$  and is ideal relative to the trivial  $\sigma$ -algebra  $\mathcal{A}_2 = \{\emptyset, \Omega\}$ .

More elaborate notions of prediction spaces are feasible. In particular, one might include a covariate or feature vector  $Z$  and consider random tuples of the form  $(Z, F, Y, U)$ . Indeed, the transdisciplinary scientific literature has considered reliability relative to covariate information, under labels such as *strong* ([Van Calster et al., 2016](#)) or *individual* ([Chung et al., 2021](#); [Zhao, Ma and Ermon, 2020](#)) calibration. We refrain from doing so as our simple setting adheres to the *sequential principle* posited by [Dawid \(1984\)](#), according to which predictive performance needs to be evaluated on the basis of the tuple  $(F, Y)$  only, without consideration of the forecast-generating mechanism. The aforementioned extensions become critical in studies of cross-calibration ([Strähl and Ziegel, 2017](#)), stratification ([Ehm and Ovcharov, 2017](#); [Allen, 2021](#); [Bashaykh, 2022](#)), sensitivity ([Fissler and Pesenti, 2023](#)), and fairness ([Pleiss et al., 2017](#); [Mitchell et al., 2021](#)).

## 2.2. Traditional notions of unconditional calibration

Let us recall the classical notions of calibration of predictive distributions for real-valued outcomes. In order to do so, we define the *probability integral transform* (PIT)

$$Z_F = F(Y-) + U(F(Y) - F(Y-)) \tag{1}$$

of the CDF-valued random quantity  $F$ , where  $F(y-) = \lim_{x \uparrow y} F(x)$  denotes the left-hand limit of  $F$  at  $y \in \mathbb{R}$ , and the random variable  $U$  is standard

uniform and independent of  $F$  and  $Y$ . The PIT of a continuous CDF  $F$  is simply  $Z_F = F(Y)$ . The predictive distribution  $F$  is *probabilistically calibrated* or *PIT calibrated* if  $Z_F$  is uniformly distributed on the unit interval. The use of the probabilistic calibration criterion was suggested by Dawid (1984) and popularized by Diebold, Gunther and Tay (1998), who proposed the use of PIT histograms as a diagnostic tool. Importantly, in continuous settings probabilistic calibration implies that prediction intervals bracketed by quantiles capture the outcomes at the respective nominal level.

Furthermore, the predictive distribution  $F$  is *marginally calibrated* (Gneiting, Balabdaoui and Raftery, 2007) if

$$\mathbb{E}_{\mathbb{Q}}[F(y)] = \mathbb{Q}(Y \leq y) \quad \text{for all } y \in \mathbb{R}.$$

Thus, for a marginally calibrated predictive distribution, the frequency of (not) exceeding a threshold value matches the posited unconditional probability.

**Example 2.2.** In the setting of Example 2.1, let  $\eta$  attain the values  $\pm\eta_0$  with equal probability, independently of  $\mu$  and  $Y$ , where  $\eta_0 > 0$ . Then the *unfocused* forecast with CDF

$$F(y) = \frac{1}{2} (\Phi(y - \mu) + \Phi(y - \mu - \eta)) \quad (2)$$

is probabilistically calibrated but fails to be marginally calibrated (Gneiting, Balabdaoui and Raftery, 2007). Similarly, let  $\delta$  take the values  $\pm\delta_0$  with equal probability, independently of  $\mu$  and  $Y$ , where  $\delta_0 \in (0, 1)$ . Then the *lopsided* forecast  $F$  with density

$$f(y) = (1 - \delta)\varphi(y - \mu)\mathbb{1}(y < \mu) + (1 + \delta)\varphi(y - \mu)\mathbb{1}(y > \mu) \quad (3)$$

is marginally calibrated but fails to be probabilistically calibrated. For details see Appendix A.2.

It is well known that an ideal forecast is both probabilistically calibrated and marginally calibrated (Gneiting and Ranjan, 2013, Theorem 2.8; Song et al., 2019, Theorem 1). Reformulated in terms of auto-calibration the following holds.

**Theorem 2.3.** *Auto-calibration implies marginal and probabilistic calibration.*

Auto-calibration thus is a stronger requirement than either marginal or probabilistic calibration, and the latter are logically independent. However, in the special case of a binary outcome, probabilistic calibration and auto-calibration are equivalent (Gneiting and Ranjan, 2013, Theorem 2.11), and auto-calibration serves as a universal notion of calibration. In the case of three or more distinct outcomes, Gneiting and Ranjan (2013) conjectured that auto-calibration is stronger than simultaneous marginal and probabilistic calibration. Example 10 of Tsyplakov (2014) proves the conjecture by constructing continuous forecast distributions that are both marginally and probabilistically calibrated but fail to be auto-calibrated. The following example resolves the questions raised by Gneiting and Ranjan (2013) in full detail.

**Example 2.4.** We begin by considering continuous CDFs and then discuss a discrete example with three distinct outcomes only.

- (a) Suppose that  $\mu$  is normal with mean 0 and variance  $c^2$ . Conditionally on  $\mu$ , let the *piecewise uniform* predictive distribution  $F$  be a mixture of uniform measures on  $[\mu, \mu + 1]$ ,  $[\mu + 1, \mu + 2]$ , and  $[\mu + 2, \mu + 3]$  with weights  $p_1$ ,  $p_2$  and  $p_3$ , respectively, and let the outcome  $Y$  be drawn from a mixture with weights  $q_1$ ,  $q_2$  and  $q_3$  on these intervals. Finally, let the tuple  $(p_1, p_2, p_3; q_1, q_2, q_3)$  attain each of the values

$$\left(\frac{1}{2}, \frac{1}{4}, \frac{1}{4}; \frac{5}{10}, \frac{1}{10}, \frac{4}{10}\right), \quad \left(\frac{1}{4}, \frac{1}{2}, \frac{1}{4}; \frac{1}{10}, \frac{8}{10}, \frac{1}{10}\right), \quad \text{and} \quad \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{2}; \frac{4}{10}, \frac{1}{10}, \frac{5}{10}\right),$$

with  $\mathbb{Q}$ -probability  $\frac{1}{3}$ . Evidently,  $F$  fails to be auto-calibrated. However,  $F$  is marginally calibrated as, conditionally on  $\mu$ , it assigns the same mass  $\frac{1}{3}$  to each of the intervals, in agreement with the conditional distribution of  $Y$ . As for the PIT  $Z_F$ , conditionally on  $\mu$  its CDF is piecewise linear on the partition induced by  $0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}$ , and 1. Thus, in order to establish probabilistic calibration it suffices to verify that  $\mathbb{Q}(Z_F \leq x) = x$  for  $x \in \{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}\}$ , as confirmed by elementary calculations. Integration over  $\mu$  completes the argument.

- (b) For a full resolution of the aforementioned conjecture by [Gneiting and Ranjan \(2013\)](#), we fix  $\mu = 0$  and replace the intervals by fixed numbers  $y_1 < y_2 < y_3$ . Thus,  $F$  assigns mass  $p_j$  to  $y_j$ , whereas the event  $Y = y_j$  realizes with probability  $q_j$  for  $j = 1, 2$ , and 3. The forecast remains probabilistically and marginally calibrated, and fails to be auto-calibrated.

### 2.3. Conditional calibration

While checks for probabilistic calibration have become a cornerstone of predictive distribution evaluation ([Dawid, 1984](#); [Diebold, Gunther and Tay, 1998](#); [Gneiting, Balabdaoui and Raftery, 2007](#)), both marginal and probabilistic calibration concern unconditional facets of predictive performance, which is increasingly being considered insufficient (e.g., [Levi et al., 2022](#)). Stronger conditional notions of calibration, which condition on facets of the predictive distribution, have emerged in various strands of the scientific literature. For example, [Mason et al. \(2007\)](#) used conditional (non) exceedance probabilities (CEP) to assess the calibration of ensemble weather forecasts. These were used by [Held, Rufibach and Balabdaoui \(2010\)](#) and [Strähl and Ziegel \(2017\)](#) to derive calibration tests, which operate under the hypothesis that the forecast  $F$  is *CEP calibrated* in the sense that

$$\mathbb{Q}(Z_F \leq \alpha \mid q_\alpha^-(F)) = \alpha \quad \text{almost surely, for all } \alpha \in (0, 1), \quad (4)$$

where  $q_\alpha^-(F) = \inf\{x \in \mathbb{R} : F(x) \geq \alpha\}$  denotes the (lower)  $\alpha$ -quantile of  $F$ . Similarly, [Henzi, Ziegel and Gneiting \(2021\)](#) introduced the notion of a *threshold calibrated* forecast  $F$ , which stipulates that

$$\mathbb{Q}(Y \leq t \mid F(t)) = F(t) \quad \text{almost surely, for all } t \in \mathbb{R}. \quad (5)$$

Essentially, CEP calibration is a conditional version of probabilistic calibration, and threshold calibration is conditional marginal calibration.

**Theorem 2.5.** *CEP calibration implies probabilistic calibration, and threshold calibration implies marginal calibration.*

*Proof.* Immediate by taking unconditional expectations, as noted by [Henzi, Ziegel and Gneiting \(2021\)](#).  $\square$

Variants of these concepts can be found scattered in the literature. Notably, [Sahoo et al. \(2021\)](#) introduce a notion of calibration for continuous predictive distributions, which requires that

$$\mathbb{Q}(Z_F \leq \alpha \mid F(t)) = \alpha \quad \text{almost surely, for all } \alpha \in (0, 1), t \in \mathbb{R}. \quad (6)$$

As in [Figure 1](#), we refer to this property as *strong threshold calibration*. The notion is weaker than auto-calibration, but it implies both CEP calibration and threshold calibration, subject to conditions that we discuss below.

We proceed to the general notion of conditional T-calibration in terms of a statistical functional T as introduced by [Arnold \(2020\)](#) and [Bashaykh \(2022\)](#). Other authors ([Pohle, 2020](#); [Krüger and Ziegel, 2021](#)) refer to this notion or special cases thereof as auto-calibration with respect to T. A statistical functional on some class  $\mathcal{F}$  of probability measures is a measurable function  $T: \mathcal{F} \rightarrow \mathcal{T}$  into a (typically, finite-dimensional) space  $\mathcal{T}$  with Borel- $\sigma$ -algebra  $\mathcal{B}(\mathcal{T})$ . Technically, we work in the prediction space setting under a natural measurability condition that is not restrictive ([Fissler and Holzmann, 2022](#)).

**Assumption 2.6.** The class  $\mathcal{F}$  and the functional T are such that  $F \in \mathcal{F}$ , the mapping  $T(F): (\Omega, \mathcal{A}) \rightarrow (\mathcal{T}, \mathcal{B}(\mathcal{T}))$  is measurable, and  $\mathcal{L}(Y \mid T(F)) \in \mathcal{F}$  almost surely.

**Definition 2.7.** Under [Assumption 2.6](#), the predictive distribution  $F$  is *conditionally T-calibrated*, or simply *T-calibrated*, if

$$T(\mathcal{L}(Y \mid T(F))) = T(F) \quad \text{almost surely.}$$

Essentially, under a T-calibrated predictive distribution  $F$ , we can take  $T(F)$  at face value. Perhaps surprisingly, an auto-calibrated forecast is not necessarily T-calibrated, as noted by [Arnold \(2020, Section 3.2\)](#). For a simple counterexample, consider the perfect forecast from [Example 2.1](#), which fails to be T-calibrated when T is the variance, the standard deviation, the interquartile range, or a related measure of dispersion.

We proceed to show that this startling issue does not occur with *identifiable* functionals, i.e., functionals induced by an identification function ([Theorem 2.11](#)). Similar to the classical procedure in  $M$ -estimation ([Huber, 1964](#)), an identification function weighs negative values in the case of underprediction against positive values in the case of overprediction, and the corresponding functional maps to the possibly set-valued argument at which an associated expectation changes sign. Following [Jordan, Mühlemann and Ziegel \(2022\)](#), a

TABLE 1

Key examples of identifiable functionals with associated parameters, identification function, and generic type. For a similar listing see Table 1 in Jordan, Mühlemann and Ziegel (2022).

Functional	Parameters	Identification function	Type
Threshold (non) exceedance	$t \in \mathbb{R}$	$V(x, y) = x - \mathbb{1}\{y \leq t\}$	singleton
Mean		$V(x, y) = x - y$	singleton
Median		$V(x, y) = \mathbb{1}\{y < x\} - \frac{1}{2}$	interval
Moment of order $n$ ( $m_n$ )	$n = 1, 2, \dots$	$V(x, y) = x - y^n$	singleton
$\alpha$ -Expectile ( $e_\alpha$ )	$\alpha \in (0, 1)$	$V(x, y) =  \mathbb{1}\{y < x\} - \alpha (x - y)$	singleton
$\alpha$ -Quantile ( $q_\alpha$ )	$\alpha \in (0, 1)$	$V(x, y) = \mathbb{1}\{y < x\} - \alpha$	interval
Huber	$\alpha \in (0, 1),$ $a, b > 0$	$V(x, y) =  \mathbb{1}\{y < x\} - \alpha  \kappa_{a,b}(x - y)$	interval

measurable function  $V: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$  is an *identification function* if  $V(\cdot, y)$  is increasing and left-continuous for all  $y \in \mathbb{R}$ . We operate under Assumption 2.6 with the implicit understanding that  $V(x, \cdot)$  is quasi-integrable with respect to all  $F \in \mathcal{F}$  for all  $x \in \mathbb{R}$ . Then, for any probability measure  $F$  in the class  $\mathcal{F}$ , the functional  $T(F)$  induced by  $V$  is defined as

$$T(F) = [T^-(F), T^+(F)] \subseteq [-\infty, +\infty] = \bar{\mathbb{R}},$$

where the lower and upper bounds are given by the random variables

$$T^-(F) = \sup \left\{ x : \int V(x, y) dF(y) < 0 \right\} \quad (7)$$

and

$$T^+(F) = \inf \left\{ x : \int V(x, y) dF(y) > 0 \right\}. \quad (8)$$

An identifiable functional  $T$  is of *singleton type* if  $T(F)$  is a singleton for every  $F \in \mathcal{F}$ . Otherwise,  $T$  is of *interval type*. Table 1 lists key examples, such as threshold-defined event probabilities, quantiles, expectiles, and moments. The definition of the Huber functional involves the clipping function  $\kappa_{a,b}(t) = \max(\min(t, b), -a)$  with parameters  $a, b > 0$  (Taggart, 2022). In the limiting cases as  $a = b \rightarrow 0$  and  $a = b \rightarrow \infty$ , the Huber functional recovers the  $\alpha$ -quantile ( $q_\alpha$ ) and the  $\alpha$ -expectile ( $e_\alpha$ ), respectively.

For identifiable functionals we can define an unconditional notion of  $T$ -calibration as well. Note that in contrast to traditional settings, where  $F$  is fixed, we work in the prediction space setting, where  $F$  is a random CDF. In principle, the subsequent Definitions 2.9 and 2.24 depend on the choice of the identification function. However, as we demonstrate in Appendix A.4, the following condition ensures that the identification function is unique, up to a positive constant, so that ambiguities are avoided.

**Assumption 2.8.** The identification function  $V$  induces the functional  $T$  on a convex class  $\mathcal{F}_0 \supseteq \mathcal{F}$  of probability measures, which contains the Dirac measures  $\delta_y$  for all  $y \in \mathbb{R}$ . The identification function  $V$  is

- (i) of *prediction error form*, i.e., there exists an increasing, left-continuous function  $v: \mathbb{R} \rightarrow \mathbb{R}$  such that  $V(x, y) = v(x - y)$  with  $v(-r) < 0$  and  $v(r) > 0$  for some  $r > 0$ , or
- (ii) of the form  $V(x, y) = x - T(\delta_y)$  for a functional  $T$  of singleton type.

The examples in Table 1 all satisfy Assumption 2.8.

**Definition 2.9.** Suppose that the functional  $T$  is generated by an identification function  $V$ , and let Assumptions 2.6 and 2.8 hold. Then the predictive distribution  $F$  is *unconditionally T-calibrated* if

$$\mathbb{E}_{\mathbb{Q}}[V(T^+(F) - \varepsilon, Y)] \leq 0 \quad \text{and} \quad \mathbb{E}_{\mathbb{Q}}[V(T^-(F) + \varepsilon, Y)] \geq 0 \quad \text{for all } \varepsilon > 0. \quad (9)$$

For the interval-valued  $\alpha$ -quantile functional,  $q_{\alpha}(F) = [q_{\alpha}^-(F), q_{\alpha}^+(F)]$ , condition (9) reduces to the traditional unconditional coverage condition

$$\mathbb{Q}(Y \leq q_{\alpha}^-(F)) \geq \alpha \quad \text{and} \quad \mathbb{Q}(Y \geq q_{\alpha}^+(F)) \geq 1 - \alpha, \quad (10)$$

with the latter being equivalent to  $\mathbb{Q}(Y < q_{\alpha}^+(F)) \leq \alpha$ . Probabilistic calibration implies unconditional  $\alpha$ -quantile calibration at every level  $\alpha \in (0, 1)$ , as hinted at by Kuleshov, Fenner and Ermon (2018, Section 3.1), and proved in our Appendix A.5. Under technical assumptions, condition (9) simplifies to

$$\mathbb{E}_{\mathbb{Q}}[V(T(F), Y)] = 0, \quad (11)$$

with the classical unbiasedness condition  $\mathbb{E}_{\mathbb{Q}}[m_1(F)] = \mathbb{E}_{\mathbb{Q}}[Y]$  arising in the case of the mean or expectation functional.

**Example 2.10.** Let  $T$  be the mean functional or a quantile. Then the unfocused forecast from Example 2.2 is unconditionally  $T$ -calibrated but fails to be conditionally  $T$ -calibrated. For details see Figure 4 and Appendix A.1.

Importantly, for any identifiable functional auto-calibration implies both unconditional and conditional  $T$ -calibration, as we demonstrate now. Note that Assumption 2.6 is a minimal condition, which is required to define conditional  $T$ -calibration in the first place, whereas Assumption 2.8 ensures that unconditional  $T$ -calibration is a well-defined concept.

**Theorem 2.11.** *Suppose that the functional  $T$  is generated by an identification function and Assumption 2.6 holds. Then auto-calibration implies conditional  $T$ -calibration, and, subject to Assumption 2.8, conditional  $T$ -calibration implies unconditional  $T$ -calibration.*

*Proof.* The statements in this proof are understood to hold almost surely. By Theorem 4.34 and Proposition 4.35 of Breiman (1992) in concert with auto-calibration,  $F$  is a regular conditional distribution of  $Y$  given  $F$ , and we conclude that

$$\mathbb{E}[V(x, Y) | F] = \int V(x, y) dF(y).$$

Furthermore, a regular conditional distribution  $F_T = \mathcal{L}(Y \mid T(F))$  of  $Y$  given  $T(F)$  exists, and the tower property of conditional expectation implies that

$$\begin{aligned} \int V(x, y) dF_T(y) &= \mathbb{E}[V(x, Y) \mid T(F)] = \mathbb{E}[\mathbb{E}[V(x, Y) \mid F] \mid T(F)] \\ &= \mathbb{E}\left[\int V(x, y) dF(y) \mid T(F)\right]. \end{aligned}$$

Let  $T(F) = [T^-(F), T^+(F)]$  and  $T(F_T) = [T^-(F_T), T^+(F_T)]$ , where the boundaries are random variables. The proof of the first part is complete if we can show that  $T^-(F_T) = T^-(F)$  and  $T^+(F_T) = T^+(F)$ .

Let  $\varepsilon > 0$ . By the definition of  $T^+(F)$ , we know that  $\int V(T^+(F), y) dF(y) \leq 0$  and  $\int V(T^+(F) + \varepsilon, y) dF(y) > 0$ . Using nested conditional expectations as above, the same inequalities hold almost surely when integrating with respect to  $F_T$ . Hence, by the definition of  $T^+(F_T)$ , we obtain  $T^+(F) \leq T^+(F_T) < T^+(F) + \varepsilon$ . An analogous argument shows that  $T^-(F) - \varepsilon \leq T^-(F_T) \leq T^-(F) + \varepsilon$ , which completes the proof of the first part and shows that  $F$  is conditionally  $T$ -calibrated.

Finally, if  $F$  is conditionally  $T$ -calibrated, unconditional  $T$ -calibration follows by taking nested expectations in the terms in the defining inequalities.  $\square$

An analogous result is easily derived for CEP calibration.

**Theorem 2.12.** *Under Assumption 2.6 for quantiles, auto-calibration implies CEP calibration.*

*Proof.* It holds that

$$\mathbb{Q}(Z_F \leq \alpha \mid q_\alpha^-(F)) = \mathbb{E}[\mathbb{1}\{Z_F \leq \alpha\} \mid q_\alpha^-(F)] = \mathbb{E}[\mathbb{E}[\mathbb{1}\{Z_F \leq \alpha\} \mid F] \mid q_\alpha^-(F)]$$

almost surely for  $\alpha \in (0, 1)$ . As  $F$  is a version of  $\mathcal{L}(Y \mid F)$ , the nested expectation equals  $\alpha$  almost surely by Proposition 2.1 of Rüschemdorf (2009), which implies CEP calibration.  $\square$

When evaluating full predictive distributions, it is natural to consider families of functionals as in the subsequent definition, where part (a) is compatible with the extant notion in (5).

**Definition 2.13.** A predictive distribution  $F$  is

- (a) *threshold calibrated* if it is conditionally  $F(t)$ -calibrated for all  $t \in \mathbb{R}$ ;
- (b) *quantile calibrated* if it is conditionally  $q_\alpha$ -calibrated for all  $\alpha \in (0, 1)$ ;
- (c) *expectile calibrated* if it is conditionally  $e_\alpha$ -calibrated for all  $\alpha \in (0, 1)$ ;
- (d) *moment calibrated* if it is conditionally  $n$ -th moment calibrated for all integers  $n = 1, 2, \dots$

While CEP, quantile, and threshold calibration are closely related notions, they generally are not equivalent. For illustration, we consider predictive CDFs in the spirit of Example 2.4.

**Example 2.14.**

- (a) Let  $\mu \sim \mathcal{N}(0, c^2)$ . Conditionally on  $\mu$ , let  $F$  be a mixture of uniform distributions on the intervals  $[\mu, \mu + 1]$ ,  $[\mu + 1, \mu + 2]$ ,  $[\mu + 2, \mu + 3]$ , and  $[\mu + 3, \mu + 4]$  with weights  $p_1, p_2, p_3$ , and  $p_4$ , respectively, and let  $Y$  be from a mixture with weights  $q_1, q_2, q_3$ , and  $q_4$ . Furthermore, let the tuple  $(p_1, p_2, p_3, p_4; q_1, q_2, q_3, q_4)$  attain each of the values

$$\begin{aligned} & \left(\frac{1}{2}, 0, \frac{1}{2}, 0; \frac{3}{4}, 0, \frac{1}{4}, 0\right), \quad \left(\frac{1}{2}, 0, 0, \frac{1}{2}; \frac{1}{4}, 0, 0, \frac{3}{4}\right), \\ & \left(0, \frac{1}{2}, \frac{1}{2}, 0; 0, \frac{1}{4}, \frac{3}{4}, 0\right), \quad \left(0, \frac{1}{2}, 0, \frac{1}{2}; 0, \frac{3}{4}, 0, \frac{1}{4}\right) \end{aligned}$$

with equal probability. Then the *continuous* forecast  $F$  is threshold calibrated and CEP calibrated but fails to be quantile calibrated.

- (b) Let the tuple  $(p_1, p_2, p_3; q_1, q_2, q_3)$  attain each of the values

$$\left(\frac{1}{2}, \frac{1}{4}, \frac{1}{4}; \frac{5}{10}, \frac{4}{10}, \frac{1}{10}\right), \quad \left(\frac{1}{4}, \frac{1}{2}, \frac{1}{4}; \frac{1}{10}, \frac{5}{10}, \frac{4}{10}\right), \quad \text{and} \quad \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{2}; \frac{4}{10}, \frac{1}{10}, \frac{5}{10}\right)$$

with equal probability. Let  $F$  assign mass  $p_j$  to numbers  $y_j$  for  $j = 1, 2, 3$ , where  $y_1 < y_2 < y_3$ , and let the event  $Y = y_j$  realize with probability  $q_j$ . The resulting *discrete* forecast is quantile and threshold calibrated. However, it fails to be CEP calibrated or even PIT calibrated.

Under the following conditions CEP, quantile, and threshold calibration coincide.

**Assumption 2.15.** In addition to Assumption 2.6 for quantiles and threshold (non) exceedances at all levels and thresholds, respectively, let the following hold.

- (i) The CDFs in the class  $\mathcal{F}$  are continuous and strictly increasing on a common support interval.
- (ii) There exists a countable set  $\mathcal{G} \subseteq \mathcal{F}$  such that  $\mathbb{Q}(F \in \mathcal{G}) = 1$ .

**Theorem 2.16.**

- (a) Under Assumption 2.15(i) CEP and quantile calibration are equivalent and imply probabilistic calibration.
- (b) Under Assumptions 2.15(i)–(ii) CEP, quantile and threshold calibration are equivalent and imply both probabilistic and marginal calibration.

*Proof.* By Assumption 2.15(i) the CDFs  $F \in \mathcal{F}$  are invertible on the common support with the quantile function  $\alpha \mapsto q_\alpha^-(F)$  as inverse. Hence, for every  $\alpha \in (0, 1)$  the functional  $q_\alpha$  is of singleton-type and  $q_\alpha(F) = \{q_\alpha^-(F)\}$ . In this light, the almost-sure identity

$$\mathbb{Q}(Z_F \leq \alpha \mid q_\alpha^-(F)) = \mathbb{Q}(Y \leq q_\alpha^-(F) \mid q_\alpha(F))$$

implies part (a). To prove part (b), let  $\mathcal{G}$  be as in Assumption 2.15(ii) and assume without loss of generality that  $\mathbb{Q}(F = G) > 0$  for all  $G \in \mathcal{G}$ . If  $\alpha \in (0, 1)$  and  $t \in \mathbb{R}$  are such that  $\mathbb{Q}(F(t) = \alpha) > 0$ , then

$$\mathbb{Q}(Y \leq t \mid F(t) = \alpha) = \mathbb{Q}(Y \leq q_\alpha^-(F) \mid q_\alpha^-(F) = t),$$

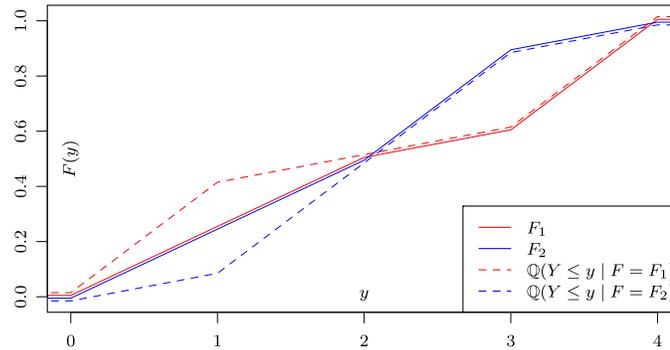


FIG 3. The equiprobable predictive distribution  $F$  picks the piecewise linear, partially (namely, for  $y \leq 2$ ) identical CDFs  $F_1$  and  $F_2$  with equal probability. It is jointly CEP, quantile, and threshold calibrated but fails to be auto-calibrated. For a very similar construction see Example 10 of [Tsyplakov \(2014\)](#).

where Assumption 2.15(ii) ensures that the events conditioned on have positive probability. Hence, quantile and threshold calibration are equivalent. The remaining implications are immediate from Theorem 2.5.  $\square$

We conjecture that the statement of part (b) holds under Assumption 2.15(i) alone but are unaware of a measure theoretic argument that serves to generalize the discrete reasoning in our proof. As indicated in panel (b) of Figure 1, we also conjecture that CEP or quantile calibration imply threshold calibration in general, though we have not been able to prove this implication, nor can we show that CEP or quantile calibration imply marginal calibration in general. Strong threshold calibration as defined in (6) implies both CEP and threshold calibration under Assumption 2.15, by arguments similar to those in the above proof. The following result thus demonstrates that the hierarchies in panel (a) and, with the aforementioned exceptions, in panel (b) of Figure 1 are complete, with the caveat that hierarchies may collapse if the class  $\mathcal{F}$  is sufficiently small, as exemplified by Theorem 2.11 of [Gneiting and Ranjan \(2013\)](#).

**Proposition 2.17.** *Even under Assumption 2.15 the following hold:*

- (a) *Strong threshold calibration does not imply auto-calibration.*
- (b) *Joint CEP, quantile, and threshold calibration does not imply strong threshold calibration.*
- (c) *Joint probabilistic and marginal calibration does not imply threshold calibration.*
- (d) *Probabilistic calibration does not imply marginal calibration.*
- (e) *Marginal calibration does not imply probabilistic calibration.*

*Proof.* We establish the claims in a series of (counter) examples, starting with part (b), where we present an example based on two equiprobable, partially overlapping CDFs in Figure 3. A similar example based on four equiprobable,

TABLE 2

Properties of the forecasts in our examples. We note whether they are auto-calibrated (AC), CEP calibrated (CC), quantile calibrated (QC), threshold calibrated (TC), probabilistically calibrated (PC), or marginally calibrated (MC), and whether the involved distributions are continuous and strictly increasing on a common support (CSI) as in Assumption 2.15(i).

Except for the auto-calibrated cases, the forecasts fail to be moment calibrated.

Source	Forecast Type	CSI	AC	CC	QC	TC	PC	MC
Example 2.1	Perfect	✓	✓	✓	✓	✓	✓	✓
Example 2.1	Unconditional	✓	✓	✓	✓	✓	✓	✓
Figure 3	Equiprobable	✓		✓	✓	✓	✓	✓
Example 2.4	Piecewise uniform	as $c \rightarrow 0$					✓	✓
Example 2.2	Unfocused	✓					✓	
Example 2.2	Lopsided	✓						✓
Example 2.14	Continuous			✓		✓	✓	✓
Example 2.14	Discrete				✓	✓		✓

partially overlapping CDFs in Appendix A.6 yields part (a). As for part (c), we return to the piecewise uniform forecast in Example 2.4, where for simplicity we fix  $\mu = 0$ . This forecast is probabilistically and marginally calibrated, but it fails to be threshold calibrated because

$$\mathbb{Q}\left(Y \leq \frac{3}{2} \mid F\left(\frac{3}{2}\right) = \frac{5}{8}\right) = \frac{5}{10} + \frac{1}{2} \cdot \frac{1}{10} = \frac{11}{20} \neq \frac{5}{8}.$$

As for parts (d) and (e), we refer to the unfocused and lopsided forecasts from Example 2.2 with  $\mu = 0$  fixed. □

Clearly, further hierarchical relations are immediate. For example, given that probabilistic calibration does not imply marginal calibration, it does not imply threshold calibration nor auto-calibration. We leave further discussion to future work but note that moment calibration does not imply probabilistic nor marginal calibration, as follows easily from classical results on the moment problem (e.g., Stoyanov, 2000). For an overview of calibration properties in our examples, see Table 2.

### 2.4. Reliability diagrams

As we proceed to define reliability diagrams, it is useful to restrict attention to single-valued functionals. To this end, if an identifiable functional  $T$  is of interval type, we instead consider its single-valued lower or upper bound,  $T^-(F)$  or  $T^+(F)$ , which we call the *lower* and *upper version* of  $T$ , or simply the *lower* and *upper functional*, without explicit reference to the original functional  $T$ .

The following result demonstrates that  $T$ -calibration implies calibration of the upper and lower functional.

**Proposition 2.18.** *Suppose that the functional  $T$  is generated by an identification function  $V$ , and let Assumption 2.6 hold. Then conditional  $T$ -calibration implies conditional  $T^-$ - and  $T^+$ -calibration, and, subject to Assumption 2.8, unconditional  $T$ -calibration implies unconditional  $T^-$ - and  $T^+$ -calibration.*

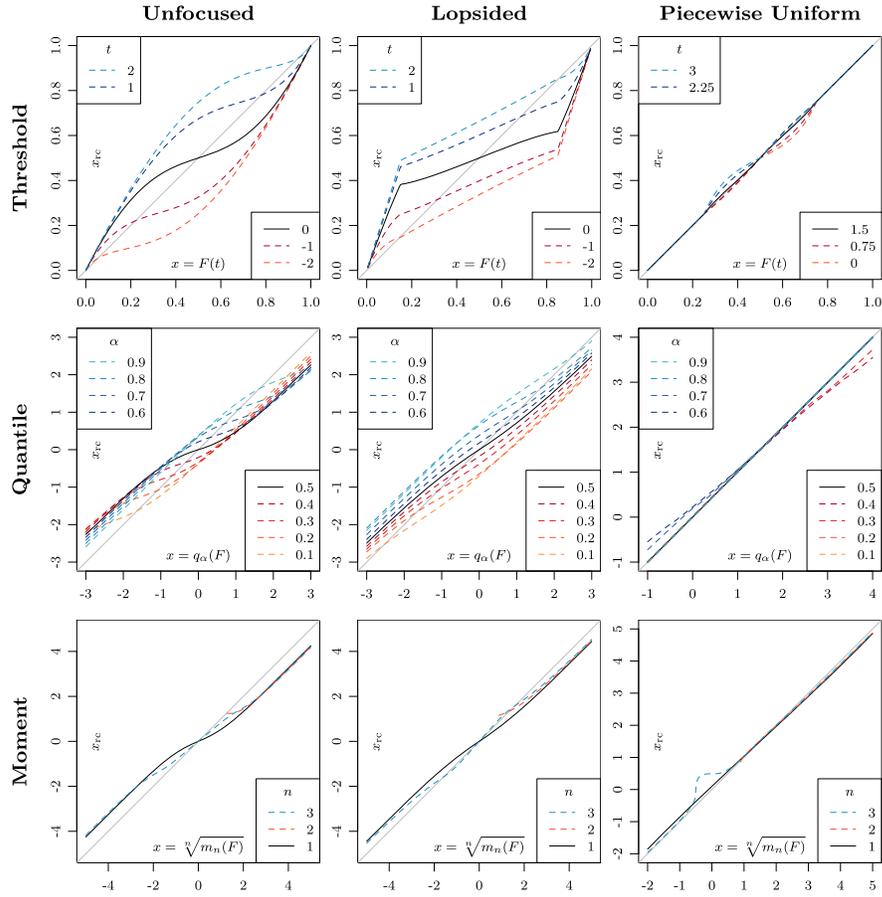


FIG 4. Threshold (top), quantile (middle), and moment (bottom) reliability diagrams for point forecasts induced by (left) the unfocused forecast with  $\eta_0 = 1.5$  and (middle) the lopsided forecast with  $\delta_0 = 0.7$  from Example 2.2, and (right) the piecewise uniform forecast with  $c = 0.5$  from Example 2.4. Each display plots recalibrated against original values. Deviations from the diagonal indicate violations of T-calibration. For details see Appendix A.

*Proof.* Suppose that  $T^*$  is the lower or upper version of a functional  $T$  generated by the identification function  $V$ . As  $\sigma(T^*(F)) \subseteq \sigma(T(F))$ , we find that

$$\mathbb{E}[V(x, Y) \mid T^*(F)] = \mathbb{E}[\mathbb{E}[V(x, Y) \mid T(F)] \mid T^*(F)]$$

is almost surely  $\leq 0$  if  $x < T^*(F)$ , and almost surely  $\geq 0$  if  $x > T^*(F)$ . Hence,  $T^*(F) \in T(\mathcal{L}(Y \mid T^*(F)))$ . If  $T^*$  is the lower functional, the former inequality is strict and hence  $T^-(F) = \min T(\mathcal{L}(Y \mid T^-(F)))$ , whereas if  $T^*$  is the upper functional, the latter is strict and hence  $T^+(F) = \max T(\mathcal{L}(Y \mid T^+(F)))$ .

Unconditional  $T^*$ -calibration is an immediate consequence of unconditional  $T$ -calibration.  $\square$

In this light, we restrict attention to single-valued functionals that are lower or upper versions of identifiable functionals, or identifiable functionals of singleton type. Any such functional can be associated with a random variable  $X = T(F)$ , and we call any random variable  $X_{rc}$ , for which

$$X_{rc} = T(\mathcal{L}(Y | X)) \quad (12)$$

almost surely, a *recalibrated* version of  $X$ . Clearly, we can also define  $X_{rc}$  for a stand-alone point forecast  $X$ , based on conceptualized distributions, by resorting to the joint distribution of the random tuple  $(X, Y)$ , provided the right-hand side of (12) is well defined and finite almost surely. The point forecast  $X$  is *conditionally T-calibrated*, or simply *T-calibrated*, if  $X = X_{rc}$  almost surely. Subject to Assumption 2.8,  $X$  is *unconditionally T-calibrated* if

$$\mathbb{E}[V(X - \varepsilon, Y)] \leq 0 \quad \text{and} \quad \mathbb{E}[V(X + \varepsilon, Y)] \geq 0 \quad \text{for all } \varepsilon > 0. \quad (13)$$

For recent discussions of the particular cases of the mean or expectation and quantile functionals see, e.g., [Nolde and Ziegel \(2017, Sections 2.1–2.2\)](#), [Patton \(2020, Proposition 2\)](#), [Krüger and Ziegel \(2021, Definition 3.1\)](#) and [Satopää \(2021, Section 2\)](#).

To compare the posited functional  $X$  with its recalibrated version  $X_{rc}$ , we introduce the T-reliability diagram.

**Assumption 2.19.** The functional  $T$  is a lower or upper version of an identifiable functional, or an identifiable functional of singleton type. The point forecast  $X$  is a random variable, and the recalibrated forecast  $X_{rc} = T(\mathcal{L}(Y | X))$  is well defined and finite almost surely.

**Definition 2.20.** Under Assumption 2.19, the *T-reliability diagram* is the graph of a mapping  $x \mapsto T(\mathcal{L}(Y | X = x))$  on the support of  $X$ .

While technically the T-reliability diagram depends on the choice of a regular conditional distribution for the outcome  $Y$ , this issue is not a matter of practical relevance. Evidently, for a T-calibrated forecast the T-reliability diagram is concentrated on the diagonal. Conversely, deviations from the diagonal indicate violations of T-calibration and can be interpreted diagnostically, as illustrated in Figure 4 for threshold, quantile, and moment calibration. For a similar display in the specific case of mean calibration see Figure 1 of [Pohle \(2020\)](#).

In the setting of fully specified predictive distributions, the distinction between unconditional and conditional T-calibration is natural. Perhaps surprisingly, the distinction vanishes in the setting of stand-alone point forecasts if the associated identification function is of prediction error form and the forecast and the residual are independent.

**Theorem 2.21.** *Let Assumption 2.19 hold, and suppose that the underlying identification function  $V$  satisfies Assumption 2.8. Suppose furthermore that the point forecast  $X$  and the generalized residual  $T(\delta_Y) - X$  are independent. Then  $X$  is conditionally T-calibrated if, and only if, it is unconditionally T-calibrated.*

*Proof.* In the case of an identification function  $V$  of type (i) in Assumption 2.8, let  $b(y) = T(\delta_y) - y$  for  $y \in \mathbb{R}$ . Then  $b = b(y)$  is constant as  $v(x - y') = v(x - (y' - y) - y)$  implies  $T(\delta_{y'}) = T(\delta_y) + y' - y$ . Given any constant  $c \in \mathbb{R}$  it holds that

$$\mathbb{E}[V(X + c, Y) \mid X] = \mathbb{E}[v(X + c - T(\delta_Y) + b) \mid X] = \mathbb{E}[v(c + b - (T(\delta_Y) - X))].$$

In view of (12) and (13), conditional and unconditional T-calibration are equivalent. In the case of an identification function of type (ii) in Assumption 2.8, the above arguments continue to hold if we take  $v$  to be the identity map  $x \mapsto x$  and  $b = 0$ .  $\square$

For quantiles, expectiles, and Huber functionals, the identification function  $V$  is of prediction error form and the generalized residual reduces to the standard residual,  $X - Y$ . In particular, this observation applies in the case of least squares regression, where  $T$  is the mean functional, and the forecast and the residual have typically been assumed to be independent in the literature. We discuss the statistical implications of Theorem 2.21 in Appendix B.

### 2.5. Score decompositions

We now revisit a score decomposition into measures of miscalibration (MCB), discrimination (DSC), and uncertainty (UNC) based on consistent scoring functions. Specifically, suppose that  $S$  is a *consistent* loss or scoring function for the functional  $T$  on the class  $\mathcal{F}$  in the sense that

$$\mathbb{E}_F[S(t, Y)] \leq \mathbb{E}_F[S(x, Y)]$$

for all  $F \in \mathcal{F}$ , all  $t \in T(F) = [T^-(F), T^+(F)]$  and all  $x \in \mathbb{R}$  (Savage, 1971; Gneiting, 2011). If the inequality is strict unless  $x \in T(F)$ , then  $S$  is *strictly consistent*. Consistent scoring functions serve as all-purpose performance measures that elicit fair and honest assessments and reward the utilization of broad information bases (Holzmann and Eulert, 2014). If the functional  $T$  is of interval type, a consistent scoring function  $S$  is consistent for both  $T^-$  and  $T^+$ , but strict consistency is lost when  $T$  is replaced by its lower or upper version and  $S$  is strictly consistent for  $T$ . For prominent examples of consistent scoring functions, see Table 3.

A functional is *elicitable* if it admits a strictly consistent scoring function (Gneiting, 2011). Under general conditions, elicibility is equivalent to identifiability (Steinwart et al., 2014, Theorem 5). The respective functionals allow for both principled relative forecast evaluation through the use of consistent scoring functions, and principled absolute forecast evaluation via T-reliability diagrams and score decompositions, as discussed in what follows.

Let  $\mathcal{L}(Y)$  denote the unconditional distribution of the outcome and suppose that  $x_0 = T(\mathcal{L}(Y))$  is well defined. As before, we operate under Assumption 2.19 and work with  $X = T(F)$ , its recalibrated version  $X_{rc}$ , and the reference forecast  $x_0$ . Again, the simplified notation accommodates stand-alone point forecasts,

and it suffices to consider the joint distribution of the tuple  $(X, Y)$ . Following the lead of Dawid (1986) in the case of binary outcomes, and Ehm and Ovcharov (2017) and Pohle (2020) in the setting of point forecasts for real-valued outcomes, we consider the expected scores

$$\bar{S} = \mathbb{E}_{\mathbb{Q}}[S(X, Y)], \quad \bar{S}_{\text{rc}} = \mathbb{E}_{\mathbb{Q}}[S(X_{\text{rc}}, Y)], \quad \text{and} \quad \bar{S}_{\text{mg}} = \mathbb{E}_{\mathbb{Q}}[S(x_0, Y)] \quad (14)$$

for the forecast at hand, its recalibrated version, and the marginal reference forecast  $x_0$ , respectively.

**Definition 2.22.** Let Assumption 2.19 hold, and let  $x_0 = T(\mathcal{L}(Y))$  and the expectations  $\bar{S}$ ,  $\bar{S}_{\text{rc}}$ , and  $\bar{S}_{\text{mg}}$  in (14) be well defined and finite. Then we refer to

$$\text{MCB}_S = \bar{S} - \bar{S}_{\text{rc}}, \quad \text{DSC}_S = \bar{S}_{\text{mg}} - \bar{S}_{\text{rc}}, \quad \text{and} \quad \text{UNC}_S = \bar{S}_{\text{mg}},$$

as *miscalibration*, *discrimination*, and *uncertainty*, respectively.

Following Dimitriadis, Gneiting and Jordan (2021), our terminology differs from terms typically used to describe score decompositions for probabilistic forecasts under proper scoring rules, where miscalibration and discrimination are commonly referred to as reliability and resolution, respectively (e.g., Bröcker, 2009; Siegert, 2017). The following result decomposes the expected score  $\bar{S}$  for the forecast at hand into miscalibration ( $\text{MCB}_S$ ), discrimination ( $\text{DSC}_S$ ), and uncertainty ( $\text{UNC}_S$ ) components.

**Theorem 2.23** (Dawid, 1986; Pohle, 2020). *In the setting of Definition 2.22, suppose that the scoring function  $S$  is consistent for the functional  $T$ . Then it holds that*

$$\bar{S} = \text{MCB}_S - \text{DSC}_S + \text{UNC}_S, \quad (15)$$

where  $\text{MCB}_S \geq 0$  with equality if  $X$  is conditionally  $T$ -calibrated, and  $\text{DSC}_S \geq 0$  with equality if  $X_{\text{rc}} = x_0$  almost surely. If  $S$  is strictly consistent then  $\text{MCB}_S = 0$  only if  $X$  is conditionally  $T$ -calibrated, and  $\text{DSC}_S = 0$  only if  $X_{\text{rc}} = x_0$  almost surely.

A remaining question is what consistent scoring function  $S$  ought to be used in practice. To address this issue, we resort to mixture or Choquet representations of consistent loss functions, as introduced by Ehm et al. (2016) for quantiles and expectiles and developed in full generality by Dawid (2016), Ziegel (2016) and Jordan, Mühlemann and Ziegel (2022). Specifically, we rely on an obvious generalization of Proposition 2.6 of Jordan, Mühlemann and Ziegel (2022), as noted at the start of their Section 2. Let  $T$  be identifiable with identification function  $V$  satisfying Assumption 2.8, and let  $\eta \in \mathbb{R}$ . Then the *elementary* loss function  $S_\eta$ , given by

$$S_\eta(x, y) = (\mathbb{1}\{\eta \leq x\} - \mathbb{1}\{\eta \leq y\}) V(\eta, y), \quad (16)$$

is consistent for  $T$ . As an immediate consequence, any well-defined function of the form

$$S(x, y) = \int_{\mathbb{R}} S_\eta(x, y) dH(\eta), \quad (17)$$

TABLE 3  
*Canonical loss functions in the sense of Definition 2.24.*

Functional	Parameter	Canonical Loss
Moment of order $n$	$n = 1, 2, \dots$	$S(x, y) = (x - y^n)^2$
$\alpha$ -Expectile	$\alpha \in (0, 1)$	$S(x, y) = 2  \mathbb{1}\{x \geq y\} - \alpha  (x - y)^2$
$\alpha$ -Quantile	$\alpha \in (0, 1)$	$S(x, y) = 2 (\mathbb{1}\{x \geq y\} - \alpha) (x - y)$

TABLE 4  
*Components of the decomposition (15) for the mean squared error (MSE) under mean-forecasts induced by the predictive distributions in Examples 2.1 and 2.2. Uncertainty (UNC) equals 2 irrespective of the forecast at hand. The term  $I(\eta_0)$  is in integral form and can be evaluated numerically. For details see Appendix A.*

Predictive Distribution	Mean-Forecast	MSE	MCB	DSC
Perfect	$\mu$	1	0	1
Unconditional	0	2	0	0
Unfocused	$\mu + \frac{1}{2}\eta$	$1 + \frac{1}{4}\eta_0^2$	$(\frac{1}{4} - I(\eta_0))\eta_0^2$	$1 - I(\eta_0)\eta_0^2$
Lopsided	$\mu + \frac{\sqrt{2}}{\sqrt{\pi}}\delta$	$1 + \frac{2}{\pi}\delta_0^2$	$(\frac{1}{4} - I(\sqrt{\frac{8}{\pi}}\delta_0))\frac{8}{\pi}\delta_0^2$	$1 - I(\sqrt{\frac{8}{\pi}}\delta_0)\frac{8}{\pi}\delta_0^2$

where  $H$  is a locally finite measure on  $\mathbb{R}$ , is consistent for  $T$ . If  $T$  is a quantile, an expectile, an event probability or a moment, then the construction includes all consistent scoring functions, subject to standard conditions, and agrees with suitably adapted classes of generalized piecewise linear (GPL) and Bregman functions, respectively (Gneiting, 2011; Ehm et al., 2016).

We now formalize what Ehm and Ovcharov (2017, p. 477) call the “most prominent” choice, namely, scoring functions for which the mixing measure  $H$  in the representation (17) is uniform.

**Definition 2.24.** Suppose that the functional  $T$  is generated by an identification function  $V$  satisfying Assumption 2.8, with elementary loss functions  $S_\eta$  as defined in (16). Then a loss function  $S$  is *canonical* for  $T$  if it is nonnegative and admits a representation of the form

$$S(x, y) = a \int_{\mathbb{R}} S_\eta(x, y) d\lambda(\eta) + b(y), \quad (18)$$

where  $\lambda$  is the Lebesgue measure,  $a > 0$  is a constant, and  $b$  is a measurable function.

Clearly, any canonical loss function is a consistent scoring function for  $T$ . Furthermore, if the identification function is of the prediction error form, then any canonical loss function has score differentials that are invariant under translation in the sense that  $S(x_1 + c, y + c) - S(x_2 + c, y + c) = S(x_1, y) - S(x_2, y)$ . Conversely, we note from Section 5.1 of Savage (1971) that for the mean functional the canonical loss functions are the only consistent scoring functions of this type.

Typically, one chooses the constant  $a > 0$  and the measurable function  $b(y)$  in (18) such that the canonical loss admits a concise closed form, as exemplified

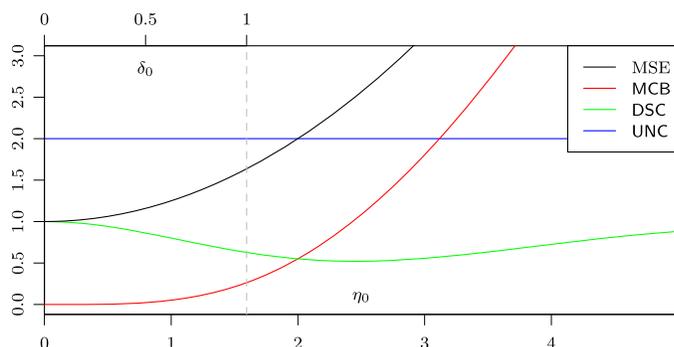


FIG 5. Components of the decomposition (15) for the mean squared error (MSE) under mean-forecasts induced by the unfocused and lopsided predictive distributions from Example 2.2 and Table 4, as functions of  $\eta_0 \geq 0$  and  $\delta_0 \in (0, 1)$ , respectively.

in Table 3. Since any selection incurs the same point forecast ranking, we refer to the choice in Table 3 as *the* canonical loss function. The most prominent example arises when  $T$  is the mean functional, where the ubiquitous *quadratic* or *squared error* scoring function,

$$S(x, y) = (x - y)^2, \tag{19}$$

is canonical. In this case, the UNC component equals the unconditional variance of  $Y$  as  $x_0$  is simply the marginal mean  $\mu_Y$  of  $Y$ , and the MCB and DSC components of the general score decomposition (15) are

$$\text{MCB} = \mathbb{E}(X - X_{rc})^2 \quad \text{and} \quad \text{DSC} = \mathbb{E}(X_{rc} - \mu_Y)^2,$$

respectively. Note that here and in the following, we drop the subscript  $S$  whenever we use a canonical loss. Table 4 and Figure 5 provide explicit examples.

In the nested case of a binary outcome  $Y$ , where  $X$  and  $X_{rc}$  specify event probabilities, the quadratic loss function reduces to the Brier score (Gneiting and Raftery, 2007), and we refer to Dimitriadis, Gneiting and Jordan (2021) and references therein for details on score decompositions. In the case of threshold calibration, the point forecast  $x = F(t)$  is induced by a predictive distribution, and the Brier score can be written as

$$S(x, y) = (F(t) - \mathbb{1}\{y \leq t\})^2. \tag{20}$$

For both real-valued and binary outcomes, it is often preferable to use the square root of the miscalibration component ( $\text{MCB}^{1/2}$ ) as a measure of calibration error that can be interpreted on natural scales (e.g., Roelofs et al., 2022).

A canonical loss function for the Huber functional (Table 1) is given by

$$S(x, y) = 2 |\mathbb{1}\{x \geq y\} - \alpha| \begin{cases} 2a|x - y| - a^2, & x - y < -a, \\ (x - y)^2, & -a \leq x - y \leq b, \\ 2b|x - y| - b^2, & x - y > b; \end{cases}$$

cf. Taggart (2022, Definition 4.2). In the limiting case as  $a = b \rightarrow \infty$ , we recover the canonical loss functions for the  $\alpha$ -expectile, which include the quadratic loss in (19). Similarly, if we rescale suitably and take the limit as  $a = b \rightarrow 0$ , we recover the asymmetric *piecewise linear* or *pinball* loss, as listed in Table 3, which lies at the heart of quantile regression.

We move on to a remarkable property of canonical loss functions. In a nutshell, the point forecast  $X$  is unconditionally T-calibrated if, and only if, the expected canonical loss deteriorates under translation. This property, which nests classical results in regression theory, as we demonstrate at the end of Section 3.3, does not hold under consistent scoring functions in general. For a counterexample see Appendix A.8.

**Assumption 2.25.** The point forecast  $X$ , the functional T and the identification function  $V$  satisfy Assumptions 2.8 and 2.19, and S is a canonical loss for T. Furthermore,  $\mathbb{E}[S(X + \eta, Y)]$  and  $\mathbb{E}[V(X + \eta, Y)]$  are well defined and locally bounded as functions of  $\eta \in \mathbb{R}$ .

**Theorem 2.26.** Under Assumption 2.25, the point forecast  $X$  is unconditionally T-calibrated if, and only if,

$$\mathbb{E}[S(X + c, Y)] \geq \mathbb{E}[S(X, Y)] \quad \text{for all } c \in \mathbb{R}.$$

*Proof.* If  $X$  is unconditionally T-calibrated and  $c > 0$ , then

$$\begin{aligned} & \mathbb{E}[S(X + c, Y)] - \mathbb{E}[S(X, Y)] \\ &= \mathbb{E} \left[ \int (\mathbb{1}\{\eta \leq X + c\} - \mathbb{1}\{\eta \leq X\}) V(\eta, Y) d\eta \right] \\ &= \mathbb{E} \left[ \int_{(0, c]} V(X + \eta, Y) d\eta \right] = \int_{(0, c]} \mathbb{E}[V(X + \eta, Y)] d\eta \end{aligned} \tag{21}$$

is nonnegative by the second part of the unconditional T-calibration criterion (13). Conversely, if the score difference in (21) is nonnegative for all  $c > 0$ , then so is

$$\mathbb{E}[V(X + c, Y)] = \frac{1}{c} \int_{(0, c]} \mathbb{E}[V(X + c, Y)] d\eta \geq \frac{1}{c} \int_{(0, c]} \mathbb{E}[V(X + \eta, Y)] d\eta$$

as the identification function  $V$  is non-decreasing in its first argument. Hence, the second part of (13) is satisfied.

An analogous argument shows that the score difference (21) is nonnegative for all  $c < 0$  if, and only if, the first inequality in (13) is satisfied.  $\square$

As a consequence, under a canonical loss function the MCB component in the score decomposition (15) of Theorem 2.23 decomposes into nonnegative *unconditional* and *conditional* components  $\text{MCB}_u$  and  $\text{MCB}_c$ , respectively, subject to the mild condition that unconditional recalibration via translation is feasible.

**Theorem 2.27.** *Let Assumption 2.25 hold, and suppose there is a constant  $c$  such that  $X + c$  is unconditionally T-calibrated. Let  $X_{\text{urc}} = X + c$  and  $\bar{S}_{\text{urc}} = \mathbb{E}[S(X_{\text{urc}}, Y)]$ , and define*

$$\text{MCB}_{\text{u}} = \bar{S} - \bar{S}_{\text{urc}} \quad \text{and} \quad \text{MCB}_{\text{c}} = \bar{S}_{\text{urc}} - \bar{S}_{\text{rc}}.$$

Then

$$\text{MCB} = \text{MCB}_{\text{u}} + \text{MCB}_{\text{c}},$$

where  $\text{MCB}_{\text{u}} \geq 0$  with equality if  $X$  is unconditionally T-calibrated, and  $\text{MCB}_{\text{c}} \geq 0$  with equality if  $X_{\text{rc}} = X_{\text{urc}}$  almost surely. If  $S$  is strictly consistent, then  $\text{MCB}_{\text{u}} = 0$  only if  $X$  is unconditionally T-calibrated, and  $\text{MCB}_{\text{c}} = 0$  only if  $X_{\text{rc}} = X_{\text{urc}}$  almost surely.

*Proof.* Immediate from Theorems 2.23 and 2.26, and the fact that conditional recalibration of  $X$  and  $X + c$  yields the same  $X_{\text{rc}}$ .  $\square$

In settings that are equivariant under translation, such as for expectiles, quantiles, and Huber functionals when both  $X$  and  $Y$  are supported on the real line,  $X$  can always be unconditionally recalibrated by adding a constant. Under any canonical loss function  $S$ , the basic decomposition (15) then extends to

$$\bar{S} = \text{MCB}_{\text{u}} + \text{MCB}_{\text{c}} - \text{DSC} + \text{UNC}. \quad (22)$$

For instance, when  $S(x, y) = (x - y)^2$  is the canonical loss for the mean functional,  $\text{MCB}_{\text{u}} = c^2$  is the squared unconditional bias. The forecasts in Figure 5 and Table 4 are free of unconditional bias, so  $\text{MCB}_{\text{u}} = 0$  and  $\text{MCB}_{\text{c}} = \text{MCB}$ .

In all cases studied thus far, canonical loss functions are strictly consistent (Ehm et al., 2016), and so  $\text{MCB}_{\text{u}} = 0$  if and only if the forecast is unconditionally T-calibrated, and  $\text{MCB}_{\text{c}} = 0$  if and only if  $X_{\text{urc}} = X_{\text{rc}}$  almost surely. While in other settings, such as when the outcomes are bounded, unconditional recalibration by translation might be counterintuitive (in principle) or impossible (in practice), the statement of Theorem 2.27 continues to hold, and the above results can be refined to admit more general forms of unconditional recalibration. We leave these and other ramifications to future work.

### 3. Empirical reliability diagrams and score decompositions: the CORP approach

We turn to empirical settings, where calibration checks, scores, and score decompositions address critical practical problems in both model diagnostics and forecast evaluation. The most direct usage is in the evaluation of out-of-sample predictive performance, where forecasts may either take the form of fully specified predictive distributions, or be single-valued point forecasts that arise, implicitly or explicitly, as functionals of predictive distributions. Similarly, in model diagnostics, where in-sample goodness-of-fit is of interest, the model might supply fully specified, parametric or non-parametric conditional distributions, or

single-valued regression output that is interpreted as a functional of an underlying, implicit or explicit, probability distribution. Prominent examples for the latter setting include ordinary least squares regression, where the mean or expectation functional is sought, and quantile regression.

In the case of fully specified predictive distributions, we work with tuples of the form

$$(F_1, y_1), \dots, (F_n, y_n), \quad (23)$$

where  $F_i$  is a posited conditional CDF for the real-valued observation  $y_i$  for  $i = 1, \dots, n$ , which we interpret as a sample from an underlying population  $\mathbb{Q}$  in the prediction space setting of Section 2. In the case of stand-alone point forecasts or regression output, we assume throughout that the functional  $T$  is of the type stated in Assumption 2.19 and work with tuples of the form

$$(x_1, y_1), \dots, (x_n, y_n), \quad (24)$$

where  $x_i = T(F_i) \in \mathbb{R}$  derives explicitly or implicitly from a predictive distribution  $F_i$  for  $i = 1, \dots, n$ .

In the remainder of the section, we introduce empirical versions of  $T$ -reliability diagrams (Definition 2.20) and score components (Definition 2.22) for samples of the form (23) or (24), which allow for both diagnostic checks and inference about an underlying population  $\mathbb{Q}$ . While practitioners may think of our empirical versions exclusively from diagnostic perspectives, we emphasize that they can be interpreted as estimators of the population quantities and be analyzed as such. A key feature of our approach is the use of nonparametric isotonic regression via the pool-adjacent-violators algorithm, as proposed by Dimitriadis, Gneiting and Jordan (2021) in the particular case of binary outcomes. The generalization that we discuss here is hinted at in the discussion section of their paper.

### 3.1. The $T$ -pool-adjacent-violators ( $T$ -PAV) algorithm

Our key tool and workhorse is a very general version of the classical pool-adjacent-violators (PAV) algorithm for nonparametric isotonic regression (Ayer et al., 1955; van Eeden, 1958). Historically, work on the PAV algorithm has focused on the mean functional, as reviewed by Barlow et al. (1972), Robertson and Wright (1980), and de Leeuw, Hornik and Mair (2009), among others. In contrast, Jordan, Mühlemann and Ziegel (2022) study the PAV algorithm in very general terms that accommodate our setting.

We rely on their work and describe the  $T$ -pool-adjacent-violators algorithm based on tuples  $(x_1, y_1), \dots, (x_n, y_n)$  of the form (24), where without loss of generality we may assume that  $x_1 \leq \dots \leq x_n$ . Furthermore, we let  $\delta_i$  denote the point measure in the outcome  $y_i$ . More generally, for  $1 \leq k \leq l \leq n$  we let

$$\delta_{k:l} = \frac{1}{l - k + 1} \sum_{i=k}^l \delta_i$$

be the associated empirical measure. Algorithm 1 describes the generation of an increasing sequence  $\hat{x}_1 \leq \dots \leq \hat{x}_n$  of recalibrated values, which by construction are conditionally T-calibrated with respect to the empirical measure associated with  $(\hat{x}_1, y_1), \dots, (\hat{x}_n, y_n)$ . The algorithm rests on partitions of the index set  $\{1, \dots, n\}$  into groups  $G_{k:l} = \{k, \dots, l\}$  of consecutive integers. Successive groups are pooled iteratively if the value of the functional applied to the empirical measure associated with the preceding group exceeds the value associated with the subsequent group.

---

**Algorithm 1:** General T-PAV algorithm based on data of the form (24)

---

**Input:**  $(x_1, y_1), \dots, (x_n, y_n) \in \mathbb{R}^2$  where  $x_1 \leq \dots \leq x_n$

**Output:** T-calibrated values  $\hat{x}_1, \dots, \hat{x}_n$

partition into groups  $G_{1:1}, \dots, G_{n:n}$  and let  $\hat{x}_i = T(\delta_i)$  for  $i = 1, \dots, n$

**while** there are groups  $G_{k:i}$  and  $G_{(i+1):l}$  such that  $\hat{x}_1 \leq \dots \leq \hat{x}_i$  and

$\hat{x}_i > \hat{x}_{i+1}$  **do**

  | merge  $G_{k:i}$  and  $G_{(i+1):l}$  into  $G_{k:l}$  and let  $\hat{x}_i = T(\delta_{k:l})$  for  $i = k, \dots, l$

**end**

---

The following result summarizes the remarkable properties of the T-PAV algorithm, as proved in Section 3.2 of Jordan, Mühlemann and Ziegel (2022).

**Theorem 3.1** (Jordan, Mühlemann and Ziegel, 2022). *Suppose that the functional T is as stated in Assumption 2.19. Then Algorithm 1 generates a sequence  $\hat{x}_1, \dots, \hat{x}_n$  such that the empirical measure associated with  $(\hat{x}_1, y_1), \dots, (\hat{x}_n, y_n)$  is conditionally T-calibrated. This sequence is optimal with respect to any scoring function S of the form (17), in that*

$$\frac{1}{n} \sum_{i=1}^n S(\hat{x}_i, y_i) \leq \frac{1}{n} \sum_{i=1}^n S(t_i, y_i) \quad (25)$$

for any non-decreasing sequence  $t_1 \leq \dots \leq t_n$ .

We note that for a functional T of interval type, the minimum on the left-hand side of (25) is the same under the lower functional  $T^-$  and the upper functional  $T^+$ , respectively, as defined in Section 2.4. For customary functionals, such as threshold (non) exceedance probabilities, quantiles, expectiles, and moments, the optimality is universal, as functions of the form (17) exhaust the class of the T-consistent scoring functions subject to mild conditions (Ehm et al., 2016). While the PAV algorithm has been used extensively for the recalibration of probabilistic classifiers (e.g., Flach, 2012), we are unaware of any extant work that uses Algorithm 1 for forecast recalibration, forecast evaluation, or model diagnostics in non-binary settings.

### 3.2. Empirical T-reliability diagrams

Recently, Dimitriadis, Gneiting and Jordan (2021) introduced the CORP approach for the estimation of reliability diagrams and score decompositions in

the case of probability forecasts for binary outcomes. In a nutshell, the acronym CORP refers to an estimator that is Consistent under the assumption of isotonicity for the population recalibration function and Optimal in both finite sample and asymptotic settings, while facilitating Reproducibility, and being based on the PAV algorithm. Here, we extend the CORP approach and employ nonparametric isotonic T-regression via the T-PAV algorithm under Assumption 2.19, where T is the lower or upper version of an identifiable functional, or an identifiable singleton functional.

We begin by defining the empirical T-reliability diagram, which is a sample version of the population diagram in Definition 2.20.

**Definition 3.2.** Let the functional T be as stated in Assumption 2.19, and suppose that  $\hat{x}_1, \dots, \hat{x}_n$  originate from tuples  $(x_1, y_1), \dots, (x_n, y_n)$  with  $x_1 \leq \dots \leq x_n$  via Algorithm 1. Then the CORP *empirical T-reliability diagram* is the graph of the piecewise linear function that connects the points  $(x_1, \hat{x}_1), \dots, (x_n, \hat{x}_n)$  in the Euclidean plane.

A few scattered references in the literature on forecast evaluation have proposed displays of recalibrated against original values for functionals other than binary event probabilities: Figures 3 and 7 of Bentzien and Friederichs (2014) and Figure 8 of Pohle (2020) consider quantiles, and Figures 2–5 of Satopää and Ungar (2015) concern the mean functional. However, none of these papers employ the PAV algorithm, and the resulting diagrams are subject to issues of stability and efficiency, as illustrated by Dimitriadis, Gneiting and Jordan (2021) in the case of binary outcomes.

For the CORP empirical T-reliability diagram to be consistent in the sense of large sample convergence to the population version of Definition 2.20, the assumption of isotonicity of the population recalibration function needs to be invoked. As argued by Roelofs et al. (2022) and Dimitriadis, Gneiting and Jordan (2021), such an assumption is natural, and practitioners tend to dismiss non-isotonic recalibration functions as artifacts. Evidently, these arguments transfer to arbitrary functionals, and any violations of the isotonicity assumption entail horizontal segments in CORP reliability diagrams, thereby indicating a lack of reliability. Large sample theory for CORP estimates of the recalibration function and the T-reliability diagram depends on the functional T, the type — discrete or continuous — of the marginal distribution of the point forecast  $X$ , and smoothness conditions. Mösching and Dümbgen (2020) establish rates of uniform convergence in the cases of threshold (non) exceedance and quantile functionals that complement classical theory (Barlow et al., 1972; Casady and Cryer, 1976; Wright, 1984; Robertson, Wright and Dykstra, 1988; El Barmi and Mukerjee, 2005; Guntuboyina and Sen, 2018).

In the case of binary outcomes, Bröcker and Smith (2007, p. 651) argue that reliability diagrams ought to be supplemented by *consistency bars* for “immediate visual evaluation as to just how likely the observed relative frequencies are under the assumption that the predicted probabilities are reliable.” Dimitriadis, Gneiting and Jordan (2021) develop asymptotic and Monte Carlo based methods for the generation of consistency bands to accompany a CORP rela-

bility diagram for dichotomous outcomes, and provide code in the form of the `reliabilitydiag` package (Dimitriadis and Jordan, 2021) for R (R Core Team, 2021). The consistency bands quantify and visualize the variability of the empirical reliability diagram under the respective null hypothesis, i.e., they show the pointwise range of the CORP T-reliability diagram that we expect to see under a calibrated forecast. Algorithms 2 and 3 in Appendix B generalize this approach to produce consistency bands from data of the form (23) under the assumption of auto-calibration. In the specific case of threshold calibration, where the induced outcome is dichotomous, the assumptions of auto-calibration (in the binary setting) and T-calibration (for the non-exceedance functional) coincide (Gneiting and Ranjan, 2013, Theorem 2.11), and we use the aforementioned algorithms to generate consistency bands (Figure 6, top row). Generally, auto-calibration is a strictly stronger assumption than T-calibration, with ensuing issues, which we discuss in Appendix B.1. Furthermore, to generate consistency bands from data of the form (24), we cannot operate under the assumption of auto-calibration.

As a crude yet viable alternative, we propose in Appendix B.2 a Monte Carlo technique for the generation of consistency bands that is based on resampling residuals. As in traditional regression diagnostics, the approach depends on the assumption of independence between point forecasts and residuals. Figure 6 shows examples of T-reliability diagrams with associated residual-based 90% consistency bands for the perfect, unfocused, and lopsided forecasts from Section 2 for the mean functional (middle row) and the lower quantile functional at level 0.10 (bottom row). For further discussion see Appendix B.2. In the case of the mean functional, we add the scatter diagram for the original data of the form (24), whereas in the other two cases, inset histograms visualize the marginal distribution of the point forecast.

We encourage follow-up work on both Monte Carlo and asymptotic methods for the generation of consistency and confidence bands that are tailored to specific functionals of interest, similar to the analysis by Dimitriadis, Gneiting and Jordan (2021) in the basic case of probability forecasts for binary outcomes.

### 3.3. Empirical score decompositions

In this section, we consider data  $(x_1, y_1), \dots, (x_n, y_n)$  of the form (24), where implicitly or explicitly  $x_i = T(F_i)$  for a single-valued functional T. Let  $\hat{x}_1, \dots, \hat{x}_n$  denote the respective T-PAV recalibrated values, and let  $\hat{x}_0 = T(\hat{F}_0)$ , where  $\hat{F}_0$  is the empirical CDF of the outcomes  $y_1, \dots, y_n$ . Let

$$\hat{S} = \frac{1}{n} \sum_{i=1}^n S(x_i, y_i), \quad \hat{S}_{\text{rc}} = \frac{1}{n} \sum_{i=1}^n S(\hat{x}_i, y_i), \quad \text{and} \quad \hat{S}_{\text{mg}} = \frac{1}{n} \sum_{i=1}^n S(\hat{x}_0, y_i) \quad (26)$$

denote the mean score of the point forecast at hand, the recalibrated point forecast, and the functional T applied to the unconditional, marginal distribution

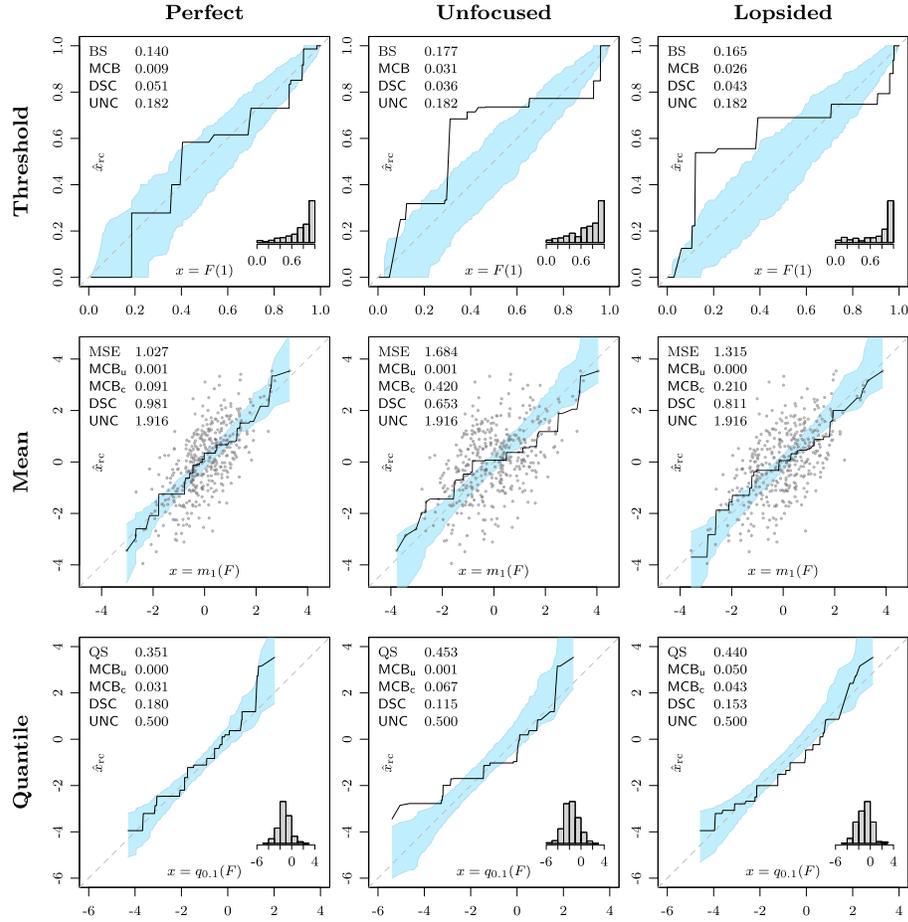


FIG 6. CORP empirical threshold (top,  $t = 1$ ), mean (middle) and quantile (bottom,  $\alpha = 0.10$ ) reliability diagrams for the perfect (left), unfocused (middle), and lopsided (right) forecast from Examples 2.1 and 2.2 with 90% consistency bands and CORP score components under the associated canonical loss function based on samples of size 400.

of the outcome, respectively. If all quantities in (26) are finite, we refer to

$$\widehat{\text{MCB}}_S = \widehat{S} - \widehat{S}_{\text{rc}}, \quad \widehat{\text{DSC}}_S = \widehat{S}_{\text{mg}} - \widehat{S}_{\text{rc}}, \quad \text{and} \quad \widehat{\text{UNC}}_S = \widehat{S}_{\text{mg}} \quad (27)$$

as the *miscalibration*, *discrimination* and *uncertainty* components of the mean score  $\widehat{S}$ . Our next result generalizes Theorem 1 of Dimitriadis, Gneiting and Jordan (2021) and decomposes the mean score  $\widehat{S}$  into a signed sum of nonnegative, readily interpretable components.

**Theorem 3.3.** *Suppose that the functional  $T$  satisfies the conditions in Assumption 2.19. Let the scoring function  $S$  be of the form (17), suppose that  $\widehat{x}_1, \dots, \widehat{x}_n$  originate from tuples  $(x_1, y_1), \dots, (x_n, y_n)$  via Algorithm 1, and let all*

terms in (26) be finite. Then

$$\widehat{S} = \widehat{MCB}_S - \widehat{DSC}_S + \widehat{UNC}_S, \tag{28}$$

where  $\widehat{MCB}_S \geq 0$  with equality if  $\widehat{x}_i = x_i$  for  $i = 1, \dots, n$ , and  $\widehat{DSC}_S \geq 0$  with equality if  $\widehat{x}_i = \widehat{x}_0$  for  $i = 1, \dots, n$ .

If  $S$  is strictly consistent, then  $\widehat{MCB}_S = 0$  only if  $\widehat{x}_i = x_i$  for  $i = 1, \dots, n$  and  $\widehat{DSC}_S = 0$  only if  $\widehat{x}_i = \widehat{x}_0$  for  $i = 1, \dots, n$ .

*Proof.* Immediate from Theorem 3.1. □

Thus, CORP estimates of score components enjoy the same properties as the respective population quantities (Theorem 2.23, eq. (15)). This agreement is not to be taken for granted, as the nonnegativity of the estimated components cannot be guaranteed if approaches other than the T-PAV algorithm are used for recalibration (Dimitriadis, Gneiting and Jordan, 2021, Supplementary Section S5).

Recently, the estimation of calibration error has seen a surge of interest in machine learning (Guo et al., 2017; Kuleshov, Fenner and Ermon, 2018; Kumar, Liang and Ma, 2019; Nixon et al., 2019; Roelofs et al., 2022). Under the natural assumption of isotonicity of the population recalibration function,  $\widehat{MCB}_S$  is a consistent estimate of the population quantity  $MCB_S$ , with canonical loss functions being natural choices for  $S$ . As noted, it is often preferable to use the square root of the miscalibration component under squared error as a measure of calibration error that can be interpreted on natural scales. Asymptotic distributions for our estimators depend on the functional  $T$ , the scoring function  $S$ , and regularity conditions. Large sample theory can leverage extant theory for nonparametric isotonic regression, as hinted at in the previous section, though score components might show distinct asymptotic behavior. Further development is beyond the scope of the present paper and strongly encouraged.

In the remainder of the section, we assume that  $S$  is a canonical score and drop the subscript in the score components. If there is a constant  $\widehat{c} \in \mathbb{R}$  such that the empirical measure in  $(x_1 + \widehat{c}, y_1), \dots, (x_n + \widehat{c}, y_n)$  is unconditionally  $T$ -calibrated, let

$$\widehat{S}_{\text{urc}} = \frac{1}{n} \sum_{i=1}^n S(x_i + \widehat{c}, y_i). \tag{29}$$

We then refer to

$$\widehat{MCB}_u = \widehat{S} - \widehat{S}_{\text{urc}} \quad \text{and} \quad \widehat{MCB}_c = \widehat{S}_{\text{urc}} - \widehat{S}_{\text{rc}}$$

as the CORP *unconditional* and *conditional* miscalibration components of the mean canonical score, respectively. Under mild conditions, these estimates are nonnegative and share properties of the respective population quantities in Theorem 2.27.

**Theorem 3.4.** *Let the conditions of Theorem 3.3 hold, and let  $S$  be a canonical loss function for  $T$ . Suppose there is a constant  $\hat{c} \in \mathbb{R}$  such that the empirical measure in  $(x_1 + \hat{c}, y_1), \dots, (x_n + \hat{c}, y_n)$  is unconditionally  $T$ -calibrated, and suppose that all terms in (29) are finite. Then*

$$\widehat{MCB} = \widehat{MCB}_u + \widehat{MCB}_c,$$

where  $\widehat{MCB}_u \geq 0$  and  $\widehat{MCB}_c \geq 0$ .

*Proof.* Immediate from Theorems 2.26 and 3.1 and the trivial fact that the addition of a constant is a special case of an isotonic mapping.  $\square$

In the middle row of Figure 6, the extended CORP decomposition,

$$\widehat{S} = \widehat{MCB}_u + \widehat{MCB}_c - \widehat{DSC} + \widehat{UNC}, \quad (30)$$

which estimates the population decomposition (22), is applied to the mean squared error (MSE). Likewise, the extended CORP decomposition of the canonical score for quantiles, i.e., the piecewise linear quantile score (QS) from Table 3, is shown in the bottom row. The top row concerns threshold calibration, and we report the standard CORP decomposition (28) of the Brier score (BS) from (20). While the assumptions of Theorem 3.3 are satisfied in this setting, the addition of the constant  $\hat{c}$  may yield forecast values outside the unit interval, whence we refrain from considering the refined decomposition in (30).

In this context, the distinction between out-of-sample forecast evaluation and in-sample model diagnostics is critical. When evaluating out-of-sample forecasts, both unconditional and conditional miscalibration are relevant. In contrast, in-sample model fits frequently enforce unconditional calibration. For example, if we fit a regression model with intercept by minimizing the canonical loss for a functional  $T$ , Theorem 2.26 applied to the associated empirical measure guarantees in-sample unconditional  $T$ -calibration. As special cases, this line of reasoning yields classical results in ordinary least squares regression, and the partitioning inequalities of quantile regression in Theorem 3.4 of Koenker and Bassett (1978).

### 3.4. Skill scores and a universal coefficient of determination

Let us revisit the mean scores in (26) under the natural assumption that the terms in  $\widehat{S}$  and  $\widehat{S}_{mg}$  are finite and that  $\widehat{S}_{mg}$  is strictly positive. In out-of-sample forecast evaluation, the quantity

$$\widehat{S}_{skill} = 1 - \frac{\widehat{S}}{\widehat{S}_{mg}} = \frac{\widehat{S}_{mg} - \widehat{S}}{\widehat{S}_{mg}} = \frac{\widehat{DSC}_S - \widehat{MCB}_S}{\widehat{UNC}_S} \quad (31)$$

is known as *skill score* (Murphy and Epstein, 1989; Murphy, 1996; Gneiting and Raftery, 2007; Jolliffe and Stephenson, 2012) and may attain both positive and negative values. In particular, when  $S(x, y) = (x - y)^2$  is the canonical loss

function for the mean functional,  $\widehat{S}_{\text{skill}}$  coincides with the popular Nash-Sutcliffe model efficiency coefficient (NSE; Nash and Sutcliffe, 1970; Moriasi et al., 2007). A positive skill score indicates predictive performance better than the simplistic unconditional reference forecast, whereas a negative skill score suggests that we are better off using the simple reference forecast. Of course, it is possible, and frequently advisable, to base skill scores on reference standards that are more sophisticated than an unconditional, constant point forecast (Hyndman and Koehler, 2006).

In contrast, if the goal is in-sample model diagnostics, the quantity in (31) typically is nonnegative. As we demonstrate now, it constitutes a powerful generalization of the coefficient of determination,  $R^2$ , or variance explained in least squares regression, and its close cousin, the  $R^1$ -measure in quantile regression (Koenker and Machado, 1999). Specifically, we propose the use of

$$R^* = \frac{\widehat{DSC}_S - \widehat{MCB}_S}{\widehat{UNC}_S}, \quad (32)$$

as a universal *coefficient of determination*. In practice, one takes  $S$  to be a canonical loss for the functional  $T$  at hand, and we drop the subscripts in this case. The classical  $R^2$  measure arises when  $S(x, y) = (x - y)^2$  is the canonical squared error loss function for the mean functional, and the  $R^1$  measure of Koenker and Machado (1999) emerges when  $S(x, y) = 2(\mathbb{1}\{x \geq y\} - \alpha)(x - y)$  is the canonical piecewise linear loss under the  $\alpha$ -quantile functional. Of course, in the case  $\alpha = \frac{1}{2}$  of the median, the piecewise linear loss reduces to the absolute error.

In Figure 7, we present a numerical example on the toy data from Figure 1 in Kvålseth (1985). The straight lines show the linear (ordinary least squares) mean and linear (Laplace) median regression fits, which Kvålseth (1985) sought to compare. The piecewise linear broken curves illustrate the nonparametric isotonic regression fits, as realized by the T-PAV algorithm, where  $T$  is the mean and the lower and the upper median, respectively. As the linear regression fits induce the same ranking of the point forecasts, they yield the same PAV-recalibrated values that enter the terms in the score decomposition (27), and thus they have identical discrimination components in (28), which equal 10.593 under squared error and 2.333 under absolute error, regardless of which isotonic median is used. The uncertainty components, which equal 12.000 under squared error, and 2.889 under absolute error, are identical as well, since they depend on the observations only. Thus, the differences in  $R^2$  respectively  $R^1$  in Figure 7 stem from distinct miscalibration components. Of course, linear mean regression is preferred under squared error, and linear median regression is preferred under absolute error.

Various authors have discussed desiderata for a generally applicable definition of a coefficient of determination (Kvålseth, 1985; Nakagawa and Schielzeth, 2013) for the assessment of in-sample fit. In particular, such a coefficient ought to be dimensionless and take values in the unit interval, with a value of 1 indicating a perfect fit, and a value of 0 representing a complete lack of fit. The

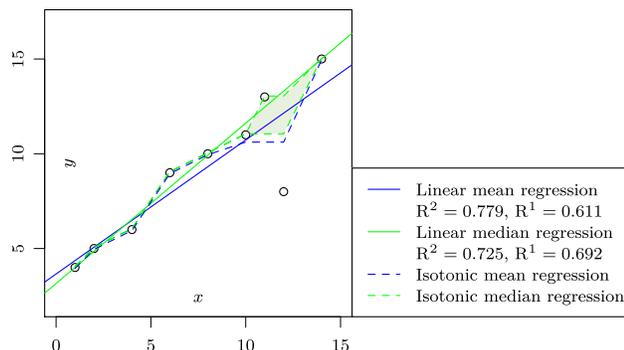


FIG 7. Linear mean and linear median regression lines for toy example from [Kvålseth \(1985, Figure 1\)](#), along with nonparametric isotonic mean and median regression fits. The isotonic median regression fit is not unique and framed by the respective lower and upper functional.

universal coefficient of determination  $R^*$  enjoys these properties under modest conditions.

**Assumption 3.5.** Suppose that the functional  $T$  is as stated in Assumption 2.19 with associated identification function  $V$ . Let the scoring function  $S$  be of the form (17), and suppose that  $\hat{x}_1, \dots, \hat{x}_n$  in (26) originate from tuples  $(x_1, y_1), \dots, (x_n, y_n)$  via Algorithm 1. Furthermore, let the following hold.

- (i) The terms contributing to  $\hat{S}$  and  $\hat{S}_{\text{mg}}$  in (26) are finite, and  $\hat{S}_{\text{mg}} > 0$ .
- (ii) The values  $x_1, \dots, x_n$  have been fitted to  $y_1, \dots, y_n$  by in-sample empirical loss minimization with respect to  $S$ , with any constant fit  $x_1 = \dots = x_n$  being admissible.

For example, suppose that  $T$  is the mean functional and  $S$  is the canonical squared error scoring function. Then condition (i) is satisfied with the exception of the trivial case where  $y_1 = \dots = y_n$ , and condition (ii) is satisfied under linear (ordinary least squares) mean regression with intercept. Similarly, if  $T$  is a quantile and  $S$  is the canonical piecewise linear loss function, then (i) is satisfied except when  $y_1 = \dots = y_n$ , and (ii) is satisfied under linear quantile regression with intercept. In this light, the following theorem covers the classical settings for the  $R^2$  and  $R^1$  measures.

**Theorem 3.6.** Under Assumption 3.5 it holds that

$$R^* \in [0, 1]$$

with  $R^* = 0$  if  $x_i = \hat{x}_0$  for  $i = 1, \dots, n$ , and  $R^* = 1$  if  $x_i = T(\delta_i)$  for  $i = 1, \dots, n$ .

*Proof.* The claim follows from Theorem 3.1, the trivial fact that a constant fit is a special case of an isotonic mapping, and the assumed form (17) of the scoring function.  $\square$

We emphasize that Assumption 3.5 and Theorem 3.6 are concerned with, and tailored to, in-sample model diagnostics. At the expense of technicalities, the regularity conditions can be relaxed, but the details are tedious and we leave them to subsequent work. The condition that any constant fit  $x_1 = \dots = x_n$  be admissible is critical and cannot be relaxed.

### 3.5. Empirical examples

We now illustrate the use of reliability diagrams, score decompositions, skill scores, and the coefficient of determination  $R^*$  for the purposes of forecast evaluation and model diagnostics.

In the basic setting of tuples  $(x_1, y_1), \dots, (x_n, y_n)$  of the form (24), the point forecast  $x_i$  represents the functional  $T$  of a posited distribution for  $y_i$ . The most prominent case of the mean functional and canonical squared error loss (19) is illustrated in Figure 2, where point forecasts by Tredennick et al. (2021) of (log-transformed) butterfly population size are assessed. The CORP mean reliability diagram along with 90% consistency bands under the hypothesis of mean calibration complements the scatter plot provided by Tredennick et al. (2021, Figure 6). With a mean squared error (MSE) of 0.224, ridge regression performs much better than the null model with an MSE of 0.262. The CORP score decomposition shown in Figure 2 refines and supports the analysis.

We move on to discuss the more complex setting of tuples  $(F_1, y_1), \dots, (F_n, y_n)$  of the form (24), where  $F_i$  is a posited distribution for  $y_i$  ( $i = 1, \dots, n$ ). As discussed in Section 2 and visualized in Figure 1, the traditional unconditional notions of calibration, namely, probabilistic and marginal calibration, constitute weak forms of reliability. For this very reason, we recommend that checks for probabilistic and marginal calibration are given priority in this setting, much in line with current practice. Typically, probabilistic calibration is checked by plotting histograms of empirical probability integral transform (PIT) values (Diebold, Gunther and Tay, 1998; Gneiting, Balabdaoui and Raftery, 2007), though this practice is hindered by the need for binning. In Appendix B.3, we discuss the *PIT reliability diagram*, a rarely used alternative that avoids binning and retains the spirit of our CORP approach by plotting the CDF of the empirical PIT values. Similarly, as we also discuss in Appendix B.3, the *marginal reliability diagram* can be used to assess marginal calibration in the spirit of the CORP approach. If the analysis indicates gross violations of probabilistic or marginal calibration, we note from Section 2 and Figure 1 that key notions of conditional calibration must be violated as well. Otherwise, we might proceed to check stronger conditional notions of calibration, such as threshold, mean, and quantile calibration.

To illustrate this process, we consider quarterly Bank of England forecasts of consumer price index (CPI) inflation rates, as issued since 2004. The forecast distributions, for which we give details and refer to extant analyses in Appendix C.2, are two-piece normal distributions that are communicated to the public via fan charts. The forecasts are at prediction horizons up to six quarters

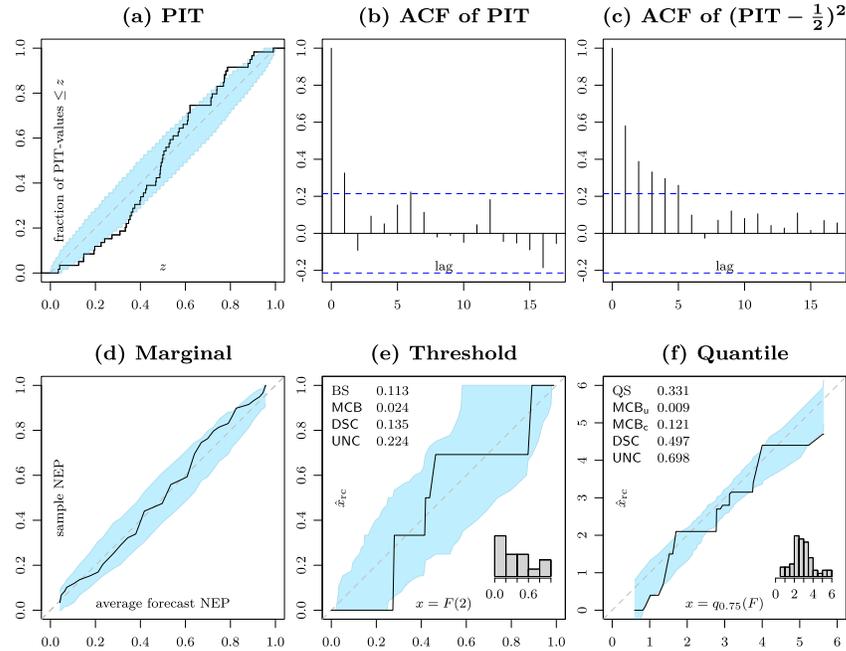


FIG 8. Calibration diagnostics for Bank of England forecasts of CPI inflation at a prediction horizon of one quarter: (a) PIT reliability diagram, along with the empirical autocorrelation functions of (b) original and (c) squared, centered PIT values, (d) marginal, and (f) 75%-quantile reliability diagram. If applicable, we show 90% consistency bands and CORP score components under the associated canonical loss function, namely, the Brier score (BS) and the piecewise linear quantile score (QS), respectively.

ahead in the time series setting, where  $k$  step ahead forecasts that are ideal with respect to the canonical filtration show PIT values that are independent at lags  $\geq k + 1$  in addition to being uniformly distributed (Diebold, Gunther and Tay, 1998). However, as discussed in Appendix C.1, independent, uniformly distributed PIT values do not imply auto-calibration, except in a special case. Thus, calibration diagnostics beyond checks of the uniformity and independence of the PIT are warranted.

In Figure 8, we consider forecasts one quarter ahead and show PIT and marginal reliability diagrams, along with empirical autocorrelation functions (ACFs) for the first two moments of the PIT. In part, the PIT reliability diagram and the ACFs lie outside the respective 90% consistency bands. For a closer look, we also plot the threshold reliability diagram at the policy target of 2% and the lower  $\alpha$ -quantile reliability diagram for  $\alpha = 0.75$ . The deviations from reliability remain minor, in stark contrast to calibration diagnostics at prediction horizons  $k \geq 4$ , for which we refer to Appendix C.2.

Figure 9 shows the standard CORP decomposition (28) of the Brier score (BS) for the induced probability forecasts at the 2% target and the extended

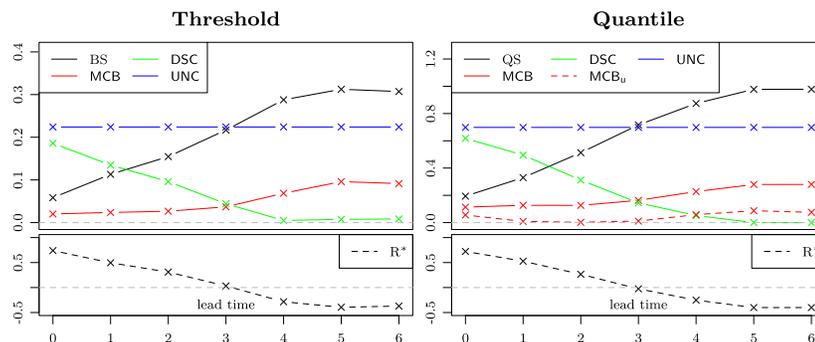


FIG 9. Score decomposition (28) respectively (30) and skill score (31) for probability forecasts of not exceeding the 2% inflation target (left) and 75%-quantile forecasts (right) induced by Bank of England fan charts for CPI inflation, under the associated canonical scoring function.

CORP decomposition (30) of the piecewise linear quantile score for  $\alpha$ -quantile forecasts at level  $\alpha = 0.75$  and lead times up to six quarters ahead. In the latter case, the difference between MCB and  $MCB_u$  equals the  $MCB_c$  component. Generally, the miscalibration components increase while the discrimination components decrease with the lead time. Related results for the quantile functional can be found in Pohle (2020, Table 5, Figures 7 and 8), where there is a notable increase in the discrimination (resolution) component at the largest two lead times, which is caused by counterintuitive decays in the recalibration functions. In contrast, the regularizing constraint of isotonicity prevents overfitting in the CORP approach.

The coefficient of determination or skill score  $R^*$  decays with the prediction horizon and becomes negative at lead times  $k \geq 4$ . This observation suggests that forecasts remain informative at lead times up to at most three quarters ahead, in line with the substantive findings in Pohle (2020) and other extant work, as hinted at in Appendix C.2.

#### 4. Discussion

We have developed a comprehensive theoretical and methodological framework for the analysis of calibration and reliability, serving the purposes of both (out-of-sample) forecast evaluation and (in-sample) model diagnostics. A common principle is that fitted or predicted distributions ought to be calibrated or reliable, ideally in the sense of auto-calibration, which stipulates that the outcomes are random draws from the posited distributions. For general real-valued outcomes, we have seen that auto-calibration is stronger than both classical unconditional and recently proposed conditional notions of calibration. We have developed hierarchies of calibration in the spirit of Van Calster et al. (2016), as highlighted in Figure 1, and proposed a generic notion of conditional calibration in terms of statistical functionals. Specifically, a posited distribution

is conditionally T-calibrated if the induced point forecast for the functional T can be taken at face value. This concept continues to apply when stand-alone point forecasts or regression output in terms of the functional T are considered, so T-reliability diagrams and associated score decompositions can be used in these settings. Importantly, our tools apply regardless of how forecasts are generated, be it through the use of traditional statistical regression models, modern machine learning techniques, or even subjective human judgment.

We have adopted and generalized the nonparametric approach of [Dimitriadis, Gneiting and Jordan \(2021\)](#), who obtained consistent, optimally binned, reproducible, and PAV based (CORP) estimators of T-reliability diagrams and score components in the case of probability forecasts for binary outcomes. While our tools apply in the much broader setting of identifiable functionals and real-valued outcomes, the arguments put forth by [Dimitriadis, Gneiting and Jordan \(2021\)](#) continue to apply, in that CORP estimators are bound to, simultaneously, improve statistical efficiency, reproducibility ([Stodden et al., 2016](#)), and stability ([Yu and Kumbier, 2020](#)). In a nutshell, the CORP approach is flexible, due to its use of nonparametric regression for recalibration, and yet it avoids overfitting, owing to the regularizing constraint of isotonicity. Notably, the CORP score decomposition yields a new, universal coefficient of determination,  $R^*$ , that nests and generalizes the classical  $R^2$  in ordinary least squares (mean) regression, and its cousin  $R^1$  in quantile regression. In independent work, [Allen \(2021\)](#) also observes the link between skill scores, score decompositions, and the coefficient of determination. We have illustrated the CORP approach on Bank of England forecasts of inflation, along with a brief ecological example. [Gneiting et al. \(2023\)](#) provide an in depth review of the particular case of forecasts in the form of (one or multiple) quantiles, accompanied by case studies. Code in R ([R Core Team, 2021](#)) for reproducing our results is available at [https://github.com/resinj/replication\\_GR21](https://github.com/resinj/replication_GR21) ([Gneiting and Resin, 2023](#)).

Follow-up work on the CORP approach for specific functionals T is essential, including but not limited to the ubiquitous cases of quantiles and the mean functional, where the newly developed tools can supplement classical approaches to regression diagnostics, as hinted at in the ecological example. In particular, we have applied a crude, all-purpose, residual-based permutation approach to generate consistency bands for T-reliability diagrams under the hypothesis of T-calibration. Clearly, this approach can be refined, and we anticipate vigorous work on consistency and confidence bands, based on either resampling or large sample theory, akin to the developments in [Dimitriadis, Gneiting and Jordan \(2021\)](#) for probability forecasts of binary outcomes. Similarly, CORP estimates of miscalibration components under canonical loss functions are natural candidates for the quantification of calibration error in empirical work. Reliability and discrimination ability are complementary attributes of point forecasts and regression output, and discrimination can be assessed quantitatively via the respective score component. When many forecasts are to be compared with each other, scatter plots of CORP miscalibration (MCB) and discrimination (DSC) components admit succinct visual displays of predictive performance. In this type of display, forecasts with the same score or, equivalently, identical coeffi-

cient of determination,  $R^*$ , gather on lines with unit slope, and multiple facets of forecast quality can be assessed simultaneously, for a general alternative to the widely used Taylor (2001) diagram.

Formal tests of hypotheses of calibration are critical in both specific applications, such as banking regulation (e.g., Nolde and Ziegel, 2017), and in generic tasks, such as the assessment of goodness of fit in regression (Dimitriadis, Gneiting and Jordan, 2021, Section S2). In Appendix B.4, we comment on this problem from the perspective of the theoretical and methodological advances presented here. While specific developments need to be deferred to future work, it is our belief that the progress in our understanding of notions and hierarchies of calibration, paired with the CORP approach to estimating reliability diagrams and score components, can spur a wealth of new and fruitful developments in these directions.

## Appendix A: Supporting calculations for Section 2

Here, we provide supporting computations and discussion for Examples 2.2 and 2.4, Definitions 2.9 and 2.24, Figures 4 and 5, and Table 4, along with a discussion of the relation between probabilistic calibration and unconditional quantile calibration, and a counterexample hinted at in the main text. For subsequent use, the first three (non-centered) moments of the normal distribution  $\mathcal{N}(\mu, \sigma^2)$  are  $\mu$ ,  $\mu^2 + \sigma^2$ , and  $\mu^3 + 3\mu\sigma^2$ . As in the main text, we let  $\varphi$  and  $\Phi$  denote the density and the cumulative distribution function (CDF), respectively, of a standard normal variable.

### A.1. Unfocused forecast

For fixed  $a, b \in \mathbb{R}$ , the function

$$y \mapsto \Phi_a(y - b) = \frac{1}{2} (\Phi(y - b) + \Phi(y - a - b))$$

is a CDF. The random CDF (2) for the unfocused forecast in Example 2.2 can be written as  $F(y) = \Phi_\eta(y - \mu)$ , where  $\eta$  and  $\mu$  are independent random variables and  $\eta = \pm\eta_0$  for some constant  $\eta_0 > 0$ . Then the conditional CDF for the outcome  $Y$  given the posited (non) exceedance probability  $F(t)$  at any fixed threshold  $t \in \mathbb{R}$  or, equivalently, given the quantile forecast  $F^{-1}(\alpha)$  at any fixed level  $\alpha \in (0, 1)$  is

$$\begin{aligned} \mathbb{Q}(Y \leq y \mid F(t) = \alpha) &= \mathbb{Q}(Y \leq y \mid F^{-1}(\alpha) = t) = \mathbb{Q}(Y \leq y \mid \mu = t - \Phi_\eta^{-1}(\alpha)) \\ &= \frac{1}{\sum_{s=\pm 1} \varphi(t - \Phi_{s\eta_0}^{-1}(\alpha))} \sum_{s=\pm 1} \varphi(t - \Phi_{s\eta_0}^{-1}(\alpha)) \Phi(y - (t - \Phi_{s\eta_0}^{-1}(\alpha))). \end{aligned}$$

As  $F$  is symmetric, conditioning on the mean is the same as conditioning on the median. The second moment is  $m_2(F) = 1 + \mu^2 + \mu\eta + \frac{1}{2}\eta^2 \geq 1 + \frac{1}{4}\eta^2$ , so that

$$\mathbb{Q}(Y \leq y \mid m_2(F) = m) = \mathbb{Q}\left(Y \leq y \mid \mu = -\frac{1}{2}\eta \pm \sqrt{m - 1 - \frac{1}{4}\eta^2}\right)$$

is a mixture of normal distributions. Similarly, the third moment is  $m_3(F) = \mu^3 + \frac{3}{2}\eta\mu^2 + 3\left(\frac{1}{2}\eta^2 + 1\right)\mu + \frac{1}{2}\eta(\eta^2 + 3) = f(\mu; \eta)$ , so that  $\mathbb{Q}(Y \leq y \mid m_3(F) = m) = \mathbb{Q}(Y \leq y \mid f(\mu; \eta) = m)$  also is a mixture of normal distributions. To compute the roots of the mapping  $x \mapsto f(x; \eta)$ , we fix  $\eta$  at  $\pm\eta_0$  and use a numeric solver (`polyroot` in R).

As regards the score decomposition (15) with  $S(x, y) = (x - y)^2$  for the implied mean-forecast,  $m_1(F) = \mu + \frac{1}{2}\eta$ , the expected score of the recalibrated mean-forecast is

$$\begin{aligned} \bar{S}_{\text{rc}} &= \mathbb{E} \left[ \frac{\sum_{s=\pm 1} \varphi(m_1(F) + \frac{s}{2}\eta_0) (m_1(F) + \frac{s}{2}\eta_0)}{\sum_{s=\pm 1} \varphi(m_1(F) + \frac{s}{2}\eta_0)} - Y \right]^2 \\ &= \mathbb{E} \left[ \frac{\sum_{s=\pm 1} \varphi(\mu + \frac{1}{2}\eta + \frac{s}{2}\eta_0) 0 (\frac{1}{2}\eta + \frac{s}{2}\eta_0)}{\sum_{s=\pm 1} \varphi(\mu + \frac{1}{2}\eta + \frac{s}{2}\eta_0)} - (Y - \mu) \right]^2 \\ &= \eta_0^2 \mathbb{E} \left[ \frac{\varphi(\mu + \eta)}{\varphi(\mu) + \varphi(\mu + \eta)} \right]^2 + \mathbb{E}[Y - \mu]^2 \\ &= \eta_0^2 \mathbb{E}[\Psi_{\eta_0}^2(\mu)] + 1, \end{aligned}$$

where we define  $\Psi_a(x) = \varphi(x + a)/(\varphi(x) + \varphi(x + a))$  for  $a \in \mathbb{R}$  and note that  $\mathbb{E}[\Psi_{\eta}^2(\mu) \mid \eta] = \mathbb{E}[\Psi_{\eta_0}^2(\mu)]$ . The associated integral

$$I(\eta_0) = \mathbb{E}[\Psi_{\eta_0}^2(\mu)] = \int_{-\infty}^{\infty} \left( \frac{\varphi(x + \eta_0)}{\varphi(x) + \varphi(x + \eta_0)} \right)^2 \varphi(x) dx$$

needs to be evaluated numerically.

## A.2. Lopsided forecast

We proceed in analogy to the development for the unfocused forecast. For fixed  $a \in [0, 1]$  and  $b \in \mathbb{R}$ , the function

$$y \mapsto \Phi_a(y - b) = (1 - a)\Phi(y - b)\mathbb{1}\{y \leq b\} + ((1 + a)\Phi(y - b) - a)\mathbb{1}\{y > b\}$$

is a CDF. The CDF for the lopsided forecast with random density (3) from Example 2.2 can be written as  $F(y) = \Phi_{\delta}(y - \mu)$ , where  $\delta$  and  $\mu$  are independent random variables and  $\delta = \pm\delta_0$  for some  $\delta_0 \in (0, 1)$ . As  $\mathbb{E}[\Phi_{\delta}(y - \mu) \mid \mu] = \Phi(y - \mu)$ , the lopsided forecast is marginally calibrated. It fails to be probabilistically calibrated since  $Z_F = \Phi_{\delta}(Y - \mu)$  has CDF

$$\mathbb{Q}(Z_F \leq u) = \frac{1}{2} \sum_{s=\pm 1} \left( \frac{u}{1 - s\delta_0} \mathbb{1} \left\{ \frac{u}{1 - s\delta_0} \leq \frac{1}{2} \right\} + \frac{u + s\delta_0}{1 + s\delta_0} \mathbb{1} \left\{ \frac{u}{1 - s\delta_0} > \frac{1}{2} \right\} \right)$$

for  $u \in (0, 1)$  by the law of total probability.

The conditional CDF for the outcome  $Y$  given the posited (non) exceedance probability  $F(t)$  at any fixed threshold  $t \in \mathbb{R}$  or, equivalently, given the quantile forecast  $F^{-1}(\alpha)$  at any fixed level  $\alpha \in (0, 1)$  is

$$\begin{aligned} \mathbb{Q}(Y \leq y \mid F(t) = \alpha) &= \mathbb{Q}(Y \leq y \mid F^{-1}(\alpha) = t) = \mathbb{Q}(Y \leq y \mid \mu = t - \Phi_{\delta}^{-1}(\alpha)) \\ &= \frac{1}{\sum_{s=\pm 1} \varphi(t - \Phi_{s\delta_0}^{-1}(\alpha))} \sum_{s=\pm 1} \varphi(t - \Phi_{s\delta_0}^{-1}(\alpha)) \Phi(y - (t - \Phi_{s\delta_0}^{-1}(\alpha))), \end{aligned}$$

where  $\Phi_a^{-1}(\alpha) = \Phi^{-1}(\alpha/(1-a))$  if  $\alpha \leq \frac{1}{2}(1-a)$  and  $\Phi_a^{-1}(\alpha) = \Phi^{-1}((a+\alpha)/(a+1))$  otherwise.

As  $F$  is a mixture of truncated normal distributions, its moments are mixtures of the component moments, for which we refer to Orjebini (2014). The first moment is  $m_1(F) = \mu + 2\delta\varphi(0)$ , so that

$$\begin{aligned} \mathbb{Q}(Y \leq y \mid m_1(F) = m) &= \mathbb{Q}(Y \leq y \mid \mu = m - 2\delta\varphi(0)) \\ &= \frac{1}{\sum_{s=\pm 1} \varphi(m - 2s\delta_0\varphi(0))} \sum_{s=\pm 1} \varphi(m - 2s\delta_0\varphi(0)) \Phi(y - (m - 2s\delta_0\varphi(0))) \end{aligned}$$

is a mixture of normal distributions. Similarly, the second and third moments are  $m_2(F) = \mu^2 + 1 + 4\delta\varphi(0)\mu \geq 1 - 4\delta^2\varphi(0)^2$  and  $m_3(F) = \mu^3 + 3\mu + 2\delta\varphi(0)(3\mu^2 + 2) = f(\mu; \delta)$ , respectively, so that

$$\begin{aligned} \mathbb{Q}(Y \leq y \mid m_2(F) = m) &= \mathbb{Q}(Y \leq y \mid \mu = -2\delta\varphi(0) \pm \sqrt{4\delta^2\varphi(0)^2 - 1 + m}), \\ \mathbb{Q}(Y \leq y \mid m_3(F) = m) &= \mathbb{Q}(Y \leq y \mid f(\mu; \delta) = m) \end{aligned}$$

also admit expressions as mixtures of normal distributions. Again, we use a numeric solver to find the roots of  $x \mapsto f(x; \pm\delta_0)$ .

As the implied mean-forecast,  $m_1(F) = \mu + 2\delta\varphi(0)$ , agrees with the implied mean-forecast of the unfocused forecast with  $\eta = (8/\pi)^{1/2}\delta$ , the terms in the score decomposition (15) with  $S(x, y) = (x - y)^2$  derive from the respective terms in the score decomposition for the unfocused forecast, as illustrated in Figure 5.

### A.3. Piecewise uniform forecast

Given any fixed index  $i \in \{1, 2, 3\}$ , let the tuple  $(p_1^{(i)}, p_2^{(i)}, p_3^{(i)}; q_1^{(i)}, q_2^{(i)}, q_3^{(i)})$  attain the value  $(\frac{1}{2}, \frac{1}{4}, \frac{1}{4}; \frac{5}{10}, \frac{1}{10}, \frac{4}{10})$  if  $i = 1$ , the value  $(\frac{1}{4}, \frac{1}{2}, \frac{1}{4}; \frac{1}{10}, \frac{8}{10}, \frac{1}{10})$  if  $i = 2$ , and the value  $(\frac{1}{4}, \frac{1}{4}, \frac{1}{2}; \frac{4}{10}, \frac{1}{10}, \frac{5}{10})$  if  $i = 3$ . Furthermore, let  $P_i$  be the CDF of a mixture of uniform measures on  $[0, 1]$ ,  $[1, 2]$ , and  $[2, 3]$  with weights  $p_1^{(i)}$ ,  $p_2^{(i)}$ , and  $p_3^{(i)}$ , respectively. Similarly, let  $Q_i$  be the CDF of the respective mixture with weights  $q_1^{(i)}$ ,  $q_2^{(i)}$ , and  $q_3^{(i)}$ , respectively.

The random CDF for the piecewise uniform forecast in Example 2.4 can then be written as  $F(x) = P_{\iota}(x - \mu)$ , where the random variables  $\iota$  and  $\mu$  are independent, and the integer-valued variable  $\iota$  is such that

$$(p_1, p_2, p_3; q_1, q_2, q_3) = \left( p_1^{(\iota)}, p_2^{(\iota)}, p_3^{(\iota)}; q_1^{(\iota)}, q_2^{(\iota)}, q_3^{(\iota)} \right).$$

The conditional CDF for the outcome  $Y$  given the posited (non) exceedance probability  $F(t)$  at any fixed threshold  $t \in \mathbb{R}$  or, equivalently, given the quantile forecast  $F^{-1}(\alpha)$  at any fixed level  $\alpha \in (0, 1)$  then is

$$\begin{aligned} \mathbb{Q}(Y \leq y \mid F(t) = \alpha) &= \mathbb{Q}(Y \leq y \mid F^{-1}(\alpha) = t) = \mathbb{Q}(Y \leq y \mid \mu = t - P_t^{-1}(\alpha)) \\ &= \frac{1}{\sum_{i=1,2,3} \varphi\left(\frac{t - P_i^{-1}(\alpha)}{c}\right)} \sum_{i=1,2,3} \varphi\left(\frac{t - P_i^{-1}(\alpha)}{c}\right) Q_i(y - (t - P_i^{-1}(\alpha))), \end{aligned}$$

where  $c$  is the standard deviation of  $\mu$ , as defined in Example 2.4. The first moment of  $F$  is  $m_1(F) = \mu + 1 + \frac{1}{4}\iota$ , so that

$$\begin{aligned} \mathbb{Q}(Y \leq y \mid m_1(F) = m) &= \mathbb{Q}(Y \leq y \mid \mu = m - 1 - \frac{1}{4}\iota) \\ &= \frac{1}{\sum_{i=1,2,3} \varphi\left(\frac{m - 1 - \frac{1}{4}\iota}{c}\right)} \sum_{i=1,2,3} \varphi\left(\frac{m - 1 - \frac{1}{4}\iota}{c}\right) Q_i(y - (m - 1 - \frac{1}{4}\iota)) \end{aligned}$$

is a mixture of shifted versions of  $Q_1$ ,  $Q_2$ , and  $Q_3$ . The associated first moment is the respective mixture of  $m + \frac{3}{20}$ ,  $m$ , and  $m - \frac{3}{20}$ .

Given any integer  $k \geq 0$ , let  $\beta_k = \sum_{j=1,2,3} (j^{k+1} - (j-1)^{k+1}) p_j^{(\iota)}$ . The second moment of  $F$  is  $m_2(F) = \mu^2 + \beta_1\mu + \frac{1}{3}\beta_2$ , whence

$$\mathbb{Q}(Y \leq y \mid m_2(F) = m) = \mathbb{Q}\left(Y \leq y \mid \mu = -\frac{1}{2}\beta_1 \pm \sqrt{\frac{1}{4}\beta_1^2 - \frac{1}{3}\beta_2 + m}\right)$$

also admits an expression in terms of mixtures of shifted versions of  $Q_1$ ,  $Q_2$ , and  $Q_3$ . Finally, the third moment of  $F$  is  $m_3(F) = \mu^3 + \frac{3}{2}\beta_1\mu^2 + \beta_2\mu + \frac{1}{4}\beta_3 = f(\mu; \iota)$ , so that the conditional distribution  $\mathbb{Q}(Y \leq y \mid m_3(F) = m) = \mathbb{Q}(Y \leq y \mid f(\mu; \iota) = m)$  and the associated third moment can be computed analogously.

#### A.4. Identification functions, unconditional calibration, and canonical loss

In this section, we demonstrate that Definitions 2.9 and 2.24 are unambiguous and do not depend on the choice of the identification function, which is essentially unique. To this end, we first contrast the notions of identification functions in Fissler and Ziegel (2016) and Jordan, Mühlemann and Ziegel (2022). Fissler and Ziegel (2016) call  $V: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$  a (strict  $\mathcal{F}$ -)identification function if  $V(x, \cdot)$  is integrable with respect to all  $F \in \mathcal{F}$  for all  $x \in \mathbb{R}$ . Jordan, Mühlemann and Ziegel (2022) additionally require  $V$  to be increasing and left-continuous in its first argument. Furthermore, there is a subtle difference in the way that the functional is induced. While Fissler and Ziegel (2016) define the induced functional as the set

$$\mathbb{T}_0(F) = \left\{ x \in \mathbb{R} : \int V(x, y) dF(y) = 0 \right\},$$

Jordan, Mühlemann and Ziegel (2022) define it to be the closed interval  $T(F) = [T^-(F), T^+(F)]$ , where  $T^-(F)$  and  $T^+(F)$  are defined in (7) and (8), respectively. The approach by Jordan, Mühlemann and Ziegel (2022) allows for quantiles to be treated in full generality and ensures that the interval  $T(F)$  coincides with the closure of  $T_0(F)$  if the latter is nonempty.

In the setting of Fissler and Ziegel (2016), if  $V$  is an identification function, then so is  $(x, y) \mapsto h(x)V(x, y)$  whenever  $h(x) \neq 0$  for all  $x \in \mathbb{R}$ . If the class  $\mathcal{F}$  is sufficiently rich, then any two locally bounded identification functions  $V$  and  $\tilde{V}$  that induce a functional  $T_0$  of singleton type relate to each other in the stated form almost everywhere on the interior of  $T_0(\mathcal{F}) \times \mathbb{R}$  (Dimitriadis, Fissler and Ziegel, 2023, Theorem 4), which implies that increasing identification functions of prediction error form are unique up to a positive constant. The following proposition provides an elementary proof under slightly different conditions that are tailored to our setting. Notably, identification functions of prediction error form induce functionals that are equivariant under translation by Proposition 4.7 of Fissler and Ziegel (2019), a result which can easily be transferred to the setting of Jordan, Mühlemann and Ziegel (2022).

**Proposition A.1.** *Let  $\mathcal{F}$  be a convex class of probability measures such that  $\delta_y \in \mathcal{F}$  for all  $y \in \mathbb{R}$ . If the functional  $T$  is induced on  $\mathcal{F}$  by an identification function  $V(x, y) = v(x - y)$  of prediction error form, where  $v$  is increasing and left-continuous with  $v(-r) < 0$  and  $v(r) > 0$  for some  $r > 0$ , then any other identification function of the stated form that induces  $T$  is a positive multiple of  $V$ .*

*Proof.* Let  $V: (x, y) \mapsto v(x - y)$  and  $\tilde{V}: (x, y) \mapsto \tilde{v}(x - y)$  induce the functionals  $T$  and  $\tilde{T}$ , respectively. We proceed to show that  $T = \tilde{T}$  implies  $\tilde{v} = h_0 \cdot v$  for some constant  $h_0 > 0$ .

To this end, suppose that  $T = \tilde{T}$ , and let

$$\begin{aligned} r^- &= \sup\{r : \tilde{v}(r) < 0\} = \tilde{T}^-(\delta_0) = T^-(\delta_0) = \sup\{r : v(r) < 0\} > -\infty, \\ r^+ &= \inf\{r : \tilde{v}(r) > 0\} = \tilde{T}^+(\delta_0) = T^+(\delta_0) = \inf\{r : v(r) > 0\} < \infty. \end{aligned}$$

By left-continuity and monotonicity of  $v$  and  $\tilde{v}$ , it follows that  $v(r) = \tilde{v}(r) = 0$  for  $r \in (r^-, r^+]$ ,  $v(r) < 0$  and  $\tilde{v}(r) < 0$  for  $r < r^-$ , and  $v(r) > 0$  and  $\tilde{v}(r) > 0$  for  $r > r^+$ .

Let  $h(r) = \tilde{v}(r)/v(r) > 0$  for  $r \in \mathbb{R} \setminus [r^-, r^+]$ . If  $r < r^- \leq r^+ < s$ , then  $\tilde{v}(r) = h(r)v(r) < 0$  and  $\tilde{v}(s) = h(s)v(s) > 0$ . Assume  $h(r) < h(s)$ , and let  $p \in (0, 1)$  be such that

$$\left(1 - \frac{h(s)v(s)}{h(r)v(r)}\right)^{-1} < p < \left(1 - \frac{v(s)}{v(r)}\right)^{-1}.$$

Then  $(1 - p)h(r)v(r) + ph(s)v(s) > 0 > (1 - p)v(r) + pv(s)$  and  $\tilde{T}^+(p\delta_{-s} + (1 - p)\delta_{-r}) < 0 \leq T^-(p\delta_{-s} + (1 - p)\delta_{-r})$ , a contradiction. An analogous argument applies if we assume that  $h(r) > h(s)$ , and we conclude that  $h(r) = h(s)$ .

If  $r, s < r^-$ , then  $h(r) = h(s) = h(t)$  for any  $t > r^+$  by the above line of reasoning. An analogous argument yields  $h(r) = h(s)$  for  $r, s > r^+$ . Therefore, the function  $h$  is constant and  $v(r) = h_0 \cdot \tilde{v}(r)$  for a constant  $h_0 > 0$  and all  $r \in \mathbb{R} \setminus \{r^-\}$ . Finally, we obtain  $v(r^-) = \lim_{r \uparrow r^-} v(r) = \lim_{r \uparrow r^-} h_0 \cdot \tilde{v}(r) = h_0 \cdot \tilde{v}(r^-)$  by left-continuity.  $\square$

Hence, if we assume an identification function of type (i) in Assumption 2.8, Definitions 2.9 and 2.24 do not depend on the choice of the identification function, as it is unique up to a positive constant. Trivially, the same holds true for type (ii). To complete the argument that the definitions are unambiguous, the following technical argument is needed.

**Remark A.2.** If a functional  $T$  of singleton type is identified by both an identification function  $V(x, y) = v(x - y)$  of type (i) and an identification function  $\tilde{V}(x, y) = x - T(\delta_y)$  of type (ii), then  $\tilde{V}$  is also of type (i). To confirm this claim, let  $z$  denote the unique value at which the sign of  $v$  changes, and note that  $z = T(\delta_y) - y$  for all  $y$  since  $V$  induces the functional  $T$  for each Dirac measure  $\delta_y$ . Hence,  $T(\delta_y) = y + z$  and  $\tilde{V}(x, y) = x - y - z$  is of type (i).

We close this section with comments on the role of the class  $\mathcal{F}$ . As expressed by Assumption 2.8, we prefer to work with identification functions that elicit the target functional  $T$  on a large, convex class  $\mathcal{F}$  of probability measures, to avoid unnecessary constraints on forecast(er)s. Furthermore, when evaluating stand-alone point forecasts, the underlying predictive distributions typically are implicit, and assumptions other than the existence of the functional at hand are unwarranted and contradict the prequential principle. Evidently, if the class  $\mathcal{F}$  is sufficiently restricted, additional identification functions arise. For example, the piecewise constant identification function associated with the median can be used to identify the mean within any class of symmetric distributions.

### A.5. Probabilistic calibration and unconditional quantile calibration

As noted in the main text, probabilistic calibration implies the unconditional  $\alpha$ -quantile calibration condition (10) at every level  $\alpha \in (0, 1)$ . To verify this implication, it suffices to note that if  $F$  is probabilistically calibrated, then  $\alpha = \mathbb{Q}(Z_F \leq \alpha) \leq \mathbb{Q}(F(Y^-) \leq \alpha) = \mathbb{Q}(Y \leq q_\alpha^-(F))$  and  $1 - \alpha = \mathbb{Q}(Z_F > \alpha) \leq \mathbb{Q}(F(Y) > \alpha) = \mathbb{Q}(Y \geq q_\alpha^+(F))$ . As Example 2.14(b) demonstrates, the reverse implication does not hold in general. However, Assumption 2.15 ensures the equivalence of probabilistic calibration and unconditional  $\alpha$ -quantile calibration at every level  $\alpha \in (0, 1)$ .

### A.6. Counterexample (Proposition 2.17(a))

As pointed out by Sahoo et al. (2021, p. 5), strong threshold calibration does not imply auto-calibration. Here, we provide a simple example illustrating this fact as Sahoo et al. (2021) do not present such. The example is similar in spirit to the

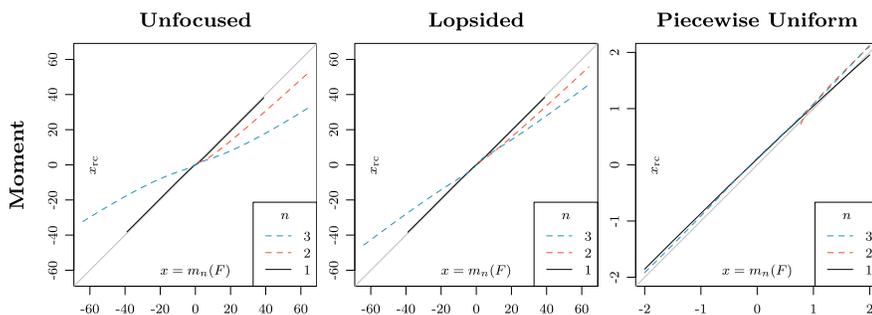


FIG 10. Same as the lower row of Figure 4 but with displays on original (rather than root-transformed) scales: Moment reliability diagrams for point forecasts induced by (left) the unfocused forecast with  $\eta_0 = 1.5$  and (middle) the lopsided forecast with  $\delta_0 = 0.7$  from Example 2.2, and (right) the piecewise uniform forecast with  $c = 0.5$  from Example 2.4.

continuous forecast of Example 2.14(a) (as  $c \rightarrow 0$ ) but with strictly increasing distribution functions satisfying Assumption 2.15.

Let  $F$  be a mixture of uniform distributions on the intervals  $[0, 1]$ ,  $[1, 2]$ ,  $[2, 3]$ , and  $[3, 4]$  with weights  $p_1, p_2, p_3$ , and  $p_4$ , respectively, and let  $Y$  be from a mixture with weights  $q_1, q_2, q_3$ , and  $q_4$ . Furthermore, let the tuple  $(p_1, p_2, p_3, p_4; q_1, q_2, q_3, q_4)$  attain each of the values

$$\left(\frac{4}{10}, \frac{1}{10}, \frac{4}{10}, \frac{1}{10}; \frac{16}{25}, \frac{4}{25}, \frac{4}{25}, \frac{1}{25}\right), \quad \left(\frac{1}{10}, \frac{4}{10}, \frac{1}{10}, \frac{4}{10}; \frac{4}{25}, \frac{16}{25}, \frac{1}{25}, \frac{4}{25}\right),$$

$$\left(\frac{4}{10}, \frac{1}{10}, \frac{1}{10}, \frac{4}{10}; \frac{4}{25}, \frac{1}{25}, \frac{4}{25}, \frac{16}{25}\right), \quad \left(\frac{1}{10}, \frac{4}{10}, \frac{4}{10}, \frac{1}{10}; \frac{1}{25}, \frac{4}{25}, \frac{16}{25}, \frac{4}{25}\right)$$

with equal probability. The equal average of the distribution of the PIT conditional on either forecast from the top row, and either forecast from the bottom row, is uniform. As any nontrivial conditioning in terms of a threshold yields a combination of two forecast cases, one from the top row and one from the bottom row, the forecast  $F$  is strongly threshold calibrated.

### A.7. Remarks on Figure 4

The root transforms in the moment reliability diagrams in the bottom row of Figure 4 bring the first, second, and third moment to the same scale. The peculiar dent in the reliability curve for the (third root of the) third moment of the piecewise uniform forecast results from the transform, which magnifies small deviations between  $x = m_3(F)$  and  $x_{rc}$  when  $x$  is close to zero. For comparison, Figure 10 shows moment reliability diagrams for all three forecasts without applying the root transform.

### A.8. Counterexample (Theorem 2.26)

The statement in Theorem 2.26 does not hold under consistent scoring functions in general. For a counterexample, consider the empirical distribution of

$(x_1, y_1), \dots, (x_n, y_n)$ , where  $x_i = i$  and  $y_i = x_i + \frac{10}{9}$  for  $i = 1, \dots, 9$ , and  $x_{10} = y_{10} = -10$ . The respective mean-forecast  $X$  fails to be unconditionally mean calibrated, whereas the shifted version  $X + 1$  is unconditionally mean calibrated. Nonetheless, the expected elementary score (16) for the mean functional (i.e.,  $V(x, y) = x - y$ ) with index  $\eta = -\frac{19}{2}$  increases when  $X$  gets replaced with  $X + 1$ .

## Appendix B: Consistency resamples and calibration tests

Monte Carlo based consistency bands for T-reliability diagrams can be generated from resamples, at any desired nominal level. The consistency bands then show the pointwise range of the resampled calibration curves. For now, let us assume that we have data  $(x_1, y_1), \dots, (x_n, y_n)$  of the form (24) along with  $m$  resamples at hand, and defer the critical question of how to generate the resamples.

---

**Algorithm 2:** Consistency bands for T-reliability curves based on resamples

---

**Input:** resamples  $(x_1, \tilde{y}_1^{(j)}), \dots, (x_n, \tilde{y}_n^{(j)})$  for  $j = 1, \dots, m$

**Output:**  $\alpha \times 100\%$  consistency band

**for**  $j \in \{1, \dots, m\}$  **do**

| apply Algorithm 1 to obtain  $\hat{x}_1^{(j)}, \dots, \hat{x}_n^{(j)}$  from  
    $(x_1, \tilde{y}_1^{(j)}), \dots, (x_n, \tilde{y}_n^{(j)})$

**end**

**for**  $i \in \{1, \dots, n\}$  **do**

| let  $l_i$  and  $u_i$  be the empirical quantiles of  $\hat{x}_i^{(1)}, \dots, \hat{x}_i^{(m)}$  at level  $\frac{\alpha}{2}$   
   and  $1 - \frac{\alpha}{2}$

**end**

interpolate the point sets  $(x_1, l_1), \dots, (x_n, l_n)$  and  $(x_1, u_1), \dots, (x_n, u_n)$  linearly, to obtain the lower and upper bound of the consistency band, respectively

---

Complementary to consistency bands, tests for the assumed type of calibration, as quantified by the functional T and a generic miscalibration measure MCB, can be performed as usual. Specifically, we compute  $\text{MCB}_j$  for each resample  $j = 1, \dots, m$ , and, if  $r$  of the resampled measures  $\text{MCB}_1, \dots, \text{MCB}_m$  are less than or equal to the miscalibration measure computed from the original data, we declare a Monte Carlo  $p$ -value of  $1 - \frac{r}{m+1}$ .

### B.1. Consistency resamples under the hypothesis of auto-calibration

When working with original data of the form (23), we can generate resamples under the hypothesis of auto-calibration in the obvious way, as follows.

---

**Algorithm 3:** Consistency resamples under the hypothesis of auto-calibration

---

**Input:**  $(F_1, y_1), \dots, (F_n, y_n)$   
**Output:** resamples  $(x_1, \tilde{y}_1^{(j)}), \dots, (x_n, \tilde{y}_n^{(j)})$  for  $j = 1, \dots, m$   
**for**  $i \in \{1, \dots, n\}$  **do**  
  | let  $x_i = T(F_i)$   
**end**  
**for**  $j \in \{1, \dots, m\}$  **do**  
  | **for**  $i \in \{1, \dots, n\}$  **do**  
    | sample  $\tilde{y}_i^{(j)}$  from  $F_i$   
  | **end**  
**end**

---

As noted, in the case of threshold calibration, the induced outcome is binary, whence the assumptions of auto-calibration and T-calibration coincide. For other types of functionals, auto-calibration is a strictly stronger assumption than T-calibration, and it is important to note that the resulting inferential procedures may be confounded by forecast attributes other than T-calibration. For illustration, let us return to the setting of Example 2.1 and suppose that, conditionally on a standard normal variate  $\mu$ , the outcome  $Y$  is normal with mean  $\mu$  and variance 1. Given any fixed  $\sigma > 0$ , the forecast  $F_\sigma = \mathcal{N}(\mu, \sigma^2)$  is auto-calibrated if, and only if,  $\sigma = 1$ . However, if  $T$  is the mean or median functional, then  $F_\sigma$  is T-calibrated under any  $\sigma > 0$ . Clearly, if we use Algorithm 3 to generate resamples, then the consistency bands generated by Algorithm 2 might be misleading with regard to the assessment of T-calibration. For example, if  $\sigma < 1$  the confidence bands tend to be narrow and might erroneously suggest a lack of T-calibration, despite the forecast being T-calibrated.

### B.2. Consistency resamples under the hypothesis of T-calibration

The issues just described call for an alternative to Algorithm 3. Residual-based approaches can be used to generate resamples under the weaker hypothesis of T-calibration. In developing such a method, we restrict the discussion to single-valued functionals  $T$  under which  $y_i = T(\delta_i)$ , which covers all cases of key interest, such as the mean functional, lower or upper quantiles, and expectiles. As is standard in regression diagnostics, residual-based approaches operate on the basis of tuples  $(x_1, y_1), \dots, (x_n, y_n)$  of the form (24) under the assumptions of independence between the point forecast,  $x_i$ , and the residual,  $y_i - x_i$ , and exchangeability of the residuals. For a discussion in the context of backtests in banking regulation, see Example 3 of Nolde and Ziegel (2017).

Interestingly, Theorem 2.21 demonstrates that under these assumptions a forecast is conditionally T-calibrated if, and only if, it is unconditionally T-calibrated. Thus, we draw resamples in a two-stage procedure. First, we find the constant  $c$  from Theorem 2.27 such that the empirical distribution of  $(x_1 + c, y_1), \dots, (x_n + c, y_n)$  or, equivalently,  $(x_1, y_1 - c), \dots, (x_n, y_n - c)$ , is uncon-

ditionally T-calibrated, and then we resample from the respective residuals, as follows.

---

**Algorithm 4:** Consistency resamples under the joint hypothesis of T-calibration and independence between point forecasts and residuals

---

**Input:**  $(x_1, y_1), \dots, (x_n, y_n)$   
**Output:** resamples  $(x_1, \tilde{y}_1^{(j)}), \dots, (x_n, \tilde{y}_n^{(j)})$  for  $j = 1, \dots, m$   
**for**  $i \in \{1, \dots, n\}$  **do**  
  | let  $r_i = y_i - x_i$   
**end**  
find  $c$  such that  $(x_1 + c, y_1), \dots, (x_n + c, y_n)$  is unconditionally T-calibrated  
**for**  $j \in \{1, \dots, m\}$  **do**  
  | sample  $\tilde{r}_1, \dots, \tilde{r}_n$  from  $\{r_1, \dots, r_n\}$  with replacement  
    **for**  $i \in \{1, \dots, n\}$  **do**  
      | let  $\tilde{y}_i = x_i + \tilde{r}_i - c$   
    **end**  
**end**

---

As noted in the main text, the consistency bands for the threshold reliability diagrams in Figures 6 and 8 have been generated by Algorithms 2 and 4. This approach is similar to the Monte Carlo technique proposed by Dimitriadis, Gneiting and Jordan (2021) that applies in the case of (induced) binary outcomes (only). However, unlike Dimitriadis, Gneiting and Jordan (2021), we do not resample the forecasts themselves. To generate consistency bands for the mean and quantile reliability diagrams in these figures, we apply Algorithm 2 to  $m = 1000$  resamples generated by Algorithm 4. Evidently, this procedure is crude and relies on classical assumptions. Nonetheless, we believe that in many practical settings, where visual tools for diagnostic checks of calibration are sought, the consistency bands thus generated provide useful guidance.

Further methodological development on consistency and confidence bands needs to be tailored to the specific functional T of interest, and follow-up work on Monte Carlo techniques and large sample theory is strongly encouraged. Extant asymptotic theory for nonparametric isotonic regression, as implemented by Algorithm 1, is available for quantiles and the mean or expectation functional, as developed and reviewed by Barlow et al. (1972), Casady and Cryer (1976), Wright (1984), Robertson, Wright and Dykstra (1988), El Barmi and Mukerjee (2005), and Mösching and Dümbgen (2020), and can be leveraged, though with hurdles, as rates of convergence depend on distributional assumptions and limit distributions involve nuisance parameters that need to be estimated, whereas the use of bootstrap methods might be impacted by the issues described by Sen, Banerjee and Woodroffe (2010).

### B.3. Reliability diagrams and consistency bands for probabilistic and marginal calibration

For the classical notions of unconditional calibration in Section 2.2, the CORP approach does not apply directly, but its spirit can be retained and adapted.

As for probabilistic calibration, the prevalent practice is to plot histograms of empirical probability integral transform (PIT) values, as proposed by Diebold, Gunther and Tay (1998), Gneiting, Balabdaoui and Raftery (2007), and Czado, Gneiting and Held (2009), though this practice is hindered by the necessity for binning, as analyzed by Heinrich (2021) in the nearly equivalent setting of rank histograms. The population version of our suggested alternative is the *PIT reliability diagram*, which is simply the graph of the CDF of the PIT  $Z_F$  in (1). The PIT reliability diagram coincides with the diagonal in the unit square if, and only if,  $F$  is probabilistically calibrated. For tuples of the form (23) the *empirical PIT reliability diagram* shows the empirical CDF of the (potentially randomized) PIT values. This approach does not require binning and can be interpreted in much the same way as a PIT diagram: An inverse S-shape corresponds to a U-shape in histograms and indicates underdispersion of the forecast, as typically encountered in practice. Evidently, this idea is not new and extant implementations can be found in work by Pinson and Hagedorn (2012) and Henzi, Ziegel and Gneiting (2021).

As regards marginal calibration, we define the population version of the *marginal reliability diagram* as the point set

$$\{(\mathbb{E}_{\mathbb{Q}}[F(y)], \mathbb{Q}(Y \leq y)) \in [0, 1]^2 : y \in \mathbb{R}\}.$$

The marginal reliability diagram is concentrated on the diagonal in the unit square if, and only if,  $F$  is marginally calibrated. For tuples of the form (23) the *empirical marginal reliability diagram* is a plot of the empirical non-exceedance probability (NEP)  $\widehat{F}_0(y) = \frac{1}{n} \sum_{i=1}^n \mathbb{1}\{y \geq y_i\}$  against the average forecast NEP  $\bar{F}(y) = \frac{1}{n} \sum_{i=1}^n F_i(y)$  at the unique values  $y$  of the outcomes  $y_1, \dots, y_n$ , and interpolated linearly in between. Of course, this idea is not new either and the resulting diagram can be interpreted as a P-P plot.

For marginal calibration diagrams, we obtain consistency bands under the assumption of marginal calibration by drawing resamples  $y_1^{(j)}, \dots, y_n^{(j)}$  from  $\bar{F} = \frac{1}{n} \sum_{i=1}^n F_i$ , computing the respective marginal reliability curve, and repeating over Monte Carlo replicates  $j = 1, \dots, m$ . Then we find consistency bands in the spirit of Algorithm 2. For PIT reliability diagrams, a trivial technique applies as we may obtain consistency bands under the assumption of probabilistic calibration by (re)sampling  $n$  independent standard uniform variates, computing the respective empirical CDF, and repeating over Monte Carlo replicates. Evidently, there are alternatives based on empirical process theory (Shorack and Wellner, 2009).

Figure 11 illustrates PIT and marginal reliability diagrams on our customary examples, along with 90% consistency bands based on  $m = 1000$  Monte Carlo replicates.

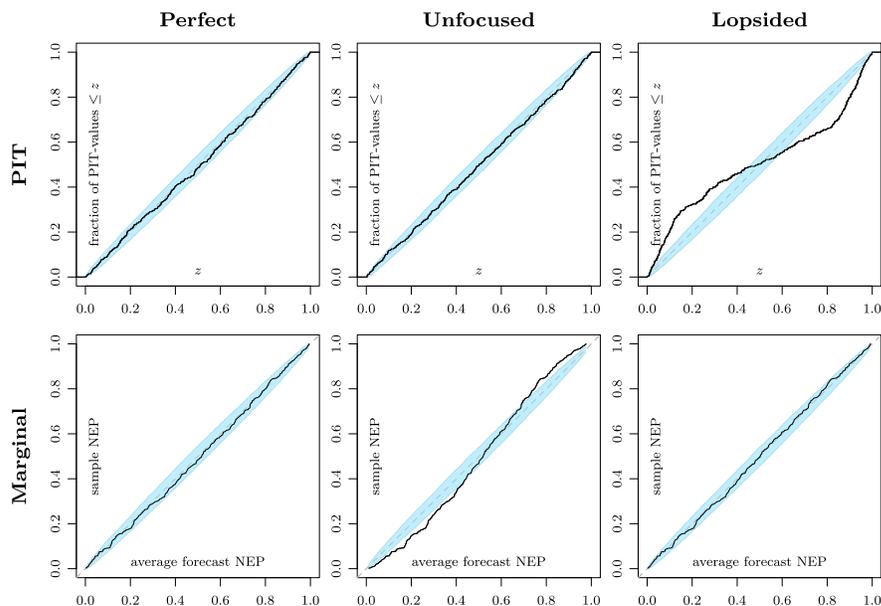


FIG 11. PIT (top) and marginal (bottom) reliability diagrams for the perfect (left), unfocused (middle), and lopsided (right) forecast from Examples 2.1 and 2.2, along with 90% consistency bands based on samples of size 400.

#### B.4. Testing hypotheses of calibration

While the explicit development of calibration tests exceeds the scope of our paper, we believe that the results and discussion in Section 2 convey an important general message: It is critical that the assessed notion of calibration be carefully and explicitly specified. Throughout, we consider tests under the assumption of independent, identically distributed data from a population. For extensions to dependent samples, we refer to Strähl and Ziegel (2017), who generalized the prediction space concept to allow for serial dependence, and point at methods introduced by, e.g., Corradi and Swanson (2007), Knüppel (2015), and Bröcker and Ben Bouallègue (2020).

The most basic case is that of tuples  $(x_1, y_1), \dots, (x_n, y_n)$  of the form (24), where implicitly or explicitly  $x_i = T(F_i)$  for a single-valued functional  $T$ . We first discuss tests of unconditional calibration. If the simplified condition (11) is sufficient, a two-sided  $t$ -test based on  $\hat{v} = \frac{1}{n} \sum_{i=1}^n V(x_i, y_i)$  can be used to test for unconditional calibration. In the general case, two one-sided  $t$ -tests can be used along with a Bonferroni correction. In the special case of quantiles, there is no need to resort to the approximate  $t$ -tests, and exact binomial tests can be used instead. Essentially, this special case is the setting of backtests for value-at-risk reports in banking regulation, for which we refer to Nolde and Ziegel (2017, Sections 2.1–2.2).

As noted earlier in the section, resamples generated under the hypothesis of

conditional T-calibration can readily be used to perform Monte Carlo tests for the respective hypothesis, based on CORP score components that are computed on the surrogate data. Alternatively, one might leverage extant large sample theory for nonparametric isotonic regression (Barlow et al., 1972; Casady and Cryer, 1976; Wright, 1984; Robertson, Wright and Dykstra, 1988; El Barmi and Mukerjee, 2005; Mösching and Dümbgen, 2020). Independently of the use of resampling or asymptotic theory, CORP based tests avoid the issues and instabilities incurred by binning (Dimitriadis, Gneiting and Jordan, 2021, Section S2) and may simultaneously improve efficiency and stability. In passing, we hint at relations to the null hypothesis of Mincer-Zarnowitz regression (Krüger and Ziegel, 2021) and tests of predictive content (Galbraith, 2003; Breitung and Knüppel, 2021).

We move on to the case of fully specified distributions, where we work with tuples  $(F_1, y_1), \dots, (F_n, y_n)$  of the form (23), where  $F_i$  is a posited conditional CDF for  $y_i$  ( $i = 1, \dots, n$ ). Tests for probabilistic calibration then amount to tests for the uniformity of the (potentially, randomized) PIT values. Wallis (2003) and Wilks (2019, p. 769) suggest chi-square tests for this purpose, which depend on binning, and thus are subject to the aforementioned instabilities. To avoid binning, we recommend the use of test statistics that operate on the empirical CDF of the PIT values, such as the classical Kolmogorov–Smirnov statistic, as suggested and used to test for PIT calibration by Noceti, Smith and Hodges (2003) and Knüppel (2015), or, more generally, tests based on distance measures between the empirical CDF of the PIT values, and the CDF of the standard uniform distribution that arises under the hypothesis of probabilistic calibration. Recently proposed alternatives arise via  $e$ -values (Henzi and Ziegel, 2022). Similarly, tests for marginal calibration can be based on resamples and distance measures between  $\bar{F}$  and  $\hat{F}_0$ , or leverage asymptotic theory.

In the distributional setting, arbitrarily many types of reliability can be tested for, and all of the aforementioned tests for unconditional or conditional T-calibration apply. Multiple testing needs to be accounted for properly, and the development of simultaneous tests for various types of calibration would be useful. In this context, let us recall from Theorem 2.16 that, subject to technical conditions, CEP, threshold, and quantile calibration are equivalent and tests for CEP calibration (Held, Rufibach and Balabdaoui, 2010; Strähl and Ziegel, 2017), quantile and threshold calibration assess identical hypotheses.

### Appendix C: Time series settings and the Bank of England example

In typical time series settings, as exemplified by our analysis of Bank of England forecasts in Section 3, the assumption of independent replicates of forecasts and observations is too restrictive. While the diagnostic methods proposed in our paper continue to apply, statistical inference requires care, as discussed by Corradi and Swanson (2007) and Knüppel (2015), among other authors. Here, we elucidate the role of uniform and independent probability integral transform (PIT) values for calibration in time series settings, and give further details and results for the Bank of England example.

### C.1. The role of uniform and independent PITs

In a landmark paper, Diebold, Gunther and Tay (1998, p. 867) showed that a sequence of continuous predictive distributions  $F_t$  for a sequence  $Y_t$  of observations at time  $t = 0, 1, \dots$  results in a sequence of independent, uniformly distributed PITs if  $F_t$  is ideal relative to the  $\sigma$ -algebra generated by past observations,  $\mathcal{A}_t = \sigma(Y_0, Y_1, \dots, Y_{t-1})$ . This property does not depend on the continuity of  $F_t$  and continues to hold under general predictive CDFs and the randomized definition (1) of the PIT (Rüschendorf and de Valk, 1993, Theorem 3).

In the case of continuous predictive distributions, Tsyplakov (2011, Section 2) noted without proof that if the forecasts  $F_t$  are based only on past observations, i.e., if  $F_t$  is  $\mathcal{A}_t$ -measurable, then the converse holds, namely, uniform and independent PITs arise only if  $F_t$  is ideal relative to  $\mathcal{A}_t$ . The following result formalizes Tsyplakov's claim and proves it in the general setting, without any assumption of continuity.

**Theorem C.1.** *Let  $(Y_t)_{t=0,1,\dots}$  be a sequence of random variables, and let  $\mathcal{A}_t = \sigma(Y_0, \dots, Y_{t-1})$  for  $t = 0, 1, \dots$ . Furthermore, let  $(F_t)_{t=0,1,\dots}$  be a sequence of CDFs, such that  $F_t$  is  $\mathcal{A}_t$ -measurable for  $t = 0, 1, \dots$ , and let  $(U_t)_{t=0,1,\dots}$  be a sequence of independent, uniformly distributed random variables, independent of the sequence  $(Y_t)$ . Then the sequence of randomized PITs,  $(Z_t) = (F_t(Y_t-) + U_t(F_t(Y_t) - F_t(Y_t-)))$  is an independent sequence of uniform random variables on the unit interval if, and only if,  $F_t$  is ideal relative to  $\mathcal{A}_t$ , i.e.,  $F_t = \mathcal{L}(Y_t | \mathcal{A}_t)$  almost surely for  $t = 0, 1, \dots$ .*

The proof utilizes the following simple lemma.

**Lemma C.2.** *Let  $X, Y, Z$  be random variables. If  $X = Z$  almost surely, then  $\mathbb{E}[Y | X] = \mathbb{E}[Y | Z]$  almost surely.*

*Proof.* Problem 14 of Breiman (1992, Chapter 4), which is proved by Schmidt (2011, Satz 18.2.10), states that for random variables  $X_1$  and  $X_2$  such that  $\sigma(Y, X_1)$  is independent of  $\sigma(X_2)$ ,  $\mathbb{E}[Y | X_1, X_2] = \mathbb{E}[Y | X_1]$  almost surely. The statement of the lemma follows as  $\mathbb{E}[Y | X] = \mathbb{E}[Y | X, X - Z] = \mathbb{E}[Y | Z, X - Z] = \mathbb{E}[Y | Z]$  almost surely.  $\square$

*Proof of Theorem C.1.* Since  $F_t$  is measurable with respect to  $\mathcal{A}_t$ , there exists a measurable function  $f_t: \mathbb{R}^t \rightarrow \mathcal{F}$  such that  $F_t = f_t(Y_0, \dots, Y_{t-1})$  for each  $t$  by the Doob–Dynkin Lemma (Schmidt, 2011, Satz 7.1.16).<sup>2</sup> We define

$$G_t := f_t(G_0^{-1}(Z_0), \dots, G_{t-1}^{-1}(Z_{t-1}))$$

<sup>2</sup>Note that  $f_0$  is constant, and  $f_t$  is not a random quantity but a fixed function that encodes how the predictive distributions are generated from past observations. The  $\sigma$ -algebra on  $\mathcal{F}$ , which is implicitly used throughout, is given by

$$\mathcal{A}_{\mathcal{F}} = \sigma(\{F \in \mathcal{F} : F(x) \in B\} : x \in \mathbb{Q}, B \in \mathcal{B}(\mathbb{R})\}),$$

where  $\mathcal{B}(\mathbb{R})$  denotes the Borel  $\sigma$ -algebra on  $\mathbb{R}$ . For each  $x \in \mathbb{Q}$  there exists a measurable function  $f_{x,t}$  such that  $F_t(x) = f_{x,t}(Y_0, \dots, Y_{t-1})$  by the Doob–Dynkin Lemma, and  $f_t$  is essentially the countable (and hence measurable) collection  $(f_{x,t})_{x \in \mathbb{Q}}$ .

recursively for all  $t$ , and show the “only if” assertion by induction.

To this end, let  $t \geq 0$  and assume the induction hypothesis that  $F_i$  is ideal relative to  $\mathcal{A}_i$  for  $i = 0, \dots, t-1$ . By Rüschenendorf and de Valk (1993, Theorem 3(a)) and the construction of  $G_t$ , the induction hypothesis implies

$$(Y_0, \dots, Y_{t-1}) = (F_0^{-1}(Z_0), \dots, F_{t-1}^{-1}(Z_{t-1})) = (G_0^{-1}(Z_0), \dots, G_{t-1}^{-1}(Z_{t-1}))$$

almost surely, where the last vector is  $\sigma(Z_0, \dots, Z_{t-1})$ -measurable. By Lemma C.2, it follows that

$$\mathcal{L}(Z_t | \mathcal{A}_t) = \mathcal{L}(Z_t | \sigma(G_0^{-1}(Z_0), \dots, G_{t-1}^{-1}(Z_{t-1}))) = U([0, 1])$$

almost surely, where the second equality stems from the fact that  $Z_t$  is independent of  $\sigma(G_0^{-1}(Z_0), \dots, G_{t-1}^{-1}(Z_{t-1})) \subset \sigma(Z_0, \dots, Z_{t-1})$ . This independence implies that  $F_t$  is ideal relative to  $\mathcal{A}_t$  because

$$F_t(y) = \mathbb{Q}(Z_t < F_t(y) | \mathcal{A}_t) \leq \mathbb{Q}(Y_t \leq y | \mathcal{A}_t) \leq \mathbb{Q}(Z_t \leq F_t(y) | \mathcal{A}_t) = F_t(y)$$

almost surely, and hence  $F_t(y) = \mathbb{Q}(Y_t \leq y | \mathcal{A}_t)$  almost surely for all  $y \in \mathbb{Q}$ , thereby completing both the induction step and the claim for the base case  $t = 0$ .  $\square$

Evidently, the assumption that no information other than the history of the time series itself has been utilized to construct the forecasts is very limiting. In this light, it is not surprising that, while the “if” part of Theorem C.1 is robust, the “only if” claim fails if  $F_t$  is allowed to use information beyond the canonical filtration, even if that information is uninformative. A simple counterexample is given by the unfocused forecast from Example 2.2, which is probabilistically calibrated but fails to be auto-calibrated. Its PIT nevertheless is uniform and independent even for autoregressive variants (Tsyplakov, 2011, Section 6).

## C.2. Details and further results for the Bank of England example

Bank of England forecasts of inflation rates are available within the data accompanying the quarterly Monetary Policy Report (formerly Inflation Report), which is available online at <https://www.bankofengland.co.uk/sitemap/monetary-policy-report>. The forecasts are visualized and communicated in the form of *fan charts* that span prediction intervals at increasing forecast horizons, and derive from two-piece normal forecast distributions. A detailed account of the parametrizations for the two-piece normal distribution used by the Bank of England can be found in Julio (2006), and we have implemented the formulas in this reference. Historical quarterly CPI inflation rates are published by the UK Office for National Statistics (ONS) and available online at <https://www.ons.gov.uk/economy/inflationandpriceindices/timeseries/d7g7>.

We consider forecasts of consumer price index (CPI) inflation based on market expectations for future interest rates at prediction horizons of zero to six quarters ahead, valid for the third quarter of 2005 up to the first quarter of 2020, for a total of  $n = 59$  quarters. These and earlier Bank of England forecasts of inflation rates have been checked for reliability by Wallis (2003), who considered

probabilistic calibration, by Clements (2004) in terms of probabilistic, mean, and threshold calibration, by Galbraith and van Norden (2012), who considered probabilistic and mean calibration, by Strähl and Ziegel (2017) with focus on conditional exceedance probability (CEP) calibration, and by Pohle (2020), who considered quantile calibration. The 2% inflation target is discussed on the Bank of England website at <https://www.bankofengland.co.uk/monetary-policy/inflation>.

Figures 12–17 show calibration diagnostics for inflation forecasts at prediction horizons of  $k \in \{0, 2, 3, 4, 5, 6\}$  quarters ahead, in the same format as Figure 8 in the main text, which concerns forecasts at a lead time of one quarter.

### Acknowledgments

The authors would like to thank Sebastian Arnold, Fadoua Balabdaoui-Mohr, Jonas Brehmer, Frank Diebold, Timo Dimitriadis, Uwe Ehret, Andreas Fink, Tobias Fissler, Rafael Frongillo, Norbert Henze, Alexander Henzi, Alexander I. Jordan, Kristof Kraus, Fabian Krüger, Sebastian Lerch, Michael Maier-Gerber, Anja Mühlemann, Jim Pitman, Marc-Oliver Pohle, Roopesh Ranjan, Benedikt Schulz, Ville Satopää, Daniel Wolfram and Johanna F. Ziegel, as well as anonymous reviewers, for helpful comments and discussion.

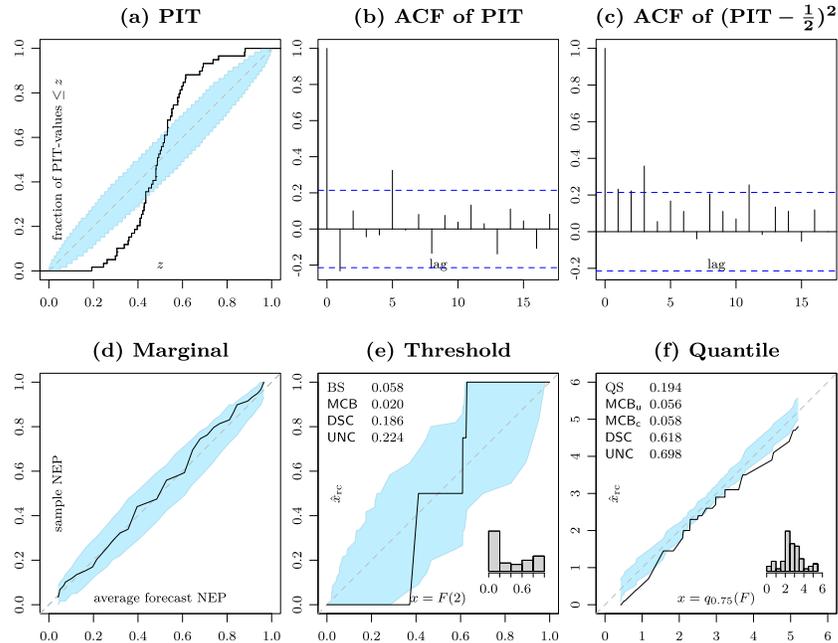


FIG 12. Same as Figure 8 in the main text but at a prediction horizon of zero quarters.

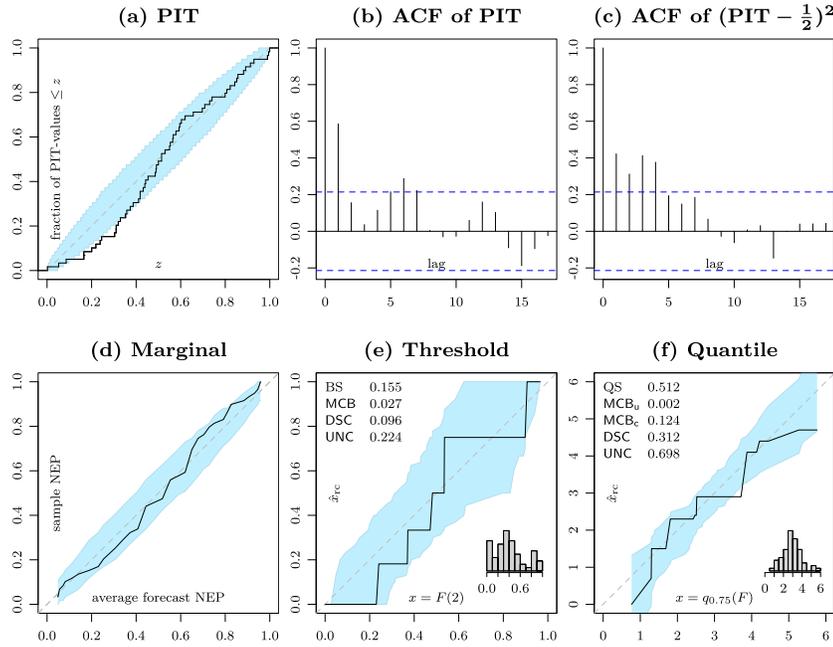


FIG 13. Same as Figure 12 but at a prediction horizon of two quarters.

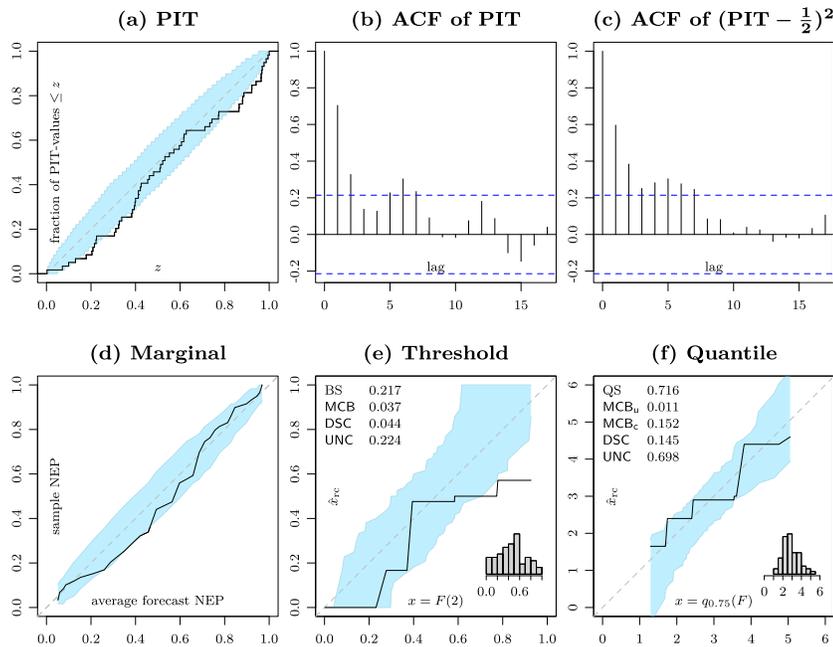


FIG 14. Same as Figure 12 but at a prediction horizon of three quarters.

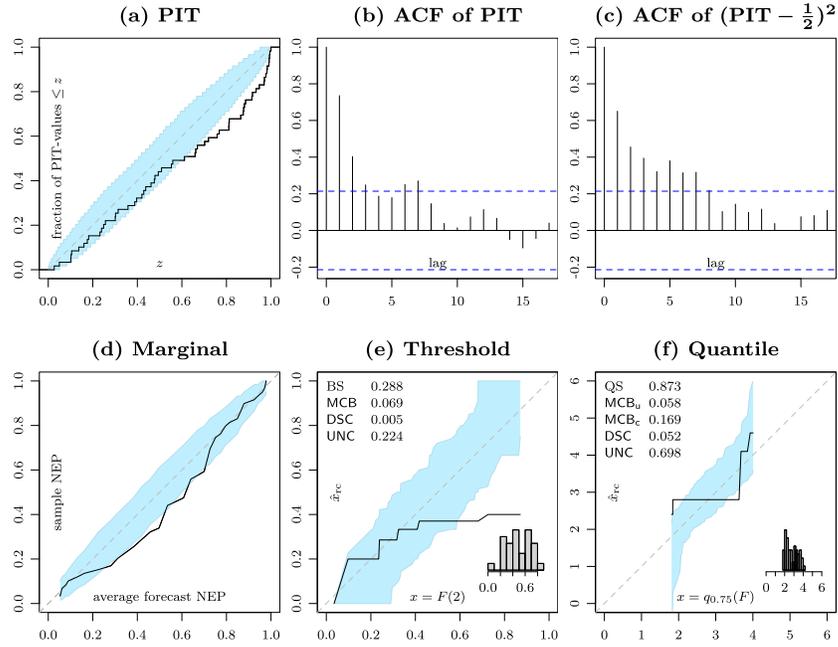


FIG 15. Same as Figure 12 but at a prediction horizon of four quarters.

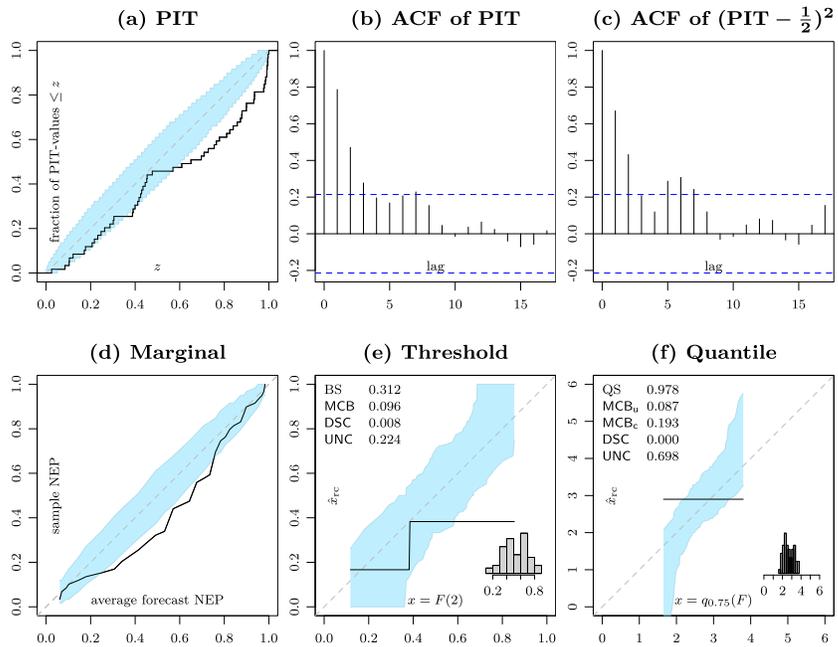


FIG 16. Same as Figure 12 but at a prediction horizon of five quarters.

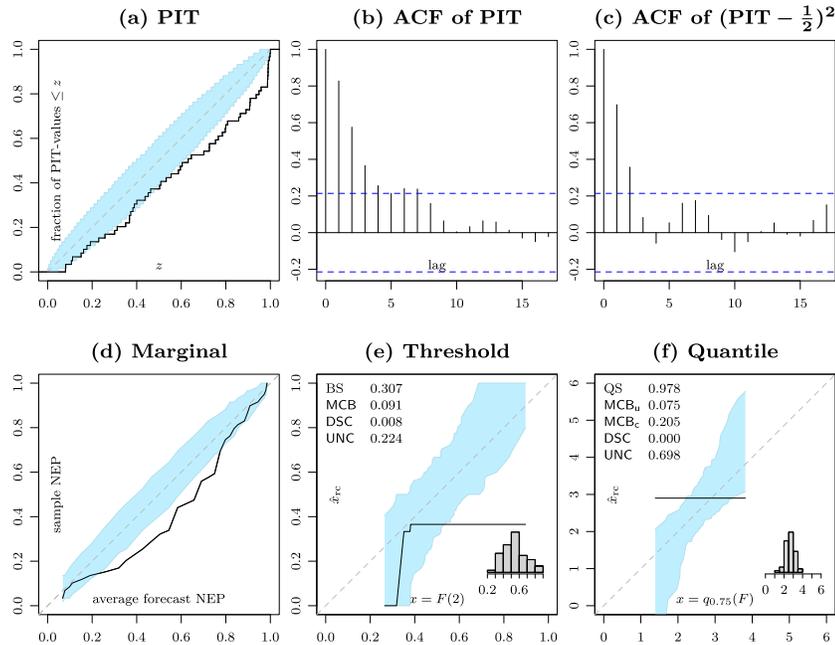


FIG 17. Same as Figure 12 but at a prediction horizon of six quarters.

## Funding

Our research has been funded by the Klaus Tschira Foundation. Johannes Resin gratefully acknowledges support from the German Research Foundation (DFG) through grant number 502572912.

## Supplementary Material

### Replication code

(doi: [10.1214/23-EJS2180SUPP](https://doi.org/10.1214/23-EJS2180SUPP); .zip). R code for replication purposes.

## References

- ALLEN, S. (2021). Advanced Statistical Post-Processing of Ensemble Weather Forecasts, PhD thesis, University of Exeter, UK.
- ARNOLD, S. (2020). Isotonic Distributional Approximation, Master's thesis, Universität Bern, Switzerland.
- AYER, M., BRUNK, H. D., EWING, G. M., REID, W. T. and SILVERMANN, E. (1955). An empirical distribution function for sampling with incomplete information. *Annals of Mathematical Statistics* **26** 641–647. [MR0073895](https://doi.org/10.1080/00036815508839000)

- BARLOW, R. E., BARTHOLOMEW, D. J., BREMNER, J. M. and BRUNK, H. D. (1972). *Statistical Inference Under Order Restrictions: The Theory and Application of Isotonic Regression*. Wiley, New York. [MR0326887](#)
- EL BARMÍ, H. and MUKERJEE, H. (2005). Inferences under a stochastic ordering constraint: The  $k$ -sample case. *Journal of the American Statistical Association* **100** 252–261. [MR2156835](#)
- BASHAYKH, H. (2022). Statistical Assessment of Forecast Calibration, PhD thesis, University of Exeter, UK.
- BENTZIEN, S. and FRIEDERICH, P. (2014). Decomposition and graphical portrayal of the quantile score. *Quarterly Journal of the Royal Meteorological Society* **140** 1924–1934.
- BREIMAN, L. (1992). *Probability*, SIAM Classics ed. Society for Industrial and Applied Mathematics (SIAM), Philadelphia. [MR1163370](#)
- BREITUNG, J. and KNÜPPEL, M. (2021). How far can we forecast? Statistical tests of the predictive content. *Journal of Applied Econometrics* **36** 369–392. [MR4309589](#)
- BRÖCKER, J. (2009). Reliability, sufficiency, and the decomposition of proper scores. *Quarterly Journal of the Royal Meteorological Society* **135** 1512–1519.
- BRÖCKER, J. and BEN BOUALLÈGUE, Z. (2020). Stratified rank histograms for ensemble forecast verification under serial dependence. *Quarterly Journal of the Royal Meteorological Society* **146** 1976–1990.
- BRÖCKER, J. and SMITH, L. A. (2007). Increasing the reliability of reliability diagrams. *Weather and Forecasting* **22** 651–661.
- CASADY, R. J. and CRYER, J. D. (1976). Monotone percentile regression. *Annals of Statistics* **4** 532–541. [MR0403050](#)
- CHUNG, Y., NEISWANGER, W., CHAR, I. and SCHNEIDER, J. (2021). Beyond pinball loss: Quantile methods for calibrated uncertainty quantification. In *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS)*.
- CLEMENTS, M. P. (2004). Evaluating the Bank of England density forecasts of inflation. *The Economic Journal* **114** 844–866.
- CORRADI, V. and SWANSON, N. R. (2007). Predictive density and conditional confidence interval accuracy tests. *Journal of Econometrics* **135** 187–228. [MR2328400](#)
- CZADO, C., GNEITING, T. and HELD, L. (2009). Predictive model assessment for count data. *Biometrics* **65** 1254–1261. [MR2756513](#)
- DAWID, A. P. (1984). Statistical theory: The prequential approach. *Journal of the Royal Statistical Society Series A* **147** 278–292. [MR0763811](#)
- DAWID, A. P. (1986). Probability forecasting. In *Encyclopedia of Statistical Sciences*, **7** 210–218. Wiley-Interscience.
- DAWID, A. P. (2016). Contribution to the discussion of “Of quantiles and expectiles: Consistent scoring functions, Choquet representations and forecast rankings” by W. Ehm, T. Gneiting, A. Jordan and F. Krüger. *Journal of the Royal Statistical Society Series B* **78** 505–562. [MR3506792](#)
- DE LEEUW, J., HORNIK, K. and MAIR, P. (2009). Isotone optimization in R: Pool-adjacent-violators algorithm (PAVA) and active set methods. *Journal of*

- Statistical Software* **32** 1–24.
- DIEBOLD, F. X., GUNTHER, T. A. and TAY, A. S. (1998). Evaluating density forecasts with applications to financial risk management. *International Economic Review* **39** 863–883.
- DIMITRIADIS, T., FISSLER, T. and ZIEGEL, J. F. (2023). Osband’s principle for identification functions. *Statistical Papers*. In press, <https://doi.org/10.1007/s00362-023-01428-x>.
- DIMITRIADIS, T., GNEITING, T. and JORDAN, A. I. (2021). Stable reliability diagrams for probabilistic classifiers. *Proceedings of the National Academy of Sciences of the United States of America* **118** e2016191118. [MR4275118](#)
- DIMITRIADIS, T. and JORDAN, A. I. (2021). reliabilitydiag: Reliability diagrams using isotonic regression. R package version 0.2.0, <https://cran.r-project.org/package=reliabilitydiag>.
- EHM, W. and OVCHAROV, E. Y. (2017). Bias-corrected score decomposition for generalized quantiles. *Biometrika* **104** 473–480. [MR3698267](#)
- EHM, W., GNEITING, T., JORDAN, A. and KRÜGER, F. (2016). Of quantiles and expectiles: Consistent scoring functions, Choquet representations and forecast rankings. *Journal of the Royal Statistical Society Series B* **78** 505–562. [MR3506792](#)
- FISSLER, T. and HOLZMANN, H. (2022). Measurability of functionals and of ideal point forecasts. *Electronic Journal of Statistics* **16** 5019–5034. [MR4490414](#)
- FISSLER, T. and PESENTI, S. M. (2023). Sensitivity measures based on scoring functions. *European Journal of Operational Research* **307** 1408–1423. [MR4543545](#)
- FISSLER, T. and ZIEGEL, J. F. (2016). Higher order elicibility and Osband’s principle. *Annals of Statistics* **44** 1680–1706. [MR3519937](#)
- FISSLER, T. and ZIEGEL, J. F. (2019). Order-sensitivity and equivariance of scoring functions. *Electronic Journal of Statistics* **13** 1166–1211. [MR3935847](#)
- FLACH, P. (2012). *Machine Learning: The Art and Science of Algorithms that Make Sense of Data*. Cambridge University Press, Cambridge. [MR3088204](#)
- GALBRAITH, J. W. (2003). Content horizons for univariate time-series forecasts. *International Journal of Forecasting* **19** 43–55.
- GALBRAITH, J. W. and VAN NORDEN, S. (2012). Assessing gross domestic product and inflation probability forecasts derived from Bank of England fan charts. *Journal of the Royal Statistical Society Series A* **175** 713–727. [MR2948371](#)
- GNEITING, T. (2011). Making and evaluating point forecasts. *Journal of the American Statistical Association* **106** 746–762. [MR2847988](#)
- GNEITING, T., BALABDAOUI, F. and RAFTERY, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society Series B* **69** 243–268. [MR2325275](#)
- GNEITING, T. and KATZFUSS, M. (2014). Probabilistic forecasting. *Annual Review of Statistics and Its Application* **1** 125–151.
- GNEITING, T. and RAFTERY, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* **102**

- 359–378. [MR2345548](#)
- GNEITING, T. and RANJAN, R. (2013). Combining predictive distributions. *Electronic Journal of Statistics* **7** 1747–1782. [MR3080409](#)
- GNEITING, T., WOLFFRAM, D., RESIN, J., KRAUS, K., BRACHER, J., DIMITRIADIS, T., HAGENMEYER, V., JORDAN, A. I., LERCH, S., PHIPPS, K. and SCHIENLE, M. (2023). Model diagnostics and forecast evaluation for quantiles. *Annual Review of Statistics and Its Application* **10** 597–621.
- GNEITING, T. and RESIN, J., (2023). Supplement to “Regression diagnostics meets forecast evaluation: conditional calibration, reliability diagrams, and coefficient of determination”. DOI: [10.1214/23-EJS2180SUPP](#).
- GUNTUBOYINA, A. and SEN, B. (2018). Nonparametric shape-restricted regression. *Statistical Science* **33** 568–594. [MR3881209](#)
- GUO, C., PLEISS, G., SUN, Y. and WEINBERGER, K. Q. (2017). On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*.
- GUPTA, C., PODKOPAEV, A. and RAMDAS, A. (2020). Distribution-free binary classification: Prediction sets, confidence intervals and calibration. In *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS)*.
- HEINRICH, C. (2021). On the number of bins in a rank histogram. *Quarterly Journal of the Royal Meteorological Society* **147** 544–556.
- HELD, L., RUFIBACH, K. and BALABDAOUI, F. (2010). A score regression approach to assess calibration of continuous probabilistic predictions. *Biometrics* **66** 1295–1305. [MR2758518](#)
- HENZI, A., ZIEGEL, J. F. and GNEITING, T. (2021). Isotonic distributional regression. *Journal of the Royal Statistical Society Series B* **83** 963–993. [MR4349124](#)
- HENZI, A. and ZIEGEL, J. F. (2022). Valid sequential inference on probability forecast performance. *Biometrika* **109** 647–663. [MR4472840](#)
- HOLZMANN, H. and EULERT, M. (2014). The role of the information set for forecasting—with applications to risk management. *Annals of Applied Statistics* **8** 595–621. [MR3192004](#)
- HOTHORN, T., KNEIB, T. and BÜHLMANN, P. (2014). Conditional transformation models. *Journal of the Royal Statistical Society Series B* **76** 3–27. [MR3153931](#)
- HUBER, P. J. (1964). Robust estimation of a location parameter. *Annals of Mathematical Statistics* **35** 73–101. [MR0161415](#)
- HYNDMAN, R. J. and KOEHLER, A. B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting* **22** 679–688.
- JOLLIFFE, I. T. and STEPHENSON, D. B. (2012). *Forecast Verification: A Practitioner’s Guide in Atmospheric Science*, second ed. Wiley, Chichester.
- JORDAN, A. I., MÜHLEMANN, A. and ZIEGEL, J. F. (2022). Characterizing the optimal solutions to the isotonic regression problem for identifiable functionals. *Annals of the Institute of Statistical Mathematics* **74** 489–514. [MR4417369](#)
- JULIO, J. M. (2006). The fan chart: The technical details of the new implemen-

- tation. Banco de la República Colombia Bogotá, Borradores de Economía, 468.
- KNÜPPEL, M. (2015). Evaluating the calibration of multi-step-ahead density forecasts using raw moments. *Journal of Business & Economic Statistics* **33** 270–281. [MR3337062](#)
- KOENKER, R. and BASSETT, G. (1978). Regression quantiles. *Econometrica* **46** 33–50. [MR0474644](#)
- KOENKER, R. and MACHADO, J. A. F. (1999). Goodness of fit and related inference processes for quantile regression. *Journal of the American Statistical Association* **94** 1296–1310. [MR1731491](#)
- KRÜGER, F. and ZIEGEL, J. F. (2021). Generic conditions for forecast dominance. *Journal of Business & Economic Statistics* **39** 972–983. [MR4319685](#)
- KULESHOV, V., FENNER, N. and ERMON, S. (2018). Accurate uncertainties for deep learning using calibrated regression. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*.
- KUMAR, A., LIANG, P. S. and MA, T. (2019). Verified uncertainty calibration. In *Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS)*.
- KVÅLSETH, T. (1985). Cautionary note about  $R^2$ . *American Statistician* **39** 279–285.
- LEVI, D., GISPAN, L., GILADI, N. and FETAYA, E. (2022). Evaluating and calibrating uncertainty prediction in regression tasks. *Sensors* **22** 5540.
- MASON, S. J., GALPIN, J. S., GODDARD, L., GRAHAM, N. E. and RAJARTNAM, B. (2007). Conditional exceedance probabilities. *Monthly Weather Review* **135** 363–372.
- MITCHELL, S., POTASH, E., BAROCAS, S., D’AMOUR, A. and LUM, K. (2021). Algorithmic fairness: Choices, assumptions, and definitions. *Annual Review of Statistics and Its Application* **8** 141–163. [MR4243544](#)
- MORIASI, D. N., ARNOLD, J. G., VAN LIEW, M. W., BINGNER, R. L., HARMEL, R. D. and VEITH, T. L. (2007). Model evaluation guidelines for systematic quantification of accuracy in watershed simulations. *Transactions of the ASABE* **50** 885–900.
- MÖSCHING, A. and DÜMBGEN, L. (2020). Monotone least squares and isotonic quantiles. *Electronic Journal of Statistics* **14** 24–49. [MR4047593](#)
- MURPHY, A. H. (1996). General decomposition of MSE-based skill scores: Measures of some basic aspects of forecast quality. *Monthly Weather Review* **124** 2353–2369.
- MURPHY, A. H. and EPSTEIN, E. S. (1989). Skill scores and correlation coefficients in model verification. *Monthly Weather Review* **117** 572–581.
- MURPHY, A. H. and WINKLER, R. L. (1987). A General Framework for Forecast Verification. *Monthly Weather Review* **115** 1330–1338.
- NAKAGAWA, S. and SCHIELZETH, H. (2013). A general and simple method for obtaining  $R^2$  from generalized linear mixed-effects models. *Methods in Ecology and Evolution* **4** 133–142.
- NASH, J. E. and SUTCLIFFE, J. V. (1970). River flow forecasting through conceptual models. Part I – A discussion of principles. *Journal of Hydrology*

- 10** 282–290.
- NIXON, J., DUSENBERRY, M. W., ZHANG, L., JERFEL, G. and TRAN, D. (2019). Measuring calibration in deep learning. In *Proceedings of Computer Vision and Pattern Recognition (CVPR) Conference Workshops*.
- NOCETI, P., SMITH, J. and HODGES, S. (2003). An evaluation of tests of distributional forecasts. *Journal of Forecasting* **22** 447–455.
- NOLDE, N. and ZIEGEL, J. F. (2017). Elicitability and backtesting: Perspectives for banking regulation. *Annals of Applied Statistics* **11** 1833–1874. [MR3743276](#)
- ORJEBIN, E. (2014). A recursive formula for the moments of a truncated univariate normal distribution. Working paper, [https://people.smp.uq.edu.au/YoniNazarathy/teaching\\_projects/studentWork/EricOrjebin\\_TruncatedNormalMoments.pdf](https://people.smp.uq.edu.au/YoniNazarathy/teaching_projects/studentWork/EricOrjebin_TruncatedNormalMoments.pdf).
- PATTON, A. J. (2020). Comparing possibly misspecified forecasts. *Journal of Business & Economic Statistics* **38** 796–809. [MR4154889](#)
- PINSON, P. and HAGEDORN, R. (2012). Verification of the ECMWF ensemble forecasts of wind speed against analyses and observations. *Meteorological Applications* **19** 484–500.
- PLEISS, G., RAGHAVAN, M., WU, F., KLEINBERG, J. and WEINBERGER, K. J. (2017). On fairness and calibration. In *Proceedings of the 31st Conference on Neural Information Processing Systems (NIPS)*.
- POHLE, M. O. (2020). The Murphy decomposition and the calibration-resolution principle: A new perspective on forecast evaluation. Preprint, [arXiv:2005.01835](#).
- ROBERTSON, T. and WRIGHT, F. T. (1980). Algorithms in order restricted statistical inference and the Cauchy mean value property. *Annals of Statistics* **8** 645–651. [MR0568726](#)
- ROBERTSON, T., WRIGHT, F. T. and DYKSTRA, R. L. (1988). *Order Restricted Statistical Inference*. Wiley, Chichester. [MR0961262](#)
- ROELOFS, R., CAIN, N., SHLENS, J. and MOZER, M. C. (2022). Mitigating bias in calibration error estimation. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- RÜSCHENDORF, L. (2009). On the distributional transform, Sklar’s theorem, and the empirical copula process. *Journal of Statistical Planning and Inference* **139** 3921–3927. [MR2553778](#)
- RÜSCHENDORF, L. and DE VALK, V. (1993). On regression representations of stochastic processes. *Stochastic Processes and their Applications* **46** 183–198. [MR1226406](#)
- SAHOO, R., ZHAO, S., CHEN, A. and ERMON, S. (2021). Reliable decisions with threshold calibration. In *Advances in Neural Information Processing Systems*.
- SATOPÄÄ, V. and UNGAR, L. (2015). Combining and extremizing real-valued forecasts. Preprint, [arXiv:1506.06405](#).
- SATOPÄÄ, V. A. (2021). Improving the wisdom of crowds with analysis of variance of predictions of related outcomes. *International Journal of Forecasting* **37** 1728–1747.
- SAVAGE, L. J. (1971). Elicitation of personal probabilities and expectations.

- Journal of the American Statistical Association* **66** 783–801. [MR0331571](#)
- SCHMIDT, K. D. (2011). *Maß und Wahrscheinlichkeit*, revised ed. Springer, Heidelberg. [MR3524794](#)
- SEN, B., BANERJEE, M. and WOODROOFE, M. (2010). Inconsistency of bootstrap: The Grenander estimator. *Annals of Statistics* **38** 1953–1977. [MR2676880](#)
- SHORACK, G. R. and WELLNER, J. A. (2009). *Empirical Processes with Applications to Statistics*, SIAM Classics ed. Society for Industrial and Applied Mathematics (SIAM), Philadelphia. [MR3396731](#)
- SIEGERT, S. (2017). Simplifying and generalising Murphy’s Brier score decomposition. *Quarterly Journal of the Royal Meteorological Society* **143** 1178–1183.
- SONG, H., DIETHE, T., KULL, M. and FLACH, P. (2019). Distribution calibration for regression. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*.
- STEINWART, I., PASIN, C., WILLIAMSON, R. and ZHANG, S. (2014). Elicitation and identification of properties. *Journal of Machine Learning Research: Workshop and Conference Proceedings* **35** 1–45.
- STODDEN, V., MCNUTT, M., BAILEY, D. H., DEELMAN, E., GIL, Y., HANSON, B., HEROUX, M. A., IOANNIDIS, J. P. A. and TAUFER, M. (2016). Enhancing reproducibility for computational methods. *Science* **354** 1240–1241.
- STOYANOV, J. (2000). Krein condition in probabilistic moment problems. *Bernoulli* **6** 939–949. [MR1791909](#)
- STRÄHL, C. and ZIEGEL, J. (2017). Cross-calibration of probabilistic forecasts. *Electronic Journal of Statistics* **11** 608–639. [MR3619318](#)
- TAGGART, R. (2022). Point forecasting and forecast evaluation with generalized Huber loss. *Electronic Journal of Statistics* **16** 201–231. [MR4359360](#)
- TAYLOR, K. E. (2001). Summarizing multiple aspects of model performance in a single diagram. *Journal of Geophysical Research* **106** 7183–7192.
- R CORE TEAM (2021). R: A Language and Environment for Statistical Computing R Foundation for Statistical Computing, Vienna, Austria, <https://www.R-project.org/>.
- TREDENNICK, A. T., HOOKER, G., ELLNER, S. P. and ADLER, P. B. (2021). A practical guide to selecting models for exploration, inference, and prediction in ecology. *Ecology* **102** e03336.
- TSYPLAKOV, A. (2011). Evaluating density forecasts: A comment. Preprint, <http://dx.doi.org/10.2139/ssrn.1907799> .
- TSYPLAKOV, A. (2013). Evaluation of probabilistic forecasts: Proper scoring rules and moments. Preprint, <http://dx.doi.org/10.2139/ssrn.2236605>.
- TSYPLAKOV, A. (2014). Theoretical guidelines for a partially informed forecast examiner. Preprint, <https://mpra.ub.uni-muenchen.de/67333/>.
- VAN CALSTER, B., NIEBOER, D., VERGOUWE, Y., DE COCK, B., PENCINA, M. J. and STEYERBERG, E. W. (2016). A calibration hierarchy for risk models was defined: From utopia to empirical data. *Journal of Clinical Epidemiology* **74** 167–176.
- VAN EEDEN, C. (1958). Testing and Estimating Ordered Parameters of Proba-

- bility Distributions, PhD thesis, University of Amsterdam, Netherlands.
- WALLIS, K. F. (2003). Chi-squared tests of interval and density forecasts, and the Bank of England's fan charts. *International Journal of Forecasting* **19** 165–175.
- WILKS, D. S. (2019). Indices of rank histogram flatness and their sampling properties. *Monthly Weather Review* **147** 763–769.
- WRIGHT, F. T. (1984). The asymptotic behavior of monotone regression estimates. *Canadian Journal of Statistics* **12** 229–236.
- YU, B. and KUMBIER, K. (2020). Veridical data science. *Proceedings of the National Academy of Sciences of the United States of America* **117** 3920–3929.
- ZHAO, S., MA, T. and ERMON, S. (2020). Individual calibration with randomized forecasting. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*.
- ZIEGEL, J. F. (2016). Contribution to the discussion of “Of quantiles and expectiles: Consistent scoring functions, Choquet representations and forecast rankings” by W. Ehm, T. Gneiting, A. Jordan and F. Krüger. *Journal of the Royal Statistical Society Series B: Methodological* **78** 505–562.