Beyond Strictly Proper Scoring Rules: The Importance of Being Local

HAILIANG DU^{a,b}

^a *Department of Mathematical Sciences, Durham University, Durham, United Kingdom*

^b *Centre for the Analysis of Time Series, London School of Economics, London, United Kingdom*

(Manuscript received 4 October 2019, in final form 9 December 2020)

ABSTRACT: The evaluation of probabilistic forecasts plays a central role both in the interpretation and in the use of forecast systems and their development. Probabilistic scores (scoring rules) provide statistical measures to assess the quality of probabilistic forecasts. Often, many probabilistic forecast systems are available while evaluations of their performance are not standardized, with different scoring rules being used to measure different aspects of forecast performance. Even when the discussion is restricted to strictly proper scoring rules, there remains considerable variability between them; indeed strictly proper scoring rules need not rank competing forecast systems in the same order when none of these systems are perfect. The locality property is explored to further distinguish scoring rules. The nonlocal strictly proper scoring rules considered are shown to have a property that can produce “unfortunate” evaluations, particularly the fact that the continuous rank probability score prefers the outcome close to the median of the forecast distribution regardless of the probability mass assigned to the value at/near the median raises concern to its use. The only local strictly proper scoring rule, the logarithmic score, has direct interpretations in terms of probabilities and bits of information. The nonlocal strictly proper scoring rules, on the other hand, lack meaningful direct interpretation for decision support. The logarithmic score is also shown to be invariant under smooth transformation of the forecast variable, while the nonlocal strictly proper scoring rules considered may, however, change their preferences due to the transformation. It is therefore suggested that the logarithmic score always be included in the evaluation of probabilistic forecasts.

KEYWORDS: Forecast verification/skill; Probability forecasts/models/distribution; Model evaluation/performance; Decision support

1. Introduction

Forecast evaluation has a long history of being a crucial topic for model development and decision support. The outputs from a stochastic model can be naturally interpreted in the form of a probabilistic forecast. Given a deterministic model, uncertainty in the initial state due to the observational noise, limited computational power, and model discrepancy prevent one from making a perfect deterministic forecast of the future or even identifying the truth in the past. To account for all sorts of uncertainties, the model outputs are often interpreted as probabilistic forecasts with the aim of providing useful information for decision support. Probabilistic forecasts have been widely adopted in various fields including meteorology, social science, pharmacology, economics, and finance, and have become common in operational forecasting over the last quarter century.

The evaluation of probabilistic forecasts plays a central role both in the interpretation and in the use of forecast systems and their development. Such evaluation has not yet been standardized, with many different probabilistic scoring rules (Gneiting and Raftery 2007; Jolliffe and Stephenson 2003; Roulston and Smith 2002; Wilks 1995) being used. As probabilistic forecasts become more common, the need to select (probabilistic) scoring rule(s) for constructing probabilistic forecasts, calibrating forecast systems, ranking competing forecast systems and quantifying forecast improvement has led to the research work presented in this paper.

The importance of using strictly proper scoring rules has been noted in the literature (Bröcker and Smith 2007), as only strictly proper scoring rules encourage the forecaster to be honest, i.e., reporting a forecast probability distribution gives an optimal expected score only when the verification is, in fact, drawn from that probability distribution. When the discussion is restricted to strictly proper scoring rules, however, there remains considerable variability between scoring rules (there are, in fact, an infinite number of strictly proper scoring rules). And strictly proper scoring rules need not rank competing forecast systems in the same order when none of these systems are perfect.

The locality property is explored to further distinguish various strictly proper scoring rules. A property that reflects “unfortunate”¹ evaluations is introduced. Nonlocal strictly proper scoring rules considered are shown to have a mathematical property, named implausible, that could produce unfortunate evaluations. A few striking examples of the potential issues that result from the use of nonlocal scoring rules are presented. The only local strictly proper scoring rule, the logarithmic score (also known as “ignorance”), has direct interpretations in terms of probabilities and bits of information. The nonlocal strictly proper scoring rules are found to lack meaningful direct interpretation. The logarithmic score is also shown to be invariant under smooth transformation of the forecast

¹ Fortuna was the goddess of fortune (luck) in Roman religion. “Unfortunate” in this paper refers to bad advice, which is a disaster in terms of decision support.

Corresponding author: Hailiang Du, hailiang.du@durham.ac.uk

variable, while the nonlocal strictly proper scoring rules considered may, however, change their preferences due to the transformation.

This paper emphasizes the fact that being strictly proper is not sufficient in decision support when measuring the difference between imperfect forecast systems and suggests that the only local scoring rule, ignorance, should always be included in the evaluation of probabilistic forecasts.

The definition of a scoring rule for probabilistic forecast and the importance of using strictly proper scoring rules are presented in [section 2](#). A number of strictly proper scoring rules are defined in [section 3](#). A common example of strictly proper scoring rules ranking forecast systems differently without the presence of True underlying distribution is given in [section 4](#). The locality property is defined and discussed in [section 5](#). [Section 6](#) introduces a mathematical property that reflects “unfortunate” evaluations and shows nonlocal scoring rules considered have such property. The interpretation of local and nonlocal scoring rules is discussed in [section 7](#). [Section 8](#) investigates the behavior of proper scoring rules when smooth transformation is applied to the forecast variable. [Section 9](#) provides discussion and conclusions.

2. Probabilistic scoring rules and importance of being strictly proper

While the true value of a forecast is most clearly reflected in its utility to the end user, probabilistic scores are fundamental to the performance analysis of probabilistic forecasts. Ideally they provide a general measure of future forecast quality, independent of any specific end user ([Bröcker and Smith 2007](#)). A probabilistic score (scoring rule) is a function $S[p(x), Y]$, where $p(x)$ is a probability density function and Y is the outcome. In this paper, probabilistic forecasts in the form of probability density functions (PDFs) $p(x)$ are considered.² The notation $p(x)$ denotes the entire function, while $p(Y)$ always denotes the value of the function at the particular outcome Y . By convention, a lower score is taken to reflect a better forecast. Analytically, the *expected score*,

$$E\{S[p(x), Y]\} = \int S[p(x), Y]Q(Y) dY, \quad (1)$$

which takes the expectation of the scoring rule under the true underlying distribution Q from which the outcome Y is drawn, quantifies the quality of a forecast system. In practice, an archive of forecast–outcome pairs is required to evaluate the quality of a forecast system. It contains a large number N of forecasts $\{p_i(x), i = 1, \dots, N\}$ and corresponding outcomes $\{Y_i, i = 1, \dots, N\}$. The forecast system yields an *empirical score*:

$$S_{\text{emp}} = \frac{1}{N} \sum_{i=1}^N S[p_i(x), Y_i]. \quad (2)$$

² Results and conclusions presented in this paper also apply to probabilistic forecasts in the form of probability mass functions in the context of categorical variables.

Note the size of the forecast archive can play a major role in determining the significance of the result ([Machete and Smith 2016](#)), regardless of which scoring rule is employed ([Bröcker and Smith 2007](#)).

Several scoring rules are widely used for the evaluation of probabilistic forecasts ([Brier 1950](#); [Epstein 1969](#); [Gneiting and Raftery 2007](#); [Good 1952](#); [Jolliffe and Stephenson 2003](#); [Mason and Weigel 2009](#); [Selten 1998](#); [Wilks 1995](#)); different scoring rules might quantify different attributes of the forecast. Note, however, that [Good \(1952\)](#)’s logarithmic score, also known as ignorance, (defined below in [section 3d](#)) is the only scoring rule consistent with the use of (log) likelihoods to evaluate assessors or Bayesian inference ([Winkler 1969, 1996](#)).

Since any functional form based on $p(x)$ and Y could be considered as a scoring rule, one may introduce and use a scoring rule that favors to particular forecast system which might lead to dishonest and misleading evaluations where the scoring rule encourages the forecaster to select a probabilistic forecast distribution that the forecaster knows is not correct [For example the well-known Finley (1884) tornado forecasts ([Murphy 1996](#); [Stephenson 2000](#)).] To avoid such dishonest evaluation, strictly proper ([Roby 1965](#); [Toda 1963](#)) scoring rules (defined in the following) are preferred. The term, *proper*, was first introduced by [Winkler and Murphy \(1968\)](#), while the general idea goes back to [Brier \(1950\)](#) and [Good \(1952\)](#). A scoring rule, $S[p(x), Y]$, is said to be *proper* if inequality (3) holds for any pair of forecast PDFs, and *strictly proper* when equality implies $p = q$:

$$\int q(z)S[p(x), z] dz \geq \int q(z)S[q(x), z] dz. \quad (3)$$

For a given forecast p , a scoring rule evaluated at the outcome is a random variable with values that depend on the outcome Y . Note being strictly proper is a property of the functional form of the scoring rule alone, not of the particular distribution $p(x)$ or $q(x)$. Strictly proper scoring rules give a probabilistic forecast distribution an optimal expected score only when the outcome is, in fact, drawn from that probability distribution ([Bröcker and Smith 2007](#)). In expectation, a strictly proper scoring rule does not judge any other forecast p to score better than q as a forecast of q itself. Note that the interpretation of strictly proper does not, however, require one to believe that the true underlying distribution Q exists. Strictly proper is a property of the scoring rule; it is neither necessary to assume that Y is drawn from any kind of true distribution nor that any kind of data are to hand.

The question of whether the employed scoring rule is strictly proper or not can be answered independently of any data being considered ([Bröcker and Smith 2007](#)). Although concerns of hedging are often mentioned ([Selten 1998](#)), strictly proper scoring rules are preferred even when there is no human involvement, as in parameter selection ([Du and Smith 2012](#)).

While the importance of using strictly proper scoring rules is well recognized ([Bröcker and Smith 2007](#); [Brown 1970](#); [Fricker et al. 2013](#)), researchers often face requests to present results under a variety of scoring rules, both proper and nonproper scoring rules. The fact that nonproper scoring rules like root mean squared error (RMSE) are still widely used in forecast

evaluation often leads to confusion and poorly optimized forecast systems. There have been many discussions regarding the evil of RMSE in the literature [see Bröcker and Smith (2007), McSharry and Smith (1999), Smith et al. (2015), and Wilks (1995)], therefore RMSE, which is in fact not a strictly proper scoring rule, will not be considered in this paper.

3. Strictly proper scoring rules

A variety of strictly proper scoring rules have been introduced since the 1950s. Some of those widely used are listed below:

a. Energy score family

Gneiting and Raftery (2007) introduced the energy score family based on (Székely 2003) statistical energy perspective. The energy score family S_{ES} is defined as follows:³

$$S_{ES}[p(\mathbf{x}), \mathbf{Y}] = E_p \|\mathbf{x} - \mathbf{Y}\|^\beta - \frac{1}{2} E_p \|\mathbf{x} - \mathbf{x}'\|^\beta, \quad (4)$$

where $\beta \in (0, 2)$ is a real number; $\mathbf{x}$ and $\mathbf{x}'$ are independent copies of a random vector with distribution p ; and $\|\cdot\|$ denotes the Euclidean norm. Székely (2003) and Gneiting and Raftery (2007) show that the energy score is strictly proper relative to the class $\mathbb{P}_\beta$, where $\mathbb{P}_\beta$ denotes the class of the Borel probability measures p such that $E_p \|\mathbf{x}\|^\beta$ is finite. When $\beta = 1$, one obtains

$$S_{CRPS}[p(x), Y] = E_p \|x - Y\| - \frac{1}{2} E_p \|x - x'\|. \quad (5)$$

This is equivalent to the well-known continuous ranked probability score⁴ (CRPS) (Matheson and Winkler 1976; Unger 1995) (and see Baringhaus and Franz (2004), Székely and Rizzo (2005), and Gneiting and Raftery (2007) for the proof of equivalence), where it is the integral of the square of the L^2 distance between the cumulative distribution function (CDF) of the forecast p and a step function at the outcome (Epstein 1969):

$$S_{CRPS}[p(x), Y] = \int \left[\int_{-\infty}^x p(z) dz - H(x - Y) \right]^2 dx, \quad (6)$$

where the Heaviside (step) function H is defined as follows:

$$H(x) = \begin{cases} 0, & \text{if } x < 0 \\ 1, & \text{if } x \geq 0. \end{cases} \quad (7)$$

The CRPS was to our knowledge first published by Brown (1974). It can also be considered as a generalization of the Brier score (Brier 1950) [the Brier score only applies to binary

outcomes (Matheson and Winkler 1976; Murphy and Winkler 1970)]. For a point forecast, the CRPS is equal to the mean absolute error. In the past decade, the CRPS has been widely used by the atmospheric sciences community (Goddard et al. 2013; Scheuerer 2014; Zhang et al. 2014).

b. Power score family

Let α be a real number with $\alpha > 1$. The power score family (Selten 1998) S_{PS} is defined as follows:

$$S_{PS}[p(x), Y] = -\alpha p(Y)^{\alpha-1} + (\alpha - 1) \int p^\alpha(z) dz. \quad (8)$$

The power score family is also strictly proper; this can simply be derived from the derivatives of the expected score (Selten 1998). When $\alpha = 2$, one obtains the proper linear score (PLS) [also called the quadratic score (Brier 1950)]:

$$S_{PLS}[p(x), Y] = -2p(Y) + \int p^2(z) dz. \quad (9)$$

PLS derives from the (naive) linear score (Staël von Holstein 1970b), $S_{LS}[p(x), Y] = -p(Y)$, which is not a proper scoring rule as the (naive) linear score favors a $p(x)$ featuring a very small spread and which is centered at the point x^* for which $Q(x^*)$ is very large (Bröcker and Smith 2007).

c. Pseudospherical score family

Good (1971) introduced the pseudospherical score family S_{PSS} (β is a real number with $\beta > 1$), defined as follows:

$$S_{PSS}[p(x), Y] = -\frac{p(Y)^{\beta-1}}{\left[\int p^\beta(z) dz \right]^{1/\beta}}. \quad (10)$$

The pseudospherical score family is strictly proper; this can be derived using the Hölder and Minkowski inequality. When $\beta = 2$, one obtains the traditional spherical score (SPS):

$$S_{SPS}[p(x), Y] = -\frac{p(Y)}{\left[\int p^2(z) dz \right]^{1/2}}. \quad (11)$$

d. Ignorance

Good (1952) introduced the logarithmic score [also known as ignorance (Roulston and Smith 2002)] given by⁵

$$S[p(x), Y] = -\log_2[p(Y)], \quad (12)$$

where $p(Y)$ is the density assigned to the outcome Y . Ignorance (IGN) is a strictly proper scoring rule; this can be derived using Kullback–Leibler inequality (Kullback and Leibler 1951). The expected (with respect to p) IGN is also a famous information

³Negative orientation is applied to the original energy score defined by Gneiting and Raftery (2007) so that it is consistent with the convention that a lower score reflects a better forecast.

⁴The continuous ranked probability score is generalized from the ranked probability score (RPS) (Epstein 1969; Murphy 1969), which is widely used to evaluate discrete (categorical) probabilistic forecasts.

⁵Note that defining the logarithmic score in terms of $\log_2$ is equivalent to the alternative definition in terms of $\ln$ up to the factor $1/\ln 2$, which does not affect rankings of different forecast systems.

measure, Shannon entropy. In addition, the expected IGN of p relative to a distribution q becomes the classical Kullback–Leibler divergence (Kullback 1959).

4. Different strictly proper scoring rules rank forecast systems differently

Obviously the energy score family, power score family, and pseudospherical score family contain an infinite number of strictly proper scoring rules. Jose et al. (2008) have introduced weighted scoring rules by blending the power score with the pseudospherical score; the weighted scoring rule is shown to be strictly proper too. Furthermore, Toda (1963) proved that a linear transformation of a strictly proper scoring rule is also strictly proper. Given a strictly proper scoring rule, a forecast system providing Q will always be preferred whenever it is included among those under consideration. When none of the competing forecast systems are perfect, then even strictly proper scoring rules may rank two forecast systems differently, making it impossible to provide definitive statements regarding the relative merit of imperfect forecast systems without considering an additional measure of forecast quality.

Consider the case where outcomes are independent random draws from a standard Gaussian distribution. Two forecast systems are constructed, where forecast system A uses $N(0, \sigma^2)$ and forecast system B uses $N(0, 1/\sigma^2)$ where $\sigma > 1$. Obviously neither of the forecast systems is perfect; forecast system A represents a wider distribution around 0 with larger standard deviation while forecast system B represents a narrower distribution with smaller standard deviation. Figure 1 shows the expectation [under the true distribution, $N(0, 1)$] of various scoring rules (ignorance,⁶ continuous rank probability score, proper linear score, and spherical score) of forecast system A relative to forecast system B as a function of σ . If the relative score (also known as *skill score*) is negative, it indicates forecast system A outperforms forecast system B. Both IGN and PLS prefer wider forecast distribution (forecast system A) than narrower forecast distribution (forecast system B). The CRPS, on the contrary, ranks forecast system B higher than A. Interestingly SPS considers both imperfect forecast system as having the same forecast quality with the expected relative score being zero. (Note given any finite sample of forecasts, there is a 50% chance that the empirical SPS prefers forecast A (or B) to the other.) A more thorough investigation in contrasting how certain scoring rules would rank competing forecasts of specified departures from the target distribution can be found in Machete (2013).

FIG. 1. The expectation of various scoring rules of forecast system A, $N(0, \sigma^2)$, relative to forecast system B, $N(0, 1/\sigma^2)$, where the outcome is drawn from a standard Gaussian distribution. A negative relative score suggests system A outperforms system B.

5. Locality

To distinguish between strictly proper scoring rules, the locality property is explored here. A scoring rule is *local* if the probabilistic forecast is evaluated only at the actual outcome, which means that the scoring rule depends solely on the probability assigned to the outcome, rather than being rewarded for other features of the forecast distribution, such as its shape. Shuford (1966) and Bernardo (1979) show that every local, smooth, and proper scoring rule for continuous variables is equivalent to (an affine function of) IGN, which makes IGN the only proper local scoring rule for continuous variables. Thus all other proper scoring rules, including those listed in section 3, are nonlocal. The locality property itself does not suggest whether local or nonlocal scoring rules should be preferred, although it might seem unreasonable that features of the forecast other than the value it assigned to the outcome should matter at all. In the following sections, the preference of a local scoring rule is supported based on both mathematical properties and interpretation of the scoring rule.

6. Implausible

A mathematical property called implausible is introduced in this section. Nonlocal scoring rules listed in section 3 are shown to have such undesirable property; striking unfortunate evaluation examples that result from the use of nonlocal scoring rules are presented below.

A scoring rule is *implausible*,⁸ if for any $r > 1$, $r \in \mathbb{R}$, there exist two forecast systems $p_1(x)$ and $p_2(x)$, and Y , where

⁶ Ignorance is downscaled by a factor of 20 in order to have a similar scale with other scoring rules in Fig. 1.

⁷ This particular example, based on Gaussian distributions, is designed to show that different scoring rules may rank forecast systems differently and not to indicate whether each of the scoring rules considered here prefers wider or narrower forecast distributions in general, which is not true either.

⁸ Smith et al. (2015) defined “perverse” scoring rules to be those that systematically prefer forecasts that place a lower probability on the outcome. Maynard (2016) considered a scoring rule to be “not feasible” when a probable event scores worse than an improbable one. These definitions are elevated here.

$p_1(Y)/p_2(Y) = r$ while $S(p_1, Y) > S(p_2, Y)$. In other words, an implausible scoring rule means that for all $r > 1$ it is possible to find $p_1(x)$, $p_2(x)$, and Y (which may all vary with r) such that $p_1(Y)/p_2(Y) = r$ and $S(p_1, Y) > S(p_2, Y)$. Ignorance is clearly not implausible as given $p_1(Y)/p_2(Y) = r$, $S(p_1, Y)$ would always be smaller than $S(p_2, Y)$ by $\log_2 r$. The energy score family is implausible; this can be shown via investigating an undesirable mathematical property of the energy score. Take the derivative of the energy score respect to the outcome Y (where Y is a realization of the random variable x):

$$\frac{\partial S_{\text{ES}}[p(x), Y]}{\partial Y} = \int_{-\infty}^Y \beta(Y-x)^{\beta-1} p(x) dx - \int_Y^{\infty} \beta(x-Y)^{\beta-1} p(x) dx. \quad (13)$$

The zero solution of the RHS of Eq. (13) only relies on the location of Y regardless the value of $p(Y)$. For the CRPS, where $\beta = 1$, $\min_Y S(p, Y)$ is achieved when $\int_{-\infty}^Y p(x) dx - \int_Y^{\infty} p(x) dx = 0$, which gives Y as being the median of $p(x)$. Such mathematical property may lead to “unfortunate” results as illustrated in Fig. 2.⁹ The blue line and red line represent two forecast systems A and B (each based on a bimodal distribution with the same shape but different centers). Intuitively, one would expect that if the outcome lands between -0.5 and 0.5 (or more generally that the outcome is drawn from some PDF which is bounded between -0.5 and 0.5) forecast system B shall be preferred as system B would assign significantly more probability mass to the outcome than system A (especially when the outcome lands around 0); similarly if the outcome lands between 0.5 and 1.5 forecast system A shall be preferred. The green line represents the CRPS of system A relative to system B, a negative (below the dotted zero line) relative score suggests system A outperforms system B according to the CRPS. It appears that if the outcome lands between -0.5 and 0.5 , the CRPS would prefer system A over B even when system B assigns significantly more probability mass to the outcome than system A. This is due to the fact that the CRPS prefers the outcome to be close to the median of the forecast distribution no matter how much probability mass is around the median. Obviously the CRPS is implausible, as shown in Fig. 2, when $Y = 1$, $p_A(Y)/p_B(Y) = \infty$ while $S_{\text{CRPS}}(p_A, Y) > S_{\text{CRPS}}(p_B, Y)$ (similar examples can be found to show all members of the energy score family are implausible). Ironically, if the forecast system A is delivered to the user, the developer of forecast system A would hope the outcome lands at 0 in order to achieve the best CRPS despite the fact the forecast system A assigns 0 probability to the outcome. Considering a parameter estimation scenario, if the observed outcomes are drawn from a delta function or a sharp Gaussian distribution centered at 0 and the forecast distribution is a bimodal distribution with its center to be tuned, tuning the parameter based on the CRPS would converge to a bimodal distribution centered at

FIG. 2. Example showing that the continuous rank probability score produces “unfortunate” results. The blue line and the red line represent PDFs of forecast systems A and B based on bimodal distributions with the same shape but different centers. The green line represents the CRPS of system A relative to system B. A negative relative score suggests system A outperforms system B. The dashed vertical lines enclose the regions where unfortunate results occur.

0 where the probability mass assign to the outcome would always be near 0.

The example shown in Fig. 2 contradicts the claim (Kohonen and Suomela 2006; Boero et al. 2011; Tödter and Ahrens 2012) that the CRPS/RPS gives credit for assigning high probabilities to the values near but not identical to the outcome. This kind of claims is mostly originated from (Staël von Holstein 1970a), where Staël von Holstein showed that the RPS is “sensitive to distance” from the “true” outcome. Actually the sensitive to distance defined by Staël von Holstein is based on his definition of “more distant from the true event” [P360 of (Staël von Holstein 1970a)], which is, however, NOT equivalent to assigning high probabilities to the values near but not identical to the outcome. It was, in fact, noted by Staël von Holstein himself [section 6 of (Staël von Holstein 1970a)] that his definition of more distant from the true event is rather restrictive and changing the definition to an alternative definition (Murphy 1970) will lead to the RPS not being sensitive to distance, which is consistent with the example shown in Fig. 2.

The power score and spherical score are also implausible. This can be shown in the case, where $p_1(x)$ and $p_2(x)$ are both Gaussian distributions.

Let $p_1(x)$ be a Gaussian distribution with mean u_1 and standard deviation σ_1 , then

$$\begin{aligned} \int_{-\infty}^{\infty} p_1^\alpha(z) dz &= \int_{-\infty}^{\infty} \left\{ \frac{1}{\sqrt{2\pi}\sigma_1} \exp \left[-\frac{(z-u_1)^2}{2\sigma_1^2} \right] \right\}^\alpha dz \\ &= (2\pi)^{\frac{1-\alpha}{2}} \alpha^{-1/2} \sigma_1^{1-\alpha}. \end{aligned} \quad (14)$$

Let $p_2(x)$ be a Gaussian distribution with mean u_2 and standard deviation σ_2 . To prove the power score family is implausible,

⁹ It is usual to analyze how scores change as the forecast varies, while this and the following examples investigate how the values of scores change as the outcome varies.

one needs to find $p_1(\cdot)$, $p_2(\cdot)$, and Y so that $p_1(Y) = rp_2(Y)$ but $S_{PS}[p_1(x), Y] > S_{PS}[p_2(x), Y]$, which requires

$$\begin{aligned} & -\alpha p_1(Y)^{\alpha-1} + (\alpha-1) \int_{-\infty}^{\infty} p_1^{\alpha}(z) dz > -\alpha p_2(Y)^{\alpha-1} \\ & + (\alpha-1) \int_{-\infty}^{\infty} p_2^{\alpha}(z) dz \\ & -\alpha p_1(Y)^{\alpha-1} + (\alpha-1)(2\pi)^{\frac{1-\alpha}{2}} \alpha^{-1/2} \times \sigma_1^{1-\alpha} > -\alpha p_2(Y)^{\alpha-1} \\ & + (\alpha-1)(2\pi)^{\frac{1-\alpha}{2}} \alpha^{-1/2} \sigma_2^{1-\alpha}. \end{aligned} \quad (15)$$

Note that even if $p_2(Y) = 0$, it is still possible that $S_{PS}[p_1(x), Y] > S_{PS}[p_2(x), Y]$, as long as $S_{PS}[p_1(x), Y] > 0$, as one can always find σ_2 large enough so that $(\alpha-1) \int_{-\infty}^{\infty} p_2^{\alpha}(z) dz$ is smaller than $S_{PS}[p_1(x), Y]$. To have $S_{PS}[p_1(x), Y] > 0$:

$$\begin{aligned} p_1(Y) & < (\alpha-1)^{\frac{1}{\alpha-1}} \alpha^{-\frac{3}{\alpha-1}} \frac{1}{\sqrt{2\pi}\sigma_1} \\ p_1(Y) & < (\alpha-1)^{\frac{1}{\alpha-1}} \alpha^{-\frac{3}{\alpha-1}} \frac{1}{2(\alpha-1)} p_1(u_1). \end{aligned} \quad (16)$$

This condition also defines a vulnerable subspace where the evaluation using power scores might be misinformative. Figure 3 gives an example where PLS may produce unfortunate results. The blue line and red line represent the PDFs of two forecast systems A and B. Intuitively one would expect that if the outcome is less than -4 (or between -2 and -1), system A shall be preferred as system A would assign significantly more probability mass around the outcome than system B. On the contrary, the relative PLS (the green line) prefers system B instead as positive relative PLS would be observed as shown in Fig. 3. Ironically, if the forecast systems A and B are delivered to the user, the developer of forecast system B would hope for the outcome being smaller than -4 in order to “outperform” forecast system A by achieving better PLS despite the fact that forecast system B assigns ~ 0 probability to the outcome.

Similarly to prove the pseudospherical score family is implausible, one needs to find $p_1(\cdot)$, $p_2(\cdot)$ and Y so that $p_1(Y) = rp_2(Y)$ but $S_{PSS}[p_1(x), Y] > S_{PSS}[p_2(x), Y]$, which requires

$$\begin{aligned} & -\frac{p_1(Y)^{\beta-1}}{\left[\int_{-\infty}^{\infty} p_1^{\beta}(z) dz\right]^{1/\beta}} > -\frac{p_2(Y)^{\beta-1}}{\left[\int_{-\infty}^{\infty} p_2^{\beta}(z) dz\right]^{1/\beta}}, \\ & -\frac{p_1(Y)^{\beta-1}}{(2\pi)^{\frac{1-\beta}{2}} \frac{1}{2\beta} \sigma_1^{\frac{1-\beta}{\beta}}} > -\frac{p_2(Y)^{\beta-1}}{(2\pi)^{\frac{1-\beta}{2}} \frac{1}{2\beta} \sigma_2^{\frac{1-\beta}{\beta}}}, \\ & -\frac{[rp_2(Y)]^{\beta-1}}{\sigma_1^{\frac{1-\beta}{\beta}}} > -\frac{p_2(Y)^{\beta-1}}{\sigma_2^{\frac{1-\beta}{\beta}}}, \\ & \sigma_2 > r^{\beta} \sigma_1. \end{aligned} \quad (17)$$

FIG. 3. Example showing that the proper linear score produces unfortunate results due to the fact it is implausible. The blue line and the red line represent PDFs of forecast system A $[N(-3, 0.5^2)]$ and B $[N(3, 1^2)]$. The green line represents the PLS of system A relative to system B. A negative relative score suggests system A outperforms system B. The left side of the dashed vertical line at -4 and the region between dashed vertical lines at -2 and -1 define the regions where unfortunate results occur.

Note the condition in Eq. (17) also places a restriction onto Y , as σ_2 gets larger, the maximum value of $p_2(x)$ can be smaller than $p_1(Y)/r$. Therefore Y has to be chosen so that $p_1(Y)/r \leq p_2(u_2)$, i.e., $p_1(Y)/r \leq 1/\sqrt{2\pi}\sigma_2$ and as $\sigma_2 > r^{\beta}\sigma_1$, it requires $p_1(Y) < (1/\sqrt{2\pi}\sigma_1)r^{1-\beta}$. This condition also defines a vulnerable subspace (given $r > 1$) where the evaluation using the pseudospherical score might be misinformative.

Figure 4 gives an example where SPS may produce unfortunate results. Consider two forecast systems based on Gaussian distributions, where the PDF of system A (blue line) is standard Gaussian and system B (red line) being $N(0, 5^2)$. Intuitively, one would expect that if the outcome lands in the two regions bounded by the black dashed vertical lines, system A shall be preferred as system A would assign significantly more probability mass around the outcome than system B. On the contrary, the relative SPS (the green line) prefers system B instead as positive relative SPS would be observed in both regions.

7. Score interpretation

The difference between two forecast systems is reflected by the difference between their scores. This provides a rank ordering, and thus a preference. Without any reference, a single score of a forecast system hardly provides any evaluation information, which is why score interpretation should be considered base on the relative score between forecast systems. It is also helpful if the relative score has some meaningful interpretation that relates to the benefit of the users. Otherwise it only indicates which forecast system is better, without answering the question of how much better one is in a meaningful way that adds value to decision support. For example, Fig. 1

FIG. 4. Example showing that the spherical score produces unfortunate results due to the fact it is implausible. The blue line and the red line represent PDFs of forecast system A (standard Gaussian) and B $[N(0, 5^2)]$. The green line represents the SPS of system A relative to system B. A negative relative score suggests system A outperforms system B. The dashed vertical lines enclose the regions where unfortunate results occur.

does show the IGN, PLS, and CRPS preferences between the two forecast systems; however, the interpretation of the relative score (the y axis in Fig. 1) is also important to decision-makers.

A number of meaningful interpretations to proper scoring rules have been identified in the literature. IGN can be interpreted in terms of gambling returns (Good 1952; Hagedorn and Smith 2009; Kelly 1956; Roulston and Smith 2002). Under a Kelly betting scenario,¹⁰ IGN describes the rate at which the forecaster's fortune increases with time. A house setting fair odds (Frigg et al. 2014) based on a forecast system with a lower value of a nonlocal scoring rule is expected to lose money to a gambler who places bets based on a different forecast system with a lower IGN. Through its close relation to Shannon's information entropy, IGN is related to the amount of information expected from a forecast (Roulston and Smith 2002). IGN can also be easily communicated as an effective interest rate (Hagedorn and Smith 2009). Jose et al. (2008) show that the pseudospherical score and power score families can be interpreted as profits in certain decision problems. Note all the interpretations listed above are based on some specific scenario in which one can, in fact, define a corresponding utility function to replace the scoring rule. For example in a Kelly betting contest, one can define a utility function that reflects the rate (at which the forecasters'

fortune increases with time) and use such utility function to replace IGN for forecast evaluation. In practice, it is usually not easy to define a relevant utility function based on probabilistic forecasts for the use of decision support. It is therefore desirable for a scoring rule to have a rather direct and generic interpretation.

The expected IGN can be written as

$$E\{S_{\text{IGN}}[p(x), Y]\} = \int [-\log_2 p(Y)] Q(Y) dY. \quad (18)$$

And the expected relative IGN between two probabilistic forecast system p_1 and p_2 is

$$\int \left[-\log_2 \frac{p_1(Y)}{p_2(Y)} \right] Q(Y) dY. \quad (19)$$

Therefore the empirical relative IGN score, $(1/N) \sum -\log_2 [p_1(Y)/p_2(Y)]$, reflects the (average) increase in probability mass that the model forecast p_1 placed on the outcome relative to that of the reference forecast p_2 . Note that although p_1 and p_2 are probability density functions, $-\log_2 [p_1(Y)/p_2(Y)]$ can be interpreted as increase/decrease in probability,¹¹ which gives the ignorance score a meaningful direct interpretation.¹² The relative IGN of two forecast systems also quantifies the information gain (in terms of bits) the model forecast system provides over the reference system. A relative IGN of 1 bit means that, on average, forecasts from the system assign twice the probability to the outcome compared to the reference forecast (Roulston and Smith 2002).

Nonlocal scoring rules include contributions from the entire PDF; the scoring rule may be largely determined by outcomes that did NOT occur, making a meaningful direct interpretation somewhat challenging. For example, the empirical relative PLS between two forecast system p_1 and p_2 based on a large number N of forecast–outcome pairs is

$$\left[\int p_1^2(z) dz - \int p_2^2(z) dz \right] + \frac{2}{N} \sum_i [p_2(Y_i) - p_1(Y_i)]. \quad (20)$$

The interpretation of Eq. (20) is clearly more sophisticated than that of the relative IGN. In the second term of Eq. (20), $p_2(Y) - p_1(Y)$, which ranges $(-\infty, \infty)$, is the difference between two probability density functions rather than two probabilities. In the context of decision support, it is unclear how to interpret the probability density function(s) meaningfully other than by using $\log p_2(Y) - \log p_1(Y)$ to reflect the increase/decrease in probability mass placed on Y (this is in fact the approach used by relative IGN). The first term of Eq. (20) being a function of the entire PDF of forecast systems (not depending on the outcome Y) clearly makes it even more challenge for interpretation. Similar interpretation challenges apply to the CRPS and SPS. There are better ways to interpret these nonlocal

¹⁰ In a Kelly betting contest (Kelly 1956), one bets all of one's wealth on every outcome in proportion to the forecast probability of that outcome. More precisely, a fraction ω_i of one's wealth, where ω_i is the forecast probability of event E_i occurring, should be wagered on the i th outcome.

¹¹ For small values of δ , one can write $P(x < X \leq x + \delta) \approx p(X)\delta$.

¹² Note the interpretation of the empirical relative ignorance score does not require the knowledge of the true underlying distribution Q .

scoring rules by using the true underlying distribution as a reference.

For example, assuming the true underlying distribution Q exists then the expectation of PLS is

$$E\{S_{\text{PLS}}[p(x), Y]\} = \int \left[-2p(Y) + \int p^2(z) dz \right] Q(Y) dY. \quad (21)$$

PLS is based on the idea that the scoring rule should reflect “nearness” of the predicted probability distribution to the true underlying distribution. By straightforward manipulation, it comes to the following representation:

$$E\{S_{\text{PLS}}[p(x), Y]\} = \int [Q(Y) - p(Y)]^2 dY - \int Q^2(Y) dY. \quad (22)$$

The second term in the RHS of Eq. (22) will vanish when comparing two forecast systems using the expected relative score, which gives the expected relative PLS between two probabilistic forecast system p_1 and p_2 :

$$\int [Q(Y) - p_1(Y)]^2 dY - \int [Q(Y) - p_2(Y)]^2 dY. \quad (23)$$

Therefore the expected relative PLS between two forecast systems can be interpreted with regard to the mean square difference between the forecast distribution and the true underlying distribution Q .

Similarly, the expectation of CRPS can be written as

$$E\{S_{\text{CRPS}}[p(x), Y]\} = \int [G(Y) - F(Y)]^2 dY - \int G(Y)[1 - G(Y)] dY, \quad (24)$$

where $F(\cdot)$ is the CDF of the forecast distribution and $G(\cdot)$ is the true underlying CDF. The expectation of the relative CRPS between two forecast systems can be interpreted with regard to the mean square difference between the forecast CDF and the CDF of the truth.

The expected SPS, $E\{S_{\text{SPS}}[p(x), Y]\}$, can be written as

$$\left[\int Q(Y)^2 dY \right]^{1/2} \frac{\int p(Y)Q(Y) dY}{\left(\int Q(Y)^2 dY \right)^{1/2} \left[\int p(Y)^2 dY \right]^{1/2}}. \quad (25)$$

It can be interpreted regarding the interior angle of deviation between the forecast distribution p and the true underlying distribution Q .

In some cases it makes sense to consider an integration over the true underlying distribution Q . The interpretation of the expected relative score with respect to Q is cloudy in reality, for example in weather-like forecasting scenarios, where the same Q distribution is never seen twice over the lifetime of the system. In practice, the true underlying distribution is rarely (if ever) available to provide such an interpretation, and were it to be the use of imperfect probabilistic forecast is mute. Furthermore, functions of the entire forecast distribution, as in Eqs. (22), (24), and (25) can hardly be interpreted in a meaningful way for decision support. Therefore, even if the true underlying distribution

were available, it is unclear the interpretations of relative scores derived from Eqs. (23)–(25) are informative to the decision-maker except providing their preference between two forecast systems.

8. Scoring rules under transformation

In practice, it is common that the variable of interest is not the variable observed but a function of the observed variable. For example, wind power is a function of wind speed cubed; wave power is principally a function of wave height squared (Savenkov 2009). It is desirable for a scoring rule to provide coherent evaluations before and after a smooth transformation being applied to the forecast variable. Consider $x^* = \phi(x)$ as a smooth one-to-one (transformation) function of a random variable x . The forecast PDF of x , $p(x)$, becomes $p[\phi^{-1}(x^*)][d\phi^{-1}(x^*)/dx^*]$ for the random variable x^* after the transformation and the scoring rule $S[p(x), Y]$ becomes $S\{p[\phi^{-1}(x^*)][d\phi^{-1}(x^*)/dx^*], Y^*\}$, where $Y^* = \phi(Y)$. It is almost certain that the value of a scoring rule will change after the transformation. Note score interpretation should always be based on the relative score between forecast systems instead of a single score of a forecast system in order to provide useful information for decision support. It is therefore of interest to investigate whether the relative score will change after taking the transformation and if so, will the scoring rule's preference change as well. Given the relative score between two probabilistic forecast system p_1 and p_2 by

$$S[p_1(x), Y] - S[p_2(x), Y], \quad (26)$$

after taking a transformation $\phi(x)$ it becomes

$$S\{p_1[\phi^{-1}(x^*)][d\phi^{-1}(x^*)/dx^*], Y^*\} - S\{p_2[\phi^{-1}(x^*)][d\phi^{-1}(x^*)/dx^*], Y^*\}. \quad (27)$$

Note that Y and Y^* are one-to-one and $p(x)$ and $p[\phi^{-1}(x^*)][d\phi^{-1}(x^*)/dx^*]$ reflect the same information of a forecast system. Therefore if (26) does not equal to (27) based on some scoring rule S , the scoring rule will have a nonunique interpretation of the relative skill between two competing forecast systems. Furthermore if $(26) \times (27) < 0$ for some Y and ϕ , it indicates such scoring rule might also change its preference due to the transformation, then the use of such scoring rule as an evaluation tool for decision support is questionable.

Ignorance score is *invariant* under smooth transformation as (26) and (27) are equal for any smooth transformation, proved in the following:

$$\begin{aligned} & S\left\{p_1[\phi^{-1}(x^*)]\frac{d\phi^{-1}(x^*)}{dx^*}, Y^*\right\} - S\left\{p_2[\phi^{-1}(x^*)]\frac{d\phi^{-1}(x^*)}{dx^*}, Y^*\right\} \\ &= -\log p_1[\phi^{-1}(x^*)]\frac{d\phi^{-1}(x^*)}{dx^*}\bigg|_{Y^*} + \log p_2[\phi^{-1}(x^*)]\frac{d\phi^{-1}(x^*)}{dx^*}\bigg|_{Y^*} \\ &= -\log p_1(Y) + \log p_2(Y) = S[p_1(x), Y] - S[p_2(x), Y]. \end{aligned} \quad (28)$$

For nonlocal scoring rules like proper linear score, spheric score, and continuous rank probability score, smooth transformation not only have impact on the value of the relative scores but also may cause the change of their preference. Figure 5 gives an example where the CRPS may contradict itself by changing its preference under transformation (similar examples can be found for PLS and SPS). Following Fig. 2, Fig. 5a compares two forecast systems based on a bimodal distribution with the same shape but different centers. The green line represents the CRPS of system A relative to system B, a negative relative score suggests system A outperforms system B according to the CRPS. The black dashed vertical line in Fig. 5a corresponds to the threshold $Y = 11.5$, where the CRPS prefers forecast system A when $Y < 11.5$ and prefers forecast system B when $Y > 11.5$. Figure 5b compares the same two forecast systems after cubic transformation being applied to the forecast variable. Clearly the relative CRPS has changed after the cubic transformation. Let x refer to wind speed, then x^3 reflects wind power. When the observed wind speed is 10, the relative CRPS (forecast system A relative to forecast system B) is roughly -0.9 as in Fig. 5a. Comparing the same¹³ forecast system A and B in terms of wind power under the same observation (wind speed 10 corresponds to wind power 1000), however, the relative CRPS¹⁴ as in Fig. 5b becomes roughly 340, which indicates the interpretation of CRPS evaluation for comparing forecast system A and B based on a unique observed wind speed is not unique. Furthermore, the CRPS may even change its preference after the transformation as the threshold (the black solid vertical line in Fig. 5b) that distinguishes the CRPS preference is $Y^3 = 1700$ rather than $Y^3 = 11.5 = 1520.875$ (dashed vertical line). If the underlying distribution of the wind speed were bounded between the black dashed line and black solid line, the CRPS would prefer forecast system A before the cubic transformation when the wind speed is evaluated directly, while it prefer forecast system B after the cubic transformation when the wind power (which corresponds to the same wind speed) is evaluated.

9. Discussion and summary

Measures of skill play a critical role in the development, deployment, and application of probabilistic forecasts. The properties of some common strictly proper scoring rules have been discussed. Given a strictly proper scoring rule, the true forecast system will always be preferred whenever it is included among those under consideration. In practice, to correctly measure the difference between imperfect forecast schemes, being strictly proper is not enough, as strictly proper scoring rules need not rank competing forecast systems in the same order when none of the forecast systems

FIG. 5. Example showing that the continuous rank probability score changes its preference under transformation. The blue line and the red line represent PDFs of forecast systems A and B. The green line represents the CRPS of system A relative to system B. A negative relative score suggests system A outperforms system B. (a) Before the cubic transformation and (b) after the cubic transformation. The black dashed vertical line and solid line in (a), where $Y = 11.5$ and $Y = 11.935$, respectively, correspond to those in (b), where $Y = 11.5^3$ and $Y = 11.935^3$.

are perfect. In general, any scoring rules can be presented with the form:

$$S[p(x), Y] = s_1[p(x)] + s_2[p(x), Y] + s_3[p(Y)]. \quad (29)$$

For local scoring rules, the first two terms in the RHS of Eq. (28) are both zero with only the presence of $s_3[p(Y)]$, for example the only local proper scoring rule, the logarithmic score (ignorance). Nonlocal scoring rules contain at least one of the first two terms, for example the energy scores consist of s_1 and s_2 , the power scores s_1 and s_3 , and the pseudosphere scores only s_2 . The presence of s_1 or s_2 or both allows the

¹³ In fact the information presented by forecast systems A and B remain the same although their pdf changed when the forecast variable wind speed transformed to wind power.

¹⁴ The relative CRPS (the green curve) in Fig. 5b is scaled down by 1.6×10^5 in order to have similar magnitude as the PDFs.

scoring rule to give extra credit to the structure of the forecast PDF. Note such extra credit is not necessarily given for assigning high probabilities to the values near the outcome as the examples in Figs. 2–4 show. Without knowing the true underlying distribution, the justification of giving such extra credit is untenable. Nonlocal strictly proper scoring rules considered¹⁵ are shown to have property that can produce unfortunate evaluations due to the fact that contributions from the entire shape of the PDF may overwhelm that from the probability assigned to the outcome. Particularly the fact that continuous rank probability score prefers the outcome close to the median of the forecast distribution regardless the probability mass assigned to the value at/near the median raises concern to the use of continuous rank probability score. Ignorance has direct interpretations in terms of probabilities and bits of information while the direct interpretation of nonlocal strictly proper scoring rules on the other hand relies on information regarding the unknown (if it even exists) true underlying distribution as a reference. The nonlocal strictly proper scoring rules considered may also contradict themselves when a smooth transformation is applied to the forecast variable while IGN is shown to be invariant under smooth transformation. It is suggested that ignorance should always be included in the evaluation of probabilistic forecasts.

One of the reasons for using nonlocal scoring rules is to address particular problems where a local scoring rule is not considered “suitable.” For example, ignorance is infinity if the forecast assigns vanishing probability to an event that obtains. Selten (1998) emphasizes that the use of ignorance implies the value judgment that small differences between small probabilities should be taken very seriously and that wrongly describing something extremely improbable as having zero probability is “an unforgivable sin.” Roulston and Smith (2002) pointed out that forecasters should replace zero forecast probabilities with small probabilities based on the uncertainties in the forecast PDF. Not to do so means reporting the improbable as the impossible. Within the Bayesian framework, Cromwell’s rule states that the use of prior probabilities of 0 or 1 should be avoided. Assigning zero probability to events that are possible also contradicts Laplace’s rule of succession (Jaynes 2003). In the insurance sector, the premium is inversely proportional to the probability of an event occurring; zero probability would suggest free insurance.

In this manuscript, the value outcome is assumed to be certain. In the presence of uncertainty in the value of the outcome (e.g., due to measurement error), one may obtain benefit by assigning probability to the events that do not match the outcome exactly. Note again this does not imply one should use nonlocal scoring rules, as for nonlocal scoring rules the

contributions from the entire shape of the PDF are not designed to account for the uncertainty in the value of the outcome. For a local scoring rule, the evaluation can still be considered over the observational uncertainty distribution of the outcome, for example by coupling the forecast distribution with the distribution of observational noise (Bröcker and Smith 2007).

Scoring rules are designed to assess (probabilistic) forecast performance, which hopefully leads to better decision-making. Bernardo and Smith (2000) argue that a local proper score should be preferred for “pure inference” problems in which the outcome is the sole arbiter of forecast quality, yet there are other forms of scoring rules that would typically be appropriate in more directly practical contexts [see stock control example in Bernardo and Smith (2000)]. Note that in such “more directly practical contexts,” if a utility function based on probabilistic forecasts can be conveniently defined according to the practical objective (which often is not the case in practice), there is no need for any kind of scoring rules (using the utility function directly will serve the purpose of forecast evaluation sufficiently). Any scoring rules can be directly considered as a utility function, yet the meaning of the corresponding utility function relies on the direct interpretation of the skill of the scoring rule. It is questionable whether nonlocal scoring rules can provide any meaningful direct interpretation. L. A. Smith (2020, personal communication) claims that interpretation is a critical aspect in accepting a scoring rule for use in practice; he uses valued property of probabilistic forecasts to support this assertion.

Acknowledgments. This research was supported the EPSRC-funded Uncertainty analysis of hierarchical energy systems models: Models versus real energy systems (EP/K03832X/1) and Centre for Energy Systems Integration (EP/P001173/1). Additional support was also provided by evaluating probability scores for the Insurance Sector funded by LSE KEI and Lighthill Risk Network. The author would like to thank Leonard A. Smith and Edward Wheatcroft for reading earlier versions of this article and giving useful feedback; and two anonymous reviewers whose comments and suggestions helped improve and clarify this article.

REFERENCES

- Baringhaus, L., and C. Franz, 2004: On a new multivariate two-sample test. *J. Multivariate Anal.*, **88**, 190–206, [https://doi.org/10.1016/S0047-259X\(03\)00079-4](https://doi.org/10.1016/S0047-259X(03)00079-4).
- Bernardo, J. M., 1979: Expected information as expected utility. *Ann. Stat.*, **7**, 686–690, <https://doi.org/10.1214/aos/1176344689>.
- , and A. F. M. Smith, 2000: *Bayesian Theory*. Wiley, 608 pp.
- Boero, G., J. Smith, and K. F. Wallis, 2011: Scoring rules and survey density forecasts. *Int. J. Forecasting*, **27**, 379–393, <https://doi.org/10.1016/j.ijforecast.2010.04.003>.
- Brier, G. W., 1950: Verification of forecasts expressed in terms of probability. *Mon. Wea. Rev.*, **78**, 1–3, [https://doi.org/10.1175/1520-0493\(1950\)078<0001:VOFEIT>2.0.CO;2](https://doi.org/10.1175/1520-0493(1950)078<0001:VOFEIT>2.0.CO;2).
- Bröcker, J., and L. A. Smith, 2007: Scoring probabilistic forecasts: The importance of being proper. *Wea. Forecasting*, **22**, 382–388, <https://doi.org/10.1175/WAF966.1>.

¹⁵ The author does believe that all nonlocal scoring rules are implausible, but as there is no general mathematical function of nonlocal scoring rules, only the popular nonlocal scoring rules (and their families) are considered and shown to be implausible in this manuscript.

- Brown, T. A., 1970: Probabilistic forecasts and reproducing scoring systems. Tech. Rep. RM-6299-ARPA, RAND Corporation, 65 pp.
- , 1974: Admissible scoring systems for continuous distributions. Manuscript P-5235, The Rand Corporation, 24 pp.
- Du, H., and L. A. Smith, 2012: Parameter estimation using ignorance. *Phys. Rev. E*, **86**, 016213, <https://doi.org/10.1103/PhysRevE.86.016213>.
- Epstein, E. S., 1969: A scoring system for probability forecasts of ranked categories. *J. Appl. Meteor.*, **8**, 985–987, [https://doi.org/10.1175/1520-0450\(1969\)008<0985:ASSFPF>2.0.CO;2](https://doi.org/10.1175/1520-0450(1969)008<0985:ASSFPF>2.0.CO;2).
- Fricker, T. E., C. A. T. Ferro, and D. B. Stephenson, 2013: Three recommendations for evaluating climate prediction. *Meteor. Appl.*, **20**, 246–255, <https://doi.org/10.1002/met.1409>.
- Frigg, R., S. Bradley, H. Du, and L. Smith, 2014: Laplace's demon and the adventures of his apprentices. *Philos. Sci.*, **81**, 31–59, <https://doi.org/10.1086/674416>.
- Gneiting, T., and A. E. Raftery, 2007: Strictly proper scoring rules, prediction and estimation. *J. Amer. Stat. Assoc.*, **102**, 359–378, <https://doi.org/10.1198/016214506000001437>.
- Goddard, L., and Coauthors, 2013: A verification framework for interannual-to-decadal predictions experiments. *Climate Dyn.*, **40**, 245–272, <https://doi.org/10.1007/s00382-012-1481-2>.
- Good, I. J., 1952: Rational decisions. *J. Roy. Stat. Soc.*, **14A**, 107–114.
- , 1971: Comment on “Measuring information and uncertainty” by Robert J. Buehler. *Foundations of Statistical Inference*, Holt, Rinehart and Winston, 337–339.
- Hagedorn, R., and L. Smith, 2009: Communicating the value of probabilistic forecasts with weather roulette. *Meteor. Appl.*, **16**, 143–155, <https://doi.org/10.1002/met.92>.
- Jaynes, E. T., 2003: *Probability Theory: The Logic of Science*. Cambridge University Press, 727 pp.
- Jolliffe, I. T., and D. B. Stephenson, 2003: *Forecast Verification: A Practitioner's Guide in Atmospheric Science*. Wiley, 254 pp.
- Jose, V. R. R., R. F. Nau, and R. L. Winkler, 2008: Scoring rules, generalized entropy, and utility maximization. *Oper. Res.*, **56**, 1146–1157, <https://doi.org/10.1287/opre.1070.0498>.
- Kelly, J. L., 1956: A new interpretation of information rate. *Bell Syst. Tech. J.*, **35**, 917–926, <https://doi.org/10.1002/j.1538-7305.1956.tb03809.x>.
- Kohonen, J., and J. Suomela, 2006: Lessons learned in the challenge: Making predictions and scoring them. *Machine Learning Challenges: Evaluating Predictive Uncertainty, Visual Object Classification, and Recognising Textual Entailment*, J. Quiñero-Candela et al., Eds., Springer, 95–116.
- Kullback, S., 1959: *Information Theory and Statistics*. Wiley, 395 pp.
- , and R. A. Leibler, 1951: On information and sufficiency. *Ann. Math. Stat.*, **22**, 79–86, <https://doi.org/10.1214/aoms/1177729694>.
- Machete, R., 2013: Contrasting probabilistic scoring rules. *J. Stat. Plan. Inference*, **143**, 1781–1790, <https://doi.org/10.1016/j.jspi.2013.05.012>.
- , and L. Smith, 2016: Demonstrating the value of larger ensembles in forecasting physical systems. *Tellus*, **68A**, 28393, <https://doi.org/10.3402/tellusa.v68.28393>.
- Mason, S. J., and A. P. Weigel, 2009: A generic forecast verification framework for administrative purposes. *Mon. Wea. Rev.*, **137**, 331–349, <https://doi.org/10.1175/2008MWR2553.1>.
- Matheson, J. E., and R. L. Winkler, 1976: Scoring rules for continuous probability distributions. *Manage. Sci.*, **22**, 1087–1096, <https://doi.org/10.1287/mnsc.22.10.1087>.
- Maynard, T., 2016: Extreme insurance and the dynamics of risk. Ph.D. thesis, The London School of Economics and Political Science, 416 pp.
- McSharry, P., and L. A. Smith, 1999: Better nonlinear models from noisy data: Attractors with maximum likelihood. *Phys. Rev. Lett.*, **83**, 4285–4288, <https://doi.org/10.1103/PhysRevLett.83.4285>.
- Murphy, A. H., 1969: On the “ranked probability score.” *J. Appl. Meteor.*, **8**, 988–989, [https://doi.org/10.1175/1520-0450\(1969\)008<0988:OTPS>2.0.CO;2](https://doi.org/10.1175/1520-0450(1969)008<0988:OTPS>2.0.CO;2).
- , 1970: The ranked probability score and the probability score: A comparison. *Mon. Wea. Rev.*, **98**, 917–924, [https://doi.org/10.1175/1520-0493\(1970\)098<0917:TRPSAT>2.3.CO;2](https://doi.org/10.1175/1520-0493(1970)098<0917:TRPSAT>2.3.CO;2).
- , 1996: The Finley affair: A signal event in the history of forecast verification. *Wea. Forecasting*, **11**, 3–20, [https://doi.org/10.1175/1520-0434\(1996\)011<0003:TFAASE>2.0.CO;2](https://doi.org/10.1175/1520-0434(1996)011<0003:TFAASE>2.0.CO;2).
- , and R. L. Winkler, 1970: Scoring rules in probability assessment and evaluation. *Acta Psychol.*, **34**, 273–286, [https://doi.org/10.1016/0001-6918\(70\)90023-5](https://doi.org/10.1016/0001-6918(70)90023-5).
- Roby, T. B., 1965: Belief states: A preliminary empirical study. *Behav. Sci.*, **10**, 255–270.
- Roulston, M., and L. Smith, 2002: Evaluating probabilistic forecasts using information theory. *Mon. Wea. Rev.*, **130**, 1653–1660, [https://doi.org/10.1175/1520-0493\(2002\)130<1653:EPFUIT>2.0.CO;2](https://doi.org/10.1175/1520-0493(2002)130<1653:EPFUIT>2.0.CO;2).
- Savenkov, M., 2009: On the truncated Weibull distribution and its usefulness in evaluating the theoretical capacity factor of potential wind (or wave) energy sites. *Univ. J. Eng. Technol.*, **1**, 21–25.
- Scheuerer, M., 2014: Probabilistic quantitative precipitation forecasting using ensemble model output statistics. *Quart. J. Roy. Meteor. Soc.*, **140**, 1086–1096, <https://doi.org/10.1002/qj.2183>.
- Selten, R., 1998: Axiomatic characterization of the quadratic scoring rule. *Exp. Econ.*, **1**, 43–61, <https://doi.org/10.1023/A:1009957816843>.
- Shuford, E. H., Jr., A. Albert, and H. E. Massengill, 1966: Admissible probability measurement procedures. *Psychometrika*, **31**, 125–145, <https://doi.org/10.1007/BF02289503>.
- Smith, L. A., E. B. Suckling, E. L. Thompson, T. Maynard, and H. Du, 2015: Towards improving the framework for probabilistic forecast evaluation. *Climatic Change*, **132**, 31–45, <https://doi.org/10.1007/s10584-015-1430-2>.
- Stäel von Holstein, C.-A. S., 1970a: A family of strictly proper scoring rules which are sensitive to distance. *J. Appl. Meteor.*, **9**, 360–364, [https://doi.org/10.1175/1520-0450\(1970\)009<0360:AFOSPS>2.0.CO;2](https://doi.org/10.1175/1520-0450(1970)009<0360:AFOSPS>2.0.CO;2).
- , 1970b: Measurement of subjective probability. *Acta Psychol.*, **34**, 146–159, [https://doi.org/10.1016/0001-6918\(70\)90013-2](https://doi.org/10.1016/0001-6918(70)90013-2).
- Stephenson, D. B., 2000: Use of the “odds ratio” for diagnosing forecast skill. *Wea. Forecasting*, **15**, 221–232, [https://doi.org/10.1175/1520-0434\(2000\)015<0221:UOTORF>2.0.CO;2](https://doi.org/10.1175/1520-0434(2000)015<0221:UOTORF>2.0.CO;2).
- Székel, G. J., 2003: E-statistics: The energy of statistical samples. Tech. Rep. 2003-16, Bowling Green State University, 20 pp., <https://doi.org/10.13140/RG.2.1.5063.9761>.
- , and M. L. Rizzo, 2005: A new test for multivariate normality. *J. Multivariate Anal.*, **93**, 58–80, <https://doi.org/10.1016/j.jmva.2003.12.002>.
- Toda, M., 1963: Measurement of subjective probability distributions. Tech. Doc. esd-tdr-63-407, Decision Sciences Laboratory, Electronic Systems Division, Air Force Systems Command, Vol. 86, 42 pp.
- Tödter, J., and B. Ahrens, 2012: Generalization of the ignorance score: Continuous ranked version and its decomposition. *Mon. Wea. Rev.*, **140**, 2005–2017, <https://doi.org/10.1175/MWR-D-11-00266.1>.

- Unger, D. A., 1995: A method to estimate the continuous ranked probability score. *Proc. Ninth Conf. on Probability and Statistics*, Boston, MA, Amer. Meteor. Soc., 206–213.
- Wilks, D. S., 1995: *Statistical Methods in the Atmospheric Sciences: An Introduction*. International Geophysics Series, Vol. 59, Elsevier, 467 pp.
- Winkler, R. L., 1969: Scoring rules and the evaluation of probability assessors. *J. Amer. Stat. Assoc.*, **64**, 1073–1078, <https://doi.org/10.1080/01621459.1969.10501037>.
- , 1996: Scoring rules and the evaluation of probabilities. *Test*, **5**, 1–60, <https://doi.org/10.1007/BF02562681>.
- , and A. H. Murphy, 1968: “Good” probability assessors. *J. Appl. Meteor.*, **7**, 751–758, [https://doi.org/10.1175/1520-0450\(1968\)007<0751:PA>2.0.CO;2](https://doi.org/10.1175/1520-0450(1968)007<0751:PA>2.0.CO;2).
- Zhang, Y., J. Wang, and X. Wang, 2014: Review on probabilistic forecasting of wind power generation. *Renew. Sustain. Energy Rev.*, **32**, 255–270, <https://doi.org/10.1016/j.rser.2014.01.033>.