http://eprints.lse.ac.uk

Contrasting Probabilistic Scoring Rules

Reason L. Machete*

Dept. of Mathematics and Statistics, P. O. Box 220, Reading, RG6 6AX, UK

Dated: March 15, 2013

Abstract

There are several scoring rules that one can choose from in order to score prob-
abilistic forecasting models or estimate model parameters. Whilst it is generally
agreed that proper scoring rules are preferable, there is no clear criterion for pre-
ferring one proper scoring rule above another. This manuscript compares and con-
trasts some commonly used proper scoring rules and provides guidance on scoring
rule selection. In particular, it is shown that the logarithmic scoring rule prefers
erring with more uncertainty, the spherical scoring rule prefers erring with lower
uncertainty, whereas the other scoring rules are indifferent to either option.

Keywords: estimation; forecast evaluation; probabilistic forecasting; utility function

1 Introduction

Issuing probabilistic forecasts is meant to express uncertainty about the future evolu-
tion of some quantity of interest. Such forecasts arise in many applications such as
macroeconomics, finance, weather and climate forecasting. There are several scoring
rules that one can choose from in order to elicit probabilistic forecasts, rank competing
forecasting models or estimate forecast distribution parameters. It is generally agreed
that one should select scoring rules that encourage a forecaster to state his 'best' judge-
ment of the distribution, the so called *proper* scoring rules (Friedman, 1983; Nau, 1985;
Gneiting and Raftery, 2007), but which one to use is generally an open question. We
shall take scoring rules to be loss functions that a forecaster wishes to minimise. Scoring
rules that are minimised if and only if the issued forecasts coincide with the forecaster's
best judgement are said to be *strictly proper* (Gneiting and Raftery, 2007; Brocker and
Smith, 2007). We shall restrict our attention to strictly proper scoring rules.

Nonetheless using scoring rules to rank competing forecasting models poses a prob-
lem; scoring rules do not provide a universally acceptable ranking of performance. In
estimation, different scoring rules will yield different parameter estimates (Gneiting and
Raftery, 2007; Johnstone and Lin, 2011). Moreover, a forecaster's best judgement may
depart from the ideal; the ideal is a distribution that nature or the data generating pro-
cess would give (Gneiting *et al.*, 2007). Although strictly proper scoring rules encourage
experts to issue their best judgements, such judgements may yet differ from each other

---

*email: r.l.machete@lse.ac.uk

and the ideal. Which scoring rule should one use to choose between two experts? Savage (1971) made the instructive statement that "any criteria for distinguishing among scoring rules must arise out of departures of actual subjects from the ideal." There have been some efforts to contrast scoring rules, but none seem to have followed this insight.

Bickel (2007) made empirical comparisons of the quadratic, spherical and logarithmic scoring rules and found them to yield different rankings of competing forecasts but failed to see why. Considering a concave nonlinear utility function that explicitly depends on the scoring rule, he also found the logarithmic scoring rule to yield the least departures from honest opinions at maximal utility, a point he claimed favours it as a rule of choice. But a utility function need not be exponential nor explicitly depend on the scoring rule. Jose *et al.* (2008) considered weighted scoring rules and showed that they correspond to different utility functions. A limiting feature of the utility functions considered is that they are defined on bounded intervals; there are many applications in which the variable of interest is unbounded. Their motivation for weighted scoring rules is based on betting arguments, but it is not clear what the betting strategies (if any) are. Recently, Boero *et al.* (2011) empirically compared the Quadratic Probability Score (QPS), Ranked Probability Score (RPS) and the logarithmic scoring rule on UK inflation forecasts by the Monetary Policy Committee and the Survey of External Forecasters (SEF). They found the scoring rules to rank the two sets of distributions similarly. Upon ranking individual forecasters from the SEF, they found the RPS to have better discriminatory power than the QPS, a feature they attributed to the RPS's sensitivity to distance. Despite the foregoing efforts, there is lacking a theoretical assessment of what the preferences of the commonly used scoring rules are with respect to the ideal.

This paper contrasts how different scoring rules would rank competing forecasts of specified departures from ideal forecasts and provides guidance on scoring rule selection. It focuses upon those scoring rules that are commonly used in the forecasting literature, including econometrics and meteorology. More specifically, we contrast the relative information content of forecasts preferred by different scoring rules. Implications of the results on decision making are then suggested, noting that it may be desirable to be more or less uncertain when communicating probabilistic forecasts. We realise that an appropriate utility function may be unknown (Bickel, 2007) and expected utility theory may not even be appropriate (Kahneman and Tversky, 1979).

In section 2, we consider three scoring rules of categorical forecasts, which then inspires our study of density forecasts in section 3, where we consider four scoring rules: the Quadratic Score (Gneiting and Raftery, 2007), Logarithmic Score (Good, 1952), Spherical Score (Friedman, 1983) and Continuous Ranked Probability Score (Epstein, 1969). We conclude with a discussion of the results in section 4.

2   Categorical Forecasts

In this section, we consider the scoring of categorical forecasts. The scoring rules considered are Brier score (Brier, 1950), the logarithmic scoring rule and the spherical scoring rule (Friedman, 1983). In order to aid intuition in the next section, here we focus on the binary case. Another commonly used scoring rule for categorical forecasts is the Ranked Probability Score (RPS) (Epstein, 1969). In the binary case, the RPS score reduces to the Brier score.

It will be useful to be aware of the following basics. Given any vectors $\boldsymbol{f}, \boldsymbol{g} \in \Re^m$,

the *inner product* between the two vectors is

$$\langle \boldsymbol{f}, \boldsymbol{g} \rangle = \sum_{i=1}^{m} f_i g_i,$$

from which the L_2-norm is defined by $||\boldsymbol{f}||_2 = \langle \boldsymbol{f}, \boldsymbol{f} \rangle^{1/2}$.

2.1 The Brier score

Consider a probabilistic forecast $\{f_i\}_{i=1}^{m}$ of m categorical events. Suppose the true distribution is $\{p_i\}_{i=1}^{m}$. If the actual outcome is the jth category, the Brier score is given by (Brier, 1950)

$$BS(\boldsymbol{f}, j) = \frac{1}{m} \sum_{i=1}^{m} (f_i - \delta_{ij})^2,$$

where $\delta_{ij} = 0$ if $i \neq j$ and $\delta_{ij} = 1$ if $i = j$. If follows that if we expand out the bracket we get

$$BS(\boldsymbol{f}, j) = \frac{1}{m} \left(\sum_{i=1}^{m} f_i^2 - 2f_j + 1 \right).$$

The expected Brier score is then given by

$$\begin{aligned}
\mathbb{E}[BS(\boldsymbol{f}, J)] &= \sum_{j=1}^{m} p_j BS(f, j) \\
&= \frac{1}{m} \sum_{i=1}^{m} \left(f_i^2 - 2p_i f_i + p_i \right) \\
&= \frac{1}{m} \sum_{i=1}^{m} \left[(f_i - p_i)^2 + p_i - p_i^2 \right] \\
&= \frac{1}{m} \left\{ ||\boldsymbol{\gamma}||_2^2 + \sum_{i=1}^{m} p_i(1 - p_i) \right\},
\end{aligned}$$

where $\boldsymbol{\gamma}$ is a vector with components $\gamma_i = f_i - p_i$ for all $i = 1, \ldots, m$. It is evident from the last expression on the right hand side that the Brier score is effective with respect to the metric $d_2(\boldsymbol{f}, \boldsymbol{g}) = ||\boldsymbol{f} - \boldsymbol{g}||_2$. When $m = 2$, we can put $f_1 = p + \gamma$, $p_1 = p$ and $p_2 = q$ and obtain

$$\mathbb{E}[BS(\boldsymbol{f}, J)] = \gamma^2 + pq.$$

It follows that $\pm \gamma$ will yield the same Brier score. This means the Brier score does not discriminate between over-estimating and under-estimating the probabilities with the same amount. Further more, for any two forecasts $\boldsymbol{f}_i = (p + \gamma_i, q - \gamma_i)$, $i = 1, 2$, with $|\gamma_1| < |\gamma_2|$, the Brier score would prefer the forecast corresponding to γ_1.

2.2 Logarithmic scoring rule

The logarithmic scoring rule was proposed by Good (1952). It was later termed Ignorance by Roulston and Smith (2002) when they introduced it to the meteorological community. Given a probabilistic forecast $\boldsymbol{f} = (f_1, f_2, \ldots, f_m)$, the logarithmic scoring

rule is given by $\mathrm{LS}(\boldsymbol{f}, j) = -\log f_j$, where j denotes the category that materialises. Let us consider the expected logarithmic score of the forecasting scheme $\boldsymbol{f} = (p + \gamma, q - \gamma)$:

$$\mathbb{E}[\mathrm{LS}(\boldsymbol{f}, J)] = -p\log(p + \gamma) - q\log(q - \gamma), \tag{1}$$

where $J \in \{1, 2\}$ is a random variable. The above expectation is also referred to as the Kullback-Leibler Information Criterion (Corradi and Swanson, 2006). As noted by Friedman (1983), this scoring rule is not effective.

If we let $\boldsymbol{f}_+ = (p + \gamma, q - \gamma)$ and $\boldsymbol{f}_- = (p - \gamma, q + \gamma)$, then we can define $\mathbb{E}[\mathrm{LS}]_\pm = \mathbb{E}[\mathrm{LS}(\boldsymbol{f}_+, J)] - \mathbb{E}[\mathrm{LS}(\boldsymbol{f}_-, J)]$. Then, assuming that $\gamma > 0$ without loss of generality,

$$\mathbb{E}[\mathrm{LS}]_\pm = p\log\left(\frac{p - \gamma}{p + \gamma}\right) + q\log\left(\frac{q + \gamma}{q - \gamma}\right) \tag{2}$$

Note that when $p = q = 0.5$, then $\mathbb{E}[\mathrm{LS}]_\pm = 0$, otherwise $\mathbb{E}[\mathrm{LS}]_\pm \neq 0$. Differentiating (2) with respect to γ yields

$$\frac{\mathrm{d}}{\mathrm{d}\gamma}\mathbb{E}[\mathrm{LS}]_\pm = \frac{2\gamma^2(p^2 - q^2)}{(p^2 - \gamma^2)(q^2 - \gamma^2)} \tag{3}$$

Expressions (2) and (3) are well defined provided $\gamma < \min(p, q)$.

$$\frac{\mathrm{d}}{\mathrm{d}\gamma}\mathbb{E}[\mathrm{LS}]_\pm > 0, \quad \text{if } p > q$$

$$\frac{\mathrm{d}}{\mathrm{d}\gamma}\mathbb{E}[\mathrm{LS}]_\pm < 0, \quad \text{if } p < q$$

It follows that $\mathbb{E}[\mathrm{LS}]_\pm > 0$ if $p > q$ and $\mathbb{E}[\mathrm{LS}]_\pm < 0$ if $p < q$. In other words, the logarithmic score penalises over confidence on the likely outcome and rewards erring on the side of caution. Given forecasting schemes that are equally calibrated, the logarithmic score will prefer the one with a higher entropy. To explain this further, let us denote the entropy of the forecast corresponding to γ by $h(\gamma)$, i.e.

$$h(\gamma) = -(p + \gamma)\log(p + \gamma) - (q - \gamma)\log(q - \gamma). \tag{4}$$

We now define the function $G(\gamma) = h(\gamma) - h(-\gamma)$ and claim that $G(\gamma) < 0$ for $0 < \gamma < q < p$. To prove this claim, we first note that $G(0) = 0$. It then suffices to show that $G'(0) < 0$. Note that

$$\begin{aligned} G'(\gamma) &= -\log(p + \gamma) + \log(q - \gamma) - \log(p - \gamma) + \log(q + \gamma) \\ &= -\log\left(\frac{p + \gamma}{q + \gamma}\right) + \log\left(\frac{q - \gamma}{p - \gamma}\right). \end{aligned}$$

The condition $p > q$ implies that $G'(\gamma) < 0$ for all $\gamma \in (0, q)$. Therefore, it is evident that, of the two forecasts, the logarithmic score prefers the one with a higher entropy. We have thus proved the following proposition:

Proposition 2.1 *Given two forecasts, $\boldsymbol{f}_+ = (p + \gamma, q - \gamma)$ and $\boldsymbol{f}_- = (p - \gamma, q + \gamma)$, where $0 < \gamma < q < p$, the logarithmic scoring rule prefers $\boldsymbol{f}_-$. Moreover, $\boldsymbol{f}_-$ has a higher entropy than $\boldsymbol{f}_+$.*

4

What about when there are two forecasts $\boldsymbol{f}_i = (p+\gamma_i, q-\gamma_i)$, $i = 1, 2$ with $0 < \gamma_1 < \gamma_2 < q$ and $p > q$? It is obvious that the Brier score will prefer $\boldsymbol{f}_1$ over $\boldsymbol{f}_2$. The question is, which of the two forecasts will the logarithmic scoring rule prefer? We answer this question by stating the following proposition:

Proposition 2.2 *Given two forecasts $\boldsymbol{f}_i = (p+\gamma_i, q-\gamma_i)$, $i = 1, 2$ with $0 < \gamma_1 < \gamma_2 < q$ and $p > q$, the logarithmic scoring rule prefers $\boldsymbol{f}_1$ over $\boldsymbol{f}_2$.*

Proof. In order to prove this proposition, it is sufficient to consider the expected logarithmic score of the forecast $\boldsymbol{f} = (p + \gamma, q - \gamma)$, which is given by equation (1). Differentiating the equation with respect to γ yields

$$\frac{\mathrm{d}}{\mathrm{d}\gamma}\mathbb{E}[\mathrm{LS}(\boldsymbol{f}, J)] = \frac{\gamma}{(p+\gamma)(q-\gamma)} \tag{5}$$

Equation (5) implies that, if $q > \gamma > 0$, $\mathbb{E}[\mathrm{LS}(\boldsymbol{f}, J)]$ is an increasing function of γ. Hence, the logarithmic scoring rule prefers the forecast $\boldsymbol{f}_1$.

On the other hand, if $\gamma < 0$ with $|\gamma| < p$, then equation (5) implies that $\mathbb{E}[\mathrm{LS}(\boldsymbol{f}, J)]$ is a decreasing function of γ. It then follows that, given $\gamma_2 < \gamma_1 < 0$ with $|\gamma_2| < p$, the logarithmic scoring rule will prefer the forecast $\boldsymbol{f}_1$. ■

Finally, let us consider the case of two forecasts $\boldsymbol{f}_1 = (p + \gamma_1, q - \gamma_1)$ and $\boldsymbol{f}_2 = (p - \gamma_2, q + \gamma_2)$, where $0 < \gamma_1 < \gamma_2 < q < p$. Again, it is clear that the Brier score will prefer the forecast $\boldsymbol{f}_1$ over $\boldsymbol{f}_2$. It remains to be seen which forecast the logarithmic scoring rule will prefer. This may be determined by considering the function $H(\gamma_1, \gamma_2)$, where

$$H(\gamma_1, \gamma_2) = p \log \left(\frac{p - \gamma_2}{p + \gamma_1}\right) + q \log \left(\frac{q + \gamma_2}{q - \gamma_1}\right) \tag{6}$$

Note that $H(\gamma_1, \gamma_2) = \mathbb{E}[\mathrm{LS}(\boldsymbol{f}_1, J)] - \mathbb{E}[\mathrm{LS}(\boldsymbol{f}_2, J)]$. The forecast $\boldsymbol{f}_1$ is preferred if $H(\gamma_1, \gamma_2) < 0$. The following proposition gives insights of relative forecast performance in the parameter space.

Proposition 2.3 *Given that $0 < \gamma_2 < q < p$, there exists $\gamma^* \in (0, \gamma_2)$ such that (a) $H(\gamma*, \gamma_2) = 0$, (b) $H(\gamma_1, \gamma_2) > 0$ for $\gamma_1 \in (\gamma^*, \gamma_2)$ and (c) $H(\gamma_1, \gamma_2) < 0$ for $\gamma_1 \in (0, \gamma^*)$.*

Before proving the above proposition, we remark that $H(\gamma_1, \gamma_2) < 0$ if and only if the logarithmic scoring rule prefers the forecast $\boldsymbol{f}_1$. This proposition implies that the logarithmic scoring rule and the Brier score prefer different forecasts when $\gamma_1 \in (\gamma^*, \gamma_2)$. Let us now consider the proof of this proposition.

Proof. In proving this proposition, it is useful to bear in mind that $H(\gamma_2, \gamma_2) > 0$. The partial derivatives of equation (6) are given by

$$\frac{\partial H}{\partial \gamma_1} = \frac{\gamma_1}{(p + \gamma_1)(q - \gamma_1)} \quad \text{and} \quad \frac{\partial H}{\partial \gamma_2} = \frac{-\gamma_2}{(p + \gamma_2)(q - \gamma_2)}. \tag{7}$$

Further more, we can differentiate equations (7) to obtain

$$\frac{\partial^2 H}{\partial \gamma_1^2} = \frac{pq + \gamma_1^2}{(p + \gamma_1)^2 (q - \gamma_1)^2} \quad \text{and} \quad \frac{\partial^2 H}{\partial \gamma_2^2} = \frac{-(pq + \gamma_2^2)}{(p - \gamma_2)^2 (q + \gamma_2)^2}. \tag{8}$$

It follows from equations (7) that $\partial H/\partial \gamma_1 = 0$ at $\gamma_1 = 0$ and $\partial H/\partial \gamma_2 = 0$ at $\gamma_2 = 0$. Since $\partial^2 H/\partial \gamma_1^2 > 0$ for all γ_1, $H(\gamma_1, \cdot)$ has a global minimum at $\gamma_1 = 0$. Similarly,

$H(\cdot, \gamma_2)$ has a global maximum at $\gamma_2 = 0$ since $\partial^2 H/\partial\gamma_2^2 < 0$ for all γ_2 and the first partial derivative with respect to γ_2 vanishes there. In particular, $H(0, \gamma_2) \leq H(0, 0) = 0$, i.e. $H(0, \gamma_2) \leq 0$. For $\gamma_2 > 0$, we have the strict inequality, $H(0, \gamma_2) < 0$. But we also have $H(\gamma_2, \gamma_2) > 0$ from Proposition 2.1. It, therefore, follows from the intermediate value theorem that $H(\gamma_1, \gamma_2) = 0$ for some $\gamma_1 = \gamma^* \in (0, \gamma_2)$, which completes the proof. ∎

Proposition 2.4 *For positive γ_1 and γ_2 such that $\gamma_1 < q < p$ and $\gamma_2 < p$, the entropy of the forecast $\boldsymbol{f}_1 = (p + \gamma_1, q - \gamma_1)$ is lower than that of the forecast $\boldsymbol{f}_2 = (p - \gamma_2, q + \gamma_2)$ whenever $\gamma_2 \leq (p - q)/2$.*

A consequence of this proposition is that the forecast corresponding to $\gamma_1 = \gamma^*$ is more informative than $\boldsymbol{f}_2$ provided $\gamma_2 \leq (p - q)/2$. Otherwise, either forecast could be more informative than the other. We now give the proof of this proposition.

Proof. To prove the above proposition, we consider the derivative of equation (4):

$$\frac{\mathrm{d}h}{\mathrm{d}\gamma} = -\log\left(\frac{p + \gamma}{q - \gamma}\right).$$

We then note that $\mathrm{d}h/\mathrm{d}\gamma < 0$ provided that $(p - q) > -2\gamma$. If $\gamma > 0$, this inequality is trivially satisfied. On the other hand, if $\gamma < 0$, then the inequality is satisfied provided $|\gamma| < (p - q)/2$. If $\gamma_2 < (p - q)/2$, then $h(\gamma)$ is a strictly decreasing function for all $\gamma \in [-\gamma_2, \gamma_2]$, which implies that $h(\gamma_1) > h(\gamma_2)$. If $\gamma_2 > (p - q)/2$, then $h(\gamma)$ is an increasing function for all $\gamma \in (-\gamma_2, -(p-q)/2)$ (provided $p > 3q$) and strictly decreasing function in $(-(p - q)/2, \gamma_1)$, which implies that $h(-(p - q)/2) > \max\{h(\gamma_1), h(-\gamma_2)\}$. Hence, in this case, we cannot determine which of $h(\gamma_1)$ and $h(-\gamma_2)$ is lower. ∎

2.3 The Spherical Scoring Rule

The *spherical scoring rule* is given by

$$S(\boldsymbol{f}, j) = -\frac{f_j}{||\boldsymbol{f}||_2}.$$

Define $\boldsymbol{f}_- = \boldsymbol{p} - \boldsymbol{\gamma}$ and $\boldsymbol{f}_+ = \boldsymbol{p} + \boldsymbol{\gamma}$. Which of the two forecasts $\boldsymbol{f}_-$ and $\boldsymbol{f}_+$ does the spherical scoring rule prefer? In order to address this question, we appeal to geometry. Considering $\boldsymbol{f} = \boldsymbol{p} + \boldsymbol{\gamma}$, the dot product rule yields

$$||\boldsymbol{p}||_2||\boldsymbol{f}||_2 \cos\theta = \langle\boldsymbol{f}, \boldsymbol{p}\rangle$$

where θ is the angle between $\boldsymbol{f}$ and $\boldsymbol{p}$. The above formula may be rewritten as

$$\cos\theta = \frac{||\boldsymbol{p}||_2^2 + \langle\boldsymbol{\gamma}, \boldsymbol{p}\rangle}{||\boldsymbol{p}||_2||\boldsymbol{f}||_2}. \tag{9}$$

We then state the following proposition:

Proposition 2.5 *If $\boldsymbol{p} = (p, q)$ and $\boldsymbol{\gamma} = (\gamma, -\gamma)$ and if we denote the right hand side of equation (9) by $C(\gamma)$, then*

$$\frac{\mathrm{d}C(\gamma)}{\mathrm{d}\gamma} = \frac{-\gamma}{||\boldsymbol{p}||_2||\boldsymbol{f}||_2^3}.$$

6

Proof. First note that $||\boldsymbol{f}||_2^2 = ||\boldsymbol{p}||_2^2 + 2\langle\boldsymbol{\gamma},\boldsymbol{p}\rangle + ||\boldsymbol{\gamma}||_2^2$ and

$$\frac{\mathrm{d}||\boldsymbol{f}||_2}{\mathrm{d}\gamma} = \frac{(p-q)+2\gamma}{||\boldsymbol{f}||_2}.$$

Using the quotient rule, we then differentiate $C(\gamma)$ with respect to γ to obtain

$$
\begin{aligned}
\frac{\mathrm{d}C(\gamma)}{\mathrm{d}\gamma} &= \frac{||\boldsymbol{f}||_2\frac{\mathrm{d}}{\mathrm{d}\gamma}(||\boldsymbol{p}||_2^2 + \langle\boldsymbol{\gamma},\boldsymbol{p}\rangle) - (||\boldsymbol{p}||_2^2 + \langle\boldsymbol{\gamma},\boldsymbol{p}\rangle)\frac{\mathrm{d}}{\mathrm{d}\gamma}||\boldsymbol{f}||_2}{||\boldsymbol{p}||_2||\boldsymbol{f}||_2^2} \\[2mm]
&= \frac{||\boldsymbol{f}||_2^2(p-q) - (||\boldsymbol{p}||_2^2 + \langle\boldsymbol{\gamma},\boldsymbol{p}\rangle)[(p-q)+2\gamma]}{||\boldsymbol{p}||_2||\boldsymbol{f}||_2^3} \\[2mm]
&= \frac{(p-q)(||\boldsymbol{p}||_2^+2\langle\boldsymbol{\gamma},\boldsymbol{p}\rangle + ||\boldsymbol{\gamma}||_2^2) - (||\boldsymbol{p}||_2^2 + \langle\boldsymbol{\gamma},\boldsymbol{p}\rangle)[(p-q)+2\gamma]}{||\boldsymbol{p}||_2||\boldsymbol{f}||_2^3} \\[2mm]
&= \frac{-2\gamma||\boldsymbol{p}||_2^2 + \gamma(p-q)^2}{||\boldsymbol{p}||_2||\boldsymbol{f}||_2^3} \\[2mm]
&= \frac{-2\gamma||\boldsymbol{p}||_2^2 + \gamma(||\boldsymbol{p}||_2^2 - 2pq)}{||\boldsymbol{p}||_2||\boldsymbol{f}||_2^3} \\[2mm]
&= \frac{-\gamma(||\boldsymbol{p}||_2^2 + 2pq)}{||\boldsymbol{p}||_2||\boldsymbol{f}||_2^3} \\[2mm]
&= \frac{-\gamma(p+q)^2}{||\boldsymbol{p}||_2||\boldsymbol{f}||_2^3}.
\end{aligned}
$$

The desired result follows from noting that $p+q=1$. ∎

Proposition 2.6 *Suppose that $p > q$ and $\gamma \in (0,q)$. Then the spherical scoring rule prefers the lower entropy forecast, $\boldsymbol{f}_+$, instead of $\boldsymbol{f}_-$.*

Proof. Since the spherical scoring rule is effective, it suffices for us to show that $d_*(\boldsymbol{f}_+,\boldsymbol{p}) < d_*(\boldsymbol{f},\boldsymbol{p})$. Suppose the angles that each of $\boldsymbol{f}_+$ and $\boldsymbol{f}_-$ makes with $\boldsymbol{p}$ are respectively θ_+ and θ_-. It is then true that $d_*(\boldsymbol{f}_+,\boldsymbol{p}) < d_*(\boldsymbol{f},\boldsymbol{p})$ if and only $\theta_+ < \theta_-$ since each distance is the length of a chord on a unit circle. Note that $C(0) = 1$ and $C'(0) = 0$. From Proposition 2.5, $-C'(\tau) < C'(-\tau)$ for all $\tau \in (0,\gamma)$, which implies that

$$
\begin{aligned}
-\int_0^{\gamma} C'(\tau)\mathrm{d}\tau < \int_0^{\gamma} C'(-\tau)\mathrm{d}\tau \quad &\Rightarrow \quad -\int_0^{\gamma} C'(\tau)\mathrm{d}\tau < -\int_0^{-\gamma} C'(\tau)\mathrm{d}\tau \\
&\Rightarrow \quad \int_0^{\gamma} C'(\tau)\mathrm{d}\tau > \int_0^{-\gamma} C'(\tau)\mathrm{d}\tau \\
&\Rightarrow \quad C(\tau)|_0^{\gamma} > C(\tau)|_0^{-\gamma} \\
&\Rightarrow \quad C(\gamma) - C(0) > C(-\gamma) - C(0) \\
&\Rightarrow \quad C(\gamma) > C(-\gamma).
\end{aligned}
$$

But $C(\gamma) > C(-\gamma)$ implies that $\theta_+ < \theta_-$. ∎

3   Density Forecasts

This section considers scoring rules for forecasts of continuous variables. It is in some sense a generalisation of the previous section. As before, we consider how each scoring

rule would rank two competing predictive distributions of fairly good quality. In the case of the logarithmic scoring rule and the Continuous Ranked Probability Score, we consider errors of each predictive distribution, $f(x)$, from the target distribution, $p(x)$, that are odd functions, i.e. $\gamma(x) = f(x) - p(x)$ with $\gamma(-x) = -\gamma(x)$.

Familiarity with the following notation and definitions will be useful. Given two functions $f(x)$ and $g(x)$ that are bounded, an *inner product* is defined by

$$\langle f, g \rangle = \int_{-\infty}^{\infty} f(x)g(x)\mathrm{d}x.$$

Then the L_2 norm is defined to be $||f||_2 = \langle f, f \rangle^{1/2}$.

3.1 The Quadratic scoring rule

A continuous counterpart of the Brier score is the quadratic score (Gneiting and Raftery, 2007), given by

$$QS(f, X) = ||f||_2^2 - 2f(X),$$

where X is a random variable. Taking the expectation yields

$$\mathbb{E}[QS(f, X)] = ||f - p||_2^2 - ||p||_2^2. \tag{10}$$

We can now write $f(x) = p(x) + \gamma(x)$, where $\int \gamma(x)\mathrm{d}x = 0$, and substitute it into (10) to obtain

$$\mathbb{E}[QS(f, X)] = ||\gamma||_2^2 - ||p||_2^2$$

As was the case with the Brier score, the functions $\pm\gamma(x)$ yield the same quadratic score. For any two forecasts, $f_i(x) = p(x) + \gamma_i(x)$, $i = 1, 2$ with $||\gamma_1||_2 < ||\gamma_2||_2$, the quadratic scoring rule would prefer $f_1(x)$. Further more, $||\gamma_1|| = ||\gamma_2||$ implies that $\mathbb{E}[QS(f_1, X)] = \mathbb{E}[QS(f_2, X)]$.

3.2 The Logarithmic scoring rule

The expectation of the logarithmic scoring rule for the forecast is

$$\mathbb{E}[LS(f, X))] = -\int p(x)\log(p(x) + \gamma(x))\mathrm{d}x.$$

As in the discrete case, we introduce the pdfs $f_+(x) = p(x) + \gamma(x)$, $f_-(x) = p(x) - \gamma(x)$ so that we can define $\mathbb{E}[LS]_{\pm} = \mathbb{E}[LS(f_+, X))] - \mathbb{E}[LS(f_-, X))]$. It follows that

$$\mathbb{E}[LS]_{\pm} = \int p(x)\log\left(\frac{p(x) - \gamma(x)}{p(x) + \gamma(x)}\right)\mathrm{d}x. \tag{11}$$

It is necessary that $|\gamma(x)| \leq p(x)$ for (11) to be well defined. Consider the case when $p(x) = p(-x)$. If, in addition, $\gamma(x)$ is an odd function, i.e. $\gamma(-x) = -\gamma(x)$, then equation (11) yields $\mathbb{E}[LS]_{\pm} = 0$.

When $\gamma(-x) = -\gamma(x)$ and $\int_{\infty}^{0} p(x)\mathrm{d}x > 0.5$, we state the following proposition:

Proposition 3.1 *Given that $\gamma(-x) = -\gamma(x)$ with $\gamma(|x|) < 0$ and $p(|x|) \leq p(x)$, then $\mathbb{E}[LS]_{\pm} \geq 0$.*

8

The above proposition gives conditions under which the forecast $f_-(x)$ is preferred by the logarithmic scoring rule over $f_+(x)$.

Proof. The proof proceeds as follows:

$$
\begin{aligned}
\mathbb{E}[\mathrm{LS}]_\pm &= \int_{-\infty}^{\infty} p(x) \log\left(\frac{p(x) - \gamma(x)}{p(x) + \gamma(x)}\right) \mathrm{d}x \\
&= \int_{-\infty}^{0} p(x) \log\left(\frac{p(x) - \gamma(x)}{p(x) + \gamma(x)}\right) \mathrm{d}x + \int_{0}^{\infty} p(x) \log\left(\frac{p(x) - \gamma(x)}{p(x) + \gamma(x)}\right) \mathrm{d}x
\end{aligned}
$$

If we now perform a change of variable $u = -x$ in the right hand integral and then replace u by x, we obtain

$$
\begin{aligned}
\mathbb{E}[\mathrm{LS}]_\pm &= \int_{-\infty}^{0} p(x) \log\left(\frac{p(x) - \gamma(x)}{p(x) + \gamma(x)}\right) \mathrm{d}x - \int_{0}^{-\infty} p(-x) \log\left(\frac{p(-x) - \gamma(-x)}{p(-x) + \gamma(-x)}\right) \mathrm{d}x \\
&= \int_{-\infty}^{0} p(x) \log\left(\frac{p(x) - \gamma(x)}{p(x) + \gamma(x)}\right) \mathrm{d}x + \int_{-\infty}^{0} p(-x) \log\left(\frac{p(-x) + \gamma(x)}{p(-x) - \gamma(x)}\right) \mathrm{d}x \\
&\geq \int_{-\infty}^{0} p(x) \log\left(\frac{p(x) - \gamma(x)}{p(x) + \gamma(x)}\right) \mathrm{d}x + \int_{-\infty}^{0} p(x) \log\left(\frac{p(x) + \gamma(x)}{p(x) - \gamma(x)}\right) \mathrm{d}x = 0,
\end{aligned}
$$

where we used $p(|x|) \leq p(x)$ to obtain the last inequality. To justify the use of this inequality, we need to show that the function

$$
\Phi(p) = p \log\left(\frac{p + \gamma}{p - \gamma}\right)
$$

is a decreasing function for $\gamma \in (0, p)$. Differentiating Φ with respect to p yields

$$
\Phi'(p) = \log\left(\frac{p + \gamma}{p - \gamma}\right) - \frac{2p\gamma}{p^2 - \gamma^2}.
$$

It now suffices to show that $\Phi'(p) < 0$ for all p. Let us introduce the notation $W(p) = \log[(p + \gamma)/(p - \gamma)]$ and $Y(p) = 2p\gamma/(p^2 - \gamma^2)$ so that $\Phi'(p) = W(p) - Y(p)$. Note that $W(2\gamma) = \log 2$ and $Y(2\gamma) = 4/3 = \log(e^{4/3})$. Hence $W(2\gamma) < Y(2\gamma)$, which implies that $\Phi'(2\gamma) < 0$. Differentiating $W(p)$ and $Y(p)$ with respect to p yields

$$
W'(p) = \frac{-2\gamma}{p^2 - \gamma^2} \quad \text{and} \quad Y'(p) = \frac{-2\gamma(p^2 + \gamma^2)}{(p^2 - \gamma^2)^2}.
$$

It is now clear that $W'(p) < 0$ and $Y'(p) < 0$ for all p. Further more, $Y'(p) < W'(p)$. Hence $W(p) < Y(p)$ for all $p \in (\gamma, 2\gamma]$, which implies that $\Phi'(p) < 0$ for all $p \in (\gamma, 2\gamma]$. It now remains to be shown that $\Phi'(p) < 0$ for all $p \in (2\gamma, \infty)$. It suffices to consider the asymptotic behaviour as $p \to \infty$. Applying L'Hopital's rule, we obtain

$$
\lim_{p \to \infty} \frac{|W(p)|}{|Y(p)|} = \lim_{p \to \infty} \frac{|W'(p)|}{|Y'(p)|} = \lim_{p \to \infty} \frac{p^2 + \gamma^2}{p^2 - \gamma^2} = 1.
$$

Hence, $\lim_{p \to \infty} W(p) = \lim_{p \to \infty} Y(p)$, i.e. $W(\infty) = Y(\infty)$. With this result in mind, for all $p \in [2\gamma, \infty)$, we have

$$
\begin{aligned}
\int_{p}^{\infty} Y'(\tau)\mathrm{d}\tau < \int_{p}^{\infty} W'(\tau)\mathrm{d}\tau &\Rightarrow Y(\tau)|_{p}^{\infty} < W(\tau)|_{p}^{\infty} \\
&\Rightarrow Y(\infty) - Y(p) < W(\infty) - W(p) \\
&\Rightarrow W(p) < Y(p),
\end{aligned}
$$

which completes the proof. The condition $p(|x|) < p(x)$ implies that $\int_{-\infty}^{0} p(x)\mathrm{d}x \geq 0.5$. It corresponds to the case $p > q$ in the discrete case. $\blacksquare$

We now want to compare the entropies of the forecasts $f(x) = p(x) \pm \gamma(x)$ when $\gamma(-x) = -\gamma(x)$ and $\gamma(|x|) \leq \gamma(x)$. The entropy of the function $f(x) = p(x) + \gamma(x)$ is then given by

$$h(\gamma) = -\int (p(x) + \gamma(x)) \log(p(x) + \gamma(x))\mathrm{d}x. \tag{12}$$

The functional derivative of $h(\gamma)$ with respect to γ is then given by

$$
\begin{aligned}
\frac{\delta h(\gamma)}{\delta \gamma(x)} &= -\frac{\partial}{\partial \gamma(x)} \left\{ (p(x) + \gamma(x)) \log(p(x) + \gamma(x)) \right\} \\
&= -[\log(p + \gamma) + 1]. \tag{13}
\end{aligned}
$$

The order $O(\varepsilon)$ part of $h(\gamma + \varepsilon \delta\gamma) - h(\gamma)$ is given by (see Stone and Goldbart (2008) for further insights)

$$\delta h(\gamma) = \int_{-\infty}^{\infty} \frac{\delta h(\gamma)}{\delta \gamma(x)} \delta\gamma(x)\mathrm{d}x. \tag{14}$$

Plugging (13) into (14) yields

$$
\begin{aligned}
\delta h(\gamma) &= -\int_{-\infty}^{\infty} [\log(p(x) + \gamma(x)) + 1]\delta\gamma(x)\mathrm{d}x \\
&= -\int_{-\infty}^{0} [\log(p(x) + \gamma(x)) + 1]\delta\gamma(x)\mathrm{d}x - \int_{0}^{\infty} [\log(p(x) + \gamma(x)) + 1]\delta\gamma(x)\mathrm{d}x \\
&= -\int_{-\infty}^{0} [\log(p(x) + \gamma(x)) + 1]\delta\gamma(x)\mathrm{d}x + \int_{0}^{-\infty} [\log(p(-x) + \gamma(-x)) + 1]\delta\gamma(-x)\mathrm{d}x \\
&= -\int_{-\infty}^{0} [\log(p(x) + \gamma(x)) + 1]\delta\gamma(x)\mathrm{d}x - \int_{-\infty}^{0} [\log(p(-x) + \gamma(-x)) + 1]\delta\gamma(-x)\mathrm{d}x \\
&= -\int_{-\infty}^{0} [\log(p(x) + \gamma(x)) + 1]\delta\gamma(x)\mathrm{d}x + \int_{-\infty}^{0} [\log(p(-x) - \gamma(x)) + 1]\delta\gamma(x)\mathrm{d}x \\
&= -\int_{-\infty}^{0} \log\left(\frac{p(x) + \gamma(x)}{p(-x) - \gamma(x)} \right) \delta\gamma(x)\mathrm{d}x,
\end{aligned}
$$

where we have applied a change of variable $x \to -x$ in the second integral of the third line and assumed $\delta\gamma(-x) = -\delta\gamma(x)$ in the fifth line. In particular,

$$\delta h(\gamma)|_{\gamma=0} = -\int_{-\infty}^{0} \log\left(\frac{p(x)}{p(-x)} \right) \delta\gamma(x)\mathrm{d}x.$$

Using the assumption that $p(x) \geq p(-x)$ whenever $x < 0$, we consequently obtain

$$\delta h(\gamma)|_{\gamma=0} \leq 0, \tag{15}$$

if $\delta\gamma(x) > 0$ for all $x < 0$. In effect, we have just proved the following proposition:

Proposition 3.2 *Given that $\gamma(-x) = -\gamma(x)$, $\int \gamma(x)\mathrm{d}x = 0$, $\gamma(|x|) \leq 0$, $p(|x|) \leq p(x)$ and $|\gamma(x)| < p(x)$, then the entropy of the forecast density $f_+(x) = p(x) + \gamma(x)$ is lower than that of the forecast density $f_-(x) = p(x) - \gamma(x)$.*

10

Propositions 3.1 and 3.2 imply that the logarithmic scoring rule prefers the forecast density that is less informative, which is in agreement with the categorical case considered in the previous section.

Proposition 3.3 *Given two forecasts $f_i(x) = p(x) + \gamma_i(x)$, $i = 1, 2$, with (i) $|\gamma_1(x)| < |\gamma_2(x)|$, (ii) $\gamma_i(|x|) \leq 0$, (iii) $\gamma_i(-x) = -\gamma_i(x)$, (iv) $|\gamma_i(x)| \leq p(x)$ and (v) $p(|x|) \leq p(x)$, then the logarithmic scoring rule prefers forecast $f_1(x)$ over forecast $f_2(x)$.*

Proof. To prove the above proposition, we consider the functional derivative of the expected logarithmic scoring rule, $\mathbb{E}[\text{LS}] = -\int_{-\infty}^{\infty} p(x) \log(p(x) + \gamma(x)) \mathrm{d}x$. The functional derivative with respect to $\gamma(x)$ is

$$\frac{\delta}{\delta\gamma} \mathbb{E}[\text{LS}] = \frac{-p(x)}{p(x) + \gamma(x)}.$$

Using this result, we obtain the first variation of $\mathbb{E}[\text{LS}]$ as

$$
\begin{aligned}
\delta\mathbb{E}[\text{LS}] &= \int_{-\infty}^{\infty} \frac{\delta\mathbb{E}[\text{LS}]}{\delta\gamma(x)} \delta\gamma(x) \mathrm{d}x \\
&= \int_{-\infty}^{\infty} \frac{-p(x)}{p(x) + \gamma(x)} \delta\gamma(x) \mathrm{d}x \\
&= \int_{-\infty}^{0} \frac{-p(x)}{p(x) + \gamma(x)} \delta\gamma(x) \mathrm{d}x + \int_{0}^{\infty} \frac{-p(x)}{p(x) + \gamma(x)} \delta\gamma(x) \mathrm{d}x \\
&= \int_{-\infty}^{0} \frac{-p(x)}{p(x) + \gamma(x)} \delta\gamma(x) \mathrm{d}x + \int_{0}^{-\infty} \frac{p(-x)}{p(-x) + \gamma(-x)} \delta\gamma(-x) \mathrm{d}x \\
&= \int_{-\infty}^{0} \frac{-p(x)}{p(x) + \gamma(x)} \delta\gamma(x) \mathrm{d}x + \int_{-\infty}^{0} \frac{p(-x)}{p(-x) - \gamma(x)} \delta\gamma(x) \mathrm{d}x \\
&= \int_{-\infty}^{0} \left[\frac{p(-x)}{p(-x) - \gamma(x)} - \frac{p(x)}{p(x) + \gamma(x)} \right] \delta\gamma(x) \mathrm{d}x \\
&= \int_{-\infty}^{0} \frac{[p(-x) + p(x)]\gamma(x)}{[p(-x) - \gamma(x)][p(x) + \gamma(x)]} \delta\gamma(x) \mathrm{d}x \\
&\geq 0,
\end{aligned}
$$

provided $\delta\gamma(x) > 0$ for all $x < 0$, $\delta\gamma(-x) = -\delta\gamma(x)$, $\gamma(-x) = -\gamma(x)$ and $\gamma(|x|) \leq 0$. What been shown is that as $\gamma(x)$ changes by $\delta\gamma(x)$, the expected logarithmic score changes by a positive amount. In particular, if we start at $\gamma(x) = \gamma_1(x)$, and progressively move towards $\gamma(x) = \gamma_2(x)$ by making successive additions of $\delta\gamma(x)$, the expected logarithmic score can only increase. Hence the expected logarithmic score of $\gamma_2(x)$ will be higher than that of $\gamma_2(x)$, which yields the result. ∎

We shall now consider two forecasts, $f_1(x) = p(x) + \gamma_1(x)$ and $f_2(x) = p(x) - \gamma_2(x)$ with $|\gamma_1(x)| \leq |\gamma_2(x)| \leq p(x)$. In this case, the quadratic scoring rule would prefer $f_1(x)$ over $f_2(x)$. In order to determine which forecast the logarithmic scoring would prefer, we consider the functional

$$\mathcal{H}(\gamma_1, \gamma_2) = \int_{-\infty}^{\infty} p(x) \log\left(\frac{p(x) - \gamma_2(x)}{p(x) + \gamma_1(x)} \right) \mathrm{d}x. \tag{16}$$

Then the following proposition holds

Proposition 3.4 *Given that $|\gamma_1(x)| \leq |\gamma_2(x)| \leq p(x)$ and $\gamma_i(-x) = -\gamma_i(x)$, $i = 1,2$, there exists $\gamma*(x)$ satisfying the inequalities $\gamma^*(x)\gamma_2(x) \geq 0$ and $|\gamma^*(x)| \leq |\gamma_2(x)|$ such that (a) $\mathcal{H}(\gamma^*, \gamma_2) = 0$, (b) $\mathcal{H}(\gamma_1, \gamma_2) > 0$ for $|\gamma^*| < |\gamma_1|$ and (c) $\mathcal{H}(\gamma_1, \gamma_2) < 0$ for $|\gamma^*| > |\gamma_1|$.*

Proof. It is helpful to first note that Proposition 3.1 implies that $\mathcal{H}(\gamma_2, \gamma_2) > 0$ when $\gamma_2 \neq 0$. Thinking of $\gamma_1(x)$ as fixed, the first variation of $\mathcal{H}(\cdot, \gamma_2)$ with respect to $\gamma_2(x)$ is given by

$$\delta\mathcal{H}(\cdot, \gamma_2) = \int_{-\infty}^{\infty} \frac{\delta\mathcal{H}(\cdot, \gamma_2)}{\delta\gamma_2(x)}\delta\gamma_2(x)\mathrm{d}x$$

$$= \int_{-\infty}^{\infty} \frac{-p(x)}{p(x) - \gamma_2(x)}\delta\gamma_2(x)\mathrm{d}x$$

$$= \int_{-\infty}^{0} \frac{-p(x)}{p(x) - \gamma_2(x)}\delta\gamma_2(x)\mathrm{d}x + \int_{0}^{\infty} \frac{-p(x)}{p(x) - \gamma_2(x)}\delta\gamma_2(x)\mathrm{d}x$$

$$= \int_{-\infty}^{0} \frac{-p(x)}{p(x) - \gamma_2(x)}\delta\gamma_2(x)\mathrm{d}x - \int_{0}^{-\infty} \frac{p(-x)}{p(-x) + \gamma_2(x)}\delta\gamma_2(x)\mathrm{d}x$$

$$= \int_{-\infty}^{0} \frac{-p(x)}{p(x) - \gamma_2(x)}\delta\gamma_2(x)\mathrm{d}x + \int_{-\infty}^{0} \frac{p(-x)}{p(-x) + \gamma_2(x)}\delta\gamma_2(x)\mathrm{d}x$$

$$= \int_{-\infty}^{0} \left[\frac{p(-x)}{p(-x) + \gamma_2(x)} - \frac{p(x)}{p(x) - \gamma_2(x)} \right] \delta\gamma_2(x)\mathrm{d}x$$

$$= \int_{-\infty}^{0} \frac{-[p(-x) + p(x)]\gamma_2(x)}{[p(-x) + \gamma_2(x)][p(x) - \gamma_2(x)]}\delta\gamma_2(x)\mathrm{d}x$$

$$\leq 0,$$

provided $\delta\gamma_2 > 0$ and $\delta\gamma_2(-x) = -\delta\gamma_2(x)$. In the fourth line a change of variable $x = -\tau$ was applied and then τ was replaced with x since it is a dummy variable. It follows that $\mathcal{H}(\cdot, \gamma_2)$ has a maximum when $\gamma_2 = 0$, i.e. $\mathcal{H}(\cdot, \gamma_2) \leq \mathcal{H}(\cdot, 0)$. In particular, $\mathcal{H}(0, \gamma_2) \leq \mathcal{H}(0, 0) = 0$. For $\gamma_2 \neq 0$, we have the strict inequality, $\mathcal{H}(0, \gamma_2) < 0$. Since $\mathcal{H}(\gamma_2, \gamma_2) > 0$, continuity implies that $\mathcal{H}(\gamma_1, \gamma_2) = 0$ for some $\gamma_1(x) = \gamma^*(x)$ such that $|\gamma^*| < |\gamma_2|$, and this completes the proof. ∎

3.3   The Spherical Scoring Rule

Given a forecast $f(x)$, the spherical scoring rule is given by

$$S(f, X) = -\frac{f(X)}{||f||_2}.$$

If we define the operator $\rho f = f(x)/||f||_2$, the expected spherical score is the inner product

$$\mathbb{E}[S(f, X)] = -\langle \rho f, p \rangle.$$

The minimum of this expectation is achieved if and only $f = p$ since it is a strictly proper scoring rule (Friedman, 1983). We now state the following proposition:

Proposition 3.5 *Given that $\gamma(-x) = -\gamma(x)$ with $\gamma(|x|) < 0$, $|\gamma(x)| < p(x)$ and $p(|x|) \leq p(x)$, then the spherical scoring rule prefers the forecast $f_+(x)$ over $f_-(x)$, i.e. $\mathbb{E}[S(f_+, X)] \leq \mathbb{E}[S(f_-, X)]$.*

Proof. The aim here is to show that $\mathbb{E}[S(f_+, X)] \leq \mathbb{E}[S(f_-, X)]$, which is equivalent to $\langle \rho f_+, p \rangle \geq \langle \rho f_-, p \rangle$. Note that each of these inner products is non-negative since

$$
\begin{aligned}
\langle \rho f_\pm, p \rangle &= \frac{\langle f_\pm, p \rangle}{||f_\pm||_2} \\
&= \frac{\langle p \pm \gamma, p \rangle}{||f_\pm||_2} \\
&= \frac{||p||_2^2 \pm \langle \gamma, p \rangle}{||f_\pm||_2} \\
&\geq 0,
\end{aligned}
$$

due to Cauchy Schwartz's inequality, $\langle \pm\gamma, p \rangle \leq ||\gamma||_2 ||p||_2$, and the hypothesis, $|\gamma(x)| \leq p(x) \Rightarrow ||\gamma||_2 \leq ||p||_2$. Therefore, $\langle \rho f_+, p \rangle \geq \langle \rho f_-, p \rangle$ is equivalent to $\langle \rho f_+, p \rangle^2 \geq \langle \rho f_-, p \rangle^2$. It therefore suffices to show that the latter inequality holds.

$$
\begin{aligned}
\langle \rho f_+, p \rangle^2 - \langle \rho f_-, p \rangle^2 &= \frac{\langle f_+, p \rangle^2}{||f_+||_2^2} - \frac{\langle f_-, p \rangle^2}{||f_-||_2^2} \\
&= \frac{\langle p + \gamma, p \rangle^2}{||f_+||_2^2} - \frac{\langle p - \gamma, p \rangle^2}{||f_-||_2^2} \\
&= \frac{(||p||_2^2 + \langle \gamma, p \rangle)^2}{||f_+||_2^2} - \frac{(||p||_2^2 - \langle \gamma, p \rangle)^2}{||f_-||_2^2} \\
&= \frac{||f_-||_2^2 (||p||_2^2 + \langle \gamma, p \rangle)^2 - ||f_+||_2^2 (||p||_2^2 - \langle \gamma, p \rangle)^2}{||f_+||_2^2 ||f_-||_2^2}.
\end{aligned}
$$

Plugging in $||f_+||_2^2 = ||p||_2^2 + 2\langle \gamma, p \rangle + ||\gamma||_2^2$ and $||f_-||_2^2 = ||p||_2^2 - 2\langle \gamma, p \rangle + ||\gamma||_2^2$ into the numerator of the last expression, removing brackets and collecting like terms yield

$$
\langle \rho f_+, p \rangle^2 - \langle \rho f_-, p \rangle^2 = \frac{\langle \gamma, p \rangle (||p||_2^2 ||\gamma||_2^2 - \langle \gamma, p \rangle^2)}{||f_+||_2^2 ||f_-||_2^2}. \tag{17}
$$

As a consequence of Cauchy-Schwartz's inequality, $||p||_2^2 ||\gamma||_2^2 - \langle \gamma, p \rangle^2 \geq 0$. It will now be shown that under the hypothesis of the proposition, $\langle \gamma, p \rangle \geq 0$.

$$
\begin{aligned}
\langle \gamma, p \rangle &= \int_{-\infty}^{\infty} \gamma(x) p(x) \mathrm{d}x \\
&= \int_{-\infty}^{0} \gamma(x) p(x) \mathrm{d}x + \int_{0}^{\infty} \gamma(x) p(x) \mathrm{d}x \\
&= \int_{-\infty}^{0} \gamma(x) p(x) \mathrm{d}x - \int_{0}^{-\infty} \gamma(-x) p(-x) \mathrm{d}x \\
&= \int_{-\infty}^{0} \gamma(x) p(x) \mathrm{d}x + \int_{-\infty}^{0} \gamma(-x) p(-x) \mathrm{d}x \\
&= \int_{-\infty}^{0} \gamma(x) p(x) \mathrm{d}x - \int_{-\infty}^{0} \gamma(x) p(-x) \mathrm{d}x \\
&= \int_{-\infty}^{0} \gamma(x) [p(x) - p(-x)] \mathrm{d}x \\
&\geq 0,
\end{aligned}
$$

since $p(x) \geq p(-x)$ and $\gamma(x) \geq 0$ for all $x \leq 0$. Hence, the right hand side of equation (17) is non-negative. $\blacksquare$

The distribution preferred by the spherical scoring rule is already known through Proposition 3.2 to be of lower entropy. As was the case in the binary case, the spherical scoring rule prefers an opposite distribution to the logarithmic scoring rule.

3.4 Continuous Ranked Probability Score

Finally, we consider the *Continuous Ranked Probability Score* (CRPS) of the density forecast $f(x)$ whose cumulative distribution is $F(x)$. The CRPS is a function of F and the verification X and is defined by (Gneiting and Raftery, 2007)

$$\mathrm{CRPS}(F, X) = \int_{-\infty}^{\infty} (F(\tau) - \mathbb{I}\{\tau \geq X\})^2 \mathrm{d}\tau.$$

The above score may equivalently be written as

$$\mathrm{CRPS}(F, X) = \int_{-\infty}^{X} F^2(\tau) \mathrm{d}\tau + \int_{X}^{\infty} (F(\tau) - 1)^2 \mathrm{d}\tau. \tag{18}$$

It follows from (18) that

$$\mathbb{E}[\mathrm{CRPS}(F, X)] = \int_{-\infty}^{\infty} p(x) \int_{-\infty}^{x} F^2(\tau) \mathrm{d}\tau \mathrm{d}x + \int_{-\infty}^{\infty} p(x) \int_{x}^{\infty} (F(\tau) - 1)^2 \mathrm{d}\tau \mathrm{d}x, \tag{19}$$

where $p(x)$ is the true (or target) density function. If $P(x) = \int_{-\infty}^{x} p(\tau) \mathrm{d}\tau$, we can then apply the integration by parts formula to each term on the right hand side of (19) to obtain

$$\int_{-\infty}^{\infty} p(x) \int_{-\infty}^{x} F^2(\tau) \mathrm{d}\tau \mathrm{d}x = P(x) \int_{-\infty}^{x} F^2(\tau) \mathrm{d}\tau \Big|_{-\infty}^{\infty} - \int_{-\infty}^{\infty} P(x) F^2(x) \mathrm{d}x$$

$$= \int_{-\infty}^{\infty} F^2(x) \mathrm{d}x - \int_{-\infty}^{\infty} P(x) F^2(x) \mathrm{d}x$$

and

$$\int_{-\infty}^{\infty} p(x) \int_{x}^{\infty} (F(\tau) - 1) \mathrm{d}\tau \mathrm{d}x = P(x) \int_{x}^{\infty} (F(\tau) - 1)^2 \mathrm{d}\tau \Big|_{-\infty}^{\infty} + \int_{-\infty}^{\infty} P(x)(F(x) - 1)^2 \mathrm{d}x$$

$$= 0 + \int_{-\infty}^{\infty} P(x)(F(x) - 1)^2 \mathrm{d}x,$$

whence

$$\mathbb{E}[\mathrm{CRPS}(F, X)] = \int_{-\infty}^{\infty} P(x)(1 - P(x)) \mathrm{d}x + \int_{-\infty}^{\infty} (F(x) - P(x))^2 \mathrm{d}x, \tag{20}$$

after some algebraic manipulation. Define $F(x) = P(x) + \Gamma(x)$, where $\Gamma(x) = \int_{-\infty}^{x} \gamma(\tau) \mathrm{d}\tau$ and $f(x) = p(x) + \gamma(x)$. If we also define $G(P) = \int P(x)(1 - P(x)) \mathrm{d}x$, then equation (20) can be re-written as $\mathbb{E}[\mathrm{CRPS}(F, X)] = G(P) + ||\Gamma||_2^2$. We have thus proved the following proposition:

Proposition 3.6 *The Continuous Ranked Probability Score does not distinguish between distributions whose cumulative errors from the target distribution are equal in the sense of the L^2 norm.*

Consequently, the CRPS does not distinguish between density forecasts whose errors from the target density differ by the sign. To see this, consider two forecasts which whose errors from the target density are $\gamma_1(x)$ and $\gamma_2(x)$ respectively, with $\gamma_1(x) = -\gamma_2(x)$. It then follows that

$$
\begin{aligned}
||\Gamma_1||_2^2 - ||\Gamma_2||_2^2 &= \int_{-\infty}^{\infty} \Gamma_1^2(x)\mathrm{d}x - \int_{-\infty}^{\infty} \Gamma_2^2(x)\mathrm{d}x \\
&= \int_{-\infty}^{\infty} \left(\Gamma_1^2(x) - \Gamma_2^2(x)\right)\mathrm{d}x \\
&= \int_{-\infty}^{\infty} \left(\Gamma_1(x) + \Gamma_2(x)\right)\left(\Gamma_1(x) - \Gamma_2(x)\right)\mathrm{d}x \\
&= 0,
\end{aligned}
$$

since $\gamma_1(x) = -\gamma_2(x) \Rightarrow \Gamma_1(x) + \Gamma_2(x) = 0$.

As a final remark, we note that the second term in the expectation of CRPS in (20) somewhat resembles the *mean squared error criterion* discussed in Corradi and Swanson (2006). The mean squared error of the forecast $F(x)$ is $\mathbb{E}[\Gamma^2(X)] = \int p(x)\Gamma^2(x)\mathrm{d}x$. Likewise, the mean squared error criterion does not distinguish between forecasts whose errors from the target density differ by a sign (i.e. $\gamma_1(x) = -\gamma_2(x)$) because

$$
\begin{aligned}
\mathbb{E}[\Gamma_1^2(X)] - \mathbb{E}[\Gamma_2^2(X)] &= \int_{-\infty}^{\infty} p(x)\left(\Gamma_1^2(x) - \Gamma_2^2(x)\right)\mathrm{d}x \\
&= \int_{-\infty}^{\infty} p(x)\left(\Gamma_1(x) + \Gamma_2(x)\right)\left(\Gamma_1(x) - \Gamma_2(x)\right)\mathrm{d}x \\
&= 0.
\end{aligned}
$$

4   Discussion and Conclusions

This manuscript contrasted how certain scoring rules would rank competing forecasts of specified departures from the target distribution. We considered the Brier Score, the logarithmic scoring rule and the spherical scoring rule, in categorical case, focusing on the binary case. In the continuous case, we considered the Quadratic Score, Logarithmic scoring rule, spherical scoring rule and the Continuous Ranked Probability Score (CRPS). Given two forecasts whose errors from the target distribution differ only by the sign, we found that the logarithmic scoring rule prefers the higher entropy distribution whilst the spherical scoring rule prefers the lower entropy distribution: bear in mind that higher entropy corresponds to more uncertainty (Shannon, 1948). Both the Brier score and the Quadratic Score do not distinguish two distributions with equal L^2 norms of their errors from the target distribution. The logarithmic scoring rule selects a lower entropy forecast only if it is nearer to the target distribution in the sense of the L^2 norm and vice versa for the spherical scoring rule. The CRPS is indifferent to forecasts whose errors from the target density differ by a sign.

Some have criticised the logarithmic scoring rule for placing a heavy penalty on assigning zero probability to events that materialise (e.g. Boero *et al.*, 2011; Gneiting and Raftery, 2007); but assigning zero probability to events that are possible is also discouraged by *Laplace's rule of succession* (Jaynes, 2003). What has been shown here is that the logarithmic scoring rule is good at highlighting forecasts that are less uncertain than ideal forecasts. Such forecasts may have to be dealt with appropriately. One way of

dealing with such forecasts is to blend with the unconditional distribution as discussed in Machete (2012). Nonetheless, given two density forecasts, the logarithmic scoring rule does not just reject the more extreme in the sense of entropy: If both forecasts are more uncertain than the ideal forecast, the logarithmic scoring rule will tend to prefer the less uncertain of the two.

Note that in the continuous case, we made the assumption that the true distribution satisfies $p(|x|) \leq p(x)$, which mean for every negative x, $p(-x) \leq p(x)$. Further research would be required to determine the effect of relaxing this assumption. We conjecture that the general result should hold when this assumption is relaxed to saying the density should have more mass overall in the negative side than in the positive side (or vice versa), but not necessarily point by point. The assumptions mean that the case of skewed distributions is accounted for, which is useful because many distributions are skewed. Looking beyond the scoring rules considered, we envisage that scoring rules may be split into three categories: those that tend to favour more uncertain forecasts, those that favour less uncertain forecasts and those that are indifferent to either option. Discriminating between scoring rules within each of the three sets would then be an issue to confront.

Does our consideration of departures from ideal forecasts amount to advocating for dishonesty by forecasters? Not at all. We are merely making an observation that forecasters can honestly report predictive distributions that have departures from ideal forecasts. Although strictly proper scoring rules encourage forecasters to be honest when they report their best judgements, they do not guarantee that the reported forecasts will coincide with ideal forecasts. Our point then is that using a given scoring rule may inherently favour departures from ideal forecasts in one direction more than in another. Therefore, when one selects a scoring rule to estimate distribution parameters or choose between two competing experts, it amounts to deciding preferred departures.

Which scoring rule one should choose will depend on the application at hand. Combining insights of scoring rules set forth in this paper with an understanding of the situation at hand can help decide which scoring rule is most appropriate. The issue to consider may be decisions associated with high impact, low probability events. To illustrate our point, let us consider inflation forecasting. It is undesirable to over estimate the probability for extreme inflation because of the panic it can create as buyers rush to spend now before prices rise. In order to manage peoples expectations better, the spherical scoring rule is preferable in this case. As another example, consider seasonal forecasts of drought in the UK, which is arguably a rare event. Under estimating the probability of this event could result in water shortages since water companies might not be stringent on water usage. In this case, the logarithmic scoring rule is preferable.

Johnstone *et al.* (2011) made a good effort of proposing scoring tules tailored to "the decision problem and the utility function of the decision maker". The scoring rules they proposed depend linearly on the utility function. Whilst taloring scoring rules to decision problems is a step in the right direction, an acknowledged limitation is when the forecasts are issued to multiple users who may not even have a utility function. Our work extends to situations where the assumptions made by Johnstone *et al.* (2011) are not valid.

Acknowledgements

References

Bickel JE, 2007. Some Comparisons among Quadratic, Spherical, and Logarithmic Scoring Rules. *Decision Analysis* **4**:49–65.

Boero G, Smith J, Wallis KF, 2011. Scoring rules and survey density forecasts. *International Journal of Forecasting* **27**:379–393.

Brier GW, 1950. Verification of forecasts expressed in terms of probability. *Monthly Weather Review* **78**:1–3.

Brocker J, Smith LA, 2007. Scoring Probabilistic Forecasts: The importance of Being proper. *Weather and Forecasting* **22**:382–388.

Corradi V, Swanson NR, 2006. Predictive density evaluation. In Elliott G, Granger CWJ, Timmermann A (editors), *Handbook for Econometric Forecasting*, volume 1, 197–284. North-Holland.

Epstein ES, 1969. A Scoring System for Probability Forecasts of Ranked Categories. *Journal of Applied Meteorology* **8**:985–987.

Friedman D, 1983. Effective Scoring Rules for Probabilistic Forecasts. *Management Science* **29**:447–454.

Gneiting T, Balabdaoui F, Raftery AE, 2007. Probabilistic forecasts, calibration and sharpness. *J. R. Statist. Soc. B* **69**:243–268.

Gneiting T, Raftery AE, 2007. Strictly proper scoring rules, prediction and estimation. *J. Amer. Stat. Ass.* **102**:359–378.

Good IJ, 1952. Rational decisions. *Journal of the Royal Statistical Society. Series B (Methodological)* **14**:107–114.

Jaynes ET, 2003. *Probability Theory: The Logic of Science*. Cambridge University Press.

Johnstone D, Lin Y, 2011. Fitting probability forecasting models by scoring rules and maximum likelihood. *Journal of Statistical Planning and Inference* **141**:1832–1837.

Johnstone DJ, Jose VRR, Winkler RL, 2011. Tailored Scoring Rules for Probabilities. *Decision Analysis* **8**:256–268.

Jose VRR, Nau RF, Winkler RL, 2008. Scoring Rules, Generalized Entropy, and Utility Maximisation. *Operations Research* **56**:1146–1157.

Kahneman D, Tversky A, 1979. Prospect Theory: An analysis of decision under risk. *Economterica* **47**:263–291.

Machete RL, 2012. Early warning with calibrated and sharper probabilistic forecasts. *Journal of Forecasting* doi:10.1002/for.2242.

Nau RF, 1985. Should scoring rules be éffective? *Management Science* **31**:527–535.

Roulston MS, Smith LA, 2002. Evaluating Probabilistic Forecasts Using Information Theory. *Monthly Weather Review* **130**:1653–1660.

Savage LJ, 1971. Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association* **66**:783–801.

Shannon CE, 1948. A Mathematical theory of communication. *Bell Systems Technology Journal* **27**:379–423,623–656.

Stone M, Goldbart P, 2008. *Mathematics for Physics I*. PIMANDER-CASAUBON.