

NECESSARY AND SUFFICIENT CONDITIONS FOR DOMINATION RESULTS FOR PROPER SCORING RULES

ALEXANDER R. PRUSS

Department of Philosophy, Baylor University

Abstract. Scoring rules measure the deviation between a forecast, which assigns degrees of confidence to various events, and reality. Strictly proper scoring rules have the property that for any forecast, the mathematical expectation of the score of a forecast p by the lights of p is strictly better than the mathematical expectation of any other forecast q by the lights of p . Forecasts need not satisfy the axioms of the probability calculus, but Predd et al. [9] have shown that given a finite sample space and any strictly proper *additive and continuous* scoring rule, the score for any forecast that does not satisfy the axioms of probability is strictly dominated by the score for some probabilistically consistent forecast. Recently, this result has been extended to non-additive continuous scoring rules. In this paper, a condition weaker than continuity is given that suffices for the result, and the condition is proved to be optimal.

§1. The main results. Scoring rules measure the deviation between a forecast, which assigns degrees of confidence or credence to various events, and reality. Strictly proper scoring rules have the property that for any forecast, the mathematical expectation of the score of a forecast p by the lights of p is strictly better than the mathematical expectation of any other forecast q by the lights of p . Forecasts need not satisfy the axioms of probability, but under some conditions whose discussion is the main purpose of this paper, the score of a forecast that does not satisfy the axioms of probability is strictly dominated by the score of a forecast by that does satisfy these axioms. This result has been interpreted by epistemologists as supporting the idea that reasonable forecasts will always be probabilistically consistent (e.g., [4, 5, 7]).

To be precise, let Ω be a finite sample space, encoding the possible situations that the forecasts concern. Let \mathcal{C} be the set of all functions from the power set of Ω to the reals: these we call credence functions. Let \mathcal{P} be the subset of \mathcal{C} which consists of the functions satisfying the axioms of probability. Members of \mathcal{C} can be thought of as forecasts regarding Ω . An *accuracy* (respectively, *inaccuracy*) *scoring rule* is a function s from a set $\mathcal{F} \supseteq \mathcal{P}$ of credence function to $[-\infty, M]^\Omega$ (respectively, $[M, \infty]^\Omega$) for some finite M , where A^B is the set of functions from B to A . Then $s(c)(\omega)$ for $c \in \mathcal{F}$ measures the accuracy of the forecast c when we are in fact at $\omega \in \Omega$, with higher (respectively, lower) values being more accurate. Since inaccuracy and accuracy scoring rules differ merely by a sign, we shall now assume that scoring rules are accuracy scoring rules, translating results from the literature as needed.

Received: September 8, 2021.

2020 *Mathematics Subject Classification*: Primary 60A99, 62A01, Secondary 94D99.

Key words and phrases: domination, scoring rules, probabilism, probability, forecasts, accuracy, continuity.

© The Author(s), 2023. Published by Cambridge University Press on behalf of The Association for Symbolic Logic. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.



Given a probability $p \in \mathcal{P}$ and an extended real function f on Ω , let $E_p f$ be the expected value with respect to p defined in the following way to avoid multiplying infinity by zero:

$$E_p f = \sum_{\omega \in \Omega, p(\{\omega\}) \neq 0} p(\{\omega\}) f(\omega).$$

We then say that a scoring rule s is *proper* on $\mathcal{F} \supseteq \mathcal{P}$ provided that for every $p \in \mathcal{P}$ and every $c \in \mathcal{F}$, we have $E_p s(p) \geq E_p s(c)$, that it is *strictly proper* there provided the inequality is always strict, and that it is *quasi-strictly proper* there provided that it is proper and the inequality is strict when $p \in \mathcal{P}$ and $c \in \mathcal{F} \setminus \mathcal{P}$.

Propriety captures the idea that if an agent adopts a probability function p as their forecast, then by the agent’s lights there can be no improvement in the expected score from switching to a different forecast. Strict propriety captures the idea that an agent who has adopted a probability function p as their forecast will think other forecasts to be inferior. Proper and strictly proper scoring rules have been widely studied: for a few examples, see [1, 3, 7, 9, 12].

A scoring rule is said to be *additive* provided that $\mathcal{F} = \mathcal{C}$ and there is a collection of extended-real functions $(s_A)_{A \subseteq \Omega}$ on $\mathbb{R} \times \{0, 1\}$ such that for all $c \in \mathcal{F}$ and $\omega \in \Omega$:

$$s(c)(\omega) = \sum_{A \subseteq \Omega} s_A(c(A), 1_A(\omega)).$$

The set of probabilities \mathcal{P} can be equipped with the topology resulting from its natural embedding ψ in the product space $[0, 1]^\Omega$, where $\psi(p)(\omega) = p(\{\omega\})$. Thus, a sequence of probabilities (p_n) converges to a probability p just in case $p_n(\{\omega\}) \rightarrow p(\{\omega\})$ for all $\omega \in \Omega$.

A scoring rule is *probability-continuous* provided that the restriction of s to \mathcal{P} is a continuous function to $[-\infty, M]^\Omega$ equipped with the Euclidean topology.

Say that a score $s(c_1)$ is *weakly dominated* by a score $s(c_2)$ provided that $s(c_2)(\omega) \geq s(c_1)(\omega)$ for all $\omega \in \Omega$, and *strictly dominated* if the inequality is strict.

Predd et al. [9] showed that if s is a probability-continuous additive strictly proper scoring rule, then for any non-probability c , there is a probability p such that $s(c)$ is strictly dominated by $s(p)$. In other words, any forecaster whose forecast fails to be a probability can find a forecast that is a probability and that is strictly better no matter what. Recently, Pettigrew [8] announced that this result holds without the assumption of additivity, merely assuming probability-continuity. His proof was shown to have flaws [6], but correct proofs were found by Nielsen [6] and Pruss [10]. Nielsen’s proof also extended the result to the quasi-proper case. A philosophical upshot of these results is that we can get an argument in favor of probabilistic consistency in one’s credence assignments under much weaker conditions than the additivity assumed by Predd et al.

However, the Pettigrew–Nielsen–Pruss theorem still assumes the continuity of the scoring rule. The purpose of the present paper is to identify what is the weakest possible assumption on a strictly proper scoring rule as restricted to the probability functions that guarantees the strict domination property.

We need some definitions to state our main result.

Say that a score $s \in [-\infty, M]^\Omega$ is finite provided that $|s(\omega)| < \infty$ for all ω . Note that \mathbb{R}^Ω is an n -dimensional vector space. Let $\langle \cdot, \cdot \rangle$ be the usual scalar product on

$\mathbb{R}^\Omega: \langle f, g \rangle = \sum_{\omega} f(\omega)g(\omega)$ for any $f, g \in \mathbb{R}^\Omega$. We say that a boundary point z of a set $G \subseteq \mathbb{R}^\Omega$ is positive-facing provided that G has a normal v at z all of whose components are positive, i.e., there is a $v \in (0, \infty)^\Omega$ such that $\langle v, w \rangle \leq \langle v, z \rangle$ for all $w \in G$. Write $\partial^+ G$ for the set of all the positive-facing boundary points.

Write $\text{Conv } F$ for the convex hull of F , i.e., the union of all the line segments with endpoints in F . A set A is dense in a set B in a topological space if every open set that intersects B also intersects A (e.g., the rational numbers are dense in the reals).

THEOREM 1. *Consider a proper scoring rule s on \mathcal{P} . Then the following conditions are equivalent:*

- (a) *For every extension of s to a quasi-strictly proper scoring rule $s' : \mathcal{C} \rightarrow [-\infty, M]^\Omega$, if $c \in \mathcal{C} \setminus \mathcal{P}$, then there is a $p \in \mathcal{P}$ such that $s'(p) = s(p)$ strictly dominates $s'(c)$.*
- (b) *Either $E_p s(p)$ is infinite for some $p \in \mathcal{P}$ or both:*
 - (i) *For any sequence (p_n) in \mathcal{P} that converges to some probability function p such that $s(p_n)$ is finite for all n while $s(p)$ is not finite, we have $\lim_n E_{p_n} s(p_n) = E_p s(p)$.*
 - (ii) *If $F = s[\mathcal{P}] \cap \mathbb{R}^\Omega$ is the set of finite scores, then F is dense in $\partial^+ \text{Conv } F$.*

Combining the above with the following will yield a new proof of the Pettigrew–Nielsen–Pruss theorem.

PROPOSITION 1. *Suppose that s is quasi-strictly proper and continuous on \mathcal{C} . Then condition (b) of Theorem 1 is fulfilled.*

It follows from Lemma 3, below, that for any proper scoring rule s , the function $p \mapsto E_p s(p)$ is continuous on the probabilities with finite score. Thus, in condition (b)(i) of Theorem 1 we can drop the restriction that $s(p)$ is not finite.

The proofs of the theorem and proposition will be given in Section 2.

Theorem 1 becomes simpler in the special case where all the values of the proper scoring rule s are finite, because in that case condition (b)(i) is always satisfied, and (b)(ii) is necessary and sufficient for the domination condition (a).

To visualize geometrically what the crucial condition (b)(ii) says, identify our space Ω with the set $\{1, 2, \dots, n\}$. Then a probability p can be thought of as a vector in n -dimensional Euclidean space \mathbb{R}^n whose k th coordinate is $p(\{k\})$, with all the coordinates non-negative and summing to one, and a score $s(p)$ can be thought of as an extended-real “vector” whose k th coordinate is $s(p)(k)$. By abuse of notation, in this geometrical explanation we won’t distinguish probabilities and the vectors corresponding to them, or scores and the vectors corresponding to them (we will be a little more careful when we get to proofs). Then F is the set of scores that lie in \mathbb{R}^n . The set $\partial^+ \text{Conv } F$ consists of those points z on the boundary of the convex hull of F such that some ray starting at z whose direction is positive in every coordinate immediately leaves the convex hull of F . Condition (b)(ii) then says that any such z is the limit of some sequence $s(p_i)$ of finite scores of probabilities p_i .

When we form the convex hull of F , we are adding to F various line segments to obtain a convex set (a set that contains the line segment joining every pair of points in it). Doing this may increase the positive-facing boundary of F to include some additional line segments. Condition (b)(ii) then says that any point on any of these additional line segments has a point of F arbitrarily close to it. In some sense, this means that F doesn’t have any positive-facing open gaps.

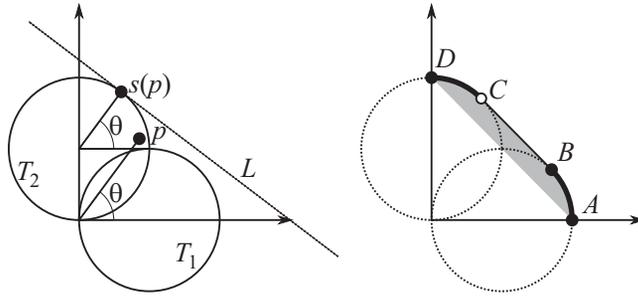


Fig. 1. Left: Calculating the score of $p = (3/7, 4/7)$ with $\theta = \theta(p) = 0.927 > \pi/4$. Right: Thick lines and filled nodes indicate the set F . Shading indicates the convex hull. The set $\partial^+ \text{Conv } F$ consists of the thick lines, the thin line segment BC , and the filled and unfilled nodes.

We can illustrate this by describing a finite strictly proper scoring rule that does not satisfy (b)(ii). Suppose $\Omega = \{1, 2\}$, so our probabilities and scores are identified with points in the plane \mathbb{R}^2 . Given a probability p , i.e., a non-zero vector with both coordinates non-negative, let $\theta(p)$ be the angle that p makes with the x -axis. Then $\theta(p)$ ranges between 0 and $\pi/2$ radians. If $\theta(p) \leq \pi/4$, let $s(p)$ be the point at angle $\theta(p)$ on the circle T_1 of radius 1 with center $(1, 0)$. Thus, $s(p) = (1 + \cos \theta(p), \sin \theta(p))$. If $\pi/4 < \theta(p)$, let $s(p)$ be the point at angle $\theta(p)$ on the circle T_2 of radius 1 with center $(0, 1)$, so $s(p) = (\cos \theta(p), 1 + \sin \theta(p))$ (see Figure 1, left). This is a strictly proper scoring rule.¹ It's worth noting if T_1 and T_2 were both the unit circle (so that the $\pi/4$ switchover was irrelevant), the resulting scoring rule would have been the spherical scoring rule.

In Figure 1 (right), the set F of values (all finite) of s consists of an arc AB of T_1 from angle 0 inclusive to $\pi/4$ inclusive together with an arc CD of T_2 from angle $\pi/4$ to $\pi/2$. The convex hull of F consists of the shaded region in the figure as well as some parts of the boundary of the shaded region². The positive-facing boundary of $\text{Conv } F$ consists of the arcs AB and CD as well as the line segment BC , together with all four endpoints. We can now see that F is not dense in $\partial^+ \text{Conv } F$, because F has a gap consisting of the interior of the line segment BC , while $\text{Conv } F$ has positive-facing boundary there.

¹ Here is a quick geometric sketch why. Consider a probability p at angle $\theta(p)$. Let L be a line through $s(p)$ tangent to the circle that $s(p)$ is on (Figure 1, left). The line L is tangent to one of our circles at the point $s(p)$, which is at angle $\theta(p)$ on the relevant circle. Thus, L has angle $\theta(p) + \pi/2$ and is thus perpendicular to the vector p , which has angle $\theta(p) \in [0, \pi/2]$. The line L cuts the plane into two open half-spaces. Let H be the half-space that contains the origin (this half-space is below/left of L —see the figure). Then H is the set of points z such that $\langle p, z \rangle < \langle p, s(p) \rangle$, where $\langle \cdot, \cdot \rangle$ is the dot product. But $\langle p, z \rangle = E_p z$ and $\langle p, s(p) \rangle = E_p s(p)$. Now, all possible scores lie on the arcs AB and CD , with the point C not being a possible score, and so it is easy to see (considering the case where $s(p)$ is B separately) that all the possible scores $s(q)$ for $q \neq p$ lie below and/or to the left of the line L . Hence they all satisfy $E_p s(q) < E_p s(p)$, and we have strict propriety.

² The set $\text{Conv } F$ includes all of its topological boundary except for the point C and the interior of the line segment BC , but it doesn't matter for the definition of $\partial^+ \text{Conv } F$ which parts of the boundary are members of $\text{Conv } F$.

Finally, we give a family of examples of proper scoring rules that satisfy (b)(i) and (b)(ii) but are not continuous. The family will even include some scores that are finite everywhere on Ω . Let \mathcal{R} be the set of probabilities that are *regular* in the Bayesian sense, i.e., probability functions p such that $p(\{\omega\}) > 0$ for all ω . Note that if p is regular and s is strictly proper, then we must have $E_p(s(p)) > -\infty$ or else we couldn't have $E_p(s(q)) < E_p(s(p))$ for a different probability q . Then by regularity of p it follows that $s(p)$ is finite.

Let s be any continuous strictly proper everywhere-finite accuracy scoring rule. Choose any $\alpha \in [-\infty, M]$ such that $\alpha \leq s(p)(\omega)$ for all $p \in \mathcal{P}$ and $\omega \in \Omega$. Define a scoring rule s_α as follows. If either $p \notin \mathcal{P}$ or $p \in \mathcal{R}$, let $s_\alpha(p) = s(p)$. If $p \in \mathcal{P} \setminus \mathcal{R}$, then let $s_\alpha(p)(\omega) = s(p)(\omega)$ if $p(\{\omega\}) > 0$, and $s_\alpha(p)(\omega) = \alpha$ if $p(\{\omega\}) = 0$. In other words, we have tweaked s so that in the case where the forecast assigns zero probability to some outcome ω , we assign α there. It is easy to see that s_α is strictly proper because $E_p s_\alpha(p) = E_p s(p)$ and $E_p s_\alpha(q) \leq E_p s(q)$ for all $p, q \in \mathcal{P}$.

For a fixed ω , the set of possible values of $s(p)(\omega)$ is contained in a finite interval, since a continuous function on a compact set takes on a compact set of values, and \mathcal{P} is compact. Thus we can choose α so that $\alpha < s(c)(\omega)$ for all ω , letting α be $-\infty$ or a finite value as we wish. Then s_α will be discontinuous everywhere on $\mathcal{P} \setminus \mathcal{R}$, since s_α is discontinuous wherever it differs from s . And if α is finite, then s_α will be everywhere finite as well.

We now show that s_α satisfies the conditions of Theorem 1. First, s satisfies condition (b) by Proposition 1, and hence it satisfies condition (a) by the theorem. Now let s'_α be a quasi-strictly proper extension of s_α to \mathcal{C} . Define $s'(c) = s(c)$ for $c \in \mathcal{P}$ and $s'(c) = s'_\alpha(c)$ for $c \notin \mathcal{P}$. Then s' is quasi-strictly proper, because for any $c \notin \mathcal{P}$ and $p \in \mathcal{P}$ we have

$$E_p s'(c) = E_p s'_\alpha(c) < E_p s_\alpha(p) = E_p s(p) = E_p s'(p),$$

where the first inequality follows from the quasi-strict propriety of s'_α . Thus, because s satisfies condition (a), for any $c \in \mathcal{P}$ there is a $p \in \mathcal{P}$ such that $s'(c)$ is strictly dominated by $s'(p)$. Since s' is continuous on \mathcal{P} and \mathcal{R} is dense in \mathcal{P} , we can choose the probability p to be in \mathcal{R} . Then $s'_\alpha(c) = s'(c)$ will be strictly dominated by $s'(p) = s(p) = s_\alpha(p)$, and so s_α will satisfy condition (a) of the theorem, and hence also condition (b).

Note that in the case where $\alpha = -\infty$ the modified scoring rule $s_{-\infty}$ has some intuitive plausibility for measuring the accuracy of a credence assignment or forecast. For if I assign credence zero to an outcome ω of our sample space, when in fact ω is the actual outcome, then I make an error that in an important sense is infinitely bad, because no amount of Bayesian conditionalization on events with non-zero probability will get me out of the error. The modified scoring rule $s_{-\infty}$ thus makes forecasts that are not regular be much more risky than an everywhere finite scoring rule s does. I do not, of course, propose that scoring rules like s_α be generally adopted, but only want to note that someone sufficiently attached to regularity might have some reason to do so.

§2. Proofs. Given a probability function p , let \hat{p} be the function from Ω to $[0, 1]$ defined by $\hat{p}(\omega) = p(\{\omega\})$. Define $\hat{\mathcal{P}} = \{\hat{p} : p \in \mathcal{P}\}$ and $\hat{\mathcal{R}} = \{\hat{p} : p \in \mathcal{R}\}$ (recall that \mathcal{R} is the set of regular probabilities, i.e., ones that are non-zero on every singleton).

Thus, $\hat{\mathcal{P}}$ is the set of non-negative functions on Ω the sum of whose values is 1, and $\hat{\mathcal{R}}$ is the subset of the strictly positive ones.

Given $p \in \hat{\mathcal{P}}$, let \check{p} be the probability function such that $\check{p}(\{\omega\}) = p(\omega)$ for all $\omega \in \Omega$.

Given a scoring rule s defined on \mathcal{P} , by abuse of notation let $s(p) = s(\check{p})$ for $p \in \hat{\mathcal{P}}$.

Given two functions f and g from Ω to the extended reals, say that g strictly (weakly) dominates f provided that $f < g$ ($f \leq g$) everywhere.

Let \cdot be a multiplication operation on the extended reals with the stipulation that $a \cdot 0 = 0 \cdot a = 0$ for any a , finite or not. With this stipulation, multiplication is upper semicontinuous on $[-\infty, \infty) \times [0, 1]$, and continuous at all (x, y) where x is finite or y is positive. Now define the *extended scalar product* on $X = [0, 1]^\Omega \times [-\infty, \infty)^\Omega$ by

$$\langle f, g \rangle = \sum_{\omega \in \Omega} f(\omega) \cdot g(\omega),$$

using the above stipulation. The following properties are easy to check.

LEMMA 1. (a) *The extended scalar product is an upper semicontinuous function from X to $[-\infty, \infty)$. Moreover, (b) it is continuous at any (f, g) such that either f is strictly positive everywhere or g is finite everywhere. Finally, (c) for a fixed $f \in [0, 1]^\Omega$, the function $g \mapsto \langle f, g \rangle$ is continuous on $[-\infty, \infty)^\Omega$.*

Observe that $E_p f = \langle \hat{p}, f \rangle$. This fact will allow us to go back and forth between probabilistic concepts and geometric concepts.

For a subset F of \mathbb{R}^Ω and a vector $v \in \mathbb{R}^\Omega$, let

$$\sigma_F(v) = \sup_{z \in F} \langle v, z \rangle$$

be the support function of F at v .

LEMMA 2. *Assume s is a proper scoring rule on \mathcal{P} with $E_p s(p)$ finite for all p . Let F be the set of finite scores. Suppose that for every convergent sequence (p_n) of members of $\hat{\mathcal{P}}$ with $s(p_n) \in F$ for all n , we have $\lim_n \langle p_n, s(p_n) \rangle = \langle p, s(p) \rangle$, where $p = \lim_n p_n$. Then $\sigma_F(p) = \langle p, s(p) \rangle$ for all $p \in \hat{\mathcal{P}}$.*

Proof of Lemma 2. First, suppose $s(p) \in F$. Then for every $z \in F$ we have $\langle p, z \rangle \leq \langle p, s(p) \rangle$ by propriety. Since $s(p) \in F$, it follows that $\sigma_F(p) = \langle p, s(p) \rangle$.

Now suppose $s(p)$ is not finite. Let (p_n) be a sequence in $\hat{\mathcal{R}}$ converging to p . For $q \in \hat{\mathcal{R}}$, the fact that $E_q s(q)$ is finite implies that $s(q)$ is finite, so $s(p_n) \in F$ for all n . Using compactness and passing to a subsequence if necessary, assume that $s(p_n)$ converges to some value $z \in [-\infty, M]^\Omega$. Then

$$\sigma_F(p) \geq \lim_n \langle p, s(p_n) \rangle = \langle p, z \rangle \geq \limsup_n \langle p_n, s(p_n) \rangle = \langle p, s(p) \rangle,$$

where the relations follow respectively by definition of σ_F , the continuity of the extended scalar product for a fixed first argument (Lemma 1(c)), the upper semicontinuity of the extended scalar product on X (Lemma 1(a)), and the assumptions of our present lemma.

On the other hand:

$$\langle p, s(p) \rangle \geq \langle p, s(q) \rangle$$

by propriety for all $q \in \hat{\mathcal{P}}$. Hence $\langle p, s(p) \rangle \geq \sigma_F(p)$. Thus, $\sigma_F(p) = \langle p, s(p) \rangle$. □

LEMMA 3. Let s be a proper scoring rule on \mathcal{P} . Let $H = \{p \in \hat{\mathcal{P}} : s(p) \text{ is finite}\}$. For any p in the closure \bar{H} of H , the limit of $\langle q, s(q) \rangle$ as q tends to p within H exists. If $s(p)$ is finite, the limit equals $\langle p, s(p) \rangle$. Finally, if p_n is a sequence in H converging to $p \in \bar{H}$ such that $\lim_n s(p_n) = r$, then $\langle p, r \rangle = \lim_n \langle p_n, s(p_n) \rangle$.

Proof. Fix $p \in \bar{H}$. To show that a sequence converges to some member of a compact set, it suffices to show that every convergent subsequence of it converges to the same point. The set $[-\infty, M]$ is compact and $\langle q, s(q) \rangle \in [-\infty, M]$ for all q . Thus to show the existence of our limit, all we need to show is that if (p_n) and (p'_n) are two sequences in H converging to p with $\langle p_n, s(p_n) \rangle \rightarrow L$ and $\langle p'_n, s(p'_n) \rangle \rightarrow L'$, then $L = L'$. Moreover, if we can show this, then letting $p'_n = p$ for all n , it will follow that $L = s(p)$ if $s(p)$ is finite.

Thus, suppose (p_n) and (p'_n) are two sequences in H converging to p with $\langle p_n, s(p_n) \rangle$ and $\langle p'_n, s(p'_n) \rangle$ convergent. Passing to subsequences if necessary, assume that $s(p_n)$ and $s(p'_n)$ converge respectively to r and r' in $[-\infty, M]^\Omega$.

I now claim that $\langle p, r \rangle = \lim_n \langle p_n, s(p_n) \rangle$. First, note that

$$\limsup_n \langle p_n, s(p_n) \rangle \leq \langle p, r \rangle$$

by Lemma 1(a). Next observe that for any fixed m we have

$$\langle p_n, s(p_n) \rangle \geq \langle p_n, s(p_m) \rangle$$

by propriety. Since $s(p_m)$ is finite as $p_m \in H$, by Lemma 1(b) the right-hand side converges to $\langle p, s(p_m) \rangle$ as $n \rightarrow \infty$. Thus,

$$\liminf_n \langle p_n, s(p_n) \rangle \geq \langle p, s(p_m) \rangle.$$

Taking the limit as $m \rightarrow \infty$ and using Lemma 1(c) we get

$$\liminf_n \langle p_n, s(p_n) \rangle \geq \langle p, r \rangle.$$

Thus we have $\langle p, r \rangle = \lim_n \langle p_n, s(p_n) \rangle$, which is the final remark in the statement of our lemma.

Thus,

$$\begin{aligned} \langle p, r' \rangle &= \lim_m \langle p, s(p'_m) \rangle \\ &= \lim_m \lim_n \langle p_n, s(p'_m) \rangle \\ &\leq \lim_m \lim_n \langle p_n, s(p_n) \rangle \\ &= \lim_m \langle p, r \rangle = \langle p, r \rangle, \end{aligned}$$

where the first equality was due to Lemma 1(c), the second due to Lemma 1(b) and the finiteness of $s(p'_m)$, and the inequality was due to propriety. In exactly the same way, $\langle p, r \rangle \leq \langle p, r' \rangle$. Thus, $\langle p, r \rangle = \langle p, r' \rangle$. But $\langle p, r' \rangle = \lim_n \langle p'_n, s(p'_n) \rangle$, just as we proved in the unprimed case. Thus, $\lim_n \langle p_n, s(p_n) \rangle = \lim_n \langle p'_n, s(p'_n) \rangle$. \square

LEMMA 4. Let F be a nonempty subset of $[-\infty, \infty)^\Omega$. Suppose $z_0 \in [-\infty, \infty)^\Omega$ is such that $\langle p, z_0 \rangle < \sigma_F(p)$ for all $p \in \hat{\mathcal{P}}$. Then z_0 is strictly dominated by some member of $\text{Conv}(F)$.

Proof of Lemma 4. Let z_0 be as in the statement of the lemma. Let $Q = \{z \in \mathbb{R}^\Omega : \forall \omega (z(\omega) > z_0(\omega))\}$ be the set of points of \mathbb{R}^Ω strictly dominating z_0 . We need to show that $Q \cap \text{Conv}(F) \neq \emptyset$.

Suppose that Q does not intersect $\text{Conv}(F)$. Both Q and $\text{Conv}(F)$ are convex sets. Thus by hyperplane separation, there is a non-zero $p \in \mathbb{R}^\Omega$ and an $\alpha \in \mathbb{R}$ such that $\langle p, z \rangle \geq \alpha$ for all $z \in Q$ and $\langle p, z \rangle \leq \alpha$ for all $z \in \text{Conv}(F)$.

I claim that $p(\omega) \geq 0$ for all ω . To see this, suppose $p(\omega_0) < 0$ for some $\omega_0 \in \Omega$. Let z be any member of Q . Let $\beta = \langle p, z \rangle$. Let $z'(\omega) = z(\omega)$ for $\omega \neq \omega_0$ and $z'(\omega_0) = z(\omega_0) + (\alpha - \beta - 1)/p(\omega_0)$. Then $z' \in Q$ since $z \in Q$ while $\beta \geq \alpha$ and $p(\omega_0) < 0$. Observe that

$$\langle p, z' \rangle = \langle p, z \rangle + (\alpha - \beta - 1) = \alpha - 1,$$

which is impossible as $\langle p, z' \rangle \geq \alpha$ since $z' \in Q$.

Rescaling if necessary, we may assume that $\sum_{\omega} p(\omega) = 1$ and hence $p \in \hat{\mathcal{P}}$.

Since z_0 is on the boundary of $Q \subseteq [-\infty, \infty)^\Omega$, by the upper semicontinuity of the extended scalar product (Lemma 1(a)) we have $\langle p, z_0 \rangle \geq \alpha$. Moreover, since $\langle p, z \rangle \leq \alpha$ for all $z \in \text{Conv}(F)$, we must have $\sigma_F(p) \leq \alpha$. Thus, $\langle p, z_0 \rangle \geq \sigma_F(p)$, which contradicts the assumptions of the lemma. \square

Say that a vector v is normal to a convex set G at a point $z_1 \in G$ provided that $\langle v, z \rangle \leq \langle v, z_1 \rangle$ for all $z \in G$. The following lemma is due to a *MathOverflow* user [2]. Given a point $z \in \mathbb{R}^n$, let Q_z be the positive orthant $\mathbb{R}_+^n = (0, \infty)^n$ translated to put its vertex at z , i.e., $Q_z = \{z + w : w \in \mathbb{R}_+^n\}$.

LEMMA 5. *Fix $z_1 \in \mathbb{R}^n$. Let G be a closed convex subset of \mathbb{R}^n whose intersection with Q_{z_1} is non-empty and bounded. Then there is a vector v in the positive orthant \mathbb{R}_+^n that is normal to G at some point $z \in G \cap Q_{z_1}$.*

Proof. Translating G as needed, assume without loss of generality that $z_1 = 0$ and $Q_{z_1} = \mathbb{R}_+^n$.

Let $f(x_1, \dots, x_n) = x_1 \cdots x_n$ be the product-of-coordinates function. Let $U = G \cap \mathbb{R}_+^n$. Then f is a continuous function on the closure \bar{U} of U in \mathbb{R}^n . Moreover, \bar{U} is compact by the boundedness requirement, so f attains a maximum on \bar{U} . This maximum is not attained on the boundary of \mathbb{R}_+^n , i.e., on any point with a zero coordinate, since f is zero at any such point, while there is at least one point in \bar{U} where f is strictly positive (since f is non-zero everywhere on $U \neq \emptyset$). Hence, the maximum is attained at a point z of U . That point cannot be an interior point (since then $z + (\varepsilon, \dots, \varepsilon)$ would be in U and f would be bigger there than at z), so it must be a boundary point of U that isn't a boundary point of \mathbb{R}_+^n . Thus, z must be a boundary point of G . The basic normal cone condition for optimality tells us that the gradient of f must be normal to U at z for a differentiable f to be optimal at z over U (see [11, theorem 6.12] in the case of minima). Since \mathbb{R}_+^n is a neighborhood of z , and the intersection of U with that neighborhood is the same as that of G with it, it follows that the gradient of f must be normal to G at z . But the gradient of f on \mathbb{R}_+^n always lies in \mathbb{R}_+^n , which completes the proof. \square

Let $\|z\|_\infty = \max_{\omega \in \Omega} |z(\omega)|$ for $z \in [-\infty, \infty)^\Omega$ and let

$$B(z, r) = \{z' : \|z' - z\|_\infty < r\}$$

be the open ball of radius r around a finite z in this norm.

LEMMA 6. Fix $p \in \hat{\mathcal{R}}$ and $\alpha \in \mathbb{R}$, and let $K = \{y \in [-\infty, \infty)^\Omega : \langle p, y \rangle \leq \alpha\}$. Fix $\varepsilon > 0$ and $x \in \mathbb{R}^\Omega$ such that $\langle p, x \rangle = \alpha$. Then there is a $\delta \in (0, \varepsilon)$ such that no point of $K \cap B(x, \delta)$ is weakly dominated by any point of $K \setminus B(x, \varepsilon)$.

Proof of Lemma 6. Translating if necessary, without loss of generality assume $x = 0$ and $\alpha = 0$. Fix $y \in K \cap B(0, \delta)$. Suppose that y is weakly dominated by $z \in K$.

Then $-z(\omega) \leq -y(\omega) \leq \delta$ for all ω and $\sum_\omega z(\omega)p(\omega) \leq 0$ as well as $\sum_\omega y(\omega)p(\omega) \leq 0$. Let $c = 1/\min_\omega p(\omega)$. Then for any ω :

$$z(\omega) \leq -\frac{\sum_{\omega' \neq \omega} z(\omega')p(\omega')}{p(\omega)} \leq -\frac{\sum_{\omega' \neq \omega} y(\omega')p(\omega')}{p(\omega)} \leq \frac{\sum_{\omega' \neq \omega} \delta p(\omega')}{p(\omega)} \leq c\delta.$$

Moreover $-\delta \leq z(\omega)$ and $c \geq 1$, so $\|z\|_\infty \leq c\delta$. Thus, we have shown that if $y \in K \cap B(0, \delta)$ is weakly dominated by $z \in K \setminus B(x, \varepsilon)$, then $\varepsilon \geq c\delta$. Hence, if $\varepsilon > 0$ is fixed, any choice of $\delta \in (0, c^{-1}\varepsilon)$ will complete the proof. \square

Proof of Theorem 1. If $E_p s(p)$ is infinite (i.e., equal to $-\infty$ since we are working with accuracy scores) for some probability function p , then s has no extension to a quasi-strictly proper scoring rule on \mathcal{C} , as no point $z \in [-\infty, \infty)^\Omega$ will be such that $E_p z < E_p s(p)$. Thus, we may assume for all our proofs that $E_p s(p)$ is finite for all probabilities p , and hence that so is $\langle \hat{p}, s(\hat{p}) \rangle$.

Without loss of generality, assume we have accuracy scores with ranges in $[-\infty, -1]$.

Recall that $E_p s(q) = \langle \hat{p}, s(\hat{q}) \rangle$. If (i) fails, then there is a sequence $p_n \rightarrow p$ in $\hat{\mathcal{P}}$ such that $\langle p_n, s(p_n) \rangle$ does not converge to $\langle p, s(p) \rangle$ while $s(p_n)$ is finite and $s(p)$ is infinite. By Lemma 3, $\lim_n \langle p_n, s(p_n) \rangle = L$ exists and cannot equal $\langle p, s(p) \rangle$. Passing to a subsequence if necessary, we may assume that $s(p_n)$ has a limit $r \in [-\infty, -1]^\Omega$, and then

$$\lim_n \langle p_n, s(p_n) \rangle = \langle p, r \rangle \tag{1}$$

by the same lemma.

Note that

$$\langle q, r \rangle = \lim_n \langle q, s(p_n) \rangle \leq \langle q, s(q) \rangle$$

for all $q \in \hat{\mathcal{P}}$ by Lemma 1(c) and propriety. In particular, $L = \langle p, r \rangle \leq \langle p, s(p) \rangle$. Thus, $L < \langle p, s(p) \rangle$ since $\langle p_n, s(p_n) \rangle$ does not converge to $\langle p, s(p) \rangle$. Further, $\langle p, s(p) \rangle$ is negative since our accuracy scores are negative. Choose $\alpha \in (1, L/\langle p, s(p) \rangle)$ so that

$$\alpha \langle p, s(p) \rangle > L = \langle p, r \rangle. \tag{2}$$

Let $x = \alpha s(p)$.

Infinite scores cannot strictly dominate any score. Now I claim that x is not weakly dominated by any finite score. For suppose that $s(q)$ is finite. Then by Lemma 1(b), propriety, (1) and (2) we have

$$\langle p, s(q) \rangle = \lim_n \langle p_n, s(q) \rangle \leq \lim_n \langle p_n, s(p_n) \rangle = \langle p, r \rangle < \alpha \langle p, s(p) \rangle = \langle p, x \rangle.$$

And this implies that x is not strictly dominated by $s(q)$.

On the other hand,

$$\langle q, x \rangle = \alpha \langle q, s(p) \rangle \leq \alpha \langle q, s(q) \rangle < \langle q, s(q) \rangle \tag{3}$$

for any $q \in \hat{\mathcal{P}}$ since $\alpha > 0$ and our scores are negative.

Now define $s'(c) = s(c)$ if $c \in \mathcal{P}$ and $s'(c) = x$ if $c \in \mathcal{C} \setminus \mathcal{P}$. Then s' is quasi-strictly proper by (3), but $s'(c)$ is not dominated by any score of a probability. Hence condition (a) fails.

Now suppose (ii) fails. Thus there is a $z_0 \in \partial^+ \text{Conv } F$ and $\varepsilon > 0$ such that $F \cap B(z_0, \varepsilon) = \emptyset$. Then $\text{Conv } F$ has a normal in $(0, \infty)^\Omega$ at z_0 . Rescaling if necessary, we can assume that normal is some $p \in \hat{\mathcal{R}}$. Let $K = \{z \in [-\infty, \infty)^\Omega : \langle p, z \rangle \leq \langle p, z_0 \rangle\}$, which then contains $\text{Conv } F$. By Lemma 6, there is a $\delta \in (0, \varepsilon)$ such that no point in $B(z_0, \delta)$ is strictly dominated by any point in $K \setminus B(z_0, \varepsilon)$. Choose a point z_1 in $B(z_0, \delta)$ that is strictly dominated by z_0 . Note that z_0 is a limit of convex combinations of points of F . If q is any member of $\hat{\mathcal{P}}$ and u is any point of F , then $\langle q, u \rangle \leq \langle q, s(q) \rangle$ by propriety. Thus, the same is true if u is a convex combination of points of F , and by Lemma 1(c) also if u is a limit of convex combinations of points of F .

Hence, $\langle q, z_0 \rangle \leq \langle q, s(q) \rangle$ for all $q \in \hat{\mathcal{P}}$. Since z_1 is strictly dominated by z_0 , it follows that $\langle q, z_1 \rangle < \langle q, s(q) \rangle$ for all $q \in \hat{\mathcal{P}}$.

On the other hand, since z_1 is not strictly dominated by anything in $K \setminus B(z_0, \varepsilon)$, it's not strictly dominated by anything in F as $\text{Conv } F \subseteq K$. Much as before, let $s'(c) = s(c)$ if c is a probability function and $s'(c) = z_1$ if c is not a probability function. As before, we have strict quasi-proprity and yet no credence that isn't a probability is strictly s' -dominated by any probability function.

Now suppose that (i) and (ii) hold. Let s' be an extension of s to a quasi-strictly proper scoring rule defined for all credences. Fix a non-probability credence c and let $z_0 = s'(c)$. By (i) and Lemma 2, $\langle p, s(p) \rangle = \sigma_F(p)$ for all p . By Lemma 4 and quasi-strict propriety, z_0 is strictly dominated by some member of $\text{Conv}(F)$.

Let G be the closure of $\text{Conv}(F)$ in \mathbb{R}^Ω . Then there is a point $z_1 \in (-\infty, \infty)^\Omega$ such that z_0 is strictly dominated by z_1 and z_1 is strictly dominated by some member of G (e.g., if z_2 is a point of G that strictly dominates z_0 , then let $z_1(\omega) = z_2(\omega) - 1$ if $z_0(\omega)$ is infinite, and $z_1(\omega) = (z_0(\omega) + z_2(\omega))/2$ otherwise). Let $Q = Q_{z_1}$ be the set of points of $(-\infty, \infty)^\Omega$ that strictly dominate z_1 . Then $Q \cap G$ is non-empty. Note that every point z of F satisfies

$$\langle u, z \rangle \leq \langle u, s(u) \rangle \tag{4}$$

by propriety, where $u(\omega) = 1/|\Omega|$ for all ω , and hence every point z of G satisfies (4) as well (a convex combination of points z satisfying (4) will satisfy it as well, and by Lemma 1(c), so will every point in the closure of $\text{Conv } F$). Moreover, $\langle u, s(u) \rangle$ is finite. The set of points $z \in Q$ such that $\langle u, z \rangle \leq \langle u, s(u) \rangle$ is bounded, and hence $Q \cap G$ is bounded.

Since $Q \cap G$ is bounded and non-empty, by Lemma 5 (letting $n = |\Omega|$ so that \mathbb{R}^Ω and \mathbb{R}^n are isomorphic as vector spaces), there is a $z_3 \in Q \cap G$ such that $z_3 \in G$ has a normal in the positive orthant. Thus, $z_3 \in \partial^+ G = \partial^+ \text{Conv } F$. By condition (ii), there are points of F arbitrarily close to z_2 . Since z_2 strictly dominates z_0 , so do points that are sufficiently close to z_2 , and hence some point of F strictly dominates z_0 . \square

Proof of Proposition 1. Suppose s is continuous. We must have $E_p s(p)$ finite for all probability functions p or else quasi-strict propriety is impossible. It follows that the score of any regular probability is finite.

By Lemma 1(a), $p \mapsto \langle p, s(p) \rangle$ is upper semicontinuous on $\hat{\mathcal{P}}$ if the scoring rule is continuous.

Moreover, $\langle p, s(p) \rangle = \sup_{q \in \mathcal{P}} \langle p, s(q) \rangle$. Since $\hat{\mathcal{R}}$ is dense in $\hat{\mathcal{P}}$ and $q \mapsto \langle p, s(q) \rangle$ is continuous for a fixed p (using Lemma 1(c)), it follows that $\sup_{q \in \hat{\mathcal{P}}} \langle p, s(q) \rangle = \sup_{q \in \hat{\mathcal{R}}} \langle p, s(q) \rangle$. Moreover, since every score of a regular probability is finite, $p \mapsto \langle p, s(q) \rangle$ is continuous by Lemma 1(b) for $q \in \hat{\mathcal{R}}$. Thus, $p \mapsto \sup_{q \in \hat{\mathcal{R}}} \langle p, s(q) \rangle$ is the supremum of continuous functions and hence it is lower semicontinuous at these points. (This observation is due to Nielsen (2021).)

Hence, $p \mapsto \langle p, s(p) \rangle$ is continuous at every point of $\hat{\mathcal{P}}$, and so we have (i).

It remains to prove (ii). Assume first that s is strictly proper. Let F be the set of finite scores of probabilities. Then F considered as a subset of \mathbb{R}^Ω is closed, since it is the intersection with \mathbb{R}^Ω of the set of scores of probability functions, and the set of probability functions is compact while s is continuous. Moreover, for any $p \in \hat{\mathcal{P}}$, we have $\langle p, s(p) \rangle > \langle p, z \rangle$ for all $z \in F \setminus \{s(p)\}$. It follows that for any p , the only point z in the closed convex hull of F such that $\langle p, s(p) \rangle = \langle p, z \rangle$ is $s(p)$ itself, from which (ii) follows.

Now, suppose s is merely quasi-strictly proper. Let $b : \mathcal{C} \rightarrow [-1, 0]$ be any strictly proper continuous accuracy score, for instance -1 plus the Brier score. Let $s_\varepsilon = s + \varepsilon b$ for any $\varepsilon > 0$. Then s_ε is a strictly proper continuous score, and (ii) must hold for it. Let $F_\varepsilon = s_\varepsilon[\mathcal{P}] \cap \mathbb{R}^\Omega$.

Now, consider a point $z_0 \in \partial^+ \text{Conv } F$. Fix $\varepsilon > 0$. We will show that there is a point of F within distance ε of z_0 . Let $\delta = \varepsilon/4$. Thus there is a $p \in \hat{\mathcal{R}}$ (since any vector in the positive orthant $(0, \infty)^\Omega$ can be rescaled to get a vector in $\hat{\mathcal{R}}$) such that $\langle p, z \rangle \leq \langle p, z_0 \rangle$ for all $z \in F$. The point z_0 is a convex combination $z_0 = c_1 w_1 + \dots + c_n w_n$ of points of F , where $c_i > 0$ and $\sum_i c_i = 1$. We then have to have $\langle p, w_i \rangle = \langle p, z_0 \rangle$ for each i . Choose $p_i \in \hat{\mathcal{P}}$ such that $w_i = s(p_i)$. Since s_δ is continuous and $\hat{\mathcal{R}}$ is dense in $\hat{\mathcal{P}}$, for each i choose p'_i of $\hat{\mathcal{R}}$ such that $\|s_\delta(p'_i) - s_\delta(p_i)\|_\infty \leq \delta$. Then $\|w_i - s_\delta(p_i)\|_\infty \leq 2\delta$ since s_δ is never more than δ away from s .

Let $z'_0 = \sum_i c_i s_\delta(p'_i)$. Then, $\|z'_0 - z_0\|_\infty \leq 2\delta$. Note that $s_\delta(p'_i) \in \partial^+ F_\varepsilon$, since $p'_i \in (0, \infty)^\Omega$ is normal to F_ε at $s_\delta(p'_i)$ by the propriety of s_δ . Applying the strictly proper case of our proposition to s_δ , we conclude that there is a $p \in \hat{\mathcal{P}}$ such that $\|s_\delta(p) - z'_0\|_\infty \leq \delta$ and $s_\delta(p) \in F_\varepsilon$. Then, $\|s(p) - z'_0\|_\infty \leq 2\delta$, and so $\|s(p) - z_0\|_\infty \leq 4\delta = \varepsilon$, which completes the proof. \square

Acknowledgments. I am grateful to Michael Nielsen for many discussions and much encouragement, and to two anonymous readers for a number of comments that helped with clarity, as well as for discovering a gap in my initial proofs.

BIBLIOGRAPHY

- [1] Dawid, A. P. & Musio, M. (2014). Theory and applications of proper scoring rules. *Metron*, **72**, 169–183.
- [2] fedja (2021) Answer to: If a Convex Compact Set Meets the Positive Orthant Does It Meet It at a Point with a Normal in the Positive Orthant? Mathoverflow. Available from: <https://mathoverflow.net/questions/384140>.
- [3] Gneiting, T. & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, **102**, 359–378.

- [4] Greaves, H. & Wallace, D. (2006). Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind*, **115**, 607–632.
- [5] Joyce, J. M. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science*, **65**, 575–603.
- [6] Nielsen, M. (2021). On the best accuracy arguments for probabilism. *Philosophy of Science*, **89**, 621–630.
- [7] Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
- [8] ———. (2022). Accuracy-first epistemology without the additivity axiom. *Philosophy of Science*, **89**, 128–151.
- [9] Predd, J. B., Seiringer, R., Lieb, E. H., Osherson, D. N., Vincent Poor, H., & Kulkarni, S. R. (2009). Probabilistic coherence and proper scoring rules. *IEEE Transactions on Information Theory*, **55**, 4786–4792.
- [10] Pruss, A. R. (2021). Proper scoring rules and domination. Preprint, [arXiv:2102.02260](https://arxiv.org/abs/2102.02260).
- [11] Tyrrell Rockafellar, R. & Wets, R. J. B. (1998). *Variational Analysis*. Berlin: Springer.
- [12] Winkler, R. L., Munoz, J., Cervera, J. L., Bernardo, J. M., Blattenberger, G., Kadane, J. B., Lindley, D. V., Murphy, A. H., Oliver, R. M., & Ríos-Insua, D. (1996). Scoring rules and the evaluation of probabilities. *Test*, **5**, 1–60.

DEPARTMENT OF PHILOSOPHY

BAYLOR UNIVERSITY

ONE BEAR PLACE #97273 WACO, TX 76706, USA

E-mail: alexander_pruss@baylor.edu