

Constant Regret, Generalized Mixability, and Mirror Descent

Zakaria Mhammedi
Robert C. Williamson

The Australian National University & Data61
Canberra, ACT, Australia

ZAK.MHAMMEDI@ANU.EDU.AU
BOB.WILLIAMSON@ANU.EDU.AU

Abstract

We consider the setting of prediction with expert advice; a learner makes predictions by aggregating those of a group of experts. Under this setting, and with the right choice of loss function and “mixing” algorithm, it is possible for the learner to achieve constant regret regardless of the number of prediction rounds. For example, constant regret can be achieved with *mixable* losses using the *Aggregating Algorithm* (AA). The *Generalized Aggregating Algorithm* (GAA) is a name for a family of algorithms parameterized by convex functions on simplices (entropies), which reduce to the AA when using the *Shannon entropy*. For a given entropy Φ , losses for which constant regret is possible using the GAA are called Φ -mixable. Which losses are Φ -mixable was previously left as an open question. We fully characterize Φ -mixability, and answer other open questions posed by Reid et al. (2015). We also elaborate on the tight link between the GAA and the *mirror descent algorithm* which minimizes the weighted loss of experts.

1. Introduction

Two fundamental problems in learning are how to aggregate information and under what circumstances can one learn fast. In this paper, we consider the problems jointly, extending the understanding and characterization of exponential mixing due to Vovk (1998), who showed that not only does the “*aggregating algorithm*” learn quickly when the loss is suitably chosen, but that it is in fact a generalization of classical Bayesian updating, to which it reduces when the loss is log-loss. We consider a general class of aggregating schemes, going beyond Vovk’s exponential mixing, and provide a complete characterization of the mixing behavior for general losses and general mixing schemes parameterized by an arbitrary entropy function.

In the *game of prediction with expert advice* a learner predicts the outcome of a random variable (outcome of the *environment*) by aggregating the predictions of a pool of experts. At the end of each prediction round t , the outcome of the environment is announced and the learner and experts suffer losses based on their predictions. We are interested in algorithms that the learner can use to “aggregate” the experts’ predictions and minimize the *regret* at the end of the game. In this case, the regret is defined as the difference between the cumulative loss of the learner and that of the best expert (in hindsight) after T rounds; if $\ell_t(\mathbf{a}_\theta^t)$ [resp. $\ell_t(\mathbf{a}_*^t)$] is the loss suffered by expert $\theta \in [k]$ [resp. the learner] who makes prediction \mathbf{a}_θ^t [resp. \mathbf{a}_*^t] at round t , then after T rounds the regret can be expressed as

$$R_\ell(T) = \sum_{1 \leq t \leq T} \ell_t(\mathbf{a}_*^t) - \inf_{\theta \in [k]} \sum_{1 \leq t \leq T} \ell_t(\mathbf{a}_\theta^t).$$

The *Aggregating Algorithm* (AA) (Vovk, 1998) achieves constant regret — a precise notion of “fast learning” — for *mixable* losses (formally defined later); that is, $R_\ell(T)$ is bounded from above by a constant R_ℓ independently of the number of rounds T . In this case, it is said that the algorithm achieves a *fast rate* for the corresponding loss. Reid et al. (2015) introduced the *Generalized Aggregating Algorithm* (GAA), going beyond the AA. The GAA is parameterized by the choice of a convex function Φ on the simplex (entropy) and reduces to the AA when Φ is the Shannon entropy. The GAA can achieve fast rates for losses satisfying a certain mixability condition (Φ -mixability). In particular, when a loss ℓ is Φ -mixable, the GAA achieves a constant regret R_ℓ^Φ which depends jointly on the *generalized mixability constant* η_ℓ^Φ — the largest η such that ℓ is $(\frac{1}{\eta}\Phi)$ -mixable — and the divergence $D_\Phi(e_\theta, \mathbf{q})$, where $\mathbf{q} \in \Delta_k$ is a prior distribution over experts (indexed by $\theta \in [k]$) and e_θ is the θ 'th standard basis element (Reid et al., 2015). Characterizing when losses are Φ -mixable was left as an open problem.

At each prediction round, the GAA can be divided into two steps; a *substitution step* where the learner picks a prediction from a set specified by the Φ -mixability condition; and an *update step* where a new distribution \mathbf{q} over experts is computed depending on their performance on the new outcome of the environment and the previous distribution. The set of predictions specified by the Φ -mixability condition is non-empty when the loss is Φ -mixable. Interestingly, the update step of the GAA is exactly the *Mirror Descent Algorithm* (MDA) (Steinhardt and Liang, 2014; Orabona et al., 2015) which minimizes the weighted loss of experts. In fact, both the MDA and the GAA use a divergence “measure” D_Φ , generated by a convex function Φ , as a regularizer when updating the distribution \mathbf{q} .

Contributions. In this paper, we answer the questions presented by Reid et al. (2015) around the notion of *generalized mixability* using entropic duality. For an entropy Φ and a loss ℓ , we derive a necessary and sufficient condition (Theorems 15 and 16) for ℓ to be Φ -mixable. In particular, if ℓ and Φ satisfy some regularity conditions, then ℓ is Φ -mixable if and only if $\eta_\ell\Phi - S$ is convex on the simplex, where S is the Shannon entropy and η_ℓ the largest η such that ℓ is η -mixable (Vovk, 1998; van Erven et al., 2012).

We derive an explicit expression for η_ℓ^Φ (Corollary 18), and hence, for the regret bound of the GAA. This allows us to compare the regret bound R_ℓ^Φ of any entropy Φ (such that ℓ is Φ -mixable) with that of the Shannon entropy S . In this case, we show (Theorem 19) that $R_\ell^S \leq R_\ell^\Phi$; that is, the Shannon entropy achieves the lowest worst-case regret when using the GAA. This fact is similar to Vovk’s result regarding the fundamental nature of log-loss (Vovk, 2015). Nevertheless, by leveraging the connection between the GAA and MDA, we discuss possible modifications to the GAA to enable better regret bounds in practice.

In Section 2, we introduce notations used throughout the paper and give some background on key notions in convex analysis which will be crucial in our proofs. We also present loss functions considered in the paper and state some new results (Theorem 3 and 4) along with corrected proofs for old ones (Theorem 5). In Section 3, we introduce the notions of classical and generalized mixability and derive some useful properties and characterizations. In particular, we derive a necessary and sufficient condition for Φ -mixability (Theorems 15 and 16). We also elucidate the tight relationship between the GAA and the MDA first observed by Reid et al. (2015). We conclude this paper by a general discussion and direction for future work.

2. Preliminaries

In this section, we introduce the notation used throughout the paper and present some key concepts in convex analysis. We also define entropies and describe loss functions considered in this paper. Furthermore, we present some new results regarding the latter.

2.1. Notations

For $n \in \mathbb{N}$, we define $\tilde{n} = n - 1$. We denote $[n] := \{1, \dots, n\}$ the set of integers between 1 and n . Let $\langle \cdot, \cdot \rangle$ denote the standard inner product in \mathbb{R}^n and $\|\cdot\|$ the corresponding norm. Let I_n and $\mathbf{1}_n$ denote the $n \times n$ identity matrix and the vector of all ones in \mathbb{R}^n . Let e_1, \dots, e_n denote the *standard basis* for \mathbb{R}^n . For a set $\mathcal{I} \subsetneq \mathbb{N}$ and $\mathbf{r}_1, \dots, \mathbf{r}_n \in \mathbb{R}^k$, we denote $[\mathbf{r}_i]_{i \in \mathcal{I}} := [\mathbf{r}_{i_1}, \dots, \mathbf{r}_{i_k}] \in \mathbb{R}^{n \times k}$, where $\mathcal{I} = \{i_1, \dots, i_k\}$ and $i_1 \leq \dots \leq i_k$. We denote its transpose by $[\mathbf{r}_i]_{i \in \mathcal{I}}^\top \in \mathbb{R}^{k \times n}$. For two vectors $\mathbf{p}, \mathbf{q} \in \mathbb{R}^n$, we write $\mathbf{p} \leq \mathbf{q}$ [resp. $\mathbf{p} < \mathbf{q}$], if $\forall i \in [n], p_i \leq q_i$ [resp. $p_i < q_i$]. We also denote $\mathbf{p} \odot \mathbf{q} = [p_i q_i]_{1 \leq i \leq n}^\top \in \mathbb{R}^n$ the *Hadamard product* of \mathbf{p} and \mathbf{q} . If (\mathbf{c}_k) is a sequence of vectors in $\mathcal{C} \subseteq \mathbb{R}^n$, we simply write $(\mathbf{c}_k) \subset \mathcal{C}$. For a sequence $(\mathbf{v}_m) \subset \mathbb{R}^n$, we write $\mathbf{v}_m \xrightarrow{m \rightarrow \infty} \mathbf{v}$ or $\lim_{m \rightarrow \infty} \mathbf{v}_m = \mathbf{v}$, if $\forall i \in [n], \lim_{m \rightarrow \infty} [\mathbf{v}_m]_i = v_i$. For a square matrix $A \in \mathbb{R}^{n \times n}$, $\lambda_{\min}(A)$ [resp. $\lambda_{\max}(A)$] denote its minimum [resp. maximum] eigenvalue. For $k \geq 1$, $\mathbf{u} \in [0, +\infty[^k$ and $\mathbf{w} \in \mathbb{R}^k$, we define $\log \mathbf{u} := [\log u_i]_{1 \leq i \leq k}^\top \in \mathbb{R}^k$ and $\exp \mathbf{w} := [\exp w_i]_{1 \leq i \leq k}^\top \in \mathbb{R}^k$.

Let $\Delta_n := \{\mathbf{p} \in [0, +\infty[^n : \langle \mathbf{p}, \mathbf{1}_n \rangle = 1\}$ be the *probability simplex* in \mathbb{R}^n . We also define $\tilde{\Delta}_n := \{\tilde{\mathbf{p}} \in [0, +\infty[^{\tilde{n}} : \langle \tilde{\mathbf{p}}, \mathbf{1}_{\tilde{n}} \rangle \leq 1\}$. We will use the notations $\Delta_n^k := (\Delta_n)^k$ and $\tilde{\Delta}_n^k := (\tilde{\Delta}_n)^k$. For $\mathcal{I} \subseteq [n]$, the set $\Delta_{\mathcal{I}} = \{\mathbf{q} \in \Delta_n : q_i = 0, \forall i \in [n] \setminus \mathcal{I}\}$ is a $|\mathcal{I}|$ -*face* of Δ_n . We denote $\Pi_{\mathcal{I}}^n : \mathbb{R}^n \rightarrow \mathbb{R}^{|\mathcal{I}|}$ the linear projection operator satisfying $\Pi_{\mathcal{I}}^n \mathbf{u} = [u_i]_{i \in \mathcal{I}}^\top$. It is easy to verify that $\Pi_{\mathcal{I}}^n [\Pi_{\mathcal{I}}^n]^\top = I_{|\mathcal{I}|}$ and that $\mathbf{q} \mapsto \Pi_{\mathcal{I}}^n \mathbf{q}$ is a bijection from $\Delta_{\mathcal{I}}$ to $\Delta_{|\mathcal{I}|}$. In the special case where $\mathcal{I} = [\tilde{n}]$, we write $\Pi_n := \Pi_{[\tilde{n}]}$ and we define the affine operator $\Pi_n : \mathbb{R}^{\tilde{n}} \rightarrow \mathbb{R}^n$ by $\Pi_n(\mathbf{u}) := [u_1, \dots, u_{\tilde{n}}, 1 - \langle \mathbf{u}, \mathbf{1}_{\tilde{n}} \rangle]^\top = J_n \mathbf{u} + \mathbf{e}_n$, where $J_n := \begin{bmatrix} I_{\tilde{n}} \\ -\mathbf{1}_{\tilde{n}}^\top \end{bmatrix} \in \mathbb{R}^{n \times \tilde{n}}$.

For $\mathbf{u} \in \mathbb{R}^n$ and $c \in \mathbb{R}$, we denote $\mathcal{H}_{\mathbf{u}, c} := \{\mathbf{y} \in \mathbb{R}^n : \langle \mathbf{y}, \mathbf{u} \rangle \leq c\}$ and $\mathcal{B}(\mathbf{u}, c) := \{\mathbf{v} \in \mathbb{R}^n : \|\mathbf{u} - \mathbf{v}\| \leq c\}$. $\mathcal{H}_{\mathbf{u}, c}$ is a closed half space and $\mathcal{B}(\mathbf{u}, c)$ is the c -ball in \mathbb{R}^n centered at \mathbf{u} . Let $\mathcal{C} \subseteq \mathbb{R}^n$ be a non-empty set. We denote $\text{int } \mathcal{C}$, $\text{ri } \mathcal{C}$, $\text{bd } \mathcal{C}$, and $\text{rbd } \mathcal{C}$ the *interior*, *relative interior*, *boundary*, and *relative boundary* of a set $\mathcal{C} \in \mathbb{R}^n$, respectively (Hiriart-Urruty and Lemaréchal, 2001). We denote the *indicator function* of \mathcal{C} by $\iota_{\mathcal{C}}$, where for $\mathbf{u} \in \mathcal{C}$, $\iota_{\mathcal{C}}(\mathbf{u}) = 0$, otherwise $\iota_{\mathcal{C}}(\mathbf{u}) = +\infty$. The *support function* of \mathcal{C} is defined by

$$\sigma_{\mathcal{C}}(\mathbf{u}) := \sup_{\mathbf{s} \in \mathcal{C}} \langle \mathbf{u}, \mathbf{s} \rangle, \quad \mathbf{u} \in \mathbb{R}^n.$$

Let $f : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$. We denote $\text{dom } f := \{\mathbf{u} \in \mathbb{R}^n : f(\mathbf{u}) < +\infty\}$ the *effective domain* of f . The function f is *proper* if $\text{dom } f \neq \emptyset$. The function f is *convex* if $\forall (\mathbf{u}, \mathbf{v}) \in \mathbb{R}^n$ and $\lambda \in]0, 1[$, $f(\lambda \mathbf{u} + (1 - \lambda)\mathbf{v}) \leq \lambda f(\mathbf{u}) + (1 - \lambda)f(\mathbf{v})$. When the latter inequality is strict for all $\mathbf{u} \neq \mathbf{v}$, f is *strictly convex*. When f is convex, it is *closed* if it is *lower semi-continuous*; that is, for all $\mathbf{u} \in \mathbb{R}^n$, $\liminf_{\mathbf{v} \rightarrow \mathbf{u}} f(\mathbf{v}) \geq f(\mathbf{u})$. The function f is said to be *1-homogeneous* if $\forall (\mathbf{u}, \alpha) \in \mathbb{R}^n \times]0, +\infty[$, $f(\alpha \mathbf{u}) = \alpha f(\mathbf{u})$, and it is said to be *1-coercive* if $\frac{f(\mathbf{u})}{\|\mathbf{u}\|} \rightarrow +\infty$ as $\|\mathbf{u}\| \rightarrow \infty$. Let f be proper. The *sub-differential* of f is defined by

$$\partial f(\mathbf{u}) := \{\mathbf{s}^* \in \mathbb{R}^n : f(\mathbf{v}) \geq f(\mathbf{u}) + \langle \mathbf{s}^*, \mathbf{v} - \mathbf{u} \rangle, \forall \mathbf{v} \in \mathbb{R}^n\}.$$

Any element $s \in \partial f(\mathbf{u})$ is called a *sub-gradient* of f at \mathbf{u} . We say that f is *directionally differentiable* if for all $(\mathbf{u}, \mathbf{v}) \in \text{dom } f \times \mathbb{R}^n$ the limit $\lim_{t \downarrow 0} \frac{f(\mathbf{u}+t\mathbf{v})-f(\mathbf{u})}{t}$ exists in $[-\infty, +\infty]$. In this case, we denote the limit by $f'(\mathbf{u}; \mathbf{v})$. When f is convex, it is directionally differentiable (Rockafellar, 1997). Let f be proper and directionally differentiable. The *divergence* generated by f is the map $D_f: \mathbb{R}^n \times \text{dom } f \rightarrow [0, +\infty]$ defined by

$$D_f(\mathbf{v}, \mathbf{u}) := \begin{cases} f(\mathbf{v}) - f(\mathbf{u}) - f'(\mathbf{u}; \mathbf{v} - \mathbf{u}), & \text{if } \mathbf{v} \in \text{dom } f; \\ +\infty, & \text{otherwise.} \end{cases}$$

For $\mathcal{I} \subset [n]$ and $f^{\mathcal{I}} := f \circ [\Pi_{\mathcal{I}}^n]^\top$, it is easy to verify that $(f^{\mathcal{I}})'(\Pi_{\mathcal{I}}^n \mathbf{p}; \Pi_{\mathcal{I}}^n \mathbf{q} - \Pi_{\mathcal{I}}^n \mathbf{p}) = f'(\mathbf{p}; \mathbf{q} - \mathbf{p})$, $\forall (\mathbf{p}, \mathbf{q}) \in \Delta_{\mathcal{I}}$. In this case, it holds that $D_f(\mathbf{q}, \mathbf{p}) = D_{f^{\mathcal{I}}}(\Pi_{\mathcal{I}}^n \mathbf{q}, \Pi_{\mathcal{I}}^n \mathbf{p})$. If f is differentiable [resp. twice differentiable] at $\mathbf{u} \in \text{int dom } f$, we denote $\nabla f(\mathbf{u}) \in \mathbb{R}^n$ [resp. $\text{H}f(\mathbf{u}) \in \mathbb{R}^{n \times n}$] its *gradient* vector [resp. *Hessian* matrix] at \mathbf{u} . A vector-valued function $g: \mathbb{R}^n \rightarrow \mathbb{R}^m$ is differentiable at \mathbf{u} if for all $i \in [m]$, g_i is differentiable at \mathbf{u} . In this case, the *differential* of g at \mathbf{u} is the linear operator $\text{D}g(\mathbf{u}): \mathbb{R}^n \rightarrow \mathbb{R}^m$ defined by $\text{D}g(\mathbf{u}) := [\nabla g_i(\mathbf{u})]_{1 \leq i \leq m}^\top$. If f has k continuous derivatives on a set $\Omega \subset \mathbb{R}^k$, we write $f \in C^k(\Omega)$.

We define $\tilde{f}: \mathbb{R}^{\tilde{n}} \rightarrow \mathbb{R} \cup \{+\infty\}$ by $\tilde{f} := f \circ \Pi_n + \iota_{\tilde{\Delta}_n}$. That is,

$$\tilde{f}(\tilde{\mathbf{u}}) := \begin{cases} f(J_n \tilde{\mathbf{u}} + \mathbf{e}_n), & \text{for } \tilde{\mathbf{u}} \in \tilde{\Delta}_n; \\ +\infty, & \text{for } \tilde{\mathbf{u}} \in \mathbb{R}^{n-1} \setminus \tilde{\Delta}_n. \end{cases} \quad (1)$$

If \tilde{f} is directionally differentiable, then $f'(\mathbf{p}, \mathbf{q} - \mathbf{p}) = \tilde{f}'(\tilde{\mathbf{p}}, \tilde{\mathbf{q}} - \tilde{\mathbf{p}})$, for $\mathbf{p}, \mathbf{q} \in \Delta_n$. If \tilde{f} is differentiable at $\tilde{\mathbf{p}} = \Pi_n(\mathbf{p})$, then $\tilde{f}'(\tilde{\mathbf{p}}, \tilde{\mathbf{q}} - \tilde{\mathbf{p}}) = \langle \nabla \tilde{f}(\tilde{\mathbf{p}}), \tilde{\mathbf{q}} - \tilde{\mathbf{p}} \rangle$. If, additionally, f is differentiable at $\mathbf{p} \in \text{ri } \Delta_k$, the chain rule yields $\nabla \tilde{f}(\tilde{\mathbf{p}}) = J_n^\top \nabla f(\mathbf{p})$. Since $J_n(\tilde{\mathbf{p}} - \tilde{\mathbf{q}}) = \Pi_n(\tilde{\mathbf{p}} - \tilde{\mathbf{q}}) = \mathbf{p} - \mathbf{q}$, it also follows that $\langle \tilde{\mathbf{p}} - \tilde{\mathbf{q}}, \nabla \tilde{f}(\tilde{\mathbf{p}}) \rangle = \langle \mathbf{p} - \mathbf{q}, \nabla f(\mathbf{p}) \rangle$.

The *Fenchel dual* of a (proper) function f is defined by $f^*(\mathbf{v}) := \sup_{\mathbf{u} \in \text{dom } f} \langle \mathbf{u}, \mathbf{v} \rangle - f(\mathbf{u})$, and it is a closed, convex function on \mathbb{R}^n (Hiriart-Urruty and Lemaréchal, 2001). Some useful properties of the Fenchel dual are given in Appendix A.

2.2. Entropies on the simplex

A function $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ is an *entropy* if it is closed, convex, and $\Delta_k \subseteq \text{dom } \Phi$. Its *entropic dual* $\Phi^*: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ is defined by $\Phi^*(\mathbf{z}) := \sup_{\mathbf{q} \in \Delta_k} \langle \mathbf{q}, \mathbf{z} \rangle - \Phi(\mathbf{q})$, $\mathbf{z} \in \mathbb{R}^k$. For the remainder of this paper, we consider entropies defined on \mathbb{R}^k , where $k \geq 2$.

Let $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ be an entropy and $\Phi_\Delta := \Phi + \iota_{\Delta_k}$. In this case, $\Phi^* = \Phi_\Delta^*$. It is clear that Φ_Δ is 1-coercive, and therefore, $\text{dom } \Phi^* = \text{dom } \Phi_\Delta^* = \mathbb{R}^k$ (Hiriart-Urruty and Lemaréchal, 2001, Prop. E.1.3.8). The entropic dual of Φ can also be expressed using the Fenchel dual of $\tilde{\Phi}: \mathbb{R}^{k-1} \rightarrow \mathbb{R} \cup \{+\infty\}$ defined by (1) after substituting f by Φ and n by k . In fact,

$$\begin{aligned} \Phi^*(\mathbf{z}) &= \sup_{\tilde{\mathbf{q}} \in \tilde{\Delta}_k} \langle J_k \tilde{\mathbf{q}} + \mathbf{e}_k, \mathbf{z} \rangle - \Phi(J_k \tilde{\mathbf{q}} + \mathbf{e}_k), \\ &= \langle \mathbf{e}_k, \mathbf{z} \rangle + \sup_{\tilde{\mathbf{q}} \in \tilde{\Delta}_k} \left\langle \tilde{\mathbf{q}}, J_k^\top \mathbf{z} \right\rangle - \tilde{\Phi}(\tilde{\mathbf{q}}), \\ &= \langle \mathbf{e}_k, \mathbf{z} \rangle + \tilde{\Phi}^*(J_k^\top \mathbf{z}), \end{aligned} \quad (2)$$

where (2) follows from the fact that $\text{dom } \tilde{\Phi} = \tilde{\Delta}_k$. Note that when Φ is an entropy, $\tilde{\Phi}$ is a closed convex function on \mathbb{R}^{k-1} . Hence, it holds that $\tilde{\Phi}^{**} = \tilde{\Phi}$ (Rockafellar, 1997).

The *Shannon entropy* by $S(\mathbf{q}) := \sum_{i \in [k]: q_i \neq 0} q_i \log q_i$,¹ if $\mathbf{q} \in [0, +\infty[^k$; and $+\infty$ otherwise. In the next proposition, we give the expressions of the entropic dual of S as well as the Fenchel dual \tilde{S} . The proof is in Appendix C.1.

Proposition 1 *For the Shannon entropy S , it holds that $\tilde{S}^*(\mathbf{v}) = \log(\langle \exp(\mathbf{v}), \mathbf{1}_k \rangle + 1)$, $\forall \mathbf{v} \in \mathbb{R}^{k-1}$, and $S^*(\mathbf{z}) = \log(\langle \exp(\mathbf{z}), \mathbf{1}_k \rangle)$, $\forall \mathbf{z} \in \mathbb{R}^k$.*

2.3. Loss Functions

In general, a loss function is a map $\ell: \mathcal{X} \times \mathcal{A} \rightarrow [0, +\infty]$ where \mathcal{X} is an outcome set and \mathcal{A} is a prediction set. In this paper, we consider the case of n possible outcomes; $\mathcal{X} = [n]$. Overloading notation slightly, we define the mapping $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$ by $[\ell(\mathbf{a})]_x = \ell(x, \mathbf{a})$, $\forall x \in [n]$ and denote $\ell_x(\cdot) := [\ell(\cdot)]_x$. We further extend the new definition of ℓ to the set $\bigcup_{k \geq 1} \mathcal{A}^k$ such that for $x \in [n]$ and $A := [\mathbf{a}_\theta]_{\theta \in [k]}^\top \in \mathcal{A}^k$, we have $\ell_x(A) := [\ell_x(\mathbf{a}_\theta)]_{\theta \in [k]}^\top \in [0, +\infty]^k$. We define the *effective domain* of ℓ by $\text{dom } \ell := \{\mathbf{a} \in \mathcal{A} : \ell(\mathbf{a}) \in [0, +\infty]^n\}$, and the *loss surface* by $\mathcal{S}_\ell := \{\ell(\mathbf{a}) : \mathbf{a} \in \text{dom } \ell\}$. We say that ℓ is *closed* if \mathcal{S}_ℓ is closed. The *superprediction* set of ℓ is defined by $\mathcal{S}_\ell^\oplus := \mathcal{S}_\ell \oplus [0, +\infty]^n$, where \oplus denotes the Minkowski sum. The η -*exponentiated* superprediction set of ℓ is defined by $\exp(-\eta \mathcal{S}_\ell^\oplus) := \{\exp(-\eta \mathbf{s}) : \mathbf{s} \in \mathcal{S}_\ell^\oplus\}$. Let $\mathbf{a}_0, \mathbf{a}_1 \in \mathcal{A}$. The prediction \mathbf{a}_0 is said to be *better* than \mathbf{a}_1 if $\ell(\mathbf{a}_0) \leq \ell(\mathbf{a}_1)$ and for some $x \in [n]$, $\ell_x(\mathbf{a}_0) < \ell_x(\mathbf{a}_1)$ (Williamson et al., 2016). A loss ℓ is *admissible* if for any $\mathbf{a} \in \mathcal{A}$ there are no better predictions. A loss ℓ is said to be *trivial* if there exists a prediction $\mathbf{a} \in \mathcal{A}$ which minimizes ℓ over all outcomes $x \in [n]$. Thus, a closed loss is *non-trivial* if there exists $(x_0, x_1, \mathbf{a}_0, \mathbf{a}_1) \in [n] \times [n] \times \mathcal{A} \times \mathcal{A}$ such that

$$\mathbf{a}_1 \in \text{argmin}\{\ell_{x_0}(\mathbf{a}) : \ell_{x_1}(\mathbf{a}) = \inf_{\hat{\mathbf{a}} \in \mathcal{A}} \ell_{x_1}(\hat{\mathbf{a}})\} \text{ and } \inf_{\mathbf{a} \in \mathcal{A}} \ell_{x_0}(\mathbf{a}) = \ell_{x_0}(\mathbf{a}_0) < \ell_{x_0}(\mathbf{a}_1). \quad (3)$$

In the rest of this paper, we will often make use of the following condition on losses.

Condition I ℓ is a closed, admissible, and non-trivial loss such that $\text{dom } \ell \neq \emptyset$.

In this paragraph let $\mathcal{A} = \Delta_n$. We define the *conditional risk* $L_\ell: \Delta_n \times \Delta_n \rightarrow \mathbb{R}$ by $L_\ell(\mathbf{p}, \mathbf{q}) = \mathbb{E}_{x \sim \mathbf{p}}[\ell_x(\mathbf{q})] = \langle \mathbf{p}, \ell(\mathbf{q}) \rangle$ and the *Bayes risk* by $\underline{L}_\ell(\mathbf{p}) := \inf_{\mathbf{q} \in \Delta_n} L_\ell(\mathbf{p}, \mathbf{q})$. In this case, the loss ℓ is *proper* [resp. *strictly proper*] if $\underline{L}_\ell(\mathbf{p}) = \langle \mathbf{p}, \ell(\mathbf{p}) \rangle \leq \langle \mathbf{p}, \ell(\mathbf{q}) \rangle$ [resp. $\langle \mathbf{p}, \ell(\mathbf{p}) \rangle < \langle \mathbf{p}, \ell(\mathbf{q}) \rangle$] for all $\mathbf{p} \neq \mathbf{q}$ in Δ_n . For example, the *log-loss* $\ell_{\log}: \Delta_n \rightarrow [0, +\infty]^n$ is defined by $\ell_{\log}(\mathbf{p}) = -\log \mathbf{p}$, where the “log” of a vector applies coordinate-wise. One can easily check that ℓ_{\log} is strictly proper. We denote \underline{L}_{\log} the corresponding Bayes risk.

The above definition of the Bayes risk is restricted to losses defined on the simplex. For a loss $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$, we use the following definition of the Bayes risk.

Definition 2 (Bayes Risk (Williamson, 2014)) *Let $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$ be a loss such that $\mathcal{S}_\ell^\oplus \neq \emptyset$. The Bayes risk $\underline{L}_\ell: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{-\infty\}$ is defined by*

$$\forall \mathbf{u} \in \mathbb{R}^n, \quad \underline{L}_\ell(\mathbf{u}) := \inf_{\mathbf{z} \in \mathcal{S}_\ell^\oplus} \langle \mathbf{u}, \mathbf{z} \rangle = -\sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{u}).$$

1. The Shannon entropy is usually defined with a minus sign. However, it will be more convenient for us to work without it.

Note that since \mathcal{S}_ℓ^\oplus is a non-empty subset of $[0, +\infty]^n$, $\mathbf{u} \mapsto \sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{u})$ is convex and finite on $[0, +\infty]^n$; it is bounded from above by 0. In particular, $\partial\sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{p}) \neq \emptyset$ for all $\mathbf{p} \in \text{ri } \Delta_n \subset \text{int dom } \sigma_{\mathcal{S}_\ell^\oplus} =]0, +\infty[^n$ (Rockafellar, 1997, Thm. 23.4).

We call $\underline{\ell} : \Delta_n \rightarrow [0, +\infty]^n$ a *support loss* of ℓ if $\forall \mathbf{p} \in \text{ri } \Delta_n$, $\underline{\ell}(\mathbf{p}) \in \partial\sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{p})$, and $\forall \mathbf{p} \in \text{rbd } \Delta_n$, there exists a sequence $(\mathbf{p}_m) \subset \text{ri } \Delta_n$, such that $\underline{\ell}(\mathbf{p}_m) \xrightarrow{m \rightarrow \infty} \underline{\ell}(\mathbf{p})$.

Theorem 3 *Any loss $\ell : \mathcal{A} \rightarrow [0, +\infty]^n$ such that $\text{dom } \ell \neq \emptyset$, has a proper support loss $\underline{\ell}$ whose Bayes risk is \underline{L}_ℓ .*

The proof of the proposition is in Appendix C.2. Note that when the Bayes risk is differentiable on $]0, +\infty[^n$, or equivalently, when $\mathbf{u} \mapsto \sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{u})$ is differentiable on $]0, +\infty[^n$, the support loss of ℓ is uniquely defined on $\text{ri } \Delta_n$. This is because when $\sigma_{\mathcal{S}_\ell^\oplus}$ is differentiable at $-\mathbf{p}$, $\partial\sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{p}) = \{\nabla\sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{p})\} = \{\underline{\ell}(\mathbf{p})\}$ (Hiriart-Urruty and Lemaréchal, 2001, Cor. D.2.1.4).

Theorem 4 *Let $\ell : \mathcal{A} \rightarrow [0, +\infty]^n$ be a loss satisfying Condition I and $\underline{\ell}$ a proper support loss of ℓ . If \underline{L}_ℓ is not differentiable at \mathbf{p} then there exist $\mathbf{a}_0, \mathbf{a}_1 \in \text{dom } \ell$, such that $\ell(\mathbf{a}_0) \neq \ell(\mathbf{a}_1)$ and $\underline{L}_\ell(\mathbf{p}) = \langle \mathbf{p}, \underline{\ell}(\mathbf{p}) \rangle = \langle \mathbf{p}, \ell(\mathbf{a}_0) \rangle = \langle \mathbf{p}, \ell(\mathbf{a}_1) \rangle$.*

If the Bayes risk \underline{L}_ℓ is differentiable on $]0, +\infty[^n$, then

$$\begin{aligned} \forall \mathbf{p} \in \text{dom } \underline{\ell}, \quad \exists \mathbf{a}_* \in \text{dom } \ell, \quad \ell(\mathbf{a}_*) &= \underline{\ell}(\mathbf{p}), \\ \forall \mathbf{a} \in \text{dom } \ell, \quad \exists (\mathbf{p}_m) \subset \text{ri } \Delta_n, \quad \underline{\ell}(\mathbf{p}_m) &\xrightarrow{m \rightarrow \infty} \ell(\mathbf{a}). \end{aligned}$$

The proof is in Appendix C.3. Note that we always have $\text{ri } \Delta_n \subset \text{dom } \underline{\ell}$. Consistent with the notations introduced in the previous subsection, we write $\tilde{\ell} := \underline{\ell} \circ \Pi_n$ and $\tilde{L}_\ell := \underline{L}_\ell \circ \Pi_n$.

In the literature, many theoretical results involving loss functions relied on the fact that the superprediction set of a proper loss is convex (Williamson et al., 2016; Dawid, 2007). An earlier proof of this result by Williamson et al. (2016) was incomplete². In the next theorem, we restate this result and prove it in Appendix C.4.

Theorem 5 *If $\ell : \Delta_n \rightarrow [0, +\infty]^n$ is a continuous proper loss, then $\mathcal{S}_\ell^\oplus = \bigcap_{\mathbf{p} \in \Delta_n} \mathcal{H}_{-\mathbf{p}, -\underline{\ell}(\mathbf{p})}$. In particular, \mathcal{S}_ℓ^\oplus is convex.*

3. Mixability in the Game of Prediction with Expert Advice

Here, we consider the setting of prediction with expert advice (Vovk, 1998); there is pool of k experts, parameterized by $\theta \in [k]$, which make predictions $\mathbf{a}_\theta^t \in \mathcal{A}$ at each round t . In the same round, the learner predicts $\mathbf{a}_* \in F(A^t) \in \mathcal{A}$, where $A^t = [\mathbf{a}_\theta^t]_{1 \leq \theta \leq k} \in \mathcal{A}^k$ is the matrix of experts' predictions and F is set valued *aggregating* function on \mathcal{A}^k . At the end of the round, the outcome $x^t \in [n]$ is announced and each expert θ [resp. learner] suffers a loss $\ell_{x^t}(\mathbf{a}_\theta)$ [resp. $\ell_{x^t}(\mathbf{a}_*)$], where $\ell : \mathcal{A} \rightarrow [0, +\infty]^n$. We will refer to such a game by $\mathcal{G}_\ell(k, T)$.

Vovk (1998) introduced the *Aggregating Algorithm* (AA) which specifies the aggregating function F of the learner. When a loss ℓ is η -mixable (see below) the AA achieves *constant regret* in

2. It was claimed that if \mathcal{S}_ℓ^\oplus is non-convex, there exists a point \mathbf{s}_0 on the loss surface \mathcal{S}_ℓ such that no hyperplane supports \mathcal{S}_ℓ^\oplus at \mathbf{s}_0 . The non-convexity of a set by itself is not sufficient to make such a claim; the continuity of the loss ℓ is required.

the $\mathcal{G}_\ell(k, T)$ game; the difference between the cumulative loss of the learner and the best expert is upper bounded by a constant independent of the number of rounds T .

The *Generalized Aggregating Algorithm* (GAA) (Reid et al., 2015) uses an entropy function $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ to specify the learner's aggregating function F . The GAA also achieves a constant regret when a certain condition is satisfied (Φ -mixability). The GAA reduces to the AA when Φ is a scaled *Shannon entropy* S .

Definition 6 (η -mixability) For $\eta > 0$, a loss $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$ is said to be η -mixable, if $\forall \mathbf{q} \in \Delta_k, \forall A := [\mathbf{a}_\theta]_{\theta \in [k]} \in \mathcal{A}^k, \exists \mathbf{a}_* \in \mathcal{A}$, such that

$$\forall x \in [n], \quad \ell_x(\mathbf{a}_*) \leq -\eta^{-1} \log \langle \exp(-\eta \ell_x(A)), \mathbf{q} \rangle. \quad (4)$$

Chernov et al. (2010) showed that the η -mixability condition (4) is equivalent to the convexity of the η -exponentiated superprediction set $\exp(-\eta \mathcal{S}_\ell^\oplus)$. The largest η such that ℓ is η -mixable is denoted η_ℓ . If $\eta_\ell > 0$, we say that ℓ is *classically mixable*.

For a strictly proper loss $\ell: \Delta_n \rightarrow [0, +\infty]^n$ whose Bayes risk satisfies $\underline{L}_\ell \in C^2([0, +\infty]^n)$, van Erven et al. (2012) showed that the mixability constant η_ℓ is equal to

$$\underline{\eta}_\ell := \inf_{\tilde{\mathbf{p}} \in \text{int } \tilde{\Delta}_n} (\lambda_{\max}([\mathbf{H}\tilde{L}_{\log}(\tilde{\mathbf{p}})]^{-1} \mathbf{H}\tilde{L}_\ell(\tilde{\mathbf{p}})))^{-1}, \quad (5)$$

The next theorem extends this result by showing that the mixability constant η_ℓ of any loss $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$ satisfying condition I and such that \underline{L}_ℓ is twice differentiable is lower bounded by $\underline{\eta}_\ell$.

Theorem 7 Let $\eta > 0$, $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$ a loss satisfying Condition I, and $\underline{\ell}$ a proper support loss of ℓ . Suppose that $\text{dom } \ell = \mathcal{A}$ and that \underline{L}_ℓ is twice differentiable on $]0, +\infty[^n$. Then for $\underline{\eta}_\ell$ as in (5), we have $\underline{\eta}_\ell > 0 \implies \ell$ is $\underline{\eta}_\ell$ -mixable. In particular, $\eta_\ell \geq \underline{\eta}_\ell$.

The proof of the theorem is in Appendix C.5. We will show later (Corollary 17) that, under the same conditions of Theorem 7, we actually have $\eta_\ell = \underline{\eta}_\ell$.

Definition 8 (Φ -mixability) ³ Let $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ be an entropy. A loss $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$ is Φ -mixable if $\forall \mathbf{q} \in \Delta_k, \forall A := [\mathbf{a}_\theta]_{1 \leq \theta \leq k} \in \mathcal{A}^k, \exists \mathbf{a}_* \in \mathcal{A}$, such that

$$\forall x \in [n], \quad \ell_x(\mathbf{a}_*) \leq M_\Phi(\ell_x(A), \mathbf{q}) := \inf_{\hat{\mathbf{q}} \in \Delta_k} \langle \hat{\mathbf{q}}, \ell_x(A) \rangle + D_\Phi(\hat{\mathbf{q}}, \mathbf{q}). \quad (6)$$

As stated earlier, when ℓ is $(\frac{1}{\eta}\Phi)$ -mixable, the GAA (see Figure 1) can achieve a constant regret in the $\mathcal{G}_\ell(k, T)$ game, where the constant does not depend on T . In fact, Reid et al. (2015) showed that for all $T \geq 1$ and $\theta \in [k]$,

$$\sum_{1 \leq i \leq T} \ell_{x^i}(\mathbf{a}_*^i) \leq \sum_{1 \leq i \leq T} \ell_{x^i}(\mathbf{a}_\theta^i) + R_\ell^\Phi,$$

where $R_\ell^\Phi := \inf_{\mathbf{q} \in \Delta_k} \max_{\theta \in [k]} D_\Phi(\mathbf{e}_\theta, \mathbf{q})/\eta_\ell^\Phi$ is the worst-case regret, η_ℓ^Φ is the generalized mixability constant (defined formally in Corollary 18), and (\mathbf{a}_*^i) are the outputs of the GAA with initial distribution over experts $\mathbf{q}^0 = \text{argmin}_{\mathbf{q} \in \Delta_k} \max_{\theta \in [k]} D_\Phi(\mathbf{e}_\theta, \mathbf{q})$.

3. Our definition of Φ -mixability is slightly different than that introduced by Reid et al. (2015); we use the directional derivative to define the divergence D_Φ .

<p>Aggregating Algorithm</p> <p>Input: $\mathbf{q}^0 \in \Delta_k, \eta \leq \eta_\ell, T \geq 1$.</p> <p>for $t = 1$ to T</p> <p> Observe $A^t = [\mathbf{a}_\theta^t]_{1 \leq \theta \leq k} \in \mathcal{A}^k$.</p> <p> Find $\mathbf{a}_*^t \in \mathcal{A}$ such that $\forall x \in [n]$,</p> <p> $\ell_x(\mathbf{a}_*^t) \leq -\eta^{-1} \log(\langle \exp[-\eta \ell_x(A)], \mathbf{q}^{t-1} \rangle)$.</p> <p> Observe outcome $x^t \in [n]$.</p> <p> Set $q_\theta^t = \frac{q_\theta^{t-1} \exp(-\eta \ell_{x^t}(\mathbf{a}_\theta^t))}{\sum_{\theta \in [k]} q_\theta^{t-1} \exp(-\eta \ell_{x^t}(\mathbf{a}_\theta^t))}, \forall \theta \in [k]$.</p> <p>end for</p>	<p>Generalized Aggregating Algorithm</p> <p>Input: $\mathbf{q}^0 \in \Delta_k, \eta \leq \eta_\ell^\Phi, T \geq 1$.</p> <p>for $t = 1$ to T</p> <p> Observe $A^t = [\mathbf{a}_\theta^t]_{1 \leq \theta \leq k} \in \mathcal{A}^k$.</p> <p> Find $\mathbf{a}_*^t \in \mathcal{A}$ such that $\forall x \in [n]$,</p> <p> $\ell_x(\mathbf{a}_*^t) \leq M_{\Phi_\eta}(\ell_x(A^t), \mathbf{q}^{t-1})$.</p> <p> Observe outcome $x^t \in [n]$.</p> <p> Set $\mathbf{q}^t = \operatorname{argmin}_{\mu \in \Delta_k} \langle \mu, \ell_{x^t}(A^t) \rangle + D_\Phi(\mu, \mathbf{q}^{t-1})$.</p> <p>end for</p>
---	--

Figure 1: The AA versus the GAA in the $\mathcal{G}_\ell(k, T)$ game. A^t and x^t are the experts' predictions and environment outcome at round t , respectively. \mathbf{q}^0 is a prior distribution over experts and η is the learning rate. We used the notation $\Phi_\eta := \eta^{-1}\Phi$. When ℓ is η -mixable [resp. Φ -mixable] the AA [resp. GAA] achieves constant regret using the learner's predictions $(\mathbf{a}_*^t)_{1 \leq t \leq T}$, which are guaranteed to exist. The GAA reduces to the AA when $\Phi = \eta^{-1}S$ and S is the Shannon entropy.

Remark 9 *In order for the update distribution \mathbf{q}^t of the GAA (Figure 1) to be well defined, the infimum of the map $\mu \mapsto \langle \mu, \mathbf{d} \rangle + D_\Phi(\mu, \mathbf{q})$ must be attained at some $\mathbf{q}_* \in \Delta_k$ for any given $(\mathbf{d}, \mathbf{q}) \in [0, +\infty]^n \times \Delta_k$. We verify this for $\mathbf{q} \in \operatorname{ri} \Delta_k$; since $\tilde{\mathbf{q}} \in \operatorname{int} \operatorname{dom} \tilde{\Phi} = \operatorname{int} \tilde{\Delta}_k$, the function $\tilde{\mu} \mapsto -\tilde{\Phi}'(\tilde{\mathbf{q}}; \tilde{\mu} - \tilde{\mathbf{q}})$ is lower semicontinuous (Rockafellar, 1997, Cor. 24.5.1). Given that $\tilde{\mu} \mapsto \langle \Pi_k(\tilde{\mu}), \mathbf{d} \rangle + \tilde{\Phi}(\tilde{\mu}) - \tilde{\Phi}(\tilde{\mathbf{q}})$ is a closed convex function, it is also lower semi-continuous. Therefore, the function $\tilde{\mu} \mapsto \langle \Pi_k(\tilde{\mu}), \mathbf{d} \rangle + \tilde{\Phi}(\tilde{\mu}) - \tilde{\Phi}(\tilde{\mathbf{q}}) - \tilde{\Phi}'(\tilde{\mathbf{q}}; \tilde{\mu} - \tilde{\mathbf{q}})$ is lower semicontinuous, and thus attains its minimum on the compact set $\tilde{\Delta}_k$ at some point $\tilde{\mathbf{q}}_*$ (Holder, 2005, Thm. 1.13). Using the fact that $D_\Phi(\mu, \mathbf{q}) = D_{\tilde{\Phi}}(\tilde{\mu}, \tilde{\mathbf{q}})$, we get that $\mathbf{q}_* := \Pi_k(\tilde{\mathbf{q}}_*) = \operatorname{argmin}_{\mu \in \Delta_k} \langle \mu, \mathbf{d} \rangle + D_\Phi(\mu, \mathbf{q})$.*

When \mathbf{q} in Remark 9 is in $\operatorname{ri} \Delta_k$, then either \mathbf{q} is a vertex or there exists $\mathcal{I} \subset [k]$, with $|\mathcal{I}| > 1$, such that $\mathbf{q} \in \operatorname{ri} \Delta_{\mathcal{I}}$. In either case, if we require the existence of a Φ -mixable loss one can use Proposition 11 below and $\Phi^{\mathcal{I}} := \Phi \circ [\Pi_{\mathcal{I}}^k]^T$ to show that $\mu \mapsto \langle \mu, \mathbf{d} \rangle + D_\Phi(\mu, \mathbf{q})$ attains a minimum in $\Delta_{\mathcal{I}}$ by following the steps of Remark 9. In this case, the GAA's update step is well defined.

From Figure 1, it is clear that the GAA is divided into two steps; 1) a *substitution step* which consists of finding the prediction $\mathbf{a}_* \in \mathcal{A}$ satisfying the mixability condition (5); and 2) an *update step* where a new distribution over experts is computed. Note that the substitution step is not well defined in the sense that there is not a unique choice of \mathbf{a}_* . One systematic way of choosing \mathbf{a}_* is through inverse losses serving as *substitution functions* (Williamson, 2014). Kamalaruban et al. (2015) discuss other alternatives depending on the curvature of the Bayes risk. Note, however, that the choice of the substitution function does not affect the regret bound of the GAA when the loss ℓ is Φ -mixable. On the other hand, the update step is well defined and we show next that this is exactly the MDA for a certain sequence of losses.

Example 10 (GAA versus MDA) *Let $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$ be a loss and $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ an entropy such that $\tilde{\Phi}$ is differentiable on $\operatorname{int} \tilde{\Delta}_k$. Let \mathbf{q}^t be the update distribution of the GAA at round t and $\tilde{\mathbf{q}}^t = \Pi_k(\mathbf{q}^t)$. From Remark 9, it holds that*

$$\begin{aligned}
 \tilde{\mathbf{q}}^t &= \operatorname{argmin}_{\tilde{\mathbf{q}} \in \tilde{\Delta}_k} \langle \Pi_k(\tilde{\mathbf{q}}), \ell_{x^t}(A^t) \rangle + \eta^{-1} D_{\tilde{\Phi}}(\tilde{\mathbf{q}}, \tilde{\mathbf{q}}^{t-1}), \\
 &= \operatorname{argmin}_{\tilde{\mathbf{q}} \in \tilde{\Delta}_k} \left\langle \tilde{\mathbf{q}}, J_k^\top \ell_{x^t}(A^t) \right\rangle + \eta^{-1} D_{\tilde{\Phi}}(\tilde{\mathbf{q}}, \tilde{\mathbf{q}}^{t-1}), \\
 &= \operatorname{argmin}_{\tilde{\mathbf{q}} \in \tilde{\Delta}_k} \langle \tilde{\mathbf{q}}, \nabla l_t(\tilde{\mathbf{q}}^{t-1}) \rangle + \eta^{-1} D_{\tilde{\Phi}}(\tilde{\mathbf{q}}, \tilde{\mathbf{q}}^{t-1}), \tag{7}
 \end{aligned}$$

where $l_t(\tilde{\boldsymbol{\mu}}) := \langle \Pi_k(\tilde{\boldsymbol{\mu}}), \ell_{x^t}(A^t) \rangle = \langle \boldsymbol{\mu}, \ell_{x^t}(A^t) \rangle$. Update (7) is, by definition (Beck and Teboulle, 2003), the MDA with the sequence of losses l_t on $\operatorname{int} \tilde{\Delta}_k$, “distance” function $D_{\tilde{\Phi}}(\cdot, \cdot)$, and learning rate η . Therefore, the MDA is exactly the update step of the GAA.

3.1. Useful Properties

Given a differentiable entropy on Δ_k , Reid et al. (2015) showed that for a non-trivial loss to be Φ -mixable, Φ must have gradients whose norms diverge to infinity near the relative boundary of the simplex. However, the proof of this result assumed differentiability of Φ on the boundary of Δ_k , which is imprecise given that in their definition, Φ is a function on Δ_k — a set with empty interior. In the next proposition, we show that a variant of their result holds when working with directional derivatives. This will be crucial for subsequent results. Before stating the next proposition, note that given an entropy $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ and $\mathcal{I} \subseteq [k]$ with $|\mathcal{I}| > 1$, the function $\Phi^\mathcal{I} := \Phi \circ [\Pi_{\mathcal{I}}^k]^\top: \mathbb{R}^{|\mathcal{I}|} \rightarrow \mathbb{R} \cup \{+\infty\}$ is also an entropy.

Proposition 11 *Let $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ be an entropy and $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$ a closed non-trivial loss. If ℓ is Φ -mixable, then $\forall \mathcal{I} \subseteq [k]$ with $|\mathcal{I}| > 1$, ℓ is $\Phi^\mathcal{I}$ -mixable and*

$$\forall \mathbf{q} \in \operatorname{rbd} \Delta_\mathcal{I}, \forall \hat{\mathbf{q}} \in \operatorname{ri} \Delta_\mathcal{I}, \Phi'(\mathbf{q}; \hat{\mathbf{q}} - \mathbf{q}) = -\infty. \tag{8}$$

The proof of Proposition 11 is in Appendix C.6. Contrary to what was claimed previously (Reid et al., 2015), condition (8) together with the strict convexity and differentiability of Φ on $\operatorname{ri} \Delta_k$ is not sufficient for the existence of a Φ -mixable loss (see Appendix D for a counter-example).

In the next proposition, we show that the Bayes risk and the Fenchel dual of $\tilde{\Phi}$ need to be differentiable in the interior of their respective domains for the existence of a Φ mixable loss. As we will argue later (Lemma 35), the differentiability of $\tilde{\Phi}^*$ implies the strict convexity of Φ on $\operatorname{ri} \Delta_k$. The proof of the next proposition is in Appendix C.7.

Proposition 12 *Let $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ be an entropy and $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$ a loss satisfying Condition 1. If ℓ is Φ -mixable, then the Bayes risk satisfies $\underline{L}_\ell \in C^1(]0, +\infty[^n)$. If, additionally, \underline{L}_ℓ is twice differentiable on $]0, +\infty[^n$, then $\tilde{\Phi}^* \in C^1(\mathbb{R}^{k-1})$.*

In the next lemma (proved in Appendix C.8), we provide a useful expression of $M_\Phi(\mathbf{d}, \mathbf{q})$, which will be used to prove a sufficient condition for Φ -mixability (Theorem 15).

Proposition 13 *Let $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ be an entropy. Let $x \in [n]$, $\mathbf{d} \in \mathbb{R}^k$, $\mathbf{q} \in \operatorname{ri} \Delta_k$, and $\mathbf{q}_* = \operatorname{argmin}_{\boldsymbol{\mu} \in \Delta_k} \langle \boldsymbol{\mu}, \mathbf{d} \rangle + D_\Phi(\boldsymbol{\mu}, \mathbf{q})$. Then any $\mathbf{s}_q^* \in \operatorname{argmax}\{\langle \mathbf{s}, \tilde{\mathbf{q}}_* - \tilde{\mathbf{q}} \rangle : \mathbf{s} \in \partial \tilde{\Phi}(\tilde{\mathbf{q}})\}$ satisfies*

$$\tilde{\mathbf{q}}_* \in \partial \tilde{\Phi}^*(\mathbf{s}_q^* - J_k^\top \mathbf{d}), \tag{9}$$

$$M_\Phi(\mathbf{d}, \mathbf{q}) = d_k + \tilde{\Phi}^*(\mathbf{s}_q^*) - \tilde{\Phi}^*(\mathbf{s}_q^* - J_k^\top \mathbf{d}). \tag{10}$$

3.2. The generalized aggregating algorithm using the Shannon entropy S

The purpose of this subsection is to show that the GAA reduces to the AA when the former uses the Shannon entropy. In this case, generalized and classical mixability are equivalent.

Let $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$ be a loss and Φ be as in Proposition 13 and suppose that Φ and $\tilde{\Phi}^*$ are differentiable on $\text{ri } \Delta_k$ and \mathbb{R}^{k-1} , respectively. If we substitute \mathbf{d} for $\ell_x(A)$ in (9) and (10), where $A \in (\text{dom } \ell)^k$, it is possible to show that

$$\nabla \Phi^*(\nabla \Phi(\mathbf{q}) - \ell_x(A)) = \underset{\boldsymbol{\mu} \in \Delta_k}{\text{argmin}} \langle \boldsymbol{\mu}, \ell_x(A) \rangle + D_{\Phi}(\boldsymbol{\mu}, \mathbf{q}), \quad (11)$$

$$M_{\Phi}(\ell_x(A), \mathbf{q}) = \Phi^*(\nabla \Phi(\mathbf{q})) - \Phi^*(\nabla \Phi(\mathbf{q}) - \ell_x(A)). \quad (12)$$

This result was already established by Reid et al. (2015).

Let $\mathbf{q} \in \text{ri } \Delta_k$. By definition of S, $\nabla S(\mathbf{q}) = \log \mathbf{q} + \mathbf{1}_k$, and due to Proposition 1, $S^*(\mathbf{z}) = \log \langle \exp \mathbf{z}, \mathbf{1}_k \rangle$, $\mathbf{z} \in \mathbb{R}^k$. Therefore, $\nabla S(\mathbf{q}) - \eta \ell_x(A) = \log(\exp(-\eta \ell_x(A)) \odot \mathbf{q}) + \mathbf{1}_k$ and $\nabla S^*(\mathbf{z}) = \frac{\exp \mathbf{z}}{\langle \exp \mathbf{z}, \mathbf{1}_k \rangle}$, $\forall (x, A) \in [n] \times (\text{dom } \ell)^k$. Thus,

$$\nabla S^*(\nabla S(\mathbf{q}) - \eta \ell_x(A)) = \frac{\exp(-\eta \ell_x(A)) \odot \mathbf{q}}{\langle \exp(-\eta \ell_x(A)), \mathbf{q} \rangle}. \quad (13)$$

Let $S_{\eta} := \eta^{-1} S$. Then $\nabla S = \eta \nabla S_{\eta}$ and $\forall \mathbf{z} \in \mathbb{R}^k$, $\nabla S_{\eta}^*(\mathbf{z}) = \nabla S^*(\eta \mathbf{z})$ (Reid et al., 2015).⁴ Then the left hand side of (13) can be written as $\nabla S_{\eta}^*(\nabla S_{\eta}(\mathbf{q}) - \ell_x(A))$. Using this fact, (11) and (13) show that the update distribution \mathbf{q}^t of the GAA (Figure 1) coincides with that of the AA after substituting \mathbf{q} , x , and A by \mathbf{q}^{t-1} , x^t , and $A^t := [\mathbf{a}_{\theta}]_{\theta \in [k]}$, respectively.

Now using the fact that $M_{S_{\eta}}(\ell_x(A), \mathbf{q}) = \eta^{-1} M_S(\eta \ell_x(A), \mathbf{q})$ (Reid et al., 2015) and (12), we get $M_{S_{\eta}}(\ell_x(A), \mathbf{q}) = \eta^{-1} [S^*(\nabla S(\mathbf{q})) - S^*(\nabla S(\mathbf{q}) - \eta \ell_x(A))] = -\eta^{-1} \log \langle \exp(-\eta \ell_x(A)), \mathbf{q} \rangle$. Suppose now that \mathbf{q} belongs to relative interior of a face $\Delta_{\mathcal{I}} \subsetneq \Delta_k$ with $\mathcal{I} \subsetneq [k]$. We can repeat the argument above for $S^{\mathcal{I}} := S \circ [\Pi_{\mathcal{I}}^k]^{\top}$ to show that

$$\begin{aligned} M_{S_{\eta}^{\mathcal{I}}}(\Pi_{\mathcal{I}}^k \ell_x(A), \Pi_{\mathcal{I}}^k \mathbf{q}) &= -\eta^{-1} \log \langle \exp(-\eta \Pi_{\mathcal{I}}^k \ell_x(A)), \Pi_{\mathcal{I}}^k \mathbf{q} \rangle, \\ &= -\eta^{-1} \log \langle \exp(-\eta \ell_x(A)), \mathbf{q} \rangle. \end{aligned} \quad (14)$$

Equation 14 will now allow us to prove the equivalence between classical and generalized mixability when using the Shannon entropy. The proof of Theorem 14 is in Appendix C.8.

Theorem 14 *Let $\eta > 0$ and $S_{\eta} := \eta^{-1} S$, where S is the Shannon entropy. For a loss $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$, ℓ is η -mixable if and only if ℓ is S_{η} -mixable.*

3.3. Necessary and Sufficient Conditions for Φ -Mixability

In this subsection, we show that given an entropy $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ and a loss $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$ satisfying certain regularity conditions, ℓ is Φ -mixable if and only if

$$\boxed{\eta \ell \Phi - S \text{ is convex on } \Delta_k.} \quad (15)$$

4. Reid et al. (2015) showed the equality $\nabla \Phi_{\eta}^*(\mathbf{u}) = \nabla \Phi^*(\eta \mathbf{u})$, $\forall \mathbf{u} \in \text{dom } \Phi^*$, for any entropy differentiable on Δ_k - not just for the Shannon Entropy.

Theorem 15 *Let $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ be an entropy and $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$ be a loss. If $\eta_\ell \Phi - S$ is convex on Δ_k , then ℓ is Φ -mixable.*

As consequence of Theorem 15 (whose proof is in Appendix C.10), if a loss ℓ is not classically mixable (that is, $\eta_\ell = 0$) it cannot be Φ -mixable for any entropy Φ . This is because $\eta_\ell \Phi - S = -S$ is not convex. The converse of Theorem 15 also holds under additional smoothness conditions on Φ and ℓ as stated in the next theorem. The proof is in Appendix C.11.

Theorem 16 *Let $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ be an entropy such that $\tilde{\Phi} := \Phi \circ \Pi_k$ is twice differentiable on $\text{int } \tilde{\Delta}_k$, and $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$ a loss satisfying Condition I and such that \underline{L}_ℓ is twice differentiable on $]0, +\infty[^n$. Then ℓ is Φ -mixable only if $\underline{\eta}_\ell \Phi - S$ is convex on Δ_k .*

The twice differentiability condition of \underline{L}_ℓ in Theorem 16 is not as strong as it may first appear; we showed in Proposition 12 that ℓ is Φ -mixable only if $\underline{L}_\ell \in C^1(]0, +\infty[^n)$. Furthermore, since \underline{L}_ℓ is concave, Alexandrov's Theorem (see e.g. (Borwein et al., 2010, Thm. 6.7)) guarantees that \underline{L}_ℓ is twice differentiable almost everywhere in $]0, +\infty[^n$. A version of Theorem 16 which does not assume the twice differentiability of the Bayes risk is given in Appendix D (Theorem 38).

Corollary 17 *Let ℓ and Φ be as in Theorem 16. If $\text{dom } \ell = \mathcal{A}$, then $\eta_\ell = \underline{\eta}_\ell$. Furthermore, ℓ is Φ -mixable if and only if $\underline{\eta}_\ell \Phi - S$ is convex on Δ_k .*

Proof We already know from Theorem 7 that $\underline{\eta}_\ell \leq \eta_\ell$. Suppose now that ℓ is classically mixable. By definition of η_ℓ , ℓ is η_ℓ -mixable, and thus from Theorem 14, ℓ is $(\eta_\ell^{-1} S)$ -mixable. Substituting Φ for $\eta_\ell^{-1} S$ in Theorem 16 implies that $(\underline{\eta}_\ell/\eta_\ell - 1)S$ is convex on $\text{ri } \Delta_k$. Consequently, $\eta_\ell \leq \underline{\eta}_\ell$, and hence $\eta_\ell = \underline{\eta}_\ell$. From this fact and Theorems 15 and 16, it follows that ℓ is Φ -mixable if and only if $\underline{\eta}_\ell \Phi - S$ is convex on Δ_k . ■

Corollary 17 suggests that when the Bayes risk is twice differentiable on the interior of its domain it contains all necessary information for the characterization of classical mixability.

Corollary 18 (The Generalized Mixability Constant) *Let ℓ and Φ be as in Theorem 16. If $\text{dom } \ell = \mathcal{A}$, then the largest $\eta \geq 0$ such that ℓ is Φ_η -mixable is given by*

$$\eta_\ell^\Phi = \underline{\eta}_\ell \inf_{\tilde{q} \in \text{int } \tilde{\Delta}_k} \lambda_{\min}(\mathbf{H}\tilde{\Phi}(\tilde{q})(\mathbf{H}\tilde{S}(\tilde{q}))^{-1}). \quad (16)$$

The proof is in Appendix C.12. Observe that if we substitute Φ by S in (14), we get $\eta_\ell^S = \underline{\eta}_\ell$.

4. Discussion and Future Work

So far, we showed that the Φ -mixability of losses satisfying Condition I are characterized by the convexity of $\eta^* \Phi - S$, where $\eta^* \in \{\eta_\ell, \underline{\eta}_\ell\}$ (see Theorems 15 and 16). As a consequence of this, and contrary to what was conjectured previously (Reid et al., 2015), the generalized mixability condition does not induce a correspondence between losses and entropies; for a given loss ℓ , there is no particular entropy Φ^ℓ — specific to the choice of ℓ — which minimizes the regret of the GAA. Rather, the Shannon entropy S minimizes the regret regardless of the choice of ℓ (see Theorem 19

for the precise statement). This is consistent with Vovk’s result regarding the fundamental nature of the log-loss (Vovk, 2015).

Nevertheless, given a loss ℓ and entropy Φ , the curvature of the loss surface \mathcal{S}_ℓ determines the maximum learning rate η_ℓ^Φ of the GAA; in other words, it determines for which scaled entropies $\Phi_\eta := \eta^{-1}\Phi$, ℓ is Φ_η -mixable. In fact, the curvature of \mathcal{S}_ℓ is linked to $\underline{\eta}_\ell$ through the Hessian of the Bayes risk (see Theorem 41 in Appendix E.1 for the precise statement). The magnitude of $\underline{\eta}_\ell$ determines the learning rate η_ℓ^Φ as per Corollary 18.

Given a loss ℓ satisfying Condition I, we will now use the expression of the generalized mixability constant η_ℓ^Φ to explicitly compare the regret bounds achieved using the GAA with different entropies. For an entropy $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$, we define $R_\ell^\Phi(\mathbf{q}) := \max_\theta D_\Phi(\mathbf{e}_\theta, \mathbf{q})/\eta_\ell^\Phi$, $\mathbf{q} \in \text{ri } \Delta_k$. The optimal regret bound achieved using the GAA is given by $R_\ell^\Phi := \inf_{\mathbf{q} \in \Delta_k} R_\ell^\Phi(\mathbf{q})$ (Reid et al., 2015). The next theorem shows that $R_\ell^S \leq R_\ell^\Phi$. The proof is in Appendix C.13.

Theorem 19 *Let $S, \Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$, where S is the Shannon entropy and Φ is an entropy such that $\tilde{\Phi}$ is twice differentiable on $\text{int } \tilde{\Delta}_k$. A loss $\ell: \mathcal{A} \rightarrow [0, +\infty]^n$, satisfying Condition I and with \underline{L}_ℓ twice differentiable on $]0, +\infty[^n$, is Φ -mixable only if $R_\ell^S \leq R_\ell^\Phi$.*

We conjecture that the result of Theorem 19 still holds even if one relaxes the smoothness requirement on $\tilde{\Phi}$ and \underline{L}_ℓ .

The result of Theorem 19 only concerns the worst case bound of the GAA. Since the update step of the GAA is exactly the MDA one can perhaps benefit from varying the entropy and the learning rate at each round to improve the regret bounds (Steinhardt and Liang, 2014; Joulani et al., 2017). In this scenario, at each round t the (adaptive) GAA would use a different entropy Φ^t computed according to the past performance of experts. Corollary 18 would still give an upper bound on the learning rate η^t at each round t as a function of the Hessian of Φ^t . Furthermore, this upper bound could be made larger by only requiring the Φ -mixability condition (5) to be satisfied locally; that is, ensuring (5) for $\mathbf{q} = \mathbf{q}^{t-1}$ and $A = A^t \in \mathcal{A}^k$, where \mathbf{q}^{t-1} is the previous distribution over experts and A^t is their prediction matrix. This would allow higher learning rates and potentially faster convergence.

Acknowledgments

This work was supported by the Australian Research Council of Data61.

References

- Ravi P. Agarwal, Maria Meehan, and Donal O’Regan. *Fixed point theory and applications*. Cambridge university press, 2001.
- Amir Beck and Marc Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
- Dennis S. Bernstein. *Matrix Mathematics: Theory, Facts, and Formulas*. Princeton University Press, second edition, 2011.
- Jonathan M. Borwein, Jon D. Vanderwerff, et al. *Convex functions: constructions, characterizations and counterexamples*. Cambridge University Press, 2010.

- Alexey Chernov, Yuri Kalnishkan, Fedor Zhdanov, and Vladimir Vovk. Supermartingales in prediction with expert advice. *Theoretical Computer Science*, 411(29-30):2647–2669, 2010.
- A. Philip Dawid. The geometry of proper scoring rules. *Annals of the Institute of Statistical Mathematics*, 59(1):77–93, 2007.
- Godfrey Harold Hardy. *A course of pure mathematics*. Cambridge University Press, 2008.
- J-B. Hiriart-Urruty and C. Lemaréchal. *Fundamentals of convex analysis*. Springer, 2001.
- Allen Holder. *An introduction to mathematical optimization*, 2005.
- Pooria Joulani, András György, and Csaba Szepesvári. A modular analysis of adaptive (non-)convex optimization: Optimism, composite objectives, and variational bounds. In *International Conference on Algorithmic Learning Theory*, pages 681–720, 2017.
- Parameswaran Kamalaruban, Robert Williamson, and Xinhua Zhang. Exp-concavity of proper composite losses. In *Conference on Learning Theory*, pages 1035–1065, 2015.
- Nelson Merentes and Kazimierz Nikodem. Remarks on strongly convex functions. *Aequationes mathematicae*, 80(1):193–199, 2010.
- Francesco Orabona, Koby Crammer, and Nicolò Cesa-Bianchi. A generalized online mirror descent with applications to classification and regression. *Machine Learning*, 99(3):411–435, 2015.
- Mark D. Reid, Rafael M. Frongillo, Robert C. Williamson, and Nishant Mehta. Generalized mixability via entropic duality. In *Conference on Learning Theory*, pages 1501–1522, 2015.
- Joel W. Robbin and Dietmar A. Salamon. *Introduction to differential geometry. ETH, Lecture Notes, preliminary version, January, 2011*.
- R. Tyrrell Rockafellar. *Convex analysis*. Princeton University Press, Princeton, NJ, 1997.
- Walter Rudin. *Principles of mathematical analysis*. McGraw-Hill, 1964.
- Jacob Steinhardt and Percy Liang. Adaptivity and optimism: An improved exponentiated gradient algorithm. In *International Conference on Machine Learning*, pages 1593–1601, 2014.
- John A. Thorpe. *Elementary Topics in Differential Geometry*. Springer Science & Business Media, 1994.
- Tim van Erven, Mark D. Reid, and Robert C. Williamson. Mixability is Bayes risk curvature relative to log loss. *Journal of Machine Learning Research*, 13:1639–1663, 2012.
- Vladimir Vovk. A game of prediction with expert advice. *Journal of Computer and System Sciences*, 56(2):153–173, 1998.
- Vladimir Vovk. The fundamental nature of the log loss function. In *Fields of Logic and Computation II*, pages 307–318. Springer, 2015.
- Robert C. Williamson. The geometry of losses. In *Conference on Learning Theory*, pages 1078–1108, 2014.

Robert C. Williamson, Elodie Vernet, and Mark D. Reid. Composite multiclass losses. *Journal of Machine Learning Research*, 17:223:1–223:52, 2016.

Appendix A. Additional Background Results

The following proposition gives some useful properties of the Fenchel dual which will be used in several proofs.

Proposition 20 (Hiriart-Urruty and Lemaréchal (2001)) *Let $f, h : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$. If f and h are proper and there are affine functions minorizing them on \mathbb{R}^n , then for $\mathbf{v} \in \mathbb{R}^n$*

$$\begin{aligned}
 (i) \quad & g(\mathbf{u}) = f(\mathbf{u}) + r & \implies & g^*(\mathbf{v}) = f^*(\mathbf{v}) - r, \\
 (ii) \quad & g(\mathbf{u}) = f(\mathbf{u}) + \langle \mathbf{v}_0, \mathbf{u} \rangle & \implies & g^*(\mathbf{v}) = f^*(\mathbf{v} - \mathbf{v}_0), \\
 (iii) \quad & f \leq h & \implies & f^* \geq h^*, \\
 (iv) \quad & \mathbf{s} \in \partial f^*(\mathbf{v}) & \implies & f^*(\mathbf{v}) = \langle \mathbf{v}, \mathbf{s} \rangle - f(\mathbf{s}), \\
 (v) \quad & g(\mathbf{u}) = f(t\mathbf{u}), t > 0 & \implies & g^*(\mathbf{v}) = f^*(\mathbf{v}/t),
 \end{aligned}$$

The following result due to Chernov et al. (2010) will be crucial to prove the convexity of the superprediction set (Theorem 5).

Lemma 21 (Chernov et al. (2010)) *Let $\Delta(\Omega)$ be the set of distributions over some set $\Omega \subseteq \mathbb{R}$. Let a function $Q : \Delta(\Omega) \times \Omega \rightarrow \mathbb{R}$ be such that $Q(\cdot, \omega)$ is continuous for all $\omega \in \Omega$. If for all $\pi \in \Delta(\Omega)$ it holds that $\mathbb{E}_{\omega \sim \pi} Q(\pi, \omega) \leq r$, where $r \in \mathbb{R}$ is some constant, then*

$$\exists \pi \in \Delta(\Omega), \forall \omega \in \Omega, Q(\pi, \omega) \leq r.$$

Note that when Ω in the lemma above is $[n]$, $\Delta([n]) \equiv \Delta_n$.

We make use of the following lemma due to Bernstein (2011) in proving a necessary condition for Φ -mixability (Theorem 16).

Lemma 22 $\forall m \geq 1, \forall A, B \in \mathbb{R}^{m \times m}$, $\lambda_{\max}(AB) = \lambda_{\max}(BA)$ and $\lambda_{\min}(AB) = \lambda_{\min}(BA)$.

Appendix B. Technical Lemmas

This appendix presents technical lemmas which will be needed in various proofs within this paper.

For an open convex set Ω in \mathbb{R}^n , a function $\phi : \Omega \rightarrow \mathbb{R}$ is said to be α -strongly convex if $z \mapsto \phi(\mathbf{u}) - \alpha \|\mathbf{u}\|^2$ is convex on Ω (Merentes and Nikodem, 2010). The next lemma is a generalization of α -strong convexity, where $\mathbf{u} \mapsto \|\mathbf{u}\|^2$ is replaced by any strictly convex function.

Lemma 23 *Let $\Omega \subseteq \mathbb{R}^n$ be an open convex set. Let $\phi, \psi : \Omega \rightarrow \mathbb{R}$ be twice differentiable.*

If ψ is strictly convex, then $\forall \mathbf{u} \in \Omega$, $\mathbf{H}\psi(\mathbf{u})$ is invertible, and for any $\alpha > 0$

$$\forall \mathbf{u} \in \Omega, \lambda_{\min}(\mathbf{H}\phi(\mathbf{u})(\mathbf{H}\psi(\mathbf{u}))^{-1}) \geq \alpha \iff \phi - \alpha\psi \text{ is convex}, \quad (17)$$

Furthermore, if $\alpha > 1$, then the left hand side of (17) implies that $\phi - \psi$ is strictly convex.

Proof Suppose that $\inf_{\mathbf{u} \in \Omega} \lambda_{\min}(\mathbf{H}\phi(\mathbf{u})(\mathbf{H}\psi(\mathbf{u}))^{-1}) \geq \alpha$. Since g is strictly convex and twice differentiable on Ω , $\mathbf{H}\psi(\mathbf{u})$ is symmetric positive definite, and thus invertible. Therefore, there

exists a symmetric positive definite matrix $G \in \mathbb{R}^{n \times n}$ such that $GG = H\psi(\mathbf{u})$. Lemma 22 implies

$$\begin{aligned}
 & \inf_{\mathbf{u} \in \Omega} \lambda_{\min}(H\phi(\mathbf{u})(H\psi(\mathbf{u}))^{-1}) && \geq \alpha, \\
 \iff & \inf_{\mathbf{u} \in \Omega} \lambda_{\min}(G^{-1}H\phi(\mathbf{u})G^{-1}) && \geq \alpha, \\
 \iff & \forall \mathbf{u} \in \Omega, \forall \mathbf{v} \in \mathbb{R}^n \setminus \{\mathbf{0}\}, \frac{\mathbf{v}^T G^{-1}(H\phi(\mathbf{u}))G^{-1}\mathbf{v}}{\mathbf{v}^T \mathbf{v}} && \geq \alpha, \\
 \iff & \forall \mathbf{u} \in \Omega, \forall \mathbf{w} \in \mathbb{R}^n \setminus \{\mathbf{0}\}, \mathbf{w}^T (H\phi(\mathbf{u}))\mathbf{w} && \geq \alpha \mathbf{w}^T GG\mathbf{w} = \mathbf{w}^T (\alpha H\psi(\mathbf{u}))\mathbf{w}, \\
 \iff & \forall \mathbf{u} \in \Omega, H\phi(\mathbf{u}) && \succeq \alpha H\psi(\mathbf{u}), \\
 \iff & \forall \mathbf{u} \in \Omega, H(\phi - \alpha\psi)(\mathbf{u}) && \succeq 0,
 \end{aligned}$$

where in the third and fifth lines we used the definition of minimum eigenvalue and performed the change of variable $\mathbf{w} = G^{-1}\mathbf{v}$, respectively. To conclude the proof of (17), note that the positive semi-definiteness of $H(\phi - \alpha\psi)$ is equivalent to the convexity of $\phi - \alpha\psi$ (Hiriart-Urruty and Lemaréchal, 2001, Thm B.4.3.1).

Finally, note that the equivalences established above still hold if we replace α , “ \geq ”, and “ \succeq ” by 1, “ $>$ ”, and “ \succ ”, respectively. The strict convexity of $\phi - \psi$ would then follow from the positive definiteness of $H(\phi - \psi)$ (ibid.). \blacksquare

The next crucial lemma is a slight modification of a result due to Chernov et al. (2010).

Lemma 24 *Let $f: \text{ri } \Delta_n \times [n] \rightarrow \mathbb{R}$ be a continuous function in the first argument and such that $\forall(\mathbf{q}, x) \in \text{ri } \Delta_n \times [n], -\infty < f(\mathbf{q}, x)$. Suppose that $\forall \mathbf{p} \in \text{ri } \Delta_n, \mathbb{E}_{x \sim \mathbf{p}}[f(\mathbf{p}, x)] \leq 0$, then*

$$\forall \epsilon > 0, \exists \mathbf{p}_\epsilon \in \text{ri } \Delta_n, \forall x \in [n], f(\mathbf{p}_\epsilon, x) \leq \epsilon.$$

Proof Pick any $\delta > 0$ such that $\delta(n-1) < 1$, and $c_0 < 0$ such that $\forall(\mathbf{q}, x) \in \text{ri } \Delta_n \times [n], c_0 \leq f(\mathbf{q}, x)$. We define $\Delta_n^\delta := \{\mathbf{p} \in \Delta_n : \forall x \in [n], p_x \geq \delta\}$ and $g(\mathbf{q}, \mathbf{p}) := \mathbb{E}_{x \sim \mathbf{q}}[f(\mathbf{p}, x)]$. For a fixed \mathbf{q} , $\mathbf{p} \mapsto g(\mathbf{q}, \mathbf{p})$ is continuous, since f is continuous in the first argument. For a fixed \mathbf{p} , $\mathbf{q} \mapsto g(\mathbf{q}, \mathbf{p})$ is linear, and thus concave. Since Δ_n^δ is convex and compact, g satisfies Ky Fan’s minimax Theorem (Agarwal et al., 2001, Thm. 11.4), and therefore, there exists $\mathbf{p}^\delta \in \Delta_n^\delta$ such that

$$\forall \mathbf{q} \in \Delta_n^\delta, \mathbb{E}_{x \sim \mathbf{q}}[f(\mathbf{p}^\delta, x)] = g(\mathbf{q}, \mathbf{p}^\delta) \leq \sup_{\mu \in \Delta_n^\delta} g(\mu, \mu) = \sup_{\mu \in \Delta_n^\delta} \mathbb{E}_{x \sim \mu}[f(\mu, x)] \leq 0. \quad (18)$$

For $x_0 \in [n]$, let $\hat{\mathbf{q}} \in \Delta_n^\delta$ be such that $\hat{q}_{x_0} = 1 - \delta(n-1)$ and $\hat{q}_x = \delta$ for $x \neq x_0$ (this is a legitimate distribution since $\delta(n-1) < 1$ by construction). Substituting $\hat{\mathbf{q}}$ for \mathbf{q} in (18) gives

$$\begin{aligned}
 & (1 - \delta(n-1))f(\mathbf{p}^\delta, x_0) + \delta \sum_{x \neq x_0} f(\mathbf{p}^\delta, x) && \leq 0, \\
 \implies & (1 - \delta(n-1))f(\mathbf{p}^\delta, x_0) && \leq -c_0\delta(n-1), \\
 \implies & f(\mathbf{p}^\delta, x_0) && \leq [-c_0\delta(n-1)]/[1 - \delta(n-1)].
 \end{aligned}$$

Choosing $\delta^* := \epsilon/[(-c_0 + \epsilon)(n-1)]$, and $\mathbf{p}_\epsilon := \mathbf{p}^{\delta^*}$ gives the desired result. \blacksquare

Lemma 25 *Let $f, g: I \rightarrow \mathbb{R}^n$, where $I \subseteq \mathbb{R}$ is an open interval containing 0. Suppose $g(t)$ [resp. f] is continuous [resp. differentiable] at 0. Then $t \mapsto \langle f(t), g(t) \rangle$ is differentiable at 0 if and only if $t \mapsto \langle f(0), g(t) \rangle$ is differentiable at 0, and we have*

$$\left. \frac{d}{dt} \langle f(t), g(t) \rangle \right|_{t=0} = \left\langle \left. \frac{d}{dt} f(t) \right|_{t=0}, g(0) \right\rangle + \left. \frac{d}{dt} \langle f(0), g(t) \rangle \right|_{t=0}.$$

Proof We have

$$\begin{aligned} \frac{\langle f(t), g(t) \rangle - \langle f(0), g(0) \rangle}{t} &= \frac{\langle f(t), g(t) \rangle - \langle f(0), g(t) \rangle}{t} + \frac{\langle f(0), g(t) \rangle - \langle f(0), g(0) \rangle}{t}, \\ &= \left\langle \frac{f(t) - f(0)}{t}, g(t) \right\rangle + \frac{\langle f(0), g(t) \rangle - \langle f(0), g(0) \rangle}{t}. \end{aligned}$$

But since g [resp. f] is continuous [resp. differentiable] at 0, the first term on the right hand side of the above equation converges to $\langle \frac{d}{dt}f(t)|_{t=0}, g(0) \rangle$ as $t \rightarrow 0$. Therefore, $\frac{1}{t}(\langle f(0), g(t) \rangle - \langle f(0), g(0) \rangle)$ admits a limit when $t \rightarrow 0$ if and only if $\frac{1}{t}(\langle f(t), g(t) \rangle - \langle f(0), g(0) \rangle)$ admits a limit when $t \rightarrow 0$. This shows that $t \mapsto \langle f(0), g(t) \rangle$ is differentiable at 0 if and only if $t \mapsto \langle f(t), g(t) \rangle$ is differentiable at 0, and in this case the above equation yields

$$\begin{aligned} \left. \frac{d}{dt} \langle f(t), g(t) \rangle \right|_{t=0} &= \lim_{t \rightarrow 0} \frac{\langle f(t), g(t) \rangle - \langle f(0), g(0) \rangle}{t}, \\ &= \lim_{t \rightarrow 0} \left(\left\langle \frac{f(t) - f(0)}{t}, g(t) \right\rangle + \frac{\langle f(0), g(t) \rangle - \langle f(0), g(0) \rangle}{t} \right), \\ &= \left\langle \left. \frac{d}{dt} f(t) \right|_{t=0}, g(0) \right\rangle + \left. \frac{d}{dt} \langle f(0), g(t) \rangle \right|_{t=0}. \end{aligned}$$

■

Note that the differentiability of $t \mapsto \langle f(0), g(t) \rangle$ at 0 does not necessarily imply the differentiability of g at 0. Take for example $n = 3$, $f(t) = \mathbf{1}/3$ for $t \in]-1, 1[$, and

$$g(t) = \begin{cases} -te_1 + t\frac{\mathbf{1}}{3}, & \text{if } t \in]-1, 0[; \\ -t\frac{\mathbf{1}}{3} + te_2, & \text{if } t \in [0, 1[. \end{cases}$$

Then the function $t \mapsto \langle f(0), g(t) \rangle = 0$ is differentiable at 0 but g is not. The preceding Lemma will be particularly useful in settings where it is desired to compute the derivative $\left. \frac{d}{dt} \langle f(0), g(t) \rangle \right|_{t=0}$ without any explicit assumptions on the differentiability of $g(t)$ at 0. For example, this will come up when computing $\left. \frac{d}{dt} \langle \mathbf{p}, D\tilde{\ell}(\tilde{\alpha}^t)\mathbf{v} \rangle \right|_{t=0}$, where $\mathbf{v} \in \mathbb{R}^{n-1}$ and $t \mapsto \tilde{\alpha}^t$ is smooth curve on $\text{int } \tilde{\Delta}_n$, with the only assumption that \tilde{L}_ℓ is twice differentiable at $\tilde{\alpha}^0 \in \text{int } \tilde{\Delta}_n$.

Lemma 26 *Let $\ell: \Delta_n \rightarrow [0, +\infty]^n$ be a proper loss. For any $\mathbf{p} \in \text{ri } \Delta_n$, it holds that*

$$\ell \text{ is continuous at } \mathbf{p} \stackrel{(i)}{\iff} \underline{L}_\ell \text{ is differentiable at } \mathbf{p} \stackrel{(ii)}{\iff} \partial[-\underline{L}_\ell](\mathbf{p}) = \{\nabla \underline{L}_\ell(\mathbf{p})\} = \{\ell(\mathbf{p})\}.$$

Proof [$\stackrel{(i)}{\iff}$] This equivalence has been shown before by [Williamson et al. \(2016\)](#).

[$\stackrel{(ii)}{\iff}$] Since $\underline{L}_\ell(\mathbf{p}) = -\sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{p})$, for all $\mathbf{p} \in \text{ri } \Delta_n$, it follows that \underline{L}_ℓ is differentiable at \mathbf{p} if and only if $\partial[-\underline{L}_\ell](\mathbf{p}) = \partial\sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{p}) = \{-\nabla\sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{p})\} = \{\nabla \underline{L}_\ell(\mathbf{p})\}$ ([Hiriart-Urruty and Lemaréchal, 2001](#), Cor. D.2.1.4). It remains to show that $\nabla \underline{L}_\ell(\mathbf{r}) = \ell(\mathbf{r})$ when \underline{L}_ℓ is differentiable at $\mathbf{r} \in \text{ri } \Delta_n$. Let $\alpha_x^t = \mathbf{r} + te_x$ and $\tilde{\alpha}_x^t = \Pi_n(\alpha_x^t)$, where $(e_x)_{x \in [n]}$ is the standard basis of \mathbb{R}^n . For $x \in [n]$, the functions $f_x(t) := \alpha_x^t$ and $g_x(t) := \tilde{\ell}(\tilde{\alpha}_x^t)$ satisfy the conditions of [Lemma 25](#).

Therefore, $h_x(t) := \langle f_x(0), g_x(t) \rangle = \langle \mathbf{r}, \tilde{\ell}(\tilde{\alpha}_x^t) \rangle$ is differentiable at 0 and

$$\begin{aligned} \nabla \tilde{L}(\mathbf{r}) \mathbf{e}_x &= \left. \frac{d}{dt} \tilde{L}(\tilde{\alpha}_x^t) \right|_{t=0} = \left. \frac{d}{dt} \langle f_x(t), g_x(t) \rangle \right|_{t=0}, \\ &= \left\langle \mathbf{e}_x, \tilde{\ell}(\tilde{\mathbf{r}}) \right\rangle + \left. \frac{d}{dt} h_x(t) \right|_{t=0}, \\ &= \tilde{\ell}_x(\tilde{\mathbf{r}}), \end{aligned}$$

where the last equality holds because h_x attains a minimum at 0 due to the properness of ℓ . The result being true for all $x \in [n]$ implies that $\nabla \tilde{L}(\tilde{\mathbf{r}}) = \tilde{\ell}(\tilde{\mathbf{r}}) = \ell(\mathbf{r})$. \blacksquare

The next Lemma is a restatement of earlier results due to [van Erven et al. \(2012\)](#). Our proof is more concise due to our definition of the Bayes risk in terms of the support function of the superprediction set.

Lemma 27 ([van Erven et al. \(2012\)](#)) *Let $\ell: \Delta_n \rightarrow [0, +\infty]^n$ be a proper loss whose Bayes risk is twice differentiable on $]0, +\infty[^n$ and let $X_{\mathbf{p}} = I_{\tilde{n}} - \mathbf{1}_{\tilde{n}} \tilde{\mathbf{p}}^T$. The following holds*

- (i) $\forall \mathbf{p} \in \text{ri } \Delta_n, \langle \mathbf{p}, D\tilde{\ell}(\tilde{\mathbf{p}}) \rangle = \mathbf{0}_{\tilde{n}}^T$.
- (ii) $\forall \tilde{\mathbf{p}} \in \text{int } \tilde{\Delta}_n, D\tilde{\ell}(\tilde{\mathbf{p}}) = \begin{bmatrix} X_{\mathbf{p}} \\ -\tilde{\mathbf{p}}^T \end{bmatrix} H\tilde{L}_{\ell}(\tilde{\mathbf{p}})$.
- (iii) $\forall \tilde{\mathbf{p}} \in \text{int } \tilde{\Delta}_n, H\tilde{L}_{\log}(\tilde{\mathbf{p}}) = -(X_{\mathbf{p}})^{-1} (\text{diag}(\tilde{\mathbf{p}}))^{-1}$.

Proof [We show (i) and (ii)] Let $\mathbf{p} \in \text{ri } \Delta_n$ and $f(\tilde{\mathbf{q}}) := \langle \mathbf{p}, \tilde{\ell}(\tilde{\mathbf{q}}) \rangle = \langle \mathbf{p}, \nabla \underline{L}_{\ell}(\mathbf{q}) \rangle$, where the equality is due to Lemma 26. Since \underline{L}_{ℓ} is twice differentiable $]0, +\infty[^n$, f is differentiable on $\text{int } \tilde{\Delta}_n$ and we have $Df(\tilde{\mathbf{q}}) = \langle \mathbf{p}, D\tilde{\ell}(\tilde{\mathbf{p}}) \rangle$. Since ℓ is proper, f reaches a minimum at $\tilde{\mathbf{p}} \in \text{int } \Delta_n$, and thus $\langle \mathbf{p}, D\tilde{\ell}(\tilde{\mathbf{p}}) \rangle = \mathbf{0}_{\tilde{n}}^T$ (this shows (i)). On the other hand, we have $\nabla \tilde{L}_{\ell}(\tilde{\mathbf{p}}) = J_n^T \nabla \underline{L}_{\ell}(\mathbf{p}) = J_n^T \tilde{\ell}(\tilde{\mathbf{p}})$. By differentiating and using the chain the rule, we get $H\tilde{L}_{\ell}(\tilde{\mathbf{p}}) = [D\tilde{\ell}(\tilde{\mathbf{p}})]^T J_n$. This means that $\forall i \in [\tilde{n}]$, $[H\tilde{L}_{\ell}(\tilde{\mathbf{p}})]_{\cdot, i} = \nabla \tilde{\ell}_i(\tilde{\mathbf{p}}) - \nabla \tilde{\ell}_n(\tilde{\mathbf{p}})$, and thus $\sum_{i=1}^{\tilde{n}} p_i [H\tilde{L}_{\ell}(\tilde{\mathbf{p}})]_{\cdot, i} = \sum_{i=1}^{\tilde{n}} p_i \nabla \tilde{\ell}_i(\tilde{\mathbf{p}}) - (1 - p_n) \nabla \tilde{\ell}_n(\tilde{\mathbf{p}})$. On the other hand, it follows from point (i) of the lemma that $\sum_{i=1}^{\tilde{n}} p_i \nabla \tilde{\ell}_i(\tilde{\mathbf{p}}) = \mathbf{0}_{\tilde{n}}$. Therefore, $[H\tilde{L}_{\ell}(\tilde{\mathbf{p}})] \tilde{\mathbf{p}} = -\nabla \tilde{\ell}_n(\tilde{\mathbf{p}})$ and, as a result, $\forall i \in [\tilde{n}]$, $[H\tilde{L}_{\ell}(\tilde{\mathbf{p}})]_{\cdot, i} - [H\tilde{L}_{\ell}(\tilde{\mathbf{p}})] \tilde{\mathbf{p}} = \nabla \tilde{\ell}_i(\tilde{\mathbf{p}})$. The last two equations can be combined as $D\tilde{\ell}(\tilde{\mathbf{p}}) = \begin{bmatrix} X_{\mathbf{p}} \\ -\tilde{\mathbf{p}}^T \end{bmatrix} H\tilde{L}_{\ell}(\tilde{\mathbf{p}})$.

[We show (iii)] It follows from (ii), since $\forall i \in [\tilde{n}], \nabla [\tilde{\ell}_{\log}]_i(\tilde{\mathbf{p}}) = \frac{1}{p_i} \mathbf{e}_i$, for $\tilde{\mathbf{p}} \in \text{int } \tilde{\Delta}_n$. \blacksquare

In the next lemma we state a new result for proper losses which will be crucial to prove a necessary condition for Φ -mixability (Theorem 16) — one of the main result of the paper.

Lemma 28 *Let $\ell: \Delta_n \rightarrow [0, +\infty]^n$ be a proper loss whose Bayes risk is twice differentiable on $]0, +\infty[^n$. For $\mathbf{v} \in \mathbb{R}^{n-1}$ and $\tilde{\mathbf{p}} \in \text{int } \tilde{\Delta}_n$,*

$$\left\langle \mathbf{p}, (D\tilde{\ell}(\tilde{\mathbf{p}})\mathbf{v}) \odot (D\tilde{\ell}(\tilde{\mathbf{p}})\mathbf{v}) \right\rangle = -\mathbf{v}^T H\tilde{L}_{\ell}(\tilde{\mathbf{p}}) [H\tilde{L}_{\log}(\tilde{\mathbf{p}})]^{-1} H\tilde{L}_{\ell}(\tilde{\mathbf{p}}) \mathbf{v}, \quad (19)$$

where $\mathbf{p} = \Pi_n(\tilde{\mathbf{p}})$ and \underline{L}_{\log} is the Bayes risk of the log loss.

Furthermore, if $t \mapsto \tilde{\alpha}^t$ is a smooth curve in $\text{int } \tilde{\Delta}_n$ and satisfies $\tilde{\alpha}^0 = \tilde{\mathbf{p}}$ and $\frac{d}{dt} \tilde{\alpha}^t|_{t=0} = \mathbf{v}$, then $t \mapsto \langle \mathbf{p}, D\tilde{\ell}(\tilde{\alpha}^t)\mathbf{v} \rangle$ is differentiable at 0 and we have

$$\frac{d}{dt} \left\langle \mathbf{p}, D\tilde{\ell}(\tilde{\alpha}^t)\mathbf{v} \right\rangle \Big|_{t=0} = -\mathbf{v}^\top \mathbf{H}\tilde{\underline{L}}_\ell(\tilde{\mathbf{p}})\mathbf{v}. \quad (20)$$

Proof We know from Lemma 27 that for $\tilde{\mathbf{p}} \in \text{int } \tilde{\Delta}_n$, we have $D\tilde{\ell}(\tilde{\mathbf{p}}) = \begin{bmatrix} X_{\tilde{\mathbf{p}}} \\ -\tilde{\mathbf{p}}^\top \end{bmatrix} \mathbf{H}\tilde{\underline{L}}_\ell(\tilde{\mathbf{p}})$, where $X_{\tilde{\mathbf{p}}} = I_{n-1} - \mathbf{1}_{n-1}\tilde{\mathbf{p}}^\top$. Thus, we can write

$$\begin{aligned} \left\langle \mathbf{p}, D\tilde{\ell}(\tilde{\mathbf{p}})\mathbf{v} \odot D\tilde{\ell}(\tilde{\mathbf{p}})\mathbf{v} \right\rangle &= \mathbf{v}^\top (D\tilde{\ell}(\tilde{\mathbf{p}}))^\top \text{diag}(\mathbf{p}) D\tilde{\ell}(\tilde{\mathbf{p}})\mathbf{v}, \\ &= \mathbf{v}^\top (\mathbf{H}\tilde{\underline{L}}_\ell(\tilde{\mathbf{p}}))^\top [X_{\tilde{\mathbf{p}}}^\top, -\tilde{\mathbf{p}}] \text{diag}(\mathbf{p}) \begin{bmatrix} X_{\tilde{\mathbf{p}}} \\ -\tilde{\mathbf{p}}^\top \end{bmatrix} \mathbf{H}\tilde{\underline{L}}_\ell(\tilde{\mathbf{p}})\mathbf{v}. \end{aligned} \quad (21)$$

Observe that $[X_{\tilde{\mathbf{p}}}^\top, -\tilde{\mathbf{p}}] \text{diag}(\mathbf{p}) = [I_{n-1} - \tilde{\mathbf{p}}\mathbf{1}_{n-1}^\top, -\tilde{\mathbf{p}}] \text{diag}(\mathbf{p}) = [\text{diag}(\tilde{\mathbf{p}}) - \tilde{\mathbf{p}}\tilde{\mathbf{p}}^\top, -\tilde{\mathbf{p}}p_n]$. Therefore,

$$\begin{aligned} [X_{\tilde{\mathbf{p}}}^\top, -\tilde{\mathbf{p}}] \text{diag}(\mathbf{p}) \begin{bmatrix} X_{\tilde{\mathbf{p}}} \\ -\tilde{\mathbf{p}}^\top \end{bmatrix} &= [\text{diag}(\tilde{\mathbf{p}}) - \tilde{\mathbf{p}}\tilde{\mathbf{p}}^\top, -\tilde{\mathbf{p}}p_n] \begin{bmatrix} I_{n-1} - \mathbf{1}_{n-1}\tilde{\mathbf{p}}^\top \\ -\tilde{\mathbf{p}}^\top \end{bmatrix}, \\ &= \text{diag}(\tilde{\mathbf{p}}) - \tilde{\mathbf{p}}\tilde{\mathbf{p}}^\top - \tilde{\mathbf{p}}\tilde{\mathbf{p}}^\top + \tilde{\mathbf{p}}\tilde{\mathbf{p}}^\top(1 - p_n) + p_n\tilde{\mathbf{p}}\tilde{\mathbf{p}}^\top, \\ &= \text{diag}(\tilde{\mathbf{p}}) - \tilde{\mathbf{p}}\tilde{\mathbf{p}}^\top, \\ &= \text{diag}(\tilde{\mathbf{p}}) X_{\tilde{\mathbf{p}}}, \\ &= -(\mathbf{H}\tilde{\underline{L}}_{\log}(\tilde{\mathbf{p}}))^{-1}, \end{aligned} \quad (22)$$

where the last equality is due to Lemma 27. The desired result follows by combining (21) and (22).

[We show (20)] Let $\tilde{\mathbf{p}} \in \text{int } \tilde{\Delta}_n$, we define $\tilde{\alpha}^t := \tilde{\mathbf{p}} + t\mathbf{v}$, $\alpha^t := \Pi_n(\tilde{\alpha}^t) = \mathbf{p} + tJ_n\mathbf{v}$, and $r(t) := \alpha^t / \|\alpha^t\|$, where $t \in \{s : \tilde{\mathbf{p}} + s\mathbf{v} \in \text{int } \tilde{\Delta}_n\}$. Since $t \mapsto r(t)$ is differentiable at 0 and $t \mapsto D\tilde{\ell}(\tilde{\alpha}^t)\mathbf{v}$ is continuous at 0, it follows from Lemma 23 that

$$\begin{aligned} \frac{d}{dt} \left\langle r(0), D\tilde{\ell}(\tilde{\alpha}^t)\mathbf{v} \right\rangle \Big|_{t=0} &= \frac{d}{dt} \left\langle r(t), D\tilde{\ell}(\tilde{\alpha}^t)\mathbf{v} \right\rangle \Big|_{t=0} - \left\langle \frac{d}{dt} r(t) \Big|_{t=0}, D\tilde{\ell}(\tilde{\mathbf{p}})\mathbf{v} \right\rangle, \\ &= - \left\langle \frac{d}{dt} r(t) \Big|_{t=0}, D\tilde{\ell}(\tilde{\mathbf{p}})\mathbf{v} \right\rangle, \end{aligned}$$

where the second equality holds since, according to Lemma 27, we have $\langle \alpha^t, D\tilde{\ell}(\tilde{\alpha}^t)\mathbf{v} \rangle = 0$. Since $r(0) = \mathbf{p} / \|\mathbf{p}\|$, $\frac{d}{dt} r(t) \Big|_{t=0} = \|\mathbf{p}\|^{-1} (I_n - r(0)[r(0)]^\top) J_n\mathbf{v}$, and $J_n = \begin{bmatrix} I_{n-1} \\ -\mathbf{1}_{n-1}^\top \end{bmatrix}$, we get

$$\begin{aligned} \|\tilde{\mathbf{p}}\| \frac{d}{dt} \left\langle r(0), D\tilde{\ell}(\tilde{\alpha}^t)\mathbf{v} \right\rangle \Big|_{t=0} &= - \left\langle (I_n - r(0)[r(0)]^\top) J_n\mathbf{v}, D\tilde{\ell}(\tilde{\mathbf{p}})\mathbf{v} \right\rangle, \\ &= - \left\langle J_n\mathbf{v}, D\tilde{\ell}(\tilde{\mathbf{p}})\mathbf{v} \right\rangle, \\ &= - \left\langle J_n\mathbf{v}, \begin{bmatrix} X_{\tilde{\mathbf{p}}} \\ -\tilde{\mathbf{p}}^\top \end{bmatrix} \mathbf{H}\tilde{\underline{L}}_\ell(\tilde{\mathbf{p}})\mathbf{v} \right\rangle, \\ &= -\mathbf{v}^\top \mathbf{H}\tilde{\underline{L}}_\ell(\tilde{\mathbf{p}})\mathbf{v}, \end{aligned} \quad (23)$$

where the passage to (23) is due to $r(0) = \mathbf{p}/\|\mathbf{p}\| \perp D\tilde{\ell}(\tilde{\mathbf{p}})$. In the last equality we used the fact that $J_n^\top \begin{bmatrix} X_{\mathbf{p}} \\ -\tilde{\mathbf{p}}^\top \end{bmatrix} = [I_{n-1}, -\mathbf{1}_{n-1}] \begin{bmatrix} I_{n-1} - \mathbf{1}_{n-1}\tilde{\mathbf{p}} \\ -\tilde{\mathbf{p}}^\top \end{bmatrix} = I_{n-1}$. \blacksquare

Lemma 29 For $\eta > 0$, $S_\eta := \eta^{-1} S$ satisfies (8), where S is the Shannon entropy.

Proof Let $(\hat{\mathbf{q}}, \mathbf{q}) \in (\text{rbd } \Delta_{\mathcal{I}}) \times (\text{ri } \Delta_{\mathcal{I}})$ and $\mathbf{q}^\lambda := \hat{\mathbf{q}} + \lambda(\mathbf{q} - \hat{\mathbf{q}})$, for $\lambda \in]0, 1[$. Let $\mathcal{J} := \{j \in \mathcal{I} : \hat{q}_j \neq 0\}$ and $\mathcal{H} := \mathcal{I} \setminus \mathcal{J}$. We have

$$\begin{aligned} S(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}}) &= \lim_{\lambda \downarrow 0} \lambda^{-1} \left[\sum_{\theta \in \mathcal{I}} q_\theta^\lambda \log q_\theta^\lambda - \sum_{\theta' \in \mathcal{J}} \hat{q}_{\theta'} \log \hat{q}_{\theta'} \right], \\ &= \lim_{\lambda \downarrow 0} \lambda^{-1} \left[\sum_{\theta \in \mathcal{J}} (q_\theta^\lambda \log q_\theta^\lambda - \hat{q}_\theta \log \hat{q}_\theta) + \sum_{\theta' \in \mathcal{H}} q_{\theta'}^\lambda \log q_{\theta'}^\lambda \right]. \end{aligned} \quad (24)$$

Observe that the limit of either summation term inside the bracket in (24) is equal to zero. Thus, using l'Hopital's rule we get

$$\begin{aligned} S(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}}) &= \lim_{\lambda \downarrow 0} \left[\sum_{\theta \in \mathcal{J}} [(q_\theta - \hat{q}_\theta) \log q_\theta^\lambda + (q_\theta - \hat{q}_\theta)] + \sum_{\theta' \in \mathcal{H}} [q_{\theta'} \log q_{\theta'}^\lambda + q_{\theta'}] \right], \\ &= \sum_{\theta \in \mathcal{J}} (q_\theta - \hat{q}_\theta) \log \hat{q}_\theta + \sum_{\theta' \in \mathcal{H}} q_{\theta'} \left[\lim_{\lambda \downarrow 0} \log q_{\theta'}^\lambda \right], \end{aligned} \quad (25)$$

where in (25) we used the fact that $\sum_{\theta \in \mathcal{J}} (q_\theta - \hat{q}_\theta) + \sum_{\theta' \in \mathcal{H}} q_{\theta'} = 0$. Since for all $\theta' \in \mathcal{H}$, $\lim_{\lambda \downarrow 0} q_{\theta'}^\lambda = 0$, the right hand side of (25) is equal to $-\infty$. Therefore S satisfies (8). Since $S_\eta = \eta^{-1} S$, it is clear that S_η also satisfies (8). \blacksquare

Lemma 30 Let $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ be an entropy. If Φ satisfies (8), then $\partial\tilde{\Phi}(\tilde{\mathbf{q}}) = \emptyset, \forall \tilde{\mathbf{q}} \in \text{bd } \tilde{\Delta}_k$, and $\forall \mathbf{d} \in \mathbb{R}^k, \forall \mathbf{q} \in \text{ri } \Delta_{\mathcal{I}}, M_\Phi(\mathbf{d}, \mathbf{q}) = M_{\Phi_{\mathcal{I}}}(\Pi_{\mathcal{I}}^k \mathbf{d}, \Pi_{\mathcal{I}}^k \mathbf{q})$.

Proof Let $\boldsymbol{\mu} \in \text{rbd } \Delta_k$. Since Φ satisfies (8), it follows that $\forall \mathbf{q} \in \text{ri } \Delta_k, \tilde{\Phi}(\tilde{\boldsymbol{\mu}}; \tilde{\mathbf{q}} - \tilde{\boldsymbol{\mu}}) = \Phi'(\boldsymbol{\mu}; \mathbf{q} - \boldsymbol{\mu}) = -\infty$. Therefore, $\partial\tilde{\Phi}'(\tilde{\boldsymbol{\mu}}) = \emptyset$ (Rockafellar, 1997, Thm. 23.4).

Let $\mathbf{d} \in \mathbb{R}^n, \mathcal{I} \subseteq [k]$, with $|\mathcal{I}| > 1$, and $\mathbf{q} \in \text{ri } \Delta_{\mathcal{I}}$. Then

$$\begin{aligned} M_{\Phi_{\mathcal{I}}}(\Pi_{\mathcal{I}}^k \mathbf{d}, \Pi_{\mathcal{I}}^k \mathbf{q}) &= \inf_{\boldsymbol{\pi} \in \Delta_{|\mathcal{I}|}} \langle \boldsymbol{\pi}, \Pi_{\mathcal{I}}^k \mathbf{d} \rangle + D_{\Phi_{\mathcal{I}}}(\boldsymbol{\pi}, \Pi_{\mathcal{I}}^k \mathbf{q}), \\ &= \inf_{\boldsymbol{\mu} \in \Delta_{\mathcal{I}}} \langle \boldsymbol{\mu}, \mathbf{d} \rangle + D_\Phi(\boldsymbol{\mu}, \mathbf{q}), \\ &\leq \inf_{\boldsymbol{\mu} \in \Delta_k} \langle \boldsymbol{\mu}, \mathbf{d} \rangle + D_\Phi(\boldsymbol{\mu}, \mathbf{q}), \\ &= M_\Phi(\mathbf{d}, \mathbf{q}). \end{aligned} \quad (26)$$

To complete the proof, we need to show that (26) holds with equality. For this, it suffice to prove that $\forall \boldsymbol{\mu} \in \Delta_k \setminus \Delta_{\mathcal{I}}, D_\Phi(\boldsymbol{\mu}, \mathbf{q}) = +\infty$. Let $\boldsymbol{\mu} \in \Delta_k \setminus \Delta_{\mathcal{I}}$ and $\mathcal{J} := \{\theta \in [k] : \mu_\theta \neq 0\} \cup \mathcal{I}$. In this case, we have $\mathbf{q} \in \text{rbd } \Delta_{\mathcal{J}}$ and $\mathbf{q} + 2^{-1}(\boldsymbol{\mu} - \mathbf{q}) \in \text{ri } \Delta_{\mathcal{J}}$. Thus, since Φ satisfies (8) and $\Phi'(\mathbf{q}; \cdot)$ is 1-homogeneous (Hiriart-Urruty and Lemaréchal, 2001, Prop. D.1.1.2), it follows that $2^{-1}\Phi'(\mathbf{q}; \boldsymbol{\mu} - \mathbf{q}) = \Phi'(\mathbf{q}; 2^{-1}(\boldsymbol{\mu} - \mathbf{q})) = -\infty$. Hence $D_\Phi(\boldsymbol{\mu}, \mathbf{q}) = +\infty$. \blacksquare

Lemma 31 For $\mathbf{q} \in \Delta_k$, there exists a sequence (\mathbf{d}_m) in $[0, +\infty[^k$ converging to $\mathbf{d} \in [0, +\infty]^k$ such that for any entropy $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$, we have $M_\Phi(\mathbf{d}_m, \mathbf{q}) \xrightarrow{m \rightarrow \infty} M_\Phi(\mathbf{d}, \mathbf{q})$.

Proof Let $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ be an entropy and let $\mathbf{d} \in [0, +\infty]^n$. Note that if $\mathcal{J} := \{\theta \in [k] : d_\theta < +\infty\} = \emptyset$, then by definition of M_Φ , it trivially holds that $M_\Phi(\mathbf{d}_m, \mathbf{q}) \xrightarrow{m \rightarrow \infty} M_\Phi(\mathbf{d}, \mathbf{q}) = +\infty$, for any sequence (\mathbf{d}_m) such that $\mathbf{d}_m \xrightarrow{m \rightarrow \infty} \mathbf{d}$. Similarly, if $\mathcal{I} := \{\theta \in [k] : d_\theta \neq 0\} = \emptyset$, then the zero sequence gives the desired result. Assume now that $\mathcal{J}, \mathcal{I} \neq \emptyset$ and define the function $f_\mu: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ by

$$f_\mu(\mathbf{z}) = \begin{cases} \sum_{\theta \in \mathcal{I}} \mu_\theta / z_\theta + D_\Phi(\mu, \mathbf{q}), & \text{if } \mathbf{z} \in]0, +\infty[^k; \\ +\infty, & \text{otherwise.} \end{cases}$$

Since $\forall \mu \in \Delta_k$, f_μ is upper-semicontinuous and concave, f_μ is a closed concave function. It follows that $f := \inf_{\mu \in \Delta_k} f_\mu$ is also closed and concave (Hiriart-Urruty and Lemaréchal, 2001, Prop. B.2.1.2). Let $\mathbf{b} \in [0, +\infty[^k$ be such that $b_\theta = \frac{1}{d_\theta}$, if $\theta \in \mathcal{I}$ (with the convention that $1/\infty = 0$); and $b_\theta = 0$ otherwise. Note that $f(\mathbf{b}) = M_\Phi(\mathbf{d}, \mathbf{q})$. Let $\mathbf{z} \in \mathbb{R}^k$ be such that $z_\theta = b_\theta$, if $\theta \in \mathcal{J}$; and $z_\theta = 1$ otherwise. Note that $\mathbf{z} \in \text{int dom } f$. Thus, since f is closed and concave, we have $f(\mathbf{b}) = \lim_{\lambda \downarrow 0} f(\mathbf{b} + \lambda(\mathbf{z} - \mathbf{b}))$ (ibid., Prop. B.1.2.5). Choosing $\mathbf{d}_m \in [0, +\infty[^n$ such that $[\mathbf{d}_m]_\theta = d_\theta + 1/m(z_\theta^{-1} - d_\theta)$, if $\theta \in \mathcal{I}$; and 0 otherwise, gives the desired sequence. \blacksquare

Appendix C. Proofs of Results in the Main Body

C.1. Proof of Proposition 1

Proof Given $\mathbf{v} \in \mathbb{R}^{k-1}$, we first derive the expression of the Fenchel dual $\tilde{S}^*(\mathbf{v}) := \sup_{\tilde{\mathbf{q}} \in \tilde{\Delta}_k} \langle \tilde{\mathbf{q}}, \mathbf{v} \rangle - \tilde{S}(\tilde{\mathbf{q}})$. Setting the gradient of $\tilde{\mathbf{q}} \mapsto \langle \tilde{\mathbf{q}}, \mathbf{v} \rangle - \tilde{S}(\tilde{\mathbf{q}})$ to $\mathbf{0}_{\tilde{k}}$ gives $\mathbf{v} = \nabla \tilde{S}(\tilde{\mathbf{q}})$. For $\mathbf{q} \in]0, +\infty[^k$, we have $\nabla S(\mathbf{q}) = \log \mathbf{q} + \mathbf{1}_k$, and from §2.1 we know that $\nabla \tilde{S}(\tilde{\mathbf{q}}) = J_k^\top \nabla S(\mathbf{q})$. Therefore,

$$\mathbf{v} = \nabla \tilde{S}(\tilde{\mathbf{q}}) \implies \mathbf{v} = J_k^\top \nabla S(\mathbf{q}) \implies \mathbf{v} = \log \frac{\tilde{\mathbf{q}}}{q_k},$$

where the right most equality is equivalent to $\tilde{\mathbf{q}}/q_k = \exp(\mathbf{v})$. Since $\langle \tilde{\mathbf{q}}, \mathbf{1}_{\tilde{k}} \rangle = 1 - q_k$, we get $q_k = (\langle \exp(\mathbf{v}), \mathbf{1}_{\tilde{k}} \rangle + 1)^{-1}$. Therefore, the supremum in the definition of $\tilde{S}^*(\mathbf{v})$ is attained at $\tilde{\mathbf{q}}_* = \exp(\mathbf{v})(\langle \exp(\mathbf{v}), \mathbf{1}_{\tilde{k}} \rangle + 1)^{-1}$. Hence $\tilde{S}^*(\mathbf{v}) = \langle \tilde{\mathbf{q}}_*, \mathbf{v} \rangle - \langle \tilde{\mathbf{q}}_*, \log \tilde{\mathbf{q}}_* \rangle = \log(\langle \exp(\mathbf{v}), \mathbf{1}_{\tilde{k}} \rangle + 1)$. Finally, using (2) we get $S^*(\mathbf{z}) = \log \langle \exp(\mathbf{z}), \mathbf{1}_k \rangle$, for $\mathbf{z} \in \mathbb{R}^k$. \blacksquare

C.2. Proof of Theorem 3

Proof We will construct a proper support loss $\underline{\ell}$ of ℓ .

Let $\mathbf{p} \in \text{ri } \Delta_n$ ($-\mathbf{p} \in \text{int dom } \sigma_{S_\ell^\oplus}$). Since the support function of a non-empty set is closed and convex, we have $\sigma_{S_\ell^\oplus}^{**} = \sigma_{S_\ell^\oplus}$ (Hiriart-Urruty and Lemaréchal, 2001, Prop. C.2.1.2). Pick any $\mathbf{v} \in \partial \sigma_{S_\ell^\oplus}(-\mathbf{p}) = \partial \sigma_{S_\ell^\oplus}^{**}(-\mathbf{p}) \neq \emptyset$. Since $\sigma_{S_\ell^\oplus}^* = \iota_{S_\ell^\oplus}$ (Rockafellar, 1997), we can apply Proposition 20-(iv) with f replaced by $\sigma_{S_\ell^\oplus}^*$ to obtain $\langle -\mathbf{p}, \mathbf{v} \rangle = \sigma_{S_\ell^\oplus}^*(-\mathbf{p}) + \iota_{S_\ell^\oplus}(\mathbf{v})$. The fact that

$\langle -\mathbf{p}, \mathbf{v} \rangle$ and $\sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{p})$ are both finite implies that $\iota_{\mathcal{S}_\ell^\oplus}(\mathbf{v}) = 0$. Therefore, $\mathbf{v} \in \mathcal{S}_\ell^\oplus$ and $\langle \mathbf{p}, \mathbf{v} \rangle = -\sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{p}) = \underline{L}_\ell(\mathbf{p})$. Define $\underline{\ell}(\mathbf{p}) := \mathbf{v} \in \mathcal{S}_\ell^\oplus$.

Now let $\mathbf{p} \in \text{rbd } \Delta_n$ and $\mathbf{q} := \mathbf{1}_n/n$. Since the \underline{L}_ℓ is a closed concave function and $\mathbf{q} \in \text{int dom } \underline{L}_\ell$, it follows that $\underline{L}_\ell(\mathbf{p} + m^{-1}(\mathbf{q} - \mathbf{p})) \xrightarrow{m \rightarrow \infty} \underline{L}_\ell(\mathbf{p})$ (Hiriart-Urruty and Lemaréchal, 2001, Prop. B.1.2.5). Note that $\mathbf{q}_m := \mathbf{p} + m^{-1}(\mathbf{q} - \mathbf{p}) \in \text{ri } \Delta_n, \forall m \in \mathbb{N}$. Now let $v_{x,m} := \underline{\ell}_x(\mathbf{q}_m)$, where $\underline{\ell}(\mathbf{q}_m)$ is as constructed in the previous paragraph. If $(v_{1,m})$ is bounded [resp. unbounded], we can extract a subsequence $(v_{1,\varphi_1(m)})$ which converges [resp. diverges to $+\infty$], where $\varphi_1 : \mathbb{N} \rightarrow \mathbb{N}$ is an increasing function. By repeating this process for $(v_{2,\varphi_1(m)})$ and so on, we can construct an increasing function $\varphi := \varphi_n \circ \dots \circ \varphi_1 : \mathbb{N} \rightarrow \mathbb{N}$, such that $\mathbf{v}_m := [v_{x,\varphi(m)}]_{x \in [n]}^\top$ has a well defined (coordinate-wise) limit in $[0, +\infty]^n$. Define $\underline{\ell}(\mathbf{p}) := \lim_{m \rightarrow \infty} \mathbf{v}_m$. By continuity of the inner product, we have

$$\begin{aligned} \langle \mathbf{p}, \underline{\ell}(\mathbf{p}) \rangle &= \lim_{m \rightarrow \infty} \langle \mathbf{q}_{\varphi(m)}, \mathbf{v}_m \rangle = \lim_{m \rightarrow \infty} \langle \mathbf{q}_{\varphi(m)}, \underline{\ell}(\mathbf{q}_{\varphi(m)}) \rangle, \\ &= \lim_{m \rightarrow \infty} \underline{L}_\ell(\mathbf{q}_{\varphi(m)}) = \underline{L}_\ell(\mathbf{p}). \end{aligned}$$

By construction, $\forall m \in \mathbb{N}, \mathbf{p}_m := \mathbf{q}_{\varphi(m)} \in \text{ri } \Delta_n$ and $\underline{\ell}(\mathbf{p}_m) = \mathbf{v}_m \xrightarrow{m \rightarrow \infty} \underline{\ell}(\mathbf{p})$. Therefore, $\underline{\ell}$ is support loss of ℓ .

It remains to show that it is proper; that is $\forall \mathbf{p} \in \Delta_n, \forall \mathbf{q} \in \Delta_n, \langle \mathbf{p}, \underline{\ell}(\mathbf{p}) \rangle \leq \langle \mathbf{p}, \underline{\ell}(\mathbf{q}) \rangle$. Let $\mathbf{q} \in \text{ri } \Delta_n$. We just showed that $\forall \mathbf{p} \in \Delta_n, \langle \mathbf{p}, \underline{\ell}(\mathbf{p}) \rangle = \underline{L}_\ell(\mathbf{p})$ and that $\underline{\ell}(\mathbf{q}) \in \mathcal{S}_\ell^\oplus$. Using the fact that $\underline{L}_\ell(\mathbf{p}) = \inf_{\mathbf{z} \in \mathcal{S}_\ell^\oplus} \langle \mathbf{p}, \mathbf{z} \rangle$, we obtain $\langle \mathbf{p}, \underline{\ell}(\mathbf{p}) \rangle \leq \langle \mathbf{p}, \underline{\ell}(\mathbf{q}) \rangle$.

Now let $\mathbf{q} \in \text{rbd } \Delta_k$. Since $\underline{\ell}$ is a support loss, we know that there exists a sequence $(\mathbf{q}_m) \subset \text{ri } \Delta_n$ such that $\underline{\ell}(\mathbf{q}_m) \xrightarrow{m \rightarrow \infty} \underline{\ell}(\mathbf{q})$. But as we established in the previous paragraph, $\langle \mathbf{p}, \underline{\ell}(\mathbf{p}) \rangle \leq \langle \mathbf{p}, \underline{\ell}(\mathbf{q}_m) \rangle$. By passing to the limit $m \rightarrow \infty$, we obtain $\langle \mathbf{p}, \underline{\ell}(\mathbf{p}) \rangle \leq \langle \mathbf{p}, \underline{\ell}(\mathbf{q}) \rangle$. Therefore $\underline{\ell}$ is a proper loss with Bayes risk \underline{L}_ℓ . \blacksquare

C.3. Proof of Theorem 4

For a set \mathcal{C} , we denote $\text{co } \mathcal{C}$ and $\overline{\text{co}} \mathcal{C}$ its *convex hull* and *closed convex hull*, respectively.

Definition 32 ((Hiriart-Urruty and Lemaréchal, 2001)) *Let \mathcal{C} be non-empty convex set in \mathbb{R}^n . We say that $\mathbf{u} \in \mathcal{C}$ is an extreme point of \mathcal{C} if there are no two different points \mathbf{u}_1 and \mathbf{u}_2 in \mathcal{C} and $\lambda \in]0, 1[$ such that $\mathbf{u} = \lambda \mathbf{u}_1 + (1 - \lambda) \mathbf{u}_2$.*

We denote the set of extreme points of a set \mathcal{C} by $\text{ext } \mathcal{C}$.

Lemma 33 *Let $\ell : \mathcal{A} \rightarrow [0, +\infty]^n$ be a closed loss. Then $\text{ext } \overline{\text{co}} \mathcal{S}_\ell^\oplus \subseteq \mathcal{S}_\ell$.*

Proof Since $\text{co } \mathcal{S}_\ell^\oplus \subseteq \mathbb{R}^n$ is connected, $\text{co } \mathcal{S}_\ell^\oplus = \{\mathbf{v} + \sum_{k=1}^n \alpha_k \ell(\mathbf{a}_k) : (\mathbf{a}_{k \in [n]}, \boldsymbol{\alpha}, \mathbf{v}) \in \mathcal{A}^n \times \Delta_n \times [0, +\infty]^n\}$ (Hiriart-Urruty and Lemaréchal, 2001, Prop. A.1.3.7).

We claim that $\overline{\text{co}} \mathcal{S}_\ell^\oplus = \text{co } \mathcal{S}_\ell^\oplus$. Let $(\mathbf{z}_m) := (\mathbf{v}_m + \sum_{k=1}^n \alpha_{m,k} \ell(\mathbf{a}_{m,k}))$ be a convergent sequence in $[0, +\infty]^n$, where $(\boldsymbol{\alpha}_m)$, $([\mathbf{a}_{m,k}]_{k \in [n]})$ and (\mathbf{v}_m) are sequences in Δ_n, \mathcal{A}^n , and $[0, +\infty]^n$, respectively. Since Δ_n is compact, we may assume, by extraction a subsequence if necessary, that $\boldsymbol{\alpha}_m \xrightarrow{m \rightarrow \infty} \boldsymbol{\alpha}^* \in \Delta_n$. Let $\mathcal{K} := \{k \in [n] : \alpha_k^* \neq 0\}$. Since \mathbf{z}_m converges, $([[\ell(\mathbf{a}_{m,k})]_{k \in \mathcal{K}}, \mathbf{v}_m])$

is a bounded sequence in $[0, +\infty[^{n|\mathcal{K}|+n}$. Since ℓ is closed, we may assume, by extraction a subsequence if necessary, that $\forall k \in \mathcal{K}$, $\ell(\mathbf{a}_{m,k}) \xrightarrow{m \rightarrow \infty} \ell(\mathbf{a}_k^*)$ and $\mathbf{v}_m \xrightarrow{m \rightarrow \infty} \mathbf{v}^*$, where $[\mathbf{a}_k^*]_{k \in \mathcal{K}} \in \mathcal{A}^{|\mathcal{K}|}$ and $\mathbf{v}^* \in [0, +\infty[^n$. Consequently,

$$\begin{aligned} \mathbf{v}^* + \sum_{k=1}^n \alpha_k^* \ell(\mathbf{a}_k^*) &= \lim_{m \rightarrow \infty} \left[\mathbf{v}_{m,k} + \sum_{k \in \mathcal{K}} \alpha_{m,k} \ell(\mathbf{a}_{m,k}) \right], \\ &\leq \lim_{m \rightarrow \infty} \mathbf{z}_m, \end{aligned}$$

where the last inequality is coordinate-wise. Therefore, there exists $\mathbf{v}' \in [0, +\infty[^n$ such that $\lim_{m \rightarrow \infty} \mathbf{z}_m = \mathbf{v}' + \mathbf{v}^* + \sum_{k=1}^n \alpha_k^* \ell(\mathbf{a}_k^*) \in \text{co} \mathcal{S}_\ell^\oplus$. This shows that $\overline{\text{co}} \mathcal{S}_\ell^\oplus \subset \text{co} \mathcal{S}_\ell^\oplus$, and thus $\overline{\text{co}} \mathcal{S}_\ell^\oplus = \text{co} \mathcal{S}_\ell^\oplus$ which proves our first claim.

By definition of an extreme point, $\text{ext} \overline{\text{co}} \mathcal{S}_\ell^\oplus \subseteq \overline{\text{co}} \mathcal{S}_\ell^\oplus$. Let $e \in \text{ext} \overline{\text{co}} \mathcal{S}_\ell^\oplus$ and $(\mathbf{a}_{k \in [n]}, \boldsymbol{\alpha}, \mathbf{v}) \in \mathcal{A}^n \times \Delta_n \times [0, +\infty[^n$ such that $e = \sum_{k=1}^n \alpha_k \ell(\mathbf{a}_k) + \mathbf{v}$. If there exists $i, j \in [n]$ such that $\alpha_i \alpha_j \neq 0$ or $\alpha_i v_j \neq 0$ then e would violate the definition of an extreme point. Therefore, the only possible extreme points are of the form $\{\ell(\mathbf{a}) : \mathbf{a} \in \text{dom} \ell\} = \mathcal{S}_\ell$. \blacksquare

Proof [Theorem 4] Suppose \underline{L}_ℓ is not differentiable at $\mathbf{p} \in \text{ri} \Delta_n$. Then from the definition of the Bayes risk, $\sigma_{\mathcal{S}_\ell^\oplus}$ is not differentiable at $-\mathbf{p}$. This implies that $\mathcal{F}(\mathbf{p}) := \partial \sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{p})$ has more than one element (Hiriart-Urruty and Lemaréchal, 2001, Cor. D.2.1.4). Since $\sigma_{\mathcal{S}_\ell^\oplus} = \sigma_{\overline{\text{co}} \mathcal{S}_\ell^\oplus}$ (ibid., Prop. C.2.2.1), $\mathcal{F}(\mathbf{p}) = \partial \sigma_{\overline{\text{co}} \mathcal{S}_\ell^\oplus}(-\mathbf{p})$ is a subset of $\overline{\text{co}} \mathcal{S}_\ell^\oplus$ and every extreme point of $\mathcal{F}(\mathbf{p})$ is also an extreme point of $\overline{\text{co}} \mathcal{S}_\ell^\oplus$ (ibid., Prop. A.2.3.7). Thus, from Lemma 33, we have $\text{ext} \mathcal{F}(\mathbf{p}) \subset \mathcal{S}_\ell$. On the other hand, since $-\mathbf{p} \in \text{int} \text{dom} \sigma_{\mathcal{S}_\ell^\oplus}$, $\mathcal{F}(\mathbf{p})$ is a compact, convex set (Rockafellar, 1997, Thm. 23.4), and thus $\mathcal{F}(\mathbf{p}) = \text{co}(\text{ext} \mathcal{F}(\mathbf{p}))$ (Hiriart-Urruty and Lemaréchal, 2001, Thm. A.2.3.4). Hence, the fact that $\mathcal{F}(\mathbf{p})$ has more than one element implies that there exists $\mathbf{a}_0, \mathbf{a}_1 \in \mathcal{A}$ such that $\ell(\mathbf{a}_0), \ell(\mathbf{a}_1) \in \text{ext} \mathcal{F}(\mathbf{p}) \subseteq \mathcal{F}(\mathbf{p})$ and $\ell(\mathbf{a}_0) \neq \ell(\mathbf{a}_1)$. Since $\mathcal{F}(\mathbf{p}) = \partial \sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{p})$, Proposition 20-(iv) and the fact that $\sigma_{\mathcal{S}_\ell^\oplus}^* = \iota_{\mathcal{S}_\ell^\oplus}$ imply that $\underline{L}_\ell(\mathbf{p}) = \langle \mathbf{p}, \underline{\ell}(\mathbf{p}) \rangle = \langle \mathbf{p}, \ell(\mathbf{a}_0) \rangle = \langle \mathbf{p}, \ell(\mathbf{a}_1) \rangle$.

Let $\mathbf{p} \in \text{ri} \Delta_n$ and suppose that \underline{L}_ℓ is differentiable at \mathbf{p} . In this case, $\sigma_{\mathcal{S}_\ell^\oplus}$ is differentiable at $-\mathbf{p}$, which implies that $\mathcal{F}(\mathbf{p}) = \partial \sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{p})$ is the singleton $\{\nabla \sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{p})\}$ (ibid., Cor. D.2.1.4). In this case, $\underline{\ell}(\mathbf{p}) = \nabla \sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{p})$ is the only extreme point of $\mathcal{F}(\mathbf{p}) \subset \overline{\text{co}} \mathcal{S}_\ell^\oplus$. From Lemma 33, there exists $\mathbf{a}_* \in \mathcal{A}$ such that $\ell(\mathbf{a}_*) = \underline{\ell}(\mathbf{p})$. In this paragraph, we showed the following

$$\forall \mathbf{p} \in \text{ri} \Delta_n, \exists \mathbf{a}_* \in \text{dom} \ell, \ell(\mathbf{a}_*) = \underline{\ell}(\mathbf{p}). \quad (27)$$

For the rest of this proof we will assume that \underline{L}_ℓ is differentiable on $]0, +\infty[^n$. Let $\mathbf{p} \in \text{rbd} \Delta_n \cap \text{dom} \underline{\ell}$. Since $\underline{\ell}$ is a support loss, there exists (\mathbf{p}_m) in $\text{ri} \Delta_n$ such that $(\underline{\ell}(\mathbf{p}_m))_m$ converges to $\underline{\ell}(\mathbf{p})$. From (27) it holds that $\forall \mathbf{p}_m \in \text{ri} \Delta_n, \exists \mathbf{a}_m \in \mathcal{A}, \ell(\mathbf{a}_m) = \underline{\ell}(\mathbf{p}_m)$. Since $(\ell(\mathbf{a}_m))_m$ converges and ℓ is closed, there exists $\mathbf{a}_* \in \mathcal{A}$ such that $\ell(\mathbf{a}_*) = \lim_{m \rightarrow \infty} \ell(\mathbf{a}_m) = \underline{\ell}(\mathbf{p})$.

Now let $\mathbf{a} \in \text{dom} \ell$ and $f(\mathbf{p}, x) := \underline{\ell}_x(\mathbf{p}) - \ell_x(\mathbf{a})$. Since $\ell(\mathbf{a}) \in \mathcal{S}_\ell^\oplus$ and $\underline{\ell}$ is proper, we have for all $\mathbf{p} \in \text{ri} \Delta_n, \mathbb{E}_{x \sim \mathbf{p}}[f(\mathbf{p}, x)] \leq 0$ and $-\infty < f(\mathbf{p}, x), \forall x \in [n]$. Therefore, Lemma 24 implies that for all $m \in \mathbb{N} \setminus \{0\}$ there exists $\mathbf{p}_m \in \text{ri} \Delta_n$, such that $\forall x \in [n], \underline{\ell}_x(\mathbf{p}_m) \leq \ell_x(\mathbf{a}) + 1/m$. On one hand, since $(\underline{\ell}(\mathbf{p}_m))_m$ is bounded (from the previous inequality), we may assume by extracting a subsequence if necessary, that $(\underline{\ell}(\mathbf{p}_m))_m$ converges. On the other hand, since $\mathbf{p}_m \in \text{ri} \Delta_n$, (27) implies that there exists $\mathbf{a}_m \in \text{dom} \ell$ such that $\underline{\ell}(\mathbf{p}_m) = \ell(\mathbf{a}_m)$. Since ℓ is closed and $(\ell(\mathbf{a}_m))_m$ converges, there exists $\mathbf{a}_* \in \mathcal{A}$, such that $\ell(\mathbf{a}_*) = \lim_{m \rightarrow \infty} \ell(\mathbf{a}_m) =$

$\lim_{m \rightarrow \infty} \underline{\ell}(\mathbf{p}_m) \leq \ell(\mathbf{a})$. But since ℓ is admissible, the latter component-wise inequality implies that $\ell(\mathbf{a}_*) = \ell(\mathbf{a}) = \lim_{m \rightarrow \infty} \underline{\ell}(\mathbf{p})$. \blacksquare

C.4. Proof of Theorem 5

Proof [$\mathcal{S}_\ell^\oplus \subseteq \bigcap_{\mathbf{p} \in \Delta_n} \mathcal{H}_{-\mathbf{p}, -\underline{L}_\ell(\mathbf{p})}$]: Let $\mathbf{v} \in \mathcal{S}_\ell^\oplus$, $\mathbf{u} \in [0, +\infty[^n$, and $\mathbf{q} \in \Delta_n$ such that $\mathbf{v} = \ell(\mathbf{q}) + \mathbf{u}$. Since ℓ is proper then $\forall \mathbf{p} \in \Delta_n$, $\underline{L}_\ell(\mathbf{p}) = \langle \mathbf{p}, \ell(\mathbf{p}) \rangle \leq \langle \mathbf{p}, \ell(\mathbf{q}) \rangle \leq \langle \mathbf{p}, \ell(\mathbf{q}) + \mathbf{u} \rangle = \langle \mathbf{p}, \mathbf{v} \rangle$. Therefore, $\mathbf{v} \in \bigcap_{\mathbf{p} \in \Delta_n} \mathcal{H}_{-\mathbf{p}, -\underline{L}_\ell(\mathbf{p})}$.

[$\bigcap_{\mathbf{p} \in \Delta_n} \mathcal{H}_{-\mathbf{p}, -\underline{L}_\ell(\mathbf{p})} \subseteq \mathcal{S}_\ell^\oplus$]: Let $\mathbf{v} \in \bigcap_{\mathbf{p} \in \Delta_n} \mathcal{H}_{-\mathbf{p}, -\underline{L}_\ell(\mathbf{p})}$. Let $\Omega = [n]$, $\Delta(\Omega) = \Delta_n$, and $Q(\mathbf{p}, x) = \ell_x(\mathbf{p}) - v_x$ for all $(\mathbf{p}, x) \in \Delta_n \times [n]$. Since $\mathbf{v} \in \bigcap_{\mathbf{p} \in \Delta_n} \mathcal{H}_{-\mathbf{p}, -\underline{L}_\ell(\mathbf{p})}$, $\mathbb{E}_{x \sim \mathbf{p}} Q(\mathbf{p}, x) = \langle \mathbf{p}, \ell(\mathbf{p}) \rangle - \langle \mathbf{p}, \mathbf{v} \rangle \leq 0$ for all $\mathbf{p} \in \Delta_n$. Lemma 21, implies that there exists $\mathbf{p}_* \in \Delta_n$ such that $Q(\mathbf{p}_*, x) = \ell_x(\mathbf{p}_*) - v_x \leq 0$, for all $x \in [n]$. This shows that $\mathbf{v} \in \mathcal{S}_\ell^\oplus$. \blacksquare

C.5. Proof of Theorem 7

Proof Let $\eta := \eta_\ell$. We will show that $\exp(-\eta \mathcal{S}_\ell^\oplus)$ is convex, which will imply that ℓ is η -mixable (Chernov et al., 2010).

Since $\eta_\ell = \inf_{\tilde{\mathbf{p}} \in \text{int } \tilde{\Delta}_n} (\lambda_{\max}([\mathbf{H}\tilde{L}_{\log}(\tilde{\mathbf{p}})]^{-1} \mathbf{H}\tilde{L}_\ell(\tilde{\mathbf{p}})))^{-1} > 0$, $\eta \underline{L}_\ell - \underline{L}_{\log}$ is convex on $\text{ri } \Delta_n$ (van Erven et al., 2012, Thm. 10). Let $\mathbf{p} \in \text{ri } \Delta_n$ and define

$$\Lambda(\mathbf{r}) := \underline{L}_{\log}(\mathbf{r}) + \langle \mathbf{r}, \eta \underline{\ell}(\mathbf{p}) - \underline{\ell}_{\log}(\mathbf{p}) \rangle, \quad \mathbf{r} \in \text{ri } \Delta_n.$$

Since Λ is equal to \underline{L}_{\log} plus an affine function, it follows that $\eta \underline{L}_\ell - \Lambda$ is also convex on $\text{ri } \Delta_n$. On one hand, since $\underline{\ell}$ and $\underline{\ell}_{\log}$ are proper losses, we have $\langle \mathbf{p}, \underline{\ell}(\mathbf{p}) \rangle = \underline{L}_\ell(\mathbf{p})$ and $\langle \mathbf{p}, \underline{\ell}_{\log}(\mathbf{p}) \rangle = \underline{L}_{\log}(\mathbf{p})$ which implies that

$$\eta \underline{L}_\ell(\mathbf{p}) - \Lambda(\mathbf{p}) = 0. \quad (28)$$

On the other hand, since \underline{L}_ℓ and \underline{L}_{\log} are differentiable we have $\underline{\ell}(\mathbf{p}) = \nabla \underline{L}_\ell(\mathbf{p})$ and $\nabla \underline{L}_{\log}(\mathbf{p}) = \underline{\ell}_{\log}(\mathbf{p})$, which yields $\eta \nabla \underline{L}_\ell(\mathbf{p}) - \nabla \Lambda(\mathbf{p}) = \mathbf{0}_n$. This implies that $\eta \underline{L}_\ell - \Lambda$ attains a minimum at \mathbf{p} (Hiriart-Urruty and Lemaréchal, 2001, Thm. D.2.2.1). Combining this fact and (28) gives $\eta \underline{L}_\ell(\mathbf{r}) \geq \Lambda(\mathbf{r}), \forall \mathbf{r} \in \text{ri } \Delta_n$. By Proposition 20-(iii), $-\eta \underline{L}_\ell \leq -\Lambda$ implies

$$[-\eta \underline{L}_\ell]^* \geq [-\Lambda]^*. \quad (29)$$

Using Proposition 20-(ii), we get $[-\Lambda]^*(\mathbf{s}) = [-\underline{L}_{\log}]^*(\mathbf{s} - \underline{\ell}_{\log}(\mathbf{p}) + \eta \underline{\ell}(\mathbf{p}))$ for $\mathbf{s} \in \mathbb{R}^n$. Since $-\eta \underline{L}_\ell(\mathbf{u}) = -\underline{L}_\ell(\eta \mathbf{u}) = \sigma_{\mathcal{S}_\ell^\oplus}(-\eta \mathbf{u})$ and $\sigma_{\mathcal{S}_\ell^\oplus}^* = \iota_{\mathcal{S}_\ell^\oplus}$, Proposition 20-(v) implies $[-\eta \underline{L}_\ell]^*(\mathbf{s}) = \iota_{\mathcal{S}_\ell^\oplus}(-\mathbf{s}/\eta)$. Similarly, we have $[-\underline{L}_{\log}]^*(\mathbf{s}) = \iota_{\mathcal{S}_{\log}^\oplus}(-\mathbf{s})$. Therefore, (29) implies

$$\forall \mathbf{s} \in \mathbb{R}^n, \quad \iota_{\mathcal{S}_\ell^\oplus}(-\mathbf{s}/\eta) \geq \iota_{\mathcal{S}_{\log}^\oplus}(-\mathbf{s} + \underline{\ell}_{\log}(\mathbf{p}) - \eta \underline{\ell}(\mathbf{p})).$$

This inequality implies that if $\mathbf{s} \in -\eta \mathcal{S}_\ell^\oplus$, then $\mathbf{s} \in -\mathcal{S}_{\log}^\oplus + \underline{\ell}_{\log}(\mathbf{p}) - \eta \underline{\ell}(\mathbf{p})$. In particular, if $\mathbf{u} \in e^{-\eta \mathcal{S}_\ell^\oplus}$ then

$$\mathbf{u} \in e^{-\mathcal{S}_{\log}^\oplus + \underline{\ell}_{\log}(\mathbf{p}) - \eta \underline{\ell}(\mathbf{p})} \subseteq \mathcal{H}_{\tau(\mathbf{p}), 1} = \{\mathbf{v} \in \mathbb{R}^n : \langle \mathbf{v}, \mathbf{p} \odot e^{\eta \underline{\ell}(\mathbf{p})} \rangle \leq 1\}. \quad (30)$$

To see the set inclusion in (30), consider $\mathbf{s} \in -\mathcal{S}_{\log}^{\oplus} + \underline{\ell}_{\log}(\mathbf{p}) - \eta \underline{\ell}(\mathbf{p})$, then by definition of the superprediction set $\mathcal{S}_{\log}^{\oplus}$ there exists $\mathbf{r} \in \Delta_n$ and $\mathbf{v} \in [0, +\infty]^n$, such that $\mathbf{s} = \log \mathbf{r} - \log \mathbf{p} - \eta \underline{\ell}(\mathbf{p}) - \mathbf{v}$. Thus,

$$\langle e^{\mathbf{s}}, \mathbf{p} \odot e^{\eta \underline{\ell}(\mathbf{p})} \rangle = \langle \mathbf{r}, e^{-\mathbf{v}} \rangle \leq 1, \quad (31)$$

where the inequality is true because $\mathbf{r} \in \Delta_n$ and $\mathbf{v} \in [0, +\infty]^n$. The above argument shows that $e^{-\eta \mathcal{S}_{\ell}^{\oplus}} \subseteq \mathcal{H}_{\tau(\mathbf{p}),1}$. Furthermore, $e^{-\eta \mathcal{S}_{\ell}^{\oplus}} \subseteq \mathcal{H}_{\tau(\mathbf{p}),1} \cap [0, +\infty]^n$, since all elements of $e^{-\eta \mathcal{S}_{\ell}^{\oplus}}$ have non-negative components. The latter set inclusion still holds for $\hat{\mathbf{p}} \in \text{ri} \Delta_n$. In fact, from the definition of a support loss, there exists a sequence (\mathbf{p}_m) in $\text{ri} \Delta_n$ converging to $\hat{\mathbf{p}}$ such that $\underline{\ell}(\mathbf{p}_m) \xrightarrow{m \rightarrow \infty} \underline{\ell}(\hat{\mathbf{p}})$. Equation 31 implies that for $\mathbf{u} \in e^{-\eta \mathcal{S}_{\ell}^{\oplus}}$, $\langle \mathbf{u}, \mathbf{p}_m \odot e^{\eta \underline{\ell}(\mathbf{p}_m)} \rangle \leq 1$. Since the inner product is continuous, by passage to the limit, we obtain $\langle \mathbf{u}, \hat{\mathbf{p}} \odot e^{\eta \underline{\ell}(\hat{\mathbf{p}})} \rangle \leq 1$. Therefore,

$$e^{-\eta \mathcal{S}_{\ell}^{\oplus}} \subseteq \bigcap_{\mathbf{p} \in \Delta_n} \mathcal{H}_{\tau(\mathbf{p}),1} \cap [0, +\infty]^n. \quad (32)$$

Now suppose $\mathbf{u} \in \bigcap_{\mathbf{p} \in \Delta_n} \mathcal{H}_{\tau(\mathbf{p}),1} \cap [0, +\infty]^n$; that is, for all $\mathbf{p} \in \Delta_n$,

$$\begin{aligned} 1 \geq \langle \mathbf{u}, \mathbf{p} \odot e^{\eta \underline{\ell}(\mathbf{p})} \rangle &= \langle \mathbf{p}, \mathbf{u} \odot e^{\eta \underline{\ell}(\mathbf{p})} \rangle = \langle \mathbf{p}, e^{\eta \underline{\ell}(\mathbf{p}) + \log \mathbf{u}} \rangle, \\ &\geq e^{\langle \mathbf{p}, \eta \underline{\ell}(\mathbf{p}) \rangle + \langle \mathbf{p}, \log \mathbf{u} \rangle}, \end{aligned} \quad (33)$$

where the first equality is obtained merely by expanding the expression of the inner product, and the second inequality is simply Jensen's Inequality. Since $\mathbf{u} \mapsto e^{\mathbf{u}}$ is strictly convex, the Jensen's inequality in (33) is strict unless $\exists(c, \mathbf{p}) \in \mathbb{R} \times \Delta_n$, such that

$$\eta \underline{\ell}(\mathbf{p}) + \log \mathbf{u} = c \mathbf{1}_n. \quad (34)$$

By substituting (34) into (33), we get $1 \geq \exp(c)$, and thus $c \leq 0$. Furthermore, (24) implies that $\mathbf{p} \in \text{dom} \underline{\ell}$, and thus there exists $\mathbf{a} \in \text{dom} \ell$ such that $\ell(\mathbf{a}) = \underline{\ell}(\mathbf{p})$ (Theorem 4). Using this and rearranging (34), we get $\mathbf{u} = \exp(-\eta \ell(\mathbf{a}) + c \mathbf{1})$. Since $c \leq 0$, this means that $\mathbf{u} \in \exp(-\eta \mathcal{S}_{\ell}^{\oplus})$. Suppose now that (34) does not hold. In this case, (33) must be a strict inequality for all $\mathbf{p} \in \Delta_n$. By applying the log in (33),

$$\forall \mathbf{p} \in \Delta_n, \underline{L}_{\ell}(\mathbf{p}) + \langle \mathbf{p}, \log \mathbf{u} \rangle = \langle \mathbf{p}, \eta \underline{\ell}(\mathbf{p}) \rangle + \langle \mathbf{p}, \log \mathbf{u} \rangle < 0. \quad (35)$$

Since $\mathbf{p} \mapsto \underline{L}_{\ell}(\mathbf{p}) = -\sigma_{\mathcal{S}_{\ell}^{\oplus}}(-\mathbf{p})$ is a closed concave function, the map $g: \mathbf{p} \mapsto \underline{L}_{\ell}(\mathbf{p}) + \langle \mathbf{p}, \log \mathbf{u} \rangle$ is also closed and concave, and thus upper semi-continuous. Since Δ_n is compact, the function g must attain its maximum in Δ_n (Holder, 2005, Thm. 1.13). Due to (35) this minimum is negative; there exists $c_1 > 0$ such that $\langle \mathbf{p}, \eta \underline{\ell}(\mathbf{p}) \rangle - \langle \mathbf{p}, -\log \mathbf{u} \rangle \leq -c_1$. Let $f(\mathbf{p}, x) := \underline{\ell}_x(\mathbf{p}) + \log u_x + c_1$, for $x \in [n]$. Consequently for all $\mathbf{p} \in \Delta_n$, $\mathbb{E}_{x \sim \mathbf{p}} f(\mathbf{p}, x) \leq 0$ and $\forall x \in [n], -\infty < f(\mathbf{p}, x)$. Thus, Lemma 25 applied to f with $\epsilon = c_1/2$, implies that there exists $\mathbf{p}_* \in \text{ri} \Delta_n$, such that $\eta \underline{\ell}(\mathbf{p}_*) \leq -\log \mathbf{u} - c_1/2 \leq -\log \mathbf{u}$. From this inequality, $\mathbf{p}_* \in \text{dom} \underline{\ell}$, and therefore, there exists $\mathbf{a}_* \in \text{dom} \ell$ such that $\ell(\mathbf{a}_*) = \underline{\ell}(\mathbf{p}_*)$ (Theorem 4). This shows that $\eta \ell(\mathbf{a}_*) \leq -\log \mathbf{u}$, which implies that $\mathbf{u} \in \exp -\eta \mathcal{S}_{\ell}^{\oplus}$. Therefore, $\bigcap_{\mathbf{p} \in \Delta_n} \mathcal{H}_{\tau(\mathbf{p}),1} \cap [0, +\infty]^n \subseteq e^{-\eta \mathcal{S}_{\ell}^{\oplus}}$. Combining this with (32) shows that $e^{-\eta \mathcal{S}_{\ell}^{\oplus}} = \bigcap_{\mathbf{p} \in \Delta_n} \mathcal{H}_{\tau(\mathbf{p}),1} \cap [0, +\infty]^n$. Since $e^{-\eta \mathcal{S}_{\ell}^{\oplus}}$ is the intersection of convex set, it is a convex set itself. Therefore ℓ is η -mixable. \blacksquare

C.6. Proof of Proposition 11

Given an entropy $\Phi: \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ and a loss $\ell: \mathcal{A} \rightarrow [0, +\infty]$, we define

$$\Delta m_{\Phi}(x, A, \mathbf{a}, \hat{\mathbf{q}}, \boldsymbol{\mu}) := \langle \boldsymbol{\mu}, \ell_x(A) \rangle + D_{\Phi}(\boldsymbol{\mu}, \hat{\mathbf{q}}) - \ell_x(\mathbf{a}),$$

where $x \in [n]$, $A \in \mathcal{A}^k$, $\mathbf{a} \in \mathcal{A}$, and $\hat{\mathbf{q}}, \mathbf{q} \in \Delta_k$. Reid et al. (2015) showed that ℓ is Φ mixable if and only if $\Delta M_{\Phi} := \inf_{A \in \mathcal{A}^k, \hat{\mathbf{q}} \in \Delta_k} \sup_{\mathbf{a}_* \in \mathcal{A}} \inf_{\boldsymbol{\mu} \in \Delta_k, x \in [n]} \Delta m_{\Phi}(x, A, \mathbf{a}, \hat{\mathbf{q}}, \boldsymbol{\mu}) \geq 0$.

Proof [Proposition 11]

[We show that ℓ is $\Phi^{\mathcal{I}}$ -mixable] Let $\mathcal{I} \subseteq [k]$, with $|\mathcal{I}| > 1$, $A \in \mathcal{A}^k$, and $\mathbf{q} \in \Delta_{\mathcal{I}}$. Since ℓ is Φ -mixable, the following holds

$$\exists \mathbf{a}_* \in \Delta_n, \forall x \in [n], \ell_x(\mathbf{a}_*) \leq \inf_{\hat{\mathbf{q}} \in \Delta_k} \langle \hat{\mathbf{q}}, \ell_x(A) \rangle + D_{\Phi}(\hat{\mathbf{q}}, \mathbf{q}), \quad (36)$$

$$\leq \inf_{\hat{\mathbf{q}} \in \Delta_{\mathcal{I}}} \langle \hat{\mathbf{q}}, \ell_x(A) \rangle + D_{\Phi}(\hat{\mathbf{q}}, \mathbf{q}), \quad (37)$$

$$\begin{aligned} &= \inf_{\hat{\mathbf{q}} \in \Delta_{\mathcal{I}}} \left\langle \Pi_{\mathcal{I}}^k \hat{\mathbf{q}}, \Pi_{\mathcal{I}}^k \ell_x(A) \right\rangle + D_{\Phi^{\mathcal{I}}}(\Pi_{\mathcal{I}}^k \hat{\mathbf{q}}, \Pi_{\mathcal{I}}^k \mathbf{q}), \\ &= \inf_{\hat{\boldsymbol{\mu}} \in \Delta_{|\mathcal{I}|}} \left\langle \hat{\boldsymbol{\mu}}, \ell_x(A[\Pi_{\mathcal{I}}^k]^{\top}) \right\rangle + D_{\Phi^{\mathcal{I}}}(\hat{\boldsymbol{\mu}}, \Pi_{\mathcal{I}}^k \mathbf{q}), \end{aligned} \quad (38)$$

where in (36) we used the fact that $\Phi^{\mathcal{I}}(\Pi_{\mathcal{I}}^k \mathbf{q}) = \Phi(\mathbf{q}), \forall \mathbf{q} \in \Delta_{\mathcal{I}}$. Given that $A \mapsto A[\Pi_{\mathcal{I}}^k]^{\top}$ [resp. $\mathbf{q} \mapsto \Pi_{\mathcal{I}}^k \mathbf{q}$] is onto from \mathcal{A}^k to $\mathcal{A}^{|\mathcal{I}|}$ [resp. from $\Delta_{\mathcal{I}}$ to $\Delta_{|\mathcal{I}|}$], (38) implies that ℓ is $\Phi^{\mathcal{I}}$ -mixable.

[We show (8)] Suppose that there exists $\hat{\mathbf{q}} \in \text{rbd } \Delta_k$ and $\mathbf{q} \in \text{ri } \Delta_k$ such that $|\Phi'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}})| < +\infty$. Let $f: [0, \epsilon] \rightarrow \mathbb{R}$ be defined by $f(\lambda) := \Phi(\hat{\mathbf{q}} + \lambda(\mathbf{q} - \hat{\mathbf{q}}))$, where $\epsilon > 0$ is such that $\hat{\mathbf{q}} + \epsilon(\mathbf{q} - \hat{\mathbf{q}}) \in \text{ri } \Delta_k$. The function f is closed and convex on $\text{dom } f = [0, \epsilon]$ and $\lim_{\lambda \downarrow 0} \frac{f(\lambda) - f(0)}{\lambda} = f'(0; 1) = \Phi'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}})$ which is finite by assumption. Using this and the fact that $\lambda f'(0; 1) = f'(0; \lambda)$, we have $\lim_{\lambda \downarrow 0} \lambda^{-1}(f(\lambda) - f(0) - f'(0; \lambda)) = 0$. Substituting f by its expression in terms of Φ in the latter equality gives

$$\lim_{\lambda \downarrow 0} \lambda^{-1} D_{\Phi}(\hat{\mathbf{q}} + \lambda(\mathbf{q} - \hat{\mathbf{q}}), \hat{\mathbf{q}}) = 0. \quad (39)$$

Let $\eta > 0$ and $\theta^* \in [k]$ be such that $\hat{q}_{\theta^*} = 0$. Suppose that ℓ is a non-trivial Φ -mixable loss. Let x_1, x_0, \mathbf{a}_1 , and \mathbf{a}_0 be as in the definition of non-trivial (see (3)). In particular, it holds that $\ell_{x_0}(\mathbf{a}_0) < \ell_{x_0}(\mathbf{a}_1)$. Fix $A \in \mathcal{A}^k$, such that $A_{\cdot, \theta^*} = \mathbf{a}_0$ and $A_{\cdot, \theta} = \mathbf{a}_1$ for $\theta \in [k] \setminus \{\theta^*\}$. Let

$$\mathbf{a}_* = \operatorname{argmax}_{\mathbf{a} \in \Delta_n} \inf_{\boldsymbol{\mu} \in \Delta_k, x \in [n]} \Delta m_{\Phi}(x, A, \mathbf{a}, \hat{\mathbf{q}}, \boldsymbol{\mu}),$$

with $\hat{\mathbf{q}} \in \text{rbd } \Delta_k$ as in (39). Note that \mathbf{a}_* exists since ℓ is closed.

If \mathbf{a}_* is such that $\ell_{x_1}(\mathbf{a}_*) > \ell_{x_1}(\mathbf{a}_1)$, then taking $\boldsymbol{\mu} = \hat{\mathbf{q}}$ puts all weights on experts predicting \mathbf{a}_1 , while $D_{\Phi}(\boldsymbol{\mu}, \hat{\mathbf{q}}) = 0$. Therefore,

$$\Delta M_{\Phi} \leq \inf_{\boldsymbol{\mu} \in \Delta_k, x \in [n]} \Delta m_{\Phi}(x, A, \mathbf{a}_*, \hat{\mathbf{q}}, \boldsymbol{\mu}) \leq \Delta m_{\Phi}(x_1, A, \mathbf{a}, \hat{\mathbf{q}}, \hat{\mathbf{q}}) < 0.$$

This contradicts the Φ -mixability of ℓ . Therefore, $\ell_{x_1}(\mathbf{a}_*) = \ell_{x_1}(\mathbf{a}_1)$, which by (3) implies $\ell_{x_0}(\mathbf{a}_*) \geq \ell_{x_0}(\mathbf{a}_1)$. For $\mathbf{q}^\lambda = \hat{\mathbf{q}} + \lambda(\mathbf{q} - \hat{\mathbf{q}})$, with $\mathbf{q} \in \text{ri } \Delta_k$ as in (36) and $\lambda \in [0, \epsilon]$,

$$\begin{aligned} \Delta M_\Phi &\leq \inf_{\mu \in \Delta_k, x \in [n]} \Delta m_\Phi(x, A, \mathbf{a}_*, \hat{\mathbf{q}}, \mu), \\ &\leq \Delta m_\Phi(x_0, A, \mathbf{a}, \hat{\mathbf{q}}, \mathbf{q}^\lambda), \\ &= \langle \mathbf{q}^\lambda, \ell_{x_0}(A) \rangle + D_\Phi(\mathbf{q}^\lambda, \hat{\mathbf{q}}) - \ell_{x_0}(\mathbf{a}_*), \\ &= (1 - \lambda q_{\theta^*}) \ell_{x_0}(\mathbf{a}_1) + \lambda q_{\theta^*} \ell_{x_0}(\mathbf{a}_0) + D_\Phi(\mathbf{q}^\lambda, \hat{\mathbf{q}}) - \ell_{x_0}(\mathbf{a}_*), \\ &\leq (1 - \lambda q_{\theta^*}) \ell_{x_0}(\mathbf{a}_*) + \lambda q_{\theta^*} \ell_{x_0}(\mathbf{a}_0) + D_\Phi(\mathbf{q}^\lambda, \hat{\mathbf{q}}) - \ell_{x_0}(\mathbf{a}_*), \\ &= \lambda q_{\theta^*} (\ell_{x_0}(\mathbf{a}_0) - \ell_{x_0}(\mathbf{a}_*)) + D_\Phi(\hat{\mathbf{q}} + \lambda(\mathbf{q} - \hat{\mathbf{q}}), \hat{\mathbf{q}}). \end{aligned}$$

Since $q_{\theta^*} > 0$ ($\mathbf{q} \in \text{ri } \Delta_k$) and $\ell_{x_0}(\mathbf{a}_0) < \ell_{x_0}(\mathbf{a}_1) \leq \ell_{x_0}(\mathbf{a}_*)$, (36) implies that there exists $\lambda_* > 0$ small enough such that $\lambda_* q_{\theta^*} (\ell_{x_0}(\mathbf{a}_0) - \ell_{x_0}(\mathbf{a}_*)) + D_\Phi(\hat{\mathbf{q}} + \lambda_*(\mathbf{q} - \hat{\mathbf{q}}), \hat{\mathbf{q}}) < 0$. But this implies that $\Delta M_\Phi < 0$ which contradicts the Φ -mixability of ℓ . Therefore, $\Phi'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}})$ is either equal to $+\infty$ or $-\infty$. The former case is not possible. In fact, since Φ is convex, it must have non-decreasing slopes; In particular, it holds that $\Phi'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}}) \leq \Phi(\mathbf{q} - \hat{\mathbf{q}}) - \Phi(\hat{\mathbf{q}})$. Since Φ is finite on Δ_k (by definition of an entropy), we have $\Phi'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}}) < +\infty$. Therefore, we have just shown that

$$\forall \hat{\mathbf{q}} \in \text{rbd } \Delta_k, \forall \mathbf{q} \in \text{ri } \Delta_k, \Phi'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}}) = -\infty. \quad (40)$$

Now suppose that $(\hat{\mathbf{q}}, \mathbf{q}) \in (\text{rbd } \Delta_{\mathcal{I}}) \times (\text{ri } \Delta_{\mathcal{I}})$ for $\mathcal{I} \subseteq [k]$, with $|\mathcal{I}| > 1$. Note that in this case, we have $(\Phi^{\mathcal{I}})'(\Pi_{\mathcal{I}}^k \hat{\mathbf{q}}; \Pi_{\mathcal{I}}^k (\mathbf{q} - \hat{\mathbf{q}})) = \Phi'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}})$. We showed in the first step of this proof that under the assumptions of the proposition, ℓ must be $\Phi^{\mathcal{I}}$ -mixable. Therefore, repeating the steps above that lead to (40) for Φ , $\hat{\mathbf{q}}$, and \mathbf{q} substituted by $\Phi^{\mathcal{I}}$, $\Pi_{\mathcal{I}}^k \hat{\mathbf{q}} \in \text{rbd } \Delta_{|\mathcal{I}|}$, and $\Pi_{\mathcal{I}}^k \mathbf{q} \in \text{ri } \Delta_{|\mathcal{I}|}$, we obtain $\Phi'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}}) = (\Phi^{\mathcal{I}})'(\Pi_{\mathcal{I}}^k \hat{\mathbf{q}}; \Pi_{\mathcal{I}}^k (\mathbf{q} - \hat{\mathbf{q}})) = -\infty$. This shows (8). \blacksquare

C.7. Proof of Proposition 12

Let “sgn” denote the *sign* function.

Proof [Proposition 12] Let $\mathcal{I} = \{1, 2\}$. Since ℓ is Φ -mixable, it must be $\Phi^{\mathcal{I}}$ -mixable, where $\Phi^{\mathcal{I}} := \Phi^{\mathcal{I}} \circ [\Pi_{\mathcal{I}}^k]^\top : \mathbb{R}^2 \rightarrow \mathbb{R} \cup \{+\infty\}$ (Proposition 11). Let $\Psi := \Phi^{\mathcal{I}}$.

For $w \in]0, +\infty[$ and $z \in \text{int dom } \tilde{\Psi}^* = \mathbb{R}$ (see §2.2), we define $(\tilde{\Psi}^*)'_\infty(w) := \lim_{t \rightarrow +\infty} [\tilde{\Psi}^*(z + tw) - \tilde{\Psi}^*(z)]/t$. The value of $(\tilde{\Psi}^*)'_\infty(w)$ does not depend on the choice of z , and it holds that $(\tilde{\Psi}^*)'_\infty(w) = \sigma_{\text{dom } \tilde{\Psi}}(w)$ and $(\tilde{\Psi}^*)'_\infty(-w) = \sigma_{\text{dom } \tilde{\Psi}}(-w)$ (Hiriart-Urruty and Lemaréchal, 2001, Prop. C.1.2.2). In our case, we have $\text{dom } \tilde{\Psi} = [0, 1]$ (by definition of $\tilde{\Psi}$), which implies that $\sigma_{\text{dom } \tilde{\Psi}}(1) = 1$ and $\sigma_{\text{dom } \tilde{\Psi}}(-1) = 0$. Therefore, $(\tilde{\Psi}^*)'_\infty(1) + (\tilde{\Psi}^*)'_\infty(-1) = 1$. As a result $\tilde{\Psi}^*$ cannot be affine; that is, it cannot have a constant slope between any two points in \mathbb{R} . For $\forall \delta > 0$, let $g_\delta : \mathbb{R} \times \{-1, 0, +1\} \rightarrow \mathbb{R}$ be defined by $g_\delta(s, u) := [\tilde{\Psi}^*(s + \delta(u + 1)/2) - \tilde{\Psi}^*(s + \delta(u - 1)/2)]/\delta$. Since $\tilde{\Psi}^*$ is convex it must have non-decreasing slopes (ibid., p.13). Combining this with the fact that $\tilde{\Psi}^*$ is not affine implies that

$$\exists s_\delta^* \in \mathbb{R}, g_\delta(s_\delta^*, -1) < g_\delta(s_\delta^*, +1). \quad (41)$$

The fact that $\tilde{\Psi}^*$ has non-decreasing slopes also implies that

$$g_\delta(s_\delta^*, +1) = [\tilde{\Psi}^*(s_\delta^* + \delta) - \tilde{\Psi}^*(s_\delta^*)]/\delta \leq \lim_{t \rightarrow \infty} [\tilde{\Psi}^*(s_\delta^* + t) - \tilde{\Psi}^*(s_\delta^*)]/t = (\tilde{\Psi}^*)'_\infty(1) = 1.$$

Similarly, we have $0 = -(\tilde{\Psi}^*)'_\infty(-1) \leq g_\delta(s_\delta^*, -1)$. Let $\tilde{\mu} \in \partial\tilde{\Psi}^*(s_\delta^*)$. Since $\tilde{\Psi}$ is a closed convex function the following equivalence holds $\tilde{\mu} \in \partial\tilde{\Psi}^*(s_\delta^*) \iff s_\delta^* \in \partial\tilde{\Psi}(\tilde{\mu})$ (ibid., Cor. D.1.4.4). Thus, if $\tilde{\mu} \in \{0, 1\} = \text{bd } \tilde{\Delta}_2$, then $\partial\tilde{\Psi}(\tilde{\mu}) \neq \emptyset$, which is not possible since ℓ is Ψ -mixable (Lemma 30).

[We show $\underline{L}_\ell \in C^1(]0, +\infty[^n)$] We will now show that \underline{L}_ℓ is continuously differentiable on $]0, +\infty[^n$. Since \underline{L}_ℓ is 1-homogeneous, it suffice to check the differentiability on $\text{ri } \Delta_n$. Suppose \underline{L}_ℓ is not differentiable at $\mathbf{p} \in \text{ri } \Delta_n$. From Theorem 4, there exists $\mathbf{a}_0, \mathbf{a}_1 \in \mathcal{A}$ such that $\ell(\mathbf{a}_0), \ell(\mathbf{a}_1) \in \partial\sigma_{\mathcal{S}_\ell^\oplus}(-\mathbf{p})$ and $\ell(\mathbf{a}_0) \neq \ell(\mathbf{a}_1)$. Let $A := [\mathbf{a}_0, \mathbf{a}_1] \in \mathbb{R}^{n \times 2}$, $\delta := \min\{|\ell_x(\mathbf{a}_0) - \ell_x(\mathbf{a}_1)| : x \in [n], |\ell_x(\mathbf{a}_0) - \ell_x(\mathbf{a}_1)| > 0\}$, and $s_\delta^* \in \mathbb{R}$ as in (41). We denote $g^- := g_\delta(s_\delta^*, -1)$ and $g^+ := g_\delta(s_\delta^*, +1) \in]0, 1]$. Let $\tilde{\mu} \in \partial\tilde{\Psi}^*(s_\delta^*) \in \text{int } \tilde{\Delta}_2$ and $\boldsymbol{\mu} = \Pi_2(\tilde{\mu}) \in \text{ri } \Delta_2$. From the fact that ℓ is Ψ -mixable, $J_2^\top \ell_x(A) = \ell_x(\mathbf{a}_0) - \ell_x(\mathbf{a}_1)$, and (9), there must exist $\mathbf{a}_* \in \mathcal{A}$ such that for all $x \in [n]$,

$$\begin{aligned} \ell_x(\mathbf{a}_*) &\leq M_\Psi(\ell_x(A), \boldsymbol{\mu}) = \ell_x(\mathbf{a}_1) + \tilde{\Psi}^*(s_\delta^*) - \tilde{\Psi}^*(s_\delta^* - \ell_x(\mathbf{a}_0) + \ell_x(\mathbf{a}_1)), \\ &\leq \ell_x(\mathbf{a}_1) + g_\delta(s_\delta^*, -\text{sgn}[\ell_x(\mathbf{a}_0) - \ell_x(\mathbf{a}_1)])[\ell_x(\mathbf{a}_0) - \ell_x(\mathbf{a}_1)], \end{aligned} \quad (42)$$

where in (42) we used the fact that $\tilde{\Psi}^*$ has non-decreasing slopes and the definition of δ . When $\ell_x(\mathbf{a}_0) \leq \ell_x(\mathbf{a}_1)$, (42) becomes $\ell_x(\mathbf{a}_*) \leq (1 - g^+)\ell_x(\mathbf{a}_1) + g^+\ell_x(\mathbf{a}_0)$. Otherwise, we have $\ell_x(\mathbf{a}_*) \leq (1 - g^-)\ell_x(\mathbf{a}_1) + g^-\ell_x(\mathbf{a}_0) < (1 - g^+)\ell_x(\mathbf{a}_1) + g^+\ell_x(\mathbf{a}_0)$. Since ℓ is admissible, there must exist at least one $x \in [n]$ such that $\ell_x(\mathbf{a}_0) > \ell_x(\mathbf{a}_1)$. Combining this with the fact that $p_x > 0, \forall x \in [n]$ ($\mathbf{p} \in \text{ri } \Delta_n$), implies that $\langle \mathbf{p}, \ell(\mathbf{a}_*) \rangle < \langle \mathbf{p}, (1 - g^+)\ell(\mathbf{a}_1) + g^+\ell(\mathbf{a}_0) \rangle = \underline{L}_\ell(\mathbf{p})$. This contradicts the fact that $\ell(\mathbf{a}_*) \in \mathcal{S}_\ell^\oplus$. Therefore, \underline{L}_ℓ must be differentiable at \mathbf{p} . As argued earlier, this implies that \underline{L}_ℓ must be differentiable on $]0, +\infty[^n$. Combining this with the fact that \underline{L}_ℓ is concave on $]0, +\infty[^n$, implies that \underline{L}_ℓ is continuously differentiable on $]0, +\infty[^n$ (ibid., Rmk. D.6.2.6).

[We show $\tilde{\Phi}^* \in C^1(\mathbb{R}^{k-1})$] Suppose that $\tilde{\Phi}^*$ is not differentiable at some $\mathbf{s}^* \in \mathbb{R}^{k-1}$. Then there exists $\mathbf{d} \in \mathbb{R}^{k-1} \setminus \{\mathbf{0}_{\tilde{k}}\}$ such that $(-\tilde{\Phi}^*)'(\mathbf{s}^*; -\mathbf{d}) < (\tilde{\Phi}^*)'(\mathbf{s}^*; \mathbf{d})$. Since $\mathbf{s}^* \in \text{int dom } \tilde{\Phi}^*$, $(\tilde{\Phi}^*)'(\mathbf{s}^*, \cdot)$ is finite and convex (Hiriart-Urruty and Lemaréchal, 2001, Prop. D.1.1.2), and thus it is continuous on $\text{dom } \tilde{\Phi}^* = \mathbb{R}^{k-1}$ (ibid., Rmk. B.3.1.3). Consequently, there exists $\delta^* > 0$ such that

$$\forall \hat{\mathbf{d}} \in \mathbb{R}^{n-1}, \|\hat{\mathbf{d}} - \mathbf{d}\| \leq \delta^* \implies -(\tilde{\Phi}^*)'(\mathbf{s}^*; -\hat{\mathbf{d}}) < (\tilde{\Phi}^*)'(\mathbf{s}^*; \hat{\mathbf{d}}) \quad (43)$$

Let $g: \{-1, 1\} \rightarrow \mathbb{R}$ such that $g(u) := \sup_{\|\hat{\mathbf{d}} - \mathbf{d}\| \leq \delta^*} u(\tilde{\Phi}^*)'(\mathbf{s}^*; u\hat{\mathbf{d}})$. Note that since $\tilde{\Phi}^*$ has increasing slopes ($\tilde{\Phi}^*$ is convex), $g(1) \leq \sup_{\|\hat{\mathbf{d}} - \mathbf{d}\| \leq \delta^*} (\tilde{\Phi}^*)'_\infty(\hat{\mathbf{d}}) = \sup_{\|\hat{\mathbf{d}} - \mathbf{d}\| \leq \delta^*} \sigma_{\text{dom } \tilde{\Phi}^*}(\hat{\mathbf{d}}) \leq 1$, where the last inequality holds because $\tilde{\Delta}_k \subset \mathcal{B}(\mathbf{0}_{\tilde{k}}, 1)$, and thus $\sigma_{\text{dom } \tilde{\Phi}^*}(\hat{\mathbf{d}}) = \sigma_{\tilde{\Delta}_k}(\hat{\mathbf{d}}) \leq \sigma_{\mathcal{B}(\mathbf{0}_{\tilde{k}}, 1)}(\hat{\mathbf{d}}) = 1$. Let $\Delta g := g(1) - g(-1)$. From (40), it is clear that $\Delta g > 0$.

Suppose that \underline{L}_ℓ is twice differentiable on $]0, +\infty[^n$ and let $\underline{\ell}$ be a support loss of ℓ . By definition of a support loss, $\forall \mathbf{p} \in \text{ri } \Delta_k, \underline{\ell}(\tilde{\mathbf{p}}) = \underline{\ell}(\mathbf{p}) = \nabla \underline{L}_\ell(\mathbf{p})$ (where $\underline{\ell} := \underline{\ell} \circ \Pi_n$). Thus, since \underline{L}_ℓ is twice differentiable on $]0, +\infty[^n$, $\underline{\ell}$ is differentiable on $\text{int } \tilde{\Delta}_n$. Furthermore, $\underline{\ell}$ is continuous on $\text{ri } \Delta_k$

given that $\underline{L}_\ell \in C^1(]0, +\infty[^n)$ as shown in the first part of this proof. Since ℓ is non-trivial, it is non-constant. Therefore, from Theorem 4, $\underline{\ell}$ is also non-constant. Consequently, the mean value theorem applied to $\underline{\ell}$ (see e.g. (Rudin, 1964, Thm. 5.10)) between any two points in $\text{ri } \Delta_n$ with distinct image under $\underline{\ell}$, implies that there exists $(\tilde{\mathbf{p}}_*, \mathbf{v}_*) \in \text{int } \tilde{\Delta}_n \times \mathbb{R}^{n-1}$, such that $\mathbf{D}\tilde{\underline{\ell}}(\tilde{\mathbf{p}}_*)\mathbf{v}_* \neq \mathbf{0}_{\tilde{n}}$. For the rest of the proof let $(\tilde{\mathbf{p}}, \mathbf{v}) := (\tilde{\mathbf{p}}_*, \mathbf{v}_*)$ and define $\mathcal{J} := \{x \in [n] : \mathbf{D}\tilde{\underline{\ell}}_x(\tilde{\mathbf{p}})\mathbf{v} \neq 0\}$. From Lemma 27, we have $\langle \mathbf{p}, \mathbf{D}\tilde{\underline{\ell}}(\tilde{\mathbf{p}}) \rangle = \mathbf{0}_{\tilde{n}}^\top$, which implies $\exists x \in \mathcal{J}, \mathbf{D}\tilde{\underline{\ell}}_x(\tilde{\mathbf{p}})\mathbf{v} < 0$. Let $\tilde{\mathbf{p}}^t := \tilde{\mathbf{p}} + t\mathbf{v}$. From Taylor's Theorem (see e.g. (Hardy, 2008, §151)) applied to the function $t \mapsto \tilde{\underline{\ell}}(\tilde{\mathbf{p}}^t)$, there exists $\epsilon^* > 0$ and functions $\delta_x : [-\epsilon^*, \epsilon^*] \rightarrow \mathbb{R}^n, x \in [n]$, such that $\lim_{t \rightarrow 0} t^{-1}\delta_x(t) = 0$ and

$$\forall |t| \leq \epsilon^*, \quad \underline{\ell}_x(\mathbf{p}^t) = \underline{\ell}_x(\mathbf{p}) + t\mathbf{D}\tilde{\underline{\ell}}_x(\tilde{\mathbf{p}})\mathbf{v} + \delta_x(t). \quad (44)$$

By shrinking ϵ^* if necessary, we may also assume that

$$\forall x \in \mathcal{J}, \forall \theta \in [k], \forall |t| \leq \epsilon^*, \quad t^{-1}\delta_x(td_\theta) \leq \frac{\delta^* |\mathbf{D}\tilde{\underline{\ell}}_x(\tilde{\mathbf{p}})\mathbf{v}|}{\sqrt{n} \|\mathbf{d}\|}, \quad (45)$$

$$\forall x \notin \mathcal{J}, \forall \theta \in [k], \forall |t| \leq \epsilon^*, \quad \delta_x(t) \leq \frac{\Delta g}{2} \min_{x \in \mathcal{J}} \{p_x\}. \quad (46)$$

Note that $\min_{x \in \mathcal{J}} \{p_x\} > 0$, since $\mathbf{p} \in \text{ri } \Delta_n$. Let $\lambda_\theta := t^* \frac{d_\theta}{\|\mathbf{d}\|}$, for $\theta \in [k-1]$. From Theorem 4, there exists $[\mathbf{a}_\theta^*]_{\theta \in [k]} \in \mathcal{A}^k$, such that $\ell(\mathbf{a}_k) = \underline{\ell}(\mathbf{p})$ and $\ell(\mathbf{a}_\theta) = \underline{\ell}(\mathbf{p}^{\lambda_\theta}) = \underline{\ell}(\mathbf{p}) + t^* \frac{d_\theta}{\|\mathbf{d}\|} \mathbf{D}\tilde{\underline{\ell}}(\tilde{\mathbf{p}})\mathbf{v} + \delta\left(t^* \frac{d_\theta}{\|\mathbf{d}\|}\right)$, where $[\delta(\cdot)]_x := \delta_x(\cdot)$. Let $A := [\mathbf{a}_\theta]_{\theta \in [k]}$. From the fact that ℓ is Φ -mixable, it follows that there exists $\mathbf{a}_* \in \mathcal{A}$ such that for all $x \in [n]$,

$$\ell_x(\mathbf{a}_*) \leq \mathbf{M}_\Phi(\ell_x(A), \boldsymbol{\mu}) = \ell_x(\mathbf{a}_k) + \tilde{\Phi}^*(\mathbf{s}^*) - \tilde{\Phi}^*(\mathbf{s}^* - J_k^\top \ell_x(A)). \quad (47)$$

Note that for all $x \in [n]$, $J_k^\top \ell_x(A) = [\ell_x(\mathbf{a}_\theta) - \ell_x(\mathbf{a}_k)]_{\theta \in [\tilde{k}]}$. Thus, $\forall x \in \mathcal{J}, \forall \theta \in [\tilde{k}], [J_k^\top \ell_x(A)]_\theta = \left[\frac{t^*}{\|\mathbf{d}\|} \mathbf{D}\tilde{\underline{\ell}}_x(\tilde{\mathbf{p}})\mathbf{v} \right] d_{x,\theta}$, where $d_{x,\theta} := d_\theta + \frac{\|\mathbf{d}\|}{t^* |\mathbf{D}\tilde{\underline{\ell}}_x(\tilde{\mathbf{p}})\mathbf{v}|} \delta_x\left(t^* \frac{d_\theta}{\|\mathbf{d}\|}\right)$. On the other hand, if $\theta \notin \mathcal{J}$, then $[J_k^\top \ell_x(A)]_\theta = \delta_x\left(t^* \frac{d_\theta}{\|\mathbf{d}\|}\right)$. From (45), we have $\|\mathbf{d}_x - \mathbf{d}\| \leq \delta_*, \forall x \in \mathcal{J}$, and from the monotonicity of the slopes of $\tilde{\Phi}^*$

$$\begin{aligned} \forall x \in \mathcal{J}, \tilde{\Phi}^*(\mathbf{s}^*) - \tilde{\Phi}^*(\mathbf{s}^* - J_k^\top \ell_x(A)) &\leq (\tilde{\Phi}^*)' \left(\mathbf{s}^*; \frac{t^*}{\|\mathbf{d}\|} [\mathbf{D}\tilde{\underline{\ell}}_x(\tilde{\mathbf{p}})\mathbf{v}] \mathbf{d}_x \right), \\ &\leq \frac{t^*}{\|\mathbf{d}\|} [\mathbf{D}\tilde{\underline{\ell}}_x(\tilde{\mathbf{p}})\mathbf{v}] g(\text{sgn}[\mathbf{D}\tilde{\underline{\ell}}_x(\tilde{\mathbf{p}})\mathbf{v}]), \\ &\leq \frac{t^*}{\|\mathbf{d}\|} [\mathbf{D}\tilde{\underline{\ell}}_x(\tilde{\mathbf{p}})\mathbf{v}] g(1) - \Delta g \llbracket \mathbf{D}\tilde{\underline{\ell}}_x(\tilde{\mathbf{p}})\mathbf{v} < 0 \rrbracket, \end{aligned} \quad (48)$$

where $\llbracket \cdot \rrbracket$ denotes the Iverson bracket. Equation 48 follows from the definition of Δg . Combining (48) with (47) yields

$$\begin{aligned} \langle \mathbf{p}, \ell(\mathbf{a}_*) \rangle &\leq \langle \mathbf{p}, \ell(\mathbf{a}_k) \rangle + \frac{t^*}{\|\mathbf{d}\|} \langle \mathbf{p}, \mathbf{D}\tilde{\underline{\ell}}(\tilde{\mathbf{p}})\mathbf{v} \rangle g(1) - \Delta g \min_{x' \in \mathcal{J}} \{p_{x'}\} + \sum_{x \notin \mathcal{J}} p_x \delta_x \left(t^* \frac{d_\theta}{\|\mathbf{d}\|} \right), \\ &\leq \langle \mathbf{p}, \ell(\mathbf{a}_k) \rangle - \frac{\Delta g}{2} \min_{x \in \mathcal{J}} \{p_x\}, \end{aligned} \quad (49)$$

$$< \langle \mathbf{p}, \underline{\ell}(\mathbf{p}) \rangle, \quad (50)$$

where in the first inequality we used the fact that $\exists x \in \mathcal{J}, D\tilde{\ell}_x(\tilde{\mathbf{p}})\mathbf{v} < 0$. In (49) we used (45) and the fact that $\langle \mathbf{p}, D\tilde{\ell}(\tilde{\mathbf{p}}) \rangle = \mathbf{0}_n^\top$ (Lemma 27). Finally, (50) follows from the fact that $\min_{x \in \mathcal{J}} \{p_x\} > 0$ ($\mathbf{p} \in \text{ri } \Delta_n$) and $\tilde{\ell}(\mathbf{p}) = \ell(\mathbf{a}_k)$ (by definition of \mathbf{a}_k). Equation 50 shows that $\ell(\mathbf{a}^*) \notin \mathcal{S}_\ell^\oplus$, which is a contradiction. \blacksquare

C.8. Proof of Proposition 13

Proof First note that \mathbf{q}_* is well defined; the infimum of $\boldsymbol{\mu} \mapsto \langle \boldsymbol{\mu}, \mathbf{q} \rangle + D_\Phi(\boldsymbol{\mu}, \mathbf{q})$ is attained for $\mathbf{q} \in \text{ri } \Delta_k$ (see Remark 9).

Now since $\tilde{\Phi}$ is convex and $\tilde{\mathbf{q}} = \Pi_k(\mathbf{q}) \in \text{int } \tilde{\Delta}_k = \text{int } \text{dom } \tilde{\Phi}$, we have $\partial\tilde{\Phi}(\tilde{\mathbf{q}}) \neq \emptyset$ (Rockafellar, 1997, Thm. 23.4). This means that there exists $\mathbf{s}_q^* \in \partial\tilde{\Phi}(\tilde{\mathbf{q}})$ such that $\langle \mathbf{s}_q^*, \tilde{\mathbf{q}}_* - \tilde{\mathbf{q}} \rangle = \tilde{\Phi}'(\tilde{\mathbf{q}}; \tilde{\mathbf{q}}_* - \tilde{\mathbf{q}})$ (Hiriart-Urruty and Lemaréchal, 2001, p.166). We will now show that $\mathbf{s}_q^* - J_k^\top \mathbf{d} \in \partial\tilde{\Phi}(\tilde{\mathbf{q}}_*)$, which will imply that $\tilde{\mathbf{q}}_* \in \partial\tilde{\Phi}^*(\mathbf{s}_q^* - J_k^\top \mathbf{d})$ (ibid., Cor. D.1.4.4). Let $\mathbf{q}_* = \text{argmin}_{\boldsymbol{\mu} \in \Delta_k} \langle \boldsymbol{\mu}, \mathbf{d} \rangle + D_\Phi(\boldsymbol{\mu}, \mathbf{q})$. Thus, for all $\boldsymbol{\mu} \in \Delta_k$,

$$\begin{aligned} & \langle \boldsymbol{\mu}, \mathbf{d} \rangle + \tilde{\Phi}(\tilde{\boldsymbol{\mu}}) - \tilde{\Phi}(\tilde{\mathbf{q}}) - \tilde{\Phi}'(\tilde{\mathbf{q}}; \tilde{\boldsymbol{\mu}} - \tilde{\mathbf{q}}) \geq \langle \mathbf{q}_*, \mathbf{d} \rangle + \tilde{\Phi}(\tilde{\mathbf{q}}_*) - \tilde{\Phi}(\tilde{\mathbf{q}}) - \langle \mathbf{s}_q^*, \tilde{\mathbf{q}}_* - \tilde{\mathbf{q}} \rangle, \\ \implies & \tilde{\Phi}(\tilde{\boldsymbol{\mu}}) \geq \tilde{\Phi}(\tilde{\mathbf{q}}_*) - \langle \tilde{\boldsymbol{\mu}} - \tilde{\mathbf{q}}_*, J_k^\top \mathbf{d} \rangle + \langle \mathbf{s}_q^*, \tilde{\mathbf{q}} - \tilde{\mathbf{q}}_* \rangle + \tilde{\Phi}'(\tilde{\mathbf{q}}; \tilde{\boldsymbol{\mu}} - \tilde{\mathbf{q}}), \\ \implies & \tilde{\Phi}(\tilde{\boldsymbol{\mu}}) \geq \tilde{\Phi}(\tilde{\mathbf{q}}_*) - \langle \tilde{\boldsymbol{\mu}} - \tilde{\mathbf{q}}_*, J_k^\top \mathbf{d} \rangle + \langle \mathbf{s}_q^*, \tilde{\mathbf{q}} - \tilde{\mathbf{q}}_* \rangle + \langle \mathbf{s}_q^*, \tilde{\boldsymbol{\mu}} - \tilde{\mathbf{q}} \rangle, \\ \implies & \tilde{\Phi}(\tilde{\boldsymbol{\mu}}) \geq \tilde{\Phi}(\tilde{\mathbf{q}}_*) + \langle \tilde{\boldsymbol{\mu}} - \tilde{\mathbf{q}}_*, \mathbf{s}_q^* - J_k^\top \mathbf{d} \rangle, \end{aligned}$$

where in the second line we used the fact that $\forall \mathbf{q} \in \Delta_k, \langle \mathbf{q}, \mathbf{d} \rangle = \langle \tilde{\mathbf{q}}, J_k^\top \mathbf{d} \rangle + d_k$, and in third line we used the fact that $\forall \mathbf{s} \in \partial\tilde{\Phi}(\tilde{\mathbf{q}}), \langle \mathbf{s}, \tilde{\boldsymbol{\mu}} - \tilde{\mathbf{q}} \rangle \leq \tilde{\Phi}'(\tilde{\mathbf{q}}; \tilde{\boldsymbol{\mu}} - \tilde{\mathbf{q}})$ (ibid.). This shows that $\mathbf{s}_q^* - J_k^\top \mathbf{d} \in \partial\tilde{\Phi}(\tilde{\mathbf{q}}_*)$.

Substituting $\tilde{\Phi}'(\tilde{\mathbf{q}}; \tilde{\mathbf{q}}_* - \tilde{\mathbf{q}})$ by $\langle \mathbf{s}_q^*, \mathbf{q}_* - \mathbf{q} \rangle$ in the expression of $M_\Phi(\mathbf{d}, \mathbf{q})$, we get

$$\begin{aligned} M_\Phi(\mathbf{d}, \mathbf{q}) &= d_k + \langle \tilde{\mathbf{q}}_*, J_k^\top \mathbf{d} \rangle + \tilde{\Phi}(\tilde{\mathbf{q}}_*) - \tilde{\Phi}(\tilde{\mathbf{q}}) - \langle \mathbf{s}_q^*, \tilde{\mathbf{q}}_* - \tilde{\mathbf{q}} \rangle, \\ &= d_k + \langle \mathbf{s}_q^*, \tilde{\mathbf{q}} \rangle - \tilde{\Phi}(\tilde{\mathbf{q}}) - [\langle \mathbf{s}_q^* - J_k^\top \mathbf{d}, \tilde{\mathbf{q}}_* \rangle - \tilde{\Phi}(\tilde{\mathbf{q}}_*)], \\ &= d_k + \tilde{\Phi}^*(\mathbf{s}_q^*) - \tilde{\Phi}^*(\mathbf{s}_q^* - J_k^\top \mathbf{d}), \end{aligned}$$

where in the last line we used the fact that $\tilde{\Phi}$ is a closed convex function, and thus $\forall \tilde{\mathbf{q}} \in \tilde{\Delta}_k, \mathbf{s} \in \partial\tilde{\Phi}(\tilde{\mathbf{q}}) \implies \tilde{\Phi}^*(\mathbf{s}) = \langle \mathbf{s}, \tilde{\mathbf{q}} \rangle - \tilde{\Phi}(\tilde{\mathbf{q}})$ (ibid., Cor. E.1.4.4). \blacksquare

C.9. Proof of Theorem 14

Proof Let $\mathbf{q} \in \Delta_k$ and $A := [\mathbf{a}_\theta]_{1 \leq \theta \leq k} \in \mathcal{A}^k$. We claim that for $x \in [n]$

$$-\eta^{-1} \log \langle \exp(-\eta \ell_x(A)), \mathbf{q} \rangle = M_{S_\eta}(\ell_x(A), \mathbf{q}). \quad (51)$$

Let $x \in [n]$. Since $\mathbf{q} \in \Delta_k$, \mathbf{q} is either a vertex of Δ_k or there exists $\mathcal{I} \subseteq [n]$, with $|\mathcal{I}| > 1$, such that $\mathbf{q} \in \text{ri } \Delta_{\mathcal{I}}$. Suppose the latter case holds. We showed (see (14)) that for $A' \in (\text{dom } \ell)^k$,

$$-\eta^{-1} \log \langle \exp(-\eta \ell_x(A')), \mathbf{q} \rangle = M_{S_\eta^{\mathcal{I}}}(\Pi_{\mathcal{I}}^k \ell_x(A'), \Pi_{\mathcal{I}}^k \mathbf{q}). \quad (52)$$

Since this result was obtained without any special assumptions on ℓ — we only assumed that $\ell_x(A') \in [0, +\infty[^k$ — then (52) still holds if we replace $\ell_x(A')$ by any $\mathbf{d} \in [0, +\infty[^k$; that is,

$$\begin{aligned} \forall \mathbf{d} \in [0, +\infty[^n, \quad -\eta^{-1} \log \langle \exp(-\eta \mathbf{d}), \mathbf{q} \rangle &= M_{S_{\mathcal{I}}^{\eta}}(\Pi_{\mathcal{I}}^k \mathbf{d}, \Pi_{\mathcal{I}}^k \mathbf{q}), \\ &= M_{S_{\eta}}(\mathbf{d}, \mathbf{q}), \end{aligned} \quad (53)$$

where (53) is implied by Lemma 30 since S_{η} satisfies (8) (see Lemma 29).

Fix $x \in [n]$ and let $\hat{\mathbf{d}} := \ell_x(A) \in [0, +\infty[^k$. From Lemma 31, there exists a sequence $(\hat{\mathbf{d}}_m) \subset [0, +\infty[^k$ converging to $\hat{\mathbf{d}}$ such that $M_{S_{\eta}}(\hat{\mathbf{d}}_m, \mathbf{q}) \xrightarrow{m \rightarrow \infty} M_{S_{\eta}}(\hat{\mathbf{d}}, \mathbf{q})$. Then, from (53)

$$\begin{aligned} -\eta^{-1} \log \langle \exp(-\eta \ell_x(A)), \mathbf{q} \rangle &= \lim_{m \rightarrow \infty} -\eta^{-1} \log \langle \exp(-\eta \hat{\mathbf{d}}_m), \mathbf{q} \rangle, \\ &= \lim_{m \rightarrow \infty} M_{S_{\eta}}(\hat{\mathbf{d}}_m, \mathbf{q}), \\ &= M_{S_{\eta}}(\hat{\mathbf{d}}, \mathbf{q}) = M_{S_{\eta}}(\ell_x(A), \mathbf{q}). \end{aligned} \quad (54)$$

Now suppose that \mathbf{q} is a vertex. Without loss of generality assume that $\mathbf{q} = \mathbf{e}_1$ and let $\boldsymbol{\mu} \in \Delta_k \setminus \{\mathbf{e}_1\}$. Then there exists $\mathcal{I}_* \subset [k]$, such that $(\mathbf{e}_1, \boldsymbol{\mu}) \in (\text{rbd } \Delta_{\mathcal{I}_*}) \times (\text{ri } \Delta_{\mathcal{I}_*})$ and by Lemma 29, $S'(\mathbf{e}_1; \boldsymbol{\mu} - \mathbf{e}_1) = -\infty$. Therefore, $\forall \mathbf{q} \in \Delta_k \setminus \{\mathbf{e}_1\}$, $D_{S_{\eta}}(\mathbf{q}, \mathbf{e}_1) = +\infty$, which implies

$$\begin{aligned} \forall x \in [n], M_{S_{\eta}}(\ell_x(A), \mathbf{e}_1) &= \inf_{\mathbf{q} \in \Delta_k} \langle \mathbf{q}, \ell_x(A) \rangle + D_{S_{\eta}}(\mathbf{q}, \mathbf{e}_1), \\ &= \langle \mathbf{e}_1, \ell_x(A) \rangle + D_{S_{\eta}}(\mathbf{e}_1, \mathbf{e}_1), \\ &= \langle \mathbf{e}_1, \ell_x(A) \rangle, \\ &= \ell_x(\mathbf{a}_1) = -\eta^{-1} \log \langle \exp(-\eta \ell_x(A)), \mathbf{e}_1 \rangle. \end{aligned} \quad (55)$$

Combining (55) and (54) proves the claim in (51). The desired equivalence follows trivially from the definitions of η -mixability and S_{η} -mixability. ■

C.10. Proof of Theorem 15

We need the following lemma to show Theorem 15.

Lemma 34 *Let Φ be as in Theorem 15. Then $\eta_{\ell}\Phi - S$ is convex on Δ_k only if Φ satisfies (8).*

Proof Let $\hat{\mathbf{q}} \in \text{rbd } \Delta_k$. Suppose that there exists $\mathbf{q} \in \text{ri } \Delta_k$ such that $\Phi'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}}) > -\infty$. Since Φ is convex, it must have non-decreasing slopes; in particular, it holds that $\Phi'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}}) \leq \Phi(\mathbf{q}) - \Phi(\hat{\mathbf{q}})$. Therefore, since Φ is finite on Δ_k (by definition of an entropy), we have $\Phi'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}}) < +\infty$. Since by assumption $\eta_{\ell}\Phi - S$ is convex and finite on the simplex, we can use the same argument to show that $[\eta_{\ell}\Phi - S]'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}}) = \eta_{\ell}\Phi'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}}) - S'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}}) < +\infty$. This is a contradiction since $S'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}}) = -\infty$ (Lemma 29). Therefore, it must hold that $\Phi'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}}) = -\infty$.

Suppose now that $(\hat{\mathbf{q}}, \mathbf{q}) \in (\text{rbd } \Delta_{\mathcal{I}}) \times (\text{ri } \Delta_{\mathcal{I}})$ for $\mathcal{I} \subseteq [k]$, with $|\mathcal{I}| > 1$. Let $\Phi^{\mathcal{I}} := \Phi \circ [\Pi_{\mathcal{I}}^k]^{\top}$ and $S^{\mathcal{I}} := S \circ [\Pi_{\mathcal{I}}^k]^{\top}$. Since $\eta_{\ell}\Phi - S$ is convex on Δ_k and $\Pi_{\mathcal{I}}^k$ is a linear function, $\eta_{\ell}\Phi^{\mathcal{I}} - S^{\mathcal{I}}$ is convex on $\Delta_{|\mathcal{I}|}$. Repeating the steps above for Φ and S substituted by $\Phi^{\mathcal{I}}$ and $S^{\mathcal{I}}$, respectively, we get that $(\Phi^{\mathcal{I}})'(\Pi_{\mathcal{I}}^k \hat{\mathbf{q}}; \Pi_{\mathcal{I}}^k \mathbf{q} - \Pi_{\mathcal{I}}^k \hat{\mathbf{q}}) = -\infty$. Since $(\Phi^{\mathcal{I}})'(\Pi_{\mathcal{I}}^k \hat{\mathbf{q}}; \Pi_{\mathcal{I}}^k \mathbf{q} - \Pi_{\mathcal{I}}^k \hat{\mathbf{q}}) = \Phi'(\hat{\mathbf{q}}; \mathbf{q} - \hat{\mathbf{q}})$ the proof is completed. ■

Proof [Theorem 15] Assume $\eta_\ell \Phi - S$ is convex on Δ_k . For this to hold, it is necessary that $\eta_\ell > 0$ since $-S$ is strictly concave. Let $\eta := \eta_\ell$ and $S_\eta := \eta^{-1}S$. Then $\tilde{S}_\eta = \eta^{-1}\tilde{S}$ and $\tilde{\Phi} - \tilde{S}_\eta = (\Phi - S_\eta) \circ \Pi_k$ is convex on $\tilde{\Delta}_k$, since $\Phi - S_\eta$ is convex on Δ_k and Π_k is affine.

Let $x \in [n]$, $A := [\mathbf{a}_\theta]_{\theta \in [k]}$, and $\mathbf{q} \in \Delta_k$. Suppose that $\mathbf{q} \in \text{ri } \Delta_k$ and let $\mathbf{s}_q^* \in \partial \tilde{\Phi}(\tilde{\mathbf{q}})$ be as in Proposition 13. Note that if $\ell_x(\mathbf{a}_\theta) = +\infty, \forall \theta \in [k]$, then the Φ -mixability condition (6) is trivially satisfied. Suppose, without loss of generality, that $\ell_x(\mathbf{a}_k) < +\infty$. Let $(\mathbf{d}_m) \subset [0, +\infty]^k$ be such that $\mathbf{d}_m \xrightarrow{m \rightarrow \infty} \mathbf{d} := \ell_x(A) \in [0, +\infty]^k$ and $M_\Psi(\mathbf{d}_m, \mathbf{q}) \xrightarrow{m \rightarrow \infty} M_\Psi(\mathbf{d}, \mathbf{q})$ for $\Psi \in \{\Phi, S_\eta\}$. This sequence is guaranteed to exist by Lemma 31.

Let $\tilde{\Upsilon}_q : \mathbb{R}^{k-1} \rightarrow \mathbb{R} \cup \{+\infty\}$ be defined by

$$\tilde{\Upsilon}_q(\tilde{\boldsymbol{\mu}}) := \tilde{S}_\eta(\tilde{\boldsymbol{\mu}}) + \langle \tilde{\boldsymbol{\mu}}, \mathbf{s}_q^* - \nabla \tilde{S}_\eta(\tilde{\mathbf{q}}) \rangle - \tilde{\Phi}^*(\mathbf{s}_q^*) + \tilde{S}_\eta^*(\nabla \tilde{S}_\eta(\tilde{\mathbf{q}})),$$

and it's Fenchel dual follows from Proposition 20 (i+ii):

$$\tilde{\Upsilon}_q^*(\mathbf{v}) = \tilde{S}_\eta^*(\mathbf{v} - \mathbf{s}_q^* + \nabla \tilde{S}_\eta(\tilde{\mathbf{q}})) + \tilde{\Phi}^*(\mathbf{s}_q^*) - \tilde{S}_\eta^*(\nabla \tilde{S}_\eta(\tilde{\mathbf{q}})),$$

After substituting \mathbf{v} by $\mathbf{s}_q^* - J_k^\top \mathbf{d}$ in the expression of $\tilde{\Upsilon}_q^*$ and rearranging, we get

$$\tilde{S}_\eta^*(\nabla \tilde{S}_\eta(\tilde{\mathbf{q}})) - \tilde{S}_\eta^*(\nabla \tilde{S}_\eta(\tilde{\mathbf{q}}) - J_k^\top \mathbf{d}_m) = \tilde{\Phi}^*(\mathbf{s}_q^*) - \tilde{\Upsilon}_q^*(\mathbf{s}_q^* - J_k^\top \mathbf{d}_m). \quad (56)$$

Since $\mathbf{s}_q^* \in \partial \tilde{\Phi}(\tilde{\mathbf{q}})$ and $\tilde{\Phi}$ is a closed convex function, combining Proposition 20-(iv) and the fact that $\tilde{\Phi}^{**} = \tilde{\Phi}$ (Hiriart-Urruty and Lemaréchal, 2001, Cor. E.1.3.6) yields $\langle \tilde{\mathbf{q}}, \mathbf{s}_q^* \rangle - \tilde{\Phi}^*(\mathbf{s}_q^*) = \tilde{\Phi}(\tilde{\mathbf{q}})$. Thus, after substituting $\tilde{\boldsymbol{\mu}}$ by $\tilde{\mathbf{q}}$ in the expression of $\tilde{\Upsilon}_q$, we get

$$\tilde{\Phi}(\tilde{\mathbf{q}}) = \tilde{\Upsilon}_q(\tilde{\mathbf{q}}). \quad (57)$$

On the other hand, $\tilde{\Phi} - \tilde{\Upsilon}_q$ is convex on $\tilde{\Delta}_k$, since $\tilde{\Upsilon}_q$ is equal to \tilde{S}_η plus an affine function. Thus, $\partial[\tilde{\Phi} - \tilde{\Upsilon}_q](\tilde{\mathbf{q}}) + \partial \tilde{\Upsilon}_q(\tilde{\mathbf{q}}) = \partial \tilde{\Phi}(\tilde{\mathbf{q}})$, since $\tilde{\Phi}$ and $\tilde{\Upsilon}_q$ are both convex (ibid., Thm. D.4.1.1). Since $\tilde{\Upsilon}_q$ is differentiable at $\tilde{\mathbf{q}}$, we have $\partial \tilde{\Upsilon}_q(\tilde{\mathbf{q}}) = \{\nabla \tilde{\Upsilon}_q(\tilde{\mathbf{q}})\} = \{\mathbf{s}_q^*\}$. Furthermore, since $\mathbf{s}_q^* \in \partial \tilde{\Phi}(\tilde{\mathbf{q}})$, then $\mathbf{0}_{\tilde{k}} \in \partial \tilde{\Phi}(\tilde{\mathbf{q}}) - \partial \tilde{\Upsilon}_q(\tilde{\mathbf{q}}) = \partial[\tilde{\Phi} - \tilde{\Upsilon}_q](\tilde{\mathbf{q}})$. Hence, $\tilde{\Phi} - \tilde{\Upsilon}_q$ attains a minimum at $\tilde{\mathbf{q}}$ (ibid., Thm. D.2.2.1). Due to this and (57), $\tilde{\Phi} \geq \tilde{\Upsilon}_q$, which implies that $\tilde{\Phi}^* \leq \tilde{\Upsilon}_q^*$ (Proposition 20-(iii)). Using this in (56) gives for all $m \in \mathbb{N}$

$$\begin{aligned} \tilde{S}_\eta^*(\nabla \tilde{S}_\eta(\tilde{\mathbf{q}})) - \tilde{S}_\eta^*(\nabla \tilde{S}_\eta(\tilde{\mathbf{q}}) - J_k^\top \mathbf{d}_m) &\leq \tilde{\Phi}^*(\mathbf{s}_q^*) - \tilde{\Phi}^*(\mathbf{s}_q^* - J_k^\top \mathbf{d}_m), \\ \implies M_{S_\eta}(\mathbf{d}_m, \mathbf{q}) &\leq M_\Phi(\mathbf{d}_m, \mathbf{q}), \end{aligned}$$

where the implication is obtained by adding $[\mathbf{d}_m]_k$ on both sides of the first inequality and using Proposition 13.

Suppose now that $\mathbf{q} \in \text{ri } \Delta_{\mathcal{I}}$, with $|\mathcal{I}| > 1$, and let $\Phi^{\mathcal{I}} := \Phi \circ [\Pi_{\mathcal{I}}^k]^\top$ and $S^{\mathcal{I}} := S \circ [\Pi_{\mathcal{I}}^k]^\top$. Note that since $\eta_\ell \Phi - S$ is convex on Δ_k and $\Pi_{\mathcal{I}}^k$ is a linear function, $\eta_\ell \Phi^{\mathcal{I}} - S^{\mathcal{I}}$ is convex on $\Delta_{|\mathcal{I}|}$. Repeating the steps above for Φ, S, \mathbf{q} , and A substituted by $\Phi^{\mathcal{I}}, S^{\mathcal{I}}, \Pi_{\mathcal{I}}^k \mathbf{q}$, and $A[\Pi_{\mathcal{I}}^k]^\top$, respectively, yields

$$\begin{aligned} M_{S_\eta^{\mathcal{I}}}(\Pi_{\mathcal{I}}^k \mathbf{d}_m, \Pi_{\mathcal{I}}^k \mathbf{q}) &\leq M_{\Phi^{\mathcal{I}}}(\Pi_{\mathcal{I}}^k \mathbf{d}_m, \Pi_{\mathcal{I}}^k \mathbf{q}), \\ \implies M_{S_\eta}(\mathbf{d}_m, \mathbf{q}) &\leq M_\Phi(\mathbf{d}_m, \mathbf{q}), \\ \implies M_{S_\eta}(\ell_x(A), \mathbf{q}) &\leq M_\Phi(\ell_x(A), \mathbf{q}), \end{aligned} \quad (58)$$

where the first implication follows from Lemma 30, since S_η and Φ both satisfy (8) (see Lemmas 29 and 34), and (58) is obtained by passage to the limit $m \rightarrow \infty$. Since $\eta = \eta_\ell > 0$, ℓ is η -mixable, which implies that ℓ is S_η -mixable (Theorem 14). Therefore, there exists $\mathbf{a}_* \in \mathcal{A}$, such that

$$\ell_x(\mathbf{a}_*) \leq M_{S_\eta}(\ell_x(A), \mathbf{q}) \leq M_\Phi(\ell_x(A), \mathbf{q}). \quad (59)$$

To complete the proof (that is, to show that ℓ is Φ -mixable), it remains to consider the case where \mathbf{q} is a vertex of Δ_k . Without loss of generality assume that $\mathbf{q} = \mathbf{e}_1$ and let $\boldsymbol{\mu} \in \Delta_k \setminus \{\mathbf{e}_1\}$. Thus, there exists $\mathcal{I}_* \subseteq [k]$, with $|\mathcal{I}_*| > 1$, such that $(\mathbf{e}_1, \boldsymbol{\mu}) \in (\text{rbd } \Delta_{\mathcal{I}_*}) \times (\text{ri } \Delta_{\mathcal{I}_*})$, and Lemma 34 implies that $\Phi'(\mathbf{e}_1; \boldsymbol{\mu} - \mathbf{e}_1) = -\infty$. Therefore, $\forall \mathbf{q} \in \Delta_k \setminus \{\mathbf{e}_1\}$, $D_\Phi(\mathbf{q}, \mathbf{e}_1) = +\infty$, which implies

$$\begin{aligned} \forall x \in [n], M_\Phi(\ell_x(A), \mathbf{e}_1) &= \inf_{\mathbf{q} \in \Delta_k} \langle \mathbf{q}, \ell_x(A) \rangle + D_\Phi(\mathbf{q}, \mathbf{e}_1), \\ &= \langle \mathbf{e}_1, \ell_x(A) \rangle + D_\Phi(\mathbf{e}_1, \mathbf{e}_1) = \langle \mathbf{e}_1, \ell_x(A) \rangle, \\ &= \ell_x(\mathbf{a}_1). \end{aligned} \quad (60)$$

The Φ -mixability condition (6) is trivially satisfied in this case. Combining (59) and (60) shows that ℓ is Φ -mixable. \blacksquare

C.11. Proof of Theorem 16

The following Lemma gives necessary regularity conditions on the entropy Φ under the assumptions of Theorem 16.

Lemma 35 *Let Φ and ℓ be as in Theorem 16. Then the following holds*

- (i) $\tilde{\Phi}$ is strictly concave on $\text{int } \tilde{\Delta}_k$.
- (ii) $\tilde{\Phi}^*$ is continuously differentiable on \mathbb{R}^{k-1} .
- (iii) $\tilde{\Phi}^*$ is twice differentiable on \mathbb{R}^{k-1} and $\forall \tilde{\mathbf{q}} \in \text{int } \tilde{\Delta}_k$, $\text{H}\tilde{\Phi}^*(\nabla\tilde{\Phi}(\tilde{\mathbf{q}})) = (\text{H}\tilde{\Phi}(\tilde{\mathbf{q}}))^{-1}$.
- (iv) For the Shannon entropy, we have $(\text{H}\tilde{S}(\tilde{\mathbf{q}}))^{-1} = \text{H}\tilde{S}^*(\nabla\tilde{S}(\tilde{\mathbf{q}})) = \text{diag } \tilde{\mathbf{q}} - \tilde{\mathbf{q}}\tilde{\mathbf{q}}^\top$.

Proof Since ℓ is Φ -mixable and \underline{L}_ℓ is twice differentiable on $]0, +\infty[^n$, $\tilde{\Phi}^*$ is continuously differentiable on \mathbb{R}^{n-1} (Proposition 12). Therefore, $\tilde{\Phi}$ is strictly convex on $\text{ri } \Delta_k$ (Hiriart-Urruty and Lemaréchal, 2001, Thm. E.4.1.2).

The differentiability of $\tilde{\Phi}$ and $\tilde{\Phi}^*$ implies $\nabla\tilde{\Phi}^*(\nabla\tilde{\Phi}(\tilde{\mathbf{q}})) = \tilde{\mathbf{q}}$ (ibid.). Since $\tilde{\Phi}$ is twice differentiable on $\text{int } \tilde{\Delta}_k$ (by assumption), the latter equation implies that $\tilde{\Phi}^*$ is twice differentiable on $\nabla\tilde{\Phi}(\text{int } \tilde{\Delta}_k)$. Using the chain rule, we get $\text{H}\tilde{\Phi}^*(\nabla\tilde{\Phi}(\mathbf{u}))\text{H}\tilde{\Phi}(\mathbf{u}) = I_{\tilde{\mathbf{q}}}$. Multiplying both sides of the equation by $(\text{H}\tilde{\Phi}(\mathbf{u}))^{-1}$ from the right gives the expression in (iii). Note that $\text{H}\tilde{\Phi}(\cdot)$ is in fact invertible on $\text{int } \tilde{\Delta}_k$ since $\tilde{\Phi}$ is strictly convex on $\text{int } \tilde{\Delta}_k$. It remains to show that $\nabla\tilde{\Phi}(\text{int } \tilde{\Delta}_k) = \mathbb{R}^{k-1}$. This set equality follows from 1) $[\tilde{\mathbf{q}} \in \partial\tilde{\Phi}^*(s) \iff s \in \partial\tilde{\Phi}(\tilde{\mathbf{q}})]$ (ibid., Cor. E.1.4.4); 2) $\text{dom } \tilde{\Phi}^* = \mathbb{R}^{k-1}$; and 3) $\forall \tilde{\mathbf{q}} \in \text{bd } \tilde{\Delta}_k$, $\partial\tilde{\Phi}(\tilde{\mathbf{q}}) = \emptyset$ (Lemma 30).

For the Shannon entropy, we have $\tilde{S}^*(\mathbf{v}) = \log(\langle \exp(\mathbf{v}), \mathbf{1}_{\tilde{k}} \rangle + 1)$ (Proposition 1) and $\nabla\tilde{S}(\tilde{\mathbf{q}}) = \log \frac{\tilde{\mathbf{q}}}{q_k}$, for $(\mathbf{v}, \tilde{\mathbf{q}}) \in \mathbb{R}^{k-1} \times \tilde{\Delta}_k$. Thus $(\text{H}\tilde{S}(\tilde{\mathbf{q}}))^{-1} = \text{H}\tilde{S}^*(\nabla\tilde{S}(\tilde{\mathbf{q}})) = \text{diag } \tilde{\mathbf{q}} - \tilde{\mathbf{q}}\tilde{\mathbf{q}}^\top$. \blacksquare

To show Theorem 16, we analyze a particular parameterized curve defined in the next lemma.

Lemma 36 *Let $\ell: \Delta_n \rightarrow [0, +\infty]^n$ be a proper loss whose Bayes risk \underline{L}_ℓ is twice differentiable on $]0, +\infty[^n$, and let Φ be an entropy such that $\tilde{\Phi}$ and $\tilde{\Phi}^*$ are twice differentiable on $\text{int } \tilde{\Delta}_k$ and \mathbb{R}^{k-1} , respectively. For $(\tilde{\mathbf{p}}, \tilde{\mathbf{q}}, V) \in \text{int } \tilde{\Delta}_n \times \text{int } \tilde{\Delta}_k \times \mathbb{R}^{\tilde{n} \times \tilde{k}}$, let $\beta: \mathbb{R} \rightarrow \mathbb{R}^n$ be the curve defined by*

$$\forall x \in [n], \quad \beta_x(t) = \tilde{\ell}_x(\tilde{\mathbf{p}}) + \tilde{\Phi}^*(\nabla \tilde{\Phi}(\tilde{\mathbf{q}})) - \tilde{\Phi}^*(\nabla \tilde{\Phi}(\tilde{\mathbf{q}}) - J_k^\top \tilde{\ell}_x(\tilde{P}^t)), \quad (61)$$

where $\tilde{P}^t = [\tilde{\mathbf{p}} \mathbf{1}_k^\top + tV, \tilde{\mathbf{p}}] \in \mathbb{R}^{\tilde{n} \times k}$ and $t \in \{s \in \mathbb{R} : \forall j \in [\tilde{k}], \tilde{\mathbf{p}} + sV_{\cdot,j} \in \text{int } \tilde{\Delta}_n\}$. Then

$$\begin{aligned} \beta(0) &= \tilde{\ell}(\tilde{\mathbf{p}}), \\ \dot{\beta}(0) &= D\tilde{\ell}(\tilde{\mathbf{p}})V\tilde{\mathbf{q}}, \end{aligned}$$

$$\left. \frac{d}{dt} \langle \mathbf{p}, \dot{\beta}(t) \rangle \right|_{t=0} = - \sum_{j=1}^{k-1} q_j V_{\cdot,j}^\top \mathbf{H} \tilde{L}_\ell(\tilde{\mathbf{p}}) V_{\cdot,j} - \text{tr}(\text{diag}(\mathbf{p}) D\tilde{\ell}(\tilde{\mathbf{p}}) V (\mathbf{H}\tilde{\Phi}(\tilde{\mathbf{q}}))^{-1} (D\tilde{\ell}(\tilde{\mathbf{p}}) V)^\top). \quad (62)$$

Proof Since $\tilde{P}^t = [\tilde{\mathbf{p}} \mathbf{1}_k^\top + tV, \tilde{\mathbf{p}}] \in \mathbb{R}^{\tilde{n} \times k}$, $\tilde{P}^0 = \tilde{\mathbf{p}} \mathbf{1}_k^\top$ and $\tilde{\ell}_x(\tilde{P}^0) = \tilde{\ell}_x(\tilde{\mathbf{p}}) \mathbf{1}_k$. As a result, $J_k^\top \tilde{\ell}_x(\tilde{P}^0) = \mathbf{0}_{\tilde{k}}$, and thus $\beta_x(0) = \tilde{\ell}_x(\tilde{\mathbf{p}}) + \tilde{\Phi}^*(\nabla \tilde{\Phi}(\tilde{\mathbf{q}})) - \tilde{\Phi}^*(\nabla \tilde{\Phi}(\tilde{\mathbf{q}}) - \mathbf{0}_{\tilde{k}}) = \tilde{\ell}_x(\tilde{\mathbf{p}})$. This shows that $\beta(0) = \tilde{\ell}(\tilde{\mathbf{p}})$. Let $\gamma_x(t) := \nabla \tilde{\Phi}(\tilde{\mathbf{q}}) - J_k^\top \tilde{\ell}_x(\tilde{P}^t)$. For $j \in [k-1]$,

$$\begin{aligned} \frac{d}{dt} [\gamma_x(t)]_j &= \frac{d}{dt} \left([\nabla \tilde{\Phi}(\tilde{\mathbf{q}})]_j - [J_k^\top \tilde{\ell}_x(\tilde{P}^t)]_j \right), \\ &= - \frac{d}{dt} \left(\tilde{\ell}_x(\tilde{P}^t_{\cdot,j}) - \tilde{\ell}_x(\tilde{P}^t_{\cdot,k}) \right), \\ &= - \frac{d}{dt} \left(\tilde{\ell}_x(\tilde{\mathbf{p}} + tV_{\cdot,j}) - \tilde{\ell}_x(\tilde{\mathbf{p}}) \right), \quad \left(\text{since } \frac{d}{dt} \tilde{\ell}_x(\tilde{P}^t_{\cdot,k}) = \frac{d}{dt} \tilde{\ell}_x(\tilde{\mathbf{p}}) = 0 \right) \\ &= -D\tilde{\ell}_x(\tilde{P}^t_{\cdot,j})V_{\cdot,j}. \end{aligned}$$

From the definition of \tilde{P}^t , $\tilde{P}^0_{\cdot,j} = \tilde{\mathbf{p}}$, $\forall j \in [\tilde{k}]$, and therefore, $\dot{\gamma}_x(0) = -(D\tilde{\ell}_x(\tilde{\mathbf{p}})V)^\top$. By differentiating β_x in (61) and using the chain rule, $\dot{\beta}_x(t) = -(\dot{\gamma}_x(t))^\top \nabla \tilde{\Phi}^*(\gamma_x(t))$. By setting

$t = 0$, $\dot{\beta}_x(0) = -(\dot{\gamma}_x(0))^\top \nabla \tilde{\Phi}^*(\nabla \tilde{\Phi}(\tilde{\mathbf{q}})) = D\tilde{\ell}_x(\tilde{\mathbf{p}})V\tilde{\mathbf{q}}$. Thus, $\dot{\beta}(0) = D\tilde{\ell}(\tilde{\mathbf{p}})V\tilde{\mathbf{q}}$. Furthermore,

$$\begin{aligned}
 \left. \frac{d}{dt} \langle \mathbf{p}, \dot{\beta}(t) \rangle \right|_{t=0} &= \left. \frac{d}{dt} \sum_{x=1}^n p_x \left(\sum_{j=1}^{k-1} D\tilde{\ell}_x(\tilde{P}_{\cdot,j}^t) V_{\cdot,j} [\nabla \tilde{\Phi}^*(\gamma_x(t))]_j \right) \right|_{t=0}, \\
 &= \sum_{j=1}^{k-1} \left. \frac{d}{dt} \left(\sum_{x=1}^n p_x D\tilde{\ell}_x(\tilde{P}_{\cdot,j}^t) V_{\cdot,j} [\nabla \tilde{\Phi}^*(\gamma_x(t))]_j \right) \right|_{t=0}, \\
 &= \sum_{j=1}^{k-1} \left(\left. \frac{d}{dt} \langle \mathbf{p}, D\tilde{\ell}(\tilde{P}_{\cdot,j}^t) V_{\cdot,j} q_j \rangle \right|_{t=0} + \sum_{x=1}^n p_x D\tilde{\ell}_x(\tilde{\mathbf{p}}) V_{\cdot,j} \left. \frac{d}{dt} [\nabla \tilde{\Phi}^*(\gamma_x(t))]_j \right|_{t=0} \right), \\
 &= - \sum_{j=1}^{k-1} q_j V_{\cdot,j}^\top H \tilde{L}_\ell(\tilde{\mathbf{p}}) V_{\cdot,j} - \sum_{x=1}^n \sum_{i=1}^{k-1} p_x D\tilde{\ell}_x(\tilde{\mathbf{p}}) V_{\cdot,j} [H \tilde{\Phi}^*(\nabla \tilde{\Phi}(\tilde{\mathbf{q}}))]_{j,i} D\tilde{\ell}_x(\tilde{\mathbf{p}}) V_{\cdot,i}, \\
 &= - \sum_{j=1}^{k-1} q_j V_{\cdot,j}^\top H \tilde{L}_\ell(\tilde{\mathbf{p}}) V_{\cdot,j} - \text{tr}(\text{diag}(\mathbf{p}) D\tilde{\ell}(\tilde{\mathbf{p}}) V H \tilde{\Phi}^*(\nabla \tilde{\Phi}(\tilde{\mathbf{q}})) (D\tilde{\ell}(\tilde{\mathbf{p}}) V)^\top), \\
 &= - \sum_{j=1}^{k-1} q_j V_{\cdot,j}^\top H \tilde{L}_\ell(\tilde{\mathbf{p}}) V_{\cdot,j} - \text{tr}(\text{diag}(\mathbf{p}) D\tilde{\ell}(\tilde{\mathbf{p}}) V (H \tilde{\Phi}(\tilde{\mathbf{q}}))^{-1} (D\tilde{\ell}(\tilde{\mathbf{p}}) V)^\top),
 \end{aligned}$$

where in the third equality we used Lemma 25, in the fourth equality we used Lemma 28, and in the sixth equality we used Lemma 35-(iii). ■

In next lemma, we state a necessary condition for Φ -mixability in terms of the parameterized curve β defined in Lemma 36.

Lemma 37 *Let ℓ , Φ , and β be as in Lemma 36. If $\exists(\tilde{\mathbf{p}}, \tilde{\mathbf{q}}, V) \in \text{int } \tilde{\Delta}_n \times \text{int } \tilde{\Delta}_k \times \mathbb{R}^{\tilde{n} \times \tilde{k}}$ such that the curve $\gamma(t) := \tilde{\ell}(\tilde{\mathbf{p}} + tV\tilde{\mathbf{q}})$ satisfies $\left. \frac{d}{dt} \langle \mathbf{p}, \dot{\beta}(t) - \dot{\gamma}(t) \rangle \right|_{t=0} < 0$, then ℓ is not Φ -mixable. In particular, $\exists P \in \text{ri } \Delta_n^k$, such that $[M_\Phi(\ell_x(P), \mathbf{q})]_{x \in [n]}^\top$ lies outside \mathcal{S}_ℓ^\oplus .*

Proof First note that for any triplet $(\tilde{\mathbf{p}}, \tilde{\mathbf{q}}, V) \in \text{int } \tilde{\Delta}_n \times \text{int } \tilde{\Delta}_k \times \mathbb{R}^{\tilde{n} \times \tilde{k}}$, the map $t \mapsto \langle \mathbf{p}, \dot{\beta}(t) - \dot{\gamma}(t) \rangle$ is differentiable at 0. This follows from Lemmas 25 and 36. Let $r(t) := \Pi_n(\tilde{\mathbf{p}} + tV\tilde{\mathbf{q}})$ and $\delta(t) := \langle r(t), \beta(t) - \gamma(t) \rangle$. Then

$$\dot{\delta}(t) = \langle r(t), \dot{\beta}(t) - \dot{\gamma}(t) \rangle + \langle V\tilde{\mathbf{q}}, \beta(t) - \gamma(t) \rangle.$$

Since $t \mapsto \langle \mathbf{p}, \dot{\beta}(t) - \dot{\gamma}(t) \rangle$ is differentiable at 0, it follows from Lemma 25 that $t \mapsto \dot{\delta}(t)$ is also differentiable at 0, and thus

$$\begin{aligned} \ddot{\delta}(0) &= \left. \frac{d}{dt} \langle r(t), \dot{\beta}(t) - \dot{\gamma}(t) \rangle \right|_{t=0} + \langle J_n V \tilde{\mathbf{q}}, \dot{\beta}(0) - \dot{\gamma}(0) \rangle, \\ &= \left\langle \left. \frac{d}{dt} r(t) \right|_{t=0}, \dot{\beta}(0) - \dot{\gamma}(0) \right\rangle + \left. \frac{d}{dt} \langle \mathbf{p}, \dot{\beta}(t) - \dot{\gamma}(t) \rangle \right|_{t=0}, \end{aligned} \quad (63)$$

$$\begin{aligned} &= \langle J_n V \tilde{\mathbf{q}}, \dot{\beta}(0) - \dot{\gamma}(0) \rangle + \left. \frac{d}{dt} \langle \mathbf{p}, \dot{\beta}(t) - \dot{\gamma}(t) \rangle \right|_{t=0}, \\ &= \left. \frac{d}{dt} \langle \mathbf{p}, \dot{\beta}(t) - \dot{\gamma}(t) \rangle \right|_{t=0} < 0, \end{aligned} \quad (64)$$

where (63) and (64) hold because $\dot{\beta}(0) = D\tilde{\ell}(\tilde{\mathbf{p}})V\tilde{\mathbf{q}} = \dot{\gamma}(0)$ (see Lemma 36). According to Taylor's theorem (see e.g. (Hardy, 2008, §151)), there exists $\epsilon > 0$ and $h : [-\epsilon, \epsilon] \rightarrow \mathbb{R}$ such that

$$\forall |t| \leq \epsilon, \quad \delta(t) = \delta(0) + t\dot{\delta}(0) + \frac{t^2}{2}\ddot{\delta}(0) + h(t)t^2, \quad (65)$$

and $\lim_{t \rightarrow 0} h(t) = 0$. From Lemma 36, $\beta(0) = \gamma(0) = 0$ and $\dot{\beta}(0) = \dot{\gamma}(0)$. Therefore, $\delta(0) = \dot{\delta}(0) = 0$ and (65) becomes $\delta(t) = \frac{t^2}{2}\ddot{\delta}(0) + h(t)t^2$. Due to (64) and the fact that $\lim_{t \rightarrow 0} h(t) = 0$, we can choose $\epsilon_* > 0$ small enough such that $\delta(\epsilon_*) = \frac{\epsilon_*^2}{2}\ddot{\delta}(0) + h(\epsilon_*)\epsilon_*^2 < 0$. This means that $\langle \Pi_n(\tilde{\mathbf{p}} + \epsilon_* V \tilde{\mathbf{q}}), \beta(\epsilon_*) \rangle < \langle \Pi_n(\tilde{\mathbf{p}} + \epsilon_* V \tilde{\mathbf{q}}), \tilde{\ell}(\tilde{\mathbf{p}} + \epsilon_* V \tilde{\mathbf{q}}) \rangle = \langle \Pi_n(\tilde{\mathbf{p}} + \epsilon_* V \tilde{\mathbf{q}}), \ell(\Pi_n(\tilde{\mathbf{p}} + \epsilon_* V \tilde{\mathbf{q}})) \rangle$. Therefore, $\beta(\epsilon_*)$ must lie outside the superprediction set. Thus, the mixability condition (6) does not hold for $P^{\epsilon_*} = \Pi_n[\tilde{\mathbf{p}}\mathbf{1}_k^\top + \epsilon_* V, \tilde{\mathbf{p}}] \in \text{ri } \Delta_n^k$. This completes the proof. \blacksquare

Proof [Theorem 16] We will prove the contrapositive; suppose that $\eta_\ell \Phi - S$ is not convex on Δ_k and we show that ℓ cannot be Φ -mixable. Note first that from Lemma 35-(iii), $\tilde{\Phi}^*$ is twice differentiable on \mathbb{R}^{k-1} . Thus Lemmas 36 and 37 apply. Let $\underline{\ell}$ be a proper support loss of ℓ and suppose that $\eta_\ell \Phi - S$ is not convex on Δ_k . This implies that $\eta_\ell \tilde{\Phi} - \tilde{S}$ is not convex on $\text{int } \tilde{\Delta}_k$, and by Lemma 23 there exists $\tilde{\mathbf{q}}_* \in \text{int } \tilde{\Delta}_k$, such that $1 > \underline{\eta}_\ell \lambda_{\min}(\mathbf{H}\tilde{\Phi}(\tilde{\mathbf{q}}_*)(\mathbf{H}\tilde{S}(\tilde{\mathbf{q}}_*))^{-1})$. From this and the definition of $\underline{\eta}_\ell$, there exists $\tilde{\mathbf{p}}_* \in \text{int } \tilde{\Delta}_n$ such that

$$1 > \frac{\lambda_{\min}(\mathbf{H}\tilde{\Phi}(\tilde{\mathbf{q}}_*)(\mathbf{H}\tilde{S}(\tilde{\mathbf{q}}_*))^{-1})}{\lambda_{\max}([\mathbf{H}\tilde{L}_{\log}(\tilde{\mathbf{p}}_*)]^{-1}\mathbf{H}\tilde{L}_\ell(\tilde{\mathbf{p}}_*))} = \frac{\lambda_{\min}(\mathbf{H}\tilde{\Phi}(\tilde{\mathbf{q}}_*)(\text{diag}(\tilde{\mathbf{q}}_*) - \tilde{\mathbf{q}}_*\tilde{\mathbf{q}}_*^\top))}{\lambda_{\max}([\mathbf{H}\tilde{L}_{\log}(\tilde{\mathbf{p}}_*)]^{-1}\mathbf{H}\tilde{L}_\ell(\tilde{\mathbf{p}}_*))}, \quad (66)$$

where the equality is due to Lemma 35-(iv). For the rest of this proof let $(\tilde{\mathbf{p}}, \tilde{\mathbf{q}}) = (\tilde{\mathbf{p}}^*, \tilde{\mathbf{q}}^*)$. By assumption, \tilde{L}_ℓ twice differentiable and concave on $\text{int } \tilde{\Delta}_n$, and thus $-\mathbf{H}\tilde{L}_\ell(\tilde{\mathbf{p}})$ is symmetric positive semi-definite. Therefore, there exists a symmetric positive semi-definite matrix $\Lambda_{\mathbf{p}}$ such that $\Lambda_{\mathbf{p}}\Lambda_{\mathbf{p}} = -\mathbf{H}\tilde{L}_\ell(\tilde{\mathbf{p}})$. From Lemma 35-(i), $\tilde{\Phi}$ is strictly convex on $\text{int } \tilde{\Delta}_k$, and so there exists a symmetric positive definite matrix $K_{\mathbf{q}}$ such that $K_{\mathbf{q}}K_{\mathbf{q}} = \mathbf{H}\tilde{\Phi}(\tilde{\mathbf{q}})$. Let $\mathbf{w} \in \mathbb{R}^{n-1}$ be the unit norm eigenvector of $[\mathbf{H}\tilde{L}_{\log}(\tilde{\mathbf{p}})]^{-1}\mathbf{H}\tilde{L}_\ell(\tilde{\mathbf{p}})$ associated with $\lambda_*^\ell := \lambda_{\max}([\mathbf{H}\tilde{L}_{\log}(\tilde{\mathbf{p}})]^{-1}\mathbf{H}\tilde{L}_\ell(\tilde{\mathbf{p}}))$. Suppose that $c_\ell := \mathbf{w}^\top \mathbf{H}\tilde{L}_\ell(\tilde{\mathbf{p}})\mathbf{w} = 0$. Since $\mathbf{w}^\top \Lambda_{\mathbf{p}}\Lambda_{\mathbf{p}}\mathbf{w} = -c_\ell = 0$, it follows from the positive semi-definiteness of $\Lambda_{\mathbf{p}}$ that $\Lambda_{\mathbf{p}}\mathbf{w} = \mathbf{0}_{\tilde{n}}$, and thus $\mathbf{H}\tilde{L}_\ell(\tilde{\mathbf{p}})\mathbf{w} = -\Lambda_{\mathbf{p}}\Lambda_{\mathbf{p}}\mathbf{w} = \mathbf{0}_{\tilde{n}}$. This implies that $\lambda_*^\ell = 0$, which is not possible due to (66). Therefore, $\mathbf{H}\tilde{L}_\ell(\tilde{\mathbf{p}})\mathbf{w} \neq \mathbf{0}_{\tilde{n}}$. Furthermore, the negative semi-definiteness of $\mathbf{H}\tilde{L}_\ell(\tilde{\mathbf{p}})$ implies that

$$c_\ell = \mathbf{w}^\top \mathbf{H}\tilde{L}_\ell(\tilde{\mathbf{p}})\mathbf{w} < 0. \quad (67)$$

Let $\mathbf{v} \in \mathbb{R}^{k-1}$ be the unit norm eigenvector of $K_q(\text{diag}(\tilde{\mathbf{q}}) - \tilde{\mathbf{q}}\tilde{\mathbf{q}}^\top)K_q$ associated with $\lambda_*^\Phi := \lambda_{\min}(K_q(\text{diag}(\tilde{\mathbf{q}}) - \tilde{\mathbf{q}}\tilde{\mathbf{q}}^\top)K_q) = \lambda_{\min}(\mathbf{H}\tilde{\Phi}(\tilde{\mathbf{q}})(\text{diag}(\tilde{\mathbf{q}}) - \tilde{\mathbf{q}}\tilde{\mathbf{q}}^\top))$, where the equality is due to Lemma 22. Let $\hat{\mathbf{v}} := K_q\mathbf{v}$.

We will show that for $V = \mathbf{w}\hat{\mathbf{v}}^\top$, the parametrized curve β defined in Lemma 36 satisfies $\left. \frac{d}{dt} \langle \mathbf{p}, \dot{\beta}(t) - \dot{\gamma}(t) \rangle \right|_{t=0} < 0$, where $\gamma(t) = \tilde{\ell}(\tilde{\mathbf{p}} + tV\tilde{\mathbf{q}})$. According to Lemma 37 this would imply that there exists $P \in \text{ri} \Delta_n^k$, such that $[\mathbf{M}_\Phi(\underline{\ell}_x(P), \mathbf{q})]_{x \in [n]}^\top$ lies outside \mathcal{S}_ℓ^\oplus . From Theorem 4, we know that there exists $A_* \in \mathcal{A}^k$, such that $\ell_x(A_*) = \underline{\ell}_x(P), \forall x \in [n]$. Therefore, $[\mathbf{M}_\Phi(\ell_x(A_*), \mathbf{q})]_{x \in [n]}^\top = [\mathbf{M}_\Phi(\underline{\ell}_x(P), \mathbf{q})]_{x \in [n]}^\top \notin \mathcal{S}_\ell^\oplus$, and thus ℓ is not Φ -mixable.

From Lemma 36 (Equation 62) and the fact that $V_{:,j} = \hat{v}_j \mathbf{w}$, for $j \in [k]$, we can write

$$\begin{aligned} \left. \frac{d}{dt} \langle \mathbf{p}, \dot{\beta}(t) \rangle \right|_{t=0} &= - \sum_{j=1}^{k-1} q_j \hat{v}_j^2 \mathbf{w}^\top \mathbf{H} \tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}}) \mathbf{w} - \text{tr}(\text{diag}(\mathbf{p}) \mathbf{D} \tilde{\ell}(\tilde{\mathbf{p}}) V (\mathbf{H} \tilde{\Phi}(\tilde{\mathbf{q}}))^{-1} (\mathbf{D} \tilde{\ell}(\tilde{\mathbf{p}}) V)^\top), \\ &= - \langle \tilde{\mathbf{q}}, \hat{\mathbf{v}} \odot \hat{\mathbf{v}} \rangle \mathbf{w}^\top \mathbf{H} \tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}}) \mathbf{w} - (\hat{\mathbf{v}}^\top (\mathbf{H} \tilde{\Phi}(\mathbf{q}))^{-1} \hat{\mathbf{v}}) \langle \mathbf{p}, [\mathbf{D} \tilde{\ell}(\tilde{\mathbf{p}}) \mathbf{w}] \odot [(\mathbf{D} \tilde{\ell}(\tilde{\mathbf{p}}) \mathbf{w})] \rangle, \end{aligned}$$

where the second equality is obtained by noting that 1) $(\hat{\mathbf{v}}^\top (\mathbf{H} \tilde{\Phi}(\mathbf{q}))^{-1} \hat{\mathbf{v}})$ is a scalar quantity and can be factorized out; and 2) $\text{tr}(\text{diag}(\mathbf{p}) \mathbf{D} \tilde{\ell}(\tilde{\mathbf{p}}) \mathbf{w} (\mathbf{D} \tilde{\ell}(\tilde{\mathbf{p}}) \mathbf{w})^\top) = \langle \mathbf{p}, (\mathbf{D} \tilde{\ell}(\tilde{\mathbf{p}}) \mathbf{w}) \odot (\mathbf{D} \tilde{\ell}(\tilde{\mathbf{p}}) \mathbf{w}) \rangle$.

On the other hand, from Lemma 28, $\left. \frac{d}{dt} \langle \mathbf{p}, \dot{\gamma}(t) \rangle \right|_{t=0} = - \langle \tilde{\mathbf{q}}, \hat{\mathbf{v}} \rangle^2 \mathbf{w}^\top \mathbf{H} \tilde{\mathcal{L}}_\ell(\tilde{\mathbf{q}}) \mathbf{w}$. Using (19) and the definition of c_ℓ , we get

$$\begin{aligned} \left. \frac{d}{dt} \langle \mathbf{p}, \dot{\beta}(t) - \dot{\gamma}(t) \rangle \right|_{t=0} &= [- \langle \tilde{\mathbf{q}}, \hat{\mathbf{v}} \odot \hat{\mathbf{v}} \rangle + \langle \tilde{\mathbf{q}}, \hat{\mathbf{v}} \rangle^2] c_\ell + \\ &\quad (\hat{\mathbf{v}}^\top (\mathbf{H} \tilde{\Phi}(\mathbf{q}))^{-1} \hat{\mathbf{v}}) (\mathbf{w}^\top (\mathbf{H} \tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}})) (\mathbf{H} \tilde{\mathcal{L}}_{\log}(\tilde{\mathbf{p}}))^{-1} \mathbf{H} \tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}}) \mathbf{w}), \\ &= -c_\ell [\langle \tilde{\mathbf{q}}, \hat{\mathbf{v}} \odot \hat{\mathbf{v}} \rangle - \langle \tilde{\mathbf{q}}, \hat{\mathbf{v}} \rangle^2 - \lambda_*^\ell (\hat{\mathbf{v}}^\top (\mathbf{H} \tilde{\Phi}(\mathbf{q}))^{-1} \hat{\mathbf{v}})], \\ &= -c_\ell [\hat{\mathbf{v}}^\top (\text{diag}(\tilde{\mathbf{q}}) - \tilde{\mathbf{q}}\tilde{\mathbf{q}}^\top) \hat{\mathbf{v}} - \lambda_*^\ell (\hat{\mathbf{v}}^\top (\mathbf{H} \tilde{\Phi}(\mathbf{q}))^{-1} \hat{\mathbf{v}})], \\ &= -c_\ell [\hat{\mathbf{v}}^\top (\text{diag}(\tilde{\mathbf{q}}) - \tilde{\mathbf{q}}\tilde{\mathbf{q}}^\top) \hat{\mathbf{v}} - \lambda_*^\ell (\mathbf{v}^\top K_q (K_q K_q)^{-1} K_q \mathbf{v})], \\ &= -c_\ell [\mathbf{v}^\top K_q (\text{diag}(\tilde{\mathbf{q}}) - \tilde{\mathbf{q}}\tilde{\mathbf{q}}^\top) K_q \mathbf{v} - \lambda_*^\ell], \tag{68} \\ &= -c_\ell [\lambda_*^\Phi - \lambda_*^\ell], \\ &= -c_\ell [\lambda_{\min}(\mathbf{H} \tilde{\Phi}(\mathbf{q}) (\text{diag}(\tilde{\mathbf{q}}) - \tilde{\mathbf{q}}\tilde{\mathbf{q}}^\top)) - \lambda_{\max}(\mathbf{H} \tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}}) (\mathbf{H} \tilde{\mathcal{L}}_{\log}(\tilde{\mathbf{p}}))^{-1})], \end{aligned}$$

where in (68) we used the fact that $\mathbf{v}^\top \mathbf{v} = 1$. The last equality combined with (66) and (67) shows that $\left. \frac{d}{dt} \langle \mathbf{p}, \dot{\beta}(t) - \dot{\gamma}(t) \rangle \right|_{t=0} < 0$, which completes the proof. \blacksquare

C.12. Proof of Corollary 18

Proof From Corollary 17, ℓ is Φ_η -mixable if and only if $\eta_\ell \Phi_\eta - \mathbf{S} = \eta^{-1} \eta_\ell \Phi - \mathbf{S}$ is convex on Δ_k . When this is the case, Lemma 23 implies that

$$1 \leq \eta^{-1} \eta_\ell \left(\inf_{\tilde{\mathbf{q}} \in \text{int} \tilde{\Delta}_k} \lambda_{\min}[\mathbf{H} \tilde{\Phi}(\tilde{\mathbf{q}}) [\mathbf{H} \tilde{\mathbf{S}}(\tilde{\mathbf{q}})]^{-1}] \right), \tag{69}$$

where we used the facts that $\mathbf{H}(\eta^{-1} \eta_\ell \tilde{\Phi}) = \eta^{-1} \eta_\ell \mathbf{H} \tilde{\Phi}$, $\lambda_{\min}(\cdot)$ is linear, and $\eta^{-1} \eta_\ell$ is independent of $\tilde{\mathbf{q}} \in \text{int} \tilde{\Delta}_k$. Inequality 69 shows that the largest η such that ℓ is Φ_η -mixable is given by η_ℓ^Φ in (16). \blacksquare

C.13. Proof of Theorem 19

Proof Suppose ℓ is Φ -mixable. Then from Corollary 17, $\underline{\eta}_\ell \Phi - S$ is convex on Δ_k , and thus $\underline{\eta}_\ell = \eta_\ell^S > 0$ (Corollary 18). Furthermore, $\underline{\eta}_\ell \tilde{\Phi} - \tilde{S} = [\underline{\eta}_\ell \Phi - S] \circ \Pi_k$ is convex on $\text{int } \tilde{\Delta}_k$, since Π_k is an affine function. It follows from Lemma 23 and Corollary 18 that

$$\eta_\ell^\Phi = \underline{\eta}_\ell \inf_{\tilde{\mathbf{q}} \in \text{int } \tilde{\Delta}_k} \lambda_{\min}(\mathbf{H}\tilde{\Phi}(\tilde{\mathbf{q}})(\mathbf{H}\tilde{S}(\tilde{\mathbf{q}}))^{-1}) \geq 1 > 0.$$

Let $\boldsymbol{\mu} \in \text{ri } \Delta_k$ and $\theta_* := \text{argmax}_\theta D_S(\mathbf{e}_\theta, \boldsymbol{\mu})$. By definition of an entropy and the fact that the directional derivatives $\Phi'(\boldsymbol{\mu}; \cdot)$ and $S'(\boldsymbol{\mu}; \cdot)$ are finite on Δ_k (Hiriart-Urruty and Lemaréchal, 2001, Prop. D.1.1.2), it holds that $D_\Phi(\mathbf{e}_{\theta_*}, \boldsymbol{\mu}), D_S(\mathbf{e}_{\theta_*}, \boldsymbol{\mu}) \in]0, +\infty[$. Therefore, there exists $\alpha > 0$ such that $\alpha^{-1} D_\Phi(\mathbf{e}_{\theta_*}, \boldsymbol{\mu}) = D_S(\mathbf{e}_{\theta_*}, \boldsymbol{\mu})$. If we let $\Psi := \alpha^{-1} \Phi$, we get

$$D_\Psi(\mathbf{e}_{\theta_*}, \boldsymbol{\mu}) = D_S(\mathbf{e}_{\theta_*}, \boldsymbol{\mu}). \quad (70)$$

Let $d_\Psi(\tilde{\mathbf{q}}) := \tilde{\Psi}(\tilde{\mathbf{q}}) - \tilde{\Psi}(\tilde{\boldsymbol{\mu}}) - \langle \tilde{\mathbf{q}} - \tilde{\boldsymbol{\mu}}, \nabla \tilde{\Psi}(\tilde{\boldsymbol{\mu}}) \rangle$. Observe that $d_\Psi(\tilde{\mathbf{q}}) = \Psi(\mathbf{q}) - \Psi(\boldsymbol{\mu}) - \langle \mathbf{q} - \boldsymbol{\mu}, \nabla \Psi(\boldsymbol{\mu}) \rangle = D_\Psi(\mathbf{q}, \boldsymbol{\mu})$. We define d_S similarly. Suppose that $\eta_\ell^\Psi > \eta_\ell^S = \underline{\eta}_\ell$. Then, from Corollary 18, $\forall \tilde{\mathbf{q}} \in \text{int } \tilde{\Delta}_k$, $\lambda_{\min}(\mathbf{H}\tilde{\Psi}(\tilde{\mathbf{q}})(\mathbf{H}\tilde{S}(\tilde{\mathbf{q}}))^{-1}) > 1$. This implies that $\forall \tilde{\mathbf{q}} \in \text{int } \tilde{\Delta}_k$, $\lambda_{\min}(\mathbf{H}d_\Psi(\tilde{\mathbf{q}})(\mathbf{H}d_S(\tilde{\mathbf{q}}))^{-1}) > 1$, and from Lemma 23, $d_\Psi - d_S$ must be strictly convex on $\text{int } \tilde{\Delta}_k$. We also have $\nabla d_\Psi(\tilde{\boldsymbol{\mu}}) - \nabla d_S(\tilde{\boldsymbol{\mu}}) = 0$ and $d_\Psi(\tilde{\boldsymbol{\mu}}) - d_S(\tilde{\boldsymbol{\mu}}) = 0$. Therefore, $d_\Psi - d_S$ attains a strict minimum at $\tilde{\boldsymbol{\mu}}$ (ibid., Thm. D.2.2.1); that is, $d_\Psi(\tilde{\mathbf{q}}) > d_S(\tilde{\mathbf{q}})$, $\forall \tilde{\mathbf{q}} \in \tilde{\Delta}_k \setminus \{\tilde{\boldsymbol{\mu}}\}$. In particular, for $\tilde{\mathbf{q}} = \Pi_k(\mathbf{e}_{\theta_*})$, we get $D_\Psi(\mathbf{e}_{\theta_*}, \boldsymbol{\mu}) = d_\Psi(\tilde{\mathbf{q}}) > d_S(\tilde{\mathbf{q}}) = D_S(\mathbf{e}_{\theta_*}, \boldsymbol{\mu})$, which contradicts (70). Therefore, $\eta_\ell^\Psi \leq \eta_\ell^S$, and thus

$$\begin{aligned} R_\ell^S(\boldsymbol{\mu}) &= \max_\theta D_S(\mathbf{e}_\theta, \boldsymbol{\mu}) / \eta_\ell^S = D_S(\mathbf{e}_{\theta_*}, \boldsymbol{\mu}) / \eta_\ell^S, \\ &\leq D_\Psi(\mathbf{e}_{\theta_*}, \boldsymbol{\mu}) / \eta_\ell^\Psi, \end{aligned} \quad (71)$$

$$\begin{aligned} &\leq \max_\theta D_\Psi(\mathbf{e}_\theta, \boldsymbol{\mu}) / \eta_\ell^\Psi, \\ &= R_\ell^\Psi(\boldsymbol{\mu}), \end{aligned} \quad (72)$$

where (71) is due to $D_\Psi(\mathbf{e}_{\theta_*}, \boldsymbol{\mu}) = D_S(\mathbf{e}_{\theta_*}, \boldsymbol{\mu})$ and $\eta_\ell^\Psi \leq \eta_\ell^S$. Equation 72, implies that $R_\ell^S(\boldsymbol{\mu}) \leq R_\ell^\Phi(\boldsymbol{\mu})$, since $R_\ell^\Psi(\boldsymbol{\mu}) = R_\ell^{\alpha\Phi}(\boldsymbol{\mu}) = R_\ell^\Phi(\boldsymbol{\mu})$ (Reid et al., 2015). Therefore,

$$\forall \boldsymbol{\mu} \in \text{ri } \Delta_k, R_\ell^S(\boldsymbol{\mu}) \leq R_\ell^\Phi(\boldsymbol{\mu}). \quad (73)$$

It remains to consider the case where $\boldsymbol{\mu}$ is in the relative boundary of Δ_k . Let $\boldsymbol{\mu} \in \text{rbd } \Delta_k$. There exists $\mathcal{I}_0 \subsetneq [k]$ such that $\boldsymbol{\mu} \in \Delta_{\mathcal{I}_0}$. Let $\theta^* \in [k] \setminus \mathcal{I}_0$ and $\mathcal{I} := \mathcal{I}_0 \cup \{\theta^*\}$. It holds that $\boldsymbol{\mu} \in \text{rbd } \Delta_{\mathcal{I}}$ and $\boldsymbol{\mu} + 2^{-1}(\mathbf{e}_{\theta^*} - \boldsymbol{\mu}) \in \text{ri } \Delta_{\mathcal{I}}$. Since ℓ is Φ -mixable, it follows from Proposition 11 and the 1-homogeneity of $\Phi'(\boldsymbol{\mu}; \cdot)$ (Hiriart-Urruty and Lemaréchal, 2001, Prop. D.1.1.2) that

$$\Phi'(\boldsymbol{\mu}; \mathbf{e}_{\theta^*} - \boldsymbol{\mu}) = 2\Phi'(\boldsymbol{\mu}; [\boldsymbol{\mu} + 2^{-1}(\mathbf{e}_{\theta^*} - \boldsymbol{\mu})] - \boldsymbol{\mu}) = -\infty.$$

Hence,

$$\begin{aligned} R_\ell^\Phi(\boldsymbol{\mu}) &= \max_{\theta \in [k]} D_\Phi(\mathbf{e}_\theta, \boldsymbol{\mu}), \\ &\geq D_\Phi(\mathbf{e}_{\theta^*}, \boldsymbol{\mu}) = \Phi(\mathbf{e}_{\theta^*}) - \Phi(\boldsymbol{\mu}) - \Phi'(\boldsymbol{\mu}; \mathbf{e}_{\theta^*} - \boldsymbol{\mu}) = +\infty. \end{aligned} \quad (74)$$

Inequality 74 also applies to S, since ℓ is $(\underline{\eta}_\ell^{-1} S)$ -mixable. From (74) and (73), we conclude that $\forall \boldsymbol{\mu} \in \Delta_k$, $R_\ell^S(\boldsymbol{\mu}) \leq R_\ell^\Phi(\boldsymbol{\mu})$. \blacksquare

Appendix D. Legendre Φ , but no Φ -mixable ℓ

In this appendix, we construct a *Legendre type* entropy (Rockafellar, 1997) for which there are no Φ -mixable losses satisfying a weak condition (see below).

Let $\ell : \mathcal{A} \rightarrow [0, +\infty]^n$ be a loss satisfying condition I. According to Alexandrov's Theorem, a concave function is twice differentiable almost everywhere (see e.g. (Borwein et al., 2010, Thm. 6.7)). Now we give a version of Theorem 16 which does not assume the twice differentiability of the Bayes risk. The proof is almost identical to that of Theorem 16 with only minor modifications.

Theorem 38 *Let $\Phi : \mathbb{R}^k \rightarrow \mathbb{R} \cup \{+\infty\}$ be an entropy such that $\tilde{\Phi}$ is twice differentiable on $\text{int } \tilde{\Delta}_k$, and $\ell : \mathcal{A} \rightarrow [0, +\infty]^n$ a loss satisfying Condition I and such that $\exists(\tilde{\mathbf{p}}, \mathbf{v}) \in \mathcal{D} \times \mathbb{R}^{\tilde{n}}, \mathbf{H}_{\tilde{L}_\ell}(\tilde{\mathbf{p}})\mathbf{v} \neq \mathbf{0}_{\tilde{n}}$, where $\mathcal{D} \subset \text{int } \tilde{\Delta}_n$ is a set of Lebesgue measure 1 where \tilde{L}_ℓ is twice differentiable, and define*

$$\underline{\eta}_\ell^* := \inf_{\tilde{\mathbf{p}} \in \mathcal{D}} (\lambda_{\max}([\mathbf{H}_{\tilde{L}_{\log}}(\tilde{\mathbf{p}})]^{-1} \mathbf{H}_{\tilde{L}_\ell}(\tilde{\mathbf{p}})))^{-1}. \quad (75)$$

Then ℓ is Φ -mixable only if $\underline{\eta}_\ell^* \Phi - \mathbf{S}$ is convex on Δ_k .

The new condition on the Bayes risk is much weaker than requiring \underline{L}_ℓ to be twice differentiability on $]0, +\infty[^n$. In the next example, we will show that there exists a Legendre type entropy for which there are no Φ -mixable losses satisfying the condition of Theorem 38.

Example 39 *Let $\Phi : \mathbb{R}^2 \rightarrow \mathbb{R} \cup \{+\infty\}$ be an entropy such that*

$$\forall q \in]0, 1[, \Phi(q, 1-q) = \tilde{\Phi}(q) = \int_{1/2}^q \log \left(\frac{\log(1-t)}{\log t} \right) dt.$$

$\tilde{\Phi}$ is differentiable and strictly convex on the open set $(0, 1)$. Furthermore, it satisfies (8) which makes it a function of Legendre type (Rockafellar, 1997, Lem. 26.2). In fact, (8) is satisfied due to

$$\begin{aligned} \left| \frac{d}{dq} \tilde{\Phi}(q) \right| &= \left| \log \left(\frac{\log(1-q)}{\log q} \right) \right| \xrightarrow{q \rightarrow b} +\infty, \text{ where } b \in \{0, 1\}, \\ \frac{d^2}{dq^2} \tilde{\Phi}(q) &= \frac{-1}{q \log q} + \frac{-1}{(1-q) \log(1-q)} > 0, \forall q \in]0, 1[. \end{aligned}$$

The Shannon entropy on Δ_2 is defined by $\mathbf{S}(q, 1-q) = \tilde{\mathbf{S}}(q) = q \log q + (1-q) \log(1-q)$, for $q \in]0, 1[$. Thus, $\frac{d^2}{dq^2} \tilde{\mathbf{S}}(q) = \frac{1}{q(1-q)}$.

Suppose now that there exists a Φ -mixable loss $\ell : \mathcal{A} \rightarrow [0, +\infty]^n$ satisfying condition I and such that $\exists(\tilde{\mathbf{p}}, \mathbf{v}) \in \mathcal{D} \times \mathbb{R}^{\tilde{n}}, \mathbf{H}_{\tilde{L}_\ell}(\tilde{\mathbf{p}})\mathbf{v} \neq \mathbf{0}_{\tilde{n}}$. Let $\underline{\eta}_\ell^*$ be as in (75). By definition, we have $\underline{\eta}_\ell^* < +\infty$, and thus

$$\underline{\eta}_\ell^* \left[\frac{d^2}{dq^2} \tilde{\Phi}(q) \right] \left[\frac{d^2}{dq^2} \tilde{\mathbf{S}}(q) \right]^{-1} = \underline{\eta}_\ell^* \left(\frac{q-1}{\log q} + \frac{-q}{\log(1-q)} \right) \xrightarrow{q \rightarrow b} 0, \quad (76)$$

where $b \in \{0, 1\}$. From Lemma 23, (76) implies that $\underline{\eta}_\ell^* \Phi - \mathbf{S}$ is not convex on Δ_k , which is a contradiction according to Theorem 38.

Appendix E. Loss Surface and Superprediction Set

In this appendix, we derive an expression for the curvature of the image of a proper loss function. We will need the following lemma.

Lemma 40 *Let $\sigma : [0, +\infty[^n \rightarrow \mathbb{R}$ be a 1-homogeneous, twice differentiable function on $]0, +\infty[^n$. Then σ is concave on $]0, +\infty[^n$ if and only if $\tilde{\sigma} = \sigma \circ \Pi_n$ is concave on $\text{int } \tilde{\Delta}_n$.*

Proof The forward implication is immediate; if σ is concave on $]0, +\infty[^n$, then $\sigma \circ \Pi_k$ is concave on $\text{int } \tilde{\Delta}_k$, since Π_k is an affine function.

Now assume that $\tilde{\sigma}$ is concave on $\text{int } \tilde{\Delta}_k$. Let $\lambda \in [0, 1]$ and $(\mathbf{p}, \mathbf{q}) \in [0, +\infty[^n \times [0, +\infty[^n$. We need to show that

$$\lambda\sigma(\mathbf{p}) + (1 - \lambda)\sigma(\mathbf{q}) \leq \sigma(\lambda\mathbf{p} + (1 - \lambda)\mathbf{q}). \quad (77)$$

Note that if $\mathbf{p} = \mathbf{0}$ or $\mathbf{q} = \mathbf{0}$, (77) is trivially with equality due to the 1-homogeneity of σ . Now assume that \mathbf{p} and \mathbf{q} are non-zero and let $c := \lambda \|\mathbf{p}\|_1 + (1 - \lambda) \|\mathbf{q}\|_1$. For convenience, we also denote $\mathbf{p}_1 = \mathbf{p} / \|\mathbf{p}\|_1$ and $\mathbf{q}_1 = \mathbf{q} / \|\mathbf{q}\|_1$ which are both in Δ_n . It follows that

$$\begin{aligned} \lambda\sigma(\mathbf{p}) + (1 - \lambda)\sigma(\mathbf{q}) &= cM \left(\lambda \frac{\|\mathbf{p}\|_1}{c} \sigma(\mathbf{p}_1) + (1 - \lambda) \frac{\|\mathbf{q}\|_1}{c} \sigma(\mathbf{q}_1) \right), \\ &= c \left(\lambda \frac{\|\mathbf{p}\|_1}{c} \tilde{\sigma}(\tilde{\mathbf{p}}_1) + (1 - \lambda) \frac{\|\mathbf{q}\|_1}{c} \tilde{\sigma}(\tilde{\mathbf{q}}_1) \right), \\ &\leq c\tilde{\sigma} \left(\lambda \frac{\|\mathbf{p}\|_1}{c} \tilde{\mathbf{p}}_1 + (1 - \lambda) \frac{\|\mathbf{q}\|_1}{c} \tilde{\mathbf{q}}_1 \right), \\ &= c\sigma \left(\lambda \frac{\|\mathbf{p}\|_1}{c} \mathbf{p}_1 + (1 - \lambda) \frac{\|\mathbf{q}\|_1}{c} \mathbf{q}_1 \right), \\ &= \sigma(\lambda\mathbf{p} + (1 - \lambda)\mathbf{q}), \end{aligned}$$

where the first and last equalities are due the 1-homogeneity of σ and the inequality is due to $\tilde{\sigma}$ being concave on the $\text{int } \tilde{\Delta}_n$. \blacksquare

E.1. Curvature of the Loss Surface

The *normal curvature* of a \tilde{n} -manifold \mathcal{S} (Thorpe, 1994) at a point $\mathbf{r} \in \mathcal{S}$ in the direction of $\mathbf{w} \in T_{\mathbf{r}}\mathcal{S}$, where $T_{\mathbf{r}}\mathcal{S}$ is the *tangent space* of \mathcal{S} at $\mathbf{r} \in \mathcal{S}$, is defined by

$$\kappa(\mathbf{r}, \mathbf{w}) = \frac{\langle \mathbf{w}, \text{DN}^{\mathcal{S}}(\mathbf{r})\mathbf{w} \rangle}{\langle \mathbf{w}, \mathbf{w} \rangle}, \quad (78)$$

where $\text{N}^{\mathcal{S}}(\mathbf{r})$ is the normal vector to the surface at \mathbf{r} . The *minimum principal curvature* of \mathcal{S} at \mathbf{r} is expressed as $\underline{\kappa}(\mathbf{r}) := \inf\{\kappa(\mathbf{r}, \mathbf{w}) : \mathbf{w} \in T_{\mathbf{r}}\mathcal{S} \cap \mathcal{B}(\mathbf{r}, 1)\}$.

In the next theorem, we establish a direct link between the curvature of a loss surface and the Hessian of the loss' Bayes risk.

Theorem 41 *Let $\ell : \text{ri } \Delta_n \rightarrow [0, +\infty[^n$ be a loss whose Bayes risk is twice differentiable and strictly concave on $]0, +\infty[^n$. Let $\mathbf{p} \in \text{ri } \Delta_n$, $X_{\mathbf{p}} := I_{\tilde{n}} - \tilde{\mathbf{p}}\mathbf{1}_{\tilde{n}}^{\top}$, and $\mathbf{w} \in T_{\tilde{\ell}(\tilde{\mathbf{p}})}\mathcal{S}_{\ell}$. Then*

1. $\exists \mathbf{v} \in \mathbb{R}^{n-1}$ such that $D\tilde{\ell}(\tilde{\mathbf{p}})\mathbf{v} = \mathbf{w}$.
2. \mathcal{S}_ℓ is a \tilde{n} -manifold.
3. The normal curvature of \mathcal{S}_ℓ at $\ell(\mathbf{p}) = \tilde{\ell}(\tilde{\mathbf{p}})$ in the direction \mathbf{w} is given by

$$\kappa_\ell(\ell(\mathbf{p}), \mathbf{w}) = \left\| \begin{bmatrix} X_{\mathbf{p}} \\ -\tilde{\mathbf{p}}^\top \end{bmatrix} (-H\tilde{L}_\ell(\tilde{\mathbf{p}}))^{\frac{1}{2}} \mathbf{u} \right\|^{-1}, \quad (79)$$

where $\mathbf{u} = (-H\tilde{L}_\ell(\tilde{\mathbf{p}}))^{\frac{1}{2}} \mathbf{v} / \|(-H\tilde{L}_\ell(\tilde{\mathbf{p}}))^{\frac{1}{2}} \mathbf{v}\|$.

It becomes clear from (79) that smaller eigenvalues of $-H\tilde{L}_\ell(\tilde{\mathbf{p}})$ will tend to make the loss surface more curved at $\ell(\mathbf{p})$, and vice versa.

Before proving Theorem 41, we first define parameterizations on manifolds.

Definition 42 (Local and Global Parameterization) Let $\mathcal{S} \subseteq \mathbb{R}^n$ be a \tilde{n} -manifold and \mathcal{U} an open set in $\mathbb{R}^{\tilde{n}}$. The map $\varphi : \mathcal{U} \rightarrow \mathcal{S}$ is called a local parameterization of \mathcal{S} if $D\varphi(\mathbf{u}) : \mathbb{R}^{\tilde{n}} \rightarrow T_{\varphi(\mathbf{u})}\mathcal{S}$ is injective for all $\mathbf{u} \in \mathcal{U}$, where $T_{\varphi(\mathbf{u})}\mathcal{S}$ is the tangent space of \mathcal{S} at $\varphi(\mathbf{u}) \in \mathcal{S}$. φ is called a global parameterization of \mathcal{S} if it is, additionally, onto.

Let φ be a global parameterization of \mathcal{S} and $N^\varphi := N^{\mathcal{S}} \circ \varphi$. By a direct application of the chain rule, (78) can be written as

$$\kappa(\varphi(\mathbf{u}), \mathbf{w}) = \frac{\langle \mathbf{w}, DN^\varphi(\mathbf{u})\mathbf{v} \rangle}{\langle \mathbf{w}, \mathbf{w} \rangle}, \quad (80)$$

where \mathbf{v} is such that $D\varphi(\mathbf{u})\mathbf{v} = \mathbf{w}$. The existence of such a \mathbf{v} is guaranteed by the fact that $D\varphi$ is injective and $\dim \mathbb{R}^{\tilde{n}} = \dim T_{\varphi(\mathbf{u})}\mathcal{S} = \tilde{n}$.

Proof [Theorem 41] First we show that \mathcal{S}_ℓ is a \tilde{n} -manifold. Consider the map $\tilde{\ell} : \text{int } \tilde{\Delta}_n \rightarrow \mathcal{S}_\ell$ and note that $\text{int } \tilde{\Delta}_n$ is trivially a \tilde{n} -manifold. Due to the strict concavity of the Bayes risk, $\tilde{\ell}$ is injective (van Erven et al., 2012) and from Lemmas 27 and 40, $D\tilde{\ell}(\tilde{\mathbf{p}}) : \mathbb{R}^{\tilde{n}} \rightarrow T_{\tilde{\ell}(\tilde{\mathbf{p}})}\mathcal{S}_\ell$ is also injective. Therefore, $\tilde{\ell}$ is an immersion (Robbin and Salamon, 2011). $\tilde{\ell}$ is also proper in the sense that the preimage of every compact subset of \mathcal{S}_ℓ is compact. Therefore, $\tilde{\ell}$ is a proper injective immersion, and thus it is an embedding from the \tilde{n} -manifold $\text{int } \tilde{\Delta}_n$ to \mathcal{S}_ℓ (ibid.). Hence, \mathcal{S}_ℓ is a manifold.

Now we prove (79). The map $\tilde{\ell}$ is a global parameterization of \mathcal{S}_ℓ . In fact, from Lemma 27, $D\tilde{\ell}(\tilde{\mathbf{p}})$ has rank \tilde{n} , for all $\tilde{\mathbf{p}} \in \text{int } \tilde{\Delta}_n$, which implies that $D\tilde{\ell}(\tilde{\mathbf{p}})$ is onto from $\mathbb{R}^{\tilde{n}}$ to $T_{\tilde{\ell}(\tilde{\mathbf{p}})}\mathcal{S}_\ell$. Therefore, given $\mathbf{w} \in T_{\tilde{\ell}(\tilde{\mathbf{p}})}\mathcal{S}_\ell$, there exists $\mathbf{v} \in \mathbb{R}^{\tilde{n}}$ such that $\mathbf{w} = D\tilde{\ell}(\tilde{\mathbf{p}})\mathbf{v}$. Furthermore, Lemma 27 implies that $N^{\tilde{\ell}}(\tilde{\mathbf{p}}) = \mathbf{p}$, since $\langle \mathbf{p}, D\tilde{\ell}(\tilde{\mathbf{p}}) \rangle = \mathbf{0}_{\tilde{n}}^\top$. Substituting $N^{\tilde{\ell}}$ into (80) yields

$$\begin{aligned}
 \kappa_\ell(\tilde{\ell}(\tilde{\mathbf{p}}), \mathbf{w}) &= \frac{\mathbf{v}^\top (\mathbf{D}\tilde{\ell}(\tilde{\mathbf{p}}))^\top \begin{bmatrix} I_{\tilde{n}}, \\ \mathbf{1}_{\tilde{n}} \end{bmatrix} \mathbf{v}}{\langle \mathbf{D}\tilde{\ell}(\tilde{\mathbf{p}})\mathbf{v}, \mathbf{D}\tilde{\ell}(\tilde{\mathbf{p}})\mathbf{v} \rangle}, \\
 &= \frac{\mathbf{v}^\top \mathbf{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}}) \begin{bmatrix} X_{\tilde{\mathbf{p}}}^\top, & -\tilde{\mathbf{p}} \end{bmatrix} \begin{bmatrix} I_{\tilde{n}} \\ \mathbf{1}_{\tilde{n}} \end{bmatrix} \mathbf{v}}{\langle \mathbf{D}\tilde{\ell}(\tilde{\mathbf{p}})\mathbf{v}, \mathbf{D}\tilde{\ell}(\tilde{\mathbf{p}})\mathbf{v} \rangle}, \\
 &= \frac{\mathbf{v}^\top \mathbf{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}})\mathbf{v}}{\mathbf{v}^\top \mathbf{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}}) \begin{bmatrix} X_{\tilde{\mathbf{p}}}^\top, & -\tilde{\mathbf{p}} \end{bmatrix} \begin{bmatrix} X_{\tilde{\mathbf{p}}} \\ -\tilde{\mathbf{p}}^\top \end{bmatrix} \mathbf{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}})\mathbf{v}}.
 \end{aligned} \tag{81}$$

Setting $\mathbf{u} = (-\mathbf{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}}))^{\frac{1}{2}}\mathbf{v} / \|(-\mathbf{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}}))^{\frac{1}{2}}\mathbf{v}\|$ in (81) gives the desired result. \blacksquare

Appendix F. Classical Mixability

In this appendix, we provide a more concise proof of the necessary and sufficient conditions for the convexity of the superprediction set (van Erven et al., 2012).

Theorem 43 *Let $\ell: \Delta_n \rightarrow [0, +\infty]^n$ be a strictly proper loss whose Bayes risk is twice differentiable on $]0, +\infty[^n$. The following points are equivalent;*

- (i) $\forall \tilde{\mathbf{p}} \in \text{int } \tilde{\Delta}_n, \eta \mathbf{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}}) \succeq \mathbf{H}\tilde{\mathcal{L}}_{\log}(\tilde{\mathbf{p}})$.
- (ii) $e^{-\eta \mathcal{S}_\ell^\oplus} = \bigcap_{\mathbf{p} \in \Delta_n} \mathcal{H}_{\tau(\mathbf{p}), 1} \cap]0, +\infty[^n$, where $\tau(\mathbf{p}) := \mathbf{p} \odot e^{\eta \ell(\mathbf{p})}$.
- (iii) $e^{-\eta \mathcal{S}_\ell^\oplus}$ is convex.

Proof We already showed (i) \implies (ii) \implies (iii) in the proof of Theorem 7.

We now show (iii) \implies (i). Since $e^{-\eta \mathcal{S}_\ell^\oplus}$ is convex, any point $\mathbf{s} \in \text{bd } e^{-\eta \mathcal{S}_\ell^\oplus}$ is supported by a hyperplane (Hiriart-Urruty and Lemaréchal, 2001, Lem. A.4.2.1). Since $\mathbf{u} \rightarrow e^{-\eta \mathbf{u}}$ is a homeomorphism, it maps boundaries to boundaries. From this and Lemma 33, $\text{bd } e^{-\eta \mathcal{S}_\ell^\oplus} = e^{-\eta \mathcal{S}_\ell}$. Thus, for $\mathbf{p} \in \text{ri } \Delta_n$, there exists a unit-norm vector $\mathbf{u} \in \mathbb{R}^n$ such that for all $\mathbf{s} \in \mathcal{S}_\ell^\oplus$ it either holds that $\langle \mathbf{u}, e^{-\eta \ell(\mathbf{p})} \rangle \leq \langle \mathbf{u}, e^{-\eta \mathbf{s}} \rangle$; or $\langle \mathbf{u}, e^{-\eta \ell(\mathbf{p})} \rangle \geq \langle \mathbf{u}, e^{-\eta \mathbf{s}} \rangle$. It is easy to see that it is the latter case that holds, since we can choose $\mathbf{s} = \ell(\mathbf{r}) + c\mathbf{1} \in \mathcal{S}_\ell^\oplus$, for $\mathbf{r} \in \Delta_n$, and make $\langle \mathbf{u}, e^{-\eta \mathbf{s}} \rangle$ arbitrarily small by making $c \in \mathbb{R}$ large. Therefore, $\forall \mathbf{r} \in \text{ri } \Delta_n, \langle \mathbf{u}, e^{-\eta \tilde{\ell}(\tilde{\mathbf{p}})} \rangle = \langle \mathbf{u}, e^{-\eta \ell(\mathbf{p})} \rangle \geq \langle \mathbf{u}, e^{-\eta \ell(\mathbf{r})} \rangle = \langle \mathbf{u}, e^{-\eta \tilde{\ell}(\tilde{\mathbf{r}})} \rangle$ and $\tilde{\mathbf{p}}$ is a critical point of the function $f(\tilde{\mathbf{r}}) := \langle \mathbf{u}, e^{-\eta \tilde{\ell}(\tilde{\mathbf{r}})} \rangle$ on $\text{int } \tilde{\Delta}_n$. This implies that $\nabla f(\tilde{\mathbf{p}}) = \mathbf{0}_{\tilde{n}}$; that is, $-\eta \langle \mathbf{u}, \text{diag}(e^{-\eta \tilde{\ell}(\tilde{\mathbf{p}})}) \mathbf{D}\tilde{\ell}(\tilde{\mathbf{p}}) \rangle = -\eta \langle \text{diag}(e^{-\eta \tilde{\ell}(\tilde{\mathbf{p}})}) \mathbf{u}, \mathbf{D}\tilde{\ell}(\tilde{\mathbf{p}}) \rangle = \mathbf{0}_{\tilde{n}}^\top$. From Lemma 27, there exists $\lambda \in \mathbb{R}$ such that $\text{diag}(e^{-\eta \tilde{\ell}(\tilde{\mathbf{p}})}) \mathbf{u} = \lambda \tilde{\mathbf{p}}$. Therefore, $\mathbf{u} = \lambda \tilde{\mathbf{p}} \odot e^{\eta \tilde{\ell}(\tilde{\mathbf{p}})}$, where $\lambda = \|\tilde{\mathbf{p}} \odot e^{\eta \tilde{\ell}(\tilde{\mathbf{p}})}\|^{-1}$, since $\|\mathbf{u}\| = 1$. For $\mathbf{v} \in \mathbb{R}^{n-1}$, let $\tilde{\boldsymbol{\alpha}}^t := \tilde{\mathbf{p}} + t\mathbf{v}$, where $t \in \{s : \tilde{\mathbf{p}} + s\mathbf{v} \in$

$\text{int } \tilde{\Delta}_n\}$. Since f is twice differentiable and attains a maximum at $\tilde{\mathbf{p}}$,

$$\begin{aligned}
 0 &\geq \frac{1}{\lambda\eta} \frac{d^2}{dt^2} f \circ \tilde{\boldsymbol{\alpha}}^t \Big|_{t=0} = \frac{1}{\lambda} \frac{d}{dt} \left\langle \mathbf{u}, \text{diag } e^{-\eta\tilde{\ell}(\tilde{\boldsymbol{\alpha}}^t)} \text{D}\tilde{\ell}(\tilde{\boldsymbol{\alpha}}^t)\mathbf{v} \right\rangle \Big|_{t=0}, \\
 &= \frac{d}{dt} \left\langle \mathbf{p} \odot e^{\eta\tilde{\ell}(\tilde{\mathbf{p}})}, \text{diag } e^{-\eta\tilde{\ell}(\tilde{\boldsymbol{\alpha}}^t)} \text{D}\tilde{\ell}(\tilde{\boldsymbol{\alpha}}^t)\mathbf{v} \right\rangle \Big|_{t=0} + \frac{d}{dt} \left\langle \mathbf{p}, \text{D}\tilde{\ell}(\tilde{\boldsymbol{\alpha}}^t)\mathbf{v} \right\rangle \Big|_{t=0}, \\
 &= \eta\mathbf{v}^\top \text{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}}) (\text{H}\tilde{\mathcal{L}}_{\log}(\tilde{\mathbf{p}}))^{-1} \text{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}})\mathbf{v} - \mathbf{v}^\top \text{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}})\mathbf{v}, \tag{82}
 \end{aligned}$$

where in the second equality we substituted \mathbf{u} by $\lambda\mathbf{p} \odot e^{\eta\tilde{\ell}(\tilde{\mathbf{p}})}$ and in (82) we used (19) and (20) from Lemma 28. Note that by the assumptions on ℓ it follows that the Bayes risk $\tilde{\mathcal{L}}_\ell$ is strictly concave (van Erven et al., 2012, Lemma 6) and $-\text{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}})$ is symmetric negative-definite. In particular, $\text{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}})$ is invertible. Setting $\hat{\mathbf{v}} := \text{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}})\mathbf{v}$ in (82) yields

$$0 \geq \eta\hat{\mathbf{v}} (\text{H}\tilde{\mathcal{L}}_{\log}(\tilde{\mathbf{p}}))^{-1} \hat{\mathbf{v}} - \hat{\mathbf{v}} (\text{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}}))^{-1} \hat{\mathbf{v}}.$$

Since $\mathbf{v} \in \mathbb{R}^{n-1}$ was chosen arbitrarily, $(\text{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}}))^{-1} \succeq \eta(\text{H}\tilde{\mathcal{L}}_{\log}(\tilde{\mathbf{p}}))^{-1}, \forall \tilde{\mathbf{p}} \in \text{int } \tilde{\Delta}_n$. This is equivalent to the condition $\forall \tilde{\mathbf{p}} \in \text{int } \tilde{\Delta}_n, \eta\text{H}\tilde{\mathcal{L}}_\ell(\tilde{\mathbf{p}}) \succeq \text{H}\tilde{\mathcal{L}}_{\log}(\tilde{\mathbf{p}})$. \blacksquare