

Exp-Concavity of Proper Composite Losses

Parameswaran Kamalaruban

KAMALARUBAN.PARAMESWARAN@NICTA.COM.AU

Robert C. Williamson

BOB.WILLIAMSON@NICTA.COM.AU

Xinhua Zhang

XINHUA.ZHANG@NICTA.COM.AU

Australian National University and NICTA, Canberra ACT 0200, Australia

Abstract

The goal of online prediction with expert advice is to find a decision strategy which will perform almost as well as the best expert in a given pool of experts, on any sequence of outcomes. This problem has been widely studied and $O(\sqrt{T})$ and $O(\log T)$ regret bounds can be achieved for convex losses (Zinkevich (2003)) and strictly convex losses with bounded first and second derivatives (Hazan et al. (2007)) respectively. In special cases like the Aggregating Algorithm (Vovk (1995)) with mixable losses and the Weighted Average Algorithm (Kivinen and Warmuth (1999)) with exp-concave losses, it is possible to achieve $O(1)$ regret bounds. van Erven (2012) has argued that mixability and exp-concavity are roughly equivalent under certain conditions. Thus by understanding the underlying relationship between these two notions we can gain the best of both algorithms (strong theoretical performance guarantees of the Aggregating Algorithm and the computational efficiency of the Weighted Average Algorithm). In this paper we provide a complete characterization of the exp-concavity of any proper composite loss. Using this characterization and the mixability condition of proper losses (Van Erven et al. (2012)), we show that it is possible to transform (re-parameterize) any β -mixable binary proper loss into a β -exp-concave composite loss with the same β . In the multi-class case, we propose an approximation approach for this transformation.

Keywords: proper scoring rules, link functions, composite losses, sequential prediction, regret bound, aggregating algorithm, weighted average algorithm, mixability, exp-concavity, substitution functions.

1. Introduction

Loss functions are the means by which the quality of a prediction in learning problem is evaluated. A composite loss (the composition of a class probability estimation (CPE) loss with an invertible link function which is essentially just a re-parameterization) is proper if its risk is minimized when predicting the true underlying class probability (a formal definition is given later). In Vernet et al. (2014), there is an argument that shows that there is no point in using losses that are neither proper nor proper composite as they are inadmissible. Flexibility in the choice of loss function is important to tailor the solution to a learning problem (Buja et al. (2005), Hand (1994), Hand and Vinciotti (2003)), and it could be attained by characterizing the set of loss functions using natural parameterizations.

The goal of the learner in a *game of prediction with expert advice* (which is formally described in section 2.5) is to predict as well as the best expert in the given pool of experts. The regret bound of the learner depends on the merging scheme used to merge the experts' predictions and the type of loss function used to measure the performance. It has already been shown that constant regret

bounds are achievable for mixable losses when the Aggregating Algorithm is the merging scheme (Vovk (1995)), and for exp-concave losses when the Weighted Average Algorithm is the merging scheme (Kivinen and Warmuth (1999)). We can see that the exp-concavity trivially implies mixability. Even though the converse implication is not true in general, under some re-parameterization we can make it possible. This paper discusses general conditions on proper losses under which they can be transformed to an exp-concave loss through a suitable link function. In the binary case, these conditions give two concrete formulas (Proposition 1 and Corollary 8) for link functions that can transform β -mixable proper losses into β -exp-concave, proper, composite losses. The explicit form of the link function given in Proposition 1 is derived using the same geometric construction used in van Erven (2012).

Further we extend the work by Vernet et al. (2014), to provide a complete characterization of the exp-concavity of the proper composite multi-class losses in terms of the Bayes risk associated with the underlying proper loss, and the link function. The mixability of proper losses (mixability of a proper composite loss is equivalent to the mixability of its generating proper loss) is studied in Van Erven et al. (2012). Using these characterizations (for the binary case), in Corollary 8 we derive an *exp-concavifying link* function that can also transform any β -mixable proper loss into a β -exp-concave composite loss. Since for the multi-class losses these conditions do not hold in general, we propose a geometric approximation approach (Proposition 2) which takes a parameter ϵ and transforms the mixable loss function appropriately on a subset S_ϵ of the prediction space. When the prediction space is Δ^n , any prediction belongs to the subset S_ϵ for sufficiently small ϵ . In the conclusion we provide a way to use the Weighted Average Algorithm with learning rate β for proper β -mixable but non-exp-concave loss functions to achieve $O(1)$ regret bound.

The exp-concave losses achieve $O(\log T)$ regret bound in online convex optimization algorithms, which is a more general setting of online learning problems. Thus the exp-concavity characterization of composite losses could be helpful in constructing exp-concave losses for online learning problems.

The paper is organized as follows. In Section 2 we formally introduce the loss function, several loss types, conditional risk, proper composite losses and a game of prediction with expert advice. In Section 3 we consider our main problem — whether one can always find a link function to transform β -mixable losses into β -exp-concave losses. Section 4 concludes with a brief discussion. The impact of the choice of substitution function on the regret of the learner is explored via experiments in Appendix A. In Appendix B, we discuss the mixability conditions of *probability games* with continuous outcome space. Detailed proofs are in Appendix C.

2. Preliminaries and Background

This section provides the necessary background on loss functions, conditional risks, and the sequential prediction problem.

2.1. Notation

We use the following notation throughout. A superscript prime, A' denotes transpose of the matrix or vector A , except when applied to a real-valued function where it denotes derivative (f'). We denote the matrix multiplication of compatible matrices A and B by $A \cdot B$, so the inner product of two vectors $x, y \in \mathbb{R}^n$ is $x' \cdot y$. Let $[n] := \{1, \dots, n\}$, $\mathbb{R}_+ := [0, \infty)$ and the n -simplex $\Delta^n := \{(p_1, \dots, p_n)' : 0 \leq p_i \leq 1, \forall i \in [n], \text{ and } \sum_{i \in [n]} p_i = 1\}$. If x is a n -vector, $A = \text{diag}(x)$ is

the $n \times n$ matrix with entries $A_{i,i} = x_i$, $i \in [n]$ and $A_{i,j} = 0$ for $i \neq j$. If $A - B$ is positive definite (resp. semi-definite), then we write $A \succ B$ (resp. $A \succeq B$). We use e_i^n to denote the i th n -dimensional unit vector, $e_i^n = (0, \dots, 0, 1, 0, \dots, 0)'$ when $i \in [n]$, and define $e_i^n = 0_n$ when $i > n$. The n -vector $\mathbb{1}_n := (1, \dots, 1)'$. We write $\llbracket P \rrbracket = 1$ if P is true and $\llbracket P \rrbracket = 0$ otherwise. Given a set S and a weight vector w , the *convex combination* of the elements of the set w.r.t the weight vector is denoted by $\text{co}_w S$, and the *convex hull* of the set which is the set of all possible convex combinations of the elements of the set is denoted by $\text{co} S$ (Rockafellar (1970)). If $S, T \subset \mathbb{R}^n$, then the *Minkowski sum* $S \oplus T := \{s + t : s \in S, t \in T\}$. $\mathcal{Y}^{\mathcal{X}}$ represents the set of all functions $f : \mathcal{X} \rightarrow \mathcal{Y}$. Other notation (the Kronecker product \otimes , the Jacobian D , and the Hessian H) is defined in Appendix A of Van Erven et al. (2012). $Df(v)$ and $Hf(v)$ denote the Jacobian and Hessian of $f(v)$ w.r.t. v respectively. When it is not clear from the context, we will explicitly mention the variable; for example $D_{\tilde{v}} f(v)$ where $v = h(\tilde{v})$.

2.2. Loss Functions

For a prediction problem with an instance space \mathcal{X} , outcome space \mathcal{Y} and prediction space \mathcal{V} , a loss function $\ell : \mathcal{Y} \times \mathcal{V} \rightarrow \mathbb{R}_+$ (bivariate function representation) can be defined to assign a penalty $\ell(y, v)$ for predicting $v \in \mathcal{V}$ when the actual outcome is $y \in \mathcal{Y}$. When the outcome space $\mathcal{Y} = [n]$, $n \geq 2$, the loss function ℓ is called a *multi-class loss* and it can be expressed in terms of its partial losses $\ell_i := \ell(i, \cdot)$ for any outcome $i \in [n]$, as $\ell(y, v) = \sum_{i \in [n]} \llbracket y = i \rrbracket \ell_i(v)$. The vector representation of the multi-class loss is given by $\ell : \mathcal{V} \rightarrow \mathbb{R}_+^n$, which assigns a vector $\ell(v) = (\ell_1(v), \dots, \ell_n(v))'$ to each prediction $v \in \mathcal{V}$. A loss is differentiable if all of its partial losses are differentiable. In this paper, we will use the bivariate function representation $(\ell(y, v))$ to denote a general loss function and the vector representation for multi-class loss functions.

The *super-prediction set* of a binary loss ℓ is defined as

$$S_\ell := \{x \in \mathbb{R}^n : \exists v \in \mathcal{V}, x \geq \ell(v)\},$$

where inequality is component-wise. For any dimension n and $\beta \geq 0$, the β -exponential operator $E_\beta : [0, \infty]^n \rightarrow [0, 1]^n$ is defined by $E_\beta(x) := (e^{-\beta x_1}, \dots, e^{-\beta x_n})'$. For $\beta > 0$ it is clearly invertible with inverse $E_\beta^{-1}(z) = -\beta^{-1}(\ln z_1, \dots, \ln z_n)'$. The β -exponential transformation of the super-prediction set is given by

$$E_\beta(S_\ell) := \{(e^{-\beta x_1}, \dots, e^{-\beta x_n})' \in \mathbb{R}^n : (x_1, \dots, x_n)' \in S_\ell\}, \quad \beta > 0.$$

A multi-class loss ℓ is *convex* if $f(v) = \ell_y(v)$ is convex in v for all $y \in [n]$, α -*exp-concave* (for $\alpha > 0$) if $f(v) = e^{-\alpha \ell_y(v)}$ is concave in v for all $y \in [n]$, *weakly mixable* if the super-prediction set S_ℓ is convex (Kalnishkan and Vyugin (2005)), and β -*mixable* (for $\beta > 0$) if the set $E_\beta(S_\ell)$ is convex (Vovk and Zhdanov (2009); Vovk (1995)). The *mixability constant* β_ℓ of a loss ℓ is the largest β such that ℓ is β -mixable; i.e. $\beta_\ell := \sup \{\beta > 0 : \ell \text{ is } \beta\text{-mixable}\}$. If the loss function ℓ is α -exp-concave (resp. β -mixable) then it is α' -exp-concave for any $0 < \alpha' \leq \alpha$ (resp. β' -mixable for any $0 < \beta' \leq \beta$), and its λ -scaled version $(\lambda \ell)$ for some $\lambda > 0$ is $\frac{\alpha}{\lambda}$ -exp-concave (resp. $\frac{\beta}{\lambda}$ -mixable). If the loss ℓ is α -exp-concave, then it is convex and α -mixable.

For a multi-class loss ℓ , if the prediction space $\mathcal{V} = \Delta^n$ then it is said to be *multi-class probability estimation (CPE) loss*, where the predicted values are directly interpreted as probability estimates: $\ell : \Delta^n \rightarrow \mathbb{R}_+^n$. We will say a multi-CPE loss is *fair* whenever $\ell_i(e_i^n) = 0$, for all $i \in [n]$. That is, there is no loss incurred for perfect prediction. Examples of multi-CPE losses include the

square loss $\ell_i^{\text{sq}}(q) := \sum_{j \in [n]} (\mathbb{1}[i = j] - q_j)^2$, the log loss $\ell_i^{\text{log}}(q) := -\log q_i$, the absolute loss $\ell_i^{\text{abs}}(q) := \sum_{j \in [n]} |\mathbb{1}[i = j] - q_j|$, and the 0-1 loss $\ell_i^{01}(q) := \mathbb{1}[i \in \arg \max_{j \in [n]} q_j]$.

2.3. Conditional and Full Risks

Let \mathbf{X} and \mathbf{Y} be random variables defined on the instance space \mathcal{X} and the outcome space $\mathcal{Y} = [n]$ respectively. Let D be the joint distribution of (\mathbf{X}, \mathbf{Y}) and for $x \in \mathcal{X}$, denote the conditional distribution by $p(x) = (p_1(x), \dots, p_n(x))'$ where $p_i(x) := P(\mathbf{Y} = i | \mathbf{X} = x)$, $\forall i \in [n]$, and the marginal distribution by $M(x) := P(\mathbf{X} = x)$. For any multi-CPE loss ℓ , the *conditional risk* is defined as

$$L_\ell : \Delta^n \times \Delta^n \ni (p, q) \mapsto L_\ell(p, q) = \mathbb{E}_{\mathbf{Y} \sim p}[\ell_{\mathbf{Y}}(q)] = p' \cdot \ell(q) = \sum_{i \in [n]} p_i \ell_i(q) \in \mathbb{R}_+, \quad (1)$$

where $\mathbf{Y} \sim p$ represents a Multinomial distribution with parameter $p \in \Delta^n$. The *full risk* of the estimator function $q : \mathcal{X} \rightarrow \Delta^n$ is defined as

$$\mathbb{L}_\ell(p, q, M) := \mathbb{E}_{(\mathbf{X}, \mathbf{Y}) \sim D}[\ell_{\mathbf{Y}}(q(\mathbf{X}))] = \mathbb{E}_{\mathbf{X} \sim M}[L_\ell(p(\mathbf{X}), q(\mathbf{X}))].$$

Furthermore the *Bayes risk* is defined as

$$\underline{\mathbb{L}}_\ell(p, M) := \inf_{q \in (\Delta^n)^{\mathcal{X}}} \mathbb{L}_\ell(p, q, M) = \mathbb{E}_{\mathbf{X} \sim M}[\underline{L}_\ell(p(\mathbf{X}))],$$

where $\underline{L}_\ell(p) = \inf_{q \in \Delta^n} L_\ell(p, q)$ is the *conditional Bayes risk* and is always concave (Gneiting and Raftery (2007)). If ℓ is fair, $\underline{L}_\ell(e_i^n) = \ell_i(e_i^n) = 0$. One can understand the effect of choice of loss in terms of the conditional perspective (Reid and Williamson (2011)), which allows one to ignore the marginal distribution M of \mathbf{X} which is typically unknown.

2.4. Proper and Composite Losses

A multi-CPE loss $\ell : \Delta^n \rightarrow \mathbb{R}_+^n$ is said to be *proper* if for all $p \in \Delta^n$, $\underline{L}_\ell(p) = L_\ell(p, p) = p' \cdot \ell(p)$ (Buja et al. (2005), Gneiting and Raftery (2007)), and *strictly proper* if $\underline{L}_\ell(p) < L_\ell(p, q)$ for all $p, q \in \Delta^n$ and $p \neq q$. It is easy to see that the log loss, square loss, and 0-1 loss are proper while absolute loss is not. Furthermore, both log loss and square loss are strictly proper while 0-1 loss is proper but not strictly proper.

Given a proper loss $\ell : \Delta^n \rightarrow \mathbb{R}_+^n$ with differentiable Bayes conditional risk $\underline{L}_\ell : \Delta^n \mapsto \mathbb{R}_+$, in order to be able to calculate derivatives easily, following Van Erven et al. (2012) we define

$$\tilde{\Delta}^n := \left\{ (p_1, \dots, p_{n-1})' : p_i \geq 0, \sum_{i=1}^{n-1} p_i \leq 1 \right\} \quad (2)$$

$$\Pi_\Delta : \mathbb{R}_+^n \ni p = (p_1, \dots, p_n)' \mapsto \tilde{p} = (p_1, \dots, p_{n-1})' \in \mathbb{R}_+^{n-1} \quad (3)$$

$$\Pi_\Delta^{-1} : \tilde{\Delta}^n \ni \tilde{p} = (\tilde{p}_1, \dots, \tilde{p}_{n-1})' \mapsto p = (\tilde{p}_1, \dots, \tilde{p}_{n-1}, 1 - \sum_{i=1}^{n-1} \tilde{p}_i)' \in \Delta^n \quad (4)$$

$$\tilde{\underline{L}}_\ell : \tilde{\Delta}^n \ni \tilde{p} \mapsto \underline{L}_\ell(\Pi_\Delta^{-1}(\tilde{p})) \in \mathbb{R}_+ \quad (5)$$

$$\tilde{\ell} : \tilde{\Delta}^n \ni \tilde{p} \mapsto \Pi_\Delta(\ell(\Pi_\Delta^{-1}(\tilde{p}))) \in \mathbb{R}_+^{n-1}. \quad (6)$$

Let $\tilde{\psi} : \tilde{\Delta}^n \rightarrow \mathcal{V} \subseteq \mathbb{R}_+^{n-1}$ be continuous and strictly monotone (hence invertible) for some convex set \mathcal{V} . It induces $\psi : \Delta^n \rightarrow \mathcal{V}$ via

$$\psi := \tilde{\psi} \circ \Pi_{\Delta}. \quad (7)$$

Clearly ψ is continuous and invertible with $\psi^{-1} = \Pi_{\Delta}^{-1} \circ \tilde{\psi}^{-1}$. We can now extend the notion of properness to the prediction space \mathcal{V} from Δ^n using this link function. Given a proper loss $\ell : \Delta^n \rightarrow \mathbb{R}_+^n$, a *proper composite loss* $\ell^\psi : \mathcal{V} \rightarrow \mathbb{R}_+^n$ for multi-class probability estimation is defined as $\ell^\psi := \ell \circ \psi^{-1} = \ell \circ \Pi_{\Delta}^{-1} \circ \tilde{\psi}^{-1}$. We can easily see that the conditional Bayes risks of the composite loss ℓ^ψ and the underlying proper loss ℓ are equal ($\underline{L}_\ell = \underline{L}_{\ell^\psi}$). Every continuous proper loss has a convex super-prediction set (Vernet et al. (2014)). Thus they are weakly mixable. Since by applying a link function the super-prediction set won't change (as it is just a re-parameterization), all proper composite losses are also weakly mixable.

2.5. Game of Prediction with Expert Advice

Let \mathcal{Y} be the outcome space, \mathcal{V} be the prediction space, and $\ell : \mathcal{Y} \times \mathcal{V} \rightarrow \mathbb{R}_+$ be the loss function, then a *game of prediction with expert advice* represented by the tuple $(\mathcal{Y}, \mathcal{V}, \ell)$ can be described as follows: for each trial $t = 1, \dots, T$,

- N experts make their prediction $v_t^1, \dots, v_t^N \in \mathcal{V}$
- the learner makes his own decision $v_t \in \mathcal{V}$
- the environment reveals the actual outcome $y_t \in \mathcal{Y}$

Let $S = (y_1, \dots, y_T)$ be the outcome sequence in T trials. Then the *cumulative loss* of the learner over S is given by $L_{S,\ell} := \sum_{t=1}^T \ell(y_t, v_t)$, of the i -th expert is given by $L_{S,\ell}^i := \sum_{t=1}^T \ell(y_t, v_t^i)$, and the *regret* of the learner is given by $R_{S,\ell} := L_{S,\ell} - \min_i L_{S,\ell}^i$. The goal of the learner is to predict as well as the best expert; to which end the learner tries to minimize the regret.

When using the exponential weights algorithm (which is an important family of algorithms in game of prediction with expert advice), at the end of each trial, the weight of each expert is updated as $w_{t+1}^i \propto w_t^i \cdot e^{-\eta \ell(y_t, v_t^i)}$ for all $i \in [N]$, where η is the learning rate and w_t^i is the weight of the i^{th} expert at time t (the weight vector of experts at time t is denoted by $w_t = (w_t^1, \dots, w_t^N)'$). Then based on the weights of experts, their predictions are merged using different merging schemes to make the learner's prediction. The Aggregating Algorithm and the Weighted Average Algorithm are two important algorithms in the family of exponential weights algorithm.

Consider multi-class games with outcome space $\mathcal{Y} = [n]$. In the Aggregating Algorithm with learning rate β , first the loss vectors of the experts and their weights are used to make a *generalized prediction* $g_t = (g_t(1), \dots, g_t(n))'$ which is given by

$$g_t := E_\beta^{-1} \left(\text{co}_{w_t} \{ E_\beta \left((\ell_1(v_t^i), \dots, \ell_n(v_t^i))' \right) \}_{i \in [N]} \right) = E_\beta^{-1} \left(\sum_{i \in [N]} w_t^i (e^{-\beta \ell_1(v_t^i)}, \dots, e^{-\beta \ell_n(v_t^i)})' \right).$$

Then this generalized prediction is mapped into a permitted prediction v_t via a *substitution function* such that $(\ell_1(v_t), \dots, \ell_n(v_t))' \leq c(\beta)(g_t(1), \dots, g_t(n))'$, where the inequality is element-wise and the constant $c(\beta)$ depends on the learning rate. If ℓ is β -mixable, then $E_\beta(S_\ell)$ is convex, so

$\text{co}\{E_\beta(\ell(v)) : v \in \mathcal{V}\} \subseteq E_\beta(S_\ell)$, and we can always choose a substitution function with $c(\beta) = 1$. Consequently for β -mixable losses, the learner of the Aggregating Algorithm is guaranteed to have regret bounded by $\frac{\log N}{\beta}$ (Vovk (1995)).

In the Weighted Average Algorithm with learning rate α , the experts' predictions are simply merged according to their weights to make the learner's prediction $v_t = \text{co}_{w_t}\{v_t^i\}_{i \in [N]}$, and this algorithm is guaranteed to have a $\frac{\log N}{\alpha}$ regret bound for α -exp-concave losses (Kivinen and Warmuth (1999)). In either case it is preferred to have bigger values for the constants β and α to have better regret bounds. Since an α -exp-concave loss is β -mixable for some $\beta \geq \alpha$, the regret bound of the Weighted Average Algorithm is worse than that of the Aggregating Algorithm by a small constant factor. In Vovk (2001), it is noted that $(\ell_1(\text{co}_{w_t}\{v_t^i\}_i), \dots, \ell_n(\text{co}_{w_t}\{v_t^i\}_i))' \leq g_t = E_\beta^{-1}(\text{co}_{w_t}\{E_\beta((\ell_1(v_t^i), \dots, \ell_n(v_t^i))')\}_i)$ is always guaranteed only for β -exp-concave losses. Thus for α -exp-concave losses, the Weighted Average Algorithm is equivalent to the Aggregating Algorithm with the weighted average of the experts' predictions as its substitution function and α as the learning rate for both algorithms.

Even though the choice of substitution function will not have any impact on the regret bound and the weight update mechanism of the Aggregating Algorithm, it will certainly have impact on the actual regret of the learner over a given sequence of outcomes. According to the results given in Appendix A (where we have empirically compared some substitution functions), this impact on the actual regret varies depending on the outcome sequence, and in general the regret values for practical substitution functions don't differ much — thus we can stick with a computationally efficient substitution function.

3. Exp-Concavity of Proper Composite Losses

Exp-concavity of a loss is desirable for better (logarithmic) regret bounds in online convex optimization algorithms, and for efficient implementation of exponential weights algorithms. In this section we will consider whether one can always find a link function that can transform a β -mixable proper loss into β -exp-concave composite loss — first by using the geometry of the set $E_\beta(S_\ell)$ (Section 3.1), and then by using the characterization of the composite loss in terms of the associated Bayes risk (Sections 3.2, and 3.3).

3.1. Geometric approach

In this section we will use the same construction used by van Erven (2012) to derive an explicit closed form of a link function that could re-parameterize any β -mixable loss into a β -exp-concave loss, under certain conditions which are explained below. Given a multi-class loss $\ell : \mathcal{V} \rightarrow \mathbb{R}_+^n$, define

$$\ell(\mathcal{V}) := \{\ell(v) : v \in \mathcal{V}\}, \tag{8}$$

$$\mathcal{B}_\beta := \text{co}E_\beta(\ell(\mathcal{V})). \tag{9}$$

For any $g \in \mathcal{B}_\beta$ let $c(g) := \sup\{c \geq 0 : (g + c\mathbb{1}_n) \in \mathcal{B}_\beta\}$. Then the “north-east” boundary of the set \mathcal{B}_β is given by $\partial_{\mathbb{1}_n}\mathcal{B}_\beta := \{g + c(g) : g \in \mathcal{B}_\beta\}$. The following proposition is the main result of this section.

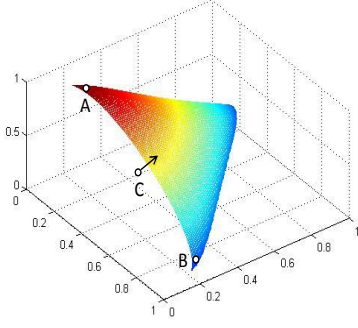


Figure 1: Ray “escaping” in $\mathbb{1}_n$ direction. More evidence in Figure 12 in Appendix D.

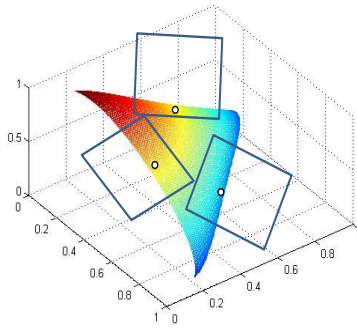


Figure 2: Adding “faces” to block rays in (almost) all positive directions.

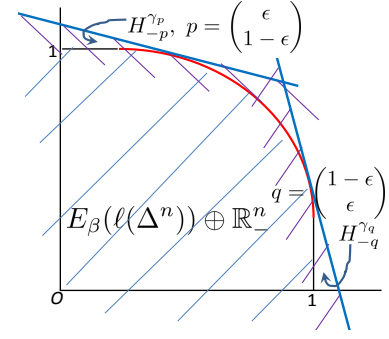


Figure 3: Sub-exp-prediction set extended by removing near axis-parallel supporting hyperplanes.

Proposition 1 *Assume ℓ is strictly proper and it satisfies the condition : $\partial_{\mathbb{1}_n} \mathcal{B}_\beta \subseteq E_\beta(\ell(\mathcal{V}))$ for some $\beta > 0$. Define $\psi(p) := JE_\beta(\ell(p))$ for all $p \in \Delta^n$, where $J = [I_{n-1}, -\mathbb{1}_{n-1}]$. Then ψ is invertible, and $\ell \circ \psi^{-1}$ is β -exp-concave over $\psi(\Delta^n)$, which is a convex set.*

The condition stated in the above proposition is satisfied by any β -mixable proper loss in the binary case ($n = 2$), but it is not guaranteed in the multi-class case where $n > 2$. In the binary case the link function can be given as $\psi(p) = e^{-\beta \ell_1(p)} - e^{-\beta \ell_2(p)}$ for all $p \in \Delta^2$.

Unfortunately, the condition that $\partial_{\mathbb{1}_n} \mathcal{B}_\beta \subseteq E_\beta(\ell(\mathcal{V}))$ is generally not satisfied; an example based on squared loss ($\beta = 1$ and $n = 3$ classes) is shown in Figure 1, where for A and B in $E_\beta(\ell(\mathcal{V}))$, the mid-point C can travel along the ray of direction $\mathbb{1}_3$ without hitting any point in the exp-prediction set $E_\beta(\ell(\mathcal{V}))$. Therefore we resort to *approximating* a given β -mixable loss by a sequence of β -exp-concave losses parameterised by positive constant ϵ , while the approximation approaches the original loss in some appropriate sense as ϵ tends to 0. Without loss of generality, we assume $\mathcal{V} = \Delta^n$.

Inspired by Proposition 1, a natural idea to construct the approximation is by adding “faces” to the exp-prediction set such that all rays in the $\mathbb{1}_n$ direction will be blocked. Technically, it turns out more convenient to add faces that block rays in (almost) all directions of positive orthant. See Figure 2 for an illustration. In particular, we extend the “rim” of the exp-prediction set by hyperplanes that are ϵ close to axis-parallel. The key challenge underlying this idea is to design an appropriate *parameterisation* of the surrogate loss $\tilde{\ell}_\epsilon$, which not only produces such an extended exp-prediction set, but also ensures that $\tilde{\ell}_\epsilon(p) = \ell(p)$ for almost all $p \in \Delta^n$ as $\epsilon \downarrow 0$.

Given a β -mixable loss ℓ , its sub-exp-prediction set defined as follows must be convex:

$$T_\ell := E_\beta(\ell(\Delta^n)) \oplus \mathbb{R}_+^n = \{E_\beta(\ell(p)) - x : p \in \Delta^n, x \in \mathbb{R}_+^n\}. \quad (10)$$

Note T_ℓ extends infinitely in any direction $p \notin \mathbb{R}_+^n$. Therefore it can be written in terms of supporting hyperplanes as

$$T_\ell = \bigcap_{p \in \Delta^n} H_{-p}^{\gamma_p}, \quad \text{where } \gamma_p = \min_{x \in T_\ell} x' \cdot (-p), \text{ and } H_{-p}^{\gamma_p} := \{x : x' \cdot (-p) \geq \gamma_p\}. \quad (11)$$

To extend the sub-exp-prediction set with “faces”, we remove some hyperplanes involved in (11) that correspond to the ϵ “rim” of the simplex (see Figure 3 for an illustration in 2-D)

$$T_\ell^\epsilon := \bigcap_{p \in \Delta_\epsilon^n} H_{-p}^{\gamma_p}, \quad \text{where } \Delta_\epsilon^n := \{p \in \Delta^n : \min_i p_i > \epsilon\}. \quad (12)$$

Since $\epsilon > 0$, for any $p \in \Delta_\epsilon^n$, $E_\beta^{-1}(H_{-p}^{\gamma_p} \cap \mathbb{R}_+^n)$ is exactly the super-prediction set of a log-loss with appropriate scaling and shifting (see proof in Appendix C). So it must be convex. Therefore $E_\beta^{-1}(T_\ell^\epsilon \cap \mathbb{R}_+^n) = \bigcap_{p \in \Delta_\epsilon^n} E_\beta^{-1}(H_{-p}^{\gamma_p} \cap \mathbb{R}_+^n)$ must be convex, and its recession cone is clearly \mathbb{R}_+^n . This guarantees that the following loss is proper over $p \in \Delta^n$ (Williamson, 2014, Proposition 2):

$$\tilde{\ell}_\epsilon(p) = \operatorname{argmin}_{z \in E_\beta^{-1}(T_\ell^\epsilon \cap \mathbb{R}_+^n)} p' \cdot z, \quad \text{where the argmin must be attained uniquely (Appendix C).} \quad (13)$$

Our next proposition states that $\tilde{\ell}_\epsilon$ meets all the requirements of approximation suggested above.

Proposition 2 *For any $\epsilon > 0$, $\tilde{\ell}_\epsilon$ must be β -exp-concave provided ℓ is β -mixable. In addition, $\tilde{\ell}_\epsilon = \ell$ over a subset $S_\epsilon \subseteq \Delta^n$, where for any p in the relative interior of Δ^n , $p \in S_\epsilon$ for sufficiently small ϵ .*

Note $\|\tilde{\ell}_\epsilon(p) - \ell(p)\|$ is not bounded for $p \notin S_\epsilon$. While the result does not show that all β -mixable losses can be made β -exp-concave, it is suggestive that such a result may be obtainable by a different argument.

3.2. Calculus approach

Proper composite losses are defined by the proper loss ℓ and the link ψ . In this section we will characterize the exp-concave proper composite losses in terms of $(\mathbf{H}\tilde{L}_\ell(\tilde{p}), \mathbf{D}\tilde{\psi}(\tilde{p}))$. The following proposition provides the identities of the first and second derivatives of the proper composite losses (Vernet et al. (2014)).

Proposition 3 *For all $i \in [n]$, $\tilde{p} \in \overset{\circ}{\Delta}^n$ (the interior of $\tilde{\Delta}^n$), and $v = \tilde{\psi}(\tilde{p}) \in \mathcal{V} \subseteq \mathbb{R}_+^{n-1}$ (so $\tilde{p} = \tilde{\psi}^{-1}(v)$),*

$$\mathbf{D}\ell_i^\psi(v) = -(e_i^{n-1} - \tilde{p})' \cdot k(\tilde{p}), \quad (14)$$

$$\mathbf{H}\ell_i^\psi(v) = -((e_i^{n-1} - \tilde{p})' \otimes I_{n-1}) \cdot \mathbf{D}_v[k(\tilde{p})] + k(\tilde{p})' \cdot [\mathbf{D}\tilde{\psi}(\tilde{p})]^{-1}, \quad (15)$$

where

$$k(\tilde{p}) := -\mathbf{H}\tilde{L}_\ell(\tilde{p}) \cdot [\mathbf{D}\tilde{\psi}(\tilde{p})]^{-1}. \quad (16)$$

The term $k(\tilde{p})$ can be interpreted as the curvature of the Bayes risk function \tilde{L}_ℓ relative to the rate of change of the link function $\tilde{\psi}$. In the binary case where $n = 2$, above proposition reduces to

$$(\ell_1^\psi)'(v) = -(1 - \tilde{p})k(\tilde{p}) \quad ; \quad (\ell_2^\psi)'(v) = \tilde{p}k(\tilde{p}), \quad (17)$$

$$(\ell_1^\psi)''(v) = \frac{-(1 - \tilde{p})k'(\tilde{p}) + k(\tilde{p})}{\tilde{\psi}'(\tilde{p})}, \quad (18)$$

$$(\ell_2^\psi)''(v) = \frac{\tilde{p}k'(\tilde{p}) + k(\tilde{p})}{\tilde{\psi}'(\tilde{p})}, \quad (19)$$

where $k(\tilde{p}) = \frac{-\tilde{L}''(\tilde{p})}{\tilde{\psi}'(\tilde{p})} \geq 0$ and so $\frac{d}{dv}k(\tilde{p}) = \frac{d}{d\tilde{p}}k(\tilde{p}) \cdot \frac{d}{dv}\tilde{p} = \frac{k'(\tilde{p})}{\tilde{\psi}'(\tilde{p})}$.

A loss $\ell : \Delta^n \rightarrow \mathbb{R}_+^n$ is α -exp-concave (i.e. $\Delta^n \ni q \mapsto \ell_y(q)$ is α -exp-concave for all $y \in [n]$) if and only if the map $\Delta^n \ni q \mapsto L_\ell(p, q) = p' \cdot \ell(q)$ is α -exp-concave for all $p \in \Delta^n$. It can be easily shown that the maps $v \mapsto \ell_y^\psi(v)$ are α -exp-concave if and only if $\mathbf{H}\ell_y^\psi(v) \succcurlyeq \alpha \mathbf{D}\ell_y^\psi(v)' \cdot \mathbf{D}\ell_y^\psi(v)$. By applying Proposition 3 we obtain the following characterization of the α -exp-concavity of the composite loss ℓ^ψ .

Proposition 4 *A proper composite loss $\ell^\psi = \ell \circ \psi^{-1}$ is α -exp-concave (with $\alpha > 0$ and $v = \tilde{\psi}(\tilde{p})$) if and only if for all $\tilde{p} \in \tilde{\Delta}^n$ and for all $i \in [n]$*

$$((e_i^{n-1} - \tilde{p})' \otimes I_{n-1}) \cdot \mathbf{D}_v[k(\tilde{p})] \preccurlyeq k(\tilde{p})' \cdot [\mathbf{D}\tilde{\psi}(\tilde{p})]^{-1} - \alpha k(\tilde{p})' \cdot (e_i^{n-1} - \tilde{p}) \cdot (e_i^{n-1} - \tilde{p})' \cdot k(\tilde{p}). \quad (20)$$

Based on this characterization, we can determine which loss functions can be exp-concavified by a chosen link function and how much a link function can exp-concavify a given loss function. In the binary case ($n = 2$), the above proposition reduces to the following.

Proposition 5 *Let $\tilde{\psi} : [0, 1] \rightarrow \mathcal{V} \subseteq \mathbb{R}$ be an invertible link and $\ell : \Delta^2 \rightarrow \mathbb{R}_+^2$ be a strictly proper binary loss with weight function $w(\tilde{p}) := -\mathbf{H}\tilde{L}_\ell(\tilde{p}) = -\tilde{L}_\ell''(\tilde{p})$. Then the binary composite loss $\ell^\psi := \ell \circ \Pi_\Delta^{-1} \circ \tilde{\psi}^{-1}$ is α -exp-concave (with $\alpha > 0$) if and only if*

$$-\frac{1}{\tilde{p}} + \alpha w(\tilde{p})\tilde{p} \leq \frac{w'(\tilde{p})}{w(\tilde{p})} - \frac{\tilde{\psi}''(\tilde{p})}{\tilde{\psi}'(\tilde{p})} \leq \frac{1}{1 - \tilde{p}} - \alpha w(\tilde{p})(1 - \tilde{p}), \quad \forall \tilde{p} \in (0, 1). \quad (21)$$

The following proposition gives an easier to check necessary condition for the binary proper losses that generate an α -exp-concave (with $\alpha > 0$) binary composite loss given a particular link function. Since scaling a loss function will not affect what a sensible learning algorithm will do, it is possible to normalize the loss functions by normalizing their weight functions by setting $w(\frac{1}{2}) = 1$. By this normalization we are scaling the original loss function by $\frac{1}{w(\frac{1}{2})}$ and the super-prediction set is scaled by the same factor. If the original loss function is β -mixable (resp. α -exp-concave), then the normalized loss function is $\beta w(\frac{1}{2})$ -mixable (resp. $\alpha w(\frac{1}{2})$ -exp-concave).

Proposition 6 *Let $\tilde{\psi} : [0, 1] \rightarrow \mathcal{V} \subseteq \mathbb{R}$ be an invertible link and $\ell : \Delta^2 \rightarrow \mathbb{R}_+^2$ be a strictly proper binary loss with weight function $w(\tilde{p}) := -\mathbf{H}\tilde{L}_\ell(\tilde{p}) = -\tilde{L}_\ell''(\tilde{p})$ normalised such that $w(\frac{1}{2}) = 1$. Then the binary composite loss $\ell^\psi := \ell \circ \Pi_\Delta^{-1} \circ \tilde{\psi}^{-1}$ is α -exp-concave (with $\alpha > 0$) only if*

$$\frac{\tilde{\psi}'(\tilde{p})}{\tilde{p}(2\tilde{\psi}'(\frac{1}{2}) - \alpha(\tilde{\psi}(\tilde{p}) - \tilde{\psi}(\frac{1}{2})))} \leq w(\tilde{p}) \leq \frac{\tilde{\psi}'(\tilde{p})}{(1 - \tilde{p})(2\tilde{\psi}'(\frac{1}{2}) + \alpha(\tilde{\psi}(\tilde{p}) - \tilde{\psi}(\frac{1}{2})))}, \quad \forall \tilde{p} \in (0, 1), \quad (22)$$

where \leq denotes \leq for $\tilde{p} \geq \frac{1}{2}$ and denotes \geq for $\tilde{p} \leq \frac{1}{2}$.

Proposition 5 provides necessary and sufficient conditions for the exp-concavity of binary composite losses, whereas Proposition 6 provides simple necessary but not sufficient conditions. By setting $\alpha = 0$ in all the above results we have obtained for exp-concavity, we recover the convexity conditions for proper and composite losses which are already derived by Reid and Williamson (2010) for the binary case and Vernet et al. (2014) for multi-class.

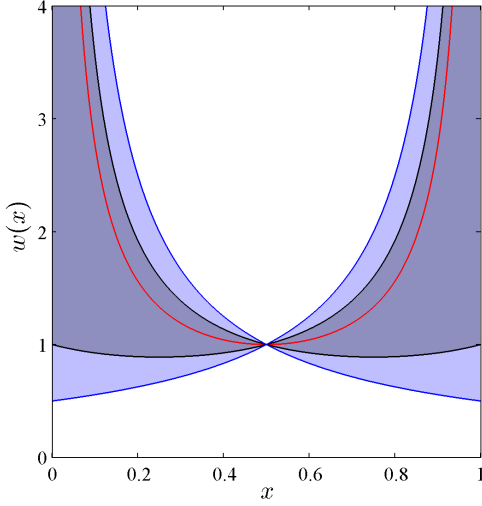


Figure 4: Necessary but not sufficient region of normalised weight functions to ensure α -exp-concavity and convexity of proper losses (— $\alpha = 4$; — $\alpha = 2$; — convexity).

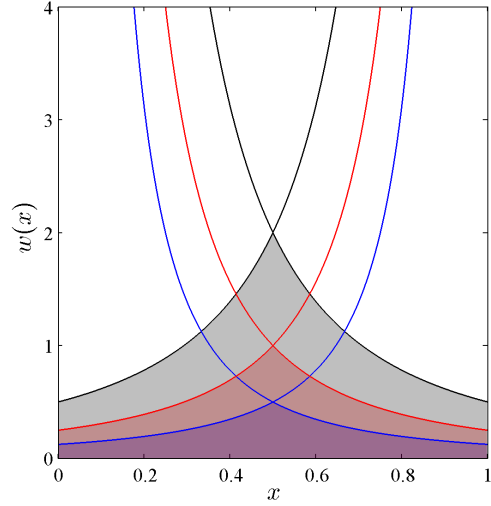


Figure 5: Necessary and sufficient region of unnormalised weight functions to ensure α -exp-concavity of composite losses with canonical link (— $\alpha = 2$; — $\alpha = 4$; — $\alpha = 8$).

3.3. Link functions

A proper loss can be exp-concavified ($\alpha > 0$) by some link function only if the loss is mixable ($\beta_\ell > 0$) and the maximum possible value for exp-concavity constant is the mixability constant of the loss (since the link function won't change the super-prediction set and an α -exp-concave loss is always β -mixable for some $\beta \geq \alpha$).

By applying the *identity link* $\tilde{\psi}(\tilde{p}) = \tilde{p}$ in (21) we obtain the necessary and sufficient conditions for a binary proper loss to be α -exp-concave (with $\alpha > 0$) as given by,

$$-\frac{1}{\tilde{p}} + \alpha w(\tilde{p})\tilde{p} \leq \frac{w'(\tilde{p})}{w(\tilde{p})} \leq \frac{1}{1-\tilde{p}} - \alpha w(\tilde{p})(1-\tilde{p}), \quad \forall \tilde{p} \in (0, 1). \quad (23)$$

By substituting $\tilde{\psi}(\tilde{p}) = \tilde{p}$ in (22) we obtain the following necessary but not sufficient (simpler) constraints for a normalized binary proper loss to be α -exp-concave

$$\frac{1}{\tilde{p}(2 - \alpha(\tilde{p} - \frac{1}{2}))} \stackrel{\leq}{\geq} w(\tilde{p}) \stackrel{\leq}{\geq} \frac{1}{(1-\tilde{p})(2 + \alpha(\tilde{p} - \frac{1}{2}))}, \quad \forall \tilde{p} \in (0, 1), \quad (24)$$

which are illustrated as the shaded region in Figure 4 for different values of α . Observe that normalized proper losses can be α -exp-concave only for $0 < \alpha \leq 4$. When $\alpha = 4$, only the normalized weight function of log loss ($w_{\log}(\tilde{p}) = \frac{1}{4\tilde{p}(1-\tilde{p})}$) will satisfy (24), and when $\alpha > 4$, the allowable (necessary) $w(\tilde{p})$ region to ensure α -exp-concavity vanishes. Thus normalized log loss is the most exp-concave normalized proper loss. Observe (from Figure 4) that normalized square loss

($w_{\ell^{\text{sq}}}(\tilde{p}) = 1$) is at most 2-exp-concave. Further from (24), if $\alpha' > \alpha$, then the allowable $w(\tilde{p})$ region to ensure α' -exp-concavity will be within the region for α -exp-concavity, and also any allowable $w(\tilde{p})$ region to ensure α -exp-concavity will be within the region for convexity, which is obtained by setting $\alpha = 0$ in (24). Here we recall the fact that, if the normalized loss function is α -exp-concave, then the original loss function is $\frac{\alpha}{w(\frac{1}{2})}$ -exp-concave. The following theorem provides sufficient conditions for the exp-concavity of binary proper losses.

Theorem 7 *A binary proper loss $\ell : \Delta^2 \rightarrow \mathbb{R}_+^2$ with the weight function $w(\tilde{p}) = -\tilde{L}_\ell''(\tilde{p})$ normalized such that $w(\frac{1}{2}) = 1$ is α -exp-concave (with $\alpha > 0$) if*

$$w(\tilde{p}) = \frac{1}{\tilde{p} \left(2 - \int_{\tilde{p}}^{1/2} a(t) dt\right)} \text{ for } a(\tilde{p}) \text{ s.t. } -\alpha + \frac{\alpha}{2\tilde{p}^2} - \frac{2}{\tilde{p}^2} \leq a(\tilde{p}) \leq -\alpha, \forall \tilde{p} \in (0, 1/2],$$

and

$$w(\tilde{p}) = \frac{1}{(1 - \tilde{p}) \left(2 - \int_{\frac{1}{2}}^{\tilde{p}} b(t) dt\right)} \text{ for } b(\tilde{p}) \text{ s.t. } \frac{\alpha\tilde{p}}{(1 - \tilde{p})} + \frac{2\alpha\tilde{p} - \alpha - 4}{2(1 - \tilde{p})^2} \leq b(\tilde{p}) \leq -\alpha, \forall \tilde{p} \in [\frac{1}{2}, 1).$$

For square loss we can find that $a(\tilde{p}) = \frac{-1}{\tilde{p}^2}$ and $b(\tilde{p}) = \frac{-1}{(1 - \tilde{p})^2}$ will satisfy the above sufficient condition with $\alpha = 2$ and for log loss $a(\tilde{p}) = b(\tilde{p}) = -4$ will satisfy the sufficient condition with $\alpha = 4$. It is also easy to see that for symmetric losses $a(\tilde{p})$ and $b(\tilde{p})$ will be symmetric.

When the canonical link function $\tilde{\psi}_\ell(\tilde{p}) := -\mathbf{D}\tilde{L}_\ell(\tilde{p})'$ is combined with a strictly proper loss to form ℓ^{ψ_ℓ} , since $\mathbf{D}\tilde{\psi}_\ell(\tilde{p}) = -\mathbf{H}\tilde{L}_\ell(\tilde{p})$, the first and second derivatives of the composite loss become considerably simpler as follows

$$\mathbf{D}\ell_i^{\psi_\ell}(v) = -(e_i^{n-1} - \tilde{p})', \quad (25)$$

$$\mathbf{H}\ell_i^{\psi_\ell}(v) = -[\mathbf{H}\tilde{L}_\ell(\tilde{p})]^{-1}. \quad (26)$$

Since a proper loss ℓ is β -mixable if and only if $\beta\mathbf{H}\tilde{L}_\ell(\tilde{p}) \succcurlyeq \mathbf{H}\tilde{L}_{\ell^{\log}}(\tilde{p})$ for all $\tilde{p} \in \overset{\circ}{\Delta}^n$ (Van Erven et al. (2012)), by applying the canonical link any β -mixable proper loss will be transformed to α -exp-concave proper composite loss (with $\beta \geq \alpha > 0$) but $\alpha = \beta$ is not guaranteed in general. In the binary case, since $\tilde{\psi}'_\ell(\tilde{p}) = -\tilde{L}_\ell''(\tilde{p}) = w(\tilde{p})$, we get

$$w(\tilde{p}) \leq \frac{1}{\alpha\tilde{p}^2} \quad \text{and} \quad w(\tilde{p}) \leq \frac{1}{\alpha(1 - \tilde{p})^2}, \quad \forall \tilde{p} \in (0, 1), \quad (27)$$

as the necessary and sufficient conditions for ℓ^{ψ_ℓ} to be α -exp-concave. In this case when $\alpha \rightarrow \infty$ the allowed region vanishes (since for proper losses $w(\tilde{p}) \geq 0$). From Figure 5 it can be seen that, if the normalized loss function satisfies

$$w(\tilde{p}) \leq \frac{1}{4\tilde{p}^2} \quad \text{and} \quad w(\tilde{p}) \leq \frac{1}{4(1 - \tilde{p})^2}, \quad \forall \tilde{p} \in (0, 1),$$

then the composite loss obtained by applying the canonical link function on the unnormalized loss with weight function $w_{\text{org}}(\tilde{p})$ is $\frac{4}{w_{\text{org}}(\frac{1}{2})}$ -exp-concave.

We now consider whether one can always find a link function that can transform a β -mixable proper loss into β -exp-concave composite loss. In the binary case, such a link function exists and is given in the following corollary.

Corollary 8 Let $w_\ell(\tilde{p}) = -\tilde{L}_\ell''(\tilde{p})$. The exp-concavifying link function $\tilde{\psi}_\ell^*$ defined via

$$\tilde{\psi}_\ell^*(\tilde{p}) = \frac{w_{\ell \log}(\frac{1}{2})}{w_\ell(\frac{1}{2})} \int_0^{\tilde{p}} \frac{w_\ell(v)}{w_{\ell \log}(v)} dv, \quad \forall \tilde{p} \in [0, 1] \quad (28)$$

(which is a valid strictly increasing link function) will always transform a β -mixable proper loss ℓ into β -exp-concave composite loss $\ell^{\psi_\ell^*}$, where $\ell_y^{\psi_\ell^*}(v) = \ell_y \circ \Pi_\Delta^{-1} \circ (\tilde{\psi}_\ell^*)^{-1}(v)$.

For log loss, the exp-concavifying link is equal to the identity link and the canonical link could be written as $\int_0^{\tilde{p}} w_\ell(v) dv$. If ℓ is a binary proper loss with weight function $w_\ell(\tilde{p})$, then we can define a new proper loss ℓ^{mix} with weight function $w_{\ell^{\text{mix}}}(\tilde{p}) = \frac{w_\ell(\tilde{p})}{w_{\ell \log}(\tilde{p})}$. Then applying the exp-concavifying link $\tilde{\psi}_\ell^*$ on the original loss ℓ is equivalent to applying the canonical link $\tilde{\psi}_\ell$ on the new loss ℓ^{mix} .

The links constructed by the geometric and calculus approaches can be completely different (see Appendix D-F). The former can be further varied by replacing $\mathbf{1}_n$ with any direction in the positive orthant, and the latter can be arbitrarily rescaled. Furthermore, as both links satisfy (21) with $\alpha = \beta$, any appropriate interpolation also works.

4. Conclusions

If a loss is β -mixable, one can run the Aggregating Algorithm with learning rate β and obtain a $\frac{\log N}{\beta}$ regret bound. Similarly a $\frac{\log N}{\alpha}$ regret bound can be attained by the Weighted Average Algorithm with learning rate α , when the loss is α -exp-concave. Vovk (2001) observed that the weighted average of the expert predictions (Kivinen and Warmuth (1999)) will be a *perfect* (in the technical sense defined in Vovk (2001)) substitution function for the Aggregating Algorithm if and only if the loss function is exp-concave. Thus if we have to use a proper, mixable but non-exp-concave loss function ℓ for a sequential prediction (online learning) problem, an $O(1)$ regret bound could be achieved by the following two approaches:

- Use the Aggregating Algorithm (Vovk (1995)) with the *inverse loss* ℓ^{-1} (Williamson (2014)) as the universal substitution function.
- Apply the exp-concavifying link ($\tilde{\psi}_\ell^*$) on ℓ , derive the β_ℓ -exp-concave composite loss $\ell^{\psi_\ell^*}$. Then use the Weighted Average Algorithm (Kivinen and Warmuth (1999)) with $\ell^{\psi_\ell^*}$ to obtain the learner's prediction in the transformed domain ($v_{\text{avg}} \in \psi_\ell^*(\tilde{\Delta}^n)$). Finally output the inverse link value of this prediction ($(\psi_\ell^*)^{-1}(v_{\text{avg}})$).

In either approach we are faced with a computational problem of evaluating an inverse function. But in the binary class case the inverse of a strictly monotone function can be efficiently evaluated using one sided bisection method (or lookup table). So in conclusion, the latter approach can be more convenient and efficient in computation than the former.

Acknowledgments

This work was supported by the ARC and NICTA, which is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program. Thanks to the referees for helpful comments.

References

- Arindam Banerjee, Xin Guo, and Hui Wang. On the optimality of conditional expectation as a Bregman predictor. *IEEE Transactions on Information Theory*, 51(7):2664–2669, 2005.
- Andreas Buja, Werner Stuetzle, and Yi Shen. Loss functions for binary class probability estimation and classification: Structure and applications, manuscript, available at www-stat.wharton.upenn.edu/~buja, 2005.
- Sever Silvestru Dragomir. Some Gronwall type inequalities and applications. *RGMIA Monographs, Victoria University, Australia*, 19, 2000.
- Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- David J. Hand. Deconstructing statistical questions. *Journal of the Royal Statistical Society, Series A (Statistics in Society)*, 157(3):317–356, 1994.
- David J. Hand and Veronica Vinciotti. Local versus global models for classification problems: Fitting models where it matters. *The American Statistician*, 57(2):124–131, 2003.
- David Haussler, Jyrki Kivinen, and Manfred K. Warmuth. Sequential prediction of individual sequences under general loss functions. *IEEE Transactions on Information Theory*, 44(5):1906–1925, 1998.
- Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.
- Yuri Kalnishkan and Michael V. Vyugin. The weak aggregating algorithm and weak mixability. In *Proc. Annual Conf. Computational Learning Theory*, pages 188–203. Springer, 2005.
- Jyrki Kivinen and Manfred K. Warmuth. Averaging expert predictions. In *Proc. Annual Conf. Computational Learning Theory*, pages 153–167. Springer, 1999.
- Mark D. Reid and Robert C. Williamson. Composite binary losses. *Journal of Machine Learning Research*, 11:2387–2422, 2010.
- Mark D. Reid and Robert C. Williamson. Information, divergence and risk for binary experiments. *Journal of Machine Learning Research*, 12:731–817, 2011.
- R. Tyrrell Rockafellar. *Convex analysis*. Princeton university press, 1970.
- Tim van Erven. From exp-concavity to mixability, 2012. <http://www.timvanerven.nl/blog/2012/12/from-exp-concavity-to-mixability/>.
- 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(1):1639–1663, 2012.
- Elodie Vernet, Mark D. Reid, and Robert C. Williamson. Composite multi-class losses. *Journal of Machine Learning Research*, 2014. Under review. <http://users.cecs.anu.edu.au/~williams/papers/P189.pdf>.

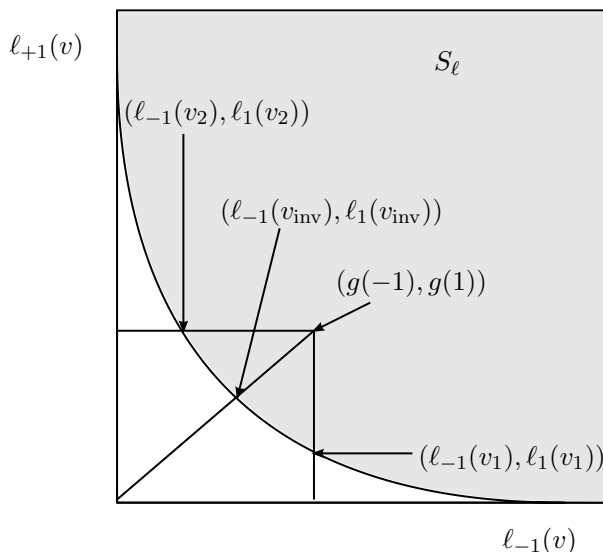


Figure 6: Super-prediction set (S_ℓ) of a binary game $(\{-1, 1\}, [0, 1], \ell)$, where a generalized prediction is represented by $(g(-1), g(1))$.

Volodya Vovk. A game of prediction with expert advice. In *Proc. Annual Conf. Computational Learning Theory*, pages 51–60. ACM, 1995.

Volodya Vovk. Competitive on-line statistics. *International Statistical Review*, 69(2):213–248, 2001.

Volodya Vovk and Fedor Zhdanov. Prediction with expert advice for the Brier game. *Journal of Machine Learning Research*, 10:2445–2471, 2009.

Robert C. Williamson. The geometry of losses. In *Proc. Annual Conf. Computational Learning Theory*, pages 1078–1108, 2014.

Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the International Conference on Machine Learning*, pages 928–936, 2003.

Appendix A. Substitution Functions

We consider the following choices of substitution functions — *Best look ahead*, *Worst look ahead*, *Inverse loss* (Williamson (2014)) and *Weighted average* (Kivinen and Warmuth (1999)). The first two choices are just hypothetical ones as they look ahead the actual outcome first and then choose their predictions to incur low and high loss respectively, still they give us an idea about how well or worse the Aggregating Algorithm performs over a given outcome sequence. The weighted average is a computationally efficient and easily implementable substitution function, but it is applicable only for exp-concave losses.

For a binary game represented by $(\{-1, 1\}, [0, 1], \ell)$ and shown in the Figure 6,

- If the outcome $y = -1$, the Best look ahead and the Worst look ahead will choose the predictions v_2 and v_1 and incur losses $\ell_{-1}(v_2)$ and $\ell_{-1}(v_1) = g(-1)$ respectively; and if $y = 1$, they will choose v_1 and v_2 and suffer losses $\ell_1(v_1)$ and $\ell_1(v_2) = g(1)$ respectively.
- The Inverse loss will choose the prediction v_{inv} such that $\frac{\ell_1(v_{\text{inv}})}{\ell_{-1}(v_{\text{inv}})} = \frac{g(1)}{g(-1)}$ and will incur a loss $\ell(y, v_{\text{inv}})$, and the Weighted average will choose $v_{\text{avg}} = \sum_i w^i v^i$ (where w^i and v^i are the weight and the prediction of the i -the expert respectively) and will incur a loss $\ell(y, v_{\text{avg}})$.

Further if the loss function ℓ is chosen to be the square loss (which is both 2-mixable and $\frac{1}{2}$ -exp-concave), then we have $v_1 = \sqrt{g(-1)}$ (since $\ell_{-1}(v_1) = v_1^2 = g(-1)$), $v_2 = 1 - \sqrt{g(1)}$ (since $\ell_1(v_2) = (1 - v_2)^2 = g(1)$), and $v_{\text{inv}} = \frac{\sqrt{g(-1)}}{\sqrt{g(-1)} + \sqrt{g(1)}}$ (since $\frac{(1 - v_{\text{inv}})^2}{v_{\text{inv}}^2} = \frac{g(1)}{g(-1)}$). Thus for a binary square loss game over an outcome sequence y_1, \dots, y_T , the cumulative losses of the Aggregating Algorithm for different choices of substitution function are given as follows:

- Best look ahead: $\sum_1^T \left(1 - \sqrt{g_t(-y_t)}\right)^2$
- Worst look ahead: $\sum_1^T g_t(y_t)$
- Inverse loss: $\sum_1^T \left(y_t - \frac{\sqrt{g_t(0)}}{\sqrt{g_t(0)} + \sqrt{g_t(1)}}\right)^2$
- Weighted average: $\sum_1^T \left(y_t - \sum_i w_i^t v_i^t\right)^2$

Some experiments are conducted on a binary square loss game to compare these substitution functions. For this, binary outcome sequences of 100 elements are generated using the Bernoulli distribution with success probabilities 0.5, 0.7, 0.9, and 1.0 (these sequences are represented by $\{y_t\}_{p=0.5}$, $\{y_t\}_{p=0.7}$, $\{y_t\}_{p=0.9}$ and $\{y_t\}_{p=1.0}$ respectively). Furthermore the following expert settings are used:

- 2 experts where one expert always make the prediction $v = 0$, and the other one always makes the prediction $v = 1$. This setting is represented by $\{E_t\}_{\text{set.1}}$.
- 3 experts where two experts are as in the previous setting, and the other one is always accurate expert. This setting is represented by $\{E_t\}_{\text{set.2}}$.
- 101 constant experts where the prediction values of the experts are from 0 to 1 with equal interval. This setting is represented by $\{E_t\}_{\text{set.3}}$.

The results of these experiments are presented in the figures 7,8,9, and 10. From these figures, it can be seen that for the expert setting $\{E_t\}_{\text{set.1}}$, the difference between the regret values of the worst look ahead and the best look ahead substitution functions relative to the theoretical regret bound is very high, whereas that relative difference is very low for the expert setting $\{E_t\}_{\text{set.3}}$. Further the performance of the Aggregating Algorithm over a real dataset is shown in the Figure 11. From

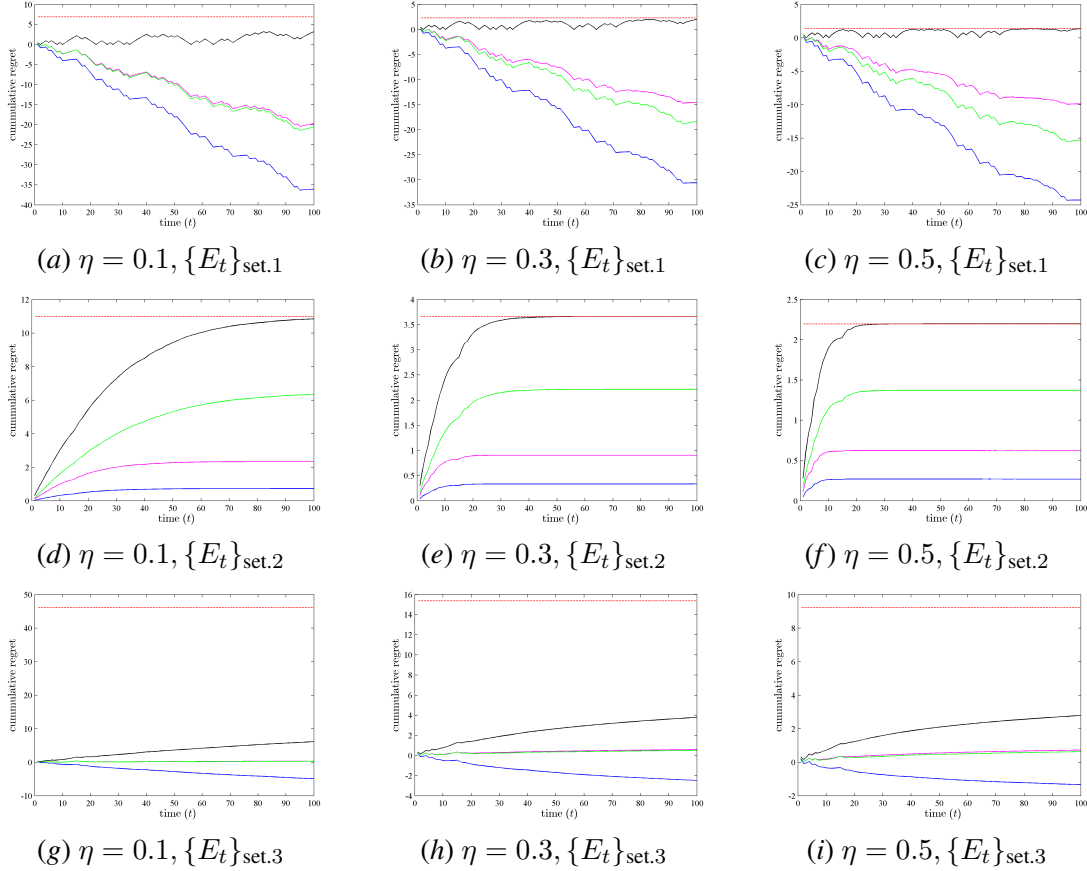


Figure 7: Cumulative regret of the Aggregating Algorithm over the outcome sequence $\{y_t\}_{p=0.5}$ for different choices of substitution functions (Best look ahead(—), Worst look ahead(—), Inverse loss(—), and Weighted average(—)) with learning rate η and expert setting $\{E_t\}_i$ (theoretical regret bound is shown by - - -).

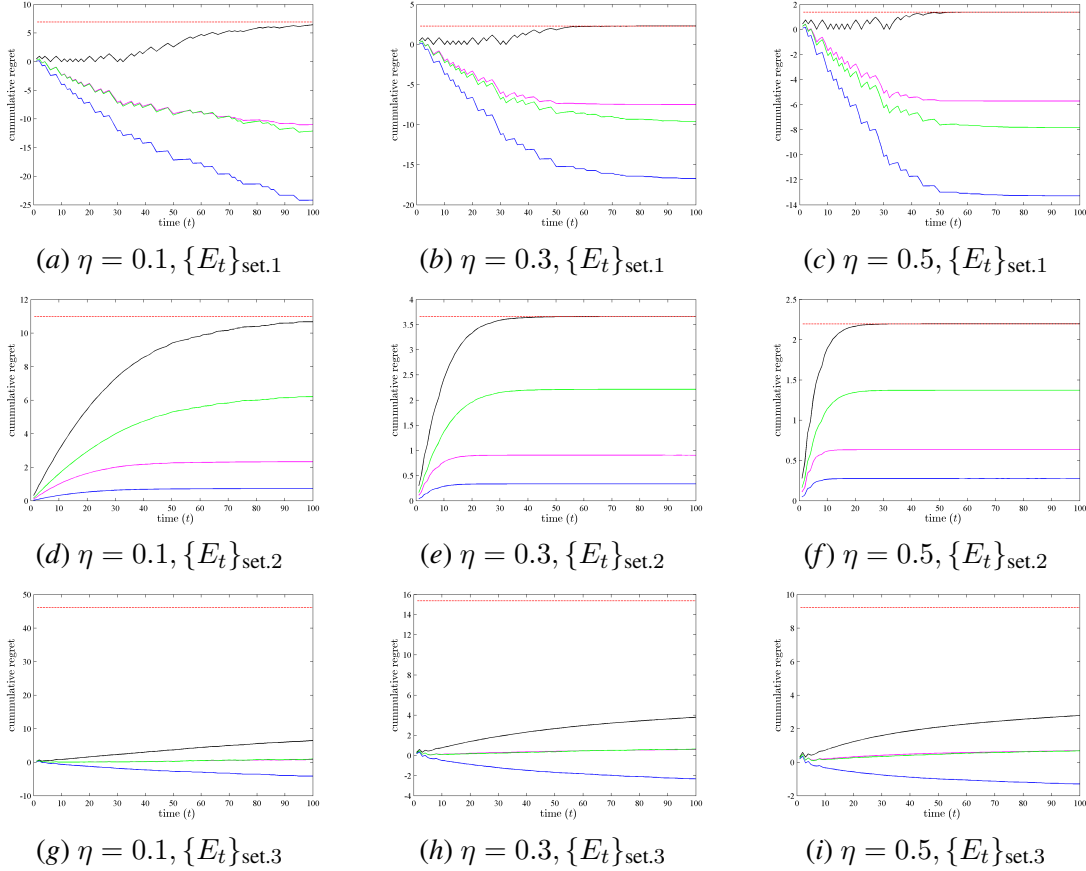


Figure 8: Cumulative regret of the Aggregating Algorithm over the outcome sequence $\{y_t\}_{p=0.7}$ for different choices of substitution functions (Best look ahead(—), Worst look ahead(—), Inverse loss(—), and Weighted average(—)) with learning rate η and expert setting $\{E_t\}_i$ (theoretical regret bound is shown by - - -).

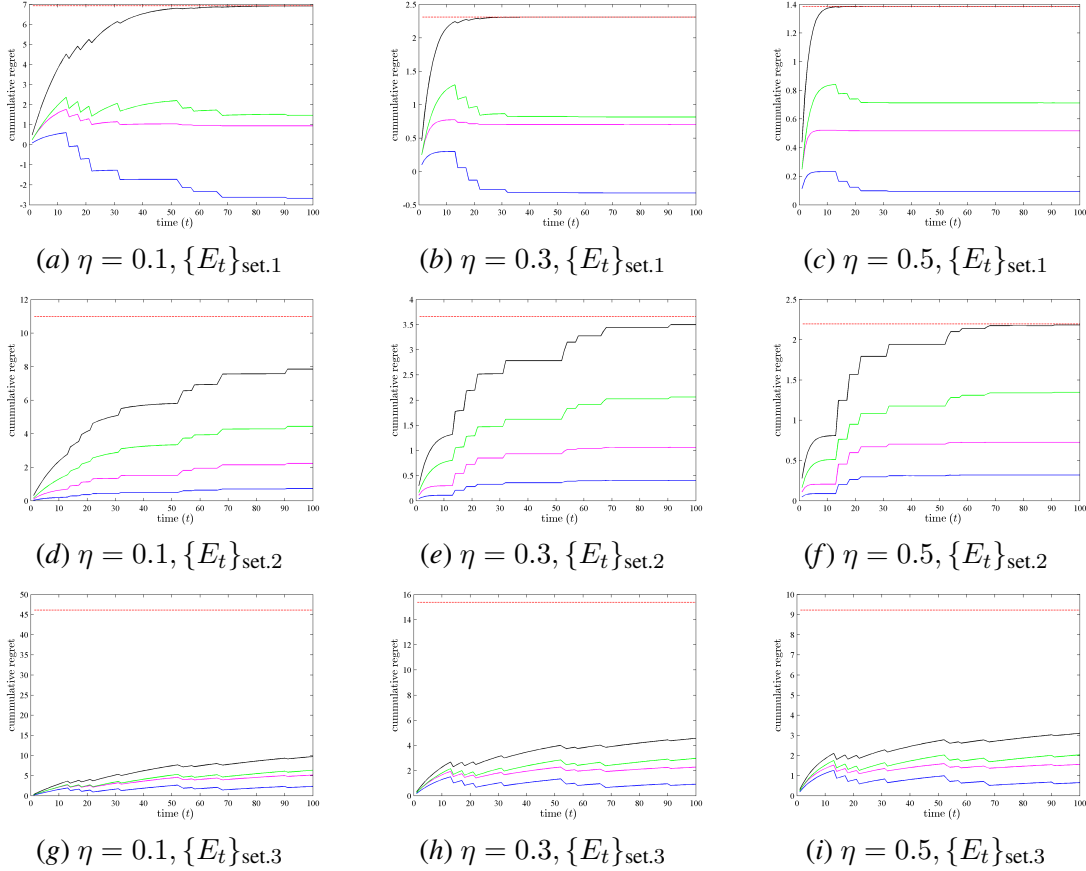


Figure 9: Cumulative regret of the Aggregating Algorithm over the outcome sequence $\{y_t\}_{p=0.9}$ for different choices of substitution functions (Best look ahead(—), Worst look ahead(—), Inverse loss(—), and Weighted average(—)) with learning rate η and expert setting $\{E_t\}_i$ (theoretical regret bound is shown by - - -).

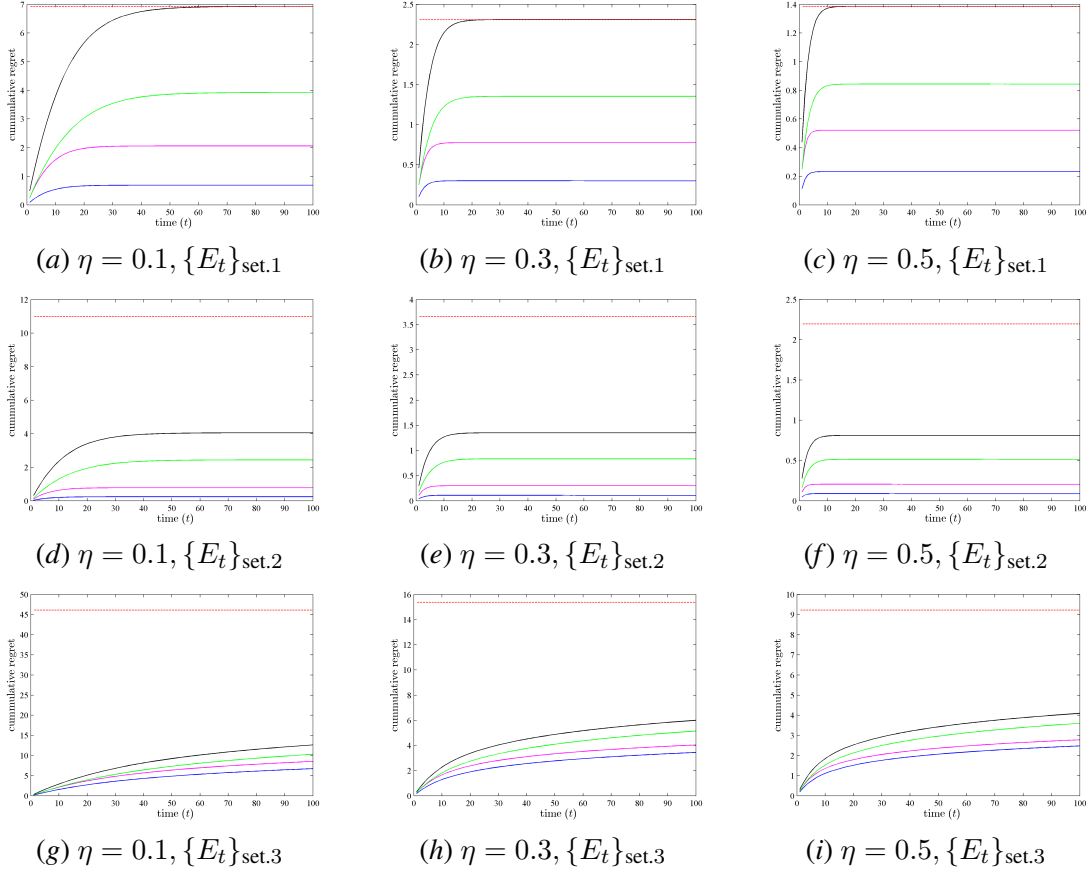


Figure 10: Cumulative regret of the Aggregating Algorithm over the outcome sequence $\{y_t\}_{p=1.0}$ for different choices of substitution functions (Best look ahead(—), Worst look ahead(—), Inverse loss(—), and Weighted average(—)) with learning rate η and expert setting $\{E_t\}_i$ (theoretical regret bound is shown by - - -).

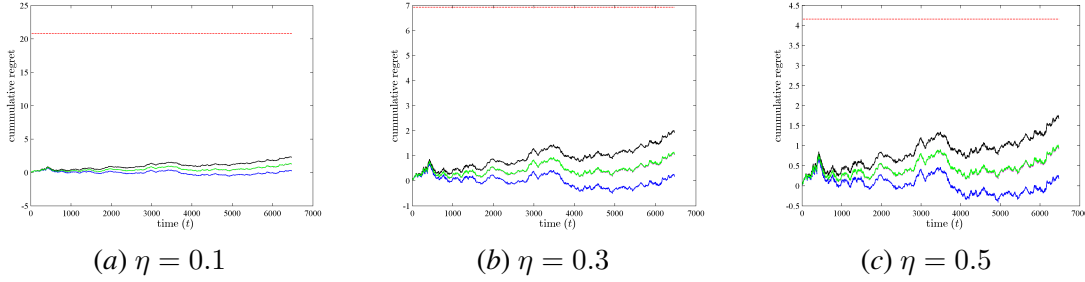


Figure 11: Cumulative regret of the Aggregating Algorithm over the football dataset as used by [Vovk and Zhdanov \(2009\)](#), for different choices of substitution functions (Best look ahead(—), Worst look ahead(—), Inverse loss(—), and Weighted average(—)) with learning rate η (theoretical regret bound is shown by - - -).

these results for both simulated dataset (for all three expert settings) and real dataset, observe that the difference between the regret values of the inverse loss and the weighted average substitution functions relative to the theoretical regret bound is very low.

Appendix B. Probability Games with Continuous outcome space

We consider an important class of prediction problem called *probability games* (as explained by [Vovk \(2001\)](#)), in which the prediction v and the outcome y are probability distributions in some set (for example a finite set of the form $[n]$). A special class of loss functions called *Bregman Loss Functions* (defined below) would be an appropriate choice for the probability games.

Given a differentiable convex function $\phi : \mathcal{S} \rightarrow \mathbb{R}$ defined on a convex set $\mathcal{S} \subset \mathbb{R}^d$ and two points $s_0, s \in \mathcal{S}$ the *Bregman divergence* of s from s_0 is defined as

$$B_\phi(s, s_0) := \phi(s) - \phi(s_0) - (s - s_0)' \cdot \mathbf{D}\phi(s_0),$$

where $\mathbf{D}\phi(s_0)$ is the gradient of ϕ at s_0 . For any strictly convex function $\phi : \tilde{\Delta}^n \rightarrow \mathbb{R}$, differentiable over the interior of $\tilde{\Delta}^n$, the *Bregman Loss Function* (BLF, [Banerjee et al. \(2005\)](#)) $\ell_\phi : \Delta^n \times \Delta^n \rightarrow \mathbb{R}_+$ with generator ϕ is given by

$$\ell_\phi(y, v) := B_\phi(\tilde{y}, \tilde{v}) = \phi(\tilde{y}) - \phi(\tilde{v}) - (\tilde{y} - \tilde{v})' \cdot \mathbf{D}\phi(\tilde{v}); \quad y, v \in \Delta^n, \quad (29)$$

where $\tilde{y} = \Pi_\Delta(y)$, and $\tilde{v} = \Pi_\Delta(v)$. Since the conditional Bayes risk of a strictly proper loss is strictly concave, any differentiable strictly proper loss $\ell : \Delta^n \rightarrow \mathbb{R}_+$ will generate a BLF ℓ_ϕ with generator $\phi = -\tilde{\mathcal{L}}_\ell$. Further if ℓ is fair, $\ell_i(v) = \ell_\phi(e_i^n, v)$; i.e. reconstruction is possible. For example the *Kullback-Leibler loss* given by $\ell_{KL}(y, v) := \sum_{i=1}^n y(i) \log \frac{y(i)}{v(i)}$, is a BLF generated by the log loss which is strictly proper.

The following lemma (multi-class extension of a result given by [Haussler et al. \(1998\)](#)) provides the mixability condition for probability games.

Lemma 9 For given $\ell : \Delta^n \times \Delta^n \rightarrow \mathbb{R}_+$, assume that for all $\tilde{y}, \tilde{v}_1, \tilde{v}_2 \in \tilde{\Delta}^n$ (let $y = \Pi_{\Delta}^{-1}(\tilde{y})$, $v_1 = \Pi_{\Delta}^{-1}(\tilde{v}_1)$, and $v_2 = \Pi_{\Delta}^{-1}(\tilde{v}_2)$), the function g defined by

$$g(\tilde{y}, \tilde{v}_1, \tilde{v}_2) = \frac{\beta}{c(\beta)} \ell(y, v_1) - \beta \ell(y, v_2) \quad (30)$$

satisfies

$$\mathbf{H}_{\tilde{y}}g(\tilde{y}, \tilde{v}_1, \tilde{v}_2) + \mathbf{D}_{\tilde{y}}g(\tilde{y}, \tilde{v}_1, \tilde{v}_2) \cdot (\mathbf{D}_{\tilde{y}}g(\tilde{y}, \tilde{v}_1, \tilde{v}_2))' \succcurlyeq 0. \quad (31)$$

If

$$\exists \tilde{v}^* \in \tilde{\Delta}^n \text{ s.t. } \ell(y, v^*) \leq -\frac{c(\beta)}{\beta} \log \int e^{-\beta \ell(y, v)} P(d\tilde{v}) \quad (32)$$

holds for the vertices $\tilde{y} = e_i^{n-1}$, $i \in [n]$, then it holds for all values $\tilde{y} \in \tilde{\Delta}^n$ (where $y = \Pi_{\Delta}^{-1}(\tilde{y})$, $v^* = \Pi_{\Delta}^{-1}(\tilde{v}^*)$ and $v = \Pi_{\Delta}^{-1}(\tilde{v})$).

Proof From (32)

$$\frac{\beta}{c(\beta)} \ell(y, v^*) + \log \int e^{-\beta \ell(y, v)} P(d\tilde{v}) \leq 0.$$

By exponentiating both sides we get

$$e^{\frac{\beta}{c(\beta)} \ell(y, v^*)} \cdot \int e^{-\beta \ell(y, v)} P(d\tilde{v}) \leq 1.$$

Denoting the left hand side of the above inequality by $f(\tilde{y})$ we have

$$f(\tilde{y}) = \int e^{g(\tilde{y}, \tilde{v}^*, \tilde{v})} P(d\tilde{v}).$$

Since the Hessian of f w.r.t. \tilde{y} given by

$$\mathbf{H}_{\tilde{y}}f(\tilde{y}) = \int e^{g(\tilde{y}, \tilde{v}^*, \tilde{v})} (\mathbf{H}_{\tilde{y}}g(\tilde{y}, \tilde{v}^*, \tilde{v}) + \mathbf{D}_{\tilde{y}}g(\tilde{y}, \tilde{v}^*, \tilde{v}) \cdot (\mathbf{D}_{\tilde{y}}g(\tilde{y}, \tilde{v}^*, \tilde{v}))') P(d\tilde{v})$$

is positive semi-definite (by (31)), $f(\tilde{y})$ is convex in \tilde{y} . So the maximum values of f for $\tilde{y} \in \tilde{\Delta}^n$ occurs for some $\tilde{y} = e_i^{n-1}$, $i \in [n]$. And by noting that, (32) is equivalent to $f(\tilde{y}) \leq 1$ for $\tilde{y} = e_i^{n-1}$, $i \in [n]$, the proof is completed. \blacksquare

The next proposition shows that the mixability and exp-concavity of a strictly proper loss is carried over to the BLF generated by it.

Proposition 10 For a strictly proper fair loss $\ell : \Delta^n \rightarrow \mathbb{R}_+$, and the BLF $\ell_{\phi} : \Delta^n \times \Delta^n \rightarrow \mathbb{R}_+$ generated by ℓ with $\phi = -\tilde{\underline{L}}_{\ell}$, if ℓ is β -mixable (resp. α -exp-concave), then ℓ_{ϕ} is also β -mixable (resp. α -exp-concave).

Proof From (30) and (29), for the BLF ℓ_{ϕ} we have

$$\begin{aligned} g(\tilde{y}, \tilde{v}_1, \tilde{v}_2) &= \frac{\beta}{c(\beta)} \{\phi(\tilde{y}) - \phi(\tilde{v}_1) - (\tilde{y} - \tilde{v}_1)' \cdot \mathbf{D}\phi(\tilde{v}_1)\} - \beta \{\phi(\tilde{y}) - \phi(\tilde{v}_2) - (\tilde{y} - \tilde{v}_2)' \cdot \mathbf{D}\phi(\tilde{v}_2)\}, \\ \mathbf{D}_{\tilde{y}}g(\tilde{y}, \tilde{v}_1, \tilde{v}_2) &= \frac{\beta}{c(\beta)} \{\mathbf{D}\phi(\tilde{y}) - \mathbf{D}\phi(\tilde{v}_1)\} - \beta \{\mathbf{D}\phi(\tilde{y}) - \mathbf{D}\phi(\tilde{v}_2)\}, \\ \mathbf{H}_{\tilde{y}}g(\tilde{y}, \tilde{v}_1, \tilde{v}_2) &= \frac{\beta}{c(\beta)} \mathbf{H}\phi(\tilde{y}) - \beta \mathbf{H}\phi(\tilde{y}). \end{aligned}$$

And since $x \cdot x' \succ 0, \forall x \in \mathbb{R}^n$, (31) is satisfied for all $\tilde{y}, \tilde{v}_1, \tilde{v}_2 \in \tilde{\Delta}^n$ when $c(\beta) = 1$, which is the mixability condition (in addition requiring $\tilde{v}^* = \int \tilde{v}P(d\tilde{v})$ in (32) is the exp-concavity condition). Then by applying Lemma 9 proof is completed. \blacksquare

As an application of Proposition 10, we can see that both Kullback-Leibler loss and log loss are 1-mixable and 1-exp-concave.

Appendix C. Proofs

Proof (Proposition 1) We first prove that the set $\psi(\Delta^n)$ is convex. For any $p, q \in \Delta^n$, by assumption there exists $c \geq 0$ and $r \in \Delta^n$ such that $\frac{1}{2}(E_\beta(\ell(p)) + E_\beta(\ell(q))) + c\mathbb{1}_n = E_\beta(\ell(r))$. Therefore $\frac{1}{2}(\psi(p) + \psi(q)) = \psi(r)$, which implies the convexity of the set $\psi(\Delta^n)$.

Let $T : \mathbb{R}^{n-1} \ni (e^{-\beta z_1} - e^{-\beta z_n}, \dots, e^{-\beta z_{n-1}} - e^{-\beta z_n})' \rightarrow (e^{-\beta z_1}, \dots, e^{-\beta z_n})' \in [0, 1]^n$. Note this mapping from low dimension to high dimension is well defined because if there are two different z and \bar{z} in $\ell(\mathcal{V})$ such that $JE_\beta(z) = JE_\beta(\bar{z})$, then there must be $c \neq 0$ such that $E_\beta(z) + c\mathbb{1} = E_\beta(\bar{z})$. This means $z > \bar{z}$ or $z < \bar{z}$, which violates the strict properness of ℓ .

Since for any $v = \psi(p) = JE_\beta(\ell(p))$ we have $p = \ell^{-1}(E_\beta^{-1}(Tv))$, the link ψ is invertible (ℓ is invertible if it is strictly proper (Vernet et al. (2014)), and E_β is invertible for $\beta > 0$).

Now $\ell \circ \psi^{-1}$ is β -exp-concave if for all $p, q \in \Delta^n$

$$E_\beta \left(\ell \circ \psi^{-1} \left(\frac{1}{2}(\psi(p) + \psi(q)) \right) \right) \geq \frac{1}{2}E_\beta(\ell \circ \psi^{-1}(\psi(p))) + \frac{1}{2}E_\beta(\ell \circ \psi^{-1}(\psi(q))). \quad (33)$$

The right-hand side is obviously $\frac{1}{2}(E_\beta(\ell(p)) + E_\beta(\ell(q)))$. Let $r = \psi^{-1}(\frac{1}{2}(\psi(p) + \psi(q))) \in \Delta^n$. Then

$$JE_\beta(\ell(r)) = \psi(r) = \frac{1}{2}(JE_\beta(\ell(p)) + JE_\beta(\ell(q))).$$

Therefore $\frac{1}{2}(E_\beta(\ell(p)) + E_\beta(\ell(q))) = E_\beta(\ell(r)) + c\mathbb{1}_n$ for some $c \in \mathbb{R}$. To establish (33), it suffices to show $c \leq 0$. But this is guaranteed by the condition assumed. \blacksquare

Proof (Proposition 2) We first show that for a half space $H_{-p}^{\gamma_p}$ defined in (11) with $p \in \Delta_\epsilon^n$, $E_\beta^{-1}(H_{-p}^{\gamma_p} \cap \mathbb{R}_+^n)$ must be the super-prediction set of a scaled and shifted log loss. In fact, as $p_i > 0$, clearly $\gamma_p = \min_{x \in T_\ell} x' \cdot (-p) < 0$. Define a new loss $\tilde{\ell}_i^{\log}(q) = -\frac{1}{\beta} \log(-\frac{\gamma_p}{p_i} q_i)$ over $q \in \Delta^n$. Then $S_{\tilde{\ell}^{\log}} \subseteq E_\beta^{-1}(H_{-p}^{\gamma_p} \cap \mathbb{R}_+^n)$ can be seen from

$$\sum_i (-p_i) \exp(-\beta \tilde{\ell}_i^{\log}(q)) = \sum_i (-p_i) \left(-\frac{\gamma_p}{p_i} q_i\right) = \gamma_p. \quad (34)$$

Conversely, for any u such that $u_i > 0$ and $(-p)' \cdot u = \gamma_p$, simply choose $q_i = -\frac{u_i p_i}{\gamma_p}$. Then $q \in \Delta^n$ and $E_\beta(\tilde{\ell}^{\log}(q)) = u$. In summary, $E_\beta^{-1}(H_{-p}^{\gamma_p} \cap \mathbb{R}_+^n)$ is the super-prediction set of $\tilde{\ell}^{\log}$.

To prove Proposition 2, we first show that for any point $a \in T_\ell^\epsilon$ and any direction d from the relative interior of the positive orthant (which includes the $\mathbb{1}_n$ direction), the ray $\{a + rd : r \geq 0\}$ will be blocked by a boundary point of T_ℓ^ϵ . This is because by the definition of T_ℓ^ϵ in (12), the largest value of r to guarantee $a + rd \in T_\ell^\epsilon$ can be computed by

$$r^* := \sup\{r \geq 0 : a + rd \in T_\ell^\epsilon\} = \sup\{r \geq 0 : (a + rd)' \cdot (-p) \geq \gamma_p, \forall p \in \Delta_\epsilon^n\} \quad (35)$$

must be finite and attained. Denote $x = a + r^*d$, which must be on the boundary of T_ℓ^ϵ because

$$-x' \cdot p \geq \gamma_p, \text{ for all } p \in \Delta_\epsilon^n, \quad (36)$$

$$\text{and } -x' \cdot p^* = \gamma_{p^*} \text{ for some } p^* \in \Delta_\epsilon^n \text{ (not necessarily unique)}. \quad (37)$$

In order to prove the first statement of Proposition 2 by using Proposition 1, it suffices to show that for any point x on the boundary of T_ℓ^ϵ , there exists a $q \in \Delta^n$ such that $E_\beta(\tilde{\ell}_\epsilon(q)) = x$. Suppose x satisfies (36) and (37). Then consider the (shifted/scaled) log loss $\tilde{\ell}^{\log}$ that corresponds to $H_{p^*}^{\gamma_{p^*}}$. Because log loss is strictly proper, there must be a unique $q \in \Delta^n$ such that the hyperplane $H_0 := \{z : q' \cdot z = q' \cdot E_\beta^{-1}(x)\}$ supports the super-prediction set of $\tilde{\ell}^{\log}$ (i.e. $E_\beta^{-1}(H_{-p^*}^{\gamma_{p^*}} \cap \mathbb{R}_+^n)$) at $E_\beta^{-1}(x)$. Since $E_\beta^{-1}(T_\ell^\epsilon \cap \mathbb{R}_+^n)$ is a convex subset of $E_\beta^{-1}(H_{-p^*}^{\gamma_{p^*}} \cap \mathbb{R}_+^n)$, this hyperplane also supports $E_\beta^{-1}(T_\ell^\epsilon \cap \mathbb{R}_+^n)$ at $E_\beta^{-1}(x)$. Therefore $E_\beta^{-1}(x)$ is an optimal solution to the problem in the definition of $\tilde{\ell}_\epsilon(q)$ in (13). Finally observe that it must be the unique optimal solution, because if there were another solution which also lies on H_0 , then by the convexity of the super-prediction set of $\tilde{\ell}_\epsilon$, the line segment between them must also lie on the prediction set of $\tilde{\ell}_\epsilon$. This violates the mixability condition of $\tilde{\ell}_\epsilon$, because by construction its sub-exp-prediction set is convex.

In order to check where $\ell(p) = \tilde{\ell}_\epsilon(p)$, a sufficient condition is that the normal direction d on the exp-prediction set evaluated at $E_\beta(\ell(p))$ satisfies $d_i / \sum_j d_j > \epsilon$. Simple calculus shows that $d_i \propto p_i \exp(\beta \ell_i(p))$. Therefore as long as p is in the relative interior of Δ^n , $d_i / \sum_j d_j > \epsilon$ can always be satisfied by choosing a sufficiently small ϵ . And for each fixed ϵ , the set S_ϵ mentioned in the theorem consists exactly of all such p that satisfies this condition. ■

Proof (Proposition 5) When $n = 2$, (17), (18) and (19) and the positivity of $\tilde{\psi}'$ simplify (20) to the two conditions:

$$\begin{aligned} (1 - \tilde{p}) k'(\tilde{p}) &\leq k(\tilde{p}) - \alpha \tilde{\psi}'(\tilde{p}) k(\tilde{p})^2 (1 - \tilde{p})^2, \\ -\tilde{p} k'(\tilde{p}) &\leq k(\tilde{p}) - \alpha \tilde{\psi}'(\tilde{p}) k(\tilde{p})^2 \tilde{p}^2, \end{aligned}$$

for all $\tilde{p} \in (0, 1)$. These two conditions can be merged as follows

$$-\frac{1}{\tilde{p}} + \alpha \tilde{\psi}'(\tilde{p}) k(\tilde{p}) \tilde{p} \leq \frac{k'(\tilde{p})}{k(\tilde{p})} \leq \frac{1}{1 - \tilde{p}} - \alpha \tilde{\psi}'(\tilde{p}) k(\tilde{p}) (1 - \tilde{p}), \quad \forall \tilde{p} \in (0, 1).$$

By noting that $k(\tilde{p}) = \frac{w(\tilde{p})}{\tilde{\psi}'(\tilde{p})}$ and $k'(\tilde{p}) = \frac{w'(\tilde{p})\tilde{\psi}'(\tilde{p}) - w(\tilde{p})\tilde{\psi}''(\tilde{p})}{\tilde{\psi}'(\tilde{p})^2}$ completes the proof. ■

Proof (Proposition 6) Let $g(\tilde{p}) = \frac{1}{w(\tilde{p})}$ and so $g'(\tilde{p}) = -\frac{1}{w(\tilde{p})^2} w'(\tilde{p})$, $g(v) = \int_{\frac{1}{2}}^v g'(\tilde{p}) d\tilde{p} + g(\frac{1}{2})$ and $g(\frac{1}{2}) = \frac{1}{w(\frac{1}{2})} = 1$. By dividing all sides of (21) by $-w(\tilde{p})$ and applying the substitution we get,

$$\frac{1}{\tilde{p}} g(\tilde{p}) - \alpha \tilde{p} \geq g'(\tilde{p}) - \Phi_{\tilde{\psi}}(\tilde{p}) g(\tilde{p}) \geq -\frac{1}{1 - \tilde{p}} g(\tilde{p}) + \alpha(1 - \tilde{p}), \quad \forall \tilde{p} \in (0, 1), \quad (38)$$

where $\Phi_{\tilde{\psi}}(\tilde{p}) := -\frac{\tilde{\psi}''(\tilde{p})}{\tilde{\psi}'(\tilde{p})}$. If we take the first inequality of (38) and rearrange it we obtain,

$$-\alpha \geq \left(g'(\tilde{p}) \frac{1}{\tilde{p}} - g(\tilde{p}) \frac{1}{\tilde{p}^2} \right) - \Phi_{\tilde{\psi}}(\tilde{p}) \frac{g(\tilde{p})}{\tilde{p}} = \left(\frac{g(\tilde{p})}{\tilde{p}} \right)' - \Phi_{\tilde{\psi}}(\tilde{p}) \left(\frac{g(\tilde{p})}{\tilde{p}} \right), \quad \forall \tilde{p} \in (0, 1). \quad (39)$$

Multiplying (39) by $e^{-\int_0^{\tilde{p}} \Phi_{\tilde{\psi}}(t)dt}$ will result in,

$$\begin{aligned} -\alpha e^{-\int_0^{\tilde{p}} \Phi_{\tilde{\psi}}(t)dt} &\geq \left(\frac{g(\tilde{p})}{\tilde{p}}\right)' e^{-\int_0^{\tilde{p}} \Phi_{\tilde{\psi}}(t)dt} + \left(\frac{g(\tilde{p})}{\tilde{p}}\right) e^{-\int_0^{\tilde{p}} \Phi_{\tilde{\psi}}(t)dt} (-\Phi_{\tilde{\psi}}(\tilde{p})) \\ &= \left(\frac{g(\tilde{p})}{\tilde{p}} e^{-\int_0^{\tilde{p}} \Phi_{\tilde{\psi}}(t)dt}\right)', \quad \forall \tilde{p} \in (0, 1). \end{aligned} \quad (40)$$

Since

$$-\int_0^{\tilde{p}} \Phi_{\tilde{\psi}}(t)dt = -\int_0^{\tilde{p}} -\frac{\tilde{\psi}''(t)}{\tilde{\psi}'(t)} dt = \int_0^{\tilde{p}} (\log \tilde{\psi}'(t))' dt = \log \frac{\tilde{\psi}'(\tilde{p})}{\tilde{\psi}'(0)},$$

(40) is reduced to

$$\begin{aligned} -\alpha \frac{\tilde{\psi}'(\tilde{p})}{\tilde{\psi}'(0)} &\geq \left(\frac{g(\tilde{p})}{\tilde{p}} \frac{\tilde{\psi}'(\tilde{p})}{\tilde{\psi}'(0)}\right)', \quad \forall \tilde{p} \in (0, 1) \\ \Rightarrow -\alpha \tilde{\psi}'(\tilde{p}) &\geq \left(\frac{g(\tilde{p})}{\tilde{p}} \tilde{\psi}'(\tilde{p})\right)', \quad \forall \tilde{p} \in (0, 1). \end{aligned}$$

For $v \geq \frac{1}{2}$ we thus have

$$\begin{aligned} -\alpha \int_{\frac{1}{2}}^v \tilde{\psi}'(\tilde{p}) d\tilde{p} &\geq \int_{\frac{1}{2}}^v \left(\frac{g(\tilde{p})}{\tilde{p}} \tilde{\psi}'(\tilde{p})\right)' d\tilde{p}, \quad \forall v \in [1/2, 1) \\ \Rightarrow -\alpha(\tilde{\psi}(v) - \tilde{\psi}(\frac{1}{2})) &\geq \left(\frac{g(v)}{v} \tilde{\psi}'(v) - \frac{g(\frac{1}{2})}{\frac{1}{2}} \tilde{\psi}'(\frac{1}{2})\right) \\ &= \left(\frac{1}{w(v)v} \tilde{\psi}'(v) - 2\tilde{\psi}'(\frac{1}{2})\right), \quad \forall v \in [1/2, 1) \\ \Rightarrow w(v) &\geq \frac{\tilde{\psi}'(v)}{v(2\tilde{\psi}'(\frac{1}{2}) - \alpha(\tilde{\psi}(v) - \tilde{\psi}(\frac{1}{2})))}, \quad \forall v \in [1/2, 1). \end{aligned}$$

Also by considering $v \leq \frac{1}{2}$ case as above, we get

$$\frac{\tilde{\psi}'(v)}{v(2\tilde{\psi}'(\frac{1}{2}) - \alpha(\tilde{\psi}(v) - \tilde{\psi}(\frac{1}{2})))} \leq w(v), \quad \forall v \in (0, 1).$$

Finally by following the similar steps for the second inequality of (38), the proof will be completed. \blacksquare

Here we provide an integral inequalities related result (without proof) due to Beesack and presented in Dragomir (2000). This result will be helpful in deriving the sufficient conditions for the exp-concavity of proper losses.

Theorem 11 *Let y and k be continuous and f and g Riemann integrable functions on $J = [\alpha, \beta]$ with g and k nonnegative on J . If*

$$y(x) \geq f(x) + g(x) \int_{\alpha}^x y(t)k(t)dt, \quad x \in J,$$

then

$$y(x) \geq f(x) + g(x) \int_{\alpha}^x f(t)k(t) \exp\left(\int_t^x g(r)k(r)dr\right)dt, \quad x \in J.$$

The result remains valid if \int_{α}^x is replaced by \int_x^{β} and \int_t^x by \int_x^t throughout.

Proof (Theorem 7) The necessary and sufficient condition for the exp-concavity of proper losses is given by (23). But from (39), we can see that (23) is equivalent to

$$\left(\frac{g(\tilde{p})}{\tilde{p}}\right)' \leq -\alpha, \quad \forall \tilde{p} \in (0, 1), \quad (41)$$

and

$$-\left(\frac{g(\tilde{p})}{1-\tilde{p}}\right)' \leq -\alpha, \quad \forall \tilde{p} \in (0, 1), \quad (42)$$

where $g(\tilde{p}) = \frac{1}{w(\tilde{p})}$ with $w(\frac{1}{2}) = 1$; i.e. (23) if and only if (41) & (42).

Now if we choose the weight function $w(\tilde{p})$ as follows

$$w(\tilde{p}) = \frac{1}{\tilde{p} \left(2 + \int_{1/2}^{\tilde{p}} a(t)dt\right)}, \quad (43)$$

such that $a(t) \leq -\alpha$, then (41) will be satisfied (since (43) $\implies 2 + \int_{1/2}^{\tilde{p}} a(t)dt = \frac{1}{w(\tilde{p})\tilde{p}} = \frac{g(\tilde{p})}{\tilde{p}} \implies a(\tilde{p}) = \left(\frac{g(\tilde{p})}{\tilde{p}}\right)'$). Similarly the weight function $w(\tilde{p})$ given by

$$w(\tilde{p}) = \frac{1}{(1-\tilde{p}) \left(2 - \int_{1/2}^{\tilde{p}} b(t)dt\right)}, \quad (44)$$

with $b(t) \leq -\alpha$ will satisfy (42) (since (44) $\implies 2 - \int_{1/2}^{\tilde{p}} b(t)dt = \frac{1}{w(\tilde{p})(1-\tilde{p})} = \frac{g(\tilde{p})}{1-\tilde{p}} \implies -b(\tilde{p}) = \left(\frac{g(\tilde{p})}{1-\tilde{p}}\right)'$). To satisfy both (41) and (42) at the same time (then obviously (23) will be satisfied), we can make the two forms of the weight function ((43) and (44)) equivalent with the appropriate choice of $a(t)$ and $b(t)$. This can be done in two cases.

In the first case, for $\tilde{p} \in (0, 1/2]$ we can fix the weight function $w(\tilde{p})$ as given by (43) and choose $a(t)$ such that,

- $a(t) \leq -\alpha$ (then (41) is satisfied) and
- (43) = (44) $\implies b(t) \leq -\alpha$ (then (42) is satisfied).

But (43) = (44) for all $\tilde{p} \in (0, 1/2]$ if and only if

$$\begin{aligned} & \tilde{p} \left(2 - \int_{\tilde{p}}^{1/2} a(t)dt\right) = (1-\tilde{p}) \left(2 + \int_{\tilde{p}}^{1/2} b(t)dt\right), \quad \tilde{p} \in (0, 1/2] \\ \iff & \frac{\tilde{p}}{1-\tilde{p}} \left(2 - \int_{\tilde{p}}^{1/2} a(t)dt\right) - 2 = \int_{\tilde{p}}^{1/2} b(t)dt, \quad \tilde{p} \in (0, 1/2] \\ \iff & \frac{\tilde{p}}{1-\tilde{p}} a(\tilde{p}) + \frac{1}{(1-\tilde{p})^2} \left(2 - \int_{\tilde{p}}^{1/2} a(t)dt\right) = -b(x), \quad \tilde{p} \in (0, 1/2], \end{aligned}$$

where the last step is obtained by differentiating both sides w.r.t \tilde{p} . Thus the constraint (43) = (44) $\implies b(t) \leq -\alpha$ can be given as

$$\begin{aligned} & \frac{\tilde{p}}{1-\tilde{p}} a(\tilde{p}) + \frac{1}{(1-\tilde{p})^2} \left(2 - \int_{\tilde{p}}^{1/2} a(t) dt \right) \geq \alpha, \quad \tilde{p} \in (0, 1/2] \\ \iff & a(\tilde{p}) \geq \left[\frac{\alpha(1-\tilde{p})}{\tilde{p}} - \frac{2}{\tilde{p}(1-\tilde{p})} \right] + \frac{1}{\tilde{p}(1-\tilde{p})} \int_{\tilde{p}}^{1/2} a(t) dt, \quad \tilde{p} \in (0, 1/2] \\ \iff & a(\tilde{p}) \geq f(\tilde{p}) + g(\tilde{p}) \int_{\tilde{p}}^{1/2} a(t) k(t) dt, \quad \tilde{p} \in (0, 1/2], \end{aligned}$$

where $f(\tilde{p}) = \left[\frac{\alpha(1-\tilde{p})}{\tilde{p}} - \frac{2}{\tilde{p}(1-\tilde{p})} \right]$, $g(\tilde{p}) = \frac{1}{\tilde{p}(1-\tilde{p})} = \frac{1}{\tilde{p}} + \frac{1}{(1-\tilde{p})}$ and $k(t) = 1$. Now by applying Theorem 11 we have

$$a(\tilde{p}) \geq f(\tilde{p}) + g(\tilde{p}) \int_{\tilde{p}}^{1/2} f(t) k(t) \exp \left(\int_{\tilde{p}}^t g(r) k(r) dr \right) dt, \quad \tilde{p} \in (0, 1/2].$$

Since

$$\int_{\tilde{p}}^t g(r) k(r) dr = \int_{\tilde{p}}^t \frac{1}{r} + \frac{1}{(1-r)} dr = [\ln r - \ln(1-r)]_{\tilde{p}}^t = \ln \left(\frac{t}{(1-t)} \frac{(1-\tilde{p})}{\tilde{p}} \right),$$

$$\begin{aligned} \int_{\tilde{p}}^{1/2} f(t) k(t) \exp \left(\int_{\tilde{p}}^t g(r) k(r) dr \right) dt &= \int_{\tilde{p}}^{1/2} f(t) \frac{t}{(1-t)} \frac{(1-\tilde{p})}{\tilde{p}} dt \\ &= \frac{(1-\tilde{p})}{\tilde{p}} \int_{\tilde{p}}^{1/2} \left[\frac{\alpha(1-t)}{t} - \frac{2}{t(1-t)} \right] \frac{t}{(1-t)} dt \\ &= \frac{(1-\tilde{p})}{\tilde{p}} \int_{\tilde{p}}^{1/2} \alpha - \frac{2}{(1-t)^2} dt \\ &= \frac{(1-\tilde{p})}{\tilde{p}} \left[\alpha t - \frac{2}{1-t} \right]_{\tilde{p}}^{1/2} \\ &= \frac{(1-\tilde{p})}{\tilde{p}} \left[\frac{\alpha}{2} - 4 - \alpha \tilde{p} + \frac{2}{1-\tilde{p}} \right], \end{aligned}$$

we get

$$\begin{aligned} a(\tilde{p}) &\geq \left[\frac{\alpha(1-\tilde{p})}{\tilde{p}} - \frac{2}{\tilde{p}(1-\tilde{p})} \right] + \frac{1}{\tilde{p}(1-\tilde{p})} \frac{(1-\tilde{p})}{\tilde{p}} \left[\frac{\alpha}{2} - 4 - \alpha \tilde{p} + \frac{2}{1-\tilde{p}} \right], \quad \tilde{p} \in (0, 1/2] \\ &= -\alpha + \frac{\alpha}{2\tilde{p}^2} - \frac{2}{\tilde{p}^2}, \quad \tilde{p} \in (0, 1/2]. \end{aligned}$$

Similarly in the second case, for $\tilde{p} \in [1/2, 1)$ we can fix the weight function $w(\tilde{p})$ as given by (44) and choose $b(t)$ such that,

- $b(t) \leq -\alpha$ (then (42) is satisfied) and
- (43) = (44) $\implies a(t) \leq -\alpha$ (then (41) is satisfied).

But (43) = (44) for all $\tilde{p} \in [1/2, 1)$ if and only if

$$\begin{aligned} & \tilde{p} \left(2 + \int_{1/2}^{\tilde{p}} a(t) dt \right) = (1 - \tilde{p}) \left(2 - \int_{1/2}^{\tilde{p}} b(t) dt \right), \quad \tilde{p} \in [1/2, 1) \\ \iff & \int_{\tilde{p}}^{1/2} a(t) dt = \frac{1 - \tilde{p}}{\tilde{p}} \left(2 - \int_{1/2}^{\tilde{p}} b(t) dt \right) - 2, \quad \tilde{p} \in [1/2, 1) \\ \iff & a(\tilde{p}) = -\frac{1 - \tilde{p}}{\tilde{p}} b(\tilde{p}) - \frac{1}{\tilde{p}^2} \left(2 - \int_{1/2}^{\tilde{p}} b(t) dt \right), \quad \tilde{p} \in [1/2, 1), \end{aligned}$$

where the last step is obtained by differentiating both sides w.r.t \tilde{p} . Thus the constraint (43) = (44) $\implies a(t) \leq -\alpha$ can be given as

$$\begin{aligned} & \frac{1 - \tilde{p}}{\tilde{p}} b(\tilde{p}) + \frac{1}{\tilde{p}^2} \left(2 - \int_{1/2}^{\tilde{p}} b(t) dt \right) \geq \alpha, \quad \tilde{p} \in [1/2, 1) \\ \iff & b(\tilde{p}) \geq \left[\frac{\alpha \tilde{p}}{(1 - \tilde{p})} - \frac{2}{\tilde{p}(1 - \tilde{p})} \right] + \frac{1}{\tilde{p}(1 - \tilde{p})} \int_{1/2}^{\tilde{p}} b(t) dt, \quad \tilde{p} \in [1/2, 1) \\ \iff & b(\tilde{p}) \geq f(\tilde{p}) + g(\tilde{p}) \int_{1/2}^{\tilde{p}} b(t) k(t) dt, \quad \tilde{p} \in [1/2, 1), \end{aligned}$$

where $f(\tilde{p}) = \left[\frac{\alpha \tilde{p}}{(1 - \tilde{p})} - \frac{2}{\tilde{p}(1 - \tilde{p})} \right]$, $g(\tilde{p}) = \frac{1}{\tilde{p}(1 - \tilde{p})} = \frac{1}{\tilde{p}} + \frac{1}{(1 - \tilde{p})}$ and $k(t) = 1$. Again by applying Theorem 11 we have

$$b(\tilde{p}) \geq f(\tilde{p}) + g(\tilde{p}) \int_{1/2}^{\tilde{p}} f(t) k(t) \exp \left(\int_t^{\tilde{p}} g(r) k(r) dr \right) dt, \quad \tilde{p} \in [1/2, 1).$$

Since

$$\begin{aligned} \int_t^{\tilde{p}} g(r) k(r) dr &= \int_t^{\tilde{p}} \frac{1}{r} + \frac{1}{(1 - r)} dr = [\ln r - \ln(1 - r)]_t^{\tilde{p}} = \ln \left(\frac{\tilde{p}}{(1 - \tilde{p})} \frac{(1 - t)}{t} \right), \\ \int_{1/2}^{\tilde{p}} f(t) k(t) \exp \left(\int_t^{\tilde{p}} g(r) k(r) dr \right) dt &= \int_{1/2}^{\tilde{p}} f(t) \frac{\tilde{p}}{(1 - \tilde{p})} \frac{(1 - t)}{t} dt \\ &= \frac{\tilde{p}}{(1 - \tilde{p})} \int_{1/2}^{\tilde{p}} \left[\frac{\alpha t}{(1 - t)} - \frac{2}{t(1 - t)} \right] \frac{(1 - t)}{t} dt \\ &= \frac{\tilde{p}}{(1 - \tilde{p})} \int_{1/2}^{\tilde{p}} \alpha - \frac{2}{t^2} dt \\ &= \frac{\tilde{p}}{(1 - \tilde{p})} \left[\alpha t + \frac{2}{t} \right]_{1/2}^{\tilde{p}} \\ &= \frac{\tilde{p}}{(1 - \tilde{p})} \left[\alpha \tilde{p} + \frac{2}{\tilde{p}} - \frac{\alpha}{2} - 4 \right], \end{aligned}$$

we get

$$\begin{aligned} b(\tilde{p}) &\geq \left[\frac{\alpha\tilde{p}}{(1-\tilde{p})} - \frac{2}{\tilde{p}(1-\tilde{p})} \right] + \frac{1}{\tilde{p}(1-\tilde{p})} \frac{\tilde{p}}{(1-\tilde{p})} \left[\alpha\tilde{p} + \frac{2}{\tilde{p}} - \frac{\alpha}{2} - 4 \right], \quad \tilde{p} \in [1/2, 1) \\ &= \frac{\alpha\tilde{p}}{(1-\tilde{p})} + \frac{\alpha\tilde{p}}{(1-\tilde{p})^2} - \frac{\alpha}{2(1-\tilde{p})^2} - \frac{2}{(1-\tilde{p})^2}, \quad \tilde{p} \in [1/2, 1). \end{aligned}$$

■

Proof (Corollary 8) We have to show that $\tilde{\psi}_\ell^*$ will satisfy (21) with $\alpha = \beta$, for all β -mixable proper loss functions. Since

$$\frac{(\tilde{\psi}_\ell^*)''(\tilde{p})}{(\tilde{\psi}_\ell^*)'(\tilde{p})} = \frac{\frac{w'_\ell(\tilde{p})w_{\ell\log}(\tilde{p}) - w_\ell(\tilde{p})w'_{\ell\log}(\tilde{p})}{w_{\ell\log}(\tilde{p})^2}}{\frac{w_\ell(\tilde{p})}{w_{\ell\log}(\tilde{p})}} = \frac{w'_\ell(\tilde{p})}{w_\ell(\tilde{p})} - \frac{w'_{\ell\log}(\tilde{p})}{w_{\ell\log}(\tilde{p})} = \frac{w'_\ell(\tilde{p})}{w_\ell(\tilde{p})} - (\log w_{\ell\log}(\tilde{p}))', \quad (45)$$

by substituting $\tilde{\psi} = \tilde{\psi}_\ell^*$ and $\alpha = \beta$ in (21) we have,

$$\begin{aligned} -\frac{1}{\tilde{p}} + \beta w_\ell(\tilde{p})\tilde{p} &\leq \frac{w'_\ell(\tilde{p})}{w_\ell(\tilde{p})} - \frac{(\tilde{\psi}_\ell^*)''(\tilde{p})}{(\tilde{\psi}_\ell^*)'(\tilde{p})} \leq \frac{1}{1-\tilde{p}} - \beta w_\ell(\tilde{p})(1-\tilde{p}), \quad \forall \tilde{p} \in (0, 1) \\ \iff -\frac{1}{\tilde{p}} + \beta w_\ell(\tilde{p})\tilde{p} &\leq (\log w_{\ell\log}(\tilde{p}))' \leq \frac{1}{1-\tilde{p}} - \beta w_\ell(\tilde{p})(1-\tilde{p}), \quad \forall \tilde{p} \in (0, 1) \\ \iff -\frac{1}{\tilde{p}} + \beta w_\ell(\tilde{p})\tilde{p} &\leq -\frac{1}{\tilde{p}} + \frac{1}{1-\tilde{p}} \leq \frac{1}{1-\tilde{p}} - \beta w_\ell(\tilde{p})(1-\tilde{p}), \quad \forall \tilde{p} \in (0, 1) \\ \iff \beta &\leq \frac{1}{\tilde{p}(1-\tilde{p})w_\ell(\tilde{p})} = \frac{w_{\ell\log}(\tilde{p})}{w_\ell(\tilde{p})}, \quad \forall \tilde{p} \in (0, 1), \end{aligned}$$

which is true for all β -mixable binary proper loss functions. From (45)

$$\begin{aligned} (\log (\tilde{\psi}_\ell^*)'(\tilde{p}))' &= (\log w_\ell(\tilde{p}))' - (\log w_{\ell\log}(\tilde{p}))' = \left(\log \frac{w_\ell(\tilde{p})}{w_{\ell\log}(\tilde{p})} \right)', \\ \Rightarrow \left[\log (\tilde{\psi}_\ell^*)'(\tilde{p}) \right]_{1/2}^{\tilde{p}} &= \left[\log \frac{w_\ell(\tilde{p})}{w_{\ell\log}(\tilde{p})} \right]_{1/2}^{\tilde{p}}, \\ \Rightarrow \log (\tilde{\psi}_\ell^*)'(\tilde{p}) - \log (\tilde{\psi}_\ell^*)' \left(\frac{1}{2} \right) &= \log \frac{w_\ell(\tilde{p})}{w_{\ell\log}(\tilde{p})} \cdot \frac{w_{\ell\log}(\frac{1}{2})}{w_\ell(\frac{1}{2})}, \end{aligned}$$

it can be seen that a design choice of $(\tilde{\psi}_\ell^*)'(\frac{1}{2}) = 1$ is made in the construction of this link function.

■

Appendix D. Squared Loss

In this section we will consider the multi-class squared loss with partial losses given by $\ell_i^{\text{sq}}(p) := \sum_{j \in [n]} (\mathbb{1}\{i = j\} - p_j)^2$. The Bayes risk of this loss is $\underline{L}_{\ell^{\text{sq}}}(\tilde{p}) = 1 - \sum_{i=1}^{n-1} p_i^2 - (1 - \sum_{i=1}^{n-1} p_i)^2$.

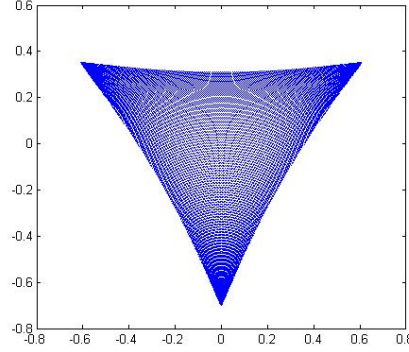


Figure 12: Projection of the exp-prediction set of square loss ($\beta = 1$) along the $\mathbb{1}_3$ direction. By the apparent lack of convexity of the projection, the condition $\partial_{\mathbb{1}_n} \mathcal{B}_\beta \subseteq E_\beta(\ell(\mathcal{V}))$ in Proposition 1 does not hold in this case.

Thus the Hessian of the Bayes risk is given by

$$\mathbf{H}_{\tilde{\mathcal{L}}_{\ell^{\text{sq}}}(\tilde{p})} = 2 \begin{pmatrix} -2 & -1 & \cdots & -1 \\ -1 & -2 & \cdots & -1 \\ \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & \cdots & -2 \end{pmatrix}.$$

For the identity link, from (16) we get $k(\tilde{p}) = -\mathbf{H}_{\tilde{\mathcal{L}}_{\ell^{\text{sq}}}(\tilde{p})}$ and $\mathbf{D}_v[k(\tilde{p})] = 0$ since $\mathbf{D}\tilde{\psi}(\tilde{p}) = I_{n-1}$. Thus from (20), the multi-class squared loss is α -exp-concave (with $\alpha > 0$) if and only if for all $\tilde{p} \in \overset{\circ}{\Delta}^n$ and for all $i \in [n]$

$$\begin{aligned} 0 &\preceq k(\tilde{p}) - \alpha k(\tilde{p}) \cdot (e_i^{n-1} - \tilde{p}) \cdot (e_i^{n-1} - \tilde{p})' \cdot k(\tilde{p}) \\ \iff k(\tilde{p})^{-1} &\succcurlyeq \alpha (e_i^{n-1} - \tilde{p}) \cdot (e_i^{n-1} - \tilde{p})'. \end{aligned} \quad (46)$$

Similarly for the canonical link, from (25) and (26), the composite loss is α -exp-concave (with $\alpha > 0$) if and only if for all $\tilde{p} \in \overset{\circ}{\Delta}^n$ and for all $i \in [n]$

$$k(\tilde{p})^{-1} = -[\mathbf{H}_{\tilde{\mathcal{L}}_{\ell^{\text{sq}}}(\tilde{p})}]^{-1} \succcurlyeq \alpha (e_i^{n-1} - \tilde{p}) \cdot (e_i^{n-1} - \tilde{p})'. \quad (47)$$

From (46) and (47), it can be seen that for the multi-class squared loss the level of exp-concavification by identity link and canonical link are same. When $n = 2$, since $k(\tilde{p}) = 4$, the condition (46) is equivalent to

$$\begin{aligned} \frac{1}{4} &\geq \alpha (e_i^{n-1} - \tilde{p}) \cdot (e_i^{n-1} - \tilde{p})', \quad i \in [2], \forall \tilde{p} \in (0, 1) \\ \iff \alpha &\leq \frac{1}{4\tilde{p}^2} \quad \text{and} \quad \alpha \leq \frac{1}{4(1-\tilde{p})^2}, \quad \forall \tilde{p} \in (0, 1) \\ \iff \alpha &\leq \frac{1}{4}. \end{aligned}$$

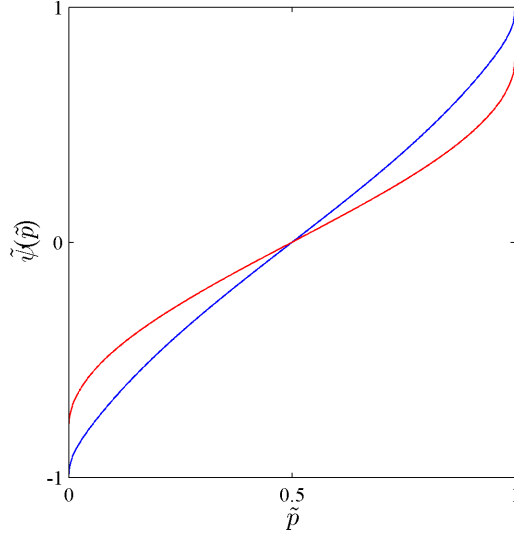


Figure 13: Exp-concavifying link functions for binary boosting loss constructed by Proposition 1 (—) and Corollary 8 (—).

When $n = 3$, using the fact that a 2×2 matrix is positive semi-definite if its trace and determinant are both non-negative, it can be easily verified that the condition (46) is equivalent to $\alpha \leq \frac{1}{12}$.

For binary squared loss, the link functions constructed by geometric (Proposition 1) and calculus (Corollary 8) approach are:

$$\tilde{\psi}(\tilde{p}) = e^{-2(1-\tilde{p})^2} - e^{-2\tilde{p}^2} \quad \text{and} \quad \tilde{\psi}_\ell^*(\tilde{p}) = \frac{4}{4} \int_0^{\tilde{p}} \frac{w_{\ell^{\text{sq}}}(v)}{w_{\ell^{\log}}(v)} dv = 4 \left(\frac{\tilde{p}^2}{2} - \frac{\tilde{p}^3}{3} \right),$$

respectively. By applying these link functions we can get 1-exp-concave composite squared loss.

Appendix E. Boosting Loss

Consider the binary “boosting loss” (Buja et al. (2005)) with partial losses given by

$$\ell_1^{\text{boost}}(\tilde{p}) = \frac{1}{2} \sqrt{\frac{1-\tilde{p}}{\tilde{p}}} \quad \text{and} \quad \ell_2^{\text{boost}}(\tilde{p}) = \frac{1}{2} \sqrt{\frac{\tilde{p}}{1-\tilde{p}}}, \quad \forall \tilde{p} \in (0, 1).$$

This loss has weight function

$$w_{\ell^{\text{boost}}}(\tilde{p}) = \frac{1}{4(\tilde{p}(1-\tilde{p}))^{3/2}}, \quad \forall \tilde{p} \in (0, 1).$$

By applying the results of Van Erven et al. (2012), we can show that this loss is mixable with mixability constant 2 (since $\beta_\ell = \inf_{\tilde{p} \in (0,1)} \frac{w_{\ell^{\log}}(\tilde{p})}{w_\ell(\tilde{p})}$).

Now we can check the level of exp-concavification of this loss for different choices of link functions. By considering the identity link $\tilde{\psi}(\tilde{p}) = \tilde{p}$, from (23)

$$\begin{aligned}
 -\frac{1}{\tilde{p}} + \alpha w_{\ell^{\text{boost}}}(\tilde{p})\tilde{p} &\leq \frac{w'_{\ell^{\text{boost}}}(\tilde{p})}{w_{\ell^{\text{boost}}}(\tilde{p})}, \quad \forall \tilde{p} \in (0, 1) \\
 \Rightarrow -\frac{1}{\tilde{p}} + \alpha w_{\ell^{\text{boost}}}(\tilde{p})\tilde{p} &\leq 6w_{\ell^{\text{boost}}}(\tilde{p})\sqrt{\tilde{p}(1-\tilde{p})}(2\tilde{p}-1), \quad \forall \tilde{p} \in (0, 1) \\
 \Rightarrow \alpha\tilde{p} - 6\sqrt{\tilde{p}(1-\tilde{p})}(2\tilde{p}-1) &\leq \frac{1}{w_{\ell^{\text{boost}}}(\tilde{p})\tilde{p}}, \quad \forall \tilde{p} \in (0, 1) \\
 \Rightarrow \alpha &\leq 8\sqrt{\frac{1-\tilde{p}}{\tilde{p}}}(\tilde{p}-1/4), \quad \forall \tilde{p} \in (0, 1) \\
 \Rightarrow \alpha &\leq 0,
 \end{aligned}$$

we see that the boosting loss is non-exp-concave. Similarly from (27)

$$\alpha \leq \frac{1}{w_{\ell^{\text{boost}}}(\tilde{p})\tilde{p}^2} = 4\sqrt{\frac{1-\tilde{p}}{\tilde{p}}}(1-\tilde{p}), \quad \forall \tilde{p} \in (0, 1) \quad (48)$$

it can be seen that the RHS of (48) approaches 0 as $p \rightarrow 1$, thus it is not possible to exp-concavify (for some $\alpha > 0$) this loss using the canonical link. For binary boosting loss, the link functions constructed by geometric (Proposition 1) and calculus (Corollary 8) approach are:

$$\tilde{\psi}(\tilde{p}) = e^{-\sqrt{\frac{1-\tilde{p}}{\tilde{p}}}} - e^{-\sqrt{\frac{\tilde{p}}{1-\tilde{p}}}} \quad \text{and} \quad \tilde{\psi}_{\ell}^*(\tilde{p}) = \frac{4}{2} \int_0^{\tilde{p}} \frac{w_{\ell^{\text{boost}}}(v)}{w_{\ell^{\log}}(v)} dv = \frac{1}{2} \arcsin(-1 + 2\tilde{p}),$$

respectively (as shown in Figure 13). By applying these link functions we can get 2-exp-concave composite boosting loss.

Appendix F. Log Loss

By using the results from this paper and Van Erven et al. (2012) one can easily verify that the multi-class log loss is both 1-mixable and 1-exp-concave. For binary log loss, the link functions constructed by geometric (Proposition 1) and calculus (Corollary 8) approach are:

$$\tilde{\psi}(\tilde{p}) = e^{\log \tilde{p}} - e^{\log 1-\tilde{p}} = 2\tilde{p} - 1 \quad \text{and} \quad \tilde{\psi}_{\ell}^*(\tilde{p}) = \frac{4}{4} \int_0^{\tilde{p}} \frac{w_{\ell^{\log}}(v)}{w_{\ell^{\log}}(v)} dv = \tilde{p},$$

respectively.