

Unified Binary and Multiclass Margin-Based Classification

Yutong Wang^{1,2}

Clayton Scott^{1,3}

¹Department of Electrical Engineering and Computer Science

²Michigan Institute of Data Science

³Department of Statistics

University of Michigan

Ann Arbor, MI 48109, USA

YUTONGW@UMICH.EDU

CLAYSCOT@UMICH.EDU

Editor: Zhihua Zhang

Abstract

The notion of margin loss has been central to the development and analysis of algorithms for binary classification. To date, however, there remains no consensus as to the analogue of the margin loss for multiclass classification. In this work, we show that a broad range of multiclass loss functions, including many popular ones, can be expressed in the *relative margin form*, a generalization of the margin form of binary losses. The relative margin form is broadly useful for understanding and analyzing multiclass losses as shown by our prior work (Wang and Scott, 2020, 2021). To further demonstrate the utility of this way of expressing multiclass losses, we use it to extend the seminal result of Bartlett et al. (2006) on classification-calibration of binary margin losses to multiclass. We then analyze the class of Fenchel-Young losses, and expand the set of these losses that are known to be classification-calibrated.

Keywords: Classification, loss functions, consistency, margins, label encodings

1. Introduction

Classification into $k \geq 2$ categories is the learning task of selecting a *classifier*, i.e., a function from the feature space \mathcal{X} to the set of labels $[k] := \{1, \dots, k\}$, given training data. Many of the most popular and successful classification methods aim to find a classifier with minimum risk, where the risk of a classifier is the expected value of a loss function that measures the quality of predictions. Indeed, logistic regression, support vector machines, boosting, and neural network methods can all be viewed as algorithms to minimize the risk associated with a certain loss. In practice, “loss-based” approaches to binary ($k = 2$) and multiclass ($k \geq 3$) cases are formulated differently. Because of this discrepancy, the theory and practice of multiclass methods often lag behind their binary counterparts. The purpose of this work is to bridge this gap by introducing a framework that unifies binary and multiclass loss-based classification.

The standard approach to binary classification is to learn a *discriminant* function $g : \mathcal{X} \rightarrow \mathbb{R}$ that maps an instance x to a *discriminant* $g(x)$, e.g., $g(x) = w^T x + b$ for linear classification. Identifying $\{1, 2\}$ with $\{-1, +1\}$, a label y is predicted by $\text{sign}(g(x))$. The loss ascribed to a discriminant g and a pair (x, y) is defined in terms of a function $\psi : \mathbb{R} \rightarrow \mathbb{R}$, with the loss being equal to $\psi((-1)^y g(x))$. Loss functions of this form are referred to as

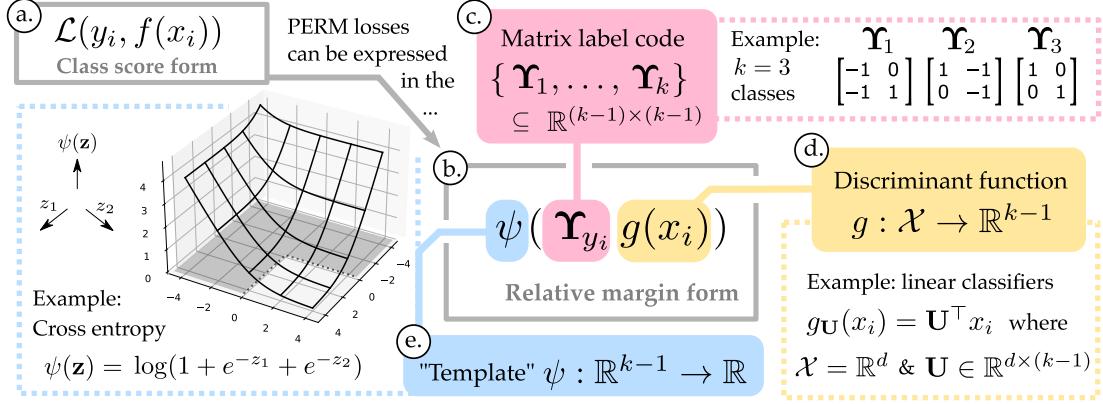


Figure 1: **Relative margin form.** *Panel (a):* Multiclass losses \mathcal{L} satisfying the permutation equivariant and relative margin-based conditions (i.e., PERM losses, Definition 2.2) can be expressed in the relative margin form as in *Panel (b)*. See Theorem 2.5. The relative margin form employs three components: *Panels (c)*. the matrix label code $\{\Upsilon_y\}_{y \in [k]}$, *(d)*. the discriminant function g , and *(e)*. the “template”, a symmetric function $\psi : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$.

binary *margin loss* functions. Examples include the logistic loss $\psi(t) = \log(1 + e^{-t})$ (logistic regression), the hinge loss $\psi(t) = \max(0, 1 - t)$ (support vector machines), the exponential loss $\psi(t) = e^{-t}$ (AdaBoost), and the sigmoid loss $\psi(t) = 1/(1 + e^t)$ (some neural networks).

The conventional approach to multiclass classification seeks to learn a *class-score function* $f = (f_1, \dots, f_k) : \mathcal{X} \rightarrow \mathbb{R}^k$, e.g., a feed-forward neural network. The label for an input x is predicted by taking the argmax of the *class-score vectors* $\mathbf{v} := f(x)$. Multiclass loss functions, e.g., cross entropy, may be viewed as functions $\mathcal{L} : \{1, \dots, k\} \times \mathbb{R}^k \rightarrow \mathbb{R}$. The loss incurred by a class-score function f on a pair $(x, y) \in \mathcal{X} \times [k]$ is then $\mathcal{L}(y, f(x))$. When a multiclass loss function is expressed in terms of a class-score function output, we say that it is in *class-score form*.

We show that a large family of multiclass loss functions in class-score form, including cross-entropy, multiclass exponential loss (Mukherjee and Schapire, 2013), multiclass hinge losses (Crammer and Singer, 2001; Weston and Watkins, 1998), Gamma-Phi losses (Beijbom et al., 2014), and Fenchel-Young losses (Blondel et al., 2020), can be expressed in what we call the *relative-margin form*. Instead of being defined in terms of a class-score function output $f(x)$, in the relative margin form a loss is expressed $\psi(\Upsilon_y g(x))$, which involves the following elements:

1. a symmetric¹ function $\psi : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$,
2. a set of $(k - 1) \times (k - 1)$ matrices $\{\Upsilon_y\}_{y \in [k]}$ that encode the label, and
3. a discriminant function of the form $g : \mathcal{X} \rightarrow \mathbb{R}^{k-1}$.

When $k = 2$, we have $\Upsilon_y = (-1)^y$, and the relative margin form coincides with the standard binary margin form. See Figure 1 for an illustration of the framework. Thus, the relative margin form generalizes the notion of margin loss from binary classification.

1. Recall that a function is symmetric if its value does not change when its input is permuted.

1.1 Our contributions

In this work, we develop the relative margin from as described above, and illustrate the utility of this framework by establishing novel results on *classification-calibration*². More specifically, our contributions are:

- **Theorem 2.5: characterization of losses expressible in the relative margin form.** We show that a multiclass loss function can be expressed in the relative margin form if and only if both of the following conditions are met: 1. *permutation equivariance*: \mathcal{L} treats all classes equally, and 2. *relative margin-based*: $\mathcal{L}(y, \bar{f}(x))$ depends only on the differences between the $f_j(x)$ ’s rather than the values of the $f_j(x)$ ’s themselves. Below, such losses are referred to as PERM losses.

Our characterization implies that many popular losses such as the cross entropy, multiclass exponential, and Fenchel-Young losses (Blondel, 2019) are PERM losses. PERM losses are characterized by their *template*, a symmetric function $\psi : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$. Thus, our result implies that for the analysis of existing and design of new PERM losses, it suffices to focus on the template ψ which can be simpler than the original \mathcal{L} .

- **Theorem 4.7: Extending a fundamental result on classification-calibration (CC) to the multiclass case.** A seminal result of Bartlett et al. (2006) in the binary case shows that a convex margin loss is classification-calibrated if and only if ψ is differentiable at 0 and has negative derivative there. We use the PERM loss framework to prove a multiclass extension of this result for \mathcal{L} that are totally regular (Definition 4.2). We use Theorem 4.7 to prove a novel sufficient condition for CC of *sums* of losses (Proposition 4.8), and, as a corollary, establish CC for sums of Gamma-Phi losses (Example 5). A key component of the totally regular property is that the gradient of the template ∇_ψ must be entry-wise negative everywhere³.
- **Theorem 5.10: Expanding previous sufficient conditions of CC for Fenchel-Young losses.** Fenchel-Young losses, defined via taking the convex conjugate of a certain *negentropy* function, have recently been proposed for structured prediction, including multiclass classification. However, an open question is whether Fenchel-Young losses are classification-calibrated under the assumption that the negentropy is *strictly convex*. Our Theorem 5.10 answers this in the affirmative.
- **The matrix label code.** The key component in our analysis is the *matrix label code* (Definition 2.4) which extends to the multiclass case the universally adopted $\{\pm 1\}$ label code for binary classification. To obtain our main results, we developed a suite of technical lemmas for working with the matrix label code in Appendix B. We expect the matrix label code to be useful for multiclass classification research beyond the scope here.

1.2 Background

In this section, we define the probabilistic setting assumed throughout this work. Moreover, we first review the background on classification-calibration.

2. For readers unfamiliar with this theory, we provide a brief review below in the “Background on Classification-Calibration” section (Section 1.2).
3. Thus, in the binary case, our result is weaker than that of Bartlett et al. (2006), which only requires negativity at $0 \in \mathbb{R}$. We leave as an open question whether this gap can be closed.

Let $\{(x_i, y_i)\}_{i=1}^n$ be drawn from a joint distribution P over $\mathcal{X} \times [k]$. Let \mathbb{I} be the indicator function and $\mathbb{I}\{y \neq j\}$ be the *01-loss* for $y, j \in [k]$. The goal of classification is to select a classifier $h : \mathcal{X} \rightarrow [k]$ minimizing the *01-risk* $R_{01,P}(h) := \mathbb{E}_{(X,Y) \sim P} [\mathbb{I}\{Y \neq h(X)\}]$ objective. A fundamental approach for learning a classifier is *empirical risk minimization* (ERM), which selects an h minimizing the empirical 01-risk $\hat{R}_{01}^n(h) := \frac{1}{n} \sum_{i=1}^n \mathbb{I}\{y_i \neq h(x_i)\}$ over a class of functions \mathcal{H} .

Directly minimizing the empirical 01-risk objective is often intractable⁴ due to the discreteness of the objective. The surrogate-based approach addresses this issue as follows. Define the \mathcal{L} -risk for a fixed continuous loss function \mathcal{L} by $R_{\mathcal{L}}(f) := \mathbb{E}_{(X,Y) \sim P} [\mathcal{L}(Y, f(X))]$ and the empirical \mathcal{L} -risk by $\hat{R}_{\mathcal{L}}^n(f) := \frac{1}{n} \sum_{i=1}^n \mathcal{L}(y_i, f(x_i))$. Minimizing $\hat{R}_{\mathcal{L}}^n(f)$ over a family of f 's, e.g., neural networks, is employed as a tractable surrogate objective. A class-score function f induces a discrete-valued classifier $\arg \max \circ f(x) := \arg \max_{j=1, \dots, k} f_j(x)$ with ties broken arbitrarily. An important question is whether good performance with respect to the surrogate \mathcal{L} -risk “transfers” back to the 01-risk, the original objective.

Definition 1.1 A surrogate loss \mathcal{L} has the *consistency transfer property*⁵ if for any distribution P over $\mathcal{X} \times [k]$ and any sequence of class-score functions $\{\hat{f}^{(n)}\}_n$, e.g., $\hat{f}^{(n)}$ is an empirical \mathcal{L} -risk minimizer over a family of functions depending on n , the following is satisfied: $\lim_{n \rightarrow \infty} R_{\mathcal{L}}(\hat{f}^{(n)}) = \inf_f R_{\mathcal{L}}(f)$ implies that $\lim_{n \rightarrow \infty} R_{01}(\arg \max \circ \hat{f}^{(n)}) = \inf_h R_{01}(h)$. The infimums are, respectively, over all measurable functions $f : \mathcal{X} \rightarrow \mathbb{R}^k$ and $h : \mathcal{X} \rightarrow [k]$.

Thus, the consistency transfer property (CTP) provides justification for \mathcal{L} -risk minimization when minimizing the 01-risk is the original objective of interest. For binary margin-based losses ψ , the seminal result of Bartlett et al. (2006, Theorem 1.3) shows that CTP is equivalent to a functional property of ψ known as *classification-calibration* (CC). For the multiclass case, Tewari and Bartlett (2007) define the CC property (reviewed in detail in Section 3 below) and show its equivalence to the CTP as well.

1.3 Related work

- **Sufficient conditions for CC of multiclass convex margin-based losses.** For binary convex loss in margin form, the sufficient conditions for CC are easy to verify and apply to common losses of interest. Indeed, Bartlett et al. (2006, Theorem 2.1) asserts that a nonnegative binary convex loss in margin form $\psi : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ is CC if and only if ψ is differentiable at 0 and $\psi'(0) < 0$.

There are several difficulties for extending the above characterization to the multiclass case. The result in the binary case relies on the expression of a binary margin-based loss as $\psi((-1)^y z)$. While proposed multiclass extensions of the expression exist (to be reviewed below), existing work on sufficient conditions for CC of multiclass losses focus on losses in the class-score form $\mathcal{L} : [k] \times \mathbb{R}^k \rightarrow \mathbb{R}$.

Most previous works studied CC of losses of a specific form such as the Gamma-Phi losses (Zhang, 2004; Beijbom et al., 2014; Wang and Scott, 2023a), Fenchel-Young losses (Duchi et al., 2018; Blondel et al., 2020), and hinge-like losses (Tan and Zhang, 2022).

4. See Bhattacharyya et al. (2018) and the references therein

5. This property appears in many works, e.g., Steinwart (2007); Bartlett et al. (2006); Tewari and Bartlett (2007); Zhang (2004). The name “consistency transfer” was used by Wang and Scott (2023a).

Notably, Tewari and Bartlett (2007) derive a characterization of CC for multiclass losses analogous to that of Bartlett et al. (2006, Theorem 2.1). However, their sufficient condition (Tewari and Bartlett, 2007, Theorem 7), while geometrically elegant, is hard to verify for common losses such as the multiclass exponential loss⁶.

• **Multiclass margins & margin-based loss.** A line of work has considered the problem of extending the binary margin $(-1)^y z$ and binary margin form $\psi((-1)^y z)$ to the multiclass setting. For the margin, Crammer and Singer (2001) and Mohri et al. (2018, §9.2) defined a notion of *scalar-valued* multiclass margin as $\min_{j \in [k]: j \neq y} v_y - v_j$ where y is the ground truth label. In other words, this notion of multiclass margin is the minimum difference of score between the ground truth label and the rest of the labels. While intuitive, according to Tewari and Bartlett (2007, §5.1), convex losses based on this notion of margin are never classification-calibrated.

Lee et al. (2004); Zou et al. (2008) proposes definitions of *vector-valued* multiclass margin vectors and associated margin losses for multiclass SVMs and boosting, respectively. These losses are known to be classification-calibrated if and only if sum-to-zero constraints are enforced on the input margin vectors (Dogan et al., 2016, Theorem 6). Lee et al. (2004) develops a multiclass SVM based on a hinge loss that leverages this notion of margin. However, in practice, enforcing these sum-to-zero constraints leads to significantly slower computations (Fu et al., 2022).

Dogan et al. (2016) developed a framework using relative margins to unify the analysis of several variants of multiclass support vector machines. This notion of relative margin⁷ is first introduced by Rosset et al. (2003, §4). Both of these prior works only established classification-calibration of specific losses for specific algorithms. By contrast, our work develops a framework that characterizes when multiclass losses can be expressed via relative margins and proves general classification-calibration results for a large family of losses.

• **Label encodings for multiclass classification.** As alluded to earlier, the multiclass SVM introduced by Lee et al. (2004) is classification-calibrated provided that certain sum-to-zero constraints are enforced. Moreover, enforcing these constraints is computationally prohibitive. The *simplex code* (Hill and Doucet, 2007) is a multiclass label code designed to address this computational issue by using a constraint-free reparametrization of the multiclass SVM dual formulation (Wu and Lange, 2010; Saberian and Vasconcelos, 2011; Mroueh et al., 2012; van den Burg and Groenen, 2016; Pouliot, 2018).

Another approach to label encoding is the multivector construction⁸ (Shalev-Shwartz and Ben-David, 2014, §17.7) which has been used to study the multiclass perceptron, logistic regression and SVMs (Duda et al., 2006, §5.12). Moreover, the multivector construction has been used to study the PAC learning sample complexity of multiclass linear classifiers (Daniely and Shalev-Shwartz, 2014).

While the simplex code and the multivector constructions are elegant techniques, the scope of these work are on specific algorithms and classifiers. In particular, a framework for

6. See Footnote 7 in Tewari and Bartlett (2007, Table 1).

7. Rosset et al. (2003) also uses the name “margin vector” in conflict with the later work by Zou et al. (2008) which defines a different notion of margin. Due to this confusion, we will refer to the earlier notion of Rosset et al. (2003) by “relative margins”.

8. Also known as “Kesler’s construction” which, according to Crammer and Singer (2003), is credited to (Kesler, 1961). A description of the construction can be found in Duda et al. (2006, §5.12) and Shalev-Shwartz and Ben-David (2014, §17.7).

general multiclass margin-based losses and analysis of their properties, e.g., classification-calibration, is lacking.

Our work fills this gap by using the framework of the relative-margin form and matrix label code to 1. characterize the set of multiclass losses expressible in this form (Theorem 2.5), 2. prove sufficient conditions for classification-calibration of a large family of multiclass losses (Theorems 4.7) and 3. apply our results to expand the previously known family of classification-calibrated Fenchel-Young losses (Theorem 5.10).

1.4 Notations

In this subsection, we briefly discuss some of the key notations. For the reader’s convenience, we include a more comprehensive reference for the notations in Table 1 in Section A of the appendix. Moreover, we tabulate the mathematical objects defined in Table 2.

Throughout this work, let $k \geq 2$ denote the number of classes. Denote the k -probability simplex by $\Delta^k = \{\mathbf{p} \in \mathbb{R}_{\geq 0}^k : \sum_{j=1}^k p_j = 1\}$.

Vectors/matrices. Let the square bracket with subscript $[\cdot]_j$ be the projection of a vector onto its j -th component, i.e., $[\mathbf{v}]_j := v_j$ where $\mathbf{v} = (v_1, \dots, v_k) \in \mathbb{R}^k$. Given two vectors $\mathbf{w}, \mathbf{v} \in \mathbb{R}^k$, we write $\mathbf{w} \succeq \mathbf{v}$ (resp. $\mathbf{w} \succ \mathbf{v}$) if $w_j \geq v_j$ (resp. $w_j > v_j$) for all $j \in [k]$. All-zeros/all-ones/ i -th elementary basis vector in \mathbb{R}^n are denoted $\mathbb{0}^{(n)}/\mathbb{1}^{(n)}/\mathbb{e}_i^{(n)}$, respectively. When the ambient dimension is clear, we drop the superscript (n) . The $n \times n$ identity matrix is denoted \mathbf{I}_n .

Permutations and permutation matrices. A bijection from $[k]$ to itself is called a *permutation* (on $[k]$). Denote by $\text{Sym}(k)$ the set of all permutations on $[k]$. For each $\sigma \in \text{Sym}(k)$, let \mathbf{S}_σ denote the permutation matrix corresponding to σ . In other words, if $\mathbf{v} \in \mathbb{R}^k$ is a vector, then $[\mathbf{S}_\sigma \mathbf{v}]_j = [\mathbf{v}]_{\sigma(j)} = v_{\sigma(j)}$.

2. Permutation equivariant relative margin (PERM) losses

In this section, we define PERM losses and prove fundamental properties used throughout the rest of the work. We begin with a discussion of the two defining properties: permutation equivariance and relative-margins.

Recall that in the introduction, we denoted loss functions as $\mathcal{L} : [k] \times \mathbb{R}^k \rightarrow \mathbb{R}$. For a class-score function $f : \mathcal{X} \rightarrow \mathbb{R}^k$, the loss is evaluated as $\mathcal{L}(y, f(x))$ on an instance (x, y) . Throughout the rest of this work, for mathematical convenience we will use the equivalent notation convention $\mathcal{L}_y(f(x))$ where y appears in the subscript. In this convention, a multiclass loss function is a vector-valued function $\mathcal{L} = (\mathcal{L}_1, \dots, \mathcal{L}_k) : \mathbb{R}^k \rightarrow \mathbb{R}^k$.

With this convention, the property that the loss treats each of the k classes equally is captured by *permutation equivariance*, a notion similar to but distinct from symmetry⁹. Denote by \mathbf{v} the class-score vector $f(x)$ on some generic instance (x, y) . The permutation equivariance property states that if the class-score vector \mathbf{v} is “relabelled” by some permutation $\sigma \in \text{Sym}(k)$, then the vector of losses $\mathcal{L}(\mathbf{v})$ should be “relabelled” in the same way, i.e., $\mathcal{L}(\mathbf{S}_\sigma \mathbf{v}) = \mathbf{S}_\sigma \mathcal{L}(\mathbf{v})$.

9. Recall that a function is symmetric if its value does not change when its input is permuted. Sometimes, this is also referred to as permutation *invariance* (Bronstein et al., 2021, §3.1). For our purposes, *equivariance* is the correct descriptor.

Relative margins¹⁰ have been recently used by Dogan et al. (2016) and by Fathony et al. (2016) in the context of multiclass SVMs that only utilize the set of differences $v_y - v_j$ over all $y, j \in [k]$ such that $y \neq j$, rather than the non-relative or “absolute” class-scores (v_1, \dots, v_k) themselves. Below, it will be notationally convenient to use a matrix that converts the vector of class-score \mathbf{v} into relative margins:

Definition 2.1 Let $\mathbf{D} := [-\mathbf{I}_{k-1} \ \mathbf{1}^{(k-1)}] \in \mathbb{R}^{(k-1) \times k}$. Observe that $[\mathbf{D}\mathbf{v}]_y = v_k - v_y$ for $y \in [k-1]$ and $\mathbf{v} \in \mathbb{R}^k$. Equivalently, $\mathbf{D}\mathbf{v} = (v_k - v_1, v_k - v_2, \dots, v_k - v_{k-1})^\top$.

We are now ready to define

Definition 2.2 (PERM losses) Let $k \geq 2$ be an integer. A k -ary multiclass loss function is a vector-valued function $\mathcal{L} = (\mathcal{L}_1, \dots, \mathcal{L}_k) : \mathbb{R}^k \rightarrow \mathbb{R}^k$. We say that \mathcal{L} is

1. permutation equivariant if $\mathcal{L}(\mathbf{S}_\sigma \mathbf{v}) = \mathbf{S}_\sigma \mathcal{L}(\mathbf{v})$ for all $\mathbf{v} \in \mathbb{R}^k$ and $\sigma \in \text{Sym}(k)$,
2. relative margin-based if for each $y \in [k]$ there exists a function $\ell_y : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$ so that

$$\mathcal{L}_y(\mathbf{v}) = \ell_y(\mathbf{D}\mathbf{v}) = \ell_y(v_k - v_1, v_k - v_2, \dots, v_k - v_{k-1}), \quad \text{for all } \mathbf{v} \in \mathbb{R}^k. \quad (1)$$

We refer to the vector-valued function $\ell := (\ell_1, \dots, \ell_k)$ as the reduced form of \mathcal{L} .

3. PERM if \mathcal{L} is both permutation equivariant and relative margin-based. In this case, the function $\psi := \ell_k$ is referred to as the template of \mathcal{L} .

While the concepts of permutation equivariance and relative margin have been studied largely in isolation in previous works, our work is the first to systematically study the properties of losses having both properties. Below in Theorem 2.5, we show that a PERM loss is completely determined by its template via what we call the *relative-margin form*.

Remark 2.3 (On the name “template”) The symbol of the template ψ is chosen intentionally to match that of Bartlett et al. (2006), where a (binary) margin loss is expressed as $\psi((-1)^y g(x))$. Treating $g(x) = z$ as an arbitrary input to ψ , the rationale behind the name is that the “positive branch” $\psi((-1)^2 z)$, of the margin loss serves as a “template” for the “negative branch” $\psi((-1)^1 z)$.

Before proceeding, let us examine some notable PERM losses:

Example 1 The cross entropy (also multinomial logistic) loss is given by

$$\mathcal{L}_y^{\text{CE}}(\mathbf{v}) = \log \left(1 + \sum_{j \in [k]: j \neq y} \exp(-(v_y - v_j)) \right), \quad \text{for all } \mathbf{v} \in \mathbb{R}^k.$$

It is easy to see that its template is the function $\psi^{\text{CE}}(\mathbf{z}) = \log \left(1 + \sum_{j=1}^{k-1} \exp(-z_j) \right)$. When $k = 2$, we have $\psi^{\text{CE}}(z) = \log(1 + \exp(-z))$ which is the binary logistic loss (also known as the binary cross entropy). When $k = 3$, we have $\psi^{\text{CE}}(\mathbf{z}) = \log(1 + \exp(-z_1) + \exp(-z_2))$ is a function defined over the 2D plane \mathbb{R}^2 , plotted in Figure 1.

10. The idea of relative margins goes back to Rosset et al. (2003). We note that Jebara and Shivaswamy (2008) also use this term in a unrelated context in *binary* SVMs.

Example 2 The Gamma-Phi loss (Beijbom et al., 2014), denoted $\mathcal{L}_y^{\gamma, \phi}(\mathbf{v})$, is a generalization of the cross entropy loss defined as follows: Let $\gamma : \mathbb{R} \rightarrow \mathbb{R}$ and $\phi : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ be functions. Define

$$\mathcal{L}_y^{\gamma, \phi}(\mathbf{v}) := \gamma \left(\sum_{j \in [k]: j \neq y} \phi(v_j - v_y) \right), \quad \text{for all } \mathbf{v} \in \mathbb{R}^k.$$

It is easy to see that its template is the function $\psi^{\gamma, \phi}(\mathbf{z}) = \gamma \left(\sum_{j=1}^{k-1} \phi(z_j) \right)$. The cross entropy \mathcal{L}^{CE} is the Gamma-Phi loss where $\phi(t) = \exp(-t)$ and $\gamma(t) = \log(1+t)$. The multiclass exponential (Mukherjee and Schapire, 2013) loss \mathcal{L}^{Exp} is the Gamma-Phi loss where $\phi(t) = \exp(-t)$ and γ is the identity.

Example 3 A notable case of the Gamma-Phi loss is when $\phi(t) = \max\{0, 1-t\}$ is the hinge loss and γ is the identity. The resulting loss is known as the Weston-Watkins hinge loss (Weston and Watkins, 1998; Bredensteiner and Bennett, 1999; Vapnik, 1998), which is well-known to be not classification-calibrated (Liu, 2007; Tewari and Bartlett, 2007). However, the Weston-Watkins hinge loss is calibrated with respect to a discrete loss called the “ordered partition loss” related to ranking with ties (Wang and Scott, 2020).

Example 4 The Crammer-Singer hinge loss (Crammer and Singer, 2001) is a well-known loss that is a PERM loss but not a Gamma-Phi loss. It is defined as

$$\mathcal{L}_y^{\text{CS}}(\mathbf{v}) := \max_{j \in [k]: j \neq y} \{\max\{0, 1 - (v_j - v_y)\}\}, \quad \text{for all } \mathbf{v} \in \mathbb{R}^k.$$

Section 5 features another example of a family of PERM losses, namely, the Fenchel-Young losses (Blondel et al., 2020).

2.1 Matrix label code and the relative margin form

This section introduces the multiclass generalization of $\{\pm 1\}$: the *matrix label code* $\{\mathbf{\Upsilon}_y\}_{y=1}^k$.

Definition 2.4 (Matrix label code) For $k \geq 2$ and $y \in [k]$, define the $(k-1) \times (k-1)$ matrix $\mathbf{\Upsilon}_y$ as follows: For $y = k$, $\mathbf{\Upsilon}_k := \mathbf{I}_{k-1}$. For $y \in [k-1]$, define $\mathbf{\Upsilon}_y$ column-wise by

$$[\mathbf{\Upsilon}_y]_{:j} := \begin{cases} \mathbf{e}_j^{(k-1)} & : j \neq y \\ -\mathbf{1}^{(k-1)} & : j = y, \end{cases} \quad \text{for each } j \in [k-1].$$

Equivalently, to construct $\mathbf{\Upsilon}_y$ for each $y \in [k-1]$, first take the identity matrix \mathbf{I}_{k-1} , then replace the y -th column by all -1 's.

Note that $\mathbf{\Upsilon}_y^2 = \mathbf{I}_{k-1}$ for all $y \in [k-1]$, i.e., $\mathbf{\Upsilon}_y$ is an *involution*. Moreover, in the binary case where $k = 2$, we have $\mathbf{\Upsilon}_1 = -1$ and $\mathbf{\Upsilon}_2 = 1$, i.e., the matrix label code reduces to the label encoding $\{\pm 1\}$.

Next, recall that a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is *symmetric* if $f(\mathbf{S}_\sigma(\cdot)) = f(\cdot)$ for all $\sigma \in \text{Sym}(n)$. We now state the main property of the PERM loss and the matrix label code:

Theorem 2.5 (Relative-margin form) Let $\mathcal{L} : \mathbb{R}^k \rightarrow \mathbb{R}^k$ be a PERM loss with template ψ , and let $\mathbf{v} \in \mathbb{R}^k$ and $y \in [k]$ be arbitrary. Then ψ is a symmetric function. Moreover,

$$\mathcal{L}_y(\mathbf{v}) = \psi(\mathbf{\Upsilon}_y \mathbf{D} \mathbf{v}). \tag{2}$$

Conversely, let $\psi : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$ be a symmetric function. Define a multiclass loss function $\mathcal{L} = (\mathcal{L}_1, \dots, \mathcal{L}_k) : \mathbb{R}^k \rightarrow \mathbb{R}^k$ according to eq. (2). Then \mathcal{L} is a PERM loss with template ψ .

Equation (2) is referred to as the *relative margin form* of \mathcal{L} . Theorem 2.5 and other results in this section are proved in Section C in the appendix.

Remark 2.6 *The relative margin form can be interpreted as a one-to-one correspondence between PERM losses $\mathcal{L} : \mathbb{R}^k \rightarrow \mathbb{R}^k$ and symmetric functions $\psi : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$. Thus, we can refer to a PERM loss by either \mathcal{L} or ψ without ambiguity.*

Remark 2.7 (Uniqueness of the matrix label code) *In Section J of the appendix, we prove¹¹ a unique-ness result: any set of matrices $\{\mathbf{\Upsilon}'_y\}_{y=1}^k$ satisfying the conclusion of Theorem 2.5 is equal to the matrix label code (Definition 2.4) up to row permutations of the $\mathbf{\Upsilon}_y$'s.*

Remark 2.8 *Observe that $\mathbf{\Upsilon}_y \mathbf{D}\mathbf{v} = (v_y - v_1, v_y - v_2, \dots, v_y - v_k)^\top \in \mathbb{R}^{k-1}$ where the $v_y - v_y$ entry is omitted. This is Lemma B.2 in the appendix. Now, note that the right hand side is exactly the relative margin as defined in Rosset et al. (2003, Eqn. (8)). The advantage of the left hand expression $\mathbf{\Upsilon}_y \mathbf{D}\mathbf{v}$ is that the label y acts via its label encoding $\mathbf{\Upsilon}_y$ by left matrix multiplication. This “disentanglement” of the label encoding and the loss can facilitate calculations and will be crucial in our key Lemma G.6.*

Remark 2.9 (Prior work) *The matrices in the definition of the matrix label code have been used by Wang and Scott (2020, 2021) for analyzing the Weston-Watkins (WW) SVMs (Weston and Watkins, 1998). Wang and Scott (2020) show that the hinge loss from the WW-SVM is calibrated with respect to the ordered partition loss, and uses this theory to explain the empirical observation made by Dogan et al. (2016) that WW-SVM performs well even under significant label noise. Wang and Scott (2021) derive a reparametrization of the Weston-Watkins dual problem that decomposes into subproblems that can be solved exactly in $O(k \log(k))$ time, leading to faster performance for the linear WW-SVMs with many classes. The “disentanglement” discussed in the previous remark is the key ingredient for deriving this reparametrization.*

This work extends Wang and Scott (2020, 2021) beyond the multiclass SVM setting, and provides a framework for using these matrices for margin-based losses.

2.2 Relationship between class-score functions and discriminant functions

As mentioned in the introduction, the “sign” function converts a real-valued discriminant function $g : \mathcal{X} \rightarrow \mathbb{R}$ to a discrete-valued classifier $h : \mathcal{X} \rightarrow \{1, 2\}$ by choosing $h(x)$ to be $y \in \{1, 2\}$ such that $(-1)^y g(x) \geq 0$ with ties broken arbitrarily. To generalize this to the multiclass case, first recall that when $k = 2$, the matrix label code is equal to $\mathbf{\Upsilon}_y = (-1)^y$. Thus, $(-1)^y g(x) \geq 0$ iff $\mathbf{\Upsilon}_y g(x) \geq 0$. More simply put, y is the predicted class if and only if $\mathbf{\Upsilon}_y$ is the sign of $g(x)$.

Next, we define the multiclass “sign” function analogously. Given a vector-valued discriminant function $g : \mathcal{X} \rightarrow \mathbb{R}^{k-1}$ we define a discrete-valued classifier $h : \mathcal{X} \rightarrow [k]$ by taking $y \in [k]$ such that $\mathbf{\Upsilon}_y g(x) \succeq \mathbf{0}$. Similar to the binary case, ties can be broken arbitrarily. Thus, $\mathbf{\Upsilon}_y$ can be viewed as the sign of $g(x)$. To relate this multiclass “sign” function method for producing a classifier to the conventional “arg max” method, we have

11. We omit the proofs of some of the technical intermediate lemmas. For the omitted proofs, see the arXiv version of our manuscript (Wang and Scott, 2023b).

Proposition 2.10 *Let $\mathbf{v} \in \mathbb{R}^k$, $\mathbf{z} := \mathbf{D}\mathbf{v}$, and $y \in [k]$. Then $y \in \arg \max_j v_j$ iff $\mathbf{\Upsilon}_y \mathbf{z} \succeq \mathbf{0}$. Furthermore, let $\mathbf{v}' := [0 \ -\mathbf{z}^\top]^\top$. Then $y \in \arg \max_j v_j$ iff $y \in \arg \max_j v'_j$.*

Proof For the first part, recall from Remark 2.8 that $\mathbf{\Upsilon}_y \mathbf{D}\mathbf{v} = (v_y - v_1, v_y - v_2, \dots, v_y - v_k)^\top$ where the $v_y - v_y$ entry is omitted. Therefore, the assertion “ $y \in \arg \max_j v_j$ iff $\mathbf{\Upsilon}_y \mathbf{z} \succeq \mathbf{0}$ ” follows immediately from the definition of \succeq and \mathbf{z} . For the “Furthermore” part, note that by construction we have $\mathbf{v}' = (0, v_2 - v_1, v_3 - v_1, \dots, v_k - v_1)^\top$. By adding the constant v_1 to all entries of \mathbf{v}' , we recover \mathbf{v} . Thus, the indices maximizing \mathbf{v}' are the same as those of $\mathbf{v} = (v_1, v_2, \dots, v_k)^\top$. \blacksquare

The first part of Proposition 2.10 gives an intuitive explanation for our earlier definition of a multiclass “sign”. The second part of Proposition 2.10 gives a simple formula for computing the “arg max” from the discriminant function.

Remark 2.11 *Proposition 2.10 can be seen as an equivalence between class-score functions and discriminant functions. Given a class-score function $f : \mathcal{X} \rightarrow \mathbb{R}^k$, we can derive a discriminant function $g : \mathcal{X} \rightarrow \mathbb{R}^{k-1}$ by defining $g(x) := \mathbf{D}f(x)$ for all $x \in \mathcal{X}$. Conversely, we can view the discriminant function g as the “given”, and derive a class-score function f by defining $f(x) := [0 \ -g(x)]^\top$ for all $x \in \mathcal{X}$. In both cases, $\arg \max f(x) = \arg \max [0 \ -g(x)]^\top$.*

3. Classification-calibration and Consistency

This section is a brief review of the core definitions and theorem from Tewari and Bartlett (2007) in preparation for our results on sufficient conditions for classification-calibration.

Below, we assume the probability setting introduced earlier in Section 1.2.

Definition 3.1 (Range and its convex hull) *Let $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a function. Denote by $\text{ran}(f) := \{f(x) : x \in \mathbb{R}^m\}$ the range of f , and $\text{ran}_c(f) := \text{conv}(\text{ran}(f))$ the convex hull of the range of f .*

Definition 3.2 (Tewari and Bartlett (2007)) *A set $S \subseteq \mathbb{R}_+^k$ is classification-calibrated if there exists a function¹² $\theta : \mathbb{R}^k \rightarrow [k]$ such that*

$$\inf\{\langle \mathbf{p}, \zeta \rangle : \zeta \in S : p_{\theta(\zeta)} < \max \mathbf{p}\} > \inf_{\zeta \in S} \langle \mathbf{p}, \zeta \rangle \quad (3)$$

for all $\mathbf{p} \in \Delta^k$. A multiclass loss function \mathcal{L} is classification-calibrated if $\text{ran}_c(\mathcal{L})$ is classification-calibrated.

Intuitively, Definition 3.2 says that the lowest achievable conditional risk when predicting the wrong label (Eqn. (3) LHS) is still strictly larger than the conditional Bayes risk (Eqn. (3) RHS). Define $\underline{\arg \max} : \mathbb{R}^k \rightarrow [k]$ by $\underline{\arg \max}(v) = \min\{i \in [k] : v_i = \max_{j \in [k]} v_j\}$. When \mathcal{L} is permutation equivariant and classification-calibrated, we can take θ in Definition 3.2 to be simply $\underline{\arg \max}$. See Tewari and Bartlett (2007, Lemma 4). As alluded to earlier, the significance of Definition 3.2 is demonstrated by the following theorem:

12. The function θ is called a *calibrated link* for S .

Theorem 3.3 (Tewari and Bartlett 2007) *Let \mathcal{L} be a PERM loss. Let \mathcal{F} be the set of all Borel functions $\mathcal{X} \rightarrow \mathbb{R}^k$. If \mathcal{L} is classification-calibrated, then \mathcal{L} has the consistency transfer property: For all sequence of function classes $\{\mathcal{F}_n\}_n$ such that $\mathcal{F}_n \subseteq \mathcal{F}$, $\bigcup_n \mathcal{F}_n = \mathcal{F}$, all $\hat{f}_n \in \mathcal{F}_n$ and all probability distributions P on $\mathcal{X} \times [k]$*

$$R_{\mathcal{L},P}(\hat{f}_n) \xrightarrow{P} \inf_f R_{\mathcal{L},P}(f) \quad \text{implies} \quad R_{01,P}(\arg \max \circ \hat{f}_n) \xrightarrow{P} \inf_h R_{01,P}(h)$$

where the infimums are taken over all Borel functions $f : \mathcal{X} \rightarrow \mathbb{R}^k$ and $h : \mathcal{X} \rightarrow [k]$, respectively.

Remark 3.4 *In applications, \hat{f}_n is often taken to be an \mathcal{L} -risk empirical minimizer over a training dataset of cardinality n . However, the above property holds for any sequence of functions $\hat{f}_n \in \mathcal{F}_n$.*

Checking that a multiclass loss \mathcal{L} is classification-calibration is a non-trivial task, requiring delicate analysis of the loss function (Tewari and Bartlett, 2007). Recently, Wang and Scott (2023a) established a sufficient condition of classification-calibration for Gamma-Phi losses (Example 2). Below, we derive a sufficient conditions for general PERM losses.

4. Regular PERM losses

To demonstrate the utility of the relative margin form and Theorem 2.5, we prove in this section Theorem 4.7, a sufficient condition for classification-calibration (CC) of multiclass loss functions. To this end, we begin with the necessary definitions.

Definition 4.1 *A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is*

1. coercive if for all $c \in \mathbb{R}$, the set $\{\mathbf{v} \in \mathbb{R}^n : f(\mathbf{v}) \leq c\}$ (i.e., the c -sublevel set) is bounded,
2. semi-coercive if for all $c \in \mathbb{R}$ there exists $b \in \mathbb{R}$ such that

$$\{\mathbf{v} \in \mathbb{R}^n : f(\mathbf{v}) \leq c\} \subseteq \{\mathbf{v} \in \mathbb{R}^n : b \leq \min_{j \in [n]} v_j\}.$$

The definition of a coercive function is well-studied (Boyd and Vandenberghe, 2004). However, semi-coercivity appears to be new. Intuitively, a function is semi-coercive if, for all $c \in \mathbb{R}$, its c -sublevel set is contained in a translation of the positive orthant.

Definition 4.2 (Regular PERM loss) *Let \mathcal{L} be a PERM loss with template ψ . We say that \mathcal{L} is regular if ψ is nonnegative, twice differentiable, strictly convex, semi-coercive, and the gradient $\nabla_\psi(\mathbf{z}) \prec \mathbf{0}$ is entrywise negative for all $\mathbf{z} \in \mathbb{R}^{k-1}$.*

Below, we give a sufficient condition for a Gamma-Phi loss (Example 2) to be a regular PERM loss. But first, we discuss the intuition behind the condition $\nabla_\psi(\mathbf{z}) \prec \mathbf{0}$. The condition $\nabla_\psi(\mathbf{z}) \prec \mathbf{0}$ is reminiscent of a condition in Bartlett et al. (2006, Theorem 6), which shows that in the binary case a convex margin loss ψ is classification-calibrated if and only if ψ is differentiable at 0 and $\dot{\psi}(0) < 0$ where the “overdot” denotes taking derivative of a univariate function. If we view $-\dot{\psi}(\cdot)$ as a “vector field” on \mathbb{R} , then this condition on the derivative can be stated as “the vector field $-\dot{\psi}(\cdot)$ near 0 points toward the positive half of the real line”.

The condition $\nabla_\psi(\mathbf{z}) \prec \mathbf{0}$ can be thought of in a similar fashion. See Figure 2 for a plot of the vector field $-\nabla_\psi(\mathbf{z})$ where $\psi \equiv \psi^{\text{CE}}$ is the template of the cross entropy (Example 1), and where $k = 3$. Note that the vectors point toward the shaded region whose interior is the positive orthant in \mathbb{R}^2 .

However, Definition 4.2 requires the strict negativity of the gradient for all of \mathbb{R}^{k-1} , whereas the derivative of ψ is only required to be negative at 0 in the binary case in Bartlett et al. (2006, Theorem 6).

Next, we define the “truncation” of a PERM loss \mathcal{L} . Intuitively, if $k \geq 3$ is the number of classes and \mathcal{L} is k -ary, then the truncation of \mathcal{L} results in a $(k-1)$ -ary loss.

Proposition 4.3 (Truncation) *Assume $k \geq 3$. Let $\mathcal{L} : \mathbb{R}^k \rightarrow \mathbb{R}_{\geq 0}^k$ be a PERM loss with template $\psi : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$. Define $\text{trunc}[\psi](\mathbf{w}) := \lim_{\lambda \rightarrow \infty} \psi(\lambda, \mathbf{w})$ for all $\mathbf{w} \in \mathbb{R}^{k-2}$. If \mathcal{L} is regular (Definition 4.2), then $\text{trunc}[\psi]$ is a well-defined symmetric function $\mathbb{R}^{k-2} \rightarrow \mathbb{R}$ (i.e., all limits exist in \mathbb{R}), referred to as the truncation of ψ .*

Proof Let $h(\lambda) := \psi(\lambda, \mathbf{w})$. The condition that $\nabla_\psi(\cdot) \prec \mathbf{0}$ implies that h is decreasing as a function of λ . Moreover, h is nonnegative since ψ is nonnegative. Thus, $\lim_{\lambda \rightarrow +\infty} h(\lambda)$ exists. The symmetry of $\text{trunc}[\psi]$ follows immediately from the symmetry of ψ . \blacksquare

By Theorem 2.5, the symmetric function $\text{trunc}[\psi]$ induces a unique loss function which we call the *truncation* of \mathcal{L} . Note that if $\psi^{\text{CE}}(\mathbf{z}) = \log(1 + \exp(-z_1) + \exp(-z_2))$ is the template of the cross entropy with $k = 3$ (Example 1), then $\text{trunc}[\psi^{\text{CE}}](w) = \lim_{\lambda \rightarrow \infty} \psi^{\text{CE}}(\lambda, w) = \log(1 + \exp(-w))$, which is just the binary cross entropy/logistic loss.

Next, we define the m -fold iterated truncation of the template of a PERM loss.

Corollary 4.4 *For each $m \in \{0, 1, \dots, k-2\}$, define $\text{trunc}^{\times m}[\psi]$ to be m -fold repeated applications of trunc to ψ , i.e., $\text{trunc}^{\times m}[\psi] := \text{trunc}[\dots \text{trunc}[\text{trunc}[\psi]] \dots]$ where trunc appears m -times. By convention, let $\text{trunc}^{\times 0}[\psi] = \psi$. Moreover, for each $n \in \{2, \dots, k\}$, define the notational shorthand $\psi^{(n)} := \text{trunc}^{\times (k-n)}[\psi]$. If, for $n \in \{2, \dots, k-1\}$, $\psi^{(n+1)} : \mathbb{R}^n \rightarrow \mathbb{R}$ is a symmetric function such that the associated PERM loss, denoted $\mathcal{L}^{(n+1)}$, is a regular PERM loss, then $\psi^{(n)} : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$ is a symmetric function.*

The n -ary truncated loss captures the behavior of ψ when the first $k-n$ inputs to the template ψ approach $+\infty$. Note that if ψ^{CE} is the template of the k -ary cross entropy, then $\text{trunc}^{\times (k-n)}[\psi^{\text{CE}}]$ is the template of the n -ary cross entropy. This follows from a calculation similar to the one right before Corollary 4.4.

Next, we define “totally regular PERM loss” which generalizes the above observation:

Definition 4.5 (Totally regular PERM loss) *Let $\mathcal{L} : \mathbb{R}^k \rightarrow \mathbb{R}^k$ be a regular PERM loss with template $\psi : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$. For each $n \in \{2, \dots, k\}$, let $\psi^{(n)} : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$ be as in*

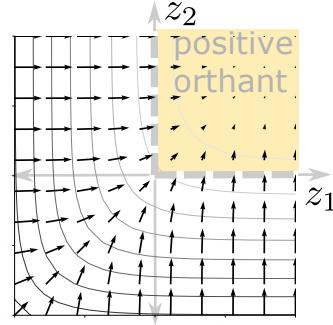


Figure 2: Vector field $-\nabla_\psi(\mathbf{z})$ where ψ is the template of the cross entropy (Example 1). The contour lines are shaded according to the value of the function ψ (darker is larger). See Figure 1.

Corollary 4.4 and let $\mathcal{L}^{(n)}$ be the unique¹³ PERM loss associated to $\psi^{(n)}$. We say that \mathcal{L} is a totally regular PERM loss if $\mathcal{L}^{(n)}$ is regular for each $n \in \{2, \dots, k\}$.

Before proving our main result, we note that Gamma Phi losses form a large family of totally regular PERM losses:

Proposition 4.6 *Let $\gamma, \phi : \mathbb{R} \rightarrow \mathbb{R}$ be functions and $\mathcal{L}^{\gamma, \phi}$ be the associated Gamma-Phi loss as defined in Example 2. Suppose that all of the following holds*

1. γ is a twice differentiable function such that $\gamma \geq 0$, $\frac{d\gamma}{dt} > 0$ and $\frac{d^2\gamma}{dt^2} \geq 0$ on all of \mathbb{R} ,
2. ϕ is a twice differentiable function such that $\frac{d\phi}{dt} < 0$, $\frac{d^2\phi}{dt^2} > 0$ on all of \mathbb{R} . Moreover, $\phi(t) \rightarrow 0$ as $t \rightarrow +\infty$.

Then $\mathcal{L}^{\gamma, \phi}$ is totally regular.

We now state our main result:

Theorem 4.7 *If \mathcal{L} is totally regular, then \mathcal{L} is classification-calibrated.*

In view of Theorem 4.7, the assumptions of Proposition 4.6 turns out to be a sufficient condition for a Gamma-Phi loss to be classification-calibrated. Theorem 4.7 and Proposition 4.6 together recover the result that the coherence loss (Zhang et al., 2009) is classification-calibrated. However, this sufficient condition is subsumed by a previous result in Wang and Scott (2023a, Theorem 3.3). Nevertheless, that $\mathcal{L}^{\gamma, \phi}$ is totally regular has nontrivial consequences as we will see next.

First, new totally regular PERM losses can be constructed from existing ones:

Proposition 4.8 *Let \mathcal{L} and $\mathcal{L}' : \mathbb{R}^k \rightarrow \mathbb{R}^k$ be totally regular PERM losses and $\lambda > 0$ be a number. Then $\lambda\mathcal{L}$ and $\mathcal{L} + \mathcal{L}'$ are also totally regular.*

Thus, $\lambda\mathcal{L}$ and $\mathcal{L} + \mathcal{L}'$ are both classification-calibrated by Theorem 4.7. Next, we apply this result to sums of Gamma-Phi losses:

Example 5 (Sum of Gamma-Phi losses) *Consider two Gamma-Phi losses each satisfying the assumptions of Proposition 4.6. For concreteness, take the cross entropy \mathcal{L}^{CE} and the multiclass exponential loss \mathcal{L}^{Exp} (See Example 2). Then by Proposition 4.8, the loss $\frac{1}{2}(\mathcal{L}^{\text{CE}} + \mathcal{L}^{\text{Exp}})$ is totally regular and thus classification-calibrated by Theorem 4.7. Note that the sum of two Gamma-Phi losses need not be another Gamma-Phi losses. Thus, Theorem 4.7 combined with Proposition 4.8 gives the first sufficient condition for when the sum of classification-calibrated losses are again classification-calibrated, in the multiclass case¹⁴.*

In the next section, we will apply Theorem 4.7 to obtain a sufficient condition for classification-calibration of Fenchel-Young losses (Theorem 5.10).

13. The existence and uniqueness is guaranteed by Theorem 2.5.

14. Note that the situation is much simpler in the binary case, due to the result of Bartlett et al. (2006, Theorem 6) that convex margin loss ψ is classification-calibrated if and only if ψ is differentiable at 0 and $\dot{\psi}(0) < 0$ where the “overdot” denotes taking derivative of a univariate function.

5. Fenchel-Young losses: sufficient conditions for classification-calibration

In this section, we consider a subset of PERM losses called the Fenchel-Young losses:

Definition 5.1 (Blondel et al. 2020) *Let $\Omega : \Delta^k \rightarrow \mathbb{R}$ be a continuous function and $\mu \in \mathbb{R}_{\geq 0}$. Define $\mathbf{c}_y := \mu(\mathbf{1}^{(k)} - \mathbf{e}_y^{(k)})$. The Fenchel-Young loss associated to Ω and μ is the loss function $\mathcal{L}^{\Omega, \mu} : \mathbb{R}^k \rightarrow \mathbb{R}_{\geq 0}^k$ whose y -th component is given by*

$$\mathcal{L}_y^{\Omega, \mu}(\mathbf{v}) := \max_{\mathbf{p} \in \Delta^k} -\Omega(\mathbf{p}) + \Omega(\mathbf{e}_y^{(k)}) + \langle \mathbf{v} + \mathbf{c}_y, \mathbf{p} - \mathbf{e}_y^{(k)} \rangle. \quad (4)$$

The maximization in Eqn. 4 is essentially the Fenchel conjugate¹⁵ of $-\Omega$. See Definition 5.4.

It is not obvious that Fenchel-Young losses are indeed PERM losses. We prove this fact in Proposition I.1 in the appendix.

Definition 5.2 (Negentropy) *A function $\Omega : \Delta^k \rightarrow \mathbb{R}$ is a negentropy if :*

1. Ω is closed (maps closed sets to closed sets) and convex,
2. Ω is symmetric, i.e., $\Omega(\mathbf{S}_\sigma \mathbf{p}) = \Omega(\mathbf{p})$ for all $\mathbf{p} \in \Delta^k$ and $\sigma \in \text{Sym}(k)$,
3. $-\Omega(\mathbf{p}) \geq 0$ for all $\mathbf{p} \in \Delta^k$ and $\Omega(\mathbf{e}_i^{(k)}) = 0$ for all $i \in [k]$.

The negative entropy is more convenient to work with due to its convexity. We use the term “negentropy” in the interest of brevity. This term has been previously used in statistics (Hyvärinen and Oja, 2000) and in machine learning (Mensch et al., 2019).

Example 6 When $\Omega(\mathbf{p}) := \sum_{y=1}^k p_y \log(p_y)$ is the negative Shannon entropy and $\mu = 0$, we have that $\mathcal{L}^{\Omega, 0}$ is the cross entropy (Example 1).

Below, we use the notation that $\tilde{\mathbf{p}} = [p_1, \dots, p_{k-1}]^\top \in \mathbb{R}^{k-1}$ is the sub-vector of \mathbf{p} with the last element dropped. Later in Proposition I.1, we show that the Fenchel-Young loss Eqn. (4) is a PERM loss with template

$$\psi^{\Omega, \mu}(\mathbf{z}) := \max_{\mathbf{p} \in \Delta^k} -\Omega(\mathbf{p}) + \mu \mathbf{1}^\top \tilde{\mathbf{p}} - \langle \tilde{\mathbf{p}}, \mathbf{z} \rangle. \quad (5)$$

Remark 5.3 Blondel et al. (2020) allow the vector $\mathbf{c}_y \in \mathbb{R}^k$ to be arbitrary, in which case the resulting loss is known as cost-sensitive Fenchel-Young loss. However, known calibration results (Blondel, 2019; Nowak-Vila et al., 2019) are only for \mathbf{c}_y as in Definition 5.1.

In order to state our sufficient condition for classification-calibration (Theorem 5.10), we briefly review two key notions from the theory of convex analysis and Legendre transformation following Rockafellar (1970, Section 26):

Definition 5.4 (Convex conjugates) *Let $D \subseteq \mathbb{R}^n$ be a closed convex set. Let $f : D \rightarrow \mathbb{R}$ be a function. Define $D^* := \{y \in \mathbb{R}^n : \sup_{x \in D} \langle y, x \rangle - f(x) < \infty\}$. The convex conjugate of a function $f : D \rightarrow \mathbb{R}$ is the function $f^* : D^* \rightarrow \mathbb{R}$ given by $f^*(y) = \sup_{x \in D} \langle y, x \rangle - f(x)$.*

Definition 5.5 (Convex functions of Legendre type) *Let $D \subseteq \mathbb{R}^n$ be a closed convex set. A convex function $f : D \rightarrow \mathbb{R}$ is said to be of Legendre type if*

15. Also known as the convex conjugate.

1. $C := \text{int}(D)$ is an open convex subset of \mathbb{R}^n ,
2. f is strictly convex and differentiable on C ,
3. for all $\{x_i\}_{i=1}^\infty \subseteq C$ with $\lim_{i \rightarrow \infty} x_i \in \text{bdry}(D)$, we have $\lim_{i \rightarrow \infty} \|\nabla f(x_i)\| = +\infty$.

For example, when $D = \Delta^k$ and $f = -H$ is the negative Shannon entropy, then $f : D \rightarrow \mathbb{R}_{\leq 0}$ is of Legendre type. See the paragraph immediately following Blondel et al. (2020, Definition 3).

Definition 5.6 (Regular negentropies) *A negentropy $\Omega : \Delta^k \rightarrow \mathbb{R}$ is a regular negentropy if Ω is of Legendre type and twice differentiable.*

The term “negentropy” was previously used by Mensch et al. (2019). To the best of our knowledge, the term “regular negentropy” is new.

Proposition 5.7 (Truncation of a negentropy) *Assume $k \geq 3$. Let $\text{pad} : \Delta^{k-1} \rightarrow \Delta^k$ be the “zero-padding” operator, i.e., $\text{pad}(\mathbf{q}) := [0, q_1, \dots, q_{k-1}] \in \Delta^k$ for all $\mathbf{q} = [q_1, \dots, q_{k-1}] \in \Delta^{k-1}$. Let $\Omega : \Delta^k \rightarrow \mathbb{R}$ be a negentropy. Define the truncation of Ω , denoted $\text{trunc}[\Omega] : \Delta^{k-1} \rightarrow \mathbb{R}$, by $\text{trunc}[\Omega](\cdot) := \Omega(\text{pad}(\cdot))$. Then $\text{trunc}[\Omega]$ is a negentropy.*

The following is analogous to the earlier iterated truncation construction in Corollary 4.4:

Corollary 5.8 *For each $m \in \{0, 1, \dots, k-2\}$, define $\text{trunc}^{\times m}[\Omega]$ to be m -fold repeated applications of trunc to Ω , i.e., $\text{trunc}^{\times m}[\Omega] := \text{trunc}[\dots \text{trunc}[\text{trunc}[\Omega]] \dots]$ where trunc appears m -times. By convention, let $\text{trunc}^{\times 0}[\Omega] = \Omega$. Moreover, for each $n \in \{2, \dots, k\}$, define the notational shorthand $\Omega^{(n)} := \text{trunc}^{\times (k-n)}[\Omega]$. Then $\Omega^{(n)} : \Delta^n \rightarrow \mathbb{R}$ is a negentropy.*

Definition 5.9 (Totally regular negentropy) *Let $\Omega : \Delta^k \rightarrow \mathbb{R}$ be a negentropy. We say that Ω is a totally regular negentropy if $\Omega^{(n)}$ is a regular negentropy for each $n \in \{2, \dots, k\}$.*

To the best of our knowledge, the definition of a totally regular negentropy is new. As suggested by the name, totally regular negentropies induce Fenchel-Young losses that are totally regular. This is a critical step¹⁶ towards proving the main result of this section below. The result leverages Theorem 4.7 to establish a sufficient condition for the classification-calibration of Fenchel-Young losses:

Theorem 5.10 *Let Ω be a totally regular negentropy, $\mu \in \mathbb{R}_{\geq 0}$ be fixed, and $\mathcal{L}^{\Omega, \mu}$ be the Fenchel-Young loss associated to Ω and μ . Then $\mathcal{L}^{\Omega, \mu}$ is classification-calibrated.*

In light of Theorem 3.3, if Ω satisfies the conditions of Theorem 5.10, then $\mathcal{L}^{\Omega, \mu}$ satisfies the consistency transfer property (Definition 1.1). Theorem 5.10 recovers all known classification-calibration sufficient conditions (Nowak-Vila et al., 2019; Blondel, 2019) which require that Ω be *strongly*-convex whereas our result only requires *strict* convexity. The following proposition and example show that there is always a nonempty set of losses to which the result applies but the prior results don’t.

¹⁶ See the proof of Theorem I.5 in the appendix.

Proposition 5.11 *Let $\Omega : \Delta^k \rightarrow \mathbb{R}$ be a totally regular negentropy and $g : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ be convex, twice differentiable and strictly increasing. Let $\mathbf{u} := (1/k)\mathbb{1}^{(k)} \in \Delta^k$ be the uniform probability vector. Define $\Theta : \Delta^k \rightarrow \mathbb{R}$ by $\Theta(\mathbf{p}) := g(\Omega(\mathbf{p}) - \Omega(\mathbf{u})) - g(-\Omega(\mathbf{u}))$ for $\mathbf{p} \in \Delta^k$. Then Θ is a totally regular negentropy. Furthermore, if $g'(0) = 0$, then Θ is not strongly-convex.*

Example 7 *For a concrete application of Proposition 5.11, take Ω to be the negative Shannon entropy (Example 6) and $g(x) = x^2$ the square function. Then Θ as defined in Proposition 5.11 is a totally regular negentropy that is not strongly-convex. Moreover, $\mathcal{L}^{\Theta, \mu}$ is classification-calibration for any $\mu \in \mathbb{R}_{\geq 0}$ by Theorem 5.10.*

Note that previous sufficient conditions (Nowak-Vila et al., 2019; Blondel, 2019) do not cover our example above.

6. Discussion and Future Directions

We proposed the relative margin form as a multiclass extension to the popular (binary) margin loss framework. We characterized the set of losses that can be expressed in the relative margin form — the PERM losses. Our Theorem 4.7 generalizes to the multiclass case the seminal result of (Bartlett et al., 2006, Theorem 6). Central to our analysis is the matrix label code (Definition 2.4) which we expect to be useful beyond the scope here. Utilizing our framework, we extended the existing set of known sufficient conditions for the classification-calibration of Fenchel-Young losses. We expect the framework to be useful for research on multiclass classification, e.g., for \mathcal{H} -consistency (Long and Servedio, 2013; Zhang and Agarwal, 2020; Awasthi et al., 2022).

- **Characterization of classification-calibration.** One weakness of our work is that, compared to the characterization of CC in the binary case (Bartlett et al., 2006, Theorem 6), we are only able to prove an analogous result (Theorem 4.7) in the multiclass case under additional assumptions. Future work will investigate whether this can be improved.
- **Cost-weighted classification.** Our results are limited to the traditional notion of classification-calibration under the assumption that the cost of misclassifying all classes are the same. The more general notion of \mathbf{c} -classification-calibration described in Williamson et al. (2016) allows the cost of misclassification to be class-dependent according to some class-specific weight vector \mathbf{c} . An interesting future direction is to consider a setting where a general \mathbf{c} -classification-calibration sufficient condition similar to our Theorem 4.7 can be derived for non-symmetric losses. More specifically, given a classification-calibrated PERM loss, is there a principled way to modify the loss to make it \mathbf{c} -classification calibrated?
- **Combining losses.** Hui et al. (2023) demonstrate the effectiveness of summing existing losses (the cross entropy and the square loss) to design new losses for classification. Unfortunately, our result in Proposition 4.8 does not apply to their new loss since the square loss is not totally regular. One approach may be to connect our technique with that of Williamson and Cranko (2023, §6) for studying the geometric properties of sums of losses.
- **Beyond classification.** While this work is concerned with classification-calibration, there are many works on calibration for more general discrete supervised learning tasks. Ramaswamy and Agarwal (2016) developed theory for multiclass classification with abstain option and, more generally, losses defined over finite sets i.e., discrete losses. Finocchiaro

et al. (2019) showed that there exists polyhedral losses that are calibrated with respect to arbitrary discrete losses. An interesting question is if the PERM loss framework can be useful for analyzing such surrogate losses. Beyond classification-calibration, a potentially fruitful direction is to employ the PERM loss framework to study optimization and learning theory.

Acknowledgments

The authors were supported in part by the National Science Foundation under awards 1838179 and 2008074, and by the Department of Defense, Defense Threat Reduction Agency under award HDTRA1-20-2-0002. YW is supported by the Eric and Wendy Schmidt AI in Science Postdoctoral Fellowship, a Schmidt Futures program.

References

Pranjal Awasthi, Anqi Mao, Mehryar Mohri, and Yutao Zhong. Multi-Class H -Consistency Bounds. In *Advances in Neural Information Processing Systems*, 2022.

Peter L Bartlett, Michael I Jordan, and Jon D McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.

Oscar Beijbom, Mohammad Saberian, David Kriegman, and Nuno Vasconcelos. Guess-averse loss functions for cost-sensitive multiclass boosting. In *International Conference on Machine Learning*, pages 586–594. PMLR, 2014.

M Beltagy and S Shenawy. On the boundary of closed convex sets in E^n . *arXiv preprint arXiv:1301.0688*, 2013.

Arnab Bhattacharyya, Suprovat Ghoshal, and Rishi Saket. Hardness of learning noisy halfspaces using polynomial thresholds. In *Conference On Learning Theory*, pages 876–917. PMLR, 2018.

David A Blackwell and Meyer A Girshick. *Theory of games and statistical decisions*. Courier Corporation, 1979.

Mathieu Blondel. Structured prediction with projection oracles. *Advances in neural information processing systems*, 32, 2019.

Mathieu Blondel, André FT Martins, and Vlad Niculae. Learning with Fenchel-Young losses. *Journal of Machine Learning Research*, 21(35):1–69, 2020.

Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

Erin J Bredensteiner and Kristin P Bennett. Multicategory classification by support vector machines. *Computational Optimization: A Tribute to Olvi Mangasarian Volume I*, pages 53–79, 1999.

Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.

Koby Crammer and Yoram Singer. On the algorithmic implementation of multiclass kernel-based vector machines. *Journal of machine learning research*, 2(Dec):265–292, 2001.

Koby Crammer and Yoram Singer. Ultraconservative online algorithms for multiclass problems. *Journal of Machine Learning Research*, 3(Jan):951–991, 2003.

Amit Daniely and Shai Shalev-Shwartz. Optimal learners for multiclass problems. In *Conference on Learning Theory*, pages 287–316. PMLR, 2014.

Ürün Dogan, Tobias Glasmachers, and Christian Igel. A unified view on multi-class support vector classification. *Journal of Machine Learning Research*, 17(45):1–32, 2016.

John Duchi, Khashayar Khosravi, and Feng Ruan. Multiclass classification, information, divergence and surrogate risk. *The Annals of Statistics*, 46(6B):3246–3275, 2018.

Richard O Duda, Peter E Hart, and David G Stork. *Pattern classification*. John Wiley & Sons, 2006.

Rizal Fathony, Anqi Liu, Kaiser Asif, and Brian Ziebart. Adversarial multiclass classification: A risk minimization perspective. *Advances in Neural Information Processing Systems*, 29, 2016.

Jessica Finocchiaro, Rafael Frongillo, and Bo Waggoner. An embedding framework for consistent polyhedral surrogates. *Advances in Neural Information Processing Systems*, 32, 2019.

Sheng Fu, Piao Chen, and Zhisheng Ye. Simplex-based proximal multicategory support vector machine. *IEEE Transactions on Information Theory*, 2022.

Simon I Hill and Arnaud Doucet. A framework for kernel-based multi-category classification. *Journal of Artificial Intelligence Research*, 30:525–564, 2007.

Like Hui, Mikhail Belkin, and Stephen Wright. Cut your losses with squentropy. In *Proceedings of the 40th International Conference on Machine Learning*, pages 14114–14131. PMLR, 2023.

Aapo Hyvärinen and Erkki Oja. Independent component analysis: algorithms and applications. *Neural Networks*, 13(4-5):411–430, 2000.

Tony Jebara and Pannagadatta Shivaswamy. Relative margin machines. *Advances in Neural Information Processing Systems*, 21, 2008.

Yuri Kalnishkan and Michael V Vyugin. The weak aggregating algorithm and weak mixability. *Journal of Computer and System Sciences*, 74(8):1228–1244, 2008.

Carl Kesler. *Preliminary Experiments on Perceptron: Application to Bubble Chamber Event Recognition*. Cornell University, 1961.

John M Lee. Smooth manifolds. In *Introduction to smooth manifolds*, pages 1–31. Springer, 2013.

Yoonkyung Lee, Yi Lin, and Grace Wahba. Multicategory support vector machines: Theory and application to the classification of microarray data and satellite radiance data. *Journal of the American Statistical Association*, 99(465):67–81, 2004.

Yufeng Liu. Fisher consistency of multicategory support vector machines. In *Artificial intelligence and statistics*, pages 291–298. PMLR, 2007.

Phil Long and Rocco Servedio. Consistency versus realizable h-consistency for multiclass classification. In *International Conference on Machine Learning*, pages 801–809. PMLR, 2013.

Arthur Mensch, Mathieu Blondel, and Gabriel Peyré. Geometric losses for distributional learning. In *International Conference on Machine Learning*, pages 4516–4525. PMLR, 2019.

Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. MIT press, 2018.

Youssef Mroueh, Tomaso Poggio, Lorenzo Rosasco, and Jean-Jacques Slotine. Multiclass learning with simplex coding. *Advances in Neural Information Processing Systems*, 25, 2012.

Indraneel Mukherjee and Robert E Schapire. A theory of multiclass boosting. *Journal of Machine Learning Research*, 2013.

Alex Nowak-Vila, Francis Bach, and Alessandro Rudi. A general theory for structured prediction with smooth convex surrogates. *arXiv preprint arXiv:1902.01958*, 2019.

Robert J Plemmons. M-matrix Characterizations. i. Nonsingular M-matrices. *Linear Algebra and its applications*, 18(2):175–188, 1977.

Guillaume Pouliot. Equivalence of multicategory SVM and simplex cone SVM: Fast computations and statistical theory. In *International conference on machine learning*, pages 4133–4140. PMLR, 2018.

Harish G Ramaswamy and Shivani Agarwal. Convex calibration dimension for multiclass loss matrices. *The Journal of Machine Learning Research*, 17(1):397–441, 2016.

R Tyrrell Rockafellar. *Convex analysis*, volume 18. Princeton University Press, 1970.

Saharon Rosset, Ji Zhu, and Trevor Hastie. Margin maximizing loss functions. *Advances in Neural Information Processing Systems*, 16, 2003.

Mohammad Saberian and Nuno Vasconcelos. Multiclass boosting: Theory and algorithms. *Advances in Neural Information Processing Systems*, 24, 2011.

Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.

Ingo Steinwart. How to compare different loss functions and their risks. *Constructive Approximation*, 26(2):225–287, 2007.

Zhiqiang Tan and Xinwei Zhang. On loss functions and regret bounds for multi-category classification. *IEEE Transactions on Information Theory*, 2022.

Ambuj Tewari and Peter L Bartlett. On the consistency of multiclass classification methods. *Journal of Machine Learning Research*, 8(5), 2007.

Gertjan van den Burg and Patrick Groenen. GenSVM: A generalized multiclass support vector machine. *Journal of Machine Learning Research*, 17:1–42, 2016.

Vladimir Vapnik. *Statistical Learning Theory*. Wiley-Interscience, 1998.

Yutong Wang and Clayton Scott. Weston-Watkins hinge loss and ordered partitions. *Advances in Neural Information Processing Systems*, 33:19873–19883, 2020.

Yutong Wang and Clayton Scott. An exact solver for the Weston-Watkins SVM subproblem. In *International Conference on Machine Learning*, pages 10894–10904. PMLR, 2021.

Yutong Wang and Clayton Scott. On classification-calibration of gamma-phi losses. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 4929–4951. PMLR, 2023a.

Yutong Wang and Clayton Scott. Unified binary and multiclass margin-based classification. *arXiv preprint arXiv:2311.17778*, 2023b.

Jason Weston and Chris Watkins. Multi-class support vector machines. Technical Report CSD-TR-98-04, Department of Computer Science, Royal Holloway, University of London, Egham, UK, 1998.

Robert C Williamson and Zac Cranko. The geometry and calculus of losses. *Journal of Machine Learning Research*, 24(342):1–72, 2023.

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

Tong Tong Wu and Kenneth Lange. Multicategory vertex discriminant analysis for high-dimensional data. *The Annals of Applied Statistics*, 4(4):1698–1721, 2010.

Mingyuan Zhang and Shivani Agarwal. Bayes consistency vs. H -consistency: The interplay between surrogate loss functions and the scoring function class. *Advances in Neural Information Processing Systems*, 33:16927–16936, 2020.

Tong Zhang. Statistical analysis of some multi-category large margin classification methods. *Journal of Machine Learning Research*, 5(Oct):1225–1251, 2004.

Zhihua Zhang, Michael Jordan, Wu-Jun Li, and Dit-Yan Yeung. Coherence functions for multicategory margin-based classification methods. In *Artificial Intelligence and Statistics*, pages 647–654. PMLR, 2009.

Hui Zou, Ji Zhu, and Trevor Hastie. New multicategory boosting algorithms based on multicategory Fisher-consistent losses. *The Annals of Applied Statistics*, 2(4):1290, 2008.

Appendix A. Notational conventions and definitions

For the reader's convenience, we tabulate the notational convention used in this work in Table 1. Moreover, we tabulate the mathematical objects defined in Table 2. Finally, we require the following additional notations:

Permutations. A bijection from $[k]$ to itself is called a *permutation* (on $[k]$). Recall we denote by $\text{Sym}(k)$ the set of all permutations on $[k]$. Two permutations $\sigma, \sigma' \in \text{Sym}(k)$ can be composed resulting in another permutation $\sigma\sigma' \in \text{Sym}(k)$ defined by function composition: for $y \in [k]$, we have $\sigma\sigma'(y) := \sigma(\sigma'(y))$. For $i, j \in [k]$, let $\tau_{(i,j)} \in \text{Sym}(k)$ denote the *transposition* which swaps i and j , leaving all other elements unchanged. More precisely, $\tau_{(i,j)}(i) = j$, $\tau_{(i,j)}(j) = i$ and $\tau_{(i,j)}(y) = y$ for $y \in [k] \setminus \{i, j\}$. Define the notational shorthand $\tau_i := \tau_{(k,i)}$, the transposition that swaps k and i .

Permutation matrices. Recall that for each $\sigma \in \text{Sym}(k)$, $\mathbf{S}_\sigma \in \mathbb{R}^{k \times k}$ denotes the permutation matrix corresponding to σ . In other words, if $\mathbf{v} \in \mathbb{R}^k$ is a vector, then $[\mathbf{S}_\sigma \mathbf{v}]_j = [\mathbf{v}]_{\sigma(j)} = v_{\sigma(j)}$. Note that if $\sigma, \sigma' \in \text{Sym}(k)$, then $\mathbf{S}_{\sigma\sigma'} = \mathbf{S}_{\sigma'} \mathbf{S}_\sigma$ where the order of compositions are reversed (Lemma B.6). Define the notational shorthand $\mathbf{T}_{(i,j)} := \mathbf{S}_{\tau_{(i,j)}}$ the matrix corresponding to the transposition of i and j . Likewise, define $\mathbf{T}_i := \mathbf{T}_{(k,i)}$. When k is ambiguous, we disambiguate it with the superscript notation and write $\mathbf{S}_{(i,j)}^{(k)}$ and $\mathbf{T}_{(i,j)}^{(k)}$.

Topology. Let S be a subset of a topological space. Let $\text{int}(S)$ and $\text{bdry}(S)$ denote the interior and the boundary of the set S , respectively. See Table 1 for the full list of symbols.

Mathematical object	Description of notation	Example of notation
Set of all permutations		$\text{Sym}(k)$
Permutations	Lower case sigma or tau	σ, τ
Transpositions	tau with subscripts	$\tau_{(i,j)}$
Vector	Bold lower case	\mathbf{v}, \mathbf{w}
Entries of vector	Normal font lower case	v_1, v_2, \dots
Special vector	Blackboard font	
All zeros/ones vector in \mathbb{R}^n		$\mathbb{0}^{(n)}, \mathbb{1}^{(n)}$
i -th elem. basis vector in \mathbb{R}^n		$\mathbf{e}_i^{(n)}$
Matrix	Bold upper case	\mathbf{A}
j -th Column		$[\mathbf{A}]_{:,j}$
$n \times n$ Identity		\mathbf{I}_n

Table 1: Notation convention used throughout this work.

Appendix B. Properties of the matrix label code

As mentioned in the main text, the definition of matrix label code has already been introduced in Wang and Scott (2020, Definition S3.14). However, that work, the matrix label code is used to analyze a particular optimization problem resulting from the Weston-

Symbol	Description	Defined in...
\mathcal{L}	Multiclass loss function	Definition 2.2
ℓ	Reduced form of \mathcal{L}	Definition 2.2 item 2
ψ	Template of \mathcal{L}	Definition 2.2 item 3
Ω	Negentropy	Definition 5.2
$\mathcal{L}^{\Omega, \mu}$	Fenchel-Young loss assoc. to Ω, μ	Definition 5.1
$\text{trunc}[\psi]$	truncation of ψ	Proposition 4.3
$\psi^{(n)}$	trunc applied $k - n$ times to ψ	Corollary 4.4
$\mathcal{L}^{(n)}$	PERM loss corresponding to $\psi^{(n)}$	Definition 4.5
$\text{trunc}[\ell]$	truncation of ℓ	Definition H.3
pad	“zero-padding” operator	Proposition 5.7
$\text{trunc}[\Omega]$	truncation of Ω	Proposition 5.7
$\Omega^{(n)}$	trunc applied $k - n$ times to Ω	Corollary 5.8
\mathbf{D}	relative margins conversion matrix	Definition 2.1
\mathbf{M}_σ	$\mathbf{D}\mathbf{S}_\sigma\mathbf{D}^\dagger$	Definition B.4
$\mathbf{\Upsilon}$	matrix label code	Definition 2.4
\mathbf{T}	Transposition matrix	Section A
τ	Transposition permutation	Section A
$\mathbf{A}(\mathbf{z})$	gradient of $\ell(\mathbf{z})$	Lemma G.6
\mathbf{P} (resp. \mathbf{Q})	drop last (resp. first) coordinate	Lemma G.25 (resp. G.26)
$\text{ran}(\cdot)$ and $\text{ran}_c(\cdot)$	range of convex hull of range	Definition 3.1
$\text{ran}_c^\bullet(f)$ and $\text{ran}_c^\circ(f)$	closure and interior of $\text{ran}_c(f)$	Definition G.27
$\mathcal{N}(\zeta; S)$	Positive normals	Definition G.35

Table 2: Mathematical objects and where they are defined.

Watkins SVM (Weston and Watkins, 1998). Here, we develop a comprehensive theory for using matrix label code with PERM loss. With the exception of Lemma B.15, all results in this section have already appeared in Wang and Scott (2020, 2021) using slightly different notation. For the reader’s convenience, the proofs are omitted here but are all included in the arXiv version of this manuscript (Wang and Scott, 2023b).

First, we note that for $y \in \{1, \dots, k - 1\}$, the matrix $\mathbf{\Upsilon}_y^{(k)}$ acts on a vector $\mathbf{z} = (z_1, \dots, z_{k-1})^\top \in \mathbb{R}^{k-1}$ by

$$\forall j \in [k-1], [\mathbf{\Upsilon}_y^{(k)} \mathbf{z}]_j := \begin{cases} z_j - z_y & : j \neq y \\ -z_y & : j = y. \end{cases} \quad (6)$$

Remark B.1 Throughout this section, we consider k fixed and write $\mathbf{\Upsilon}_1, \dots, \mathbf{\Upsilon}_k$ instead of $\mathbf{\Upsilon}_1^{(k)}, \dots, \mathbf{\Upsilon}_k^{(k)}$. Moreover, whenever \mathbf{T}_y does not have a superscript, we implicitly assume that $\mathbf{T}_y = \mathbf{T}_y^{(k)}$. In some results, we will work with $(k-1) \times (k-1)$ permutation matrices in which case we will write, for example, $\mathbf{T}_{(y,y')}^{(k-1)}$ where $y, y' \in [k-1]$.

Lemma B.2 For all $y \in [k]$, we have $\mathbf{DT}_y = \mathbf{\Upsilon}_y \mathbf{D}$. In particular, for all $i > 1$ and $j \in [k-1]$, we have

$$[\mathbf{\Upsilon}_y \mathbf{Dv}]_j = \begin{cases} v_y - v_j & : y \neq j \\ v_y - v_k & : y = j. \end{cases}$$

Remark B.3 Let \mathbf{D}^\dagger denote the Moore-Penrose inverse of \mathbf{D} . Since \mathbf{D} contains a copy of the identity matrix, \mathbf{D} has full column rank. Hence, \mathbf{DD}^\dagger is the identity.

Definition B.4 For each $\sigma \in \text{Sym}(k)$, define the $(k-1) \times (k-1)$ matrix $\mathbf{M}_\sigma := \mathbf{DS}_\sigma \mathbf{D}^\dagger$.

Lemma B.5 For all $y \in [k]$, we have $\mathbf{M}_{\tau_y} = \mathbf{\Upsilon}_y$.

Lemma B.6 If $\sigma, \sigma' \in \text{Sym}(k)$, then $\mathbf{S}_{\sigma\sigma'} = \mathbf{S}_{\sigma'} \mathbf{S}_\sigma$.

Lemma B.7 For all $\sigma, \sigma' \in \text{Sym}(k)$, we have $\mathbf{M}_{\sigma\sigma'} = \mathbf{M}_{\sigma'} \mathbf{M}_\sigma$.

Lemma B.8 For all $y \in [k]$, $\mathbf{\Upsilon}_y^2$ is the identity.

Lemma B.9 Let $y_1, y_2 \in [k-1]$ be distinct, i.e., $y_1 \neq y_2$. Then $\tau_{y_1} \tau_{y_2} \tau_{y_1} = \tau_{(y_1, y_2)}$, as elements of $\text{Sym}(k)$. Moreover, $\mathbf{T}_{y_1} \mathbf{T}_{y_2} \mathbf{T}_{y_1} = \mathbf{T}_{(y_1, y_2)}$, as elements of $\mathbb{R}^{k \times k}$.

Corollary B.10 Every $\sigma \in \text{Sym}(k)$ can be written as a product $\sigma = \tau_{y_1} \tau_{y_2} \cdots \tau_{y_n}$ for some integer $n \geq 0$ and $y_i \in [k-1]$ for each $i \in [n]$.

Lemma B.11 For all $y_1, y_2 \in [k-1]$, we have $\mathbf{T}_{(y_1, y_2)}^{(k-1)} \mathbf{D} = \mathbf{D} \mathbf{T}_{(y_1, y_2)}^{(k)}$.

Lemma B.12 Let $y_1, y_2 \in [k-1]$ be distinct. Then $\mathbf{T}_{(y_1, y_2)}^{(k-1)} = \mathbf{\Upsilon}_{y_1} \mathbf{\Upsilon}_{y_2} \mathbf{\Upsilon}_{y_1}$.

Lemma B.13 Let $y, j \in [k]$. Then

$$\mathbf{\Upsilon}_{\tau_y(j)} = \begin{cases} \mathbf{T}_{(j, y)}^{(k-1)} \mathbf{\Upsilon}_j \mathbf{\Upsilon}_y & : \text{both } y \text{ and } j \in [k-1] \\ \mathbf{\Upsilon}_j \mathbf{\Upsilon}_y & : \text{otherwise.} \end{cases} \quad (7)$$

Proposition B.14 For an arbitrary $\sigma \in \text{Sym}(k-1)$, let $\sigma' \in \text{Sym}(k)$ denote “the permutation that extends σ to $[k]$ ”, i.e., $\sigma'(k) := k$ and $\sigma'(y) := \sigma(y)$ for $y \in [k-1]$. Then we have $\mathbf{S}_\sigma \mathbf{D} = \mathbf{D} \mathbf{S}_{\sigma'}$.

Lemma B.15 Let $\mathbb{1} := \mathbb{1}^{(k-1)}$ denote¹⁷ the $(k-1)$ -dimensional (column) vector of all ones. Let $\mathbf{e}_y := \mathbf{e}_y^{(k-1)}$ denote the $(k-1)$ -dimensional y -th elementary basis (column) vector. Let $y, j \in [k]$ be such that $y \neq j$. Then we have the following identities

$$\mathbf{\Upsilon}_y - \mathbf{T}_{(y, j)} \mathbf{\Upsilon}_j = (\mathbb{1} + \mathbf{e}_j)(\mathbf{e}_j - \mathbf{e}_y)^\top \quad \text{if } y, j \in [k-1], \quad (8)$$

$$\mathbf{\Upsilon}_y - \mathbf{\Upsilon}_j = (\mathbb{1} + \mathbf{e}_j)(\mathbf{e}_j)^\top \quad \text{if } y = k, \quad (9)$$

$$\mathbf{\Upsilon}_y - \mathbf{\Upsilon}_j = -(\mathbb{1} + \mathbf{e}_y)(\mathbf{e}_y)^\top \quad \text{if } j = k. \quad (10)$$

¹⁷ We drop the superscript $(k-1)$ for convenience.

Appendix C. Proof of Theorem 2.5

We first recall the definition of a template from Theorem 2.5. For the reader's convenience, we restate it as a definition below:

Lemma C.1 *Let \mathcal{L} be a PERM loss with template ψ . Let \mathbf{D} be as in Definition 2.1. Then for all $\mathbf{v} \in \mathbb{R}^k$ we have $\mathcal{L}_k(\mathbf{v}) := [\mathcal{L}(\mathbf{v})]_k = \psi(\mathbf{D}\mathbf{v})$.*

Proof Recall from Definition 2.1 that $\mathbf{D}\mathbf{v} = (v_k - v_1, v_k - v_2, \dots, v_k - v_{k-1})^\top$. Thus, the result now follows immediately from Definition 2.2. \blacksquare

For the reader's convenience, the following is a restatement of Theorem 2.5.

Theorem C.2 *Let $\mathcal{L} : \mathbb{R}^k \rightarrow \mathbb{R}^k$ be a PERM loss with template ψ , and let $\mathbf{v} \in \mathbb{R}^k$ and $y \in [k]$ be arbitrary. Then ψ is a symmetric function and \mathcal{L} can be expressed as*

$$\mathcal{L}_y(\mathbf{v}) = \psi(\mathbf{T}_y \mathbf{D}\mathbf{v}). \quad (11)$$

Conversely, given a symmetric function $\psi : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$, the function $\mathcal{L} = (\mathcal{L}_1, \dots, \mathcal{L}_k) : \mathbb{R}^k \rightarrow \mathbb{R}^k$ defined componentwise via eq. (11) is a PERM loss with template ψ .

Proof [Proof of Theorem C.2] We begin with the “ \implies ” direction of the proof. First, we check that $\psi : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$ is symmetric, i.e., $\psi(\mathbf{S}_\sigma \mathbf{z}) = \psi(\mathbf{z})$ for all $\mathbf{z} \in \mathbb{R}^{k-1}$ and all $\sigma \in \text{Sym}(k-1)$. Below, fix such \mathbf{z} and σ .

Pick $\mathbf{v} \in \mathbb{R}^k$ such that $\mathbf{D}\mathbf{v} = \mathbf{z}$. For instance, let $\mathbf{v} := [-\mathbf{z}^\top \ 0]^\top$. Define $\sigma' \in \text{Sym}(k)$ to be the permutation that extends σ to $[k]$ as in Proposition B.14. Then

$$\begin{aligned} \psi(\mathbf{S}_\sigma \mathbf{z}) &= \psi(\mathbf{S}_\sigma \mathbf{D}\mathbf{v}) = \psi(\mathbf{D}\mathbf{S}_{\sigma'} \mathbf{v}) && \because \text{Proposition B.14} \\ &= [\mathcal{L}(\mathbf{S}_{\sigma'} \mathbf{v})]_k && \because \text{Definition 2.2 of the template } \psi \text{ of a PERM loss} \\ &= [\mathbf{S}_{\sigma'} \mathcal{L}(\mathbf{v})]_k && \because \mathcal{L} \text{ is permutation equivariant (Definition 2.2)} \\ &= [\mathcal{L}(\mathbf{v})]_{\sigma'(k)} && \because \mathbf{S}_{\sigma'} \text{ is the permutation matrix of } \sigma' \\ &= [\mathcal{L}(\mathbf{v})]_k && \because \text{Definition of } \sigma' \text{ in Proposition B.14} \\ &= \psi(\mathbf{D}\mathbf{v}). && \because \text{Lemma C.1} \end{aligned}$$

Since by definition $\psi(\mathbf{D}\mathbf{v}) = \psi(\mathbf{z})$, we have by the above that $\psi(\mathbf{S}_\sigma \mathbf{z}) = \psi(\mathbf{z})$ for all \mathbf{z} . This proves that ψ is symmetric. Next, we prove eq. (11), i.e., $[\mathcal{L}(\mathbf{v})]_y = \psi(\mathbf{T}_y \mathbf{D}\mathbf{v})$ for all $y \in [k]$. Again, this is a straight forward computation:

$$\begin{aligned} [\mathcal{L}(\mathbf{v})]_y &= [\mathcal{L}(\mathbf{v})]_{\tau_y(k)} = [\mathbf{T}_y \mathcal{L}(\mathbf{v})]_k && \because \text{definitions of } \tau_y \text{ and } \mathbf{T}_y \\ &= [\mathcal{L}(\mathbf{T}_y \mathbf{v})]_k && \because \mathcal{L} \text{ is permutation equivariant} \\ &= \psi(\mathbf{D}\mathbf{T}_y \mathbf{v}) && \because \text{Lemma C.1} \\ &= \psi(\mathbf{T}_y \mathbf{D}\mathbf{v}) && \because \text{Lemma B.2.} \end{aligned}$$

This proves eq. (11) and thus the “ \implies ” direction follows. The proof of the “ \Leftarrow ” direction, i.e., the “Conversely” part of the theorem, proceeds similarly. For the reader's convenience, it is omitted here but is included in the arXiv version of this manuscript (Wang and Scott, 2023b). \blacksquare

Lemma C.3 Let $\mathcal{L} : \mathbb{R}^k \rightarrow \mathbb{R}^k$ be a function such that $\mathbf{T}_y \mathcal{L}(\mathbf{v}) = \mathcal{L}(\mathbf{T}_y \mathbf{v})$ for all $y \in [k]$ and $\mathbf{v} \in \mathbb{R}^k$. Then \mathcal{L} is permutation-equivariant, i.e., $\mathcal{L}(\mathbf{S}_\sigma \mathbf{v}) = \mathbf{S}_\sigma \mathcal{L}(\mathbf{v})$ for all $y \in [k]$ and $\mathbf{v} \in \mathbb{R}^k$.

Proof For an arbitrary $\sigma \in \text{Sym}(k)$, write $\sigma = \tau_{y_1} \tau_{y_2} \cdots \tau_{y_n}$ as a product of transposition where $n \geq 0$ is an integer and $y_i \in [k-1]$ for each $i \in [n]$. This is possible because of Corollary B.10. By Lemma B.6, we have $\mathbf{S}_\sigma = \mathbf{T}_{y_n} \cdots \mathbf{T}_{y_2} \mathbf{T}_{y_1}$. Thus

$$\mathcal{L}(\mathbf{S}_\sigma \mathbf{v}) = \mathcal{L}(\mathbf{T}_{y_n} \cdots \mathbf{T}_{y_1} \mathbf{v}) = \mathbf{T}_{y_n} \mathcal{L}(\mathbf{T}_{y_{n-1}} \cdots \mathbf{T}_{y_1} \mathbf{v}) = \cdots = \mathbf{T}_{y_n} \cdots \mathbf{T}_{y_1} \mathcal{L}(\mathbf{v}) = \mathbf{S}_\sigma \mathcal{L}(\mathbf{v})$$

where “ $= \cdots =$ ” denotes repeated application of $\mathcal{L}(\mathbf{T}_y \mathbf{v}) = \mathbf{T}_y \mathcal{L}(\mathbf{v})$ for all $y \in [k]$. \blacksquare

Remark C.4 For convenience, we now summarize the relationship between \mathcal{L} , ℓ and ψ from Definition 2.2 in terms of the input to ψ , typically denoted below by $\mathbf{z} \in \mathbb{R}^{k-1}$. Let $y \in [k]$ and $\mathbf{z} \in \mathbb{R}^{k-1}$ be arbitrary. Define $\mathbf{v} \in \mathbb{R}^k$ be such that $\mathbf{D}\mathbf{v} = \mathbf{z}$, where \mathbf{D} is as defined in Definition 2.1. Since $\mathbf{D}\mathbf{D}^\dagger$ is the identity (Remark B.3), we can for instance let $\mathbf{v} = \mathbf{D}^\dagger \mathbf{z}$. Then we have

$$\ell_y(\mathbf{z}) = [\mathcal{L}(\mathbf{v})]_y = \psi(\mathbf{T}_y \mathbf{z}). \quad (12)$$

The first equality is eq. (1) (rewritten using $\mathbf{D}\mathbf{v} = \mathbf{z}$) and the second equality is eq. (11). A useful corollary of the above identity is the following:

$$\mathbf{S}_\sigma \ell(\mathbf{z}) = \ell(\mathbf{M}_\sigma \mathbf{z}), \quad \text{for all } \sigma \in \text{Sym}(k). \quad (13)$$

To prove eq. (13), we let $\mathbf{v} := \mathbf{D}^\dagger \mathbf{z}$ as discussed earlier. Then

$$\mathbf{S}_\sigma \ell(\mathbf{z}) \stackrel{1}{=} \mathbf{S}_\sigma \mathcal{L}(\mathbf{v}) \stackrel{2}{=} \mathcal{L}(\mathbf{S}_\sigma \mathbf{v}) \stackrel{3}{=} \ell(\mathbf{D}\mathbf{S}_\sigma \mathbf{v}) \stackrel{4}{=} \ell(\mathbf{D}\mathbf{S}_\sigma \mathbf{D}^\dagger \mathbf{z}) \stackrel{5}{=} \ell(\mathbf{M}_\sigma \mathbf{z}).$$

For 1 and 3, we used eq. (11). For 2, we used permutation equivariance of \mathcal{L} . For 4, we used the definition of \mathbf{v} . For 5, we used Definition B.4 of \mathbf{M}_σ .

Appendix D. Well-incentivized losses

This section discusses the “well-incentivized” property of a loss that should be satisfied for a multiclass classification loss to be useful. This will be made clear after Remark D.2.

Definition D.1 (Well-incentivized losses) Let \mathcal{L} be a multiclass loss function (Definition 2.2). We say that \mathcal{L} is well-incentivized (resp. strictly well-incentivized) if for all $\mathbf{v} \in \mathbb{R}^k$ and distinct $y, j \in [k]$, $v_j \leq v_y$ (resp. $v_j < v_y$) implies $\mathcal{L}_j(\mathbf{v}) \geq \mathcal{L}_y(\mathbf{v})$ (resp. $\mathcal{L}_j(\mathbf{v}) > \mathcal{L}_y(\mathbf{v})$).

Remark D.2 Suppose that \mathcal{L} is well-incentivized. Let $\mathbf{v} \in \mathbb{R}^k$ and $y \in \arg \max_{j \in [k]} v_j$. Then $\mathcal{L}_y(\mathbf{v}) = \min_{j \in [k]} \mathcal{L}_j(\mathbf{v})$.

Remark D.2 implies that for well-incentivized losses, the $\arg \max$ predictor is correct, i.e., $\mathcal{L}_y(\mathbf{v})$ is minimized when the class score $[\mathbf{v}]_y$ is the highest. Next, we define a condition on the template ψ such that the corresponding PERM loss is well-incentivized:

Definition D.3 A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is non-increasing (resp. decreasing) if for all $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ such that $\mathbf{v} \succeq \mathbf{w}$ (resp. $\mathbf{v} \succ \mathbf{w}$) we have $f(\mathbf{v}) \leq f(\mathbf{w})$ (resp. $f(\mathbf{v}) < f(\mathbf{w})$).

Proposition D.4 Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a continuously differentiable function. If $\nabla f(\cdot) \preceq \mathbb{0}$ (resp. $\nabla f(\cdot) \prec \mathbb{0}$) then f is non-increasing (resp. decreasing).

Proposition D.5 Let \mathcal{L} be a PERM loss whose template ψ is non-increasing (resp. decreasing). Then \mathcal{L} is (resp. strictly) well-incentivized.

The proofs of Propositions D.4 and D.5 are straightforward calculations involving the Fundamental Theorem of Calculus. For the reader's convenience, they are omitted here but are included in the arXiv version of this manuscript (Wang and Scott, 2023b).

Appendix E. Proof of Proposition 4.6

The proof is a straightforward calculation involving the chain rule and properties of the functions ϕ and γ . For the reader's convenience, it is omitted here but is included in the arXiv version of this manuscript (Wang and Scott, 2023b).

Appendix F. Proof of Proposition 4.8

We first prove that $\mathcal{L} + \mathcal{L}'$ is regular. Let ψ and ψ' be the templates of \mathcal{L} and \mathcal{L}' respectively. Then $\mathcal{L} + \mathcal{L}'$ has template $\psi + \psi'$. All conditions in Definition 4.2 clearly holds, except for the semi-coercive property which we now check. Let $c, b \in \mathbb{R}$ be arbitrary and satisfy

$$\{\mathbf{z} \in \mathbb{R}^n : \psi(\mathbf{z}) \leq c\} \subseteq \{\mathbf{z} \in \mathbb{R}^n : b \leq \min_{j \in [n]} z_j\}.$$

Since ψ and ψ' are both non-negative, we have $\psi(\mathbf{z}) + \psi'(\mathbf{z}) \leq c$ implies $\psi(\mathbf{z}) \leq c$. Thus

$$\{\mathbf{z} \in \mathbb{R}^n : \psi(\mathbf{z}) + \psi'(\mathbf{z}) \leq c\} \subseteq \{\mathbf{z} \in \mathbb{R}^n : \psi(\mathbf{z}) \leq c\} \subseteq \{\mathbf{z} \in \mathbb{R}^n : b \leq \min_{j \in [n]} z_j\},$$

which shows that $\psi + \psi'$ is semi-coercive, as desired. The fact that $\mathcal{L} + \mathcal{L}'$ is *totally* regular follows immediately from the definition of the truncation in Proposition 4.3. That $\lambda\mathcal{L}$ is totally regular is completely straightforward and thus we omit its proof.

Appendix G. Properties of Regular PERM losses

In this section, we will prove several key properties of regular PERM losses which were introduced in Definition 4.2. Recall that a regular PERM loss has a template ψ such that ψ is nonnegative, twice differentiable, strictly convex, semi-coercive, and the gradient $\nabla_\psi(\mathbf{z}) \prec \mathbb{0}$ is entrywise negative for all $\mathbf{z} \in \mathbb{R}^{k-1}$. Section G.1 focuses on consequences of the semi-coercivity condition. Sections G.2 and G.3 focus on consequences of the other aforementioned conditions. Finally, Section H presents the proof of our main result Theorem 4.7.

G.1 Semi-coercive functions

In this section, we study the properties of a PERM loss \mathcal{L} whose template ψ is semi-coercive (Definition 4.1).

Lemma G.1 *Let $\mathcal{L} : \mathbb{R}^k \rightarrow \mathbb{R}_+^k$ be a PERM loss whose template ψ is semi-coercive. Let ℓ be the reduced form of \mathcal{L} . Then, for all $\zeta \in \mathbb{R}^k$, the set $\{\mathbf{z} \in \mathbb{R}^{k-1} : \ell(\mathbf{z}) \preceq \zeta\}$ is bounded.*

Proof Below, let \mathbf{z} denote an arbitrary element of \mathbb{R}^{k-1} . In set-builder notations, we write $\{\mathbf{z} : \text{condition}\}$ to mean $\{\mathbf{z} \in \mathbb{R}^{k-1} : \text{condition}\}$. Now, by eq. (12), we have that $\ell(\mathbf{z}) \preceq \zeta$ if and only if $\psi(\mathbf{Y}_y \mathbf{z}) \leq \zeta_y$ for all $y \in [k]$. Thus, we get

$$\{\mathbf{z} : \ell(\mathbf{z}) \preceq \zeta\} = \bigcap_{y \in [k]} \{\mathbf{z} : \psi(\mathbf{Y}_y \mathbf{z}) \leq \zeta_y\} = \bigcap_{y \in [k]} \mathbf{Y}_y(\{\mathbf{z} : \psi(\mathbf{z}) \leq \zeta_y\}) \quad (14)$$

For the last equality above, we used the fact that $\mathbf{Y}_y = \mathbf{Y}_y^{-1}$ (Lemma B.8) and that

$$\{\mathbf{z} : \psi(\mathbf{Y}_y \mathbf{z}) \leq \zeta_y\} = \mathbf{Y}_y^{-1}(\{\mathbf{z} : \psi(\mathbf{z}) \leq \zeta_y\})$$

for all $y \in [k]$. Moreover, by the semi-coercivity assumption on ψ we have that for all $y \in [k]$ there exists $b_y \in \mathbb{R}$ such that $\{\mathbf{z} : \psi(\mathbf{z}) \leq \zeta_y\} \subseteq \{\mathbf{z} : \min \mathbf{z} \geq b_y\}$. Putting it all together, we have

$$\{\mathbf{z} : \ell(\mathbf{z}) \preceq \zeta\} \subseteq \bigcap_{y \in [k]} \mathbf{Y}_y(\{\mathbf{z} : \min \mathbf{z} \geq b_y\}) =: B \quad (15)$$

Thus, it suffices to show that B is bounded. Below, we prove this.

Since the empty set is bounded, we assume below that B is nonempty. Let $\mathbf{z}' \in B \subseteq \mathbb{R}^{k-1}$ be a fixed arbitrary point. First, recall the infinity-norm: $\|\mathbf{z}'\|_\infty := \max\{|z'_i| : i \in [k-1]\}$. To prove that B is bounded, it suffices to show that there exists some number M such that $\|\mathbf{z}'\|_\infty \leq M$. Note that we can express $\|\mathbf{z}'\|_\infty$ alternatively as

$$\|\mathbf{z}'\|_\infty = \max \{|\max \mathbf{z}'|, |\min \mathbf{z}'|\}, \quad (16)$$

where $\max \mathbf{z}' := \max_{j \in [k-1]} z'_j$ and $\min \mathbf{z}' := \min_{j \in [k-1]} z'_j$. Define $M_1 = \max\{|b_y| : y \in [k]\}$ and $M_2 = \max\{|b_y + b_j| : y \in [k], j \in [k]\}$. Finally, define $M = \max\{M_1, M_2\}$. We claim that $\|\mathbf{z}'\|_\infty \leq M$.

First, we note that $\min \mathbf{z}' \geq b_k$. To see this, first recall that \mathbf{Y}_k is the identity (Definition 2.4). From eq. (15) we have $B \subseteq \mathbf{Y}_y(\{\mathbf{z} : \min \mathbf{z} \geq b_y\})$ for each $y \in [k]$. Now, recall from Definition 2.4 that \mathbf{Y}_k is the identity matrix. Thus in particular $B \subseteq \{\mathbf{z} : \min \mathbf{z} \geq b_k\}$ and so $\min \mathbf{z}' \geq b_k$ holds.

Next, let $y \in \arg \min_{j \in [k-1]} z'_j$ and thus $z'_y = \min \mathbf{z}'$. From eq. (15), we have $\mathbf{z}' \in \mathbf{Y}_y(\{\mathbf{z} : \min \mathbf{z} \geq b_y\})$. Thus, $\mathbf{Y}_y \mathbf{z}' \in \{\mathbf{z} : \min \mathbf{z} \geq b_y\}$ and in particular, $[\mathbf{Y}_y \mathbf{z}']_y \geq b_y$. Moreover, by eq. (6), we have $[\mathbf{Y}_y \mathbf{z}']_y = -z'_y = -\min \mathbf{z}'$, and thus $\min \mathbf{z}' \leq -b_y$. Combining with the result from the previous paragraph, we now have that $\min \mathbf{z}' \in [b_k, -b_y]$ and, in particular, $|\min \mathbf{z}'| \leq M_1$. Note that here, we used the fact that if $\alpha, \beta, \gamma \in \mathbb{R}$ are such that $\gamma \in [\alpha, \beta]$, then $|\gamma| \leq \max\{|\alpha|, |\beta|\}$.

Next, let $l \in \arg \max_{j \in [k-1]} z'_j$ (and y be still the same as before). First consider the case when $l = y$. Then \mathbf{z}' is a constant vector and $\|\mathbf{z}'\|_\infty = |\min \mathbf{z}'|$ in which case $\|\mathbf{z}'\|_\infty \leq M_1 \leq M$ holds.

Next, consider the case when $l \neq y$. Then we have $[\mathbf{\Upsilon}_l \mathbf{z}']_y = z'_l - z'_y = (\min \mathbf{z}') - (\max \mathbf{z}')$ by eq. (6). Similar as in the previous case, $\mathbf{z}' \in B \subseteq \mathbf{\Upsilon}_l(\{\mathbf{z} : \min \mathbf{z} \geq b_l\})$ which implies that $[\mathbf{\Upsilon}_l \mathbf{z}']_y \geq b_l$. Thus, $\max \mathbf{z}' \leq \min \mathbf{z}' - b_l \leq -(b_y + b_l)$. Furthermore, $\max \mathbf{z}' \geq \min \mathbf{z}' \geq b_k$. Thus, we've shown that $\max \mathbf{z}' \in [b_k, -(b_y + b_l)]$. This implies that $|\max \mathbf{z}'| \leq \max\{M_1, M_2\} = M$, where again we used the fact that $\gamma \in [\alpha, \beta]$ implies $|\gamma| \leq \max\{|\alpha|, |\beta|\}$, for arbitrary $\alpha, \beta, \gamma \in \mathbb{R}$. To conclude, since $|\min \mathbf{z}'| \leq M$, by Equation (16), we have $\|\mathbf{z}'\|_\infty \leq M$. \blacksquare

Proposition G.2 *Let \mathcal{L} be a nonnegative PERM loss with template ψ . If ψ is semi-coercive (Definition 4.1), then $C_p^{\mathcal{L}}$ is coercive for all $\mathbf{p} \in \text{int}(\Delta^k)$.*

Proof Let $c \in \mathbb{R}$ be arbitrary and $S := \{\mathbf{z} \in \mathbb{R}^{k-1} : C_p^{\mathcal{L}}(\mathbf{z}) = \langle \mathbf{p}, \ell(\mathbf{z}) \rangle \leq c\}$ be the c -sublevel set. To show that $C_p^{\mathcal{L}}$ is coercive, we show that $S \subseteq \mathbb{R}^{k-1}$ is a bounded set. Observe that for all $\mathbf{z} \in S$ we have $c \geq \langle \mathbf{p}, \ell(\mathbf{z}) \rangle = \sum_{y \in [k]} p_y \psi(\mathbf{\Upsilon}_y \mathbf{z}) \geq p_y \psi(\mathbf{\Upsilon}_y \mathbf{z})$ for all $y \in [k]$. We remark that the preceding inequality crucially uses $\psi(\cdot) \geq 0$. Thus, $S \subseteq \bigcap_{y \in [k]} \{\mathbf{z} \in \mathbb{R}^{k-1} : \psi(\mathbf{\Upsilon}_y \mathbf{z}) \leq c/p_y\}$, where the right hand side is a bounded set (Lemma G.1). Hence, S is also bounded. \blacksquare

G.2 The link function

In this section, we study the set of minimizers of the conditional risk of a PERM loss \mathcal{L} , i.e., the set $\arg \min_{\mathbf{z} \in \mathbb{R}^{k-1}} C_p^{\mathcal{L}}(\mathbf{z})$. When \mathcal{L} is the multinomial cross entropy (Example 1), this argmin is a singleton set for all $\mathbf{p} \in \text{int}(\Delta^k)$ and the mapping from $\text{int}(\Delta^k) \ni \mathbf{p}$ to this unique minimizer recovers the logit function.

For a general loss \mathcal{L} , this mapping is sometimes referred to as the *link function* (Nowak-Vila et al., 2019; Williamson et al., 2016). See Definition G.5 below. This section will study the properties of the link function, culminating in a sufficient condition for when the link function is a bijection (Proposition G.9).

Proposition G.3 *Let \mathcal{L} be a nonnegative PERM loss with template ψ . If ψ is convex, then $C_p^{\mathcal{L}}$ is convex for all $\mathbf{p} \in \Delta^k$. Furthermore, if ψ is strictly convex, then $C_p^{\mathcal{L}}$ is strictly convex for all $\mathbf{p} \in \text{int}(\Delta^k)$.*

Proof Recall that $C_p^{\mathcal{L}}(\mathbf{z}) = \sum_{y \in [k]} p_y \psi(\mathbf{\Upsilon}_y \mathbf{z})$ where $\mathbf{\Upsilon}_y$ is an invertible matrix by Lemma B.8. Thus, if ψ is (strictly) convex, then $\mathbf{z} \mapsto \psi(\mathbf{\Upsilon}_y \mathbf{z})$ is (strictly) convex for each $y \in [k]$. For each $\mathbf{p} \in \Delta^k$, $C_p^{\mathcal{L}}$ is a convex combination of convex function and is thus convex. Furthermore, if $\mathbf{p} \in \text{int}(\Delta^k)$, then $C_p^{\mathcal{L}}$ is a convex combination of strictly convex function and is thus strictly convex. See Boyd and Vandenberghe (2004, Section 3.2.1) for instance. \blacksquare

An easy consequence of the above result is the following:

Corollary G.4 *Let $\mathbf{p} \in \text{int}(\Delta^k)$ be arbitrary and \mathcal{L} be a nonnegative PERM loss whose ψ is semi-coercive. If ψ is convex, then the infimum $\inf_{\mathbf{z} \in \mathbb{R}^{k-1}} C_p^{\mathcal{L}}(\mathbf{z})$ is attained. Furthermore, if ψ is strictly convex, then the infimum is attained by a unique minimizer, i.e., $\arg \min_{\mathbf{z} \in \mathbb{R}^{k-1}} C_p^{\mathcal{L}}(\mathbf{z})$ is a singleton set.*

Proof By Proposition G.2, $C_{\mathbf{P}}^{\mathcal{L}}$ is coercive. By Proposition G.3, $C_{\mathbf{P}}^{\mathcal{L}}$ is strictly convex. By the Extreme Value Theorem, a continuous and coercive function has at least one global minimum. Furthermore, a strictly convex functions have at most one global minimum. For a reference of these result standards, see Boyd and Vandenberghe (2004, Section 4.2). ■

In view of Corollary G.4, we define:

Definition G.5 Let \mathcal{L} be a PERM loss whose template ψ is nonnegative, strictly convex and semi-coercive. Define the link function $\text{link}^{\mathcal{L}} : \text{int}(\Delta^k) \rightarrow \mathbb{R}^{k-1}$ by letting $\text{link}^{\mathcal{L}}(\mathbf{p})$ be the unique element of $\arg \min_{\mathbf{z} \in \mathbb{R}^{k-1}} C_{\mathbf{P}}^{\mathcal{L}}(\mathbf{z})$.

In this section, we give a sufficient condition on \mathcal{L} for $\text{link}^{\mathcal{L}}$ of Definition G.5 to be a bijection. We will need the concept of an M -matrix, which is reviewed in Section K.1.

Lemma G.6 Let $\mathcal{L} : \mathbb{R}^k \rightarrow \mathbb{R}_+^k$ be a regular PERM loss (Definition 4.2) with reduced form ℓ and template ψ . For all $\mathbf{z} \in \mathbb{R}^{k-1}$, the $(k-1) \times (k-1)$ matrix

$$\mathbf{A}(\mathbf{z}) := [\nabla_{\ell_1}(\mathbf{z}) \quad \cdots \quad \nabla_{\ell_{k-1}}(\mathbf{z})]$$

is a non-singular M -matrix with strictly positive diagonal elements. Furthermore, both $\mathbf{A}(\mathbf{z})$ and $\mathbf{A}(\mathbf{z})^\top$ are strictly monotone.

While our usage of the operator ∇ is standard, for completeness we refer the reader to the “Vector Calculus” section in the appendix of the arXiv version of this manuscript (Wang and Scott, 2023b) for our notations on the derivative operator ∇ .

The definition of a (*strictly*) monotone matrix is discussed in Section K.1.

Proof First, we compute the gradient of $\ell = (\ell_1, \dots, \ell_k) : \mathbb{R}^{k-1} \rightarrow \mathbb{R}^k$. For each $y \in [k]$, we have by definition that $\ell_y(\mathbf{z}) = \psi(\mathbf{\Upsilon}_y \mathbf{z})$ (eq. (12)). Thus, by the chain rule, we have

$$\nabla_{\ell_y}(\mathbf{z}) = \mathbf{\Upsilon}_y^\top \nabla_\psi(\mathbf{\Upsilon}_y \mathbf{z}). \quad (17)$$

Next, fix $y \in [k-1]$ and $\mathbf{z} \in \mathbb{R}^{k-1}$. Let $\mathbf{w} := \nabla_\psi(\mathbf{\Upsilon}_y \mathbf{z})$. Then by the assumption that $\nabla_\psi(\cdot) \prec \mathbf{0}$, we have \mathbf{w} is a entrywise negative column vector, i.e., $\mathbf{w} \prec \mathbf{0}$. Moreover, by eq. (17) we have $\nabla_{\ell_y}(\mathbf{z}) = \mathbf{\Upsilon}_y^\top \mathbf{w}$. Thus, for each $j \in [k-1]$, we have $[\mathbf{\Upsilon}_y^\top \mathbf{w}]_j = ([\mathbf{\Upsilon}_y]_{:j})^\top \mathbf{w}$. Recall from Definition 2.4 that

$$[\mathbf{\Upsilon}_y]_{:j} = \begin{cases} \mathbf{e}_j^{(k-1)} & : j \neq y \\ -\mathbf{1}^{(k-1)} & : j = y \end{cases}, \quad \text{which implies that} \quad [\mathbf{\Upsilon}_y^\top \mathbf{w}]_j = \begin{cases} w_j & : j \neq y \\ -\sum_{l \in [k-1]} w_l & : j = y. \end{cases}$$

In particular, $[\mathbf{\Upsilon}_y^\top \mathbf{w}]_j \leq 0$ for all $j \neq y$ which proves that $\mathbf{A}(\mathbf{z})$ is a Z-matrix. Furthermore, note that the fact $\mathbf{w} < 0$ and $[\mathbf{\Upsilon}_y^\top \mathbf{w}]_y = -\sum_{l \in [k-1]} w_l$ implies that the diagonals of $\mathbf{A}(\mathbf{z})$ are positive. Observe that $\mathbf{\Upsilon}_y^\top \mathbf{w}$ has the property that

$$|[\mathbf{\Upsilon}_y^\top \mathbf{w}]_y| = -\sum_{l \in [k-1]} w_l > -\sum_{l \in [k-1]:l \neq y} w_l = \sum_{l \in [k-1]:l \neq y} |[\mathbf{\Upsilon}_y^\top \mathbf{w}]_l|.$$

Note that the strict inequality above follows from the fact that $\mathbf{w} \prec \mathbf{0}$ and so in particular $-w_y > 0$. This proves that $\mathbf{A}(\mathbf{z})$ is strictly diagonally dominant. By Corollary K.3, we

have that $\mathbf{A}(\mathbf{z})$ is a non-singular M-matrix. For the ‘‘Furthermore’’ part, we can apply Lemma K.5 since the diagonal elements of $\mathbf{A}(\mathbf{z})$ are positive. By Corollary K.6, $\mathbf{A}(\mathbf{z})^\top$ is also strictly monotone. \blacksquare

Lemma G.7 *Let $\mathcal{L} : \mathbb{R}^k \rightarrow \mathbb{R}_+^k$ be a regular PERM loss. Let $\mathbf{z} \in \mathbb{R}^{k-1}$ and $\mathbf{p} \in \Delta^k$ be arbitrary. Then \mathbf{z} minimizes $C_{\mathbf{p}}^{\mathcal{L}}$ if and only if*

$$-p_k \nabla_\psi(\mathbf{z}) = \mathbf{A}(\mathbf{z}) [p_1 \ \cdots \ p_{k-1}]^\top \quad (18)$$

where $\mathbf{A}(\mathbf{z})$ is as in Lemma G.6. Furthermore, $C_{\mathbf{p}}^{\mathcal{L}}$ has a minimizer, then $\mathbf{p} \in \text{int}(\Delta^k)$.

Proof Proposition G.3 asserts that $C_{\mathbf{p}}^{\mathcal{L}}$ is convex. For a differentiable convex function, recall that the gradient-vanishing condition is necessary and sufficient for optimality of unconstrained optimization (Boyd and Vandenberghe, 2004). Thus, \mathbf{z} minimizes $C_{\mathbf{p}}^{\mathcal{L}}$ iff

$$\mathbf{0} = \nabla_{C_{\mathbf{p}}^{\mathcal{L}}}(\mathbf{z}) = \sum_{j \in [k]} p_j \nabla_{\ell_j}(\mathbf{z}) = p_k \nabla_\psi(\mathbf{z}) + \mathbf{A}(\mathbf{z}) [p_1 \ \cdots \ p_{k-1}]^\top. \quad (19)$$

which is equivalent to eq. (18).

Next, for the ‘‘Furthermore’’ part, first note that Lemma G.6 says $\mathbf{A}(\mathbf{z})$ is a non-singular M-matrix. Now, we first show that $p_k \neq 0$. If $p_k = 0$, then Equation (18) reduces to $\mathbf{0} = \mathbf{A}(\mathbf{z}) [p_1 \ \cdots \ p_{k-1}]^\top$. Since $\mathbf{A}(\mathbf{z})$ is non-singular, we must have $p_1 = \cdots = p_{k-1} = 0$ as well which implies that $\mathbf{p} = \mathbf{0}$. But this contradicts that $\mathbf{p} \in \Delta^k$. Thus, $p_k > 0$ and so $-p_k \nabla_\psi(\mathbf{z}) \succ 0$. From Lemma G.6, we have that $\mathbf{A}(\mathbf{z})$ is strictly monotone. Combined with eq. (18), we can conclude that $p_y > 0$ for each $y \in [k-1]$ as well. \blacksquare

The ‘‘Furthermore’’ part of Lemma G.7 immediately implies the following.

Corollary G.8 *If $\mathbf{p} \in \Delta^k \setminus \text{int}(\Delta^k)$, then $\arg \min_{\mathbf{z} \in \mathbb{R}^{k-1}} C_{\mathbf{p}}^{\mathcal{L}}(\mathbf{z}) = \emptyset$.*

Proposition G.9 *Let \mathcal{L} be a regular PERM loss. Then $\text{link}^{\mathcal{L}}$ (Definition G.5) is a bijection.*

Proof First, we prove that link is injective. Suppose that $\mathbf{p}, \mathbf{q} \in \text{int}(\Delta^k)$ are such that $\text{link}^{\mathcal{L}}(\mathbf{p}) = \text{link}^{\mathcal{L}}(\mathbf{q}) =: \mathbf{z}$. Then by eq. (18), we have that

$$-\nabla_\psi(\mathbf{z}) \mathbf{A}(\mathbf{z})^{-1} = p_k^{-1} [p_1 \ \cdots \ p_{k-1}]^\top = q_k^{-1} [q_1 \ \cdots \ q_{k-1}]^\top. \quad (20)$$

Thus, $(1 - p_k)/p_k = (p_1 + \cdots + p_{k-1})/p_k = (q_1 + \cdots + q_{k-1})/q_k = (1 - q_k)/q_k$ implies that $p_k = q_k$. Therefore, eq. (20) implies that $p_y = q_y$ for each $y \in [k-1]$ as well. Thus, $\mathbf{p} = \mathbf{q}$ which proves that link is injective.

Next, we prove that link is surjective. Pick $\mathbf{z} \in \mathbb{R}^{k-1}$. From Lemma G.6, we have that $\mathbf{A}(\mathbf{z})$ is non-singular and strictly monotone. By non-singular-ness, there exists $\mathbf{v} \in \mathbb{R}^{k-1}$ such that $-\nabla_\psi(\mathbf{z}) = \mathbf{A}(\mathbf{z})^\top \mathbf{v}$. Furthermore, since $-\nabla_\psi(\mathbf{z}) \succ 0$ and $\mathbf{A}(\mathbf{z})$ is strictly monotone, we have $\mathbf{v} \succ 0$. Define p_1, \dots, p_k by $p_k := (v_1 + \cdots + v_{k-1} + 1)^{-1}$ and $p_y := v_y p_k$ for each $y \in [k-1]$. Clearly, we have $\mathbf{p} \succ 0$. Furthermore,

$$p_1 + p_2 + \cdots + p_k = p_k(v_1 + \cdots + v_{k-1} + 1) = 1.$$

Thus, we have $\mathbf{p} \in \text{int}(\Delta^k)$. By construction, \mathbf{z} and \mathbf{p} satisfy eq. (18). Thus $\text{link}^{\mathcal{L}}(\mathbf{p}) = \mathbf{z}$ follows from the definition of $\text{link}^{\mathcal{L}}$. \blacksquare

Remark G.10 *Before proceeding, we remark that Proposition G.9 gives theoretical support to the conjectural observation in Nowak-Vila et al. (2019, Remark 3.1) regarding the injectivity of the link function.*

G.3 Geometry of the loss surface

Recall from Definition 3.1 and Theorem 3.3 that the classification-calibration of the set $\text{ran}_c(\mathcal{L})$ implies the classification-calibration of the loss \mathcal{L} . In general, the set $\text{ran}_c(\mathcal{L})$ may be difficult to compute. In this section, we study the geometry of the set $\text{ran}(\mathcal{L})$ when \mathcal{L} is a regular PERM loss which enables us to compute the convex hull $\text{ran}_c(\mathcal{L})$ of $\text{ran}(\mathcal{L})$. One of the main tools is the mapping defined below:

Corollary G.11 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ . Then $\ell : \mathbb{R}^{k-1} \rightarrow \mathbb{R}^k$ is injective.*

Proof Suppose that $\mathbf{z}, \mathbf{w} \in \mathbb{R}^{k-1}$ are such that $\ell(\mathbf{z}) = \ell(\mathbf{w})$. By Proposition G.9, there exists $\mathbf{p} \in \text{int}(\Delta^k)$ such that $\text{link}^{\mathcal{L}}(\mathbf{p}) = \mathbf{z}$. Now, $\langle \mathbf{p}, \ell(\mathbf{z}) \rangle = \langle \mathbf{p}, \ell(\mathbf{w}) \rangle$ implies that both \mathbf{z}, \mathbf{w} minimize $C_{\mathbf{p}}^{\mathcal{L}}$. By Corollary G.4, we have $\mathbf{z} = \mathbf{w}$ and so ℓ is injective. \blacksquare

Definition G.12 *Given a PERM loss \mathcal{L} with reduced form ℓ , we define two functions F and G mapping from $\mathbb{R}^{k-1} \times \mathbb{R}$ to \mathbb{R}^k by $F(\mathbf{z}, \lambda) = \ell(\mathbf{z}) + \lambda \mathbf{1}$ and $G(\mathbf{z}, t) = \ell(\mathbf{z}) + t \mathbf{e}_k^{(k)}$, where $\mathbf{z} \in \mathbb{R}^{k-1}$ and $\lambda \in \mathbb{R}$. For computing ∇F , we view the tuple \mathbf{z}, λ as a (column) vector $[\mathbf{z}^\top \ \lambda]^\top$. Hence, $\nabla F(\mathbf{z}, \lambda) = [\nabla \ell(\mathbf{z})^\top \ \mathbf{1}^\top]^\top$. Likewise for the tuple \mathbf{z}, t and ∇G .*

Remark G.13 *In the context of Definition G.12, it is perhaps more precise to write $F([\mathbf{z}^\top \ \lambda]^\top)$. However, this notation is cumbersome. Below, we will always use the tuple notation. The reason we discuss the vector notation is so that ∇F and ∇G can be more simply treated as gradients of vector input-valued functions (rather than matrix input-valued).*

Below, we will study the properties of the two functions from Definition G.12.

G.3.1 PROPERTIES OF THE F FUNCTION

The main result of this section is to prove that F is a homeomorphism from $\mathbb{R}^{k-1} \times \mathbb{R}$ to \mathbb{R}^k (Corollary G.19).

Lemma G.14 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ and F be as in Definition G.12. Then F is injective. In other words, if $\ell(\mathbf{z}) + \lambda \mathbf{1} = \ell(\mathbf{w}) + \mu \mathbf{1}$ for some $\mathbf{z}, \mathbf{w} \in \mathbb{R}^{k-1}$ and $\lambda, \mu \in \mathbb{R}$, then both $\mathbf{z} = \mathbf{w}$ and $\lambda = \mu$.*

Proof Let $\mathbf{z}, \mathbf{w} \in \mathbb{R}^{k-1}$ and $\lambda, \mu \in \mathbb{R}$ be as in the statement of the lemma. Our goal is to show that both $\mathbf{z} = \mathbf{w}$ and $\lambda = \mu$. First, consider the case that $\mathbf{z} = \mathbf{w}$. Then $\ell(\mathbf{z}) = \ell(\mathbf{w})$ and so $\lambda = \mu$. Thus, $\mathbf{z} = \mathbf{w}$ implies $\lambda = \mu$.

Next, consider the case that $\lambda = \mu$. Then we have $\ell(\mathbf{z}) = \ell(\mathbf{w})$. By Corollary G.11, we have $\mathbf{z} = \mathbf{w}$. Therefore, $\lambda = \mu$ implies $\mathbf{z} = \mathbf{w}$.

Thus, it only remains to show that if both $\mathbf{z} \neq \mathbf{w}$ and $\lambda \neq \mu$, then we have a contradiction. Without loss of generality, suppose that $\lambda > \mu$. Then we have $\ell(\mathbf{z}) + (\lambda - \mu)\mathbf{1} = \ell(\mathbf{w})$. Thus, for all $\mathbf{p} \in \Delta^k$, we have $\langle \mathbf{p}, \ell(\mathbf{w}) \rangle = \langle \mathbf{p}, \ell(\mathbf{z}) + (\lambda - \mu)\mathbf{1} \rangle > \langle \mathbf{p}, \ell(\mathbf{z}) \rangle$. Thus, there does not exist $\mathbf{p} \in \Delta^k$ such that \mathbf{w} is the minimizer of $C_{\mathbf{p}}^{\mathcal{L}}$. But this contradicts since Proposition G.9 implies that *link* is surjective. \blacksquare

Lemma G.15 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ . Let F be as in Definition G.12. Then for all $(\mathbf{z}, \lambda) \in \mathbb{R}^{k-1} \times \mathbb{R}$, $\nabla_F(\mathbf{z}, \lambda)$ is non-singular.*

Proof Recall that $\nabla_{\ell}(\mathbf{z}) \in \mathbb{R}^{(k-1) \times k}$ and $\mathbf{1}^{(k)} \in \mathbb{R}^k$. Thus, $\nabla_F(\mathbf{z}, \lambda) = [\nabla_{\ell}(\mathbf{z})^{\top} \quad (\mathbf{1}^{(k)})^{\top}]^{\top}$ is a $k \times k$ square matrix. To show that it is non-singular, first pick $\mathbf{v} \in \mathbb{R}^k$ arbitrary. It suffices to check that if $\nabla_F(\mathbf{z}, \lambda)\mathbf{v} = \mathbf{0}$ then $\mathbf{v} = \mathbf{0}$.

Towards this, first note that $\nabla_F(\mathbf{z}, \lambda)\mathbf{v} = \mathbf{0}$ can be equivalently stated as both $\nabla_{\ell}(\mathbf{z})\mathbf{v} = \mathbf{0}$ and $(\mathbf{1}^{(k)})^{\top}\mathbf{v} = v_1 + \dots + v_k = 0$. Replacing \mathbf{v} by $-\mathbf{v}$ if necessary, we assume that $v_k \geq 0$. Recall $\mathbf{A}(\mathbf{z})$ from Lemma G.6. Then the identity $\nabla_{\ell}(\mathbf{z})\mathbf{v} = \mathbf{0}$ can be rewritten as $-v_k \nabla_{\psi}(\mathbf{z}) = \mathbf{A}(\mathbf{z}) [v_1 \ \dots \ v_{k-1}]^{\top}$. Since $\mathbf{A}(\mathbf{z})$ is monotone and $-v_k \nabla_{\psi}(\mathbf{z}) \succeq \mathbf{0}$, we get that $v_y \geq 0$ for each $y \in [k-1]$. Combined with the fact that $\mathbf{v}^{\top}\mathbf{1} = v_1 + \dots + v_k = 0$, we have that $\mathbf{v} = \mathbf{0}$, as desired. \blacksquare

Now, by applying the inverse function theorem¹⁸, we immediately have the following.

Corollary G.16 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ . Let F be as in Definition G.12. For all $(\mathbf{z}, \lambda) \in \mathbb{R}^{k-1} \times \mathbb{R}$, there exist open neighborhoods $U \ni (\mathbf{z}, \lambda)$ and $V \ni F(\mathbf{z}, \lambda)$ such that $F|_U : U \rightarrow V$ is a diffeomorphism.*

Proposition G.17 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ . Let F be as in Definition G.12. The map F is a bijection.*

Proof Lemma G.14 shows that F is injective. To show that F is surjective, we prove that $\text{ran}(F)$ is both open and closed as a subset of \mathbb{R}^k . This would imply that $\text{ran}(F) = \mathbb{R}^k$ since the only subsets of \mathbb{R}^k that are both open and closed are \emptyset and \mathbb{R}^k . Now, Lemma G.16 shows that $\text{ran}(F)$ is an open subset of \mathbb{R}^k . It remains to prove that $\text{ran}(F)$ is closed.

To this end, consider a sequence $\{(\mathbf{z}^{(i)}, \lambda^{(i)})\}_{i=1}^{\infty}$ such that $F(\mathbf{z}^{(i)}, \lambda^{(i)}) = \ell(\mathbf{z}^{(i)}) + \lambda^{(i)}\mathbf{1}$ converges to $\zeta \in \mathbb{R}^k$. Our goal is to show that $\zeta \in \text{ran}(F)$.

We begin by first picking $\epsilon > 0$. Let $\mathbf{1} := \mathbf{1}^{(k)}$ (without the superscript (k)) denotes the k -dimensional vector of all ones. Since the sequence converges, there exists M such that

$$\zeta - \epsilon\mathbf{1} \preceq \ell(\mathbf{z}^{(i)}) + \lambda^{(i)}\mathbf{1} \preceq \zeta + \epsilon\mathbf{1}$$

18. The inverse function theorem is a standard result in multivariate calculus. See the arXiv version of this manuscript (Wang and Scott, 2023b) for the result's statement and a textbook reference.

for all $i \geq M$. Before proceeding, we prove a helper lemma.

Lemma G.18 (Helper lemma) *Let \mathcal{L} be a PERM loss with reduced form ℓ and template ψ such that $\nabla_\psi(\cdot) \preceq 0$. Then for all $\mathbf{z} \in \mathbb{R}^{k-1}$, we have that $\min_{y \in [k-1]} \ell_y(\mathbf{z}) \leq \psi(\mathbf{0}^{(k-1)}) =: C$, where $\mathbf{0}^{(k-1)}$ is the $(k-1)$ -dimensional all-zeros vector.*

Proof [Proof of helper lemma] Let $\mathbf{v} \in \mathbb{R}^k$ be such that $\mathbf{D}\mathbf{v} = \mathbf{z}$, where we recall from Definition 2.1 that $\mathbf{D} = [-\mathbf{I}_{k-1} \ \mathbf{1}^{(k-1)}]$. As in Remark C.4, we can take $\mathbf{v} = [-\mathbf{z}^\top \ 0]^\top$. Next, let $y \in \arg \max_{j \in [k-1]} v_j$. Recall from Section A that $\min(\cdot)$ over vector-valued inputs denotes entrywise minimum. Then by Remark D.2 and Proposition D.5, we have that $\min(\mathcal{L}(\mathbf{z})) = \mathcal{L}_y(\mathbf{v})$. Next by eq. (12) from Remark C.4, we have $[\mathcal{L}(\mathbf{v})]_y = \ell_y(\mathbf{z})$.

Let $\mathbf{w} := \sigma_y(\mathbf{v})$. (Recall from Section A that $\sigma_y \in \text{Sym}(k)$ is the transposition that swaps k and y .) Note that by construction we have $k \in \arg \max \mathbf{w}$. By permutation-equivariance, we have $[\mathcal{L}(\mathbf{v})]_y = [\mathcal{L}(\mathbf{v})]_{\sigma_y(k)} = [\mathcal{L}(\sigma_y(\mathbf{v}))]_k = [\mathcal{L}(\mathbf{w})]_k$. Again by eq. (12) from Remark C.4, we have $[\mathcal{L}(\mathbf{w})]_k = \psi(\mathbf{D}\mathbf{w})$.

Since $k \in \arg \max \mathbf{w}$, we have that $\mathbf{D}\mathbf{w} \in \mathbb{R}_{\geq 0}^{k-1}$ belongs to the non-negative orthant. In other words, $\mathbf{D}\mathbf{w} \geq \mathbf{0}^{(k-1)}$. Finally, since $\nabla_\psi(\cdot) \preceq 0$, we have that $\psi(\mathbf{D}\mathbf{w}) \leq \psi(\mathbf{0}^{(k-1)})$. ■

We now return to the proof of the proposition. Let C be as in the helper lemma. Then

$$\min(\ell(\mathbf{z}^{(i)}) + \lambda^{(i)} \mathbf{1}) = \min(\ell(\mathbf{z}^{(i)})) + \lambda^{(i)} \leq C + \lambda^{(i)}.$$

Thus, we have $\min(\zeta) - \epsilon \leq C + \lambda^{(i)}$ and which implies that $-\lambda^{(i)} \leq C + \epsilon - \min(\zeta) =: D$. From this, we get that

$$\ell(\mathbf{z}^{(i)}) \preceq \zeta + \epsilon \mathbf{1} - \lambda^{(i)} \mathbf{1} \preceq \zeta + (\epsilon + D) \mathbf{1}$$

Thus, for all $i \geq M$ we have $\mathbf{z}^{(i)} \in \{\mathbf{z} \in \mathbb{R}^{k-1} : \ell(\mathbf{z}) \preceq \zeta + (\epsilon + D) \mathbf{1}\}$ which is a bounded set by Lemma G.1. By passing to a subsequence, we may assume that $\mathbf{z}^{(i)}$ converges to some $\mathbf{z}^* \in \mathbb{R}^{k-1}$. Thus, we have $\lambda^{(i)} \mathbf{1}$ converges to $\ell(\mathbf{z}^*) + \zeta$, which implies in particular that $\lambda^{(i)}$ converges to some λ^* . Putting it all together, we have shown that $F(\mathbf{z}^{(i)}, \lambda^{(i)})$ converges to $\zeta = F(\mathbf{z}^*, \lambda^*)$ and so $\text{ran}(F)$ is closed. ■

Corollary G.19 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ . Let F be as in Definition G.12. The map F is a diffeomorphism, i.e., F is a differentiable bijection with a differentiable inverse. In particular, F is a homeomorphism.*

Proof Lee (2013, Proposition 4.6 (f)) states that every bijective local diffeomorphism is a (global) diffeomorphism. Thus, the result follows in view of the facts that F is a bijection (Proposition G.17) and that F is a local diffeomorphism (Corollary G.16). ■

Proposition G.20 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ . Let F be as in Definition G.12. Consider arbitrary $\mathbf{v}, \mathbf{x} \in \mathbb{R}^k$ and $t \in \mathbb{R}$. Define¹⁹ $\alpha(t) \in \mathbb{R}^{k-1}$ and*

¹⁹ Note that by Corollary G.19, such $\alpha(t)$ and $\beta(t)$ exist and are unique with respect to this property.

$\beta(t) \in \mathbb{R}$ to be such that $t\mathbf{v} + \mathbf{x} = F(\alpha(t), \beta(t)) = \ell(\alpha(t)) + \beta(t)\mathbf{1}^{(k)}$. Then for all $t \in \mathbb{R}$, we have

1. $\alpha : \mathbb{R} \rightarrow \mathbb{R}^{k-1}$ and $\beta : \mathbb{R} \rightarrow \mathbb{R}$ are differentiable,
2. If $\mathbf{v} \succ 0$, then $\beta'(t) = \frac{d\beta}{dt}(t) > 0$,
3. β is concave, i.e., $\beta''(t) = \frac{d^2\beta}{dt^2}(t) \leq 0$.

Remark G.21 We note that α and β in Proposition G.20 implicitly depend on \mathbf{x} and \mathbf{v} .

Proof Below, let $\mathbf{1} := \mathbf{1}^{(k)}$ be the k -dimensional vector of all ones. To prove the first part, first note that $(\alpha(t), \beta(t)) = F^{-1}(t\mathbf{v} + \mathbf{x})$. Hence, α and β are differentiable.

Next, we prove the second part. Let $y \in [k]$ be arbitrary, to be specified later. Now, the y -th coordinate of $t\mathbf{v} + \mathbf{x} = \ell(\alpha(t)) + \beta(t)\mathbf{1}$ is

$$v_y t + x_y = \ell_y(\alpha(t)) + \beta(t). \quad (21)$$

Differentiating (21) on both sides with respect to t and applying the chain rule²⁰, we get

$$v_y = \alpha'(t)^\top \nabla_{\ell_y}(\alpha(t)) + \beta'(t). \quad (22)$$

For convenience, we adopt the notation $f'(\cdot) := (\nabla_f(\cdot))^\top$ for functions with a scalar input. Now, we claim that $\alpha'(t)^\top \nabla_{\ell_y}(\alpha(t)) \leq 0$ for some $y \in [k]$. In fact, we prove this claim by proving a slightly more general statement that will be used again later.

Lemma G.22 Let \mathcal{L} be a regular PERM loss with reduced form ℓ . Let $\mathbf{z}, \mathbf{w} \in \mathbb{R}^{k-1}$ be arbitrary. Then there exists $y \in [k]$ such that $\mathbf{w}^\top \nabla_{\ell_y}(\mathbf{z}) \leq 0$.

Proof Suppose for the sake of contradiction that $\mathbf{w}^\top \nabla_{\ell_y}(\mathbf{z}) > 0$ for all $y \in [k]$. Then

$$0 \prec [\nabla_{\ell_1}(\mathbf{z}) \quad \cdots \quad \nabla_{\ell_{k-1}}(\mathbf{z}) \quad \nabla_{\ell_k}(\mathbf{z})]^\top \mathbf{w} = [\mathbf{A}(\mathbf{z}) \quad \nabla_\psi(\mathbf{z})]^\top \mathbf{w}.$$

In other words, we have $\mathbf{A}(\mathbf{z})^\top \mathbf{w} \succ 0$ and $\mathbf{w}^\top \nabla_\psi(\mathbf{z}) > 0$. By Lemma G.6, $\mathbf{A}(\mathbf{z})^\top$ is strictly monotone. Thus, $\mathbf{A}(\mathbf{z})^\top \mathbf{w} \succ 0$ implies²¹ $\mathbf{w} \succ 0$. But $\nabla_\psi(\mathbf{z}) \prec 0$ by assumption that \mathcal{L} is a regular PERM loss. Hence, $\mathbf{w}^\top \nabla_\psi(\mathbf{z}) < 0$, a contradiction. \blacksquare

Applying Lemma G.22 with $\mathbf{w} = \alpha'(t)$, we get the desired claim. Now, pick $y \in [k]$ such that $\alpha'(t)^\top \nabla_{\ell_y}(\alpha(t)) \leq 0$. Thus, from Eqn. (22) we have $\beta'(t) = v_y - \alpha'(t)^\top \nabla_{\ell_y}(\alpha(t)) > 0$. This proves the second part of Proposition G.20.

Finally, we prove the last part of Proposition G.20. Pick $y \in [k]$ (unrelated to the earlier choice) such that $\alpha''(t)^\top \nabla_{\ell_y}(\alpha(t)) \geq 0$. Such a $y \in [k]$ exists by Lemma G.22 by setting $\mathbf{w} = -\alpha''(t)$. Differentiating eq. (22) with respect to t , we get by the product rule for curves²² that

$$0 = \alpha''(t)^\top \nabla_{\ell_y}(\alpha(t)) + \alpha'(t)^\top \frac{d}{dt}(\nabla_{\ell_y}(\alpha(t))) + \beta''(t)$$

20. The chain rule is a standard result in multivariate calculus. See the arXiv version of this manuscript (Wang and Scott, 2023b) for the result's statement and a textbook reference.

21. This is the definition of "strictly monotone matrices". See Section K.1

22. See the "Vector Calculus" section in the appendix of the arXiv version of this manuscript (Wang and Scott, 2023b).

By the chain rule for curves²³, we have $\frac{d}{dt}(\nabla_{\ell_y}(\alpha(t))) = \nabla_{\ell_y}^2(\alpha(t))\alpha'(t)$ which combined with the above implies

$$-\beta''(t) = \alpha''(t)^\top \nabla_{\ell_y}(\alpha(t)) + \alpha'(t)^\top \nabla_{\ell_y}^2(\alpha(t))\alpha'(t).$$

Since ℓ_y is convex, $\nabla_{\ell_y}^2(\alpha(t))$ is positive semidefinite. Therefore, $\alpha'(t)^\top \nabla_{\ell_y}^2(\alpha(t))\alpha'(t) \geq 0$. Putting it all together, $-\beta''(t) \geq 0$, i.e., β is concave. \blacksquare

G.3.2 PROPERTIES OF THE G FUNCTION

Lemma G.23 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ . Let G be as in Definition G.12. Then for all $(\mathbf{z}, t) \in \mathbb{R}^{k-1} \times \mathbb{R}$, the gradient (matrix) $\nabla_G(\mathbf{z}, t)$ is non-singular.*

Proof The proof proceeds similarly as in Lemma G.15. Recall that $\nabla_\ell(\mathbf{z}) \in \mathbb{R}^{(k-1) \times k}$ and $\mathbf{e}_k^{(k)} \in \mathbb{R}^k$. Below, we suppress this superscript and simply write $\mathbf{e}_k := \mathbf{e}_k^{(k)}$. Thus, $\nabla_G(\mathbf{z}, \lambda) = [\nabla_\ell(\mathbf{z})^\top \ \mathbf{e}_k^\top]^\top$ is a $k \times k$ square matrix. To show that it is non-singular, first pick $\mathbf{v} \in \mathbb{R}^k$ arbitrary. It suffices to check that if $\nabla_G(\mathbf{z}, \lambda)\mathbf{v} = \mathbf{0}$ then $\mathbf{v} = \mathbf{0}$.

Towards this, first note that $\nabla_G(\mathbf{z}, \lambda)\mathbf{v} = \mathbf{0}$ can be equivalently stated as both $\nabla_\ell(\mathbf{z})\mathbf{v} = \mathbf{0}$ and $\mathbf{e}_k^\top \mathbf{v} = v_k = 0$. Recall $\mathbf{A}(\mathbf{z})$ from Lemma G.6. Then the identity $\nabla_\ell(\mathbf{z})\mathbf{v} = \mathbf{0}$ can be rewritten as $\mathbf{0} = -v_k \nabla_\psi(\mathbf{z}) = \mathbf{A}(\mathbf{z}) [v_1 \ \cdots \ v_{k-1}]^\top$. Since $\mathbf{A}(\mathbf{z})$ is non-singular, we get that $v_y = 0$ for each $y \in [k-1]$. Thus, we have that $\mathbf{v} = \mathbf{0}$, as desired. \blacksquare

Lemma G.24 *Suppose that $\mathbf{A} \in \mathbb{R}^{n \times n}$ and $\mathbf{v} \in \mathbb{R}^n$ are such that 1) \mathbf{A} is strictly monotone (Definition K.1), 2) $\mathbf{v} \prec \mathbf{0}$ has strictly negative entries, 3) the matrix $[\begin{smallmatrix} \mathbf{A} & \mathbf{v} \\ \mathbf{1}^\top & 1 \end{smallmatrix}]$ is invertible, and 4) $\mathbf{M} := [\begin{smallmatrix} \mathbf{A} & \mathbf{v} \\ \mathbf{0}^\top & 1 \end{smallmatrix}] ([\begin{smallmatrix} \mathbf{A} & \mathbf{v} \\ \mathbf{1}^\top & 1 \end{smallmatrix}])^{-1} \in \mathbb{R}^{(n+1) \times (n+1)}$ is also invertible. Then $M_{n+1, n+1} > 0$, i.e., the bottom right entry of \mathbf{M} is positive.*

Proof Define $\mathbf{B} \in \mathbb{R}^{n \times n}$, $\mathbf{w}, \mathbf{u} \in \mathbb{R}^n$ and $c \in \mathbb{R}$ such that $([\begin{smallmatrix} \mathbf{A} & \mathbf{v} \\ \mathbf{1}^\top & 1 \end{smallmatrix}])^{-1} = [\begin{smallmatrix} \mathbf{B} & \mathbf{w} \\ \mathbf{u}^\top & c \end{smallmatrix}]$. Then by definition of the matrix inverse, we have

$$\begin{bmatrix} \mathbf{I}_n & \mathbf{0} \\ \mathbf{0}^\top & 1 \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{v} \\ \mathbf{1}^\top & 1 \end{bmatrix} \begin{bmatrix} \mathbf{B} & \mathbf{w} \\ \mathbf{u}^\top & c \end{bmatrix} = \begin{bmatrix} \mathbf{AB} + \mathbf{vu}^\top & \mathbf{Aw} + c\mathbf{v} \\ \mathbf{1}^\top \mathbf{B} + \mathbf{u}^\top & \mathbf{1}^\top \mathbf{w} + c \end{bmatrix}. \quad (23)$$

Now, we observe that

$$\mathbf{M} = \begin{bmatrix} \mathbf{A} & \mathbf{v} \\ \mathbf{0}^\top & 1 \end{bmatrix} \begin{bmatrix} \mathbf{B} & \mathbf{w} \\ \mathbf{u}^\top & c \end{bmatrix} = \begin{bmatrix} \mathbf{AB} + \mathbf{vu}^\top & \mathbf{Aw} + c\mathbf{v} \\ \mathbf{u}^\top & c \end{bmatrix} = \begin{bmatrix} \mathbf{I}_n & \mathbf{0} \\ \mathbf{u}^\top & c \end{bmatrix}$$

Since \mathbf{M} is invertible by assumption, we have that $c \neq 0$. To finish the proof, it suffices to show that $c > 0$. Below, we assume that $c < 0$ and derive a contradiction.

The top right block of eq. (23) implies that $\mathbf{Aw} + c\mathbf{v} = \mathbf{0}$, i.e., $\mathbf{A}(-\mathbf{w}) = c\mathbf{v}$. Since $c < 0$ and $\mathbf{v} \prec \mathbf{0}$, we have $\mathbf{A}(-\mathbf{w}) \succ \mathbf{0}$. By strict monotonicity of \mathbf{A} , we have that

23. See the “Vector Calculus” section in the appendix of the arXiv version of this manuscript (Wang and Scott, 2023b).

$-\mathbf{w} \succ 0$, or equivalently, $\mathbf{w} \prec \mathbf{0}$. Finally, the bottom right entry of eq. (23) implies that $1 = \mathbf{1}^\top \mathbf{w} + c < 0$, which is a contradiction. \blacksquare

Lemma G.25 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ and $\mathbf{z} \in \mathbb{R}^{k-1}$ be arbitrary. Define \mathbf{P} to be the projection $\mathbb{R}^k \rightarrow \mathbb{R}^{k-1}$ that drops the last coordinate, i.e., $\mathbf{P}([v_1, \dots, v_k]^\top) = [v_1, \dots, v_{k-1}]^\top$. Then there exists $\mathbf{z}^* \in \mathbb{R}^{k-1}$ and $t^* \in \mathbb{R}_{>0}$ such that $\mathbf{P}\ell(\mathbf{z}^*) + t^* \mathbf{1} = \mathbf{P}(\ell(\mathbf{z}))$.*

Proof Define functions $\zeta : \mathbb{R}^{k-1} \times \mathbb{R} \rightarrow \mathbb{R}^{k-1}$ and $\tau : \mathbb{R}^{k-1} \times \mathbb{R} \rightarrow \mathbb{R}$ by

$$[\zeta(\mathbf{z}, t)^\top \quad \tau(\mathbf{z}, t)]^\top := F^{-1} \circ G(\mathbf{z}, t).$$

We first show that $\tau(\mathbf{z}, 0) = 0$. Observe that

$$F^{-1}(G(\mathbf{z}, 0)) = F^{-1}(\ell(\mathbf{z}) + 0 \cdot \mathbf{e}_k) = F^{-1}(\ell(\mathbf{z}) + 0 \cdot \mathbf{1}) = [\mathbf{z}^\top \quad 0]^\top. \quad (24)$$

Therefore, $\zeta(\mathbf{z}, t) = \mathbf{z}$ and $\tau(\mathbf{z}, 0) = 0$. Next, we claim that $\frac{\partial \tau}{\partial t}(\mathbf{z}, 0) \neq 0$. By definition,

$$\nabla_{F^{-1} \circ G}(\mathbf{z}, t) = \begin{bmatrix} \frac{\partial \zeta}{\partial \mathbf{z}}(\mathbf{z}, t) & \frac{\partial \tau}{\partial \mathbf{z}}(\mathbf{z}, t) \\ \frac{\partial \zeta}{\partial t}(\mathbf{z}, t) & \frac{\partial \tau}{\partial t}(\mathbf{z}, t) \end{bmatrix} \quad (25)$$

Let “.” be a notational shorthand for the input (\mathbf{z}, t) to F and G . The gradient of $F^{-1} \circ G$

$$\nabla_{F^{-1} \circ G}(\cdot) = \nabla_G(\cdot) \nabla_{F^{-1}}(G(\cdot)) = \nabla_G(\cdot) (\nabla_F(F^{-1} \circ G(\cdot)))^{-1}$$

Below, let $\mathbf{e}_k := \mathbf{e}_k^{(k)}$, i.e., we drop the superscript. Now, from the proof of Lemma G.23, we have that $\nabla_G(\mathbf{z}, t) = [\nabla_\ell(\mathbf{z})^\top \quad \mathbf{e}_k]^\top$ and from the proof of Lemma G.15 that $\nabla_F(\mathbf{z}, t) = [\nabla_\ell(\mathbf{z})^\top \quad \mathbf{1}^{(k)}]^\top$. By eq. (24), we have

$$\nabla_G(\mathbf{z}, 0) (\nabla_F(F^{-1} \circ G(\mathbf{z}, 0)))^{-1} = \nabla_G(\mathbf{z}, 0) (\nabla_F(\mathbf{z}, 0))^{-1}$$

The above is invertible because $\nabla_G(\mathbf{z}, 0)$ is invertible by Lemma G.23. Moreover,

$$\nabla_G(\mathbf{z}, 0) (\nabla_F(\mathbf{z}, 0))^{-1} = \begin{bmatrix} \nabla_\ell(\mathbf{z}) \\ \mathbf{e}_k^\top \end{bmatrix} \left(\begin{bmatrix} \nabla_\ell(\mathbf{z}) \\ (\mathbf{1}^{(k)})^\top \end{bmatrix} \right)^{-1} \quad (26)$$

$$= \begin{bmatrix} \mathbf{A}(\mathbf{z}) & \nabla_\psi(\mathbf{z}) \\ (\mathbf{0}^{(k-1)})^\top & 1 \end{bmatrix} \left(\begin{bmatrix} \mathbf{A}(\mathbf{z}) & \nabla_\psi(\mathbf{z}) \\ (\mathbf{1}^{(k-1)})^\top & 1 \end{bmatrix} \right)^{-1}. \quad (27)$$

By eq. (25), $\frac{\partial \tau}{\partial t}(\mathbf{z}, 0)$ is the bottom right element of $\nabla_G(\mathbf{z}, 0) (\nabla_F(\mathbf{z}, 0))^{-1}$. By applying Lemma G.24 to the RHS of eq. (27) with $\mathbf{A} = \mathbf{A}(\mathbf{z})$ and $\mathbf{v} = \nabla_\psi(\mathbf{z})$, we see that the bottom right entry of $\nabla_G(\mathbf{z}, 0) (\nabla_F(\mathbf{z}, 0))^{-1}$ is nonzero. This proves that $\frac{\partial \tau}{\partial t}(\mathbf{z}, 0) \neq 0$. Note that Lemmas G.6 and G.23 together guarantee that the requirements of Lemma G.24 are all met.

Next, we claim that there exists some $t^* \in \mathbb{R}$ such that $\tau(\mathbf{z}, t^*) > 0$. To see this, assume the contrary. Then $\tau(\mathbf{z}, t) \leq 0$ for all $t \in \mathbb{R}$. In particular, $t = 0$ is a (global) maximizer of $t \mapsto \tau(\mathbf{z}, t)$ which implies that $\frac{\partial \tau}{\partial t}(\mathbf{z}, 0) = 0$, a contradiction. Thus, the claim follows.

Now, fix $t^\circ \in \mathbb{R}$ such that $\tau(\mathbf{z}, t^\circ) > 0$. Define $\mathbf{z}^* := \zeta(\mathbf{z}, t^\circ)$ and $t^* = \tau(\mathbf{z}, t^\circ)$. Then by the definition of ζ and τ , we have $F(\mathbf{z}^*, t^*) = G(\mathbf{z}, t^\circ)$ by applying F to both side of

$$[(\mathbf{z}^*)^\top \quad t^*]^\top = [\zeta(\mathbf{z}, t^\circ)^\top \quad \tau(\mathbf{z}, t^\circ)]^\top = F^{-1}(G(\mathbf{z}, t^\circ)).$$

Unwinding the definition of F and G (Definition G.12), the identity $F(\mathbf{z}^*, t^*) = G(\mathbf{z}, t^\circ)$ implies

$$\ell(\mathbf{z}^*) + t^* \mathbf{1} = F(\mathbf{z}^*, t^*) = G(\mathbf{z}, t^\circ) = \ell(\mathbf{z}) + t^\circ \mathbf{e}_k.$$

Now, applying $\mathbf{P}(\cdot)$ to both side, we have $\mathbf{P}(\ell(\mathbf{z}^*) + t^* \mathbf{1}) = \mathbf{P}(\ell(\mathbf{z}) + t^\circ \mathbf{e}_k) = \mathbf{P}(\ell(\mathbf{z}))$ where we used the fact that $\mathbf{P}(\cdot)$ is linear and that $\mathbf{P}(\mathbf{e}_k) = \mathbf{0}$. \blacksquare

Corollary G.26 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ and $\mathbf{z} \in \mathbb{R}^{k-1}$ be arbitrary. Define \mathbf{Q} to be the projection $\mathbb{R}^k \rightarrow \mathbb{R}^{k-1}$ that projs the first coordinate, i.e., $\mathbf{Q}([v_1, \dots, v_k]^\top) = [v_2, \dots, v_k]^\top$. Then there exists $\mathbf{z}^* \in \mathbb{R}^{k-1}$ and $t^* \in \mathbb{R}_{>0}$ such that $\mathbf{Q}(\ell(\mathbf{z}^*) + t^* \mathbf{1}) = \mathbf{Q}\ell(\mathbf{z})$.*

Proof Recall \mathbf{P} from Lemma G.25. Let $\sigma \in \text{Sym}(k)$ be such that $\mathbf{Q} = \mathbf{P}\mathbf{S}_\sigma$. Recall \mathbf{M}_σ from Definition B.4. Then $\mathbf{S}_\sigma \ell(\mathbf{z}) = \ell(\mathbf{M}_\sigma \mathbf{z})$ by eq. (13). Applying Lemma G.25 to $\mathbf{M}_\sigma \mathbf{z}$, there exists $\tilde{\mathbf{z}} \in \mathbb{R}^{k-1}$ and $t^* \in \mathbb{R}_{>0}$ such that $\mathbf{P}\ell(\mathbf{M}_\sigma \mathbf{z}) = \mathbf{P}(\ell(\tilde{\mathbf{z}}^*) + t^* \mathbf{1})$. Putting it all together, we have

$$\mathbf{Q}\ell(\mathbf{z}) = \mathbf{P}\mathbf{S}_\sigma \ell(\mathbf{z}) = \mathbf{P}\ell(\mathbf{M}_\sigma \mathbf{z}) = \mathbf{P}(\ell(\tilde{\mathbf{z}}^*) + t^* \mathbf{1}).$$

Next, $\mathbf{Q} = \mathbf{P}\mathbf{S}_\sigma$ implies $\mathbf{Q}\mathbf{S}_{\sigma^{-1}} = \mathbf{P}$. Hence, again applying eq. (13), we have

$$\mathbf{P}(\ell(\tilde{\mathbf{z}}^*) + t^* \mathbf{1}) = \mathbf{Q}\mathbf{S}_{\sigma^{-1}}(\ell(\tilde{\mathbf{z}}^*) + t^* \mathbf{1}) = \mathbf{Q}(\ell(\mathbf{M}_{\sigma^{-1}} \tilde{\mathbf{z}}^*) + t^* \mathbf{1}).$$

Letting $\mathbf{z}^* := \mathbf{M}_{\sigma^{-1}} \tilde{\mathbf{z}}^*$, we have $\mathbf{Q}\ell(\mathbf{z}) = \mathbf{Q}(\ell(\mathbf{z}^*) + t^* \mathbf{1})$, as desired. \blacksquare

Definition G.27 *Let $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a function. Define the following sets:*

1. $\text{ran}_c^\bullet(f) := \{\zeta + \lambda \mathbf{1} : \zeta \in \text{ran}(f), \lambda \in [0, \infty)\}$
2. $\text{ran}_c^\circ(f) := \{\zeta + \lambda \mathbf{1} : \zeta \in \text{ran}(f), \lambda \in (0, \infty)\}$

When $f = \ell$ is the reduced form of a PERM loss, the above two sets are closely related to $\text{ran}_c(\ell)$ (Definition 3.1), as the following lemma and Proposition G.31 show. These sets are convenient alternative characterizations.

Lemma G.28 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ . Then we have the following:*

1. $\text{ran}(\ell)$ is closed.
2. $\text{ran}_c^\bullet(\ell)$ is closed and $\text{bdry}(\text{ran}_c^\bullet(\ell)) = \text{ran}(\ell)$.
3. $\text{ran}_c^\circ(\ell) = \text{int}(\text{ran}_c^\bullet(\ell))$ and $\text{bdry}(\text{ran}_c^\circ(\ell)) = \text{ran}(\ell)$.

Proof For item 1, define $C := \mathbb{R}^{k-1} \times \{0\}$ which is a closed subset of $\mathbb{R}^{k-1} \times \mathbb{R}$. Now, note that $\text{ran}(\ell) = F(C)$ where F is as in Definition G.12. Since F is a homeomorphism (Corollary G.19), $F(C)$ is closed as well.

For item 2, define $D := \mathbb{R}^{k-1} \times [0, \infty)$ which is a closed subset of $\mathbb{R}^{k-1} \times \mathbb{R}$. Then $\text{ran}_c^\bullet(\ell) = F(D)$. Thus, as in the previous case $\text{ran}_c^\bullet(\ell)$ is closed. Next, we have

$$\text{bdry}(\text{ran}_c^\bullet(\ell)) = \text{bdry}(F(D)) = F(\text{bdry}(D)) = F(C) = \text{ran}(\ell),$$

where the second equality from the left follows from F being a homeomorphism.

For item 3, let $E = \mathbb{R}^{k-1} \times (0, \infty)$. Then similar to the above, we have

$$\text{int}(\text{ran}_c^\bullet(\ell)) = \text{int}(F(D)) = F(\text{int}(D)) = F(E) = \text{ran}_c^\circ(\ell).$$

To conclude, note that $\text{bdry}(\text{ran}_c^\circ(\ell)) = \text{bdry}(F(E)) = F(\text{bdry}(E)) = F(C) = \text{ran}(\ell)$. \blacksquare

Proposition G.29 *Let \mathcal{L} be a regular PERM loss. Then $\text{ran}_c^\bullet(\ell)$ is convex.*

Proof Let $\zeta, \xi \in \text{ran}_c^\bullet(f)$. Write $\zeta = \ell(\mathbf{z}) + \lambda \mathbb{1}$ and $\xi = \ell(\mathbf{w}) + \mu \mathbb{1}$, where $\mathbf{z}, \mathbf{w} \in \mathbb{R}^{k-1}$ and $\lambda, \mu \in [0, \infty)$. Let $\mathbf{v} = \xi - \zeta \in \mathbb{R}^k$ and $\mathbf{x} = \zeta$. Take α and β as defined in Proposition G.20, i.e., we have for all $t \in \mathbb{R}$ that

$$t\mathbf{v} + \mathbf{x} = F(\alpha(t), \beta(t)) = \ell(\alpha(t)) + \beta(t)\mathbb{1}. \quad (28)$$

Plugging in $t = 0$ into (28), we get $\zeta = \ell(\alpha(0)) + \beta(0)$. Thus, $\alpha(0) = \mathbf{z}$ and $\beta(0) = \lambda$. Likewise, plugging in $t = 1$, we get $\alpha(1) = \mathbf{w}$ and $\beta(1) = \mu$. In particular, we have $\beta(0) \geq 0$ and $\beta(1) \geq 0$. By Proposition G.20, β is concave. Thus, $\beta(t) \geq 0$ for all $t \in [0, 1]$, i.e.,

$$t\mathbf{v} + \mathbf{w} = t\xi + (1-t)\zeta = \ell(\alpha(t)) + \beta(t)\mathbb{1} \in \text{ran}_c^\bullet(\ell)$$

for all $t \in [0, 1]$. This proves that $\text{ran}_c^\bullet(\ell)$ is convex. \blacksquare

The following result is a restatement of Beltagy and Shenawy (2013, Theorem 9):

Theorem G.30 (Beltagy and Shenawy (2013)) *Let C be a nonempty closed convex subset of \mathbb{R}^n . If C contains no hyperplane, then $C = \text{conv}(\text{bdry}(C))$.*

Proposition G.31 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ . Then $\text{ran}_c(\ell) = \text{ran}_c^\bullet(\ell)$ and $\text{bdry}(\text{ran}_c(\ell)) = \text{ran}(\ell)$.*

Proof Clearly, $\text{ran}_c^\bullet(\ell)$ is nonempty. Furthermore, by Lemma G.28 and Lemma G.29, $\text{ran}_c^\bullet(\ell)$ is closed and convex. Next, note that $\text{ran}_c^\bullet(\ell)$ lies in the nonnegative quadrant $[0, \infty)^k$. Since no hyperplane lies entirely inside the nonnegative quadrant, $\text{ran}_c^\bullet(\ell)$ cannot contain any hyperplane. Hence, we have verified that $\text{ran}_c^\bullet(\ell)$ satisfies the condition of Theorem G.30. To finish the proof, we have

$$\text{ran}_c(\ell) = \text{conv}(\text{ran}(\ell)) \quad \because \text{Definition of } \text{ran}_c(\ell) \quad (29)$$

$$= \text{conv}(\text{bdry}(\text{ran}_c^\bullet(\ell))) \quad \because \text{Lemma G.29} \quad (30)$$

$$= \text{ran}_c^\bullet(\ell) \quad \because \text{Theorem G.30} \quad (31)$$

This proves the first part. For the second part, note that $\text{bdry}(\text{ran}_c(\ell)) = \text{bdry}(\text{ran}_c^\bullet(\ell)) = \text{ran}(\ell)$ by Lemma G.29. \blacksquare

Before moving on, we summarize the results on $\text{ran}_c(\ell)$ we have thus obtained below:

Corollary G.32 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ . Recall from Definition G.27*

$$\text{ran}_c^\circ(\ell) := \{\zeta + \lambda \mathbf{1} : \zeta \in \text{ran}(\ell), \lambda \in (0, \infty)\}.$$

Then $\text{ran}_c(\ell)$ is a closed and convex set with the following properties:

1. $\text{ran}_c(\ell) = \{\zeta + \lambda \mathbf{1} : \zeta \in \text{ran}(\ell), \lambda \in [0, \infty)\}$
2. $\text{int}(\text{ran}_c(\ell)) = \text{ran}_c^\circ(\ell)$ (see Definitions 3.1 and G.27)
3. $\text{bdry}(\text{ran}_c(\ell)) = \text{bdry}(\text{ran}_c^\circ(\ell)) = \text{ran}(\ell)$.

We state one more result about the set $\text{ran}_c^\circ(\ell)$ which will be useful later.

Lemma G.33 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ . Then $\text{ran}_c^\circ(\ell) = \{\zeta \in \mathbb{R}^k : \exists \mathbf{z} \in \mathbb{R}^{k-1} \text{ such that } \zeta \succ \ell(\mathbf{z})\}$.*

Proof Recall that by definition we have $\text{ran}_c^\circ(\ell) = \{\ell(\mathbf{z}) + \lambda \mathbf{1} : \mathbf{z} \in \mathbb{R}^{k-1}, \lambda \in (0, \infty)\}$. Thus, the “ \subseteq ” direction is immediate. For the other inclusion, take $\zeta = \ell(\mathbf{z}) + \mathbf{v}$ where $\mathbf{z} \in \mathbb{R}^{k-1}$ and $\mathbf{v} \succ 0$. Let $\mathbf{x} = \ell(\mathbf{z})$. Take α and β as defined in Proposition G.20, i.e., we have for all $t \in \mathbb{R}$ that

$$t\mathbf{v} + \mathbf{x} = F(\alpha(t), \beta(t)) = \ell(\alpha(t)) + \beta(t)\mathbf{1}. \quad (32)$$

Plugging in $t = 0$ into (32), we get $\mathbf{x} = \ell(\mathbf{z}) = \ell(\alpha(0)) + \beta(0)$. Thus, $\alpha(0) = \mathbf{z}$ and $\beta(0) = 0$. Recall that $\mathbf{v} = \zeta - \ell(\mathbf{z}) \succ 0$ by assumption. Hence, by Proposition G.20, β is strictly increasing. In particular, $\beta(1) > \beta(0) = 0$. Now, plugging in $t = 1$ into (32), we get $\mathbf{v} + \mathbf{x} = \mathbf{v} + \ell(\mathbf{z}) = \zeta = \ell(\alpha(1)) + \beta(1)\mathbf{1}$. This shows that $\zeta \in \text{ran}_c^\circ(\ell)$, as desired. \blacksquare

Remark G.34 *From basic point-set topology²⁴, we know that a closed set is the union of its interior and its boundary. Thus, $\text{ran}_c(\ell) = \text{int}(\text{ran}_c(\ell)) \cup \text{bdry}(\text{ran}_c(\ell))$. Hence, a consequence of Lemma G.33 and Lemma G.28 is that $\text{ran}_c(\ell)$ is precisely the superprediction set of ℓ (see Williamson et al. (2016, Definition 15) and Kalnishkan and Vyugin (2008)):*

$$\text{ran}_c(\ell) = \{\zeta \in \mathbb{R}^k : \exists \mathbf{z} \in \mathbb{R}^{k-1} \text{ such that } \zeta \succeq \ell(\mathbf{z})\}.$$

Recall (Tewari and Bartlett, 2007, Definition 5):

Definition G.35 (Admissible sets; Tewari and Bartlett (2007)) *Let $S \subseteq \mathbb{R}_+^k$ be a set and $\zeta \in \mathbb{R}_+^k$. Define the set $\mathcal{N}(\zeta; S) := \{\mathbf{p} \in \Delta^k : \langle \zeta - \zeta, \mathbf{p} \rangle \geq 0, \forall \xi \in S\}$. We say that S is admissible if for all $\zeta \in \text{bdry}(S)$ and $\mathbf{p} \in \mathcal{N}(\zeta; S)$ we have $\arg \min(\zeta) \subseteq \arg \max(\mathbf{p})$.*

24. The boundary of a set A is defined as the closure of A minus the interior of A .

Proposition G.36 (Tewari and Bartlett (2007)) *Let $S \subseteq \mathbb{R}_+^k$ be a symmetric set. If $|\mathcal{N}(\zeta; S)| = 1$ for all $\zeta \in \text{bdry}(S)$, then S is admissible.*

Lemma G.37 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ . Let $\zeta \in \text{ran}(\ell)$. Then we have $\mathcal{N}(\zeta; \text{ran}(\ell)) = \mathcal{N}(\zeta; \text{ran}_c(\ell)) = \mathcal{N}(\zeta; \text{ran}_c^\circ(\ell))$.*

Proof We first prove that $\mathcal{N}(\zeta; \text{ran}(\ell)) = \mathcal{N}(\zeta; \text{ran}_c(\ell))$. Since $\text{ran}_c(\ell) \supseteq \text{ran}(\ell)$, we immediately have $\mathcal{N}(\zeta; \text{ran}(\ell)) \supseteq \mathcal{N}(\zeta; \text{ran}_c(\ell))$. For the other inclusion, we first note that $\text{ran}_c(\ell) = \text{ran}_c^\bullet(\ell)$ by Proposition G.31. Thus, every $\xi \in \text{ran}_c(\ell)$ can be written as $\xi = \alpha + \beta \mathbb{1}$ for some $\alpha \in \text{ran}(\ell)$ and $\beta \geq 0$. Now, let $\mathbf{p} \in \mathcal{N}(\zeta; \text{ran}(\ell))$ and let $\xi \in \text{ran}_c(\ell)$ be decomposed as in the preceding sentence. Then

$$\langle \xi - \zeta, \mathbf{p} \rangle = \langle \alpha + \beta \mathbb{1} - \zeta, \mathbf{p} \rangle = \langle \alpha - \zeta, \mathbf{p} \rangle + \beta \langle \mathbb{1}, \mathbf{p} \rangle \geq 0$$

where the last inequality holds since (1) $\langle \alpha - \zeta, \mathbf{p} \rangle \geq 0$ because $\mathbf{p} \in \mathcal{N}(\zeta; \text{ran}(\ell))$, and (2) $\beta \geq 0$. Hence, such a \mathbf{p} satisfies $\langle \xi - \zeta, \mathbf{p} \rangle \geq 0, \forall \xi \in \text{ran}_c(\ell)$ as well which implies that $\mathbf{p} \in \mathcal{N}(\zeta; \text{ran}_c(\ell))$, as desired.

Next, we prove $\mathcal{N}(\zeta; \text{ran}_c(\ell)) = \mathcal{N}(\zeta; \text{ran}_c^\circ(\ell))$. Again, since $\text{ran}_c(\ell) \supseteq \text{ran}_c^\circ(\ell)$, we immediately have $\mathcal{N}(\zeta; \text{ran}_c^\circ(\ell)) \supseteq \mathcal{N}(\zeta; \text{ran}_c(\ell))$. For the other inclusion, we first note that $\text{clos}(\text{ran}_c^\circ(\ell)) = \text{ran}_c(\ell)$. Suppose $\mathbf{p} \in \Delta^k$ is such that $\langle \xi - \zeta, \mathbf{p} \rangle \geq 0$ for all $\xi \in \text{ran}_c^\circ(\ell)$. Then by continuity, we must have that $\langle \xi - \zeta, \mathbf{p} \rangle \geq 0$ for all $\xi \in \text{clos}(\text{ran}_c^\circ(\ell)) = \text{ran}_c(\ell)$. ■

Proposition G.38 *Let \mathcal{L} be a regular PERM loss with reduced form ℓ . Then $\text{ran}_c(\ell)$ and $\text{ran}_c^\circ(\ell)$ are both admissible.*

Proof By Proposition G.36, it suffices to check the following two claims hold:

1. for all $\zeta \in \text{bdry}(\text{ran}_c(\ell))$ we have $|\mathcal{N}(\zeta; \text{ran}_c(\ell))| = 1$, and
2. for all $\zeta \in \text{bdry}(\text{ran}_c^\circ(\ell))$ we have $|\mathcal{N}(\zeta; \text{ran}_c^\circ(\ell))| = 1$.

By Corollary G.32, we have $\text{bdry}(\text{ran}_c(\ell)) = \text{bdry}(\text{ran}_c^\circ(\ell)) = \text{ran}(\ell)$. Hence, by Lemma G.37, to show both above claims it suffices to show that $|\mathcal{N}(\zeta; \text{ran}(\ell))| = 1$ for all $\zeta \in \text{ran}(\ell)$. Note that here we can replace $\mathcal{N}(\zeta; \text{ran}_c(\ell))$ and $\mathcal{N}(\zeta; \text{ran}_c^\circ(\ell))$ by $\mathcal{N}(\zeta; \text{ran}(\ell))$ because of Lemma G.37. Below, fix $\zeta = \ell(\mathbf{z}) \in \text{ran}(\ell)$ where $\mathbf{z} \in \mathbb{R}^{k-1}$ is arbitrary. Then

$$\begin{aligned} \mathcal{N}(\zeta; \text{ran}(\ell)) &= \{\mathbf{p} \in \Delta^k : \langle \xi - \zeta, \mathbf{p} \rangle \geq 0, \forall \xi \in \text{ran}(\ell)\} \quad \because \text{Definition G.35} \\ &= \{\mathbf{p} \in \Delta^k : \langle \ell(\mathbf{w}) - \ell(\mathbf{z}), \mathbf{p} \rangle \geq 0, \forall \mathbf{w} \in \mathbb{R}^{k-1}\} \quad \because \text{Corollary G.11} \\ &= \left\{ \mathbf{p} \in \Delta^k : \mathbf{z} \in \arg \min_{\mathbf{w} \in \mathbb{R}^{k-1}} C_{\mathbf{p}}^{\mathcal{L}}(\mathbf{w}) \right\} \quad \because \text{definition of arg min} \\ &= \left\{ \mathbf{p} \in \text{int}(\Delta^k) : \mathbf{z} \in \arg \min_{\mathbf{w} \in \mathbb{R}^{k-1}} C_{\mathbf{p}}^{\mathcal{L}}(\mathbf{w}) \right\} \quad \because \text{Corollary G.8} \\ &= \{\mathbf{p} \in \text{int}(\Delta^k) : \mathbf{z} = \text{link}^{\mathcal{L}}(\mathbf{p})\} \quad \because \text{Definition G.5 and Corollary G.4} \end{aligned}$$

By Proposition G.9, $\text{link}^{\mathcal{L}}$ is an injection. Thus, $|\{\mathbf{p} \in \text{int}(\Delta^k) : \mathbf{z} = \text{link}^{\mathcal{L}}(\mathbf{p})\}| = 1$. ■

Appendix H. Proof of Theorem 4.7

Lemma H.1 Suppose that $k \geq 3$ and $y \in \{2, \dots, k\}$. Let \mathbf{Q} be as in Corollary G.26. Then $\mathbf{Q}\Upsilon_y^{(k)} = \Upsilon_{y-1}^{(k-1)}\mathbf{Q}$.

The proof of Lemma H.1 is similar to that of Lemma B.2 and is thus omitted here. For the proof, see the arXiv version of this work (Wang and Scott, 2023b).

Lemma H.2 Assume $k \geq 3$. Let $\mathcal{L} : \mathbb{R}^k \rightarrow \mathbb{R}^k$ be a regular PERM loss with template $\psi : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$. Let \mathbf{Q} be as in Corollary G.26. Recall from Proposition 4.3 the truncation of ψ denoted by $\text{trunc}[\psi] : \mathbb{R}^{k-2} \rightarrow \mathbb{R}$. Let $\mathbf{z} \in \mathbb{R}^{k-1}$ and $\mathbf{w} \in \mathbb{R}^{k-2}$ be arbitrary and such that $\mathbf{Q}\mathbf{z} = \mathbf{w}$. Then $\lim_{\lambda \rightarrow \infty} \psi(\mathbf{z} + \lambda \mathbf{e}_1^{(k-1)}) = \text{trunc}[\psi](\mathbf{w})$.

Proof Write $\mathbf{z} = [z_1, \dots, z_{k-1}]$ and $\mathbf{w} = [w_1, \dots, w_{k-2}]$. Since $\mathbf{Q}\mathbf{z} = \mathbf{w}$, we have $z_2 = w_1$, $z_3 = w_2$ and so on. Now, define $g : \mathbb{R} \rightarrow \mathbb{R}$ by $g(\lambda) := \psi(\mathbf{z} + \lambda \mathbf{e}_1^{(k-1)}) = \psi(z_1 + \lambda, z_2, \dots, z_{k-1})$ and $h : \mathbb{R} \rightarrow \mathbb{R}$ by

$$h(\lambda) := \psi(\lambda, \mathbf{w}) = \psi(\lambda, w_1, w_2, \dots, w_{k-2}) = \psi(\lambda, z_2, z_3, \dots, z_{k-1}) = g(\lambda - z_1).$$

As argued in the proof of Proposition 4.3, h is decreasing and nonnegative. Thus, g is also decreasing and nonnegative. Moreover, $\lim_{\lambda \rightarrow \infty} g(\lambda) = \lim_{\lambda \rightarrow \infty} g(\lambda - z_1) = \lim_{\lambda \rightarrow \infty} h(\lambda)$. The right hand side is equal to $\text{trunc}[\psi](\mathbf{w})$, as in the proof of Proposition 4.3. \blacksquare

Earlier in Proposition 4.3 and Corollary 4.4, we defined $\text{trunc}[\psi]$, $\psi^{(n)}$ and $\mathcal{L}^{(n)}$. We now define the analogous notation for the reduced form ℓ :

Definition H.3 (Truncation of ℓ) Assume $k \geq 3$. Let $\mathcal{L} : \mathbb{R}^k \rightarrow \mathbb{R}_{\geq 0}^k$ be a regular PERM loss with template $\psi : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$. As in Proposition 4.3, define $\text{trunc}[\ell]$ to be the reduced form of $\text{trunc}[\mathcal{L}]$ (whose template is $\text{trunc}[\psi]$).

Lemma H.4 Assume $k \geq 3$. Let $\mathcal{L} : \mathbb{R}^k \rightarrow \mathbb{R}^k$ be a regular PERM loss with template $\psi : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$. Let \mathbf{Q} be as in Corollary G.26. Recall from Proposition 4.3 the truncation of ψ denoted by $\text{trunc}[\psi] : \mathbb{R}^{k-2} \rightarrow \mathbb{R}$. Let $\text{trunc}[\ell]$ be as in Definition H.3. Let \mathbf{Q} be as in Corollary G.26. Let $\mathbf{z} \in \mathbb{R}^{k-1}$ and $\mathbf{x} \in \mathbb{R}_{\geq 0}^k$ be arbitrary. For brevity, let $\mathbf{e}_1 := \mathbf{e}_1^{(k-1)}$. Then we have

$$\lim_{\lambda \rightarrow +\infty} \mathbf{Q}(\ell(\mathbf{z} + \lambda \mathbf{e}_1) + \mathbf{x}) = \text{trunc}[\ell](\mathbf{Q}\mathbf{z}) + \mathbf{Q}\mathbf{x} \quad (33)$$

and

$$\mathbf{Q}(\ell(\mathbf{z}) + \mathbf{x}) \succ \text{trunc}[\ell](\mathbf{Q}\mathbf{z}). \quad (34)$$

Proof Throughout this proof, let $y \in \{2, \dots, k\}$ be arbitrary. First, we claim that $\Upsilon_y^{(k)} \mathbf{e}_1 = \mathbf{e}_1$ for each $y \in \{2, \dots, k\}$. If $y = k$, then this is clearly true since $\Upsilon_k^{(k)}$ is, by definition, the identity matrix. Now, for $y \in \{2, \dots, k-1\}$, recall that $\Upsilon_y^{(k)}$ is defined as the matrix obtained by replacing the y -th column of the identity matrix with the all -1 's vector. Thus,

since $y > 1$, the first column of $\mathbf{\Upsilon}_y^{(k)}$ is equal to that of the identity matrix. Therefore, $\mathbf{\Upsilon}_y^{(k)}\mathbf{e}_1 = \mathbf{e}_1$ as well. Next, still assuming $y \in \{2, \dots, k-1\}$, we have

$$\ell_y(\mathbf{z} + \lambda\mathbf{e}_1) = \psi(\mathbf{\Upsilon}_y^{(k)}(\mathbf{z} + \lambda\mathbf{e}_1)) = \psi(\mathbf{\Upsilon}_y^{(k)}\mathbf{z} + \lambda\mathbf{e}_1).$$

Hence, we have

$$\begin{aligned} \lim_{\lambda \rightarrow +\infty} \ell_y(\mathbf{z} + \lambda\mathbf{e}_1) &= \lim_{\lambda \rightarrow +\infty} \psi(\mathbf{\Upsilon}_y^{(k)}\mathbf{z} + \lambda\mathbf{e}_1) \quad \because \text{Theorem C.2.} \\ &= \text{trunc}[\psi](\mathbf{Q}\mathbf{\Upsilon}_y^{(k)}\mathbf{z}) \quad \because \text{Proposition 4.3 and Lemma H.2} \\ &= \text{trunc}[\psi]\left(\mathbf{\Upsilon}_{y-1}^{(k-1)}\mathbf{Q}\mathbf{z}\right) \quad \because \text{Lemma H.1} \\ &= [\text{trunc}[\ell]]_{y-1}(\mathbf{Q}\mathbf{z}) \quad \because \text{Theorem C.2 and explanation below.} \end{aligned}$$

Application of Theorem C.2 to the last equality requires a bit more explanation. For said equality, we used the fact that $\text{trunc}[\ell]$ is the reduced form of $\mathcal{L}^{(k-1)}$ (Definition H.3), whose template is $\text{trunc}[\psi]$. Thus, applying Equation (12) (which is a corollary of Equation (11) from Theorem C.2) to $\text{trunc}[\mathcal{L}]$, we have $[\text{trunc}[\ell]]_{y-1}(\mathbf{z}) = \text{trunc}[\psi](\mathbf{\Upsilon}_{y-1}^{(k-1)}\mathbf{z})$. Thus,

$$\lim_{\lambda \rightarrow +\infty} \mathbf{Q}(\ell(\mathbf{z} + \lambda\mathbf{e}_1) + \mathbf{x}) = \text{trunc}[\ell](\mathbf{Q}\mathbf{z}) + \mathbf{Q}\mathbf{x}.$$

Next, for every $y \in \{2, \dots, k\}$, we note that the function

$$g_y(\lambda) := \ell_y(\mathbf{z} + \lambda\mathbf{e}_1) = \psi(\mathbf{\Upsilon}_y\mathbf{z} + \lambda\mathbf{e}_1)$$

is strictly decreasing. To see this, by the chain rule for curves²⁵, we have

$$g'_y(\lambda) = \nabla_\psi(\mathbf{\Upsilon}_y\mathbf{z} + \lambda\mathbf{e}_1)^\top \mathbf{e}_1 < 0.$$

Thus, $\ell_y(\mathbf{z}) = g_y(0) > \lim_{\lambda \rightarrow +\infty} g_y(\lambda) = \text{trunc}[\ell]_{y-1}(\mathbf{Q}\mathbf{z})$, which proves that

$$\mathbf{Q}(\ell(\mathbf{z}) + \mathbf{x}) \succ \text{trunc}[\ell](\mathbf{Q}\mathbf{z}) + \mathbf{Q}\mathbf{x} \succeq \text{trunc}[\ell](\mathbf{Q}\mathbf{z})$$

as desired. ■

Lemma H.5 *Assume $k \geq 3$. Let $\mathcal{L} : \mathbb{R}^k \rightarrow \mathbb{R}^k$ be a regular PERM loss with template $\psi : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$. Let \mathbf{Q} be as in Corollary G.26. Recall from Proposition 4.3 the truncation of ψ denoted by $\text{trunc}[\psi] : \mathbb{R}^{k-2} \rightarrow \mathbb{R}$. Let $\text{trunc}[\ell]$ be as in Definition H.3. Let \mathbf{Q} be as in Corollary G.26. Then $\mathbf{Q}(\text{ran}_c(\ell)) \subseteq \text{ran}_c^\circ(\text{trunc}[\ell])$ and $\text{clos}[\mathbf{Q}(\text{ran}_c(\ell))] = \text{ran}_c(\text{trunc}[\ell])$.*

Proof Let $C := \mathbf{Q}(\text{ran}_c(\ell))$ and take $\zeta \in C$. We first prove $C \subseteq \text{ran}_c^\circ(\text{trunc}[\ell])$. By the characterization of $\text{ran}_c(\ell)$ from Corollary G.32 item 1, there exists $\mathbf{z} \in \mathbb{R}^{k-1}$ and $\mathbf{x} \in [0, \infty)^k$ such that $\zeta = \mathbf{Q}(\ell(\mathbf{z}) + \mathbf{x})$. Now let $\tilde{\mathbf{z}} := \mathbf{Q}(\mathbf{z})$. Applying Eqn. (34) from Lemma H.4, we get

$$\zeta = \mathbf{Q}(\ell(\mathbf{z}) + \mathbf{x}) \succ \text{trunc}[\ell](\tilde{\mathbf{z}}).$$

25. See the “Vector Calculus” section in the appendix of the arXiv version of this manuscript (Wang and Scott, 2023b).

In particular, by the characterization of $\text{ran}_c^\circ(\text{trunc}[\ell])$ from Lemma G.33, we have that $\zeta \in \text{ran}_c^\circ(\text{trunc}[\ell])$. This proves that $C \subseteq \text{ran}_c^\circ(\text{trunc}[\ell])$.

Next, we prove $\text{clos}[C] = \text{ran}_c(\text{trunc}[\ell])$. We first show that $\text{clos}[C] \supseteq \text{ran}_c(\text{trunc}[\ell])$ by proving that every point $\text{ran}_c(\text{trunc}[\ell])$ is a limit point of C .

Let $\zeta \in \text{ran}_c(\text{trunc}[\ell])$. By the characterization of $\text{ran}_c(\text{trunc}[\ell])$ as in Corollary G.32, there exists $\bar{\mathbf{z}} \in \mathbb{R}^{k-2}$ and $\bar{\mathbf{x}} \in [0, \infty)^{k-1}$ such that $\zeta = \text{trunc}[\ell](\bar{\mathbf{z}}) + \bar{\mathbf{x}}$. Now, pick $\mathbf{z} \in \mathbb{R}^{k-1}$ and $\mathbf{x} \in [0, \infty)^k$ such that $\bar{\mathbf{z}} = \mathbf{Q}\mathbf{z}$ and $\bar{\mathbf{x}} = \mathbf{Q}\mathbf{x}$. Applying Lemma H.4 Eqn. (33), we get that ζ is a limit point of S , which proves the desired claim. This proves that $\text{clos}(C) \supseteq \text{ran}_c(\text{trunc}[\ell])$. By the first part, we know that $C \subseteq \text{ran}_c^\circ(\text{trunc}[\ell])$. By Corollary G.32, $\text{ran}_c^\circ(\text{trunc}[\ell]) = \text{int}(\text{ran}_c(\text{trunc}[\ell])) \subseteq \text{ran}_c(\text{trunc}[\ell])$. Putting it all together, we have

$$C \subseteq \text{ran}_c^\circ(\text{trunc}[\ell]) \subseteq \text{ran}_c(\text{trunc}[\ell]) \subseteq \text{clos}(C).$$

From Corollary G.32, we have that $\text{ran}_c(\text{trunc}[\ell])$ is closed. Since by definition $\text{clos}(C)$ is the smallest closed set containing C , we get that $\text{ran}_c(\text{trunc}[\ell]) = \text{clos}(C)$, as desired. \blacksquare

Theorem H.6 (Blackwell and Girshick (1979)) *Let $C \subseteq \mathbb{R}^n$ be a convex set. Then $\text{int}(C) = \text{int}(\text{clos}(C))$.*

Proposition H.7 *Assume $k \geq 3$. Let $\mathcal{L} : \mathbb{R}^k \rightarrow \mathbb{R}^k$ be a regular PERM loss with template $\psi : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$. Let \mathbf{Q} be as in Corollary G.26. Recall from Proposition 4.3 the truncation of ψ denoted by $\text{trunc}[\psi] : \mathbb{R}^{k-2} \rightarrow \mathbb{R}$. Let $\text{trunc}[\ell]$ be as in Definition H.3. Let \mathbf{Q} be as in Corollary G.26. Then we have $\mathbf{Q}(\text{ran}_c(\ell)) = \text{ran}_c^\circ(\text{trunc}[\ell])$*

Proof For brevity, let $C := \mathbf{Q}(\text{ran}_c(\ell))$. By Corollary G.32, $\text{ran}_c(\text{trunc}[\ell])$ is convex. Since convexity is preserved under projection, we have that C is convex as well. Now,

$$\text{int}(C) = \text{int}(\text{clos}(C)) \quad \because \text{Theorem H.6} \quad (35)$$

$$= \text{int}(\text{ran}_c(\text{trunc}[\ell])) \quad \because \text{Lemma H.5} \quad (36)$$

$$= \text{ran}_c^\circ(\text{trunc}[\ell]) \quad \because \text{Lemma G.28} \quad (37)$$

$$\supseteq C \quad \because \text{Lemma H.5} \quad (38)$$

Since $C \supseteq \text{int}(C)$ by definition, we conclude that $C = \text{ran}_c^\circ(\text{trunc}[\ell])$. \blacksquare

Proposition H.8 *Assume $k \geq 3$. Let $\mathcal{L} : \mathbb{R}^k \rightarrow \mathbb{R}^k$ be a regular PERM loss with template $\psi : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$. Let \mathbf{Q} be as in Corollary G.26. Recall from Proposition 4.3 the truncation of ψ denoted by $\text{trunc}[\psi] : \mathbb{R}^{k-2} \rightarrow \mathbb{R}$. Let $\text{trunc}[\ell]$ be as in Definition H.3. Let \mathbf{Q} be as in Corollary G.26. Then we have $\mathbf{Q}(\text{ran}_c^\circ(\ell)) = \text{ran}_c^\circ(\text{trunc}[\ell])$.*

Proof By the preceding Proposition H.7, we have $\mathbf{Q}(\text{ran}_c(\ell)) = \text{ran}_c^\circ(\text{trunc}[\ell])$. Since $\text{ran}_c^\circ(\ell) \subseteq \text{ran}_c(\ell)$ we have $\mathbf{Q}(\text{ran}_c^\circ(\ell)) \subseteq \mathbf{Q}(\text{ran}_c(\ell))$. Thus, to prove the result we only have to show $\mathbf{Q}(\text{ran}_c^\circ(\ell)) \supseteq \text{ran}_c^\circ(\text{trunc}[\ell])$.

To this end, let $\text{trunc}[\ell](\mathbf{w}) \in \text{ran}_c^\circ(\text{trunc}[\ell])$ and $\mathbf{z} \in \text{ran}_c(\ell)$ be such that $\mathbf{Q}\ell(\mathbf{z}) = \text{trunc}[\ell](\mathbf{w})$. By Corollary G.26, there exist $\mathbf{z}^* \in \mathbb{R}^{k-1}$ and $t^* \in \mathbb{R}$ such that $t^* > 0$

and $\mathbf{Q}(\ell(\mathbf{z}^*) + t^* \mathbf{1}) = \mathbf{Q}\ell(\mathbf{z}) = \text{trunc}[\ell](\mathbf{w})$. Since $\ell(\mathbf{z}^*) + t^* \mathbf{1} \in \text{ran}_c^\circ(\ell)$, we get that $\text{trunc}[\ell](\mathbf{w}) \in \mathbf{Q}(\text{ran}_c^\circ(\ell))$ as desired. \blacksquare

Definition H.9 For each $m \in \{0, 1, \dots, k-2\}$, define $\text{trunc}^{\times m}[\ell]$ to be m -fold repeated applications of trunc to ℓ , i.e., $\text{trunc}^{\times m}[\ell] := \text{trunc}[\dots \text{trunc}[\text{trunc}[\ell]] \dots]$ where trunc appears m -times. By convention, let $\text{trunc}^{\times 0}[\ell] = \ell$. Moreover, for each $n \in \{2, \dots, k\}$, define the notational shorthand $\ell^{(n)} := \text{trunc}^{\times(k-n)}[\ell]$.

Remark H.10 It follows tautologically from the definition of $\psi^{(n)}$ in Corollary 4.4, that define $\ell^{(n)}$ is the reduced form of $\mathcal{L}^{(n)}$ (whose template is $\psi^{(n)}$).

Next, let $m \geq 1$ be an integer. Below, let $\mathbf{Q}^{\times m}$ denote the m -fold iterated composition of \mathbf{Q} . In other words, $\mathbf{Q}^{\times m} := \mathbf{Q} \circ \dots \circ \mathbf{Q}$ repeated m times.

Proposition H.11 Suppose that \mathcal{L} is totally regular. Then Let $m \in \{1, \dots, k-2\}$. Then $\mathbf{Q}^{\times m}(\text{ran}_c(\ell)) = \text{ran}_c^\circ(\text{trunc}^{\times m}[\ell])$.

Proof We prove by induction. The case when $m = 1$ is simply Proposition H.8. Now, suppose that the result holds for m where $1 < m < k-2$. Then

$$\begin{aligned} \mathbf{Q}^{\times m+1}(\text{ran}_c(\ell)) &= \mathbf{Q}(\mathbf{Q}^{\times m}(\text{ran}_c(\ell))) \\ &= \mathbf{Q}(\text{ran}_c^\circ(\text{trunc}^{\times m}[\ell])) \quad \because \text{Induction hypothesis} \\ &= \text{ran}_c^\circ(\text{trunc}[\text{trunc}^{\times m}[\ell]]) \quad \because \text{Proposition H.8} \\ &= \text{ran}_c^\circ(\text{trunc}^{\times m+1}[\ell]). \end{aligned}$$

This completes the induction step and the desired result follows. \blacksquare

Before presenting the proof of Theorem 4.7, we recall (Tewari and Bartlett, 2007, Theorem 7):

Theorem H.12 (Tewari and Bartlett (2007)) Let $S \subseteq \mathbb{R}_+^k$ be a symmetric convex set. Then S is classification calibrated if and only if 1. S is admissible and 2. $\mathbf{Q}^{\times m}(S)$ is admissible for all $m \in \{1, \dots, k-2\}$.

Finally, we conclude with the

Proof of Theorem 4.7 Assume that we are in the situation stated at the beginning of Section H. Let $S = \text{ran}_c(\ell^{(k)})$. By Theorem H.12, it suffices to prove that 1. S is admissible and 2. $\mathbf{Q}^{\times m}(S)$ is admissible for all $m \in \{1, \dots, k-2\}$. First, from Proposition G.38, $S = \text{ran}_c(\ell^{(k)})$ is admissible. Next, for each $m \in \{1, \dots, k-2\}$, we have by Proposition H.11 that $\mathbf{Q}^{\times m}(\text{ran}_c(\ell^{(k)})) = \text{ran}_c^\circ(\ell^{(k-m)})$. Since $\mathcal{L}^{(k-m)}$ is a regular PERM loss with reduced form $\ell^{(k-m)}$, we have again by Proposition G.38 that $\text{ran}_c^\circ(\ell^{(k-m)})$ is admissible. \blacksquare

Appendix I. Classification-Calibration of Fenchel-Young losses

The goal of this section is two fold. The first subsection (Section I.1) presents the proof of Theorem 5.10 on a sufficient condition for the classification-calibration of Fenchel-Young losses. The second subsection (Section I.2) shows the existence of a totally regular negentropy that is strictly convex, but not strongly convex.

To prepare for the later proofs, we introduce a more convenient parametrization of the negentropy and its Fenchel-Young loss. Define the *reduced k -probability simplex* as

$$\tilde{\Delta}^k := \{\tilde{\mathbf{p}} := (p_1, \dots, p_{k-1}) \in [0, 1]^{k-1} : \sum_{i=1}^{k-1} p_i \leq 1\}. \quad (39)$$

In other words, $\tilde{\Delta}^k$ is simply Δ^k without the first coordinate. To every function $\Omega : \Delta^k \rightarrow \mathbb{R}$ with domain on the k -simplex, we define a corresponding function $\tilde{\Omega} : \tilde{\Delta}^k \rightarrow \mathbb{R}$ called the *reduced form* of Ω , defined by

$$\tilde{\Omega}(\tilde{\mathbf{p}}) := \Omega\left(p_1, \dots, p_{k-1}, 1 - \sum_{i=1}^{k-1} p_i\right), \quad \text{for all } \tilde{\mathbf{p}} = (p_1, \dots, p_{k-1})^\top \in \tilde{\Delta}^k. \quad (40)$$

Equation (40) induces a one-to-one correspondence between functions $\Omega : \Delta^k \rightarrow \mathbb{R}$ on the simplex Δ^k and functions $\tilde{\Omega} : \tilde{\Delta}^k \rightarrow \mathbb{R}$ on the reduced simplex $\tilde{\Delta}^k$.

I.1 Proof of Theorem 5.10

Before proceeding with the proof, we establish two key results. Recall that $\tilde{\Delta}^k$ is defined in Eqn. (39) and $\tilde{\Omega}$ in Eqn. 40.

Proposition I.1 *Let Ω be a negentropy (Definition 5.2) and $\mu \in \mathbb{R}$. Then the Fenchel-Young loss $\mathcal{L}^{\Omega, \mu}$ associated to Ω and μ (Definition 5.1) is a PERM loss that is closed, convex, and non-negative. The template $\psi^{\Omega, \mu}$ of $\mathcal{L}^{\Omega, \mu}$ is semi-coercive and is given by*

$$\psi^{\Omega, \mu}(\mathbf{z}) = \max_{\tilde{\mathbf{p}} = [p_1, \dots, p_{k-1}] \in \tilde{\Delta}^k} -\tilde{\Omega}(\tilde{\mathbf{p}}) + \mu \mathbf{1}^\top \tilde{\mathbf{p}} - \langle \tilde{\mathbf{p}}, \mathbf{z} \rangle. \quad (41)$$

Furthermore, if Ω is a regular negentropy, then \mathcal{L} is a regular PERM loss.

Remark I.2 We note that results from Blondel et al. (2020) may be used to prove portions of Proposition I.1. Combining their Proposition 1-“Order preservation” part with the expression immediately following their Definition 2 can be used to prove that $\mathcal{L}^{\Omega, \mu}$ is permutation equivariant. Moreover, the expression in our Eqn. (41) is related to their Eqn. (15). However, there are key technical differences since we work with the ($k - 1$ -dimensional) relative margin, while Blondel et al. (2020) use the class-score formulation. For this reason and for the reader’s convenience, we prove Proposition I.1 without leveraging their results.

Proof In this proof, all elementary basis vectors are implicitly assumed to be k -dimensional, i.e., we write \mathbf{e}_y instead of $\mathbf{e}_y^{(k)}$. First, recall that the Fenchel conjugate of a closed convex function is again closed convex (Rockafellar, 1970). Next, we show that \mathcal{L} is permutation equivariant. First, we show that \mathcal{L} is symmetric. By Lemma C.3, it suffices to prove the claim that $\mathcal{L}(\mathbf{T}_j \mathbf{v}) = \mathbf{T}_j \mathcal{L}(\mathbf{v})$ for all $j \in [k]$ and $\mathbf{v} \in \mathbb{R}^k$. To this end, let $y \in [k]$ be arbitrary. Recall that the Fenchel-Young loss is defined as

$$[\mathcal{L}(\mathbf{v})]_y = \max_{\mathbf{p} \in \Delta^k} -\Omega(\mathbf{p}) + \langle \mathbf{c}_y, \mathbf{p} - \mathbf{e}_y \rangle + \langle \mathbf{v}, \mathbf{p} - \mathbf{e}_y \rangle \quad (42)$$

Thus,

$$\begin{aligned}
[\mathbf{T}_j \mathcal{L}(\mathbf{v})]_y &= [\mathcal{L}(\mathbf{v})]_{\tau_j(y)} = \max_{\mathbf{p} \in \Delta^k} -\Omega(\mathbf{p}) + \langle \mathbf{v} + \mathbf{c}_{\tau_j(y)}, \mathbf{p} - \mathbf{e}_{\tau_j(y)} \rangle \quad \because \text{definition of } \mathcal{L} \\
&= \max_{\mathbf{p} \in \Delta^k} -\Omega(\mathbf{p}) + \langle \mathbf{T}_j(\mathbf{v} + \mathbf{c}_{\tau_j(y)}), \mathbf{T}_j(\mathbf{p} - \mathbf{e}_{\tau_j(y)}) \rangle \quad \because \mathbf{T}_j \text{ is an isometry} \\
&= \max_{\mathbf{p} \in \Delta^k} -\Omega(\mathbf{p}) + \langle \mathbf{T}_j \mathbf{v} + \mathbf{c}_{\tau_j(\tau_j(y))}, \mathbf{T}_j \mathbf{p} - \mathbf{e}_{\tau_j(\tau_j(y))} \rangle \\
&= \max_{\mathbf{p} \in \Delta^k} -\Omega(\mathbf{p}) + \langle \mathbf{T}_j \mathbf{v} + \mathbf{c}_y, \mathbf{T}_j \mathbf{p} - \mathbf{e}_y \rangle \quad \because \tau_j \circ \tau_j \text{ is the identity} \\
&= \max_{\mathbf{p} \in \Delta^k} -\Omega(\mathbf{T}_j \mathbf{p}) + \langle \mathbf{T}_j \mathbf{v} + \mathbf{c}_y, \mathbf{T}_j \mathbf{p} - \mathbf{e}_y \rangle \quad \because \Omega \text{ is symmetric} \\
&= \max_{\mathbf{p} \in \Delta^k} -\Omega(\mathbf{p}) + \langle \mathbf{T}_j \mathbf{v} + \mathbf{c}_y, \mathbf{p} - \mathbf{e}_y \rangle \quad \because \mathbf{T}_j|_{\Delta^k} \text{ is a bijection} \\
&= [\mathcal{L}(\mathbf{T}_j \mathbf{v})]_y \quad \because \text{definition of } \mathcal{L}.
\end{aligned}$$

Since y is arbitrary, we have proven the claim.

Next, we show that \mathcal{L} is relative margin-based, i.e., eq. (42) depends on \mathbf{v} only through $\mathbf{D}\mathbf{v}$. Recall that $\mathbf{D}\mathbf{v} = [v_k - v_1 \ v_k - v_2 \ \cdots \ v_k - v_{k-1}]^\top$. See Definition 2.1. Now, only the “ $\langle \mathbf{v}, \mathbf{p} - \mathbf{e}_y \rangle$ ” term of eq. (42) depends on \mathbf{v} . Thus, it suffices to prove that this term only depends on $\mathbf{D}\mathbf{v}$. To this end, observe that

$$\begin{aligned}
\langle \mathbf{v}, \mathbf{p} \rangle &= p_1 v_1 + \cdots + p_k v_k = p_1 v_1 + \cdots + p_{k-1} v_{k-1} + (1 - (p_1 + \cdots + p_{k-1})) v_k \\
&= v_k - (p_1(v_k - v_1) + \cdots + p_{k-1}(v_k - v_{k-1})) = v_k - \langle \tilde{\mathbf{p}}, \mathbf{D}\mathbf{v} \rangle
\end{aligned}$$

where we write $\tilde{\mathbf{p}}$ to denote the vector $[p_1, \dots, p_{k-1}]^\top$. Thus, for each $y \in [k-1]$, we have

$$\langle \mathbf{v}, \mathbf{p} - \mathbf{e}_y \rangle = \langle \mathbf{v}, \mathbf{p} \rangle - v_y = v_k - v_y - \langle \tilde{\mathbf{p}}, \mathbf{D}\mathbf{v} \rangle = [\mathbf{D}\mathbf{v}]_y - \langle \tilde{\mathbf{p}}, \mathbf{D}\mathbf{v} \rangle$$

This shows that \mathcal{L} is relative margin-based. Now the template of \mathcal{L} is

$$\psi(\mathbf{z}) = \ell_k(\mathbf{z}) = \max_{\mathbf{p} \in \Delta^k} -\Omega(\mathbf{p}) + \langle \mathbf{c}_k, \mathbf{p} - \mathbf{e}_k \rangle - \langle \tilde{\mathbf{p}}, \mathbf{z} \rangle.$$

Since $\mathbf{e}_k \in \Delta^k$ and $[\mathbf{e}_k]_y = 0$ for $y \in [k-1]$, we have by construction that

$$\psi(\mathbf{z}) \geq -\Omega(\mathbf{e}_1) + \langle \mathbf{c}_1, \mathbf{e}_1 - \mathbf{e}_1 \rangle - \langle \mathbf{0}, \mathbf{z} \rangle = 0.$$

When $\mathbf{c}_k = \mu(\mathbf{1} - \mathbf{e}_k)$, we have

$$\psi(\mathbf{z}) = \max_{\tilde{\mathbf{p}} \in \tilde{\Delta}^k} -\tilde{\Omega}(\tilde{\mathbf{p}}) + \mu \mathbf{1}^\top \tilde{\mathbf{p}} - \langle \tilde{\mathbf{p}}, \mathbf{z} \rangle, \quad \text{for all } \mathbf{z} \in \mathbb{R}^{k-1}. \quad (43)$$

Finally, we prove that ψ is semi-coercive (Definition 4.1), i.e., for all $c \in \mathbb{R}$ there exists $b \in \mathbb{R}$ such that $\{\mathbf{z} \in \mathbb{R}^{k-1} : \psi(\mathbf{z}) \leq c\} \subseteq \{\mathbf{z} \in \mathbb{R}^{k-1} : b \leq \min \mathbf{z}\}$. To this end, let $c \in \mathbb{R}$ and $\mathbf{z} \in \mathbb{R}^{k-1}$ be such that $c \geq \psi(\mathbf{z})$. Let $j \in \arg \min \mathbf{z}$. Then since $\mathbf{e}_j = \mathbf{e}_j^{k-1} \in \tilde{\Delta}^k$, we have

$$c \geq \psi(\mathbf{z}) = \sup_{\mathbf{p} \in \tilde{\Delta}^k} -\tilde{\Omega}(\mathbf{p}) - \langle \mathbf{p}, \mathbf{z} \rangle \geq -\tilde{\Omega}(\mathbf{e}_j) + \mu - \langle \mathbf{e}_j, \mathbf{z} \rangle \geq -z_j = -\min \mathbf{z}.$$

Thus, we have shown that $\{\mathbf{z} \in \mathbb{R}^{k-1} : \psi(\mathbf{z}) \leq c\} \subseteq \{\mathbf{z} \in \mathbb{R}^{k-1} : \min \mathbf{z} \geq -c\}$.

Next, we prove the ‘‘Furthermore’’ part. By the first part, it remains to show that ψ is strictly convex, twice differentiable and $\nabla\psi(\mathbf{z}) \prec \mathbf{0}$ for all $\mathbf{z} \in \mathbb{R}^{k-1}$. Define $\varphi(\tilde{\mathbf{p}}) := \tilde{\Omega}(\tilde{\mathbf{p}}) - \mu\mathbf{1}^\top\tilde{\mathbf{p}}$. Then $\varphi : \tilde{\Delta}^k \rightarrow \mathbb{R}$ is also of Legendre type. Moreover, note that

$$\begin{aligned}\psi(\mathbf{z}) &= \max_{\tilde{\mathbf{p}} \in \tilde{\Delta}^k} \langle \tilde{\mathbf{p}}, -\mathbf{z} \rangle - \tilde{\Omega}(\tilde{\mathbf{p}}) + \mu\mathbf{1}^\top\tilde{\mathbf{p}} \quad \because \text{Eqn. (43)} \\ &= \max_{\tilde{\mathbf{p}} \in \tilde{\Delta}^k} \langle \tilde{\mathbf{p}}, -\mathbf{z} \rangle - \varphi(\tilde{\mathbf{p}}) = \varphi^*(-\mathbf{z}) \quad \because \text{definition of Fenchel conjugate}\end{aligned}$$

Recall the following fundamental theorem regarding convex conjugates (Rockafellar, 1970).

Theorem I.3 (Rockafellar (1970)) *If (C, f) is a convex function of Legendre type, then (C^*, f^*) is a convex function of Legendre type. The map $\nabla_f : C \rightarrow C^*$ is a homeomorphism and $\nabla_{f^*} = (\nabla_f)^{-1}$.*

By Theorem I.3, we have

1. The function φ^* , and hence ψ , is of Legendre type. In particular, ψ is strictly convex.
2. The derivative $\nabla\varphi : \text{int}(\tilde{\Delta}^k) \rightarrow \mathbb{R}^{k-1}$ is a bijection and the derivative of φ^* satisfies $\nabla_{\varphi^*} = (\nabla\varphi)^{-1} : \mathbb{R}^{k-1} \rightarrow \text{int}(\tilde{\Delta}^k)$.

It follows that if φ is twice differentiable, then so is φ^* . Finally, by the chain rule, we have $\nabla\psi(\mathbf{z}) = -\nabla\varphi^*(\mathbf{z})$. Since $\nabla\varphi^*(\mathbf{z}) \in \text{int}(\tilde{\Delta}^k)$ for all \mathbf{z} , we have in particular that $\nabla\varphi^*(\mathbf{z}) \succ \mathbf{0}$. Thus, $\nabla\psi(\mathbf{z}) \prec \mathbf{0}$ for all \mathbf{z} . \blacksquare

Proposition I.4 *Let $k \geq 3$ be an integer, $\Omega : \Delta^k \rightarrow \mathbb{R}$ a regular negentropy and $\mu \geq 0$. Let $\mathcal{L}^{\Omega, \mu}$ be the Fenchel-Young loss corresponding to Ω and μ . Then $\text{trunc}[\mathcal{L}^{\Omega, \mu}] = \mathcal{L}^{\text{trunc}[\Omega], \mu}$.*

Proof Throughout this proof, let $\Theta := \text{trunc}[\Omega]$. Let $\tilde{\Omega}$ and $\tilde{\Theta}$ be the reduced versions (see eq. (40)) of Ω and Θ , respectively. We remark that the notation used in this proof slightly departs from that of eq. (40) in the following way: Here, elements of $\tilde{\Delta}^k$ will be denoted as \mathbf{p} instead of $\tilde{\mathbf{p}}$. Likewise, elements of $\tilde{\Delta}^{k-1}$ will be denoted as \mathbf{q} rather than $\tilde{\mathbf{q}}$. Moreover, throughout this proof, we abuse notation and denote by $\mathbf{1}$ either the $(k-1)$ or the $(k-2)$ -dimensional vector. The dimension of $\mathbf{1}$ will be clear from the context.

Now, from eq. (41) in Proposition I.1, recall that the template of $\mathcal{L}^{\Omega, \mu}$ is

$$\psi^{\Omega, \mu} : \mathbb{R}^{k-1} \rightarrow \mathbb{R}, \text{ where } \psi^{\Omega, \mu}(\mathbf{z}) = \max_{\mathbf{p} = [p_1, \dots, p_{k-1}] \in \tilde{\Delta}^k} -\tilde{\Omega}(\mathbf{p}) + \mu\mathbf{1}^\top\mathbf{p} - \langle \mathbf{z}, \mathbf{p} \rangle. \quad (44)$$

Similarly, the template of $\mathcal{L}^{\Theta, \mu}$ is

$$\psi^{\Theta, \mu} : \mathbb{R}^{k-2} \rightarrow \mathbb{R}, \text{ where } \psi^{\Theta, \mu}(\mathbf{w}) = \max_{\mathbf{q} = [q_1, \dots, q_{k-2}] \in \tilde{\Delta}^{k-1}} -\tilde{\Theta}(\mathbf{q}) + \mu\mathbf{1}^\top\mathbf{q} - \langle \mathbf{w}, \mathbf{q} \rangle. \quad (45)$$

By the definition of the truncation of a PERM loss (given in Proposition 4.3), the template of $\text{trunc}[\mathcal{L}^{\Omega, \mu}]$ is $\text{trunc}[\psi^{\Omega, \mu}]$. Since PERM losses are uniquely defined by their templates (Theorem 2.5), the result of Proposition I.4 can be equivalently stated as $\text{trunc}[\psi^{\Omega, \mu}] = \psi^{\Theta, \mu}$. Below, our proof will focus on proving this identity.

Both $\text{trunc}[\psi^{\Omega,\mu}]$ and $\psi^{\Theta,\mu}$ have \mathbb{R}^{k-2} as domain. Fix $\mathbf{w} \in \mathbb{R}^{k-2}$ arbitrarily. Our goal is to show that $\text{trunc}[\psi^{\Omega,\mu}](\mathbf{w}) = \psi^{\Theta,\mu}(\mathbf{w})$. Pick $\mathbf{z} \in \mathbb{R}^{k-1}$ such that $\mathbf{Qz} = \mathbf{w}$. For instance, we can pad \mathbf{w} with a zero, i.e., take $\mathbf{z} := [0, \mathbf{w}] = \text{pad}(\mathbf{w})$. By Lemma H.2, we have

$$\text{trunc}[\psi^{\Omega,\mu}](\mathbf{w}) = \lim_{\lambda \rightarrow +\infty} \psi^{\Omega,\mu}(\mathbf{z} + \lambda \mathbf{e}_1^{(k-1)})$$

To simplify notations, let $\mathbf{e}_1 := \mathbf{e}_1^{(k-1)}$. In view of eq. (44) and eq. (45), establishing our goal, i.e., the identity $\text{trunc}[\psi^{\Omega,\mu}](\mathbf{w}) = \psi^{\Theta,\mu}(\mathbf{w})$, is equivalent to proving

$$\lim_{\lambda \rightarrow +\infty} \max_{\mathbf{p} \in \tilde{\Delta}^k} -\tilde{\Omega}(\mathbf{p}) + \mu \mathbb{1}^\top \mathbf{p} - \langle \mathbf{z} + \lambda \mathbf{e}_1, \mathbf{p} \rangle = \max_{\mathbf{q} \in \tilde{\Delta}^{k-1}} -\tilde{\Theta}(\mathbf{q}) + \mu \mathbb{1}^\top \mathbf{q} - \langle \mathbf{Qz}, \mathbf{q} \rangle. \quad (46)$$

For brevity, we define

$$g(\lambda) := \max_{\mathbf{p} \in \tilde{\Delta}^k} -\tilde{\Omega}(\mathbf{p}) + \mu \mathbb{1}^\top \mathbf{p} - \langle \mathbf{z} + \lambda \mathbf{e}_1, \mathbf{p} \rangle. \quad (47)$$

For each $\lambda \in \mathbb{R}$, fix arbitrarily an element $\mathbf{p}^\lambda \in \arg \max_{\mathbf{p} \in \tilde{\Delta}^k} -\tilde{\Omega}(\mathbf{p}) + \mu \mathbb{1}^\top \mathbf{p} - \langle \mathbf{z} + \lambda \mathbf{e}_1, \mathbf{p} \rangle$. We note that the maximization is over a compact domain with a continuous objective. We will prove the result via a series of claims.

Claim 1: $g : \mathbb{R} \rightarrow \mathbb{R}_+$, defined at eq. (47), is monotone non-increasing.

Claim 2: For all $\lambda \in \mathbb{R}$, $g(\lambda) \geq -\tilde{\Theta}(\hat{\mathbf{q}}) + \mu \mathbb{1}^\top \hat{\mathbf{q}} - \langle \mathbf{w}, \hat{\mathbf{q}} \rangle$.

Claim 3: $\bar{p}_1 = 0$.

Claim 4: $\lim_{t \rightarrow \infty} g(\lambda_t) = -\tilde{\Omega}(\hat{\mathbf{q}}) + \mu \mathbb{1}^\top \hat{\mathbf{q}} - \langle \mathbf{w}, \hat{\mathbf{q}} \rangle$

The proofs of these claims proceed via straightforward but tedious computations and applications of inequalities. They are omitted here but appear in the arXiv version of this work (Wang and Scott, 2023b). Now we finish the proof of eq. (46), which we argued earlier suffices to establish Proposition I.4. By Claims 1 and 4, we have that $\lim_{\lambda \rightarrow +\infty} g(\lambda) = -\tilde{\Theta}(\hat{\mathbf{q}}) + \mu \mathbb{1}^\top \hat{\mathbf{q}} - \langle \mathbf{w}, \hat{\mathbf{q}} \rangle$. But $\lim_{\lambda \rightarrow +\infty} g(\lambda)$ is exactly the left hand side of eq. (46), by the construction of g in eq. (47). By the definition of $\hat{\mathbf{q}}$, moreover we have that $-\tilde{\Theta}(\hat{\mathbf{q}}) + \mu \mathbb{1}^\top \hat{\mathbf{q}} - \langle \mathbf{w}, \hat{\mathbf{q}} \rangle$ is precisely the right hand side of eq. (46). ■

Theorem I.5 *Let $\Omega : \Delta^k \rightarrow \mathbb{R}$ be a negentropy (Definition 5.2). For each $n \in \{2, \dots, k\}$, let $\Omega^{(n)}$ be the n -ary truncated negentropy of Ω (Definition 5.9). Let $\mathcal{L}^{\Omega,\mu}$ be the Fenchel-Young loss corresponding to Ω and μ (Definition 5.1). Likewise, for each $n \in \{2, \dots, k\}$, let $\mathcal{L}^{(n)}$ be the n -ary truncation of $\mathcal{L}^{\Omega,\mu}$ (Definition 4.5). Then $\mathcal{L}^{(n)} = \mathcal{L}^{\Omega^{(n)},\mu}$. In other words, the n -ary truncated loss $\mathcal{L}^{(n)}$ of $\mathcal{L}^{\Omega,\mu}$ is equal to the Fenchel-Young loss of the n -ary truncated negentropy $\Omega^{(n)}$ and μ .*

The proof of Theorem I.5 is omitted here since it largely mirrors that of Proposition H.11. See the arXiv version of manuscript (Wang and Scott, 2023b) for the proof.

Proof [Proof of Theorem 5.10] By Theorem 4.7, it suffices to show that $\mathcal{L}^{\Omega,\mu}$ is totally regular (Definition 4.5). Unwinding the definition, our goal is to show for each $n \in \{2, \dots, k\}$ that the n -ary truncated loss of $\mathcal{L}^{\Omega,\mu}$, denoted $\mathcal{L}^{(n)}$, is regular.

Next, let $\Omega^{(n)}$ be the n -ary truncated negentropy of Ω (Definition 5.9) and let $\mathcal{L}^{\Omega^{(n)}, \mu}$ be the Fenchel-Young loss corresponding to $\Omega^{(n)}$ and μ . By Theorem I.5, we have $\mathcal{L}^{(n)} = \mathcal{L}^{\Omega^{(n)}, \mu}$. By the assumption of Theorem 5.10, $\Omega^{(n)}$ is a regular negentropy. Thus, by Proposition I.1, $\mathcal{L}^{\Omega^{(n)}, \mu}$ is a regular PERM loss. Since $\mathcal{L}^{(n)} = \mathcal{L}^{\Omega^{(n)}, \mu}$, we are done. \blacksquare

I.2 Totally regular negentropy that is not strongly convex

In this section, we prove Proposition 5.11 which is used to show that there exists totally regular entropies that are not strongly convex. See Example 7. Thus, the associated Fenchel-Young loss is calibrated by Theorem 5.10. Moreover, this calibration result is outside of the purview of previously established results (Blondel, 2019; Nowak-Vila et al., 2019) which requires strong convexity

The next four results serve as the technical tools for proving Proposition 5.11. Their proofs are omitted here and are in the arXiv version of this work (Wang and Scott, 2023b).

Proposition I.6 *Let $f : D \rightarrow \mathbb{R}_{\geq 0}$ be of Legendre type and $g : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ be convex, differentiable and strictly increasing. Let $C = \text{int}(D)$. Suppose that D is compact and there exists $\mathbf{x}^* \in C$ such that $\inf_{\mathbf{x} \in D} f(\mathbf{x}) = f(\mathbf{x}^*)$. Then $g \circ f : D \rightarrow \mathbb{R}$ is of Legendre type.*

Lemma I.7 *Let f and g be as in Proposition I.6. If $g'(0) = 0$ and $\mathbf{x}^* \in \text{int}(D)$ is such that $\inf_{\mathbf{x} \in D} f(\mathbf{x}) = f(\mathbf{x}^*) = 0$, then the Hessian of $g \circ f$ vanishes at \mathbf{x}^* , i.e., $\nabla^2_{g \circ f}(\mathbf{x}^*) = 0$.*

Corollary I.8 *In the situation of Lemma I.7, $g \circ f$ is not α -strongly convex for any $\alpha > 0$.*

Proposition I.9 *Let $\Omega : \Delta^k \rightarrow \mathbb{R}$ be a regular negentropy and let $g : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ be as in Proposition I.6. Furthermore, suppose that g is twice differentiable. Let $a \in \mathbb{R}$ be a negative number such that $a \leq \Omega(\mathbf{p})$ for all $\mathbf{p} \in \Delta^k$. Define $\Theta : \Delta^k \rightarrow \mathbb{R}$ by*

$$\Theta(\mathbf{p}) := g(\Omega(\mathbf{p}) - a) - g(-a), \quad \forall \mathbf{p} \in \Delta^k.$$

Then Θ is a regular negentropy.

Proof [Proof of Proposition 5.11] First we note that the assumptions on g is in the setting of Proposition I.6. Next, by Definition 5.9, we must show that $\Theta^{(n)}$ is a regular negentropy for each $n \in \{2, \dots, k\}$. Let $a = \Omega(\mathbf{u})$. Note that since Ω is symmetric and convex, we must have that $a \leq \Omega(\mathbf{p})$ for all $\mathbf{p} \in \Delta^k$. Furthermore, it is easy to see that $\tilde{\Theta}^{(n)}(\tilde{\mathbf{q}}) := g(\tilde{\Omega}^{(n)}(\tilde{\mathbf{q}}) - a) - g(-a)$ for all $\mathbf{q} \in \tilde{\Delta}^n$. Now, apply Proposition I.9 to $\Theta^{(n)}$ and $a = \Omega(\mathbf{u})$, we get the desired result. The part of the proposition assuming $g'(0) = 0$ follows immediately from Corollary I.8. \blacksquare

Appendix J. Uniqueness of the matrix label code

In this section, we prove Theorem J.1. The proof appears at the end of this section.

Theorem J.1 Let \mathcal{L} be the cross entropy loss and ψ be its template as in Example 1. Suppose that $\{\mathbf{A}_y\}_{y=1}^k$ is a set of $(k-1) \times (k-1)$ matrices satisfying $\mathcal{L}_y(\mathbf{v}) = \psi(\mathbf{A}_y \mathbf{D}\mathbf{v})$ for all $y \in [k]$ and all $\mathbf{v} \in \mathbb{R}^k$. Then for each $y \in [k]$ there exists a permutation $\sigma_y \in \text{Sym}(k)$ so that $\mathbf{A}_y = \mathbf{S}_{\sigma_y} \boldsymbol{\Upsilon}_y$.

Lemma J.2 Let ψ be the template of the multinomial logistic loss (Example 1). Let $\mathbf{a} \in \mathbb{R}^{k-1}$ be a vector. Suppose that $\psi(t\mathbf{a}) = \psi(t\mathbf{e}_1)$ for all $t \in \mathbb{R}$. Then $\mathbf{a} = \mathbf{e}_i$ for some $i \in [k-1]$.

Lemma J.2 serves as the main tool for Lemma J.3 below. The proofs of both Lemmas J.2 and J.3 appear in the arXiv version of our manuscript (Wang and Scott, 2023b).

Lemma J.3 Let ψ be the template of the multinomial logistic loss (Example 1). Suppose that \mathbf{A} is such that $\psi(\mathbf{z}) = \psi(\mathbf{A}\mathbf{z})$ for all $\mathbf{z} \in \mathbb{R}^{k-1}$, then \mathbf{A} is the identity matrix up to row permutation. In other words, there exists $\sigma \in \text{Sym}(k)$ such that $\mathbf{A} = \mathbf{S}_\sigma$.

Lemma J.4 Let ψ be the template of the multinomial logistic loss (Example 1). Let $j \in [k]$ and $\mathbf{A} \in \mathbb{R}^{(k-1) \times (k-1)}$ be arbitrary. Suppose that $\psi(\boldsymbol{\Upsilon}_j \mathbf{z}) = \psi(\mathbf{A}\mathbf{z})$ for all $\mathbf{z} \in \mathbb{R}^{k-1}$. Then \mathbf{A} is equal to the $\boldsymbol{\Upsilon}_j$ up to row permutation. In other words, there exists $\sigma \in \text{Sym}(k)$ such that $\mathbf{A} = \mathbf{S}_\sigma \boldsymbol{\Upsilon}_j$.

Proof [Proof of Lemma J.4] We first prove that the assumption of Lemma J.4 implies that of the previous Lemma J.3: $\psi(\mathbf{z}) = \psi(\mathbf{A}\mathbf{z})$ for all $\mathbf{z} \in \mathbb{R}^{k-1}$. Let $\mathbf{u} \in \mathbb{R}^{k-1}$ be arbitrary and let $\mathbf{z} := \boldsymbol{\Upsilon}_j \mathbf{u}$. Since $\boldsymbol{\Upsilon}_j^2 = \mathbf{I}$, we have $\psi(\mathbf{u}) = \psi(\boldsymbol{\Upsilon}_j \boldsymbol{\Upsilon}_j \mathbf{u}) = \psi(\boldsymbol{\Upsilon}_j \mathbf{z}) = \psi(\mathbf{A}\mathbf{z}) = \psi(\mathbf{A}\boldsymbol{\Upsilon}_j \mathbf{u})$ where for the third equality from the left we used the assumption of Lemma J.4.

Now, by the previous Lemma J.3, we have that $\mathbf{A}\boldsymbol{\Upsilon}_j = \mathbf{S}_\sigma$ for some $\sigma \in \text{Sym}(k)$. Using $\boldsymbol{\Upsilon}_j^2 = \mathbf{I}$, we get $\mathbf{S}_\sigma \boldsymbol{\Upsilon}_j = \mathbf{A}\boldsymbol{\Upsilon}_j \boldsymbol{\Upsilon}_j = \mathbf{A}$. That is to say \mathbf{A} is equal to $\boldsymbol{\Upsilon}_j$ up to permutations of the rows. \blacksquare

Proof [Proof of Theorem J.1] Let $y \in [k]$ and $\mathbf{v} \in \mathbb{R}^k$ be fixed. Let $\mathbf{z} := \mathbf{D}\mathbf{v}$. By Theorem 2.5, we have $\psi(\boldsymbol{\Upsilon}_y \mathbf{z}) = \psi(\mathbf{A}_y \mathbf{z})$. Since \mathbf{D} is surjective, the preceding identity holds for all $\mathbf{z} \in \mathbb{R}^{k-1}$. Thus by Lemma J.4, $\mathbf{A}_y = \mathbf{S}_\sigma \boldsymbol{\Upsilon}_y$ for some $\sigma \in \text{Sym}(k)$. \blacksquare

Appendix K. Mathematical Background

We review mathematical background on fundamental topics that are used throughout this work.

K.1 Non-singular M-matrix

We recall some definitions from linear algebra.

Definition K.1 Let $\mathbf{A} = (a_{ij}) \in \mathbb{R}^{n \times n}$ be a matrix. We say that \mathbf{A} is a

1. Z-matrix if $a_{ij} \leq 0$ whenever $i \neq j$.

2. M-matrix if \mathbf{A} is a Z-matrix and all eigenvalues of \mathbf{A} have nonnegative real parts.
3. strictly diagonally dominant matrix if $|a_{ii}| > \sum_{j \in [n]: j \neq i} |a_{ij}|$ for all $i \in [n]$.
4. monotone matrix if for all $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{Ax} \succeq 0$ implies $\mathbf{x} \succeq 0$. If, in addition, $\mathbf{Ax} > 0$ implies $\mathbf{x} > 0$, then \mathbf{A} is said to be strictly monotone.

Theorem K.2 (Levy-Desplanques/Gershgorin circle) *Let \mathbf{A} be a strictly diagonally dominant matrix. Then \mathbf{A} is non-singular whose eigenvalues all have nonnegative real parts.*

The above result immediately implies the following:

Corollary K.3 *A strictly diagonally dominant Z-matrix is a non-singular M-matrix.*

Non-singular M-matrices have many equivalent characterizations. The one relevant to us is the following:

Theorem K.4 (Plemmons (1977)) *Let \mathbf{A} be a Z-matrix. Then \mathbf{A} is a non-singular M-matrix if and only \mathbf{A} is a monotone matrix.*

Lemma K.5 *Let $\mathbf{A} = (a_{ij}) \in \mathbb{R}^{n \times n}$ be a non-singular M-matrix. If the diagonals of \mathbf{A} are positive, then \mathbf{A} is strictly monotone.*

Proof Let $\mathbf{x} \in \mathbb{R}^n$ be arbitrary such that $\mathbf{Ax} > 0$. Our goal is to show that $\mathbf{x} > 0$. First, from Theorem K.4, we have that \mathbf{A} is monotone. Thus, $\mathbf{Ax} > 0$ implies $\mathbf{x} \succeq 0$. We only have to check additionally that $\mathbf{x} > 0$. Since \mathbf{A} is a Z-matrix, the off-diagonals are non-positive, i.e., $a_{ij} \leq 0$ for all $i, j \in [n]$ such that $i \neq j$. Now, let $i \in [n]$. We need to check that $x_i > 0$. To this end, note that $0 < [\mathbf{Ax}]_i = \sum_{j=1}^n a_{ij}x_j = a_{ii}x_i + \sum_{j \neq i} a_{ij}x_j \leq a_{ii}x_i$. Note that $\sum_{j \neq i} a_{ij}x_j \leq 0$ because $\mathbf{x} \succeq 0$ and $a_{ij} \leq 0$. Finally, $x_i > 0$ since $a_{ii} > 0$. \blacksquare

Corollary K.6 *Let $\mathbf{A} = (a_{ij}) \in \mathbb{R}^{n \times n}$ be a non-singular M-matrix. Then \mathbf{A}^\top is also a non-singular M-matrix. Moreover, if the diagonals of \mathbf{A} are positive, then \mathbf{A}^\top is strictly monotone.*

Proof Note that \mathbf{A} and \mathbf{A}^\top are both non-singular and have the same set of eigenvalues. Moreover, \mathbf{A} being a Z-matrix implies that \mathbf{A}^\top is a Z-matrix. Thus, \mathbf{A}^\top is a non-singular M-matrix. Finally, Lemma K.5 implies that \mathbf{A}^\top is a strictly monotone matrix. \blacksquare