

# Bias-variance decompositions: the exclusive privilege of Bregman divergences

**Tom Heskes**

TOM.HESKES@RU.NL

*Institute for Computing and Information Sciences  
Radboud University  
Nijmegen, The Netherlands*

**Editor:** Mladen Kolar

## Abstract

Bias-variance decompositions are widely used to understand the generalization performance of machine learning models. While the squared error loss permits a straightforward decomposition, other loss functions – such as zero-one loss or  $L_1$  loss – either fail to sum bias and variance to the expected loss or rely on definitions that lack the essential properties of meaningful bias and variance. Recent research has shown that clean decompositions can be achieved for the broader class of Bregman divergences, with the cross-entropy loss as a special case. However, the necessary and sufficient conditions for these decompositions remain an open question.

In this paper, we address this question by studying continuous, nonnegative loss functions that satisfy the identity of indiscernibles (zero loss if and only if the two arguments are identical). We prove that so-called  $g$ -Bregman or rho-tau divergences are the only such loss functions that have a clean bias-variance decomposition. A  $g$ -Bregman divergence can be transformed into a standard Bregman divergence through an invertible change of variables. This makes the squared Euclidean distance, up to such a variable transformation, the only symmetric loss function with a clean bias-variance decomposition. Consequently, common metrics such as 0-1 and  $L_1$  losses cannot admit a clean bias-variance decomposition, explaining why previous attempts have failed. We also examine the impact of relaxing the restrictions on the loss functions and how this affects our results.

**Keywords:** Bias-variance decomposition, Bregman divergence

## 1. Introduction

Bias-variance decompositions have been introduced to the machine learning community by Geman et al. (1992). They decompose the expected loss over a family of machine learning models into a bias term, depending on a suitable central prediction (defined more specifically below, and corresponding to the expected prediction in the case of squared loss), and a variance term, independent of the (distribution of the) actual labels. These decompositions are often used to study and better understand the generalization performance of machine learning approaches, also in more recent work that questions the standard adagium that bias decreases and variance increases as a function of model complexity (Neal et al., 2018; Belkin et al., 2019).

The bias-variance decomposition for the squared error loss has a very appealing form and is also relatively easy to derive. However, generalizing this decomposition to other loss

functions, such as the zero-one loss or the  $L_1$  loss, has proven to be extremely difficult, if not impossible (Wolpert, 1997; Domingos, 2000; James, 2003). Attempts to define bias and variance for these other loss functions often result in decompositions that either do not sum to the expected loss, or require relaxing requirements, such as allowing the bias to depend on the distribution of predictions or the variance to depend on the label (Domingos, 2000; James, 2003).

Fortunately, there are loss functions beyond the squared error that do have clean bias-variance decompositions. For instance, bias-variance decompositions have been derived for loss functions related to the log likelihood of probabilistic models, such as the cross-entropy loss (Heskes, 1998). More recent work has shown that these clean decompositions apply to the entire class of Bregman divergences, which includes the squared error and cross-entropy loss as special cases (Pfau, 2025; Gupta et al., 2022; Wood et al., 2023). However, it is still unclear whether these are the only loss functions with such decompositions. As noted in recent literature, “the necessary and sufficient conditions for such a [clean bias-variance] decomposition are an open research question” (Brown and Ali, 2024).

In this paper, we aim to answer this question. We focus on loss functions that are non-negative and satisfy the identity of indiscernibles (zero loss if and only if the two arguments are identical). These restrictions are specified in Section 2. In Section 3, we discuss the bias-variance decomposition and what we, following earlier work (Neal et al., 2018), refer to as a ‘clean’ bias-variance decomposition. Bregman divergences, and their relation to probability distributions within the exponential family, are introduced in Section 4. Here we also consider a slight generalization of the standard Bregman divergence, the so-called  $g$ -Bregman divergence (Nishiyama, 2018) or rho-tau divergence (Zhang, 2004; Naudts and Zhang, 2018), which can be turned into a standard Bregman divergence through an invertible change of variables. Section 5 proves our main contribution, which is reflected in less subtle form in the title of this paper:  $g$ -Bregman divergences are the only nonnegative loss functions with the identity of indiscernibles that have a clean bias-variance decomposition. In Section 6 we discuss the extent to which we can relax the restrictions on the loss functions and how this affects our results. All main proofs can be found in the appendix. The examples presented throughout this paper are intended to illustrate the ideas in the main text.

## 2. Loss Functions

We consider loss functions  $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$ , not necessarily symmetric, where  $\ell(\mathbf{t}, \mathbf{y})$  measures the error when predicting  $\mathbf{y}$  for the actual label (target)  $\mathbf{t}$ . We assume that  $\mathbf{t}$  and  $\mathbf{y}$  belong to the same domain  $\mathcal{Y} \subseteq \mathbb{R}^d$ , which is a subset of the  $d$ -dimensional Euclidean space and treat  $\mathbf{t}$  and  $\mathbf{y}$  as  $d$ -dimensional column vectors.

Note that our notation slightly deviates from recent machine learning literature, where  $y$  often denotes the labels and  $\hat{y} = f(x)$  the predictions. There the dependency of the prediction on the model  $f$  and inputs  $x$  is made explicit, but for the analysis here this dependency is irrelevant: what matters is the distribution over predictions (and possibly labels), not its origin. Using  $y$  for predictions allows concise notation such as  $\bar{y}$  for central predictions and avoids overloading  $f$ .

Unless specified otherwise, we will assume that  $\ell$  is nonnegative ( $\ell(\mathbf{t}, \mathbf{y}) \geq 0$ ), satisfies the identity of indiscernibles ( $\ell(\mathbf{t}, \mathbf{y}) = 0$  if and only if  $\mathbf{t} = \mathbf{y}$ ), and is continuous everywhere.

### Example 1.

Arguably the best known example of such a loss function is the squared Euclidean distance

$$\ell(\mathbf{t}, \mathbf{y}) = \|\mathbf{t} - \mathbf{y}\|^2 = \sum_{i=1}^d |t_i - y_i|^2. \quad (1)$$

Its domain  $\mathcal{Y}$  is typically the whole of  $\mathbb{R}^d$ . For  $d = 1$  it simplifies to the squared error

$$\ell(t, y) = (t - y)^2.$$

If the labels  $t_i$  and predictions  $y_i$  can be interpreted as probabilities, the squared Euclidean distance becomes the Brier score. With  $d$  probabilities, we can set  $\mathcal{Y}$  to the probability simplex:

$$\mathcal{Y} = \Delta^{d-1} = \left\{ \mathbf{y} \in [0, 1]^d \mid \sum_{i=1}^d y_i = 1 \right\}.$$

When optimizing loss functions, we need to take into account such equality constraints. An alternative to the Brier score is the cross-entropy loss or Kullback-Leibler divergence, described below in Example 5.

The squared Euclidean distance is smooth and differentiable in both its arguments. It can be generalized to a broader class of Minkowski-type losses, as in the following example, which are not necessarily differentiable everywhere.

### Example 2.

We can generalize the squared Euclidean distance to loss functions of the form

$$\ell(\mathbf{t}, \mathbf{y}) = \|\mathbf{t} - \mathbf{y}\|_\epsilon^\epsilon = \sum_{i=1}^d |t_i - y_i|^\epsilon.$$

For  $\epsilon = 2$ , we get back the squared Euclidean loss and for  $\epsilon = 1$  the  $L_1$  norm. For  $0 < \epsilon < 2$ , this loss function is infinitely differentiable everywhere, except for the set  $S = \{(\mathbf{t}, \mathbf{y}) \in \mathcal{Y} \times \mathcal{Y} \mid \exists i \text{ such that } t_i = y_i\}$ .

We explicitly do not require the loss function  $\ell$  to be a metric: it need not be symmetric, nor does it need to satisfy the triangle inequality.

### 3. Bias-Variance Decomposition

Suppose that we have some distribution of predictions  $\mathbf{Y}$ , for example resulting from machine learning models trained on different training sets and/or started from different initializations. We write  $\mathbb{E}_{\mathbf{Y}}$  for the expectation with respect to such a distribution. This expectation can be taken over any relevant quantity, whether it involves a probability density over continuous predictions, a probability mass distribution over discrete predictions, or a finite ensemble of models generating predictions. For a fixed vector label  $\mathbf{t}$ , the expected squared Euclidean distance can be decomposed as

$$\underbrace{\mathbb{E}_{\mathbf{Y}} \|\mathbf{t} - \mathbf{Y}\|^2}_{\text{expected loss}} = \underbrace{\|\mathbf{t} - \mathbb{E}_{\mathbf{Y}} \mathbf{Y}\|^2}_{\text{bias}} + \underbrace{\mathbb{E}_{\mathbf{Y}} \|\mathbb{E}_{\mathbf{Y}} \mathbf{Y} - \mathbf{Y}\|^2}_{\text{variance}} .$$

We will refer to the first term on the right-hand side as the bias and the second term as the variance. The important observation here is that the variance is independent of the label  $\mathbf{t}$ , whereas the bias only measures the error between the label and the expected prediction. If we now take into account that labels  $\mathbf{T}$  are randomly drawn from some distribution and then also consider the expectation  $\mathbb{E}_{\mathbf{T}}$  with respect to this distribution, we can further decompose

$$\underbrace{\mathbb{E}_{\mathbf{T}, \mathbf{Y}} \|\mathbf{T} - \mathbf{Y}\|^2}_{\text{expected loss}} = \underbrace{\mathbb{E}_{\mathbf{T}} \|\mathbf{T} - \mathbb{E}_{\mathbf{T}} \mathbf{T}\|^2}_{\text{intrinsic noise}} + \underbrace{\|\mathbb{E}_{\mathbf{T}} \mathbf{T} - \mathbb{E}_{\mathbf{Y}} \mathbf{Y}\|^2}_{\text{bias}} + \underbrace{\mathbb{E}_{\mathbf{Y}} \|\mathbb{E}_{\mathbf{Y}} \mathbf{Y} - \mathbf{Y}\|^2}_{\text{variance}} . \quad (2)$$

The first term on the right-hand side is referred to as the intrinsic noise or irreducible error: it provides a lower bound on the loss that can be obtained even with the best possible predictions.

Following Neal et al. (2018), we refer to (2) as a ‘clean’ bias-variance decomposition that, for general loss functions, we define as follows.

**Definition 1.** *A loss function  $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  is said to have a clean bias-variance decomposition if, for all valid distributions of labels  $\mathbf{T}$  and predictions  $\mathbf{Y}$ , the expected loss decomposes as*

$$\underbrace{\mathbb{E}_{\mathbf{T}, \mathbf{Y}} \ell(\mathbf{T}, \mathbf{Y})}_{\text{expected loss}} = \underbrace{\mathbb{E}_{\mathbf{T}} \ell(\mathbf{T}, \mathbf{t}^*)}_{\text{intrinsic noise}} + \underbrace{\ell(\mathbf{t}^*, \mathbf{y}^*)}_{\text{bias}} + \underbrace{\mathbb{E}_{\mathbf{Y}} \ell(\mathbf{y}^*, \mathbf{Y})}_{\text{variance}} , \quad (3)$$

with

$$\mathbf{t}^* = \arg \min_{\mathbf{y} \in \mathcal{Y}} \mathbb{E}_{\mathbf{T}} \ell(\mathbf{T}, \mathbf{y}) \quad \text{and} \quad \mathbf{y}^* = \arg \min_{\mathbf{t} \in \mathcal{Y}} \mathbb{E}_{\mathbf{Y}} \ell(\mathbf{t}, \mathbf{Y}) . \quad (4)$$

In the decomposition (3), all three terms have a natural interpretation. The irreducible error is independent of the distribution of the predictions, the bias only depends on the distributions of the labels and the predictions through summary statistics, and the variance is independent of the distribution of the labels. Most importantly, these terms collectively sum to the expected loss. Wood et al. (2023) refer to such bias-variance decompositions as ‘additive’ decompositions. In this paper, we adopt the term ‘clean’ to highlight that the definitions of bias and variance have a natural interpretation, viewing additivity as an inherent requirement of any decomposition.

By definition,  $\mathbf{t}^*$  is the prediction closest to the distribution of the labels and, vice versa,  $\mathbf{y}^*$  the label closest to the distribution of the predictions. Here “closest” is interpreted relative to the argument being optimized:  $\mathbf{t}^*$  minimizes the loss over the first argument, while  $\mathbf{y}^*$  minimizes the loss over the second argument. In the literature on Bregman divergences,  $\mathbf{t}^*$  and  $\mathbf{y}^*$  are referred to as right(-sided) and left(-sided) centroid (Brown and Ali, 2024; Nielsen and Nock, 2009), or the central label and prediction (Gupta et al., 2022), respectively.

The specific ordering of  $\mathbf{t}^*$  and  $\mathbf{y}^*$  in (3) is intentional and important: any other ordering is doomed to fail for all asymmetric loss functions (and does not make a difference for symmetric loss functions).

**Proposition 2.** *For asymmetric loss functions  $\ell$  with the identity of indiscernibles, the specific ordering of arguments in the loss terms on the right-hand side of (3) is the only one that yields a clean bias-variance decomposition.*

Furthermore, for the decomposition (3) to hold,  $\mathbf{y}^*$  must be the unique minimizer of the variance term and  $\mathbf{t}^*$  the unique minimizer of the intrinsic noise term. Thus, their definition in (4) is not merely a convenient requirement but in fact a necessary consequence of having a clean bias-variance decomposition.

**Proposition 3.** *Suppose the bias-variance decomposition (3) holds for a loss function  $\ell$  with the identity of indiscernibles, using some  $\mathbf{t}^*$  independent of the distribution of predictions  $\mathbf{Y}$  and some  $\mathbf{y}^*$  independent of the distribution of labels  $\mathbf{T}$ . Then  $\mathbf{t}^*$  and  $\mathbf{y}^*$  must coincide with the central label and central prediction from (4), minimizing the intrinsic noise and variance terms, respectively. Moreover, these minimizers are necessarily unique.*

The uniqueness of the minimizer of the variance term results from the loss function’s clean bias-variance decomposition, rather than assumptions like (quasi-)convexity. Similarly, we do not need to specify whether the domain  $\mathcal{Y}$  is convex, open, or closed: the existence of a clean bias-variance decomposition and hence of the bias  $\ell(\mathbf{t}^*, \mathbf{y}^*)$  implies that the central label and central prediction must be within the same domain  $\mathcal{Y}$ . If they are not, we can try to extend the domain  $\mathcal{Y}$  such that they are.

## 4. Bregman Divergences

Bregman divergences (Bregman, 1967) form a specific class of loss functions and have become quite popular in the recent machine learning literature (Siahkamari et al., 2020; Cilinçir et al., 2020; Amid et al., 2019). The Bregman divergence  $\mathcal{D}_A : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$ , with  $\mathcal{Y}$  a convex subset of  $\mathbb{R}^d$ , is fully determined by its generating function  $A : \mathcal{Y} \rightarrow \mathbb{R}$ :

$$\mathcal{D}_A(\mathbf{t}, \mathbf{y}) = (\mathbf{y} - \mathbf{t})^\top \nabla_{\mathbf{y}} A(\mathbf{y}) - A(\mathbf{y}) + A(\mathbf{t}). \quad (5)$$

$A$  is strictly convex and differentiable. A Bregman divergence is strictly convex in its first argument, but not necessarily in its second argument (Bauschke and Borwein, 2001). Banerjee et al. (2005b) provide many other properties as well as examples of Bregman divergences.

**Example 3.**

The squared Euclidean distance (1) from Example 1 follows by taking  $A(\mathbf{y}) = \|\mathbf{y}\|^2$ . Generalizing this to  $A(\mathbf{y}) = \mathbf{y}^\top K \mathbf{y}$  for a positive definite matrix  $K$ , we obtain the squared Mahalanobis distance

$$\ell(\mathbf{t}, \mathbf{y}) = (\mathbf{t} - \mathbf{y})^\top K (\mathbf{t} - \mathbf{y}) .$$

The squared Mahalanobis distance is the only symmetric Bregman divergence: all other Bregman divergences are asymmetric (Boissonnat et al., 2010).

Bregman divergences relate to distributions in the exponential family as follows. In canonical form, any distribution in the exponential family can be written

$$p_{\mathbf{y},A}(\mathbf{z}) = h(\mathbf{z}) \exp[\mathbf{y}^\top \boldsymbol{\phi}(\mathbf{z}) - A(\mathbf{y})] ,$$

with  $h(\mathbf{z})$  the base function,  $\boldsymbol{\phi}(\mathbf{z})$  the sufficient statistics,  $A(\mathbf{y})$  the log-partition function, and  $\mathbf{y}$  the canonical (or natural) parameters. A Bregman divergence  $\mathcal{D}_A$  between parameters  $\mathbf{t}$  and  $\mathbf{y}$  equals the Kullback-Leibler divergence between two distributions in the exponential family with log-partition function  $A$  and canonical parameters  $\mathbf{y}$  and  $\mathbf{t}$  (for example, Chowdhury et al., 2023):

$$\text{KL}(p_{\mathbf{y},A}, p_{\mathbf{t},A}) = (\mathbf{y} - \mathbf{t})^\top \nabla_{\mathbf{y}} A(\mathbf{y}) - A(\mathbf{y}) + A(\mathbf{t}) = \mathcal{D}_A(\mathbf{t}, \mathbf{y}) ,$$

where we used

$$\mathbb{E}_{\mathbf{Z}} \boldsymbol{\phi}(\mathbf{Z}) = \nabla_{\mathbf{y}} A(\mathbf{y}) .$$

**Example 4.**

A univariate Gaussian distribution with mean  $m_y$  and variance  $v_y$  can be written in the form

$$p_{\mathbf{y},A}(z) = \frac{1}{\sqrt{2\pi v_y}} \exp \left[ -\frac{1}{2v_y} (m_y - z)^2 \right] = h(z) \exp \left[ \mathbf{y}^\top \boldsymbol{\phi}(z) - A(\mathbf{y}) \right] ,$$

by choosing  $\mathbf{y} = (m_y/v_y, -1/(2v_y))$ ,  $h(z) = 1$ ,  $\boldsymbol{\phi}(z) = (z, z^2)$ , and

$$A(\mathbf{y}) = -\frac{y_1^2}{4y_2} - \frac{1}{2} \log \left( -\frac{y_2}{\pi} \right) ,$$

where  $y_1 \in \mathbb{R}$  and  $y_2 \in \mathbb{R}_{<0}$ . Plugging this into the definition of the Bregman divergence, we obtain for the KL divergence between two Gaussian distributions with canonical parameters  $\mathbf{y}$  and  $\mathbf{t}$  the intriguing form

$$\ell(\mathbf{t}, \mathbf{y}) = -\frac{1}{2} + \frac{y_1 t_1}{2y_2} - \frac{y_1^2 t_2}{4y_2^2} + \frac{t_2}{2y_2} - \frac{t_1^2}{4t_2} - \frac{1}{2} \log \left( \frac{t_2}{y_2} \right) , \quad (6)$$

which may look more familiar when translated back to means and variances:

$$\ell(\mathbf{t}, \mathbf{y}) = \frac{(m_y - m_t)^2}{2v_t} - \frac{1}{2} \left[ 1 - \frac{v_y}{v_t} + \log \left( \frac{v_y}{v_t} \right) \right] . \quad (7)$$

Another, more direct connection between the exponential family and Bregman divergences goes through the log likelihood. Consider a member of the exponential family with canonical parameter  $\tilde{\mathbf{y}}$ , sufficient statistic  $\phi(\mathbf{z})$ , and log-partition function  $B(\tilde{\mathbf{y}})$ . Suppose that our model outputs an estimate of the canonical parameter  $\tilde{\mathbf{y}}$ , and receives observation  $\mathbf{z}$ . Then a natural, and often used loss function to measure the quality of our estimate is minus the log likelihood, which, up to irrelevant constants independent of  $\tilde{\mathbf{y}}$ , reads

$$-\log p_B(\mathbf{z}; \tilde{\mathbf{y}}) = -\tilde{\mathbf{y}}^\top \phi(\mathbf{z}) + B(\tilde{\mathbf{y}}) + \text{const} . \quad (8)$$

This loss function does not yet satisfy our requirements of nonnegativity and identity of indiscernibles. However, we can turn it into one by rewriting the observation  $\mathbf{z}$  in terms of the sufficient statistic  $\mathbf{t} = \phi(\mathbf{z})$  and (Legendre) transforming the canonical parameter  $\tilde{\mathbf{y}}$ :

$$\mathbf{y} = \nabla_{\tilde{\mathbf{y}}} B(\tilde{\mathbf{y}}) ,$$

and, correspondingly,  $B$  to its convex conjugate

$$A(\mathbf{y}) = \tilde{\mathbf{y}}^\top \nabla_{\tilde{\mathbf{y}}} B(\tilde{\mathbf{y}}) - B(\tilde{\mathbf{y}}) .$$

such that

$$\tilde{\mathbf{y}} = \nabla_{\mathbf{y}} B(\mathbf{y}) \quad \text{and} \quad B(\tilde{\mathbf{y}}) = \mathbf{y}^\top \nabla_{\mathbf{y}} A(\mathbf{y}) - A(\mathbf{y}) .$$

Substituting these into (8), we obtain, again up to constants independent of  $\tilde{\mathbf{y}}$  and  $\mathbf{y}$ ,

$$-\log p_B(\mathbf{z}; \tilde{\mathbf{y}}) = (\mathbf{y} - \mathbf{t})^\top \nabla_{\mathbf{y}} A(\mathbf{y}) - A(\mathbf{y}) + \text{const} . \quad (9)$$

By adding  $A(\mathbf{t})$ , we arrive at the Bregman divergence  $\mathcal{D}_A(\mathbf{t}, \mathbf{y})$  from (5).

### Example 5.

Let us consider the case of probabilistic binary classification with Bernoulli likelihood function

$$p(t; y) = y^t (1 - y)^{1-t} ,$$

with  $t \in \{0, 1\}$  and  $y \in [0, 1]$ . Taking the negative log likelihood, we obtain the cross-entropy loss

$$-\log p(t; y) = -t \log y - (1 - t) \log(1 - y) ,$$

with the convention  $0 \log 0 = 0$ . Thus, when  $y \in \{0, 1\}$ , the loss is finite if  $t = y$ , and infinite otherwise. This is indeed of the form (9) if we set  $A(y) = y \log y + (1 - y) \log(1 - y)$ . The cross-entropy loss by itself may not be a Bregman divergence, but can be turned into one by adding  $A(t)$  to arrive at the Kullback-Leibler divergence

$$\ell(t, y) = t \log \left( \frac{t}{y} \right) + (1 - t) \log \left( \frac{1 - t}{1 - y} \right) ,$$

which works fine for any  $t \in [0, 1]$  and is equal to the cross-entropy loss at the extremes  $t = 0$  and  $t = 1$ .

Generalizing this to the multinomial, multi-class case, for  $\mathbf{t}, \mathbf{y} \in \Delta^{d-1}$  (see Example 1), we obtain the KL divergence

$$\ell(\mathbf{t}, \mathbf{y}) = \sum_{i=1}^d t_i \log \left( \frac{t_i}{y_i} \right) + \sum_{i=1}^d y_i - \sum_{i=1}^d t_i, \quad (10)$$

which is a Bregman divergence with generating function  $A(\mathbf{y}) = \sum_{i=1}^d y_i \log y_i - \sum_{i=1}^d y_i$ . Adding  $\sum_{i=1}^d y_i - \sum_{i=1}^d t_i$  to the loss function makes the Kullback-Leibler divergence a proper Bregman divergence that satisfies (5). Expressed in this form, the KL divergence guarantees the identity of indiscernibles also when we extend  $\mathcal{Y}$  from the probability simplex  $\Delta^{d-1}$  to the unit cube  $[0, 1]^d$ , that is, when we ignore the equality constraint  $\sum_i y_i = 1$ :

$$\arg \min_{\mathbf{y} \in \Delta^{d-1}} \ell(\mathbf{t}, \mathbf{y}) = \arg \min_{\mathbf{y} \in [0, 1]^d} \ell(\mathbf{t}, \mathbf{y}) = \mathbf{t} \quad \text{and} \quad \arg \min_{\mathbf{t} \in \Delta^{d-1}} \ell(\mathbf{t}, \mathbf{y}) = \arg \min_{\mathbf{t} \in [0, 1]^d} \ell(\mathbf{t}, \mathbf{y}) = \mathbf{y}.$$

We will refer to a Bregman divergence that has the property that the identity of indiscernibles still holds when we ignore equality constraints as a Bregman divergence in proper form.

For now, we will assume that the observations  $\mathbf{z}$  and the sufficient statistics/labels  $\mathbf{t}$  (and hence parameters/predictions  $\mathbf{y}$ ) have the same dimensionality. In Section 6, we examine scenarios where the observation  $\mathbf{z}$  has a lower dimensionality than the label  $\mathbf{t}$ .

We consider a slightly generalized variant of the Bregman divergence, by Nishiyama (2018) referred to as the  $g$ -Bregman divergence and by Naudts and Zhang (2018) and Brekelmans and Nielsen (2024) as the rho-tau divergence. We write  $\mathbf{g}(\mathcal{Y})$  for the image of  $\mathcal{Y}$  under mapping  $\mathbf{g}$  and define the set  $\mathcal{Y} \setminus \mathcal{Y}_{\text{eq}}$  as the original  $\mathcal{Y}$ , but then without taking into account any (linear) equality constraints. So, in Example 5 above,  $\mathcal{Y}$  refers to the probability simplex  $\Delta^{d-1}$ ,  $\mathcal{Y}_{\text{eq}}$  is the hyperplane  $\{y \in \mathcal{Y} : \sum_i y_i = 1\}$  enforcing the equality constraint, and  $\mathcal{Y} \setminus \mathcal{Y}_{\text{eq}}$  then corresponds to the  $d$ -dimensional cube  $[0, 1]^d$ .

**Definition 4.** Given an invertible (bijective) function  $\mathbf{g} : \mathcal{Y} \rightarrow \mathbf{g}(\mathcal{Y})$  and strictly convex generating function  $A : \mathbf{g}(\mathcal{Y}) \rightarrow \mathbb{R}$ , with  $\mathcal{Y}$  and  $\mathbf{g}(\mathcal{Y} \setminus \mathcal{Y}_{\text{eq}})$  both convex subsets of  $\mathbb{R}^d$ , the  $g$ -Bregman divergence  $\mathcal{D}_{A, \mathbf{g}} : \mathcal{Y} \times \mathcal{Y}$  is defined as

$$\mathcal{D}_{A, \mathbf{g}}(\mathbf{t}, \mathbf{y}) = \mathcal{D}_A(\mathbf{g}(\mathbf{t}), \mathbf{g}(\mathbf{y})) = (\mathbf{g}(\mathbf{y}) - \mathbf{g}(\mathbf{t}))^\top \nabla_{\mathbf{g}(\mathbf{y})} A(\mathbf{g}(\mathbf{y})) - A(\mathbf{g}(\mathbf{y})) + A(\mathbf{g}(\mathbf{t})), \quad (11)$$

Loosely speaking, a  $g$ -Bregman divergence  $\mathcal{D}_{A, \mathbf{g}}$  between  $\mathbf{t}$  and  $\mathbf{y}$  is a standard Bregman divergence between  $\tilde{\mathbf{t}} = \mathbf{g}(\mathbf{t})$  and  $\tilde{\mathbf{y}} = \mathbf{g}(\mathbf{y})$ . It can be interpreted as the Kullback-Leibler divergence between two distributions in the exponential family in noncanonical form, with canonical parameter function  $\mathbf{g}$  and log-partition function  $A \circ \mathbf{g}$ . A  $g$ -Bregman divergence need not be convex in its first, nor in its second argument. It is nonnegative and by construction satisfies the identity of indiscernibles.

**Example 6.**

For  $\mathbf{t}, \mathbf{y} \in \Delta^{d-1}$ , the reverse KL divergence,

$$\ell(\mathbf{t}, \mathbf{y}) = \sum_{i=1}^d y_i \log \left( \frac{y_i}{t_i} \right) + \sum_{i=1}^d t_i - \sum_{i=1}^d y_i, \quad (12)$$

is a  $g$ -Bregman divergence with mapping  $\mathbf{g}$  such that  $g_i(\mathbf{y}) = \log y_i$  and  $A(\mathbf{g}) = \sum_{i=1}^d \exp(g_i)$ .

The forward KL divergence (10) and the reverse KL divergence (12) can be considered special cases of the so-called  $\alpha$ -divergence (Amari, 1985)

$$\ell(\mathbf{t}, \mathbf{y}) = -\frac{1}{\alpha(1-\alpha)} \sum_{i=1}^d t_i^\alpha y_i^{1-\alpha} + \frac{1}{1-\alpha} \sum_{i=1}^d t_i + \frac{1}{\alpha} \sum_{i=1}^d y_i, \quad (13)$$

where we follow the convention of Zhu and Rohwer (1995). For  $0 < \alpha < 1$ , this  $\alpha$ -divergence is a  $g$ -Bregman divergence with mapping  $\mathbf{g}$  such that  $g_i(\mathbf{y}) = \frac{1}{1-\alpha} y_i^\alpha$  and  $A(\mathbf{g}) = (1-\alpha)^{(1-\alpha)/\alpha} \sum_{i=1}^d g_i^{1/\alpha}$ . The forward KL divergence (10) is obtained in the limit  $\alpha \rightarrow 1$  and the reverse KL divergence (12) in the limit  $\alpha \rightarrow 0$ .

**Proposition 5.** *The reverse of a  $g$ -Bregman divergence, with  $\mathbf{t}$  and  $\mathbf{y}$  interchanged, is another  $g$ -Bregman divergence:*

$$\mathcal{D}_{A,\mathbf{g}}(\mathbf{t}, \mathbf{y}) = \mathcal{D}_{B,\mathbf{f}}(\mathbf{y}, \mathbf{t}) \quad \text{with} \quad \mathbf{f}(\mathbf{y}) = \nabla_{\mathbf{g}(\mathbf{y})} A(\mathbf{g}(\mathbf{y})) \quad \text{and} \quad B(\mathbf{f}(\mathbf{y})) = \mathbf{g}(\mathbf{y})^\top \mathbf{f}(\mathbf{y}) - A(\mathbf{g}(\mathbf{y})).$$

With  $\mathcal{D}_{A,\mathbf{g}}(\mathbf{t}, \mathbf{y}) = \mathcal{D}_{B,\mathbf{f}}(\mathbf{y}, \mathbf{t})$ , we can write the  $g$ -Bregman divergence in the concise form

$$\mathcal{D}_{A,\mathbf{g}}(\mathbf{t}, \mathbf{y}) = A(\mathbf{g}(\mathbf{t})) - \mathbf{f}(\mathbf{y})^\top \mathbf{g}(\mathbf{t}) + B(\mathbf{f}(\mathbf{y})). \quad (14)$$

This form is sometimes referred to as the mixed parameterization or canonical form (Breklemans and Nielsen, 2024) and is reminiscent of the Fenchel-Young divergence (Blondel et al., 2020).  $A$  and  $B$  are each other's convex conjugates, in the sense that for all  $\mathbf{y} \in \mathcal{Y}$ ,

$$A^*(\mathbf{f}(\mathbf{y})) = \sup_{\mathbf{t} \in \mathcal{Y}} \mathbf{g}(\mathbf{t})^\top \mathbf{f}(\mathbf{y}) - A(\mathbf{g}(\mathbf{t})) = B(\mathbf{f}(\mathbf{y})). \quad (15)$$

We will therefore refer to  $\{A, \mathbf{g}\}$  and  $\{B, \mathbf{f}\}$  as each other's dual pair. The duality (15) still applies for  $\mathbf{y} \in \mathcal{Y}$  when equality constraints are involved, but then we are no longer guaranteed that  $A^*(\mathbf{f}) = B(\mathbf{f})$  for arbitrary  $\mathbf{f}$ . The form (14) is very convenient to directly read off the dual pairs  $\{A, \mathbf{g}\}$  and  $\{B, \mathbf{f}\}$  from any  $g$ -Bregman divergence in proper form.

Given a  $g$ -Bregman divergence  $\mathcal{D}_{A,\mathbf{g}}$  with dual pair  $\{B, \mathbf{f}\}$ , we will call  $\mathbf{g}(\mathbf{t})$  the label  $\mathbf{t}$  in canonical (natural) form and  $\mathbf{f}(\mathbf{y})$  the prediction  $\mathbf{y}$  in moment (expectation) form. We define

$$\bar{\mathbf{t}} = \mathbf{g}^{-1}(\mathbb{E}_{\mathbf{T}} \mathbf{g}(\mathbf{T})) \quad \text{and} \quad \bar{\mathbf{y}} = \mathbf{f}^{-1}(\mathbb{E}_{\mathbf{Y}} \mathbf{f}(\mathbf{Y})),$$

the  $\mathbf{g}$ -mean label and  $\mathbf{f}$ -mean prediction, respectively, that follow by averaging the labels in canonical form and predictions in moment form and transforming the resulting averages

back to the original space. These means play a role in the next lemmas and theorem, which demonstrate that, in the absence of equality constraints, all  $g$ -Bregman divergences have a clean bias-variance decomposition.

**Lemma 6.** *Consider a  $g$ -Bregman divergence  $\mathcal{D}_{A,\mathbf{g}} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  with dual pair  $\{B, \mathbf{f}\}$ . If the  $\mathbf{g}$ -mean label  $\bar{\mathbf{t}}$  satisfies any constraints  $\mathcal{Y}$  may impose, that is, if  $\bar{\mathbf{t}} \in \mathcal{Y}$ , then it minimizes the intrinsic noise term:*

$$\bar{\mathbf{t}} = \mathbf{t}^* = \arg \min_{\mathbf{y} \in \mathcal{Y}} \mathbb{E}_{\mathbf{T}} \mathcal{D}_{A,\mathbf{g}}(\mathbf{T}, \mathbf{y}),$$

yielding

$$\text{intrinsic noise} = \mathbb{E}_{\mathbf{T}} \mathcal{D}_{A,\mathbf{g}}(\mathbf{T}, \mathbf{t}^*) = \mathbb{E}_{\mathbf{T}} A(\mathbf{g}(\mathbf{T})) - A(\mathbf{g}(\mathbf{t}^*)).$$

Along the same lines, if the  $\mathbf{f}$ -mean prediction  $\bar{\mathbf{y}}$  satisfies any constraints  $\mathcal{Y}$  may impose, that is, if  $\bar{\mathbf{y}} \in \mathcal{Y}$ , then it minimizes the variance term:

$$\bar{\mathbf{y}} = \mathbf{y}^* = \arg \min_{\mathbf{t} \in \mathcal{Y}} \mathbb{E}_{\mathbf{Y}} \mathcal{D}_{A,\mathbf{g}}(\mathbf{t}, \mathbf{Y}), \quad (16)$$

yielding

$$\text{variance} = \mathbb{E}_{\mathbf{Y}} \mathcal{D}_{B,\mathbf{f}}(\mathbf{Y}, \mathbf{y}^*) = \mathbb{E}_{\mathbf{Y}} B(\mathbf{f}(\mathbf{Y})) - B(\mathbf{f}(\mathbf{y}^*)).$$

**Lemma 7.** *Consider a  $g$ -Bregman divergence  $\mathcal{D}_{A,\mathbf{g}} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  with generating function  $A : \mathbf{g}(\mathcal{Y}) \rightarrow \mathbb{R}$  and invertible mapping  $\mathbf{g} : \mathcal{Y} \rightarrow \mathbf{g}(\mathcal{Y})$ . Assume that the central label  $\mathbf{t}^*$  minimizing the intrinsic noise term is equal to the  $\mathbf{g}$ -mean label  $\mathbf{g}^{-1}(\mathbb{E}_{\mathbf{T}} \mathbf{g}(\mathbf{T}))$ , and that the central prediction  $\mathbf{y}^*$  minimizing the variance term is equal to the  $\mathbf{f}$ -mean prediction  $\mathbf{f}^{-1}(\mathbb{E}_{\mathbf{Y}} \mathbf{f}(\mathbf{Y}))$ , where  $\mathbf{f}(\mathbf{y}) = \nabla_{\mathbf{g}(\mathbf{y})} A(\mathbf{g}(\mathbf{y}))$ . Then  $\mathcal{D}_{A,\mathbf{g}}$  has a clean bias-variance decomposition:*

$$\underbrace{\mathbb{E}_{\mathbf{T}, \mathbf{Y}} \mathcal{D}_{A,\mathbf{g}}(\mathbf{T}, \mathbf{Y})}_{\text{expected loss}} = \underbrace{\mathbb{E}_{\mathbf{T}} \mathcal{D}_{A,\mathbf{g}}(\mathbf{T}, \mathbf{t}^*)}_{\text{intrinsic noise}} + \underbrace{\mathcal{D}_{A,\mathbf{g}}(\mathbf{t}^*, \mathbf{y}^*)}_{\text{bias}} + \underbrace{\mathbb{E}_{\mathbf{Y}} \mathcal{D}_{A,\mathbf{g}}(\mathbf{y}^*, \mathbf{Y})}_{\text{variance}}. \quad (17)$$

### Example 7.

To illustrate the bias-variance decomposition of a  $g$ -Bregman divergence, we return to Example 4. Here we derived a (standard) Bregman divergence between the parameters of two univariate Gaussian distributions. The loss function (6) has a quite peculiar form, essentially because for a Bregman divergence we need to fix a parameterization such that the additive term in the loss function that depends on both  $\mathbf{t}$  and  $\mathbf{y}$  is linear in  $\mathbf{t}$ .

The more natural form (7) can be viewed as a  $g$ -Bregman divergence between parameters  $\mathbf{t} = (m_t, v_t)$  and  $\mathbf{y} = (m_y, v_y)$  with

$$\mathbf{g}(m, v) = \left( \frac{m}{v}, -\frac{1}{2v} \right) \quad \text{and} \quad \mathbf{f}(m, v) = (m, m^2 + v).$$

In this parameterization, the bias-variance decomposition is relatively straightforward to derive. The central label follows from averaging the canonical parameters

and transforming these back to moment form:

$$v_t^* = (\mathbb{E}_{V_t}(V_t^{-1}))^{-1} \quad \text{and} \quad m_t^* = v_t^* \mathbb{E}_{M_t, V_t} \left( \frac{M_t}{V_t} \right),$$

To arrive at the central prediction we need to average the moment parameters

$$m_y^* = \mathbb{E}_{M_y} M_y \quad \text{and} \quad v_y^* = \mathbb{E}_{M_y, V_y} (M_y^2 + V_y) - m_y^{*2}.$$

The bias-variance decomposition then reads

$$\underbrace{\mathbb{E}_{\mathbf{T}, \mathbf{Y}} \ell(\mathbf{T}, \mathbf{Y})}_{\text{expected loss}} = \underbrace{\mathbb{E}_{\mathbf{T}} \tilde{A}(\mathbf{T}) - \tilde{A}(\mathbf{t}^*)}_{\text{intrinsic noise}} + \underbrace{\ell(\mathbf{t}^*, \mathbf{y}^*)}_{\text{bias}} + \underbrace{\mathbb{E}_{\mathbf{Y}} \tilde{B}(\mathbf{Y}) - \tilde{B}(\mathbf{y}^*)}_{\text{variance}},$$

with

$$\tilde{A}(\mathbf{t}) = A(\mathbf{g}(m_t, v_t)) = \frac{m_t^2}{2v_t} + \frac{1}{2} \log(2\pi v_t) \quad \text{and} \quad \tilde{B}(\mathbf{y}) = B(\mathbf{f}(m_y, v_y)) = -\frac{1}{2} \log(2\pi v_y) - \frac{1}{2}.$$

**Theorem 8.** *Consider a  $g$ -Bregman divergence  $\mathcal{D}_{A,g} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  with generating function  $A : \mathbf{g}(\mathcal{Y}) \rightarrow \mathbb{R}$  and invertible mapping  $\mathbf{g} : \mathcal{Y} \rightarrow \mathbf{g}(\mathcal{Y})$ , and assume that  $\mathbf{g}(\mathcal{Y}) = \text{dom}(A)$  is a convex feasible set defined by inequality constraints only. Then this  $g$ -Bregman divergence has a clean bias-variance decomposition of the form (17).*

When  $\mathcal{Y} = \text{dom}(\mathbf{g})$  includes equality constraints, the situation becomes considerably more complicated. If  $\mathbf{g}$  is a nonlinear function, the linear constraints in  $\mathcal{Y}$  turn into nonlinear constraints in  $\mathbf{g}(\mathcal{Y})$ . Consequently, this transformation can lead to a nonconvex optimization problem, as the feasible region defined by the nonlinear constraints may no longer be convex. We therefore first consider the case of a standard Bregman divergence, that is, with  $\mathbf{g}$  the identity function.

**Theorem 9.** *Consider a Bregman divergence  $\mathcal{D}_A : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  with generating function  $A : \mathcal{Y} \rightarrow \mathbb{R}$ , and dual pair  $\{B, \mathbf{f}\}$ . The feasible set  $\mathcal{Y}$  contains equality constraints that can be written in the form  $W\mathbf{y} = \mathbf{b}$ , with  $W \in \mathbb{R}^{\kappa \times d}$ ,  $\mathbf{b} \in \mathbb{R}^{\kappa}$ , and  $\kappa$  the number of constraints. Then this Bregman divergence has a clean bias-variance decomposition of the form*

$$\underbrace{\mathbb{E}_{\mathbf{T}, \mathbf{Y}} \mathcal{D}_A(\mathbf{T}, \mathbf{Y})}_{\text{expected loss}} = \underbrace{\mathbb{E}_{\mathbf{T}} A(\mathbf{T}) - A(\mathbf{t}^*)}_{\text{intrinsic noise}} + \underbrace{\mathcal{D}_A(\mathbf{t}^*, \mathbf{y}^*)}_{\text{bias}} + \underbrace{\boldsymbol{\lambda}^\top \mathbf{b} + \mathbb{E}_{\mathbf{Y}} B(\mathbf{f}(\mathbf{Y})) - B(\mathbf{f}(\mathbf{y}^*))}_{\text{variance}},$$

where

$$\mathbf{t}^* = \mathbb{E}_{\mathbf{T}} \mathbf{T} \quad \text{and} \quad \mathbf{y}^* = \mathbf{f}^{-1} \left( \mathbb{E}_{\mathbf{Y}} \mathbf{f}(\mathbf{Y}) + \boldsymbol{\lambda}^\top W \right) \quad \text{with } \boldsymbol{\lambda} \text{ such that } W\mathbf{y}^* = \mathbf{b}.$$

Because of the symmetry between  $\mathbf{t}$  and  $\mathbf{y}$  in our analysis, every statement about Bregman divergences applies equally to reverse Bregman divergences. This leads to the following theorem.

**Corollary 10.** Consider a reverse Bregman divergence, that is, a  $g$ -Bregman divergence  $\mathcal{D}_{A,g} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  with generating function  $A : \mathcal{Y} \rightarrow \mathbb{R}$ , and dual pair  $\{B, \mathbf{f}\}$  such that  $\mathbf{f}(\mathbf{y}) = \mathbf{y}$ . The feasible set  $\mathcal{Y}$  contains equality constraints that can be written in the form  $W\mathbf{y} = \mathbf{b}$ , with  $W \in \mathbb{R}^{\kappa \times d}$ ,  $\mathbf{b} \in \mathbb{R}^{\kappa}$ , and  $\kappa$  the number of constraints. Then this  $g$ -Bregman divergence has a clean bias variance decomposition of the form

$$\underbrace{\mathbb{E}_{\mathbf{T}, \mathbf{Y}} \mathcal{D}_{A,g}(\mathbf{T}, \mathbf{Y})}_{\text{expected loss}} = \underbrace{\mathbb{E}_{\mathbf{T}} A(\mathbf{g}(\mathbf{T})) - A(\mathbf{g}(\mathbf{t}^*)) + \boldsymbol{\lambda}^\top \mathbf{b}}_{\text{intrinsic noise}} + \underbrace{\mathcal{D}_{A,g}(\mathbf{t}^*, \mathbf{y}^*)}_{\text{bias}} + \underbrace{\mathbb{E}_{\mathbf{Y}} B(\mathbf{Y}) - B(\mathbf{y}^*)}_{\text{variance}},$$

where

$$\mathbf{t}^* = \mathbf{g}^{-1} \left( \mathbb{E}_{\mathbf{T}} \mathbf{g}(\mathbf{T}) + \boldsymbol{\lambda}^\top W \right) \quad \text{with } \boldsymbol{\lambda} \text{ such that } W\mathbf{t}^* = \mathbf{b}, \text{ and } \mathbf{y}^* = \mathbb{E}_{\mathbf{Y}} \mathbf{Y}.$$

The proof follows directly by interchanging the roles of  $\mathbf{t}$  and  $\mathbf{y}$  in Theorem 9.

### Example 8.

Let us revisit the KL divergence (10) from Example 5. This is a (standard) Bregman divergence, so  $\mathbf{g}(\mathbf{t}) = \mathbf{t}$ , with generating function  $A(\mathbf{g}) = \sum_{i=1}^d g_i \log g_i - \sum_{i=1}^d g_i$ . Its dual pair  $\{B, \mathbf{f}\}$  is such that  $f_i(\mathbf{y}) = \log y_i$  and  $B(\mathbf{f}) = \sum_{i=1}^d \exp(f_i)$ . With  $\mathbf{1}$  the  $d$ -dimensional column vector of all ones,  $W = \mathbf{1}^\top$ , and  $b = 1$ , the conditions for Theorem 9 are satisfied and we get the bias-variance decomposition, in correspondence with Heskes (1998),

$$\underbrace{\mathbb{E}_{\mathbf{T}, \mathbf{Y}} \sum_{i=1}^d t_i \log \left( \frac{t_i}{y_i} \right)}_{\text{expected loss}} = \underbrace{\mathbb{E}_{\mathbf{T}} \sum_{i=1}^d t_i \log t_i - \sum_{i=1}^d t_i^* \log t_i^*}_{\text{intrinsic noise}} + \underbrace{\sum_{i=1}^d t_i^* \log \left( \frac{t_i^*}{y_i^*} \right)}_{\text{bias}} + \underbrace{\log Z}_{\text{variance}},$$

with

$$t_i^* = \mathbb{E}_{\mathbf{T}} T_i, \quad y_i^* = \frac{1}{Z} \exp(\mathbb{E}_{\mathbf{Y}} \log Y_i), \quad \text{and} \quad Z = \sum_{i=1}^d \exp(\mathbb{E}_{\mathbf{Y}} \log Y_i).$$

We obtain the bias-variance decomposition for the reverse KL divergence (12) from Example 6 by interchanging the role of labels and predictions:

$$\underbrace{\mathbb{E}_{\mathbf{T}, \mathbf{Y}} \sum_{i=1}^d y_i \log \left( \frac{y_i}{t_i} \right)}_{\text{expected loss}} = \underbrace{\log Z}_{\text{intrinsic noise}} + \underbrace{\sum_{i=1}^d y_i^* \log \left( \frac{y_i^*}{t_i^*} \right)}_{\text{bias}} + \underbrace{\mathbb{E}_{\mathbf{Y}} \sum_{i=1}^d y_i \log y_i - \sum_{i=1}^d y_i^* \log y_i^*}_{\text{variance}},$$

with

$$y_i^* = \mathbb{E}_{\mathbf{Y}} Y_i, \quad t_i^* = \frac{1}{Z} \exp(\mathbb{E}_{\mathbf{T}} \log T_i), \quad \text{and} \quad Z = \sum_{i=1}^d \exp(\mathbb{E}_{\mathbf{T}} \log T_i).$$

So, where for a standard KL divergence the central prediction follows from taking the geometric mean, the central prediction for the reverse KL divergence is the arithmetic

mean. For both the forward and the reverse KL divergence, we still arrive at a clean bias-variance decomposition when incorporating the normalization constraint on the probabilities.

The  $\alpha$ -divergence (13) for  $0 < \alpha < 1$  is not a standard Bregman divergence, but can be interpreted as a  $g$ -Bregman divergence. Without the normalization constraint on the probabilities, so for  $\mathcal{Y} = [0, 1]^d$ , it admits a clean bias-variance decomposition. However, when incorporating the normalization constraint, so for  $\mathcal{Y} = \Delta^{d-1}$ , this is no longer the case. The minimizers of the intrinsic noise term and the variance term follow from taking the power mean

$$t_i^* = \frac{(\mathbb{E}_{\mathbf{T}} T_i^\alpha)^{1/\alpha}}{\sum_{j=1}^d (\mathbb{E}_{\mathbf{T}} T_j^\alpha)^{1/\alpha}} \quad \text{and} \quad y_i^* = \frac{(\mathbb{E}_{\mathbf{Y}} Y_i^{1-\alpha})^{1/(1-\alpha)}}{\sum_{j=1}^d (\mathbb{E}_{\mathbf{Y}} Y_j^{1-\alpha})^{1/(1-\alpha)}},$$

We may still plug these into our definitions for the intrinsic noise, bias, and variance on the right-hand side of (3), but these then do not add up to the expected loss on the left-hand side. An alternative attempt could be to explicitly incorporate the normalization constraint and for example write

$$\ell(\mathbf{t}, \mathbf{y}) = -\frac{1}{\alpha(1-\alpha)} \left( \sum_{i=1}^{d-1} t_i^\alpha y_i^{1-\alpha} + \phi(\mathbf{t})^\alpha \phi(\mathbf{y})^{1-\alpha} \right) + \frac{1}{\alpha(1-\alpha)},$$

with  $\phi(\mathbf{t}) = 1 - \sum_{i=1}^{d-1} t_i$  and where now  $\mathcal{Y} = \{\mathbf{y} \in [0, 1]^{d-1} \mid \sum_i y_i \leq 1\}$ . However, this resulting loss function no longer has the form of a  $g$ -Bregman divergence and then also does not admit a clean bias-variance decomposition.

In summary,  $g$ -Bregman divergences have a clean bias-variance decomposition in the absence of equality constraints. This property extends to both (standard) Bregman and reverse Bregman divergences, even when linear equality constraints are present. Our results are a generalization of those provided by Pfau (2025) and Gupta et al. (2022) for standard Bregman divergences without equality constraints. Beyond relatively simple cases – such as those involving reversible linear transformations  $\mathbf{g}$  or  $\mathbf{f}$  – it appears challenging to identify other  $g$ -Bregman divergences that admit a clean bias-variance decomposition under equality constraints.

## 5. Uniqueness

In the previous section, we have introduced  $g$ -Bregman divergences with a clean bias-variance decomposition. In this section we will show that  $g$ -Bregman divergences in fact are the only (nonnegative, with the identity of indiscernibles) loss functions that can have this property.

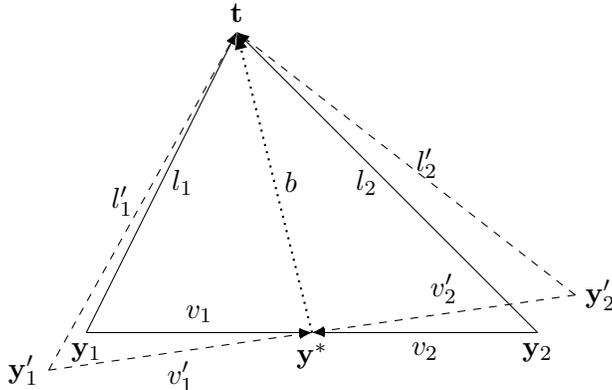


Figure 1: Sketch of the construction used in the proof of Lemma 11. We start from a distribution with just two predictions  $\mathbf{y}_1$  and  $\mathbf{y}_2$ , here each with probability  $1/2$ , leading to the central prediction  $\mathbf{y}^*$ . Each line segment symbolically represents a loss – for example,  $v_1 = \ell(\mathbf{y}^*, \mathbf{y}_1)$  – with the arrow pointing from the second to the first argument. Specializing the bias-variance decomposition (18) to this two prediction case gives  $(l_1 + l_2)/2 = b + (v_1 + v_2)/2$ . The parallelogram indicated by the solid lines is essentially the same construction as in Figure 10 of Nielsen (2021).

Next, we consider a slight, infinitesimal change in the predictions, yielding  $\mathbf{y}'_1$  and  $\mathbf{y}'_2$ , that keeps the central prediction  $\mathbf{y}^*$  and hence the bias  $b$  invariant. This invariance implies  $(l'_1 + l'_2) - (v'_1 + v'_2) = 2b = (l_1 + l_2) - (v_1 + v_2)$ . Requiring this to hold for any such change that keeps the central prediction intact – even when  $\mathbf{y}^*$  is not close to the label  $\mathbf{t}$  – places a strong constraint on the mixed second derivative of the loss function.

To simplify the analysis, we will for now ignore the distribution over labels and consider the case of a single, fixed label vector  $\mathbf{t}$ , for which we get

$$\underbrace{\mathbb{E}_{\mathbf{Y}} \ell(\mathbf{t}, \mathbf{Y})}_{\text{expected loss}} = \underbrace{\ell(\mathbf{t}, \mathbf{y}^*)}_{\text{bias}} + \underbrace{\mathbb{E}_{\mathbf{Y}} \ell(\mathbf{y}^*, \mathbf{Y})}_{\text{variance}}, \quad (18)$$

for some  $\mathbf{y}^*$  dependent on the distribution of the labels, but independent of the label  $\mathbf{t}$ . In the following lemma, we show that the second derivative of loss functions that satisfy (18) factorizes in a special manner. Our line of reasoning is inspired by the analysis of Banerjee et al. (2005a).

**Lemma 11.** *Consider a loss function  $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  that has a clean bias-variance decomposition. Then its mixed second derivative w.r.t. prediction  $\mathbf{y}$  and label  $\mathbf{t}$ ,  $\mathcal{L}_{\mathbf{y}, \mathbf{t}}(\mathbf{t}, \mathbf{y}) = \nabla_{\mathbf{y}} \nabla_{\mathbf{t}} \ell(\mathbf{t}, \mathbf{y})$ , necessarily factorizes as*

$$\mathcal{L}_{\mathbf{y}, \mathbf{t}}(\mathbf{t}, \mathbf{y}) = H_2(\mathbf{y}) H_1^\top(\mathbf{t}), \quad (19)$$

for some matrix functions  $H_1 : \mathcal{Y} \rightarrow \mathbb{R}^{d \times d}$  and  $H_2 : \mathcal{Y} \rightarrow \mathbb{R}^{d \times d}$  for any  $(\mathbf{t}, \mathbf{y}) \in \mathcal{Y} \times \mathcal{Y}$ .

The proof is given in the appendix and uses the construction sketched in Figure 1. We start from a distribution over labels  $\mathbf{Y}$  that corresponds to some central prediction  $\mathbf{y}^*$ . Next

we will make a tiny (infinitesimal) change in the distribution, which we deliberately choose such that  $\mathbf{y}^*$  and hence the bias  $\ell(\mathbf{t}, \mathbf{y}^*)$  stay the same. The expected loss and the variance will change, but, for loss functions with a clean bias-variance distribution, the difference between these two equals the bias and must therefore be invariant under such a change for any choice of  $\mathbf{t}$  and  $\mathbf{y}^*$ .

The difference between the expected loss and the variance stays the same for any sufficiently smooth  $\ell$ , as long as  $\mathbf{t}$  and  $\mathbf{y}^*$  are (infinitesimally) close. However, the requirement that this invariance also holds for  $\mathbf{t}$  and  $\mathbf{y}^*$  further apart, imposes a serious constraint on the loss function. In particular, invariance under all perturbations preserving  $\mathbf{y}^*$  implies that differences of prediction gradients must lie in a fixed linear subspace determined by the mixed derivative of the loss. This enforces a separable structure of the mixed derivative and yields the factorization in (19).

With the help of Lemma 11, we can now prove our main theorem.

**Theorem 12.** *Consider a continuous, nonnegative loss function  $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  that satisfies the identity of indiscernibles and has a clean bias-variance decomposition. Then  $\ell$  must be a  $g$ -Bregman divergence  $\mathcal{D}_{A, \mathbf{g}}$  of the form (11) for some invertible mapping  $\mathbf{g} : \mathcal{Y} \rightarrow \mathbf{g}(\mathcal{Y})$  and strictly convex generating function  $A : \mathbf{g}(\mathcal{Y}) \rightarrow \mathbb{R}$ .*

The proof can again be found in the appendix. Integrating the mixed derivative  $\mathcal{L}_{\mathbf{y}, \mathbf{t}}$  from Lemma 11 twice, first with respect to  $\mathbf{t}$  and then with respect to  $\mathbf{y}$ , we return to the original loss function  $\ell$ . The requirement of the identity of indiscernibles imposes constraints on the various components of this loss function, which are only satisfied by  $g$ -Bregman divergences.

**Corollary 13.** *Within the class of continuous, nonnegative loss functions with the identity of indiscernibles, the ‘ $g$ -generalized’ squared Euclidean distance*

$$\ell(\mathbf{t}, \mathbf{y}) = \|\mathbf{g}(\mathbf{t}) - \mathbf{g}(\mathbf{y})\|^2,$$

*with  $g : \mathcal{Y} \rightarrow \mathbf{g}(\mathcal{Y})$  an invertible mapping, is the only symmetric loss function that has a clean bias-variance decomposition.*

The proof follows directly from Theorem 12 and the fact that the squared Mahalanobis distance is the only symmetric Bregman divergence (Boissonnat et al., 2010). The squared Mahalanobis distance can be written as a ‘ $g$ -generalized’ squared Euclidean distance via an invertible linear mapping, and since the composition of invertible mappings is again invertible, this yields the stated form.

Since the squared Euclidean distance does not satisfy the triangle inequality, this also implies that no metric can admit a clean bias-variance decomposition. In particular, well-studied losses such as the 0-1 loss or the  $L_1$  (Manhattan) distance cannot be decomposed in this way: there exists no central prediction that separates the bias term from a variance term completely independent of the labels. This provides a rigorous explanation of why previous attempts, such as (Wolpert, 1997; Domingos, 2000; James, 2003), could not yield a decomposition with a clean separation of bias and variance.

**Corollary 14.** *Consider the class of continuous, nonnegative loss functions with the identity of indiscernibles. If a loss function within this class has a clean bias-variance decomposition,*

then the central prediction (assuming no equality constraints) corresponds to an  $\mathbf{f}$ -mean prediction, that is,

$$\mathbf{y}^* = \arg \min_{\mathbf{t} \in \mathcal{Y}} \mathbb{E}_{\mathbf{Y}} \ell(\mathbf{t}, \mathbf{Y}) = \mathbf{f}^{-1}(\mathbb{E}_{\mathbf{Y}} \mathbf{f}(\mathbf{Y})) ,$$

for some invertible mapping  $\mathbf{f} : \mathcal{Y} \rightarrow \mathbf{f}(\mathcal{Y})$ . And vice versa, if the central prediction derived from a loss function within the class boils down to an  $\mathbf{f}$ -mean prediction, then this loss function must have a clean bias-variance decomposition.

The proof from loss function to moment averaging follows because such a loss function must be a  $g$ -Bregman divergence, which has this property of moment averaging as shown in the proof of Theorem 8. The proof from moment averaging to loss functions follows from Banerjee et al. (2005a), who showed that Bregman divergences are the only loss functions for which (in our terms) the central label  $\mathbf{t}^*$  is always equal to the average label  $\mathbb{E}_{\mathbf{T}} \mathbf{T}$ .

Our uniqueness result in Theorem 12 can be viewed as a generalization of earlier work by Hansen and Heskes (2000). They already argued that loss functions with a clean bias-variance decomposition must be of the form (18). However, the analysis was restricted to loss functions that measure the error between univariate labels and predictions. Furthermore, although the analysis did start from distributions in the exponential family and followed the reasoning leading to (9), it did not establish the connection to Bregman divergences.

## 6. Relaxing the Assumptions

We imposed some restrictions on the loss functions under consideration, most notably the identity of indiscernibles. In this section, we discuss whether these are necessary and how the bias-variance decomposition changes, and can to some extent be saved, when we relax them.

The identity of indiscernibles,  $\ell(\mathbf{t}, \mathbf{y}) = 0$  if and only if  $\mathbf{t} = \mathbf{y}$ , is quite heavily used throughout this paper. It can be relaxed to  $\ell(\mathbf{t}, \mathbf{y}) = 0$  if and only if  $\mathbf{t} = \mathbf{c}(\mathbf{y})$ , where  $\mathbf{c} : \mathcal{Y} \rightarrow \mathcal{Y}'$  is an involution: a function such that  $\mathbf{c} \circ \mathbf{c}$  is the identity. In other words,  $\mathbf{c}(\mathbf{c}(\mathbf{t})) = \mathbf{t}$  or equivalently  $\mathbf{c}^{-1}(\mathbf{t}) = \mathbf{c}(\mathbf{t})$  for all  $\mathbf{t}$ . Obvious examples of such functions include  $\mathbf{c}(\mathbf{t}) = -\mathbf{t}$  and  $\mathbf{c}((t_1, t_2)) = (t_2, t_1)$ . If  $\mathbf{c}$  is an invertible function but not an involution, we can transform the loss function into one with the identity of indiscernibles through a change of variables in either  $\mathbf{t}$  or  $\mathbf{y}$ .

Related to the identity of indiscernibles is our assumption that labels  $\mathbf{t}$  and predictions  $\mathbf{y}$  have the same domain and hence the same dimensionality. We can relax this assumption, which is particularly relevant for loss functions derived from the log likelihood when the dimensionality of the observations is lower than that of the predictions.

### Example 9.

A prototypical example of a loss function with unequal domains for the labels and the predictions is minus the log likelihood of the form (8) for a Gaussian distribution with canonical parameters  $\tilde{\mathbf{y}} = (m/v, -1/(2v))$  and sufficient statistics  $\mathbf{t} = \phi(z) = (z, z^2)$ . This loss function is popular for fitting so-called mean-variance estimation neural

networks (Nix and Weigend, 1994; Sluijterman et al., 2024) that output predictions for both the mean  $m$  and the variance  $v$ , while receiving observation  $z$ .

Following the line of reasoning that led to (9), we can rework minus the log likelihood, up to irrelevant constants independent of  $\tilde{\mathbf{y}}$  and  $\mathbf{y}$ , to the form

$$\begin{aligned} -\log p_B(z; \tilde{\mathbf{y}}) &= -\mathbf{t}^\top \tilde{\mathbf{y}} + B(\tilde{\mathbf{y}}) + \text{const} = A(\mathbf{t}) - \mathbf{t}^\top \mathbf{f}(\mathbf{y}) + B(\mathbf{f}(\mathbf{y})) + \text{const} = \\ &= \mathcal{D}_A(\mathbf{t}, \mathbf{y}) + \text{const}, \end{aligned} \quad (20)$$

with appropriate choices for  $A$  and  $\mathbf{f}$ , while for the moment ignoring the constraint that  $t_2 = t_1^2$ , that is, for arbitrary  $\mathbf{t} \in \mathcal{Y} \setminus \mathcal{Y}_{\text{eq}}$ . We can then define the central prediction through

$$\mathbf{y}^* = \arg \min_{\mathbf{t} \in \mathcal{Y} \setminus \mathcal{Y}_{\text{eq}}} \mathbb{E}_{\mathbf{Y}} \mathcal{D}_A(\mathbf{t}, \mathbf{Y}) = \mathbf{f}^{-1}(\mathbb{E}_{\mathbf{Y}} \mathbf{f}(\mathbf{Y})),$$

yielding the usual

$$\text{variance} = \min_{\mathbf{t} \in \mathcal{Y} \setminus \mathcal{Y}_{\text{eq}}} \mathbb{E}_{\mathbf{Y}} \mathcal{D}_A(\mathbf{t}, \mathbf{Y}) = B(\mathbf{f}(\mathbf{y}^*)) - \mathbb{E}_{\mathbf{Y}} B(\mathbf{f}(\mathbf{Y})).$$

When the constraint  $t_2 = t_1^2$  is not ignored, the construction in (20) fails because  $A(\mathbf{t})$  goes to infinity on the manifold  $\mathbf{t} = (z, z^2)$ . Luckily, not adding  $A(\mathbf{t})$  in (20) still gives a perfectly valid bias-variance decomposition for a single observation  $z$ :

$$\underbrace{-\mathbb{E}_{\tilde{\mathbf{Y}}} \log p_B(z; \tilde{\mathbf{Y}})}_{\text{expected loss}} = \underbrace{-\log p_B(z; \tilde{\mathbf{y}}^*)}_{\text{bias}} + \underbrace{B(\tilde{\mathbf{y}}^*) - \mathbb{E}_{\tilde{\mathbf{Y}}} B(\tilde{\mathbf{Y}})}_{\text{variance}}. \quad (21)$$

The first term on the right-hand side can be interpreted as the bias and the combination of the second and third term as the variance.

This bias-variance decomposition, which applies to any log likelihood derived from a probability distribution within the exponential family, is completely in line with the bias-variance decomposition for likelihood-based estimators (Heskes, 1998). Viewed as a loss function,  $-\log p_B(z; \tilde{\mathbf{y}})$  is no longer guaranteed to be nonnegative. It is also not easy to see how to properly define the central observation  $z^*$  from a distribution of observations, nor how to decompose the expectation  $\mathbb{E}_Z -\log p_B(Z; \tilde{\mathbf{y}}^*)$  into an intrinsic error and a bias.

The decomposition of Heskes (1998) in principle applies to arbitrary sets of distributions. The central prediction in that decomposition is the geometric average of the predictive distributions. For exponential families, this geometric average is again a member of the same family and the Kullback-Leibler divergence corresponds to a Bregman divergence, where the parameters of the exponential family play the role of labels and predictions (see Section 4). For distributions outside the exponential family, the Kullback-Leibler divergence itself still defines a Bregman divergence, but now expressed in terms of the probabilities themselves rather than a lower-dimensional parameterization; in this case, the decomposition still holds with the probabilities as labels and predictions.

Many classification losses similarly involve predictions and labels from different domains. Margin losses, which depend on the product of the label  $t \in \{-1, 1\}$  and the prediction score  $y \in \mathbb{R}$  written as  $\ell(ty)$  for some  $\ell : \mathbb{R} \rightarrow \mathbb{R}^+$ , do not satisfy the identity of indiscernibles for the loss itself. Nevertheless, Wood et al. (2022) show that a subclass – the gradient-symmetric margin losses – can be expressed as Bregman divergences, either via their conditional excess risk or by anchoring the second argument to a label-dependent constant, and then do admit a clean bias-variance decomposition similar to (21). This connection to Bregman divergences is crucial: it explains why these losses admit a clean bias-variance decomposition, and suggests that without such a connection, no margin loss would have this property. In this sense, the results of Wood et al. (2022) are perfectly in line with ours, providing a complementary perspective on how Bregman-type structure enables decomposability.

## 7. Conclusion

Our analysis reveals that loss functions permitting a clean bias-variance decomposition can be expressed in the simplified form

$$\ell(\mathbf{t}, \mathbf{y}) = \mathbf{h}_1(\mathbf{y})^\top \mathbf{h}_2(\mathbf{t}) + H_3(\mathbf{y}) + H_4(\mathbf{t}) \quad (22)$$

for certain choices of  $\mathbf{h}_{1,2} : \mathcal{Y} \rightarrow \mathbb{R}^d$  and  $H_{3,4} : \mathcal{Y} \rightarrow \mathbb{R}$ . This compact representation summarizes the same structure that appears, in more elaborate form, in the proof of our main Theorem 12, and in the more specific forms given in (11) and (14).

The key characteristic is the factorization of the interaction term between labels  $\mathbf{t}$  and predictions  $\mathbf{y}$  into an inner product of two separate functions. By imposing the natural requirements of nonnegativity and the identity of indiscernibles, we show that such loss functions must necessarily take the form of  $g$ -Bregman divergences.  $g$ -Bregman divergences are then the only loss functions with a clean bias-variance decomposition.

The utility of bias-variance decompositions for studying the generalization of machine learning models has already been called into question, for example by Brown and Ali (2024), who highlighted that the central prediction may not always be realizable by any machine learning model within the class under consideration. For instance, with single decision stumps, the central prediction  $\mathbf{y}^*$  for squared error is an average over stumps, which cannot itself be represented by a single stump. Our finding, that clean bias-variance decompositions only apply to (minor variations of) Bregman divergences, adds to these concerns. While Bregman divergences are widely used and include prominent examples like squared error and cross-entropy loss, they are inherently sensitive to outliers due to their convexity in the label space (Vemuri et al., 2011; Zhang et al., 2023). This convexity (or quasi-convexity when considering  $g$ -Bregman divergences) appears to be a necessary property for a clean bias-variance decomposition, implying that such decompositions do not extend to more robust loss functions that are less affected by outliers.

Although much broader than just the squared Euclidean distance, the class of loss functions for which clean bias-variance decompositions can be obtained is more special than perhaps previously believed. Future work could consider the generalization of these results to bias-variance decompositions in function spaces (Frigyik et al., 2008; Gruber and Buettner, 2023) or bias-variance-diversity decompositions for ensembles (Wood et al., 2023).

## Acknowledgments

Special thanks to the reviewer for excellent feedback on an earlier version of this manuscript and to Rob Brekelmans for directing attention to the work on rho-tau divergences.

## Appendix A.

In this appendix, we provide the remaining proofs for the propositions, lemmas, and theorems in the main text.

**Proposition 2.** *For asymmetric loss functions  $\ell$  with the identity of indiscernibles, the specific ordering of arguments in the loss terms on the right-hand side of (3) is the only one that yields a clean bias-variance decomposition.*

**Proof** Let us first consider interchanging  $\mathbf{t}^*$  and  $\mathbf{y}^*$  in  $\ell(\mathbf{t}^*, \mathbf{y}^*)$  to try

$$\mathbb{E}_{\mathbf{T}, \mathbf{Y}} \ell(\mathbf{T}, \mathbf{Y}) = \mathbb{E}_{\mathbf{T}} \ell(\mathbf{T}, \mathbf{t}^*) + \ell(\mathbf{y}^*, \mathbf{t}^*) + \mathbb{E}_{\mathbf{Y}} \ell(\mathbf{y}^*, \mathbf{Y}).$$

This must also hold in the special case of an arbitrary single label  $\mathbf{t}$  and arbitrary single prediction  $\mathbf{y}$ , which simplifies to

$$\mathbb{E}_{\mathbf{T}, \mathbf{Y}} \ell(\mathbf{T}, \mathbf{Y}) = \ell(\mathbf{t}, \mathbf{y}) = \ell(\mathbf{t}, \mathbf{t}^*) + \ell(\mathbf{y}^*, \mathbf{t}^*) + \ell(\mathbf{y}, \mathbf{y}^*) = \ell(\mathbf{y}, \mathbf{t}),$$

since, by identity of indiscernibles,  $\mathbf{y}^* = \mathbf{y}$  and  $\mathbf{t}^* = \mathbf{t}$ . That is, the bias-variance decomposition with the ordering  $\ell(\mathbf{y}^*, \mathbf{t}^*)$  instead of  $\ell(\mathbf{t}^*, \mathbf{y}^*)$  forces  $\ell$  to be symmetric, in contradiction with being appropriate for asymmetric loss functions.

Similar reasoning rules out interchanging  $\mathbf{T}$  and  $\mathbf{t}^*$  in the intrinsic noise term  $\mathbb{E}_{\mathbf{T}} \ell(\mathbf{T}, \mathbf{t}^*)$ . In case of a single prediction  $\mathbf{y} = \mathbf{y}^*$ , the bias-variance decomposition with a reverse intrinsic noise term would give

$$\mathbb{E}_{\mathbf{T}, \mathbf{Y}} \ell(\mathbf{T}, \mathbf{Y}) = \mathbb{E}_{\mathbf{T}} \ell(\mathbf{T}, \mathbf{y}^*) = \mathbb{E}_{\mathbf{T}} \ell(\mathbf{t}^*, \mathbf{T}) + \ell(\mathbf{t}^*, \mathbf{y}^*).$$

This must also apply when  $\mathbf{y}^*$  happens to coincide with  $\mathbf{t}^*$ , which yields  $\mathbb{E}_{\mathbf{T}} \ell(\mathbf{T}, \mathbf{t}^*) = \mathbb{E}_{\mathbf{T}} \ell(\mathbf{t}^*, \mathbf{T})$  and must be valid for any distribution of labels  $\mathbf{T}$ . Again, we conclude that  $\ell$  then must be symmetric for the bias-variance decomposition to uphold. The exact same reasoning applies to interchanging  $\mathbf{y}^*$  and  $\mathbf{Y}$  in  $\mathbb{E}_{\mathbf{Y}} \ell(\mathbf{y}^*, \mathbf{Y})$  as well as to changing the order in two or even all three terms. ■

**Proposition 3.** *Suppose the bias-variance decomposition (3) holds for a loss function  $\ell$  with the identity of indiscernibles, using some  $\mathbf{t}^*$  independent of the distribution of predictions  $\mathbf{Y}$  and some  $\mathbf{y}^*$  independent of the distribution of labels  $\mathbf{T}$ . Then  $\mathbf{t}^*$  and  $\mathbf{y}^*$  must coincide with the central label and central prediction from (4), minimizing the intrinsic noise and variance terms, respectively. Moreover, these minimizers are necessarily unique.*

**Proof** To prove that  $\mathbf{y}^*$  must be the central prediction, we consider the special case of a single vector label  $\mathbf{t}$  and suppose that the decomposition holds for some  $\tilde{\mathbf{y}}$ , that is,

$$\mathbb{E}_{\mathbf{Y}} \ell(\mathbf{t}, \mathbf{Y}) = \ell(\mathbf{t}, \tilde{\mathbf{y}}) + \mathbb{E}_{\mathbf{Y}} \ell(\tilde{\mathbf{y}}, \mathbf{Y}).$$

Substituting  $\mathbf{t} = \mathbf{y}^*$ , with  $\mathbf{y}^*$  by definition the label minimizing the variance term, we get

$$0 \leq \ell(\mathbf{y}^*, \tilde{\mathbf{y}}) = \mathbb{E}_{\mathbf{Y}} \ell(\mathbf{y}^*, \mathbf{Y}) - \mathbb{E}_{\mathbf{Y}} \ell(\tilde{\mathbf{y}}, \mathbf{Y}) \leq 0,$$

where the first inequality follows from the nonnegativity of the loss function, and the second inequality follows from  $\mathbf{y}^*$  minimizing the variance term. By the identity of indiscernibles,  $\ell(\mathbf{y}^*, \tilde{\mathbf{y}}) = 0$  implies  $\tilde{\mathbf{y}} = \mathbf{y}^*$ . This argument and the outcome  $\mathbf{y}^* = \tilde{\mathbf{y}}$  apply, even when  $\tilde{\mathbf{y}}$  is a minimizer of the variance term. Hence the variance term must have a unique minimizer. The proof for the central label  $\mathbf{t}^*$  follows the exact same steps, simply interchanging the roles of predictions and labels.  $\blacksquare$

**Proposition 5.** *The reverse of a  $g$ -Bregman divergence, with  $\mathbf{t}$  and  $\mathbf{y}$  interchanged, is another  $g$ -Bregman divergence:*

$$\mathcal{D}_{A,\mathbf{g}}(\mathbf{t}, \mathbf{y}) = \mathcal{D}_{B,\mathbf{f}}(\mathbf{y}, \mathbf{t}) \quad \text{with } \mathbf{f}(\mathbf{y}) = \nabla_{\mathbf{g}(\mathbf{y})}A(\mathbf{g}(\mathbf{y})) \quad \text{and} \quad B(\mathbf{f}(\mathbf{y})) = \mathbf{g}(\mathbf{y})^\top \mathbf{f}(\mathbf{y}) - A(\mathbf{g}(\mathbf{y})).$$

**Proof** See Theorem 6 of Zhang (2004). A  $g$ -Bregman divergence is a sum of three terms, each depending on both  $\mathbf{y}$  and  $\mathbf{t}$ , only on  $\mathbf{y}$ , or only on  $\mathbf{t}$ :

$$\begin{aligned} \mathcal{D}_{A,\mathbf{g}}(\mathbf{t}, \mathbf{y}) &= \underbrace{-\mathbf{g}(\mathbf{t})^\top \nabla_{\mathbf{g}(\mathbf{y})}A(\mathbf{g}(\mathbf{y}))}_{\text{both } \mathbf{y} \text{ and } \mathbf{t}} + \underbrace{\mathbf{g}(\mathbf{y})^\top \nabla_{\mathbf{g}(\mathbf{y})}A(\mathbf{g}(\mathbf{y})) - A(\mathbf{g}(\mathbf{y}))}_{\text{only } \mathbf{y}} + \underbrace{A(\mathbf{g}(\mathbf{t}))}_{\text{only } \mathbf{t}} \\ &= \underbrace{-\mathbf{f}(\mathbf{y})^\top \nabla_{\mathbf{f}(\mathbf{t})}B(\mathbf{f}(\mathbf{t}))}_{\text{both } \mathbf{y} \text{ and } \mathbf{t}} + \underbrace{B(\mathbf{f}(\mathbf{y}))}_{\text{only } \mathbf{y}} + \underbrace{\mathbf{f}(\mathbf{t})^\top \nabla_{\mathbf{f}(\mathbf{t})}B(\mathbf{f}(\mathbf{t})) - B(\mathbf{f}(\mathbf{t}))}_{\text{only } \mathbf{t}} \\ &= \mathcal{D}_{B,\mathbf{f}}(\mathbf{y}, \mathbf{t}). \end{aligned}$$

Equating the first term depending on both  $\mathbf{y}$  and  $\mathbf{t}$  leads to  $\mathbf{f}(\mathbf{y}) = \nabla_{\mathbf{g}(\mathbf{y})}A(\mathbf{g}(\mathbf{y}))$ . Equating the second term only depending on  $\mathbf{y}$  gives  $B(\mathbf{f}(\mathbf{y})) = \mathbf{g}(\mathbf{y})^\top \mathbf{f}(\mathbf{y}) - A(\mathbf{g}(\mathbf{y}))$ . The other terms precisely match with the reverse mapping from  $\{B, \mathbf{f}\}$  to  $\{A, \mathbf{g}\}$ .  $\blacksquare$

**Lemma 6.** *Consider a  $g$ -Bregman divergence  $\mathcal{D}_{A,\mathbf{g}} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  with dual pair  $\{B, \mathbf{f}\}$ . If the  $\mathbf{g}$ -mean label  $\bar{\mathbf{t}}$  satisfies any constraints  $\mathcal{Y}$  may impose, that is, if  $\bar{\mathbf{t}} \in \mathcal{Y}$ , then it minimizes the intrinsic noise term:*

$$\bar{\mathbf{t}} = \mathbf{t}^* = \arg \min_{\mathbf{y} \in \mathcal{Y}} \mathbb{E}_{\mathbf{T}} \mathcal{D}_{A,\mathbf{g}}(\mathbf{T}, \mathbf{y}),$$

yielding

$$\text{intrinsic noise} = \mathbb{E}_{\mathbf{T}} \mathcal{D}_{A,\mathbf{g}}(\mathbf{T}, \mathbf{t}^*) = \mathbb{E}_{\mathbf{T}} A(\mathbf{g}(\mathbf{T})) - A(\mathbf{g}(\mathbf{t}^*)).$$

Along the same lines, if the  $\mathbf{f}$ -mean prediction  $\bar{\mathbf{y}}$  satisfies any constraints  $\mathcal{Y}$  may impose, that is, if  $\bar{\mathbf{y}} \in \mathcal{Y}$ , then it minimizes the variance term:

$$\bar{\mathbf{y}} = \mathbf{y}^* = \arg \min_{\mathbf{t} \in \mathcal{Y}} \mathbb{E}_{\mathbf{Y}} \mathcal{D}_{A,\mathbf{g}}(\mathbf{t}, \mathbf{Y}),$$

yielding

$$\text{variance} = \mathbb{E}_{\mathbf{Y}} \mathcal{D}_{B,\mathbf{f}}(\mathbf{Y}, \mathbf{y}^*) = \mathbb{E}_{\mathbf{Y}} B(\mathbf{f}(\mathbf{Y})) - B(\mathbf{f}(\mathbf{y}^*)).$$

**Proof** We follow the same line of reasoning as in Proposition 1 of Banerjee et al. (2005b) for a standard Bregman divergence and as in Theorem 3 of Brekelmans and Nielsen (2024) for the rho-tau divergence. Let  $\{B, \mathbf{f}\}$  be the dual pair of  $\{A, \mathbf{g}\}$ . For any  $\mathbf{t} \in \mathcal{Y}$ , we consider the difference

$$\begin{aligned} \mathbb{E}_{\mathbf{T}} \mathcal{D}_{A, \mathbf{g}}(\mathbf{T}, \mathbf{t}) - \mathbb{E}_{\mathbf{T}} \mathcal{D}_{A, \mathbf{g}}(\mathbf{T}, \bar{\mathbf{t}}) &= B(\mathbf{f}(\mathbf{t})) - B(\mathbf{f}(\bar{\mathbf{t}})) + (\mathbf{f}(\bar{\mathbf{t}}) - \mathbf{f}(\mathbf{t}))^\top \mathbb{E}_{\mathbf{T}} \mathbf{g}(\mathbf{T}) \\ &= B(\mathbf{f}(\mathbf{t})) - B(\mathbf{f}(\bar{\mathbf{t}})) + (\mathbf{f}(\bar{\mathbf{t}}) - \mathbf{f}(\mathbf{t}))^\top \mathbf{g}(\bar{\mathbf{t}}) \\ &= \mathcal{D}_{B, \mathbf{f}}(\bar{\mathbf{t}}, \mathbf{t}) \geq 0. \end{aligned}$$

That is, when the  $\mathbf{g}$ -mean  $\bar{\mathbf{t}}$  indeed exists and satisfies any constraints  $\mathcal{Y}$  may impose, there can be no other  $\mathbf{t} \in \mathcal{Y}$  that leads to a lower intrinsic noise. Plugging  $\mathbf{t}^* = \bar{\mathbf{t}}$  into the definition of the intrinsic noise, we get

$$\begin{aligned} \text{intrinsic noise} &= \mathbb{E}_{\mathbf{T}} \mathcal{D}_{A, \mathbf{g}}(\mathbf{T}, \bar{\mathbf{t}}) \\ &= (\mathbf{g}(\bar{\mathbf{t}}) - \mathbb{E}_{\mathbf{T}} \mathbf{g}(\mathbf{T}))^\top \nabla_{\mathbf{g}(\bar{\mathbf{t}})} A(\mathbf{g}(\bar{\mathbf{t}})) + \mathbb{E}_{\mathbf{T}} A(\mathbf{g}(\mathbf{T})) - A(\mathbf{g}(\bar{\mathbf{t}})) \\ &= \mathbb{E}_{\mathbf{T}} A(\mathbf{g}(\mathbf{T})) - A(\mathbf{g}(\bar{\mathbf{t}})). \end{aligned}$$

The proof for the  $\mathbf{f}$ -mean prediction  $\bar{\mathbf{y}}$  follows the exact same lines, interchanging labels for predictions and  $\{A, \mathbf{g}\}$  for  $\{B, \mathbf{f}\}$ .  $\blacksquare$

**Lemma 7.** Consider a  $g$ -Bregman divergence  $\mathcal{D}_{A, g} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  with generating function  $A : \mathbf{g}(\mathcal{Y}) \rightarrow \mathbb{R}$  and invertible mapping  $\mathbf{g} : \mathcal{Y} \rightarrow \mathbf{g}(\mathcal{Y})$ . Assume that the central label  $\mathbf{t}^*$  minimizing the intrinsic noise term is equal to the  $\mathbf{g}$ -mean label  $\mathbf{g}^{-1}(\mathbb{E}_{\mathbf{T}} \mathbf{g}(\mathbf{T}))$ , and that the central prediction  $\mathbf{y}^*$  minimizing the variance term is equal to the  $\mathbf{f}$ -mean prediction  $\mathbf{f}^{-1}(\mathbb{E}_{\mathbf{Y}} \mathbf{f}(\mathbf{Y}))$ , where  $\mathbf{f}(\mathbf{y}) = \nabla_{\mathbf{g}(\mathbf{y})} A(\mathbf{g}(\mathbf{y}))$ . Then  $\mathcal{D}_{A, \mathbf{g}}$  has a clean bias-variance decomposition:

$$\underbrace{\mathbb{E}_{\mathbf{T}, \mathbf{Y}} \mathcal{D}_{A, \mathbf{g}}(\mathbf{T}, \mathbf{Y})}_{\text{expected loss}} = \underbrace{\mathbb{E}_{\mathbf{T}} \mathcal{D}_{A, \mathbf{g}}(\mathbf{T}, \mathbf{t}^*)}_{\text{intrinsic noise}} + \underbrace{\mathcal{D}_{A, \mathbf{g}}(\mathbf{t}^*, \mathbf{y}^*)}_{\text{bias}} + \underbrace{\mathbb{E}_{\mathbf{Y}} \mathcal{D}_{A, \mathbf{g}}(\mathbf{y}^*, \mathbf{Y})}_{\text{variance}}.$$

**Proof** The proof follows by making the change of variables from  $\mathbf{t}$  and  $\mathbf{y}$  to  $\mathbf{g}(\mathbf{t})$  and  $\mathbf{g}(\mathbf{y})$  in the bias-variance decomposition for (standard) Bregman divergences provided by Pfau (2025) and Gupta et al. (2022).

For completeness and illustration, we here provide an alternative proof that makes use of the symmetry following from Proposition 5. With  $\{B, \mathbf{f}\}$  the dual pair of  $\{A, \mathbf{g}\}$ , taking expectations in the symmetric form (14) yields

$$\begin{aligned} \mathbb{E}_{\mathbf{T}, \mathbf{Y}} \mathcal{D}_{A, \mathbf{g}}(\mathbf{T}, \mathbf{Y}) &= \mathbb{E}_{\mathbf{T}} A(\mathbf{g}(\mathbf{T})) - \mathbb{E}_{\mathbf{Y}} \mathbf{f}(\mathbf{Y})^\top \mathbb{E}_{\mathbf{T}} \mathbf{g}(\mathbf{T}) + \mathbb{E}_{\mathbf{Y}} B(\mathbf{f}(\mathbf{Y})) \\ &= \underbrace{\mathbb{E}_{\mathbf{T}} A(\mathbf{g}(\mathbf{T})) - A(\mathbf{g}(\mathbf{t}^*))}_{\text{intrinsic noise}} + \underbrace{A(\mathbf{g}(\mathbf{t}^*)) - \mathbf{f}(\mathbf{y}^*)^\top \mathbf{g}(\mathbf{t}^*) + B(\mathbf{f}(\mathbf{y}^*))}_{\text{bias}} + \underbrace{\mathbb{E}_{\mathbf{Y}} B(\mathbf{f}(\mathbf{Y})) - B(\mathbf{f}(\mathbf{y}^*))}_{\text{variance}}. \end{aligned}$$

$\blacksquare$

**Theorem 8.** Consider a  $g$ -Bregman divergence  $\mathcal{D}_{A, g} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  with generating function  $A : \mathbf{g}(\mathcal{Y}) \rightarrow \mathbb{R}$  and invertible mapping  $\mathbf{g} : \mathcal{Y} \rightarrow \mathbf{g}(\mathcal{Y})$ , and assume that  $\mathbf{g}(\mathcal{Y}) = \text{dom}(A)$  is a convex feasible set defined by inequality constraints only. Then this  $g$ -Bregman divergence has a clean bias-variance decomposition of the form (17).

**Proof** Since  $\mathbf{g}$  is invertible and  $\mathbf{g}(\mathcal{Y})$  is assumed to be convex with only inequality constraints (if any), we can pull the expectation and the optimization through the mapping and do not need to worry about any equality constraints:

$$\mathbf{t}^* = \arg \min_{\mathbf{y} \in \mathcal{Y}} \mathbb{E}_{\mathbf{T}} \mathcal{D}_{A, \mathbf{g}}(\mathbf{T}, \mathbf{y}) = \mathbf{g}^{-1} \left( \arg \min_{\mathbf{g}(\mathbf{y}) \in \mathbf{g}(\mathcal{Y})} \mathbb{E}_{\mathbf{g}(\mathbf{T})} \mathcal{D}_A(\mathbf{g}(\mathbf{T}), \mathbf{g}(\mathbf{y})) \right) = \mathbf{g}^{-1}(\mathbb{E}_{\mathbf{T}} \mathbf{g}(\mathbf{T})) = \bar{\mathbf{t}}.$$

With  $\{B, \mathbf{f}\}$  the dual pair of  $\{A, \mathbf{g}\}$ ,  $B$  is the convex conjugate of  $A$  and its domain  $\mathbf{f}(\mathcal{Y})$  is a convex set. The exact same reasoning then applies to the central prediction  $\mathbf{y}^*$ , by interchanging labels for predictions and  $\{A, \mathbf{g}\}$  for  $\{B, \mathbf{f}\}$ :

$$\mathbf{y}^* = \arg \min_{\mathbf{t} \in \mathcal{Y}} \mathbb{E}_{\mathbf{Y}} \mathcal{D}_{A, \mathbf{g}}(\mathbf{t}, \mathbf{Y}) = \mathbf{f}^{-1}(\mathbb{E}_{\mathbf{Y}} \mathbf{f}(\mathbf{Y})) = \bar{\mathbf{y}}.$$

Since the conditions of Lemma 7 are now satisfied, this  $g$ -Bregman divergence has a clean bias-variance decomposition.  $\blacksquare$

**Theorem 9.** Consider a Bregman divergence  $\mathcal{D}_A : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  with generating function  $A : \mathcal{Y} \rightarrow \mathbb{R}$ , and dual pair  $\{B, \mathbf{f}\}$ . The feasible set  $\mathcal{Y}$  contains equality constraints that can be written in the form  $W\mathbf{y} = \mathbf{b}$ , with  $W \in \mathbb{R}^{\kappa \times d}$ ,  $\mathbf{b} \in \mathbb{R}^{\kappa}$ , and  $\kappa$  the number of constraints. Then this Bregman divergence has a clean bias-variance decomposition of the form

$$\underbrace{\mathbb{E}_{\mathbf{T}, \mathbf{Y}} \mathcal{D}_A(\mathbf{T}, \mathbf{Y})}_{\text{expected loss}} = \underbrace{\mathbb{E}_{\mathbf{T}} A(\mathbf{T}) - A(\mathbf{t}^*)}_{\text{intrinsic noise}} + \underbrace{\mathcal{D}_A(\mathbf{t}^*, \mathbf{y}^*)}_{\text{bias}} + \underbrace{\boldsymbol{\lambda}^\top \mathbf{b} + \mathbb{E}_{\mathbf{Y}} B(\mathbf{f}(\mathbf{Y})) - B(\mathbf{f}(\mathbf{y}^*))}_{\text{variance}},$$

where

$$\mathbf{t}^* = \mathbb{E}_{\mathbf{T}} \mathbf{T} \quad \text{and} \quad \mathbf{y}^* = \mathbf{f}^{-1} \left( \mathbb{E}_{\mathbf{Y}} \mathbf{f}(\mathbf{Y}) + \boldsymbol{\lambda}^\top W \right) \quad \text{with } \boldsymbol{\lambda} \text{ such that } W\mathbf{y}^* = \mathbf{b}.$$

**Proof** The central label is defined as

$$\mathbf{t}^* = \arg \min_{\mathbf{y} \in \mathcal{Y}} \mathbb{E}_{\mathbf{T}} \mathcal{D}_A(\mathbf{T}, \mathbf{y}) = \arg \min_{\mathbf{y} \in \mathcal{Y}} \left[ B(\mathbf{f}(\mathbf{y})) - \mathbf{f}(\mathbf{y})^\top \mathbb{E}_{\mathbf{T}} \mathbf{T} \right].$$

Adding Lagrange multipliers  $\boldsymbol{\lambda} \in \mathbb{R}^{\kappa}$  to incorporate the linear constraints  $W\mathbf{y} - \mathbf{b} = 0$ , and setting the derivative with respect to  $\mathbf{f}(\mathbf{y})$  to zero, while making use of  $\nabla_{\mathbf{f}(\mathbf{y})} B(\mathbf{f}(\mathbf{y})) = \mathbf{y}$ , we obtain

$$\mathbf{t}^* = \mathbb{E}_{\mathbf{Y}} \mathbf{Y} + \boldsymbol{\lambda}^\top W.$$

With all labels  $\mathbf{Y} \in \mathcal{Y}$ , so does  $\mathbf{t}^*$  for  $\boldsymbol{\lambda} = \mathbf{0}$ , yielding the usual expression

$$\text{intrinsic noise} = \mathbb{E}_{\mathbf{T}} \mathcal{D}_A(\mathbf{T}, \mathbf{t}^*) = \mathbb{E}_{\mathbf{T}} A(\mathbf{T}) - A(\mathbf{t}^*).$$

When solving for the central prediction, the constraints do matter. We have

$$\mathbf{y}^* = \arg \min_{\mathbf{t} \in \mathcal{Y}} \mathbb{E}_{\mathbf{Y}} \ell(\mathbf{t}, \mathbf{Y}) = \arg \min_{\mathbf{t} \in \mathcal{Y}} \left[ A(\mathbf{t}) - \mathbf{t}^\top \mathbb{E}_{\mathbf{Y}} \mathbf{f}(\mathbf{Y}) \right].$$

Adding Lagrange multipliers  $\boldsymbol{\lambda}$  and setting the derivative with respect to  $\mathbf{t}$  to zero, while making use of  $\nabla_{\mathbf{t}}A(\mathbf{t}) = \mathbf{f}(\mathbf{t})$ , we obtain

$$\mathbf{f}(\mathbf{y}^*) = \mathbb{E}_{\mathbf{Y}}\mathbf{f}(\mathbf{Y}) + \boldsymbol{\lambda}^\top W ,$$

where  $\boldsymbol{\lambda}$  must be chosen such that indeed  $W\mathbf{y}^* = \mathbf{b}$ . Plugging this solution in the expression for the variance, we arrive at

$$\begin{aligned} \text{variance} &= \mathbb{E}_{\mathbf{Y}}\mathcal{D}_{B,\mathbf{f}}(\mathbf{Y},\mathbf{y}^*) = \boldsymbol{\lambda}^\top W\mathbf{y}^* + \mathbb{E}_{\mathbf{Y}}B(\mathbf{f}(\mathbf{Y})) - B(\mathbf{f}(\mathbf{y}^*)) \\ &= \boldsymbol{\lambda}^\top \mathbf{b} + \mathbb{E}_{\mathbf{Y}}B(\mathbf{f}(\mathbf{Y})) - B(\mathbf{f}(\mathbf{y}^*)) . \end{aligned}$$

■

**Lemma 11.** *Consider a loss function  $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  that has a clean bias-variance decomposition. Then its mixed second derivative w.r.t. prediction  $\mathbf{y}$  and label  $\mathbf{t}$ ,  $\mathcal{L}_{\mathbf{y},\mathbf{t}}(\mathbf{t},\mathbf{y}) = \nabla_{\mathbf{y}}\nabla_{\mathbf{t}}\ell(\mathbf{t},\mathbf{y})$ , necessarily factorizes as*

$$\mathcal{L}_{\mathbf{y},\mathbf{t}}(\mathbf{t},\mathbf{y}) = H_2(\mathbf{y})H_1^\top(\mathbf{t}) ,$$

for some matrix functions  $H_1 : \mathcal{Y} \rightarrow \mathbb{R}^{d \times d}$  and  $H_2 : \mathcal{Y} \rightarrow \mathbb{R}^{d \times d}$  for any  $(\mathbf{t},\mathbf{y}) \in \mathcal{Y} \times \mathcal{Y}$ .

**Proof** At first, we do not worry about potential irregularities and assume that the loss function  $\ell$  is differentiable with respect to the label  $\mathbf{t}$  and the prediction  $\mathbf{y}$  across the whole domain  $\mathcal{Y} \times \mathcal{Y}$ . We will justify below that this regularity is in fact enforced by the requirement that the clean bias-variance decomposition holds everywhere.

Let us consider a fixed label  $\mathbf{t}$  and a distribution with just two labels,  $\mathbf{y}_1$  and  $\mathbf{y}_2$ , that have probability  $p$  and  $1 - p$ , respectively. Our line of reasoning is sketched in Figure 1 for  $p = 1/2$ .

The central prediction  $\mathbf{y}^*$  reads

$$\mathbf{y}^* = \arg \min_{\mathbf{t} \in \mathcal{Y}} p\ell(\mathbf{t},\mathbf{y}_1) + (1 - p)\ell(\mathbf{t},\mathbf{y}_2) ,$$

which, since  $\ell$  is differentiable with respect to  $\mathbf{t}$ , implies

$$p\boldsymbol{\ell}_t(\mathbf{y}^*,\mathbf{y}_1) + (1 - p)\boldsymbol{\ell}_t(\mathbf{y}^*,\mathbf{y}_2) = \mathbf{0} ,$$

with convention  $[\boldsymbol{\ell}_t(\mathbf{t},\mathbf{y})]_i = \nabla_{t_i}\ell(\mathbf{t},\mathbf{y})$ . Now we slightly change to  $\mathbf{y}'_1 = \mathbf{y}_1 + \boldsymbol{\epsilon}_1$  and  $\mathbf{y}'_2 = \mathbf{y}_2 + \boldsymbol{\epsilon}_2$  with  $\boldsymbol{\epsilon}_1$  and  $\boldsymbol{\epsilon}_2$  constructed such that  $\mathbf{y}^*$  stays the same. That is, for infinitesimal  $\boldsymbol{\epsilon}_1$  and  $\boldsymbol{\epsilon}_2$ , we enforce the constraint

$$\begin{aligned} \mathbf{0} &= p\boldsymbol{\ell}_t(\mathbf{y}^*,\mathbf{y}_1 + \boldsymbol{\epsilon}_1) + (1 - p)\boldsymbol{\ell}_t(\mathbf{y}^*,\mathbf{y}_2 + \boldsymbol{\epsilon}_2) \\ &= p\boldsymbol{\ell}_t(\mathbf{y}^*,\mathbf{y}_1) + p\mathcal{L}_{\mathbf{y},\mathbf{t}}(\mathbf{y}^*,\mathbf{y}_1)^\top \boldsymbol{\epsilon}_1 + (1 - p)\boldsymbol{\ell}_t(\mathbf{y}^*,\mathbf{y}_2) + (1 - p)\mathcal{L}_{\mathbf{y},\mathbf{t}}(\mathbf{y}^*,\mathbf{y}_2)^\top \boldsymbol{\epsilon}_2 \\ &= p\mathcal{L}_{\mathbf{y},\mathbf{t}}(\mathbf{y}^*,\mathbf{y}_1)^\top \boldsymbol{\epsilon}_1 + (1 - p)\mathcal{L}_{\mathbf{y},\mathbf{t}}(\mathbf{y}^*,\mathbf{y}_2)^\top \boldsymbol{\epsilon}_2 , \end{aligned} \tag{23}$$

with  $[\mathcal{L}_{\mathbf{y},\mathbf{t}}(\mathbf{t},\mathbf{y})]_{j,i} = \nabla_{y_j}\nabla_{t_i}\ell(\mathbf{t},\mathbf{y})$ .

If we indeed have a clean bias-variance decomposition, the bias is the expected loss minus the variance, which must be the same before and after the transformation:

$$\begin{aligned} p\ell(\mathbf{t}, \mathbf{y}_1 + \boldsymbol{\epsilon}_1) + (1-p)\ell(\mathbf{t}, \mathbf{y}_2 + \boldsymbol{\epsilon}_2) - p\ell(\mathbf{y}^*, \mathbf{y}_1 + \boldsymbol{\epsilon}_1) - (1-p)\ell(\mathbf{y}^*, \mathbf{y}_2 + \boldsymbol{\epsilon}_2) = \\ = p\ell(\mathbf{t}, \mathbf{y}_1) + (1-p)\ell(\mathbf{t}, \mathbf{y}_2) - p\ell(\mathbf{y}^*, \mathbf{y}_1) - (1-p)\ell(\mathbf{y}^*, \mathbf{y}_2). \end{aligned}$$

For infinitesimal  $\boldsymbol{\epsilon}_1$  and  $\boldsymbol{\epsilon}_2$ , this boils down to

$$p(\boldsymbol{\ell}_y(\mathbf{t}, \mathbf{y}_1) - \boldsymbol{\ell}_y(\mathbf{y}^*, \mathbf{y}_1))^\top \boldsymbol{\epsilon}_1 + (1-p)(\boldsymbol{\ell}_y(\mathbf{t}, \mathbf{y}_2) - \boldsymbol{\ell}_y(\mathbf{y}^*, \mathbf{y}_2))^\top \boldsymbol{\epsilon}_2 = 0, \quad (24)$$

with  $[\boldsymbol{\ell}_y(\mathbf{t}, \mathbf{y})]_j = \nabla_{y_j} \ell(\mathbf{t}, \mathbf{y})$ .

Combining (23) and (24), we obtain the two equations

$$\begin{aligned} \mathbf{0} &= p\mathcal{L}_{y,t}(\mathbf{y}^*, \mathbf{y}_1)^\top \boldsymbol{\epsilon}_1 + (1-p)\mathcal{L}_{y,t}(\mathbf{y}^*, \mathbf{y}_2)^\top \boldsymbol{\epsilon}_2 \\ 0 &= p[\boldsymbol{\ell}_y(\mathbf{t}, \mathbf{y}_1) - \boldsymbol{\ell}_y(\mathbf{y}^*, \mathbf{y}_1)]^\top \boldsymbol{\epsilon}_1 + (1-p)[\boldsymbol{\ell}_y(\mathbf{t}, \mathbf{y}_2) - \boldsymbol{\ell}_y(\mathbf{y}^*, \mathbf{y}_2)]^\top \boldsymbol{\epsilon}_2. \end{aligned}$$

The first equation tells us how to choose  $\boldsymbol{\epsilon}_2$  if we make an infinitesimal change  $\boldsymbol{\epsilon}_1$  and want  $\mathbf{y}^*$  to stay the same. The second equation must be valid for a loss function with a clean bias-variance decomposition, for any combination of  $\boldsymbol{\epsilon}_1$  and  $\boldsymbol{\epsilon}_2$  that satisfies the first one. When  $\mathbf{t}$  is close to  $\mathbf{y}^*$ , that is, up to first order in the difference  $\mathbf{t} - \mathbf{y}^*$ , this second equation directly follows from the first. However, for loss functions with a clean bias-variance decomposition, this should not just hold for small differences between the label and the central prediction, but for any difference between the two.

For the second equation to be valid, while still satisfying the first one, we must be able to write the second equation as a linear combination of the first one. That is, there must exist some vector function  $\mathbf{q} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}^d$ , which may still depend on  $\mathbf{t}$  and  $\mathbf{y}^*$ , but cannot otherwise depend on  $\mathbf{y}_1$  and  $\mathbf{y}_2$ , such that

$$\mathcal{L}_{y,t}(\mathbf{y}^*, \mathbf{y}_1)\mathbf{q}(\mathbf{t}, \mathbf{y}^*) = \boldsymbol{\ell}_y(\mathbf{t}, \mathbf{y}_1) - \boldsymbol{\ell}_y(\mathbf{y}^*, \mathbf{y}_1),$$

and the same with  $\mathbf{y}_1$  replaced with  $\mathbf{y}_2$ . Since this must hold for arbitrary  $\mathbf{y}_1, \mathbf{y}_2$ , and  $p$ , we may as well rewrite

$$[\boldsymbol{\ell}_y(\mathbf{t}', \mathbf{y}) - \boldsymbol{\ell}_y(\mathbf{t}, \mathbf{y})]_j = \sum_k [\mathcal{L}_{y,t}(\mathbf{t}, \mathbf{y})]_{jk} [\mathbf{q}(\mathbf{t}', \mathbf{t})]_k, \quad (25)$$

which must be valid for some function  $\mathbf{q}$  and arbitrary  $\mathbf{t}', \mathbf{t}$ , and  $\mathbf{y}$ .

The left-hand side gives the difference in gradients  $\boldsymbol{\ell}_y(\mathbf{t}', \mathbf{y}) - \boldsymbol{\ell}_y(\mathbf{t}, \mathbf{y}) \in \mathbb{R}^d$ . The right-hand side expresses this vector as a linear combination of the columns of  $\mathcal{L}_{y,t}(\mathbf{t}, \mathbf{y})$ , with coefficients  $\mathbf{q}(\mathbf{t}', \mathbf{t})$  that depend only on the labels. For fixed  $\mathbf{y}$  and  $\mathbf{t}$ , the right-hand side therefore ranges precisely over the column space of  $\mathcal{L}_{y,t}(\mathbf{t}, \mathbf{y})$ . Since (25) must hold for arbitrary pairs of labels  $\mathbf{t}'$  and  $\mathbf{t}$  at fixed  $\mathbf{y}$ , every gradient difference  $\boldsymbol{\ell}_y(\mathbf{t}', \mathbf{y}) - \boldsymbol{\ell}_y(\mathbf{t}, \mathbf{y})$  must lie in this column space. Because this requirement holds for all choices of the reference label  $\mathbf{t}$ , the column space of  $\mathcal{L}_{y,t}(\mathbf{t}, \mathbf{y})$  cannot depend on  $\mathbf{t}$ .

It follows that, for each fixed  $\mathbf{y}$ , there exists a matrix  $H_2(\mathbf{y})$  whose columns form a basis of this common column space, such that each column of  $\mathcal{L}_{y,t}(\mathbf{t}, \mathbf{y})$  lies in the span of  $H_2(\mathbf{y})$ . Collecting the corresponding coefficients into a matrix  $H_1(\mathbf{t})$  yields the factorization

$$\mathcal{L}_{y,t}(\mathbf{t}, \mathbf{y}) = H_2(\mathbf{y}) H_1(\mathbf{t})^\top,$$

where  $H_2(\mathbf{y})$  encodes a fixed set of directions that span all possible gradient differences at prediction  $\mathbf{y}$ , and  $H_1(\mathbf{t})$  determines how these directions are combined for a given label  $\mathbf{t}$ .

Since  $\mathcal{L}_{y,t}(\mathbf{t}, \mathbf{y}) \in \mathbb{R}^{d \times d}$ , its rank is at most  $d$ , so the factorization involves at most  $d$  independent directions. Without loss of generality, we may therefore take  $H_1, H_2 \in \mathbb{R}^{d \times d}$  in (19). Choosing a larger intermediate dimension would be redundant, while choosing a smaller one would require  $\mathcal{L}_{y,t}(\mathbf{t}, \mathbf{y})$  to be rank-deficient everywhere.

Finally, the differentiability assumptions made at the beginning of the proof are not just technical but are enforced by the clean bias-variance decomposition itself. The decomposition is an exact identity that must hold for all choices of labels and predictions, and therefore also under arbitrarily small perturbations of  $\mathbf{y}_1$  and  $\mathbf{y}_2$  that leave the central prediction  $\mathbf{y}^*$  unchanged. Without local differentiability with respect to  $\mathbf{t}$  and  $\mathbf{y}$ , the change of the loss under infinitesimal perturbations cannot, in general, be expressed linearly as in (23) and (24). Hence, for a clean bias-variance decomposition to hold locally, the loss must have well-defined first-order derivatives with respect to  $\mathbf{t}$  and  $\mathbf{y}$ , so that the bias remains invariant under all perturbations preserving  $\mathbf{y}^*$ .  $\blacksquare$

**Theorem 12.** *Consider a continuous, nonnegative loss function  $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  that satisfies the identity of indiscernibles and has a clean bias-variance decomposition. Then  $\ell$  must be a  $g$ -Bregman divergence  $\mathcal{D}_{A,g}$  of the form (11) for some invertible mapping  $\mathbf{g} : \mathcal{Y} \rightarrow \mathbf{g}(\mathcal{Y})$  and strictly convex generating function  $A : \mathbf{g}(\mathcal{Y}) \rightarrow \mathbb{R}$ .*

**Proof** From the factorization (19) in Lemma 11, we know that the second derivative of the loss function  $\ell$  with respect to  $\mathbf{t}$  and  $\mathbf{y}$ ,  $\mathcal{L}_{t,y}(\mathbf{t}, \mathbf{y})$ , factorizes as

$$[\mathcal{L}_{y,t}(\mathbf{t}, \mathbf{y})]_{ij} = \sum_k [H_2(\mathbf{y})]_{ik} [H_1(\mathbf{t})]_{jk} = [H_2(\mathbf{y})H_1(\mathbf{t})^\top]_{ij},$$

for some  $d \times d$  matrices  $H_1(\mathbf{t})$  and  $H_2(\mathbf{y})$ . To arrive at the loss function itself, we perform two integrations, first with respect to  $\mathbf{t}$  and then with respect to  $\mathbf{y}$ .

Integration with respect to  $\mathbf{t}$  yields

$$[\ell_y(\mathbf{t}, \mathbf{y})]_i = \sum_k [H_2(\mathbf{y})]_{ik} [\mathbf{h}_1(\mathbf{t})]_k + [\mathbf{h}_3(\mathbf{y})]_i = [H_2(\mathbf{y})\mathbf{h}_1(\mathbf{t}) + \mathbf{h}_3(\mathbf{y})]_i, \quad (26)$$

for some vector functions  $\mathbf{h}_1 : \mathcal{Y} \rightarrow \mathbb{R}^d$  and  $\mathbf{h}_3 : \mathcal{Y} \rightarrow \mathbb{R}^d$ , with

$$\nabla_{t_j} [\mathbf{h}_1(\mathbf{t})]_k = [H_1(\mathbf{t})]_{jk}.$$

Since the minimum of  $\ell(\mathbf{t}, \mathbf{y})$  w.r.t.  $\mathbf{y}$  is obtained for  $\mathbf{y} = \mathbf{t}$ , we need to have  $\ell_y(\mathbf{t}, \mathbf{t}) = \mathbf{0}$  for all  $\mathbf{t}$ , and thus

$$\mathbf{h}_3(\mathbf{y}) = -H_2(\mathbf{y})\mathbf{h}_1(\mathbf{y}).$$

Substitution into (26) gives

$$\nabla_{y_i} \ell(\mathbf{t}, \mathbf{y}) = [\ell_y(\mathbf{t}, \mathbf{y})]_i = \sum_k [H_2(\mathbf{y})]_{ik} ([\mathbf{h}_1(\mathbf{t})]_k - [\mathbf{h}_1(\mathbf{y})]_k) = [H_2(\mathbf{y}) (\mathbf{h}_1(\mathbf{t}) - \mathbf{h}_1(\mathbf{y}))]_i.$$

Integrating out  $\mathbf{y}$  by parts, directly incorporating that  $\ell(\mathbf{y}, \mathbf{y}) = 0$  for all  $\mathbf{y}$ , yields

$$\begin{aligned} \ell(\mathbf{t}, \mathbf{y}) &= \sum_k [\mathbf{h}_2(\mathbf{y})]_k ([\mathbf{h}_1(\mathbf{t})]_k - [\mathbf{h}_1(\mathbf{y})]_k) + H_4(\mathbf{y}) - H_4(\mathbf{t}) \\ &= \mathbf{h}_2(\mathbf{y})^\top (\mathbf{h}_1(\mathbf{t}) - \mathbf{h}_1(\mathbf{y})) + H_4(\mathbf{y}) - H_4(\mathbf{t}), \end{aligned}$$

with vector function  $\mathbf{h}_2 : \mathcal{Y} \rightarrow \mathbb{R}^d$  satisfying

$$\nabla_{y_i} [\mathbf{h}_2(\mathbf{y})]_k = [H_2(\mathbf{y})]_{ik},$$

and scalar function  $H_4 : \mathcal{Y} \rightarrow \mathbb{R}$  then following from our choice of  $h_1$  and  $h_2$  through

$$\nabla_{y_i} H_4(\mathbf{y}) = \sum_k [H_1(\mathbf{y})]_{ik} [\mathbf{h}_2(\mathbf{y})]_k. \quad (27)$$

To arrive at the form (11), we choose  $\mathbf{h}_1(\mathbf{y}) = \mathbf{g}(\mathbf{y})$ ,  $H_4(\mathbf{y}) = A(\mathbf{g}(\mathbf{y}))$  and  $h_2(\mathbf{y}) = -\nabla_{\mathbf{g}(\mathbf{y})} A(\mathbf{g}(\mathbf{y}))$  for some scalar function  $A : \mathbf{g}(\mathcal{Y}) \rightarrow \mathbb{R}$ , and see that these choices indeed satisfy (27).

Fixing  $\mathbf{t}$ , the solution for  $\mathbf{y}$  such that  $\ell(\mathbf{t}, \mathbf{y}) = 0$  follows from  $\mathbf{g}(\mathbf{y}) = \mathbf{g}(\mathbf{t})$ . Our requirement that the solution  $\mathbf{y} = \mathbf{t}$  is unique for any choice of  $\mathbf{t}$  implies that the mapping  $\mathbf{g}$  is invertible. The strict convexity of  $A$  then follows from the requirement that  $\ell(\mathbf{t}, \mathbf{y}) \geq 0$ .

We only enforced the requirement that the loss function  $\ell$  has a clean bias-variance decomposition for the special case of a single fixed label  $\mathbf{t}$ . This requirement already suffices to nail down  $\ell$  to the class of  $g$ -Bregman divergences. The clean bias-variance decomposition for a distribution of labels  $\mathbf{T}$  we then get ‘for free’, since it applies to any  $g$ -Bregman divergence: it does not imply any further restrictions. ■

## References

- Shun-ichi Amari. *Differential-Geometrical Methods in Statistics*, volume 28 of *Lecture Notes in Statistics*. Springer-Verlag, 1985.
- Ehsan Amid, Manfred Warmuth, Rohan Anil, and Tomer Koren. Robust bi-tempered logistic loss based on Bregman divergences. *Advances in Neural Information Processing Systems*, 32, 2019.
- Arindam Banerjee, Xin Guo, and Hui Wang. On the optimality of conditional expectation as a Bregman predictor. *IEEE Transactions on Information Theory*, 51(7):2664–2669, 2005a.
- Arindam Banerjee, Srujana Merugu, Inderjit Dhillon, Joydeep Ghosh, and John Lafferty. Clustering with Bregman divergences. *Journal of Machine Learning Research*, 6(10), 2005b.
- Heinz Bauschke and Jonathan Borwein. Joint and separate convexity of the Bregman distance. In *Inherently Parallel Algorithms in Feasibility and Optimization and their Applications*, volume 8 of *Studies in Computational Mathematics*, pages 23–36. Elsevier, 2001.

- Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.
- Mathieu Blondel, André Martins, and Vlad Niculae. Learning with Fenchel-Young losses. *Journal of Machine Learning Research*, 21:1–69, 2020.
- Jean-Daniel Boissonnat, Frank Nielsen, and Richard Nock. Bregman Voronoi diagrams. *Discrete & Computational Geometry*, 44:281–307, 2010.
- Lev Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 7(3):200–217, 1967.
- Rob Brekelmans and Frank Nielsen. Variational representations of annealing paths: Bregman information under monotonic embedding. *Information Geometry*, 7:193–228, 2024.
- Gavin Brown and Riccardo Ali. Bias/variance is not the same as approximation/estimation. *Transactions on Machine Learning Research*, 2024.
- Sayak Ray Chowdhury, Patrick Saux, Odalric Maillard, and Aditya Gopalan. Bregman deviations of generic exponential families. In *Conference on Learning Theory*, pages 394–449, 2023.
- Hatice Kubra Cilingir, Rachel Manzelli, and Brian Kulis. Deep divergence learning. In *International Conference on Machine Learning*, pages 2027–2037, 2020.
- Pedro Domingos. A unified bias-variance decomposition. In *International Conference on Machine Learning*, pages 231–238, 2000.
- Béla Frigyik, Santosh Srivastava, and Maya Gupta. Functional Bregman divergence and Bayesian estimation of distributions. *IEEE Transactions on Information Theory*, 54(11):5130–5139, 2008.
- Stuart Geman, Elie Bienenstock, and René Doursat. Neural networks and the bias/variance dilemma. *Neural Computation*, 4(1):1–58, 1992.
- Sebastian Gruber and Florian Buettner. Uncertainty estimates of predictions via a general bias-variance decomposition. In *International Conference on Artificial Intelligence and Statistics*, pages 11331–11354, 2023.
- Neha Gupta, Jamie Smith, Ben Adlam, and Zelda Mariet. Ensembles of classifiers: a bias-variance perspective. *Transactions on Machine Learning Research*, 2022.
- Jakob Vogdrup Hansen and Tom Heskes. General bias/variance decomposition with target independent variance of error functions derived from the exponential family of distributions. In *International Conference on Pattern Recognition (ICPR-2000)*, volume 2, pages 207–210, 2000.
- Tom Heskes. Bias/variance decompositions for likelihood-based estimators. *Neural Computation*, 10(6):1425–1433, 1998.

- Gareth James. Variance and bias for general loss functions. *Machine Learning*, 51:115–135, 2003.
- Jan Naudts and Jun Zhang. Rho-tau embedding of statistical models. In *Geometric Structures of Information*, pages 1–13. Springer, 2018.
- Brady Neal, Sarthak Mittal, Aristide Baratin, Vinayak Tantia, Matthew Scicluna, Simon Lacoste-Julien, and Ioannis Mitliagkas. A modern take on the bias-variance tradeoff in neural networks. *\*arXiv\**, 2018. arXiv:1810.08591.
- Frank Nielsen. On geodesic triangles with right angles in a dually flat space. In Frank Nielsen, editor, *Progress in Information Geometry: Theory and Applications*, pages 153–190. Springer International Publishing, Cham, 2021.
- Frank Nielsen and Richard Nock. Sided and symmetrized Bregman centroids. *IEEE Transactions on Information Theory*, 55(6):2882–2904, 2009.
- Tomohiro Nishiyama. Generalized Bregman and Jensen divergences which include some f-divergences. *\*arXiv\**, 2018. arXiv:1808.06148.
- David Nix and Andreas Weigend. Estimating the mean and variance of the target probability distribution. In *International Conference on Neural Networks (ICNN'94)*, volume 1, pages 55–60, 1994.
- David Pfau. A generalized bias-variance decomposition for Bregman divergences. *\*arXiv\**, 2025. arXiv:2511.08789.
- Ali Siahkamari, Xide Xia, Venkatesh Saligrama, David Castañón, and Brian Kulis. Learning to approximate a Bregman divergence. *Advances in Neural Information Processing Systems*, 33:3603–3612, 2020.
- Laurens Sluijterman, Eric Cator, and Tom Heskes. Optimal training of mean variance estimation neural networks. *Neurocomputing*, 597:127929, 2024.
- Baba Vemuri, Meizhu Liu, Shun-Ichi Amari, and Frank Nielsen. Total Bregman divergence and its applications to DTI analysis. *IEEE Transactions on Medical Imaging*, 30(2):475–483, 2011.
- David Wolpert. On bias plus variance. *Neural Computation*, 9(6):1211–1243, 1997.
- Danny Wood, Tingting Mu, and Gavin Brown. Bias-variance decompositions for margin losses. In *International Conference on Artificial Intelligence and Statistics*, pages 1975–2001. PMLR, 2022.
- Danny Wood, Tingting Mu, Andrew Webb, Henry Reeve, Mikel Luján, and Gavin Brown. A unified theory of diversity in ensemble learning. *Journal of Machine Learning Research*, 24(359):1–49, 2023.
- Chunming Zhang, Lixing Zhu, and Yanbo Shen. Robust estimation in regression and classification methods for large dimensional data. *Machine Learning*, 112(9):3361–3411, 2023.

Jun Zhang. Divergence function, duality, and convex analysis. *Neural Computation*, 16: 159–195, 2004.

Huaiyu Zhu and Richard Rohwer. Information geometric measurements of generalisation. Technical Report NCRG/95/005, Aston University, Birmingham, UK, 1995.