

Informative Label Missingness in Multiclass Classification: Information Geometry and Excess Risk

Fariborz Setoudehtazangi

University of Padua

Geoffrey McLachlan

g.mclachlan@uq.edu.au

The University of Queensland

Research Article

Keywords: Bayes classifier, classification risk, Fisher information, informative missingness, missing class labels, multiclass classification, semi-supervised learning

Posted Date: September 23rd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10892916/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Informative Label Missingness in Multiclass Classification: Information Geometry and Excess Risk

Fariborz Setoudehtazangi¹ and Geoffrey J. McLachlan^{*2}

¹Dipartimento di Scienze Statistiche, Università di Padova, Padova, Italy

²School of Mathematics and Physics, The University of Queensland, Australia

Abstract

Informative label missingness can change the usual efficiency ordering between completely and partially labelled classifiers because the pattern of missing labels may itself carry information about the classification model. We develop a general likelihood-based theory for this phenomenon in parametric multiclass classification. An efficient-information decomposition separates information lost through unavailable class memberships from information contributed by the missing-label mechanism. We then derive a quadratic expansion of plug-in excess risk over the active pairwise faces of the multiclass Bayes boundary, showing that classification efficiency depends on how information gains and losses align with directions that perturb the decision boundary. This yields a classification-weighted generalized-eigenvalue criterion under which informative partial classification may have smaller asymptotic classification risk without globally dominating complete classification in Fisher information. Near missing completely at random, with the marginal missing-label proportion fixed, redistribution of missing labels changes lost class-label information at first order, whereas efficient information from the missingness pattern appears only at second order. Three-class quadratic discriminant calculations, finite-sample experiments, and a semi-synthetic multiclass application illustrate the resulting regime-dependent behaviour.

Keywords: Bayes classifier; classification risk; Fisher information; informative missingness; missing class labels; multiclass classification; semi-supervised learning.

1 Introduction

Classification from partially labelled data arises when the feature vectors in a training sample are observed but the class memberships of some observations are unavailable. This problem has a long history in statistical discriminant analysis and is now commonly viewed as a form of semi-supervised learning (McLachlan, 1975; Dempster et al., 1977; Chapelle et al., 2006;

^{*}Corresponding author. Email: g.mclachlan@uq.edu.au

Van Engelen and Hoos, 2020; Ahfock and McLachlan, 2020). From a likelihood perspective, the unavailable class memberships constitute missing data, so labelled and unlabelled observations can be combined through the likelihood of the corresponding finite-mixture model. When label availability is unrelated to the observed data, an observation whose class membership is unknown is generally less informative for estimating a classification rule than the same observation with its label observed. The resulting loss of information and its implications for discriminant efficiency were studied in early work on partially classified samples (O’Neill, 1978; Ganesalingam and McLachlan, 1978; McLachlan and Scot, 1995); related questions concerning the statistical value of labelled and unlabelled observations have also been considered from a learning theory perspective (Castelli and Cover, 1996).

The comparison changes when label availability is related to the classification problem. In many applications, labels are assigned selectively. For example, a human expert may be more likely to leave an observation unclassified when its features place it in a region of substantial class overlap. The observed pattern of label availability can then contain information about parameters governing the classifier. Such situations are naturally described within a missing-data framework (Rubin, 1976; Mealli and Rubin, 2015). If M denotes the indicator that the class label Z is unavailable and $\mathbf{Y}$ is the observed feature vector, we consider mechanisms satisfying

$$M \perp Z \mid \mathbf{Y}, \quad \Pr(M = 1 \mid \mathbf{Y} = \mathbf{y}) = q(\mathbf{y}; \boldsymbol{\theta}, \boldsymbol{\xi}),$$

where $\boldsymbol{\theta}$ parameterizes the classification model and $\boldsymbol{\xi}$ contains parameters specific to the missing-label mechanism. Although M is conditionally independent of the latent class given the observed features, the mechanism need not be ignorable for likelihood inference on $\boldsymbol{\theta}$, because its distribution may itself depend on $\boldsymbol{\theta}$.

A counterintuitive consequence of this dependence was developed by Ahfock and McLachlan (2020). For two homoscedastic Gaussian classes, they modelled the probability of a missing label as a function of classification uncertainty and showed that the information carried by the missing-label indicators can exceed the information lost through the unavailable class memberships. Consequently, an estimated Bayes rule based on a partially classified sample can have a smaller asymptotic expected error rate than the corresponding rule based on a completely classified sample. Their formulation was motivated by the observation that unclassified observations may concentrate near class boundaries, and related the missingness probability to classification entropy or, in the binary Gaussian setting, to the squared discriminant function (McLachlan and Ahfock, 2019; Ahfock and McLachlan, 2020). This phenomenon and its relation to statistical semi-supervised learning have subsequently been reviewed by Ahfock and McLachlan (2023).

The phenomenon is not confined numerically to the homoscedastic binary model. In particular, Lyu (2024) studied Bayes-rule estimation under a specified missing-data mechanism for Gaussian mixture models and considered two- and three-component settings with unequal covariance matrices through simulation and real-data illustrations. More recent work has continued to emphasize the implications of informative missingness in semi-supervised learning and the roles of class overlap, the proportion of missing labels, and the strength of the association between missingness and classification uncertainty (Wu et al., 2026). Thus there is evidence that favourable informative missingness can arise beyond the particular

Gaussian configuration in which the phenomenon was first analysed.

What remains less understood is the *classification-risk geometry* of this phenomenon in a general multiclass problem. Existing analytic results explaining efficiency gains from informative label missingness have been developed primarily through binary discriminant structures. In a multiclass problem, however, the Bayes rule is determined by several pairwise equality surfaces, only portions of which form active decision boundaries, and these surfaces may meet at triple or higher-order class junctions. More importantly, global improvement in Fisher information is stronger than what classification itself requires. Information gained in a parameter direction that barely moves the Bayes boundary may have little effect on classification risk, whereas a gain concentrated in a strongly boundary-relevant direction may be substantially more important.

This issue connects informative label missingness with the classical asymptotic theory of estimated discriminant rules. Previous work has studied the effect of parameter estimation on classification error and developed large-sample and higher-order risk expansions for discriminant procedures (O’Neill, 1980; Taniguchi, 1994; Ducinkas and Saltyte, 2002). A central implication of this literature is that parameter-estimation precision and classification performance are not interchangeable: classification risk depends on how estimation error perturbs the decision rule. For informative partial classification, a comparison of Fisher information alone therefore does not reveal which information gains and losses are relevant to classification.

The aim of this paper is to connect these two strands of theory. We first derive an efficient-information decomposition for a general parametric multiclass model with informative label missingness. After eliminating the parameters specific to the missing-label mechanism, the information for the classification model is

$$\mathbf{I}_{\text{PC}}^{\text{eff}} = \mathbf{I}_{\text{CC}} - E[q(\mathbf{Y})\mathbf{I}_{Z|\mathbf{Y}}(\mathbf{Y})] + \mathbf{I}_M^{\text{eff}}, \quad (1.1)$$

where the second term is the information lost through unavailable class memberships and $\mathbf{I}_M^{\text{eff}} \succeq \mathbf{0}$ is the efficient information contributed by the missing-label indicators. The weighting by $q(\mathbf{Y})$ is important: the information cost of missing labels depends not only on how many labels are unavailable but also on where in the feature space they are missing. Missing completely at random is recovered as a special case in which the missingness indicators contribute no information about the classification parameters.

We then relate this information structure to classification performance by deriving a local quadratic excess-risk expansion for regular multiclass Bayes classifiers. The leading curvature is determined by the *active pairwise Bayes faces*, namely the portions of pairwise equality surfaces on which the corresponding two classes jointly attain the Bayes maximum. Under regularity and transversality conditions, triple and higher-order Bayes junctions do not contribute to the leading quadratic term. If $\hat{\boldsymbol{\theta}}$ is asymptotically normal with asymptotic covariance matrix $\mathbf{V}$, then

$$E\left\{R(\hat{\boldsymbol{\theta}})\right\} - R^* = \frac{1}{2n} \text{tr}(\mathbf{H}_R \mathbf{V}) + o(n^{-1}), \quad (1.2)$$

where $\mathbf{H}_R$ is the curvature matrix induced by the active Bayes faces.

Combining the information and risk calculations yields the main classification-specific comparison. Let $\mathbf{A}$ and $\mathbf{J}$ denote the relevant efficient information matrices under complete classification and informative partial classification, respectively. The leading classification advantage of the latter is governed by

$$\Delta_R = \text{tr} [\mathbf{H}_R (\mathbf{A}^{-1} - \mathbf{J}^{-1})]. \quad (1.3)$$

Thus a smaller asymptotic classification risk does not require $\mathbf{J} \succeq \mathbf{A}$. A partially classified experiment may lose information in some parameter directions and nevertheless improve classification if its information gains are sufficiently concentrated in directions that move the active Bayes boundary. A generalized-eigenvalue representation makes this distinction explicit by separating relative information gain or loss from the classification relevance of the corresponding parameter direction. Favourable informative missingness is therefore a directional alignment phenomenon rather than simply a question of global Fisher-information dominance.

We also examine how this advantage emerges as a missing-label mechanism departs from missing completely at random while the marginal proportion of missing labels is held fixed. Locally, redistributing the missing labels according to classification uncertainty alters the loss of conditional class-label information at first order, whereas the efficient information carried by the missingness indicators appears only at second order. Hence a weakly informative mechanism need not initially improve classification and may instead make it worse. Whether a favourable regime subsequently emerges depends on the global information geometry and its alignment with the active decision boundaries.

The theory is illustrated using a three-class Gaussian quadratic discriminant model with unequal covariance matrices. This is the smallest setting that allows curved pairwise Bayes faces and a genuine multiclass junction to arise simultaneously while the decision geometry remains directly interpretable. We examine how the classification-risk comparison varies with the proportion of missing labels, the strength of uncertainty dependence, class separation, prior imbalance, covariance heterogeneity, and the choice of entropy- or Gini-based uncertainty measures. The population calculations reveal face-specific gains and losses and geometry-dependent transitions between unfavourable and favourable regimes. A finite-sample Monte Carlo experiment then assesses the asymptotic predictions under likelihood estimation; in the reference configuration, the theoretical classification-risk coefficients, empirical quadratic approximations, and directly evaluated scaled excess risks are in close agreement.

The contribution of this paper is therefore not to establish that informative label missingness can sometimes be beneficial—that phenomenon is already known—nor simply to demonstrate it for another class-conditional distribution. Rather, we develop a general multiclass framework that explains how information gained and lost through informative label missingness is translated into classification risk through the geometry of the active Bayes boundary. This perspective identifies when and where the missingness pattern can improve classification and shows why favourable partial classification does not require global information dominance.

The remainder of the paper is organized as follows. Section 2 introduces the general par-

tially labelled model and derives the efficient-information decomposition. Section 3 considers efficient classifier information in the presence of nuisance parameters. Section 4 develops the multiclass excess-risk geometry, and Section 5 gives the classification-weighted spectral characterization and studies local departures from missing completely at random. Section 6 specializes the framework to quadratic discriminant analysis. Sections 7 and 8 present the population and finite-sample investigations, respectively. Section 9 provides a semi-synthetic application to the Vertebral Column data. Section 10 discusses the main findings, limitations, and directions for further research, and Section 11 concludes.

2 General model and information decomposition

Let $\mathbf{Y} \in \mathcal{Y} \subseteq \mathbb{R}^p$ denote the observed feature vector and let $Z \in \{1, \dots, g\}$ denote the corresponding class label. We assume a parametric class model of the form

$$\Pr_{\boldsymbol{\theta}}(Z = k) = \pi_k, \quad \mathbf{Y} \mid Z = k \sim f_k(\mathbf{y}; \boldsymbol{\vartheta}_k), \quad k = 1, \dots, g,$$

where $\boldsymbol{\theta} = (\boldsymbol{\pi}^\top, \boldsymbol{\vartheta}_1^\top, \dots, \boldsymbol{\vartheta}_g^\top)^\top$ denotes an identifiable finite-dimensional parameter vector, with a suitable parameterization of the class probabilities $\boldsymbol{\pi} = (\pi_1, \dots, \pi_g)^\top$.

The joint class-feature density is

$$p_{\boldsymbol{\theta}}(\mathbf{y}, Z = k) = \pi_k f_k(\mathbf{y}; \boldsymbol{\vartheta}_k),$$

and the marginal feature density is

$$p_{\boldsymbol{\theta}}(\mathbf{y}) = \sum_{k=1}^g \pi_k f_k(\mathbf{y}; \boldsymbol{\vartheta}_k).$$

The posterior class probabilities are therefore

$$\tau_k(\mathbf{y}; \boldsymbol{\theta}) = \Pr(Z = k \mid \mathbf{Y} = \mathbf{y}) = \frac{\pi_k f_k(\mathbf{y}; \boldsymbol{\vartheta}_k)}{p_{\boldsymbol{\theta}}(\mathbf{y})}.$$

The corresponding Bayes classifier assigns $\mathbf{y}$ to a class maximizing $\tau_k(\mathbf{y}; \boldsymbol{\theta})$, or equivalently, $\pi_k f_k(\mathbf{y}; \boldsymbol{\vartheta}_k)$.

Let

$$M = \begin{cases} 1, & \text{if the class label is unavailable,} \\ 0, & \text{if the class label is observed.} \end{cases}$$

We assume throughout that

$$M \perp Z \mid \mathbf{Y}, \tag{2.1}$$

while allowing the probability of label absence to depend on both the observed features and the parameters of the classification model:

$$\Pr_{\boldsymbol{\theta}, \boldsymbol{\xi}}(M = 1 \mid \mathbf{Y} = \mathbf{y}) = q(\mathbf{y}; \boldsymbol{\theta}, \boldsymbol{\xi}), \tag{2.2}$$

where $\boldsymbol{\xi}$ denotes parameters specific to the missing-label mechanism and

$$0 < q(\mathbf{y}; \boldsymbol{\theta}, \boldsymbol{\xi}) < 1$$

on the relevant support.

The observed datum for one individual is

$$\mathcal{O} = (\mathbf{Y}, M, (1 - M)Z).$$

Under (2.1), its density can be written as

$$\begin{aligned} p_{\boldsymbol{\theta}, \boldsymbol{\xi}}(\mathcal{O}) &= p_{\boldsymbol{\theta}}(\mathbf{Y}) \tau_Z(\mathbf{Y}; \boldsymbol{\theta})^{1-M} \\ &\times q(\mathbf{Y}; \boldsymbol{\theta}, \boldsymbol{\xi})^M \{1 - q(\mathbf{Y}; \boldsymbol{\theta}, \boldsymbol{\xi})\}^{1-M}. \end{aligned} \quad (2.3)$$

Consequently, the individual observed-data log-likelihood is

$$\begin{aligned} \ell_{\text{obs}}(\boldsymbol{\theta}, \boldsymbol{\xi}) &= \log p_{\boldsymbol{\theta}}(\mathbf{Y}) + (1 - M) \log \tau_Z(\mathbf{Y}; \boldsymbol{\theta}) \\ &+ M \log q(\mathbf{Y}; \boldsymbol{\theta}, \boldsymbol{\xi}) + (1 - M) \log \{1 - q(\mathbf{Y}; \boldsymbol{\theta}, \boldsymbol{\xi})\}. \end{aligned} \quad (2.4)$$

Expression (2.4) separates three sources of information: the marginal distribution of the features, the observed class labels, and the pattern of label availability. The last component is absent from an ignorable analysis when the missing-label mechanism is treated as ancillary to the classification model.

To quantify the contribution of an observed class label at a given feature value, define the conditional class score

$$\mathbf{S}_{Z|\mathbf{Y}} = \nabla_{\boldsymbol{\theta}} \log \tau_Z(\mathbf{Y}; \boldsymbol{\theta}),$$

and its conditional Fisher information

$$\mathbf{I}_{Z|\mathbf{Y}}(\boldsymbol{\theta}; \mathbf{y}) = E_{\boldsymbol{\theta}} [\mathbf{S}_{Z|\mathbf{Y}} \mathbf{S}_{Z|\mathbf{Y}}^\top \mid \mathbf{Y} = \mathbf{y}]. \quad (2.5)$$

This matrix measures the additional information provided by observing the class membership after the feature vector is known.

For the missing-label mechanism, write

$$q_{\boldsymbol{\theta}} = \frac{\partial q}{\partial \boldsymbol{\theta}}, \quad q_{\boldsymbol{\xi}} = \frac{\partial q}{\partial \boldsymbol{\xi}},$$

and define

$$\mathbf{B}_{ab} = E \left[\frac{q_a q_b^\top}{q(1 - q)} \right], \quad a, b \in \{\boldsymbol{\theta}, \boldsymbol{\xi}\}, \quad (2.6)$$

where the arguments of q are suppressed for notational simplicity.

The following result gives the central information decomposition.

Theorem 2.1. *Assume that the model is locally identifiable and satisfies the standard regularity conditions for Fisher-information calculations, including common support in a neighborhood of the true parameter, differentiability with respect to $(\boldsymbol{\theta}, \boldsymbol{\xi})$, interchange of differentiation and integration, finite second moments of the relevant score functions, and nonsingularity of $\mathbf{B}_{\boldsymbol{\xi}\boldsymbol{\xi}}$. Then the efficient Fisher information for $\boldsymbol{\theta}$, after eliminating the missingness-specific nuisance parameter $\boldsymbol{\xi}$, is*

$$\mathbf{I}_{\text{PC}}^{\text{eff}}(\boldsymbol{\theta}) = \mathbf{I}_{\text{CC}}(\boldsymbol{\theta}) - \mathbf{D}(\boldsymbol{\theta}, \boldsymbol{\xi}) + \mathbf{I}_M^{\text{eff}}(\boldsymbol{\theta}, \boldsymbol{\xi}), \quad (2.7)$$

where

$$\mathbf{D}(\boldsymbol{\theta}, \boldsymbol{\xi}) = E \left[q(\mathbf{Y}; \boldsymbol{\theta}, \boldsymbol{\xi}) \mathbf{I}_{Z|\mathbf{Y}}(\boldsymbol{\theta}; \mathbf{Y}) \right], \quad (2.8)$$

and

$$\mathbf{I}_M^{\text{eff}} = \mathbf{B}_{\boldsymbol{\theta}\boldsymbol{\theta}} - \mathbf{B}_{\boldsymbol{\theta}\boldsymbol{\xi}} \mathbf{B}_{\boldsymbol{\xi}\boldsymbol{\xi}}^{-1} \mathbf{B}_{\boldsymbol{\xi}\boldsymbol{\theta}} \succeq \mathbf{0}. \quad (2.9)$$

Here $\mathbf{I}_{\text{CC}}(\boldsymbol{\theta})$ denotes the Fisher information that would be available if all class labels were observed.

The proof is given in Supplementary Section S1.

The decomposition in (2.7) separates two effects that are otherwise confounded in a partially labelled sample. The matrix $\mathbf{D}$ is the information lost because some class memberships are not observed. Importantly, it is not determined solely by the marginal missing-label proportion. Instead, each feature value is weighted by both its probability of losing the label and the conditional information that the label would have provided. In contrast, $\mathbf{I}_M^{\text{eff}}$ is the efficient information conveyed by the missing-label indicators after the missingness-specific parameters have been removed.

For the finite mixture model considered here, $\mathbf{I}_{Z|\mathbf{Y}}$ has a particularly useful form. Define the class-specific score

$$\mathbf{s}_k(\mathbf{y}) = \nabla_{\boldsymbol{\theta}} \log \{ \pi_k f_k(\mathbf{y}; \boldsymbol{\vartheta}_k) \},$$

and

$$\bar{\mathbf{s}}(\mathbf{y}) = \sum_{k=1}^g \tau_k(\mathbf{y}; \boldsymbol{\theta}) \mathbf{s}_k(\mathbf{y}).$$

Since

$$\nabla_{\boldsymbol{\theta}} \log \tau_k(\mathbf{y}; \boldsymbol{\theta}) = \mathbf{s}_k(\mathbf{y}) - \bar{\mathbf{s}}(\mathbf{y}),$$

we obtain

$$\begin{aligned} \mathbf{I}_{Z|\mathbf{Y}}(\boldsymbol{\theta}; \mathbf{y}) &= \sum_{k=1}^g \tau_k(\mathbf{y}; \boldsymbol{\theta}) \{ \mathbf{s}_k(\mathbf{y}) - \bar{\mathbf{s}}(\mathbf{y}) \} \\ &\quad \times \{ \mathbf{s}_k(\mathbf{y}) - \bar{\mathbf{s}}(\mathbf{y}) \}^{\top}. \end{aligned} \quad (2.10)$$

Thus the information supplied by the class label is the posterior covariance of the class-specific score vectors. This representation will be useful below when we examine how uncertainty-dependent missingness redistributes label information across the feature space.

A direct consequence of Theorem 2.1 is the standard missing-completely-at-random benchmark.

Corollary 2.2. *Suppose*

$$q(\mathbf{Y}; \boldsymbol{\theta}, \boldsymbol{\xi}) = \gamma, \quad 0 < \gamma < 1,$$

where γ is variation-independent of $\boldsymbol{\theta}$. Then

$$\mathbf{I}_M^{\text{eff}} = \mathbf{0},$$

and

$$\mathbf{I}_{\text{PC}}^{\text{eff}} = \mathbf{I}_{\text{CC}} - \gamma E [\mathbf{I}_{Z|Y}(\boldsymbol{\theta}; \mathbf{Y})] \preceq \mathbf{I}_{\text{CC}}. \quad (2.11)$$

Hence randomly removing class labels cannot increase Fisher information about the classification model. Any favourable information effect must arise because the observed pattern of missing labels is itself informative about $\boldsymbol{\theta}$.

Theorem 2.1 concerns the complete identifiable data-model parameter $\boldsymbol{\theta}$, after eliminating only the missingness-specific parameter $\boldsymbol{\xi}$. In many classification models, however, some components of $\boldsymbol{\theta}$ may be nuisance parameters for a particular classifier parameter or functional of interest. Eliminating such nuisance parameters does not generally preserve the simple additive form in (2.7). The resulting efficient classifier information is considered in the next section.

3 Efficient classifier information in the presence of nuisance parameters

The decomposition in (2.7) concerns the full identifiable parameter $\boldsymbol{\theta}$, after eliminating only the missingness-specific parameter $\boldsymbol{\xi}$. In some classification models, however, the inferential target depends on only part of $\boldsymbol{\theta}$, with the remaining components acting as nuisance parameters. Eliminating these additional nuisance parameters need not preserve the additive information decomposition of Section 2.

Write

$$\boldsymbol{\theta} = \begin{pmatrix} \boldsymbol{\beta} \\ \boldsymbol{\lambda} \end{pmatrix}, \quad \boldsymbol{\beta} \in \mathbb{R}^r, \quad \boldsymbol{\lambda} \in \mathbb{R}^s,$$

where $\boldsymbol{\beta}$ is the parameter of interest and $\boldsymbol{\lambda}$ is nuisance. Here “nuisance” is meant in the inferential sense: the target depends on $\boldsymbol{\beta}$, whereas $\boldsymbol{\lambda}$ is eliminated by efficient-score projection. Parameters that determine the Bayes decision boundary are therefore not nuisance when classification risk itself is the target.

Let

$$\mathbf{A} = \mathbf{I}_{\text{CC}}(\boldsymbol{\theta}), \quad \mathbf{K} = \mathbf{I}_M^{\text{eff}} - \mathbf{D}, \quad \mathbf{J} = \mathbf{I}_{\text{PC}}^{\text{eff}} = \mathbf{A} + \mathbf{K},$$

and partition $\mathbf{A}$ and $\mathbf{K}$ conformably with $(\boldsymbol{\beta}^\top, \boldsymbol{\lambda}^\top)^\top$:

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}_{\beta\beta} & \mathbf{A}_{\beta\lambda} \\ \mathbf{A}_{\lambda\beta} & \mathbf{A}_{\lambda\lambda} \end{pmatrix}, \quad \mathbf{K} = \begin{pmatrix} \mathbf{K}_{\beta\beta} & \mathbf{K}_{\beta\lambda} \\ \mathbf{K}_{\lambda\beta} & \mathbf{K}_{\lambda\lambda} \end{pmatrix}.$$

Under complete classification, the efficient information for $\boldsymbol{\beta}$ is the Schur complement

$$\mathbf{A}_{\text{eff}}(\boldsymbol{\beta}) = \mathbf{A}_{\beta\beta} - \mathbf{A}_{\beta\lambda} \mathbf{A}_{\lambda\lambda}^{-1} \mathbf{A}_{\lambda\beta}. \quad (3.1)$$

To express the corresponding information under partial classification, define $\mathbf{R} = \mathbf{A}_{\beta\lambda}\mathbf{A}_{\lambda\lambda}^{-1}$ and

$$\widetilde{\mathbf{K}}_{\beta\beta} = \mathbf{K}_{\beta\beta} - \mathbf{R}\mathbf{K}_{\lambda\beta} - \mathbf{K}_{\beta\lambda}\mathbf{R}^\top + \mathbf{R}\mathbf{K}_{\lambda\lambda}\mathbf{R}^\top, \quad (3.2)$$

$$\widetilde{\mathbf{K}}_{\beta\lambda} = \mathbf{K}_{\beta\lambda} - \mathbf{R}\mathbf{K}_{\lambda\lambda}. \quad (3.3)$$

These quantities describe the perturbation induced by partial classification after orthogonalization with respect to the complete-data nuisance score.

Proposition 3.1. *Assume that $\mathbf{A}_{\lambda\lambda}$ and $\mathbf{J}_{\lambda\lambda} = \mathbf{A}_{\lambda\lambda} + \mathbf{K}_{\lambda\lambda}$ are nonsingular. Then the efficient Fisher information for $\boldsymbol{\beta}$ under informative partial classification is*

$$\mathbf{J}_{\text{eff}}(\boldsymbol{\beta}) = \mathbf{A}_{\text{eff}}(\boldsymbol{\beta}) + \widetilde{\mathbf{K}}_{\beta\beta} - \widetilde{\mathbf{K}}_{\beta\lambda}\mathbf{J}_{\lambda\lambda}^{-1}\widetilde{\mathbf{K}}_{\lambda\beta}. \quad (3.4)$$

Moreover,

$$\mathbf{C}_K = \widetilde{\mathbf{K}}_{\beta\lambda}\mathbf{J}_{\lambda\lambda}^{-1}\widetilde{\mathbf{K}}_{\lambda\beta} \succeq \mathbf{0}. \quad (3.5)$$

The proof is given in Supplementary Section S2. Proposition 3.1 shows that eliminating additional nuisance parameters introduces a nonnegative coupling penalty. In particular,

$$\mathbf{J}_{\text{eff}}(\boldsymbol{\beta}) - \mathbf{A}_{\text{eff}}(\boldsymbol{\beta}) = \widetilde{\mathbf{K}}_{\beta\beta} - \mathbf{C}_K. \quad (3.6)$$

Thus $\widetilde{\mathbf{K}}_{\beta\beta}$ represents the net information perturbation in the complete-data efficient direction for $\boldsymbol{\beta}$, whereas $\mathbf{C}_K$ measures the loss arising from the coupling of that direction with the nuisance score under partial classification. A positive perturbation in the classifier-relevant block is therefore not sufficient for an information gain; it must also dominate this coupling penalty.

If $\widetilde{\mathbf{K}}_{\beta\lambda} = \mathbf{0}$, then $\mathbf{C}_K = \mathbf{0}$ and

$$\mathbf{J}_{\text{eff}}(\boldsymbol{\beta}) = \mathbf{A}_{\text{eff}}(\boldsymbol{\beta}) + \widetilde{\mathbf{K}}_{\beta\beta}. \quad (3.7)$$

In this case, informative partial classification does not recouple the complete-data efficient score for $\boldsymbol{\beta}$ with the nuisance score, and the additive form is recovered after nuisance elimination.

This distinction is relevant when inference concerns a lower-dimensional classifier parameter. In the multiclass risk analysis that follows, however, all parameters determining the Bayes decision boundary are retained in the classifier-relevant vector. Thus, for quadratic discriminant analysis, the class probabilities, means, and covariance matrices are not eliminated as nuisance parameters.

4 Excess-risk geometry for multiclass Bayes classification

The information decompositions in Sections 2 and 3 describe how informative label missingness affects the precision with which the parameters of a classification model can be estimated.

To determine whether these changes improve classification, however, the information matrix must be related to the geometry of the Bayes decision boundary. We develop this connection here.

For $k = 1, \dots, g$, define the prior-weighted class density $r_k(\mathbf{y}; \boldsymbol{\theta}) = \pi_k f_k(\mathbf{y}; \boldsymbol{\theta}_k)$, and write $r_k^0(\mathbf{y}) = r_k(\mathbf{y}; \boldsymbol{\theta}_0)$ at the true parameter $\boldsymbol{\theta}_0$. The Bayes classifier is

$$C_0(\mathbf{y}) = \arg \max_{1 \leq k \leq g} r_k^0(\mathbf{y}), \quad (4.1)$$

with an arbitrary fixed rule for breaking ties on sets of probability zero. For $\boldsymbol{\theta}$ in a neighborhood of $\boldsymbol{\theta}_0$, let

$$C_{\boldsymbol{\theta}}(\mathbf{y}) = \arg \max_k r_k(\mathbf{y}; \boldsymbol{\theta})$$

denote the corresponding plug-in classifier, so that $C_{\boldsymbol{\theta}_0} = C_0$. For a local perturbation $\mathbf{h}$, write

$$C_{\mathbf{h}}(\mathbf{y}) := C_{\boldsymbol{\theta}_0 + \mathbf{h}}(\mathbf{y}).$$

For $k \neq l$, define the pairwise weighted-density contrast

$$g_{kl}(\mathbf{y}; \boldsymbol{\theta}) = r_k(\mathbf{y}; \boldsymbol{\theta}) - r_l(\mathbf{y}; \boldsymbol{\theta}). \quad (4.2)$$

Not every equality surface $\{g_{kl} = 0\}$ contributes to the Bayes decision boundary. Only the portion on which classes k and l jointly attain the largest prior-weighted density is relevant.

Definition 4.1. The active Bayes face separating classes k and l is

$$\mathcal{F}_{kl} = \left\{ \mathbf{y} : r_k^0(\mathbf{y}) = r_l^0(\mathbf{y}) > \max_{m \notin \{k, l\}} r_m^0(\mathbf{y}) \right\}. \quad (4.3)$$

The distinction between a pairwise equality surface and an active face is essential in the multiclass setting. If $r_k^0(\mathbf{y}) = r_l^0(\mathbf{y})$ while a third class has a strictly larger weighted density, perturbing the k -versus- l comparison does not change the Bayes decision locally and therefore does not contribute to the leading excess risk.

For $\mathbf{s} \in \mathcal{F}_{kl}$, let $\mathbf{b}_{kl}(\mathbf{s}) = \nabla_{\boldsymbol{\theta}} g_{kl}(\mathbf{s}; \boldsymbol{\theta}_0)$. This vector describes the first-order change in the k -versus- l contrast induced by a local parameter perturbation at the boundary point $\mathbf{s}$.

We impose regularity conditions excluding singular decision geometry. For each nonempty active face, assume

$$\nabla_{\mathbf{y}} g_{kl}(\mathbf{s}; \boldsymbol{\theta}_0) \neq \mathbf{0}, \quad \mathbf{s} \in \mathcal{F}_{kl}. \quad (4.4)$$

The regular part of $\mathcal{F}_{kl}$ is then a $(p - 1)$ -dimensional smooth hypersurface. At points where $r \geq 3$ classes tie at the Bayes maximum, we additionally assume transversality: $r - 1$ independent pairwise contrast gradients span the normal space of the corresponding tie stratum. Ties of positive Lebesgue measure are excluded. When the feature support is unbounded, we additionally assume the tail condition stated in Supplementary Section S4. When one or more active faces are noncompact, we also assume the regular-exhaustion and boundary-integrability conditions stated there.

Let

$$R(\boldsymbol{\theta}) = \Pr_{\boldsymbol{\theta}_0} \{C_{\boldsymbol{\theta}}(\mathbf{Y}) \neq Z\}$$

denote the classification risk of the rule determined by $\boldsymbol{\theta}$, evaluated under the true distribution $\boldsymbol{\theta}_0$, and let $R^* = R(\boldsymbol{\theta}_0)$ denote the Bayes risk. The local geometry of the excess risk is characterized by the following result.

Theorem 4.2. *Suppose that the active Bayes boundary lies in the interior of the common feature support and that the class-weighted densities are twice continuously differentiable in $(\mathbf{y}, \boldsymbol{\theta})$ in a neighborhood of the active Bayes boundary. Assume that the active pairwise faces satisfy (4.4), that higher-order active tie sets satisfy the stated transversality condition, and that ties of positive Lebesgue measure are absent. When the feature support is unbounded, we additionally assume the tail condition stated in Supplementary Section S4. When one or more active faces are noncompact, we also assume the regular-exhaustion and boundary-integrability conditions stated there. Then, as $\mathbf{h} \rightarrow \mathbf{0}$,*

$$R(\boldsymbol{\theta}_0 + \mathbf{h}) - R^* = \frac{1}{2} \mathbf{h}^\top \mathbf{H}_R \mathbf{h} + o(\|\mathbf{h}\|^2), \quad (4.5)$$

where

$$\mathbf{H}_R = \sum_{1 \leq k < l \leq g} \mathbf{H}_{kl}, \quad \mathbf{H}_{kl} = \int_{\mathcal{F}_{kl}} \frac{\mathbf{b}_{kl}(\mathbf{s}) \mathbf{b}_{kl}(\mathbf{s})^\top}{\|\nabla_{\mathbf{y}} g_{kl}(\mathbf{s}; \boldsymbol{\theta}_0)\|} dS(\mathbf{s}). \quad (4.6)$$

In particular, $\mathbf{H}_R \succeq \mathbf{0}$. Generic triple and higher-order Bayes junctions do not contribute to the quadratic term.

A proof is given in Supplementary Sections S3–S4. The argument localizes classification disagreement to a thin neighborhood of the Bayes boundary and uses the pairwise contrast g_{kl} as a coordinate normal to each active face. The loss incurred by crossing the k -versus- l boundary is then $|g_{kl}|$, and integration across the displaced boundary yields the quadratic contribution in (4.6). Under transversality, higher-order junctions are lower-dimensional and contribute only to higher-order terms.

The matrix $\mathbf{H}_R$ therefore defines a local classification-relevance metric on the parameter space. For any direction $\mathbf{v}$,

$$\mathbf{v}^\top \mathbf{H}_R \mathbf{v} = \sum_{k < l} \int_{\mathcal{F}_{kl}} \frac{\{\mathbf{b}_{kl}(\mathbf{s})^\top \mathbf{v}\}^2}{\|\nabla_{\mathbf{y}} g_{kl}(\mathbf{s}; \boldsymbol{\theta}_0)\|} dS(\mathbf{s}). \quad (4.7)$$

Thus a direction $\mathbf{v}$ is locally irrelevant to classification if and only if $\mathbf{b}_{kl}(\mathbf{s})^\top \mathbf{v} = 0$ for dS -almost every $\mathbf{s}$ on every active face $\mathcal{F}_{kl}$; equivalently, it produces no first-order displacement of the active Bayes boundary.

An equivalent representation in terms of pairwise log-contrasts, which is convenient for the QDA calculations in Section 6, is given in Supplementary Section S4.

Theorem 4.2 immediately yields an asymptotic classification-risk expansion for regular parameter estimators.

Corollary 4.3. *Suppose that*

$$\sqrt{n} \left(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0 \right) \xrightarrow{d} N(\mathbf{0}, \mathbf{V}),$$

and that, for some $\delta > 0$,

$$\sup_n E \left\| \sqrt{n} \left(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0 \right) \right\|^{2+\delta} < \infty.$$

Then

$$n \left\{ R(\hat{\boldsymbol{\theta}}_n) - R^* \right\} \xrightarrow{d} \frac{1}{2} \mathbf{Z}^\top \mathbf{H}_R \mathbf{Z}, \quad \mathbf{Z} \sim N(\mathbf{0}, \mathbf{V}), \quad (4.8)$$

and

$$E \left\{ R(\hat{\boldsymbol{\theta}}_n) \right\} - R^* = \frac{1}{2n} \text{tr}(\mathbf{H}_R \mathbf{V}) + o(n^{-1}). \quad (4.9)$$

The proof and the corresponding weighted- χ^2 representation of the limiting quadratic form are given in Supplementary Section S4. For an efficient likelihood estimator, $\mathbf{V} = \mathbf{I}_{\text{eff}}^{-1}$, and hence

$$E \left\{ R(\hat{\boldsymbol{\theta}}_n) \right\} - R^* = \frac{1}{2n} \text{tr}(\mathbf{H}_R \mathbf{I}_{\text{eff}}^{-1}) + o(n^{-1}). \quad (4.10)$$

This expression provides the required link between the information calculations of Sections 2–3 and classification performance: information gains matter only insofar as they occur in directions to which $\mathbf{H}_R$ assigns classification relevance.

Let $\mathbf{A}$ and $\mathbf{J}$ denote the relevant efficient information matrices under complete and informative partial classification, respectively. Their leading expected excess-risk difference is determined by

$$\Delta_R = \text{tr} \left[\mathbf{H}_R (\mathbf{A}^{-1} - \mathbf{J}^{-1}) \right], \quad (4.11)$$

since

$$E \left\{ R(\hat{\boldsymbol{\theta}}_{\text{CC}}) \right\} - E \left\{ R(\hat{\boldsymbol{\theta}}_{\text{PC}}) \right\} = \frac{\Delta_R}{2n} + o(n^{-1}). \quad (4.12)$$

Thus $\Delta_R > 0$ characterizes favourable informative missingness at the leading n^{-1} order in expected classification risk. The next section examines this criterion through the relative information geometry of the two experiments.

5 Classification-weighted information and favourable missingness

The previous section shows that, for a regular estimator with asymptotic covariance matrix $\mathbf{V}/n$, the leading excess classification risk is governed by

$$\text{tr}(\mathbf{H}_R \mathbf{V}).$$

Thus the comparison between complete and partial classification depends not only on the amount of information available, but also on the directions in which that information is gained or lost relative to the geometry of the active Bayes boundary.

Let $\mathbf{A}$ and $\mathbf{J}$ denote the relevant efficient information matrices under complete and informative partial classification, respectively, and assume throughout that $\mathbf{A} \succ \mathbf{0}$ and $\mathbf{J} \succ \mathbf{0}$. The leading difference in expected excess classification risk is

$$\Delta_R = \text{tr} \left[\mathbf{H}_R (\mathbf{A}^{-1} - \mathbf{J}^{-1}) \right]. \quad (5.1)$$

By (4.12), $\Delta_R > 0$ implies that informative partial classification has the smaller expected excess classification risk at the leading n^{-1} order.

The corresponding asymptotic relative efficiency is

$$\text{ARE}_R = \frac{\text{tr}(\mathbf{H}_R \mathbf{A}^{-1})}{\text{tr}(\mathbf{H}_R \mathbf{J}^{-1})}, \quad (5.2)$$

provided the denominator is positive. Hence $\text{ARE}_R > 1$ if and only if $\Delta_R > 0$. We use Δ_R as the primary criterion because it admits an additive directional decomposition.

Let $\mathbf{A}^{1/2}$ denote the symmetric positive-definite square root of $\mathbf{A}$, with inverse $\mathbf{A}^{-1/2}$, and define

$$\mathbf{C} = \mathbf{A}^{-1/2} \mathbf{J} \mathbf{A}^{-1/2}, \quad \mathbf{W} = \mathbf{A}^{-1/2} \mathbf{H}_R \mathbf{A}^{-1/2}. \quad (5.3)$$

Write

$$\mathbf{C} = \mathbf{Q} \mathbf{\Lambda} \mathbf{Q}^\top, \quad \mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_r),$$

where $\mathbf{Q}^\top \mathbf{Q} = \mathbf{I}$ and $\lambda_j > 0$. For the j th eigenvector $\mathbf{q}_j$, define

$$w_j = \mathbf{q}_j^\top \mathbf{W} \mathbf{q}_j \geq 0. \quad (5.4)$$

Theorem 5.1. *Under the conditions above,*

$$\Delta_R = \sum_{j=1}^r w_j \frac{\lambda_j - 1}{\lambda_j}. \quad (5.5)$$

Consequently, if $\mathcal{P} = \{j : \lambda_j > 1\}$ and $\mathcal{N} = \{j : \lambda_j < 1\}$, then $\Delta_R > 0$ if and only if

$$\sum_{j \in \mathcal{P}} w_j \frac{\lambda_j - 1}{\lambda_j} > \sum_{j \in \mathcal{N}} w_j \frac{1 - \lambda_j}{\lambda_j}. \quad (5.6)$$

The proof is given in Supplementary Section S5.

The generalized eigenvalues in Theorem 5.1 have a direct statistical interpretation. If $\mathbf{v}_j = \mathbf{A}^{-1/2} \mathbf{q}_j$, then

$$\mathbf{J} \mathbf{v}_j = \lambda_j \mathbf{A} \mathbf{v}_j, \quad (5.7)$$

and, under the normalization $\mathbf{v}_j^\top \mathbf{A} \mathbf{v}_j = 1$,

$$\lambda_j = \mathbf{v}_j^\top \mathbf{J} \mathbf{v}_j.$$

Thus $\lambda_j > 1$ indicates greater information under partial classification than under complete classification in direction $\mathbf{v}_j$, whereas $\lambda_j < 1$ indicates less information in that direction.

Moreover,

$$w_j = \mathbf{v}_j^\top \mathbf{H}_R \mathbf{v}_j. \quad (5.8)$$

Hence w_j measures the classification relevance of the same direction. By (4.7), it is large when perturbation in $\mathbf{v}_j$ moves one or more active Bayes faces substantially.

Theorem 5.1 therefore characterizes favourable informative missingness as an alignment phenomenon. Global Fisher-information dominance is not required: $\mathbf{J} \not\preceq \mathbf{A}$ is compatible with $\Delta_R > 0$ whenever information gains occur primarily in directions receiving large classification weights.

For interpretation, define

$$G_R = \sum_{\lambda_j > 1} w_j \frac{\lambda_j - 1}{\lambda_j}, \quad L_R = \sum_{\lambda_j < 1} w_j \frac{1 - \lambda_j}{\lambda_j}. \quad (5.9)$$

Then

$$\Delta_R = G_R - L_R. \quad (5.10)$$

Thus informative partial classification is favourable precisely when its classification-weighted information gains exceed the corresponding losses.

Two limiting cases follow immediately. If $\mathbf{J} \succeq \mathbf{A}$, then all $\lambda_j \geq 1$ and hence $\Delta_R \geq 0$. Conversely, if $\mathbf{J} \preceq \mathbf{A}$, then $\Delta_R \leq 0$. Neither matrix ordering is required in the general case.

5.1 Local departures from missing completely at random

We now examine how favourable informative missingness can emerge as the missing-label mechanism departs continuously from missing completely at random. Let $U_{\boldsymbol{\theta}}(\mathbf{Y})$ be a smooth scalar measure of classification uncertainty and consider

$$q_t(\mathbf{y}; \boldsymbol{\theta}) = \text{expit} \{ \alpha(t) + t U_{\boldsymbol{\theta}}(\mathbf{y}) \}, \quad t \geq 0. \quad (5.11)$$

For comparisons across t , the intercept $\alpha(t)$ is calibrated at the true parameter $\boldsymbol{\theta}_0$ so that

$$E_{\boldsymbol{\theta}_0} [q_t(\mathbf{Y}; \boldsymbol{\theta}_0)] = \gamma, \quad 0 < \gamma < 1. \quad (5.12)$$

Thus t changes the locations at which labels are unavailable while keeping their marginal proportion fixed. At $t = 0$, $q_0(\mathbf{y}) = \gamma$, and the mechanism reduces to MCAR.

Equation (5.12) defines a sequence of population experiments; it is not imposed as a constraint on the likelihood for arbitrary $\boldsymbol{\theta}$. In estimation, the missingness intercept and slope remain ordinary nuisance parameters.

Write $U(\mathbf{Y}) = U_{\boldsymbol{\theta}_0}(\mathbf{Y})$ and

$$\mathbf{G}(\mathbf{Y}) = \nabla_{\boldsymbol{\theta}} U_{\boldsymbol{\theta}}(\mathbf{Y})|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0}.$$

The following proposition describes the local behavior of the information decomposition around MCAR.

Proposition 5.2. *Assume that*

$$\text{Var}\{U(\mathbf{Y})\} > 0, \quad E\{U(\mathbf{Y})^2\} < \infty, \quad E\{\|\mathbf{G}(\mathbf{Y})\|^2\} < \infty,$$

and that the required differentiations with respect to t and $\boldsymbol{\theta}$ may be interchanged with expectation in a neighborhood of $(t, \boldsymbol{\theta}) = (0, \boldsymbol{\theta}_0)$, with sufficient local smoothness for the Taylor expansions below. Then

$$\alpha(0) = \text{logit}(\gamma), \quad \alpha'(0) = -E\{U(\mathbf{Y})\}, \quad (5.13)$$

and

$$q_t(\mathbf{Y}) = \gamma + t\gamma(1 - \gamma) [U(\mathbf{Y}) - E\{U(\mathbf{Y})\}] + O(t^2). \quad (5.14)$$

The label-information loss

$$\mathbf{D}(t) = E [q_t(\mathbf{Y}) \mathbf{I}_{Z|\mathbf{Y}}(\mathbf{Y})]$$

satisfies

$$\mathbf{D}'(0) = \gamma(1 - \gamma) E [[U(\mathbf{Y}) - E\{U(\mathbf{Y})\}] \mathbf{I}_{Z|\mathbf{Y}}(\mathbf{Y})]. \quad (5.15)$$

By contrast, the efficient information contributed by the missing-label mechanism satisfies

$$\mathbf{I}_M^{\text{eff}}(t) = t^2 \gamma(1 - \gamma) \boldsymbol{\mathcal{V}}_U + o(t^2), \quad (5.16)$$

where

$$\boldsymbol{\mathcal{V}}_U = E [\mathbf{G}\mathbf{G}^\top] - E [\mathbf{G}\mathbf{X}^\top] E [\mathbf{X}\mathbf{X}^\top]^{-1} E [\mathbf{X}\mathbf{G}^\top] \succeq \mathbf{0}, \quad (5.17)$$

with

$$\mathbf{X} = \begin{pmatrix} 1 \\ U(\mathbf{Y}) \end{pmatrix}.$$

Finally, if

$$\mathbf{J}_0 = \mathbf{A} - \gamma E [\mathbf{I}_{Z|\mathbf{Y}}(\mathbf{Y})],$$

then the derivative of the classification advantage at MCAR is

$$\Delta'_R(0) = -\text{tr} [\mathbf{H}_R \mathbf{J}_0^{-1} \mathbf{D}'(0) \mathbf{J}_0^{-1}]. \quad (5.18)$$

The proof is given in Supplementary Section S6.

Proposition 5.2 reveals an asymmetry in how informative missingness first enters the experiment. Redistributing missing labels changes the conditional class-label information loss at order t , whereas the efficient information carried by the missingness indicators appears only at order t^2 . Weak dependence between label absence and classification uncertainty therefore need not improve classification.

In particular, if $\mathbf{D}'(0) \succeq \mathbf{0}$ and

$$\text{tr} [\mathbf{H}_R \mathbf{J}_0^{-1} \mathbf{D}'(0) \mathbf{J}_0^{-1}] > 0,$$

then $\Delta'_R(0) < 0$. A small departure from MCAR then initially worsens classification. This can occur, for example, when observations with above-average uncertainty are also those for which the true class label carries greater classification-relevant information.

The result is local: it does not imply that $\Delta_R(t)$ is monotone, nor that it must eventually become positive. Whether a favourable regime emerges for larger t depends on the global information geometry of the missingness mechanism. This behavior is examined numerically in Section 7.

6 Three-class quadratic discriminant analysis

We illustrate the preceding theory using Gaussian quadratic discriminant analysis (QDA). The purpose is not to introduce a new classification model, but to examine the information–risk framework in a setting where the Bayes geometry is genuinely multiclass and nonlinear. Unequal covariance matrices produce curved pairwise decision boundaries, all distributional parameters may affect the classifier, and different active faces may meet at multiclass junctions.

Let $Z \in \{1, 2, 3\}$, with $\Pr(Z = k) = \pi_k$, and

$$\mathbf{Y} \mid Z = k \sim N_p(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \quad k = 1, 2, 3, \quad (6.1)$$

where $\pi_k > 0$, $\sum_{k=1}^3 \pi_k = 1$, and $\boldsymbol{\Sigma}_k \succ \mathbf{0}$. The prior-weighted class densities are

$$r_k(\mathbf{y}) = \pi_k \phi_p(\mathbf{y}; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \quad (6.2)$$

and the Bayes rule assigns $\mathbf{y}$ to the class maximizing $r_k(\mathbf{y})$.

Although the theory in Sections 2–5 is stated for arbitrary g and p , the numerical analysis uses $g = 3$ and $p = 2$. This is the smallest configuration that allows curved active pairwise faces and a genuine three-class junction to occur simultaneously while remaining directly visualizable.

6.1 Bayes geometry and risk curvature

For $k \neq l$, define the log weighted-density contrast

$$d_{kl}(\mathbf{y}) = \log \frac{r_k(\mathbf{y})}{r_l(\mathbf{y})}. \quad (6.3)$$

Under (6.1),

$$\begin{aligned} d_{kl}(\mathbf{y}) = & \log \frac{\pi_k}{\pi_l} - \frac{1}{2} \log \frac{|\boldsymbol{\Sigma}_k|}{|\boldsymbol{\Sigma}_l|} \\ & - \frac{1}{2} (\mathbf{y} - \boldsymbol{\mu}_k)^\top \boldsymbol{\Sigma}_k^{-1} (\mathbf{y} - \boldsymbol{\mu}_k) \\ & + \frac{1}{2} (\mathbf{y} - \boldsymbol{\mu}_l)^\top \boldsymbol{\Sigma}_l^{-1} (\mathbf{y} - \boldsymbol{\mu}_l). \end{aligned} \quad (6.4)$$

Thus $\{\mathbf{y} : d_{kl}(\mathbf{y}) = 0\}$ is a quadratic hypersurface. For three classes, the corresponding active face is

$$\mathcal{F}_{kl} = \{\mathbf{y} : d_{kl}(\mathbf{y}) = 0, r_k(\mathbf{y}) = r_l(\mathbf{y}) > r_m(\mathbf{y})\}, \quad (6.5)$$

where $m \notin \{k, l\}$.

Differentiation with respect to the feature vector gives

$$\nabla_{\mathbf{y}} d_{kl}(\mathbf{y}) = -\boldsymbol{\Sigma}_k^{-1} (\mathbf{y} - \boldsymbol{\mu}_k) + \boldsymbol{\Sigma}_l^{-1} (\mathbf{y} - \boldsymbol{\mu}_l), \quad (6.6)$$

so a regular point of $\mathcal{F}_{kl}$ satisfies

$$\Sigma_k^{-1}(\mathbf{y} - \boldsymbol{\mu}_k) \neq \Sigma_l^{-1}(\mathbf{y} - \boldsymbol{\mu}_l).$$

At a three-class junction $\mathbf{y}^\dagger$,

$$r_1(\mathbf{y}^\dagger) = r_2(\mathbf{y}^\dagger) = r_3(\mathbf{y}^\dagger), \quad (6.7)$$

or equivalently $d_{12}(\mathbf{y}^\dagger) = d_{13}(\mathbf{y}^\dagger) = 0$. For $p = 2$, transversality requires

$$\det \begin{pmatrix} \nabla_{\mathbf{y}} d_{12}(\mathbf{y}^\dagger)^\top \\ \nabla_{\mathbf{y}} d_{13}(\mathbf{y}^\dagger)^\top \end{pmatrix} \neq 0. \quad (6.8)$$

Under this condition the junction is isolated and, by Theorem 4.2, contributes only a higher-order term to the local excess risk.

For the class probabilities, we use class 3 as the baseline and write

$$\alpha_1 = \log \frac{\pi_1}{\pi_3}, \quad \alpha_2 = \log \frac{\pi_2}{\pi_3}. \quad (6.9)$$

A convenient theoretical parameter vector is

$$\boldsymbol{\theta} = (\boldsymbol{\alpha}^\top, \boldsymbol{\mu}_1^\top, \boldsymbol{\mu}_2^\top, \boldsymbol{\mu}_3^\top, \text{vech}(\Sigma_1)^\top, \text{vech}(\Sigma_2)^\top, \text{vech}(\Sigma_3)^\top)^\top.$$

For numerical optimization, the covariance matrices are represented instead through unconstrained log-Cholesky coordinates, which guarantee positive definiteness. Information and risk quantities are transformed consistently between the two parameterizations; the scalar criteria in Section 5 are invariant under smooth nonsingular reparameterization.

The derivatives of the pairwise discriminant with respect to the two class means are

$$\nabla_{\boldsymbol{\mu}_k} d_{kl}(\mathbf{y}) = \Sigma_k^{-1}(\mathbf{y} - \boldsymbol{\mu}_k), \quad \nabla_{\boldsymbol{\mu}_l} d_{kl}(\mathbf{y}) = -\Sigma_l^{-1}(\mathbf{y} - \boldsymbol{\mu}_l). \quad (6.10)$$

The corresponding covariance derivatives and their transformation to vech coordinates are given in Supplementary Section S7, while the log-Cholesky parameterization used for numerical optimization is described in Supplementary Section S9.

Using the equivalent log-contrast representation given in Supplementary Section S4, the contribution of the active k -versus- l face to the classification-risk curvature is

$$\mathbf{H}_{kl} = \int_{\mathcal{F}_{kl}} \frac{c_{kl}(\mathbf{s})}{\|\nabla_{\mathbf{y}} d_{kl}(\mathbf{s})\|} \mathbf{a}_{kl}(\mathbf{s}) \mathbf{a}_{kl}(\mathbf{s})^\top dS(\mathbf{s}), \quad (6.11)$$

where $c_{kl}(\mathbf{s}) = r_k(\mathbf{s}) = r_l(\mathbf{s})$ and $\mathbf{a}_{kl}(\mathbf{s}) = \nabla_{\boldsymbol{\theta}} d_{kl}(\mathbf{s})$. For three classes,

$$\mathbf{H}_R = \mathbf{H}_{12} + \mathbf{H}_{13} + \mathbf{H}_{23}. \quad (6.12)$$

This face-wise decomposition later allows classification gains and losses to be attributed to individual pairwise decision boundaries.

6.2 Information and uncertainty-dependent missingness

For QDA, the complete-classification Fisher information has a convenient block structure. For $\boldsymbol{\alpha} = (\alpha_1, \alpha_2)^\top$,

$$\mathbf{I}_\alpha = \begin{pmatrix} \pi_1(1 - \pi_1) & -\pi_1\pi_2 \\ -\pi_1\pi_2 & \pi_2(1 - \pi_2) \end{pmatrix}, \quad (6.13)$$

while the information for the mean and covariance parameters of class k is

$$\mathbf{I}_{\mu_k} = \pi_k \boldsymbol{\Sigma}_k^{-1}, \quad \mathbf{I}_{\Sigma_k} = \frac{\pi_k}{2} \mathbf{D}_p^\top (\boldsymbol{\Sigma}_k^{-1} \otimes \boldsymbol{\Sigma}_k^{-1}) \mathbf{D}_p, \quad (6.14)$$

where $\mathbf{D}_p$ is the duplication matrix. The prior score is orthogonal to the within-class distributional scores, and the mean and covariance scores are orthogonal in expectation within each class. Hence the complete-classification information is block diagonal in these coordinates. Further details are given in Supplementary Section S7.

The general theory does not require a particular missing-label mechanism. For the numerical analysis, we consider mechanisms driven by posterior classification uncertainty, with

$$\tau_k(\mathbf{y}) = \frac{r_k(\mathbf{y})}{\sum_{l=1}^3 r_l(\mathbf{y})}.$$

Our primary uncertainty measure is normalized Shannon entropy,

$$U_H(\mathbf{y}) = \frac{-\sum_{k=1}^3 \tau_k(\mathbf{y}) \log \tau_k(\mathbf{y})}{\log 3}, \quad 0 \leq U_H(\mathbf{y}) \leq 1, \quad (6.15)$$

with missing-label probability

$$q_H(\mathbf{y}) = \text{expit} \{ \xi_0 + \xi_1 U_H(\mathbf{y}) \}. \quad (6.16)$$

Positive ξ_1 makes labels more likely to be unavailable in regions of greater posterior uncertainty.

To assess sensitivity to the uncertainty functional, we also consider normalized Gini uncertainty,

$$U_G(\mathbf{y}) = \frac{3}{2} \left\{ 1 - \sum_{k=1}^3 \tau_k(\mathbf{y})^2 \right\}, \quad 0 \leq U_G(\mathbf{y}) \leq 1, \quad (6.17)$$

with

$$q_G(\mathbf{y}) = \text{expit} \{ \zeta_0 + \zeta_1 U_G(\mathbf{y}) \}. \quad (6.18)$$

Because both uncertainty measures are normalized to $[0, 1]$, their slope parameters have a more comparable interpretation than under their unscaled forms.

For a specified marginal missing-label proportion γ , the intercept in either mechanism is calibrated so that

$$E\{q(\mathbf{Y})\} = \gamma. \quad (6.19)$$

Changing the slope therefore redistributes the unavailable labels across feature space while preserving their expected proportion, corresponding directly to the fixed- γ comparison in Proposition 5.2.

For a generic uncertainty functional $U_{\theta}(\mathbf{y})$,

$$\nabla_{\theta}\tau_k(\mathbf{y}) = \tau_k(\mathbf{y}) \{ \mathbf{s}_k(\mathbf{y}) - \bar{\mathbf{s}}(\mathbf{y}) \}, \quad (6.20)$$

where $\mathbf{s}_k$ and $\bar{\mathbf{s}}$ are defined in Section 2. These derivatives enter the missingness-information blocks in Theorem 2.1; explicit entropy and Gini derivatives are given in Supplementary Section S7.

6.3 Population criteria for the numerical analysis

For a specified QDA configuration and missing-label mechanism, the population comparison requires the complete-classification information $\mathbf{A} = \mathbf{I}_{CC}$, the partially classified information

$$\mathbf{J} = \mathbf{A} - \mathbf{D} + \mathbf{I}_M^{\text{eff}}, \quad \mathbf{D} = E \left[q(\mathbf{Y}) \mathbf{I}_{Z|Y}(\mathbf{Y}) \right], \quad (6.21)$$

and the classification-risk curvature $\mathbf{H}_R$ from (6.11)–(6.12). Here

$$\mathbf{I}_M^{\text{eff}} = \mathbf{B}_{\theta\theta} - \mathbf{B}_{\theta\xi} \mathbf{B}_{\xi\xi}^{-1} \mathbf{B}_{\xi\theta},$$

as in Theorem 2.1.

The corresponding leading excess-risk coefficients are

$$\mathcal{E}_{CC} = \text{tr}(\mathbf{H}_R \mathbf{A}^{-1}), \quad \mathcal{E}_{PC} = \text{tr}(\mathbf{H}_R \mathbf{J}^{-1}), \quad (6.22)$$

with

$$\Delta_R = \mathcal{E}_{CC} - \mathcal{E}_{PC}, \quad \text{ARE}_R = \frac{\mathcal{E}_{CC}}{\mathcal{E}_{PC}}. \quad (6.23)$$

For $p = 2$, each $\mathbf{H}_{kl}$ is a one-dimensional integral along the active portion of the corresponding quadratic boundary. These integrals are evaluated numerically after identifying the active contour segments. Population expectations entering $\mathbf{D}$ and $\mathbf{I}_M^{\text{eff}}$ are evaluated independently using deterministic quadrature or high-accuracy Monte Carlo integration. Numerical implementation and validation details are provided in Supplementary Section S9.

7 Numerical investigation

We use the three-class QDA model to examine the implications of the information–risk theory developed above. The numerical analysis has three main objectives: to verify the regular multiclass geometry required by Theorem 4.2 in a nontrivial unequal-covariance setting; to examine the transition between unfavourable and favourable informative missingness as the amount and location of missing labels vary; and to assess the sensitivity of this transition to class separation, prior imbalance, covariance heterogeneity, and the choice of uncertainty functional.

All quantities in this section are population quantities. In particular, $\mathbf{A}$, $\mathbf{D}$, $\mathbf{I}_M^{\text{eff}}$, $\mathbf{J}$, and $\mathbf{H}_R$ are evaluated at the generating parameter rather than estimated from finite samples.

7.1 Reference configuration and baseline comparison

The reference model has class probabilities

$$(\pi_1, \pi_2, \pi_3) = (0.35, 0.35, 0.30),$$

means

$$\boldsymbol{\mu}_1 = \begin{pmatrix} -1.5 \\ 0 \end{pmatrix}, \quad \boldsymbol{\mu}_2 = \begin{pmatrix} 1.5 \\ 0 \end{pmatrix}, \quad \boldsymbol{\mu}_3 = \begin{pmatrix} 0 \\ 2 \end{pmatrix}, \quad (7.1)$$

and covariance matrices

$$\boldsymbol{\Sigma}_1 = \begin{pmatrix} 1 & 0.30 \\ 0.30 & 0.70 \end{pmatrix}, \quad \boldsymbol{\Sigma}_2 = \begin{pmatrix} 0.80 & -0.20 \\ -0.20 & 1.20 \end{pmatrix}, \quad \boldsymbol{\Sigma}_3 = \begin{pmatrix} 1.30 & 0.40 \\ 0.40 & 0.80 \end{pmatrix}. \quad (7.2)$$

All three covariance matrices are positive definite, and each class has a nonempty Bayes region.

The corresponding Bayes geometry contains a genuine three-class junction. Solving $d_{12}(\mathbf{y}) = d_{13}(\mathbf{y}) = 0$ gives

$$\mathbf{y}^\dagger \approx (0.149, 0.824)^\top. \quad (7.3)$$

At this point, $r_1(\mathbf{y}^\dagger) = r_2(\mathbf{y}^\dagger) = r_3(\mathbf{y}^\dagger)$. The associated spatial gradients are

$$\nabla_{\mathbf{y}} d_{12}(\mathbf{y}^\dagger) \approx (-3.070, -0.117)^\top, \quad \nabla_{\mathbf{y}} d_{13}(\mathbf{y}^\dagger) \approx (-0.817, -2.345)^\top.$$

Their absolute determinant is approximately 7.10, confirming that the junction is transversal rather than tangential and hence satisfies (6.8). The reference configuration therefore provides a genuinely multiclass setting with unequal covariance matrices, curved active Bayes boundaries, and a regular three-class junction.

The classification-risk curvature matrix was obtained by numerical integration over the three active Bayes faces. Under complete classification, the leading excess-risk coefficient is

$$\mathcal{E}_{\text{CC}} = \text{tr}(\mathbf{H}_R \mathbf{A}^{-1}) = 1.5581, \quad (7.4)$$

so that $E\{R(\hat{\boldsymbol{\theta}}_{\text{CC}})\} - R^* \approx 0.7791/n$ to first order.

We next consider entropy-dependent missingness with marginal missing-label proportion $\gamma = 0.30$ and slope $t_H = 4 \log 3$. Calibrating the intercept so that $E\{q_H(\mathbf{Y})\} = 0.30$ gives $\xi_0 = -2.5583$. The corresponding population coefficients are

$$\mathcal{E}_{\text{CC}} = 1.5581, \quad \mathcal{E}_{\text{MCAR}} = 2.0541, \quad \mathcal{E}_{\text{PC}} = 1.3645,$$

which yield

$$\text{ARE}_{R,\text{PC:CC}} = 1.1419, \quad \text{ARE}_{R,\text{MCAR:CC}} = 0.7586,$$

and

$$\Delta_R = \mathcal{E}_{\text{CC}} - \mathcal{E}_{\text{PC}} = 0.1936.$$

Thus the same 30% marginal rate of unavailable labels is detrimental under MCAR but favourable when missingness depends sufficiently strongly on classification uncertainty and the mechanism is incorporated into the likelihood. Numerical integration and reparameterization checks are reported in Supplementary Section S9.

Table 1: Critical normalized entropy slope $t_H^*(\gamma)$ for the reference QDA configuration. For $t_H < t_H^*$, complete classification has the smaller leading excess risk; for $t_H > t_H^*$, informative partial classification is favourable.

Missing-label proportion γ	Critical slope t_H^*
0.10	2.90
0.20	3.17
0.30	3.52
0.40	4.03
0.50	4.87
0.60	6.63

7.2 Phase transition under uncertainty-dependent missingness

To examine how this favourable regime emerges, we vary the strength of dependence on normalized entropy while holding the marginal missing-label proportion fixed. For each pair (t_H, γ) , the intercept in

$$q_H(\mathbf{y}) = \text{expit} \{ \xi_0 + t_H U_H(\mathbf{y}) \}$$

is recalibrated so that $E\{q_H(\mathbf{Y})\} = \gamma$.

At $t_H = 0$, the mechanism is MCAR and, consistently with Corollary 2.2, $\text{ARE}_R < 1$ for every positive missing-label rate considered. As t_H increases, missing labels become increasingly concentrated in regions of high posterior uncertainty and the relative efficiency eventually crosses one. We define the critical slope by

$$t_H^*(\gamma) = \inf \{ t_H : \Delta_R(\gamma, t_H) > 0 \}. \quad (7.5)$$

The resulting values are reported in Table 1.

The critical slope increases monotonically over the range considered and rises more rapidly at larger missing-label proportions. Stronger dependence between missingness and classification uncertainty is therefore required to offset the greater loss of class-label information as γ increases.

Figure 1 displays the corresponding phase boundary. The numerically determined critical slopes satisfy $\Delta_R = 0$, equivalently $\text{ARE}_R = 1$, and separate the unfavourable and favourable regimes over the range considered.




Figure 1: Population phase boundary for the reference QDA model under normalized entropy-dependent missingness. Points denote the numerically determined critical slopes satisfying $\Delta_R = 0$, equivalently $ARE_R = 1$; the connecting curve displays the evolution of the boundary over the range considered.

This phase structure is consistent with Proposition 5.2. Near MCAR, changing the uncertainty slope first redistributes the loss of class-label information, whereas the efficient information supplied by the missingness indicators enters only at second order. The dependence must therefore become sufficiently strong before the latter contribution can offset the former.

7.3 Sensitivity to model geometry and uncertainty specification

We next examine how the favourable regime changes when the geometry of the classification problem is altered. Three perturbations are considered: class separation, prior imbalance, and covariance heterogeneity. Throughout these comparisons, unless otherwise stated, the missing-label proportion is $\gamma = 0.30$ and the normalized entropy slope is $t_H = 4 \log 3$.

Class overlap. We first scale the reference means according to

$$\boldsymbol{\mu}_k(s) = s \boldsymbol{\mu}_k^{(0)}, \quad s > 0, \quad (7.6)$$

while holding the covariance matrices and class probabilities fixed. Values $s < 1$ increase overlap, whereas $s > 1$ increase separation.

Under severe overlap ($s = 0.40$), the missingness pattern does not compensate for the information carried by the unavailable labels. The comparison becomes favourable between $s = 0.40$ and $s = 0.50$, reaches its largest relative advantage near the reference configuration,

Table 2: Population classification efficiency along the class-separation path.

s	$\mathcal{E}_{\text{CC}}$	$\mathcal{E}_{\text{PC}}$	ARE_R
0.40	4.121	4.313	0.955
0.50	3.083	3.039	1.014
0.60	2.523	2.379	1.061
0.75	2.042	1.835	1.113
0.90	1.722	1.511	1.140
1.00	1.558	1.365	1.142
1.15	1.352	1.196	1.130
1.30	1.162	1.046	1.111
1.50	0.926	0.847	1.093
1.70	0.698	0.652	1.071
2.00	0.406	0.386	1.051

and then gradually weakens as the classes become more separated. This nonmonotone behaviour reflects the competing information sources. Under strong overlap, observing the true class membership is highly informative, so removing labels is costly. At moderate separation, posterior uncertainty remains concentrated around active decision boundaries, allowing the missingness pattern to provide information in classification-relevant directions. Under strong separation, both experiments have increasingly small excess-risk coefficients, and the relative advantage diminishes.

Prior imbalance. We next vary the class-3 prior according to

$$\pi_1(\omega) = \pi_2(\omega) = \frac{1 - \omega}{2}, \quad \pi_3(\omega) = \omega, \quad (7.7)$$

while retaining the reference means, covariance matrices, missing-label rate, and uncertainty slope.

Informative partial classification is favourable over a broad range of class probabilities but ceases to be favourable when class 3 becomes sufficiently rare or dominant. The largest relative advantage occurs for moderately balanced priors. The face-specific calculations clarify the reversal under strong imbalance. At $\pi_3 = 0.05$, the relative efficiencies associated with faces 12, 13, and 23 are approximately 1.040, 0.968, and 0.856, respectively. Thus the gain associated with the 12 boundary is insufficient to compensate for losses on the two boundaries involving the rare third class. At the reference prior $\pi_3 = 0.30$, the corresponding efficiencies are 1.134, 1.139, and 1.151, so all three active faces contribute favourably. Additional face-specific results are reported in Supplementary Section S8.

Covariance heterogeneity. To examine the role of nonlinear boundary geometry, define

$$\bar{\Sigma} = \sum_{k=1}^3 \pi_k \Sigma_k^{(0)}$$

and

$$\Sigma_k(\rho) = (1 - \rho)\bar{\Sigma} + \rho\Sigma_k^{(0)}, \quad 0 \leq \rho \leq 1,$$

Table 3: Population classification efficiency as the prior probability of class 3 varies.

π_3	$\mathcal{E}_{\text{CC}}$	$\mathcal{E}_{\text{PC}}$	ARE_R
0.03	1.634	1.862	0.877
0.05	1.633	1.742	0.937
0.08	1.620	1.602	1.011
0.10	1.615	1.546	1.044
0.20	1.578	1.407	1.121
0.30	1.558	1.365	1.142
0.40	1.542	1.345	1.147
0.50	1.529	1.338	1.143
0.60	1.513	1.338	1.131
0.70	1.487	1.343	1.107
0.80	1.435	1.354	1.060
0.90	1.322	1.367	0.968

where $\Sigma_k^{(0)}$ denotes the reference covariance matrix. Thus $\rho = 0$ gives a common-covariance model with linear Bayes boundaries, whereas $\rho = 1$ recovers the reference QDA model. Detailed results for this path are reported in Supplementary Table S1. The classification advantage is stable along the path: the relative efficiency decreases only modestly from 1.160 under common covariance matrices to 1.142 in the reference QDA model. Thus, within this controlled family, the favourable regime persists as the Bayes geometry changes continuously from linear to curved. The ignored-mechanism coefficient remains substantially larger than both the complete-classification and correctly modelled partial-classification coefficients throughout the path.

Taken together, the three geometric perturbations show that the classification advantage is sensitive to the underlying Bayes geometry but not uniformly so. Within the ranges examined here, class overlap and prior imbalance can change the sign of Δ_R , whereas covariance heterogeneity primarily changes its magnitude. Additional face-specific and geometry-dependent phase-boundary calculations are reported in Supplementary Section S8.

Finally, replacing normalized Shannon entropy with normalized Gini uncertainty produces the same qualitative phase structure. For both uncertainty measures, the critical slope increases with the marginal missing-label proportion. In the reference configuration, normalized Gini uncertainty reaches the favourable region at a smaller slope than normalized Shannon entropy for each value of γ examined. This comparison is a robustness check rather than an ordering of the two uncertainty measures; it shows that the phase-transition phenomenon is not specific to entropy-based missingness. Detailed values are reported in Supplementary Table S4.

7.4 Directional information anatomy

We conclude the population investigation by using the spectral representation in Theorem 5.1 to examine why the reference informative mechanism is favourable. Recall that

$$\Delta_R = \sum_j w_j \frac{\lambda_j - 1}{\lambda_j}.$$

The generalized eigenvalues are not uniformly larger than one. Hence the partially classified experiment does not dominate complete classification in the Loewner order. Instead, several directions with $\lambda_j > 1$ also carry substantial classification weights w_j , and their combined contribution outweighs the losses in directions with $\lambda_j < 1$.

The four largest positive directional contributions are approximately

$$0.0877, \quad 0.0571, \quad 0.0536, \quad 0.0352,$$

whereas the two largest negative contributions are approximately

$$-0.0313 \quad \text{and} \quad -0.0303.$$

Summing over all generalized directions gives

$$\sum_j w_j \frac{\lambda_j - 1}{\lambda_j} = 0.193601,$$

which agrees to numerical precision with the direct calculation

$$\Delta_R = \mathcal{E}_{\text{CC}} - \mathcal{E}_{\text{PC}} = 0.193601.$$

This decomposition provides the numerical counterpart of Theorem 5.1. The favourable result does not arise from a uniform increase in Fisher information: some parameter directions lose information. The improvement occurs because the gains are concentrated sufficiently strongly in directions that receive substantial classification-risk weight.

8 Finite-sample validation

We next examine whether the population information–risk comparison is reflected in finite samples. The purpose is not to compare alternative classification procedures broadly, but to assess three implications of the theory: whether the empirical covariance of the estimators approaches its information-based limit, whether the excess classification risk exhibits the predicted n^{-1} behaviour, and whether the population criterion correctly predicts the relative classification performance of complete and partially classified samples. In this section, IPC denotes informative partial classification.

8.1 Simulation design and estimation

Data are generated from the reference three-class QDA configuration of Section 7.1. For each observation,

$$Z_i \sim \text{Multinomial}(1; \pi_1, \pi_2, \pi_3), \quad \mathbf{Y}_i \mid Z_i = k \sim N_2(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k).$$

We consider $n \in \{250, 500, 1000\}$, with $B = 500$ Monte Carlo replications at each sample size.

For every generated complete sample, three observation experiments are constructed. Under complete classification (CC), all class labels are observed. Under MCAR, each label is independently unavailable with probability $\gamma = 0.30$. The third experiment uses informative partial classification with

$$q(\mathbf{y}) = \text{expit} \{ \xi_0 + t_H U_H(\mathbf{y}) \}, \quad (8.1)$$

where $\gamma = 0.30$, $t_H = 4 \log 3$, and U_H is normalized posterior entropy. The intercept is calibrated under the generating model so that $E\{q(\mathbf{Y})\} = 0.30$, giving $\xi_0 = -2.5583$. Thus MCAR and IPC have the same marginal proportion of unavailable labels but differ in their location: under IPC, missing labels are concentrated more strongly in regions of high posterior uncertainty.

The population coefficients for this configuration are

$$\mathcal{K}_{\text{CC}} = 1.5581, \quad \mathcal{K}_{\text{MCAR}} = 2.0541, \quad \mathcal{K}_{\text{IPC}} = 1.3645,$$

corresponding to $\text{ARE}_{R, \text{IPC:CC}} = 1.1419$ and $\text{ARE}_{R, \text{MCAR:CC}} = 0.7586$. The theory therefore predicts a classification advantage for IPC and an efficiency loss under MCAR.

The same generated complete sample is used to construct all three experiments within each replication, reducing Monte Carlo variation in their comparison. Under IPC, the missingness parameters $\boldsymbol{\xi} = (\xi_0, t_H)^\top$ are estimated jointly with $\boldsymbol{\theta}$; the population calibration $E\{q(\mathbf{Y})\} = \gamma$ is used only to define the generating experiment and is not imposed during estimation. Covariance matrices are parameterized through lower log-Cholesky factors to ensure positive definiteness. Numerical optimization uses multiple starting values and standard convergence and admissibility checks. Further computational details are provided in Supplementary Section S9.

8.2 Evaluation criteria

For an estimate $\hat{\boldsymbol{\theta}}$, let $\hat{C}(\mathbf{y}) = \arg \max_k r_k(\mathbf{y}; \hat{\boldsymbol{\theta}})$ denote the corresponding plug-in Bayes classifier. We evaluate its excess population risk using

$$\mathcal{X}(\hat{\boldsymbol{\theta}}) = E_{\boldsymbol{\theta}_0} \left[\max_{1 \leq k \leq 3} \tau_{0k}(\mathbf{Y}) - \tau_{0, \hat{C}(\mathbf{Y})}(\mathbf{Y}) \right], \quad (8.2)$$

where τ_{0k} is the posterior class probability under the generating parameter. This criterion evaluates the fitted decision rule under the true population distribution rather than through empirical error on a finite test sample.

The expectation in (8.2) is evaluated using a common stratified Monte Carlo population sample of size 600,000, with 200,000 observations generated conditionally from each class and class-specific averages weighted by the true mixing proportions. The same population sample is used for every fitted classifier.

For method m , define the scaled excess-risk estimate

$$\hat{\mathcal{K}}_{m,n} = 2n \frac{1}{B_m} \sum_{b \in \mathcal{C}_m} \mathcal{X} \left(\hat{\boldsymbol{\theta}}_m^{(b)} \right), \quad (8.3)$$

where $\mathcal{C}_m$ is the set of usable replications and $B_m = |\mathcal{C}_m|$. The theory shows that

$$\hat{\mathcal{K}}_{m,n} \longrightarrow \mathcal{K}_m = \text{tr}(\mathbf{H}_R \mathbf{V}_m), \quad (8.4)$$

where $\mathbf{V}_{\text{CC}} = \mathbf{A}^{-1}$, $\mathbf{V}_{\text{IPC}} = \mathbf{J}^{-1}$, and $\mathbf{V}_{\text{MCAR}}$ is obtained from the corresponding partially classified information matrix.

As an independent covariance check, let

$$\hat{\mathbf{V}}_{m,n} = n \widehat{\text{Cov}}(\hat{\boldsymbol{\theta}}_m).$$

We assess the classification-relevant discrepancy through

$$D_{R,m}(n) = \frac{\left| \text{tr}(\mathbf{H}_R \hat{\mathbf{V}}_{m,n}) - \text{tr}(\mathbf{H}_R \mathbf{V}_m) \right|}{\text{tr}(\mathbf{H}_R \mathbf{V}_m)}. \quad (8.5)$$

8.3 Finite-sample results

Table 4 reports the scaled excess classification risks. The theoretical ordering is

$$\mathcal{K}_{\text{IPC}} < \mathcal{K}_{\text{CC}} < \mathcal{K}_{\text{MCAR}}, \quad (8.6)$$

and this ordering is reproduced at every sample size considered.

The finite-sample results approach the population predictions as the sample size increases. At $n = 1000$, the scaled risks are 1.5247, 2.0294, and 1.3682 for CC, MCAR, and IPC, respectively, compared with theoretical limits 1.5581, 2.0541, and 1.3645. In particular, the IPC result differs from its first-order limit by only about 0.004.

The corresponding finite-sample IPC-to-CC relative efficiencies are approximately 1.052, 1.034, and 1.114 for $n = 250, 500$, and 1000, respectively, compared with the population value 1.142. Thus the favourable population comparison is already clearly visible at the largest sample size.

Table 4: Finite-sample scaled excess classification risk. Values are Monte Carlo averages over usable fits of $2n\{R(\hat{\boldsymbol{\theta}}) - R^*\}$, with Monte Carlo standard errors in parentheses. The final column gives the corresponding asymptotic coefficient $\mathcal{K}_m = \text{tr}(\mathbf{H}_R \mathbf{V}_m)$.

n	Method	Usable fits	$2n\{R(\hat{\boldsymbol{\theta}}) - R^*\}$	$\mathcal{K}_m$
250	CC	500	1.5756 (0.0466)	1.5581
	MCAR	500	2.2367 (0.0697)	2.0541
	IPC	463	1.4971 (0.0497)	1.3645
500	CC	500	1.5008 (0.0416)	1.5581
	MCAR	500	2.0217 (0.0583)	2.0541
	IPC	485	1.4521 (0.0416)	1.3645
1000	CC	500	1.5247 (0.0424)	1.5581
	MCAR	499	2.0294 (0.0559)	2.0541
	IPC	491	1.3682 (0.0398)	1.3645




Figure 2: Finite-sample scaled excess classification risk for complete classification (CC), missing completely at random (MCAR), and informative partial classification (IPC). Points show Monte Carlo averages of $2n\{R(\hat{\boldsymbol{\theta}}) - R^*\}$, with error bars representing one Monte Carlo standard error across usable fits. Dashed horizontal lines show the corresponding asymptotic coefficients $\text{tr}(\mathbf{H}_R \mathbf{V}_m)$.

Numerical convergence also improves with sample size. All CC fits were usable, while the IPC convergence rates were 92.6%, 97.0%, and 98.2% for $n = 250, 500$, and 1000 , respectively.

The covariance calculation provides an independent validation of the information approximation; detailed covariance diagnostics are reported in Supplementary Table S5. At

$n = 1000$, the classification-weighted covariance discrepancies are approximately 1.4%, 0.5%, and 0.7% for CC, MCAR, and IPC, respectively. A further check based directly on the quadratic risk approximation gives 1.5341, 2.0407, and 1.3735, respectively. These values are close both to the directly evaluated scaled risks 1.5247, 2.0294, and 1.3682, and to their theoretical limits. The agreement among direct population risk, the local quadratic approximation, and the information-based covariance calculation provides a numerical validation of the mechanism underlying (8.4).

The paired comparisons require a different interpretation. At $n = 1000$, IPC has smaller excess risk than CC in approximately 55.6% of paired usable replications, while the average paired difference favours IPC. This is not inconsistent with the theory, which concerns expected excess classification risk rather than the probability that IPC outperforms CC in each individual sample.

Overall, the finite-sample experiment reproduces the central population prediction. Removing 30% of labels under MCAR increases classification risk, whereas under the informative mechanism the missingness indicators provide additional information about classifier-relevant parameters. In the configuration considered here, this contribution is sufficiently aligned with the active Bayes boundary that IPC has smaller expected excess classification risk than complete classification, both asymptotically and in the finite samples examined.

9 Application to the Vertebral Column data

To examine how informative label missingness and explicit modelling of its mechanism behave in a real multiclass classification geometry, we considered the three-class Vertebral Column data set from the UCI Machine Learning Repository (Barreto and Neto, 2005). The data contain 310 observations classified as disk hernia (DH), spondylolisthesis (SL), or normal (NO), with six continuous biomechanical measurements. Because the original data are fully labelled, the observed feature vectors and class labels were retained, while label availability was generated according to the uncertainty-dependent mechanism considered in the preceding sections. The analysis is therefore semi-synthetic: the classification problem is observed, whereas the missing-label process is imposed.

The QDA implementation in Section 6 is two-dimensional. Variable selection was therefore carried out within each outer training sample to avoid selecting the most favourable pair using observations on which predictive performance was subsequently evaluated. Specifically, for every outer cross-validation fold, all 15 pairs formed from the six biomechanical variables were compared by an inner five-fold stratified cross-validation, and the pair with the largest complete-classification accuracy was retained. The purpose of this step was to define a common two-dimensional classification representation before generating the semi-synthetic missing-label indicators. Feature standardization, variable selection, generation of the missing-label indicators, and model fitting were all performed using the outer training sample only. Across the 50 outer folds of the main experiment, sacral slope together with degree of spondylolisthesis was selected in 72% of the folds, pelvic radius together with degree of spondylolisthesis in 18%, and lumbar lordosis angle together with degree of spondylolisthesis in the remaining 10%.

For a feature vector y , let

$$U_H(y; \theta) = -\frac{1}{\log 3} \sum_{k=1}^3 \tau_k(y; \theta) \log\{\tau_k(y; \theta)\}$$

denote normalized classification entropy. Within each outer training sample, a complete-classification QDA fit with estimate $\hat{\theta}_{CC}$ was first obtained and used only to construct the semi-synthetic missing-label mechanism. Writing

$$U_H^{\text{ref}}(y) = U_H(y; \hat{\theta}_{CC}),$$

labels were made unavailable according to

$$\text{logit}\{q_{\text{gen}}(y)\} = \xi_0 + \xi_1 U_H^{\text{ref}}(y), \quad (9.1)$$

where ξ_0 was calibrated within the outer training sample so that the average missing-label probability equalled a prespecified value γ . The main setting used $\gamma = 0.30$ and $\xi_1 = 4 \log 3$, matching the reference mechanism used in the population and finite-sample investigations. The complete-classification estimate was used only to generate the missing-label indicators. In the IPC fit, the classification and missingness parameters were subsequently estimated jointly.

Four procedures were compared. CC used all training labels. MCAR removed labels independently of the features at the same nominal missing-label proportion and treated the missingness mechanism as ignorable. IG used the informatively incomplete training sample but ignored the missingness mechanism. IPC used exactly the same informatively incomplete training sample as IG while jointly modelling the classification model and the uncertainty-dependent missing-label mechanism. Predictive performance was evaluated on the untouched outer test folds. The main analysis used ten repetitions of five-fold stratified cross-validation. All four procedures converged in all 50 outer fits.

Table 5 summarizes the resulting predictive performance. Complete classification had the smallest mean misclassification rate, 0.2119, followed by MCAR, IPC, and IG. IPC therefore did not improve the 0–1 error over CC in the reference setting. The comparison between IG and IPC is more directly informative about the value of modelling the missing-label mechanism because the two procedures use the same informatively incomplete training samples. Relative to IG, IPC reduced the mean log loss from 0.5185 to 0.5015 and the mean Brier score from 0.2899 to 0.2822, whereas the reduction in mean misclassification error was much smaller, from 0.2248 to 0.2235.

At the repeat level, the mean paired differences $R_{IG} - R_{IPC}$ were 0.0013 for misclassification error, 0.0171 for log loss, and 0.0077 for Brier score. Thus, under the reference mechanism, explicit modelling of the informative missing-label process had a clearer effect on probabilistic prediction than on the resulting 0–1 decision rule.

The generated missing-label pattern was clearly uncertainty-selective. Under the reference mechanism, the empirical missing-label proportion was 0.300 on average. The mean normalized entropy among observations whose labels were unavailable was 0.641, compared

Table 5: Repeated nested cross-validation results for the Vertebral Column data under the reference missingness mechanism $\gamma = 0.30$ and $\xi_1 = 4 \log 3$. Values are averages over ten repeated five-fold outer cross-validations.

Method	Error rate	Log loss	Brier score
CC	0.2119	0.4803	0.2749
MCAR	0.2152	0.5064	0.2813
IG	0.2248	0.5185	0.2899
IPC	0.2235	0.5015	0.2822

with 0.249 among labelled observations, corresponding to an average fold-level difference of approximately 0.392. Thus the mechanism systematically concentrated missing labels in regions of substantially greater classification uncertainty. The average fitted slope of the missingness model was 4.86, compared with the generating value $4 \log 3 \simeq 4.39$.

To examine the role of informativeness more directly, we repeated the analysis over the prespecified grid

$$\gamma \in \{0.20, 0.30, 0.40\}, \quad \xi_1 \in \{2 \log 3, 4 \log 3, 6 \log 3\}.$$

Five repeated outer cross-validations were used for each configuration, with the same nested variable-selection procedure. Table 6 reports the CC, IG, and IPC error rates, together with the paired repeat-level difference

$$\Delta_{\text{IG,IPC}} = R_{\text{IG}} - R_{\text{IPC}},$$

so that positive values favour explicit modelling of the informative missing-label mechanism.

Across the sensitivity grid, the benefit of modelling the informative mechanism became more consistent as the dependence of missingness on classification uncertainty strengthened. At $\xi_1 = 2 \log 3$, the paired IPC improvement over IG was small for all three missing-label proportions. At $\xi_1 = 4 \log 3$, IPC again improved on IG throughout the grid, although the magnitude varied with γ . Under the strongest mechanism, $\xi_1 = 6 \log 3$, the paired reductions in error were approximately 0.0194, 0.0185, and 0.0194 for $\gamma = 0.20$, 0.30, and 0.40, respectively. In the corresponding repeat-level comparisons, IPC had lower error than IG in all paired repetitions for $\gamma = 0.20$ and $\gamma = 0.30$, and in four of the five repetitions for $\gamma = 0.40$.

Numerical convergence was stable across the sensitivity analysis. All methods converged in every fit for eight of the nine configurations. The only exception was IPC at $\gamma = 0.30$ and $\xi_1 = 6 \log 3$, for which 24 of the 25 outer fits were usable.

This progression is displayed in Figure 3. The vertical bars summarize repeat-level variation and are intended as descriptive stability intervals rather than independent-sample confidence intervals.

Table 6: Sensitivity to the missing-label proportion γ and informativeness strength ξ_1 . Positive values of $\Delta_{\text{IG,IPC}}$ favour IPC.

γ	$\xi_1/\log 3$	CC	IG	IPC	$\Delta_{\text{IG,IPC}}$
0.20	2	0.2110	0.2129	0.2071	0.0058
0.30	2	0.2110	0.2226	0.2142	0.0084
0.40	2	0.2110	0.2290	0.2284	0.0006
0.20	4	0.2110	0.2200	0.2077	0.0123
0.30	4	0.2110	0.2303	0.2258	0.0045
0.40	4	0.2110	0.2361	0.2194	0.0168
0.20	6	0.2110	0.2213	0.2019	0.0194
0.30	6	0.2110	0.2406	0.2323	0.0185 [†]
0.40	6	0.2110	0.2445	0.2252	0.0194

[†] For $\gamma = 0.30$ and $\xi_1 = 6 \log 3$, one IPC fold did not satisfy the convergence criterion. The displayed IG and IPC values are method-specific averages over complete repetitions, whereas $\Delta_{\text{IG,IPC}}$ is the paired mean difference over the four repetitions complete for both methods. Consequently, for this setting, $\Delta_{\text{IG,IPC}}$ need not equal the difference between the two displayed marginal means.




Figure 3: Paired test-error difference between ignoring (IG) and modelling (IPC) the informative missing-label mechanism in the Vertebral Column application, plotted against the normalized informativeness slope $\xi_1/\log 3$. Positive values favour IPC. Points are mean paired differences across complete cross-validation repetitions; vertical bars are descriptive 95% t -intervals across the repeat-level paired differences and are not interpreted as independent-sample confidence intervals.

The comparison with complete classification remained regime-dependent. IPC did not

uniformly outperform CC across the configurations considered. For example, at $\gamma = 0.20$ and $\xi_1 = 6 \log 3$, the mean IPC error was 0.2019, compared with 0.2110 for CC, whereas CC had the smaller mean error at $\gamma = 0.30$ and $\gamma = 0.40$ under the same uncertainty slope. These finite-sample comparisons are not interpreted as evidence of general dominance of partial over complete classification.

Rather, the semi-synthetic application is consistent with the regime-dependent behavior emphasized by the theoretical development. Weak uncertainty dependence provides only limited benefit from modelling the missing-label process, whereas stronger dependence can make the observed pattern of label availability increasingly useful. At the same time, informative missingness alone does not guarantee that partial classification will outperform complete classification. The population and Monte Carlo investigations in Sections 7 and 8 provide the direct evaluation of the classification-weighted information criterion and its asymptotic risk implications; the present application provides a complementary finite-sample illustration on an observed multiclass classification geometry.

10 Discussion

This paper has examined informative label missingness from the perspective of classification risk rather than parameter estimation alone. The central issue is not simply whether a partially classified experiment contains more or less Fisher information than complete classification, but whether the information that is lost and gained lies in parameter directions that affect the Bayes decision boundary.

The efficient-information decomposition separates two effects of informative label missingness: the loss of conditional class-label information when labels are unavailable and the efficient information carried by the missing-label indicators themselves. This distinction also makes clear that the cost of missing labels depends on where they occur in feature space, rather than only on their marginal frequency. Under missing completely at random, the missingness indicators carry no information about the classification parameters, and only the loss of class-label information remains.

The excess-risk analysis provides the corresponding decision-theoretic interpretation. To second order, classification risk is governed by perturbations of the active pairwise Bayes faces. Pairwise equality surfaces that are not active decision boundaries do not contribute, and, under the regularity and transversality conditions considered here, generic triple and higher-order junctions do not contribute to the leading quadratic term. The resulting matrix $\mathbf{H}_R$ therefore defines a classification-specific weighting of parameter uncertainty: estimation error matters according to the extent to which it perturbs the active Bayes geometry.

Combining these two results shows why a comparison based solely on the Loewner ordering of Fisher information matrices can be too strong for classification. The partially classified experiment may lose information in some directions and gain it in others while nevertheless achieving smaller leading excess classification risk. What matters is whether the gains occur in directions that receive sufficiently large classification weights. The generalized eigenvalue representation makes this alignment between information and decision-boundary geometry explicit. Thus favourable informative missingness is a directional phenomenon rather than a

consequence of global information dominance.

The local analysis around missing completely at random further qualifies this conclusion. With the marginal missing-label proportion held fixed, a small departure from MCAR redistributes the loss of class-label information at first order, whereas the efficient information supplied by the missingness pattern appears only at second order. Informativeness is therefore not automatically beneficial: weak dependence between missingness and classification uncertainty may initially worsen classification, and a favourable regime, when it exists, may arise only after that dependence becomes sufficiently strong.

The three-class QDA calculations illustrate these mechanisms in a setting with curved decision boundaries and a genuine multiclass junction. The numerical results also show that favourable missingness is not universal. Its magnitude is reduced under severe class overlap, can reverse under sufficiently strong class imbalance, and is comparatively stable along the covariance-heterogeneity path examined here. Entropy- and Gini-based missingness mechanisms yield the same qualitative phase structure, while the face-specific and spectral decompositions show that an overall classification advantage can coexist with losses on particular decision faces or information directions.

The finite-sample experiment provides a separate check of the asymptotic predictions. Under the reference configuration, the ordering implied by the population risk coefficients is reproduced over the sample sizes considered. Moreover, the directly evaluated excess risks, their quadratic approximations, and the risk-weighted empirical covariance calculations approach the corresponding theoretical quantities as the sample size increases. This agreement supports the information-geometry explanation of the observed classification advantage, rather than only the sign of a particular numerical comparison.

The real-data application provides a complementary finite-sample illustration. The observed feature vectors and class labels are retained from the Vertebral Column data, while label availability is generated according to the uncertainty-dependent mechanism. Informative label removal was concentrated among observations with substantially greater classification uncertainty, and explicitly modelling the missingness mechanism generally improved prediction relative to treating the same informative pattern as ignorable. The improvement in misclassification error was modest under weak or moderate uncertainty dependence but became appreciable under the strongest mechanism considered. At the same time, informative partial classification did not uniformly outperform complete classification. This behavior is consistent with the theoretical distinction developed above: informativeness alone is not sufficient for favourable partial classification, because the information carried by the missingness indicators must also compensate for the lost class-label information in classifier-relevant directions.

There is no paradox in the possibility that informative partial classification outperforms complete classification under the experiments compared here. The two experiments do not differ only in the number of observed labels. Under informative partial classification, the missing-label indicator is itself an observed random variable whose distribution may depend on the classification parameters. It can therefore contribute information that is absent from the complete-classification experiment, in which the corresponding missingness process is not observed.

The scope of the results is nevertheless limited by the assumptions underlying the analysis. The theory is parametric and relies on regular likelihood asymptotics and sufficiently smooth, transversal Bayes-boundary geometry. Nonregular boundaries, singular models, and high-dimensional regimes in which the parameter dimension grows with the sample size require different arguments. The conditional-independence assumption $M \perp Z \mid \mathbf{Y}$ also excludes mechanisms in which label availability depends directly on the latent class after conditioning on the features; such settings raise additional identification issues. Furthermore, any information gain from the missingness indicators depends on an adequate model for $q(\mathbf{y}; \boldsymbol{\theta}, \boldsymbol{\xi})$. Misspecification may remove that gain and can introduce bias, making sensitivity analysis and robust or semiparametric formulations important extensions.

The present analysis treats the label-missingness mechanism as given. An important extension is to reverse the problem and choose a labelling or abstention mechanism to minimize classification risk subject to a labelling budget. The classification-weighted information criterion developed here provides a natural starting point for such a design problem. Further extensions include high-dimensional classifiers, covariate-dependent class probabilities, alternative classification losses, semiparametric missingness mechanisms, and settings involving multiple forms of incomplete information.

11 Conclusion

The main conclusion is that the statistical value of a partially labelled sample cannot be determined by the proportion of observed labels or by an unweighted comparison of Fisher information matrices alone. When label missingness is informative, classification performance depends on how the resulting information gains and losses align with the parameter directions that perturb the active Bayes decision boundary.

The theory developed here formalizes this principle by linking efficient information under informative partial classification to the local geometry of multiclass excess risk. The numerical investigations and the application to the Vertebral Column data further illustrate that the resulting classification advantage is regime-dependent rather than universal. It explains how partial classification can be unfavourable near MCAR yet become favourable as the missingness mechanism becomes sufficiently informative, and why such an improvement can occur without global information dominance over complete classification. The relevant comparison is therefore not simply how much information is available, but where that information lies relative to the geometry of the classification problem.

Funding declaration

This work was supported by the Australian Research Council [DP230101671].

References

- Barreto, G. and Neto, A. (2005). Vertebral column. *UCI Machine Learning Repository*. doi:10.24432/C5K89B.
- Ahfock, D. and McLachlan, G. J. (2020). An apparent paradox: a classifier based on a partially classified sample may have smaller expected error rate than that if the sample were completely classified. *Statistics and Computing*, 30(6):1779–1790.
- Šaltyte, J. and Dučinskas, K. (2002). Comparison of two estimators of mean function in LDA of spatially correlated Gaussian data. *Mathematical Modelling and Analysis*, 7(1):169–176.
- McLachlan, G. and Ahfock, D. (2019). Estimation of classification rules from partially classified data. In *Conference of the International Federation of Classification Societies*, pages 149–157. Springer.
- Wu, J., Wang, Y.-G., and McLachlan, G. J. (2026). Informative missingness and its implications in semi-supervised learning. *The Innovation Informatics*, 2(2):100033. doi:10.59717/j.xinn-inform.2026.100033.
- O’Neill, T. J. (1980). The general distribution of the error rate of a classification procedure with application to logistic regression discrimination. *Journal of the American Statistical Association*, 75(369):154–160.
- Taniguchi, M. (1994). Higher order asymptotic theory for discriminant analysis in exponential families of distributions. *Journal of Multivariate Analysis*, 48(2):169–187.
- Ducinskas, K. and Saltyte, J. (2002). Second-order asymptotic expansion for the risk in classification of curved exponential populations. *Statistics & Probability Letters*, 59(3):271–279.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22.
- McLachlan, G. J. (1975). Iterative reclassification procedure for constructing an asymptotically optimal rule of allocation in discriminant analysis. *Journal of the American Statistical Association*, 70(350):365–369.
- Chapelle, O., Schölkopf, B., and Zien, A., editors (2006). *Semi-Supervised Learning*. MIT Press, Cambridge, MA.
- Van Engelen, J. E. and Hoos, H. H. (2020). A survey on semi-supervised learning. *Machine Learning*, 109(2):373–440.
- Ahfock, D. and McLachlan, G. J. (2023). Semi-supervised learning of classifiers from a statistical perspective: A brief review. *Econometrics and Statistics*, 26:124–138.

- O'Neill, T. J. (1978). Normal discrimination with unclassified observations. *Journal of the American Statistical Association*, 73(364):821–826.
- Ganesalingam, S. and McLachlan, G. J. (1978). The efficiency of a linear discriminant function based on unclassified initial samples. *Biometrika*, 65(3):658–665.
- McLachlan, G. J. and Scot, D. (1995). Asymptotic relative efficiency of the linear discriminant function under partial nonrandom classification of the training data. *Journal of Statistical Computation and Simulation*, 52(4):415–426.
- Castelli, V. and Cover, T. M. (1996). The relative value of labeled and unlabeled samples in pattern recognition with an unknown mixing parameter. *IEEE Transactions on Information Theory*, 42(6):2102–2117.
- Rubin, D. B. (1976). Inference and missing data. *Biometrika*, 63(3):581–592.
- Mealli, F. and Rubin, D. B. (2015). Clarifying missing at random and related definitions, and implications when coupled with exchangeability. *Biometrika*, 102(4):995–1000.
- Lyu, Z. (2024). Analysis of estimating the Bayes rule for Gaussian mixture models with a specified missing-data mechanism. *Computational Statistics*, 39(7):3727–3751. doi:10.1007/s00180-023-01447-0.

S1 Proof of the information decomposition

This section proves Theorem 2.1 and Corollary 2.2.

Proof of Theorem 2.1. For one observation, write $\mathcal{O} = (\mathbf{Y}, M, (1-M)Z)$. Under $M \perp Z \mid \mathbf{Y}$, the observed-data log-likelihood is

$$\begin{aligned} \ell_{\text{obs}}(\boldsymbol{\theta}, \boldsymbol{\xi}) &= \log p_{\boldsymbol{\theta}}(\mathbf{Y}) + (1-M) \log \tau_Z(\mathbf{Y}; \boldsymbol{\theta}) \\ &\quad + M \log q(\mathbf{Y}; \boldsymbol{\theta}, \boldsymbol{\xi}) + (1-M) \log \{1 - q(\mathbf{Y}; \boldsymbol{\theta}, \boldsymbol{\xi})\}. \end{aligned} \quad (\text{S1.1})$$

Let

$$\mathbf{S}_Y = \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}}(\mathbf{Y}), \quad \mathbf{S}_{Z|\mathbf{Y}} = \nabla_{\boldsymbol{\theta}} \log \tau_Z(\mathbf{Y}; \boldsymbol{\theta}),$$

and, for notational simplicity, write $q = q(\mathbf{Y}; \boldsymbol{\theta}, \boldsymbol{\xi})$. The score contributions arising from the Bernoulli missing-label mechanism are

$$\mathbf{S}_{M,\boldsymbol{\theta}} = \frac{M-q}{q(1-q)} q_{\boldsymbol{\theta}}, \quad \mathbf{S}_{M,\boldsymbol{\xi}} = \frac{M-q}{q(1-q)} q_{\boldsymbol{\xi}}, \quad (\text{S1.2})$$

where $q_{\boldsymbol{\theta}} = \partial q / \partial \boldsymbol{\theta}$ and $q_{\boldsymbol{\xi}} = \partial q / \partial \boldsymbol{\xi}$. Hence

$$\mathbf{S}_{\boldsymbol{\theta}} = \mathbf{S}_Y + (1-M) \mathbf{S}_{Z|\mathbf{Y}} + \mathbf{S}_{M,\boldsymbol{\theta}}, \quad \mathbf{S}_{\boldsymbol{\xi}} = \mathbf{S}_{M,\boldsymbol{\xi}}. \quad (\text{S1.3})$$

The conditional class score has mean zero:

$$\begin{aligned} E[\mathbf{S}_{Z|\mathbf{Y}} \mid \mathbf{Y} = \mathbf{y}] &= \sum_{k=1}^g \tau_k(\mathbf{y}; \boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}} \log \tau_k(\mathbf{y}; \boldsymbol{\theta}) \\ &= \sum_{k=1}^g \nabla_{\boldsymbol{\theta}} \tau_k(\mathbf{y}; \boldsymbol{\theta}) = \mathbf{0}. \end{aligned} \quad (\text{S1.4})$$

Also, since $M \mid \mathbf{Y} \sim \text{Bernoulli}(q)$,

$$E(M - q \mid \mathbf{Y}) = 0,$$

and therefore, by (S1.2),

$$E[\mathbf{S}_{M,\boldsymbol{\theta}} \mid \mathbf{Y}] = \mathbf{0}, \quad E[\mathbf{S}_{M,\boldsymbol{\xi}} \mid \mathbf{Y}] = \mathbf{0}. \quad (\text{S1.5})$$

These conditional-mean identities imply the orthogonality needed below. First,

$$E[\mathbf{S}_Y \{(1-M) \mathbf{S}_{Z|\mathbf{Y}}\}^{\top}] = \mathbf{0},$$

because, conditional on $\mathbf{Y}$, M and Z are independent and $E(\mathbf{S}_{Z|\mathbf{Y}} \mid \mathbf{Y}) = \mathbf{0}$. Similarly,

$$E[\mathbf{S}_Y \mathbf{S}_{M,\boldsymbol{\theta}}^{\top}] = E[\mathbf{S}_Y \mathbf{S}_{M,\boldsymbol{\xi}}^{\top}] = \mathbf{0}.$$

Finally, conditioning on $\mathbf{Y}$ and using $M \perp Z \mid \mathbf{Y}$,

$$\begin{aligned} E \left[(1 - M) \mathbf{S}_{Z|\mathbf{Y}} \mathbf{S}_{M,\theta}^\top \mid \mathbf{Y} \right] \\ = E \left[(1 - M) \mathbf{S}_{M,\theta}^\top \mid \mathbf{Y} \right] E \left[\mathbf{S}_{Z|\mathbf{Y}} \mid \mathbf{Y} \right] = \mathbf{0}, \end{aligned}$$

with the same conclusion for $\mathbf{S}_{M,\xi}$. Thus the marginal feature score, the observed-label score, and the missingness score contribute orthogonally to Fisher information.

If all labels were observed, the score for θ would be $\mathbf{S}_{\text{CC}} = \mathbf{S}_Y + \mathbf{S}_{Z|\mathbf{Y}}$. Using (S1.4),

$$\mathbf{I}_{\text{CC}} = \mathbf{I}_Y + E \left[\mathbf{I}_{Z|\mathbf{Y}}(\theta; \mathbf{Y}) \right], \quad \mathbf{I}_Y = E \left[\mathbf{S}_Y \mathbf{S}_Y^\top \right]. \quad (\text{S1.6})$$

Under partial classification, the observed-label score is $(1 - M) \mathbf{S}_{Z|\mathbf{Y}}$. Its conditional Fisher-information contribution is therefore

$$E \left[(1 - M)^2 \mathbf{S}_{Z|\mathbf{Y}} \mathbf{S}_{Z|\mathbf{Y}}^\top \mid \mathbf{Y} \right].$$

Since $(1 - M)^2 = 1 - M$, conditional independence gives

$$\begin{aligned} E \left[(1 - M) \mathbf{S}_{Z|\mathbf{Y}} \mathbf{S}_{Z|\mathbf{Y}}^\top \mid \mathbf{Y} \right] \\ = E(1 - M \mid \mathbf{Y}) E \left[\mathbf{S}_{Z|\mathbf{Y}} \mathbf{S}_{Z|\mathbf{Y}}^\top \mid \mathbf{Y} \right] \\ = (1 - q) \mathbf{I}_{Z|\mathbf{Y}}(\theta; \mathbf{Y}). \end{aligned} \quad (\text{S1.7})$$

Hence the information lost relative to complete classification is

$$\mathbf{D} = E \left[q(\mathbf{Y}; \theta, \xi) \mathbf{I}_{Z|\mathbf{Y}}(\theta; \mathbf{Y}) \right]. \quad (\text{S1.8})$$

It remains to determine the information supplied by the missing-label indicators. From (S1.2) and

$$E \left[(M - q)^2 \mid \mathbf{Y} \right] = q(1 - q),$$

their Fisher-information blocks are

$$\mathbf{B}_{ab} = E \left[\frac{q_a q_b^\top}{q(1 - q)} \right], \quad a, b \in \{\theta, \xi\}. \quad (\text{S1.9})$$

Combining these blocks with the preceding orthogonality relations, the joint observed-data Fisher information for $(\theta^\top, \xi^\top)^\top$ is

$$\mathbf{I}_{\text{obs}} = \begin{pmatrix} \mathbf{I}_{\text{CC}} - \mathbf{D} + \mathbf{B}_{\theta\theta} & \mathbf{B}_{\theta\xi} \\ \mathbf{B}_{\xi\theta} & \mathbf{B}_{\xi\xi} \end{pmatrix}. \quad (\text{S1.10})$$

The efficient Fisher information for θ , after eliminating the missingness-specific nuisance parameter ξ , is the Schur complement of $\mathbf{B}_{\xi\xi}$:

$$\begin{aligned} \mathbf{I}_{\text{PC}}^{\text{eff}} &= \mathbf{I}_{\text{CC}} - \mathbf{D} + \mathbf{B}_{\theta\theta} \\ &\quad - \mathbf{B}_{\theta\xi} \mathbf{B}_{\xi\xi}^{-1} \mathbf{B}_{\xi\theta}. \end{aligned} \quad (\text{S1.11})$$

Therefore, with

$$\mathbf{I}_M^{\text{eff}} = \mathbf{B}_{\theta\theta} - \mathbf{B}_{\theta\xi} \mathbf{B}_{\xi\xi}^{-1} \mathbf{B}_{\xi\theta}, \quad (\text{S1.12})$$

we obtain

$$\mathbf{I}_{\text{PC}}^{\text{eff}} = \mathbf{I}_{\text{CC}} - \mathbf{D} + \mathbf{I}_M^{\text{eff}}. \quad (\text{S1.13})$$

It remains only to verify that $\mathbf{I}_M^{\text{eff}}$ is positive semidefinite. The Bernoulli information matrix

$$\mathbf{B} = \begin{pmatrix} \mathbf{B}_{\theta\theta} & \mathbf{B}_{\theta\xi} \\ \mathbf{B}_{\xi\theta} & \mathbf{B}_{\xi\xi} \end{pmatrix}$$

is the covariance matrix of the joint missingness score

$$(\mathbf{S}_{M,\theta}^\top, \mathbf{S}_{M,\xi}^\top)^\top,$$

and is therefore positive semidefinite. By assumption, $\mathbf{B}_{\xi\xi}$ is nonsingular; being a nonsingular principal submatrix of a positive semidefinite matrix, it is positive definite. Its Schur complement is therefore positive semidefinite, which gives

$$\mathbf{I}_M^{\text{eff}} \succeq \mathbf{0}.$$

This proves Theorem 2.1. □

roof of Corollary 2.2. Under missing completely at random,

$$q(\mathbf{Y}; \boldsymbol{\theta}, \boldsymbol{\xi}) = \gamma, \quad 0 < \gamma < 1,$$

where γ is variation-independent of $\boldsymbol{\theta}$. Hence $q_{\boldsymbol{\theta}} = \mathbf{0}$, and consequently

$$\mathbf{B}_{\theta\theta} = \mathbf{0}, \quad \mathbf{B}_{\theta\xi} = \mathbf{0}.$$

It follows that $\mathbf{I}_M^{\text{eff}} = \mathbf{0}$. Moreover, the label-information loss reduces to

$$\mathbf{D} = \gamma E [\mathbf{I}_{Z|\mathbf{Y}}(\boldsymbol{\theta}; \mathbf{Y})].$$

Substituting these expressions into the information decomposition established above gives

$$\mathbf{I}_{\text{PC}}^{\text{eff}} = \mathbf{I}_{\text{CC}} - \gamma E [\mathbf{I}_{Z|\mathbf{Y}}(\boldsymbol{\theta}; \mathbf{Y})]. \quad (\text{S1.14})$$

Since $\mathbf{I}_{Z|\mathbf{Y}}(\boldsymbol{\theta}; \mathbf{Y})$ is positive semidefinite for every $\mathbf{Y}$, its expectation is positive semidefinite, and therefore

$$\mathbf{I}_{\text{PC}}^{\text{eff}} \preceq \mathbf{I}_{\text{CC}}.$$

This proves Corollary 2.2. □

S2 Proof of the nuisance-parameter information result

This section proves Proposition 3.1.

Proof of Proposition 3.1. Partition the complete-classification information matrix according to $\boldsymbol{\theta} = (\boldsymbol{\beta}^\top, \boldsymbol{\lambda}^\top)^\top$ as

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}_{\beta\beta} & \mathbf{A}_{\beta\lambda} \\ \mathbf{A}_{\lambda\beta} & \mathbf{A}_{\lambda\lambda} \end{pmatrix},$$

and suppose that $\mathbf{A}_{\lambda\lambda}$ is nonsingular. Since $\mathbf{A}$ is a Fisher-information matrix and $\mathbf{A}_{\lambda\lambda}$ is a nonsingular principal block, $\mathbf{A}_{\lambda\lambda}$ is positive definite. Define

$$\mathbf{R} = \mathbf{A}_{\beta\lambda} \mathbf{A}_{\lambda\lambda}^{-1}.$$

If $\mathbf{S}_\beta$ and $\mathbf{S}_\lambda$ are the corresponding complete-classification score components, then the efficient score for $\boldsymbol{\beta}$ relative to $\boldsymbol{\lambda}$ is

$$\mathbf{S}_\beta^{\text{eff}} = \mathbf{S}_\beta - \mathbf{R} \mathbf{S}_\lambda.$$

Indeed,

$$E[\mathbf{S}_\beta^{\text{eff}} \mathbf{S}_\lambda^\top] = \mathbf{A}_{\beta\lambda} - \mathbf{R} \mathbf{A}_{\lambda\lambda} = \mathbf{0},$$

and its covariance is the usual Schur complement

$$\mathbf{A}_{\text{eff}}(\boldsymbol{\beta}) = \mathbf{A}_{\beta\beta} - \mathbf{A}_{\beta\lambda} \mathbf{A}_{\lambda\lambda}^{-1} \mathbf{A}_{\lambda\beta}. \quad (\text{S2.1})$$

After eliminating the missingness-specific parameter $\boldsymbol{\xi}$, Theorem 2.1 gives

$$\mathbf{J} = \mathbf{A} + \mathbf{K}, \quad \mathbf{K} = \mathbf{I}_M^{\text{eff}} - \mathbf{D}.$$

Write

$$\mathbf{K} = \begin{pmatrix} \mathbf{K}_{\beta\beta} & \mathbf{K}_{\beta\lambda} \\ \mathbf{K}_{\lambda\beta} & \mathbf{K}_{\lambda\lambda} \end{pmatrix}$$

and introduce the nonsingular block-triangular matrix

$$\mathbf{T} = \begin{pmatrix} \mathbf{I}_r & -\mathbf{R} \\ \mathbf{0} & \mathbf{I}_s \end{pmatrix}.$$

Multiplication of the score vector by $\mathbf{T}$ replaces $\mathbf{S}_\beta$ by $\mathbf{S}_\beta - \mathbf{R} \mathbf{S}_\lambda$ while leaving $\mathbf{S}_\lambda$ unchanged. The corresponding information matrix is therefore transformed by congruence.

For the complete-classification information,

$$\mathbf{T} \mathbf{A} \mathbf{T}^\top = \begin{pmatrix} \mathbf{A}_{\text{eff}}(\boldsymbol{\beta}) & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\lambda\lambda} \end{pmatrix}. \quad (\text{S2.2})$$

The vanishing off-diagonal block follows from $\mathbf{A}_{\beta\lambda} - \mathbf{R} \mathbf{A}_{\lambda\lambda} = \mathbf{0}$, while the upper-left block is (S2.1).

Applying the same transformation to the perturbation $\mathbf{K}$ gives

$$\mathbf{T} \mathbf{K} \mathbf{T}^\top = \begin{pmatrix} \widetilde{\mathbf{K}}_{\beta\beta} & \widetilde{\mathbf{K}}_{\beta\lambda} \\ \widetilde{\mathbf{K}}_{\lambda\beta} & \mathbf{K}_{\lambda\lambda} \end{pmatrix}, \quad (\text{S2.3})$$

where

$$\widetilde{\mathbf{K}}_{\beta\beta} = \mathbf{K}_{\beta\beta} - \mathbf{R}\mathbf{K}_{\lambda\beta} - \mathbf{K}_{\beta\lambda}\mathbf{R}^\top + \mathbf{R}\mathbf{K}_{\lambda\lambda}\mathbf{R}^\top, \quad (\text{S2.4})$$

$$\widetilde{\mathbf{K}}_{\beta\lambda} = \mathbf{K}_{\beta\lambda} - \mathbf{R}\mathbf{K}_{\lambda\lambda}, \quad (\text{S2.5})$$

and $\widetilde{\mathbf{K}}_{\lambda\beta} = \widetilde{\mathbf{K}}_{\beta\lambda}^\top$ because $\mathbf{K}$ is symmetric.

Combining (S2.2) and (S2.3) yields

$$\mathbf{J}^* := \mathbf{T}\mathbf{J}\mathbf{T}^\top = \begin{pmatrix} \mathbf{A}_{\text{eff}}(\boldsymbol{\beta}) + \widetilde{\mathbf{K}}_{\beta\beta} & \widetilde{\mathbf{K}}_{\beta\lambda} \\ \widetilde{\mathbf{K}}_{\lambda\beta} & \mathbf{J}_{\lambda\lambda} \end{pmatrix}, \quad (\text{S2.6})$$

where

$$\mathbf{J}_{\lambda\lambda} = \mathbf{A}_{\lambda\lambda} + \mathbf{K}_{\lambda\lambda}.$$

Because $\mathbf{T}$ is nonsingular and its transformation replaces the score for $\boldsymbol{\beta}$ by that score minus a linear combination of the nuisance score while leaving the nuisance-score space unchanged, the efficient information for $\boldsymbol{\beta}$ is invariant under this transformation. Equivalently, it is obtained as the Schur complement of the lower-right block of $\mathbf{J}^*$. Therefore

$$\mathbf{J}_{\text{eff}}(\boldsymbol{\beta}) = \mathbf{A}_{\text{eff}}(\boldsymbol{\beta}) + \widetilde{\mathbf{K}}_{\beta\beta} - \widetilde{\mathbf{K}}_{\beta\lambda}\mathbf{J}_{\lambda\lambda}^{-1}\widetilde{\mathbf{K}}_{\lambda\beta}. \quad (\text{S2.7})$$

This is the first assertion of Proposition 3.1.

Now define

$$\mathbf{C}_K = \widetilde{\mathbf{K}}_{\beta\lambda}\mathbf{J}_{\lambda\lambda}^{-1}\widetilde{\mathbf{K}}_{\lambda\beta}. \quad (\text{S2.8})$$

The matrix $\mathbf{J}_{\lambda\lambda}$ is a principal block of the efficient Fisher-information matrix $\mathbf{J}$. Under the assumed nonsingularity, it is therefore positive definite, and so is $\mathbf{J}_{\lambda\lambda}^{-1}$. Consequently, for every $\mathbf{x} \in \mathbb{R}^r$,

$$\begin{aligned} \mathbf{x}^\top \mathbf{C}_K \mathbf{x} &= \left(\widetilde{\mathbf{K}}_{\lambda\beta} \mathbf{x} \right)^\top \mathbf{J}_{\lambda\lambda}^{-1} \left(\widetilde{\mathbf{K}}_{\lambda\beta} \mathbf{x} \right) \\ &\geq 0. \end{aligned} \quad (\text{S2.9})$$

Thus

$$\mathbf{C}_K \succeq \mathbf{0},$$

which proves the second assertion.

For completeness, the sequential nuisance elimination used above is equivalent to eliminating $\boldsymbol{\lambda}$ and $\boldsymbol{\xi}$ jointly from the full observed-data information matrix. If that matrix is partitioned according to $(\boldsymbol{\beta}^\top, \boldsymbol{\lambda}^\top, \boldsymbol{\xi}^\top)^\top$, the quotient identity for Schur complements gives, under the required invertibility conditions,

$$\text{Schur}_{(\lambda, \xi)}(\mathbf{I}_{\text{obs}}) = \text{Schur}_{\lambda} \{ \text{Schur}_{\xi}(\mathbf{I}_{\text{obs}}) \}. \quad (\text{S2.10})$$

Thus eliminating the missingness-specific nuisance parameter first and then the data-model nuisance parameter gives exactly the same efficient information for $\boldsymbol{\beta}$ as their joint elimination.

This completes the proof of Proposition 3.1. $\square$

S3 Proof of the multiclass excess-risk expansion

This section establishes the local boundary calculation underlying Theorem 4.2.

Proof of Theorem 4.2. For an arbitrary deterministic classifier C , evaluated under the true parameter $\boldsymbol{\theta}_0$,

$$R(C) = 1 - \int r_{C(\mathbf{y})}^0(\mathbf{y}) d\mathbf{y},$$

where $r_k^0(\mathbf{y}) = r_k(\mathbf{y}; \boldsymbol{\theta}_0)$. Since the Bayes classifier C_0 maximizes the prior-weighted class density pointwise,

$$R^* = 1 - \int \max_{1 \leq k \leq g} r_k^0(\mathbf{y}) d\mathbf{y}.$$

Consequently,

$$R(C) - R^* = \int \left\{ \max_k r_k^0(\mathbf{y}) - r_{C(\mathbf{y})}^0(\mathbf{y}) \right\} d\mathbf{y}. \quad (\text{S3.1})$$

For the classifier $C_{\mathbf{h}}$ determined by $\boldsymbol{\theta}_0 + \mathbf{h}$, this becomes

$$R(\boldsymbol{\theta}_0 + \mathbf{h}) - R^* = \int \left\{ r_{C_0(\mathbf{y})}^0(\mathbf{y}) - r_{C_{\mathbf{h}}(\mathbf{y})}^0(\mathbf{y}) \right\} d\mathbf{y}, \quad (\text{S3.2})$$

up to the immaterial choice of classifier on sets of Bayes ties of probability zero.

We first localize the region in which $C_{\mathbf{h}}$ can differ from C_0 . Let

$$m(\mathbf{y}) = r_{(1)}^0(\mathbf{y}) - r_{(2)}^0(\mathbf{y}),$$

where $r_{(1)}^0$ and $r_{(2)}^0$ denote the largest and second-largest values among $\{r_1^0(\mathbf{y}), \dots, r_g^0(\mathbf{y})\}$. On any compact set separated from the Bayes boundary, continuity implies $m(\mathbf{y}) \geq c$ for some $c > 0$. Joint differentiability in $(\mathbf{y}, \boldsymbol{\theta})$ gives, uniformly on such a set,

$$\max_k |r_k(\mathbf{y}; \boldsymbol{\theta}_0 + \mathbf{h}) - r_k^0(\mathbf{y})| = O(\|\mathbf{h}\|).$$

It follows that, for sufficiently small $\mathbf{h}$, a change in the winning class is possible only where

$$m(\mathbf{y}) = O(\|\mathbf{h}\|).$$

Thus the disagreement between the true and perturbed classifiers is confined to a shrinking neighborhood of the Bayes boundary.

Consider now an interior point $\mathbf{s}$ of an active face $\mathcal{F}_{kl}$. By definition,

$$r_k^0(\mathbf{s}) = r_l^0(\mathbf{s}) > r_m^0(\mathbf{s}), \quad m \notin \{k, l\}.$$

The strict inequality implies that, in a sufficiently small neighborhood of $\mathbf{s}$, classes other than k and l remain separated from the two leading classes. By continuity in the parameter, this remains true for all sufficiently small $\mathbf{h}$. Hence, locally, both C_0 and $C_{\mathbf{h}}$ are determined solely by the sign of

$$g_{kl}(\mathbf{y}; \boldsymbol{\theta}) = r_k(\mathbf{y}; \boldsymbol{\theta}) - r_l(\mathbf{y}; \boldsymbol{\theta}).$$

Write $g_{kl}^0(\mathbf{y}) = g_{kl}(\mathbf{y}; \boldsymbol{\theta}_0)$. At a regular point of the face, $\nabla_{\mathbf{y}} g_{kl}^0(\mathbf{s}) \neq \mathbf{0}$. On a sufficiently small tubular neighborhood of a compact regular portion of $\mathcal{F}_{kl}$, define

$$\mathbf{V}_{kl}(\mathbf{y}) = \frac{\nabla_{\mathbf{y}} g_{kl}^0(\mathbf{y})}{\|\nabla_{\mathbf{y}} g_{kl}^0(\mathbf{y})\|^2}.$$

Let $\boldsymbol{\Phi}_{kl}(\mathbf{s}, u)$ be the associated local flow, initialized at $\boldsymbol{\Phi}_{kl}(\mathbf{s}, 0) = \mathbf{s}$. Along this flow,

$$\begin{aligned} \frac{\partial}{\partial u} g_{kl}^0 \{ \boldsymbol{\Phi}_{kl}(\mathbf{s}, u) \} &= \nabla_{\mathbf{y}} g_{kl}^0 \{ \boldsymbol{\Phi}_{kl}(\mathbf{s}, u) \}^\top \mathbf{V}_{kl} \{ \boldsymbol{\Phi}_{kl}(\mathbf{s}, u) \} \\ &= 1. \end{aligned}$$

Since $g_{kl}^0(\mathbf{s}) = 0$, it follows that

$$g_{kl}^0 \{ \boldsymbol{\Phi}_{kl}(\mathbf{s}, u) \} = u. \quad (\text{S3.3})$$

Thus u is the pairwise contrast itself and provides a normal coordinate to the true boundary.

The corresponding volume element satisfies, uniformly on compact regular portions of the face,

$$d\mathbf{y} = \left\{ \frac{1}{\|\nabla_{\mathbf{y}} g_{kl}^0(\mathbf{s})\|} + O(|u|) \right\} dS(\mathbf{s}) du. \quad (\text{S3.4})$$

This follows from the coarea formula, or equivalently from the Jacobian of the local tubular coordinate map.

Define

$$\mathbf{b}_{kl}(\mathbf{s}) = \nabla_{\boldsymbol{\theta}} g_{kl}(\mathbf{s}; \boldsymbol{\theta}_0).$$

A Taylor expansion in the parameter, together with smoothness of the coordinate map, gives

$$\begin{aligned} g_{kl} \{ \boldsymbol{\Phi}_{kl}(\mathbf{s}, u); \boldsymbol{\theta}_0 + \mathbf{h} \} &= u + \mathbf{b}_{kl}(\mathbf{s})^\top \mathbf{h} \\ &\quad + O(|u| \|\mathbf{h}\| + \|\mathbf{h}\|^2), \end{aligned} \quad (\text{S3.5})$$

uniformly on compact regular portions of $\mathcal{F}_{kl}$. Let $u_{\mathbf{h}}(\mathbf{s})$ denote the normal coordinate of the perturbed k -versus- l boundary. The implicit-function theorem applied to (S3.5) yields

$$u_{\mathbf{h}}(\mathbf{s}) = -\mathbf{b}_{kl}(\mathbf{s})^\top \mathbf{h} + O(\|\mathbf{h}\|^2), \quad (\text{S3.6})$$

again uniformly on such compact portions.

Within this local binary neighborhood, disagreement between the true and perturbed rules occurs precisely in the strip between the two boundaries. The excess loss at $\mathbf{y}$ is

$$|r_k^0(\mathbf{y}) - r_l^0(\mathbf{y})| = |g_{kl}^0(\mathbf{y})| = |u|,$$

where the last equality follows from (S3.3). Therefore, for a fixed boundary point $\mathbf{s}$, the excess-risk contribution per unit surface measure from the displaced strip is

$$\begin{aligned} \Delta R_{kl, \mathbf{s}}(\mathbf{h}) &= \int_{\min\{0, u_{\mathbf{h}}(\mathbf{s})\}}^{\max\{0, u_{\mathbf{h}}(\mathbf{s})\}} |u| \left\{ \frac{1}{\|\nabla_{\mathbf{y}} g_{kl}^0(\mathbf{s})\|} + O(|u|) \right\} du \\ &= \frac{1}{2} \frac{u_{\mathbf{h}}(\mathbf{s})^2}{\|\nabla_{\mathbf{y}} g_{kl}^0(\mathbf{s})\|} + O(\|\mathbf{h}\|^3). \end{aligned} \quad (\text{S3.7})$$

Using (S3.6),

$$\Delta R_{kl,s}(\mathbf{h}) = \frac{1}{2} \frac{\{\mathbf{b}_{kl}(\mathbf{s})^\top \mathbf{h}\}^2}{\|\nabla_{\mathbf{y}} g_{kl}^0(\mathbf{s})\|} + O(\|\mathbf{h}\|^3). \quad (\text{S3.8})$$

Integrating over any compact regular portion $\mathcal{F}_{kl}^{(0)} \subset \mathcal{F}_{kl}$ that is bounded away from higher-order Bayes tie sets gives

$$\Delta R_{kl}^{(0)}(\mathbf{h}) = \frac{1}{2} \mathbf{h}^\top \mathbf{H}_{kl}^{(0)} \mathbf{h} + o(\|\mathbf{h}\|^2), \quad (\text{S3.9})$$

where

$$\mathbf{H}_{kl}^{(0)} = \int_{\mathcal{F}_{kl}^{(0)}} \frac{\mathbf{b}_{kl}(\mathbf{s}) \mathbf{b}_{kl}(\mathbf{s})^\top}{\|\nabla_{\mathbf{y}} g_{kl}^0(\mathbf{s})\|} dS(\mathbf{s}). \quad (\text{S3.10})$$

It remains to control neighborhoods of higher-order Bayes tie sets and to justify extending the preceding facewise calculation to the whole active boundary. Let $\varepsilon = \|\mathbf{h}\|$. Consider a point at which $r \geq 3$ classes tie at the Bayes maximum. By the transversality assumption, in a neighborhood of such a point one may choose $r - 1$ independent pairwise contrasts as local normal coordinates, say

$$\mathbf{u} = (u_1, \dots, u_{r-1})^\top, \quad u_j = g_{k_1 k_{j+1}}^0(\mathbf{y}).$$

The constant-rank theorem then provides local coordinates $(\mathbf{s}, \mathbf{u})$, where $\mathbf{s}$ parameterizes the tie stratum and the associated Jacobian is bounded above and below on compact coordinate neighborhoods.

A perturbation $\boldsymbol{\theta}_0 + \mathbf{h}$ changes each of these pairwise contrasts by $O(\varepsilon)$, uniformly on such a compact neighborhood. Consequently, a change in the genuinely r -way ordering can occur only when

$$\|\mathbf{u}\| = O(\varepsilon).$$

Since there are $r - 1$ independent normal coordinates, the volume of this region is $O(\varepsilon^{r-1})$ per unit measure of the tie stratum. Moreover, because the competing prior-weighted densities agree on the tie stratum and are continuously differentiable, their differences throughout this region are $O(\varepsilon)$. Hence the excess-risk contribution of the neighborhood of an r -way transversal tie is

$$O(\varepsilon^{r-1}) O(\varepsilon) = O(\varepsilon^r) = o(\varepsilon^2), \quad r \geq 3.$$

The same argument controls the portions of the adjacent active pairwise faces removed when forming regular tubular neighborhoods. For a triple tie, for example, the excluded portion of an adjacent face has surface measure $O(\varepsilon)$, while its displaced-strip contribution per unit surface measure is $O(\varepsilon^2)$, giving $O(\varepsilon^3)$. More generally, the corresponding contribution near an r -way transversal tie is of order $O(\varepsilon^r)$.

Therefore neighborhoods of higher-order tie strata contribute only $o(\varepsilon^2)$. Away from these neighborhoods, the active pairwise faces are regular and the preceding tubular-coordinate calculation applies uniformly on compact portions. Moreover, by transversality and smoothness, the integrand

$$\frac{\mathbf{b}_{kl}(\mathbf{s}) \mathbf{b}_{kl}(\mathbf{s})^\top}{\|\nabla_{\mathbf{y}} g_{kl}^0(\mathbf{s})\|}$$

is locally bounded near each higher-order tie stratum. The portion of an active pairwise face lying within an $O(\varepsilon)$ neighborhood of a triple-tie stratum has surface measure $O(\varepsilon)$, and the corresponding difference between the truncated and full quadratic coefficients is therefore $O(\varepsilon)$. After multiplication by $\|\mathbf{h}\|^2 = \varepsilon^2$, this contributes $O(\varepsilon^3) = o(\varepsilon^2)$. Higher-order tie strata give still smaller orders. Consequently, the truncated face integrals may be replaced by the integrals over the full active faces at quadratic order. Thus, for each active pair $k < l$,

$$\Delta R_{kl}(\mathbf{h}) = \frac{1}{2} \mathbf{h}^\top \mathbf{H}_{kl} \mathbf{h} + o(\|\mathbf{h}\|^2), \quad (\text{S3.11})$$

where

$$\mathbf{H}_{kl} = \int_{\mathcal{F}_{kl}} \frac{\mathbf{b}_{kl}(\mathbf{s}) \mathbf{b}_{kl}(\mathbf{s})^\top}{\|\nabla_{\mathbf{y}} g_{kl}^0(\mathbf{s})\|} dS(\mathbf{s}). \quad (\text{S3.12})$$

Thus only the regular active pairwise faces contribute to the quadratic term.

For compact active boundaries on bounded feature support, summing the preceding face-wise expansions over all active pairs $k < l$ gives

$$R(\theta_0 + h) - R^* = \frac{1}{2} h^\top H_R h + o(\|h\|^2), \quad H_R = \sum_{k < l} H_{kl}.$$

This establishes the quadratic expansion directly in the bounded-support case. Supplementary Section S4 extends the argument to unbounded feature supports under the stated tail condition and, when one or more active faces are noncompact, the additional regular-exhaustion and boundary-integrability conditions.

Finally, for any vector $\mathbf{v}$,

$$\begin{aligned} \mathbf{v}^\top \mathbf{H}_R \mathbf{v} &= \sum_{k < l} \int_{\mathcal{F}_{kl}} \frac{\{\mathbf{b}_{kl}(\mathbf{s})^\top \mathbf{v}\}^2}{\|\nabla_{\mathbf{y}} g_{kl}^0(\mathbf{s})\|} dS(\mathbf{s}) \\ &\geq 0. \end{aligned}$$

Hence $\mathbf{H}_R \succeq \mathbf{0}$, completing the proof of Theorem 4.2. $\square$

S4 Noncompact boundaries and tail conditions

This section states the tail condition used when the feature support is unbounded and the additional regular-exhaustion and boundary-integrability conditions required when one or more active Bayes faces are noncompact. It then completes the corresponding extension of Theorem 4.2.

Let

$$\mathcal{K}_L = \{\mathbf{y} \in \mathcal{Y} : \|\mathbf{y}\| \leq L\}, \quad \mathcal{F}_{kl}^{(L)} = \mathcal{F}_{kl} \cap \mathcal{K}_L.$$

For each active pair $k < l$, define the truncated curvature matrix

$$\mathbf{H}_{kl}^{(L)} = \int_{\mathcal{F}_{kl}^{(L)}} \frac{\mathbf{b}_{kl}(\mathbf{s}) \mathbf{b}_{kl}(\mathbf{s})^\top}{\|\nabla_{\mathbf{y}} g_{kl}(\mathbf{s}; \boldsymbol{\theta}_0)\|} dS(\mathbf{s}), \quad (\text{S4.1})$$

where

$$\mathbf{b}_{kl}(\mathbf{s}) = \nabla_{\boldsymbol{\theta}} g_{kl}(\mathbf{s}; \boldsymbol{\theta}_0).$$

For noncompact active faces, we assume the following regular-exhaustion condition in addition to boundary integrability. Tail control is imposed separately whenever the feature support is unbounded. There exists an increasing sequence $L_m \rightarrow \infty$ such that, for every m , each regular active face intersects $\partial\mathcal{K}_{L_m}$ transversely whenever the intersection is nonempty.

We further assume the boundary-integrability condition

$$\int_{\mathcal{F}_{kl}} \frac{\|\mathbf{b}_{kl}(\mathbf{s})\|^2}{\|\nabla_{\mathbf{y}} g_{kl}(\mathbf{s}; \boldsymbol{\theta}_0)\|} dS(\mathbf{s}) < \infty \quad \text{for every nonempty } \mathcal{F}_{kl}. \quad (\text{S4.2})$$

This condition guarantees that the matrix integral defining $\mathbf{H}_{kl}$ is finite and that

$$\mathbf{H}_{kl}^{(L)} \longrightarrow \mathbf{H}_{kl} \quad \text{as } L \rightarrow \infty.$$

A separate condition is required to control classification disagreement in the tails. Let

$$\Delta_{\mathbf{h}}(\mathbf{y}) = r_{C_0(\mathbf{y})}^0(\mathbf{y}) - r_{C_{\mathbf{h}}(\mathbf{y})}^0(\mathbf{y}) \geq 0$$

denote the pointwise excess loss. We assume

$$\lim_{L \rightarrow \infty} \limsup_{\mathbf{h} \rightarrow \mathbf{0}} \frac{1}{\|\mathbf{h}\|^2} \int_{\mathcal{K}_L^c} \Delta_{\mathbf{h}}(\mathbf{y}) d\mathbf{y} = 0. \quad (\text{S4.3})$$

Thus, after scaling by the quadratic order $\|\mathbf{h}\|^2$, the contribution to excess risk from sufficiently far into the tails is uniformly negligible for small perturbations.

To complete the proof of Theorem 4.2, fix $L < \infty$ such that the active faces intersect $\partial\mathcal{K}_L$ transversely whenever the intersections are nonempty. For the noncompact extension we use an increasing sequence of such regular truncation radii tending to infinity. On $\mathcal{K}_L$, the regular portions of the active faces are compact after excluding arbitrarily small neighborhoods of transversal higher-order junctions and of their intersections with $\partial\mathcal{K}_L$. The latter intersections are codimension two in the ambient feature space. Their $O(\|\mathbf{h}\|)$ neighborhoods within an active face have surface measure $O(\|\mathbf{h}\|)$, while the displaced-strip excess risk per unit surface measure is $O(\|\mathbf{h}\|^2)$. Their total contribution is therefore $O(\|\mathbf{h}\|^3) = o(\|\mathbf{h}\|^2)$. Hence the local argument of Supplementary Section S3 gives

$$\int_{\mathcal{K}_L} \Delta_{\mathbf{h}}(\mathbf{y}) d\mathbf{y} = \frac{1}{2} \mathbf{h}^\top \mathbf{H}_R^{(L)} \mathbf{h} + o_L(\|\mathbf{h}\|^2), \quad (\text{S4.4})$$

where

$$\mathbf{H}_R^{(L)} = \sum_{k < l} \mathbf{H}_{kl}^{(L)}.$$

For each fixed L , the remainder in (S4.4) is $o(\|\mathbf{h}\|^2)$ as $\mathbf{h} \rightarrow \mathbf{0}$.

By (S4.2),

$$\mathbf{H}_R^{(L)} \longrightarrow \mathbf{H}_R = \sum_{k < l} \mathbf{H}_{kl} \quad \text{as } L \rightarrow \infty.$$

Moreover, (S4.3) implies that, for every $\varepsilon > 0$, L can be chosen sufficiently large so that

$$\limsup_{\mathbf{h} \rightarrow \mathbf{0}} \frac{1}{\|\mathbf{h}\|^2} \int_{\mathcal{K}_L^\varepsilon} \Delta_{\mathbf{h}}(\mathbf{y}) d\mathbf{y} < \varepsilon.$$

To make the limiting argument explicit, write

$$\mathcal{R}(\mathbf{h}) = R(\boldsymbol{\theta}_0 + \mathbf{h}) - R^*.$$

Then

$$\begin{aligned} & \frac{|\mathcal{R}(\mathbf{h}) - \frac{1}{2} \mathbf{h}^\top \mathbf{H}_R \mathbf{h}|}{\|\mathbf{h}\|^2} \\ & \leq \frac{\left| \int_{\mathcal{K}_L} \Delta_{\mathbf{h}}(\mathbf{y}) d\mathbf{y} - \frac{1}{2} \mathbf{h}^\top \mathbf{H}_R^{(L)} \mathbf{h} \right|}{\|\mathbf{h}\|^2} \\ & \quad + \frac{1}{\|\mathbf{h}\|^2} \int_{\mathcal{K}_L^\varepsilon} \Delta_{\mathbf{h}}(\mathbf{y}) d\mathbf{y} + \frac{1}{2} \left\| \mathbf{H}_R^{(L)} - \mathbf{H}_R \right\|_{\text{op}}. \end{aligned}$$

For each fixed L , the first term tends to zero as $\mathbf{h} \rightarrow \mathbf{0}$ by (S4.4). The second term can be made arbitrarily small by choosing L sufficiently large, by (S4.3), while the third term tends to zero as $L \rightarrow \infty$ by (S4.2). Hence

$$R(\boldsymbol{\theta}_0 + \mathbf{h}) - R^* = \frac{1}{2} \mathbf{h}^\top \mathbf{H}_R \mathbf{h} + o(\|\mathbf{h}\|^2). \quad (\text{S4.5})$$

Hence the quadratic excess-risk expansion established in Supplementary Section S3 remains valid for noncompact active Bayes boundaries under (S4.2) and (S4.3).

Equivalent log-contrast representation. When the prior-weighted class densities are strictly positive in a neighborhood of an active face, define

$$d_{kl}(\mathbf{y}; \boldsymbol{\theta}) = \log \frac{r_k(\mathbf{y}; \boldsymbol{\theta})}{r_l(\mathbf{y}; \boldsymbol{\theta})}.$$

On $\mathcal{F}_{kl}$, let

$$c_{kl}(\mathbf{s}) = r_k^0(\mathbf{s}) = r_l^0(\mathbf{s}), \quad \mathbf{a}_{kl}(\mathbf{s}) = \nabla_{\boldsymbol{\theta}} d_{kl}(\mathbf{s}; \boldsymbol{\theta}_0).$$

Because

$$g_{kl} = r_k - r_l,$$

differentiation on the active face gives

$$\mathbf{b}_{kl}(\mathbf{s}) = c_{kl}(\mathbf{s}) \mathbf{a}_{kl}(\mathbf{s})$$

and

$$\nabla_{\mathbf{y}} g_{kl}(\mathbf{s}; \boldsymbol{\theta}_0) = c_{kl}(\mathbf{s}) \nabla_{\mathbf{y}} d_{kl}(\mathbf{s}; \boldsymbol{\theta}_0).$$

Since $c_{kl}(\mathbf{s}) > 0$, substitution into the direct-contrast representation yields

$$\mathbf{H}_{kl} = \int_{\mathcal{F}_{kl}} \frac{c_{kl}(\mathbf{s})}{\|\nabla_{\mathbf{y}} d_{kl}(\mathbf{s}; \boldsymbol{\theta}_0)\|} \mathbf{a}_{kl}(\mathbf{s}) \mathbf{a}_{kl}(\mathbf{s})^\top dS(\mathbf{s}).$$

Thus the direct-contrast representation requires no positivity assumption, whereas the equivalent log-contrast form is often more convenient for the QDA calculations.

Proof of Corollary 4.3. Let

$$\mathbf{h}_n = \widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0.$$

By Theorem 4.2,

$$R(\widehat{\boldsymbol{\theta}}_n) - R^* = \frac{1}{2} \mathbf{h}_n^\top \mathbf{H}_R \mathbf{h}_n + o_p(n^{-1}),$$

because $\mathbf{h}_n = O_p(n^{-1/2})$. Hence

$$n \left\{ R(\widehat{\boldsymbol{\theta}}_n) - R^* \right\} = \frac{1}{2} (\sqrt{n} \mathbf{h}_n)^\top \mathbf{H}_R (\sqrt{n} \mathbf{h}_n) + o_p(1). \quad (\text{S4.6})$$

By assumption,

$$\sqrt{n} \mathbf{h}_n \xrightarrow{d} \mathbf{Z}, \quad \mathbf{Z} \sim N(\mathbf{0}, \mathbf{V}).$$

Since the map $\mathbf{x} \mapsto \mathbf{x}^\top \mathbf{H}_R \mathbf{x}$ is continuous, the continuous mapping theorem and Slutsky's theorem applied to (S4.6) give

$$n \left\{ R(\widehat{\boldsymbol{\theta}}_n) - R^* \right\} \xrightarrow{d} \frac{1}{2} \mathbf{Z}^\top \mathbf{H}_R \mathbf{Z}. \quad (\text{S4.7})$$

If the nonzero eigenvalues of $\mathbf{V}^{1/2} \mathbf{H}_R \mathbf{V}^{1/2}$ are $\rho_1, \dots, \rho_m$, then

$$\frac{1}{2} \mathbf{Z}^\top \mathbf{H}_R \mathbf{Z} \stackrel{d}{=} \frac{1}{2} \sum_{j=1}^m \rho_j \chi_{1,j}^2,$$

where the $\chi_{1,j}^2$ variables are independent.

To justify convergence of expectations, let

$$\mathbf{X}_n = \sqrt{n} (\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0).$$

By Theorem 4.2, there exist $\varepsilon > 0$ and $C < \infty$ such that

$$0 \leq R(\boldsymbol{\theta}_0 + \mathbf{h}) - R^* \leq C \|\mathbf{h}\|^2$$

whenever $\|\mathbf{h}\| \leq \varepsilon$. Hence, on the event $\{\|\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0\| \leq \varepsilon\}$,

$$n \left\{ R(\widehat{\boldsymbol{\theta}}_n) - R^* \right\} \leq C \|\mathbf{X}_n\|^2.$$

Let

$$Y_n = n \left\{ R(\widehat{\boldsymbol{\theta}}_n) - R^* \right\}, \quad A_n = \left\{ \|\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0\| \leq \varepsilon \right\}.$$

On A_n ,

$$0 \leq Y_n \mathbf{1}_{A_n} \leq C \|\mathbf{X}_n\|^2.$$

Since

$$\sup_n E \|\mathbf{X}_n\|^{2+\delta} < \infty,$$

the family $\{\|\mathbf{X}_n\|^2\}$ is uniformly integrable, and hence so is $\{Y_n \mathbf{1}_{A_n}\}$.

On A_n^c , since

$$0 \leq R(\hat{\boldsymbol{\theta}}_n) - R^* \leq 1,$$

we have

$$E(Y_n \mathbf{1}_{A_n^c}) \leq n \Pr(A_n^c).$$

Moreover, Markov's inequality gives

$$n \Pr(A_n^c) = n \Pr\{\|\mathbf{X}_n\| > \varepsilon \sqrt{n}\} \leq \frac{\sup_m E \|\mathbf{X}_m\|^{2+\delta}}{\varepsilon^{2+\delta} n^{\delta/2}} \longrightarrow 0.$$

Thus

$$Y_n \mathbf{1}_{A_n^c} \longrightarrow 0 \quad \text{in } L^1.$$

Since an L^1 -convergent family is uniformly integrable, $\{Y_n \mathbf{1}_{A_n^c}\}$ is uniformly integrable. Consequently, since

$$Y_n = Y_n \mathbf{1}_{A_n} + Y_n \mathbf{1}_{A_n^c},$$

the family $\{Y_n\}$ is uniformly integrable. Combining this with the convergence in distribution in (S4.7) yields

$$\begin{aligned} n \left[E \left\{ R(\hat{\boldsymbol{\theta}}_n) \right\} - R^* \right] &\longrightarrow \frac{1}{2} E(\mathbf{Z}^\top \mathbf{H}_R \mathbf{Z}) \\ &= \frac{1}{2} \text{tr}(\mathbf{H}_R \mathbf{V}), \end{aligned} \tag{S4.8}$$

where the final equality follows from $E(\mathbf{Z}\mathbf{Z}^\top) = \mathbf{V}$. Equivalently,

$$E \left\{ R(\hat{\boldsymbol{\theta}}_n) \right\} - R^* = \frac{1}{2n} \text{tr}(\mathbf{H}_R \mathbf{V}) + o(n^{-1}).$$

This proves Corollary 4.3. □

S5 Proof of the classification-weighted spectral criterion

This section proves Theorem 5.1.

Proof of Theorem 5.1. Since $\mathbf{A} \succ \mathbf{0}$ and $\mathbf{J} \succ \mathbf{0}$, the relative information matrix

$$\mathbf{C} = \mathbf{A}^{-1/2} \mathbf{J} \mathbf{A}^{-1/2}$$

is symmetric positive definite. Hence

$$\mathbf{J} = \mathbf{A}^{1/2} \mathbf{C} \mathbf{A}^{1/2}, \quad \mathbf{J}^{-1} = \mathbf{A}^{-1/2} \mathbf{C}^{-1} \mathbf{A}^{-1/2},$$

and therefore

$$\mathbf{A}^{-1} - \mathbf{J}^{-1} = \mathbf{A}^{-1/2} (\mathbf{I} - \mathbf{C}^{-1}) \mathbf{A}^{-1/2}. \quad (\text{S5.1})$$

Substituting (S5.1) into

$$\Delta_R = \text{tr} [\mathbf{H}_R (\mathbf{A}^{-1} - \mathbf{J}^{-1})]$$

and using cyclic invariance of the trace gives

$$\Delta_R = \text{tr} [\mathbf{W} (\mathbf{I} - \mathbf{C}^{-1})], \quad \mathbf{W} = \mathbf{A}^{-1/2} \mathbf{H}_R \mathbf{A}^{-1/2}. \quad (\text{S5.2})$$

Because $\mathbf{H}_R \succeq \mathbf{0}$, the matrix $\mathbf{W}$ is also positive semidefinite.

Let

$$\mathbf{C} = \mathbf{Q} \mathbf{\Lambda} \mathbf{Q}^\top, \quad \mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_r),$$

where $\mathbf{Q}^\top \mathbf{Q} = \mathbf{I}$ and $\lambda_j > 0$ for all j . Then

$$\mathbf{C}^{-1} = \mathbf{Q} \mathbf{\Lambda}^{-1} \mathbf{Q}^\top.$$

Using (S5.2) and cyclic invariance once more,

$$\begin{aligned} \Delta_R &= \text{tr} [\mathbf{W} \mathbf{Q} (\mathbf{I} - \mathbf{\Lambda}^{-1}) \mathbf{Q}^\top] \\ &= \text{tr} [\mathbf{Q}^\top \mathbf{W} \mathbf{Q} (\mathbf{I} - \mathbf{\Lambda}^{-1})]. \end{aligned} \quad (\text{S5.3})$$

Since $\mathbf{I} - \mathbf{\Lambda}^{-1}$ is diagonal, only the diagonal elements of $\mathbf{Q}^\top \mathbf{W} \mathbf{Q}$ contribute to the trace. Writing

$$w_j = \mathbf{q}_j^\top \mathbf{W} \mathbf{q}_j \geq 0,$$

where $\mathbf{q}_j$ is the j th column of $\mathbf{Q}$, we obtain

$$\Delta_R = \sum_{j=1}^r w_j \left(1 - \frac{1}{\lambda_j}\right) = \sum_{j=1}^r w_j \frac{\lambda_j - 1}{\lambda_j}. \quad (\text{S5.4})$$

This proves the spectral representation in Theorem 5.1.

Separating the terms corresponding to information gains and losses gives

$$\Delta_R = \sum_{\lambda_j > 1} w_j \frac{\lambda_j - 1}{\lambda_j} - \sum_{\lambda_j < 1} w_j \frac{1 - \lambda_j}{\lambda_j}. \quad (\text{S5.5})$$

Terms for which $\lambda_j = 1$ vanish. Since every $w_j \geq 0$, equation (S5.5) shows that $\Delta_R > 0$ if and only if

$$\sum_{\lambda_j > 1} w_j \frac{\lambda_j - 1}{\lambda_j} > \sum_{\lambda_j < 1} w_j \frac{1 - \lambda_j}{\lambda_j},$$

which is the classification-weighted gain-loss criterion stated in Theorem 5.1.

If an eigenvalue λ has multiplicity greater than one, the individual eigenvectors within its eigenspace are not uniquely determined. The total contribution of that eigenspace to Δ_R , however, is invariant. If $\mathbf{P}_\lambda$ denotes the orthogonal projector onto the eigenspace associated with λ , then

$$\sum_{j:\lambda_j=\lambda} w_j = \sum_{j:\lambda_j=\lambda} \mathbf{q}_j^\top \mathbf{W} \mathbf{q}_j = \text{tr}(\mathbf{W} \mathbf{P}_\lambda).$$

Hence the total contribution of this eigenspace is

$$\left(1 - \frac{1}{\lambda}\right) \text{tr}(\mathbf{W} \mathbf{P}_\lambda), \quad (\text{S5.6})$$

which is independent of the particular orthonormal basis chosen within the eigenspace.

This completes the proof. $\square$

S6 Proof of the local departure from MCAR result

This section proves Proposition 5.2. Throughout this section, expectations are taken under the true data-generating parameter $\boldsymbol{\theta}_0$.

Proof of Proposition 5.2. Write

$$U = U_{\boldsymbol{\theta}_0}(\mathbf{Y}), \quad \mathbf{G} = \nabla_{\boldsymbol{\theta}} U_{\boldsymbol{\theta}}(\mathbf{Y})|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0},$$

and let

$$q_t = \text{expit}\{\alpha(t) + tU\}.$$

The calibration condition is

$$E(q_t) = \gamma, \quad 0 < \gamma < 1. \quad (\text{S6.1})$$

To justify the local dependence of the calibrating intercept on t , define

$$F(a, t) = E[\text{expit}\{a + tU\}] - \gamma.$$

At $t = 0$, the equation $F(a, 0) = 0$ gives

$$\text{expit}(a) = \gamma,$$

and hence

$$\alpha(0) = \text{logit}(\gamma). \quad (\text{S6.2})$$

Moreover,

$$\left. \frac{\partial F(a, t)}{\partial a} \right|_{(a,t)=(\alpha(0),0)} = \gamma(1 - \gamma) > 0.$$

Under the stated regularity conditions, the implicit-function theorem therefore yields a locally differentiable solution $a = \alpha(t)$ of $F\{\alpha(t), t\} = 0$ in a neighborhood of $t = 0$.

Differentiating the calibration identity with respect to t , with differentiation under the expectation justified by the stated regularity conditions, gives

$$0 = E[q_t(1 - q_t) \{\alpha'(t) + U\}].$$

Evaluating at $t = 0$, where $q_0 = \gamma$, yields

$$0 = \gamma(1 - \gamma) [\alpha'(0) + E(U)],$$

and therefore

$$\alpha'(0) = -E(U). \quad (\text{S6.3})$$

A first-order Taylor expansion of q_t around $t = 0$ now gives

$$\begin{aligned} q_t &= q_0 + t \left. \frac{dq_t}{dt} \right|_{t=0} + O(t^2) \\ &= \gamma + t\gamma(1 - \gamma) \{\alpha'(0) + U\} + O(t^2) \\ &= \gamma + t\gamma(1 - \gamma) \{U - E(U)\} + O(t^2). \end{aligned} \quad (\text{S6.4})$$

This establishes the first expansion in Proposition 5.2.

For the label-information loss, write

$$\mathbf{I}_{Z|\mathbf{Y}} = \mathbf{I}_{Z|\mathbf{Y}}(\boldsymbol{\theta}_0; \mathbf{Y}).$$

Along the sequence of population experiments only the missing-label mechanism varies with t , whereas the data-generating classification model remains fixed at $\boldsymbol{\theta}_0$. Thus

$$\mathbf{D}(t) = E[q_t \mathbf{I}_{Z|\mathbf{Y}}].$$

Under the stated regularity conditions, differentiation may be passed under the expectation. Since

$$\left. \frac{dq_t}{dt} \right|_{t=0} = \gamma(1 - \gamma) \{U - E(U)\},$$

it follows that

$$\mathbf{D}'(0) = \gamma(1 - \gamma) E[\{U - E(U)\} \mathbf{I}_{Z|\mathbf{Y}}]. \quad (\text{S6.5})$$

We next consider the efficient information supplied by the missing-label indicators. For likelihood inference, the intercept and uncertainty slope are ordinary nuisance parameters. Write

$$\boldsymbol{\xi} = (\xi_0, \xi_1)^\top, \quad q(\mathbf{y}; \boldsymbol{\theta}, \boldsymbol{\xi}) = \text{expit} \{\xi_0 + \xi_1 U_{\boldsymbol{\theta}}(\mathbf{y})\}.$$

Along the population path considered here, the true nuisance value is

$$\boldsymbol{\xi}(t) = \{\alpha(t), t\}^\top.$$

At $(\boldsymbol{\theta}_0, \boldsymbol{\xi}(t))$, define

$$\mathbf{X} = \begin{pmatrix} 1 \\ U \end{pmatrix}.$$

The derivatives of the missingness probability are

$$q_{\theta} = t q_t(1 - q_t)\mathbf{G}, \quad q_{\xi} = q_t(1 - q_t)\mathbf{X}. \quad (\text{S6.6})$$

Therefore the Bernoulli information blocks of Theorem 2.1 are

$$\mathbf{B}_{\theta\theta}(t) = t^2 E [q_t(1 - q_t)\mathbf{G}\mathbf{G}^{\top}], \quad (\text{S6.7})$$

$$\mathbf{B}_{\theta\xi}(t) = t E [q_t(1 - q_t)\mathbf{G}\mathbf{X}^{\top}], \quad (\text{S6.8})$$

$$\mathbf{B}_{\xi\xi}(t) = E [q_t(1 - q_t)\mathbf{X}\mathbf{X}^{\top}]. \quad (\text{S6.9})$$

As $t \rightarrow 0$,

$$q_t(1 - q_t) \longrightarrow \gamma(1 - \gamma) \quad \text{almost surely,}$$

and

$$0 \leq q_t(1 - q_t) \leq \frac{1}{4}.$$

The assumed second-moment conditions imply

$$E\|\mathbf{G}\|^2 < \infty, \quad E\|\mathbf{X}\|^2 < \infty,$$

and, by the Cauchy-Schwarz inequality,

$$E(\|\mathbf{G}\| \|\mathbf{X}\|) < \infty.$$

Dominated convergence therefore gives

$$\begin{aligned} E [q_t(1 - q_t)\mathbf{G}\mathbf{G}^{\top}] &= \gamma(1 - \gamma)E [\mathbf{G}\mathbf{G}^{\top}] + o(1), \\ E [q_t(1 - q_t)\mathbf{G}\mathbf{X}^{\top}] &= \gamma(1 - \gamma)E [\mathbf{G}\mathbf{X}^{\top}] + o(1), \\ E [q_t(1 - q_t)\mathbf{X}\mathbf{X}^{\top}] &= \gamma(1 - \gamma)E [\mathbf{X}\mathbf{X}^{\top}] + o(1). \end{aligned}$$

Consequently,

$$\mathbf{B}_{\theta\theta}(t) = t^2 \gamma(1 - \gamma)E [\mathbf{G}\mathbf{G}^{\top}] + o(t^2), \quad (\text{S6.10})$$

$$\mathbf{B}_{\theta\xi}(t) = t \gamma(1 - \gamma)E [\mathbf{G}\mathbf{X}^{\top}] + o(t), \quad (\text{S6.11})$$

$$\mathbf{B}_{\xi\xi}(t) = \gamma(1 - \gamma)E [\mathbf{X}\mathbf{X}^{\top}] + o(1). \quad (\text{S6.12})$$

The assumption $\text{Var}(U) > 0$ implies that $E(\mathbf{X}\mathbf{X}^{\top})$ is positive definite. Indeed, for $\mathbf{a} = (a_0, a_1)^{\top} \neq \mathbf{0}$,

$$\mathbf{a}^{\top} E [\mathbf{X}\mathbf{X}^{\top}] \mathbf{a} = E [(a_0 + a_1 U)^2] > 0,$$

because a nonzero affine function of a nondegenerate random variable cannot vanish almost surely. Hence $\mathbf{B}_{\xi\xi}(t)$ is nonsingular for all sufficiently small t , and

$$\mathbf{B}_{\xi\xi}(t)^{-1} = \frac{1}{\gamma(1 - \gamma)} E [\mathbf{X}\mathbf{X}^{\top}]^{-1} + o(1). \quad (\text{S6.13})$$

Using the efficient-information formula

$$\mathbf{I}_M^{\text{eff}}(t) = \mathbf{B}_{\theta\theta}(t) - \mathbf{B}_{\theta\xi}(t)\mathbf{B}_{\xi\xi}(t)^{-1}\mathbf{B}_{\xi\theta}(t),$$

and substituting (S6.10)–(S6.13), we obtain

$$\mathbf{I}_M^{\text{eff}}(t) = t^2\gamma(1 - \gamma)\mathbf{V}_U + o(t^2), \quad (\text{S6.14})$$

where

$$\mathbf{V}_U = E[\mathbf{G}\mathbf{G}^\top] - E[\mathbf{G}\mathbf{X}^\top] E[\mathbf{X}\mathbf{X}^\top]^{-1} E[\mathbf{X}\mathbf{G}^\top]. \quad (\text{S6.15})$$

To verify that $\mathbf{V}_U \succeq \mathbf{0}$, define

$$\mathbf{C} = E[\mathbf{G}\mathbf{X}^\top] E[\mathbf{X}\mathbf{X}^\top]^{-1}.$$

Then

$$E[(\mathbf{G} - \mathbf{C}\mathbf{X})(\mathbf{G} - \mathbf{C}\mathbf{X})^\top] = \mathbf{V}_U. \quad (\text{S6.16})$$

The left-hand side is a second-moment matrix and is therefore positive semidefinite. Hence

$$\mathbf{V}_U \succeq \mathbf{0}.$$

It remains to obtain the derivative of the classification advantage. Let

$$\mathbf{J}(t) = \mathbf{A} - \mathbf{D}(t) + \mathbf{I}_M^{\text{eff}}(t)$$

denote the efficient information under the partially classified experiment. At $t = 0$,

$$\mathbf{J}(0) = \mathbf{A} - \gamma E[\mathbf{I}_{Z|Y}] = \mathbf{J}_0.$$

Since $\mathbf{I}_M^{\text{eff}}(t) = O(t^2)$,

$$\left. \frac{d}{dt} \mathbf{I}_M^{\text{eff}}(t) \right|_{t=0} = \mathbf{0},$$

and therefore

$$\mathbf{J}'(0) = -\mathbf{D}'(0). \quad (\text{S6.17})$$

Now

$$\Delta_R(t) = \text{tr}[\mathbf{H}_R \{\mathbf{A}^{-1} - \mathbf{J}(t)^{-1}\}].$$

Using the matrix derivative identity

$$\frac{d}{dt} \mathbf{J}(t)^{-1} = -\mathbf{J}(t)^{-1} \mathbf{J}'(t) \mathbf{J}(t)^{-1},$$

we obtain

$$\begin{aligned} \Delta'_R(0) &= \text{tr}[\mathbf{H}_R \mathbf{J}_0^{-1} \mathbf{J}'(0) \mathbf{J}_0^{-1}] \\ &= -\text{tr}[\mathbf{H}_R \mathbf{J}_0^{-1} \mathbf{D}'(0) \mathbf{J}_0^{-1}], \end{aligned} \quad (\text{S6.18})$$

where the second equality follows from (S6.17). This is the final assertion of the proposition.

Thus the label-information loss may change at first order in t , whereas the efficient information contributed by the missing-label mechanism begins only at order t^2 . This completes the proof of Proposition 5.2. $\square$

S7 Quadratic discriminant calculations

This section gives the score, Fisher-information, and uncertainty derivatives used for the three-class QDA illustration in Section 6. The parameterization is

$$\boldsymbol{\theta} = (\boldsymbol{\alpha}^\top, \boldsymbol{\mu}_1^\top, \boldsymbol{\mu}_2^\top, \boldsymbol{\mu}_3^\top, \boldsymbol{\sigma}_1^\top, \boldsymbol{\sigma}_2^\top, \boldsymbol{\sigma}_3^\top)^\top,$$

where $\boldsymbol{\alpha} = (\alpha_1, \alpha_2)^\top$ parameterizes the class probabilities relative to class 3 and $\boldsymbol{\sigma}_k = \text{vech}(\boldsymbol{\Sigma}_k)$.

Class probabilities and QDA discriminants. Using class 3 as the baseline, let

$$\alpha_1 = \log \frac{\pi_1}{\pi_3}, \quad \alpha_2 = \log \frac{\pi_2}{\pi_3}.$$

Then, with $D_\alpha = 1 + \exp(\alpha_1) + \exp(\alpha_2)$,

$$\pi_1 = \frac{\exp(\alpha_1)}{D_\alpha}, \quad \pi_2 = \frac{\exp(\alpha_2)}{D_\alpha}, \quad \pi_3 = \frac{1}{D_\alpha}.$$

If $\mathbf{e}_1 = (1, 0)^\top$, $\mathbf{e}_2 = (0, 1)^\top$, and $\mathbf{e}_3 = (0, 0)^\top$, then

$$\nabla_{\boldsymbol{\alpha}} \log \frac{\pi_k}{\pi_l} = \mathbf{e}_k - \mathbf{e}_l. \quad (\text{S7.1})$$

The complete-classification score and information for $\boldsymbol{\alpha}$ are

$$\mathbf{S}_\alpha = \begin{pmatrix} \mathbb{I}(Z=1) - \pi_1 \\ \mathbb{I}(Z=2) - \pi_2 \end{pmatrix}, \quad \mathbf{I}_\alpha = \begin{pmatrix} \pi_1(1 - \pi_1) & -\pi_1\pi_2 \\ -\pi_1\pi_2 & \pi_2(1 - \pi_2) \end{pmatrix}.$$

For $\mathbf{x}_k = \mathbf{y} - \boldsymbol{\mu}_k$, the pairwise QDA discriminant is

$$\begin{aligned} d_{kl}(\mathbf{y}) &= \log \frac{\pi_k}{\pi_l} - \frac{1}{2} \log \frac{|\boldsymbol{\Sigma}_k|}{|\boldsymbol{\Sigma}_l|} - \frac{1}{2} \mathbf{x}_k^\top \boldsymbol{\Sigma}_k^{-1} \mathbf{x}_k \\ &\quad + \frac{1}{2} \mathbf{x}_l^\top \boldsymbol{\Sigma}_l^{-1} \mathbf{x}_l. \end{aligned} \quad (\text{S7.2})$$

Its spatial derivative is

$$\nabla_{\mathbf{y}} d_{kl}(\mathbf{y}) = -\boldsymbol{\Sigma}_k^{-1} \mathbf{x}_k + \boldsymbol{\Sigma}_l^{-1} \mathbf{x}_l, \quad (\text{S7.3})$$

while the derivatives with respect to the class means are

$$\nabla_{\boldsymbol{\mu}_k} d_{kl} = \boldsymbol{\Sigma}_k^{-1} \mathbf{x}_k, \quad \nabla_{\boldsymbol{\mu}_l} d_{kl} = -\boldsymbol{\Sigma}_l^{-1} \mathbf{x}_l.$$

All mean blocks corresponding to classes other than k and l are zero.

For the covariance terms, the standard identities

$$d \log |\boldsymbol{\Sigma}| = \text{tr}(\boldsymbol{\Sigma}^{-1} d\boldsymbol{\Sigma}), \quad d\boldsymbol{\Sigma}^{-1} = -\boldsymbol{\Sigma}^{-1} (d\boldsymbol{\Sigma}) \boldsymbol{\Sigma}^{-1}$$

give, under symmetric covariance perturbations,

$$\frac{\partial d_{kl}}{\partial \Sigma_k} = \frac{1}{2} [\Sigma_k^{-1} \mathbf{x}_k \mathbf{x}_k^\top \Sigma_k^{-1} - \Sigma_k^{-1}], \quad (\text{S7.4})$$

$$\frac{\partial d_{kl}}{\partial \Sigma_l} = -\frac{1}{2} [\Sigma_l^{-1} \mathbf{x}_l \mathbf{x}_l^\top \Sigma_l^{-1} - \Sigma_l^{-1}]. \quad (\text{S7.5})$$

If $\boldsymbol{\sigma}_k = \text{vech}(\Sigma_k)$ and $\mathbf{D}_p$ denotes the duplication matrix, then

$$\nabla_{\boldsymbol{\sigma}_k} d_{kl} = \mathbf{D}_p^\top \text{vec} \left(\frac{\partial d_{kl}}{\partial \Sigma_k} \right), \quad (\text{S7.6})$$

with the analogous expression for class l . Together with (S7.1), these derivatives form

$$\mathbf{a}_{kl}(\mathbf{y}) = \nabla_{\boldsymbol{\theta}} d_{kl}(\mathbf{y}),$$

the log-contrast boundary-sensitivity vector. On an active face $\mathcal{F}_{kl}$, where

$$c_{kl}(\mathbf{s}) = r_k(\mathbf{s}) = r_l(\mathbf{s}),$$

the relation between the density contrast $g_{kl} = r_k - r_l$ and the log contrast d_{kl} gives

$$\nabla_{\boldsymbol{\theta}} g_{kl}(\mathbf{s}) = c_{kl}(\mathbf{s}) \mathbf{a}_{kl}(\mathbf{s}), \quad \nabla_{\mathbf{y}} g_{kl}(\mathbf{s}) = c_{kl}(\mathbf{s}) \nabla_{\mathbf{y}} d_{kl}(\mathbf{s}).$$

Hence the corresponding surface-integral contribution is

$$\mathbf{H}_{kl} = \int_{\mathcal{F}_{kl}} \frac{c_{kl}(\mathbf{s})}{\|\nabla_{\mathbf{y}} d_{kl}(\mathbf{s})\|} \mathbf{a}_{kl}(\mathbf{s}) \mathbf{a}_{kl}(\mathbf{s})^\top dS(\mathbf{s}).$$

Complete-classification Fisher information. For class k , the contribution to the mean score is

$$\mathbf{S}_{\mu_k} = \mathbb{I}(Z = k) \Sigma_k^{-1} (\mathbf{Y} - \boldsymbol{\mu}_k),$$

and therefore

$$\mathbf{I}_{\mu_k} = E [\mathbf{S}_{\mu_k} \mathbf{S}_{\mu_k}^\top] = \pi_k \Sigma_k^{-1}. \quad (\text{S7.7})$$

For $\boldsymbol{\sigma}_k = \text{vech}(\Sigma_k)$, the covariance-information block is

$$\mathbf{I}_{\sigma_k} = \frac{\pi_k}{2} \mathbf{D}_p^\top (\Sigma_k^{-1} \otimes \Sigma_k^{-1}) \mathbf{D}_p. \quad (\text{S7.8})$$

The mean and covariance scores within a Gaussian class are orthogonal because the relevant centered third moments vanish. Scores for distributional parameters belonging to different classes are orthogonal because their class indicators are mutually exclusive. The prior score is also orthogonal to the within-class mean and covariance scores, since each latter score has conditional mean zero given Z . Consequently, in the parameterization above,

$$\mathbf{I}_{\text{CC}} = \text{blockdiag} (\mathbf{I}_\alpha, \mathbf{I}_{\mu_1}, \mathbf{I}_{\mu_2}, \mathbf{I}_{\mu_3}, \mathbf{I}_{\sigma_1}, \mathbf{I}_{\sigma_2}, \mathbf{I}_{\sigma_3}). \quad (\text{S7.9})$$

Posterior probabilities and uncertainty derivatives. Let

$$r_k(\mathbf{y}; \boldsymbol{\theta}) = \pi_k f_k(\mathbf{y}), \quad \tau_k(\mathbf{y}) = \frac{r_k(\mathbf{y})}{\sum_{j=1}^3 r_j(\mathbf{y})},$$

and define

$$\mathbf{s}_k(\mathbf{y}) = \nabla_{\boldsymbol{\theta}} \log r_k(\mathbf{y}), \quad \bar{\mathbf{s}}(\mathbf{y}) = \sum_{j=1}^3 \tau_j(\mathbf{y}) \mathbf{s}_j(\mathbf{y}).$$

Differentiating the normalized posterior probability gives

$$\nabla_{\boldsymbol{\theta}} \tau_k(\mathbf{y}) = \tau_k(\mathbf{y}) \{ \mathbf{s}_k(\mathbf{y}) - \bar{\mathbf{s}}(\mathbf{y}) \}. \quad (\text{S7.10})$$

For Shannon entropy

$$H(\mathbf{y}) = - \sum_{k=1}^3 \tau_k(\mathbf{y}) \log \tau_k(\mathbf{y}),$$

the identity $\sum_k \nabla_{\boldsymbol{\theta}} \tau_k(\mathbf{y}) = \mathbf{0}$ and (S7.10) yield

$$\nabla_{\boldsymbol{\theta}} H(\mathbf{y}) = - \sum_{k=1}^3 \tau_k(\mathbf{y}) \log \tau_k(\mathbf{y}) \{ \mathbf{s}_k(\mathbf{y}) - \bar{\mathbf{s}}(\mathbf{y}) \}. \quad (\text{S7.11})$$

Thus, for the normalized entropy uncertainty $U_H = H / \log 3$,

$$\nabla_{\boldsymbol{\theta}} U_H(\mathbf{y}) = \frac{1}{\log 3} \nabla_{\boldsymbol{\theta}} H(\mathbf{y}). \quad (\text{S7.12})$$

For Gini uncertainty

$$G(\mathbf{y}) = 1 - \sum_{k=1}^3 \tau_k(\mathbf{y})^2,$$

we similarly obtain

$$\begin{aligned} \nabla_{\boldsymbol{\theta}} G(\mathbf{y}) &= -2 \sum_{k=1}^3 \tau_k(\mathbf{y}) \nabla_{\boldsymbol{\theta}} \tau_k(\mathbf{y}) \\ &= -2 \sum_{k=1}^3 \tau_k(\mathbf{y})^2 \{ \mathbf{s}_k(\mathbf{y}) - \bar{\mathbf{s}}(\mathbf{y}) \}. \end{aligned} \quad (\text{S7.13})$$

Since $U_G = (3/2)G$ for three classes,

$$\nabla_{\boldsymbol{\theta}} U_G(\mathbf{y}) = -3 \sum_{k=1}^3 \tau_k(\mathbf{y})^2 \{ \mathbf{s}_k(\mathbf{y}) - \bar{\mathbf{s}}(\mathbf{y}) \}. \quad (\text{S7.14})$$

Finally, for the logistic missing-label mechanism

$$q(\mathbf{y}; \boldsymbol{\theta}, \boldsymbol{\xi}) = \text{expit} \{ \xi_0 + \xi_1 U_{\boldsymbol{\theta}}(\mathbf{y}) \}, \quad \boldsymbol{\xi} = (\xi_0, \xi_1)^\top,$$

the required derivatives are

$$q_{\theta} = q(1 - q) \xi_1 \nabla_{\theta} U_{\theta}(\mathbf{y}), \quad (\text{S7.15})$$

$$q_{\xi} = q(1 - q) \begin{pmatrix} 1 \\ U_{\theta}(\mathbf{y}) \end{pmatrix}. \quad (\text{S7.16})$$

Substitution of (S7.15)–(S7.16) into the Bernoulli information formula of Theorem 2.1 gives

$$\mathbf{B}_{\theta\theta} = E \left[q(1 - q) \xi_1^2 \nabla_{\theta} U_{\theta}(\mathbf{Y}) \nabla_{\theta} U_{\theta}(\mathbf{Y})^{\top} \right], \quad (\text{S7.17})$$

$$\mathbf{B}_{\theta\xi} = E \left[q(1 - q) \xi_1 \nabla_{\theta} U_{\theta}(\mathbf{Y}) \begin{pmatrix} 1 & U_{\theta}(\mathbf{Y}) \end{pmatrix} \right], \quad (\text{S7.18})$$

$$\mathbf{B}_{\xi\xi} = E \left[q(1 - q) \begin{pmatrix} 1 \\ U_{\theta}(\mathbf{Y}) \end{pmatrix} \begin{pmatrix} 1 & U_{\theta}(\mathbf{Y}) \end{pmatrix} \right]. \quad (\text{S7.19})$$

These expressions, together with the conditional label-information loss, provide the matrices used to evaluate $\mathbf{I}_M^{\text{eff}}$ and $\mathbf{I}_{\text{PC}}^{\text{eff}}$ in the QDA experiments.

S8 Additional population robustness results

This section reports additional population calculations complementing the numerical investigation in Section 7. The results examine the stability of the classification comparison under covariance heterogeneity, its decomposition across active Bayes faces under prior imbalance, the dependence of the phase boundary on the classification geometry, and the robustness of that phase structure to the choice of posterior-uncertainty functional.

Covariance heterogeneity. Let

$$\overline{\Sigma} = \sum_{k=1}^3 \pi_k \Sigma_k^{(0)}$$

and consider the path

$$\Sigma_k(\rho) = (1 - \rho) \overline{\Sigma} + \rho \Sigma_k^{(0)}, \quad 0 \leq \rho \leq 1.$$

The model therefore varies continuously from common-covariance linear discrimination at $\rho = 0$ to the reference QDA configuration at $\rho = 1$. Table S1 reports the corresponding population risk coefficients.

Informative partial classification remains favourable throughout this path. The relative efficiency decreases only modestly, from 1.160 under common covariance matrices to 1.142 in the reference QDA model. Thus, within this controlled family, increasing covariance heterogeneity weakens the magnitude of the advantage slightly but does not change its sign. The substantially larger values of $\mathcal{E}_{\text{IG}}$ also show that ignoring the informative missingness mechanism performs markedly worse throughout the path.

Face-specific efficiency under prior imbalance. The global classification comparison may conceal different behavior across the active pairwise boundaries. For an active pair $k < l$, define the face-specific asymptotic relative efficiency by

$$\text{ARE}_{kl} = \frac{\text{tr}(\mathbf{H}_{kl} \mathbf{A}^{-1})}{\text{tr}(\mathbf{H}_{kl} \mathbf{J}^{-1})},$$

Table S1: Classification efficiency along the covariance-heterogeneity path $\Sigma_k(\rho)$.

ρ	$\mathcal{E}_{CC}$	$\mathcal{E}_{IG}$	$\mathcal{E}_{PC}$	ARE_R
0.00	1.652	2.723	1.425	1.160
0.20	1.651	2.720	1.425	1.159
0.40	1.644	2.712	1.422	1.156
0.60	1.627	2.688	1.412	1.152
0.80	1.599	2.648	1.395	1.147
1.00	1.558	2.585	1.365	1.142

Table S2: Critical normalized entropy slopes under selected class-separation levels.

s	$\gamma = 0.10$	$\gamma = 0.30$	$\gamma = 0.50$
0.75	3.04	3.57	4.62
1.00	2.90	3.52	4.87
1.30	2.90	3.67	5.38

provided the denominator is positive. Thus $ARE_{kl} > 1$ means that informative partial classification has the smaller leading excess-risk contribution associated with the k -versus- l active Bayes face.

In the rare-class configuration $\pi_3 = 0.05$, the approximate face-specific relative efficiencies are

$$ARE_{12} = 1.040, \quad ARE_{13} = 0.968, \quad ARE_{23} = 0.856.$$

Thus the informative mechanism remains favourable for the boundary separating the two common classes but loses efficiency along both faces involving the rare third class. The combined effect yields $ARE_R < 1$.

For the reference prior $\pi_3 = 0.30$, the corresponding values are

$$ARE_{12} = 1.134, \quad ARE_{13} = 1.139, \quad ARE_{23} = 1.151,$$

so all three active faces contribute favourably to the global result. At the opposite extreme, with $\pi_3 = 0.90$,

$$ARE_{12} = 0.931, \quad ARE_{13} = 0.938, \quad ARE_{23} = 1.051.$$

The global deterioration under strong prior imbalance therefore again reflects heterogeneous face-specific contributions: only the 2-versus-3 boundary remains favourable.

Geometry-dependent phase boundaries. The critical normalized entropy slope also varies with the underlying Bayes geometry. Table S2 reports the thresholds for selected class-separation values, while Table S3 gives the corresponding thresholds for selected class-3 prior probabilities.

These thresholds confirm that the existence and location of a favourable regime depend jointly on the missing-label proportion and the classification geometry. In particular, for

Table S3: Critical normalized entropy slopes under selected class-3 prior probabilities.

π_3	$\gamma = 0.10$	$\gamma = 0.30$	$\gamma = 0.50$
0.10	3.07	4.07	no crossing observed
0.30	2.90	3.52	4.87
0.80	3.05	3.95	6.54

Table S4: Critical normalized uncertainty slopes for entropy- and Gini-dependent missingness in the reference QDA configuration.

γ	Entropy t_H^*	Gini t_G^*
0.10	2.90	2.35
0.30	3.52	2.84
0.50	4.87	3.88

$\pi_3 = 0.10$ and $\gamma = 0.50$, no crossing was observed over the mechanism range examined. The classification relative efficiency increased toward one at moderate uncertainty slopes but remained below one and subsequently declined as the uncertainty dependence became stronger. Thus stronger informativeness alone does not guarantee a favourable classification regime.

Entropy versus Gini uncertainty. To assess whether the phase-transition behavior is specific to normalized Shannon entropy, we repeated the reference population calculation using the normalized Gini uncertainty defined in Section 6. Both uncertainty functionals take values in $[0, 1]$, so their logistic slopes operate on the same normalized range. For each value of γ , the mechanism intercept was recalibrated to preserve the specified marginal missing-label proportion.

The two uncertainty measures produce the same qualitative phase structure: the critical slope increases with the marginal missing-label proportion under both mechanisms. For the reference QDA configuration, the Gini-based mechanism reaches the favourable region at a smaller normalized slope than the entropy-based mechanism for each value of γ examined. Because the two uncertainty functionals have different shapes even after normalization, these numerical thresholds should not be interpreted as a general efficiency ordering between entropy and Gini uncertainty. Their role here is instead to show that the transition between unfavourable and favourable informative missingness is not specific to the entropy specification.

S9 Additional computational details for the finite-sample study

This section gives implementation details and numerical diagnostics for the finite-sample experiment reported in Section 8.

Parameterization, optimization, and convergence. For numerical optimization, each covariance matrix was represented through a lower log-Cholesky factor,

$$\boldsymbol{\Sigma}_k = \mathbf{L}_k \mathbf{L}_k^\top, \quad \mathbf{L}_k = \begin{pmatrix} \exp(\ell_{k1}) & 0 \\ \ell_{k2} & \exp(\ell_{k3}) \end{pmatrix},$$

which guarantees positive definiteness throughout the optimization.

The complete-classification estimator was obtained in closed form. The MCAR and IPC estimators were obtained by numerical maximization of their respective observed-data likelihoods. Three starting values were used for each numerical fit in the final Monte Carlo experiment. For IPC, the optimization vector contains the 17 QDA parameters together with the two missingness parameters (ξ_0, t_H) .

A numerical fit was retained only when the optimizer reported successful termination, the estimated parameter vector was not numerically located at an optimization bound, and the relative finite-difference gradient was below the prespecified convergence tolerance. The resulting numbers of usable fits were

n	CC	MCAR	IPC
250	500	500	463
500	500	500	485
1000	500	499	491

corresponding to IPC convergence rates of 92.6%, 97.0%, and 98.2%, respectively.

Population quantities and parameterization checks. The information matrices entering the asymptotic comparison were evaluated under the true generating model. The active-face curvature matrix $\mathbf{H}_R$ was obtained by numerical contour integration over the three active pairwise Bayes boundaries.

In the log-Cholesky coordinates used for numerical optimization,

$$\text{tr}(\mathbf{H}_R) = 0.7227935085,$$

whereas in the $(\boldsymbol{\alpha}, \boldsymbol{\mu}, \text{vech}(\boldsymbol{\Sigma}))$ coordinates used in the theoretical presentation,

$$\text{tr}(\mathbf{H}_R) = 0.6035616902.$$

The difference is expected because the trace of $\mathbf{H}_R$ itself is parameterization dependent. By contrast, the classification-risk coefficients $\text{tr}(\mathbf{H}_R \mathbf{V}_m)$ are invariant under the corresponding smooth coordinate transformation.

For the reference configuration, the resulting asymptotic coefficients are

$$\mathcal{K}_{\text{CC}} = 1.5581470885, \quad \mathcal{K}_{\text{MCAR}} = 2.0540758514, \quad \mathcal{K}_{\text{IPC}} = 1.3645461995.$$

Hence

$$\text{ARE}_{R,\text{IPC:CC}} = 1.1418793215, \quad \text{ARE}_{R,\text{MCAR:CC}} = 0.7585635591.$$

Numerical evaluation of population excess risk. For a fitted parameter vector $\hat{\boldsymbol{\theta}}$, excess risk was evaluated as

$$\mathcal{X}(\hat{\boldsymbol{\theta}}) = E_{\boldsymbol{\theta}_0} \left[\max_k \tau_{0k}(\mathbf{Y}) - \tau_{0,\hat{C}(\mathbf{Y})}(\mathbf{Y}) \right].$$

The expectation was approximated with a common stratified population sample of 600,000 observations, consisting of 200,000 draws conditionally from each class. Class-specific averages were then weighted by the true class probabilities. The same population sample was used for every converged fitted classifier, thereby reducing numerical integration noise in comparisons across methods and Monte Carlo replications.

The integration variability was small relative to the sampling variability of the fitted classifiers. At $n = 1000$, the scaled integration Monte Carlo standard errors were approximately

$$0.0082, \quad 0.0099, \quad 0.0073$$

for CC, MCAR, and IPC, respectively, compared with replication Monte Carlo standard errors

$$0.0424, \quad 0.0559, \quad 0.0398.$$

Thus the uncertainty in the reported finite-sample comparisons is dominated by variation across fitted training samples rather than by the numerical population-risk calculation.

Quadratic risk and covariance validation. Since the quadratic expansion gives

$$R(\hat{\boldsymbol{\theta}}) - R^* = \frac{1}{2}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)^\top \mathbf{H}_R(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) + o_p(n^{-1}),$$

the directly evaluated excess risk was compared on the scale

$$2n \mathcal{X}(\hat{\boldsymbol{\theta}}),$$

for which the asymptotic coefficient is $\text{tr}(\mathbf{H}_R \mathbf{V}_m)$. Correspondingly, the local quadratic approximation was checked by evaluating

$$n(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)^\top \mathbf{H}_R(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)$$

for each usable fit. At $n = 1000$, the Monte Carlo averages were

$$1.5341, \quad 2.0407, \quad 1.3735$$

for CC, MCAR, and IPC, respectively. The corresponding directly evaluated values of $2n \mathcal{X}(\hat{\boldsymbol{\theta}})$ were

$$1.5247, \quad 2.0294, \quad 1.3682,$$

Table S5: Covariance diagnostics for the finite-sample Monte Carlo experiment.

n	Method	$D_{V,m}(n)$	$D_{R,m}(n)$	Empirical trace	Theoretical trace
250	CC	0.1573	0.0134	1.5373	1.5581
	MCAR	0.1608	0.0530	2.1629	2.0541
	IPC	0.1965	0.0480	1.4300	1.3645
500	CC	0.1550	0.0476	1.4840	1.5581
	MCAR	0.1777	0.0261	2.0005	2.0541
	IPC	0.1773	0.0442	1.4248	1.3645
1000	CC	0.1249	0.0142	1.5361	1.5581
	MCAR	0.1205	0.0051	2.0435	2.0541
	IPC	0.1309	0.0068	1.3739	1.3645

while the theoretical coefficients were

$$1.5581, \quad 2.0541, \quad 1.3645.$$

The agreement among these three calculations provides a direct numerical check of the quadratic risk approximation and of its information-based asymptotic limit.

For the covariance calculation, define

$$\widehat{\mathbf{V}}_{m,n} = n \widehat{\text{Cov}}(\widehat{\boldsymbol{\theta}}_m).$$

Let $\mathbf{V}_m$ denote the asymptotic covariance matrix for method m . The empirical covariance matrices were compared with their theoretical limits using the relative Frobenius discrepancy

$$D_{V,m}(n) = \frac{\|\widehat{\mathbf{V}}_{m,n} - \mathbf{V}_m\|_F}{\|\mathbf{V}_m\|_F},$$

and the classification-weighted discrepancy

$$D_{R,m}(n) = \frac{|\text{tr}(\mathbf{H}_R \widehat{\mathbf{V}}_{m,n}) - \text{tr}(\mathbf{H}_R \mathbf{V}_m)|}{\text{tr}(\mathbf{H}_R \mathbf{V}_m)}.$$

Accordingly, the “Empirical trace” and “Theoretical trace” columns below refer respectively to

$$\text{tr}(\mathbf{H}_R \widehat{\mathbf{V}}_{m,n}) \quad \text{and} \quad \text{tr}(\mathbf{H}_R \mathbf{V}_m).$$

The full covariance discrepancy is reduced by $n = 1000$, although the improvement is not monotone across all intermediate sample sizes because it weights all parameter directions equally. The classification-weighted discrepancy is substantially smaller at $n = 1000$: approximately 1.4%, 0.5%, and 0.7% for CC, MCAR, and IPC, respectively. This latter comparison is the more directly relevant diagnostic for the excess-risk theory, since only covariance error in directions weighted by $\mathbf{H}_R$ contributes to the leading classification risk.

Paired comparisons and reproducibility. Because CC and IPC were constructed from the same complete sample within each Monte Carlo replication, their excess risks can also be compared pairwise. Among replications in which both estimators were usable, the mean scaled paired difference

$$2n \{ \mathcal{X}_{\text{CC}} - \mathcal{X}_{\text{IPC}} \}$$

was

$$0.1018, \quad 0.0462, \quad 0.1497$$

for $n = 250, 500$, and 1000 , respectively. The proportions of paired replications in which IPC had the smaller excess risk were

$$0.538, \quad 0.509, \quad 0.556.$$

These proportions are not expected to approach one: the theoretical result concerns the difference in expected excess risk, not samplewise stochastic dominance.

The final simulation used $B = 500$ independent complete samples at each sample size and was implemented in R. The same generated complete sample was used to construct CC, MCAR, and IPC within each replication.