
Predicting Rare Events by Shrinking Towards Proportional Odds

Gregory Faletto¹, Jacob Bien¹

Abstract

Training classifiers is difficult with severe class imbalance, but many rare events are the culmination of a sequence with much more common intermediate outcomes. For example, in online marketing a user first sees an ad, then may click on it, and finally may make a purchase; estimating the probability of purchases is difficult because of their rarity. We show both theoretically and through data experiments that the more abundant data in earlier steps may be leveraged to improve estimation of probabilities of rare events. We present PRESTO, a relaxation of the proportional odds model for ordinal regression. Instead of estimating weights for one separating hyperplane that is shifted by separate intercepts for each of the estimated Bayes decision boundaries between adjacent pairs of categorical responses, we estimate separate weights for each of these transitions. We impose an L1 penalty on the differences between weights for the same feature in adjacent weight vectors in order to shrink towards the proportional odds model. We prove that PRESTO consistently estimates the decision boundary weights under a sparsity assumption. Synthetic and real data experiments show that our method can estimate rare probabilities in this setting better than both logistic regression on the rare category, which fails to borrow strength from more abundant categories, and the proportional odds model, which is too inflexible.

1. Introduction

Estimating probabilities of rare events is known to be difficult due to class imbalance. However, sometimes these events are the culmination of a sequential process with intermediate outcomes. For example:

¹Department of Data Sciences and Operations, University of Southern California, Los Angeles, CA, USA. Correspondence to: Gregory Faletto <gregory.faletto@marshall.usc.edu>.

1. In online marketing, a customer is first served an ad, then may click on it, then may indicate interest in making a purchase (by “liking” the product, for example), and finally may make a purchase.
2. In health and medicine, many outcomes can be encoded as ordered categorical variables, like reported quality of life and disease progression (Norris et al., 2006).
3. Sales of high-price durable goods typically follow a *sales funnel* (Duncan & Elkan, 2015). For example, when buying a car often a potential buyer first comes in to see a car, may take a test drive, and finally may buy the car.

In many of these cases, the intermediate events are much more common than the rare events. Though these intermediate events may not be of direct interest, if the features that contribute to the probability of advancing through earlier classes also contribute to the probability of advancing through later classes, then the more abundant intermediate events can be leveraged to improve estimation of the rare event probabilities.

The *proportional odds model* (McCullagh, 1980), also called the *ordered logit model* (Cameron & Trivedi, 2005, Section 15.9.1), satisfies, for ordinal outcomes $k \in \{1, \dots, K-1\}$,

$$\log \left(\frac{\mathbb{P}(y \leq k | \mathbf{x})}{\mathbb{P}(y > k | \mathbf{x})} \right) = \alpha_k + \boldsymbol{\beta}^\top \mathbf{x}, \quad (1)$$

where $\boldsymbol{\beta} \in \mathbb{R}^p$ is a vector of weights and $\mathbf{x} \in \mathbb{R}^p$ is a vector of features. This implies that for all $k \in \{1, \dots, K-1\}$

$$p_k(\mathbf{x}) := \mathbb{P}(y \leq k | \mathbf{x}) = F \left(\alpha_k + \boldsymbol{\beta}^\top \mathbf{x} \right), \quad (2)$$

where $F(\cdot)$ is the logistic cumulative distribution function, $F(t) = \exp\{t\}/[1 + \exp\{t\}]$. Notice that $\alpha_k + \boldsymbol{\beta}^\top \mathbf{x}$ is the Bayes decision boundary for the binary random variable $\mathbb{1}\{y \leq k\} | \mathbf{x}$. This problem could instead be cast as $K-1$ binary classification problems of the form (2) for adjacent classes:

$$\log \left(\frac{\mathbb{P}(y \leq k | \mathbf{x})}{\mathbb{P}(y > k | \mathbf{x})} \right) = \alpha_k + \boldsymbol{\beta}_k^\top \mathbf{x}, \quad k \in \{1, \dots, K-1\}. \quad (3)$$

The condition that the weight vectors $\boldsymbol{\beta}_k$ of the separating hyperplanes in (3) are all equal, as in (1), has been called

the *proportional odds assumption* (McCullagh, 1980) or the *parallel regression assumption* (Greene, 2012, Section 18.3.2). One way to motivate this model is by supposing that the response is driven by a latent (unobserved) variable U ,

$$U = \beta^\top \mathbf{x} + \epsilon, \quad (4)$$

where ϵ has a standard logistic distribution and is independent of \mathbf{x} . Response k is observed if and only if $-\alpha_k \leq U_i < -\alpha_{k-1}$ (where we define $\alpha_0 := -\infty$ and $\alpha_K := \infty$). This model leads to (2). (See Section 3.3.2 of Agresti 2010 for a more detailed explanation.)

Because the proportional odds model assumes that the decision boundaries between adjacent classes are all governed by the same hyperplane defined by β (separated only by different intercepts α_k), it assumes that the decision boundary between any two classes perfectly explains the decision boundary between any two other classes, other than an intercept term. If a rare event has much more common intermediate events before it, this model can therefore be very useful for better estimating the parameters of the model, and therefore better estimating the rare event probabilities. However, it could be that the proportional odds assumption is too rigid to be realistic, because observed features may have varying influence at different decision boundaries. For example:

1. In online marketing, users may click on an ad only to realize that the product is not what they were expecting, resulting in a particularly low probability of purchase.
2. For expensive goods like a home or car, potential buyers may express interest by going on a tour or taking a test drive purely out of curiosity; this may be distinct from their level of interest in actually making a purchase.
3. Students may place weights on different factors when deciding whether to apply to graduate school than they did when deciding whether to apply to an undergraduate program—they may have more appealing alternatives to additional schooling, they may face new financial or personal constraints because they are older, etc.

In each of these settings, if specific features vary in relevance for different decision boundaries while other features have about the same influence at every boundary, the proportional odds assumption may be too strong. Violations or relaxations of the proportional odds assumption along the lines of (3) have previously been considered by, for example, Brant (1990). Peterson & Harrell Jr (1990) developed *partial proportional odds models*, which allow the proportional odds assumption to hold for some features but not others, an idea previously mentioned by Armstrong & Sloan (1989).

(See Section 3.6.1 of Agresti 2010 for a textbook-level discussion). These relaxations have not been widely adopted because fitting separate weights for each outcome is too flexible unless $p(K-1) \ll n$ and all classes are reasonably common (and we discuss additional difficulties of this kind of model in Sections 3.1 and 4.1).

1.1. Our Contributions

In this paper we propose relaxing the proportional odds assumption as in (3), but controlling the amount of relaxation by placing ℓ_1 penalties on the differences in weights corresponding to the same features in adjacent β_k vectors, in a way that is reminiscent of the fused lasso (Tibshirani et al., 2005). This model allows us to borrow strength from outcomes where data is much more abundant to improve rare probability estimates when outcomes are much more rare without making the strong assumption that the weights in these models are exactly equal. In particular, it allows for the proportional odds model to hold for some specific features in some adjacent pairs of decision boundaries, but not others.

We formalize the intuitive argument we outline above—that the proportional odds model allows for precise estimation of the β vector as long as at least one decision boundary is surrounded by reasonably well-balanced outcomes, and this allows for improved estimation of rare probabilities at the end of the sequence—through theoretical results in Section 2. Motivated by this argument but skeptical of the proportional odds assumption holding exactly, we propose PRESTO in Section 3 and prove that it consistently estimates $\beta_1, \dots, \beta_{K-1}$ under a sparsity assumption in Section 3.1. In Section 4 we demonstrate through synthetic and real data experiments that PRESTO can outperform both logistic regression on the rare class and the proportional odds model, both in settings where the differences in adjacent β_k vectors are sparse, as PRESTO assumes, and in settings where these differences are not sparse. Before we move on from the introduction, we review related literature.

1.2. Related Work

The difficulty of classification with class imbalance has been well-known for decades. Johnson & Khoshgoftaar (2019) provide a recent review focusing on deep learning methods for handling class imbalance, and they also provide references for many other ways of dealing with class imbalance. One particularly closely related work is Owen (2007), which explores how logistic regression handles a vanishingly rare class. A particularly popular approach, SMOTE (Chawla et al., 2002), has its own recent review paper (Fernández et al., 2018).

Tutz & Gertheiss (2016) discuss the possibility of penalizing differences in weights between adjacent models, including

briefly proposing an ℓ_1 penalty between weights in corresponding categories for proportional hazard models, though this is not the focus of their article and they only mention the idea very briefly without investigating it.

Wurm et al. (2021) propose a generalization of a proportional odds model (and implement it in the R package `ordinalNet`) that allows for the possibility that adjacent categories have equal (or very close) weights, but their method differs from ours. The most closely related model Wurm et al. propose is an over-parameterized *semi-parallel* model with both a matrix of separate parameters for each level, an approach reminiscent of Peterson & Harrell Jr (1990). This results in more flexible, less structured models than our approach, which assumes similarity between adjacent β_k vectors. Further, Wurm et al. (2021) do not investigate the theoretical properties of their model, or the use of their model for improving estimates of rare event probabilities.

Ugba et al. (2021) and Pöbnecker & Tutz (2016) implement an ℓ_2 rather than ℓ_1 penalty between weights in models for adjacent decision boundaries. However, these works also focus on ordinal regression more generally, while we focus both theoretically and in simulations on leveraging common classes to improve estimated probabilities of rare events. Further, the ℓ_1 penalty, which imposes sparse differences, allows the proportional odds assumption to hold for some features and decision boundaries and not others, while the ℓ_2 ridge penalty used by Ugba et al. (2021) (and previously proposed by Tutz & Gertheiss 2016, Section 4.2.2) relaxes the proportional odds assumption for all features but regularizes the relaxation. The ℓ_2 group lasso penalty used by Pöbnecker & Tutz (2016) can impose the proportional odds assumption for a given feature either at all decision boundaries or none of them, making it less flexible than PRESTO.

Besides the fused and generalized (Tibshirani & Taylor, 2011) lasso, our work relates more specifically to the generalized fused lasso (Höfling et al., 2010; Xin et al., 2014). Xin et al. (2014) in particular propose and analyze an algorithm to solve a class of optimization problems similar to the PRESTO optimization problem, (7), with ℓ_1 fusion penalties. In contrast to the present work, Xin et al. (2014) focus almost entirely on the properties of their algorithm. Further, PRESTO lies outside their class of optimization problems because PRESTO directly penalizes only the coefficients in the first decision boundary, not all of the coefficients. This distinction is central to our proof strategy for Theorem 3.1; Xin et al. (2014) do not prove the consistency of their method. Viallon et al. (2013) provide theoretical results for the generalized fused lasso specifically in the cases of linear and logistic regression, though not for ordinal regression. Viallon et al. (2016) prove theoretical results for a

broader class of generalized linear models that still does not include the proportional odds model or a generalization like PRESTO. Lastly, Ekvall & Bottai (2022) prove theoretical results for a class of models in which PRESTO can be expressed, and indeed we leverage their results in proving our own theory, though they do not directly consider fusion penalties.

2. Motivating Theory

We present the following theoretical results to motivate PRESTO. The thrust of our motivation is as follows:

1. Logistic regression does arbitrarily badly as class imbalance worsens (Theorem A.1 in the appendix).
2. However, as one would expect, a logistic regression model’s ability to estimate probabilities improves when the parameters β are known (Theorem 2.2).
3. The proportional odds model allows for precise estimation of β as long as two adjacent classes are reasonably common, even if the remaining classes are arbitrarily rare (Theorem 2.3).
4. Taking 2 and 3 together, our conclusion is that we can better estimate probabilities of rare events by using a method that leverages data from decision boundaries between abundant classes to better estimate decision boundaries near rare classes. (Both the proportional odds model and PRESTO leverage the data in this way.)

Before we present our results, we discuss the metrics we will use in our results and some of the assumptions we will make.

2.1. Preliminaries

Our goal is to characterize and compare the prediction error of estimated conditional probabilities of a rare class from both logistic regression and the proportional odds model. There are many settings where estimating rare probabilities accurately (as opposed to, for example, predicting class labels accurately) is important. For example, in online advertising, advertisers bid on the price to display an ad to a given user. Advertisers could bid optimally if they knew the true probability each user would click a given ad, so they’d like to estimate these probabilities as precisely as possible (He et al., 2014; Zhang et al., 2014). Another example is public policy, where scarce resources may be allocated based on estimated probability of bad outcomes (Von Wachter et al., 2019). To prioritize optimally, precisely estimated probabilities are needed, not just accurate labels.

A natural metric in an estimation setting is mean squared error, $\mathbb{E} \left[(\hat{\pi}(\mathbf{x}) - \pi(\mathbf{x}))^2 \right]$, where $\pi(\mathbf{x})$ is the actual probability of a rare event conditional on \mathbf{x} and $\hat{\pi}(\mathbf{x})$ is an estimate. Further, we leverage asymptotic statistics and present results for *large-sample* estimators. We define the notions of asymptotic mean squared error we will use below:

Definition 2.1. Let $\hat{\theta}_n$ be a maximum likelihood estimator for a parameter $\theta \in \mathbb{R}$ from a sample size of n . Under regularity conditions, the sequence of random variables $\{\sqrt{n} \cdot (\hat{\theta}_n - \theta)\}$ converges in distribution to a Gaussian random variable. Then we define the **asymptotic mean squared error** of $\hat{\theta}_n$ to be (suppressing n from the notation)

$$\text{Asym.MSE}(\hat{\theta}) := \mathbb{E} \left[\left(\lim_{n \rightarrow \infty} \sqrt{n} [\hat{\theta}_n - \theta] \right)^2 \right].$$

Asymptotic metrics are commonly used to compare the performance of estimators. The *asymptotic relative efficiency* of two estimators is the ratio of their asymptotic variances,

$$\text{Asym.Var}(\hat{\theta}) := \text{Var} \left(\lim_{n \rightarrow \infty} \sqrt{n} [\hat{\theta}_n - \theta] \right),$$

which is equal to $\text{Asym.MSE}(\hat{\theta}_n)$ for the (asymptotically unbiased) maximum likelihood estimators we consider. See Section 10.1.3 of Casella & Berger (2021), Section 8.2 of van der Vaart (2000), or Section 4.4.5 of Greene (2012) for textbook-level discussions. The asymptotic MSE could also be used as an estimator of the MSE for large (but finite) n , under the heuristic reasoning that for large n ,

$$\begin{aligned} \text{MSE}(\hat{\theta}) &= \frac{1}{n} \mathbb{E} \left[\left(\sqrt{n} \cdot [\hat{\theta} - \theta] \right)^2 \right] \\ &\approx \frac{1}{n} \mathbb{E} \left[\left(\lim_{n \rightarrow \infty} \sqrt{n} \cdot [\hat{\theta} - \theta] \right)^2 \right] \\ &= \frac{1}{n} \text{Asym.MSE}(\hat{\theta}). \end{aligned}$$

See Section 4.4 of Greene (2012), Section 7.3 of Hansen (2022), or Section 3.5 of Wooldridge (2010) for more discussion of this kind of finite-sample estimation using asymptotic quantities.

We briefly present and discuss some of our assumptions.

- **Assumption $X(\mathcal{A})$:** The random vectors $\mathbf{x}_i \in \mathbb{R}^p$ are independent and identically distributed (iid) for $i \in \{1, \dots, n\}$, each with probability measure $dF(\mathbf{x})$ with measurable, bounded support $\mathcal{S} \subset \mathcal{A} \subseteq \mathbb{R}^p$, with $\text{Cov}(\mathbf{X})$ positive definite.
- **Assumption $Y(K)$:** The response $y_i \in \{1, \dots, K\}$ is distributed conditionally on \mathbf{x}_i as in the proportional odds model (1). (Note that if $K = 2$, this is equivalent to the logistic regression model.) All

classes have positive probability for all \mathbf{x} on the support of \mathbf{x}_i (equivalently, the intercepts strictly differ: $\alpha_1 < \dots < \alpha_{K-1}$.)

Assumption $X(\mathcal{A})$ allows a very broad class of distributions, including both discrete and continuous random variables. Notice that the boundedness assumption within $X(\mathcal{A})$ implies that the matrix $\tilde{\mathbf{X}} := (\mathbf{1}, \mathbf{X})$ (where $\mathbf{1}$ is an n -vector of ones) has a finite maximum eigenvalue. When we will refer to it, we call it λ_{\max} and write **Assumption $X(\mathcal{A}, \lambda_{\max})$** .

From (2) we see that in the proportional odds model if the intercepts strictly differ ($\alpha_1 < \dots < \alpha_{K-1}$) then for any \mathbf{x} all of the classes have conditional probability strictly between 0 and 1. That said, if the support of \mathbf{X} is unbounded then all of the probabilities for individual classes can become arbitrarily close to 0 or 1. Under Assumption $X(\mathcal{A})$, however, we can strictly bound quantities like $\sup_{\mathbf{x} \in \mathcal{S}} \{\pi_k(\mathbf{x})\}$ (where $\pi_k(\mathbf{x}) := p_k(\mathbf{x}) - p_{k-1}(\mathbf{x}) = \mathbb{P}(y = k \mid \mathbf{x})$) away from 0 or 1.

Theorem 2.2 holds under Assumption $X([0, \infty)^p)$, though for any bounded $\mathcal{S} \subseteq \mathbb{R}^p$, there is some finite a one could add to each coordinate to shift \mathcal{S} to a subset of $[0, \infty)^p$; Theorem 2.2 would then apply to these translated features.

2.2. Theorem 2.2

Theorem 2.2 suggests a possible way to circumvent the problem of class imbalance. We compare the typical logistic regression intercept estimate $\hat{\alpha}$ to the *quasi-estimated* estimator $\hat{\alpha}_q$ obtained when one estimates only the intercept of the logistic regression model with a known β . We also compare the resulting estimators of conditional probabilities for any $\mathbf{z} \in \mathbb{R}^p$: the usual logistic regression estimator $\hat{\pi}(\mathbf{z})$ and $\hat{\pi}_q(\mathbf{z})$, the estimator when β is known. Theorem 2.2 proves the reasonable intuition that $\hat{\alpha}_q$ must be a better estimator than $\hat{\alpha}$, and likewise for $\hat{\pi}_q(\mathbf{z})$ and $\hat{\pi}(\mathbf{z})$.

Theorem 2.2. Assume $X([0, \infty)^p, \lambda_{\max})$ and $Y(2)$ hold. Let $\pi(\mathbf{x}) := \mathbb{P}(y = 2 \mid \mathbf{x})$, and let $\pi_{\min} := \inf_{\mathbf{x} \in \mathcal{S}} \{\pi(\mathbf{x}) \wedge 1 - \pi(\mathbf{x})\}$. Then

1.

$$\frac{\text{Asym.MSE}(\hat{\alpha}) - \text{Asym.MSE}(\hat{\alpha}_q)}{[\text{Asym.MSE}(\hat{\alpha}_q)]^2} \geq \Delta$$

where

$$\Delta := \frac{4\pi_{\min}^2(1 - \pi_{\min})^2 \|\mathbb{E}[\mathbf{X}]\|_2^2}{\lambda_{\max}},$$

and

2. For any $\mathbf{z} \in \mathbb{R}^p \setminus \{\mathbf{z}^*\}$, where

$$\mathbf{z}^* := \frac{\mathbb{E}[\mathbf{X}\pi(\mathbf{X})[1 - \pi(\mathbf{X})]]}{\mathbb{E}[\pi(\mathbf{X})[1 - \pi(\mathbf{X})]]},$$

it holds that

$$\text{Asym.MSE}(\hat{\pi}_q(\mathbf{z})) < \text{Asym.MSE}(\hat{\pi}(\mathbf{z})).$$

(For \mathbf{z}^* , the above holds with \leq rather than $<$.)

Examining the first result, it is sensible that the lower bound for the gap between the asymptotic variances of the two estimators vanishes as π_{\min} vanishes because if $\min\{\pi_1(\mathbf{x}), 1 - \pi_1(\mathbf{x})\}$ becomes very small on the bounded support, then the imbalance between the two classes potentially becomes very large, and estimating the intercept becomes difficult regardless of whether or not β is known. As the class balance improves (π_{\min} becomes closer to its upper bound $1/2$), the guaranteed gap between $\text{Asym.MSE}(\hat{\alpha})$ and $\text{Asym.MSE}(\hat{\alpha}_q)$ becomes larger.

In addition to formally verifying intuition, Theorem 2.2 also quantifies the estimation gap between the rare intercept estimators in terms of noteworthy parameters and shows that the assumptions needed for this intuition to hold are minimal.

2.3. Theorem 2.3

Theorem 2.2 suggests that if only we could estimate β very well, we could improve our estimated probabilities even in the face of class imbalance. Theorem 2.3 suggests a way to leverage abundant data among other classes to do this.

In the proportional odds model (1), \mathbb{R}^p is partitioned into regions with separating hyperplanes defined by β , which we note are Bayes decision boundaries: for $\mathbf{x} \in \mathbb{R}^p$ such that $\alpha_k + \beta^\top \mathbf{x} = 0$, we have $p_k(\mathbf{x}) = 1/2$.

Consider the setting of ordered categorical data generated by the proportional odds model with categories 1 and 2 similarly common over the support of a bounded distribution of \mathbf{x}_i and categories $3, \dots, K$ all rare. In this setting, for many of the observed values of \mathbf{x}_i , the probabilities of being in class 1 or 2 will both be close to $1/2$. Intuitively it should be relatively easy to estimate β and α_1 , the parameters that define the Bayes decision boundary between classes 1 and 2, and therefore $p_1(\mathbf{x}_i)$. Theorem 2.2 suggests this should help us in estimating the rare class probabilities. In Theorem 2.3, we prove that even if class K becomes arbitrarily rare, as long as the first two classes are reasonably well balanced, the proportional odds model still learns β quite well.

Theorem 2.3. *Assume $X(\mathbb{R}^p)$ and $Y(3)$ hold. Assume for all $\mathbf{x} \in \mathcal{S}$ it holds that $|\pi_k(\mathbf{x}) - 1/2| \leq \Delta$ for $k \in \{1, 2\}$ for some $\Delta \in (0, 1/2)$ and let $M := \sup_{\mathbf{x} \in \mathcal{S}} \|\mathbf{x}\|_2$ (notice that $X(\mathbb{R}^p)$ ensures that $M < \infty$). Suppose $\sup_{\mathbf{x} \in \mathcal{S}} \{\pi_3(\mathbf{x})\} =$*

π_{rare} , where π_{rare} is no greater than

$$\min \left\{ \frac{1}{2} \left(\frac{1}{2} - \Delta \right) \left(\frac{1}{2} + \Delta \right), \frac{\lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right)}{3M^2(M+2)} \right\}, \quad (5)$$

where $\lambda_{\min}(\cdot)$ denotes the minimum eigenvalue of \cdot and $I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}}$ is a symmetric matrix composed of terms from the Fisher information matrix for the proportional odds model (see the definitions of these terms in Equations 8, 9, and 10 in the appendix). Then there exists $C < \infty$ not depending on π_{rare} such that for any fixed $\mathbf{v} \in \mathbb{R}^p$,

$$\frac{1}{\|\mathbf{v}\|_2^2} \text{Asym.MSE} \left(\mathbf{v}^\top \hat{\beta}^{\text{prop. odds}} \right) \leq C.$$

To give an example of applying this result, consider the choice $\mathbf{v} = (1, 0, \dots, 0)$. Then we have that $\text{Asym.MSE} \left(\hat{\beta}_1^{\text{prop. odds}} \right) \leq C$, so $\hat{\beta}_1^{\text{prop. odds}}$ (or any other estimated coefficient) has bounded asymptotic mean squared error. In contrast to logistic regression, the proportional odds model still learns β within a fixed precision even as π_{rare} vanishes.

Remark 2.4. We briefly discuss the upper bound (5). For this bound to make sense, it must hold that the symmetric matrix $I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}}$ is positive definite so that its minimum eigenvalue is strictly positive. The matrix $\mathbf{S} := I_{\beta\beta} - \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}}$ is the Schur complement of $I_{\alpha_1\alpha_1} = M_1$ in the submatrix

$$\begin{pmatrix} I_{\alpha_1\alpha_1} & I_{\beta\alpha_1}^\top \\ I_{\beta\alpha_1} & I_{\beta\beta} \end{pmatrix} \quad (6)$$

of the Fisher information matrix $I^{\text{prop. odds}}(\alpha, \beta)$ for the proportional odds model (see Lemma D.1 in the appendix). Note (6) is a principal submatrix of the positive definite $I^{\text{prop. odds}}(\alpha, \beta)$, so is positive definite by Observation 7.1.2 in Horn & Johnson (2012). From (8) we also know that $I_{\alpha_1\alpha_1} > 0$, so \mathbf{S} is positive definite by Theorem 1.12 in Zhang (2005). It seems plausible that

$$I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} = \mathbf{S} - \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}}$$

is also positive definite because $I_{\beta\beta}$ is the inverse of the asymptotic covariance matrix of $\hat{\beta}^{\text{ideal}}$, the maximum likelihood estimator of β when α_1 and α_2 are known. We expect that $\text{Cov}(\hat{\beta}^{\text{ideal}})$ would be small (and the eigenvalues of $I_{\beta\beta}$ would be large) in this setting because we can estimate β well due to the abundant observations in classes 1 and 2 (ensured if Δ is not too large), so we should be able to learn the decision boundary between these classes well. If the

eigenvalues of $I_{\beta\beta}$ are indeed large, it might be reasonable to expect $I_{\beta\beta} - 2\frac{I_{\beta\alpha_1}I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}}$ to be positive definite. In Sections E.1 and E.2 in the appendix, we present more detailed analysis as well as the results of synthetic experiments that indicate that it is plausible both that $I_{\beta\beta} - 2\frac{I_{\beta\alpha_1}I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}}$ is positive definite and that the upper bound (5) is reasonable.

3. Predicting Rare Events by Shrinking Towards proportional Odds (PRESTO)

Theorems 2.2 and 2.3 suggest a path to improve estimated probabilities for a rare event that is at the end of an ordered sequence: use the more common events that come before it to improve the estimation of the decision boundary affecting the rare class. In practice, however, the proportional odds model assumption is strong and unlikely to hold in many settings. PRESTO allows for this assumption to be relaxed; instead of assuming the β vectors governing the decision boundaries are identical, we assume they are in general different, but with differences that are (approximately) sparse.

One concrete model to motivate this is a relaxation of (4) along the lines of (3). Suppose that $U_1 := U$ as defined in (4) (with $\beta_1 = \beta$), and it still holds that an observation is in class 1 if $U_1 \geq -\alpha_1$. However, for $k \in \{2, \dots, K-1\}$, outcome k is observed if and only if $-\alpha_k \leq U_k < -\alpha_{k-1} + \psi_k^\top \mathbf{x}$ for sparse vectors $\psi_2, \dots, \psi_{K-1} \in \mathbb{R}^p$ satisfying $\psi_k = \beta_k - \beta_{k-1}$, so $U_k = U_{k-1} + \psi_k^\top \mathbf{x}$ for $k \in \{2, \dots, K-1\}$. Note that this is within the scope of (3), but we assume a structure on the differing β_k vectors rather than allowing for arbitrary differences.

Assuming sparse differences in adjacent β_k vectors in this way suggests the following optimization problem for data $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top$ and $\mathbf{y} = (y_1, \dots, y_n)$:

$$\begin{aligned} \arg \min_{\beta, \alpha} \left\{ -\frac{1}{n} \sum_{i=1}^n \log \left[F \left(\alpha_{y_i} + \beta_{y_i}^\top \mathbf{x}_i \right) \right. \right. \\ \left. \left. - F \left(\alpha_{y_i-1} + \beta_{y_i-1}^\top \mathbf{x}_i \right) \right] \right. \\ \left. + \lambda_n \left(\sum_{j=1}^p |\beta_{j1}| + \sum_{j=1}^p \sum_{k=2}^{K-1} |\beta_{jk} - \beta_{j,k-1}| \right) \right\}, \end{aligned} \quad (7)$$

where we define $\alpha_K := \infty, \alpha_0 := -\infty$ and $\beta_0 := \mathbf{0}$. The penalties on the $|\beta_{j1}|$ terms are sufficient to regularize all of the weights given the penalties on the difference terms starting from the β_1 vector, improving parameter estimation. Like the proportional odds model and the generalized lasso (Tibshirani & Taylor, 2011) optimization problem, (7) is strictly convex if and only if $\alpha_{y_i} + \beta_{y_i}^\top \mathbf{x}_i > \alpha_{y_i-1} + \beta_{y_i-1}^\top \mathbf{x}_i$ for all i (Pratt, 1981).

This can be violated if the decision boundaries, which are not parallel, cross in the support of \mathbf{X} . In Section 4.1, we discuss the practical issues this presents when implementing relaxed proportional odds models like PRESTO, and in the next section, we prove PRESTO is consistent relying in part on an assumption that these decision boundaries do not cross in the support of \mathbf{X} . See Appendix H for details on how we estimate PRESTO in practice.

3.1. Theoretical Analysis

In this section, we present Theorem 3.1, which shows that PRESTO is a consistent estimator of $\beta_1, \dots, \beta_{K-1}$ under suitable assumptions. Before stating Theorem 3.1, we present and briefly discuss some of the new assumptions we will make.

- **Assumption $S(s, c)$:** The distribution of $y_i \mid \mathbf{x}_i$ is distributed according to the PRESTO likelihood (7), where the true coefficients $\theta_* = (\beta_1^\top, \psi_2^\top, \dots, \psi_{K-1}^\top)^\top \in \mathbb{R}^{p(K-1)}$ are s -sparse (have s nonzero entries for a fixed s not increasing in n or p). Further, $\|\theta_*\|_\infty \leq c$ for a fixed c .
- **Assumption $T(c)$:** For all small enough $\rho > 0$, for all $\theta \in \mathbb{R}^{p(K-1)}$ with $\|\theta - \theta_*\|_1 \leq \rho$ it holds that none of the decision boundaries defined by θ and the true $\alpha_1, \dots, \alpha_{K-1}$ cross in \mathcal{S} . Also, $\max_{k \in \{1, \dots, K-1\}} |\alpha_k| \leq c$.

The fixed sparsity assumption $S(s, c)$ is helpful theoretically and also because without it in higher dimensions it becomes increasingly difficult to have nonparallel decision boundaries that do not cross. The first part of Assumption $T(c)$ can be interpreted to mean that none of the decision boundaries cross “too closely” to \mathcal{S} . Other than these aspects, Assumptions $S(s, c)$ and $T(c)$ are mild.

Theorem 3.1. *In a setting with fixed $K \geq 3$ and $p = p_n \rightarrow \infty$ as $n \rightarrow \infty$ and satisfying $p_n \leq C_1 n^{C_2}$ for some $C_1 > 0$ and $C_2 \in (0, 1)$, consider estimating PRESTO with penalty $\lambda_n = C_3 \log(p_n[K-1])/n$ for some $C_3 > 0$. Suppose Assumption $X(\mathbb{R}^{p_n})$ holds and there is some $C_4 < \infty$ such that $\sup_{\mathbf{x} \in \mathcal{S}} \|\mathbf{x}\|_\infty \leq C_4$ and Assumptions $S(s, C_4)$ and $T(C_4)$ hold. Assume for some fixed $b > 0$ it holds that $\lambda_{\min}^* := \min_{k \in \{1, \dots, K\}} \lambda_{\min}(\Sigma_k) > b$, where $\Sigma_k := \mathbb{E}[\mathbf{x}_i \mathbf{x}_i^\top \mid y_i = k]$. Then PRESTO is a consistent estimator of $\beta_1, \dots, \beta_{K-1}$.*

Theorem 3.1 shows that under fairly mild regularity conditions and a sparsity assumption in a high-dimensional setting, PRESTO consistently estimates all of the decision boundaries. That is, it is consistent both if the proportional odds assumption holds and in more flexible settings, where the proportional odds model is unrealistic, under sparsity.

Theorem 2.2 suggests this should be helpful for estimating rare class probabilities. The proof of Theorem 3.1 leverages recent theory developed for ℓ_1 -penalized ordinal regression (Ekvall & Bottai, 2022).

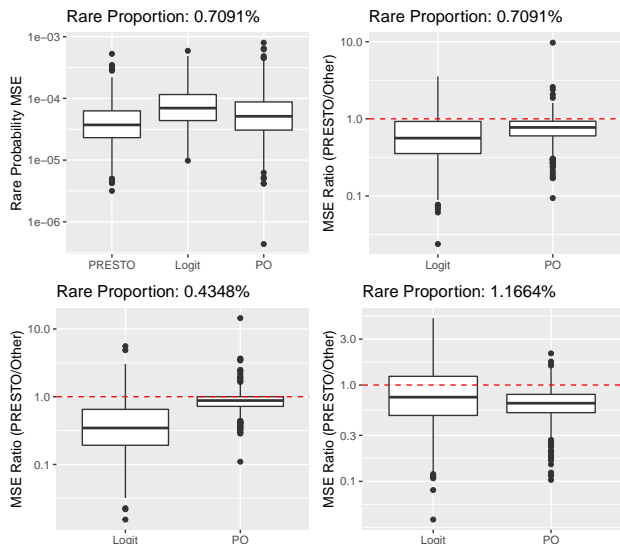


Figure 1. Top left: MSE of estimated rare class probabilities for each method across all $n = 2500$ observations, across 700 simulations, in sparse differences simulation setting of Section 4.1, for intercept setting yielding rare class proportions of about 0.71% on average and sparsity $1/2$. Remaining plots: ratios of MSE for PRESTO divided by MSE of each other method for each of three sets of intercepts with sparsity $1/2$ (PRESTO performs better if ratio is less than 1). All plots on log scale.

4. Experiments

To illustrate the efficacy of PRESTO, we conduct two synthetic experiments and also examine two real data sets. In Section 4.1, we generate random \mathbf{y} that have conditional probabilities based on a relaxation of the proportional odds model with sparse differences between adjacent decision boundary parameter vectors, rather than parameterizing all decision boundaries with the same β . This setting is well-suited to PRESTO. In Section 4.2, we show that PRESTO also performs well in a less favorable setting, where the differences between adjacent decision boundaries are instead dense; nonetheless, PRESTO still outperforms logistic regression and proportional odds models. In Section 4.3 we compare the performance of PRESTO to logistic regression and the proportional odds model at estimating rare probabilities in a real data experiment. Finally, in Section B of the appendix we conduct a second real data experiment on a data set of patients diagnosed with diabetes, where we vary the rarity of the outcome of interest. See Section H of the appendix for all implementation de-

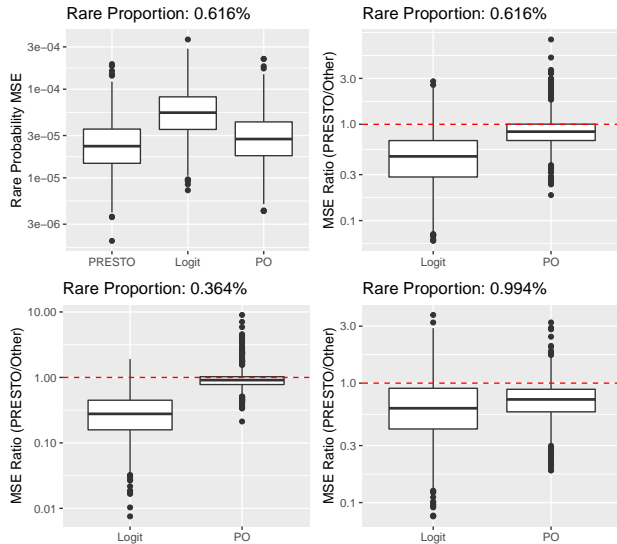


Figure 2. Same plots as in Figure 1, but for uniform differences synthetic experiment in Section 4.2.

tails. The code generating all plots and tables is available at <https://github.com/gregfaletto/presto>.

4.1. Synthetic Data: Sparse Differences Setting

We repeat the following procedure for 700 simulations. First we generate data using $n = 2500$, $p = 10$, and $K = 4$. We draw a random $\mathbf{X} \in [-1, 1]^{n \times p}$, where $X_{ij} \sim \text{Uniform}(-1, 1)$ for all $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, p\}$. Then $\mathbf{y} \in \{1, \dots, K\}^n$ is generated according to a relaxation of the proportional odds model; instead of (1), we generate probabilities according to (3) where the β_k are generated in the following way for sparsity settings of $\eta \in \{1/3, 1/2\}$: first, we generate β_1 by taking the vector $(0.5, \dots, 0.5)^\top$, but setting all of the entries equal to 0 randomly with probability $1 - \eta$ for each entry independently. Then we set $\beta_k = \beta_{k-1} + \psi_k$, $k \in \{2, \dots, K - 1\}$, where $\psi_k \in \mathbb{R}^p$ are iid random vectors for each $k \in \{2, \dots, K - 1\}$ generated according to the following distribution:

$$\psi_{kj} = \begin{cases} 0, & \text{with probability } 1 - \eta, \\ 0.5, & \text{with probability } \eta/2, \\ -0.5, & \text{with probability } \eta/2, \end{cases} \quad j \in \{1, \dots, p\}.$$

We consider three possible sets of intercepts: $\alpha = (0, 3, 5)$, $(0, 3.5, 5.5)$, and $(0, 4, 6)$, so that the first two categories are common and the remaining categories are rare. The final rare class is the one of interest; in the three settings, the average proportions of observations falling in the rare class are 1.00%, 0.62%, and 0.37%, respectively, for

the $\eta = 1/3$ setting and 1.17%, 0.71%, and 0.43% for the $\eta = 1/2$ setting.

The fact that the decision boundaries may cross in the support of \mathbf{X} , which would mean that for such \mathbf{x} some class probabilities are defined to be negative, puts practical limits on the magnitude of ψ_k in simulations. (See Section 3.6.1 of Agresti 2010 for a discussion of this point.) Also, for this reason, in each simulation we check whether or not the conditional probabilities are positive for each class for every sampled \mathbf{x} ; if not, we generate new $\psi_2, \dots, \psi_{K-1}$ for a limited number of iterations, ending the simulation study in failure if no suitable ψ_k can be found in a reasonable number of attempts. The parameters we used generated positive probabilities for all observations across all simulations.

We then estimate a model for each method; for logistic regression, we estimate the binary classification problem of whether or not each observation is in class K , and for proportional odds and PRESTO, we fit a full model on all K responses. For PRESTO, we use 5-fold cross-validation to choose a value of λ_n among 20 choices, selecting the λ_n with the best out-of-fold Brier score (other metrics, like negative log likelihood, failed because some values of λ_n in some folds resulted in models yielding negative probabilities, so these other metrics were undefined). The 20 candidate values of λ_n are generated in the following way: the largest λ_n value, $\lambda_n^{(20)}$, is the smallest λ_n for which all of the estimated sparse differences equal 0; the smallest λ_n value is set to $\lambda_n^{(1)} = 0.01 \cdot \lambda_n^{(20)}$, and the remaining λ_n values are generated at equal intervals on a logarithmic scale between these two values.

Each of these models yields estimated probabilities that each observation lies in class K . In the final step of each simulation run, we compute the mean squared error of these estimated probabilities for each method.

In Figure 1, we show boxplots of the empirical mean squared errors for each method in the setting where the rare class is observed in 0.71% of observations when $\eta = 1/2$. In order to see how the methods compare pairwise on each simulation, we also show boxplots of the ratio between the mean squared error of PRESTO and the other two methods in each of the three simulation settings. We also conduct one-tailed paired t -tests of the alternative hypothesis that the mean MSE for PRESTO is lower than each of the competitor methods in each setting; all 12 of the p -values (provided in Table 3 of Appendix C) are below 0.01. Finally, in Appendix C we also provide the means and standard errors for the MSE of each method in each simulation setting in Table 4, as well as boxplots like the one in the top left corner of Figure 1 for the other two intercept settings and all boxplots for the $\eta = 1/3$ setting.

We see that PRESTO typically estimates these rare proba-

bilities better than logistic regression, which despite being correctly specified struggles with class imbalance and does not draw strength from estimating the easier decision boundary between classes 1 and 2, and the proportional odds model, whose assumptions are not satisfied in this setting.

4.2. Synthetic Data: Dense Differences Setting

In real data sets the differences between adjacent decision boundary parameter vectors may not always be exactly sparse, so we conduct another synthetic experiment in the same way as in Section 4.1, except $\beta_{1j} \sim \text{Uniform}(-.5, .5)$ and each $\psi_{kj} \sim \text{Uniform}(-.5, .5)$, iid across $j \in \{1, \dots, p\}$ and $k \in \{2, \dots, K - 1\}$. We also add an extra intercept setting of (0, 2.5, 4.5). This yields average rare class proportions of 0.99%, 0.62%, and 0.36% using the same intercepts as the experiments in Section 4.1 and 1.60% in the new intercept setting. The uniformly distributed differences can be considered “approximately” sparse in the sense that while no deviations will exactly equal 0, some will be large and important to estimate, and some will be essentially negligible.

Figure 2 and Table 1 summarize the results, along with additional figures and tables in Appendix C. We again see that PRESTO outperforms both competitor methods by statistically significant margins.

Table 1. Calculated p -values for one-tailed paired t -tests for uniform differences simulation setting of Section 4.2 testing the alternative hypothesis that PRESTO’s rare probability MSE is less than each competitor method in each rarity setting. (Statistically significant p -values indicate better performance for PRESTO).

Rare Class Proportion	Logit p -value	PO p -value
1.6%	< 1e-10	< 1e-10
0.99%	< 1e-10	< 1e-10
0.62%	< 1e-10	< 1e-10
0.36%	< 1e-10	0.000242

4.3. Real Data Experiment

We conduct a real data experiment using the `soup` data set from the R `ordinal` package (R. H. B. Christensen, 2019). The data come from a study (Christensen et al., 2011) of participants who tasted soups and responded whether they thought each soup was a reference product they had previously been familiarized with or a new test product. The respondents also stated how sure they were in their response on a three-level scale, yielding a total of $K = 6$ possible ordered outcomes for $n = 1847$ observations. The outcome of interest corresponds to the respondent being sure the tasted soup was the reference and is observed in 228 observations (about 12% of the total). All of the features

are categorical, and after one-hot encoding we have $p = 22$ binary features related to the soup, the respondent, and the testing environment¹. This may be a promising setting for PRESTO because, while the responses have a well-defined ordering, it’s plausible that different features could have varying impacts at different levels of respondent certainty.

We complete the following procedure 350 times: first, we randomly split the data into training (90% of the data) and test (10%) sets. We estimate models using PRESTO, logistic regression, and the proportional odds model on the training data and evaluate on the test set.

We are interested in the accuracy of the rare class probabilities, but we can’t evaluate rare probability MSE directly since we don’t observe the true probabilities. Brier score could be a reasonable proxy, but it is known to be a poor metric in the presence of class imbalance (Benedetti, 2010). Instead we estimate rare probability MSE using the following procedure. For each method, we sort the estimated test set rare class probabilities in ascending order and assign the observations into 10 bins: the first 1/10 observations go in the first bin, and so on. Then we estimate the mean squared error of the estimated probabilities by $\frac{1}{n} \sum_{i=1}^n (\hat{\pi}_1^{(i)} - o_{b(i)})^2$, where $\hat{\pi}_1^{(i)}$ is the estimated rare class probability for observation i and $o_{b(i)}$ is the observed rare class proportion in the bin containing observation i . This is similar to *expected calibration error* (Naeini et al., 2015), though we use squared error rather than absolute error. 10 equal frequency bins follows the default of the R `CalibratR` package that implements expected calibration error (Schwarz & Heider, 2018).

By this metric, the mean error for PRESTO is 0.0096, 0.0157 for logistic regression and 0.0135 for proportional odds. Figure 3 displays boxplots of the results as in the synthetic experiments which indicate that PRESTO typically outperforms the other methods. We do not report p -values or standard errors since the observed samples are dependent (random splits of the same data set).

5. Conclusion

By leveraging data from earlier decision boundaries, but relaxing the rigid proportional odds assumption, PRESTO can substantially improve estimation of the probability of rare events, even when the assumption of sparse differences between adjacent decision boundary weight vectors does not exactly hold. Future work could explore ℓ_1 penalties for the coefficients themselves, not just the differences between the coefficients, to allow for simultaneous feature selection

¹The categorical predictors `PRODID` and `RESP` are omitted because in some splits not all levels of these features are observed in the training set, making it impossible to estimate parameters for these features.




Figure 3. Left: Estimated MSEs of estimated rare class probabilities for each method across 350 random draws of training and test sets in real data experiment from Section 4.3. Right: ratios of estimated MSE for PRESTO divided by MSE of each other method (PRESTO performs better if ratio is less than 1).

and model estimation. Inference for PRESTO could also be possible by extending the method for exact post-selection inference for the generalized lasso path by Hyun et al. (2018), or similar work on the fused lasso by Chen et al. (2022), to our generalized linear model setting. Future work could also explore the empirical performance of PRESTO in even more depth, perhaps by using large-scale real world data sets like those used in Duncan & Elkan (2015).

There are other possible extensions that could improve estimation. For example, we set the first decision boundary as the one that is directly penalized, with differences from this boundary assumed to be sparse. This makes sense if the classes become increasingly rare and the first decision boundary is the most balanced. However, it may make more sense to directly penalize whichever decision boundary has the best balance of observed responses on each side. Penalizing the differences from this boundary might improve estimation since this decision boundary ought to be the easiest to estimate. This could improve estimation in settings like the real data experiment from Section 4.3 where the most balanced decision boundary is closer to the center of the responses.

Also, in cases where the final categories are very rare, a better bias/variance tradeoff might be achieved by reimposing the proportional odds assumption, imposing an exact equality constraint for the last few decision boundaries. In these settings, data might be too rare to hope for better estimation by relaxing the proportional odds assumption even with regularization.

Lastly, in Section H of the appendix we discuss possible faster approaches than the one used in the present work for solving the PRESTO optimization problem.

References

- Agresti, A. *Analysis of ordinal categorical data*, volume 656. John Wiley & Sons, 2010.
- Armstrong, B. G. and Sloan, M. Ordinal regression models for epidemiologic data. *American Journal of Epidemiology*, 129(1):191–204, 1989.
- Arnold, T. B. and Tibshirani, R. J. Efficient Implementations of the Generalized Lasso Dual Path Algorithm. *Journal of Computational and Graphical Statistics*, 25(1):1–27, 2016. ISSN 15372715. doi: 10.1080/10618600.2015.1008638.
- Benedetti, R. Scoring rules for forecast verification. *Monthly Weather Review*, 138(1):203–211, 2010.
- Bickel, P. J., Ritov, Y., and Tsybakov, A. B. Simultaneous Analysis of LASSO and Dantzig Selector. *The Annals of Statistics*, 37(4):1705–1732, 2009. doi: 10.1214/08-AOS620. URL https://projecteuclid.org/libproxy1.usc.edu/download/pdfview_1/euclid.aos/1245332830.
- Bien, J. The simulator: an engine to streamline simulations. *arXiv preprint arXiv:1607.00021*, 2016.
- Brant, R. Assessing proportionality in the proportional odds model for ordinal logistic regression. *Biometrics*, pp. 1171–1178, 1990.
- Cameron, A. C. and Trivedi, P. K. *Microeconometrics: methods and applications*. Cambridge university press, 2005.
- Casella, G. and Berger, R. L. *Statistical inference*. Cengage Learning, 2021.
- Chawla, N. V., Bowyer, K. W., Hall, L. O., and Kegelmeyer, W. P. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.
- Chen, Y., Jewell, S., and Witten, D. More powerful selective inference for the graph fused lasso. *Journal of Computational and Graphical Statistics*, pp. 1–11, 2022.
- Christensen, R. H. B., Cleaver, G., and Brockhoff, P. B. Statistical and Thurstonian models for the A-not A protocol with and without sureness. *Food Quality and Preference*, 22(6):542–549, 2011. ISSN 09503293. doi: 10.1016/j.foodqual.2011.03.003. URL <http://dx.doi.org/10.1016/j.foodqual.2011.03.003>.
- Cordeiro, G. M. and McCullagh, P. Bias Correction in Generalized Linear Models. *Journal of the Royal Statistical Society: Series B (Methodological)*, 53(3):629–643, 1991. doi: 10.1111/j.2517-6161.1991.tb01852.x.
- Duncan, B. A. and Elkan, C. P. Probabilistic modeling of a sales funnel to prioritize leads. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '15, pp. 1751–1758, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450336642. doi: 10.1145/2783258.2788578. URL <https://doi.org/10.1145/2783258.2788578>.
- Ekvall, K. and Bottai, M. Concave likelihood-based regression with finite-support response variables. *Biometrics*, (March):1–12, 2022. ISSN 0006-341X. doi: 10.1111/biom.13760.
- Fernández, A., Garcia, S., Herrera, F., and Chawla, N. V. Smote for learning from imbalanced data: progress and challenges, marking the 15-year anniversary. *Journal of artificial intelligence research*, 61:863–905, 2018.
- Greene, W. H. *Econometric Analysis*. Pearson Education, 7th edition, 2012.
- Hansen, B. *Econometrics*. Princeton University Press, 2022. ISBN 9780691235899. URL <https://books.google.com/books?id=Pte7zgEACAAJ>.
- Hastie, T., Tibshirani, R., and Wainwright, M. Statistical learning with sparsity. *Monographs on statistics and applied probability*, 143:143, 2015.
- He, X., Pan, J., Jin, O., Xu, T., Liu, B., Xu, T., Shi, Y., Atallah, A., Herbrich, R., Bowers, S., et al. Practical lessons from predicting clicks on ads at facebook. In *Proceedings of the eighth international workshop on data mining for online advertising*, pp. 1–9, 2014.
- Hlavac, M. *stargazer: Well-Formatted Regression and Summary Statistics Tables*. Social Policy Institute, Bratislava, Slovakia, 2022. URL <https://CRAN.R-project.org/package=stargazer>. R package version 5.2.3.
- Höfling, H., Binder, H., and Schumacher, M. A coordinate-wise optimization algorithm for the fused lasso. *arXiv preprint arXiv:1011.6409*, 2010.
- Horn, R. A. and Johnson, C. R. *Matrix Analysis*. Cambridge University Press, 2 edition, 2012. doi: 10.1017/9781139020411.
- Hutson, G., Laldin, A., and Velásquez, I. *MLDataR: Collection of Machine Learning Datasets for Supervised Machine Learning*, 2022. URL <https://CRAN.R-project.org/package=MLDataR>. R package version 0.1.3.
- Hyun, S., G’Sell, M., and Tibshirani, R. J. Exact post-selection inference for the generalized lasso path. *Electronic Journal of Statistics*, 12(1):1053–1097, 2018.

- Johnson, J. M. and Khoshgoftaar, T. M. Survey on deep learning with class imbalance. *Journal of Big Data*, 6(1):27, 2019. doi: 10.1186/s40537-019-0192-5. URL <https://doi.org/10.1186/s40537-019-0192-5>.
- Ko, S., Yu, D., and Won, J.-H. Easily parallelizable and distributable class of algorithms for structured sparsity, with optimal acceleration. *Journal of Computational and Graphical Statistics*, 28(4):821–833, 2019.
- Lehmann, E. L. *Elements of large-sample theory*. Springer, 1999.
- McCullagh, P. Regression models for ordinal data. *Journal of the Royal Statistical Society: Series B (Methodological)*, 42(2):109–127, 1980.
- Naeini, M. P., Cooper, G., and Hauskrecht, M. Obtaining well calibrated probabilities using bayesian binning. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- Norris, C. M., Ghali, W. A., Saunders, L. D., Brant, R., Galbraith, D., Faris, P., Knudtson, M. L., Investigators, A., et al. Ordinal regression model and the linear regression model were superior to the logistic regression models. *Journal of clinical epidemiology*, 59(5):448–456, 2006.
- Owen, A. B. Infinitely imbalanced logistic regression. *Journal of Machine Learning Research*, 8(4), 2007.
- Peterson, B. and Harrell Jr, F. E. Partial proportional odds models for ordinal response variables. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 39(2):205–217, 1990.
- Pöbnecker, W. and Tutz, G. A general framework for the selection of effect type in ordinal regression, 2016. URL <http://nbn-resolving.de/urn/resolver.pl?urn=nbn:de:bvb:19-epub-26912-0>.
- Pratt, J. W. Concavity of the log likelihood. *Journal of the American Statistical Association*, 76(373):103–106, 1981. ISSN 1537274X. doi: 10.1080/01621459.1981.10477613.
- R. H. B. Christensen. *ordinal—Regression Models for Ordinal Data*, 2019. URL <https://CRAN.R-project.org/package=ordinal>.
- Schwarz, J. and Heider, D. GUESS: projecting machine learning scores to well-calibrated probability estimates for clinical decision-making. *Bioinformatics*, 35(14):2458–2465, 11 2018. ISSN 1367-4803. doi: 10.1093/bioinformatics/bty984. URL <https://doi.org/10.1093/bioinformatics/bty984>.
- Serfling, R. *Approximation theorems of mathematical statistics*. Wiley series in probability and mathematical statistics : Probability and mathematical statistics. Wiley, New York, NY [u.a.], [nachdr.] edition, 1980. ISBN 0471024031. URL http://gso.gbv.de/DB=2.1/CMD?ACT=SRCHA&SRT=YOP&IKT=1016&TRM=pfn+024353353&sourceid=fbw_bibsonomy.
- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., and Knight, K. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1):91–108, 2005.
- Tibshirani, R. J. and Taylor, J. The solution path of the generalized lasso. *The annals of statistics*, 39(3):1335–1371, 2011.
- Tutz, G. and Gertheiss, J. Regularized regression for categorical data. *Statistical Modelling*, 16(3):161–200, 2016.
- Ugba, E. R., Mörlein, D., and Gertheiss, J. Smoothing in ordinal regression: An application to sensory data. *Stats*, 4(3):616–633, 2021.
- van der Vaart, A. *Asymptotic Statistics*. Asymptotic Statistics. Cambridge University Press, 2000. ISBN 9780521784504. URL <https://books.google.com/books?id=UEuQEM5RjWgC>.
- Venables, W. N. and Ripley, B. D. *Modern Applied Statistics with S*. Springer, New York, fourth edition, 2002. URL <http://www.stats.ox.ac.uk/pub/MASS4>. ISBN 0-387-95457-0.
- Vershynin, R. *Introduction to the non-asymptotic analysis of random matrices*, pp. 210–268. Cambridge University Press, 2012. doi: 10.1017/CBO9780511794308.006.
- Vershynin, R. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018. ISBN 9781108415194. URL <https://books.google.com/books?id=J-VjswEACAAJ>.
- Viallon, V., Lambert-Lacroix, S., Höfling, H., and Picard, F. Adaptive generalized fused-lasso: Asymptotic properties and applications. 2013.
- Viallon, V., Lambert-Lacroix, S., Hoeffling, H., and Picard, F. On the robustness of the generalized fused lasso to prior specifications. *Statistics and Computing*, 26(1-2):285–301, 2016.
- Von Wachter, T., Bertrand, M., Pollack, H., Rountree, J., and Blackwell, B. Predicting and preventing homelessness in los angeles. *California Policy Lab and University of Chicago Poverty Lab*, 2019.

- Wickham, H. *ggplot2 Elegant Graphics for Data Analysis*. Use R! Springer International Publishing, Cham, 2016. ISBN 978-3-319-24275-0. doi: 10.1007/978-3-319-24277-4. URL <http://dx.doi.org/10.1007/978-3-319-24277-4>.
- Wilke, C. cowplot: streamlined plot theme and plot annotations for 'ggplot2'. r package version 1.1.1. <https://CRAN.R-project.org/package=cowplot>, 2020.
- Wooldridge, J. M. *Econometric analysis of cross section and panel data*. MIT press, 2010.
- Wurm, M. J., Rathouz, P. J., and Hanlon, B. M. Regularized ordinal regression and the ordinalnet r package. *arXiv preprint arXiv:1706.05003*, 2017.
- Wurm, M. J., Rathouz, P. J., and Hanlon, B. M. Regularized ordinal regression and the ordinalNet R package. *Journal of Statistical Software*, 99(6):1–42, 2021.
- Xin, B., Kawahara, Y., Wang, Y., and Gao, W. Efficient generalized fused lasso and its application to the diagnosis of alzheimer’s disease. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 28, 2014.
- Zhang, F. *The Schur Complement and Its Applications*. Numerical Methods and Algorithms. Springer, 2005. ISBN 9780387242712. URL https://books.google.com/books?id=Wjd8_AwjiIIC.
- Zhang, W., Yuan, S., and Wang, J. Optimal real-time bidding for display advertising. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1077–1086, 2014.
- Zhu, Y. An augmented admm algorithm with application to the generalized lasso problem. *Journal of Computational and Graphical Statistics*, 26(1):195–204, 2017.

We state and briefly discuss Theorem A.1, which shows that logistic regression does arbitrarily badly at estimating rare class probabilities as the class imbalance increases, in Section A. We present a second real data experiment on data from patients diagnosed with diabetes in Section B. In Section C, we display summary statistics and additional figures for the observed mean squared errors (MSEs) for each method from the synthetic data experiments from Sections 4.1 and 4.2. We also briefly investigate the effect of implementing PRESTO with a squared ℓ_2 (ridge) penalty rather than an ℓ_1 penalty in Section C.1. We provide the proofs of Theorems 2.2 and A.1 in Section D. In Section E, we present synthetic data experiments and analysis justifying the validity of one of the assumptions of Theorem 2.3 in Sections E.1 and E.2, and we then prove Theorem 2.3. Theorems 2.2, A.1, and 2.3 depend on Lemma D.1, which is stated at the beginning of Section D and proven in Section F. We prove Theorem 3.1 in Section G. Finally, in Section H we provide implementation details for estimating PRESTO.

A. Theorem A.1

It is well-known that class imbalance poses a major challenge for classifiers. Theorem A.1 exhibits this concretely for logistic regression.

Theorem A.1. *Assume $X(\mathbb{R}^p, \lambda_{\max})$ and $Y(2)$ hold. Let $\pi(\mathbf{x}) := \mathbb{P}(y = 2 \mid \mathbf{x})$, and assume that $\sup_{\mathbf{x} \in \mathcal{S}} \pi(\mathbf{x}) = \pi_{\text{rare}}$ for some $\pi_{\text{rare}} \leq 1/2$. Then*

1. for any fixed $\mathbf{v} \in \mathbb{R}^{p+1}$,

$$\frac{1}{\|\mathbf{v}\|_2^2} \text{Asym.MSE} \left((\hat{\alpha}, \hat{\beta}^\top) \mathbf{v} \right) \geq \frac{1}{\lambda_{\max} \pi_{\text{rare}}},$$

and

2. for any fixed $\mathbf{z} \in \mathcal{S}$,

$$\text{Asym.MSE} \left(\frac{\hat{\pi}(\mathbf{z})}{\pi(\mathbf{z})} \right) \geq \frac{1 - \pi_{\text{rare}}}{\pi_{\text{rare}}} \frac{1}{\lambda_{\max}}.$$

Proof. Provided in Section D.2. □

To give an example of applying part 1 of this result, consider the choice $\mathbf{v} = (0, 1, 0, \dots, 0)$. Then we have that $\text{Asym.MSE}(\hat{\beta}_1) \geq 1/(\lambda_{\max} \pi_{\text{rare}})$, so $\hat{\beta}_1$ (or any other estimated coefficient) has arbitrarily large asymptotic mean squared error as π_{rare} vanishes. Part 2 shows that the same thing happens to the asymptotic mean squared error for the estimated probabilities of the logistic regression estimator, when scaled by $\pi(\mathbf{z})$.

B. Real Data Experiment 2: Diabetes

We present another real data experiment using the data set `PreDiabetes` from the `R MLDataR` package (Hutson et al., 2022). This data set contains $n = 3059$ observations of patients who were eventually diagnosed with diabetes. Each observation consists of the age at which the patient was diagnosed with prediabetes and diabetes as well as $p = 5$ covariates. Given an age a , we form an ordinal variable based on the patient’s status of non-diabetic, prediabetic, or diabetic at age $a - 1$. We do this for ages $a \in \{30, 35, 40, \dots, 65\}$. The number of patients diagnosed with diabetes increases with a , so varying a allows us to change the rarity of the rarest class in a natural way. 0.92% of patients in the data were diagnosed with diabetes before age $a = 30$, and 50.93% of the patients were diagnosed with diabetes before age $a = 65$.

We use PRESTO, logistic regression, and the proportional odds model to estimate the probability that each patient was diagnosed with diabetes before age a for each a . Much like our soup tasting data application, in each setting we take repeated random splits of the data, using 90% of the data selected at random for training and 10% for testing. In each iteration we again evaluate each method on the test data using the same estimator for mean squared error of the estimated rare class probabilities. We repeat this procedure 49 times in each of the 8 settings.

We display the results in a plot in Figure 4. We also provide the mean MSEs for each method at each age cutoff in Table 2. We see that PRESTO outperforms both logistic regression and the proportional odds model in all of these settings. (For age cutoffs $a = 29$ and below we were unable to estimate the proportional odds model on all subsamples because of the

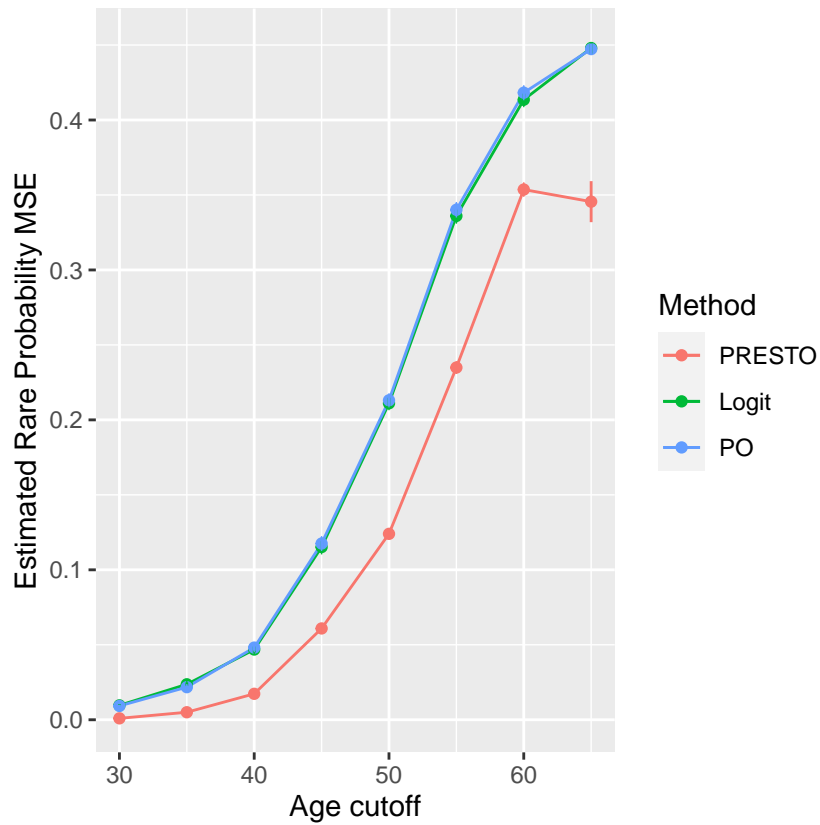


Figure 4. Estimated MSEs of estimated rare class probabilities for each method and each age cutoff across 49 random draws of training and test sets in real data experiment from Section B.

Table 2. Estimated rare class MSE for each method at each age cutoff in prediabetes real data experiment from Section B.

Age cutoff	PRESTO	Logit	PO
30	0.000943	0.009609	0.009217
35	0.005013	0.023658	0.021740
40	0.017307	0.046828	0.048123
45	0.060896	0.115189	0.117525
50	0.124000	0.211083	0.213059
55	0.234906	0.336009	0.340130
60	0.353615	0.413534	0.418179
65	0.345535	0.448015	0.447290

difficulty of having at least one observation from each class in both the training and test sets for 49 random draws.) PRESTO seems to outperform the other methods at all class rarities, and the absolute performance gap increases as the rare class becomes less rare.

C. More Simulation Results

For more results from the synthetic experiments, see Tables 3, 4, and 5, along with Figures 5, 6, 7, 8, and 9.

Table 3. Similar to Table 1; calculated p -values for one-tailed paired t -tests for sparse differences simulation setting of Section 4.1 (statistically significant p -values indicate better performance for PRESTO).

Rare Prop.	Sparsity	Logit p -value	PO p -value
1%	1/3	1.69e-33	6.42e-41
1.17%	1/2	1.61e-15	2.78e-66
0.61%	1/3	5.19e-74	4.21e-19
0.71%	1/2	8.68e-48	3.38e-35
0.37%	1/3	3.08e-61	0.00165
0.43%	1/2	3.75e-64	2.57e-11

Table 4. Means and standard errors of empirical MSEs for each method in each of three intercept settings in the sparse differences synthetic experiment setting of Section 4.1.

Rare Class Proportion	Sparsity	PRESTO	Logistic Regression	Proportional Odds
1%	1/3	6.05e-05 (2.1e-06)	9.38e-05 (2.5e-06)	8.62e-05 (3.1e-06)
1.17%	1/2	9.87e-05 (2.9e-06)	1.25e-04 (3.3e-06)	1.66e-04 (5.5e-06)
0.61%	1/3	3.03e-05 (1.1e-06)	6.90e-05 (2.1e-06)	3.64e-05 (1.4e-06)
0.71%	1/2	5.22e-05 (1.9e-06)	8.89e-05 (2.5e-06)	7.50e-05 (3e-06)
0.37%	1/3	1.40e-05 (6e-07)	5.66e-05 (2.4e-06)	1.49e-05 (6.1e-07)
0.43%	1/2	2.63e-05 (1e-06)	7.21e-05 (2.7e-06)	3.17e-05 (1.4e-06)

Table 5. Means and standard errors of empirical MSEs for each method in each of four intercept settings in the uniform differences synthetic experiment setting of Section 4.2.

Rare Class Proportion	PRESTO	Logistic Regression	Proportional Odds
1.6%	1.13e-04 (2.7e-06)	1.35e-04 (2.9e-06)	1.85e-04 (5.2e-06)
0.99%	5.82e-05 (1.7e-06)	9.25e-05 (2.3e-06)	8.53e-05 (2.6e-06)
0.62%	2.86e-05 (8.3e-07)	6.51e-05 (1.7e-06)	3.43e-05 (9.8e-07)
0.36%	1.33e-05 (4.4e-07)	5.36e-05 (2.2e-06)	1.43e-05 (5e-07)

C.1. Ridge PRESTO

We briefly investigate the effect of implementing PRESTO with a ridge penalty instead of an ℓ_1 penalty, similarly to proposals by Tutz & Gertheiss (2016, Section 4.2.2) and Ugba et al. (2021, Equation 8),

$$\lambda_n \left(\sum_{j=1}^p \beta_{j1}^2 + \sum_{j=1}^p \sum_{k=2}^{K-1} (\beta_{jk} - \beta_{j,k-1})^2 \right).$$

We implement this method (“PRESTO.L2”) in the sparse differences synthetic data experiment of Section 4.1 on the same simulated data that was used for the other methods in the intercept setting (0, 3, 5) for both sparsity levels. The




Figure 5. MSE of predicted rare class probabilities for each method across all $n = 2500$ observations, across 700 simulations, in sparse differences synthetic experiment setting of Section 4.1 with sparsity $1/3$. (These plots are for the two intercept settings that weren't shown in the main text for the sparsity setting of $1/3$)

implementation is identical to PRESTO in every way except for the ridge penalty—the method is implemented using our modification of the `ordinalNet` R package and the tuning parameter is selected in the same way.

Figures 10 and 11 display the results. (The results for all methods but PRESTO.L2 are identical to previous plots and are only displayed for reference.) We also present the means and standard deviations of the MSEs for each method in these settings in Table 6, and p -values for one-tailed paired t -tests of the alternative hypothesis that PRESTO has a lower MSE than the competitor methods in Table 7. We see that in practice, PRESTO and PRESTO.L2 seem to perform similarly in our setting, though the t -tests show that PRESTO does outperform PRESTO.L2 at a 5% significance level in both settings. We might expect PRESTO to better outperform PRESTO.L2 in settings where p/n is larger and where the sparsity is lower. We also note it is not clear if PRESTO.L2 enjoys a high-dimensional consistency guarantee similar to Theorem 2.3.

Table 6. Means and standard errors of empirical MSEs for each method in each of three intercept settings in the sparse differences synthetic experiment setting of Section 4.1 for intercept setting of $(0, 3, 5)$, with PRESTO.L2 implemented as well, as described in Section C.1.

Rare Prop.	Sparsity	PRESTO	Logistic Regression	Proportional Odds	PRESTO.L2
1%	$1/3$	$6.05e-05$ ($2.1e-06$)	$9.38e-05$ ($2.5e-06$)	$8.62e-05$ ($3.1e-06$)	$6.25e-05$ ($2.4e-06$)
1.17%	$1/2$	$9.87e-05$ ($2.9e-06$)	$1.25e-04$ ($3.3e-06$)	$1.66e-04$ ($5.5e-06$)	$1.17e-04$ ($4.2e-06$)

Table 7. Calculated p -values for one-tailed paired t -tests for sparse differences simulation setting of Section 4.1 for intercept setting of $(0, 3, 5)$, with PRESTO.L2 implemented as well, as described in Section C.1. (Statistically significant p -values indicate better performance for PRESTO).

Rare Class Proportion	Sparsity	Logistic Regression	Proportional Odds	PRESTO.L2
1%	$1/3$	$1.69e-33$	$6.42e-41$	0.0358
1.17%	$1/2$	$1.61e-15$	$2.78e-66$	$1.51e-14$

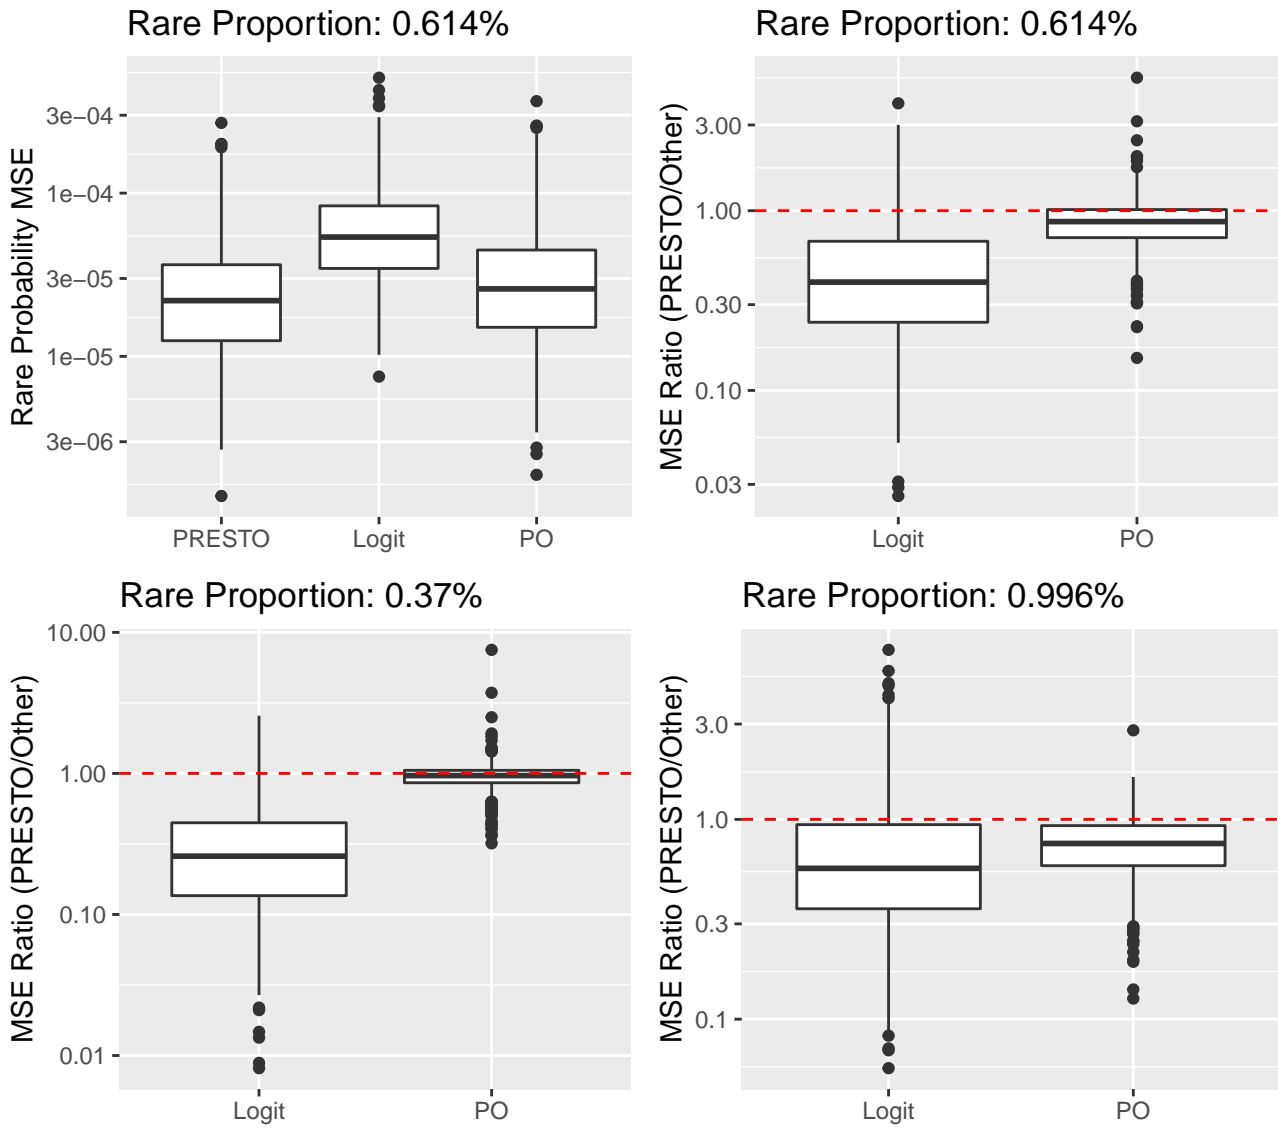


Figure 6. Same as Figure 1, but for the simulations with sparsity 1/2.




Figure 7. Same as Figure 5, but for the simulations with sparsity 1/2.




Figure 8. MSE of predicted rare class probabilities for each method across all $n = 2500$ observations, across 700 simulations, in uniform differences synthetic experiment setting of Section 4.2. (These plots are for two of the intercept settings that weren't shown in the main text.)




Figure 9. MSE of predicted rare class probabilities for each method across all $n = 2500$ observations, across 700 simulations, in uniform differences synthetic experiment setting of Section 4.2 for intercept setting of $(0, 2.5, 4.5)$.




Figure 10. MSE of predicted rare class probabilities for each method across all $n = 2500$ observations, across 700 simulations, in sparse differences synthetic experiment setting of Section 4.1 with sparsity $1/3$ for intercept setting of $(0, 3, 5)$, with PRESTO_L2 implemented as well, as described in Section C.1.




Figure 11. MSE of predicted rare class probabilities for each method across all $n = 2500$ observations, across 700 simulations, in sparse differences synthetic experiment setting of Section 4.1 with sparsity $1/2$ for intercept setting of $(0, 3, 5)$, with PRESTO.L2 implemented as well, as described in Section C.1.

D. Statement of Lemma D.1 and Proofs of Theorems A.1 and 2.2

In Section D.1 we state Lemma D.1, and we prove Theorems A.1 and 2.2 in Section D.2.

D.1. Statement of Lemma D.1

Theorems A.1, 2.2, and 2.3 relate to the asymptotic covariance matrices of the maximum likelihood estimators of the parameters of the proportional odds and logistic regression models. Under mild regularity conditions, the asymptotic covariance matrix of any maximum likelihood estimator (when scaled by \sqrt{n}) is known to be the inverse of the Fisher information matrix

$$-\mathbb{E} \left[\frac{\partial^2}{\partial \boldsymbol{\theta} \boldsymbol{\theta}^\top} \mathcal{L}(\boldsymbol{\theta}) \right],$$

where $\boldsymbol{\theta}$ are the parameters estimated by the model and $\mathcal{L}(\boldsymbol{\theta})$ is the log likelihood (Serfling, 1980, Section 4.2.2). In the proof of Lemma D.1, we calculate these Fisher information matrices for the proportional odds and logistic regression models and verify the needed regularity conditions.

Lemma D.1. *Assume that no class has probability 0 for any $\boldsymbol{x} \in \mathcal{S}$ (equivalently, assume that all of the intercepts in the proportional odds model (1) are not equal, so $\alpha_1 < \dots < \alpha_{K-1}$). Assume that $dF(\boldsymbol{x})$ has bounded support.*

1. *The Fisher information matrix for the maximum likelihood estimator of the proportional odds model (1) is*

$$I^{prop. odds}(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \begin{pmatrix} I_{\boldsymbol{\alpha}\boldsymbol{\alpha}}^{prop. odds} & \left(I_{\boldsymbol{\beta}\boldsymbol{\alpha}}^{prop. odds} \right)^\top \\ I_{\boldsymbol{\beta}\boldsymbol{\alpha}}^{prop. odds} & I_{\boldsymbol{\beta}\boldsymbol{\beta}}^{prop. odds} \end{pmatrix} \in \mathbb{R}^{(K-1+p) \times (K-1+p)}$$

where

$$I_{\alpha\alpha}^{prop. odds}(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \begin{pmatrix} M_1 & -\tilde{M}_2 & 0 & \cdots & 0 & 0 \\ -\tilde{M}_2 & M_2 & -\tilde{M}_3 & \cdots & 0 & 0 \\ 0 & -\tilde{M}_3 & M_3 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & M_{K-2} & -\tilde{M}_{K-1} \\ 0 & 0 & 0 & \cdots & -\tilde{M}_{K-1} & M_{K-1} \end{pmatrix}, \quad (8)$$

$$I_{\beta\alpha}^{prop. odds}(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \begin{pmatrix} J_1^{\mathbf{x}} + \tilde{J}_2^{\mathbf{x}} \\ J_2^{\mathbf{x}} + \tilde{J}_3^{\mathbf{x}} \\ \vdots \\ J_{K-1}^{\mathbf{x}} + \tilde{J}_K^{\mathbf{x}} \end{pmatrix}, \quad \text{and} \quad (9)$$

$$I_{\beta\beta}^{prop. odds}(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \sum_{k=1}^K \left(J_k^{\mathbf{x}\mathbf{x}^\top} + \tilde{J}_k^{\mathbf{x}\mathbf{x}^\top} \right), \quad (10)$$

where

$$M_k := \int [p_k(\mathbf{x})(1 - p_k(\mathbf{x}))]^2 \left(\frac{1}{\pi_k(\mathbf{x})} + \frac{1}{\pi_{k+1}(\mathbf{x})} \right) dF(\mathbf{x}), \quad k \in \{1, \dots, K-1\}, \quad (11)$$

$$\tilde{M}_k := \int p_k(\mathbf{x})(1 - p_k(\mathbf{x}))p_{k-1}(\mathbf{x})(1 - p_{k-1}(\mathbf{x})) \cdot \frac{1}{\pi_k(\mathbf{x})} dF(\mathbf{x}), \quad k \in \{2, \dots, K-1\} \quad (12)$$

and

$$\begin{aligned} J_k &:= \int \pi_k(\mathbf{x})p_k(\mathbf{x})[1 - p_k(\mathbf{x})] dF(\mathbf{x}) \in \mathbb{R}, \\ J_k^{\mathbf{x}} &:= \int \mathbf{x}\pi_k(\mathbf{x})p_k(\mathbf{x})[1 - p_k(\mathbf{x})] dF(\mathbf{x}) \in \mathbb{R}^p, \\ J_k^{\mathbf{x}\mathbf{x}^\top} &:= \int \mathbf{x}\mathbf{x}^\top \pi_k(\mathbf{x})p_k(\mathbf{x})[1 - p_k(\mathbf{x})] dF(\mathbf{x}) \in \mathbb{R}^{p \times p}, \\ \tilde{J}_k &:= \int \pi_k(\mathbf{x})p_{k-1}(\mathbf{x})[1 - p_{k-1}(\mathbf{x})] dF(\mathbf{x}) \in \mathbb{R}, \\ \tilde{J}_k^{\mathbf{x}} &:= \int \mathbf{x}\pi_k(\mathbf{x})p_{k-1}(\mathbf{x})[1 - p_{k-1}(\mathbf{x})] dF(\mathbf{x}) \in \mathbb{R}^p, \quad \text{and} \\ \tilde{J}_k^{\mathbf{x}\mathbf{x}^\top} &:= \int \mathbf{x}\mathbf{x}^\top \pi_k(\mathbf{x})p_{k-1}(\mathbf{x})[1 - p_{k-1}(\mathbf{x})] dF(\mathbf{x}) \in \mathbb{R}^{p \times p} \end{aligned}$$

for all $k \in [K]$.

2. The Fisher information matrix for the maximum likelihood estimator of the logistic regression model predicting whether or not each observation is in class 1 is

$$I^{logistic}(\alpha_1, \boldsymbol{\beta}) = \begin{pmatrix} I_{\alpha\alpha}^{logistic} & \left(I_{\beta\alpha}^{logistic} \right)^\top \\ I_{\beta\alpha}^{logistic} & I_{\beta\beta}^{logistic} \end{pmatrix} = \mathbb{E} \left[\pi_1(\mathbf{X})[1 - \pi_1(\mathbf{X})] \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\top \right] \in \mathbb{R}^{(p+1) \times (p+1)} \quad (13)$$

where $\tilde{\mathbf{X}} := (\mathbf{1} \ \mathbf{X})$ (an n -vector of all ones followed by \mathbf{X}) and

$$I_{\alpha\alpha}^{logistic}(\alpha_1, \boldsymbol{\beta}) = M_1^{logistic}, \quad (14)$$

$$I_{\beta\alpha}^{logistic}(\alpha_1, \boldsymbol{\beta}) = J_1^{\mathbf{x}; logistic} + \tilde{J}_2^{\mathbf{x}; logistic}, \quad \text{and} \quad (15)$$

$$I_{\beta\beta}^{logistic}(\alpha_1, \boldsymbol{\beta}) = J_1^{\mathbf{x}\mathbf{x}^\top; logistic} + \tilde{J}_2^{\mathbf{x}\mathbf{x}^\top; logistic}, \quad (16)$$

where we define

$$M_1^{\text{logistic}} := \int \pi_1(\mathbf{x})(1 - \pi_1(\mathbf{x})) dF(\mathbf{x}) \quad (17)$$

and

$$\begin{aligned} J_1^{\mathbf{x}; \text{logistic}} &:= \int \mathbf{x}\pi_1(\mathbf{x})^2[1 - \pi_1(\mathbf{x})] dF(\mathbf{x}) \in \mathbb{R}^p = J_1^{\mathbf{x}}, \\ J_1^{\mathbf{x}\mathbf{x}^\top; \text{logistic}} &:= \int \mathbf{x}\mathbf{x}^\top \pi_1(\mathbf{x})^2[1 - \pi_1(\mathbf{x})] dF(\mathbf{x}) \in \mathbb{R}^{p \times p} = J_1^{\mathbf{x}\mathbf{x}^\top}, \\ \tilde{J}_2^{\mathbf{x}; \text{logistic}} &:= \int \mathbf{x}\pi_1(\mathbf{x})[1 - \pi_1(\mathbf{x})]^2 dF(\mathbf{x}) \in \mathbb{R}^p, \quad \text{and} \\ \tilde{J}_2^{\mathbf{x}\mathbf{x}^\top; \text{logistic}} &:= \int \mathbf{x}\mathbf{x}^\top \pi_1(\mathbf{x})[1 - \pi_1(\mathbf{x})]^2 dF(\mathbf{x}) \in \mathbb{R}^{p \times p}. \end{aligned}$$

3. The information matrices $I^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta})$ and $I^{\text{logistic}}(\alpha_1, \boldsymbol{\beta})$ are finite and positive definite, and the following convergences hold:

$$\sqrt{n} \cdot \left(\hat{\boldsymbol{\theta}}_1^{\text{prop. odds}} - \boldsymbol{\theta}_1 \right) \xrightarrow{d} \mathcal{N} \left(\mathbf{0}, (I^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta}))^{-1} \right)$$

and

$$\sqrt{n} \cdot \left(\hat{\boldsymbol{\theta}}_1^{\text{logistic}} - \boldsymbol{\theta}_1 \right) \xrightarrow{d} \mathcal{N} \left(\mathbf{0}, (I^{\text{logistic}}(\alpha_1, \boldsymbol{\beta}))^{-1} \right).$$

Further, because these information matrices are symmetric and positive definite, by Observation 7.1.2 in [Horn & Johnson \(2012\)](#) the principal submatrices $I_{\alpha\alpha}^{\text{prop. odds}} := I_{\alpha\alpha}^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta})$, $I_{\beta\beta}^{\text{prop. odds}} := I_{\beta\beta}^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta})$, and $I_{\beta\beta}^{\text{logistic}} := I_{\beta\beta}^{\text{logistic}}(\alpha_1, \boldsymbol{\beta})$ are all positive definite. Finally, we can characterize the finite-sample bias: for any $\mathbf{v} \in \mathbb{R}^{p+1}$, $[(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\beta}}^\top) - (\boldsymbol{\alpha}, \boldsymbol{\beta})^\top] \mathbf{v} = \mathcal{O}(1/n)$.

Proof. Provided in Section F. □

Remark D.2. Note that $J_K = 0$ because $1 - p_K(\mathbf{x}) = 1 - \mathbb{P}(y(\mathbf{x}) \leq K | \mathbf{x}) = 0$ for all \mathbf{x} , and similarly for $J_K^{\mathbf{x}}$ and $J_K^{\mathbf{x}\mathbf{x}^\top}$. Likewise, $\tilde{J}_1 = 0$ because $p_0(\mathbf{x}) = \mathbb{P}(y(\mathbf{x}) \leq 0 | \mathbf{x}) = 0$ for all \mathbf{x} , and similarly for $\tilde{J}_1^{\mathbf{x}}$ and $\tilde{J}_1^{\mathbf{x}\mathbf{x}^\top}$.

We also take a moment to briefly establish some identities we will use later. For any $k \in \{1, \dots, K-1\}$,

$$\begin{aligned} J_k^{\mathbf{x}} + \tilde{J}_{k+1}^{\mathbf{x}} &= \int \mathbf{x}\pi_k(\mathbf{x})p_k(\mathbf{x})[1 - p_k(\mathbf{x})] dF(\mathbf{x}) + \int \mathbf{x}\pi_{k+1}(\mathbf{x})p_k(\mathbf{x})[1 - p_k(\mathbf{x})] dF(\mathbf{x}) \\ &= \int \mathbf{x} [\pi_k(\mathbf{x}) + \pi_{k+1}(\mathbf{x})] p_k(\mathbf{x})[1 - p_k(\mathbf{x})] dF(\mathbf{x}). \end{aligned} \quad (18)$$

Similarly,

$$J_k^{\mathbf{x}\mathbf{x}^\top} + \tilde{J}_{k+1}^{\mathbf{x}\mathbf{x}^\top} = \int \mathbf{x}\mathbf{x}^\top [\pi_k(\mathbf{x}) + \pi_{k+1}(\mathbf{x})] p_k(\mathbf{x})[1 - p_k(\mathbf{x})] dF(\mathbf{x}),$$

so from (10) we have

$$\begin{aligned}
 I_{\beta\beta}^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) &= \sum_{k=1}^K \left(J_k^{\boldsymbol{x}\boldsymbol{x}^\top} + \tilde{J}_k^{\boldsymbol{x}\boldsymbol{x}^\top} \right) \\
 &= \tilde{J}_1^{\boldsymbol{x}\boldsymbol{x}^\top} + \sum_{k=1}^{K-1} \left(J_k^{\boldsymbol{x}\boldsymbol{x}^\top} + \tilde{J}_{k+1}^{\boldsymbol{x}\boldsymbol{x}^\top} \right) + J_K^{\boldsymbol{x}\boldsymbol{x}^\top} \\
 &= \int \boldsymbol{x}\boldsymbol{x}^\top \pi_1(\boldsymbol{x}) \underbrace{p_0(\boldsymbol{x})}_{=\mathbb{P}(y \leq 0 | \boldsymbol{x})=0} [1 - p_0(\boldsymbol{x})] dF(\boldsymbol{x}) + \sum_{k=1}^{K-1} \left(J_k^{\boldsymbol{x}\boldsymbol{x}^\top} + \tilde{J}_{k+1}^{\boldsymbol{x}\boldsymbol{x}^\top} \right) \\
 &\quad + \int \boldsymbol{x}\boldsymbol{x}^\top \pi_K(\boldsymbol{x}) p_K(\boldsymbol{x}) \left[1 - \underbrace{p_K(\boldsymbol{x})}_{=\mathbb{P}(y \leq K | \boldsymbol{x})=1} \right] dF(\boldsymbol{x}) \\
 &= \sum_{k=1}^{K-1} \int \boldsymbol{x}\boldsymbol{x}^\top [\pi_k(\boldsymbol{x}) + \pi_{k+1}(\boldsymbol{x})] p_k(\boldsymbol{x}) [1 - p_k(\boldsymbol{x})] dF(\boldsymbol{x}). \tag{19}
 \end{aligned}$$

D.2. Proofs of Theorems A.1 and 2.2

Equipped with the results of Lemma D.1, we proceed to prove Theorems A.1 and 2.2. (Recall that for square matrices \boldsymbol{A} and \boldsymbol{B} of equal dimension p , we say $\boldsymbol{A} \preceq \boldsymbol{B}$ if $\boldsymbol{B} - \boldsymbol{A}$ is positive semidefinite.)

Proof of Theorem A.1. 1. Note that the assumptions of Lemma D.1 are satisfied. Since the inverse of the asymptotic covariance matrix is

$$\mathbb{E} \left[\pi(\boldsymbol{X}) [1 - \pi(\boldsymbol{X})] \tilde{\boldsymbol{X}} \tilde{\boldsymbol{X}}^\top \right] \preceq \pi_{\text{rare}} (1 - \pi_{\text{rare}}) \mathbb{E} \left[\tilde{\boldsymbol{X}} \tilde{\boldsymbol{X}}^\top \right]$$

(where the second step is valid because $t \mapsto t(1-t)$ is monotone increasing in t for $t \in [0, 1/2]$), by Corollary 7.7.4 in Horn & Johnson (2012) the largest eigenvalue of the inverse of the asymptotic covariance matrix is no larger than $\pi_{\text{rare}}(1 - \pi_{\text{rare}})\lambda_{\max}$. Therefore

$$\text{Asym.Cov} \left((\sqrt{n} \cdot \hat{\boldsymbol{\alpha}}, \sqrt{n} \cdot \hat{\boldsymbol{\beta}}) \right) = \left(\mathbb{E} \left[\pi(\boldsymbol{X}) [1 - \pi(\boldsymbol{X})] \tilde{\boldsymbol{X}} \tilde{\boldsymbol{X}}^\top \right] \right)^{-1} \succeq \left(\pi_{\text{rare}} (1 - \pi_{\text{rare}}) \mathbb{E} \left[\tilde{\boldsymbol{X}} \tilde{\boldsymbol{X}}^\top \right] \right)^{-1}$$

has smallest eigenvalue at least $1/[\lambda_{\max} \pi_{\text{rare}} (1 - \pi_{\text{rare}})]$, which is larger than $1/[\lambda_{\max} \pi_{\text{rare}}]$, again using Corollary 7.7.4 in Horn & Johnson (2012). So for any $\boldsymbol{v} \in \mathbb{R}^{p+1}$, by Theorem 5.1.8 in Lehmann (1999)

$$\text{Asym.Var} \left((\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\beta}}^\top) \boldsymbol{v} \right) = \boldsymbol{v}^\top \text{Asym.Cov} \left((\sqrt{n} \cdot \hat{\boldsymbol{\alpha}}, \sqrt{n} \cdot \hat{\boldsymbol{\beta}}) \right) \boldsymbol{v} \geq \frac{\boldsymbol{v}^\top \boldsymbol{v}}{\lambda_{\max} \pi_{\text{rare}}}.$$

Finally, since we have already shown that $(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\beta}}^\top)$ is asymptotically unbiased, the asymptotic MSE is equal to this asymptotic variance:

$$\begin{aligned}
 \text{Asym.MSE}((\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\beta}}^\top) \boldsymbol{v}) &= \mathbb{E} \left[\lim_{n \rightarrow \infty} \left(\sqrt{n} \cdot \left[(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\beta}}^\top) \boldsymbol{v} - (\boldsymbol{\alpha}, \boldsymbol{\beta}^\top) \boldsymbol{v} \right] \right)^2 \right] \\
 &= \mathbb{E} \left[\lim_{n \rightarrow \infty} \left(\sqrt{n} \cdot \left[\mathbb{E} \left[(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\beta}}^\top) \boldsymbol{v} \right] - (\boldsymbol{\alpha}, \boldsymbol{\beta}^\top) \boldsymbol{v} \right] \right)^2 \right. \\
 &\quad \left. + \lim_{n \rightarrow \infty} \left(\sqrt{n} \cdot \left[(\boldsymbol{\alpha}, \boldsymbol{\beta}^\top) \boldsymbol{v} - \mathbb{E}[(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\beta}}^\top) \boldsymbol{v}] \right] \right)^2 \right] \\
 &= \text{Asym.Var} \left(\sqrt{n} \left[(\boldsymbol{\alpha}, \boldsymbol{\beta}^\top) \boldsymbol{v} - (\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\beta}}^\top) \boldsymbol{v} \right] \right) + 0 \\
 &\geq \frac{\boldsymbol{v}^\top \boldsymbol{v}}{\lambda_{\max} \pi_{\text{rare}}},
 \end{aligned}$$

where in the second-to-last step we used $[(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\beta}}^\top) - (\boldsymbol{\alpha}, \boldsymbol{\beta}^\top)] \boldsymbol{v} = \mathcal{O}(1/n)$ from Lemma D.1.

2. Because $(\hat{\alpha}, \hat{\beta}) \mapsto \hat{\pi}(z)$ is differentiable for all $z \in \mathbb{R}^p$, by the delta method (Theorem 3.1 in van der Vaart 2000)

$$\sqrt{n} \cdot [\hat{\pi}(z) - \pi(z)] \xrightarrow{d} \mathcal{N}\left(0, \pi(z)^2 [1 - \pi(z)]^2 (1, z^\top) (I^{\text{logistic}}(\alpha, \beta))^{-1} (1, z^\top)^\top\right)$$

for any $z \in \mathbb{R}^p$. Therefore

$$\begin{aligned} \text{Asym.Var}(\sqrt{n} \cdot \hat{\pi}(z)) &= \pi(z)^2 [1 - \pi(z)]^2 (1, z^\top) (I^{\text{logistic}}(\alpha, \beta))^{-1} (1, z^\top)^\top \\ &\geq \pi(z)^2 [1 - \pi(z)]^2 \|(1, z^\top)\|_2^2 \lambda_{\min}\left((I^{\text{logistic}}(\alpha, \beta))^{-1}\right) \\ &= \frac{\pi(z)^2 [1 - \pi(z)]^2 \|(1, z^\top)\|_2^2}{\|I^{\text{logistic}}(\alpha, \beta)\|_{\text{op}}} \\ &\geq \frac{\pi(z)^2 [1 - \pi(z)]^2 \|(1, z^\top)\|_2^2}{\pi_{\text{rare}}(1 - \pi_{\text{rare}}) \left\| \mathbb{E} \left[\tilde{\mathbf{X}} \tilde{\mathbf{X}}^\top \right] \right\|_{\text{op}}} \\ &\stackrel{(*)}{\geq} \frac{\pi(z)^2 [1 - \pi_{\text{rare}}]^2}{\pi_{\text{rare}}(1 - \pi_{\text{rare}}) \lambda_{\max}} \\ &= \frac{\pi(z)^2 [1 - \pi_{\text{rare}}]}{\pi_{\text{rare}} \lambda_{\max}}, \end{aligned}$$

where $\lambda_{\min}(\cdot)$ denotes the minimum eigenvalue of \cdot and $(*)$ uses $\|(1, z^\top)\|_2^2 \geq 1$ and $\pi(z) \leq \pi_{\text{rare}}$ for all $z \in \mathcal{S}$. This yields

$$\text{Asym.Var}\left(\sqrt{n} \frac{\pi(z) - \hat{\pi}_n(z)}{\pi(z)}\right) = \frac{1}{\pi(z)^2} \text{Asym.Var}(\sqrt{n} \cdot \hat{\pi}(z)) \geq \frac{1 - \pi_{\text{rare}}}{\pi_{\text{rare}} \lambda_{\max}}.$$

Similarly to the previous result, $\hat{\pi}(z)$ is asymptotically unbiased, its finite sample bias $\pi(z) - \mathbb{E}[\hat{\pi}_n(z)]$ is $\mathcal{O}(1/n)$ by standard maximum likelihood theory (Cordeiro & McCullagh, 1991) since it is a maximum likelihood estimator by the functional equivariance of maximum likelihood estimators, and its asymptotic MSE is equal to its asymptotic variance:

$$\begin{aligned} \text{Asym.MSE}(\hat{\pi}(z)) &= \mathbb{E} \left[\lim_{n \rightarrow \infty} \left(\sqrt{n} \cdot \frac{\pi(z) - \hat{\pi}_n(z)}{\pi(z)} \right)^2 \right] \\ &= \mathbb{E} \left[\lim_{n \rightarrow \infty} \left(\sqrt{n} \cdot \frac{\mathbb{E}[\hat{\pi}_n(z)] - \hat{\pi}_n(z)}{\pi(z)} \right)^2 + \lim_{n \rightarrow \infty} \left(\sqrt{n} \cdot \frac{\pi(z) - \mathbb{E}[\hat{\pi}_n(z)]}{\pi(z)} \right)^2 \right] \\ &= \text{Asym.Var} \left(\sqrt{n} \frac{\pi(z) - \hat{\pi}_n(z)}{\pi(z)} \right) + 0 \\ &\geq \frac{1 - \pi_{\text{rare}}}{\pi_{\text{rare}} \lambda_{\max}}. \end{aligned}$$

□

Proof of Theorem 2.2. 1. Again, the assumptions of Lemma D.1 are satisfied. Lemma D.1 shows that the asymptotic covariance matrix of the scaled maximum likelihood estimates of the parameters of logistic regression (a special case of the proportional odds model with $K = 2$ categories) is

$$\text{Asym.Cov} \left(\sqrt{n} \cdot (\hat{\alpha}, \hat{\beta})^\top \right) = (I^{\text{logistic}}(\alpha, \beta))^{-1} = \begin{pmatrix} I_{\alpha\alpha}^{\text{logistic}} & (I_{\beta\alpha}^{\text{logistic}})^\top \\ I_{\beta\alpha}^{\text{logistic}} & I_{\beta\beta}^{\text{logistic}} \end{pmatrix}^{-1},$$

so in the case that β is known, we have

$$\text{Asym.Var}(\sqrt{n} \cdot \hat{\alpha}_q) = (I_{\alpha\alpha}^{\text{logistic}})^{-1} = \frac{1}{I_{\alpha\alpha}^{\text{logistic}}}.$$

If β is not known, then if $I_{\beta\beta}^{\text{logistic}}$ is positive definite (and therefore invertible) the formula for block matrix inversion yields

$$\text{Asym.Var}(\sqrt{n} \cdot \hat{\alpha}) = \frac{1}{I_{\alpha\alpha}^{\text{logistic}} - \left(I_{\beta\alpha}^{\text{logistic}}\right)^\top \left(I_{\beta\beta}^{\text{logistic}}\right)^{-1} I_{\beta\alpha}^{\text{logistic}}}. \quad (20)$$

We know that $I_{\beta\beta}^{\text{logistic}}$ is positive definite because $I^{\text{logistic}}(\alpha, \beta)$ is finite and positive definite from Lemma D.1, so the principal submatrix $I_{\beta\beta}^{\text{logistic}}$ is positive definite by Observation 7.1.2 in Horn & Johnson (2012). Further, since we know from Lemma D.1 that the covariance matrix of $(\hat{\alpha}, \hat{\beta})$ is finite and positive definite under our conditions, this also implies that

$$0 < I_{\alpha\alpha}^{\text{logistic}} - \left(I_{\beta\alpha}^{\text{logistic}}\right)^\top \left(I_{\beta\beta}^{\text{logistic}}\right)^{-1} I_{\beta\alpha}^{\text{logistic}} < \infty. \quad (21)$$

Now we seek a lower bound for $\text{Asym.Var}(\sqrt{n} \cdot \hat{\alpha})$. We see from (20) that we can get such a bound by lower-bounding $\left(I_{\beta\alpha}^{\text{logistic}}\right)^\top \left(I_{\beta\beta}^{\text{logistic}}\right)^{-1} I_{\beta\alpha}^{\text{logistic}}$. Because $t \mapsto t(1-t)$ is upper-bounded by $1/4$ for all $t \in [0, 1]$,

$$I_{\beta\beta}^{\text{logistic}} = \int \mathbf{x}\mathbf{x}^\top \pi_2(\mathbf{x})[1 - \pi_2(\mathbf{x})] dF(\mathbf{x}) \preceq \int \mathbf{x}\mathbf{x}^\top \cdot \frac{1}{4} dF(\mathbf{x}) = \frac{1}{4} \mathbb{E}[\mathbf{X}\mathbf{X}^\top].$$

Then

$$\begin{aligned} \left(I_{\beta\alpha}^{\text{logistic}}\right)^\top \left(I_{\beta\beta}^{\text{logistic}}\right)^{-1} I_{\beta\alpha}^{\text{logistic}} &\geq \left(I_{\beta\alpha}^{\text{logistic}}\right)^\top \left(\frac{1}{4} \mathbb{E}[\mathbf{X}\mathbf{X}^\top]\right)^{-1} I_{\beta\alpha}^{\text{logistic}} \\ &\geq 4\lambda_{\min}\left(\left(\mathbb{E}[\mathbf{X}\mathbf{X}^\top]\right)^{-1}\right) \left\|I_{\beta\alpha}^{\text{logistic}}\right\|_2^2 \\ &= \frac{4 \left\|I_{\beta\alpha}^{\text{logistic}}\right\|_2^2}{\left\|\mathbb{E}[\mathbf{X}\mathbf{X}^\top]\right\|_{\text{op}}} \\ &\geq \frac{4\pi_{\min}^2(1 - \pi_{\min})^2 \left\|\mathbb{E}[\mathbf{X}]\right\|_2^2}{\lambda_{\max}} =: \Delta, \end{aligned} \quad (22)$$

where $\lambda_{\min}(\cdot)$ denotes the minimum eigenvalue of \cdot and the last step follows because

$$\begin{aligned} \left\|I_{\beta\alpha}^{\text{logistic}}\right\|_2 &= \left\|\int \mathbf{x}\pi_2(\mathbf{x})[1 - \pi_2(\mathbf{x})] dF(\mathbf{x})\right\|_2 \\ &\geq \left\|\int \mathbf{x}\pi_{\min}(1 - \pi_{\min}) dF(\mathbf{x})\right\|_2 \\ &= \pi_{\min}(1 - \pi_{\min}) \left\|\mathbb{E}[\mathbf{X}]\right\|_2 \end{aligned}$$

(where we used the fact that \mathbf{X} has support only over nonnegative numbers). Therefore (20) and (22) yield

$$\text{Asym.Var}(\sqrt{n} \cdot \hat{\alpha}) \geq \left(\frac{1}{\text{Asym.Var}(\sqrt{n} \cdot \hat{\alpha}_q)} - \Delta\right)^{-1}. \quad (23)$$

The remainder of the argument is similar to the end of the proof of Theorem A.1: the asymptotic unbiasedness and $\mathcal{O}(1/n)$ finite-sample bias of these estimators yields

$$\text{Asym.MSE}(\hat{\alpha}) = \text{Asym.Var}(\sqrt{n} \cdot [\alpha - \hat{\alpha}])$$

and

$$\text{Asym.MSE}(\hat{\alpha}_q) = \text{Asym.Var}(\sqrt{n} \cdot [\alpha - \hat{\alpha}_q]).$$

Then from (21) we know that

$$I_{\alpha\alpha}^{\text{logistic}} > \left(I_{\beta\alpha}^{\text{logistic}}\right)^\top \left(I_{\beta\beta}^{\text{logistic}}\right)^{-1} I_{\beta\alpha}^{\text{logistic}} \geq \frac{4\pi_{\min}^2(1 - \pi_{\min})^2 \left\|\mathbb{E}[\mathbf{X}]\right\|_2^2}{\lambda_{\max}}. \quad (24)$$

Making the appropriate substitutions into (24) yields

$$\frac{1}{\text{Asym.MSE}(\hat{\alpha}_q)} - \Delta > 0,$$

and then substituting into (23) yields

$$\begin{aligned} \text{Asym.MSE}(\hat{\alpha}) &\geq \left(\frac{1}{\text{Asym.MSE}(\hat{\alpha}_q)} - \Delta \right)^{-1} \\ &= \frac{\text{Asym.MSE}(\hat{\alpha}_q)}{1 - \Delta \cdot \text{Asym.MSE}(\hat{\alpha}_q)} \\ &\stackrel{(*)}{\geq} \text{Asym.MSE}(\hat{\alpha}_q) \cdot (1 + \Delta \cdot \text{Asym.MSE}(\hat{\alpha}_q)) \\ \Leftrightarrow \frac{\text{Asym.MSE}(\hat{\alpha}) - \text{Asym.MSE}(\hat{\alpha}_q)}{[\text{Asym.MSE}(\hat{\alpha}_q)]^2} &\geq \Delta, \end{aligned}$$

where in (*) we used the inequality $c/(1 - ct) \leq c(1 + ct)$ for any $c > 0, t < \frac{1}{c}$.

2. Because $(\hat{\alpha}, \hat{\beta}) \mapsto \hat{\pi}(z)$ is differentiable for all $z \in \mathbb{R}^p$, by the delta method (Theorem 3.1 in van der Vaart 2000)

$$\sqrt{n} \cdot [\hat{\pi}(z) - \pi(z)] \xrightarrow{d} \mathcal{N}\left(0, \pi(z)^2 [1 - \pi(z)]^2 (1, z^\top) (I^{\text{logistic}}(\alpha, \beta))^{-1} (1, z^\top)^\top\right)$$

for any $z \in \mathbb{R}^p$, and similarly

$$\sqrt{n} \cdot [\hat{\pi}_q(z) - \pi(z)] \xrightarrow{d} \mathcal{N}\left(0, \frac{\pi(z)^2 [1 - \pi(z)]^2}{I_{\alpha\alpha}^{\text{logistic}}}\right).$$

We can find $(I^{\text{logistic}}(\alpha, \beta))^{-1}$ using the formula for block matrix inversion if

$$D := I_{\beta\beta}^{\text{logistic}} - \frac{I_{\beta\alpha}^{\text{logistic}} (I_{\beta\alpha}^{\text{logistic}})^\top}{I_{\alpha\alpha}^{\text{logistic}}}$$

is positive definite (and therefore invertible; note that D is symmetric) and $I_{\alpha\alpha}^{\text{logistic}} > 0$. We have

$$I_{\alpha\alpha}^{\text{logistic}} = \int \pi_2(\mathbf{x}) [1 - \pi_2(\mathbf{x})] dF(\mathbf{x}) \geq \int \pi_{\min} [1 - \pi_{\min}] dF(\mathbf{x}) = \pi_{\min} [1 - \pi_{\min}] > 0,$$

and by Theorem 1.12 in Zhang (2005), we then know D is positive definite (and invertible) since $I^{\text{logistic}}(\alpha, \beta)$ is by Lemma D.1 and $I_{\alpha\alpha}^{\text{logistic}} > 0$. Let $\lambda_{\max}^D := \|D\|_{\text{op}}$ be the largest eigenvalue of D ; then $1/\lambda_{\max}^D$ is the smallest eigenvalue of D^{-1} . Then for any $z \in \mathbb{R}^p$, we have

$$\begin{aligned} &(1 \quad z^\top) (I^{\text{logistic}}(\alpha, \beta))^{-1} \begin{pmatrix} 1 \\ z \end{pmatrix} \\ &= (1 \quad z^\top) \begin{pmatrix} \frac{1}{I_{\alpha\alpha}^{\text{logistic}}} + \frac{1}{(I_{\alpha\alpha}^{\text{logistic}})^2} (I_{\beta\alpha}^{\text{logistic}})^\top D^{-1} I_{\beta\alpha}^{\text{logistic}} & -\frac{1}{I_{\alpha\alpha}^{\text{logistic}}} (I_{\beta\alpha}^{\text{logistic}})^\top D^{-1} \\ -\frac{1}{I_{\alpha\alpha}^{\text{logistic}}} D^{-1} I_{\beta\alpha}^{\text{logistic}} & D^{-1} \end{pmatrix} \begin{pmatrix} 1 \\ z \end{pmatrix} \\ &= \frac{1}{I_{\alpha\alpha}^{\text{logistic}}} + \frac{1}{(I_{\alpha\alpha}^{\text{logistic}})^2} (I_{\beta\alpha}^{\text{logistic}})^\top D^{-1} I_{\beta\alpha}^{\text{logistic}} + z^\top D^{-1} z - 2 \frac{1}{I_{\alpha\alpha}^{\text{logistic}}} I_{\beta\alpha}^{\text{logistic}} D^{-1} z \\ &\stackrel{(a)}{=} \frac{1}{I_{\alpha\alpha}^{\text{logistic}}} + \frac{1}{(I_{\alpha\alpha}^{\text{logistic}})^2} \left\| D^{-1/2} (I_{\beta\alpha}^{\text{logistic}} - I_{\alpha\alpha}^{\text{logistic}} z) \right\|_2^2 \\ &\stackrel{(b)}{\geq} \frac{1}{I_{\alpha\alpha}^{\text{logistic}}} + \frac{1}{(I_{\alpha\alpha}^{\text{logistic}})^2 \lambda_{\max}^D} \left\| I_{\beta\alpha}^{\text{logistic}} - I_{\alpha\alpha}^{\text{logistic}} z \right\|_2^2 \\ &\geq \frac{1}{I_{\alpha\alpha}^{\text{logistic}}}, \end{aligned} \tag{25}$$

where (a) follows from

$$\begin{aligned} & \frac{1}{\left(I_{\alpha\alpha}^{\text{logistic}}\right)^2} \left\| \mathbf{D}^{-1/2} \left(I_{\beta\alpha}^{\text{logistic}} - I_{\alpha\alpha}^{\text{logistic}} \mathbf{z} \right) \right\|_2^2 \\ &= \frac{1}{\left(I_{\alpha\alpha}^{\text{logistic}}\right)^2} \left(\left(I_{\beta\alpha}^{\text{logistic}} \right)^\top \mathbf{D}^{-1} I_{\beta\alpha}^{\text{logistic}} + \left(I_{\alpha\alpha}^{\text{logistic}} \right)^2 \mathbf{z}^\top \mathbf{D}^{-1} \mathbf{z} - 2 I_{\alpha\alpha}^{\text{logistic}} I_{\beta\alpha}^{\text{logistic}} \mathbf{D}^{-1} \mathbf{z} \right) \end{aligned}$$

and (b) uses the fact that $1/\sqrt{\lambda_{\max}^{\mathbf{D}}}$ is the smallest eigenvalue of $\mathbf{D}^{-1/2}$. If we can show that $\left\| I_{\beta\alpha}^{\text{logistic}} - I_{\alpha\alpha}^{\text{logistic}} \mathbf{z} \right\|_2 \neq 0$, then (25) is enough to establish the strict inequality in the result. Using (13), note that

$$\begin{aligned} 0 &= I_{\beta\alpha}^{\text{logistic}} - I_{\alpha\alpha}^{\text{logistic}} \mathbf{z} \\ &= \int \mathbf{x} \pi(\mathbf{x}) [1 - \pi(\mathbf{x})] dF(\mathbf{x}) - \mathbf{z} \int \pi(\mathbf{x}) [1 - \pi(\mathbf{x})] dF(\mathbf{x}) \\ \iff \mathbf{z} &= \frac{\int \mathbf{x} \pi(\mathbf{x}) [1 - \pi(\mathbf{x})] dF(\mathbf{x})}{\int \pi(\mathbf{x}) [1 - \pi(\mathbf{x})] dF(\mathbf{x})}, \end{aligned}$$

so for all $\mathbf{z} \neq \mathbb{E}[\mathbf{X}\pi(\mathbf{X})[1 - \pi(\mathbf{X})]] / \mathbb{E}[\pi(\mathbf{X})[1 - \pi(\mathbf{X})]]$, we have $\left\| I_{\beta\alpha}^{\text{logistic}} - I_{\alpha\alpha}^{\text{logistic}} \mathbf{z} \right\|_2^2 > 0$. So for any $\mathbf{z} \in \mathbb{R}^p \setminus \mathbb{E}[\mathbf{X}\pi(\mathbf{X})[1 - \pi(\mathbf{X})]] / \mathbb{E}[\pi(\mathbf{X})[1 - \pi(\mathbf{X})]]$, we have

$$\begin{aligned} \pi(\mathbf{z})^2 [1 - \pi(\mathbf{z})]^2 \left(\mathbf{1} \quad \mathbf{z}^\top \right) \begin{pmatrix} I_{\alpha\alpha}^{\text{logistic}} & \left(I_{\beta\alpha}^{\text{logistic}} \right)^\top \\ I_{\beta\alpha}^{\text{logistic}} & I_{\beta\beta}^{\text{logistic}} \end{pmatrix}^{-1} \begin{pmatrix} \mathbf{1} \\ \mathbf{z} \end{pmatrix} &> \frac{\pi(\mathbf{z})^2 [1 - \pi(\mathbf{z})]^2}{I_{\alpha\alpha}^{\text{logistic}}} \\ \iff \text{Asym.Var}(\sqrt{n} \cdot \hat{\pi}_q(\mathbf{z})) &< \text{Asym.Var}(\sqrt{n} \cdot (\hat{\pi}(\mathbf{z}))). \end{aligned}$$

□

E. Investigation of Plausibility of Assumption (5) and Proof of Theorem 2.3

Before we prove Theorem 2.3, we begin by investigating the plausibility of Assumption (5). We start by investigating whether $\lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right)$ is bounded away from 0 in Section E.1. Then in Section E.2 we directly investigate whether Assumption (5) seems to hold in a variety of contexts. We prove Theorem 2.3 in Section E.3, and we prove the supporting result Lemma E.1 in Section E.4.

E.1. Investigating Whether $\lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right)$ is Bounded Away From Zero

To investigate the plausibility of the assumption that $I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}}$ is positive definite (and that the upper bound for π_{rare} in Equation 5 can hold) empirically, we perform two simulation studies, with setups similar to those of our synthetic data experiments in Section 4 of the paper. We repeat Simulation Study A 25 times. Using $n = 10^6$, $p = 10$, and $K = 3$, we generate $\mathbf{X} \in \mathbb{R}^{n \times p}$ with $\mathbf{X}_{ij} \sim \text{Uniform}(-1, 1)$ iid for all $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, p\}$. We then generate $y_i \in \{1, 2, 3\}$ from \mathbf{x}_i for each $i \in \{1, \dots, n\}$ according to the proportional odds model (1), using $\boldsymbol{\beta} = (1, \dots, 1)^\top$ and $\boldsymbol{\alpha} = (0, 20)$ (so that class 3 is very rare, with $\pi_{\text{rare}} \approx 4.54 \cdot 10^{-5}$). We estimate $I_{\beta\beta}$, $I_{\beta\alpha_1}$, and $I_{\alpha_1\alpha_1}$ using empirical (“plug-in”) estimates of the expressions in (8), (9), and (10); for example, using (19) we estimated $I_{\beta\beta}$ by

$$\frac{1}{n} \sum_{i=1}^n \sum_{k=1}^{K-1} \mathbf{x}_i \mathbf{x}_i^\top [\pi_k(\mathbf{x}_i) + \pi_{k+1}(\mathbf{x}_i)] p_k(\mathbf{x}_i) [1 - p_k(\mathbf{x}_i)].$$

(Note that we used the true $\pi_k(\cdot)$ and $p_k(\cdot)$ functions in these estimates.) Finally, we use these estimated quantities to estimate $I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}}$, and we calculate the minimum eigenvalue of this estimated matrix. Across all 25 simulations, the

sample mean of this minimum eigenvalue is 0.02361, and the minimum is 0.02359. The standard error is $2.94 \cdot 10^{-6}$, and the 95% confidence interval for the mean of the minimum eigenvalue is (0.02360, 0.02362). See Figure 12 for a boxplot of the 25 estimated minimum eigenvalues. These results seem to suggest that the assumption that $\lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right) > 0$ is reasonable under the assumptions of Theorem 2.3 for \mathbf{X} with this iid uniform distribution.




Figure 12. Boxplot of the estimated minimum eigenvalues of $I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}}$ in Simulation Study A as described in Section E.1.

Next, we conduct another simulation study to investigate whether $\lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right)$ is bounded away from 0 if the features in \mathbf{X} are correlated, and over a wider range of π_{rare} . In Simulation Study B, we generate matrices of standard multivariate Gaussian data truncated entrywise between -3 and 3 , with $n = 10^6$ and $p = 10$. We generate matrices with covariance matrices

$$\begin{pmatrix} 1 & \rho & \cdots & \rho \\ \rho & 1 & \cdots & \rho \\ \vdots & \vdots & \ddots & \vdots \\ \rho & \rho & \cdots & 1 \end{pmatrix}$$

for $\rho \in \{0, 0.25, 0.5, 0.75\}$. We then generate ordinal responses from the proportional odds model with $K = 3$ and $\beta = (1, 1, \dots, 1)^\top$. The intercept for the first separating hyperplane is 0 (so that the expected class probabilities for classes 1 and 2 are both close to $1/2$) and we choose three different values for the second intercept by analytically solving for intercepts such that $\pi_{\text{rare}} = \sup_{x \in [-3, 3]^p} \{\pi_3(\mathbf{x})\}$ equals 10^{-5} , 10^{-6} , or 10^{-7} . To see the behavior as π_{rare} vanishes, we also investigate a fourth setting where all of the $\pi_3(\mathbf{x})$ terms in the Fisher information are simply set to 0 (and $\pi_2(\mathbf{x})$ is coerced to equal $1 - \pi_1(\mathbf{x})$ pointwise), as if the intercept for the second separating hyperplane were infinity. In each setting, we generate 7 random matrices and estimate $\lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right)$ in the same way as in Simulation Study A.

The results are displayed in Figures 13, 14, 15, and 16 for correlations of 0, 0.25, 0.5, and 0.75, respectively. In each plot the rare class probability is displayed on the horizontal axis on a log scale (with a line break after the far left result for the asymptotic case with $\pi_{\text{rare}} = 0$), and the average minimum eigenvalue for each of the seven random matrices is displayed on the vertical axis. We see that the means of the minimum eigenvalues are well above 0; we also note that in every generated matrix in every setting, the minimum eigenvalue was strictly positive.




Figure 13. Plot of the average estimated minimum eigenvalues of $I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}}$ in Simulation Study B as described in Section E.1 with each column of \mathbf{X} having correlation 0 with all of the others.

Finally, we also investigate analytically whether $\lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right)$ is bounded away from 0 as π_{rare} becomes arbitrarily small. To see that this seems reasonable, note that from (19) we have

$$\begin{aligned}
 I_{\beta\beta} &= \sum_{k=1}^2 \int \mathbf{x}\mathbf{x}^\top [\pi_k(\mathbf{x}) + \pi_{k+1}(\mathbf{x})] p_k(\mathbf{x}) [1 - p_k(\mathbf{x})] dF(\mathbf{x}) \\
 &= \int \mathbf{x}\mathbf{x}^\top ([\pi_1(\mathbf{x}) + \pi_2(\mathbf{x})] \pi_1(\mathbf{x}) [1 - \pi_1(\mathbf{x})] + [\pi_2(\mathbf{x}) + \pi_3(\mathbf{x})] \pi_3(\mathbf{x}) [1 - \pi_3(\mathbf{x})]) dF(\mathbf{x}) \\
 &= \int \mathbf{x}\mathbf{x}^\top ([1 - \pi_3(\mathbf{x})] \pi_1(\mathbf{x}) [1 - \pi_1(\mathbf{x})] + [1 - \pi_1(\mathbf{x})] \pi_3(\mathbf{x}) [1 - \pi_3(\mathbf{x})]) dF(\mathbf{x}) \\
 &= \int \mathbf{x}\mathbf{x}^\top [1 - \pi_1(\mathbf{x})] [1 - \pi_3(\mathbf{x})] (\pi_1(\mathbf{x}) + \pi_3(\mathbf{x})) dF(\mathbf{x}) \\
 &= \int \mathbf{x}\mathbf{x}^\top [\pi_2(\mathbf{x}) + \pi_3(\mathbf{x})] [1 - \pi_3(\mathbf{x})] [\pi_1(\mathbf{x}) + \pi_3(\mathbf{x})] dF(\mathbf{x}),
 \end{aligned}$$

which is non-vanishing in π_{rare} . (In particular, there seems to be no reason to suspect that the eigenvalues of $I_{\beta\beta}$ change drastically for, say $\pi_3(\mathbf{x}) \leq 10^{-7}$ for all $\mathbf{x} \in \mathcal{S}$, as in Simulation Study B, versus $\pi_3(\mathbf{x}) \leq 10^{-20}$ for all $\mathbf{x} \in \mathcal{S}$.) Further, (32) below shows that

$$\left\| \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right\|_{\text{op}} = \frac{\|I_{\beta\alpha_1}\|_2^2}{I_{\alpha_1\alpha_1}}$$

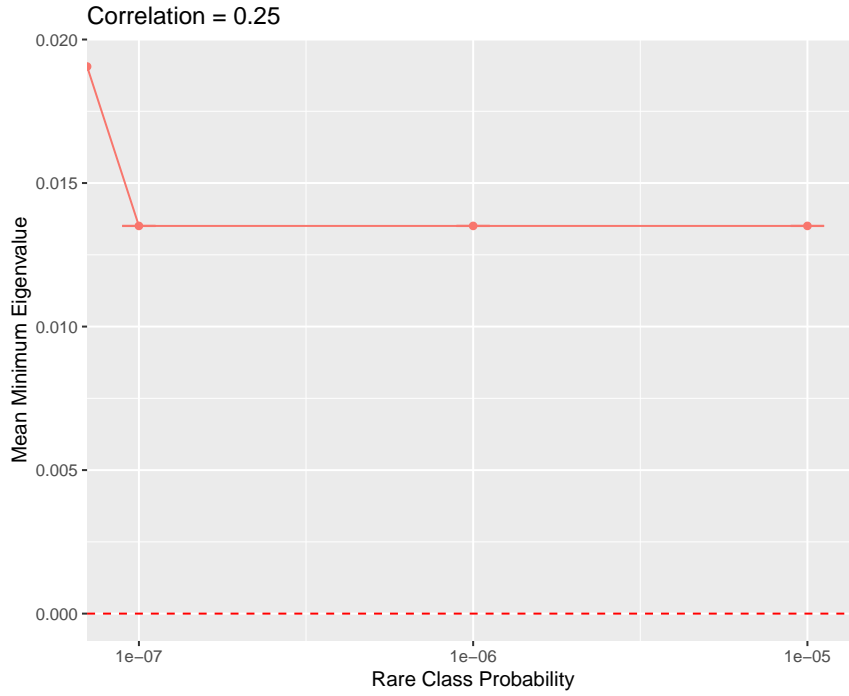


Figure 14. Plot of the average estimated minimum eigenvalues of $I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^T}{I_{\alpha_1\alpha_1}}$ in Simulation Study B as described in Section E.1 with each column of \mathbf{X} having correlation 0.25 with all of the others.




Figure 15. Plot of the average estimated minimum eigenvalues of $I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^T}{I_{\alpha_1\alpha_1}}$ in Simulation Study B as described in Section E.1 with each column of \mathbf{X} having correlation 0.5 with all of the others.




Figure 16. Plot of the average estimated minimum eigenvalues of $I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}}$ in Simulation Study B as described in Section E.1 with each column of \mathbf{X} having correlation 0.75 with all of the others.

(where $\|\cdot\|_{\text{op}}$ is the operator norm) is bounded from above by a constant not depending on π_{rare} . This also suggests that $\lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right)$ should be insensitive to π_{rare} becoming vanishingly small. Taken together, this suggests that Assumption (26) does not become more implausible as π_{rare} becomes arbitrarily small.

E.2. Investigating Whether Assumption (5) is Plausible

Having developed evidence that $\lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right)$ is bounded away from 0, we now directly investigate the plausibility of the assumption

$$\pi_{\text{rare}} \leq \frac{\lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right)}{3M^2(2+M)}. \quad (26)$$

First we examine this in the context of Simulation Study A. Note that $M = \|(1, \dots, 1)\|_2 = \sqrt{p}$, $\mathcal{S} = [-1, 1]^p$ and

$$\sup_{\mathbf{x} \in \mathcal{S}} \{\pi_3(\mathbf{x})\} = \pi_3(\mathbf{x}^*) = 1 - \frac{1}{1 + \exp\{-(20 + \sum_{j=1}^p x_j^*)\}},$$

for $\mathbf{x}^* = (-1, \dots, -1)$. So all of the quantities in (26) are known except the minimum eigenvalue, which we were able to estimate with seemingly high precision. It turns out that (26) is satisfied even when we use the minimum value of λ_{\min} across all 25 simulations as our estimate; in this case, the left side of (26) is $4.54 \cdot 10^{-5}$ and the right side is $1.52 \cdot 10^{-4}$.

Next we investigate (26) in the context of Simulation Study B. In each of the four correlation settings, we calculate the proportion of generated random matrices in which (26) is satisfied at each rarity level (according to our minimum eigenvalue estimates). The results are displayed in Figures 17, 18, 19, and 20. We see that (26) is not satisfied for any of the random matrices where $\pi_{\text{rare}} = 10^{-5}$, but it is satisfied for all of the matrices with all of the remaining values of π_{rare} regardless of correlation of the features. This suggests that assumption (5) is indeed reasonable for π_{rare} small enough.




Figure 17. Plot of the proportion of random matrices in Simulation Study B (as described in Section E.1) in which (26) is satisfied in the setting where each column of \mathbf{X} has correlation 0 with all of the others.




Figure 18. Plot of the proportion of random matrices in Simulation Study B (as described in Section E.1) in which (26) is satisfied in the setting where each column of \mathbf{X} has correlation 0.25 with all of the others.




Figure 19. Plot of the proportion of random matrices in Simulation Study B (as described in Section E.1) in which (26) is satisfied in the setting where each column of \mathbf{X} has correlation 0.5 with all of the others.

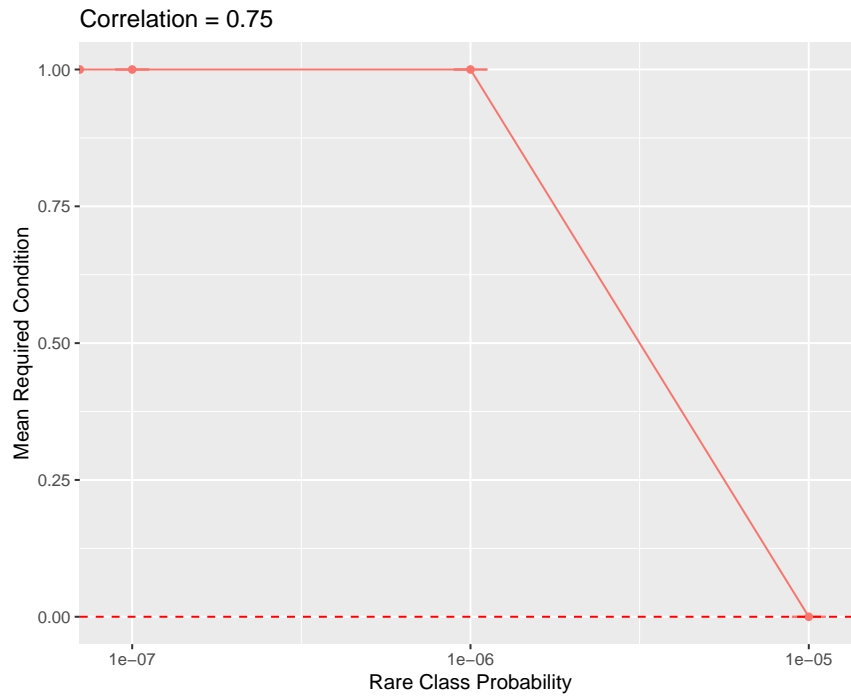


Figure 20. Plot of the proportion of random matrices in Simulation Study B (as described in Section E.1) in which (26) is satisfied in the setting where each column of \mathbf{X} has correlation 0.75 with all of the others.

E.3. Proof of Theorem 2.3

Before we proceed with the proof of Theorem 2.3, we state a lemma with inequalities we will use.

Lemma E.1. *The following inequalities hold under the assumptions of Theorem 2.3:*

$$I_{\alpha_2\alpha_2} \leq \pi_{rare} I_{\alpha_1\alpha_1} \cdot \frac{1}{1/4 - \Delta^2}, \quad (27)$$

$$I_{\alpha_2\alpha_2} \leq \pi_{rare} \left(1 + \frac{\pi_{rare}}{1/2 - \Delta} \right), \quad (28)$$

$$|I_{\alpha_1\alpha_2}| \leq I_{\alpha_2\alpha_2}, \quad (29)$$

$$\|I_{\beta\alpha_2}\|_2 \leq M \cdot I_{\alpha_2\alpha_2}, \quad (30)$$

$$\frac{\|I_{\beta\alpha_2}\|_2^2}{I_{\alpha_2\alpha_2}} \leq M^2 \pi_{rare} \left(1 + \frac{\pi_{rare}}{1/2 - \Delta} \right), \quad \text{and} \quad (31)$$

$$\frac{\|I_{\beta\alpha_1}\|_2^2}{I_{\alpha_1\alpha_1}} \leq \frac{M^2}{4}. \quad (32)$$

Proof. Provided immediately after the proof of Theorem 2.3, in Section E.4. □

By an argument analogous to the one used at the end of the proof of Theorem A.1, it is enough to upper-bound

$$\left\| \text{Cov} \left(\lim_{n \rightarrow \infty} \sqrt{n} \left[\hat{\beta}^{\text{prop. odds}} - \beta \right] \right) \right\|_{\text{op}} = \left\| \text{Asym.Cov} \left(\sqrt{n} \cdot \hat{\beta}^{\text{prop. odds}} \right) \right\|_{\text{op}}.$$

Using Lemma D.1 and the block matrix inversion formula,

$$\begin{aligned} \text{Asym.Cov} \left(\sqrt{n} \cdot \hat{\beta}^{\text{prop. odds}} \right) &= (I_{\beta\beta} - I_{\alpha\beta}^\top I_{\alpha\alpha}^{-1} I_{\alpha\beta})^{-1} \\ \implies \left\| \text{Asym.Cov} \left(\sqrt{n} \cdot \hat{\beta}^{\text{prop. odds}} \right) \right\|_{\text{op}} &= \frac{1}{\lambda_{\min} \left(I_{\beta\beta} - I_{\alpha\beta}^\top I_{\alpha\alpha}^{-1} I_{\alpha\beta} \right)}, \end{aligned} \quad (33)$$

where $\lambda_{\min}(\cdot)$ is the minimum eigenvalue of \cdot and $\|\cdot\|_{\text{op}}$ is the operator norm. $I_{\alpha\alpha}$ is a 2×2 matrix, so

$$I_{\alpha\alpha}^{-1} = \frac{1}{\det(I_{\alpha\alpha})} \begin{pmatrix} I_{\alpha_2\alpha_2} & -I_{\alpha_1\alpha_2} \\ -I_{\alpha_2\alpha_1} & I_{\alpha_1\alpha_1} \end{pmatrix},$$

and

$$\begin{aligned} I_{\alpha\beta}^\top I_{\alpha\alpha}^{-1} I_{\alpha\beta} &= \frac{1}{\det(I_{\alpha\alpha})} \begin{pmatrix} I_{\beta\alpha_1} & I_{\beta\alpha_2} \end{pmatrix} \begin{pmatrix} I_{\alpha_2\alpha_2} & -I_{\alpha_1\alpha_2} \\ -I_{\alpha_2\alpha_1} & I_{\alpha_1\alpha_1} \end{pmatrix} \begin{pmatrix} I_{\alpha_1\beta} \\ I_{\alpha_2\beta} \end{pmatrix} \\ &= \frac{1}{\det(I_{\alpha\alpha})} (I_{\alpha_2\alpha_2} I_{\beta\alpha_1} I_{\beta\alpha_1}^\top + I_{\alpha_1\alpha_1} I_{\beta\alpha_2} I_{\beta\alpha_2}^\top - I_{\alpha_1\alpha_2} I_{\beta\alpha_1} I_{\beta\alpha_2}^\top - I_{\alpha_1\alpha_2} I_{\beta\alpha_2} I_{\beta\alpha_1}^\top) \\ &= \frac{1}{\det(I_{\alpha\alpha})} (I_{\alpha_2\alpha_2} I_{\beta\alpha_1} I_{\beta\alpha_1}^\top + I_{\alpha_1\alpha_1} I_{\beta\alpha_2} I_{\beta\alpha_2}^\top + |I_{\alpha_1\alpha_2}| I_{\beta\alpha_1} I_{\beta\alpha_2}^\top + |I_{\alpha_1\alpha_2}| I_{\beta\alpha_2} I_{\beta\alpha_1}^\top), \end{aligned} \quad (34)$$

where in the last step we used that $I_{\alpha_1\alpha_2} = -\tilde{M}_2 < 0$, which is clear from (9) and (12). Because we know from Lemma D.1 that $I_{\alpha\alpha}$ (and therefore also $I_{\alpha\alpha}^{-1}$) is positive definite, by Observation 7.1.6 in Horn & Johnson (2012) $I_{\alpha\beta}^\top I_{\alpha\alpha}^{-1} I_{\alpha\beta}$ is

positive definite as well. Therefore we can use an upper bound on $1/\det(I_{\alpha\alpha})$ to upper bound $I_{\alpha\beta}^\top I_{\alpha\alpha}^{-1} I_{\alpha\beta}$. We have

$$\begin{aligned}
 \det(I_{\alpha\alpha}) &= I_{\alpha_1\alpha_1} I_{\alpha_2\alpha_2} - I_{\alpha_1\alpha_2}^2 \\
 &\stackrel{(a)}{\geq} I_{\alpha_1\alpha_1} I_{\alpha_2\alpha_2} - I_{\alpha_2\alpha_2}^2 \\
 &= I_{\alpha_2\alpha_2} (I_{\alpha_1\alpha_1} - I_{\alpha_2\alpha_2}) \\
 &\stackrel{(b)}{\geq} I_{\alpha_2\alpha_2} \left(I_{\alpha_1\alpha_1} - \pi_{\text{rare}} I_{\alpha_1\alpha_1} \cdot \frac{1}{1/4 - \Delta^2} \right) \\
 &= \left(1 - \frac{\pi_{\text{rare}}}{1/4 - \Delta^2} \right) I_{\alpha_1\alpha_1} I_{\alpha_2\alpha_2} \\
 &\stackrel{(c)}{\geq} \frac{1}{2} I_{\alpha_1\alpha_1} I_{\alpha_2\alpha_2},
 \end{aligned}$$

where in (a) we used (29), in (b) we used (27), and (c) uses that from the upper bound (5) for π_{rare} we have

$$\begin{aligned}
 \pi_{\text{rare}} &\leq \frac{1}{2} \left(\frac{1}{2} - \Delta \right) \left(\frac{1}{2} + \Delta \right) = \frac{1}{2} \left(\frac{1}{4} - \Delta^2 \right) \\
 \iff \frac{\pi_{\text{rare}}}{1/4 - \Delta^2} &\leq \frac{1}{2} \\
 \iff 1 - \frac{\pi_{\text{rare}}}{1/4 - \Delta^2} &\geq \frac{1}{2}.
 \end{aligned}$$

Now we can bound $I_{\alpha\beta}^\top I_{\alpha\alpha}^{-1} I_{\alpha\beta}$ using (34):

$$\begin{aligned}
 I_{\alpha\beta}^\top I_{\alpha\alpha}^{-1} I_{\alpha\beta} &\preceq \frac{2}{I_{\alpha_1\alpha_1} I_{\alpha_2\alpha_2}} (I_{\alpha_2\alpha_2} I_{\beta\alpha_1} I_{\beta\alpha_1}^\top + I_{\alpha_1\alpha_1} I_{\beta\alpha_2} I_{\beta\alpha_2}^\top + |I_{\alpha_1\alpha_2}| I_{\beta\alpha_1} I_{\beta\alpha_2}^\top + |I_{\alpha_1\alpha_2}| I_{\beta\alpha_2} I_{\beta\alpha_1}^\top) \\
 &= 2 \left(\frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} + \frac{I_{\beta\alpha_2} I_{\beta\alpha_2}^\top}{I_{\alpha_2\alpha_2}} + \frac{|I_{\alpha_1\alpha_2}|}{I_{\alpha_1\alpha_1} I_{\alpha_2\alpha_2}} (I_{\beta\alpha_1} I_{\beta\alpha_2}^\top + I_{\beta\alpha_2} I_{\beta\alpha_1}^\top) \right),
 \end{aligned}$$

and

$$I_{\beta\beta} - I_{\alpha\beta}^\top I_{\alpha\alpha}^{-1} I_{\alpha\beta} \succeq I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} - 2 \left(\frac{I_{\beta\alpha_2} I_{\beta\alpha_2}^\top}{I_{\alpha_2\alpha_2}} + \frac{|I_{\alpha_1\alpha_2}|}{I_{\alpha_1\alpha_1} I_{\alpha_2\alpha_2}} (I_{\beta\alpha_1} I_{\beta\alpha_2}^\top + I_{\beta\alpha_2} I_{\beta\alpha_1}^\top) \right)$$

Note that $I_{\beta\beta}$, $I_{\beta\alpha_1} I_{\beta\alpha_1}^\top$, $I_{\beta\alpha_2} I_{\beta\alpha_2}^\top$, and $I_{\beta\alpha_1} I_{\beta\alpha_2}^\top + I_{\beta\alpha_2} I_{\beta\alpha_1}^\top$ are all symmetric. By Weyl's theorem, it holds that for symmetric matrices with matching dimensions \mathbf{A} and \mathbf{B} ,

$$\lambda_{\min}(\mathbf{A} - \mathbf{B}) \geq \lambda_{\min}(\mathbf{A}) - \|\mathbf{B}\|_{\text{op}},$$

because for any \mathbf{v}

$$(\mathbf{A} - \mathbf{B})\mathbf{v} = \mathbf{A}\mathbf{v} - \mathbf{B}\mathbf{v} \geq \lambda_{\min}(\mathbf{A})\mathbf{v} - \|\mathbf{B}\|_{\text{op}}\mathbf{v} = (\lambda_{\min}(\mathbf{A}) - \|\mathbf{B}\|_{\text{op}})\mathbf{v}.$$

So

$$\begin{aligned}
 & \lambda_{\min} \left(I_{\beta\beta} - I_{\alpha\beta}^\top I_{\alpha\alpha}^{-1} I_{\alpha\beta} \right) \\
 & \geq \lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right) - \frac{2}{I_{\alpha_2\alpha_2}} \|I_{\beta\alpha_2} I_{\beta\alpha_2}^\top\|_{\text{op}} - 2 \frac{|I_{\alpha_1\alpha_2}|}{I_{\alpha_1\alpha_1} I_{\alpha_2\alpha_2}} \|I_{\beta\alpha_1} I_{\beta\alpha_2}^\top + I_{\beta\alpha_2} I_{\beta\alpha_1}^\top\|_{\text{op}} \\
 & \stackrel{(a)}{\geq} \lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right) - \frac{2}{I_{\alpha_2\alpha_2}} \|I_{\beta\alpha_2} I_{\beta\alpha_2}^\top\|_{\text{op}} - 2 \frac{|I_{\alpha_1\alpha_2}|}{I_{\alpha_1\alpha_1} I_{\alpha_2\alpha_2}} \left(\|I_{\beta\alpha_1} I_{\beta\alpha_2}^\top\|_{\text{op}} + \|I_{\beta\alpha_2} I_{\beta\alpha_1}^\top\|_{\text{op}} \right) \\
 & \stackrel{(b)}{\geq} \lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right) - 2 \frac{\|I_{\beta\alpha_2}\|_2^2}{I_{\alpha_2\alpha_2}} - 2 \frac{|I_{\alpha_1\alpha_2}|}{I_{\alpha_1\alpha_1} I_{\alpha_2\alpha_2}} (2 \|I_{\beta\alpha_1}\|_2 \|I_{\beta\alpha_2}\|_2) \\
 & \stackrel{(c)}{\geq} \lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right) - 2M^2 \pi_{\text{rare}} \left(1 + \frac{\pi_{\text{rare}}}{1/2 - \Delta} \right) - 2 \frac{I_{\alpha_2\alpha_2}}{I_{\alpha_1\alpha_1} I_{\alpha_2\alpha_2}} (2 \|I_{\beta\alpha_1}\|_2 \cdot M \cdot I_{\alpha_2\alpha_2}) \\
 & = \lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right) - 2M \left(M \pi_{\text{rare}} \left(1 + \frac{\pi_{\text{rare}}}{1/2 - \Delta} \right) + 2 \frac{I_{\alpha_2\alpha_2}}{I_{\alpha_1\alpha_1}} \|I_{\beta\alpha_1}\|_2 \right) \\
 & \stackrel{(d)}{\geq} \lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right) - 2M \left(M \pi_{\text{rare}} \left(1 + \frac{\pi_{\text{rare}}}{1/2 - \Delta} \right) + 2 \frac{1}{I_{\alpha_1\alpha_1}} \|I_{\beta\alpha_1}\|_2 \cdot \pi_{\text{rare}} \left(1 + \frac{\pi_{\text{rare}}}{1/2 - \Delta} \right) \right) \\
 & = \lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right) - 2M \pi_{\text{rare}} \left(1 + \frac{\pi_{\text{rare}}}{1/2 - \Delta} \right) \left(M + 2 \frac{\|I_{\beta\alpha_1}\|_2}{I_{\alpha_1\alpha_1}} \right) \\
 & \stackrel{(e)}{\geq} \lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right) - 2M \pi_{\text{rare}} \left(1 + \frac{1}{2} \right) \left(M + \frac{M^2}{2} \right) \\
 & \stackrel{(f)}{\geq} \frac{1}{2} \lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right),
 \end{aligned}$$

where in (a) we used the triangle inequality, in (b) we used that for any $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$ it holds that $\|\mathbf{a}\mathbf{b}^\top\|_{\text{op}} = |\mathbf{b}^\top \mathbf{a}| \leq \|\mathbf{a}\|_2 \|\mathbf{b}\|_2$ (note that $\mathbf{a}\mathbf{b}^\top$ is rank one with eigenvector \mathbf{a}) as well as the triangle inequality, in (c) we used (29), (30), and (31), in (d) we used (28), (e) follows from (32) and

$$\frac{\pi_{\text{rare}}}{1/2 - \Delta} \leq \frac{1/2(1/2 - \Delta)(1/2 + \Delta)}{1/2 - \Delta} = \frac{1}{2} \left(\frac{1}{2} + \Delta \right) \leq \frac{1}{2} \left(\frac{1}{2} + \frac{1}{2} \right) \leq \frac{1}{2}$$

(since $\Delta \leq 1/2$), and in (f) we used our assumptions that

$$\lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right) > 0$$

and

$$\begin{aligned}
 \pi_{\text{rare}} & \leq \frac{\lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right)}{3M^2(2+M)} \\
 \iff 3M\pi_{\text{rare}} \left(M + \frac{M^2}{2} \right) & \leq \frac{1}{2} \lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right).
 \end{aligned}$$

Substituting this into (33) shows that

$$\left\| \text{Cov} \left(\lim_{n \rightarrow \infty} \sqrt{n} \left[\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}^{\text{prop. odds}} \right] \right) \right\|_{\text{op}} \leq \frac{2}{\lambda_{\min} \left(I_{\beta\beta} - 2 \frac{I_{\beta\alpha_1} I_{\beta\alpha_1}^\top}{I_{\alpha_1\alpha_1}} \right)}.$$

E.4. Proof of Lemma E.1

We omit \mathbf{x} in integrals, when it is clear (e.g., π_1 stands for $\pi_1(\mathbf{x})$).

Proof of (27):

From (9) and (11), using $\pi_1(\mathbf{x}) + \pi_2(\mathbf{x}) + \pi_3(\mathbf{x}) = 1$ we have

$$\begin{aligned}
 I_{\alpha_2\alpha_2} &= M_2 \\
 &= \int [\pi_3(\mathbf{x})(1 - \pi_3(\mathbf{x}))]^2 \left(\frac{1}{\pi_2(\mathbf{x})} + \frac{1}{\pi_3(\mathbf{x})} \right) dF(\mathbf{x}) \\
 &= \int [\pi_3(\mathbf{x})(1 - \pi_3(\mathbf{x}))]^2 \left(\frac{\pi_2(\mathbf{x}) + \pi_3(\mathbf{x})}{\pi_2(\mathbf{x})\pi_3(\mathbf{x})} \right) dF(\mathbf{x}) \\
 &= \int \pi_3(\mathbf{x})[1 - \pi_3(\mathbf{x})]^2 \left(\frac{1 - \pi_1(\mathbf{x})}{\pi_2(\mathbf{x})} \right) dF(\mathbf{x}) \\
 &\leq \pi_{\text{rare}} \int (1 - \pi_3)^2(1 - \pi_1)/\pi_2 dF \\
 &= \pi_{\text{rare}} \int \pi_1(1 - \pi_1)^2 \frac{(1 - \pi_3)^2}{\pi_2} \left[\frac{1}{\pi_1(1 - \pi_1)} \right] dF \\
 &\stackrel{(a)}{\leq} \pi_{\text{rare}} \int \pi_1(1 - \pi_1)^2 \frac{1 - \pi_3}{\pi_2} \left[\frac{1}{(1/2 - \Delta)(1/2 + \Delta)} \right] dF \\
 &= \pi_{\text{rare}} \int \pi_1^2(1 - \pi_1)^2 \frac{\pi_1 + \pi_2}{\pi_1\pi_2} \frac{1}{1/4 - \Delta^2} \\
 &= \pi_{\text{rare}} \int \pi_1^2(1 - \pi_1)^2 \left(\frac{1}{\pi_1} + \frac{1}{\pi_2} \right) \frac{1}{1/4 - \Delta^2} \\
 &= \pi_{\text{rare}} M_1 \frac{1}{1/4 - \Delta^2} \\
 &= \pi_{\text{rare}} J_{\alpha_1\alpha_1} \frac{1}{1/4 - \Delta^2},
 \end{aligned} \tag{35}$$

where in (a) we used the fact that $t \mapsto 1/[t(1-t)]$ is nonincreasing in t for $t \in (0, 1/2]$, so it is largest when t is small as possible, and by assumption $\inf_{\mathbf{x} \in \mathcal{S}} \{\pi_1 \wedge 1 - \pi_1\} \geq 1/2 - \Delta$.

Proof of (28): Using (35) we have

$$\begin{aligned}
 I_{\alpha_2\alpha_2} &= \int \pi_3(\mathbf{x})[1 - \pi_3(\mathbf{x})]^2 \left(\frac{\pi_2(\mathbf{x}) + \pi_3(\mathbf{x})}{\pi_2(\mathbf{x})} \right) dF(\mathbf{x}) \\
 &= \int \pi_3(1 - \pi_3)^2(1 + \pi_3/\pi_2) dF \\
 &\leq \pi_{\text{rare}} \int (1 - 0)^2 \left(1 + \frac{\pi_{\text{rare}}}{1/2 - \Delta} \right) dF.
 \end{aligned}$$

Proof of (29): Using (35) along with (8) and (12) we have

$$\begin{aligned}
 I_{\alpha_2\alpha_2} &= \int \pi_3(\mathbf{x})[1 - \pi_3(\mathbf{x})]^2 \left(\frac{1 - \pi_1(\mathbf{x})}{\pi_2(\mathbf{x})} \right) dF(\mathbf{x}) \\
 &= \int \pi_3(\mathbf{x})[1 - \pi_3(\mathbf{x})][1 - \pi_1(\mathbf{x})] \left(\frac{\pi_1(\mathbf{x}) + \pi_2(\mathbf{x})}{\pi_2(\mathbf{x})} \right) dF(\mathbf{x}) \\
 &= \int \pi_3(1 - \pi_3)(1 - \pi_1) \frac{\pi_1}{\pi_2} dF + \int \pi_3(1 - \pi_3)(1 - \pi_1) dF \\
 &= \tilde{M}_2 + \int \pi_3(1 - \pi_3)(1 - \pi_1) dF \\
 &= |I_{\alpha_1\alpha_2}| + \int \pi_3(1 - \pi_3)(1 - \pi_1) dF \\
 &\geq |I_{\alpha_1\alpha_2}|.
 \end{aligned}$$

Proof of (30): From (9) and (18) we have

$$\begin{aligned}
 \|I_{\beta\alpha_2}\|_2 &= \left\| \int \mathbf{x} [\pi_2(\mathbf{x}) + \pi_3(\mathbf{x})] \pi_3(\mathbf{x}) [1 - \pi_3(\mathbf{x})] dF(\mathbf{x}) \right\|_2 \\
 &= \left\| \int \mathbf{x} [1 - \pi_1(\mathbf{x})] \pi_3(\mathbf{x}) [1 - \pi_3(\mathbf{x})] dF(\mathbf{x}) \right\|_2 \\
 &\leq \int \|\mathbf{x}\|_2 \pi_3(\mathbf{x}) [1 - \pi_3(\mathbf{x})] [1 - \pi_1(\mathbf{x})] dF \\
 &\leq M \int \pi_3(\mathbf{x}) [1 - \pi_3(\mathbf{x})] [1 - \pi_1(\mathbf{x})] dF \\
 &\leq M \cdot \int \pi_3(\mathbf{x}) [1 - \pi_3(\mathbf{x})] [1 - \pi_1(\mathbf{x})] \left(1 + \frac{\pi_3(\mathbf{x})}{\pi_2(\mathbf{x})} \right) dF(\mathbf{x}) \\
 &= M \cdot \int \pi_3(\mathbf{x}) [1 - \pi_3(\mathbf{x})] [1 - \pi_1(\mathbf{x})] \left(\frac{\pi_2(\mathbf{x}) + \pi_3(\mathbf{x})}{\pi_2(\mathbf{x})} \right) dF(\mathbf{x}) \\
 &= M \cdot \int \pi_3(\mathbf{x}) [1 - \pi_3(\mathbf{x})] [1 - \pi_1(\mathbf{x})] \left(\frac{1 - \pi_1(\mathbf{x})}{\pi_2(\mathbf{x})} \right) dF(\mathbf{x}) \\
 &\leq M \cdot \int \pi_3(\mathbf{x}) [1 - \pi_3(\mathbf{x})]^2 \left(\frac{1 - \pi_1(\mathbf{x})}{\pi_2(\mathbf{x})} \right) dF(\mathbf{x}) \\
 &= M \cdot I_{\alpha_2\alpha_2},
 \end{aligned}$$

where in the last inequality we used $\pi_3(\mathbf{x}) \leq \pi_{\text{rare}} < 1/2 - \Delta \leq \pi_1(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{S}$ and in the last step we used (35).

Proof of (31): This follows from (28) and (30).

Proof of (32): Using from (9)

$$\begin{aligned}
 \|I_{\beta\alpha_1}\|_2 &= \left\| \int \mathbf{x} \pi_1^2(\mathbf{x}) [1 - \pi_1(\mathbf{x})] dF(\mathbf{x}) + \int \mathbf{x} \pi_2(\mathbf{x}) \pi_1(\mathbf{x}) [1 - \pi_1(\mathbf{x})] dF(\mathbf{x}) \right\|_2 \\
 &= \left\| \int \mathbf{x} [\pi_1(\mathbf{x}) + \pi_2(\mathbf{x})] \pi_1(\mathbf{x}) [1 - \pi_1(\mathbf{x})] dF(\mathbf{x}) \right\|_2 \\
 &\leq \int \|\mathbf{x}\|_2 [\pi_1(\mathbf{x}) + \pi_2(\mathbf{x})] \pi_1(\mathbf{x}) [1 - \pi_1(\mathbf{x})] dF(\mathbf{x}) \\
 &\leq M \int [\pi_1(\mathbf{x}) + \pi_2(\mathbf{x})] \pi_1(\mathbf{x}) [1 - \pi_1(\mathbf{x})] dF(\mathbf{x})
 \end{aligned}$$

and from (8)

$$\begin{aligned}
 I_{\alpha_1 \alpha_1} &= M_1 = \int (\pi_1(\mathbf{x})[1 - \pi_1(\mathbf{x})])^2 \left(\frac{1}{\pi_1(\mathbf{x})} + \frac{1}{\pi_2(\mathbf{x})} \right) dF(\mathbf{x}) \\
 &= \int (\pi_1(\mathbf{x})[1 - \pi_1(\mathbf{x})])^2 \frac{\pi_1(\mathbf{x}) + \pi_2(\mathbf{x})}{\pi_1(\mathbf{x})\pi_2(\mathbf{x})} dF(\mathbf{x}) \\
 &= \int \pi_1(\mathbf{x})[1 - \pi_1(\mathbf{x})][\pi_1(\mathbf{x}) + \pi_2(\mathbf{x})] \frac{1 - \pi_1(\mathbf{x})}{\pi_2(\mathbf{x})} dF(\mathbf{x}) \\
 &\geq \int \pi_1(\mathbf{x})[1 - \pi_1(\mathbf{x})][\pi_1(\mathbf{x}) + \pi_2(\mathbf{x})] dF(\mathbf{x}),
 \end{aligned}$$

we have

$$\frac{\|I_{\beta \alpha_1}\|_2^2}{I_{\alpha_1 \alpha_1}} \leq M^2 \int [\pi_1(\mathbf{x}) + \pi_2(\mathbf{x})] \pi_1(\mathbf{x})[1 - \pi_1(\mathbf{x})] dF(\mathbf{x}) \leq M^2 \int \frac{1}{4} [\pi_1(\mathbf{x}) + \pi_2(\mathbf{x})] dF(\mathbf{x}) \leq \frac{M^2}{4},$$

where in the second-to-last step we used that $t \mapsto t(1 - t) \leq 1/4$ for all $t \in [0, 1]$.

F. Proof of Lemma D.1

First we will calculate the Fisher information matrices, then we will show the convergence results.

Since the logistic regression model can be considered a special case of the proportional odds model with $K = 2$ categories, we will mostly focus our calculations on the proportional odds model.

F.1. Calculating the Log Likelihood and Gradients

In the proportional odds model, the likelihood can be expressed as

$$\prod_{i=1}^n \prod_{k=1}^K \left(\frac{1}{1 + \exp\{-\alpha_k + \boldsymbol{\beta}^\top \mathbf{x}_i\}} - \frac{1}{1 + \exp\{-\alpha_{k-1} + \boldsymbol{\beta}^\top \mathbf{x}_i\}} \right)^{\mathbb{1}\{y_i=k\}}, \quad (36)$$

so the log likelihood can be expressed as

$$\begin{aligned}
 \mathcal{L}^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) &= \sum_{i=1}^n \sum_{k=1}^K \mathbb{1}\{y_i = k\} \log \left(\frac{1}{1 + \exp\{-\alpha_k + \boldsymbol{\beta}^\top \mathbf{x}_i\}} - \frac{1}{1 + \exp\{-\alpha_{k-1} + \boldsymbol{\beta}^\top \mathbf{x}_i\}} \right) \\
 &= \sum_{i=1}^n \sum_{k=1}^K \mathbb{1}\{y_i = k\} \log(p_{ik} - p_{i,k-1}) \\
 &= \sum_{i=1}^n \sum_{k=1}^K \mathbb{1}\{y_i = k\} \log(\pi_{ik}),
 \end{aligned}$$

where

$$\begin{aligned}
 p_{ik} &:= p_k(\mathbf{x}_i) = \mathbb{P}(y_i \leq k \mid \mathbf{x}_i), \\
 \pi_{ik} &:= \mathbb{P}(y_i = k \mid \mathbf{x}_i) = p_{ik} - p_{i,k-1} = p_k(\mathbf{x}_i) - p_{k-1}(\mathbf{x}_i),
 \end{aligned}$$

$\alpha_0 := -\infty$, and $\alpha_K := \infty$ (while $\boldsymbol{\alpha} := (\alpha_1, \dots, \alpha_{K-1})^\top$ are parameters to be estimated). Using

$$\frac{\partial}{\partial t} \frac{1}{1 + \exp\{-t\}} = \frac{1}{1 + \exp\{-t\}} \left(1 - \frac{1}{1 + \exp\{-t\}} \right),$$

the gradient has entries corresponding to β equal to

$$\begin{aligned}\nabla_{\beta} \mathcal{L}^{\text{prop. odds}}(\alpha, \beta) &= \sum_{i=1}^n \sum_{k=1}^K \mathbb{1}\{y_i = k\} \mathbf{x}_i \left(\frac{p_{ik}(1 - p_{ik}) - p_{i,k-1}(1 - p_{i,k-1})}{p_{ik} - p_{i,k-1}} \right) \\ &= \sum_{i=1}^n \sum_{k=1}^K \mathbb{1}\{y_i = k\} \mathbf{x}_i \left(\frac{p_{ik} - p_{i,k-1} - (p_{ik}^2 - p_{i,k-1}^2)}{p_{ik} - p_{i,k-1}} \right) \\ &= \sum_{i=1}^n \sum_{k=1}^K \mathbb{1}\{y_i = k\} \mathbf{x}_i (1 - p_{ik} - p_{i,k-1}),\end{aligned}$$

and using

$$\begin{aligned}\frac{\partial}{\partial \alpha_k} \pi_{ik} &= \frac{\partial}{\partial \alpha_k} (p_{ik} - p_{i,k-1}) = p_{ik}(1 - p_{ik}) \quad \text{and} \\ \frac{\partial}{\partial \alpha_k} \pi_{i,k+1} &= \frac{\partial}{\partial \alpha_k} (p_{i,k+1} - p_{ik}) = -p_{ik}(1 - p_{ik}),\end{aligned}$$

the entries corresponding to α equal

$$\begin{aligned}\frac{\partial}{\partial \alpha_k} \mathcal{L}^{\text{prop. odds}}(\alpha, \beta) &= \sum_{i=1}^n \frac{\partial}{\partial \alpha_k} (\mathbb{1}\{y_i = k\} \log(\pi_{ik}) + \mathbb{1}\{y_i = k+1\} \log(\pi_{i,k+1})) \\ &= \sum_{i=1}^n \left(\mathbb{1}\{y_i = k\} \frac{p_{ik}(1 - p_{ik})}{\pi_{ik}} - \mathbb{1}\{y_i = k+1\} \frac{p_{ik}(1 - p_{ik})}{\pi_{i,k+1}} \right) \\ \implies \nabla_{\alpha} \mathcal{L}^{\text{prop. odds}}(\alpha, \beta) &= \sum_{i=1}^n \mathbf{e}_k p_{ik} (1 - p_{ik}) \left(\frac{\mathbb{1}\{y_i = k\}}{\pi_{ik}} - \frac{\mathbb{1}\{y_i = k+1\}}{\pi_{i,k+1}} \right),\end{aligned}$$

where $\mathbf{e}_k \in \{0, 1\}^{K-1}$ has the k^{th} entry equal to 1 and the rest equal to 0. (Note that since $\alpha_0 = -\infty$, $p_{i0} = 0$, and similarly $p_{iK} = 1$ as expected.)

F.2. Calculating the Hessian Matrices

The entries of the Hessian corresponding to β , $H_{\beta\beta}^{\text{prop. odds}}$ are

$$\begin{aligned}\nabla_{\beta}^2 \mathcal{L}^{\text{prop. odds}}(\alpha, \beta) &= - \sum_{i=1}^n \sum_{k=1}^K \mathbb{1}\{y_i = k\} \mathbf{x}_i \mathbf{x}_i^{\top} [p_{ik}(1 - p_{ik}) + p_{i,k-1}(1 - p_{i,k-1})] \\ &= - \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^{\top} \sum_{k=1}^K (\mathbb{1}\{y_i = k+1\} + \mathbb{1}\{y_i = k\}) \text{Var}(\mathbb{1}\{y_i \leq k\}).\end{aligned}$$

Using

$$\frac{\partial}{\partial \alpha_k} p_{ik}(1 - p_{ik}) = p_{ik}(1 - p_{ik}) - 2p_{ik}^2(1 - p_{ik}) = p_{ik}(1 - p_{ik})(1 - 2p_{ik}),$$

the entries corresponding to the α block of the Hessian $H_{\alpha\alpha}^{\text{prop. odds}}$ are as follows:

$$\begin{aligned}
 \frac{\partial^2}{\partial \alpha_k^2} \mathcal{L}^{\text{prop. odds}}(\alpha, \beta) &= \sum_{i=1}^n \frac{\partial}{\partial \alpha_k} \left(p_{ik}(1-p_{ik}) \left(\frac{\mathbb{1}\{y_i = k\}}{\pi_{ik}} - \frac{\mathbb{1}\{y_i = k+1\}}{\pi_{i,k+1}} \right) \right) \\
 &= \sum_{i=1}^n \left(\mathbb{1}\{y_i = k\} \frac{\pi_{ik} p_{ik} (1-p_{ik}) (1-2p_{ik}) - p_{ik}^2 (1-p_{ik})^2}{\pi_{ik}^2} \right. \\
 &\quad \left. - \mathbb{1}\{y_i = k+1\} \frac{\pi_{i,k+1} p_{ik} (1-p_{ik}) (1-2p_{ik}) + p_{ik}^2 (1-p_{ik})^2}{\pi_{i,k+1}^2} \right) \\
 &= \sum_{i=1}^n p_{ik}(1-p_{ik}) \left(\mathbb{1}\{y_i = k\} \frac{\pi_{ik} (1-2p_{ik}) - p_{ik} (1-p_{ik})}{\pi_{ik}^2} \right. \\
 &\quad \left. - \mathbb{1}\{y_i = k+1\} \frac{\pi_{i,k+1} (1-2p_{ik}) + p_{ik} (1-p_{ik})}{\pi_{i,k+1}^2} \right), \\
 \frac{\partial^2}{\partial \alpha_k \partial \alpha_{k-1}} \mathcal{L}^{\text{prop. odds}}(\alpha, \beta) &= \sum_{i=1}^n p_{ik}(1-p_{ik}) \frac{\partial}{\partial \alpha_{k-1}} \left(\frac{\mathbb{1}\{y_i = k\}}{\pi_{ik}} - \frac{\mathbb{1}\{y_i = k+1\}}{\pi_{i,k+1}} \right) \\
 &= - \sum_{i=1}^n \mathbb{1}\{y_i = k\} p_{ik}(1-p_{ik}) \left(\frac{-p_{i,k-1} (1-p_{i,k-1})}{\pi_{ik}^2} \right) \\
 &= \sum_{i=1}^n \mathbb{1}\{y_i = k\} \cdot \frac{p_{ik}(1-p_{ik}) p_{i,k-1} (1-p_{i,k-1})}{\pi_{ik}^2}, \quad \text{and} \\
 \frac{\partial^2}{\partial \alpha_k \partial \alpha_{k'}} \mathcal{L}^{\text{prop. odds}}(\alpha, \beta) &= 0, \quad \text{all other } k \neq k',
 \end{aligned}$$

where $\pi_{i,K+1} = 0$. (Note that $\frac{\partial^2}{\partial \alpha_k \partial \alpha_{k+1}} \mathcal{L}^{\text{prop. odds}}(\alpha, \beta)$ is nonzero as well, but it matches the expression for $\frac{\partial^2}{\partial \alpha_{\tilde{k}} \partial \alpha_{\tilde{k}-1}} \mathcal{L}^{\text{prop. odds}}(\alpha, \beta)$ with $\tilde{k} := k+1$.)

Finally, the entries corresponding to the α and β mixed blocks $H_{\alpha\beta}^{\text{prop. odds}}$ of the Hessian are

$$\begin{aligned}
 \frac{\partial}{\partial \alpha_k} \nabla_{\beta} \mathcal{L}^{\text{prop. odds}}(\alpha, \beta) &= \sum_{i=1}^n \left(\mathbb{1}\{y_i = k\} \mathbf{x}_i \frac{\partial}{\partial \alpha_k} (1-p_{ik} - p_{i,k-1}) \right. \\
 &\quad \left. + \mathbb{1}\{y_i = k+1\} \mathbf{x}_i \frac{\partial}{\partial \alpha_k} (1-p_{i,k+1} - p_{i,k}) \right) \\
 &= - \sum_{i=1}^n \mathbf{x}_i p_{ik} (1-p_{ik}) (\mathbb{1}\{y_i = k\} + \mathbb{1}\{y_i = k+1\}), \quad k \in [K-1].
 \end{aligned}$$

F.3. Calculation of the Fisher Information Matrices

Now we find the Fisher information matrices

$$I^{\text{prop. odds}}(\alpha, \beta) = \begin{pmatrix} I_{\alpha\alpha}^{\text{prop. odds}} & I_{\beta\alpha}^{\text{prop. odds}} \\ \left(I_{\beta\alpha}^{\text{prop. odds}} \right)^{\top} & I_{\beta\beta}^{\text{prop. odds}} \end{pmatrix}$$

and

$$I^{\text{logistic}}(\alpha, \beta) = \begin{pmatrix} I_{\alpha\alpha}^{\text{logistic}} & \left(I_{\beta\alpha}^{\text{logistic}} \right)^{\top} \\ I_{\beta\alpha}^{\text{logistic}} & I_{\beta\beta}^{\text{logistic}} \end{pmatrix}$$

by taking the negative expectation of each block (using a single observation). For the α block, we have

$$\begin{aligned}
 -\mathbb{E} \left[\frac{\partial^2}{\partial \alpha_k^2} \mathcal{L}^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) \right] &= -\mathbb{E} \left[p_{ik}(1-p_{ik}) \left(\mathbb{1}\{y_i = k\} \frac{\pi_{ik}(1-2p_{ik}) - p_{ik}(1-p_{ik})}{\pi_{ik}^2} \right. \right. \\
 &\quad \left. \left. - \mathbb{1}\{y_i = k+1\} \frac{\pi_{i,k+1}(1-2p_{ik}) + p_{ik}(1-p_{ik})}{\pi_{i,k+1}^2} \right) \right] \\
 &= -\mathbb{E} \left[p_{ik}(1-p_{ik}) \left(\frac{\pi_{ik}(1-2p_{ik}) - p_{ik}(1-p_{ik})}{\pi_{ik}^2} \right. \right. \\
 &\quad \left. \left. - \frac{\pi_{i,k+1}(1-2p_{ik}) + p_{ik}(1-p_{ik})}{\pi_{i,k+1}^2} \right) \right] \\
 &= \mathbb{E} \left[p_{ik}^2(1-p_{ik})^2 \left(\frac{1}{\pi_{ik}} + \frac{1}{\pi_{i,k+1}} \right) \right] \\
 &= M_k, \quad k \in [K-1], \\
 -\mathbb{E} \left[\frac{\partial^2}{\partial \alpha_k \partial \alpha_{k-1}} \mathcal{L}^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) \right] &= -\mathbb{E} \left[\mathbb{1}\{y_i = k\} \cdot \frac{p_{ik}(1-p_{ik})p_{i,k-1}(1-p_{i,k-1})}{\pi_{ik}^2} \right] \\
 &= -\mathbb{E} \left[\frac{p_{ik}(1-p_{ik})p_{i,k-1}(1-p_{i,k-1})}{\pi_{ik}^2} \right] \\
 &= -\tilde{M}_k, \quad k \in \{2, \dots, K-1\}, \quad \text{and} \\
 \mathbb{E} \left[\frac{\partial^2}{\partial \alpha_k \partial \alpha_{k'}} \mathcal{L}^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) \right] &= 0, \quad \text{all other } k \neq k',
 \end{aligned}$$

where we used the definitions of M_k and \tilde{M}_k in (11) and (12). Therefore $I_{\alpha\alpha}^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) \in \mathbb{R}^{(K-1) \times (K-1)}$ has tridiagonal form

$$I_{\alpha\alpha}^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \begin{pmatrix} M_1 & -\tilde{M}_2 & 0 & \cdots & 0 & 0 \\ -\tilde{M}_2 & M_2 & -\tilde{M}_3 & \cdots & 0 & 0 \\ 0 & -\tilde{M}_3 & M_3 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & M_{K-2} & -\tilde{M}_{K-1} \\ 0 & 0 & 0 & \cdots & -\tilde{M}_{K-1} & M_{K-1} \end{pmatrix},$$

verifying (8). Observe that in the case of logistic regression ($K = 2$), we have

$$\begin{aligned}
 I_{\alpha\alpha}^{\text{logistic}}(\alpha_1, \boldsymbol{\beta}) &= M_1 = \int [p_1(\mathbf{x})(1-p_1(\mathbf{x}))]^2 \left(\frac{1}{\pi_1(\mathbf{x})} + \frac{1}{\pi_2(\mathbf{x})} \right) dF(\mathbf{x}) \\
 &= \int [\pi_1(\mathbf{x})(1-\pi_1(\mathbf{x}))]^2 \left(\frac{1}{\pi_1(\mathbf{x})} + \frac{1}{1-\pi_1(\mathbf{x})} \right) dF(\mathbf{x}) \\
 &= \int [\pi_1(\mathbf{x})(1-\pi_1(\mathbf{x}))]^2 \cdot \frac{1}{\pi_1(\mathbf{x})(1-\pi_1(\mathbf{x}))} dF(\mathbf{x}) \\
 &= \int \pi_1(\mathbf{x})(1-\pi_1(\mathbf{x})) dF(\mathbf{x}) \\
 &= M_1^{\text{logistic}},
 \end{aligned}$$

which is (14). (This is equivalent to a logistic regression predicting whether $y_i = 1$ regardless of K .)

F.4. Mixed block

For the α - β mixed block, we have for all $k \in \{1, \dots, K-1\}$

$$-\mathbb{E} \left[\frac{\partial^2}{\partial \alpha_k \partial \boldsymbol{\beta}} \mathcal{L}^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) \right] = \mathbb{E} [\mathbf{X}_1 p_k(\mathbf{X}_1)(1-p_k(\mathbf{X}_1))(\pi_k(\mathbf{X}_1) + \pi_{k+1}(\mathbf{X}_1))].$$

Then

$$-\mathbb{E} \left[\frac{\partial^2}{\partial \alpha_k \partial \beta} \mathcal{L}^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) \right] = J_k^{\mathbf{x}} + \tilde{J}_{k+1}^{\mathbf{x}}, \quad k \in \{1, \dots, K-1\},$$

so $I_{\beta\alpha}^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) \in \mathbb{R}^{(K-1) \times p}$ has the form

$$I_{\beta\alpha}^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \begin{pmatrix} J_1^{\mathbf{x}} + \tilde{J}_2^{\mathbf{x}} \\ J_2^{\mathbf{x}} + \tilde{J}_3^{\mathbf{x}} \\ \vdots \\ J_{K-1}^{\mathbf{x}} + \tilde{J}_K^{\mathbf{x}} \end{pmatrix},$$

as in (9). In the case of logistic regression predicting whether $y_i = 1$,

$$I_{\beta\alpha}^{\text{logistic}}(\alpha_1, \boldsymbol{\beta}) = J_1^{\mathbf{x}; \text{logistic}} + \tilde{J}_2^{\mathbf{x}; \text{logistic}} = \int \mathbf{x} \pi_1(\mathbf{x}) [1 - \pi_1(\mathbf{x})] dF(\mathbf{x}),$$

verifying (15).

F.5. Beta block

Finally, for the β block, we have

$$\begin{aligned} I_{\beta\beta}^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) &= -\mathbb{E} \left[\nabla_{\boldsymbol{\beta}}^2 \mathcal{L}^{\text{prop. odds}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\mathbf{X}_1 \mathbf{X}_1^\top \sum_{k=1}^K (\mathbb{1}\{y_i = k+1\} + \mathbb{1}\{y_i = k\}) p_k(\mathbf{X}_1) (1 - p_k(\mathbf{X}_1)) \mid \mathbf{X} \right] \right] \\ &= \mathbb{E} \left[\mathbf{X}_1 \mathbf{X}_1^\top \sum_{k=1}^K (\pi_{k+1}(\mathbf{X}_{1\cdot}) + \pi_k(\mathbf{X}_{1\cdot})) p_k(\mathbf{X}_1) (1 - p_k(\mathbf{X}_1)) \right] \\ &= \sum_{k=1}^K \left(J_k^{\mathbf{x}\mathbf{x}^\top} + \tilde{J}_k^{\mathbf{x}\mathbf{x}^\top} \right), \end{aligned}$$

as in (10) (recall that $J_K^{\mathbf{x}\mathbf{x}^\top} = 0$ and $\tilde{J}_1^{\mathbf{x}\mathbf{x}^\top} = 0$ for all \mathbf{x}). In the case of logistic regression predicting whether $y_i = 1$ (for any K),

$$I_{\beta\beta}^{\text{logistic}}(\alpha_1, \boldsymbol{\beta}) = J_1^{\mathbf{x}\mathbf{x}^\top; \text{logistic}} + \tilde{J}_2^{\mathbf{x}\mathbf{x}^\top; \text{logistic}} = \int \mathbf{x} \mathbf{x}^\top \pi_1(\mathbf{x}) [1 - \pi_1(\mathbf{x})] dF(\mathbf{x}),$$

matching (16).

F.6. Verifying the Asymptotic Distribution of Each Estimator

By standard maximum likelihood theory (for example, the theorem on p. 145 of [Serfling \(1980, Section 4.2.2, multidimensional generalization on p. 148\)](#), or Theorem 13.1 in [Wooldridge 2010](#)), the result holds if we can verify three regularity conditions. Before we do, we will define the set Θ of feasible parameters $(\boldsymbol{\alpha}, \boldsymbol{\beta})$. Since

$$\alpha_1 < \alpha_2 < \dots < \alpha_{K-1},$$

where the strict inequality follows from our assumption that no class has probability 0 for any $x \in [-1, 1]$, define the set $\mathcal{A} \subset \mathbb{R}^{K-1}$ of points satisfying this constraint. Then define $\Theta := \mathcal{A} \times \mathbb{R}^p$, so that $(\boldsymbol{\alpha}, \boldsymbol{\beta}) \in \Theta$.

Now we state and verify the needed regularity conditions.

1. (R1) *The third derivatives of the log likelihood with respect to each parameter $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ exist for all $\mathbf{x} \in \mathcal{S}$. This condition holds for both the proportional odds model and logistic regression because every entry of the Hessian matrices in Section (F.2) are differentiable in every parameter for any $K \geq 2$.*

2. (R2) For each $(\alpha_0, \beta_0) \in \Theta$, for all (α, β) in a neighborhood of (α_0, β_0) it holds that (i) the element-wise absolute values of the gradients and Hessians of the likelihood are bounded by functions of \mathbf{x} with finite integral over $\mathbf{x} \in \mathcal{S}$, and (ii) the element-wise absolute values of the third derivatives of the log likelihood is bounded by a function of \mathbf{x} with finite expectation with respect to \mathbf{X} . Because \mathbf{X} has bounded support, for these integrals and expectations to be finite it is enough for the bounding functions to over \mathcal{S} to be finite constants—that is, it is enough to find upper bounds on the absolute values of the gradients and Hessians of the likelihoods and the third derivatives of the log likelihoods. The logistic regression likelihood

$$\prod_{i=1}^n \frac{\exp\left\{-\mathbb{1}\{y_i = 1\}(\alpha_1 + \beta^\top \mathbf{x}_i)\right\}}{1 + \exp\{-(\alpha_1 + \beta^\top \mathbf{x}_i)\}}$$

has continuous second derivatives and therefore its gradient and Hessian both have finite element-wise absolute value on the bounded support. The same is true of the proportional odds likelihood (36) when all outcomes have positive probability for all $\mathbf{x} \in \mathcal{S}$, that is, $\alpha_1 < \dots < \alpha_{K-1}$. Finally, examining again the Hessian matrices in Section F.2, we see that they have continuous derivatives in every parameter for any $K \geq 2$, so the third derivatives of the log likelihoods are bounded for any $(\alpha_0, \beta_0) \in \Theta$ for all $\mathbf{x} \in \mathcal{S}$.

3. (R3) The Fisher information matrices exist and are finite and positive definite. One can see that both of the Fisher information matrices are finite entrywise for all $(\alpha, \beta) \in \Theta$ by examining the matrices and noting that the probabilities for all of the classes are strictly greater than 0 over \mathcal{S} by assumption. To verify the positive definiteness of the Fisher information matrices

$$-\mathbb{E}\left[\frac{\partial^2}{\partial\theta\theta^\top}\mathcal{L}(\theta)\right],$$

it is enough to show that the log likelihood for each model is strictly concave, which implies that $\frac{\partial^2}{\partial\theta\theta^\top}\mathcal{L}(\theta)$ is almost surely negative definite (since the log likelihood is twice differentiable). Strict concavity of the logistic regression log likelihood

$$\sum_{i=1}^n \left[-\mathbb{1}\{y_i = 1\}(\alpha_1 + \beta^\top \mathbf{x}_i) - \log\left(1 + e^{-(\alpha_1 + \beta^\top \mathbf{x}_i)}\right)\right]$$

is easily seen, and Pratt (1981) provides a proof that the log likelihood for the proportional odds model is strictly concave when the intercepts $\alpha_1, \dots, \alpha_{K-1}$ are not equal.

Lastly, the asymptotic $\mathcal{O}(1/n)$ bias is a consequence of standard maximum likelihood theory; see, for example, Cordeiro & McCullagh (1991).

G. Proof of Theorem 3.1

We prove Theorem 3.1 in Section G.1, and Section G.2 contains proofs of the supporting lemmas.

G.1. Proof of Theorem 3.1

The proof will proceed as follows. First we will show that PRESTO is a member of a class of models described by Ekvall & Bottai (2022), which means we can bound the estimation error of the parameters of PRESTO in a finite sample under their Theorem 3 once we show that its assumptions are satisfied. Their result depends on the probability of a particular event $\mathcal{C}_{\kappa, n, p_n(K-1)}$, and in Proposition G.1 we prove a lower bound on $\mathbb{P}(\mathcal{C}_{\kappa, n, p_n(K-1)})$ that tends towards 1 as $n \rightarrow \infty$. This establishes the consistency of PRESTO.

In the notation of Ekvall & Bottai (2022), we can express the objective function for the PRESTO estimator from (7) as $R\{b(y_i, \mathbf{x}_i, \theta)\} - R\{a(y_i, \mathbf{x}_i, \theta)\} + \lambda_n \|\theta\|_1$, where $R(\cdot) = F(\cdot)$, the logistic cumulative distribution function;

$(a(y_i, \mathbf{x}_i, \boldsymbol{\theta}); b(y_i, \mathbf{x}_i, \boldsymbol{\theta}))^\top = \mathbf{Z}_i^\top \boldsymbol{\theta} + \mathbf{m}_i$ where

$$\mathbf{Z}_i := \begin{pmatrix} \mathbb{1}\{y_i \geq 2\} \mathbf{x}_i & \mathbb{1}\{y_i \leq K-1\} \mathbf{x}_i \\ \mathbb{1}\{y_i \geq 3\} \mathbf{x}_i & \mathbb{1}\{2 \leq y_i \leq K-1\} \mathbf{x}_i \\ \vdots & \vdots \\ \mathbb{1}\{y_i = K\} \mathbf{x}_i & \mathbb{1}\{y_i = K-1\} \mathbf{x}_i \end{pmatrix} \in \mathbb{R}^{p_n(K-1) \times 2},$$

$$\boldsymbol{\theta} := \begin{pmatrix} \beta_1 \\ \psi_2 \\ \vdots \\ \psi_{K-1} \end{pmatrix} \in \mathbb{R}^{p_n(K-1)}, \quad \text{and}$$

$$\mathbf{m}_i := \begin{pmatrix} \begin{cases} -\infty, & y_i = 1, \\ \alpha_{k-1}, & y_i = k \ (k \geq 2) \end{cases} \\ \begin{cases} \alpha_k, & y_i = k \ (k < K-1) \\ \infty, & y_i = K \end{cases} \end{pmatrix} \in \overline{\mathbb{R}}^2,$$

where $\overline{\mathbb{R}} := \mathbb{R} \cup \{-\infty, \infty\}$ denotes the extended real number system; and, as elsewhere in the paper, $\psi_k = \beta_k - \beta_{k-1}$ for $k \in \{2, \dots, K-1\}$. Observe that this is a special case of model (1) of Ekvall & Bottai (2022). Now, since $\Theta = \mathbb{R}^{p_n(K-1)}$ is open and the standard logistic density $r(t) = \exp\{-t\} / (1 + \exp\{-t\})^2$ is strictly log-concave, strictly positive, and continuously differentiable on \mathbb{R} , assumptions (a) and (b) of Theorem 3 in Ekvall & Bottai (2022) are satisfied, assumption $S(C_4)$ is sufficient for assumption (c), and assumption (e) is satisfied for $C_3 = c_2$.

We now discuss Assumption 1 from Ekvall & Bottai (2022). Note that a decision boundary crosses at some $\mathbf{x} \in \mathcal{S}$ if and only if for some $k \in \{2, \dots, K-1\}$ it holds that

$$\begin{aligned} F(\alpha_k + \beta_k^\top \mathbf{x}) - F(\alpha_{k-1} + \beta_{k-1}^\top \mathbf{x}) &\leq 0 \\ \iff b(k, \mathbf{x}, \boldsymbol{\theta}) - a(k, \mathbf{x}, \boldsymbol{\theta}) &\leq 0. \end{aligned}$$

Observe that for $k=1$ it holds that $m_{i1} = -\infty$ and for $k=K$ it holds that $m_{i2} = \infty$; that is, an element of \mathbf{m}_i is infinite. So Assumption $T(C_4)$ is sufficient for it to either hold that $\mathbf{Z}_i^\top \boldsymbol{\theta} + \mathbf{m}_i \in E$ where $E \subseteq \{t \in \mathbb{R}^2 : t_1 < t_2\}$ or an element of \mathbf{m}_i is infinite. To satisfy Assumption 1, it only remains to show that there exists such an E that is compact. Note that $\|\boldsymbol{\theta}_*\|_\infty$ is bounded under Assumption $S(s, C_4)$ and $\|\boldsymbol{\theta}_*\|_0 \leq s$ under our sparsity assumption, so $\|\boldsymbol{\theta}_*\|_1$ is bounded due to Hölder's inequality. Lastly \mathbf{m}_i is bounded because $\max_{k \in \{1, \dots, K-1\}} |\alpha_k| \leq C_4$ under Assumption $T(C_4)$, so we have shown that we can find such an E that is bounded. Choose one that is closed and we have E compact by Bolzano-Weierstrass.

We have shown that the assumptions of Theorem 3 of Ekvall & Bottai (2022) are satisfied, so we have that for n large enough that $p_n \geq K-1$, with probability at least $\mathbb{P}(\mathcal{C}_{\kappa, n, p_n(K-1)}) - [p_n(K-1)]^{-c_3} \geq \mathbb{P}(\mathcal{C}_{\kappa, n, p_n(K-1)}) - (p_n^2)^{-c_3}$

$$\|\hat{\boldsymbol{\theta}}^{\lambda_n} - \boldsymbol{\theta}_*\|_1 \leq c_5 \sqrt{\frac{\log(p_n(K-1))}{n}} \leq c_5 \sqrt{\frac{\log(p_n^2)}{n}} = c_5 \sqrt{2} \sqrt{\frac{\log p_n}{n}}, \quad (37)$$

where c_3 and c_5 are constants from Ekvall & Bottai (2022), and $\mathcal{C}_{\kappa, n, p_n(K-1)}$ is defined as follows. For a set $\mathcal{A} \subseteq \{1, \dots, p_n(K-1)\}$, define $\boldsymbol{\theta}_{\mathcal{A}} \in \mathbb{R}^{p_n(K-1)}$ to have j^{th} entry

$$(\boldsymbol{\theta}_{\mathcal{A}})_j := \begin{cases} \boldsymbol{\theta}_j, & j \in \mathcal{A}, \\ 0, & j \notin \mathcal{A}, \end{cases}$$

and define $\boldsymbol{\theta}_{\mathcal{A}^c} := \boldsymbol{\theta} - \boldsymbol{\theta}_{\mathcal{A}}$. Then for an s -sparse (in the sense of Assumption $S(s, c)$) $\boldsymbol{\theta}_*$ with support set \mathcal{S} , define

$$\mathbb{C}(\mathcal{S}) := \left\{ \boldsymbol{\theta} \in \mathbb{R}^{p_n(K-1)} : \|\boldsymbol{\theta}_{\mathcal{S}^c}\|_1 \leq 3 \|\boldsymbol{\theta}_{\mathcal{S}}\|_1 \right\}.$$

We interpret $\mathbb{C}(\mathcal{S})$ to be a set of approximately s -sparse vectors (the vectors would be exactly s -sparse if $\|\boldsymbol{\theta}_{\mathcal{S}^c}\|_1 = 0$). Then for $\kappa > 0$ and $n \in \mathbb{N}$, define

$$\mathcal{C}_{\kappa, n, p_n(K-1)} := \left\{ (\mathbf{X}, \mathbf{y}) : \inf_{\{\boldsymbol{\theta} \in \mathbb{C}(\mathcal{S}) : \|\boldsymbol{\theta}\|_2 = 1\}} \left\{ \boldsymbol{\theta}^\top \left(\frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i \mathbf{Z}_i^\top \right) \boldsymbol{\theta} \right\} \geq \kappa \right\}.$$

If the set of θ over which this condition must hold were $\mathbb{R}^{p_n(K-1)}$, this would be a minimum eigenvalue condition on $\frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i \mathbf{Z}_i^\top$. This condition is sometimes called a *restricted eigenvalue condition* (Bickel et al., 2009). We bound $\mathbb{P}(C_{\kappa, n, p_n(K-1)})$ in Proposition G.1.

Proposition G.1. *Suppose the assumptions of Theorem 3.1 hold. Let*

$$\pi_{rare, min} := \inf_{\mathbf{x} \in \mathcal{S}, k \in \{1, \dots, K\}} \{\mathbb{P}(y_i = k \mid \mathbf{x})\},$$

and observe that $\pi_{rare, min} > 0$ under Assumptions $X(\mathbb{R}^{p_n})$, $S(s, C_4)$, and $T(C_4)$. Assume n is large enough so that

$$n\pi_{rare, min} > \max \left\{ 2 \left(C\sqrt{p_n} + \sqrt{\frac{a}{\lambda_{min}^*}} \right)^2, 2 \right\} \quad (38)$$

for some $a > 0$ (recall from the statement of Theorem 3.1 that we assumed $\lambda_{min}^* := \min_{k \in \{1, \dots, K\}} \lambda_{min}(\mathbb{E}[\mathbf{x}_i \mathbf{x}_i^\top \mid y_i = k]) > b$ for a fixed $b > 0$). Then

$$\mathbb{P}(C_{a\pi_{rare, min}/4, n, p_n(K-1)}) \geq 1 - 2K \exp \left\{ -\frac{1}{2} \pi_{rare, min}^2 n \right\} - 2K \exp \left\{ -c \left(\sqrt{\frac{n\pi_{rare, min}}{2}} - C\sqrt{p_n} - \sqrt{\frac{a}{\lambda_{min}^*}} \right)^2 \right\}.$$

Proof. Provided in Section G.2. □

Observe that since $p_n \leq C_1 n^{C_2}$ for some $C_1 > 0, C_2 \in (0, 1)$ this probability tends to 1 as $n \rightarrow \infty$. Lemma G.2, below, along with (37) then shows that with probability at least

$$1 - p_n^{-C_5} - 2K \exp \left\{ -\frac{1}{2} \pi_{rare, min}^2 n \right\} - 2K \exp \left\{ -c \left(\sqrt{\frac{n\pi_{rare, min}}{2}} - C\sqrt{p_n} - \sqrt{\frac{a}{\lambda_{min}^*}} \right)^2 \right\} \quad (39)$$

it holds that

$$\left\| \hat{\beta}^{\lambda_n} - \beta \right\|_2 \leq \left\| \hat{\beta}^{\lambda_n} - \beta \right\|_1 \leq C_6 \sqrt{\frac{\log p_n}{n}},$$

where $C_5 := 2c_3$ and $C_6 := c_5 \sqrt{2}(K-1)$ (where c_5 depends on the sparsity level s).

Lemma G.2. $\left\| \hat{\beta}^{\lambda_n} - \beta \right\|_2 \leq \left\| \hat{\beta}^{\lambda_n} - \beta \right\|_1 \leq (K-1) \left\| \hat{\theta}^{\lambda_n} - \theta_* \right\|_1$.

Proof. Provided in Section G.2. □

Finally, we can now show consistency by showing that the random variable $\left\| \hat{\beta}^{\lambda_n} - \beta \right\|_2$ converges in probability to 0. For any $\epsilon > 0$,

$$\begin{aligned} & \lim_{n \rightarrow \infty} \mathbb{P} \left(\left\| \hat{\beta}^{\lambda_n} - \beta \right\|_2 < \epsilon \right) \\ & \stackrel{(*)}{\geq} \lim_{n \rightarrow \infty} \mathbb{P} \left(\left\| \hat{\beta}^{\lambda_n} - \beta \right\|_2 < C_6 \sqrt{\frac{\log p_n}{n}} \right) \\ & \geq \lim_{n \rightarrow \infty} \left(1 - p_n^{-C_5} - 2K \exp \left\{ -\frac{1}{2} \pi_{rare, min}^2 n \right\} - 2K \exp \left\{ -c \left(\sqrt{\frac{n\pi_{rare, min}}{2}} - C\sqrt{p_n} - \sqrt{\frac{a}{\lambda_{min}^*}} \right)^2 \right\} \right) \\ & = 1, \end{aligned}$$

where $(*)$ follows because for large enough n , $\epsilon > C_6 \sqrt{\log(p_n)/n}$. This establishes Theorem 3.1. All that remains is to provide proofs for the supporting lemmas and proposition, which we do in the following section.

Remark G.3. Observe that for finite n we have that for n large enough, with probability at least equal to the expression in (39) it holds that

$$\sqrt{n} \left\| \hat{\beta}^{\lambda_n} - \beta \right\|_2 \leq \sqrt{n} \left\| \hat{\beta}^{\lambda_n} - \beta \right\|_1 \leq C_6 \sqrt{\log p_n}.$$

This establishes a high-probability finite sample relationship between our theoretical guarantee and the asymptotic mean squared error from Definition 2.1.

Also note that for any $\epsilon > 0$, there exists an n_ϵ such that the probability in (39) is less than ϵ for all $n \geq n_\epsilon$. Since there also exists an n_2 large enough so that for all $n \geq n_2$ it holds that

$$\epsilon \geq C_6 \sqrt{\frac{\log p_n}{n}},$$

we have for any $\epsilon > 0$, for all $n \geq \max\{n_\epsilon, n_2\}$

$$\begin{aligned} \mathbb{P} \left(\frac{\left\| \hat{\beta}^{\lambda_n} - \beta \right\|_1}{\sqrt{\log(p_n)/n}} \geq C_6 \right) &\leq 1 - p_n^{-C_5} - 2K \exp \left\{ -\frac{1}{2} \pi_{\text{rare}, \min}^2 n \right\} - 2K \exp \left\{ -c \left(\sqrt{\frac{n \pi_{\text{rare}, \min}}{2}} - C \sqrt{p_n} - \sqrt{\frac{a}{\lambda_{\min}^*}} \right)^2 \right\} \\ &< \epsilon; \end{aligned}$$

that is, $\left\| \hat{\beta}^{\lambda_n} - \beta \right\|_1 \in \mathcal{O}_p(\sqrt{\log(p_n)/n})$, and likewise for $\left\| \hat{\beta}^{\lambda_n} - \beta \right\|_2$. Again, we emphasize that C_6 depends on the sparsity s (though the above holds since we have assumed that s is fixed).

G.2. Supporting Results for Proof of Theorem 3.1

Proof of Proposition G.1. We will show that there is a high probability event such that for $a > 0$ it holds that

$$\inf_{\{\theta \in \mathbb{C}(S) : \|\theta\|_2 = 1\}} \left\{ \theta^\top \left(\frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i \mathbf{Z}_i^\top \right) \theta \right\} > \frac{1}{2} a \pi_{\text{rare}, \min}.$$

Let

$$\mathbf{A}^{(1)} = (\mathbf{0}_{K-1} \quad \mathbf{e}_1) \in \mathbb{R}^{(K-1) \times 2};$$

$$\mathbf{A}^{(k)} = \begin{pmatrix} \mathbf{1}_{k-1} & \mathbf{1}_{k-1} \\ 0 & 1 \\ \mathbf{0}_{K-k-1} & \mathbf{0}_{K-k-1} \end{pmatrix} = \left(\sum_{k'=1}^{k-1} \mathbf{e}_{k'} \quad \sum_{k'=1}^k \mathbf{e}_{k'} \right) \in \mathbb{R}^{(K-1) \times 2}, \quad k \in \{2, \dots, K-1\};$$

and

$$\mathbf{A}^{(K)} = (\mathbf{1}_{K-1} \quad \mathbf{0}_{K-1}) = \left(\sum_{k'=1}^{K-1} \mathbf{e}_{k'} \quad \mathbf{0}_{K-1} \right) \in \mathbb{R}^{(K-1) \times 2},$$

where $\mathbf{0}_n$ and $\mathbf{1}_n$ are n -vectors of zeroes and ones (respectively) and \mathbf{e}_k is the standard basis vector in \mathbb{R}^{K-1} with a 1 in the k^{th} entry and zeroes elsewhere, and $\mathbf{A}^{(k)} \in \mathbb{R}^{(K-1) \times 2}$ for all k . Note that

$$\mathbf{Z}_i = \mathbf{A}^{(y_i)} \otimes \mathbf{x}_i. \quad (40)$$

Let

$$\mathbf{B} = \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ 0 & 1 & 1 & \cdots & 1 \\ 0 & 0 & 1 & \cdots & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix} = \left(\mathbf{e}_1 \quad \sum_{k'=1}^2 \mathbf{e}_{k'} \quad \cdots \quad \sum_{k'=1}^{K-1} \mathbf{e}_{k'} \right) \in \mathbb{R}^{(K-1) \times (K-1)};$$

that is, the columns of \mathbf{B} are $\mathbf{B}_k = \sum_{\ell=1}^k \mathbf{e}_\ell$. We will make use of the following lemmas, which we prove later in this section:

Lemma G.4.

$$\sum_{k=1}^K \frac{n_k}{n} \mathbf{A}^{(k)} \left(\mathbf{A}^{(k)} \right)^\top = \mathbf{B} \mathbf{D} \mathbf{B}^\top,$$

where

$$\mathbf{D} := \text{diag} \left(\frac{n_1 + n_2}{n}, \dots, \frac{n_{K-1} + n_K}{n} \right) \in \mathbb{R}^{(K-1) \times (K-1)}.$$

Lemma G.5. $\sigma_{\min}^2(\mathbf{B}) \geq 1/2$.

Lemma G.6. For arbitrary matrices $\bar{\mathbf{A}}$, $\bar{\mathbf{B}}$, and $\bar{\mathbf{C}}$, if $\bar{\mathbf{B}} \succeq \bar{\mathbf{C}}$ and $\bar{\mathbf{A}} \succeq \mathbf{0}$ then $\bar{\mathbf{A}} \otimes \bar{\mathbf{B}} \succeq \bar{\mathbf{A}} \otimes \bar{\mathbf{C}}$.

Lemma G.7. With probability at least $1 - K \exp \left\{ -\frac{1}{2} \pi_{\text{rare}, \min}^2 n \right\}$ it holds that $\min\{n_1, \dots, n_K\} > \pi_{\text{rare}, \min} n / 2$, where $n_k := \sum_{i=1}^n \mathbb{1}\{y_i = k\}$ is the number of observations in class k .

Lemma G.8. Under the assumptions of Proposition G.1, there exist constants $c, C > 0$ such that $\lambda_{\min} \left(\frac{1}{n_k} \sum_{i: y_i = k} \mathbf{x}_i \mathbf{x}_i^\top \right) \geq a$ for all $k \in \{1, \dots, K\}$ with probability at least

$$1 - 2K \exp \left\{ -c \left(\sqrt{\frac{n \pi_{\text{rare}, \min}}{2}} - C \sqrt{p_n} - \sqrt{\frac{a}{\lambda_{\min}^*}} \right)^2 \right\} - K \exp \left\{ -\frac{1}{2} \pi_{\text{rare}, \min}^2 n \right\},$$

where n_k is the number of observations in class k and $\lambda_{\min}^* := \min_{k \in \{1, \dots, K\}} \left\{ \lambda_{\min} \left(\mathbb{E} \left[\mathbf{X}^\top \mathbf{X} \mid y = k \right] \right) \right\}$.

Using a union bound, there exists an event with probability at least

$$1 - 2K \exp \left\{ -\frac{1}{2} \pi_{\text{rare}, \min}^2 n \right\} - 2K \exp \left\{ -c \left(\sqrt{\frac{n \pi_{\text{rare}, \min}}{2}} - C \sqrt{p_n} - \sqrt{\frac{a}{\lambda_{\min}^*}} \right)^2 \right\}$$

on which the conclusions of all of the above lemmas hold for n large enough. On this event we have for n sufficiently large

$$\begin{aligned}
 \inf_{\{\boldsymbol{\theta} \in \mathbb{C}(\mathcal{S}) : \|\boldsymbol{\theta}\|_2 = 1\}} \left\{ \boldsymbol{\theta}^\top \left(\frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i \mathbf{Z}_i^\top \right) \boldsymbol{\theta} \right\} &\geq \inf_{\{\boldsymbol{\theta} \in \mathbb{R}^{p_n(K-1)} : \|\boldsymbol{\theta}\|_2 = 1\}} \left\{ \boldsymbol{\theta}^\top \left(\frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i \mathbf{Z}_i^\top \right) \boldsymbol{\theta} \right\} \\
 &= \lambda_{\min} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i \mathbf{Z}_i^\top \right) \\
 &\stackrel{(a)}{=} \lambda_{\min} \left(\frac{1}{n} \sum_{i=1}^n \left(\mathbf{A}^{(y_i)} \otimes \mathbf{x}_i \right) \left(\mathbf{A}^{(y_i)} \otimes \mathbf{x}_i \right)^\top \right) \\
 &= \lambda_{\min} \left(\frac{1}{n} \sum_{i=1}^n \left(\mathbf{A}^{(y_i)} \left(\mathbf{A}^{(y_i)} \right)^\top \right) \otimes \left(\mathbf{x}_i \mathbf{x}_i^\top \right) \right) \\
 &= \lambda_{\min} \left(\frac{1}{n} \sum_{k=1}^K \sum_{i: y_i = k} \left(\mathbf{A}^{(y_i)} \left(\mathbf{A}^{(y_i)} \right)^\top \right) \otimes \left(\mathbf{x}_i \mathbf{x}_i^\top \right) \right) \\
 &= \lambda_{\min} \left(\frac{1}{n} \sum_{k=1}^K \left(\mathbf{A}^{(k)} \left(\mathbf{A}^{(k)} \right)^\top \right) \otimes \left(\sum_{i: y_i = k} \mathbf{x}_i \mathbf{x}_i^\top \right) \right) \\
 &= \lambda_{\min} \left(\sum_{k=1}^K \left(\frac{n_k}{n} \mathbf{A}^{(k)} \left(\mathbf{A}^{(k)} \right)^\top \right) \otimes \left(\frac{1}{n_k} \sum_{i: y_i = k} \mathbf{x}_i \mathbf{x}_i^\top \right) \right) \\
 &\stackrel{(b)}{\geq} \lambda_{\min} \left(\sum_{k=1}^K \left(\frac{n_k}{n} \mathbf{A}^{(k)} \left(\mathbf{A}^{(k)} \right)^\top \right) \otimes (a \mathbf{I}_p) \right) \\
 &\stackrel{(c)}{=} \lambda_{\min} \left(\sum_{k=1}^K \frac{n_k}{n} \mathbf{A}^{(k)} \left(\mathbf{A}^{(k)} \right)^\top \right) \lambda_{\min} (a \mathbf{I}_p) \\
 &\stackrel{(d)}{=} \lambda_{\min} \left(\mathbf{B} \mathbf{D} \mathbf{B}^\top \right) \lambda_{\min} (a \mathbf{I}_p) \\
 &\stackrel{(e)}{\geq} a \min_{k \in \{1, \dots, K-1\}} \left\{ \frac{n_k + n_{k+1}}{n} \right\} \lambda_{\min} \left(\mathbf{B} \mathbf{B}^\top \right) \\
 &\stackrel{(f)}{\geq} \frac{a}{4} \min_{k \in \{1, \dots, K-1\}} \left\{ \frac{n_k + n_{k+1}}{n} \right\} \\
 &\stackrel{(g)}{>} \frac{a \pi_{\text{rare}, \min}}{4},
 \end{aligned}$$

where in (a) we used (40), (b) holds on the high probability event by Lemmas G.6 and G.8 because $\lambda_{\min} \left(\frac{1}{n_k} \sum_{i: y_i = k} \mathbf{x}_i \mathbf{x}_i^\top \right) \geq a$ implies $\frac{1}{n_k} \sum_{i: y_i = k} \mathbf{x}_i \mathbf{x}_i^\top \geq a \mathbf{I}_p$ for all k , in (c) we used the fact that for matrices \mathbf{M}, \mathbf{N} the eigenvalues of $\mathbf{M} \otimes \mathbf{N}$ are the products of the eigenvalues of \mathbf{M} and \mathbf{N} and both of the factor matrices are positive semidefinite, in (d) we applied Lemma G.4, (e) follows from $\mathbf{B} \mathbf{D} \mathbf{B}^\top \succeq \min_{k \in \{1, \dots, K-1\}} \left\{ \frac{n_k + n_{k+1}}{n} \right\} \mathbf{B} \mathbf{B}^\top$, in (f) we applied Lemma G.5, and (g) follows from Lemma G.7. That is, on this high probability event it holds that

$$\inf_{\{\boldsymbol{\theta} \in \mathbb{C}(\mathcal{S}) : \|\boldsymbol{\theta}\|_2 = 1\}} \left\{ \boldsymbol{\theta}^\top \left(\frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i \mathbf{Z}_i^\top \right) \boldsymbol{\theta} \right\} > \frac{a \pi_{\text{rare}, \min}}{4},$$

and the conclusion follows. □

Proof of Lemma G.2. For convenience, denote $\theta_1 = \beta_1$ and $\theta_k = \psi_k$, $k \in \{2, \dots, K-1\}$. Let

$$\hat{\epsilon}_{\theta}^{\lambda_n} = \left(\left(\hat{\epsilon}_{\theta,1}^{\lambda_n} \right)^\top, \left(\hat{\epsilon}_{\theta,2}^{\lambda_n} \right)^\top, \dots, \left(\hat{\epsilon}_{\theta,K-1}^{\lambda_n} \right)^\top \right)^\top := \left(\left(\hat{\theta}_1^{\lambda_n} - \theta_1 \right)^\top, \left(\hat{\theta}_2^{\lambda_n} - \theta_2 \right)^\top, \dots, \left(\hat{\theta}_{K-1}^{\lambda_n} - \theta_{K-1} \right)^\top \right)^\top.$$

Note that $\beta_k = \sum_{k' \leq k} \theta_{k'}$. Let $\beta := (\beta_1^\top, \beta_2^\top, \dots, \beta_{K-1}^\top)^\top$, and denote by

$$\hat{\beta}_k^{\lambda_n} := \sum_{k' \leq k} \hat{\theta}_{k'}, \quad k \in \{1, \dots, K-1\}$$

the estimates of β yielded from the estimates of θ . Let

$$\begin{aligned} \hat{\epsilon}_{\beta}^{\lambda_n} &= \left(\left(\hat{\epsilon}_{\beta,1}^{\lambda_n} \right)^\top, \left(\hat{\epsilon}_{\beta,2}^{\lambda_n} \right)^\top, \dots, \left(\hat{\epsilon}_{\beta,K-1}^{\lambda_n} \right)^\top \right)^\top \\ &:= \left(\left(\hat{\beta}_1^{\lambda_n} - \beta_1 \right)^\top, \left(\hat{\beta}_2^{\lambda_n} - \beta_2 \right)^\top, \dots, \left(\hat{\beta}_{K-1}^{\lambda_n} - \beta_{K-1} \right)^\top \right)^\top, \end{aligned}$$

and observe that

$$\begin{aligned} \hat{\epsilon}_{\beta}^{\lambda_n} &= \left(\left(\hat{\beta}_1^{\lambda_n} - \beta_1 \right)^\top, \left(\hat{\beta}_2^{\lambda_n} - \beta_2 \right)^\top, \dots, \left(\hat{\beta}_{K-1}^{\lambda_n} - \beta_{K-1} \right)^\top \right)^\top \\ &= \left(\left(\hat{\theta}_1^{\lambda_n} - \theta_1 \right)^\top, \left(\sum_{k' \leq 2} \hat{\theta}_{k'} - \sum_{k' \leq 2} \theta_{k'} \right)^\top, \dots, \left(\sum_{k' \leq K-1} \hat{\theta}_{k'} - \sum_{k' \leq K-1} \theta_{k'} \right)^\top \right)^\top \\ &= \left(\left(\hat{\epsilon}_{\theta,1}^{\lambda_n} \right)^\top, \left(\sum_{k' \leq 2} \hat{\epsilon}_{\theta,k'}^{\lambda_n} \right)^\top, \dots, \left(\sum_{k' \leq K-1} \hat{\epsilon}_{\theta,k'}^{\lambda_n} \right)^\top \right)^\top. \end{aligned} \tag{41}$$

Then

$$\begin{aligned} \left\| \hat{\epsilon}_{\beta,1}^{\lambda_n} \right\|_1 &= \sum_{k=1}^{K-1} \left\| \hat{\epsilon}_{\beta,k}^{\lambda_n} \right\|_1 \stackrel{(a)}{=} \sum_{k=1}^{K-1} \left\| \sum_{k' \leq k} \hat{\epsilon}_{\theta,k'}^{\lambda_n} \right\|_1 \stackrel{(b)}{\leq} \sum_{k=1}^{K-1} \sum_{k' \leq k} \left\| \hat{\epsilon}_{\theta,k'}^{\lambda_n} \right\|_1 = \sum_{k=1}^{K-1} (K-k) \left\| \hat{\epsilon}_{\theta,k}^{\lambda_n} \right\|_1 \\ &\leq (K-1) \sum_{k=1}^{K-1} \left\| \hat{\epsilon}_{\theta,k}^{\lambda_n} \right\|_1 = (K-1) \left\| \hat{\epsilon}_{\theta}^{\lambda_n} \right\|_1. \end{aligned}$$

where in (a) we used (41) and in (b) we used the triangle inequality. Lastly, $\left\| \hat{\beta}^{\lambda_n} - \beta \right\|_2 \leq \left\| \hat{\beta}^{\lambda_n} - \beta \right\|_1$ is a property of the ℓ_1 and ℓ_2 norms. \square

Proof of Lemma G.4.

$$\begin{aligned}
 & \sum_{k=1}^K \frac{n_k}{n} \mathbf{A}^{(k)} \left(\mathbf{A}^{(k)} \right)^\top \\
 &= \frac{n_1}{n} \mathbf{A}^{(1)} \left(\mathbf{A}^{(1)} \right)^\top + \sum_{k=2}^{K-1} \frac{n_k}{n} \mathbf{A}^{(k)} \left(\mathbf{A}^{(k)} \right)^\top + \frac{n_K}{n} \mathbf{A}^{(K)} \left(\mathbf{A}^{(K)} \right)^\top \\
 &= \frac{n_1}{n} \left(\mathbf{0}_{K-1} \quad \mathbf{e}_1 \right) \begin{pmatrix} \mathbf{0}_{K-1}^\top \\ \mathbf{e}_1^\top \end{pmatrix} + \sum_{k=2}^{K-1} \frac{n_k}{n} \left(\sum_{k'=1}^{k-1} \mathbf{e}_{k'} \quad \sum_{k'=1}^k \mathbf{e}_{k'} \right) \begin{pmatrix} \sum_{k'=1}^{k-1} \mathbf{e}_{k'}^\top \\ \sum_{k'=1}^k \mathbf{e}_{k'}^\top \end{pmatrix} \\
 &\quad + \frac{n_K}{n} \left(\sum_{k'=1}^{K-1} \mathbf{e}_{k'} \quad \mathbf{0}_{K-1} \right) \begin{pmatrix} \sum_{k'=1}^{K-1} \mathbf{e}_{k'}^\top \\ \mathbf{0}_{K-1}^\top \end{pmatrix} \\
 &= \frac{n_1}{n} \left(\mathbf{0}_{K-1} \mathbf{0}_{K-1}^\top + \mathbf{e}_1 \mathbf{e}_1^\top \right) + \sum_{k=2}^{K-1} \frac{n_k}{n} \left(\sum_{k'=1}^{k-1} \mathbf{e}_{k'} \sum_{k'=1}^{k-1} \mathbf{e}_{k'}^\top + \sum_{k'=1}^k \mathbf{e}_{k'} \sum_{k'=1}^k \mathbf{e}_{k'}^\top \right) \\
 &\quad + \frac{n_K}{n} \left(\sum_{k'=1}^{K-1} \mathbf{e}_{k'} \sum_{k'=1}^{K-1} \mathbf{e}_{k'}^\top + \mathbf{0}_{K-1} \mathbf{0}_{K-1}^\top \right) \\
 &= \frac{n_1}{n} \mathbf{B}_1 \mathbf{B}_1^\top + \sum_{k=2}^{K-1} \frac{n_k}{n} \left(\mathbf{B}_{k-1} \mathbf{B}_{k-1}^\top + \mathbf{B}_k \mathbf{B}_k^\top \right) + \frac{n_K}{n} \mathbf{B}_{K-1} \mathbf{B}_{K-1}^\top \\
 &= \sum_{k=1}^{K-1} \frac{n_k + n_{k+1}}{n} \mathbf{B}_k \mathbf{B}_k^\top \\
 &= \mathbf{B} \mathbf{D} \mathbf{B}^\top.
 \end{aligned}$$

□

Proof of Lemma G.5. \mathbf{B} is full rank with inverse

$$\mathbf{B}^{-1} = \begin{pmatrix} 1 & -1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & -1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & -1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 \end{pmatrix}.$$

We have

$$\|\mathbf{B}^{-1}\|_{\text{op}} \leq \sqrt{\|\mathbf{B}^{-1}\|_1 \|\mathbf{B}^{-1}\|_\infty} = \sqrt{2 \cdot 2} = 2,$$

where we have used that the ℓ_1 norm of each column and row is at most 2. Then the result follows from $\sigma_{\min}(\mathbf{B}) = 1/\|\mathbf{B}^{-1}\|_{\text{op}}$. □

Proof of Lemma G.6. The matrix $\bar{\mathbf{A}} \otimes (\bar{\mathbf{B}} - \bar{\mathbf{C}})$ is the Kronecker product of two positive semidefinite matrices and is therefore positive semidefinite. □

Proof of Lemma G.7. First we state a lemma we will use.

Lemma G.9. *Suppose that for all $\mathbf{x} \in \mathcal{S}$ and $k \in \{1, \dots, K\}$ it holds that $\mathbb{P}(y_i = k \mid \mathbf{x}) \geq \pi_{\text{rare}, \min}$. Let n_k be the number of observations in class k from a data set of size n . Then for any $q \in \{1, \dots, \lfloor n\pi_{\text{rare}, \min} \rfloor\}$,*

$$\mathbb{P} \left(\bigcap_{k=1}^K \{n_k > q\} \right) \geq 1 - K \exp \left\{ -2n \left(\pi_{\text{rare}, \min} - \frac{q}{n} \right)^2 \right\}.$$

Proof. Provided later in this section. □

Note that $n\pi_{\text{rare},\min}/2 \leq \lfloor n\pi_{\text{rare},\min} \rfloor$ because $x/2 \leq \lfloor x \rfloor$ for all $x \geq 2$ and $n\pi_{\text{rare},\min} \geq 2$ due to assumption (38). So the assumptions of Lemma G.9 are satisfied for $q = n\pi_{\text{rare},\min}/2$, and there are more than $n\pi_{\text{rare},\min}/2$ observations in each class with probability at least

$$1 - K \exp \left\{ -2n (\pi_{\text{rare},\min} - \pi_{\text{rare},\min}/2)^2 \right\} = 1 - K \exp \left\{ -\frac{1}{2} \pi_{\text{rare},\min}^2 n \right\}.$$

□

Proof of Lemma G.8. We will prove the result by using a concentration inequality on the minimum singular value of a random matrix (which will correspond to the square root of the minimum eigenvalue of $\frac{1}{n_k} \sum_{i:y_i=k} \mathbf{x}_i \mathbf{x}_i^\top$; recall that the eigenvalues of $\mathbf{A}^\top \mathbf{A}$ are the squares of the singular values of \mathbf{A}). However, the result we use applies only to random matrices with isotropic second moment matrices, so we need to standardize $\frac{1}{n_k} \sum_{i:y_i=k} \mathbf{x}_i \mathbf{x}_i^\top$ by its inverse square root second moment matrix. We can relate this quantity to the minimum eigenvalue of $\frac{1}{n_k} \sum_{i:y_i=k} \mathbf{x}_i \mathbf{x}_i^\top$ by our claim (that we will verify later) that for a symmetric positive semidefinite matrix \mathbf{S} and symmetric positive definite $\mathbf{\Sigma}$ it holds that

$$\lambda_{\min}(\mathbf{S}) \geq \lambda_{\min}(\mathbf{\Sigma}^{-1/2} \mathbf{S} \mathbf{\Sigma}^{-1/2}) \lambda_{\min}(\mathbf{\Sigma}). \quad (42)$$

Then substituting $\mathbf{S} = \frac{1}{n_k} \sum_{i:y_i=k} \mathbf{x}_i \mathbf{x}_i^\top$ and $\mathbf{\Sigma} = \mathbf{\Sigma}_k = \mathbb{E}[\mathbf{x}_i \mathbf{x}_i^\top \mid y_i = k]$ (which we assumed has a strictly positive minimum eigenvalue and is therefore invertible) into (42) yield that on an event where $n_k = \sum_{i=1}^n \mathbb{1}\{y_i = k\} \geq 1$, we have that almost surely

$$\begin{aligned} \lambda_{\min} \left(\frac{1}{n_k} \sum_{i:y_i=k} \mathbf{x}_i \mathbf{x}_i^\top \right) &\geq \lambda_{\min} \left(\frac{1}{n_k} \mathbf{\Sigma}_k^{-1/2} \sum_{i:y_i=k} \mathbf{x}_i \mathbf{x}_i^\top \mathbf{\Sigma}_k^{-1/2} \right) \lambda_{\min}(\mathbf{\Sigma}_k) \\ &\geq \lambda_{\min} \left(\frac{1}{n_k} \mathbf{\Sigma}_k^{-1/2} \sum_{i:y_i=k} \mathbf{x}_i \mathbf{x}_i^\top \mathbf{\Sigma}_k^{-1/2} \right) \lambda_{\min}^*, \end{aligned} \quad (43)$$

where the second line uses our assumption from the statement of Theorem 3.1 and the fact that $\frac{1}{n_k} \mathbf{\Sigma}_k^{-1/2} \sum_{i:y_i=k} \mathbf{x}_i \mathbf{x}_i^\top \mathbf{\Sigma}_k^{-1/2}$ is almost surely positive semidefinite. Next we lower-bound the minimum eigenvalue of the random matrix $\frac{1}{n_k} \mathbf{\Sigma}_k^{-1/2} \sum_{i:y_i=k} \mathbf{x}_i \mathbf{x}_i^\top \mathbf{\Sigma}_k^{-1/2}$ using a concentration inequality from Vershynin (2012). Observe that (as Vershynin 2018 points out in Example 2.5.8(c)) the ψ_2 norm of a bounded random vector $\mathbf{T} \in \mathbb{R}^p$ can be upper-bounded as follows:

$$\begin{aligned} \|\mathbf{T}\|_{\psi_2} &= \sup_{\mathbf{v} \in \mathbb{R}^p, \|\mathbf{v}\|_2=1} \left\{ \|\mathbf{T}^\top \mathbf{v}\|_{\psi_2} \right\} \\ &= \sup_{\mathbf{v} \in \mathbb{R}^p, \|\mathbf{v}\|_2=1} \left\{ \inf \left\{ t > 0 : \mathbb{E} \exp \left(\frac{[\mathbf{T}^\top \mathbf{v}]^2}{t^2} \right) \leq 2 \right\} \right\} \\ &\leq \inf \left\{ t > 0 : \exp \left(\frac{\|\mathbf{T}\|_\infty^2}{t^2} \right) \leq 2 \right\} \\ &= \frac{\|\mathbf{T}\|_\infty}{\sqrt{\log 2}}. \end{aligned}$$

Therefore under our assumptions $\mathbf{x}_i \mid \mathbf{y}$ is bounded and therefore subgaussian, with ψ_2 norm at most $\|\mathbf{x}_i\|_{\psi_2} \leq \|\mathbf{x}\|_\infty / \sqrt{\log 2} = C_4 / \sqrt{\log 2}$ for all k . So for any $k \in \{1, \dots, K\}$, we have from Theorem 5.39 in Vershynin (2012) that for any $t \geq 0$ there exists $c > 0$ such that the event

$$\mathbb{P} \left(\sigma_{\min} \left(\frac{1}{n_k} \sum_{i:y_i=k} \mathbf{\Sigma}_k^{-1/2} \mathbf{x}_i \right) \geq \sqrt{n_k} - C\sqrt{p_n} - t \mid \mathbf{y} \right) \geq 1 - 2 \exp\{-ct^2\} \quad (44)$$

holds almost surely for $C = C_4^2 / (\log 2) \sqrt{\log(9)/c_1}$, where C_4 is as defined in the statement of Theorem 3.1, c_1 is a constant from Vershynin (2012), $\sigma_{\min}(\cdot)$ denotes the minimum singular value, and if the set $\{i : y_i = k\}$ is empty we define

$\frac{1}{n_k} \sum_{i:y_i=k} \Sigma_k^{-1/2} \mathbf{x}_i$ to equal $\mathbf{0}_p$ (and note that the inequality then trivially holds with probability one in this case, because $n_k = 0$ so the right side is nonpositive). For $k \in \{1, \dots, K\}$ and $t \geq 0$, define the event $\mathcal{E}_k(t)$ by

$$\mathcal{E}_k(t) := \left\{ \sigma_{\min} \left(\frac{1}{n_k} \sum_{i:y_i=k} \Sigma_k^{-1/2} \mathbf{x}_i \right) \geq \sqrt{n_k} - C\sqrt{p_n} - t \right\},$$

and observe that

$$\mathbb{P}(\mathcal{E}_k(t)) \geq 1 - 2 \exp\{-ct^2\} \quad (45)$$

by marginalizing (44) over \mathbf{y} . Now we consider a particular choice of t . From our assumption (38) we have

$$\begin{aligned} n\pi_{\text{rare, min}} &> 2 \left(C\sqrt{p_n} + \sqrt{\frac{a}{\lambda_{\min}^*}} \right)^2 \\ \iff \sqrt{\frac{n\pi_{\text{rare, min}}}{2}} - C\sqrt{p_n} - \sqrt{\frac{a}{\lambda_{\min}^*}} &> 0, \end{aligned}$$

so we can choose $t = \sqrt{n\pi_{\text{rare, min}}/2} - C\sqrt{p_n} - \sqrt{\frac{a}{\lambda_{\min}^*}}$, yielding that on

$$\mathcal{E}_k \left(\sqrt{\frac{n\pi_{\text{rare, min}}}{2}} - C\sqrt{p_n} - \sqrt{\frac{a}{\lambda_{\min}^*}} \right) \cap \{n_k \geq 1\}$$

we have

$$\begin{aligned} \sigma_{\min} \left(\frac{1}{n_k} \sum_{i:y_i=k} \Sigma_k^{-1/2} \mathbf{x}_i \right) &\geq \sqrt{n_k} - C\sqrt{p_n} - \sqrt{\frac{n\pi_{\text{rare, min}}}{2}} + C\sqrt{p_n} + \sqrt{\frac{a}{\lambda_{\min}^*}} \\ &= \sqrt{n_k} - \sqrt{\frac{n\pi_{\text{rare, min}}}{2}} + \sqrt{\frac{a}{\lambda_{\min}^*}}. \end{aligned} \quad (46)$$

That is, on this event we can lower-bound the minimum eigenvalue of $\frac{1}{n_k} \Sigma_k^{-1/2} \sum_{i:y_i=k} \mathbf{x}_i \mathbf{x}_i^\top \Sigma_k^{-1/2}$ provided that n_k is at least 1 and large enough that the right side of (46) is nonnegative. Next we work on lower-bounding n_k with high probability. Consider the event

$$\mathcal{N} := \left\{ n_k \geq \frac{n\pi_{\text{rare, min}}}{2} \quad \forall k \in \{1, \dots, K\} \right\}.$$

Notice that on \mathcal{N} we have that the lower bound in (46) is at least $\sqrt{a/\lambda_{\min}^*}$ and $n_k \geq 1$ for all k due to (38). That is, for any k , on $\mathcal{N} \cap \mathcal{E}_k \left(\sqrt{n\pi_{\text{rare, min}}/2} - C\sqrt{p_n} - \sqrt{\frac{a}{\lambda_{\min}^*}} \right)$ inequality (46) holds and yields $\sigma_{\min} \left(\frac{1}{n_k} \sum_{i:y_i=k} \Sigma_k^{-1/2} \mathbf{x}_i \right) \geq \sqrt{a/\lambda_{\min}^*}$. Since the eigenvalues of $\frac{1}{n_k} \sum_{i:y_i=k} \Sigma_k^{-1/2} \mathbf{x}_i \mathbf{x}_i^\top \Sigma_k^{-1/2}$ are the squares of the singular values of $\frac{1}{n_k} \sum_{i:y_i=k} \Sigma_k^{-1/2} \mathbf{x}_i$, on $\mathcal{N} \cap \mathcal{E}_k \left(\sqrt{n\pi_{\text{rare, min}}/2} - C\sqrt{p_n} - \sqrt{\frac{a}{\lambda_{\min}^*}} \right)$ we have

$$\lambda_{\min} \left(\frac{1}{n_k} \sum_{i:y_i=k} \Sigma_k^{-1/2} \mathbf{x}_i \mathbf{x}_i^\top \Sigma_k^{-1/2} \right) \geq \frac{a}{\lambda_{\min}^*} > 0.$$

Finally, substituting this into (43) we have that on $\left(\bigcap_{k=1}^K \mathcal{E}_k \left(\sqrt{n\pi_{\text{rare, min}}/2} - C\sqrt{p_n} - \sqrt{\frac{a}{\lambda_{\min}^*}} \right) \right) \cap \mathcal{N}$

$$\lambda_{\min} \left(\frac{1}{n_k} \sum_{i:y_i=k} \mathbf{x}_i \mathbf{x}_i^\top \right) \geq \frac{a}{\lambda_{\min}^*} \lambda_{\min}^* = a \quad \forall k \in \{1, \dots, K\}.$$

Using a union bound, this holds with probability at least

$$\begin{aligned}
 & \mathbb{P} \left(\left(\bigcap_{k=1}^K \mathcal{E}_k \left(\sqrt{\frac{n\pi_{\text{rare},\min}}{2}} - C\sqrt{p_n} - \sqrt{\frac{a}{\lambda_{\min}^*}} \right) \right) \cap \mathcal{N} \right) \\
 & \geq 1 - \sum_{k=1}^K \mathbb{P} \left(\mathcal{E}_k^c \left(\sqrt{\frac{n\pi_{\text{rare},\min}}{2}} - C\sqrt{p_n} - \sqrt{\frac{a}{\lambda_{\min}^*}} \right) \right) - \mathbb{P}(\mathcal{N}^c) \\
 & \geq 1 - 2K \exp \left\{ -c \left(\sqrt{\frac{n\pi_{\text{rare},\min}}{2}} - C\sqrt{p_n} - \sqrt{\frac{a}{\lambda_{\min}^*}} \right)^2 \right\} - K \exp \left\{ -\frac{1}{2} \pi_{\text{rare},\min}^2 n \right\},
 \end{aligned}$$

where in the last step we applied (45) and Lemma G.7.

Finally, we show (42). Observe that for any \mathbf{v} with $\|\mathbf{v}\|_2 = 1$

$$\begin{aligned}
 \mathbf{v}^\top \mathbf{S} \mathbf{v} &= \left(\boldsymbol{\Sigma}^{1/2} \mathbf{v} \right)^\top \boldsymbol{\Sigma}^{-1/2} \mathbf{S} \boldsymbol{\Sigma}^{-1/2} \left(\boldsymbol{\Sigma}^{1/2} \mathbf{v} \right) \\
 &\geq \lambda_{\min} \left(\boldsymbol{\Sigma}^{-1/2} \mathbf{S} \boldsymbol{\Sigma}^{-1/2} \right) \|\boldsymbol{\Sigma}^{1/2} \mathbf{v}\|_2^2 \\
 &= \lambda_{\min} \left(\boldsymbol{\Sigma}^{-1/2} \mathbf{S} \boldsymbol{\Sigma}^{-1/2} \right) \mathbf{v}^\top \boldsymbol{\Sigma} \mathbf{v} \\
 &\geq \lambda_{\min} \left(\boldsymbol{\Sigma}^{-1/2} \mathbf{S} \boldsymbol{\Sigma}^{-1/2} \right) \lambda_{\min}(\boldsymbol{\Sigma}),
 \end{aligned}$$

proving the claim. □

Proof of Lemma G.9. By Hoeffding's inequality we have that for any $k \in \{1, \dots, K\}$ and any $q \in \{1, \dots, \lfloor n\pi_{\text{rare},\min} \rfloor\}$, $\mathbb{P}(n_k \leq q) \leq \exp \left\{ -2n \left(\pi_{\text{rare},\min} - \frac{q}{n} \right)^2 \right\}$. Then using a union bound, we have

$$\begin{aligned}
 \mathbb{P} \left(\bigcap_{k=1}^K \{n_k > q\} \right) &= 1 - \mathbb{P} \left(\bigcup_{k=1}^K \{n_k \leq q\} \right) \\
 &\geq 1 - \sum_{k=1}^K \mathbb{P}(n_k \leq q) \\
 &\geq 1 - K \exp \left\{ -2n \left(\pi_{\text{rare},\min} - \frac{q}{n} \right)^2 \right\}.
 \end{aligned}$$

□

H. Estimating PRESTO

Code generating all plots and tables that appear in the paper is available in the public GitHub repo at <https://github.com/gregfaletto/presto>. Executing the code in that repo generates the exact plots that appear in this version of the paper.

In all data experiments, we implemented PRESTO by slightly modifying the code for the R `ordinalNet` package (version 2.12), as we discuss below.

The real data experiments from Section 4.3 and Appendix B were conducted in R Version 4.3.0 running on macOS Ventura 13.3.1 on a MacBook Pro with a 2.3 GHz Quad-Core Intel Core i5 processor and 16 GB of RAM. We used the R packages MASS (Venables & Ripley, 2002, version 7.3.58.1), simulator (Bien, 2016, version 0.2.4), ggplot2 (Wickham, 2016, version 3.3.6), cowplot (Wilke, 2020, version 1.1.1), and stargazer (Hlavac, 2022, version 5.2.3), all available for download on CRAN, as well as the base `parallel` package (version 4.3.0). As mentioned earlier, the data for the real data experiment from Section 4.3 is the `soup` data set from the R `ordinal` package (version 2022.11.16), and the data for the real data experiment from Appendix B is the `PreDiabetes` data set from the R `MLDataR` package (version 1.0.1)

The synthetic data experiments from Sections 4.1 and 4.2, as well as Simulation Studies A and B in Appendix E, were conducted in R Version 4.2.2 running on macOS 10.15.7 on an iMac with a 3.5 GHz Quad-Core Intel Core i7 processor and 32 GB of RAM. We used the R packages MASS (version 7.3.58.1), simulator (version 0.2.4), ggplot2 (version 3.3.6), cowplot (version 1.1.1), and stargazer (version 5.2.3), all available for download on CRAN, as well as the base parallel package (version 4.2.2).

In the remainder of this section, we describe our implementation of PRESTO. We will use the notation and terminology of Wurm et al. (2021), with the exception of continuing to use our convention of K total categories. We fit PRESTO by reparameterizing the efficient coordinate descent algorithm used to estimate ℓ_1 -penalized ordinal regression models in the R ordinalNet (Wurm et al., 2021) package, in much the same way that the generalized lasso can be implemented using reparameterization; see Tibshirani & Taylor (2011, Section 3) or Section 4.5.1.1 of Hastie et al. (2015) for a textbook-level discussion.

The ordinalNet package implements elastic net penalized ordinal regression, including an ℓ_1 -penalized (or ℓ_2 -penalized, as we explored in Section C.1) relaxation of the proportional odds model. The parameters we seek to model are $\boldsymbol{\eta}_i \in \mathbb{R}^{K-1}$, which model the probabilities of the K outcomes by the relation $\boldsymbol{\eta}_i = g(\mathbf{p}_i)$, where $\mathbf{p}_i \in \mathcal{S}^{K-1}$ (where $\mathcal{S}^{K-1} := \{\mathbf{p} : \mathbf{p} \in (0, 1)^{K-1}, \|\mathbf{p}\|_1 < 1\}$) are the probabilities of outcomes $\{1, \dots, K-1\}$; in particular, $p_{ik} = \mathbb{P}(y_i = k)$ for $k \in \{1, \dots, K-1\}$ and $\mathbb{P}(y_i = K) = 1 - \sum_{k=1}^{K-1} p_{ik}$) and $g : \mathcal{S}^{K-1} \rightarrow \mathbb{R}^{K-1}$ is an invertible function. For our model, the forwards cumulative probability model, where $p_{ik} = \mathbb{P}(y_i \leq k)$,

$$[g(\mathbf{p}_i)]_k = \log \left(\frac{\sum_{k'=1}^k p_{ik'}}{1 - \sum_{k'=1}^k p_{ik'}} \right), \quad k \in \{1, \dots, K\}.$$

We choose the nonparallel model $\boldsymbol{\eta}_i = \mathbf{c} + \mathbf{B}^\top \mathbf{x}_i$, where \mathbf{x}_i is the i^{th} row of \mathbf{X} , $\mathbf{c} \in \mathbb{R}^{K-1}$ is a vector of intercepts, and

$$\mathbf{B} = [\mathbf{B}_{\cdot 1} \quad \cdots \quad \mathbf{B}_{\cdot K-1}]$$

is a $p \times (K-1)$ matrix of coefficients. Observe that we can simply write $\boldsymbol{\eta}_i = \tilde{\mathbf{X}}_i \boldsymbol{\beta}$, for $\boldsymbol{\beta} \in \mathbb{R}^{(K-1)(p+1)}$ defined by

$$\boldsymbol{\beta} = \begin{pmatrix} \mathbf{c} \\ \mathbf{B}_{\cdot 1} \\ \vdots \\ \mathbf{B}_{\cdot K-1} \end{pmatrix}$$

and $\tilde{\mathbf{X}}_i \in \mathbb{R}^{(K-1) \times (K-1)(p+1)}$ defined by

$$\tilde{\mathbf{X}}_i = \begin{pmatrix} \mathbf{x}_i^\top & \mathbf{0}_p^\top & \cdots & \mathbf{0}_p^\top & \mathbf{0}_p^\top \\ \mathbf{0}_p^\top & \mathbf{x}_i^\top & \cdots & \mathbf{0}_p^\top & \mathbf{0}_p^\top \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0}_p^\top & \mathbf{0}_p^\top & \cdots & \mathbf{x}_i^\top & \mathbf{0}_p^\top \\ \mathbf{0}_p^\top & \mathbf{0}_p^\top & \cdots & \mathbf{0}_p^\top & \mathbf{x}_i^\top \end{pmatrix} = (\mathbf{I}_{K-1} \quad \mathbf{I}_{K-1} \otimes \mathbf{x}_i^\top).$$

The R ordinalNet package solves the convex (Wurm et al., 2017; Pratt, 1981) optimization problem

$$\arg \min_{\boldsymbol{\beta}} \left\{ -\frac{1}{n} \sum_{i=1}^n \ell_i \left(g^{-1} \left(\tilde{\mathbf{X}}_i^\top \boldsymbol{\beta} \right) \right) + \lambda \sum_{q=K}^{(K-1)(p+1)} |\beta_q| \right\}$$

where

$$\ell_i(\mathbf{p}_i) = \sum_{k=1}^{K-1} \mathbb{1}\{y_i = k\} \log p_{ik} + \mathbb{1}\{y_i = K\} \log \left(1 - \sum_{k=1}^{K-1} p_{ik} \right).$$

We would like to place an ℓ_1 penalty on the first differences $B_{j,k+1} - B_{j,k}$ for all $k \in \{1, \dots, K-2\}$. We can do this through the parameterization $\boldsymbol{\Psi} \in \mathbb{R}^{p \times (K-1)}$ defined by

$$\Psi_{jk} = \begin{cases} B_{j1}, & j \in [p], k = 1, \\ B_{jk} - B_{j,k-1}, & j \in [p], k \in \{2, \dots, K-1\}. \end{cases}$$

Observe that these matrices are related by

$$B_{jk} = \sum_{k'=1}^k \Psi_{jk'},$$

so

$$\eta_{ik} = c_{0k} + \mathbf{B}_{\cdot k}^\top \mathbf{x}_i = c_{0k} + \sum_{k'=1}^k \Psi_{\cdot k'}^\top \mathbf{x}_i, \quad k \in \{1, \dots, K-1\}.$$

Therefore we can simply write

$$\boldsymbol{\eta}_i = \tilde{\mathbf{X}}_i' \boldsymbol{\beta}',$$

for $\boldsymbol{\beta}' \in \mathbb{R}^{(K-1)(p+1)}$ defined by

$$\boldsymbol{\beta}' = \begin{pmatrix} \mathbf{c} \\ \Psi_{\cdot 1} \\ \vdots \\ \Psi_{\cdot K-1} \end{pmatrix}$$

and $\tilde{\mathbf{X}}_i' \in \mathbb{R}^{(K-1) \times (K-1)(p+1)}$ defined by

$$\tilde{\mathbf{X}}_i' = \begin{pmatrix} \mathbf{x}_i^\top & \mathbf{0}_p^\top & \cdots & \mathbf{0}_p^\top & \mathbf{0}_p^\top \\ \mathbf{x}_i^\top & \mathbf{x}_i^\top & \cdots & \mathbf{0}_p^\top & \mathbf{0}_p^\top \\ \mathbf{I}_{K-1} & \vdots & \ddots & \vdots & \vdots \\ \mathbf{x}_i^\top & \mathbf{x}_i^\top & \cdots & \mathbf{x}_i^\top & \mathbf{0}_p^\top \\ \mathbf{x}_i^\top & \mathbf{x}_i^\top & \cdots & \mathbf{x}_i^\top & \mathbf{x}_i^\top \end{pmatrix}. \quad (47)$$

We then seek to solve the slightly modified optimization problem (modification highlighted)

$$\arg \min_{\boldsymbol{\beta}'} \left\{ -\frac{1}{n} \sum_{i=1}^n \ell_i \left(g^{-1} \left(\underbrace{\left(\tilde{\mathbf{X}}_i' \right)^\top}_{(*)} \boldsymbol{\beta}' \right) \right) + \lambda \sum_{q=K}^{(K-1)(p+1)} |\beta'_q| \right\}$$

Though this can not be implemented within the framework of the existing `ordinalNet` package, the above modification only requires changing a handful of lines of the publicly available source code in the `ordinalNet` package. Though our implementation is simple, using the modified design matrix (the change above) could cause convergence of parameter estimation to be slow, because the resulting lasso problem effectively has highly correlated features, which slows the convergence of coordinate descent. See Section 4.5.1.1 of [Hastie et al. \(2015\)](#) for a textbook-level discussion of this point. Though this is a convenient way to estimate PRESTO, it is not the best approach for estimating PRESTO in practice, and the limitations of our estimation approach are not limitations of PRESTO itself. Known efficient strategies for solving optimization problems like PRESTO are discussed in [Hastie et al. \(2015, Section 4.5\)](#) and [Arnold & Tibshirani \(2016\)](#). Though PRESTO does not strictly lie within the class of optimization problems considered by [Xin et al. \(2014\)](#), some of their ideas might also be helpful for better estimating PRESTO. An ADMM approach like the one proposed by [\(Zhu, 2017\)](#) for generalized lasso problems may also be useful for estimating PRESTO. [Ko et al. \(2019\)](#) propose their own methods for solving similar optimization problems.