

Undersampling is a Minimax Optimal Robustness Intervention in Nonparametric Classification

Niladri S. Chatterji*

*Department of Computer Science
Stanford University*

niladri@cs.stanford.edu

Saminul Haque*

*Department of Computer Science
Stanford University*

saminulh@stanford.edu

Tatsunori B. Hashimoto

*Department of Computer Science
Stanford University*

thashim@stanford.edu

Reviewed on OpenReview: <https://openreview.net/forum?id=r6oHDYOZ6p>

Abstract

While a broad range of techniques have been proposed to tackle distribution shift, the simple baseline of training on an *undersampled* balanced dataset often achieves close to state-of-the-art-accuracy across several popular benchmarks. This is rather surprising, since undersampling algorithms discard excess majority group data. To understand this phenomenon, we ask if learning is fundamentally constrained by a lack of minority group samples. We prove that this is indeed the case in the setting of nonparametric binary classification. Our results show that in the worst case, an algorithm cannot outperform undersampling unless there is a high degree of overlap between the train and test distributions (which is unlikely to be the case in real-world datasets), or if the algorithm leverages additional structure about the distribution shift. In particular, in the case of label shift we show that there is always an undersampling algorithm that is minimax optimal. In the case of group-covariate shift we show that there is an undersampling algorithm that is minimax optimal when the overlap between the group distributions is small. We also perform an experimental case study on a label shift dataset and find that in line with our theory, the test accuracy of robust neural network classifiers is constrained by the number of minority samples.

1 Introduction

A key challenge facing the machine learning community is to design models that are robust to distribution shift. When there is a mismatch between the train and test distributions, current models are often brittle and perform poorly on rare examples (Hovy & Søgaard, 2015; Blodgett et al., 2016; Tatman, 2017; Hashimoto et al., 2018; Alcorn et al., 2019). In this paper, our focus is on group-structured distribution shifts. In the training set, we have many samples from a *majority* group and relatively few samples from the *minority* group, while during test time we are equally likely to get a sample from either group.

To tackle such distribution shifts, a naïve algorithm is one that first *undersamples* the training data by discarding excess majority group samples (Kubat & Matwin, 1997; Wallace et al., 2011) and then trains a model on this resulting dataset (see Figure 1 for an illustration of this algorithm). The samples that remain in this undersampled dataset constitute i.i.d. draws from the test distribution. Therefore, while a classifier trained on this pruned dataset cannot suffer biases due to distribution shift, this algorithm is clearly wasteful, as it discards training samples. This perceived inefficiency of undersampling has led to the design of several algorithms to combat such distribution shift (Chawla et al., 2002; Lipton et al., 2018;




Figure 1: Example with linear models and linearly separable data. On the left we have the maximum margin classifier over the entire dataset, and on the right we have the maximum margin classifier over the undersampled dataset. The undersampled classifier is less biased and aligns more closely with the true boundary.

Sagawa et al., 2020; Cao et al., 2019; Menon et al., 2020; Ye et al., 2020; Kini et al., 2021; Wang et al., 2022). In spite of this algorithmic progress, the simple baseline of training models on an undersampled dataset remains competitive. In the case of label shift, where one class label is overrepresented in the training data, this has been observed by Cui et al. (2019); Cao et al. (2019), and Yang & Xu (2020). While in the case of group-covariate shift, a study by Idrissi et al. (2022) showed that the empirical effectiveness of these more complicated algorithms is limited.

For example, Idrissi et al. (2022) showed that on the group-covariate shift CelebA dataset the worst-group accuracy of a ResNet-50 model on the undersampled CelebA dataset which *discards 97%* of the available training data is as good as methods that use all of available data such as importance-weighted ERM (Shimodaira, 2000), Group-DRO (Sagawa et al., 2020) and Just-Train-Twice (Liu et al., 2021). In Table 1, we report the performance of the undersampled classifier compared to the state-of-the-art-methods in the literature across several label shift and group-covariate shift datasets. We find that, although undersampling isn’t always the optimal robustness algorithm, it is typically a very competitive baseline and within 1–4% the performance of the best method.

Table 1: Performance of undersampled classifier compared to the best classifier across several popular label shift and group-covariate shift datasets. When reporting worst-group accuracy we denote it by a *. When available, we report the 95% confidence interval. We find that the undersampled classifier is always within 1–4% of the best performing robustness algorithm, except on the CIFAR100 and MultiNLI datasets. In Appendix F we provide more details about each of the results in the table.

Shift Type	Dataset/Paper	Test/Worst-Group* Accuracy	
		Best	Undersampled
Label	Imb. CIFAR10 (step 10) (Cao et al., 2019)	87.81	84.59
	Imb. CIFAR100 (step 10) (Cao et al., 2019)	59.46	53.08
Group-Covariate	CelebA (Idrissi et al., 2022)	$86.9 \pm 1.1^*$	$85.6 \pm 2.3^*$
	Waterbirds (Idrissi et al., 2022)	$87.6 \pm 1.6^*$	$89.1 \pm 1.1^*$
	MultiNLI (Idrissi et al., 2022)	$78.0 \pm 0.7^*$	$68.9 \pm 0.8^*$
	CivilComments (Idrissi et al., 2022)	$72.0 \pm 1.9^*$	$71.8 \pm 1.4^*$

Inspired by the strong performance of undersampling in these experiments, we ask:

Is the performance of a model under distribution shift fundamentally constrained by the lack of minority group samples?

To answer this question we analyze the *minimax excess risk*. We lower bound the minimax excess risk to prove that the performance of *any* algorithm is lower bounded only as a function of the minority samples (n_{\min}). This shows that even if a robust algorithm optimally trades off between the bias and the variance, it is fundamentally constrained by the variance on the minority group which decreases only with n_{\min} .

Our contributions. In our paper, we consider the well-studied setting of nonparametric binary classification (Tsybakov, 2010). By operating in this nonparametric regime we are able to study the properties of undersampling in rich data distributions, but are able to circumvent the complications that arise due to the optimization and implicit bias of parametric models.

We provide insights into this question in the label shift scenario, where one of the labels is overrepresented in the training data, $P_{\text{train}}(y = 1) \geq P_{\text{train}}(y = -1)$, whereas the test samples are equally likely to come from either class. Here the class-conditional distribution $P(x | y)$ is Lipschitz in x . We show that in the label shift setting there is a fundamental constraint, and that the minimax excess risk of *any robust learning method* is lower bounded by $1/n_{\min}^{1/3}$. That is, minority group samples fundamentally constrain performance under distribution shift. Furthermore, by leveraging previous results about nonparametric density estimation (Freedman & Diaconis, 1981) we show a matching upper bound on the excess risk of a standard binning estimator trained on an undersampled dataset to demonstrate that undersampling is optimal.

Further, we experimentally show in a label shift dataset (Imbalanced Binary CIFAR10) that the accuracy of popular classifiers generally follow the trends predicted by our theory. When the minority samples are increased, the accuracy of these classifiers increases drastically, whereas when the number of majority samples are increased the gains in the accuracy are marginal at best.

We also study the covariate shift case. In this setting, there has been extensive work studying the effectiveness of transfer (Kpotufe & Martinet, 2018; Hanneke & Kpotufe, 2019) from train to test distributions, often focusing on deriving specific conditions under which this transfer is possible. In this work, we demonstrate that when the overlap (defined in terms of total variation distance) between the group distributions P_a and P_b is small, transfer is difficult, and that the minimax excess risk of any robust learning algorithm is lower bounded by $1/n_{\min}^{1/3}$. While this prior work also shows the impossibility of using majority group samples in the extreme case with no overlap, our results provide a simple lower bound that shows that the amount of overlap needed to make transfer feasible is unrealistic. We also show that this lower bound is tight, by proving an upper bound on the excess risk of the binning estimator acting on the undersampled dataset.

Taken together, our results underline the need to move beyond designing “general-purpose” robustness algorithms (like importance-weighting (Cao et al., 2019; Menon et al., 2020; Kini et al., 2021; Wang et al., 2022), g-DRO (Sagawa et al., 2020), JTT (Liu et al., 2021), SMOTE (Chawla et al., 2002), etc.) that are agnostic to the structure in the distribution shift. Our worst case analysis highlights that to successfully beat undersampling, an algorithm must leverage additional structure in the distribution shift.

2 Related work

On several group-covariate shift benchmarks (CelebA, CivilComments, Waterbirds), Idrissi et al. (2022) showed that training ResNet classifiers on an undersampled dataset either outperforms or performs as well as other popular reweighting methods like Group-DRO (Sagawa et al., 2020), reweighted ERM, and Just-Train-Twice (Liu et al., 2021). They find Group-DRO performs comparably to undersampling, while both tend to outperform methods that don’t utilize group information.

One classic method to tackle distribution shift is importance weighting (Shimodaira, 2000), which reweights the loss of the minority group samples to yield an unbiased estimate of the loss. However, recent work (Byrd & Lipton, 2019; Xu et al., 2020) has demonstrated the ineffectiveness of such methods when applied to overparameterized neural networks. Many followup papers (Cao et al., 2019; Ye et al., 2020; Menon et al.,

2020; Kini et al., 2021; Wang et al., 2022) have introduced methods that modify the loss function in various ways to address this. However, despite this progress undersampling remains a competitive alternative to these importance weighted classifiers.

Our theory draws from the rich literature on non-parametric classification (Tsybakov, 2010). Apart from borrowing this setting of nonparametric classification, we also utilize upper bounds on the estimation error of the simple histogram estimator (Freedman & Diaconis, 1981; Devroye & Györfi, 1985) to prove our upper bounds in the label shift case. Finally, we note that to prove our minimax lower bounds we proceed by using the general recipe of reducing from estimation to testing (Wainwright, 2019, Chapter 15). One difference from this standard framework is that our training samples shall be drawn from a different distribution than the test samples used to define the risk.

Past work has established lower bounds on the minimax risk for binary classification without distribution shift for general VC classes (see, e.g., Massart & Nédélec, 2006). Note that, these bounds are not directly applicable in the distribution shift setting, and consequently these lower bounds scale with the total number of samples $n = n_{\text{maj}} + n_{\text{min}}$ rather than with the minority number of samples (n_{min}). There are also refinements of this lower bound to obtain minimax lower bounds for cost-sensitive losses that penalize errors on the two class classes differently (Kamalaruban & Williamson, 2018). By carefully selecting these costs it is possible to apply these results in the label shift setting. However, these lower bounds remain loose and decay with n and n_{maj} in contrast to the tighter n_{min} dependence in our lower bounds. We provide a more detailed discussion about potentially applying these lower bounds to the label shift setting after the presentation of our theorem in Section 4.1.

There is rich literature that studies domain adaptation and transfer learning under label shift (Maity et al., 2020) and covariate shift (Ben-David et al., 2006; David et al., 2010; Ben-David et al., 2010; Ben-David & Urner, 2012; 2014; Berlind & Urner, 2015; Kpotufe & Martinet, 2018; Hanneke & Kpotufe, 2019). The principal focus of this line of work was to understand the value of unlabeled data from the target domain, rather than to characterize the relative value of the number of labeled samples from the majority and minority groups. Among these papers, most closely related to our work are those in the covariate shift setting (Kpotufe & Martinet, 2018; Hanneke & Kpotufe, 2019). Their lower bound results can be reinterpreted to show that under covariate shift in the absence of overlap, the minimax excess risk is lower bounded by $1/n_{\text{min}}^{1/3}$. We provide a more detailed comparison with their results after presenting our lower bounds in Section 4.2.

Finally, we note that Arjovsky et al. (2022) recently showed that undersampling can improve the worst-class accuracy of linear SVMs in the presence of label shift. In comparison, our results hold for arbitrary classifiers with the rich nonparametric data distributions.

3 Setting

In this section, we shall introduce our problem setup and define the types of distribution shift that we consider.

3.1 Problem setup

The setting for our study is nonparametric binary classification with Lipschitz data distributions. We are given n training datapoints $\mathcal{S} := \{(x_1, y_1), \dots, (x_n, y_n)\} \in ([0, 1] \times \{-1, 1\})^n$ that are all drawn from a *train* distribution $\mathbb{P}_{\text{train}}$. During test time, the data shall be drawn from a *different* distribution \mathbb{P}_{test} . Our paper focuses on the robustness to this shift in the distribution from train to test time. To present a clean analysis, we study the case where the features x are bounded scalars, however, it is easy to extend our results to the high-dimensional setting.

Given a classifier $f : [0, 1] \rightarrow \{-1, 1\}$, we shall be interested in the test error (risk) of this classifier under the test distribution \mathbb{P}_{test} :

$$R(f; \mathbb{P}_{\text{test}}) := \mathbb{E}_{(x,y) \sim \mathbb{P}_{\text{test}}} [\mathbf{1}(f(x) \neq y)].$$

3.2 Types of distribution shift

We assume that P_{train} consists of a mixture of two groups of unequal size, and P_{test} contains equal numbers of samples from both groups. Given a majority group distribution P_{maj} and a minority group distribution P_{min} , the learner has access to n_{maj} majority group samples and n_{min} minority group samples:

$$\mathcal{S}_{\text{maj}} \sim P_{\text{maj}}^{n_{\text{maj}}} \quad \text{and} \quad \mathcal{S}_{\text{min}} \sim P_{\text{min}}^{n_{\text{min}}}.$$

Here $n_{\text{maj}} > n/2$ and $n_{\text{min}} < n/2$ with $n_{\text{maj}} + n_{\text{min}} = n$. The full training dataset is $\mathcal{S} = \mathcal{S}_{\text{maj}} \cup \mathcal{S}_{\text{min}} = \{(x_1, y_1), \dots, (x_n, y_n)\}$. We assume that the learner has access to the knowledge whether a particular sample (x_i, y_i) comes from the majority or minority group.

The test samples will be drawn from $P_{\text{test}} = \frac{1}{2}P_{\text{maj}} + \frac{1}{2}P_{\text{min}}$, a uniform mixture over P_{maj} and P_{min} . Thus, the training dataset is an imbalanced draw from the distributions P_{maj} and P_{min} , whereas the test samples are balanced draws. We let $\rho := n_{\text{maj}}/n_{\text{min}} > 1$ denote the imbalance ratio in the training data. We consider the uniform mixture during test time since the resulting test loss is of the same order as the worst-group loss.

We focus on two-types of distribution shifts: label shift and group-covariate shift that we describe below.

3.2.1 Label shift

In this setting, the imbalance in the training data comes from there being more samples from one class over another. Without loss of generality, we shall assume that the class $y = 1$ is the majority class. Then, we define the majority and the minority class distributions as

$$P_{\text{maj}}(x, y) = P_1(x)\mathbf{1}(y = 1) \quad \text{and} \quad P_{\text{min}} = P_{-1}(x)\mathbf{1}(y = -1),$$

where P_1, P_{-1} are class-conditional distributions over the interval $[0, 1]$. We assume that class-conditional distributions P_i have densities on $[0, 1]$ and that they are 1-Lipschitz: for any $x, x' \in [0, 1]$,

$$|P_i(x) - P_i(x')| \leq |x - x'|.$$

We denote the class of pairs of distributions $(P_{\text{maj}}, P_{\text{min}})$ that satisfy these conditions by \mathcal{P}_{LS} . We note that such Lipschitzness assumptions are common in the literature (see Tsybakov, 2010).

3.2.2 Group-covariate shift

In this setting, we have two groups $\{a, b\}$, and corresponding to each of these groups is a distribution (with densities) over the features $P_a(x)$ and $P_b(x)$. We let a correspond to the majority group and b correspond to the minority group. Then, we define

$$P_{\text{maj}}(x, y) = P_a(x)P(y | x) \quad \text{and} \quad P_{\text{min}}(x, y) = P_b(x)P(y | x).$$

We assume that for $y \in \{-1, 1\}$, for all $x, x' \in [0, 1]$:

$$|P(y | x) - P(y | x')| \leq |x - x'|,$$

that is, the distribution of the label given the feature is 1-Lipschitz, and it varies slowly over the domain.

To quantify the shift between the train and test distribution, we define a notion of overlap between the group distributions P_a and P_b as follows:

$$\text{Overlap}(P_a, P_b) := 1 - \text{TV}(P_a, P_b)$$

where $\text{TV}(P_a, P_b) := \sup_{E \subseteq [0, 1]} |P_a(E) - P_b(E)|$, denotes the total variation distance between P_a and P_b . Notice that when P_a and P_b have disjoint supports, $\text{TV}(P_a, P_b) = 1$ and therefore $\text{Overlap}(P_a, P_b) = 0$. On the other hand when $P_a = P_b$, $\text{TV}(P_a, P_b) = 0$ and $\text{Overlap}(P_a, P_b) = 1$. When the overlap is 1, the majority and minority distributions are identical and hence we have no shift between train and test. Observe that $\text{Overlap}(P_a, P_b) = \text{Overlap}(P_{\text{maj}}, P_{\text{min}})$ since $P(y | x)$ is shared across P_{maj} and P_{min} .

Given a level of overlap $\tau \in [0, 1]$ we denote the class of pairs of distributions $(P_{\text{maj}}, P_{\text{min}})$ with overlap at least τ by $\mathcal{P}_{\text{GS}}(\tau)$. It is easy to check that, $\mathcal{P}_{\text{GS}}(\tau) \subseteq \mathcal{P}_{\text{GS}}(0)$ at any overlap level $\tau \in [0, 1]$.

Considering a notion of overlap between the marginal distributions $P_a(x)$ and $P_b(x)$ is natural in the group covariate setting since the conditional distribution that we wish to estimate $P(y | x)$ remains constant from train to test time. Higher overlap between P_a and P_b allows a classifier to learn more about the underlying conditional distribution $P(y | x)$ when it sees samples from either group. In contrast, in the label shift setting $P(x | y)$ remains constant from train to test time and higher overlap between $P(x | 1)$ and $P(x | -1)$ does not help to estimate $P(y | x)$.

4 Lower bounds on the minimax excess risk

In this section, we shall prove our lower bounds that show that the performance of any algorithm is constrained by the number of minority samples n_{min} . Before we state our lower bounds, we need to introduce the notion of excess risk and minimax excess risk.

Excess risk and minimax excess risk. We measure the performance of an algorithm \mathcal{A} through its excess risk defined in the following way. Given an algorithm \mathcal{A} that takes as input a dataset \mathcal{S} and returns a classifier $\mathcal{A}^{\mathcal{S}}$, and a pair of distributions $(P_{\text{maj}}, P_{\text{min}})$ with $P_{\text{test}} = \frac{1}{2}P_{\text{maj}} + \frac{1}{2}P_{\text{min}}$, the *expected excess risk* is given by

$$\text{Excess Risk}[\mathcal{A}; (P_{\text{maj}}, P_{\text{min}})] := \mathbb{E}_{\mathcal{S} \sim P_{\text{maj}}^{n_{\text{maj}}} \times P_{\text{min}}^{n_{\text{min}}}} [R(\mathcal{A}^{\mathcal{S}}; P_{\text{test}}) - R(f^*(P_{\text{test}}); P_{\text{test}})], \quad (1)$$

where $f^*(P_{\text{test}})$ is the Bayes classifier that minimizes the risk $R(\cdot; P_{\text{test}})$. The first term corresponds to the expected risk for the algorithm when given n_{maj} samples from P_{maj} and n_{min} samples from P_{min} , whereas the second term corresponds to the Bayes error for the problem.

Excess risk does not let us characterize the inherent difficulty of a problem, since for any particular data distribution $(P_{\text{maj}}, P_{\text{min}})$ the best possible algorithm \mathcal{A} to minimize the excess risk would be the trivial mapping $\mathcal{A}^{\mathcal{S}} = f^*(P_{\text{test}})$. Therefore, to prove meaningful lower bounds on the performance of algorithms we need to define the notion of minimax excess risk (see Wainwright, 2019, Chapter 15). Given a class of pairs of distributions \mathcal{P} define

$$\text{Minimax Excess Risk}(\mathcal{P}) := \inf_{\mathcal{A}} \sup_{(P_{\text{maj}}, P_{\text{min}}) \in \mathcal{P}} \text{Excess Risk}[\mathcal{A}; (P_{\text{maj}}, P_{\text{min}})], \quad (2)$$

where the infimum is over all measurable estimators \mathcal{A} . The minimax excess risk is the excess risk of the “best” algorithm in the worst case over the class of problems defined by \mathcal{P} .

4.1 Label shift lower bounds

We demonstrate the hardness of the label shift problem in general by establishing a lower bound on the minimax excess risk.

Theorem 4.1. *Consider the label shift setting described in Section 3.2.1. Recall that \mathcal{P}_{LS} is the class of pairs of distributions $(P_{\text{maj}}, P_{\text{min}})$ that satisfy the assumptions in that section. The minimax excess risk over this class is lower bounded as follows:*

$$\text{Minimax Excess Risk}(\mathcal{P}_{\text{LS}}) = \inf_{\mathcal{A}} \sup_{(P_{\text{maj}}, P_{\text{min}}) \in \mathcal{P}_{\text{LS}}} \text{Excess Risk}[\mathcal{A}; (P_{\text{maj}}, P_{\text{min}})] \geq \frac{1}{600} \frac{1}{n_{\text{min}}^{1/3}}. \quad (3)$$

We establish this result in Appendix B. We show that rather surprisingly, the lower bound on the minimax excess risk scales only with the number of minority class samples $n_{\text{min}}^{1/3}$, and does not depend on n_{maj} . Intuitively, this is because any learner must predict which class-conditional distribution ($P(x | 1)$ or $P(x | -1)$) assigns higher likelihood at that x . To interpret this result, consider the extreme scenario where $n_{\text{maj}} \rightarrow \infty$ but n_{min} is finite. In this case, the learner has full information about the majority class distribution. However,

the learning task continues to be challenging since any learner would be uncertain about whether the minority class distribution assigns higher or lower likelihood at any given x . This uncertainty underlies the reason why the minimax rate of classification is constrained by the number of minority samples n_{\min} .

We briefly note that, applying minimax lower bounds from the transfer learning literature (Maity et al., 2020, Theorem 3.1 with $\alpha = 1$, $\beta = 0$ and $d = 1$) to our problem leads to a more optimistic lower bound of $1/n^{1/3}$. Our lower bounds that scale as $1/n_{\min}^{1/3}$, uncover the fact that only adding minority class samples helps reduce the risk.

As noted above in the introduction, it is possible to obtain lower bounds for the label shift setting by applying bounds from the cost-sensitive classification literature. However, as we shall argue below they are loose and predict the incorrect trend when applied in this setting. Consider the result (Kamalaruban & Williamson, 2018, Theorem 4) which is a minimax lower bound for cost sensitive binary classification that applies to VC classes (which does not capture the nonparameteric setting studied here but it is illuminating to study how that bound scales with the imbalance ratio $\rho = n_{\text{maj}}/n_{\text{min}}$). Assume that the joint distribution during training is a mixture distribution given by $\mathbf{P} = \frac{\rho}{1+\rho}\mathbf{P}_{\text{maj}} + \frac{1}{1+\rho}\mathbf{P}_{\text{min}}$ so that on average the ratio of the number of samples from the majority and minority class is equal to ρ . Then by applying their lower bound we find that it scales with $1/(n\rho)$ (see Appendix E for a detailed calculation). This scales inversely with ρ the imbalance ratio and incorrectly predicts that the problem gets easier as the imbalance is larger. In contrast, our lower bound scales with $1/n_{\min} = (1 + \rho)/n$, which correctly predicts that as the imbalance is larger, the minimax test error is higher.

4.2 Group-covariate shift lower bounds

Next, we shall state our lower bound on the minimax excess risk that demonstrates the hardness of the group-covariate shift problem.

Theorem 4.2. *Consider the group shift setting described in Section 3.2.2. Given any overlap $\tau \in [0, 1]$ recall that $\mathcal{P}_{\text{GS}}(\tau)$ is the class of distributions such that $\text{Overlap}(\mathbf{P}_{\text{maj}}, \mathbf{P}_{\text{min}}) \geq \tau$. The minimax excess risk in this setting is lower bounded as follows:*

$$\begin{aligned} \text{Minimax Excess Risk}(\mathcal{P}_{\text{GS}}(\tau)) &= \inf_{\mathcal{A}} \sup_{(\mathbf{P}_{\text{maj}}, \mathbf{P}_{\text{min}}) \in \mathcal{P}_{\text{GS}}(\tau)} \text{Excess Risk}[\mathcal{A}; (\mathbf{P}_{\text{maj}}, \mathbf{P}_{\text{min}})] \\ &\geq \frac{1}{200(n_{\min} \cdot (2 - \tau) + n_{\text{maj}} \cdot \tau)^{1/3}} \geq \frac{1}{200n_{\min}^{1/3}(\rho \cdot \tau + 2)^{1/3}}, \end{aligned} \quad (4)$$

where $\rho = n_{\text{maj}}/n_{\min} > 1$.

We prove this theorem in Appendix C.

We see that in the *low overlap* setting ($\tau \ll 1/\rho$), the minimax excess risk is lower bounded by $1/n_{\min}^{1/3}$, and we are fundamentally constrained by the number of samples in minority group. To see why this is the case, consider the extreme example with $\tau = 0$ where \mathbf{P}_a has support $[0, 0.5]$ and \mathbf{P}_b has support $[0.5, 1]$. The n_{maj} majority group samples from \mathbf{P}_a provide information about the correct label predict in the interval $[0, 0.5]$ (the support of \mathbf{P}_a). However, since the distribution $\mathbf{P}(y | x)$ is 1-Lipschitz in the worst case these samples provide very limited information about the correct predictions in $[0.5, 1]$ (the support of \mathbf{P}_b). Thus, predicting on the support of \mathbf{P}_b requires samples from the minority group and this results in the n_{\min} dependent rate. In fact, in this extreme case ($\tau = 0$) even if $n_{\text{maj}} \rightarrow \infty$, the minimax excess risk is still bounded away from zero. This intuition also carries over to the case when the overlap is small but non-zero and our lower bound shows that minority samples are much more valuable than majority samples at reducing the risk.

On the other hand, when the overlap is high ($\tau \gg 1/\rho$) the minimax excess risk is lower bounded by $1/(n_{\min}(2 - \tau) + n_{\text{maj}}\tau)^{1/3}$ and the extra majority samples are quite beneficial. This is roughly because the supports of \mathbf{P}_a and \mathbf{P}_b have large overlap and hence samples from the majority group are useful in helping make predictions even in regions where \mathbf{P}_b is large. In the extreme case when $\tau = 1$, we have that $\mathbf{P}_a = \mathbf{P}_b$ and therefore recover the classic i.i.d. setting with no distribution shift. Here, the lower bound scales with $1/n^{1/3}$, as one might expect.

Previous work on transfer learning with covariate shift has considered other more elaborate notions of *transferability* (Kpotufe & Martinet, 2018; Hanneke & Kpotufe, 2019) than overlap between group distributions considered here. In the case of no overlap ($\tau = 0$), previous results (Kpotufe & Martinet, 2018, Theorem 1 with $\alpha = 1, \beta = 0$ and $\gamma = \infty$) yield the same lower bound of $1/n_{\min}^{1/3}$. On the other extreme, applying their result (Kpotufe & Martinet, 2018, Theorem 1 with $\alpha = 1, \beta = 0$ and $\gamma = 0$) in the high transfer regime yields a lower bound on $1/n^{1/3}$. This result is aligned with the high overlap $\tau = 1$ case that we consider here.

Beyond these two edge cases of no overlap ($\tau = 0$) and high overlap ($\tau = 1$), our lower bound is key to drawing the simple complementary conclusion that even when overlap between group distributions is small as compared to $1/\rho$, minority samples alone dictate the rate of convergence.

5 Upper bounds on the excess risk for the undersampled binning estimator

We will show that an undersampled estimator matches the rates in the previous section showing that undersampling is an optimal robustness intervention. We start by defining the undersampling procedure and the undersampling binning estimator.

Undersampling procedure. Given training data $\mathcal{S} := \{(x_1, y_1), \dots, (x_n, y_n)\}$, generate a new undersampled dataset \mathcal{S}_{US} by

- including all n_{\min} samples from \mathcal{S}_{\min} and,
- including n_{\min} samples from \mathcal{S}_{maj} by sampling uniformly at random without replacement.

This procedure ensures that in the undersampled dataset \mathcal{S}_{US} , the groups are balanced, and that $|\mathcal{S}_{\text{US}}| = 2n_{\min}$.

The undersampling binning estimator defined next will first run this undersampling procedure to obtain \mathcal{S}_{US} and just uses these samples to output a classifier.

Undersampled binning estimator The undersampled binning estimator \mathcal{A}_{USB} takes as input a dataset \mathcal{S} and a positive integer K corresponding to the number of bins, and returns a classifier $\mathcal{A}_{\text{USB}}^{\mathcal{S}, K} : [0, 1] \rightarrow \{-1, 1\}$. This estimator is defined as follows:

1. First, we compute the undersampled dataset \mathcal{S}_{US} .
2. Given this dataset \mathcal{S}_{US} , let $n_{1,j}$ be the number of points with label +1 that lie in the interval $I_j = [\frac{j-1}{K}, \frac{j}{K}]$. Also, define $n_{-1,j}$ analogously. Then set

$$\mathcal{A}_j = \begin{cases} 1 & \text{if } n_{1,j} > n_{-1,j}, \\ -1 & \text{otherwise.} \end{cases}$$

3. Define the classifier $\mathcal{A}_{\text{USB}}^{\mathcal{S}, K}$ such that if $x \in I_j$ then

$$\mathcal{A}_{\text{USB}}^{\mathcal{S}, K}(x) = \mathcal{A}_j. \quad (5)$$

Essentially in each bin I_j , we set the prediction to be the majority label among the samples that fall in this bin.

Whenever the number of bins K is clear from the context we shall denote $\mathcal{A}_{\text{USB}}^{\mathcal{S}, K}$ by $\mathcal{A}_{\text{USB}}^{\mathcal{S}}$. Below we establish upper bounds on the excess risk of this simple estimator.

5.1 Label shift upper bounds

We now establish an upper bound on the excess risk of \mathcal{A}_{USB} in the label shift setting (see Section 3.2.1). Below we let $c, C > 0$ be absolute constants independent of problem parameters like n_{maj} and n_{min} .

Theorem 5.1. *Consider the label shift setting described in Section 3.2.1. For any $(\mathbb{P}_{\text{maj}}, \mathbb{P}_{\text{min}}) \in \mathcal{P}_{\text{LS}}$ the expected excess risk of the Undersampling Binning Estimator (Eq. (5)) with number of bins with $K = \lceil n_{\text{min}}^{1/3} \rceil$ is upper bounded by*

$$\text{Excess Risk}[\mathcal{A}_{\text{USB}}; (\mathbb{P}_{\text{maj}}, \mathbb{P}_{\text{min}})] = \mathbb{E}_{\mathcal{S} \sim \mathbb{P}_{\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\text{min}}^{n_{\text{min}}}} [R(\mathcal{A}_{\text{USB}}^{\mathcal{S}}; \mathbb{P}_{\text{test}}) - R(f^*; \mathbb{P}_{\text{test}})] \leq \frac{C}{n_{\text{min}}^{1/3}}.$$

We prove this result in Appendix B. This upper bound combined with the lower bound in Theorem 4.1 shows that an undersampling approach is minimax optimal up to constants in the presence of label shift.

Our analysis leaves open the possibility of better algorithms when the learner has additional information about the structure of the label shift beyond Lipschitz continuity.

5.2 Group-covariate shift upper bounds

Next, we present our upper bounds on the excess risk of the undersampled binning estimator in the group-covariate shift setting (see Section 3.2.2). In the theorem below, $C > 0$ is an absolute constant independent of the problem parameters n_{maj} , n_{min} and τ .

Theorem 5.2. *Consider the group shift setting described in Section 3.2.2. For any overlap $\tau \in [0, 1]$ and for any $(\mathbb{P}_{\text{maj}}, \mathbb{P}_{\text{min}}) \in \mathcal{P}_{\text{GS}}(\tau)$ the expected excess risk of the Undersampling Binning Estimator (Eq. (5)) with number of bins with $K = \lceil n_{\text{min}}^{1/3} \rceil$ is*

$$\text{Excess Risk}[\mathcal{A}_{\text{USB}}; (\mathbb{P}_{\text{maj}}, \mathbb{P}_{\text{min}})] = \mathbb{E}_{\mathcal{S} \sim \mathbb{P}_{\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\text{min}}^{n_{\text{min}}}} [R(\mathcal{A}_{\text{USB}}^{\mathcal{S}}; \mathbb{P}_{\text{test}}) - R(f^*; \mathbb{P}_{\text{test}})] \leq \frac{C}{n_{\text{min}}^{1/3}}.$$

We provide a proof for this theorem in Appendix C. Compared to the lower bound established in Theorem 4.2 which scales as $1/((2 - \tau)n_{\text{min}} + n_{\text{maj}}\tau)^{1/3}$, the upper bound for the undersampled binning estimator always scales with $1/n_{\text{min}}^{1/3}$ since it operates on the undersampled dataset (\mathcal{S}_{US}).

Thus, we have shown that in the absence of overlap ($\tau \ll 1/\rho = n_{\text{min}}/n_{\text{maj}}$) there is an undersampling algorithm that is minimax optimal up to constants. However when there is high overlap ($\tau \gg 1/\rho$) there is a non-trivial gap between the upper and lower bounds:

$$\frac{\text{Upper Bound}}{\text{Lower Bound}} = c(\rho \cdot \tau + 2)^{1/3}.$$

6 Minority sample dependence in practice

Inspired by our worst-case theoretical predictions in nonparametric classification, we ask: how does the accuracy of neural network classifiers trained using robust algorithms evolve as a function of the majority and minority samples?

To explore this question, we conduct a small case study using the imbalanced binary CIFAR10 dataset (Byrd & Lipton, 2019; Wang et al., 2022) that is constructed using the ‘‘cat’’ and ‘‘dog’’ classes. The test set consists of all of the 1000 cat and 1000 dog test examples. To form our initial train and validation sets, we take 2500 cat examples but only 500 dog examples from the official train set, corresponding to a 5:1 label imbalance. We then use 80% of those examples for training and the rest for validation. In our experiment, we either (a) add only minority samples; (b) add only majority samples; (c) add both majority and minority samples in a 5:1 ratio. We consider competitive robust classifiers proposed in the literature that are convolutional neural networks trained either by using (i) the importance weighted cross entropy loss, or (ii) the importance weighted VS loss (Kini et al., 2021). We early stop using the importance weighted validation loss in both cases. The additional experimental details are presented in Appendix G.

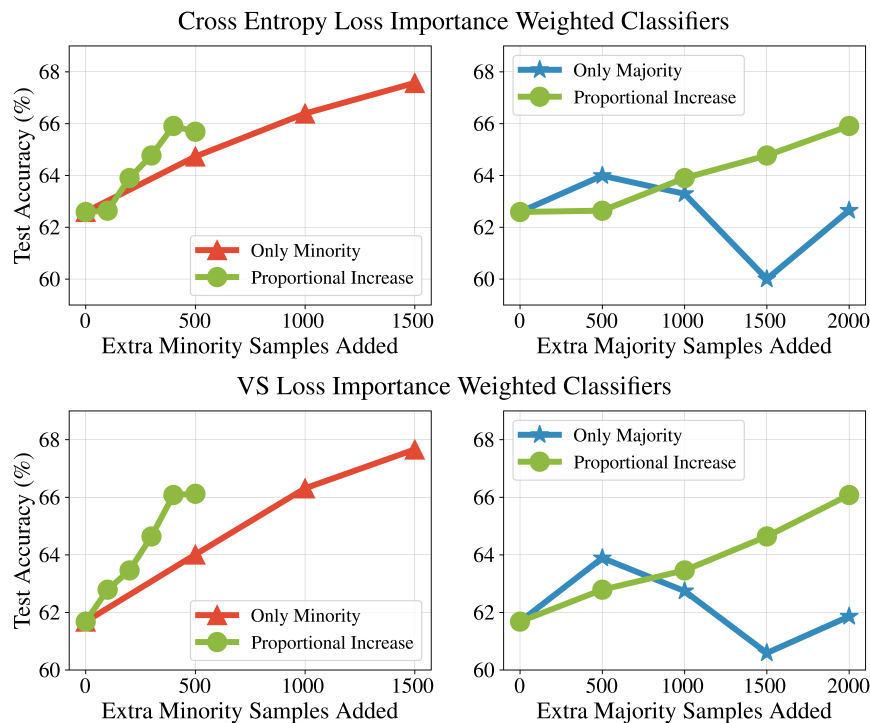


Figure 2: Convolutional neural network classifiers trained on the Imbalanced Binary CIFAR10 dataset with a 5:1 label imbalance. (Top) Models trained using the importance weighted cross entropy loss with early stopping. (Bottom) Models trained using the importance weighted VS loss (Kini et al., 2021) with early stopping. We report the average test accuracy calculated on a balanced test set over 5 random seeds. We start off with 2500 cat examples and 500 dog examples in the training dataset. We find that in accordance with our theory, for both of the classifiers adding only minority class samples (red) leads to large gain in accuracy ($\sim 6\%$), while adding majority class samples (blue) leads to little or no gain. In fact, adding majority samples sometimes hurts test accuracy due to the added bias. When we add majority and minority samples in a 5:1 ratio (green), the gain is largely due to the addition of minority samples and is only marginally higher ($< 2\%$) than adding only minority samples. The green curves correspond to the same classifiers in both the left and right panels.

Our results in Figure 2 are generally consistent with our theoretical predictions. By adding only minority class samples the test accuracy of both classifiers increases by a great extent (6%), while by adding only majority class samples the test accuracy remains constant or in some cases even decreases owing to the added bias of the classifiers. When we add samples to both groups proportionately, the increase in the test accuracy appears to largely be due to the increase in the number of minority class samples. We see this on the left panels, where the difference between adding only extra minority group samples (red) and both minority and majority group samples (green) is small. Thus, we find that the accuracy for these neural network classifiers is also constrained by the number of minority class samples. Similar conclusions hold for classifiers trained using the tilted loss (Li et al., 2020) and group-DRO objective (Sagawa et al., 2020) (see Appendix D).

7 Discussion

We showed that undersampling is an optimal robustness intervention in nonparametric classification in the absence of significant overlap between group distributions or without additional structure beyond Lipschitz continuity. We worked in one dimension for the sake of clarity and it would be interesting to extend this study to higher dimensions. We focused on Lipschitz continuous distributions here, but it is also interesting to consider other forms of regularity such as Hölder continuity.

At a high level our results highlight the need to reason about the specific structure in the distribution shift and design algorithms that are tailored to take advantage of this structure. This would require us to step away from the common practice in robust machine learning where the focus is to design “universal” robustness interventions that are agnostic to the structure in the shift. Alongside this, our results also dictate the need for datasets and benchmarks with the propensity for transfer from train to test time.

Acknowledgments

We would like to thank Ke Alexander Wang for his useful comments and feedback in the early stages of this project. We would also like to thank Shibani Santurkar and Dimitrios Tsipras for useful discussions and encouragement. Finally, we would like to thank the anonymous reviewers whose many helpful comments improved the paper. NC was supported by a SAIL Postdoctoral Fellowship and TH was supported by a gift from Open Philanthropy.

References

- Michael Alcorn, Qi Li, Zhitao Gong, Chengfei Wang, Long Mai, Wei-Shinn Ku, and Anh Nguyen. Strike (with) a pose: Neural networks are easily fooled by strange poses of familiar objects. In *Computer Vision and Pattern Recognition (CVPR)*, 2019.
- Martin Arjovsky, Kamalika Chaudhuri, and David Lopez-Paz. Throwing away data improves worst-class error in imbalanced classification. *arXiv preprint arXiv:2205.11672*, 2022.
- Shai Ben-David and Ruth Urner. On the hardness of domain adaptation and the utility of unlabeled target samples. In *Algorithmic Learning Theory (ALT)*, 2012.
- Shai Ben-David and Ruth Urner. Domain adaptation—can quantity compensate for quality? *Annals of Mathematics and Artificial Intelligence*, 2014.
- Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2006.
- Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 2010.
- Christopher Berling and Ruth Urner. Active nearest neighbors in changing environments. In *International Conference on Machine Learning (ICML)*, 2015.
- Su Lin Blodgett, Lisa Green, and Brendan O’Connor. Demographic dialectal variation in social media: A case study of african-american english. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2016.
- Jonathon Byrd and Zachary Lipton. What is the effect of importance weighting in deep learning? In *International Conference on Machine Learning (ICML)*, 2019.
- Clément Canonne. A short note on an inequality between KL and TV. *arXiv preprint arXiv:2202.07198*, 2022.
- Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Nitesh Chawla, Kevin Bowyer, Lawrence Hall, and Philip Kegelmeyer. Smote: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 2002.
- Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Computer Vision and Pattern Recognition (CVPR)*, 2019.
- Shai Ben David, Tyler Lu, Teresa Luu, and Dávid Pál. Impossibility theorems for domain adaptation. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2010.

- Luc Devroye and László Györfi. *Nonparametric density estimation: the L_1 view*. Wiley Series in Probability and Mathematical Statistics, 1985.
- David Freedman and Persi Diaconis. On the histogram as a density estimator: L_2 theory. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 1981.
- Steve Hanneke and Samory Kpotufe. On the value of target data in transfer learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Tatsunori Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. Fairness without demographics in repeated loss minimization. In *International Conference on Machine Learning (ICML)*, 2018.
- Dirk Hovy and Anders Søgaard. Tagging performance correlates with author age. In *Association for Computational Linguistics (ACL)*, 2015.
- Badr Youbi Idrissi, Martín Arjovsky, Mohammad Pezeshki, and David Lopez-Paz. Simple data balancing achieves competitive worst-group-accuracy. In *Causal Learning and Reasoning*, 2022.
- Parameswaran Kamalaruban and Robert Williamson. Minimax lower bounds for cost sensitive classification. *arXiv preprint arXiv:1805.07723*, 2018.
- Ganesh Ramachandra Kini, Orestis Paraskevas, Samet Oymak, and Christos Thrampoulidis. Label-imbalanced and group-sensitive classification under overparameterization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Samory Kpotufe and Guillaume Martinet. Marginal singularity, and the benefits of labels in covariate-shift. In *Conference On Learning Theory (COLT)*, 2018.
- Miroslav Kubat and Stan Matwin. Addressing the curse of imbalanced training sets: one-sided selection. In *International Conference on Machine Learning (ICML)*, 1997.
- Tian Li, Ahmad Beirami, Maziar Sanjabi, and Virginia Smith. Tilted empirical risk minimization. In *International Conference on Learning Representations (ICLR)*, 2020.
- Zachary Lipton, Yu-Xiang Wang, and Alexander Smola. Detecting and correcting for label shift with black box predictors. In *International Conference on Machine Learning (ICML)*, 2018.
- Evan Liu, Behzad Haghgoo, Annie Chen, Aditi Raghunathan, Pang Wei Koh, Shiori Sagawa, Percy Liang, and Chelsea Finn. Just train twice: Improving group robustness without training group information. In *International Conference on Machine Learning (ICML)*, 2021.
- Subha Maity, Yuekai Sun, and Moulinath Banerjee. Minimax optimal approaches to the label shift problem. *arXiv preprint arXiv:2003.10443*, 2020.
- Pascal Massart and Élodie Nédélec. Risk bounds for statistical learning. *Annals of Statistics*, 2006.
- Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. Long-tail learning via logit adjustment. In *International Conference on Learning Representations (ICLR)*, 2020.
- Shiori Sagawa, Pang Wei Koh, Tatsunori Hashimoto, and Percy Liang. Distributionally robust neural networks. In *International Conference on Learning Representations (ICLR)*, 2020.
- Hidetoshi Shimodaira. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of Statistical Planning and Inference*, 2000.
- Rachael Tatman. Gender and dialect bias in youtube’s automatic captions. In *ACL Workshop on Ethics in Natural Language Processing*, 2017.
- Alexandre Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 2010.

Martin Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*. Cambridge University Press, 2019.

Byron Wallace, Kevin Small, Carla Brodley, and Thomas Trikalinos. Class imbalance, redux. In *International Conference on Data Mining (ICDM)*, 2011.

Ke Alexander Wang, Niladri Chatterji, Saminul Haque, and Tatsunori Hashimoto. Is importance weighting incompatible with interpolating classifiers? In *International Conference on Learning Representations (ICLR)*, 2022.

Larry Wasserman. Lecture notes in nonparametric classification, 2019. URL <https://www.stat.cmu.edu/~larry/=sml/nonparclass.pdf>. [Online; accessed 12-May-2022].

Wikipedia contributors. Poisson binomial distribution — Wikipedia, the free encyclopedia, 2022. URL https://en.wikipedia.org/w/index.php?title=Poisson_binomial_distribution&oldid=1071847908. [Online; accessed 5-May-2022].

Da Xu, Yuting Ye, and Chuanwei Ruan. Understanding the role of importance weighting for deep learning. In *International Conference on Learning Representations (ICLR)*, 2020.

Yuzhe Yang and Zhi Xu. Rethinking the value of labels for improving class-imbalanced learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

Han-Jia Ye, Hong-You Chen, De-Chuan Zhan, and Wei-Lun Chao. Identifying and compensating for feature deviation in imbalanced deep learning. *arXiv preprint arXiv:2001.01385*, 2020.

A Technical tools

In this section we avail ourselves of some technical tools that shall be used in all of the proofs below.

A.1 Reduction to lower bounds over a finite class

The lower bound on the minimax excess risk will be established via the usual route of first identifying a “hard” finite set of problem instances and then establishing the lower bound over this finite class. One difference from the usual setup in proving such lower bounds (see Wainwright, 2019, Chapter 15) is that the training samples are drawn from an imbalanced distribution, whereas the test samples are drawn from a balanced one.

Let \mathcal{P} be a class of pairs of distributions, where each element $(P_{\text{maj}}, P_{\text{min}}) \in \mathcal{P}$ is a pair of distributions over $[0, 1] \times \{-1, 1\}$. As before, we let P_{test} denote the uniform mixture over P_{maj} and P_{min} . We let \mathcal{V} denote a finite index set. Corresponding to each element $v \in \mathcal{V}$ there is a $P_v = (P_{v,\text{maj}}, P_{v,\text{min}}) \in \mathcal{P}$ with $P_{v,\text{test}} = (P_{v,\text{maj}} + P_{v,\text{min}})/2$. Finally, also define a pair of random variables (V, S) as follows:

1. V is a uniform random variable over the set \mathcal{V} .
2. $(S \mid V = v) \sim P_{v,\text{maj}}^{n_{\text{maj}}} \times P_{v,\text{min}}^{n_{\text{min}}}$, is an independent draw of n_{maj} samples from $P_{v,\text{maj}}$ and n_{min} samples from $P_{v,\text{min}}$.

We shall let Q denote the joint distribution of the random variables (V, S) , and let Q_S denote the marginal distribution of S .

With this notation in place, we now present a lemma that lower bounds the minimax excess risk in terms of quantities defined over the finite class of “hard” instances P_v .

Lemma A.1. *Let the random variables (V, S) be as defined above. The minimax excess risk is lower bounded as follows:*

$$\begin{aligned} \text{Minimax Excess Risk}(\mathcal{P}) &= \inf_{\mathcal{A}} \sup_{(P_{\text{maj}}, P_{\text{min}}) \in \mathcal{P}} \mathbb{E}_{S \sim P_{\text{maj}}^{n_{\text{maj}}} \times P_{\text{min}}^{n_{\text{min}}}} [R(\mathcal{A}^S; P_{\text{test}}) - R(f^*(P_{\text{test}}); P_{\text{test}})] \\ &\geq \mathfrak{R}_{\mathcal{V}} - \mathfrak{B}_{\mathcal{V}}, \end{aligned}$$

where $\mathfrak{R}_{\mathcal{V}}$ and Bayes-error $\mathfrak{B}_{\mathcal{V}}$ are defined as

$$\begin{aligned} \mathfrak{R}_{\mathcal{V}} &:= \mathbb{E}_{S \sim Q_S} \left[\inf_h \Pr_{(x,y) \sim \sum_{v \in \mathcal{V}} Q(v|S) P_{v,\text{test}}} (h(x) \neq y) \right], \\ \mathfrak{B}_{\mathcal{V}} &:= \mathbb{E}_V [R(f^*(P_{V,\text{test}}); P_{V,\text{test}})]. \end{aligned}$$

Proof. By the definition of Minimax Excess Risk,

$$\begin{aligned} \text{Minimax Excess Risk} &= \inf_{\mathcal{A}} \sup_{(P_{\text{maj}}, P_{\text{min}}) \in \mathcal{P}} \mathbb{E}_{S \sim P_{\text{maj}}^{n_{\text{maj}}} \times P_{\text{min}}^{n_{\text{min}}}} [R(\mathcal{A}^S; P_{\text{test}})] - R(f^*(P_{\text{test}}); P_{\text{test}}) \\ &\geq \inf_{\mathcal{A}} \sup_{v \in \mathcal{V}} \mathbb{E}_{S|v \sim P_{v,\text{maj}}^{n_{\text{maj}}} \times P_{v,\text{min}}^{n_{\text{min}}}} [R(\mathcal{A}^S; P_{v,\text{test}})] - R(f^*(P_{v,\text{test}}); P_{v,\text{test}}) \\ &\geq \inf_{\mathcal{A}} \mathbb{E}_V \left[\mathbb{E}_{S|V \sim P_{V,\text{maj}}^{n_{\text{maj}}} \times P_{V,\text{min}}^{n_{\text{min}}}} [R(\mathcal{A}^S; P_{V,\text{test}})] - R(f^*(P_{V,\text{test}}); P_{V,\text{test}}) \right] \\ &= \inf_{\mathcal{A}} \mathbb{E}_V \left[\underbrace{\mathbb{E}_{S|V \sim P_{V,\text{maj}}^{n_{\text{maj}}} \times P_{V,\text{min}}^{n_{\text{min}}}} [R(\mathcal{A}^S; P_{V,\text{test}})] - \mathbb{E}_V [R(f^*(P_{V,\text{test}}); P_{V,\text{test}})]}_{=\mathfrak{B}_{\mathcal{V}}} \right]. \end{aligned}$$

We continue lower bounding the first term as follows

$$\begin{aligned} \inf_{\mathcal{A}} \mathbb{E}_V [\mathbb{E}_{S|V \sim P_{V,\text{maj}}^{n_{\text{maj}}} \times P_{V,\text{min}}^{n_{\text{min}}}} [R(\mathcal{A}^S; P_{V,\text{test}})]] &= \inf_{\mathcal{A}} \mathbb{E}_{(V,S) \sim Q} \left[\Pr_{(x,y) \sim P_{V,\text{test}}} (\mathcal{A}^S(x) \neq y) \right] \\ &= \inf_{\mathcal{A}} \mathbb{E}_{S \sim Q_S} \mathbb{E}_{V \sim Q(\cdot|S)} \left[\Pr_{(x,y) \sim P_{V,\text{test}}} (\mathcal{A}^S(x) \neq y) \right] \\ &\stackrel{(i)}{\geq} \mathbb{E}_{S \sim Q_S} \left[\inf_h \mathbb{E}_{V \sim Q(\cdot|S)} \left[\Pr_{(x,y) \sim P_{V,\text{test}}} (h(x) \neq y) \right] \right] \\ &= \mathbb{E}_{S \sim Q_S} \left[\inf_h \Pr_{(x,y) \sim \sum_{v \in \mathcal{V}} Q(v|S) P_{v,\text{test}}} (h(x) \neq y) \right] \\ &= \mathfrak{R}_{\mathcal{V}}, \end{aligned}$$

where (i) follows since \mathcal{A}^S is a fixed classifier given the sample set S . This, combined with the previous equation block completes the proof. \square

A.2 The hat function and its properties

In this section, we define the *hat function* and establish some of its properties. This function will be useful in defining “hard” problem instances to prove our lower bounds. Given a positive integer K the hat function is defined as

$$\phi_K(x) = \begin{cases} |x + \frac{1}{4K}| - \frac{1}{4K} & \text{for } x \in [-\frac{1}{2K}, 0], \\ \frac{1}{4K} - |x - \frac{1}{4K}| & \text{for } x \in [0, \frac{1}{2K}], \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

When K is clear from context, we omit the subscript.




Figure 3: The hat function with $K = 4$.

We first notice that this function is 1-Lipschitz and odd, so

$$\int_{-\frac{1}{2K}}^{\frac{1}{2K}} \phi_K(x) dx = 0.$$

We also compute some other key quantities for ϕ .

Lemma A.2. For any positive integer K ,

$$\int_{-\frac{1}{2K}}^{\frac{1}{2K}} |\phi_K(x)| dx = \frac{1}{8K^2}.$$

Proof. We suppress K in the notation. We have that,

$$\int_{-\frac{1}{2K}}^{\frac{1}{2K}} |\phi(x)| dx = \int_{-\frac{1}{2K}}^0 \left| \frac{1}{4K} - \left| x + \frac{1}{4K} \right| \right| dx + \int_0^{\frac{1}{2K}} \left| \left| x - \frac{1}{4K} \right| - \frac{1}{4K} \right| dx.$$

The integrand $\left| \frac{1}{4K} - \left| x + \frac{1}{4K} \right| \right|$ over $x \in [-\frac{1}{2K}, 0]$ defines a triangle with base $\frac{1}{2K}$ and height $\frac{1}{4K}$, thus it has area $\frac{1}{16K^2}$. Therefore,

$$\int_{-\frac{1}{2K}}^0 \left| \frac{1}{4K} - \left| x + \frac{1}{4K} \right| \right| dx = \frac{1}{16K^2}.$$

The same holds for the second term. Thus, by adding them up we get that $\int_{-\frac{1}{2K}}^{\frac{1}{2K}} |\phi(x)| dx = \frac{1}{8K^2}$. \square

Lemma A.3. For any positive integer K ,

$$\int_0^{\frac{1}{K}} \log \left(\frac{1 + \phi_K(x - \frac{1}{2K})}{1 - \phi_K(x - \frac{1}{2K})} \right) \left(1 + \phi_K \left(x - \frac{1}{2K} \right) \right) dx \leq \frac{1}{3K^3}$$

and

$$\int_0^{\frac{1}{K}} \log \left(\frac{1 - \phi_K(x - \frac{1}{2K})}{1 + \phi_K(x - \frac{1}{2K})} \right) \left(1 - \phi_K \left(x - \frac{1}{2K} \right) \right) dx \leq \frac{1}{3K^3}.$$

Proof. Let us suppress K in the notation. We prove the first bound below and the second bound follows by an identical argument. We have that

$$\begin{aligned}
& \int_0^{\frac{1}{2K}} \log \left(\frac{1 + \phi(x - \frac{1}{2K})}{1 - \phi(x - \frac{1}{2K})} \right) \left(1 + \phi \left(x - \frac{1}{2K} \right) \right) dx \\
&= \int_{-\frac{1}{2K}}^{\frac{1}{2K}} \log \left(\frac{1 + \phi(x)}{1 - \phi(x)} \right) (1 + \phi(x)) dx \\
&= \int_0^{\frac{1}{2K}} \log \left(\frac{1 + \phi(x)}{1 - \phi(x)} \right) (1 + \phi(x)) dx + \int_{-\frac{1}{2K}}^0 \log \left(\frac{1 + \phi(x)}{1 - \phi(x)} \right) (1 + \phi(x)) dx \\
&= \int_0^{\frac{1}{2K}} \log \left(\frac{1 + \phi(x)}{1 - \phi(x)} \right) (1 + \phi(x)) dx - \int_{\frac{1}{2K}}^0 \log \left(\frac{1 + \phi(-x)}{1 - \phi(-x)} \right) (1 + \phi(-x)) dx \\
&= \int_0^{\frac{1}{2K}} \log \left(\frac{1 + \phi(x)}{1 - \phi(x)} \right) (1 + \phi(x)) dx + \int_0^{\frac{1}{2K}} \log \left(\frac{1 - \phi(x)}{1 + \phi(x)} \right) (1 - \phi(x)) dx,
\end{aligned}$$

where the last equality follows since ϕ is an odd function. Now, we may collect the integrands to get that,

$$\begin{aligned}
& \int_0^{\frac{1}{2K}} \log \left(\frac{1 + \phi(x - \frac{1}{2K})}{1 - \phi(x - \frac{1}{2K})} \right) \left(1 + \phi \left(x - \frac{1}{2K} \right) \right) dx \\
&= 2 \int_0^{\frac{1}{2K}} \log \left(\frac{1 + \phi(x)}{1 - \phi(x)} \right) \phi(x) dx \\
&= 2 \int_0^{\frac{1}{2K}} \log \left(1 + \frac{2\phi(x)}{1 - \phi(x)} \right) \phi(x) dx \\
&\leq 2 \int_0^{\frac{1}{2K}} \frac{2\phi(x)^2}{1 - \phi(x)} dx,
\end{aligned}$$

where the last inequality follows since $\log(1+x) \leq x$ for all x . Now we observe that $\phi(x) \leq x \leq \frac{1}{2}$ for $x \in [0, \frac{1}{2K}]$, and in particular, $\frac{1}{1-\phi(x)} \leq 2$. Thus,

$$\begin{aligned}
& \int_0^{\frac{1}{2K}} \log \left(\frac{1 + \phi(x - \frac{1}{2K})}{1 - \phi(x - \frac{1}{2K})} \right) \left(1 + \phi \left(x - \frac{1}{2K} \right) \right) dx \\
&\leq 8 \int_0^{\frac{1}{2K}} \phi(x)^2 dx \\
&\leq 8 \int_0^{\frac{1}{2K}} x^2 dx \\
&= \frac{1}{3K^3}.
\end{aligned}$$

This proves the first bound. The second bound follows analogously. \square

B Proofs in the label shift setting

Throughout this section we operate in the label shift setting (Section 3.2.1).

First, in Appendix B.1 through a sequence of lemmas we prove the minimax lower bound Theorem 4.1. Next, in Appendix B.2 we prove Theorem 5.1 which is an upper bound on the excess risk of the undersampled binning estimator (see Eq. (5)) with $\lceil n_{\min} \rceil^{1/3}$ bins by invoking previous results on nonparametric density estimation (Freedman & Diaconis, 1981; Devroye & Györfi, 1985).

B.1 Proof of Theorem 4.1

In this section, we provide a proof of the minimax lower bound in the label shift setting.

We will proceed by constructing a class of distributions where the separation between any two distributions in the class is small enough such that it is hard to distinguish between them with finite minority class samples. In particular, we split the interval $[0, 1]$ into sub-intervals and each class distribution on each sub-interval either has slightly more probability mass on the left side of the sub-interval, on the right, or completely uniform. Since the minority class sample size is limited, no classifier will be able to tell which distribution the minority class is generated from, and hence will suffer high excess risk.

We construct the ‘‘hard’’ set of distributions as follows. Fix K to be an integer that will be specified in the sequel as a function of n_{\min} . Let the index set be $\mathcal{V} = \{-1, 0, 1\}^K \times \{-1, 0, 1\}^K$. For $v \in \mathcal{V}$, we will let $v_1 \in \{-1, 0, 1\}^K$ be the first K coordinates and $v_{-1} \in \{-1, 0, 1\}^K$ be the last K coordinates. That is, $v = (v_1, v_{-1})$.

For every $v \in \mathcal{P}$ we shall define pair of class-conditional distributions $P_{v,1}$ and $P_{v,-1}$ as follows: for $x \in I_j = [\frac{j-1}{K}, \frac{j}{K}]$,

$$\begin{aligned} P_{v,1}(x) &= 1 + v_{1,j}\phi\left(x - \frac{j+1/2}{K}\right) \\ P_{v,-1}(x) &= 1 + v_{-1,j}\phi\left(x - \frac{j+1/2}{K}\right), \end{aligned}$$

where ϕ is defined in Eq. 6. Notice that $P_{v,1}$ only depends on v_1 while $P_{v,-1}$ only depends on v_{-1} . We continue to define

$$\begin{aligned} P_{v,\text{maj}}(x, y) &= P_{v,1}(x)\mathbf{1}(y = 1) \\ P_{v,\text{min}}(x, y) &= P_{v,-1}(x)\mathbf{1}(y = -1), \end{aligned}$$

and

$$P_{v,\text{test}}(x, y) = \frac{P_{v,\text{maj}}(x, y) + P_{v,\text{min}}(x, y)}{2} = \frac{P_{v,1}(x)\mathbf{1}(y = 1) + P_{v,-1}(x)\mathbf{1}(y = -1)}{2}.$$

Observe that in the test distribution it is equally likely for the label to be $+1$ or -1 .

Recall that as described in Section A.1, V shall be a uniform random variable over \mathcal{V} and $S | V \sim P_{v,\text{maj}}^{n_{\text{maj}}} \times P_{v,\text{min}}^{n_{\text{min}}}$. We shall let \mathbf{Q} denote the joint distribution of (V, S) and let \mathbf{Q}_S denote the marginal over S .

With this construction in place, we first show that the minimax excess risk is lower bounded as follows.

Lemma B.1. *For any positive integers $K, n_{\text{maj}}, n_{\text{min}}$, the minimax excess risk is lower bounded as follows:*

$$\begin{aligned} &\text{Minimax Excess Risk}(\mathcal{P}_{\text{LS}}) \\ &= \inf_{\mathcal{A}} \sup_{(P_{\text{maj}}, P_{\text{min}}) \in \mathcal{P}_{\text{LS}}} \mathbb{E}_{S \sim P_{\text{maj}}^{n_{\text{maj}}} \times P_{\text{min}}^{n_{\text{min}}}} [R(\mathcal{A}^S; P_{\text{test}}) - R(f^*; P_{\text{test}})] \\ &\geq \frac{1}{36K} - \frac{1}{2} \mathbb{E}_{S \sim \mathbf{Q}_S} \left[\text{TV} \left(\sum_{v \in \mathcal{V}} \mathbf{Q}(v | S) P_{v,1}, \sum_{v \in \mathcal{V}} \mathbf{Q}(v | S) P_{v,-1} \right) \right]. \end{aligned} \quad (7)$$

Proof. By invoking Lemma A.1 we get that

$$\begin{aligned} &\text{Minimax Excess Risk}(\mathcal{P}_{\text{LS}}) \\ &\geq \underbrace{\mathbb{E}_{S \sim \mathbf{Q}_S} \left[\inf_{h(x,y) \sim \sum_{v \in \mathcal{V}} \mathbf{Q}(v|S) P_{v,\text{test}}} \Pr(h(x) \neq y) \right]}_{=: \mathfrak{R}_{\mathcal{V}}} - \underbrace{\mathbb{E}_V [R(f^*(P_{V,\text{test}}); P_{V,\text{test}})]}_{=: \mathfrak{B}_{\mathcal{V}}}. \end{aligned}$$

We proceed by calculating alternate expressions for $\mathfrak{R}_{\mathcal{V}}$ and $\mathfrak{B}_{\mathcal{V}}$ to get our desired lower bound on the minimax excess risk.

Calculation of $\mathfrak{R}_{\mathcal{V}}$: Immediately by Le Cam's lemma (Wainwright, 2019, Eq. 15.13), we get that

$$\begin{aligned}\mathfrak{R}_{\mathcal{V}} &= \mathbb{E}_{S \sim Q_S} \left[\inf_h \Pr_{(x,y) \sim \sum_{v \in \mathcal{V}} Q(v|S) P_{v,\text{test}}} (h(x) \neq y) \right] \\ &= \frac{1}{2} \mathbb{E}_{S \sim Q_S} \left[1 - \text{TV} \left(\sum_{v \in \mathcal{V}} Q(v|S) P_{v,1}, \sum_{v \in \mathcal{V}} Q(v|S) P_{v,-1} \right) \right].\end{aligned}\quad (8)$$

Calculation of $\mathfrak{B}_{\mathcal{V}}$: Again by invoking Le Cam's lemma (Wainwright, 2019, Eq. 15.13), we get that for any class conditional distributions P_1, P_{-1} ,

$$R(f^*; P_{\text{test}}) = \frac{1}{2} - \frac{1}{2} \text{TV}(P_1, P_{-1}).$$

So by taking expectations, we get that

$$\mathfrak{B}_{\mathcal{V}} = \mathbb{E}_V [R(f^*(P_{V,\text{test}}); P_{V,\text{test}})] = \mathbb{E}_V \left[\frac{1}{2} - \frac{1}{2} \text{TV}(P_{V,1}, P_{V,-1}) \right]. \quad (9)$$

We now compute $\mathbb{E}_V [\text{TV}(P_{V,1}, P_{V,-1})]$ as follows:

$$\begin{aligned}\mathbb{E}_V [\text{TV}(P_{V,1}, P_{V,-1})] &= \frac{1}{2} \mathbb{E}_V \left[\int_{x=0}^1 |P_{V,1}(x) - P_{V,-1}(x)| \, dx \right] \\ &= \frac{1}{2} \mathbb{E}_V \left[\sum_{j=1}^K \int_{\frac{j-1}{K}}^{\frac{j}{K}} |V_{1,j} - V_{-1,j}| \left| \phi \left(x - \frac{j+1/2}{K} \right) \right| \, dx \right] \\ &= \frac{1}{2} \sum_{j=1}^K \mathbb{E}_V \left[\int_{\frac{j-1}{K}}^{\frac{j}{K}} |V_{1,j} - V_{-1,j}| \left| \phi \left(x - \frac{j+1/2}{K} \right) \right| \, dx \right] \\ &\stackrel{(i)}{=} \frac{1}{16K^2} \sum_{j=1}^K \mathbb{E}_V [|V_{1,j} - V_{-1,j}|],\end{aligned}$$

where (i) follows by Lemma A.2. Observe that $V_{1,j}, V_{-1,j}$ are independent uniform random variables on $\{-1, 0, 1\}$, it is therefore straightforward to compute that

$$\mathbb{E}_V [|V_{1,j} - V_{-1,j}|] = \frac{8}{9}.$$

This yields that

$$\mathbb{E}_V [\text{TV}(P_{V,1}, P_{V,-1})] = \frac{1}{18K}.$$

Plugging this into Eq. (9) allows us to conclude that

$$\mathfrak{B}_{\mathcal{V}} = \mathbb{E}_V [R(f^*(P_{V,\text{test}}); P_{V,\text{test}})] = \frac{1}{2} \left(1 - \frac{1}{18K} \right). \quad (10)$$

Combining Eqs. (8) and (10) establishes the claimed result. \square

In light of this previous lemma we now aim to upper bound the expected total variation distance in Eq. (7).

Lemma B.2. *Suppose that v is drawn uniformly from the set $\{-1, 1\}^K$, and that $S \mid v$ is drawn from $P_{v,\text{maj}}^{n_{\text{maj}}} \times P_{v,\text{min}}^{n_{\text{min}}}$ then,*

$$\mathbb{E}_S \left[\text{TV} \left(\sum_{v \in \mathcal{V}} Q(v|S) P_{v,1}, \sum_{v \in \mathcal{V}} Q(v|S) P_{v,-1} \right) \right] \leq \frac{1}{18K} - \frac{1}{144K} \exp \left(-\frac{n_{\text{min}}}{3K^3} \right).$$

Proof. Let $\psi := \mathbb{E}_S [\text{TV}(\sum_{v \in \mathcal{V}} \mathbf{Q}(v | S) \mathbf{P}_{v,1}, \sum_{v \in \mathcal{V}} \mathbf{Q}(v | S) \mathbf{P}_{v,-1})]$. Then,

$$\begin{aligned} \psi &= \mathbb{E}_S \left[\text{TV} \left(\sum_{v \in \mathcal{V}} \mathbf{Q}(v | S) \mathbf{P}_{v,1}, \sum_{v \in \mathcal{V}} \mathbf{Q}(v | S) \mathbf{P}_{v,-1} \right) \right] \\ &= \frac{1}{2} \mathbb{E}_S \left[\int_{x=0}^1 \left| \sum_{v \in \mathcal{V}} \mathbf{Q}(v | S) (\mathbf{P}_{v,1}(x) - \mathbf{P}_{v,-1}(x)) \right| dx \right] \\ &= \frac{1}{2} \mathbb{E}_S \left[\sum_{j=1}^K \int_{x=\frac{j-1}{K}}^{\frac{j}{K}} \left| \sum_{v \in \mathcal{V}} \mathbf{Q}(v | S) (\mathbf{P}_{v,1}(x) - \mathbf{P}_{v,-1}(x)) \right| dx \right] \\ &= \frac{1}{2} \mathbb{E}_S \left[\sum_{j=1}^K \int_{x=\frac{j-1}{K}}^{\frac{j}{K}} \left| \sum_{v \in \mathcal{V}} \mathbf{Q}(v | S) (v_{1,j} - v_{-1,j}) \phi \left(x - \frac{j+1/2}{K} \right) \right| dx \right], \end{aligned}$$

where the last equality is by the definition of $\mathbf{P}_{v,1}$ and $\mathbf{P}_{v,-1}$. Continuing we get that,

$$\begin{aligned} \psi &= \frac{1}{2} \sum_{j=1}^K \left[\int_{x=\frac{j-1}{K}}^{\frac{j}{K}} \left| \phi \left(x - \frac{j+1/2}{K} \right) \right| dx \right] \mathbb{E}_S \left[\left| \sum_{v \in \mathcal{V}} \mathbf{Q}(v | S) (v_{1,j} - v_{-1,j}) \right| \right] \\ &\stackrel{(i)}{=} \frac{1}{16K^2} \mathbb{E}_S \left[\left| \sum_{j=1}^K \sum_{v \in \mathcal{V}} \mathbf{Q}(v | S) (v_{1,j} - v_{-1,j}) \right| \right] \\ &= \frac{1}{16K^2} \sum_{j=1}^K \int \left| \sum_{v \in \mathcal{V}} \mathbf{Q}(v | S) (v_{1,j} - v_{-1,j}) \right| d\mathbf{Q}_S(S) \\ &= \frac{1}{16K^2} \sum_{j=1}^K \int \left| \sum_{v \in \mathcal{V}} \mathbf{Q}(v, S) (v_{1,j} - v_{-1,j}) \right| dS \\ &\stackrel{(ii)}{=} \frac{1}{16K^2 |\mathcal{V}|} \sum_{j=1}^K \int \left| \sum_{v \in \mathcal{V}} \mathbf{Q}(S | v) (v_{1,j} - v_{-1,j}) \right| dS, \end{aligned}$$

where (i) follows by the calculation in Lemma A.2 and (ii) follows since v is a uniform random variable over the set \mathcal{V} .

The distributions $\mathbf{P}_{v,1}$ and $\mathbf{P}_{v,-1}$ are symmetrically defined over all intervals $I_j = [\frac{j-1}{K}, \frac{j}{K}]$, and hence all of the summands in the RHS above are equal. Thus,

$$\psi = \frac{1}{16K |\mathcal{V}|} \int \left| \sum_{v \in \mathcal{V}} \mathbf{Q}(S | v) (v_{1,1} - v_{-1,1}) \right| dS. \quad (11)$$

Before we continue further, let us define

$$\mathcal{V}^+ = \{v \in \mathcal{V} \mid v_{1,1} > v_{-1,1}\}.$$

For every $v \in \mathcal{V}^+$, let $\tilde{v} \in \mathcal{V}$ be such that is the same as v on all coordinates, except $\tilde{v}_{1,1} = -v_{1,1}$ and $\tilde{v}_{-1,1} = -v_{-1,1}$. Then continuing from Eq. (11) we find that,

$$\begin{aligned}
\psi &\stackrel{(i)}{=} \frac{1}{16K|\mathcal{V}|} \int \left| \sum_{v \in \mathcal{V}^+} (v_{1,1} - v_{-1,1}) (\mathbb{Q}(S | v) - \mathbb{Q}(S | \tilde{v})) \right| dS \\
&\stackrel{(ii)}{\leq} \frac{1}{16K|\mathcal{V}|} \int \sum_{v \in \mathcal{V}^+} (v_{1,1} - v_{-1,1}) |\mathbb{Q}(S | v) - \mathbb{Q}(S | \tilde{v})| dS \\
&= \frac{1}{16K|\mathcal{V}|} \sum_{v \in \mathcal{V}^+} (v_{1,1} - v_{-1,1}) \int |\mathbb{Q}(S | v) - \mathbb{Q}(S | \tilde{v})| dS \\
&= \frac{1}{8K|\mathcal{V}|} \underbrace{\sum_{v \in \mathcal{V}^+} (v_{1,1} - v_{-1,1}) \text{TV}(\mathbb{Q}(S | v), \mathbb{Q}(S | \tilde{v}))}_{=: \Xi}, \tag{12}
\end{aligned}$$

where (i) we use the definition of \mathcal{V}^+ and \tilde{v} , (ii) follows since $v_{1,1} > v_{-1,1}$ for $v \in \mathcal{V}^+$.

Now we further partition \mathcal{V}^+ into 3 sets $\mathcal{V}^{(1,0)}$, $\mathcal{V}^{(0,-1)}$, $\mathcal{V}^{(1,-1)}$ as follows

$$\begin{aligned}
\mathcal{V}^{(1,0)} &= \{v \in \mathcal{V} \mid v_{1,1} = 1, v_{-1,1} = 0\}, \\
\mathcal{V}^{(0,-1)} &= \{v \in \mathcal{V} \mid v_{1,1} = 0, v_{-1,1} = -1\}, \\
\mathcal{V}^{(1,-1)} &= \{v \in \mathcal{V} \mid v_{1,1} = 1, v_{-1,1} = -1\}.
\end{aligned}$$

Note that $\mathbb{Q}(S | v) = \mathbb{P}_{v,\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{v,\text{min}}^{n_{\text{min}}}$, and therefore

$$\begin{aligned}
\Xi &= \sum_{v \in \mathcal{V}^+} (v_{1,1} - v_{-1,1}) \text{TV}(\mathbb{P}_{v,\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{v,\text{min}}^{n_{\text{min}}}, \mathbb{P}_{\tilde{v},\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\tilde{v},\text{min}}^{n_{\text{min}}}) \\
&\stackrel{(i)}{=} \sum_{v \in \mathcal{V}^{(1,0)}} \text{TV}(\mathbb{P}_{v,\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{v,\text{min}}^{n_{\text{min}}}, \mathbb{P}_{\tilde{v},\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\tilde{v},\text{min}}^{n_{\text{min}}}) \\
&\quad + \sum_{v \in \mathcal{V}^{(0,-1)}} \text{TV}(\mathbb{P}_{v,\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{v,\text{min}}^{n_{\text{min}}}, \mathbb{P}_{\tilde{v},\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\tilde{v},\text{min}}^{n_{\text{min}}}) \\
&\quad + 2 \sum_{v \in \mathcal{V}^{(1,-1)}} \text{TV}(\mathbb{P}_{v,\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{v,\text{min}}^{n_{\text{min}}}, \mathbb{P}_{\tilde{v},\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\tilde{v},\text{min}}^{n_{\text{min}}}), \tag{13}
\end{aligned}$$

where (i) follows since $v_1, v_{-1} \in \{-1, 0, 1\}^K$ and by the definition of the sets $\mathcal{V}^{(1,0)}$, $\mathcal{V}^{(0,-1)}$ and $\mathcal{V}^{(1,-1)}$.

Now by the Bretagnolle–Huber inequality (see Canonne, 2022, Corollary 4),

$$\begin{aligned}
\text{TV}(\mathbb{P}_{v,\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{v,\text{min}}^{n_{\text{min}}}, \mathbb{P}_{\tilde{v},\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\tilde{v},\text{min}}^{n_{\text{min}}}) &= \text{TV}(\mathbb{P}_{\tilde{v},\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\tilde{v},\text{min}}^{n_{\text{min}}}, \mathbb{P}_{v,\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{v,\text{min}}^{n_{\text{min}}}) \\
&\leq 1 - \frac{1}{2} \exp(-\text{KL}(\mathbb{P}_{\tilde{v},\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\tilde{v},\text{min}}^{n_{\text{min}}} \parallel \mathbb{P}_{v,\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{v,\text{min}}^{n_{\text{min}}})) ,
\end{aligned}$$

where we flip the arguments in the first step for simplicity later.

Next, by the chain rule for KL-divergence, we have that

$$\text{KL}(\mathbb{P}_{\tilde{v},\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\tilde{v},\text{min}}^{n_{\text{min}}} \parallel \mathbb{P}_{v,\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{v,\text{min}}^{n_{\text{min}}}) = n_{\text{maj}} \text{KL}(\mathbb{P}_{\tilde{v},\text{maj}} \parallel \mathbb{P}_{v,\text{maj}}) + n_{\text{min}} \text{KL}(\mathbb{P}_{\tilde{v},\text{min}} \parallel \mathbb{P}_{v,\text{min}}).$$

Using these, let us upper bound the first term in Eq. (13) corresponding to $v \in \mathcal{V}^{(0,-1)}$. For $v \in \mathcal{V}^{(0,-1)}$, notice that $\text{KL}(\mathbb{P}_{\tilde{v},\text{maj}} \parallel \mathbb{P}_{v,\text{maj}}) = 0$ since $v_{1,j} = \tilde{v}_{1,j}$ for all $j \in \{1, \dots, K\}$. For the second term, $\text{KL}(\mathbb{P}_{\tilde{v},\text{min}} \parallel \mathbb{P}_{v,\text{min}})$,

only $v_{1,1}$ and $\tilde{v}_{1,1}$ differ, so

$$\begin{aligned} \text{KL}(\mathbf{P}_{\tilde{v},\min} \parallel \mathbf{P}_{v,\min}) &= \int_0^1 \mathbf{P}_{v,-1}(x) \log \left(\frac{\mathbf{P}_{v,-1}(x)}{\mathbf{P}_{\tilde{v},-1}(x)} \right) dx \\ &= \int_0^{\frac{1}{K}} \log \left(\frac{1 + \phi_K(x - \frac{1}{2K})}{1 - \phi_K(x - \frac{1}{2K})} \right) \left(1 + \phi_K \left(x - \frac{1}{2K} \right) \right) dx \\ &\leq \frac{1}{3K^3}, \end{aligned}$$

where the last inequality is a result of the calculation in Lemma A.3.

Therefore, we get

$$\sum_{v \in \mathcal{V}^{(0,-1)}} \text{TV}(\mathbf{P}_{v,\text{maj}}^{n_{\text{maj}}} \times \mathbf{P}_{v,\text{min}}^{n_{\text{min}}}, \mathbf{P}_{\tilde{v},\text{maj}}^{n_{\text{maj}}} \times \mathbf{P}_{\tilde{v},\text{min}}^{n_{\text{min}}}) \leq 9^{K-1} \left(1 - \frac{1}{2} \exp \left(-\frac{n_{\text{min}}}{3K^3} \right) \right).$$

For the terms in Eq. (13) corresponding to $\mathcal{V}^{(0,-1)}$, $\mathcal{V}^{(1,-1)}$, we simply take the trivial bound to get

$$\begin{aligned} \sum_{v \in \mathcal{V}^{(0,-1)}} \text{TV}(\mathbf{P}_{v,\text{maj}}^{n_{\text{maj}}} \times \mathbf{P}_{v,\text{min}}^{n_{\text{min}}}, \mathbf{P}_{\tilde{v},\text{maj}}^{n_{\text{maj}}} \times \mathbf{P}_{\tilde{v},\text{min}}^{n_{\text{min}}}) &\leq 9^{K-1}, \\ \sum_{v \in \mathcal{V}^{(1,-1)}} \text{TV}(\mathbf{P}_{v,\text{maj}}^{n_{\text{maj}}} \times \mathbf{P}_{v,\text{min}}^{n_{\text{min}}}, \mathbf{P}_{\tilde{v},\text{maj}}^{n_{\text{maj}}} \times \mathbf{P}_{\tilde{v},\text{min}}^{n_{\text{min}}}) &\leq 9^{K-1}. \end{aligned}$$

Plugging these bounds into Eq. (13) we get that,

$$\Xi \leq 4 \cdot 9^{K-1} - \frac{9^{K-1}}{2} \exp \left(-\frac{n_{\text{min}}}{3K^3} \right).$$

Now using this bound on Ξ in Eq. (12) and observing that $|\mathcal{V}| = 9^K$, we get that,

$$\begin{aligned} \psi &= \mathbb{E}_S \left[\text{TV} \left(\sum_{v \in \mathcal{V}} Q(v | S) P_{v,1}, \sum_{v \in \mathcal{V}} Q(v | S) P_{v,-1} \right) \right] \\ &\leq \frac{1}{8 \cdot 9^K K} \left(4 \cdot 9^{K-1} - \frac{9^{K-1}}{2} \exp \left(-\frac{n_{\text{min}}}{3K^3} \right) \right) \\ &= \frac{1}{18K} - \frac{1}{144K} \exp \left(-\frac{n_{\text{min}}}{3K^3} \right), \end{aligned}$$

completing the proof. \square

Finally, we combine Lemma B.1 and Lemma B.2 to establish the minimax lower bound in this label shift setting. We recall the statement of the theorem here.

Theorem 4.1. *Consider the label shift setting described in Section 3.2.1. Recall that \mathcal{P}_{LS} is the class of pairs of distributions $(\mathbf{P}_{\text{maj}}, \mathbf{P}_{\text{min}})$ that satisfy the assumptions in that section. The minimax excess risk over this class is lower bounded as follows:*

$$\text{Minimax Excess Risk}(\mathcal{P}_{\text{LS}}) = \inf_{\mathcal{A}} \sup_{(\mathbf{P}_{\text{maj}}, \mathbf{P}_{\text{min}}) \in \mathcal{P}_{\text{LS}}} \text{Excess Risk}[\mathcal{A}; (\mathbf{P}_{\text{maj}}, \mathbf{P}_{\text{min}})] \geq \frac{1}{600} \frac{1}{n_{\text{min}}^{1/3}}. \quad (3)$$

Proof. By Lemma B.1 we know that,

$$\text{Minimax Excess Risk}(\mathcal{P}_{\text{LS}}) \geq \frac{1}{36K} - \frac{1}{2} \mathbb{E}_{S \sim Q_S} \left[\text{TV} \left(\sum_{v \in \mathcal{V}} Q(v | S) P_{v,1}, \sum_{v \in \mathcal{V}} Q(v | S) P_{v,-1} \right) \right].$$

Next by the calculation in Lemma B.2 we have that

$$\begin{aligned} \text{Minimax Excess Risk}(\mathcal{P}_{\text{LS}}) &\geq \frac{1}{36K} - \frac{1}{2} \left(\frac{1}{18K} - \frac{1}{144K} \exp\left(-\frac{n_{\min}}{3K^3}\right) \right) \\ &= \frac{1}{288K} \exp\left(-\frac{n_{\min}}{3K^3}\right). \end{aligned}$$

Setting $K = \lceil n_{\min}^{1/3} \rceil$ yields the following

$$\begin{aligned} \text{Minimax Excess Risk}(\mathcal{P}_{\text{LS}}) &\geq \frac{1}{288 \lceil n_{\min}^{1/3} \rceil} \exp\left(-\frac{n_{\min}}{3 \lceil n_{\min}^{1/3} \rceil^3}\right) \\ &\geq \frac{\exp\left(-\frac{n_{\min}}{3 \lceil n_{\min}^{1/3} \rceil^3}\right)}{288} \frac{n_{\min}^{1/3}}{\lceil n_{\min}^{1/3} \rceil} \frac{1}{n_{\min}^{1/3}} \\ &\stackrel{(i)}{\geq} \frac{0.7 \exp\left(-\frac{1}{3}\right)}{288} \frac{1}{n_{\min}^{1/3}} \\ &\geq \frac{1}{600} \frac{1}{n_{\min}^{1/3}}, \end{aligned}$$

where (i) follows since $n_{\min}^{1/3} / \lceil n_{\min}^{1/3} \rceil \geq 0.7$ for $n_{\min} \geq 1$. \square

B.2 Proof of Theorem 5.1

In this section, we derive an upper bound on the excess risk of the undersampled binning estimator \mathcal{A}_{USB} (Eq. (5)) in the label shift setting. Recall that given a dataset \mathcal{S} this estimator first calculates the undersampled dataset \mathcal{S}_{US} , where the number of points from the minority group (n_{\min}) is equal to the number of points from the majority group (n_{\min}), and the size of the dataset is $2n_{\min}$. Throughout this section, $(\mathbb{P}_{\text{maj}}, \mathbb{P}_{\text{min}})$ shall be an arbitrary element of \mathcal{P}_{LS} .

To bound the excess risk of the undersampling algorithm, we will relate it to density estimation.

Recall that $n_{1,j}$ denotes the number of points in \mathcal{S}_{US} with label +1 that lie in I_j , and $n_{-1,j}$ is defined analogously.

Given a positive integer K , for $x \in I_j = [\frac{j-1}{K}, \frac{j}{K}]$, by the definition of the undersampled binning estimator (Eq. (5))

$$\mathcal{A}_{\text{USB}}^{\mathcal{S}}(x) = \begin{cases} 1 & \text{if } n_{1,j} > n_{-1,j}, \\ -1 & \text{otherwise.} \end{cases}$$

Recall that since we have undersampled, $\sum_j n_{1,j} = \sum_j n_{-1,j} = n_{\min}$. Therefore, define the simple histogram estimators for $\hat{\mathbb{P}}_1(x) = \mathbb{P}(x | y = 1)$ and $\hat{\mathbb{P}}_{-1}(x) = \mathbb{P}(x | y = -1)$ as follows: for $x \in I_j$,

$$\hat{\mathbb{P}}_1^{\mathcal{S}}(x) := \frac{n_{1,j}}{Kn_{\min}} \quad \text{and} \quad \hat{\mathbb{P}}_{-1}^{\mathcal{S}}(x) := \frac{n_{-1,j}}{Kn_{\min}}.$$

With this histogram estimator in place, we may define an estimator for $\eta(x) := \mathbb{P}_{\text{test}}(y = 1 | x)$ as follows,

$$\hat{\eta}^{\mathcal{S}}(x) := \frac{\hat{\mathbb{P}}_1^{\mathcal{S}}(x)}{\hat{\mathbb{P}}_1^{\mathcal{S}}(x) + \hat{\mathbb{P}}_{-1}^{\mathcal{S}}(x)}.$$

Observe that, for $x \in I_j$

$$\hat{\eta}^{\mathcal{S}}(x) > 1/2 \iff n_{1,j} > n_{-1,j} \iff \mathcal{A}_{\text{USB}}^{\mathcal{S}}(x) = 1.$$

Defining an estimator $\hat{\eta}^{\mathcal{S}}$ for the $\mathbb{P}_{\text{test}}(y = 1 | x)$ in this way will allow us to relate the excess risk of \mathcal{A}_{USB} to the estimation error in $\hat{\mathbb{P}}_1^{\mathcal{S}}$ and $\hat{\mathbb{P}}_{-1}^{\mathcal{S}}$.

Before proving the theorem we restate it here.

Theorem 5.1. Consider the label shift setting described in Section 3.2.1. For any $(P_{\text{maj}}, P_{\text{min}}) \in \mathcal{P}_{\text{LS}}$ the expected excess risk of the Undersampling Binning Estimator (Eq. (5)) with number of bins with $K = c \lceil n_{\text{min}}^{1/3} \rceil$ is upper bounded by

$$\text{Excess Risk}[\mathcal{A}_{\text{USB}}; (P_{\text{maj}}, P_{\text{min}})] = \mathbb{E}_{\mathcal{S} \sim P_{\text{maj}}^{n_{\text{maj}}} \times P_{\text{min}}^{n_{\text{min}}}} [R(\mathcal{A}_{\text{USB}}^{\mathcal{S}}; P_{\text{test}}) - R(f^*; P_{\text{test}})] \leq \frac{C}{n_{\text{min}}^{1/3}}.$$

Proof. By the definition of the excess risk

$$\text{Excess Risk}[\mathcal{A}_{\text{USB}}; (P_{\text{maj}}, P_{\text{min}})] := \mathbb{E}_{\mathcal{S} \sim P_{\text{maj}}^{n_{\text{maj}}} \times P_{\text{min}}^{n_{\text{min}}}} [R(\mathcal{A}_{\text{USB}}^{\mathcal{S}}; P_{\text{test}}) - R(f^*; P_{\text{test}})].$$

By invoking (Wasserman, 2019, Theorem 1) we may upper bound the excess risk given a draw of \mathcal{S} by

$$R(\mathcal{A}_{\text{USB}}^{\mathcal{S}}; P_{\text{test}}) - R(f^*; P_{\text{test}}) \leq 2 \int |\hat{\eta}^{\mathcal{S}}(x) - \eta(x)| P_{\text{test}}(x) dx.$$

Continuing using the definition of $\hat{\eta}^{\mathcal{S}}$ above and because $\eta = P_1/(P_1 + P_{-1})$ we have that,

$$\begin{aligned} & R(\mathcal{A}_{\text{USB}}^{\mathcal{S}}; P_{\text{test}}) - R(f^*; P_{\text{test}}) \\ &= 2 \int_0^1 \left| \frac{\hat{P}_1^{\mathcal{S}}(x)}{\hat{P}_1^{\mathcal{S}}(x) + \hat{P}_{-1}^{\mathcal{S}}(x)} - \frac{P_1(x)}{P_1(x) + P_{-1}(x)} \right| \left(\frac{P_1(x) + P_{-1}(x)}{2} \right) dx \\ &= \int_0^1 \left| \left(\frac{P_1(x) + P_{-1}(x)}{\hat{P}_1^{\mathcal{S}}(x) + \hat{P}_{-1}^{\mathcal{S}}(x)} \right) \hat{P}_1^{\mathcal{S}}(x) - P_1(x) \right| dx \\ &\stackrel{(i)}{\leq} \int_0^1 |\hat{P}_1^{\mathcal{S}}(x) - P_1(x)| dx + \int_0^1 \left| \frac{P_1(x) + P_{-1}(x)}{\hat{P}_1^{\mathcal{S}}(x) + \hat{P}_{-1}^{\mathcal{S}}(x)} - 1 \right| \hat{P}_1^{\mathcal{S}}(x) dx \\ &= \int_0^1 |\hat{P}_1^{\mathcal{S}}(x) - P_1(x)| dx + \int_0^1 |\hat{P}_1^{\mathcal{S}}(x) + \hat{P}_{-1}^{\mathcal{S}}(x) - P_1(x) - P_{-1}(x)| \frac{\hat{P}_1^{\mathcal{S}}(x)}{\hat{P}_1^{\mathcal{S}}(x) + \hat{P}_{-1}^{\mathcal{S}}(x)} dx \\ &\leq 2 \int_0^1 |\hat{P}_1^{\mathcal{S}}(x) - P_1(x)| dx + \int_0^1 |\hat{P}_{-1}^{\mathcal{S}}(x) - P_{-1}(x)| dx \\ &\stackrel{(ii)}{\leq} 2 \sqrt{\int_0^1 (\hat{P}_1^{\mathcal{S}}(x) - P_1(x))^2 dx} + \sqrt{\int_0^1 (\hat{P}_{-1}^{\mathcal{S}}(x) - P_{-1}(x))^2 dx}, \end{aligned}$$

where (i) follows by the triangle inequality, (ii) is by the Cauchy–Schwarz inequality.

Taking expectation over the samples \mathcal{S} and by invoking Jensen’s inequality we find that,

$$\begin{aligned} & \text{Excess Risk}(\mathcal{A}^{\mathcal{S}}; (P_{\text{maj}}, P_{\text{min}})) \\ &= \mathbb{E}_{\mathcal{S}} [R(\mathcal{A}_{\text{USB}}^{\mathcal{S}}; P_{\text{test}}) - R(f^*; P_{\text{test}})] \\ &\leq 2 \sqrt{\mathbb{E}_{\mathcal{S}} \left[\int (\hat{P}_1^{\mathcal{S}}(x) - P_1(x))^2 dx \right]} + \sqrt{\mathbb{E}_{\mathcal{S}} \left[\int (\hat{P}_{-1}^{\mathcal{S}}(x) - P_{-1}(x))^2 dx \right]}. \end{aligned}$$

We note that $\hat{P}_j^{\mathcal{S}}$ only depends on n_{min} i.i.d. draws from class j . Thus by (Freedman & Diaconis, 1981, Theorem 1.7), if $K = c \lceil n_{\text{min}}^{1/3} \rceil$ then

$$\mathbb{E}_{\mathcal{S}} \left[\int (\hat{P}_j^{\mathcal{S}}(x) - P_j(x))^2 dx \right] \leq \frac{C}{n_{\text{min}}^{2/3}}.$$

Plugging this into the previous inequality yields the desired result. \square

C Proof in the group-covariate shift setting

Throughout this section we operate in the group-covariate shift setting (Section 3.2.2).

We will proceed similarly to Section B. We shall construct a family of class-conditional distributions such that it will be necessary for adequate samples in each sub-interval of $[0, 1]$ to be able to learn the maximally likely label in that sub-interval. On the other hand, we will construct the group-covariate distributions to be separated from one another. As a consequence, sub-intervals with high probability mass under the minority group distribution will have low probability mass under the majority group distribution. Hence, these sub-intervals will not have enough training sample points for any classifier to be able to learn the maximally likely label and as a result shall suffer high excess risk.

First in Appendix C.1, we prove Theorem 4.2, the minimax lower bound through a sequence of lemmas. Second in Appendix C.2, we prove Theorem 5.2 that upper bound on the excess risk of the undersampled binning estimator with $\lceil n_{\min} \rceil^{1/3}$ bins.

C.1 Proof of Theorem 4.2

In this section, we provide a proof of the minimax lower bound in the group shift setting.

We construct the “hard” set of distributions as follows. Let the index set be $\mathcal{V} = \{-1, 1\}^K$. For every $v \in \mathcal{V}$ define a distribution as follows: for $x \in I_j = [\frac{j-1}{K}, \frac{j}{K}]$,

$$P_v(y = 1 | x) := \frac{1}{2} \left[1 + v_j \phi \left(x - \frac{j + 1/2}{K} \right) \right],$$

where ϕ is defined in Eq. 6. Given a $\tau \in [0, 1]$ we also construct the group distributions as follows:

$$P_a(x) = \begin{cases} 2 - \tau & \text{if } x \in [0, 0.5] \\ \tau & \text{if } x \in [0.5, 1], \end{cases}$$

and let

$$P_b(x) = 2 - P_a(x).$$

We can verify that

$$\text{Overlap}(P_a, P_b) = 1 - \text{TV}(P_a, P_b) = 1 - \frac{1}{2} \int_{x=0}^1 |P_a(x) - P_b(x)| dx = \tau.$$

We continue to define

$$\begin{aligned} P_{v,\text{maj}}(x, y) &= P_v(y | x) P_a(x) \\ P_{v,\text{min}}(x, y) &= P_v(y | x) P_b(x), \end{aligned}$$

and

$$P_{v,\text{test}}(x, y) = P_v(y | x) \left(\frac{P_a(x) + P_b(x)}{2} \right).$$

Observe that $(P_a(x) + P_b(x))/2 = 1$, the uniform distribution over $[0, 1]$.

Recall that as described in Section A.1, V shall be a uniform random variable over \mathcal{V} and $S | V \sim P_{v,\text{maj}}^{n_{\text{maj}}} \times P_{v,\text{min}}^{n_{\text{min}}}$. We shall let \mathbf{Q} denote the joint distribution of (V, S) and let \mathbf{Q}_S denote the marginal over S .

With this construction in place, we present the following lemma that lower bounds the minimax excess risk by a sum of $\exp(-\text{KL}(\mathbf{Q}(S | v_j = 1) \| \mathbf{Q}(S | v_j = -1)))$ over the intervals. Intuitively, $\text{KL}(\mathbf{Q}(S | v_j = 1) \| \mathbf{Q}(S | v_j = -1))$ is a measure of how difficult it is to identify whether $v_j = 1$ or $v_j = -1$ from the samples.

Lemma C.1. For any positive integers $K, n_{\text{maj}}, n_{\text{min}}$ and $\tau \in [0, 1]$, the minimax excess risk is lower bounded as follows:

$$\begin{aligned} \text{Minimax Excess Risk}(\mathcal{P}_{\text{GS}}(\tau)) &= \inf_{\mathcal{A}} \sup_{(P_{\text{maj}}, P_{\text{min}}) \in \mathcal{P}_{\text{GS}}(\tau)} \mathbb{E}_{S \sim P_{\text{maj}}^{n_{\text{maj}}} \times P_{\text{min}}^{n_{\text{min}}}} [R(\mathcal{A}^S; P_{\text{test}}) - R(f^*; P_{\text{test}})] \\ &\geq \frac{1}{32K^2} \sum_{j=1}^K \exp(-\text{KL}(\mathbb{Q}(S | v_j = 1) \| \mathbb{Q}(S | v_j = -1))). \end{aligned}$$

Proof. By invoking Lemma A.1, we know that the minimax excess risk is lower bounded by

$$\begin{aligned} &\text{Minimax Excess Risk}(\mathcal{P}_{\text{GS}}(\tau)) \\ &\geq \underbrace{\mathbb{E}_{S \sim \mathbb{Q}_S} [\inf_h \Pr_{(x,y) \sim \sum_{v \in \mathcal{V}} \mathbb{Q}(v|S) P_{v,\text{test}}} (h(x) \neq y)]}_{=\mathfrak{R}_{\mathcal{V}}} - \underbrace{\mathbb{E}_V [R(f^*(P_{V,\text{test}}); P_{V,\text{test}})]}_{=\mathfrak{B}_{\mathcal{V}}}, \end{aligned}$$

where V is a uniform random variable over the set \mathcal{V} , $S | V = v$ is a draw from $P_{v,\text{maj}}^{n_{\text{maj}}} \times P_{v,\text{min}}^{n_{\text{min}}}$, and \mathbb{Q} denotes the joint distribution over (V, S) .

We shall lower bound this minimax risk in parts. First, we shall establish a lower bound on $\mathfrak{R}_{\mathcal{V}}$, and then an upper bound on the Bayes risk $\mathfrak{B}_{\mathcal{V}}$.

Lower bound on $\mathfrak{R}_{\mathcal{V}}$. Unpacking $\mathfrak{R}_{\mathcal{V}}$ using its definition we get that,

$$\begin{aligned} \mathfrak{R}_{\mathcal{V}} &= \mathbb{E}_{S \sim \mathbb{Q}_S} [\inf_h \Pr_{(x,y) \sim \sum_{v \in \mathcal{V}} \mathbb{Q}(v|S) P_{v,\text{test}}} (h(x) \neq y)] \\ &= \mathbb{E}_{S \sim \mathbb{Q}_S} \left[\inf_h \int_0^1 P_{\text{test}}(x) \Pr_{y \sim \sum_{v \in \mathcal{V}} \mathbb{Q}(v|S) P_{v,(\cdot|x)}} [h(x) \neq y] dx \right] \\ &\stackrel{(i)}{=} \mathbb{E}_{S \sim \mathbb{Q}_S} \left[\int_0^1 P_{\text{test}}(x) \min \left\{ \sum_{v \in \mathcal{V}} \mathbb{Q}(v | S) P_v(1 | x), \sum_{v \in \mathcal{V}} \mathbb{Q}(v | S) P_v(-1 | x) \right\} dx \right] \\ &\stackrel{(ii)}{=} \frac{1}{2} - \mathbb{E}_{S \sim \mathbb{Q}_S} \left[\int_0^1 P_{\text{test}}(x) \left| \frac{1}{2} - \sum_{v \in \mathcal{V}} \mathbb{Q}(v | S) P_v(1 | x) \right| dx \right] \\ &\stackrel{(iii)}{=} \frac{1}{2} - \int_0^1 P_{\text{test}}(x) \mathbb{E}_{S \sim \mathbb{Q}_S} \left[\left| \frac{1}{2} - \sum_{v \in \mathcal{V}} \mathbb{Q}(v | S) P_v(1 | x) \right| \right] dx, \end{aligned} \tag{14}$$

where (i) follows by taking h to be the pointwise minimizer over x , (ii) follows since $P_v(-1 | x) = 1 - P_v(1 | x)$ and $\min\{s, 1 - s\} = (1 - |1 - 2s|)/2$ for all $s \in [0, 1]$, and (iii) follows by Fubini's theorem which allows us to switch the order of the integrals.

If $x \in I_j = [\frac{j-1}{K}, \frac{j}{K}]$ for some $j \in \{1, \dots, K\}$ we let j_x denote the value of this index j . With this notation in place let us continue to upper bound integrand in the second term in the RHS above as follows:

$$\begin{aligned} &\mathbb{E}_{S \sim \mathbb{Q}_S} \left[\left| \frac{1}{2} - \sum_{v \in \mathcal{V}} \mathbb{Q}(v | S) P_v(1 | x) \right| \right] \\ &\stackrel{(i)}{=} \mathbb{E}_{S \sim \mathbb{Q}_S} \left[\left| \phi \left(x - \frac{j_x + 1/2}{K} \right) \right| |\mathbb{Q}(v_{j_x} = 1 | S) - \mathbb{Q}(v_{j_x} = -1 | S)| \right] \\ &= \left| \phi \left(x - \frac{j_x + 1/2}{K} \right) \right| \mathbb{E}_{S \sim \mathbb{Q}_S} [|\mathbb{Q}(v_{j_x} = 1 | S) - \mathbb{Q}(v_{j_x} = -1 | S)|] \\ &\stackrel{(ii)}{=} \left| \phi \left(x - \frac{j_x + 1/2}{K} \right) \right| \mathbb{E}_{S \sim \mathbb{Q}_S} \left[\left| \frac{\mathbb{Q}(S | v_{j_x} = 1) \mathbb{Q}_V(v_{j_x} = 1)}{\mathbb{Q}_S(S)} - \frac{\mathbb{Q}(S | v_{j_x} = -1) \mathbb{Q}_V(v_{j_x} = -1)}{\mathbb{Q}_S(S)} \right| \right] \\ &\stackrel{(iii)}{=} \frac{1}{2} \left| \phi \left(x - \frac{j_x + 1/2}{K} \right) \right| \text{TV}(\mathbb{Q}(S | v_{j_x} = 1), \mathbb{Q}(S | v_{j_x} = -1)), \end{aligned} \tag{15}$$

where (i) follows since $P_v(1 | x) = (1 + v_{j_x} \phi(x - (j_x + 1/2)/K))/2$ and by marginalizing $Q(v | S)$ over the indices $j \neq j_x$, (ii) follows by using Bayes' rule and (iii) follows since the total-variation distance is half the ℓ_1 distance. Now by the Bretagnolle–Huber inequality (see Canonne, 2022, Corollary 4) we get that,

$$\begin{aligned} & \text{TV}(Q(S | v_{j_x} = 1), Q(S | v_{j_x} = -1)) \\ & \leq 1 - \frac{\exp(-\text{KL}(Q(S | v_{j_x} = 1) \| Q(S | v_{j_x} = -1)))}{2}. \end{aligned} \quad (16)$$

Combining Eqs. (14)-(16) we get that

$$\begin{aligned} & \mathfrak{R}_V \\ & \geq \frac{1}{2} - \frac{1}{2} \int_0^1 P_{\text{test}}(x) \left| \phi \left(x - \frac{j_x + 1/2}{K} \right) \right| dx \\ & \quad + \frac{1}{4} \int_0^1 P_{\text{test}}(x) \left| \phi \left(x - \frac{j_x + 1/2}{K} \right) \right| \exp(-\text{KL}(Q(S | v_{j_x} = 1) \| Q(S | v_{j_x} = -1))) dx. \end{aligned} \quad (17)$$

Upper bound on \mathfrak{B}_V : The Bayes error is

$$\begin{aligned} \mathfrak{B}_V &= \mathbb{E}_V[R(f^*(P_V); P_V)] \\ &= \mathbb{E}_V \left[\inf_f \mathbb{E}_{(x,y) \sim P_{v,\text{test}}} \mathbf{1}(f(x) \neq y) \right] \\ &= \mathbb{E}_V \left[\inf_f \int_{x=0}^1 \sum_{y \in \{-1,1\}} P_{\text{test}}(x) P_{V,\text{test}}(y | x) \mathbf{1}(f(x) = -y) \right] \\ &= \mathbb{E}_V \left[\int_{x=0}^1 P_{\text{test}}(x) \min_{y \in \{-1,1\}} P_{V,\text{test}}(y | x) \right] \\ &\stackrel{(i)}{=} \mathbb{E}_V \left[\frac{1}{2} \left(1 - \int_{x=0}^1 P_{\text{test}}(x) |P_{V,\text{test}}(1 | x) - P_{V,\text{test}}(-1 | x)| dx \right) \right] \\ &\stackrel{(ii)}{=} \mathbb{E}_V \left[\frac{1}{2} \left(1 - \int_{x=0}^1 P_{\text{test}}(x) \left| \phi \left(x - \frac{j_x + 1/2}{K} \right) \right| dx \right) \right] \\ &= \frac{1}{2} - \frac{1}{2} \int_{x=0}^1 P_{\text{test}}(x) \left| \phi \left(x - \frac{j_x + 1/2}{K} \right) \right| dx, \end{aligned} \quad (18)$$

where (i) follows since $P_v(1 | x) = 1 - P_v(-1 | x)$ and $\min\{s, 1 - s\} = (1 - |1 - 2s|)/2$ for all $s \in [0, 1]$, and (ii) follows by our construction of P_v above along with the fact that $P_v(1 | x) = 1 - P_v(-1 | x)$.

Putting things together: Combining Eqs. (17) and (18) allows us to conclude that

$$\begin{aligned} & \text{Minimax Excess Risk}(\mathcal{P}_{\text{GS}}(\tau)) \\ & \geq \frac{1}{4} \int_0^1 P_{\text{test}}(x) \left| \phi \left(x - \frac{j_x + 1/2}{K} \right) \right| \exp(-\text{KL}(Q(S | v_{j_x} = 1) \| Q(S | v_{j_x} = -1))) dx \\ & = \frac{1}{4} \sum_{j=1}^K \int_{\frac{j-1}{K}}^{\frac{j}{K}} P_{\text{test}}(x) \left| \phi \left(x - \frac{j + 1/2}{K} \right) \right| \exp(-\text{KL}(Q(S | v_j = 1) \| Q(S | v_j = -1))) dx \\ & = \frac{1}{4} \sum_{j=1}^K \exp(-\text{KL}(Q(S | v_j = 1) \| Q(S | v_j = -1))) \left[\int_{\frac{j-1}{K}}^{\frac{j}{K}} P_{\text{test}}(x) \left| \phi \left(x - \frac{j + 1/2}{K} \right) \right| dx \right] \\ & \stackrel{(i)}{=} \frac{1}{32K^2} \sum_{j=1}^K \exp(-\text{KL}(Q(S | v_j = 1) \| Q(S | v_j = -1))), \end{aligned}$$

where (i) follows by using Lemma A.2 along with the fact that $P_{\text{test}}(x) = 1$ in our construction to show that the integral in the square brackets is equal to $1/8K^2$. This proves the result. \square

The next lemma upper bounds the KL divergence between $\mathbb{Q}(S \mid v_j = 1)$ and $\mathbb{Q}(S \mid v_j = -1)$ for each $j \in \{1, \dots, K\}$. It shows that the KL divergence between these two posteriors is larger when the expected number of samples in that bin is larger.

Lemma C.2. *Suppose that v is drawn uniformly from the set $\{-1, 1\}^K$, and that $S \mid v$ is drawn from $\mathbb{P}_{v, \text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{v, \text{min}}^{n_{\text{min}}}$. Then for any $j \in \{1, \dots, K/2\}$ and any $\tau \in [0, 1]$,*

$$\text{KL}(\mathbb{Q}(S \mid v_j = 1) \parallel \mathbb{Q}(S \mid v_j = -1)) \leq \frac{n_{\text{maj}}(2 - \tau) + n_{\text{min}}\tau}{3K^3},$$

and for any $j \in \{K/2 + 1, \dots, K\}$

$$\text{KL}(\mathbb{Q}(S \mid v_j = 1) \parallel \mathbb{Q}(S \mid v_j = -1)) \leq \frac{n_{\text{maj}}\tau + n_{\text{min}}(2 - \tau)}{3K^3}.$$

Proof. Let us consider the case when $j = 1$. The bound for all other $j \in \{2, \dots, K\}$ shall follow analogously.

Given samples S , let $S = (S_1, \bar{S}_1)$ be a partition where S_1 are the samples that fall in the interval I_1 , and \bar{S}_1 be the other samples. Similarly, given a vector $v \in \{-1, 1\}$, let $v = (v_1, \bar{v}_1)$, where v_1 is the first component and \bar{v}_1 denotes the other components $(2, \dots, K)$ of v .

First, we will show that

$$\mathbb{Q}(S \mid v_1) = \mathbb{Q}(S_1 \mid v_1)\mathbb{Q}(\bar{S}_1).$$

To see this, observe that

$$\mathbb{Q}(S \mid v_1) = \mathbb{Q}((S_1, \bar{S}_1) \mid v_1) = \mathbb{Q}(S_1 \mid v_1)\mathbb{Q}(\bar{S}_1 \mid v_1, S_1).$$

Further, if v is chosen uniformly over the hypercube $\{-1, 1\}^K$, then

$$\begin{aligned} \mathbb{Q}(\bar{S}_1 \mid v_1, S_1) &= \sum_{\bar{v}_1} \mathbb{Q}(\bar{S}_1, \bar{v}_1 \mid v_1, S_1) \\ &= \sum_{\bar{v}_1} \mathbb{Q}(\bar{S}_1 \mid v_1, \bar{v}_1, S_1)\mathbb{Q}(\bar{v}_1 \mid v_1, S_1) \\ &\stackrel{(i)}{=} \sum_{\bar{v}_1} \mathbb{Q}(\bar{S}_1 \mid v_1, \bar{v}_1, S_1)\mathbb{Q}(\bar{v}_1) \\ &\stackrel{(ii)}{=} \sum_{\bar{v}_1} \mathbb{Q}(\bar{S}_1 \mid v_1, \bar{v}_1)\mathbb{Q}(\bar{v}_1) \\ &\stackrel{(iii)}{=} \sum_{\bar{v}_1} \mathbb{Q}(\bar{S}_1 \mid \bar{v}_1)\mathbb{Q}(\bar{v}_1) \\ &= \mathbb{Q}(\bar{S}_1), \end{aligned}$$

where (i) follows since by Bayes' rule

$$\begin{aligned} \mathbb{Q}(\bar{v}_1 \mid v_1, S_1) &= \frac{\mathbb{Q}(\bar{v}_1 \mid v_1)\mathbb{Q}(S_1 \mid v_1, \bar{v}_1)}{\mathbb{Q}(S_1 \mid v_1)} \\ &= \frac{\mathbb{Q}(\bar{v}_1)\mathbb{Q}(S_1 \mid v_1, \bar{v}_1)}{\mathbb{Q}(S_1 \mid v_1)} && \text{(since } \bar{v}_1 \text{ is independent of } v_1\text{)} \\ &= \frac{\mathbb{Q}(\bar{v}_1)\mathbb{Q}(S_1 \mid v_1)}{\mathbb{Q}(S_1 \mid v_1)} = \mathbb{Q}(\bar{v}_1) && \text{(the samples in } S_1 \text{ depend only on } v_1\text{)}. \end{aligned}$$

Inequality (ii) follows since the samples are drawn independently given $v = (v_1, \bar{v}_1)$. Finally, (iii) follows since \bar{S}_1 (the samples that lie outside the interval I_1) only depend on \bar{v}_1 since the marginal distribution of x is independent of v and the distribution of $y \mid x$ depends only on the value of v corresponding to the interval in which x lies.

Thus since, $\mathbb{Q}(S | v_1) = \mathbb{Q}(S_1 | v_1)\mathbb{Q}(\bar{S}_1)$ we have that

$$\text{KL}(\mathbb{Q}(S | v_1 = 1) \| \mathbb{Q}(S | v_1 = -1)) = \text{KL}(\mathbb{Q}(S_1 | v_1 = 1) \| \mathbb{Q}(S_1 | v_1 = -1)). \quad (19)$$

To bound this KL divergence, let us condition on the number of samples in S_1 from group a , (the majority group) $n_{1,a}$ and the number of samples from group b (the minority group), $n_{1,b}$. Now since $n_{1,a}$ and $n_{1,b}$ are independent of v_1 (which only affects the labels) we have that,

$$\begin{aligned} \mathbb{Q}(S_1 | v_1) &= \sum_{n_{1,a}, n_{1,b}} \mathbb{Q}(n_{1,a}, n_{1,b} | v_1) \mathbb{Q}(S_1 | v_1, n_{1,a}, n_{1,b}) \\ &= \sum_{n_{1,a}, n_{1,b}} \mathbb{Q}(n_{1,a}, n_{1,b}) \mathbb{Q}(S_1 | v_1, n_{1,a}, n_{1,b}) \\ &= \mathbb{E}_{n_{1,a}, n_{1,b}} [\mathbb{Q}(S_1 | v_1, n_{1,a}, n_{1,b})]. \end{aligned}$$

Therefore, by the joint convexity of the KL-divergence and by Jensen's inequality we have that,

$$\begin{aligned} &\text{KL}(\mathbb{Q}(S_1 | v_1 = 1) \| \mathbb{Q}(S_1 | v_1 = -1)) \\ &\leq \mathbb{E}_{n_{1,a}, n_{1,b}} [\text{KL}(\mathbb{Q}(S_1 | v_1 = 1, n_{1,a}, n_{1,b}) \| \mathbb{Q}(S_1 | v_1 = -1, n_{1,a}, n_{1,b}))]. \end{aligned} \quad (20)$$

Now conditioned on $v_1, n_{1,a}$ and $n_{1,b}$, samples in S_1 are composed of 2 groups of samples $(S_{1,a}, S_{1,b})$. The samples in each group $(S_{1,a}, S_{1,b})$ are drawn independently from the distributions $\mathbb{P}_a(x | x \in I_1)\mathbb{P}_v(y | x)$ and $\mathbb{P}_b(x | x \in I_1)\mathbb{P}_v(y | x)$ respectively. Therefore,

$$\begin{aligned} &\text{KL}(\mathbb{Q}(S_1 | v_1 = 1, n_{1,a}, n_{1,b}) \| \mathbb{Q}(S_1 | v_1 = -1, n_{1,a}, n_{1,b})) \\ &\stackrel{(i)}{=} n_{1,a} \text{KL}(\mathbb{P}_a(x | x \in I_1) \mathbb{P}_{v_1=1}(y | x) \| \mathbb{P}_a(x | x \in I_1) \mathbb{P}_{v_1=-1}(y | x)) \\ &\quad + n_{1,b} \text{KL}(\mathbb{P}_b(x | x \in I_1) \mathbb{P}_{v_1=1}(y | x) \| \mathbb{P}_b(x | x \in I_1) \mathbb{P}_{v_1=-1}(y | x)) \\ &\stackrel{(ii)}{=} (n_{1,a} + n_{1,b}) \mathbb{E}_{x \sim \text{Unif}(I_1)} [\text{KL}(\mathbb{P}_{v_1=1}(y | x) \| \mathbb{P}_{v_1=-1}(y | x))] \\ &\stackrel{(iii)}{=} \frac{n_{1,a} + n_{1,b}}{2} \mathbb{E}_{x \sim \text{Unif}(I_1)} \left[\sum_{y \in \{-1, 1\}} \left(1 + y \phi \left(x - \frac{1}{2K} \right) \right) \log \left(\frac{(1 + y \phi(x - \frac{1}{2K}))}{(1 + y \phi(x - \frac{1}{2K}))} \right) \right] \\ &= \frac{n_{1,a} + n_{1,b}}{2} \sum_{y \in \{-1, 1\}} \mathbb{E}_{x \sim \text{Unif}(I_1)} \left[\left(1 + y \phi \left(x - \frac{1}{2K} \right) \right) \log \left(\frac{(1 + y \phi(x - \frac{1}{2K}))}{(1 + y \phi(x - \frac{1}{2K}))} \right) \right] \\ &= \frac{n_{1,a} + n_{1,b}}{2K} \sum_{y \in \{-1, 1\}} \int_{x=0}^{\frac{1}{K}} \left[\left(1 + y \phi \left(x - \frac{1}{2K} \right) \right) \log \left(\frac{(1 + y \phi(x - \frac{1}{2K}))}{(1 + y \phi(x - \frac{1}{2K}))} \right) \right] dx \\ &\stackrel{(iv)}{\leq} \frac{n_{1,a} + n_{1,b}}{3K^2}, \end{aligned} \quad (21)$$

where in (i) we let \mathbb{P}_{v_1} denote the conditional distribution of y for $x \in I_1$ given v_1 , (ii) follows since both \mathbb{P}_a and \mathbb{P}_b are constant in the interval, (iii) follows by our construction of \mathbb{P}_v above, and finally (iv) follows by invoking Lemma A.3 that ensures that the integral is bounded by $1/3K^2$.

Using this bound in Eq. (20), along with Eq. (19) we get that

$$\text{KL}(\mathbb{Q}(S | v_1 = 1) \| \mathbb{Q}(S | v_1 = -1)) \leq \frac{\mathbb{E}[n_{1,a} + n_{1,b}]}{3K^2}.$$

Now there are n_{maj} samples from group a in S and n_{min} samples from group b . Therefore,

$$\begin{aligned} \mathbb{E}[n_{1,a}] &= n_{\text{maj}} \mathbb{P}_a(x \in I_1) = \frac{n_{\text{maj}}(2 - \tau)}{K}, \\ \mathbb{E}[n_{1,b}] &= n_{\text{min}} \mathbb{P}_b(x \in I_1) = \frac{n_{\text{min}}\tau}{K}. \end{aligned}$$

Plugging this bound into Eq. (21) completes the proof by the first interval. An identical argument holds for $j \in \{2, \dots, K/2\}$. For $j \in \{K/2 + 1, \dots, K\}$ the only change is that

$$\begin{aligned}\mathbb{E}[n_{j,a}] &= n_{\text{maj}} \mathbb{P}_a(x \in I_j) = \frac{n_{\text{maj}} \tau}{K}, \\ \mathbb{E}[n_{j,b}] &= n_{\text{min}} \mathbb{P}_b(x \in I_j) = \frac{n_{\text{min}}(2 - \tau)}{K}.\end{aligned}$$

□

Next, we combine the previous two lemmas to establish our stated lower bound. We first restate it here.

Theorem 4.2. *Consider the group shift setting described in Section 3.2.2. Given any overlap $\tau \in [0, 1]$ recall that $\mathcal{P}_{\text{GS}}(\tau)$ is the class of distributions such that $\text{Overlap}(\mathbb{P}_{\text{maj}}, \mathbb{P}_{\text{min}}) \geq \tau$. The minimax excess risk in this setting is lower bounded as follows:*

$$\begin{aligned}\text{Minimax Excess Risk}(\mathcal{P}_{\text{GS}}(\tau)) &= \inf_{\mathcal{A}} \sup_{(\mathbb{P}_{\text{maj}}, \mathbb{P}_{\text{min}}) \in \mathcal{P}_{\text{GS}}(\tau)} \text{Excess Risk}[\mathcal{A}; (\mathbb{P}_{\text{maj}}, \mathbb{P}_{\text{min}})] \\ &\geq \frac{1}{200(n_{\text{min}} \cdot (2 - \tau) + n_{\text{maj}} \cdot \tau)^{1/3}} \geq \frac{1}{200n_{\text{min}}^{1/3}(\rho \cdot \tau + 2)^{1/3}},\end{aligned}\quad (4)$$

where $\rho = n_{\text{maj}}/n_{\text{min}} > 1$.

Proof. First, by Lemma C.1 we know that

$$\text{Minimax Excess Risk}(\mathcal{P}_{\text{GS}}(\tau)) \geq \frac{1}{32K^2} \sum_{j=1}^K \exp(-\text{KL}(\mathbb{Q}(S | v_j = 1) \| \mathbb{Q}(S | v_j = -1))).$$

Next, by invoking the bound on the KL divergences in the equation above by Lemma C.2 we get that

$$\begin{aligned}\text{Minimax Excess Risk}(\mathcal{P}_{\text{GS}}(\tau)) &\geq \frac{1}{64K} \left[\exp\left(-\frac{n_{\text{maj}}(2 - \tau) + n_{\text{min}}\tau}{3K^3}\right) + \exp\left(-\frac{n_{\text{min}}(2 - \tau) + n_{\text{maj}}\tau}{3K^3}\right) \right] \\ &\geq \frac{1}{64K} \left[\exp\left(-\frac{n_{\text{min}}(2 - \tau) + n_{\text{maj}}\tau}{3K^3}\right) \right]\end{aligned}$$

Setting $K = \lceil (n_{\text{min}}(2 - \tau) + n_{\text{maj}}\tau)^{1/3} \rceil$ and recalling that $\tau \leq 1$ we get that

$$\begin{aligned}\text{Minimax Excess Risk}(\mathcal{P}_{\text{GS}}(\tau)) &\geq \frac{1}{64 \lceil (n_{\text{min}}(2 - \tau) + n_{\text{maj}}\tau)^{1/3} \rceil} \left[\exp\left(-\frac{n_{\text{min}}(2 - \tau) + n_{\text{maj}}\tau}{3 \lceil (n_{\text{min}}(2 - \tau) + n_{\text{maj}}\tau)^{1/3} \rceil^3}\right) \right] \\ &\stackrel{(i)}{\geq} \frac{\exp(-1/3)}{64} \frac{(n_{\text{min}}(2 - \tau) + n_{\text{maj}}\tau)^{1/3}}{\lceil (n_{\text{min}}(2 - \tau) + n_{\text{maj}}\tau)^{1/3} \rceil} \frac{1}{(n_{\text{min}}(2 - \tau) + n_{\text{maj}}\tau)^{1/3}} \\ &\stackrel{(ii)}{\geq} \frac{0.7 \exp(-1/3)}{64} \frac{1}{(n_{\text{min}}(2 - \tau) + n_{\text{maj}}\tau)^{1/3}} \\ &\geq \frac{1}{200} \frac{1}{(n_{\text{min}}(2 - \tau) + n_{\text{maj}}\tau)^{1/3}},\end{aligned}$$

where (i) follows since $n_{\text{min}}(2 - \tau) + n_{\text{maj}}\tau / \lceil (n_{\text{min}}(2 - \tau) + n_{\text{maj}}\tau)^{1/3} \rceil^3 \leq 1$, and (ii) follows since $0 \leq \tau \leq 1$ and $n_{\text{min}} \geq 1$ and hence $\frac{(n_{\text{min}}(2 - \tau) + n_{\text{maj}}\tau)^{1/3}}{\lceil (n_{\text{min}}(2 - \tau) + n_{\text{maj}}\tau)^{1/3} \rceil} \geq 0.7$. □

C.2 Proof of Theorem 5.2

In this section, we derive an upper bound on the excess risk of the undersampled binning estimator \mathcal{A}_{USB} (Eq. (5)). Recall that given a dataset \mathcal{S} this estimator first calculates the undersampled dataset \mathcal{S}_{US} , where the number of points from the minority group (n_{min}) is equal to the number of points from the majority group (n_{min}), and the size of the dataset is $2n_{\text{min}}$. Throughout this section, $(\mathbb{P}_{\text{maj}}, \mathbb{P}_{\text{min}})$ shall be an arbitrary element of $\mathcal{P}_{\text{GS}}(\tau)$ for any $\tau \in [0, 1]$. In this section, whenever we shall often denote $\text{Excess Risk}(\mathcal{A}; (\mathbb{P}_{\text{maj}}, \mathbb{P}_{\text{min}}))$ by simply $\text{Excess Risk}(\mathcal{A})$.

Before we proceed, we introduce some additional notation. For any $j \in \{1, \dots, K\}$ and $I_j = [\frac{j-1}{K}, \frac{j}{K}]$ let

$$q_{j,1} := \mathbb{P}_{\text{test}}(y = 1 \mid x \in I_j) = \int_{x \in I_j} \mathbb{P}(y = 1 \mid x) \mathbb{P}_{\text{test}}(x \mid x \in I_j) \, dx, \quad (22a)$$

$$q_{j,-1} := \mathbb{P}_{\text{test}}(y = -1 \mid x \in I_j) = \int_{x \in I_j} \mathbb{P}(y = -1 \mid x) \mathbb{P}_{\text{test}}(x \mid x \in I_j) \, dx. \quad (22b)$$

For the undersampled binning estimator \mathcal{A}_{USB} (defined above in Eq. (5)), define the *excess risk in an interval* I_j as follows:

$$\begin{aligned} R_j(\mathcal{A}_{\text{USB}}^{\mathcal{S}}) &:= p(y = -\mathcal{A}_j^{\mathcal{S}} \mid x \in I_j) - \min\{\mathbb{P}_{\text{test}}(y = 1 \mid x \in I_j), \mathbb{P}_{\text{test}}(y = -1 \mid x \in I_j)\} \\ &= q_{j,-\mathcal{A}_j^{\mathcal{S}}} - \min\{q_{j,1}, q_{j,-1}\}. \end{aligned}$$

The proof of the upper bound shall proceed in steps. First, in Lemma C.3 we will show that the excess risk is equal to sum the excess risk over the intervals up to a factor of $2/K$ on account of the distribution being 1-Lipschitz. Next, in Lemma C.4 we upper bound the risk over each interval. We put these two together and to upper bound the risk.

Lemma C.3. *The expected excess risk of undersampled binning estimator \mathcal{A}_{USB} can be decomposed as follows*

$$\text{Excess Risk}(\mathcal{A}_{\text{USB}}) \leq \sum_{j=0}^{K-1} \mathbb{E}_{\mathcal{S} \sim \mathbb{P}_{\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\text{min}}^{n_{\text{min}}}} [R_j(\mathcal{A}_{\text{USB}}^{\mathcal{S}})] \cdot \mathbb{P}_{\text{test}}(I_j) + \frac{2}{K},$$

where $\mathbb{P}_{\text{test}}(I_j) := \int_{x \in I_j} \mathbb{P}_{\text{test}}(x) \, dx$.

Proof. Recall that by definition, the expected excess risk is

$$\mathbb{E}_{\mathcal{S} \sim \mathbb{P}_{\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\text{min}}^{n_{\text{min}}}} [R(\mathcal{A}^{\mathcal{S}}; \mathbb{P}_{\text{test}}) - R(f^*; \mathbb{P}_{\text{test}})].$$

Let us first decompose the Bayes risk $R(f^*)$,

$$\begin{aligned} R(f^*) &= \inf_f \mathbb{E}_{(x,y) \sim \mathbb{P}_{\text{test}}} [\mathbf{1}(f(x) \neq y)] \\ &= \inf_f \int_{x=0}^1 \sum_{y \in \{-1,1\}} \mathbf{1}(f(x) \neq y) \mathbb{P}_{\text{test}}(y \mid x) \mathbb{P}_{\text{test}}(x) \, dx \\ &= \int_{x=0}^1 \inf_{f(x) \in \{-1,1\}} \sum_{y \in \{-1,1\}} \mathbf{1}(f(x) \neq y) \mathbb{P}_{\text{test}}(y \mid x) \mathbb{P}_{\text{test}}(x) \, dx \\ &= \int_{x=0}^1 \inf_{f(x) \in \{-1,1\}} \mathbb{P}_{\text{test}}(y = -f(x) \mid x) \mathbb{P}_{\text{test}}(x) \, dx \\ &= \int_{x=0}^1 \min\{\mathbb{P}_{\text{test}}(y = 1 \mid x), \mathbb{P}_{\text{test}}(y = -1 \mid x)\} \mathbb{P}_{\text{test}}(x) \, dx. \end{aligned} \quad (23)$$

The risk of the undersampled binning algorithm \mathcal{A}_{USB} is given by

$$\begin{aligned} R(\mathcal{A}_{\text{USB}}^{\mathcal{S}}) &= \int_{x=0}^1 \sum_{y \in \{-1,1\}} \mathbf{1}(\mathcal{A}_{\text{USB}}^{\mathcal{S}}(x) \neq y) \mathbf{P}_{\text{test}}(y | x) \mathbf{P}_{\text{test}}(x) \, dx \\ &= \int_{x=0}^1 \mathbf{P}_{\text{test}}(y = -\mathcal{A}_{\text{USB}}^{\mathcal{S}}(x) | x) \mathbf{P}_{\text{test}}(x) \, dx. \end{aligned}$$

Next, recall that the undersampled binning estimator is constant over the intervals I_j for $j \in \{1, \dots, K\}$ where it takes the value $\mathcal{A}_j^{\mathcal{S}}$ (to ease notation let us simply denote it by \mathcal{A}_j below), and therefore

$$R(\mathcal{A}_{\text{USB}}^{\mathcal{S}}) = \sum_{j=0}^{K-1} \int_{x \in I_j} \mathbf{P}_{\text{test}}(y = -\mathcal{A}_j | x) \mathbf{P}_{\text{test}}(x) \, dx.$$

This combined with Eq. (23) tells us that

$$\begin{aligned} &R(\mathcal{A}_{\text{USB}}^{\mathcal{S}}) - R(f^*) \\ &= \sum_{j=0}^{K-1} \int_{x \in I_j} (\mathbf{P}_{\text{test}}(y = -\mathcal{A}_j | x) - \min\{\mathbf{P}_{\text{test}}(y = 1 | x), \mathbf{P}_{\text{test}}(y = -1 | x)\}) \mathbf{P}_{\text{test}}(x) \, dx. \end{aligned} \quad (24)$$

Recall the definition of $q_{j,1}$ and $q_{j,-1}$ from Eqs. (22a)-(22b) above. For any $x \in I_j = [\frac{j-1}{K}, \frac{j}{K}]$, $|\mathbf{P}_{\text{test}}(y | x) - q_{j,y}| \leq 1/K$, since the distribution $\mathbf{P}_{\text{test}}(y | x)$ is 1-Lipschitz and $q_{j,y}$ is its conditional mean. Therefore,

$$\begin{aligned} &R(\mathcal{A}_{\text{USB}}^{\mathcal{S}}) - R(f^*) \\ &\leq \sum_{j=0}^{K-1} \int_{x \in I_j} (q_{j,-\mathcal{A}_j} - \min\{q_{j,1}, q_{j,-1}\}) \mathbf{P}_{\text{test}}(x) \, dx + \frac{2}{K} \sum_{j=0}^{K-1} \int_{x \in I_j} \mathbf{P}_{\text{test}}(x) \, dx \\ &= \sum_{j=0}^{K-1} \int_{x \in I_j} R_j(\mathcal{A}_{\text{USB}}^{\mathcal{S}}) \mathbf{P}_{\text{test}}(x) \, dx + \frac{2}{K}. \end{aligned}$$

Taking expectation over the training samples \mathcal{S} (where n_{\min} samples are drawn independently from \mathbf{P}_{\min} and n_{maj} samples are drawn independently from \mathbf{P}_{maj}) concludes the proof. \square

Next we provide an upper bound on the expected excess risk is an interval $R_j(\mathcal{A}_{\text{USB}}^{\mathcal{S}})$.

Lemma C.4. For any $j \in \{1, \dots, K\}$ with $I_j = [\frac{j-1}{K}, \frac{j}{K}]$,

$$\mathbb{E}_{\mathcal{S} \sim \mathbf{P}_{\text{maj}}^{n_{\text{maj}}} \times \mathbf{P}_{\min}^{n_{\min}}} [R_j(\mathcal{A}_{\text{USB}}^{\mathcal{S}})] \leq \frac{c}{\sqrt{n_{\min} \mathbf{P}_{\text{test}}(I_j)}} + \frac{c}{K},$$

where c is an absolute constant, and $\mathbf{P}_{\text{test}}(I_j) := \int_{x \in I_j} \mathbf{P}_{\text{test}}(x) \, dx$.

Proof. Consider an arbitrary bucket $j \in \{1, \dots, K\}$.

Let us introduce some notation that shall be useful in the remainder of the proof. Analogous to $q_{j,1}$ and $q_{j,-1}$ defined above (see Eqs. (22a)-(22b)), define $q_{j,1}^a$ and $q_{j,1}^b$ as follows:

$$q_{j,1}^a := \mathbf{P}_a(y = 1 | x \in I_j) = \int_{x \in I_j} \mathbf{P}(y = 1 | x) \mathbf{P}_a(x | x \in I_j) \, dx, \quad (25a)$$

$$q_{j,1}^b := \mathbf{P}_b(y = 1 | x \in I_j) = \int_{x \in I_j} \mathbf{P}(y = 1 | x) \mathbf{P}_b(x | x \in I_j) \, dx. \quad (25b)$$

Essentially, $q_{j,1}^a$ is the probability that a sample is from group a and has label 1, conditioned on the event that the sample falls in the interval I_j . Since

$$\mathbf{P}_{\text{test}}(x | x \in I_j) = \frac{1}{2} [\mathbf{P}_a(x | x \in I_j) + \mathbf{P}_b(x | x \in I_j)],$$

therefore

$$\begin{aligned} |q_{j,1} - q_{j,1}^a| &= \left| \int_{x \in I_j} \mathbb{P}(y = 1 | x) \mathbb{P}_{\text{test}}(x | x \in I_j) dx - \int_{x \in I_j} \mathbb{P}(y = 1 | x) \mathbb{P}_a(x | x \in I_j) dx \right| \\ &\leq \frac{1}{K}. \end{aligned} \quad (26)$$

This follows since $\mathbb{P}(y | x)$ is 1-Lipschitz and therefore can fluctuate by at most $1/K$ in the interval I_j . Of course the same bound also holds for $|q_{j,1} - q_{j,1}^b|$.

With this notation in place let us present a bound on the expected value of $R_j(\mathcal{A}_{\text{USB}}^S)$. By definition

$$R_j(\mathcal{A}_{\text{USB}}^S) = q_{j,-\mathcal{A}_j^S} - \min\{q_{j,1}, q_{j,-1}\}.$$

First, note that $q_{j,1} := \mathbb{P}_{\text{test}}(y = 1 | x \in I_j) = 1 - q_{j,-1}$. Suppose that $q_{j,1} < 1/2$ and therefore $q_{j,-1} > 1/2$ (the same bound shall hold in the other case). In this case, risk is incurred only when $\mathcal{A}_j^S = 1$. That is,

$$\begin{aligned} \mathbb{E}_{\mathcal{S} \sim \mathbb{P}_{\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\text{min}}^{n_{\text{min}}}} [R_j(\mathcal{A}_{\text{USB}}^S)] &= |q_{j,-1} - q_{j,1}| \Pr[\mathcal{A}_j^S = 1] \\ &= |1 - 2q_{j,1}| \Pr[\mathcal{A}_j^S = 1]. \end{aligned} \quad (27)$$

Now by the definition of the undersampled binning estimator (see Eq. (5)), $\mathcal{A}_j^S = 1$ only when there are more samples in the interval I_j with label 1 than -1 . However, we can bound the probability of this happening since $q_{j,1}$ is smaller than $q_{j,-1}$.

Let n_j be the number of samples in the undersampled sample set \mathcal{S}_{US} in the interval I_j . Let $n_{1,j}$ be the number of these samples with label 1, and $n_{-1,j} = n_j - n_{1,j}$ be the number of samples with label -1 . Further, let $n_{a,j}$ be the number of samples in from group a such that they fall in the interval I_j , and define $m_{b,j}$ analogously.

The probability of incurring risk is given by

$$\mathbb{P}[\mathcal{A}_j = 1] = \sum_{s=1}^{2n_{\text{min}}} \mathbb{P}[\mathcal{A}_j = 1 | n_j = s] \mathbb{P}[n_j = s], \quad (28)$$

where the sum is up to $2n_{\text{min}}$ since the size of the undersample dataset $|\mathcal{S}_{\text{US}}|$ is equal to $2n_{\text{min}}$.

Conditioned on the event that $n_j = s$ the probability of incurring risk is

$$\begin{aligned} \mathbb{P}[\mathcal{A}_j = 1 | n_j = s] &= \mathbb{P}[m_{1,j} > n_{-1,j} | n_j = s] = \mathbb{P}[n_{1,j} > n_j/2 | n_j = s] \\ &= \mathbb{P}[n_{1,j} > s/2 | n_j = s]. \end{aligned} \quad (29)$$

Now, note that $n_j = n_{a,j} + n_{b,j}$. Thus continuing, we have that

$$\begin{aligned} \mathbb{P}[n_{1,j} > s/2 | n_j = s] &= \sum_{s' \leq s} \mathbb{P}[n_{1,j} > s/2 | n_j = s, n_{b,j} = s'] \mathbb{P}[n_{b,j} = s'] \\ &= \sum_{s' \leq s} \mathbb{P}[n_{1,j} > s/2 | n_{a,j} = s - s', n_{b,j} = s'] \mathbb{P}[n_{b,j} = s']. \end{aligned}$$

In light of this previous equation, we want to control the probability that the number of samples with label 1 in the interval I_j conditioned on the event that the number of samples from group a in this interval is $s - s'$ and the number of samples from group b in this interval is s' . Recall that $q_{j,1}^a$ and $q_{j,1}^b$ the probabilities of the label of the sample being 1 conditioned the event that sample is in the interval I_j when it is group a and b respectively. So we define the random variables:

$$z_a[s - s'] \sim \text{Bin}(s - s', q_{j,1}^a), \quad z_b[s'] \sim \text{Bin}(s', q_{j,1}^b), \quad z[s] \sim \text{Bin}(s, \max\{q_{j,1}^a, q_{j,1}^b\}).$$

Then,

$$\begin{aligned}
& \mathbb{P}[n_{1,j} > s/2 \mid n_j = s] \\
&= \sum_{s' \leq s} \mathbb{P}[n_{1,j} > s/2 \mid n_{j,a} = s - s', n_{j,b} = s'] \mathbb{P}[n_{j,b} = s'] \\
&= \sum_{s' \leq s} \mathbb{P}[z_a[s - s'] + z_b[s'] > s/2 \mid n_{a,j} = s - s', n_{b,j} = s'] \mathbb{P}[n_{b,j} = s'] \\
&\leq \sum_{s' \leq s} \mathbb{P}[z[s] > s/2 \mid n_{a,j} = s - s', n_{b,j} = s'] \mathbb{P}[n_{b,j} = s'] \\
&= \sum_{s' \leq s} \mathbb{P}[z[s] > s/2] \mathbb{P}[n_{b,j} = s'] \\
&= \mathbb{P}[z[s] > s/2] \\
&\stackrel{(i)}{\leq} \exp\left(-\frac{s}{2}(1 - 2 \max\{q_{j,1}^a, q_{j,1}^b\})^2\right), \tag{30}
\end{aligned}$$

where (i) follows by invoking Hoeffding's inequality (Wainwright, 2019, Proposition 2.5). Combining this with Eqs. (28) and (29) we get that

$$\mathbb{P}[\mathcal{A}_j = 1] \leq \sum_{s=1}^{2n_{\min}} \exp\left(-\frac{s}{2}(1 - 2 \max\{q_{j,1}^a, q_{j,1}^b\})^2\right) \mathbb{P}[n_j = s].$$

Now n_j , which is the number of samples that lands in the interval I_j is equal to $n_{a,j} + n_{b,j}$. Now each of $n_{a,j}$ and $n_{b,j}$ (the number of samples in this interval from each of the groups) are random variables with distributions $\text{Bin}(n_{\min}, P_a(I_j))$ and $\text{Bin}(n_{\min}, P_b(I_j))$, where $P_a(I_j) = \int_{x \in I_j} P_a(x) dx$ and $P_b(I_j) = \int_{x \in I_j} P_b(x) dx$. Therefore, n_j is distributed as a sum of two binomial distribution and is therefore Poisson binomially distributed (Wikipedia contributors, 2022). Using the formula for the moment generating function (MGF) of a Poisson binomially distributed random variable we infer that,

$$\begin{aligned}
\mathbb{P}[\mathcal{A}_j = 1] &\leq \left(1 - P_a(I_j) + P_a(I_j) \exp\left(-\frac{(1 - 2 \max\{q_{j,1}^a, q_{j,1}^b\})^2}{2}\right)\right)^{n_{\min}} \times \\
&\quad \left(1 - P_b(I_j) + P_b(I_j) \exp\left(-\frac{(1 - 2 \max\{q_{j,1}^a, q_{j,1}^b\})^2}{2}\right)\right)^{n_{\min}}.
\end{aligned}$$

Plugging this into Eq. (28) we get that,

$$\begin{aligned}
& \mathbb{E}_{\mathcal{S} \sim P_{\text{maj}}^{n_{\text{maj}}} \times P_{\text{min}}^{n_{\text{min}}}} [R_j(\mathcal{A}_{\text{USB}}^S)] \\
&\leq |1 - 2q_{j,1}| \left[1 - P_a(I_j) + P_a(I_j) \exp\left(-\frac{(1 - 2 \max\{q_{j,1}^a, q_{j,1}^b\})^2}{2}\right)\right]^{n_{\min}} \times \\
&\quad \left[1 - P_b(I_j) + P_b(I_j) \exp\left(-\frac{(1 - 2 \max\{q_{j,1}^a, q_{j,1}^b\})^2}{2}\right)\right]^{n_{\min}} \\
&= |1 - 2q_{j,1}| \left[1 - P_a(I_j) \left(1 - \exp\left(-\frac{(1 - 2 \max\{q_{j,1}^a, q_{j,1}^b\})^2}{2}\right)\right)\right]^{n_{\min}} \times \\
&\quad \left[1 - P_b(I_j) \left(1 - \exp\left(-\frac{(1 - 2 \max\{q_{j,1}^a, q_{j,1}^b\})^2}{2}\right)\right)\right]^{n_{\min}}.
\end{aligned}$$

Since $|1 - 2 \max\{q_{j,1}^a, q_{j,1}^b\}| \leq 1$,

$$1 - \exp\left(-\frac{(1 - 2 \max\{q_{j,1}^a, q_{j,1}^b\})^2}{2}\right) \geq \frac{(1 - 2 \max\{q_{j,1}^a, q_{j,1}^b\})^2}{4},$$

and therefore

$$\begin{aligned}
\mathbb{E}_{\mathcal{S} \sim \mathbb{P}_{\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\text{min}}^{n_{\text{min}}}} [R_j(\mathcal{A}_{\text{USB}}^S)] &\leq |1 - 2q_{j,1}| \left[1 - \mathbb{P}_a(I_j) \frac{(1 - 2 \max\{q_{j,1}^a, q_{j,1}^b\})^2}{2} \right]^{n_{\text{min}}} \times \\
&\quad \left[1 - \mathbb{P}_b(I_j) \frac{(1 - 2 \max\{q_{j,1}^a, q_{j,1}^b\})^2}{2} \right]^{n_{\text{min}}} \\
&\stackrel{(i)}{\leq} |1 - 2q_{j,1}| \left[1 - \mathbb{P}_a(I_j) \frac{(1 - 2q_{j,1} - 2\gamma)^2}{2} \right]^{n_{\text{min}}} \times \\
&\quad \left[1 - \mathbb{P}_b(I_j) \frac{(1 - 2q_{j,1} - 2\gamma)^2}{2} \right]^{n_{\text{min}}} \\
&\stackrel{(ii)}{\leq} |1 - 2q_{j,1}| \exp \left(-n_{\text{min}} (\mathbb{P}_a(I_j) + \mathbb{P}_b(I_j)) \frac{(1 - 2q_{j,1} - 2\gamma)^2}{2} \right),
\end{aligned}$$

where (i) follows since $|\max\{q_{j,1}^a, q_{j,1}^b\} - q_{j,1}| \leq 1/K$ by Eq. (26) and γ is such that $|\gamma| \leq 1/K$, and (ii) follows since $(1+z)^b \leq \exp(bz)$. Now the RHS above is maximized when $(1 - 2q_{j,1} - 2\gamma)^2 = \frac{c}{n_{\text{min}}(\mathbb{P}_a(I_j) + \mathbb{P}_b(I_j))}$, for some constant c . Plugging this into the equation above we get that

$$\begin{aligned}
\mathbb{E}_{\mathcal{S} \sim \mathbb{P}_{\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\text{min}}^{n_{\text{min}}}} [R_j(\mathcal{A}_{\text{USB}}^S)] &\leq \frac{c'}{\sqrt{n_{\text{min}}(\mathbb{P}_a(I_j) + \mathbb{P}_b(I_j))}} + c'|\gamma| \\
&\leq \frac{c'}{\sqrt{n_{\text{min}}(\mathbb{P}_a(I_j) + \mathbb{P}_b(I_j))}} + \frac{c'}{K}.
\end{aligned}$$

Finally, noting that $\mathbb{P}_{\text{test}}(I_j) = (\mathbb{P}_a(I_j) + \mathbb{P}_b(I_j))/2$ completes the proof. \square

By combining the previous two lemmas we can now prove our upper bound on the risk of the undersampled binning estimator. We begin by restating it.

Theorem 5.2. *Consider the group shift setting described in Section 3.2.2. For any overlap $\tau \in [0, 1]$ and for any $(\mathbb{P}_{\text{maj}}, \mathbb{P}_{\text{min}}) \in \mathcal{P}_{\text{GS}}(\tau)$ the expected excess risk of the Undersampling Binning Estimator (Eq. (5)) with number of bins with $K = \lceil n_{\text{min}}^{1/3} \rceil$ is*

$$\text{Excess Risk}[\mathcal{A}_{\text{USB}}; (\mathbb{P}_{\text{maj}}, \mathbb{P}_{\text{min}})] = \mathbb{E}_{\mathcal{S} \sim \mathbb{P}_{\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\text{min}}^{n_{\text{min}}}} [R(\mathcal{A}_{\text{USB}}^S; \mathbb{P}_{\text{test}})] - R(f^*; \mathbb{P}_{\text{test}}) \leq \frac{C}{n_{\text{min}}^{1/3}}.$$

Proof. First by Lemma C.3 we know that

$$\text{Excess Risk}[\mathcal{A}_{\text{USB}}] \leq \sum_{j=0}^{K-1} \mathbb{E}_{\mathcal{S} \sim \mathbb{P}_{\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\text{min}}^{n_{\text{min}}}} [R_j(\mathcal{A}_{\text{USB}}^S)] \cdot \mathbb{P}_{\text{test}}(I_j) + \frac{2}{K}.$$

Next by using the bound on $\mathbb{E}_{\mathcal{S} \sim \mathbb{P}_{\text{maj}}^{n_{\text{maj}}} \times \mathbb{P}_{\text{min}}^{n_{\text{min}}}} [R_j(\mathcal{A}_{\text{USB}}^S)]$ established in Lemma C.4 we get that,

$$\begin{aligned}
\text{Excess Risk}(\mathcal{A}_{\text{USB}}) &\leq c \sum_{j=0}^{K-1} \frac{1}{\sqrt{n_{\text{min}} \mathbb{P}_{\text{test}}(I_j)}} \mathbb{P}_{\text{test}}(I_j) + \frac{c}{K} \\
&= \frac{c}{\sqrt{n_{\text{min}}}} \sum_{j=0}^{K-1} \sqrt{\mathbb{P}_{\text{test}}(I_j)} + \frac{c}{K} \\
&\stackrel{(i)}{\leq} \frac{c}{\sqrt{n_{\text{min}}}} \sqrt{K} \sum_{j=0}^{K-1} \mathbb{P}_{\text{test}}(I_j) + \frac{c}{K} \\
&= c \sqrt{\frac{K}{n_{\text{min}}}} + \frac{c}{K}.
\end{aligned}$$

where (i) follows since for any vector $z \in \mathbb{R}^K$, $\|z\|_1 \leq \sqrt{K}\|z\|_2$. Maximizing over K yields the choice $K = \lceil n_{\min}^{1/3} \rceil$, completing the proof. \square

D Additional simulations

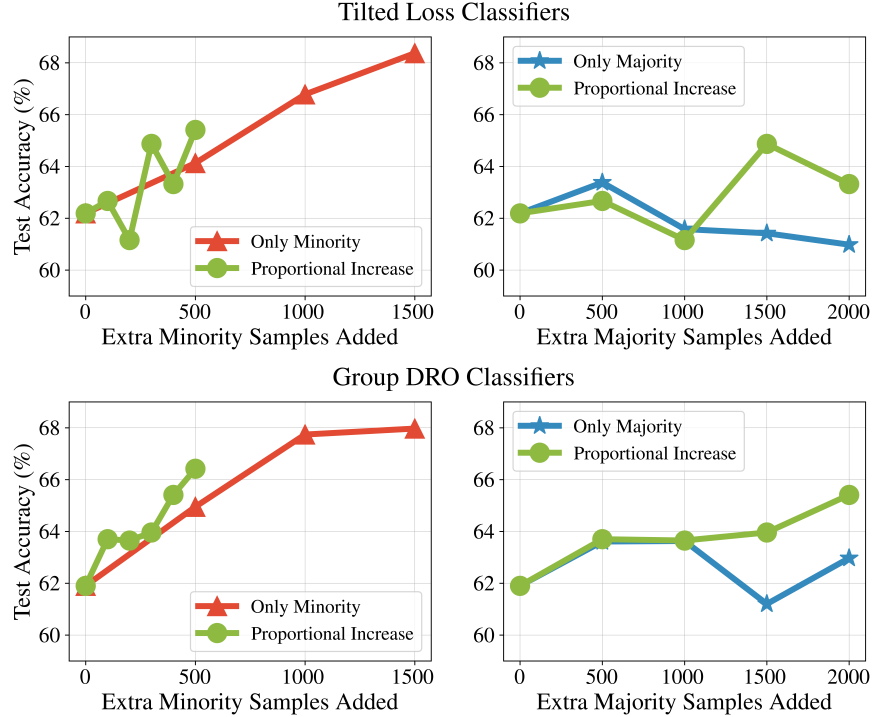


Figure 4: Convolutional neural network classifiers trained on the Imbalanced Binary CIFAR10 dataset with a 5:1 label imbalance. (Top) Models trained using the tilted loss (Li et al., 2020) with early stopping. (Bottom) Models trained using group-DRO (Sagawa et al., 2020) with early stopping. We report the average test accuracy calculated on a balanced test set over 5 random seeds. We start off with 2500 cat examples and 500 dog examples in the training dataset. We find similar trends to those obtained in Figure 2 even with these losses that are designed to optimize for the worst group accuracy.

E Discussion about minimax lower bounds for cost-sensitive losses applied to the label shift setting

We add a more detailed discussion about applying minimax cost-sensitive losses to obtain a lower bound in the presence of label shift.

Assume that P_{maj} is distribution of the covariates $x \mid y = 1$, and P_{min} is the distribution of the covariates $x \mid y = -1$. The training samples are drawn from the distribution:

$$P(x, y) = P(y = 1)P_{\text{maj}} + P(y = -1)P_{\text{min}},$$

where

$$P(y = 1) = \frac{\rho}{1 + \rho} \quad \text{and} \quad P(y = -1) = \frac{1}{1 + \rho}$$

for some imbalance ratio $\rho > 1$. On average the ratio between the number of points from the majority class to the number of points from the minority class is equal to ρ .

We set the cost of getting an incorrectly predicting the majority class label to be equal to

$$c_1 = \frac{1}{1 + \rho}$$

and the cost of incorrectly predicting the minority class label to be equal to

$$c_{-1} = \frac{\rho}{1 + \rho}.$$

Note that the costs $c_1 + c_{-1} = 1$ and that $c_1 < c_{-1}$.

The expected cost-sensitive loss is therefore equal to

$$\begin{aligned} \mathbb{E}_{(x,y) \sim \mathcal{P}} [c_y \mathbf{1}[f(x) \neq y]] &= \frac{\rho}{1 + \rho} \mathbb{E}_{(x) \sim \mathcal{P}} [c_1 \mathbf{1}[f(x) \neq 1]] + \frac{1}{1 + \rho} \mathbb{E}_{(x) \sim \mathcal{P}} [c_{-1} \mathbf{1}[f(x) \neq -1]] \\ &= \frac{\rho}{(1 + \rho)^2} \mathbb{E}_{x \sim \mathcal{P}_{\text{maj}}} [\mathbf{1}[f(x) \neq 1]] + \frac{\rho}{(1 + \rho)^2} \mathbb{E}_{x \sim \mathcal{P}_{\text{min}}} [\mathbf{1}[f(x) \neq -1]] \\ &= \frac{2\rho}{(1 + \rho)^2} \mathbb{E}_{y \sim \text{Unif}(-1,1), x \sim \mathcal{P}(x|y)} [\mathbf{1}[f(x) \neq y]]. \end{aligned}$$

Now if we invoke the minimax lower bound (Kamalaruban & Williamson, 2018, Theorem 4) we get that

$$\min_f \max_P \frac{2\rho}{(1 + \rho)^2} \mathbb{E}_{y \sim \text{Unif}(-1,1), x \sim \mathcal{P}(x|y)} [\mathbf{1}[f(x) \neq y]] \geq \frac{C}{1 + \rho} \min \left\{ \sqrt{\frac{V}{(1 + \rho)n}}, \frac{1}{1 + \rho} \frac{V}{nh} \right\},$$

where the minimum over f is over all measurable functions from the training data to binary labels, the maximum is over a data distribution that can be correctly classified with a classifier from a VC class with VC dimension at most V and h is the Massart noise margin. For more thorough definitions we urge the reader to see (Kamalaruban & Williamson, 2018). With this lower bound we get that

$$\begin{aligned} \min_f \max_P \mathbb{E}_{y \sim \text{Unif}(-1,1), x \sim \mathcal{P}(x|y)} [\mathbf{1}[f(x) \neq y]] &\geq \frac{C(1 + \rho)}{2\rho} \min \left\{ \sqrt{\frac{V}{(1 + \rho)n}}, \frac{1}{1 + \rho} \frac{V}{nh} \right\} \\ &\geq \frac{C}{2} \min \left\{ \sqrt{\frac{V}{(1 + \rho)n}}, \frac{1}{1 + \rho} \frac{V}{nh} \right\}. \end{aligned}$$

Therefore we find that this lower bound gets smaller as the imbalance ratio ρ gets larger, predicting the wrong trend for the label shift problem.

F Details about results in Table 1

In Table 1, we listed results regarding the performance of undersampled algorithms to others that are reported in the literature. Here we provide detailed references to these results.

Label shift. The results for label shift are from the paper by Cao et al. (2019). The results are reported in Table 2 of that paper. For Imb CIFAR 10 (step 10), the undersampling result corresponds to the entry CB RS from that table with accuracy 84.59% (error 15.41%), while the best method corresponds to the method LDAM-DRW with accuracy 87.81% (error 12.19%). For Imb CIFAR100 (step 10), the undersampling result again corresponds to CB RS with accuracy 53.08% (error 46.92%) while the best method corresponds to the method LDAM-DRW with accuracy 59.46% (error 40.54%).

Group-covariate shift. The results for the group-covariate shift are from Table 2 in Idrissi et al. (2022). For the CelebA dataset, the undersampled accuracy corresponds to the method SUBG and the best accuracy is for gDRO. For the Waterbirds dataset, the undersampled method is SUBG and the best competitor is RWG. For the MultiNLI dataset, the undersampled accuracy corresponds to the method SUBG and the best accuracy is for gDRO. Finally, for the CivilComments dataset, the undersampled method is SUBG and the best method is RWG.

G Experimental details for Figures 2 and 4

We construct our label shift dataset from the original CIFAR10 dataset. We create a binary classification task using the “cat” and “dog” classes. We use the official test examples as the balanced test set with 1000 cats and 1000 dogs. To form the initial train and validation sets, we use 2500 cat examples (half of the training set) and 500 dog examples, corresponding to a 5:1 label imbalance. We use 80% of those examples for training and the rest for validation. We are left with 2500 additional cat examples and 4500 dog examples from the original train set which we add into our training set to generate Figure 2.

We use the same convolutional neural network architecture as (Byrd & Lipton, 2019; Wang et al., 2022) with random initializations for this dataset. We train this model using SGD for 800 epochs with batchsize 64, a constant learning rate 0.001 and momentum 0.9. The importance weights used upweight the minority class samples in the training loss and validation loss is calculated to be $\frac{\#Cat \text{ Train Examples}}{\#Dog \text{ Train Examples}}$. We note that all of the experiments were performed on an internal cluster on 8 GPUs.

VS loss: Given a dataset $\{x_i, y_i\}_{i=1}^n$, the VS loss (Kini et al., 2021) is defined as follows

$$\mathcal{L}_{VS}(f) := \sum_{i=1}^n \log \left(1 + \exp \left(- \left(\frac{n_{g_i}}{n_{\max}} \right)^\gamma y_i f(x_i) - \frac{\tau n_{g_i}}{n} \right) \right),$$

where g_i denotes the group label, n_{g_i} corresponds to the number of samples from the group, n_{\max} is the number of samples in the largest group and n is the total number of samples. We set $\tau = 3$ and $\gamma = 0.3$, the best hyperparameters identified by Wang et al. (2022) on this dataset for this neural network architecture.

Tilted loss: The tilted loss (Li et al., 2020) is defined as

$$\mathcal{L}_{\text{Tilted}}(f) := \frac{1}{t} \log \left[\sum_{i=1}^n \exp(t\ell(y_i f(x_i))) \right],$$

where we take ℓ to be the logistic loss. In our experiments we set $t = 2$.

Group-DRO: We run group-DRO (Sagawa et al., 2020, Algorithm 1) with the logistic loss. We set adversarial step-size $\eta_q = 0.05$ which was the best hyperparameter identified by Wang et al. (2022).