## A Preliminaries

## A. 1 Norms

For a set $\mathcal{S}$, the $\ell_{1}, \ell_{2}$ and $\ell_{\infty}$ norms on $\mathbb{R}^{\mathcal{S}}$ are given respectively by

$$
\|a\|_{1}=\sum_{v \in \mathcal{S}}\left|a_{v}\right|, \quad\|a\|_{2}=\sqrt{\sum_{v \in \mathcal{S}}\left(a_{v}\right)^{2}}, \quad\|a\|_{\infty}=\max _{v \in \mathcal{S}}\left|a_{v}\right| .
$$

Given a probability distribution $\pi$ on $\mathcal{S}$, we consider the norms $\|\cdot\|_{L_{1}(\pi)}$ and $\|\cdot\|_{L_{2}(\pi)}$ on $\mathbb{R}^{\mathcal{S}}$, given by

$$
\|a\|_{L_{1}(\pi)}=\sum_{v \in \mathcal{S}} \pi(v)\left|a_{v}\right|, \quad\|a\|_{L_{2}(\pi)}=\sqrt{\sum_{v \in \mathcal{S}} \pi(v)\left(a_{v}\right)^{2}} .
$$

We also take advantage of a number of matrix norms. For norms $\|\cdot\|_{\zeta}$ and $\|\cdot\|_{\xi}$ on $\mathbb{R}^{\mathcal{S}^{\prime}}$ and $\mathbb{R}^{\mathcal{S}}$ respectively, we consider the matrix operator norm of $M \in \mathbb{R}^{\mathcal{S}} \times \mathcal{S}^{\prime}$ given by

$$
\|M\|_{\zeta \rightarrow \xi}=\max _{a \in \mathbb{R}^{\mathcal{S}} \backslash\{0\}} \frac{\|M a\|_{\xi}}{\|a\|_{\zeta}}
$$

For the special case of $\|M\|_{\ell_{s} \rightarrow \ell_{t}}$, we will simply write $\|M\|_{s \rightarrow t}$. Of particular importance are $\|M\|_{1 \rightarrow \infty}$ which corresponds to the largest entry of $M,\|M\|_{1 \rightarrow 2}$, which corresponds to the maximum $\ell_{2}$ norm of a column of $M$, and $\|M\|_{2 \rightarrow \infty}$, which corresponds to the maximum $\ell_{2}$ norm of a row of $M$.
The inner product of two matrices $M$ and $N$ in $\mathbb{R}^{\mathcal{S} \times \mathcal{S}^{\prime}}$ is defined by $M \bullet N=\operatorname{Tr}\left(M^{\top} N\right)=$ $\sum_{u \in \mathcal{S}, v \in \mathcal{S}^{\prime}} m_{u, v} n_{u, v}$.
The factorization norm known as the $\gamma_{2}$ norm is given for $M \in \mathbb{R}^{\mathcal{S} \times \mathcal{S}^{\prime}}$ by

$$
\gamma_{2}(M)=\min \left\{\|R\|_{2 \rightarrow \infty}\|A\|_{1 \rightarrow 2}: R A=M\right\}
$$

The $\gamma_{2}$ norm is, indeed, a norm, i.e., it is non-negative, $\gamma_{2}(M)=0$ if and only if $M=0$, for any real $s$ we have $\gamma_{2}(s M)=|s| \gamma_{2}(M)$, and we also have the triangle inequality $\gamma_{2}(M+N) \leq$ $\gamma_{2}(M)+\gamma_{2}(N)$.
The approximate $\gamma_{2}$ norm is the smallest $\gamma_{2}$ norm of a matrix that approximates the given matrix entrywise up to an additive $\alpha$, i.e.,

$$
\gamma_{2}(M, \alpha)=\min \left\{\gamma_{2}(\widetilde{M}):\|\widetilde{M}-M\|_{1 \rightarrow \infty} \leq \alpha\right\}
$$

The dual $\gamma_{2}$ norm of a matrix $G$ in $\mathbb{R}^{\mathcal{S} \times \mathcal{S}^{\prime}}$ is given by

$$
\gamma_{2}^{*}(k)=\max \left\{M \bullet N: \gamma_{2}(M) \leq 1\right\}=\max _{f, g} \sum_{u \in \mathcal{S}, v \in \mathcal{S}^{\prime}} n_{u, v} f(u) g(v)
$$

where the second max ranges over functions $f: \mathcal{S} \rightarrow B_{2}$ and $g: \mathcal{S}^{\prime} \rightarrow B_{2}$ that map the index sets of the rows and columns of $N$, respectively, to vectors of $\ell_{2}$ norm at most 1 .

## A. 2 Differential privacy

Let $\mathcal{X}$ denote the data universe. A generic element from $\mathcal{X}$ will be denoted by $x$. We consider datasets of the form $X=\left(x_{1}, \ldots, x_{n}\right) \in \mathcal{X}^{n}$, each of which is identified with its histogram $h \in \mathbb{Z}_{\geq 0}^{\mathcal{X}}$ where, for every $x \in \mathcal{X}, h_{x}=\left|\left\{i: x_{i}=x\right\}\right|$, so that $\|h\|_{1}=n$. To refer to a dataset, we use $\bar{X}$ and $h$ interchangeably. A pair of datasets $X=\left(x_{1}, \ldots, x_{i}, \ldots, x_{n}\right)$ and $X^{\prime}=\left(x_{1}, \ldots, x_{i}^{\prime}, \ldots, x_{n}\right)$ are called adjacent if $X^{\prime}$ is obtained from $X$ by replacing an element $x_{i}$ of $X$ with a new universe element $x_{i}^{\prime}$.
For parameters $\varepsilon>0$, an $\varepsilon$-differentially private ( $\varepsilon$-DP) mechanism [Dwork et al., 2006] is a randomized function $\mathcal{M}: \mathcal{X}^{n} \rightarrow \mathcal{Z}$ which, for all adjacent datasets $X$ and $X^{\prime}$, for all outcomes $S \subseteq \mathcal{Z}$, satisfies

$$
\underset{\mathcal{M}}{\operatorname{Pr}}[\mathcal{M}(X) \in S] \leq e^{\varepsilon} \underset{\mathcal{M}}{\operatorname{Pr}}\left[\mathcal{M}\left(X^{\prime}\right) \in S\right]
$$

Of special interest are $\varepsilon$-differentially private mechanisms $\mathcal{M}_{i}: \mathcal{X} \rightarrow \mathcal{Y}$ which take a singleton dataset $X=\{x\}$ as input. These are referred to as local randomizers. A sequence of $\varepsilon$-differentially private local randomizers $\mathcal{M}_{1}, \ldots, \mathcal{M}_{n}$, together with a post-processing function $\mathcal{A}: \mathcal{Y}^{n} \rightarrow \mathcal{Z}$, specify a (non-interactive) locally $\varepsilon$-differentially private ( $\varepsilon$-LDP) mechanism $\mathcal{M}: \mathcal{X}^{n} \rightarrow \mathcal{Z}[E v$ fimievski et al., 2003, Dwork et al., 2006, Kasiviswanathan et al., 2008]. When the local mechanism $\mathcal{M}$ is applied to a dataset $X$, we refer to $\mathcal{T}_{\mathcal{M}}(X)=\left(\mathcal{M}_{1}\left(x_{1}\right), \ldots, \mathcal{M}_{n}\left(x_{n}\right)\right)$ as the transcript of the mechanism. Then the output of the mechanism is given by $\mathcal{M}(X)=\mathcal{A}\left(\mathcal{T}_{\mathcal{M}}(X)\right)$.

A linear query is specified by a bounded function $q: \mathcal{X} \rightarrow \mathbb{R}$. Abusing notation slightly, its answer on a dataset $X$ is given by $q(X)=\frac{1}{n} \sum_{i=1}^{n} q\left(x_{i}\right)$. We also extend this notation to distributions: if $\lambda$ is a distribution on $\mathcal{X}$, then we write $q(\lambda)$ for $\underset{x \sim \lambda}{\mathbb{E}}[q(x)]$. A workload is a set of linear queries $Q=\left\{q_{1}, \ldots, q_{k}\right\}$, and $Q(X)=\left(q_{1}(X), \ldots, q_{k}(X)\right)$ is used to denote their answers. The answers on a distribution $\lambda$ on $\mathcal{X}$ are denoted by $Q(\lambda)=\left(q_{1}(\lambda), \ldots, q_{k}(\lambda)\right)$. We will often represent $Q$ by its workload matrix $W \in \mathbb{R}^{Q \times \mathcal{X}}$ with entries $w_{q, X}=q(X)$. In this notation, the answers to the queries are given by $\frac{1}{n} W h$. We will often use $Q$ and $W$ interchangeably.

## A. 3 PAC learning

A concept $c: \mathcal{U} \rightarrow\{ \pm 1\}$ from a concept class $\mathcal{C}$ identifies each sample $a \in \mathcal{U}$ with a label $c(a)$. The empirical loss of the concept $c$ on a dataset $X=\left(\left(a_{1}, b_{1}\right), \ldots,\left(a_{n}, b_{n}\right)\right) \in(\mathcal{U} \times\{ \pm 1\})^{n}$, denoted $L_{X}(c)$, is given by

$$
L_{X}(c)=\frac{1}{n} \sum_{i=1}^{n}\left(\mathbb{I}\left[c\left(a_{i}\right) \neq b_{i}\right]\right)
$$

For a distribution $\lambda$ on $\mathcal{U} \times\{ \pm 1\}$, the population loss of $c$ on $\lambda$, denoted $L_{\lambda}(c)$ is given by $L_{\lambda}(c)=$ $\mathbb{P}_{(a, b) \sim \lambda}[c(a) \neq b]$.

We will say that a mechanism $\mathcal{M}:(\mathcal{U} \times\{ \pm 1\})^{n} \rightarrow\{ \pm 1\}^{\mathcal{U}}(\alpha, \beta)$-learns $\mathcal{C}$ agnostically with $n$ samples if, for any distribution $\lambda$ over $\mathcal{U} \times\{ \pm 1\}$, given as input a random dataset $X$ drawn i.i.d. from $\lambda$, the mechanism returns some hypothesis $h \in \mathcal{U} \rightarrow\{ \pm 1\}$ which satisfies

$$
\mathbb{P}_{X, \mathcal{M}}\left[L_{\lambda}(h) \leq \min _{c \in \mathcal{C}} L_{\lambda}(c)+\alpha\right] \geq 1-\beta
$$

Realizable learning is an important special case of agnostic learning where the underlying distribution agrees with some concept. We say that $\mathcal{M}:(\mathcal{U} \times\{ \pm 1\})^{n} \rightarrow\{ \pm 1\}^{\mathcal{U}}(\alpha, \beta)$-learns $\mathcal{C}$ realizably with $n$ samples if, whenever $\lambda$ is a distribution over $\mathcal{U} \times\{ \pm 1\}$ which satisfies $\mathbb{P}_{(a, b) \sim \lambda}[c(a)=b]=1$ for some unknown $c \in \mathcal{C}$, then, given a random dataset $X$ drawn i.i.d. from $\lambda$, the mechanism returns a hypothesis $h \in \mathcal{U} \rightarrow\{ \pm 1\}$ which satisfies

$$
\mathbb{P}_{X, \mathcal{M}}\left[L_{\lambda}(h) \leq \alpha\right] \geq 1-\beta
$$

The problem of refutation asks whether the underlying distribution is well approximated by the concept class. In particular, for $\theta \in[0,1]$, we will say that $\mathcal{M}_{\theta}:(\mathcal{U} \times\{ \pm 1\})^{n} \rightarrow\{ \pm 1\}(\alpha, \beta)$ refutes $\mathcal{C}$ for threshold $\theta$ if the following two conditions are met:

1. When $\lambda$ is a distribution on $\mathcal{U} \times\{ \pm 1\}$ which satisfies $L_{\lambda}(c) \leq \theta$ for some $c \in \mathcal{C}$,

$$
\mathbb{P}_{X, \mathcal{M}}[\mathcal{M}(X)=1] \geq 1-\beta
$$

2. When $\lambda$ is a distribution on $\mathcal{U} \times\{ \pm 1\}$ which, for all $h \in\{ \pm 1\}^{\mathcal{U}}$, satisfies $L_{\lambda}(h)>\theta+\alpha$, then

$$
\mathbb{P}_{X, \mathcal{M}}[\mathcal{M}(X)=-1] \geq 1-\beta
$$

Realizable refutation is a special case of agnostic refutation where the goal is to recognize whether the underlying distribution is labeled by a concept from the concept class. We say that $\mathcal{M}_{\theta}$ : $(\mathcal{U} \times\{ \pm 1\})^{n} \rightarrow\{ \pm 1\}(\alpha, \beta)$-refutes $\mathcal{C}$ realizably if it $(\alpha, \beta)$-refutes $\mathcal{C}$ for threshold 0 .

## B Equivalence of approximate $\gamma_{2}$ norms of difference and concept matrices

We prove Lemma 6, the equivalence of the approximate $\gamma_{2}$ norm for matrices $W$ and $D$, via the following three Lemmas.
Lemma 16. Let $\mathcal{C}$ be a concept class with concept matrix $W \in \mathbb{R}^{\mathcal{C} \times \mathcal{U}}$ and difference matrix $D \in \mathbb{R}^{\mathcal{C}^{2} \times \mathcal{U}}$. Then $\gamma_{2}(D, \alpha) \leq \gamma_{2}(W, \alpha)$.
Lemma 17. Let $\mathcal{C}$ be a concept class closed under negation. Let $W \in \mathcal{C} \times \mathcal{U}$ be its concept matrix (Definition 3) and let $D \in \mathbb{R}^{\mathcal{C}^{2} \times \mathcal{U}}$ be its difference matrix (Definition 5). Then $\gamma_{2}(W, \alpha) \leq \gamma_{2}(D, \alpha)$.
Lemma 18. Let $\mathcal{C}$ be a concept class. Let $W \in \mathbb{R}^{\mathcal{C}} \times \mathcal{U}$ be its concept matrix (Definition 3) and let $D \in \mathbb{R}^{\mathcal{C}^{2} \times \mathcal{U}}$ be its difference matrix (Definition 5). Then $\gamma_{2}(W, \alpha) \leq 2 \gamma_{2}(D, \alpha / 2)+1$.

Proof of Lemma Let $\widetilde{W} \in \mathcal{C} \times \mathcal{U}$ witness $\gamma_{2}(W, \alpha)$ so that $\|W-\widetilde{W}\|_{1 \rightarrow \infty} \leq \alpha$ and $\gamma_{2}(W, \alpha)=$ $\gamma_{2}(\widetilde{W})$.
Let $W^{\prime} \in \mathbb{R}^{\mathcal{C}^{2} \times \mathcal{U}}$ be the matrix with entries $w_{\left(c, c^{\prime}\right), a}^{\prime}=\widetilde{w}_{c, a}$. Similarly, let $W^{\prime \prime} \in \mathbb{R}^{\mathcal{C}^{2} \times \mathcal{U}}$ be the matrix with entries $w_{\left(c, c^{\prime}\right), a}^{\prime \prime}=\widetilde{w}_{c^{\prime}, a}$. Since $W^{\prime}$ and $W^{\prime \prime}$ are obtained from $\widetilde{W}$ by duplicating rows,

$$
\gamma_{2}(W, \alpha)=\gamma_{2}(\widetilde{W})=\gamma_{2}\left(W^{\prime}\right)=\gamma_{2}\left(W^{\prime \prime}\right)
$$

Now consider the matrix $\widetilde{D}=\frac{1}{2}\left(W^{\prime}-W^{\prime \prime}\right)$. By subadditivity and scaling properties,

$$
\gamma_{2}(\widetilde{D}) \leq \frac{1}{2}\left(\gamma_{2}\left(W^{\prime}\right)+\gamma_{2}\left(W^{\prime \prime}\right)\right)=\gamma_{2}(\widetilde{W})
$$

Moreover, for all $c, c^{\prime} \in \mathcal{C}, a \in \mathcal{U}$, the entry $\widetilde{d}_{\left(c, c^{\prime}\right), a}$ of $\widetilde{D}$ approximates entry $d_{\left(c, c^{\prime}\right), a}$ of $D$. Specifically,

$$
\begin{aligned}
\left|\widetilde{d}_{\left(c, c^{\prime}\right), a}-d_{\left(c, c^{\prime}\right), a}\right| & =\left|\frac{\widetilde{w}_{c, a}-\widetilde{w}_{c^{\prime}, a}}{2}-\frac{c(a)-c^{\prime}(a)}{2}\right| \\
& \leq\left|\frac{\widetilde{w}_{c, a}-c(a)}{2}\right|+\left|\frac{\widetilde{w}_{c^{\prime}, a}-c^{\prime}(a)}{2}\right| \\
& \leq \alpha .
\end{aligned}
$$

Hence $\|D-\widetilde{D}\|_{1 \rightarrow \infty} \leq \alpha$. Together with $\gamma_{2}(\widetilde{D}) \leq \gamma_{2}(\widetilde{W})$, this implies

$$
\gamma_{2}(D, \alpha) \leq \gamma_{2}(\widetilde{D}) \leq \gamma_{2}(\widetilde{W})=\gamma_{2}(W, \alpha)
$$

Proof of Lemma 17 Row $c$ of $W$ is identical with row $(c,-c)$ of $D$. Hence, $W$ is obtained from $D$ by deleting some of its rows. Since the $\gamma_{2}$ norm is non-increasing under taking submatrices, it follows that $\gamma_{2}(W, \alpha) \leq \gamma_{2}(D, \alpha)$.

Proof of Lemma 18 Fix an arbitrary concept $c^{\prime} \in \mathcal{C}$. Let $D^{\prime}$ be the submatrix of $D$ which includes row $\left(c, c^{\prime}\right)$ of $D$ for each $c \in \mathcal{C}$. Then $D^{\prime}=\frac{1}{2}\left(W-\mathbf{1}\left(c^{\prime}\right)^{T}\right)$ where $\mathbf{1}$ is the all-ones vector of dimension $|\mathcal{C}|$, and we identify $c^{\prime}$ with a vector in $\mathbb{R}^{\mathcal{U}}$. Expressing our concept matrix as $W=$ $2 D^{\prime}+\mathbf{1}\left(c^{\prime}\right)^{T}$, we may apply the scaling and subadditivity properties of the approximate $\gamma_{2}$ norm to obtain

$$
\begin{aligned}
\gamma_{2}(W, \alpha) & =\gamma_{2}\left(2 D^{\prime}+\mathbf{1}\left(c^{\prime}\right)^{T}, \alpha\right) \\
& \leq \gamma_{2}\left(2 D^{\prime}, \alpha\right)+\gamma_{2}\left(\mathbf{1}\left(c^{\prime}\right)^{T}, 0\right) \\
& \leq 2 \gamma_{2}\left(D^{\prime}, \alpha / 2\right)+1
\end{aligned}
$$

## C Omitted Proofs for the Lower Bound on Agnostic Learning

In proving Lemma 10, we take advantage of the following technical lemma by a standard geometric binning argument.

Lemma 19 ([Edmonds et al. 2019]). Suppose that $a_{1}, \ldots, a_{k} \in[0,1]$ and that $\pi$ is a probability distribution over $[k]$. Then for any $\beta \in(0,1]$, there exists a set $S \subseteq[k]$ such that $\pi(S) \cdot \min _{v \in S} a_{v} \geq$ $\frac{\sum_{v=1}^{k} \pi(v) a_{v}-\beta}{O(\log (1 / \beta))}$.

Proof of Lemma 10 Let $\pi$, together with $\lambda_{c, c^{\prime}}$ and $\mu_{c, c^{\prime}}$, be defined as in (4), (5) and (6). We will apply Lemma 19 to the values given, for $c, c^{\prime} \in \mathcal{C}$, by

$$
a_{c, c^{\prime}}=L_{\lambda_{c, c^{\prime}}}(c)-L_{\lambda_{c, c^{\prime}}}(c)=L_{\mu_{c, c^{\prime}}}(c)-L_{\mu_{c, c^{\prime}}}(c)
$$

Recall that we may assume, that for all $c, c^{\prime} \in \mathcal{C}$ the following inequality holds

$$
\sum_{a \in \mathcal{U}} d_{\left(c, c^{\prime}\right), a} u_{\left(c, c^{\prime}\right), a} \geq 0
$$

Together with (7), this gives $a_{c, c^{\prime}} \geq 0$ for all $c, c^{\prime} \in \mathcal{C}$.
By Lemma 19 , there exists some $S \subset \mathcal{C}^{2}$ such that
$\pi(S) \cdot \min _{\left(c, c^{\prime}\right) \in S}\left(L_{\lambda_{c, c^{\prime}}}(c)-L_{\lambda_{c, c^{\prime}}}\left(c^{\prime}\right)\right) \geq \frac{\underset{\left(c, c^{\prime}\right) \sim \pi}{\mathbb{E}}\left[L_{\lambda_{c, c^{\prime}}}(c)-L_{\lambda_{c, c^{\prime}}}\left(c^{\prime}\right)\right]-\alpha / 4}{O(\log (1 / \alpha))}=\frac{D \bullet U-\alpha / 4}{O(\log (1 / \alpha))}$.
By applying (8), we get the similar

$$
\pi(S) \cdot \min _{\left(c, c^{\prime}\right) \in S}\left(L_{\mu_{c, c^{\prime}}}\left(c^{\prime}\right)-L_{\mu_{c, c^{\prime}}}(c)\right) \geq \frac{D \bullet U-\alpha / 4}{O(\log (1 / \alpha))}
$$

Let $\widetilde{\pi}$ be defined by

$$
\widetilde{\pi}\left(c, c^{\prime}\right)= \begin{cases}\pi\left(c, c^{\prime}\right) / \pi(S), & \text { if } c, c^{\prime} \in S \\ 0, & \text { otherwise }\end{cases}
$$

Let also $\tau=\frac{\alpha}{D \bullet U} \in(0,1)$. For $\left(c, c^{\prime}\right) \in S$, let $\widetilde{\lambda}_{c, c^{\prime}}=\lambda_{c, c^{\prime}}$ and

$$
\widetilde{\mu}_{c, c^{\prime}}=(1-\tau \pi(S)) \lambda_{c, c^{\prime}}+\tau \pi(S) \mu_{c, c^{\prime}} .
$$

It holds then, for $\left(c, c^{\prime}\right) \in S$, that

$$
\widetilde{\lambda}_{c, c^{\prime}}-\widetilde{\mu}_{c, c^{\prime}}=\tau \cdot \pi(S) \cdot\left(\lambda_{c, c^{\prime}}-\mu_{c, c^{\prime}}\right)
$$

Hence, the matrix $\widetilde{U} \in \mathbb{R}^{\mathcal{C}^{2} \times \mathcal{U}}$ with entries defined by

$$
\widetilde{u}_{v, a}=\widetilde{\pi}(v) \cdot\left(\widetilde{\lambda}_{v}(a, 1)-\widetilde{\mu}_{v}(a, 1)\right)=-\widetilde{\pi}(v) \cdot\left(\widetilde{\lambda}_{v}(a,-1)-\widetilde{\mu}_{v}(a,-1)\right)
$$

satisfies

$$
\widetilde{u}_{\left(c, c^{\prime}\right), a}= \begin{cases}\tau u_{\left(c, c^{\prime}\right), a}, & \text { if }\left(c, c^{\prime}\right) \in S \\ 0, & \text { otherwise }\end{cases}
$$

It is easy to see from the definition of $\gamma_{2}^{*}$ that this implies $\gamma_{2}^{*}(\widetilde{U}) \leq \tau \gamma_{2}^{*}(U)=\frac{\alpha \gamma_{2}^{*}(U)}{D \cdot U}$.
Proof of Lemma 12 The main observation is that, since $\lambda$ and $\mu$ share the same marginal on $\mathcal{U}$ but the labels are given by the functions $s$ and $-s$, for any hypothesis $h: \mathcal{C} \rightarrow\{ \pm 1\}$ we have $L_{\lambda}(h)+L_{\mu}(h)=1$. Therefore,

$$
\begin{aligned}
\left(L_{\lambda}(h)-L_{\lambda}\left(c^{\prime}\right)\right)+\left(L_{\mu}(h)-L_{\mu}(c)\right) & =\left(\left(L_{\lambda}(h)+L_{\mu}(h)\right)-\left(1-L_{\mu}\left(c^{\prime}\right)\right)-L_{\mu}(c)\right. \\
& =L_{\mu}\left(c^{\prime}\right)-L_{\mu}(c)>\alpha
\end{aligned}
$$

This implies that if $L_{\lambda}(h)-L_{\lambda}\left(c^{\prime}\right) \leq \frac{\alpha}{4}$, then $L_{\mu}(h)-L_{\mu}(c)>\frac{3 \alpha}{4}$, as required.
Suppose now that $\mathcal{M}(\alpha / 4, \beta)$-learns $\mathcal{C}$ agnostically with $n$ samples. Let $A \subseteq\{ \pm 1\}^{\mathcal{U}}$ be the set of hypotheses with loss at most $L_{\lambda}\left(c^{\prime}\right)+\alpha / 4$ on $\lambda^{n}$. As we just showed, every hypothesis in $A$ has loss larger than $L_{\mu}(c)+3 \alpha / 4$ under $\mu$. Since

$$
\min _{c^{\prime \prime} \in \mathcal{C}} L_{\lambda}\left(c^{\prime \prime}\right) \leq L_{\lambda}\left(c^{\prime}\right), \quad \min _{c^{\prime \prime} \in \mathcal{C}} L_{\mu}\left(c^{\prime \prime}\right) \leq L_{\mu}(c)
$$

it follows from the definition of agnostic learning that $\mathbb{P}\left[\mathcal{M}\left(\lambda^{n}\right) \in A\right] \geq 1-\beta$, and $\mathbb{P}\left[\mathcal{M}\left(\lambda^{n}\right) \in A\right] \leq \beta$. Then, by the definition of total variation,

$$
d_{\mathrm{TV}}\left(\mathcal{M}\left(\lambda^{n}\right), \mathcal{M}\left(\mu^{n}\right)\right) \geq \mathbb{P}\left[\mathcal{M}\left(\lambda^{n}\right) \in A\right]-\mathbb{P}\left[\mathcal{M}\left(\mu^{n}\right) \in A\right] \geq 1-2 \beta
$$

completing the proof of the lemma.

## D Omitted Proofs for the Realizable Refutation Lower Bound

## D. 1 Derivation of dual formulation

Proof of Lemma 15 Let $L_{\mathcal{C}}=\left\{G \in \mathbb{R}^{\mathcal{C}} \times(\mathcal{U} \times\{ \pm 1\}): \gamma_{2}(G) \leq t\right\}$. Let $K_{\mathcal{C}}$ and $K_{\mathcal{C}}^{\prime}$ be as defined by equations (9) and (10). By definition, $\eta(\mathcal{C}, \alpha)>t$ if and only if $L_{\mathcal{C}}$ and $K_{\mathcal{C}}^{\prime}$ are disjoint.

Given some $U \in \mathbb{R}^{\mathcal{C}} \times(\mathcal{U} \times\{ \pm 1\})$, we are interested in the quantities $\max \left\{U \cdot G: G \in L_{\mathcal{C}}\right\}$ and $\min \left\{U \cdot G: G \in K_{\mathcal{C}}^{\prime}\right\}$. In particular, by the hyperplane separation theorem, since $L_{\mathcal{C}}$ and $K_{\mathcal{C}}^{\prime}$ are convex and $L_{\mathcal{C}}$ is also compact, they are disjoint exactly when there exists some $U \in \mathbb{R}^{\mathcal{C}} \times(\mathcal{U} \times\{ \pm 1\})$ such that

$$
\max \left\{U \cdot G: G \in L_{\mathcal{C}}\right\}<\min \left\{U \cdot G: G \in K_{\mathcal{C}}^{\prime}\right\}
$$

By definition,

$$
\max \left\{U \cdot G: G \in L_{\mathcal{C}}\right\}=t \gamma_{2}^{*}(U)
$$

Also,

$$
\begin{aligned}
& \min \left\{U \cdot G: G \in K_{\mathcal{C}}^{\prime}\right\} \\
= & \min _{G \in K_{\mathcal{C}}^{\prime}} \sum_{c \in \mathcal{C}, a \in \mathcal{U}}\left(u_{c,(a, c(a))} g_{c,(a, c(a))}+u_{c,(a,-c(a))} g_{c,(a,-c(a))}\right) \\
= & \min _{\substack{G \in K_{\mathcal{C}} \\
\theta \in \mathbb{R}^{\mathcal{C}}}} \sum_{c \in \mathcal{C}, a \in \mathcal{U}}\left(u_{c,(a, c(a))} \cdot\left(g_{c,(a, c(a))}+\theta_{c}\right)+u_{c,(a,-c(a))} \cdot\left(g_{c,(a,-c(a))}+\theta_{c}\right)\right) \\
= & \min _{G \in K_{\mathcal{C}}} \sum_{c \in \mathcal{C}, a \in \mathcal{U}}\left(u_{c,(a, c(a))} \cdot g_{c,(a, c(a))}+u_{c,(a,-c(a))} \cdot g_{c,(a,-c(a))}\right) \\
& +\min _{\theta \in \mathbb{R}^{\mathcal{C}}} \sum_{c \in \mathcal{C}} \theta_{c} \cdot \sum_{a \in \mathcal{U}}\left(u_{c,(a, c(a))}+u_{c,(a,-c(a))}\right)
\end{aligned}
$$

If, for some $c \in \mathcal{C}$, it holds that $\sum_{a \in \mathcal{U}}\left(u_{c,(a, c(a))}+u_{c,(a,-c(a))}\right) \neq 0$, then

$$
\min _{\theta_{c} \in \mathbb{R}} \theta_{c} \cdot \sum_{a \in \mathcal{U}}\left(u_{c,(a, c(a))}+u_{c,(a,-c(a))}\right)=-\infty .
$$

Also, if there exist $c \in \mathcal{C}$ and $x \in \mathcal{U}$ such that $u_{c,(a,-c(a))}<0$, then

$$
\min _{G \in K_{\mathcal{C}}} u_{c,(a,-c(a))} g_{c,(a,-c(a))}=-\infty .
$$

However, in the remaining case where $U$ is in the set $S_{\mathcal{C}}$, then

$$
\begin{aligned}
\min \left\{U \cdot G: G \in K_{\mathcal{C}}^{\prime}\right\} & =\min _{G \in K_{\mathcal{C}}} \sum_{c \in \mathcal{C}, a \in \mathcal{U}}\left(u_{c,(a, c(a))} \cdot g_{c,(a, c(a))}+u_{c,(a,-c(a))} \cdot g_{c,(a,-c(a))}\right) \\
& =\sum_{c \in \mathcal{C}, a \in \mathcal{U}}\left(-\alpha\left|u_{c,(a, c(a))}\right|+u_{c,(a,-c(a))}\right) .
\end{aligned}
$$

With these facts at our disposal, we obtain

$$
\begin{aligned}
\eta(\mathcal{C}, \alpha)>t & \Leftrightarrow K_{\mathcal{C}}^{\prime} \cap L_{\mathcal{C}}=\emptyset \\
& \Leftrightarrow \exists U \in \mathbb{R}^{\mathcal{C} \times(\mathcal{U} \times\{ \pm 1\})}, \max \left\{U \cdot G: G \in L_{\mathcal{C}}\right\}<\min \left\{U \cdot G: G \in K_{\mathcal{C}}\right\} \\
& \Leftrightarrow \exists U \in S_{\mathcal{C}}, t \gamma_{2}^{*}(U)<\sum_{c \in \mathcal{C}, a \in \mathcal{U}}\left(-\alpha\left|u_{c,(a, c(a))}\right|+u_{c,(a,-c(a))}\right) \\
& \Leftrightarrow \max _{U \in S_{\mathcal{C}}} \frac{\sum_{c \in \mathcal{C}, a \in \mathcal{U}}\left(u_{c,(a,-c(a))}-\alpha\left|u_{c,(a, c(a))}\right|\right)}{\gamma_{2}^{*}(U)}>t
\end{aligned}
$$

Since the equivalence holds for all $t \in \mathbb{R}$, it follows that

$$
\eta(\mathcal{C}, \alpha)=\max _{U \in S_{\mathcal{C}}} \frac{\sum_{c \in \mathcal{C}, a \in \mathcal{U}}\left(u_{c,(a,-c(a))}-\alpha\left|u_{c,(a, c(a))}\right|\right)}{\gamma_{2}^{*}(U)}
$$

## D. 2 General case of realizable refutation lower bound

Section 3.1.3 proved our lower bound for realizable refutation in the single-concept case. In the general case, there are two issues to resolve:

1. Instead of equation 12 holding for each concept, it holds on average. In particular,

$$
\sum_{c \in \mathcal{C}, a \in \mathcal{U}} u_{c,(a,-c(a))}>\sum_{c \in \mathcal{C}, a \in \mathcal{U}} \alpha\left|u_{c,(a, c(a))}\right| .
$$

Equivalently,

$$
\begin{equation*}
\mathbb{E}_{c \sim \pi}\left[L_{\lambda}(c)\right]>\frac{2 \alpha}{1+\alpha} \tag{14}
\end{equation*}
$$

2. Even if we can guarantee for a concept $c \in \mathcal{C}$ that

$$
L_{\lambda_{c}}(c)>\frac{2 \alpha}{1+\alpha}
$$

it may hold, for some other $h: \mathcal{U} \rightarrow\{ \pm 1\}$, that $L_{\lambda_{c}}(h)$ is small. We need to rule out this possibility in order to give a lower bound against refutation.

The first issue is resolved in Lemma 20 by applying the binning result of Lemma 19 The second issue will be resolved in Lemma 21
Lemma 20. Suppose there exist families $\left\{\lambda_{c}\right\}_{c \in \mathcal{C}}$ and $\left\{\mu_{c}\right\}_{c \in \mathcal{C}}$ of distributions over $\mathcal{U}$, together with a parameter distribution $\pi$ over $\mathcal{C}$, such that

$$
\Delta=\mathbb{E}_{c \sim \pi}\left[L_{\lambda_{c}}(c)\right]>\frac{2 \alpha}{1+\alpha}
$$

while, for all $c \in \mathcal{C}, L_{\mu_{c}}(c)=0$. Let, further, $U \in \mathbb{R}^{\mathcal{C}} \times \mathcal{U}$ be the matrix with entries $u_{c, a}=$ $\pi(c)\left(\lambda_{c}(a)-\mu_{c}(a)\right)$.
Then there exist families $\left\{\widetilde{\lambda}_{c}\right\}_{c \in \mathcal{C}}$ and $\left\{\widetilde{\mu}_{c}\right\}_{c \in \mathcal{C}}$ of distributions over $\mathcal{U} \times\{ \pm 1\}$, together with a parameter distribution $\widetilde{\pi}$ over $\mathcal{C}$, such that, for all $c$ in the support of $\widetilde{\pi}$,

$$
L_{\tilde{\lambda}_{c}}(c) \geq \Omega\left(\frac{\alpha}{1+\alpha} / \log \left(\frac{1+\alpha}{\alpha}\right)\right)
$$

while still $L_{\widetilde{\mu}_{c}}(c)=0$ for all $c \in \mathcal{C}$. Moreover, the matrix $\widetilde{U} \in \mathbb{R}^{\mathcal{C}} \times \mathcal{U}$ with entries $\widetilde{u}_{c, a}=$ $\widetilde{\pi}(c)\left(\widetilde{\lambda}_{c}(a)-\widetilde{\mu}_{c}(a)\right)$ satisfies

$$
\gamma_{2}^{*}(\widetilde{U}) \leq \frac{2 \alpha \gamma_{2}^{*}(U)}{(1+\alpha) \Delta}
$$

Proof. Apply Lemma 19, with $a_{c}=L_{\lambda_{c}}(c)$ for all $c \in \mathcal{C}$, and $\beta=\frac{\alpha}{1+\alpha}<\frac{\Delta}{2}$, to obtain $S \subseteq \mathcal{C}$ such that

$$
\pi(S) \cdot \min _{c \in S} a_{c} \geq \frac{\Delta-\beta}{O(\log (1 / \beta)} \geq \frac{\Delta}{O(\log ((1+\alpha) / \alpha)}
$$

Let $\widetilde{\pi}$ be $\pi$ conditional on membership in $S$. Thus,

$$
\widetilde{\pi}(v)= \begin{cases}\pi(v) / \pi(S), & \text { if } v \in S \\ 0, & \text { otherwise }\end{cases}
$$

Let $\tau=\frac{2 \alpha}{(1+\alpha) \Delta} \in(0,1)$. For all $c \in \mathcal{C}$, define $\widetilde{\mu}_{c}=\mu_{c}$ and $\widetilde{\lambda}_{c}=\tau \pi(S) \lambda_{c}+(1-\tau \pi(S)) \mu_{c}$. Then, for all $c$ in the support of $\widetilde{\pi}$,

$$
\begin{gathered}
L_{\widetilde{\mu}_{c}}(c)=L_{\mu_{c}}(c)=0 \\
L_{\widetilde{\lambda}_{c}}(c)=\tau \cdot \pi(S) \cdot L_{\lambda_{c}}(c) \geq \frac{\frac{\alpha}{1+\alpha}}{O\left(\left(\log \left(\frac{1+\alpha}{\alpha}\right)\right)\right.}
\end{gathered}
$$

Moreover, the matrix $\widetilde{U} \in \mathbb{R}^{\mathcal{C} \times \mathcal{U}}$ with entries $\widetilde{u}_{c, a}=\widetilde{\pi}(v)\left(\widetilde{\lambda}_{v}(a)-\widetilde{\mu}_{v}(a)\right)$ is obtained from the matrix $\tau U$ by replacing some of its rows with the zero-vector. It follows immediately that $\gamma_{2}^{*}(\widetilde{U}) \leq \tau \gamma_{2}^{*}(U)=\frac{2 \alpha \gamma_{2}^{*}(U)}{(1+\alpha) \Delta}$.

Lemma 21. Suppose we have distributions $\lambda_{c}$ and $\mu_{c}$ on $\mathcal{U} \times\{ \pm 1\}$ for each $c \in \mathcal{C}$ where:
(a) $L_{\mu_{c}}(c)=0$;
(b) $L_{\lambda_{c}}(c)>\alpha$.

Then there exist distributions $\widetilde{\lambda}_{c}$ and $\widetilde{\mu}_{c}$ for each $c \in \mathcal{C}$ such that:
(c) $L_{\widetilde{\mu}_{c}}(c)=0$;
(d) $\forall h: \mathcal{U} \rightarrow\{ \pm 1\}, L_{\widetilde{\lambda}_{c}}(h)>\frac{\alpha}{2}$;
(e) $\widetilde{\lambda}_{c}-\widetilde{\mu}_{c}=\frac{1}{2}\left(\lambda_{c}-\mu_{c}\right)$.

Proof. For $c \in \mathcal{C}$, let $\sigma_{c}$ be the distribution on $\mathcal{U} \times\{ \pm 1\}$ which has the same marginal on $\mathcal{U}$ as does $\lambda_{c}$, and which satisfies $c(a)=b$ for all $(a, b)$ in the support of $\sigma_{c}$. Also, let $\widetilde{\lambda}_{c}=\frac{1}{2} \lambda_{c}+\frac{1}{2} \sigma_{c}$ and $\widetilde{\mu}_{c}=\frac{1}{2} \mu_{c}+\frac{1}{2} \sigma_{c}$. Properties (c) and (e) follow immediately.
To establish property (d) notice first that, for any $a \in \mathcal{U}$ in the support of $\lambda_{c}, \mathbb{P}_{\widetilde{\lambda}_{c}}[b=c(a) \mid a] \geq \frac{1}{2}$, and also $\mathbb{P}_{\widetilde{\lambda}_{c}}[b \neq c(a) \mid a]=\frac{1}{2} \cdot \mathbb{P}_{\lambda_{c}}[b \neq c(a) \mid a]$. Then, for any function $h: \mathcal{U} \rightarrow\{ \pm 1\}$.

$$
\begin{aligned}
L_{\widetilde{\lambda}_{c}}(h) & =\mathbb{P}_{\widetilde{\lambda}_{c}}[h(a) \neq b] \\
& =\sum_{(a, b) \in \mathcal{U} \times\{ \pm 1\}} \mathbb{P}_{\widetilde{\lambda}_{c}}[a] \cdot \mathbb{P}_{\widetilde{\lambda}_{c}}[h(a) \neq b \mid a] \\
& \geq \sum_{(a, b) \in \mathcal{U} \times\{ \pm 1\}} \mathbb{P}_{\widetilde{\lambda}_{c}}[a] \cdot \min \left\{\mathbb{P}_{\widetilde{\lambda}_{c}}[b=c(a) \mid a], \mathbb{P}_{\widetilde{\lambda}_{c}}[b \neq c(a) \mid a]\right\} \\
& \geq \sum_{(a, b) \in \mathcal{U} \times\{ \pm 1\}} \mathbb{P}_{\lambda_{c}}[a] \cdot \min \left\{\frac{1}{2}, \frac{1}{2} \cdot \mathbb{P}_{\lambda_{c}}[b \neq c(a) \mid a]\right\} \\
& =\frac{1}{2} \cdot \sum_{(a, b) \in \mathcal{U} \times\{ \pm 1\}} \mathbb{P}_{\lambda_{c}}[a] \cdot \mathbb{P}_{\lambda_{c}}[b \neq c(a) \mid a] \\
& =\frac{1}{2} \cdot L_{\lambda_{c}}(c) \\
& >\frac{\alpha}{2} .
\end{aligned}
$$

Equipped with Lemmas 20 and 21 , we are ready to prove our lower bound against realizable refutation.
Proof of Theorem 2 We define the parameter distribution $\pi$ over $\mathcal{C}$, and the distribution families $\left\{\lambda_{c}\right\}_{c \in \mathcal{C}}$ and $\left\{\mu_{c}\right\}_{c \in \mathcal{C}}$ over $\mathcal{U} \times\{ \pm 1\}$, as in Section 3.1.2. We denote $\Delta=\mathbb{E}_{c \sim \pi}\left[L_{\lambda_{c}}(c)\right]$. By equation (14), together with Lemmas 20 and 21, we obtain modified families of distributions $\left\{\widetilde{\lambda}_{c}\right\}_{c \in \mathcal{C}}$ and $\left\{\widetilde{\mu}_{c}\right\}_{c \in \mathcal{C}}$, together with a parameter distribution $\widetilde{\pi}$ over $\mathcal{C}$, such that, for all $c$ in the support of $\widetilde{\pi}$, and for all functions $h: \mathcal{U} \rightarrow\{ \pm 1\}$,

$$
L_{\tilde{\lambda}_{c}}(h)=\Omega\left(\frac{\alpha}{1+\alpha} / \log \left(\frac{1+\alpha}{\alpha}\right)\right)
$$

while $L_{\widetilde{\mu}_{c}}(c)=0$ for all $c \in \mathcal{C}$. By Lemmas 11,20 and 21, we may assume further that the matrix $\widetilde{M} \in \mathbb{R}^{\mathcal{C}} \times(\mathcal{U} \times\{ \pm 1\})$ with entries $m_{c,(a, b)}=\widetilde{\lambda}_{c}(a, b)-\widetilde{\mu}_{c}(a, b)$ satisfies

$$
\|\widetilde{M}\|_{\ell_{\infty} \rightarrow L_{2}(\widetilde{\pi})} \leq \frac{4 \alpha \gamma_{2}^{*}(U)}{(1+\alpha) \Delta}
$$

Now let $\mathcal{M}$ be an $\varepsilon$-LDP protocol which is able to distinguish a labeling by some $c \in \mathcal{C}$ from a distribution with which every function $h: \mathcal{U} \rightarrow\{ \pm 1\}$ disagrees with the labels with probability $\Omega\left(\frac{\alpha}{1+\alpha} / \log \left(\frac{1+\alpha}{\alpha}\right)\right)$. If this is true, then, for every $c \in \mathcal{C}$ in the support of $\widetilde{\pi}$,

$$
\mathrm{D}_{\mathrm{KL}}\left(\mathcal{T}_{\mathcal{M}}\left(\widetilde{\lambda}_{c}^{n}\right) \| \mathcal{T}_{\mathcal{M}}\left(\widetilde{\mu}_{c}^{n}\right)\right)=\Omega(1)
$$

Meanwhile, Lemma 8 guarantees

$$
\underset{c \sim \widetilde{\pi}}{\mathbb{E}}\left[\mathrm{D}_{\mathrm{KL}}\left(\mathcal{T}_{\mathcal{M}}\left(\widetilde{\lambda}_{c}^{n}\right) \| \mathcal{T}_{\mathcal{M}}\left(\widetilde{\mu}_{c}^{n}\right)\right)\right] \leq O\left(n \varepsilon^{2}\right) \cdot\|\widetilde{M}\|_{\ell_{\infty} \rightarrow L_{2}(\pi)}^{2}
$$

whereby we obtain

$$
\begin{equation*}
n=\Omega\left(\frac{1}{\varepsilon^{2} \cdot\|\widetilde{M}\|_{\ell_{\infty} \rightarrow L_{2}(\pi)}^{2}}\right)=\Omega\left(\frac{(1+\alpha)^{2} \Delta^{2}}{\varepsilon^{2} \alpha^{2} \gamma_{2}^{*}(U)^{2}}\right) \tag{15}
\end{equation*}
$$

Now we may use

$$
\begin{equation*}
\gamma_{2}^{*}(U)=\frac{\sum_{c \in \mathcal{C}, a \in \mathcal{U}}\left(u_{c,(a,-c(a))}-\alpha\left|u_{c,(a, c(a))}\right|\right)}{\eta(\mathcal{C}, \alpha)} \tag{16}
\end{equation*}
$$

Note that, for any $c \in \mathcal{C}$, since $L_{\mu_{c}}(c)=0$,

$$
\begin{aligned}
& \frac{1}{\pi(c)} \sum_{a \in \mathcal{U}} u_{c,(a,-c(a))}-\alpha\left|u_{c,(a, c(a))}\right| \\
& =\mathbb{P}_{(a, b) \sim \lambda_{c}}[c(a) \neq b]-\mathbb{P}_{(a, b) \sim \mu_{c}}[c(a) \neq b] \\
& \quad-\alpha \cdot\left(\mathbb{P}_{(a, b) \sim \lambda_{c}}[c(a)=b]+\mathbb{P}_{(a, b) \sim \mu_{c}}[c(a)=b]\right) \\
& =(1+\alpha) \cdot \mathbb{P}_{(a, b) \sim \lambda_{c}}[c(a) \neq b]-2 \alpha \\
& =(1+\alpha) \cdot L_{\lambda_{c}}(c)-2 \alpha .
\end{aligned}
$$

Taking expectations over $c \sim \pi$, we have

$$
\begin{align*}
\sum_{c \in \mathcal{C}, a \in \mathcal{U}} u_{c,(a,-c(a))}-\alpha\left|u_{c,(a, c(a))}\right| & =(1+\alpha) \cdot \underset{c \sim \pi}{\mathbb{E}}\left[L_{\lambda_{c}}(c)\right]-2 \alpha \\
& =(1+\alpha) \cdot \Delta-2 \alpha \tag{17}
\end{align*}
$$

Putting equations (15), 16, and (17) together, we have

$$
n=\Omega\left(\frac{(1+\alpha)^{2} \Delta^{2} \eta(\mathcal{C}, \alpha)^{2}}{\varepsilon^{2} \alpha^{2}((1+\alpha) \Delta-2 \alpha)^{2}}\right)=\Omega\left(\frac{\eta(\mathcal{C}, \alpha)^{2}}{\varepsilon^{2} \alpha^{2}}\right)
$$

## E Open problems

This work, together with Edmonds et al. [2019], largely completes the picture of agnostic refutability and learnability under non-interactive LDP. In the realizable setting, we have shown that refutation implies learning for non-interactive LDP. It is an interesting open problem to determine the converse - whether realizable learning implies refutation. Secondly, for an arbitrary concept class $\mathcal{C}$, can we obtain a characterization of realizable learnability in terms of a quantity which is efficiently computable from the definition of $\mathcal{C}$ ? Furthermore, for both the realizable and agnostic versions, the relationships obtained between the sample complexities of refutability and learnability in this work are indirect, via characterizations of these tasks by the approximate $\gamma_{2}$ norm. For example, although realizable refutability implies realizable learnability under non-interactive LDP, it remains open how one might obtain a non-interactive LDP protocol for learnability directly from one for refutability.

