## Supplementary Material

## A Omitted Technical Preliminaries

Here we record definitions and facts that will be used in our proofs.
Definition A. 1 (Pairwise Correlation). The pairwise correlation of two distributions with probability mass functions (pmfs) $D_{1}, D_{2}:\{0,1\}^{M} \rightarrow \mathbb{R}_{+}$with respect to a distribution with pmf $D:\{0,1\}^{M} \rightarrow \mathbb{R}_{+}$, where the support of $D$ contains the supports of $D_{1}$ and $D_{2}$, is defined as $\chi_{D}\left(D_{1}, D_{2}\right)+1 \stackrel{\text { def }}{=} \sum_{x \in\{0,1\}^{M}} D_{1}(x) D_{2}(x) / D(x)$. We say that a collection of $s$ distributions $\mathcal{D}=\left\{D_{1}, \ldots, D_{s}\right\}$ over $\{0,1\}^{M}$ is $(\gamma, \beta)$-correlated relative to a distribution $D$ if $\left|\chi_{D}\left(D_{i}, D_{j}\right)\right| \leq \gamma$ for all $i \neq j$, and $\left|\chi_{D}\left(D_{i}, D_{j}\right)\right| \leq \beta$ for $i=j$.

The following notion of dimension effectively characterizes the difficulty of the decision problem.
Definition A. 2 (SQ Dimension). For $\gamma, \beta>0$, a decision problem $\mathcal{B}(\mathcal{D}, D)$, where $D$ is fixed and $\mathcal{D}$ is a family of distributions over $\{0,1\}^{M}$, let $s$ be the maximum integer such that there exists $\mathcal{D}_{D} \subseteq \mathcal{D}$ such that $\mathcal{D}_{D}$ is $(\gamma, \beta)$-correlated relative to $D$ and $\left|\mathcal{D}_{D}\right| \geq s$. We define the Statistical Query dimension with pairwise correlations $(\gamma, \beta)$ of $\mathcal{B}$ to be $s$ and denote it by $\operatorname{SD}(\mathcal{B}, \gamma, \beta)$.

The connection between SQ dimension and lower bounds is captured by the following lemma.
Lemma A. 3 (|FGR $\left.{ }^{+} 17 \mid\right)$. Let $\mathcal{B}(\mathcal{D}, D)$ be a decision problem, where $D$ is the reference distribution and $\mathcal{D}$ is a class of distributions over $\{0,1\}^{M}$. For $\gamma, \beta>0$, let $s=\operatorname{SD}(\mathcal{B}, \gamma, \beta)$. Any SQ algorithm that solves $\mathcal{B}$ with probability at least $2 / 3$ requires at least $s \cdot \gamma / \beta$ queries to the $\operatorname{STAT}(\sqrt{2 \gamma})$ oracles.

We have the following fact about the chi-squared inner product in the discrete setting.
Fact A.4. For distributions $\mathbf{P}, \mathbf{Q}$ over $\{0,1\}^{M}$, we have that $1+\chi_{U_{M}}(\mathbf{P}, \mathbf{Q})=$ $\sum_{T \subseteq[M]} \widehat{\mathbf{P}}(T) \widehat{\mathbf{Q}}(T)$.

We will also use the following standard fact:
Fact A.5. Let $m, M \in \mathbb{Z}_{+}$with $m<M$. For any constant $0<c<1$ and $M>2 m / c$, there exists a collection $\mathcal{C}$ of $2^{\Omega_{c}(m)}$ subsets $S \subseteq[M]$ such that any pair $S, S^{\prime} \in \mathcal{C}$, with $S \neq S^{\prime}$, satisfies $\left|S \cap S^{\prime}\right|<c m$.

In fact, an appropriate size set of random subsets satisfies the above statement with high probability.

The following correlation lemma states that the distributions $\mathbf{P}_{S}^{A}$ are nearly orthogonal as long as $A$ satisfies the nearly moment-matching condition.
Lemma A. 6 (Correlation Lemma [DKS22]). Let $k, m, M \in \mathbb{Z}_{+}$with $k \leq m \leq M$. If the distribution $A$ on $[m] \cup\{0\}$ satisfies Condition 3.3 then for all $S, S^{\prime} \subseteq[M]$ with $|S|=\left|S^{\prime}\right|=m$, we have that

$$
\begin{equation*}
\left|\chi_{U_{M}}\left(\mathbf{P}_{S}^{A}, \mathbf{P}_{S^{\prime}}^{A}\right)\right| \leq\left(\left|S \cap S^{\prime}\right| / m\right)^{k+1} \chi^{2}(A, \operatorname{Bin}(m, 1 / 2))+k \nu^{2} \tag{1}
\end{equation*}
$$

## B Omitted Proofs from Section 3

## B. 1 Proof of Proposition 3.5

Let $\mathcal{C}$ be a collection of $s=2^{\Omega(m)}$ subsets $S \subseteq[M]$ with $|S|=m$ whose pairwise intersections are all less than $m / 2$. By Fact A.5 (taking the local parameter $c=1 / 2$ ), such a set is guaranteed to exist. We then need to show that for $S, S^{\prime} \in \mathcal{C}$, we have that $\left|\chi_{U_{M}^{p}}\left(\mathbf{P}_{S, a, b}^{A, B, p}, \mathbf{P}_{S^{\prime}, a, b}^{A, B, p}\right)\right|$ is small. Since $U_{M}^{p}, \mathbf{P}_{S, a, b}^{A, B, p}$, and $\mathbf{P}_{S^{\prime}, a, b}^{A, B, p}$ all assign $y=1$ with probability $p$, it is not hard to see that

$$
\begin{aligned}
\chi_{U_{M}^{p}}\left(\mathbf{P}_{S, a, b}^{A, B, p}, \mathbf{P}_{S^{\prime}, a, b}^{A, B, p}\right)= & p \chi_{U_{M}^{p} \mid y=1}\left(\left(\mathbf{P}_{S, a, b}^{A, B, p} \mid y=1\right),\left(\mathbf{P}_{S^{\prime}, a, b}^{A, B, p} \mid y=1\right)\right)+ \\
& (1-p) \chi_{U_{M}^{p} \mid y=-1}\left(\left(\mathbf{P}_{S, a, b}^{A, B, p} \mid y=-1\right),\left(\mathbf{P}_{S^{\prime}, a, b}^{A, B, p} \mid y=-1\right)\right) \\
= & p \chi_{U_{M}}\left(\mathbf{P}_{S}^{A}, \mathbf{P}_{S^{\prime}}^{A}\right)+(1-p) \chi_{U_{M}}\left(\mathbf{P}_{S}^{B}, \mathbf{P}_{S^{\prime}}^{B}\right) .
\end{aligned}
$$

By Lemma A.6, for $S, S^{\prime} \in \mathcal{C}$ with $S \neq S^{\prime}$, it holds that

$$
\chi_{U_{M}^{p}}\left(\mathbf{P}_{S, a, b}^{A, B, p}, \mathbf{P}_{S^{\prime}, a, b}^{A, B, p}\right) \leq k \nu^{2}+2^{-k}\left(\chi^{2}(A, \operatorname{Bin}(m, 1 / 2))+\chi^{2}(B, \operatorname{Bin}(m, 1 / 2))\right) \leq \tau
$$

If $S=S^{\prime}$, a similar computation shows that

$$
\chi_{U_{M}^{p}}\left(\mathbf{P}_{S, a, b}^{A, B, p}, \mathbf{P}_{S, a, b}^{A, B, p}\right)=\chi^{2}\left(\mathbf{P}_{S, a, b}^{A, B, p}, U_{M}^{p}\right) \leq \chi^{2}(A, \operatorname{Bin}(m, 1 / 2))+\chi^{2}(B, \operatorname{Bin}(m, 1 / 2))
$$

Let $\gamma=\tau$ and $\beta=\chi^{2}(A, \operatorname{Bin}(m, 1 / 2))+\chi^{2}(B, \operatorname{Bin}(m, 1 / 2))$. We have that the Statistical Query dimension of this testing problem with correlations $(\gamma, \beta)$ is at least $s$. Then applying Lemma A. 3 with $(\gamma, \beta)$ completes the proof.

## B. 2 Proof of Lemma 3.8

The conditions on $\mu$ define a linear program (LP). We will show that this LP is feasible by showing that the dual LP is infeasible. The dual LP asks for a degree at most $k$ real polynomial $q(x)$ such that

$$
|q(0)| \geq(1 / 11) \sum_{i=1-s}^{s-1}|q(i)|
$$

Consider the parameterization $p(\theta)=q(s \sin (\theta))$. We will leverage the fact that $p(\theta)$ is a degree- $k$ polynomial in $e^{\mathbf{i} \theta}$ and $e^{-\mathbf{i} \theta}$. In particular, $p(\theta)$ can be written as

$$
p(\theta)=\sum_{j=-k}^{k} a_{j} e^{\mathbf{i} j \theta}
$$

for some complex coefficients $a_{j} \in \mathbb{C}$. By normalizing, we can assume that $\sum_{j=-k}^{k}\left|a_{j}\right|^{2}=1$. Then, for any $\theta$, we have that

$$
|p(\theta)| \leq \sum_{j=-k}^{k}\left|a_{j}\right|=O(\sqrt{k})
$$

where the final inequality follows from the Cauchy-Schwarz. In particular, $|q(0)|=|p(0)|=O(\sqrt{k})$. In addition, for any $\theta$, by Cauchy-Schwarz, we have that

$$
\left|p^{\prime}(\theta)\right|=\left|\sum_{j=-k}^{k} j a_{j} e^{\mathbf{i} j \theta}\right| \leq \sum_{j=-k}^{k}|j|\left|a_{j}\right| \leq \sqrt{\sum_{j=-k}^{k} j^{2}}=O\left(k^{3 / 2}\right)
$$

Finally, we note that

$$
\frac{1}{2 \pi} \int_{0}^{2 \pi}|p(\theta)|^{2} d \theta=\sum_{j=-k}^{k}\left|a_{j}\right|^{2}=1
$$

Combining the latter with the fact that $|p(\theta)|=O(\sqrt{k})$, we obtain that

$$
\int_{0}^{2 \pi}|p(\theta)| d \theta=\Omega\left(k^{-1 / 2}\right)
$$

For any $\theta \in[0,2 \pi]$, let $n(\theta)$ be the closest $\phi \in[0,2 \pi]$ such that $s \sin (\phi)$ is an integer in $\{1-s, 2-$ $s, \ldots, s-1\}$. It is not hard to see that $|n(\theta)-\theta|=O\left(s^{-1 / 2}\right)$ for all such $\theta$. Furthermore, we have that

$$
|p(n(\theta))-p(\theta)| \leq|n(\theta)-\theta| \sup _{\theta^{\prime} \in[0,2 \pi]}\left|p^{\prime}\left(\theta^{\prime}\right)\right| \leq O\left(k^{3 / 2} s^{-1 / 2}\right)
$$

We can thus write

$$
\Omega\left(k^{-1 / 2}\right)=\int_{0}^{2 \pi}|p(\theta)| d \theta \leq \int_{0}^{2 \pi}|p(n(\theta))| d \theta+O\left(k^{3 / 2} s^{-1 / 2}\right)
$$

Therefore,

$$
\int_{0}^{2 \pi}|p(n(\theta))| d \theta \geq \Omega\left(k^{-1 / 2}\right)
$$

On the other hand, each value of $p(n(\theta))$ is equal to the value of $q$ evaluated at some integer between $1-s$ and $s-1$. Furthermore, it is not hard to see that each such integer occurs for at most a total of $O\left(s^{-1 / 2}\right)$ range of $\theta$ 's. Therefore, we get that

$$
O\left(s^{-1 / 2}\right) \sum_{i=1-s}^{s-1}|q(i)| \geq \Omega\left(k^{-1 / 2}\right)
$$

Combining with the fact that $|q(0)|=O\left(k^{1 / 2}\right)$, this shows that it is impossible that

$$
|q(0)| \geq 1 / 4 \sum_{i=1-s}^{s-1}|q(i)|
$$

This completes our proof.

## C Omitted Proofs from Section 4

## C. 1 Proof of Claim 4.2

For a $\mathbf{v}_{S}$ the vector whose $i^{t h}$ coordinate is 1 if $i \in S$ and 0 otherwise, let $g:\{0,1\}^{m^{\prime}} \rightarrow\{ \pm 1\}$ be defined as $g(\mathbf{x})=-1$ if and only if $\mathbf{v}_{S}^{T} \mathbf{x} \in J$. In this way, we are able to write $g$ as a degree- $2 d$ PTF, i.e., $g(\mathbf{x})=\operatorname{sign}\left(\prod_{z \in J}\left(\mathbf{v}_{S}^{T} \mathbf{x}-z\right)^{2}\right)$. Therefore, there exists some LTF $L: \mathbb{R}^{M} \rightarrow\{ \pm 1\}$ such that $g(\mathbf{x})=L\left(\mathbf{x}^{\prime}\right)=L\left(V_{2 d}(\mathbf{x})\right)$ for all $\mathbf{x}$. We now bound the error for LTF $L$ under the distribution $\left(\mathbf{X}^{\prime}, Y^{\prime}\right)$. By the law of total probability, we have that

$$
\begin{aligned}
& \operatorname{Pr}_{\left(\mathbf{X}^{\prime}, Y^{\prime}\right)}\left[Y^{\prime} \neq L\left(\mathbf{X}^{\prime}\right)\right]=\operatorname{Pr}_{(\mathbf{X}, Y)}[Y \neq g(\mathbf{X})] \\
\leq & \operatorname{Pr}_{(\mathbf{X}, Y)}[Y \neq g(\mathbf{X}) \mid Y=1]+\mathbf{P r}_{(\mathbf{X}, Y)}[Y \neq g(\mathbf{X}) \mid Y=-1]
\end{aligned}
$$

We note that our hard distribution returns $\left(\mathbf{x}^{\prime}, y^{\prime}\right)$ with $y^{\prime}=L\left(\mathbf{x}^{\prime}\right)$, unless it picked a sample corresponding to a sample of $\mathcal{D}_{-}$coming from $\bar{J}$, therefore,

$$
\operatorname{Pr}_{\left(\mathbf{X}^{\prime}, Y^{\prime}\right)}\left[Y^{\prime} \neq L\left(\mathbf{X}^{\prime}\right)\right] \leq \operatorname{Pr}_{(\mathbf{X}, Y)}[Y \neq g(\mathbf{X}) \mid Y=-1] \leq \zeta
$$

which implies that $\mathrm{OPT}_{\text {Mass }} \leq \zeta \leq \exp \left(-\Omega\left(\log (M)^{8 / 9}\right)\right)$. We then show that $\left(\mathbf{X}^{\prime}, Y^{\prime}\right)$ is a Massart LTF distribution with noise rate upper bound of $\eta=1 / 3$. For any fixed $\mathbf{x}^{\prime} \in \mathbb{R}^{M}$, we have that

$$
\begin{aligned}
& \frac{\operatorname{Pr}_{\left(\mathbf{X}^{\prime}, Y^{\prime}\right)}\left[Y^{\prime}=1 \mid \mathbf{X}^{\prime}=\mathbf{x}^{\prime}\right]}{\operatorname{Pr}_{\left(\mathbf{X}^{\prime}, Y^{\prime}\right)}\left[Y^{\prime}=-1 \mid \mathbf{X}^{\prime}=\mathbf{x}^{\prime}\right]}=\frac{\operatorname{Pr}_{(\mathbf{X}, Y)}[Y=1 \mid \mathbf{X}=\mathbf{x}]}{\operatorname{Pr}_{(\mathbf{X}, Y)}[Y=-1 \mid \mathbf{X}=\mathbf{x}]} \\
= & \frac{\operatorname{Pr}_{(\mathbf{X}, Y)}[Y=1] \cdot \mathbf{P r}_{(\mathbf{X}, Y)}[\mathbf{X}=\mathbf{x} \mid Y=1]}{\operatorname{Pr}_{(\mathbf{X}, Y)}[Y=-1] \cdot \operatorname{Pr}_{(\mathbf{X}, Y)}[\mathbf{X}=\mathbf{x} \mid Y=-1]}=\frac{\left\|\mathcal{D}_{+}\right\|_{1} \cdot \mathbf{P}_{S}^{\mathcal{D}_{+}}(\mathbf{x})}{\left\|\mathcal{D}_{-}\right\|_{1} \cdot \mathbf{P}_{S}^{\mathcal{D}_{-}(\mathbf{x})}}=\frac{\mathcal{D}_{+}\left(\mathbf{v}_{S}^{T} \mathbf{x}\right)}{\mathcal{D}_{-}\left(\mathbf{v}_{S}^{T} \mathbf{x}\right)} .
\end{aligned}
$$

Therefore, if $\mathbf{v}_{S}^{T} \mathbf{x} \in J$, the above ratio will be 0 and $L\left(\mathbf{x}^{\prime}\right)=-1$, which means that the noise rate $\eta\left(\mathrm{x}^{\prime}\right)=0$; otherwise the above ratio will be at least 2 (since $\mathcal{D}_{+}>2 \mathcal{D}_{-}$on $\bar{J}$ by property 1 (b) of Proposition 3.6, and $L\left(\mathbf{x}^{\prime}\right)=1$, which means that $\eta\left(\mathbf{x}^{\prime}\right) \leq 1 / 3$. This completes the proof of the claim.

## C. 2 Proof of Claim 4.5

Let $\mathbf{v}_{S}$ be the vector whose $i^{\text {th }}$ coordinate is 1 if $i \in S$ and 0 otherwise. By Lemma 4.4, there is a real univariate polynomial $p$ of degree $O(d)$ such that $p\left(\mathbf{v}_{S}^{T} \mathbf{x}\right)=1, \mathbf{v}_{S}^{T} \mathbf{x} \in J$ and $p\left(\mathbf{v}_{S}^{T} \mathbf{x}\right) \leq$ $0, \mathbf{v}_{S}^{T} \mathbf{x} \notin J$. Let $g(\mathbf{x}):=\widehat{\operatorname{ReLU}}\left(p\left(\mathbf{v}_{S}^{T} \mathbf{x}\right)\right)$. Since the absolute value of every coefficient of $p$ is at most $m^{O(d)}=\operatorname{poly}(M)$, by our definition, the total weight of the corresponding neuron $g$ is at most $m^{O(d)}=\operatorname{poly}(M)$. Therefore, there exists some $\widehat{\operatorname{ReLU}}$ function $L: \mathbb{R}^{M} \rightarrow \mathbb{R}$ such
that $g(\mathbf{x})=L\left(\mathbf{x}^{\prime}\right)=L\left(V_{O(d)}(\mathbf{x})\right)$ for all $\mathbf{x}$. We now bound the error for $L$ under the distribution $\left(\mathbf{X}^{\prime}, Y^{\prime}\right)$. By the law of total expectation, we have that

$$
\begin{aligned}
& \mathbf{E}_{\left(\mathbf{X}^{\prime}, Y^{\prime}\right)}\left[\left(Y^{\prime}-L\left(\mathbf{X}^{\prime}\right)\right)^{2}\right]=\mathbf{E}_{(\mathbf{X}, Y)}\left[(Y-g(\mathbf{X}))^{2}\right] \\
\leq & \mathbf{E}_{(\mathbf{X}, Y)}\left[(Y-g(\mathbf{X}))^{2} \mid Y=1\right]+\mathbf{E}_{(\mathbf{X}, Y)}\left[(Y-g(\mathbf{X}))^{2} \mid Y=-1\right]
\end{aligned}
$$

We note that our hard distribution returns $\left(\mathbf{X}^{\prime}, Y^{\prime}\right)$ with $Y^{\prime}=L\left(\mathbf{X}^{\prime}\right)$, unless it picked a sample corresponding to a sample of $\mathcal{D}_{-}$coming from $\bar{J}$, therefore,

$$
\mathbf{E}_{\left(\mathbf{X}^{\prime}, Y^{\prime}\right)}\left[\left(Y^{\prime}-L\left(\mathbf{X}^{\prime}\right)\right)^{2}\right] \leq \mathbf{E}_{(\mathbf{X}, Y)}\left[(Y-g(\mathbf{X}))^{2} \mid Y=1\right] \leq 4 \zeta
$$

which implies that $\mathrm{OPT}_{\text {Mass }} \leq 4 \zeta \leq \exp \left(-\Omega\left(\log (M)^{8 / 9}\right)\right)$. We then show that $\left(\mathbf{X}^{\prime}, Y^{\prime}\right)$ is a Massart single neuron distribution with $\widehat{\operatorname{ReLU}}$ activation and with noise rate upper bound of $\eta=1 / 3$. For any fixed $\mathrm{x}^{\prime} \in \mathbb{R}^{M}$, we have that

$$
\begin{aligned}
& \frac{\operatorname{Pr}_{\left(\mathbf{X}^{\prime}, Y^{\prime}\right)}\left[Y^{\prime}=-1 \mid \mathbf{X}^{\prime}=\mathbf{x}^{\prime}\right]}{\operatorname{Pr}_{\left(\mathbf{X}^{\prime}, Y^{\prime}\right)}\left[Y^{\prime}=1 \mid \mathbf{X}^{\prime}=\mathbf{x}^{\prime}\right]}=\frac{\operatorname{Pr}_{(\mathbf{X}, Y)}[Y=-1 \mid \mathbf{X}=\mathbf{x}]}{\operatorname{Pr}_{(\mathbf{X}, Y)}[Y=1 \mid \mathbf{X}=\mathbf{x}]} \\
= & \frac{\operatorname{Pr}_{(\mathbf{X}, Y)}[Y=-1] \cdot \mathbf{P r}_{(\mathbf{X}, Y)}[\mathbf{X}=\mathbf{x} \mid Y=-1]}{\operatorname{Pr}_{(\mathbf{X}, Y)}[Y=1] \cdot \mathbf{P r}_{(\mathbf{X}, Y)}[\mathbf{X}=\mathbf{x} \mid Y=1]}=\frac{\left\|\mathcal{D}_{+}\right\|_{1} \cdot \mathbf{P}_{S}^{\mathcal{D}_{+}}(\mathbf{x})}{\left\|\mathcal{D}_{-}\right\|_{1} \cdot \mathbf{P}_{S}^{\mathcal{D}_{-}}(\mathbf{x})}=\frac{\mathcal{D}_{+}\left(\mathbf{v}_{S}^{T} \mathbf{x}\right)}{\mathcal{D}_{-}\left(\mathbf{v}_{S}^{T} \mathbf{x}\right)} .
\end{aligned}
$$

Therefore, if $\mathbf{v}_{S}^{T} \mathbf{x} \in J$, the above ratio will be 0 and $L\left(\mathbf{x}^{\prime}\right)=-1$, which means that the noise rate $\eta\left(\mathbf{x}^{\prime}\right)=0$; otherwise the above ratio will be at least 2 (since $\mathcal{D}_{+}>2 \mathcal{D}_{-}$on $\bar{J}$ by property 1 (b) of Proposition 3.6 and $L\left(\mathbf{x}^{\prime}\right)=1$, which means that $\eta\left(\mathbf{x}^{\prime}\right) \leq 1 / 3$. This completes the proof of the claim.

## D SQ Hardness of Learning a Single Neuron with $L_{2}$-Massart Noise

In this section, we prove our SQ hardness result of learning a single neuron with fast convergent activations and $L_{2}$-Massart noise. Without loss of generality, we consider activations which converge on the negative side. For such an activation $f$, let $f_{-}:=f(-\infty)$ and $c_{+}$be a constant such that $f\left(c_{+}\right) \neq f_{-}$. The main theorem of this section is the following.
Theorem D. 1 (SQ Hardness of $L_{2}$-Massart Learning). Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be a fast convergent activation. Any SQ algorithm that learns a single neuron with activation $f$ on $\mathbb{R}^{M}$, in the presence of $\eta$ - $L_{2}{ }^{-}$ Massart noise with $\eta=\frac{2\left(f\left(c_{+}\right)-f_{-}\right)^{2}}{9}$, to squared error better than $1 / \operatorname{poly}(\log (M))$ requires either queries of accuracy better than $\tau:=\exp \left(-\Omega\left(\log (M)^{1.05}\right)\right)$ or at least $1 / \tau$ statistical queries. This holds even if:

1. The optimal neuron has squared error $\mathrm{OPT}_{\text {Mass-L2 }} \leq \exp \left(-\Omega\left(\log (M)^{8 / 9}\right)\right)$,
2. The $\mathbf{X}$ values are supported on $\{0,1\}^{M}$, and
3. The total weight of the neuron is poly $(M)$.

Proof. Our proof will make use of the SQ framework of Section 3.1 and will crucially rely on the one-dimensional construction of Proposition 3.6. In this section, we fix the labels $a=f_{-}, b=f\left(c_{+}\right)$, and apply the construction in Section 3.3 to obtain the joint distributions $(\mathbf{X}, Y)$ and $\left(\mathbf{X}^{\prime}, Y^{\prime}\right)$. Note that $y=y^{\prime}$ and there is a known 1-1 mapping between $\mathbf{x}$ and $\mathbf{x}^{\prime}$, therefore finding a hypothesis that predicts $y^{\prime}$ given $\mathbf{x}^{\prime}$ is equivalent to finding a hypothesis for $y$ given $\mathbf{x}$.
Claim D.2. The distribution $\left(\mathbf{X}^{\prime}, Y^{\prime}\right)$ on $\{0,1\}^{M} \times\left\{f_{-}, f\left(c_{+}\right)\right\}$is an $L_{2}$-Massart single neuron distribution with respect to activation $f$, it has optimal squared error $\mathrm{OPT}_{\mathrm{Mass}-\mathrm{L} 2} \leq$ $\exp \left(-\Omega\left(\log (M)^{8 / 9}\right)\right)$ and $L_{2}$-Massart noise rate upper bound of $\eta=\frac{2\left(f\left(c_{+}\right)-f_{-}\right)^{2}}{9}$.

Proof. We assume $M>\left|c_{+}\right|$to be sufficiently large. Let $\mathbf{v}_{S}$ be the vector whose $i^{\text {th }}$ coordinate is 1 if $i \in S$ and 0 otherwise. By Lemma 4.4 there is a real univariate polynomial $q(x)$ of degree $O(d)$ such that $q(x)=1, \forall x \in J$ and $q(x) \leq 0, \forall x \in \bar{J}$. Let $p(x)=\left(c_{+}+M\right) q(x)-M$ and $g(\mathbf{x})=f\left(p\left(\mathbf{v}_{S}^{T} \mathbf{x}\right)\right)$. By definition, we have that $p(x)=c_{+}$for $x \in J$ and $p(x) \leq-M$ for $x \in \bar{J}$.

Since the absolute value of every coefficient of $p$ is at most $m^{O(d)}=\operatorname{poly}(M)$, the weight of the corresponding neuron $g$ is at most $m^{O(d)}=\operatorname{poly}(M)$. Therefore, there exists some fast convergent activation $L: \mathbb{R}^{M} \rightarrow \mathbb{R}$ such that $g(\mathbf{x})=L\left(\mathbf{x}^{\prime}\right)=L\left(V_{O(d)}(\mathbf{x})\right)$ for all $\mathbf{x}$. We now bound the error for $L$ under the distribution $\left(\mathbf{X}^{\prime}, Y^{\prime}\right)$. We note that conditional on $Y=f_{-}$, we will always have that $\mathbf{v}_{S}^{T} \mathbf{x} \notin J$ and conditional on $Y=f\left(c_{+}\right)$, we will have that $\mathbf{v}_{S}^{T} \mathbf{x} \notin J$ with probability at most $\zeta$. Therefore, by the law of total expectation, we have that

$$
\begin{aligned}
& \mathbf{E}_{\left(\mathbf{X}^{\prime}, Y^{\prime}\right)}\left[\left(Y^{\prime}-L(\mathbf{X})\right)^{2}\right]=\mathbf{E}_{(\mathbf{X}, Y)}\left[(Y-g(\mathbf{X}))^{2}\right] \\
\leq & \mathbf{E}_{(\mathbf{X}, Y)}\left[(Y-g(\mathbf{X}))^{2} \mid Y=f_{-}\right]+\mathbf{E}_{(\mathbf{X}, Y)}\left[(Y-g(\mathbf{X}))^{2} \mid Y=f\left(c_{+}\right)\right] \\
\leq & \mathbf{E}_{(\mathbf{X}, Y)}\left[\left(f_{-}-g(\mathbf{X})\right)^{2} \mid Y=f_{-}\right]+2 \zeta \mathbf{E}_{(\mathbf{X}, Y)}\left[\left(f_{-}-f\left(c_{+}\right)\right)^{2}+\left(f_{-}-g(\mathbf{X})\right)^{2} \mid \mathbf{v}_{S}^{T} \mathbf{X} \notin J, Y=f\left(c_{+}\right)\right] \\
\leq & 1 / \operatorname{poly}(M)+2 \zeta \cdot\left(1 / \operatorname{poly}(M)+\left(f_{-}-f\left(c_{+}\right)\right)^{2}\right) \\
\leq & \exp \left(-\Omega\left(\log (M)^{8 / 9}\right)\right)+\exp \left(-\Omega\left(\log (M)^{8 / 9}\right)\right) \cdot\left(1 / \operatorname{poly}(M)+\left(f_{-}-f\left(c_{+}\right)\right)^{2}\right) \\
\leq & \exp \left(-\Omega\left(\log (M)^{8 / 9}\right)\right),
\end{aligned}
$$

where the third inequality follows from the definition of fast convergent activation. Therefore, we have that $\mathrm{OPT}_{\text {Mass-L2 }} \leq \exp \left(-\Omega\left(\log (M)^{8 / 9}\right)\right)$. We then show that $\left(\mathbf{X}^{\prime}, Y^{\prime}\right)$ is a $L_{2}$-Massart single neuron distribution with activation $f$ and with noise rate upper bound of $\eta=\frac{2\left(f\left(c_{+}\right)-f_{-}\right)^{2}}{9}$. Note that for any $\mathbf{x} \in \mathbb{R}^{m^{\prime}}$, if $\mathbf{v}_{S}^{T} \mathbf{x} \in J$, then $g(\mathbf{x})=f\left(p\left(\mathbf{v}_{S}^{T} \mathbf{x}\right)\right)=f\left(c_{+}\right)$and $Y$ will always be $f\left(c_{+}\right)$, which implies that the error will always be 0 . Hence, we assume that $\mathbf{v}_{S}^{T} \mathbf{x} \notin J$ and have that

$$
\begin{aligned}
\frac{\operatorname{Pr}_{(\mathbf{X}, Y)}\left[Y=f_{-} \mid \mathbf{X}=\mathbf{x}\right]}{\mathbf{P r}_{(\mathbf{X}, Y)}\left[Y=f\left(c_{+}\right) \mid \mathbf{X}=\mathbf{x}\right]} & =\frac{\operatorname{Pr}_{(\mathbf{X}, Y)}\left[Y=f_{-}\right] \cdot \mathbf{P r}_{(\mathbf{x}, Y)}\left[\mathbf{X}=\mathbf{x} \mid Y=f_{-}\right]}{\operatorname{Pr}_{(\mathbf{X}, Y)}\left[Y=f\left(c_{+}\right)\right] \cdot \mathbf{P r}_{(\mathbf{X}, Y)}\left[\mathbf{X}=\mathbf{x} \mid Y=f\left(c_{+}\right)\right]} \\
& =\frac{\left\|\mathcal{D}_{+}\right\|_{1} \cdot \mathbf{P}_{S}^{\mathcal{D}_{+}}(\mathbf{x})}{\left\|\mathcal{D}_{-}\right\|_{1} \cdot \mathbf{P}_{S}^{\mathcal{D}_{-}}(\mathbf{x})}=\frac{\mathcal{D}_{+}\left(\mathbf{v}_{S}^{T} \mathbf{x}\right)}{\mathcal{D}_{-}\left(\mathbf{v}_{S}^{T} \mathbf{x}\right)} \geq 2
\end{aligned}
$$

which implies that $\operatorname{Pr}_{(\mathbf{X}, Y)}\left[Y=f\left(c_{+}\right) \mid \mathbf{X}=\mathbf{x}\right] \leq 1 / 3$. Therefore,

$$
\begin{aligned}
& \mathbf{E}_{\left(\mathbf{X}^{\prime}, Y^{\prime}\right)}\left[\left(Y^{\prime}-L\left(\mathbf{X}^{\prime}\right)\right)^{2} \mid \mathbf{X}^{\prime}=\mathbf{x}^{\prime}\right]=\mathbf{E}_{(\mathbf{X}, Y)}\left[(Y-g(\mathbf{X}))^{2} \mid \mathbf{X}=\mathbf{x}\right] \\
= & \left(f\left(c_{+}\right)-g(\mathbf{x})\right)^{2} \mathbf{P r}_{(\mathbf{X}, Y)}\left[Y=f\left(c_{+}\right) \mid \mathbf{X}=\mathbf{x}\right]+\left(f_{-}-g(\mathbf{x})\right)^{2} \mathbf{P r}_{(\mathbf{X}, Y)}\left[Y=f_{-} \mid \mathbf{X}=\mathbf{x}\right] \\
\leq & \frac{\left(f\left(c_{+}\right)-g(\mathbf{x})\right)^{2}}{3}+\left(f_{-}-g(\mathbf{x})\right)^{2} \leq \frac{2\left(\left(f\left(c_{+}\right)-f_{-}\right)^{2}+\left(f_{-}-g(\mathbf{x})\right)^{2}\right)}{3}+\left(f_{-}-g(\mathbf{x})\right)^{2} \\
\leq & \frac{2\left(f\left(c_{+}\right)-f_{-}\right)^{2}}{3}+1 / \operatorname{poly}(M) \leq \frac{8\left(f\left(c_{+}\right)-f_{-}\right)^{2}}{9}
\end{aligned}
$$

where the third inequality follows from $\mathbf{v}_{S}^{T} \mathbf{x} \notin J$ and the definition of fast convergent activation. This completes the proof of the claim.

We now show that the $\left(\mathcal{D}_{+}, \mathcal{D}_{-}, f_{-}, f\left(c_{+}\right), m^{\prime}\right)$-Hidden Junta Testing Problem efficiently reduces to our learning task. In more detail, we show that any SQ algorithm that computes a hypothesis $h^{\prime}$ satisfying $\mathbf{E}_{\left(\mathbf{X}^{\prime}, Y^{\prime}\right)}\left[\left(h^{\prime}\left(\mathbf{X}^{\prime}\right)-Y^{\prime}\right)^{2}\right]<p(1-p)\left(f_{-}-f\left(c_{+}\right)\right)^{2}-2 \sqrt{2 \tau}$ can be used as a black-box to distinguish between $\mathbf{P}_{S, a, b}^{\mathcal{D}_{+}, \mathcal{D}_{-}, p}$, for some unknown subset $S \subseteq\left[m^{\prime}\right]$ with $|S|=m$, and $U_{m^{\prime}}^{p}$. Since there is a 1-1 mapping between $\mathbf{x} \in\{0,1\}^{m^{\prime}}$ and $\mathbf{x}^{\prime} \in\{0,1\}^{M}$, we denote $h:\{0,1\}^{m^{\prime}} \mapsto \mathbb{R}$ to be $h(\mathbf{x})=h^{\prime}\left(\mathbf{x}^{\prime}\right)$. We note that we can (with one additional query to estimate the $\mathbf{E}\left[\left(h^{\prime}\left(\mathbf{X}^{\prime}\right)-Y^{\prime}\right)^{2}\right]$ within error $\sqrt{2 \tau}$ ) distinguish between (i) the distribution $\mathbf{P}_{S, a, b}^{\mathcal{D}_{+}, \mathcal{D}_{-}, p}$, and (ii) the distribution $U_{m^{\prime}}^{p}$. This is because for any $h$ we have that

$$
\begin{aligned}
\mathbf{E}_{(\mathbf{X}, Y) \sim U_{m^{\prime}}^{p}}\left[(h(\mathbf{X})-Y)^{2}\right]= & \mathbf{E}_{(\mathbf{X}, Y) \sim U_{m^{\prime}}^{p}}\left[h(\mathbf{X})^{2}\right]-2 \mathbf{E}_{(\mathbf{X}, Y) \sim U_{m}^{p}}[h(\mathbf{X})] \mathbf{E}_{(\mathbf{X}, Y) \sim U^{\prime} p}[Y] \\
& +\mathbf{E}_{(\mathbf{X}, Y) \sim U_{m}^{\prime}}\left[Y^{2}\right] \\
\geq & \mathbf{E}_{(\mathbf{X}, Y) \sim U_{m^{\prime}}^{p}}[h(\mathbf{X})]^{2}-2 \mathbf{E}_{(\mathbf{X}, Y) \sim U_{m}^{p}}[h(\mathbf{X})] \mathbf{E}_{(\mathbf{X}, Y) \sim U^{\prime}}^{p}[Y] \\
& +\mathbf{E}_{(\mathbf{X}, Y) \sim U_{m^{\prime}}^{p}}\left[Y^{2}\right] \\
\geq & \mathbf{E}_{(\mathbf{X}, Y) \sim U_{m^{\prime}}^{p}}\left[Y^{2}\right]-\mathbf{E}_{(\mathbf{X}, Y) \sim U_{m^{\prime}}^{p}}[Y]^{2}=p(1-p)\left(f_{-}-f\left(c_{+}\right)\right)^{2} .
\end{aligned}
$$

Applying Proposition 3.5, we determine that any SQ algorithm which, given access to a distribution $\mathbf{P}$ so that either $\mathbf{P}=U_{m^{\prime}}^{p}$, or $\mathbf{P}$ is given by $\mathbf{P}_{S, a, b}^{\mathcal{D}_{+}, \mathcal{D}_{-}, p}$ for some unknown subset $S \subseteq\left[m^{\prime}\right]$ with $|S|=m$, correctly distinguishes between these two cases with probability at least $2 / 3$ must either make queries of accuracy better than $\sqrt{2 \tau}$ or must make at least $2^{\Omega(m)} \tau /\left(\chi^{2}(A, \operatorname{Bin}(m, 1 / 2))+\right.$ $\left.\chi^{2}(B, \operatorname{Bin}(m, 1 / 2))\right)$ statistical queries. Therefore, it is impossible for an SQ algorithm to learn a hypothesis with error better than $p(1-p)\left(f_{-}-f\left(c_{+}\right)\right)^{2}-2 \sqrt{2 \tau}=\Theta(1 / s)-\Theta(\sqrt{\tau})=1 / \operatorname{polylog}(M)$ without either using queries of accuracy better than $\tau$ or making at least $2^{\Omega(m)} \tau / \operatorname{polylog}(M)>1 / \tau$ many queries. This completes the proof of Theorem D. 1

