In Appendix A we provide heuristic justification for the scaling of the optimal minimax error rate in (77). In Appendix B we provide the proofs for Theorems 1 and 2. In Appendix Clwe provide the proofs for Theorem 4 and Proposition 6 . In Appendix D we provide the proofs for Theorem 7 . In Appendix Elwe include some useful results for the sake of completeness.

Additional notation For a matrix $A \in \mathbb{R}^{d_{1} \times d_{2}},\|A\|_{\text {op }}$ is the operator norm (with respect to Euclidean norms), and $\|A\|_{F}$ is the Frobenius norm of $A . \mathbb{S}^{d-1}:=\left\{\theta \in \mathbb{R}^{d}:\|\theta\|=1\right\}$ is the unit sphere and $\mathbb{B}^{d}:=\left\{\theta \in \mathbb{R}^{d}:\|\theta\| \leq 1\right\}$ is the unit ball in the $d$-dimensional Euclidean space. For a pair of probability distributions $P$ and $Q$ on a common alphabet $\mathcal{X}$ with densities $p$ and $q$ w.r.t. to a base measure $\nu$, we denote the total variation distance by $\mathrm{d}_{\mathrm{TV}}(P, Q):=\frac{1}{2} \int|p-q| \mathrm{d} \nu$, the Kullback-Leibler (KL) divergence by $\mathrm{D}_{\mathrm{KL}}(P \| Q):=\int p \log \frac{p}{q} \mathrm{~d} \nu$, and the chi-square divergence by $\chi^{2}(P \| Q):=\int \frac{(p-q)^{2}}{q} \mathrm{~d} \nu=\int \frac{p^{2}}{q} \mathrm{~d} \nu-1$.

## A Heuristic justification of minimax rate (7)

The main intuition behind the HMM considered in this paper comes from the correlation decay phenomenon in graphical model. Indeed, one can view the Markov chain $X_{1} \rightarrow X_{2} \rightarrow \cdots \rightarrow X_{n}$ sampled from the model (2) as a simple graphical model whose conditional independence structure is expressed by a line graph on $n$ nodes. Each node $X_{i}$ is conditionally independent of other samples given its neighbors $X_{i-1}$ and $X_{i+1}$. Furthermore, a pair of nodes $X_{i}$ and $X_{j}$ of large graph distance (i.e., $|i-j|$ ) is approximately independent. Informally, we expect that there is one sign flip (i.e., $S_{i}=-S_{i+1}$ ) per $\approx \frac{1}{\delta}$ samples. Therefore, signs $S_{i}$ and $S_{j}$ are likely to have the same value if $|i-j| \lesssim \frac{1}{\delta}$ and are approximately independent if $|i-j| \gtrsim \frac{1}{\delta}$. This is the leading guideline for the design of our estimator (14) for upper bound and the reduction to the genie-aided model (147) for lower bound.
Recall that for the Gaussian location model

$$
\begin{equation*}
X=\theta_{*}+Z \tag{28}
\end{equation*}
$$

with $Z \sim N\left(0, I_{d}\right)$, the minimax rate is given by

$$
\begin{equation*}
\mathrm{M}_{\mathrm{GLM}}(n, d, t) \asymp t \wedge \sqrt{\frac{d}{n}} \tag{29}
\end{equation*}
$$

Note that this is a compact way of writing the rate (5). Given $n$ i.i.d. samples $X_{1}^{n}$ from model (28), the estimator

$$
\begin{equation*}
\hat{\theta}=\frac{1}{n} \sum_{i=1}^{n} X_{i} \tag{30}
\end{equation*}
$$

achieves the rate (29). For the Gaussian mixture model $X=S \theta_{*}+Z$ with $S \sim \operatorname{Unif}\{-1,1\}$ and $Z \sim N\left(0, I_{d}\right)$, the minimax rate (6) [Wu and Zhou, 2019, Appendix B] can be compactly written as

$$
\begin{equation*}
\mathrm{M}_{\mathrm{GMM}}(n, d, t) \asymp\left[\frac{1}{t}\left(\sqrt{\frac{d}{n}}+\frac{d}{n}\right)+\sqrt{\frac{d}{n}}\right] \wedge t . \tag{31}
\end{equation*}
$$

The above rate is attained by the estimator

$$
\begin{equation*}
\hat{\theta}=\sqrt{\left(\lambda_{\max }(\hat{\Sigma})-1\right)_{+}} \cdot v_{\max }(\hat{\Sigma}) \tag{32}
\end{equation*}
$$

where $\hat{\Sigma}=\frac{1}{n} \sum_{i=1}^{n} X_{i} X_{i}^{\top}$. More generally, we can also consider the GMM

$$
\begin{equation*}
\tilde{X}=S \theta_{*}+\tilde{Z} \tag{33}
\end{equation*}
$$

with $S \sim \operatorname{Unif}\{-1,1\}$ and $\tilde{Z} \sim N\left(0, \sigma^{2} \cdot I_{d}\right)$, for some $\sigma>0$, which is equivalent in distribution to $\tilde{X}=\sigma\left(S \theta_{*} / \sigma+Z\right)$ where $Z \sim N\left(0, I_{d}\right)$. Generalizing (31), it is straightforward to check that the optimal minimax rate for model (33) is given by

$$
\begin{equation*}
\tilde{\mathrm{M}}_{\mathrm{GMM}}(n, d, t, \sigma) \asymp \sigma\left[\left(\frac{1}{t / \sigma}\left(\sqrt{\frac{d}{n}}+\frac{d}{n}\right)+\sqrt{\frac{d}{n}}\right) \wedge \frac{t}{\sigma}\right] \tag{34}
\end{equation*}
$$

$$
\begin{equation*}
=\left[\frac{\sigma^{2}}{t}\left(\sqrt{\frac{d}{n}}+\frac{d}{n}\right)+\sigma \sqrt{\frac{d}{n}}\right] \wedge t \tag{35}
\end{equation*}
$$

Now, for the HMM at hand, the rationale is that the original model (2) is equivalent to $\tilde{X}=\bar{S} \cdot \bar{X}$, where $\bar{S} \sim \operatorname{Unif}\{-1,1\}$ is the coherent sign of a block of $\frac{1}{\delta}$ raw samples from model $2, \bar{X}=$ $\frac{1}{1 / \delta} \sum_{i=1}^{1 / \delta} X_{i}$ represents the block-sample obtained by applying the estimator 30) to a block of $\frac{1}{\delta}$ raw samples, hence, $\bar{X} \stackrel{d}{=} \theta_{*}+N\left(0, \delta \cdot I_{d}\right)$. Furthermore, the signs across different blocks are essentially independent. Therefore

$$
\begin{equation*}
\tilde{X} \stackrel{d}{=} \bar{S} \theta_{*}+\bar{Z} \tag{36}
\end{equation*}
$$

where $\bar{Z} \sim N\left(0, \delta \cdot I_{d}\right)$, and we have $\frac{n}{1 / \delta}=n \delta$ i.i.d. block-samples from model 36). According to (35), the optimal minimax rate of the HMM (2) should be

$$
\begin{align*}
\mathrm{M}_{\mathrm{HMM}}(n, d, \delta, t) & \asymp \tilde{\mathrm{M}}_{\mathrm{GMM}}(n \delta, d, t, \sqrt{\delta})  \tag{37}\\
& \asymp\left[\frac{\delta}{t}\left(\sqrt{\frac{d}{n \delta}}+\frac{d}{n \delta}\right)+\sqrt{\delta} \sqrt{\frac{d}{n \delta}}\right] \wedge t  \tag{38}\\
& =\left[\frac{1}{t}\left(\sqrt{\frac{d \delta}{n}}+\frac{d}{n}\right)+\sqrt{\frac{d}{n}}\right] \wedge t . \tag{39}
\end{align*}
$$

Evaluating (39) immediately yields (7).
The optimal minimax error rates for GLM (cf. (5)), GMM (cf. (6), originally proved in Wu and Zhou, 2019, Appendix B]) and HMM (cf. (7), implied by our results Theorems 1 and 2) are plotted in Figure 1 .


Figure 1: Plots of minimax rates of the Gaussian Location Model, the Gaussian Mixture Model and the Hidden Markov Model. The top left, top right and bottom left figures show the scaling of the minimax rates of GLM (cf. (5)), GMM with $d \leq n$ (cf. (6)) and HMM with $d \leq \delta n$ (cf. (7)), respectively. All the above error rates are plotted in the bottom right figure in which one can clearly see how the rate varies as $\delta$ varies.

## B Proofs for Section 2; Mean estimation for known $\delta$

## B. 1 Proof of Theorem 1: Analysis of the estimator

Our proposed procedure in Theorem 1 for $\theta_{*}$ estimation with a known flip probability $\delta$ is depicted in Figure 2 .


Figure 2: A visual illustration of the construction of the estimator $\hat{\theta}_{\text {cov }}\left(X_{1}^{n} ; k\right)$ in (14) given $n$ samples $X_{1}^{n}$ from the model (2) with a known flip probability $\delta$. The block size $k$ is chosen to be $\frac{1}{8 \delta}$.

Remark 8. The condition $\delta \geq \frac{1}{n}$ in Theorem 1 is inconsequential. Indeed, if $\delta \leq \frac{1}{n}$, then also $d \geq \delta n$ and we are at the high dimension regime. In that case, the estimator can artificially increase the flip probability (by sign randomization) to $\frac{1}{n}$, which results the local minimax rate of the GLM, that is, $\mathrm{M}(n, d, \delta, t) \asymp \mathrm{M}_{\mathrm{GLM}}(n, d, t)$, which is optimal. We therefore assume $\delta \leq \frac{1}{n}$ throughout the proof.

To begin with the analysis of the estimator in Figure 2, the following lemma is a simple, yet key tool for the proof. It establishes the variance of the random gain $\bar{S}$. The proof relies on a sort of self-bounding property (cf. [Boucheron et al., 2013. Chapter 3.3]), which stems from the fact that $|\bar{S}| \leq 1$.
Lemma 9. It holds that $0 \leq \bar{S}^{2} \leq 1$ (with probability 1) and that

$$
\begin{equation*}
\mathbb{V}\left(\bar{S}^{2}\right) \leq \mathbb{E}\left[1-\bar{S}^{2}\right] \leq 4 \delta k \tag{40}
\end{equation*}
$$

Proof. Trivially $0 \leq \bar{S}^{2} \leq 1$ and so the variance is bounded as

$$
\begin{equation*}
\mathbb{V}\left(\bar{S}^{2}\right)=\mathbb{V}\left(1-\bar{S}^{2}\right) \leq \mathbb{E}\left[\left(1-\bar{S}^{2}\right)^{2}\right] \leq \mathbb{E}\left[1-\bar{S}^{2}\right] \tag{41}
\end{equation*}
$$

The expected value is then upper bounded as

$$
\begin{align*}
\mathbb{E}\left[1-\bar{S}^{2}\right] & =\mathbb{E}\left[1-\left(\frac{1}{k} \sum_{j=1}^{k} S_{j}\right)^{2}\right]  \tag{42}\\
& \left.=\mathbb{E}\left[1-\left(\frac{1}{k} \sum_{j=1}^{k} S_{j}\right)^{2}\right] S_{0}=1\right]  \tag{43}\\
& =\mathbb{E}\left[\left.\left(1-\left(\frac{1}{k} \sum_{j=1}^{k} S_{j}\right)\right)\left(1+\left(\frac{1}{k} \sum_{j=1}^{k} S_{j}\right)\right) \right\rvert\, S_{0}=1\right] \tag{44}
\end{align*}
$$

$$
\begin{align*}
& \leq 2 \cdot \mathbb{E}\left[\left.1-\left(\frac{1}{k} \sum_{j=1}^{k} S_{j}\right) \right\rvert\, S_{0}=1\right]  \tag{45}\\
& =2-\frac{2}{k} \sum_{j=1}^{k} \mathbb{E}\left[S_{j} \mid S_{0}=1\right]  \tag{46}\\
& \stackrel{(a)}{=} 2-\frac{2}{k} \sum_{j=1}^{k} \rho^{j}  \tag{47}\\
& =2-\frac{2}{k} \sum_{j=1}^{k}(1-2 \delta)^{j}  \tag{48}\\
& \stackrel{(b)}{\leq} 2-\frac{2}{k} \sum_{j=1}^{k}(1-2 j \delta)  \tag{49}\\
& =\frac{4 \delta}{k} \sum_{j=1}^{k} j  \tag{50}\\
& =\frac{2 \delta\left(k^{2}+k\right)}{k}  \tag{51}\\
& \leq 4 \delta k \tag{52}
\end{align*}
$$

where (a) follows from

$$
\begin{align*}
& \mathbb{E}\left[S_{j} \mid S_{0}=1\right] \\
& =\mathbb{P}\left[S_{j-1}=1 \mid S_{0}=1\right] \mathbb{E}\left[S_{j} \mid S_{0}=1, S_{j-1}=1\right] \\
& \quad \quad+\mathbb{P}\left[S_{j-1}=-1 \mid S_{0}=1\right] \mathbb{E}\left[S_{j} \mid S_{0}=1, S_{j-1}=-1\right]  \tag{53}\\
& =\mathbb{P}\left[S_{j-1}=1 \mid S_{0}=1\right] \mathbb{E}\left[S_{j} \mid S_{j-1}=1\right]+\mathbb{P}\left[S_{j-1}=-1 \mid S_{0}=1\right] \mathbb{E}\left[S_{j} \mid S_{j-1}=-1\right]  \tag{54}\\
& =\mathbb{P}\left[S_{j-1}=1 \mid S_{0}=1\right]\left(\frac{1+\rho}{2}-\frac{1-\rho}{2}\right)+\mathbb{P}\left[S_{j-1}=-1 \mid S_{0}=1\right]\left(-\frac{1+\rho}{2}+\frac{1-\rho}{2}\right)  \tag{55}\\
& =\mathbb{P}\left[S_{j-1}=1 \mid S_{0}=1\right] \cdot \rho-\mathbb{P}\left[S_{j-1}=-1 \mid S_{0}=1\right] \cdot \rho  \tag{56}\\
& =\rho \cdot \mathbb{E}\left[S_{j-1} \mid S_{0}=1\right]  \tag{57}\\
& =\cdots=\rho^{j} \tag{58}
\end{align*}
$$

(b) follows from Bernoulli's inequality $(1-x)^{r} \geq 1-r x$ for $x \in[0,1]$ and $r \geq 1$.

The next lemma summarizes concentration results and bounds on the expected value of various empirical quantities needed for the rest of the analysis. For a sequence of samples $\left(S_{1}^{n}, Z_{1}^{n}\right)$, let us denote the following events:

$$
\begin{gather*}
\mathcal{E}_{n, \delta, k}^{\prime}:=\left\{S_{1}^{n}:\left|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{S}_{i}^{2}-\xi_{k}\right| \leq 5 \sqrt{\frac{\delta k^{2} \log (n)}{n}}\right\},  \tag{59}\\
\mathcal{E}_{n, d, k}^{\prime \prime}:=\left\{\left(S_{1}^{n}, Z_{1}^{n}\right):\left\|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{S}_{i} \bar{Z}_{i}\right\| \leq 7 \sqrt{\frac{d}{n}}\right\},  \tag{60}\\
\mathcal{E}_{n, d, k}^{\prime \prime \prime}:=\left\{Z_{1}^{n}:\left\|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{Z}_{i} \bar{Z}_{i}^{\top}-\frac{1}{k} \cdot I_{d}\right\|_{\mathrm{op}} \leq 4 \sqrt{\frac{d}{k n}}+\frac{4 d}{n}\right\}, \tag{61}
\end{gather*}
$$

and let

$$
\begin{equation*}
\mathcal{E}_{n, d, \delta, k}:=\mathcal{E}_{n, \delta, k}^{\prime} \cap \mathcal{E}_{n, d, k}^{\prime \prime} \cap \mathcal{E}_{n, d, k}^{\prime \prime \prime} \tag{62}
\end{equation*}
$$

Note that these events depend on the choice of the block length $k$.

Lemma 10. Assume that $\delta \geq \frac{1}{n}$ and $d \leq n$. It holds that $\mathbb{P}\left[\mathcal{E}_{n, d, \delta, k}\right] \geq 1-\frac{5}{n}$ and

$$
\begin{gather*}
\mathbb{E}\left[\left|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{S}_{i}^{2}-\xi_{k}\right|\right] \leq 2 \sqrt{\frac{\delta k^{2}}{n}},  \tag{63}\\
\mathbb{E}\left[\left\|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{S}_{i} \bar{Z}_{i}\right\|\right] \leq \sqrt{\frac{d}{n}}, \tag{64}
\end{gather*}
$$

and

$$
\begin{equation*}
\mathbb{E}\left[\left\|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{Z}_{i} \bar{Z}_{i}^{\top}-\frac{1}{k} \cdot I_{d}\right\|_{o p}\right] \leq 13 \sqrt{\frac{d}{n k}}+10 \frac{d}{n} \tag{65}
\end{equation*}
$$

Proof. We analyze the expected value and concentration of each of the three terms.
First, the random variable $\bar{S}_{i}^{2}-\xi_{k}$ has zero mean, it is bounded in $[-1,1]$, and its variance is bounded by $4 \delta k$ according to Lemma 9 . Since $\left\{\bar{S}_{i}^{2}-\xi_{k}\right\}_{i=1}^{\ell}$ are i.i.d., Bernstein's inequality for bounded distributions implies that (e.g., (323) from Wainwright [2019, Proposition 2.14], by setting in the notation therein $X_{i}=\bar{S}_{i}^{2}-\xi_{k}$ which is a zero-mean random variable bounded in $[-1,1]$ )

$$
\begin{equation*}
\mathbb{P}\left[\left|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{S}_{i}^{2}-\xi_{k}\right| \geq t\right] \leq 2 \exp \left(-\frac{\ell^{2} t^{2}}{4 \ell \delta k+\ell t / 3}\right) \tag{66}
\end{equation*}
$$

Requiring that the r.h.s. is at most $\frac{2}{n}$, and using $n=\ell k$, this implies that

$$
\begin{equation*}
\left|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{S}_{i}^{2}-\xi_{k}\right| \leq \sqrt{\frac{8 \delta k^{2} \log (n)}{n}}+\frac{\frac{3}{2} k \log (n)}{n} \tag{67}
\end{equation*}
$$

with probability $1-\frac{2}{n}$. Under the assumption $\delta \geq \frac{1}{n}$, we may further upper bound

$$
\begin{align*}
\sqrt{\frac{8 \delta k^{2} \log (n)}{n}}+\frac{\frac{3}{2} k \log (n)}{n} & \leq \log (n) \cdot\left[\sqrt{\frac{8 \delta k^{2}}{n}}+\frac{\frac{3}{2} k}{n}\right]  \tag{68}\\
& \leq \log (n) \cdot\left[\sqrt{\frac{8 \delta k^{2}}{n}}+\frac{3}{2} \sqrt{\frac{\delta k^{2}}{n}}\right]  \tag{69}\\
& \leq 5 \log (n) \cdot \sqrt{\frac{\delta k^{2}}{n}} \tag{70}
\end{align*}
$$

and this implies $\mathbb{P}\left[\mathcal{E}_{n, \delta, k}^{\prime}\right] \geq 1-\frac{2}{n}$. In addition, since $\bar{S}_{i}^{2}-\xi_{k}$ are zero mean i.i.d., Lemma 9 implies that

$$
\begin{equation*}
\mathbb{E}\left[\left|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{S}_{i}^{2}-\xi_{k}\right|\right] \leq \sqrt{\mathbb{V}\left[\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{S}_{i}^{2}-\xi_{k}\right]} \leq \sqrt{\frac{4 \delta k}{\ell}}=2 \sqrt{\frac{\delta k^{2}}{n}} \tag{71}
\end{equation*}
$$

which proves 63).
Second, conditioned on any fixed $\left\{\bar{S}_{i}\right\}_{i=1}^{\ell}$, it holds that $\bar{S}_{i} \bar{Z}_{i} \sim N\left(0, \frac{\bar{S}_{i}^{2}}{k} I_{d}\right)$. Hence, $\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{S}_{i} \bar{Z}_{i} \sim$ $N\left(0, \sigma^{2} \cdot I_{d}\right)$ with

$$
\begin{equation*}
\sigma^{2}:=\frac{1}{k \ell^{2}} \sum_{i=1}^{\ell} \bar{S}_{i}^{2} \leq \frac{1}{k \ell}=\frac{1}{n} \tag{72}
\end{equation*}
$$

So, by standard concentration of norm of Gaussian (or, more generally subGaussian) random vectors (325] from [Rigollet and Hütter, 2019, Theorem 1.19] (or a degenerate case of [Wainwright, 2019, Example 6.2]), it holds with probability $1-\epsilon$ that

$$
\begin{equation*}
\left\|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{S}_{i} \bar{Z}_{i}\right\| \leq 4 \sigma \sqrt{d}+2 \sigma \sqrt{2 \log \left(\frac{1}{\epsilon}\right)} \leq 4 \sqrt{\frac{d}{n}}+\sqrt{\frac{8}{n} \log \left(\frac{1}{\epsilon}\right)} \tag{73}
\end{equation*}
$$

Since the r.h.s. does not depend on $\left\{\bar{S}_{i}\right\}_{i=1}^{\ell}$, the same bound holds unconditionally. Setting $\epsilon=\frac{1}{n}$ and further upper bounding

$$
\begin{equation*}
4 \sqrt{\frac{d}{n}}+\sqrt{\frac{8}{n} \log (n)} \leq \sqrt{\frac{32 d+16 \log (n)}{n}} \leq 7 \sqrt{\frac{d}{n}} \tag{74}
\end{equation*}
$$

using $\sqrt{a}+\sqrt{b} \leq \sqrt{2(a+b)}$ and $d \geq 2 \log (n)$, results $\mathbb{P}\left[\mathcal{E}_{n, d, k}^{\prime \prime}\right] \geq 1-\frac{1}{n}$. In addition, since $\left\{\bar{S}_{i} \bar{Z}_{i}\right\}_{i=1}^{\ell}$ are i.i.d., it holds that

$$
\begin{equation*}
\mathbb{E}\left[\left\|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{S}_{i} \bar{Z}_{i}\right\|\right] \leq \sqrt{\mathbb{E}\left[\left\|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{S}_{i} \bar{Z}_{i}\right\|^{2}\right]}=\sqrt{d \sigma^{2}} \leq \sqrt{\frac{d}{n}} \tag{75}
\end{equation*}
$$

which proves (64).
Third, by Gaussian covariance estimation from Wainwright 2019, Example 6.2], it holds for any $\eta>0$ that

$$
\begin{equation*}
V:=k\left\|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{Z}_{i} \bar{Z}_{i}^{\top}-\frac{1}{k} I_{d}\right\|_{\mathrm{op}} \leq\left[2 \sqrt{\frac{d}{\ell}}+2 \eta+\left(\sqrt{\frac{d}{\ell}}+\eta\right)^{2}\right] \tag{76}
\end{equation*}
$$

with probability larger than $1-2 e^{-\ell \eta^{2} / 2}$.
We begin with a high probability event. Setting $\epsilon=2 e^{-\ell \eta^{2} / 2}$, it holds with probability larger than $1-\epsilon$ that

$$
\begin{align*}
\left\|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{Z}_{i} \bar{Z}_{i}^{\top}-\frac{1}{k} \cdot I_{d}\right\|_{\mathrm{op}} & \leq 2 \sqrt{\frac{d}{k^{2} \ell}}+\sqrt{\frac{8 \log \left(\frac{2}{\epsilon}\right)}{k^{2} \ell}}+\left(\sqrt{\frac{d}{k \ell}}+\sqrt{\frac{2 \log \left(\frac{2}{\epsilon}\right)}{k \ell}}\right)^{2}  \tag{77}\\
& =2 \sqrt{\frac{d}{k n}}+\sqrt{\frac{8 \log \left(\frac{2}{\epsilon}\right)}{k n}}+\left(\sqrt{\frac{d}{n}}+\sqrt{\frac{2 \log \left(\frac{2}{\epsilon}\right)}{n}}\right)^{2}  \tag{78}\\
& (a)  \tag{79}\\
& \leq 2 \sqrt{\frac{d}{k n}}+\sqrt{\frac{8 \log \left(\frac{2}{\epsilon}\right)}{k n}}+\frac{2 d+4 \log \left(\frac{2}{\epsilon}\right)}{n}
\end{align*}
$$

where $(a)$ follows from $(a+b)^{2} \leq 2 a^{2}+2 b^{2}$. Setting $\epsilon=\frac{2}{n}$ and further upper bounding

$$
\begin{align*}
2 \sqrt{\frac{d}{k n}}+\sqrt{\frac{8 \log (n)}{k n}}+\frac{2 d+4 \log (n)}{n} & \leq \sqrt{\frac{8 d+16 \log (n)}{k n}}+\frac{2 d+4 \log (n)}{n}  \tag{80}\\
& \leq 4 \sqrt{\frac{d}{k n}}+\frac{4 d}{n} \tag{81}
\end{align*}
$$

using $\sqrt{a}+\sqrt{b} \leq \sqrt{2(a+b)}$ and $d \geq 2 \log (n)$, results $\mathbb{P}\left[\mathcal{E}_{n, d, k}^{\prime \prime \prime}\right] \geq 1-\frac{2}{n}$. For the bound on the expectation, we further upper bound

$$
\begin{equation*}
V \leq 2 \sqrt{\frac{d}{\ell}}+2 \eta+\frac{2 d}{\ell}+2 \eta^{2} \tag{82}
\end{equation*}
$$

Let $\alpha=2 \sqrt{\frac{d}{\ell}}+2 \frac{d}{\ell}$. Then, $U:=V-\alpha$ satisfies that $\mathbb{P}\left[U \geq 2 \eta+2 \eta^{2}\right] \leq 2 e^{-\ell \eta^{2} / 2}$. Then,

$$
\begin{align*}
\mathbb{E}[U] & =\mathbb{E}[U \cdot \mathbb{1}\{U<0\}]+\mathbb{E}[U \cdot \mathbb{1}\{U \geq 0\}]  \tag{83}\\
& \leq \mathbb{E}[U \cdot \mathbb{1}\{U \geq 0\}]  \tag{84}\\
& =\int_{0}^{\infty} \mathbb{P}[U \cdot \mathbb{1}\{U \geq 0\}>t] \mathrm{d} t  \tag{85}\\
& =\int_{0}^{\infty} \mathbb{P}[U>t] \mathrm{d} t \tag{86}
\end{align*}
$$

$$
\begin{align*}
& \stackrel{(a)}{=} \int_{0}^{\infty}(2+4 \eta) \mathbb{P}\left[U>2 \eta+2 \eta^{2}\right] \mathrm{d} \eta  \tag{87}\\
& \leq \int_{0}^{\infty}(2+4 \eta) 2 e^{-\ell \eta^{2} / 2} \mathrm{~d} \eta  \tag{88}\\
& =4 \int_{0}^{\infty} e^{-\ell \eta^{2} / 2} \mathrm{~d} \eta+8 \int_{0}^{\infty} \eta e^{-\ell \eta^{2} / 2} \mathrm{~d} \eta  \tag{89}\\
& \stackrel{(b)}{=} 4 \sqrt{\frac{2 \pi}{\ell}}+\frac{8}{\ell} \tag{90}
\end{align*}
$$

where (a) follows from the change of variables $t=2 \eta+2 \eta^{2}$, and (b) follows from Gaussian integration (the first term). Hence,

$$
\begin{equation*}
\mathbb{E}\left[\frac{V}{k}\right] \leq \frac{1}{k} \mathbb{E}[U+\alpha] \leq 13 \sqrt{\frac{d}{n k}}+10 \frac{d}{n} \tag{91}
\end{equation*}
$$

which proves $\left[65\right.$. The claim $\mathbb{P}\left[\mathcal{E}_{n, d, \delta, k}\right] \geq 1-\frac{5}{n}$ then follows from the union bound over the three events considered above.

The next lemma utilizes Lemma 10 to establish concentration on the maximal eigenvalue and the associated eigenvector.
Lemma 11. Let

$$
\begin{equation*}
\psi(n, d, \delta, k):=2 \sqrt{\frac{\delta k^{2}}{n}} \cdot\left\|\theta_{*}\right\|^{2}+2 \sqrt{\frac{d}{n}} \cdot\left\|\theta_{*}\right\|+13 \sqrt{\frac{d}{n k}}+10 \frac{d}{n} \tag{92}
\end{equation*}
$$

Then, it holds that

$$
\begin{equation*}
\mathbb{E}\left[\left|\lambda_{\max }\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right)-\left(\xi_{k}\left\|\theta_{*}\right\|^{2}+\frac{1}{k}\right)\right|\right] \leq \psi(n, d, \delta, k) \tag{93}
\end{equation*}
$$

and

$$
\begin{equation*}
\mathbb{E}\left[\operatorname{loss}\left(v_{\max }\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right), \frac{\theta_{*}}{\left\|\theta_{*}\right\|}\right)\right] \leq \frac{8 \cdot \psi(n, d, \delta, k)}{\left\|\theta_{*}\right\|^{2}} \tag{94}
\end{equation*}
$$

If $\mathcal{E}_{n, d, \delta, k}$ holds then similar bounds hold (with probability 1) with $\psi(n, d, \delta, k)$ multiplied by $7 \log (n)$.
Proof. Note that $\lambda_{\max }\left(\Sigma_{n, k}\left(\theta_{*}\right)\right)=\xi_{k} \cdot\left\|\theta_{*}\right\|^{2}+\frac{1}{k}$. Then,

$$
\begin{align*}
& \left|\lambda_{\max }\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right)-\left(\xi_{k}\left\|\theta_{*}\right\|^{2}+\frac{1}{k}\right)\right| \\
& =\left|\lambda_{\max }\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right)-\lambda_{\max }\left(\Sigma_{n, k}\left(\theta_{*}\right)\right)\right|  \tag{95}\\
& \stackrel{(a)}{\leq}\left\|\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)-\Sigma_{n, k}\left(\theta_{*}\right)\right\|_{\mathrm{op}}  \tag{96}\\
& =\left\|\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)-\xi_{k} \theta_{*} \theta_{*}^{\top}-\frac{1}{k} I_{d}\right\|_{\mathrm{op}}  \tag{97}\\
& \leq\left|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{S}_{i}^{2}-\xi_{k}\right| \cdot\left\|\theta_{*}\right\|^{2}+\left\|\frac{2}{\ell} \sum_{i=1}^{\ell} \bar{S}_{i} \bar{Z}_{i}\right\| \cdot\left\|\theta_{*}\right\|+\left\|\frac{1}{\ell} \sum_{i=1}^{\ell} \bar{Z}_{i} \bar{Z}_{i}^{\top}-\frac{1}{k} \cdot I_{d}\right\|_{\mathrm{op}} \tag{98}
\end{align*}
$$

where (a) follows from Weyl's inequality. Taking expectation and using Lemma 10 proves the claimed bound.
Next, note that $v_{\max }\left(\Sigma_{n, k}\left(\theta_{*}\right)\right)=\frac{\theta_{*}}{\left\|\theta_{*}\right\|}$, its eigenvalue is $\xi_{k}\left\|\theta_{*}\right\|^{2}+\frac{1}{k}$, and that all the $(d-1)$ other eigenvalues of $\Sigma_{n, k}\left(\theta_{*}\right)$ are $\frac{1}{k}$. Thus, the eigen-gap of the maximal eigenvalue is $\xi_{k}\left\|\theta_{*}\right\|^{2}$. Then, by Davis-Kahan's perturbation bound (328) from [Vershynin, 2018, Theorem 4.5.5]

$$
\operatorname{loss}\left(v_{\max }\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right), \frac{\theta_{*}}{\left\|\theta_{*}\right\|}\right)=\operatorname{loss}\left(v_{\max }\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right), v_{\max }\left(\Sigma_{n, k}\left(\theta_{*}\right)\right)\right)
$$

$$
\begin{aligned}
& \leq 4 \frac{\left\|\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)-\Sigma_{n, k}\left(\theta_{*}\right)\right\|_{\mathrm{op}}}{\xi_{k}\left\|\theta_{*}\right\|^{2}} \\
& \stackrel{(a)}{\leq} 8 \frac{\left\|\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)-\Sigma_{n, k}\left(\theta_{*}\right)\right\|_{\mathrm{op}}}{\left\|\theta_{*}\right\|^{2}}
\end{aligned}
$$

where (a) follows since $\xi_{k}=\mathbb{E}\left[\bar{S}^{2}\right] \geq 1-4 k \delta=\frac{1}{2}$ according to Lemma 9 and the choice $k=\frac{1}{8 \delta}$. The operator norm is upper bounded as for the maximal eigenvalue. The proof then follows by taking expectation, and utilizing Lemma 10
Finally, assuming the event $\mathcal{E}_{n, d, \delta, k}$ holds, Lemma 10 implies that the same bound holds with $7 \log (n) \cdot \psi(n, d, \delta, k)$ instead of $\psi(n, d, \delta, k)$.

The next lemma upper bounds $\hat{\theta}$ in expectation and in high probability whenever $\left\|\theta_{*}\right\|$ is below the minimax rate $\beta(n, d, \delta)=\Theta\left(\sqrt{\frac{d}{n}} \vee\left(\frac{\delta d}{n}\right)^{1 / 4}\right)$, and show that both are (roughly) at most on the scale of $\beta(n, d, \delta)$.
Lemma 12. If $\left\|\theta_{*}\right\| \leq \beta(n, d, \delta)$ and $k=\frac{1}{8 \delta}$ then

$$
\begin{equation*}
\mathbb{E}\left[\left\|\hat{\theta}_{c o v}\left(X_{1}^{n} ; k\right)\right\|\right] \leq 11 \cdot \beta(n, d, \delta) \tag{99}
\end{equation*}
$$

and if the event $\mathcal{E}_{n, d, \delta, k}$ holds then

$$
\begin{equation*}
\left\|\hat{\theta}_{c o v}\left(X_{1}^{n} ; k\right)\right\| \leq 77 \log (n) \cdot \beta(n, d, \delta) \tag{100}
\end{equation*}
$$

Proof. It holds that

$$
\begin{align*}
\left\|\hat{\theta}_{\mathrm{cov}}\left(X_{1}^{n} ; k\right)\right\|^{2} & =\frac{\lambda_{\max }\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right)-\frac{1}{k}}{\xi_{k}} \vee 0  \tag{101}\\
& \stackrel{(a)}{\leq} 2\left[\lambda_{\max }\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right)-\frac{1}{k}\right] \vee 0  \tag{102}\\
& \leq 2\left[\lambda_{\max }\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right)-\xi_{k}\left\|\theta_{*}\right\|^{2}-\frac{1}{k}\right] \vee 0+2 \xi_{k}\left\|\theta_{*}\right\|^{2}  \tag{103}\\
& \stackrel{(b)}{\leq} 2\left[\lambda_{\max }\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right)-\xi_{k}\left\|\theta_{*}\right\|^{2}-\frac{1}{k}\right] \vee 0+2\left\|\theta_{*}\right\|^{2} \tag{104}
\end{align*}
$$

where (a) follows since $\xi_{k} \geq \frac{1}{2}$ and (b) follows since $\xi_{k} \leq 1$. Taking expectation of both sides, and utilizing Lemma 11 results

$$
\begin{align*}
& \mathbb{E}\left[\left\|\hat{\theta}_{\text {cov }}\left(X_{1}^{n} ; k\right)\right\|^{2}\right] \\
& \leq 2 \cdot \psi(n, d, \delta, k)+2\left\|\theta_{*}\right\|^{2}  \tag{105}\\
& \leq 4 \sqrt{\frac{\delta k^{2}}{n}} \cdot\left\|\theta_{*}\right\|^{2}+4 \sqrt{\frac{d}{n}} \cdot\left\|\theta_{*}\right\|+26 \sqrt{\frac{d}{n k}}+20 \frac{d}{n}+2\left\|\theta_{*}\right\|^{2}  \tag{106}\\
& \stackrel{(a)}{=} 4 \sqrt{\frac{1}{64 \delta n}} \cdot\left\|\theta_{*}\right\|^{2}+4 \sqrt{\frac{d}{n}} \cdot\left\|\theta_{*}\right\|+26 \sqrt{\frac{8 \delta d}{n}}+20 \frac{d}{n}+2\left\|\theta_{*}\right\|^{2}  \tag{107}\\
& \stackrel{(b)}{\leq} 4 \sqrt{\frac{1}{64 \delta n}} \cdot\left(\frac{2 d}{n}+2 \sqrt{\frac{\delta d}{n}}\right)+4 \sqrt{\frac{d}{n}} \cdot\left(\sqrt{\frac{d}{n}}+\left(\frac{\delta d}{n}\right)^{1 / 4}\right) \\
& \quad+26 \sqrt{\frac{8 \delta d}{n}}+20 \frac{d}{n}+2\left(\frac{2 d}{n}+2 \sqrt{\frac{\delta d}{n}}\right)  \tag{108}\\
& \leq 4 \sqrt{\frac{1}{64 \delta n}} \cdot\left(\frac{2 d}{n}+2 \sqrt{\frac{\delta d}{n}}\right)+4 \sqrt{\frac{d}{n}} \cdot\left(\frac{\delta d}{n}\right)^{1 / 4}+78 \sqrt{\frac{\delta d}{n}}+28 \frac{d}{n} \tag{109}
\end{align*}
$$

$$
\begin{align*}
& =\frac{d}{\sqrt{\delta} n^{3 / 2}}+\sqrt{\frac{d}{n^{2}}}+4 \sqrt{\frac{d}{n}} \cdot\left(\frac{\delta d}{n}\right)^{1 / 4}+78 \sqrt{\frac{\delta d}{n}}+28 \frac{d}{n}  \tag{110}\\
& (c) \text { ) } 4 \sqrt{\frac{d}{n}} \cdot\left(\frac{\delta d}{n}\right)^{1 / 4}+78 \sqrt{\frac{\delta d}{n}}+30 \frac{d}{n}  \tag{111}\\
& \text { (d) } 80 \sqrt{\frac{\delta d}{n}}+32 \frac{d}{n} \tag{112}
\end{align*}
$$

where ( $a$ ) follows by setting $k=\frac{1}{8 \delta}$, and (b) follows from the assumption $\left\|\theta_{*}\right\| \leq \beta(n, d, \delta) \leq$ $\sqrt{\frac{d}{n}}+\left(\frac{\delta d}{n}\right)^{1 / 4}\left(\right.$ and $\left.(a+b)^{2} \leq 2 a^{2}+2 b^{2}\right),(c)$ follows from $\delta \geq \frac{1}{n}$, so that $\frac{d}{\sqrt{\delta} n^{3 / 2}} \leq \frac{d}{n}$, and $(d)$ follows from $a b \leq \frac{1}{2} a^{2}+\frac{1}{2} b^{2}$ applied to the first term.
The proof is then completed by Jensen's inequality

$$
\begin{align*}
\mathbb{E}\left[\left\|\hat{\theta}_{\mathrm{cov}}\left(X_{1}^{n} ; k\right)\right\|\right] & \leq \sqrt{\mathbb{E}\left[\left\|\hat{\theta}_{\mathrm{cov}}\left(X_{1}^{n} ; k\right)\right\|^{2}\right]}  \tag{113}\\
& \leq \sqrt{80 \sqrt{\frac{\delta d}{n}}+32 \frac{d}{n}}  \tag{114}\\
& \leq \sqrt{112 \cdot\left(\sqrt{\frac{\delta d}{n}} \vee \frac{d}{n}\right)}  \tag{115}\\
& \leq 11 \cdot \beta(n, d, \delta) . \tag{116}
\end{align*}
$$

If $\lambda_{\text {max }}\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right)-\frac{1}{k} \leq 0$ then $\hat{\theta}_{\text {cov }}\left(X_{1}^{n} ; k\right)=0$ and the claim is trivial. Otherwise, from the same reasoning as above, assuming that the event $\mathcal{E}_{n, d, \delta, k}$ holds, it also holds from Lemma 11 that $\left\|\hat{\theta}_{\text {cov }}\left(X_{1}^{n} ; k\right)\right\| \leq 77 \log (n) \cdot \beta(n, d, \delta)$.

We may now prove Theorem 1 .
Proof of Theorem 1] If $\left\|\theta_{*}\right\| \leq \beta(n, d, \delta)$ then Lemma 12 implies that

$$
\begin{equation*}
\mathbb{E}\left[\left\|\hat{\theta}_{\mathrm{cov}}\left(X_{1}^{n} ; k\right)\right\|\right] \leq 11 \cdot \beta(n, d, \delta) \tag{117}
\end{equation*}
$$

Otherwise, assume that $\left\|\theta_{*}\right\| \geq \beta(n, d, \delta)$. For any estimator $\tilde{\theta}$ we may write $\tilde{\theta}=\|\tilde{\theta}\| \cdot v$ where $v \in \mathbb{S}^{d-1}$, it holds that

$$
\begin{align*}
\operatorname{loss}\left(\tilde{\theta}, \theta_{*}\right) & =\| \| \tilde{\theta}\left\|\cdot v-\theta_{*}\right\|  \tag{118}\\
& =\| \| \tilde{\theta}\|\cdot v-\| \theta_{*}\|\cdot v+\| \theta_{*}\left\|\cdot v-\theta_{*}\right\|  \tag{119}\\
& \leq\left|\|\tilde{\theta}\|-\left\|\theta_{*}\right\|\right|+\left\|\theta_{*}\right\| \cdot \operatorname{loss}\left(v, \frac{\theta_{*}}{\left\|\theta_{*}\right\|}\right) \tag{120}
\end{align*}
$$

where the last inequality follows from the triangle inequality. Specifying this result to $\tilde{\theta}=\hat{\theta} \equiv$ $\hat{\theta}_{\text {cov }}\left(X_{1}^{n} ; k\right)$, results

$$
\begin{equation*}
\operatorname{loss}\left(\hat{\theta}, \theta_{*}\right) \leq\left|\|\hat{\theta}\|-\left\|\theta_{*}\right\|\right|+\left\|\theta_{*}\right\| \cdot \operatorname{loss}\left(v_{\max }\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right), \frac{\theta_{*}}{\left\|\theta_{*}\right\|}\right) \tag{121}
\end{equation*}
$$

Lemma 11 implies that the first term in 121 is bounded as

$$
\begin{align*}
\left|\|\hat{\theta}\|-\left\|\theta_{*}\right\|\right| & =\frac{\left|\|\hat{\theta}\|^{2}-\left\|\theta_{*}\right\|^{2}\right|}{\|\hat{\theta}\|+\left\|\theta_{*}\right\|}  \tag{122}\\
& \leq \frac{\|\hat{\theta}\|^{2}-\left\|\theta_{*}\right\|^{2} \mid}{\left\|\theta_{*}\right\|} \tag{123}
\end{align*}
$$

$$
\begin{align*}
& =\left|\frac{\left(\lambda_{\max }\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right)-\frac{1}{k}\right)_{+}-\xi_{k}\left\|\theta_{*}\right\|^{2}}{\xi_{k}\left\|\theta_{*}\right\|}\right|  \tag{124}\\
& \leq \frac{\left|\lambda_{\max }\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right)-\frac{1}{k}-\xi_{k}\left\|\theta_{*}\right\|^{2}\right|}{\xi_{k}\left\|\theta_{*}\right\|}  \tag{125}\\
& \leq \frac{2\left|\lambda_{\max }\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right)-\frac{1}{k}-\xi_{k}\left\|\theta_{*}\right\|^{2}\right|}{\left\|\theta_{*}\right\|} \tag{126}
\end{align*}
$$

where the last inequality follows since $\xi_{k} \geq 1-4 k \delta$ (from Lemma 9 ), and $k \leq \frac{1}{8 \delta}$. Taking expectation of both sides, and utilizing Lemma 11 results

$$
\begin{equation*}
\mathbb{E}\left[\left|\|\hat{\theta}\|-\left\|\theta_{*}\right\|\right|\right] \leq \frac{2 \cdot \psi(n, d, \delta, k)}{\left\|\theta_{*}\right\|} \tag{127}
\end{equation*}
$$

Lemma 11 further implies that the second term in 121 is bounded as

$$
\begin{equation*}
\mathbb{E}\left[\left\|\theta_{*}\right\| \cdot \operatorname{loss}\left(v_{\max }\left(\hat{\Sigma}_{n, k}\left(X_{1}^{n}\right)\right), \frac{\theta_{*}}{\left\|\theta_{*}\right\|}\right)\right] \leq \frac{8 \cdot \psi(n, d, \delta, k)}{\left\|\theta_{*}\right\|} \tag{128}
\end{equation*}
$$

Hence, (121)

$$
\begin{align*}
\operatorname{loss}\left(\hat{\theta}, \theta_{*}\right) & \leq \frac{10 \cdot \psi(n, d, \delta, k)}{\left\|\theta_{*}\right\|}  \tag{129}\\
& \leq 20 \sqrt{\frac{\delta k^{2}}{n}} \cdot\left\|\theta_{*}\right\|+20 \sqrt{\frac{d}{n}}+\frac{130}{\left\|\theta_{*}\right\|} \sqrt{\frac{d}{n k}}+100 \frac{1}{\left\|\theta_{*}\right\|} \frac{d}{n}  \tag{130}\\
& \stackrel{(a)}{=} 20 \sqrt{\frac{1}{64 \delta n}} \cdot\left\|\theta_{*}\right\|+20 \sqrt{\frac{d}{n}}+\frac{130}{\left\|\theta_{*}\right\|} \sqrt{\frac{8 \delta d}{n}}+\frac{100}{\left\|\theta_{*}\right\|} \frac{d}{n}  \tag{131}\\
& \stackrel{(b)}{\leq} 20 \sqrt{\frac{d}{n}}+\frac{130}{\left\|\theta_{*}\right\|} \sqrt{\frac{8 \delta d}{n}}+\frac{110}{\left\|\theta_{*}\right\|} \frac{d}{n} \tag{132}
\end{align*}
$$

where ( $a$ ) follows from setting $k=\frac{1}{8 \delta}$, and (b) follows since for $\left\|\theta_{*}\right\| \geq \beta(n, d, \delta)$,

$$
\begin{align*}
20 \sqrt{\frac{1}{64 \delta n}} \cdot\left\|\theta_{*}\right\| & =20 \frac{\left\|\theta_{*}\right\|^{2}}{\left\|\theta_{*}\right\|} \sqrt{\frac{1}{64 \delta n}}  \tag{133}\\
& \leq 20 \cdot \frac{\left(\sqrt{\frac{d}{n}}+\left(\frac{\delta d}{n}\right)^{1 / 4}\right)^{2}}{\left\|\theta_{*}\right\|} \cdot \sqrt{\frac{1}{64 \delta n}}  \tag{134}\\
& \leq 5 \cdot \frac{\frac{d}{n}+\sqrt{\frac{\delta d}{n}}}{\left\|\theta_{*}\right\|} \cdot \sqrt{\frac{1}{\delta n}}  \tag{135}\\
& =5 \cdot \frac{1}{\left\|\theta_{*}\right\|}\left(\frac{d}{n^{3 / 2} \sqrt{\delta}}+\frac{\sqrt{d}}{n}\right)  \tag{136}\\
& \leq \frac{5}{\left\|\theta_{*}\right\|} \frac{d}{n^{3 / 2} \sqrt{\delta}}+\frac{5}{\left\|\theta_{*}\right\|} \frac{d}{n}  \tag{137}\\
& \leq \frac{10}{\left\|\theta_{*}\right\|} \frac{d}{n} \tag{138}
\end{align*}
$$

where the last inequality follows since $\delta \geq \frac{1}{n}$.

## B. 2 Proof of Theorem 2; Impossibility lower bound

The proof's main ideas The estimation error for the Markov model cannot be better than a genieaided model for which the sign is known during a block whose size is larger than the mixing time
$\Theta\left(\frac{1}{\delta}\right)$ of the original Markov chain. The estimator can then align the signs of the samples in each block, and thus reduce the noise variance by a factor of $\delta$. This effectively reduces the problem into a Markov model whose flip probability is asymptotically close to $1 / 2$, that is, the resulting model is close to a GMM. This aforementioned closeness is quantified by a uniform bound on the ratio between the probability distributions. Based on the closeness of the models, the known lower bound for the GMM model [Wu and Zhou, 2019, Appendxi B] implies a bound on the estimation error in the Markov model.

The proof follows from an application of Fano's method [Wainwright, 2019, Section 15.3] [Yang and Barron, 1999], where, as usual, the main technical challenge follows from the bounding of the mutual information term. We first begin with a brief description of the lower bound on the estimation error, and then explain how it is solved in the Gaussian mixture case in [Wu and Zhou, 2019, Appendix B]. We then describe our analysis method for the Markov case.

Construction of a packing set The version of Fano's method that we use is based on constructing a packing set $\Theta_{M}:=\left\{\theta_{m}\right\}_{m \in[M]} \subset \mathbb{R}^{d}$ and an additional center vector $\theta_{0} \in \mathbb{R}^{d}$. Here, the set $\Theta_{M}$ packs points in a spherical cap of a fixed angle with a center at $\theta_{0}$. To construct this set, we let $\Phi_{M}:=\left\{\phi_{m}\right\}_{m \in[M]} \subset \mathbb{B}^{d-2}$ be a $\frac{1}{16}$-packing set of $\mathbb{B}^{d-2}$ of size larger than $M \geq 16^{d-2}$ in the Euclidean distance, whose existence is assured by a standard argument (e.g., 329 from Wainwright. 2019, Lemma 5.7 and Example 5.8]). We then append to each $\phi_{m}$ another coordinate to obtain $\overline{\bar{\phi}}_{m}:=\left(\phi_{m}, \sqrt{1-\left\|\phi_{m}\right\|}\right) \in \mathbb{S}^{d-2}$. Then, $\bar{\Phi}_{M}:=\left\{\bar{\phi}_{m}\right\}_{m \in[M]} \subset \mathbb{S}^{d-2}$ is a $\frac{1}{16}$-packing set of $\mathbb{S}^{d-2}$ of size $M$. Now, the packing set $\bar{\Phi}_{M}$ packs points in the Euclidean distance, however, the loss function $\operatorname{loss}(\cdot, \cdot)$ in $\sqrt{3}$ is sign-insensitive, and so the packing set requires further dilution. Specifically, there must exist an orthant in $\mathbb{R}^{d-1}$ which contains at least a $2^{-(d-1)}$ fraction of the points in $\bar{\Phi}_{M}$. So, there must also exist a rotation matrix $O \in \mathbb{R}^{(d-1) \times(d-1)}$ of $\bar{\Phi}_{M}$ so that $\tilde{\Phi}_{M}:=O \bar{\Phi}_{M} \cap \mathbb{R}_{+}^{d-1}=$ $\left\{O \bar{\phi}_{m}\right\}_{m \in[M]} \cap \mathbb{R}_{+}^{d-1}$ has at least $\frac{\left|\bar{\Phi}_{M}\right|}{2^{d-1}} \geq \frac{1}{16} \cdot 8^{d-1}$ points (with $\mathbb{R}_{+}^{d-1}$ being the positive orthant). Furthermore, for any distinct $\tilde{\phi}_{m}, \tilde{\phi}_{m^{\prime}} \in \tilde{\Phi}_{M}$ it holds that $\operatorname{loss}\left(\tilde{\phi}_{m}, \tilde{\phi}_{m^{\prime}}\right)=\left\|\tilde{\phi}_{m}-\tilde{\phi}_{m^{\prime}}\right\|$. We now choose $\epsilon \in(0,1)$, set $\tilde{\theta}_{0}=[1,0, \ldots, 0] \in \mathbb{R}^{d}$ to be the center vector (this choice is arbitrary and made for convenience), and $\tilde{\theta}_{m}=\left[\sqrt{1-\epsilon^{2}}, \epsilon \tilde{\phi}_{m}\right] \in \mathbb{R}^{d}$, for $m \in[M]$. Evidently, the angle between any $\tilde{\theta}_{m}$ and $\tilde{\theta}_{0}$ is fixed for all $m \in[M]$. Finally, given the prescribed norm $t$ in the statement of the theorem, we set the packing set that will be next used in Fano's-inequality based argument as $\Theta_{M}:=\left\{t \cdot \tilde{\theta}_{m}\right\}_{m \in[M]}$ and set $\theta_{0}=t \tilde{\theta}_{0}$. To summarize, $\Theta_{M}$ satisfies the following properties: (i) $\left|\Theta_{M}\right| \geq \frac{1}{16} \cdot 8^{d-1}$, (ii) for any $\theta_{m} \in \Theta_{M}$ it holds that $\left\|\theta_{m}\right\|=t$, (iii) for any distinct $\theta_{m}, \theta_{m^{\prime}} \in \Theta_{M}$ it holds that $\operatorname{loss}\left(\theta_{m}, \theta_{m^{\prime}}\right)=\left\|\theta_{m}-\theta_{m^{\prime}}\right\| \geq \frac{1}{16} \epsilon t$, (iv) for any $\theta_{m} \in \Theta_{m}$ it holds that

$$
\begin{equation*}
\operatorname{loss}\left(\theta_{m}, \theta_{0}\right)=\left\|\theta_{m}-\theta_{0}\right\|=t \cdot \sqrt{\left(\sqrt{1-\epsilon^{2}}-1\right)^{2}+\epsilon^{2}} \leq t \cdot\left[\sqrt{1-\epsilon^{2}}-1+\epsilon\right] \leq 2 t \epsilon \tag{139}
\end{equation*}
$$

where the first inequality is from $\sqrt{a^{2}+b^{2}} \leq a+b$ for $a, b \in \mathbb{R}_{+}$and the second inequality from $\sqrt{1-\epsilon^{2}} \leq 1+\epsilon$.

Fano's inequality based lower bound Recall that $P_{\theta}^{(n)}$ is the probability distribution of $X_{1}^{n}$ under the Markov model $X_{i}=S_{i} \theta+Z_{i}, i \in[n]$ with $S_{0}^{n}$ is as in 11 . Let $J \sim \operatorname{Unif}[M]$ and assume that given a prescribed norm $t>0$, it holds that $X_{1}^{n} \mid J=j \sim P_{\theta_{j}}^{(n)}$. Then, based on the four properties of the packing set $\Theta_{M}$, Fano's method states that [Wainwright 2019, Proposition 15.2]

$$
\begin{align*}
\mathrm{M}(n, d, \delta, t) & \geq \frac{\epsilon t}{32} \cdot\left[1-\frac{I\left(J ; X_{1}^{n}\right)+\log 2}{\log M}\right]  \tag{140}\\
& \geq \frac{\epsilon t}{32} \cdot\left[1-\frac{I\left(J ; X_{1}^{n}\right)+\log 2}{(d-1) \log 8}\right]  \tag{141}\\
& \stackrel{(a)}{\geq} \frac{\epsilon t}{32} \cdot\left[\frac{1}{2}-4 \cdot \frac{I\left(J ; X_{1}^{n}\right)}{d}\right]  \tag{142}\\
& \stackrel{(b)}{\geq} \frac{\epsilon t}{32} \cdot\left[\frac{1}{2}-4 \cdot \frac{\max _{m \in[M]} \mathrm{D}_{\mathrm{KL}}\left(P_{\theta_{m}}^{(n)} \| P_{\theta_{0}}^{(n)}\right)}{d}\right] \tag{143}
\end{align*}
$$

where $I\left(J ; X_{1}^{n}\right)$ is the mutual information between $J$ and $X_{1}^{n}$, and in $(a)$ we have used the assumption $d \geq 3$ which implies $(d-1) \log 8-\log 16 \geq \frac{d}{4}$, and in $(b)$ we have used the standard "informationradius" bound on the mutual information (e.g., (15.52) in [Wainwright, 2019. Proof of Lemma 15.21 and Exercise 15.11]). The crux of the proof is to establish a bound of the form $\mathrm{D}_{\mathrm{KL}}\left(P_{\theta_{m}}^{(n)} \| P_{\theta_{0}}^{(n)}\right) \leq$ $\frac{d}{16}$ for a given $\epsilon_{*}>0$ to obtain a lower bound of $\Omega\left(\epsilon_{*} t\right)$. Note that due to symmetry of the packing set, the last KL divergence is the same for all $m \in[M]$.
Before continuing the proof for the Markov case with $\delta<\frac{1}{2}$, we next describe the analysis of the Gaussian mixture model in [Wu and Zhou 2019]. For the Gaussian mixture model (that is, a degenerate Markov model with $\delta=\frac{1}{2}$ ), $P_{\theta}^{(n)}$ is an i.i.d. distribution, and so the tensorization property of the KL divergence immediately implies that $\mathrm{D}_{\mathrm{KL}}\left(P_{\theta_{m}}^{(n)} \| P_{\theta_{0}}^{(n)}\right)=n \cdot \mathrm{D}_{\mathrm{KL}}\left(P_{\theta_{m}} \|\right.$ $P_{\theta_{0}}$ ) (where $P_{\theta} \equiv P_{\theta}^{(1)}$ ). Then, it was established in [Wu and Zhou, 2019, Lemma 27] that $\mathrm{D}_{\mathrm{KL}}\left(P_{\theta_{m}} \| P_{\theta_{0}}\right) \lesssim t^{2} \cdot\left\|\theta_{0}-\theta_{m}\right\|^{2}$ as follows. First, it was noted that under $\theta_{0}=[t, 0, \ldots, 0] \in \mathbb{R}^{d}$, $P_{\theta}=P_{t} \otimes N\left(0, I_{d-1}\right)$, that is a product distribution, with the first coordinate being one-dimensional Gaussian mixture with means $\pm t$, and all the other $d-1$ coordinates being standard Gaussian. Based on this and the chain rule of the KL divergence, the KL divergence $\mathrm{D}_{\mathrm{KL}}\left(P_{\theta_{m}} \| P_{\theta_{0}}\right)$ was evaluated separately by the KL divergence between the distribution of the first coordinate, and the KL divergence between the distributions of the other $d-1$ coordinates, conditioned on the first coordinates. Each of these two KL divergences was bounded by the corresponding chi-square divergences, and was further shown to be $O\left(t^{2} \cdot\left\|\theta_{0}-\theta_{m}\right\|^{2}\right)$. Finally, the minimax lower bound stated in 6 was obtained by setting $\epsilon=c \cdot \min \left\{1, \frac{1}{t^{2}} \sqrt{\frac{d}{n}}\right\}$ for some small enough $c>0$ in the construction of the packing set.
For the Markov model (with $\delta<\frac{1}{2}$ ), it seems rather cumbersome to bound $\mathrm{D}_{\mathrm{KL}}\left(P_{\theta_{m}}^{(n)} \| P_{\theta_{0}}^{(n)}\right)$ since the model $P_{\theta}^{(n)}$ has memory. Thus we next propose an indirect approach, which requires three preliminary steps. At the first step, we reduce, via a genie-based argument, the original GaussianMarkov model with $n$ samples to a Gaussian-Markov model with $\ell$ samples, where each of these $\ell$ sample is a coherent average of a block of $k$ consecutive samples, $\ell=\frac{n}{k}$ (in a similar, yet not identical, form to the estimator from Section 2 . The sequences of signs underlying each of these $\ell$ samples also forms a Markov model, where the block length $k$ is judiciously chosen (roughly) as $\Theta\left(\frac{1}{\delta}\right)$ so that the dependency between the signs is much weaker, in the sense that the flip probability $\bar{\delta}$ of its signs asymptotically tends to $\frac{1}{2}$ as $\delta \rightarrow 0$. Since for Gaussian mixture models the flip probability is exactly $\frac{1}{2}$, this Markov model is approximately a Gaussian mixture model.
At the second step, we develop a change-of-measure argument. As said, the reduced model is only approximately Gaussian mixture ( $\bar{\delta} \approx \frac{1}{2}$ ); Had it was exactly a Gaussian mixture model $\left(\bar{\delta}=\frac{1}{2}\right)$ then the bound on the KL divergence from [Wu and Zhou, 2019, Lemma 27] could have been directly used. Nonetheless, since $\bar{\delta} \rightarrow \frac{1}{2}$ with our choice of $k$, the probability distribution of the samples tends to that of the memoryless Gaussian mixture. Since divergences are essentially a continuous functions of the measured distributions, we expect that the KL divergence for the Markov model with $\bar{\delta}$ is close to the KL divergence for the Gaussian mixture model. Lemma 13 forms the basis of this change-of-measure argument in terms of the probability distribution of the signs, and Lemma 14 provides the corresponding change-of-measure bound for the KL divergence.

At the third step, we bound the chi-square divergence for a $d$ dimensional Gaussian mixture model. As described above, in [Wu and Zhou, 2019, Lemma 27] the KL divergence was upper bounded for the Gaussian mixture model by first separating to the KL divergence in the first coordinate and the KL divergence in all other $d-1$ coordinates, exploiting the chain rule of the KL divergence, and then bounding each of the two KL divergences by a chi-square divergence. Here, we provide a refined proof which directly bounds the chi-square divergence, without the need to separate the first coordinate from the others. This is the result of Lemma 15 ,

The proof of Theorem 2 is then completed using the results of the three preliminary steps. We next turn to the detailed proof.
First step (genie-aided reduction): In this step we reduce the original model $P_{\theta}^{(n)}$ to a model $\tilde{P}_{\theta}^{(\ell)}$ which is aided by knowledge from a genie, and so estimation errors in the new model are only lower. Inspired by the operation of our proposed estimator in Section 2 we consider $\ell=\frac{n}{k}$ blocks of size
$k$ each, where, for simplicity of exposition, we assume that both $\ell, k$ are integers. A genie informs the estimator with the sign changes $S_{j} S_{j+1}$ at the time indices $j \in[n-1] \backslash\{i k\}_{i=1}^{\ell-1}$, that is, at all times except for $j=k, 2 k, 3 k, \cdots, n$. Hence, the estimator can align the signs within each block $\mathcal{I}_{i}:=\{(i-1) k+1,(i-1) k+2, \cdots, i k\}$ of size $k$, but not between blocks. Consequently, a statistically equivalent model to the original model with genie information is the model

$$
\begin{equation*}
Y_{i}=R_{i} \theta_{*}+Z_{i} \tag{144}
\end{equation*}
$$

where $i \in[n], Z_{i} \sim N\left(0, I_{d}\right)$ are i.i.d. exactly as in the original Markov model, and the signs $R_{i} \in\{-1,1\}$ are such that

$$
\begin{equation*}
R_{i+1}=R_{i} \tag{145}
\end{equation*}
$$

with probability 1 if $i \in[n-1] \backslash\{j k\}_{j=1}^{\ell-1}$ and

$$
R_{j k+1}= \begin{cases}R_{(j-1) k+1}, & \text { w.p. } \frac{1+\rho^{k}}{2}  \tag{146}\\ -R_{(j-1) k+1}, & \text { w.p. } \frac{1-\rho^{k}}{2}\end{cases}
$$

if $i \in\{j k\}_{j=1}^{\ell-1}$. Now, since the signs are fixed during the blocks $\left\{\mathcal{I}_{i}\right\}_{i=1}^{\ell}$ of length $k$, the mean of the samples in each block is a sufficient statistic for estimation of $\theta_{*}$. Thus, in the same spirit of the estimator in Section 2 , the genie-aided model is equivalent to

$$
\begin{equation*}
\bar{Y}_{i}:=\frac{1}{k} \sum_{j \in \mathcal{I}_{i}} Y_{i}=\bar{R}_{i} \theta_{*}+\bar{Z}_{i} \tag{147}
\end{equation*}
$$

where $i \in[\ell], \mathbb{P}\left[\bar{R}_{0}=1\right]=\frac{1}{2}$, and

$$
\bar{R}_{i+1}= \begin{cases}\bar{R}_{i} & \text { w.p. } \frac{1+\rho^{k}}{2}  \tag{148}\\ -\bar{R}_{i} & \text { w.p. } \frac{1-\rho^{k}}{2}\end{cases}
$$

is the gain (average of the signs) of the $i$ th block, and

$$
\begin{equation*}
\bar{Z}_{i}:=\frac{1}{k} \sum_{j \in \mathcal{I}_{i}} Z_{j} \sim N\left(0, \frac{1}{k} \cdot I_{d}\right) \tag{149}
\end{equation*}
$$

is a averaged Gaussian noise. Thus, there are $\ell$ samples in the equivalent model, and the flip probability is $\bar{\delta}:=\frac{1-\rho^{k}}{2}$, which is closer to $\frac{1}{2}$ compared to the flip probability in the original model, $\delta=\frac{1-\rho}{2}$. We note that there are two differences compared to the model used by the estimator in Section 2 First, due to the information supplied from the genie, here the average gain is $\bar{R}_{i} \in\{-1,1\}$ with probability 1 , to wit $\left|\bar{R}_{i}\right|$ never drops below 1 ; Second, there is no randomization between the blocks. In the definition of the estimator in Section 2 we indeed had the freedom to randomize the first sign in each block in order to make the blocks statistically independent, whereas here, in an impossibility bound, it is possible that statistical dependence between the blocks may result improved estimation rates (though as we shall see, this is essentially not the case). In what follows, we denote by $\bar{P}_{\theta}^{(\ell)}$ the probability distribution of $\bar{Y}_{1}^{\ell}$ under the model 147 with flip probability $\bar{\delta}$ (i.e., as with signs as in (148). At this point, the trade-off involved in an optimal choice of $k$ is already apparent as the block length $k$ increases, the dependence between the blocks, as reflected by the flip probability $\bar{\delta}=\frac{1-\rho^{k}}{2}$, decreases. On the other hand, the largest lower bound is attained when the genie reveals minimal information, that is, when $k$ is minimal. Our choice of $k$ is thus essentially the minimal value required so that $\frac{1-\rho^{k}}{2} \rightarrow \frac{1}{2}$, and specifically chosen as $k=\frac{\log (n)}{\delta}$.

Second step (change of measure): We denote by $p_{\delta}\left(s_{1}^{\ell}\right)=\mathbb{P}\left[S_{1}^{\ell}=s_{1}^{\ell}\right]$ the probability distribution of $\ell$ sign samples drawn according to the law of a homogeneous binary symmetric Markov chain with $\mathbb{P}\left[S_{0}=1\right]=1 / 2$ and flip probability $\delta$. Note that $p_{1 / 2}\left(s_{1}^{\ell}\right)$ is then an i.i.d. model, with uniform random signs. The next lemma uniformly bounds the ratio between $p_{\delta}\left(s_{1}^{\ell}\right)$ and $p_{1 / 2}\left(s_{1}^{\ell}\right)$ when $\delta$ is close to $1 / 2$, as was obtained in the first step.
Lemma 13. Let $S_{0}^{\ell} \in\{-1,1\}^{\ell+1}$ be a homogeneous binary symmetric Markov chain with $\mathbb{P}\left[S_{0}=\right.$ $1]=1 / 2$ and flip probability $\delta \in\left[0, \frac{1}{2}\right]$, and let $p_{\delta}\left(s_{1}^{\ell}\right)=\mathbb{P}\left[S_{1}^{\ell}=s_{1}^{\ell}\right]$. Furthermore, for $\bar{\delta}=\frac{1-\rho^{k}}{2}$ where $\rho=1-2 \delta$ and $k=\frac{\log (n)}{\delta}$ it holds that

$$
\begin{equation*}
1-\frac{1}{n} \leq \frac{p_{\bar{\delta}}\left(s_{1}^{\ell}\right)}{p_{1 / 2}\left(s_{1}^{\ell}\right)} \leq 1+\frac{2}{n} \tag{150}
\end{equation*}
$$

Proof. Note that $p_{1 / 2}\left(s_{1}^{\ell}\right)=\frac{1}{2^{\ell}}$ for all $s_{1}^{\ell} \in\{-1,1\}^{\ell}$. Then, since $\delta \in\left[0, \frac{1}{2}\right]$, letting $\rho=1-2 \delta$,

$$
\begin{equation*}
p_{\delta}\left(s_{1}^{\ell}\right) \leq(1-\delta)^{\ell}=\left(\frac{1+\rho}{2}\right)^{\ell}=p_{1 / 2}\left(s_{1}^{\ell}\right)(1+\rho)^{\ell}=p_{1 / 2}\left(s_{1}^{\ell}\right)(2-2 \delta)^{\ell} \tag{151}
\end{equation*}
$$

and

$$
\begin{equation*}
p_{\delta}\left(s_{1}^{\ell}\right) \geq \delta^{\ell}=\left(\frac{1-\rho}{2}\right)^{\ell}=p_{1 / 2}\left(s_{1}^{\ell}\right)(1-\rho)^{\ell}=p_{1 / 2}\left(s_{1}^{\ell}\right)(2 \delta)^{\ell} \tag{152}
\end{equation*}
$$

It this holds that

$$
\begin{equation*}
(2 \delta)^{\ell} \leq \frac{p_{\delta}\left(s_{1}^{\ell}\right)}{p_{1 / 2}\left(s_{1}^{\ell}\right)} \leq(2-2 \delta)^{\ell} \tag{153}
\end{equation*}
$$

Substituting $\delta=\bar{\delta}=\frac{1-\rho^{k}}{2}$, the upper bound on $\frac{p_{\bar{\delta}}\left(s_{1}^{\ell}\right)}{p_{1 / 2}\left(s_{1}^{\ell}\right)}$ is further upper bounded as

$$
\begin{align*}
\left(1+\rho^{k}\right)^{\ell} & =\left(1+(1-2 \delta)^{\frac{\log (n)}{\delta}}\right)^{\ell}  \tag{154}\\
& \stackrel{(a)}{\leq}\left(1+\frac{1}{n^{2}}\right)^{\ell}  \tag{155}\\
& \leq\left(1+\frac{1}{n^{2}}\right)^{n}  \tag{156}\\
& \stackrel{(a)}{\leq} e^{1 / n}  \tag{157}\\
& \stackrel{(b)}{\leq} 1+\frac{2}{n} \tag{158}
\end{align*}
$$

where both transitions denoted (a) follow from $1+x \leq e^{x}$ for $x \in \mathbb{R}$, and (b) follows from $e^{x} \leq 1+2 x$ for $x \in[0,1]$. Similarly, the lower bound on $\frac{p_{\bar{\delta}}\left(s_{1}^{\ell}\right)}{p_{1 / 2}\left(s_{1}^{\ell}\right)}$ is further lower bounded as

$$
\begin{align*}
\left(1-\rho^{k}\right)^{\ell} & =\left(1-(1-2 \delta)^{\frac{\log (n)}{\delta}}\right)^{\ell}  \tag{159}\\
& \stackrel{(a)}{\geq}\left(1-\frac{1}{n^{2}}\right)^{\ell}  \tag{160}\\
& \geq\left(1-\frac{1}{n^{2}}\right)^{n}  \tag{161}\\
& \stackrel{(b)}{\geq} 1-\frac{1}{n} \tag{162}
\end{align*}
$$

where (a) follows again from $1+x \leq e^{x}$ for $x \in \mathbb{R}$, and (b) follows from Bernoulli's inequality $(1-x)^{r} \geq 1-r x$ for $x \in[0,1]$ and $r \geq 1$.

Let $\tilde{P}_{\theta}^{(\ell)}$ denote the Gaussian mixture model (i.e., with flip probability $1 / 2$ ) corresponding to $\ell$ samples with means at $\pm \theta \in \mathbb{R}^{d}$, and note that it has the same $\ell$ marginal distributions as the genieaided reduced model $\bar{P}_{\theta}^{(\ell)}$ (with flip probability $\bar{\delta}$ ). As discussed, the probability distributions $\tilde{P}_{\theta}^{(\ell)}$ and $\bar{P}_{\theta}^{(\ell)}$ are close. The following lemma provides a change-of-measure bound on the KL divergence for this case.
Lemma 14. Let $U_{1}^{\ell}$ (resp. $V_{1}^{\ell}$ ) be distributed according to a probability distribution $P \equiv P_{U_{1} \cdots U_{\ell}}$ $\left(\right.$ resp. $\left.Q \equiv Q_{U_{1} \cdots U_{\ell}}\right)$ on $\left(\mathbb{R}^{d}\right)^{\ell}$, and let $\tilde{P} \equiv \tilde{P}_{U_{1} \cdots U_{\ell}}:=\prod_{i=1}^{n} P_{U_{i}}$ and $\tilde{Q} \equiv \tilde{Q}_{V_{1} \cdots V_{\ell}}:=\prod_{i=1}^{n} Q_{V_{i}}$ be the product distributions of their marginal distributions. Let

$$
\begin{equation*}
\beta_{P}:=\sup _{u_{1}^{\ell} \in\left(\mathbb{R}^{d}\right)^{\ell}}\left(\frac{P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right)}{\tilde{P}_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right)}\right) \tag{163}
\end{equation*}
$$

and

$$
\begin{equation*}
\beta_{Q}:=\sup _{u_{1}^{\ell} \in\left(\mathbb{R}^{d}\right)^{\ell}}\left(\frac{\tilde{Q}_{V_{1} \cdots V_{\ell}}\left(u_{1}^{\ell}\right)}{Q_{V_{1} \cdots V_{\ell}}\left(u_{1}^{\ell}\right)}\right) . \tag{164}
\end{equation*}
$$

Then,

$$
\begin{equation*}
\mathrm{D}_{\mathrm{KL}}(P \| Q) \leq \mathrm{D}_{\mathrm{KL}}(\tilde{P} \| \tilde{Q})+\log \left(\beta_{P} \cdot \beta_{Q}\right) \tag{165}
\end{equation*}
$$

Proof. Since $\tilde{Q}$ is a product distribution, and since $P$ and $\tilde{P}$ have the same marginal distributions on $\mathbb{R}^{d}$ (at each time point) then

$$
\begin{align*}
& \int P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\frac{1}{\tilde{Q}_{V_{1} \cdots V_{\ell}}\left(u_{1}^{\ell}\right)}\right) \mathrm{d} u_{1}^{\ell} \\
& =\sum_{i=1}^{\ell} \int P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\frac{1}{\tilde{Q}_{V_{i}}\left(u_{i}\right)}\right) \mathrm{d} u_{1}^{\ell}  \tag{166}\\
& =\sum_{i=1}^{\ell} \int P_{U_{i}}\left(u_{i}\right) \log \left(\frac{1}{\tilde{Q}_{V_{i}}\left(u_{i}\right)}\right) \mathrm{d} u_{i}  \tag{167}\\
& =\sum_{i=1}^{\ell} \int \tilde{P}_{U_{i}}\left(u_{i}\right) \log \left(\frac{1}{\tilde{Q}_{V_{i}}\left(u_{i}\right)}\right) \mathrm{d} u_{i}  \tag{168}\\
& =\sum_{i=1}^{\ell} \int \tilde{P}_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\frac{1}{\tilde{Q}_{V_{i}}\left(u_{i}\right)}\right) \mathrm{d} u_{1}^{\ell}  \tag{169}\\
& =\int \tilde{P}_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\frac{1}{\tilde{Q}_{V_{1} \cdots V_{\ell}}\left(u_{1}^{\ell}\right)}\right) \mathrm{d} u_{1}^{\ell} \tag{170}
\end{align*}
$$

Using similar reasoning

$$
\begin{align*}
& \int \tilde{P}_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\tilde{P}_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right)\right) \mathrm{d} u_{1}^{\ell} \\
& =\sum_{i=1}^{\ell} \int \tilde{P}_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\tilde{P}_{U_{i}}\left(u_{i}\right)\right) \mathrm{d} u_{1}^{\ell}  \tag{171}\\
& =\sum_{i=1}^{\ell} \int \tilde{P}_{U_{i}}\left(u_{i}\right) \log \left(\tilde{P}_{U_{i}}\left(u_{i}\right)\right) \mathrm{d} u_{i}  \tag{172}\\
& =\sum_{i=1}^{\ell} \int P_{U_{i}}\left(u_{i}\right) \log \left(\tilde{P}_{U_{i}}\left(u_{i}\right)\right) \mathrm{d} u_{i}  \tag{173}\\
& =\sum_{i=1}^{\ell} \int P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\tilde{P}_{U_{i}}\left(u_{i}\right)\right) \mathrm{d} u_{1}^{\ell}  \tag{174}\\
& =\int P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\tilde{P}_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right)\right) \mathrm{d} u_{1}^{\ell}  \tag{175}\\
& =\int P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right)\right) \mathrm{d} u_{1}^{\ell}+\int P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\frac{\tilde{P}_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right)}{P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right)}\right) \mathrm{d} u_{1}^{\ell} \tag{176}
\end{align*}
$$

Combining (170) and 176 results the following bound

$$
\begin{align*}
& \mathrm{D}_{\mathrm{KL}}(P \| Q) \\
& =\int P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\frac{P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right)}{Q_{V_{1} \cdots V_{\ell}}\left(u_{1}^{\ell}\right)}\right) \mathrm{d} u_{1}^{\ell}  \tag{177}\\
& =\int P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\frac{P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right)}{\tilde{Q}_{V_{1} \cdots V_{\ell}}\left(u_{1}^{\ell}\right)}\right) \mathrm{d} u_{1}^{\ell}+\int P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\frac{\tilde{Q}_{V_{1} \cdots V_{\ell}}\left(u_{1}^{\ell}\right)}{Q_{V_{1} \cdots V_{\ell}}\left(u_{1}^{\ell}\right)}\right) \mathrm{d} u_{1}^{\ell}  \tag{178}\\
& =\int \tilde{P}_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\tilde{P}_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right)\right) \mathrm{d} u_{1}^{\ell}-\int P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\frac{\tilde{P}_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right)}{P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right)}\right) \mathrm{d} u_{1}^{\ell}
\end{align*}
$$

$$
\begin{align*}
& +\int P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\frac{1}{\tilde{Q}_{V_{1} \cdots V_{\ell}}\left(u_{1}^{\ell}\right)}\right) \mathrm{d} u_{1}^{\ell}+\int P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right) \log \left(\frac{\tilde{Q}_{V_{1} \cdots V_{\ell}}\left(u_{1}^{\ell}\right)}{Q_{V_{1} \cdots V_{\ell}}\left(u_{1}^{\ell}\right)}\right) \mathrm{d} u_{1}^{\ell}  \tag{179}\\
\leq & \mathrm{D}_{\mathrm{KL}}(\tilde{P} \| \tilde{Q})+\sup _{u_{1}^{\ell} \in\left(\mathbb{R}^{d}\right)^{\ell}} \log \left(\frac{P_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right)}{\tilde{P}_{U_{1} \cdots U_{\ell}}\left(u_{1}^{\ell}\right)}\right)+\sup _{u_{1}^{\ell} \in\left(\mathbb{R}^{d}\right)^{\ell}} \log \left(\frac{\tilde{Q}_{V_{1} \cdots V_{\ell}}\left(u_{1}^{\ell}\right)}{Q_{V_{1} \cdots V_{\ell}}\left(u_{1}^{\ell}\right)}\right), \tag{180}
\end{align*}
$$

as claimed by the lemma.
Third step (bound on the chi-square divergence): Let $\varphi\left(y ; \theta, \sigma^{2}\right)$ be the Gaussian PDF with mean $\theta \in \mathbb{R}^{d}$ and covariance matrix $\sigma^{2} \cdot I_{d}$. The next lemma bounds the chi-square divergence between a pair of such distributions with different means. Originally, a bound of this order was established in [Wu and Zhou, 2019, proof of Lemma 27] on the KL divergence, by splitting the first coordinate (which is assumed, w.l.o.g., to contain the signal) and the other $d-1$ coordinates, using the chain rule, and then bounding each of the two KL terms with the corresponding chi-square divergence. Here we provide a direct upper bound on the chi-square divergence, which does not rely on splitting between the coordinates of the mean vector, and which might be of independent use in future works.
Lemma 15. Let $\tilde{P}_{\theta}=\frac{1}{2} N\left(\theta, \sigma^{2} \cdot I_{d}\right)+\frac{1}{2} N\left(-\theta, \sigma^{2} \cdot I_{d}\right)$ be a balanced Gaussian mixture with means at $\pm \theta \in \mathbb{R}^{d}$. Then, if $\left\|\theta_{0}\right\|=\left\|\theta_{1}\right\|=t \leq \sigma$, then

$$
\begin{equation*}
\chi^{2}\left(\tilde{P}_{\theta_{1}} \| \tilde{P}_{\theta_{0}}\right) \leq \frac{8 t^{2}}{\sigma^{4}} \cdot\left\|\theta_{0}-\theta_{1}\right\|^{2} \tag{181}
\end{equation*}
$$

Proof. For any $y \in \mathbb{R}^{d}$,

$$
\begin{align*}
\tilde{P}_{\theta}(y) & =\frac{1}{2} \frac{1}{\left(2 \pi \sigma^{2}\right)^{d / 2}} e^{-\frac{\|y-\theta\|^{2}}{2 \sigma^{2}}}+\frac{1}{2} \frac{1}{\left(2 \pi \sigma^{2}\right)^{d / 2}} e^{-\frac{\|y+\theta\|^{2}}{2 \sigma^{2}}}  \tag{182}\\
& =\varphi\left(y ; 0, \sigma^{2}\right) \cdot e^{-\frac{\|\theta\|^{2}}{2 \sigma^{2}}} \cdot \cosh \left(\frac{\theta^{\top} y}{\sigma^{2}}\right) . \tag{183}
\end{align*}
$$

Then,

$$
\begin{align*}
& \chi^{2}\left(\tilde{P}_{\theta_{1}} \| \tilde{P}_{\theta_{0}}\right) \\
& =\int \frac{\left[\varphi\left(y ; 0, \sigma^{2}\right) e^{-\frac{\left\|\theta_{1}\right\|^{2}}{2 \sigma^{2}}} \cdot \cosh \left(\frac{\theta_{1}^{\top} y}{\sigma^{2}}\right)-\varphi\left(y ; 0, \sigma^{2}\right) e^{-\frac{\left\|\theta_{0}\right\|^{2}}{2 \sigma^{2}}} \cdot \cosh \left(\frac{\theta_{0}^{\top} y}{\sigma^{2}}\right)\right]^{2}}{\varphi\left(y ; 0, \sigma^{2}\right) \cdot e^{-\frac{\left\|\theta_{0}\right\|^{2}}{2 \sigma^{2}}} \cosh \left(\frac{\theta_{0}^{\top} y}{\sigma^{2}}\right)} \cdot \mathrm{d} y  \tag{184}\\
& =e^{\frac{\left\|\theta_{0}\right\|^{2}}{2 \sigma^{2}}} \int \varphi\left(y ; 0, \sigma^{2}\right) \frac{\left[e^{-\frac{\left\|\theta_{1}\right\|^{2}}{2 \sigma^{2}}} \cdot \cosh \left(\frac{\theta_{1}^{\top} y}{\sigma^{2}}\right)-e^{-\frac{\left\|\theta_{0}\right\|^{2}}{2 \sigma^{2}}} \cdot \cosh \left(\frac{\theta_{0}^{\top} y}{\sigma^{2}}\right)\right]^{2}}{\cosh \left(\frac{\theta_{0}^{\top} y}{\sigma^{2}}\right)} \cdot \mathrm{d} y  \tag{185}\\
& \begin{array}{l}
(a) \\
\leq \\
=e^{\frac{\left\|\theta_{0}\right\|^{2}}{2 \sigma^{2}}} \cdot \int \varphi\left(y ; 0, \sigma^{2}\right)\left[e^{-\frac{\left\|\theta_{0}\right\|^{2}}{2 \sigma^{2}}} \cdot \int \varphi\left(y ; 0, \sigma^{2}\right) e^{-\frac{\left\|\theta_{1}\right\|^{2} \|^{2}}{\sigma^{2}}} \cdot \cosh ^{2}\left(\frac{\theta_{1}^{\top} y}{\sigma^{2}}\right) \cdot \mathrm{d} y\right. \\
\left.\quad-2 e^{\frac{\left\|\theta_{0}\right\|^{2}}{2 \sigma^{2}}} \cdot \int \varphi\left(y ; 0, \sigma^{2}\right) e^{-\frac{\left\|\theta_{1}\right\|^{2}+\left\|\theta_{0}\right\|^{2}}{2 \sigma^{2}}} \cosh \left(\frac{\theta_{1}^{\top} y}{\sigma^{2}}\right)-e^{-\frac{\left\|\theta_{0}\right\|^{2}}{2 \sigma^{2}}} \cdot \cosh \left(\frac{\theta_{0}^{\top} y}{\sigma^{2}}\right)\right]^{2} \cdot \mathrm{~d} y \\
\left.\sigma^{2}\right) \cosh \left(\frac{\theta_{0}^{\top} y}{\sigma^{2}}\right) \cdot \mathrm{d} y \\
\quad+e^{\frac{\left\|\theta_{0}\right\|^{2}}{2 \sigma^{2}}} \cdot \int \varphi\left(y ; 0, \sigma^{2}\right) e^{-\frac{\left\|\theta_{0}\right\|^{2}}{\sigma^{2}}} \cdot \cosh ^{2}\left(\frac{\theta_{0}^{\top} y}{\sigma^{2}}\right) \cdot \mathrm{d} y
\end{array} \tag{186}
\end{align*}
$$

where $(a)$ follows since $\cosh (x) \geq 1$ for all $x \in \mathbb{R}$. We next evaluate the integral for each of the terms. First, for any $y \in \mathbb{R}^{d}$,

$$
\int \varphi\left(y ; 0, \sigma^{2}\right) \cosh ^{2}\left(\frac{\theta_{1}^{\top} y}{\sigma^{2}}\right) \mathrm{d} y
$$

$$
\begin{align*}
& \stackrel{(a)}{=} \int \varphi(y ; 0,1) \cdot \cosh ^{2}\left(\frac{\theta_{1}^{\top} y}{\sigma}\right) \mathrm{d} y  \tag{188}\\
& \stackrel{(b)}{=} \int \varphi(t ; 0,1) \cdot \cosh ^{2}\left(\frac{\left\|\theta_{1}\right\|}{\sigma} t\right) \mathrm{d} t  \tag{189}\\
& =\int \varphi(t ; 0,1)\left[\frac{\exp \left(\frac{\left\|\theta_{1}\right\|}{\sigma} t\right)+\exp \left(-\frac{\left\|\theta_{1}\right\|}{\sigma} t\right)}{2}\right]^{2} \mathrm{~d} t  \tag{190}\\
& =\int \varphi(t ; 0,1) \frac{\exp \left(2 \frac{\left\|\theta_{1}\right\|}{\sigma} t\right)+2+\exp \left(-2 \frac{\left\|\theta_{1}\right\|}{\sigma} t\right)}{4} \mathrm{~d} t  \tag{191}\\
& \stackrel{(c)}{=} \frac{1}{2}+\frac{1}{2} \cdot e^{2 \frac{\left\|\theta_{1}\right\|^{2}}{\sigma^{2}}}  \tag{192}\\
& =e^{\frac{\left\|\theta_{1}\right\|^{2}}{\sigma^{2}}} \cosh \left(\frac{\left\|\theta_{1}\right\|^{2}}{\sigma^{2}}\right) \tag{193}
\end{align*}
$$

where $(a)$ follows from the change of variables $y \rightarrow \frac{y}{\sigma}$, $(b)$ follows from the rotational invariance of the Gaussian PDF $\varphi\left(y ; 0, \sigma^{2}\right)$, we may assume that $\theta_{1}=\left(\left\|\theta_{1}\right\|, 0,0, \ldots, 0\right)$, setting $t \in \mathbb{R}$ to be the first coordinate of $y$, and integrating over all other $d-1$ coordinates, and $(c)$ follows from the Gaussian moment-generating function formula. The third term in the integral is similarly evaluated. For the second term,

$$
\begin{align*}
& \int \varphi\left(y ; 0, \sigma^{2}\right)\left[2 \cosh \left(\frac{\theta_{1}^{\top} y}{\sigma^{2}}\right) \cosh \left(\frac{\theta_{0}^{\top} y}{\sigma^{2}}\right)\right] \cdot \mathrm{d} y \\
& \stackrel{(a)}{=} \int \varphi(y ; 0,1)\left[2 \cosh \left(\frac{\theta_{1}^{\top} y}{\sigma}\right) \cosh \left(\frac{\theta_{0}^{\top} y}{\sigma}\right)\right] \cdot \mathrm{d} y  \tag{194}\\
& \stackrel{(b)}{=} \int \varphi(y ; 0,1) \cosh \left(\frac{\left(\theta_{1}+\theta_{0}\right)^{\top} y}{\sigma}\right) \cdot \mathrm{d} y+\int \varphi(y ; 0,1) \cosh \left(\frac{\left(\theta_{1}-\theta_{0}\right)^{\top} y}{\sigma}\right) \cdot \mathrm{d} y  \tag{195}\\
& \stackrel{(c)}{=} \int \varphi(t ; 0,1) \cosh \left(\frac{\left\|\theta_{1}+\theta_{0}\right\| t}{\sigma}\right) \cdot \mathrm{d} t+\int \varphi(t ; 0,1) \cosh \left(\frac{\left\|\theta_{1}-\theta_{0}\right\| t}{\sigma}\right) \cdot \mathrm{d} t  \tag{196}\\
& \stackrel{(d)}{=} e^{\frac{\left\|\theta_{1}+\theta_{0}\right\|^{2}}{2 \sigma^{2}}}+e^{\frac{\left\|\theta_{1}-\theta_{0}\right\|^{2}}{2 \sigma^{2}}}  \tag{197}\\
& =e^{\frac{\left\|\theta_{0}\right\|^{2}+\left\|\theta_{1}\right\|^{2}}{2 \sigma^{2}}} \cdot\left[e^{\frac{\theta_{1}^{\top} \theta_{0}}{\sigma^{2}}}+e^{\frac{-\theta_{1}^{\top} \theta_{0}}{\sigma^{2}}}\right]  \tag{198}\\
& =e^{\frac{\left\|\theta_{0}\right\|^{2}+\left\|\theta_{1}\right\|^{2}}{2 \sigma^{2}}} \cdot 2 \cosh \left(\frac{\theta_{1}^{\top} \theta_{0}}{\sigma^{2}}\right) \tag{199}
\end{align*}
$$

where ( $a$ ) follows from the change of variables $y \rightarrow \frac{y}{\sigma}$, (b) follows from the identity $2 \cosh (x) \cosh (y)=\cosh (x+y)+\cosh (x-y)$, and $(c)$ follows from rotational invariance, and $(d)$ follows from

$$
\begin{equation*}
\int \varphi(t ; 0,1) \cosh (a t) \mathrm{d} t=\int \varphi(t ; 0,1) \frac{e^{a t}+e^{-a t}}{2} \mathrm{~d} t=e^{\frac{a^{2}}{2}} \tag{200}
\end{equation*}
$$

Continuing 187, we thus have

$$
\begin{equation*}
\chi^{2}\left(\tilde{P}_{\theta_{1}} \| \tilde{P}_{\theta_{0}}\right) \leq e^{\frac{\left\|\theta_{0}\right\|^{2}}{2 \sigma^{2}}} \cdot\left[\cosh \left(\frac{\left\|\theta_{1}\right\|^{2}}{\sigma^{2}}\right)-2 \cosh \left(\frac{\theta_{1}^{\top} \theta_{0}}{\sigma^{2}}\right)+\cosh \left(\frac{\left\|\theta_{0}\right\|^{2}}{\sigma^{2}}\right)\right] \tag{201}
\end{equation*}
$$

Now, $\frac{\mathrm{d}}{\mathrm{d} x} \cosh (x)=\sinh (x)$ and $\frac{\mathrm{d}^{2}}{\mathrm{~d} x^{2}} \cosh (x)=\cosh (x)$. In addition, if $0 \leq x \leq 1$ then it can be easily verified that $\sinh (x) \leq 2 x$. Thus, if $0 \leq x \leq y \leq 1$ then

$$
\begin{equation*}
\cosh (y)-\cosh (x)=\int_{x}^{y} \frac{\mathrm{~d}}{\mathrm{~d} r} \cosh (r) \mathrm{d} r=\int_{x}^{y} \sinh (r) \mathrm{d} r \leq \int_{x}^{y} 2 r \mathrm{~d} r=y^{2}-x^{2} . \tag{202}
\end{equation*}
$$

Moreover, $e^{x} \leq 1+2 x$ for $x \in[0,1]$. Thus, if we assume $\left\|\theta_{0}\right\|=\left\|\theta_{1}\right\| \leq \sigma$ we may further upper bound (201) as

$$
\chi^{2}\left(\tilde{P}_{\theta_{1}} \| \tilde{P}_{\theta_{0}}\right)
$$

$$
\begin{align*}
& \leq e^{\frac{\left\|\theta_{0}\right\|^{2}}{2 \sigma^{2}}} \cdot\left[\cosh \left(\frac{\left\|\theta_{0}\right\|^{2}}{\sigma^{2}}\right)-\cosh \left(\frac{\left|\theta_{1}^{\top} \theta_{0}\right|}{\sigma^{2}}\right)+\cosh \left(\frac{\left\|\theta_{0}\right\|^{2}}{\sigma^{2}}\right)-\cosh \left(\frac{\left|\theta_{1}^{\top} \theta_{0}\right|}{\sigma^{2}}\right)\right]  \tag{203}\\
& \stackrel{(a)}{\leq} 2 \cdot\left[\cosh \left(\frac{\left\|\theta_{0}\right\|^{2}}{\sigma^{2}}\right)-\cosh \left(\frac{\left|\theta_{1}^{\top} \theta_{0}\right|}{\sigma^{2}}\right)+\cosh \left(\frac{\left\|\theta_{0}\right\|^{2}}{\sigma^{2}}\right)-\cosh \left(\frac{\left|\theta_{1}^{\top} \theta_{0}\right|}{\sigma^{2}}\right)\right]  \tag{204}\\
& \stackrel{(b)}{\leq} 2 \cdot \frac{\left\|\theta_{1}\right\|^{4}-\left(\theta_{1}^{\top} \theta_{0}\right)^{2}+\left\|\theta_{0}\right\|^{4}-\left(\theta_{1}^{\top} \theta_{0}\right)^{2}}{\sigma^{4}}, \tag{205}
\end{align*}
$$

where $(a)$ follows since $e^{\frac{\left\|\theta_{0}\right\|^{2}}{2 \sigma^{2}}} \leq 2$ under the assumption $\left\|\theta_{0}\right\|^{2} \leq \sigma^{2},(b)$ follows from 202. Now, the numerator of 205 is further upper bounded as

$$
\begin{align*}
& \left\|\theta_{1}\right\|^{4}-2\left(\theta_{1}^{\top} \theta_{0}\right)^{2}+\left\|\theta_{0}\right\|^{4}  \tag{206}\\
& =\left(\theta_{1}^{\top} \theta_{1}\right)^{2}-\left(\theta_{1}^{\top} \theta_{0}\right)\left(\theta_{0}^{\top} \theta_{1}\right)-\left(\theta_{0}^{\top} \theta_{1}\right)\left(\theta_{1}^{\top} \theta_{0}\right)+\left(\theta_{0}^{\top} \theta_{0}\right)^{2}  \tag{207}\\
& =\operatorname{Tr}\left[\theta_{0} \theta_{0}^{\top} \theta_{0} \theta_{0}^{\top}-\theta_{0} \theta_{0}^{\top} \theta_{1} \theta_{1}^{\top}-\theta_{1} \theta_{1}^{\top} \theta_{0} \theta_{0}^{\top}+\theta_{1} \theta_{1}^{\top} \theta_{1} \theta_{1}^{\top}\right]  \tag{208}\\
& =\operatorname{Tr}\left[\left(\theta_{0} \theta_{0}^{\top}-\theta_{1} \theta_{1}^{\top}\right)^{2}\right]  \tag{209}\\
& =\left\|\theta_{0} \theta_{0}^{\top}-\theta_{1} \theta_{1}^{\top}\right\|_{F}^{2}  \tag{210}\\
& =\left\|\theta_{0} \theta_{0}^{\top}-\theta_{0} \theta_{1}^{\top}+\theta_{0} \theta_{1}^{\top}-\theta_{1} \theta_{1}^{\top}\right\|_{F}^{2}  \tag{211}\\
& \leq 2\left\|\theta_{0} \theta_{0}^{\top}-\theta_{0} \theta_{1}^{\top}\right\|_{F}^{2}+2\left\|\theta_{0} \theta_{1}^{\top}-\theta_{1} \theta_{1}^{\top}\right\|_{F}^{2}  \tag{212}\\
& =2\left\|\theta_{0}\right\|^{2} \cdot\left\|\theta_{0}-\theta_{1}\right\|^{2}+2\left\|\theta_{1}\right\|^{2} \cdot\left\|\theta_{0}-\theta_{1}\right\|^{2} \tag{213}
\end{align*}
$$

where the inequality follows from $\|A+B\|_{F}^{2} \leq 2\|A\|_{F}^{2}+2\|B\|_{F}^{2}$. Inserting this bound into 205), and using $\left\|\theta_{0}\right\|=\left\|\theta_{1}\right\|=t$ results the bound 181 .

With the results of the three steps at hand, we may complete the proof of Theorem 2
Proof of Theorem 2. Recall that we assume the low-dimension regime $3 \leq d \leq \delta n$. The condition $d \geq 3$ can be relaxed to $d \geq 2$, as promised in Theorem 2, using a different construction of the packing set. We leave this refinement to the end of this section. If $t \geq \sqrt{\delta}$ then the lower bound for the Gaussian location model $\sqrt{5}$ implies a bound of $\Theta\left(\sqrt{\frac{d}{n}}\right)$. This can be verified by separately checking the only two cases $t \geq \sqrt{\frac{d}{n}} \geq \sqrt{\delta}$ and $t \geq \sqrt{\delta} \geq \sqrt{\frac{d}{n}}$ possible in the low dimension regime. We thus henceforth may concentrate on the regime $t \leq \sqrt{\delta}$. To continue, we henceforth assume the slightly stronger requirement $t \leq \sqrt{\frac{\delta}{\log n}}=\sqrt{\frac{1}{k}}$, where $\frac{1}{k}$ is the variance in the genie-aided reduced model.
Let $\bar{\delta}=\frac{1-\rho^{k}}{2}$ and recall that $\bar{P}_{\theta_{m}}^{(\ell)}$ is the Gaussian model with Markovian signs with flip probability $\bar{\delta}$. Further let $\varphi\left(y_{1}^{\ell} ; \mu\right)$ be the Gaussian PDF for $\ell$ samples from a $d_{0}$-dimensional model with mean $\mu \in\left(\mathbb{R}^{d_{0}}\right)^{\ell}$ and covariance matrix $\Sigma=\frac{1}{k} I_{d_{0}} \otimes I_{\ell} \in \mathbb{R}_{+}^{d_{0} \ell \times d_{0} \ell}$. With this notation, it holds that

$$
\begin{equation*}
\bar{P}_{\theta_{m}}^{(\ell)}\left(y_{1}^{\ell}\right)=\sum_{r_{1}^{\ell} \in\{-1,1\}^{\ell}} p_{\bar{\delta}}\left(r_{1}^{\ell}\right) \cdot \varphi\left(y_{1}^{\ell} ; r_{1}^{\ell} \otimes \theta_{m}\right) \tag{214}
\end{equation*}
$$

Similarly, let $\tilde{P}_{\theta_{m}}^{(\ell)}$ be the Gaussian model with Markovian signs with flip probability $1 / 2$, that is, a Gaussian mixture model (which is, in fact, memoryless),

$$
\begin{align*}
\tilde{P}_{\theta_{m}}^{(\ell)}\left(y_{1}^{\ell}\right) & =\sum_{r_{1}^{\ell} \in\{-1,1\}^{\ell}} p_{1 / 2}\left(r_{1}^{\ell}\right) \cdot \varphi\left(y_{1}^{\ell} ; r_{1}^{\ell} \otimes \theta_{m}\right)  \tag{215}\\
& =\prod_{i=1}^{\ell}\left[\frac{1}{2} \varphi\left(y_{i} ; \theta_{m}\right)+\frac{1}{2} \varphi\left(y_{i} ;-\theta_{m}\right)\right] \tag{216}
\end{align*}
$$

Now, Lemma 13 implies that

$$
\begin{equation*}
1-\frac{1}{n} \leq \min _{r_{1}^{\ell} \in\{-1,1\}^{\ell}} \frac{p_{\bar{\delta}}\left(r_{1}^{\ell}\right)}{p_{1 / 2}\left(r_{1}^{\ell}\right)} \leq \frac{\bar{P}_{\theta_{m}}^{(\ell)}\left(y_{1}^{\ell}\right)}{\tilde{P}_{\theta_{m}}^{(\ell)}\left(y_{1}^{\ell}\right)} \leq \max _{r_{1}^{\ell} \in\{-1,1\}^{\ell}} \frac{p_{\bar{\delta}}\left(r_{1}^{\ell}\right)}{p_{1 / 2}\left(r_{1}^{\ell}\right)} \leq 1+\frac{2}{n} \tag{217}
\end{equation*}
$$

and hence

$$
\begin{align*}
\mathrm{D}_{\mathrm{KL}}\left(\bar{P}_{\theta_{m}}^{(\ell)} \| \bar{P}_{\theta_{0}}^{(\ell)}\right) & \stackrel{(a)}{\leq} \mathrm{D}_{\mathrm{KL}}\left(\tilde{P}_{\theta_{m}}^{(\ell)} \| \tilde{P}_{\theta_{0}}^{(\ell)}\right)+\log \left[\left(1+\frac{2}{n}\right)\left(\frac{1}{1-\frac{1}{n}}\right)\right]  \tag{218}\\
& \stackrel{(b)}{\leq} \mathrm{D}_{\mathrm{KL}}\left(\tilde{P}_{\theta_{m}}^{(\ell)} \| \tilde{P}_{\theta_{0}}^{(\ell)}\right)+2 \log \left(1+\frac{2}{n}\right)  \tag{219}\\
& \leq \mathrm{D}_{\mathrm{KL}}\left(\tilde{P}_{\theta_{m}}^{(\ell)} \| \tilde{P}_{\theta_{0}}^{(\ell)}\right)+\frac{4}{n}  \tag{220}\\
& \stackrel{(c)}{=} \ell \cdot \mathrm{D}_{\mathrm{KL}}\left(\tilde{P}_{\theta_{m}} \| \tilde{P}_{\theta_{0}}\right)+\frac{4}{n}  \tag{221}\\
& \stackrel{(d)}{\leq} \ell \cdot \chi^{2}\left(\tilde{P}_{\theta_{m}} \| \tilde{P}_{\theta_{0}}\right)+\frac{4}{n}  \tag{222}\\
& \stackrel{(e)}{\leq} 8 \frac{\ell t^{2}}{\sigma^{4}} \cdot\left\|\theta_{0}-\theta_{1}\right\|^{2}+\frac{4}{n}  \tag{223}\\
& \stackrel{(f)}{\leq} 32 \log (n) \cdot \frac{n t^{4} \epsilon^{2}}{\delta}+\frac{4}{n} \tag{224}
\end{align*}
$$

where (a) follows from Lemma 14. (b) follows since $\frac{1}{1-\frac{1}{n}} \leq 1+\frac{2}{n}$ for $n \geq 2$, (c) follows from the tensorization property of the KL divergence, $(d)$ follows from the fact (cf. [Tsybakov, 2008, Eq. (2.27)]) that $\mathrm{D}_{\mathrm{KL}}(P \| Q) \leq \chi^{2}(P \| Q)$ for any pair of probability measures $P$ and $Q,(e)$ follows from Lemma 15 , and $(f)$ follows from property (iv) of the packing set $\left\|\theta_{0}-\theta_{1}\right\|^{2} \leq 4 t^{2} \epsilon^{2}$ and $\sigma^{2}=\frac{1}{k}=\frac{\delta}{\log n}$ and $\ell=\frac{n}{k}$.
Recall that from Fano's argument 143 , the largest $\epsilon>0$ so that $\mathrm{D}_{\mathrm{KL}}\left(\bar{P}_{\theta_{m}}^{(\ell)} \| \bar{P}_{\theta_{0}}^{(\ell)}\right) \leq \frac{d}{16}$ assures the bound $\mathrm{M}(n, d, \delta, t) \geq \frac{\epsilon t}{128}$. Assuming that $\frac{4}{n} \leq \frac{d}{32}$, that is $n \geq \frac{128}{d}$, this will occur if $32 \log (n)$. $\frac{n t^{4} \epsilon^{2}}{\delta} \leq \frac{d}{32}$. This can be achieved by the choice $\epsilon=\frac{1}{32 \sqrt{\log (n)}} \cdot \min \left\{1, \frac{1}{t^{2}} \sqrt{\frac{d \delta}{n}}\right\}$. Using this value in the Fano's based bound then completes the proof of the theorem.

Relaxing $d \geq 3$ to $d \geq 2 \quad$ In the above proof, we assumed that $d \geq 3$. This condition can be relaxed to $d \geq 2$ by using a more careful construction of the packing set.
Lemma 16. Let $d \geq 2$ be an integer and $\alpha \in\left(0, \frac{\pi}{2}\right]$ be an angle. Then there exists a $\left.2 \sin \left(\frac{\alpha}{2}\right)\right)$ packing on $\mathbb{S}^{d-1}$ of size at least $\frac{\cos (\alpha)}{\sin ^{d-1}(\alpha)}$.

Proof. We will greedily construct the desired packing set. Let $\Delta=2 \sin \left(\frac{\alpha}{2}\right)$. Start with an arbitrary point on $\mathbb{S}^{d-1}$. In each of the following steps, put into the packing set another arbitrary point that is $\Delta$-far from any existing points in the packing set and their antipodal points. Repeat this process until no more points can be put without violating the distance guarantee. Note that this construction guarantees that for any pair of distinct points $\theta, \theta^{\prime}$ in the packing set,

$$
\begin{equation*}
\operatorname{loss}\left(\theta, \theta^{\prime}\right)=\min \left\{\left\|\theta-\theta^{\prime}\right\|,\left\|\theta+\theta^{\prime}\right\|\right\}>2 \sin \left(\frac{\alpha}{2}\right)=\Delta \tag{225}
\end{equation*}
$$

We then lower bound the cardinality of the above packing set. For $\alpha \in\left[0, \frac{\pi}{2}\right]$, let $S(\alpha)$ denote the surface area of a spherical cap on $\mathbb{S}^{d-1}$ of angular radius $\alpha$. For $\theta \in \mathbb{S}^{d-1}$ and $\alpha \in\left[0, \frac{\pi}{2}\right]$, let $\mathbb{K}(\theta, \alpha):=\left\{\theta^{\prime} \in \mathbb{S}^{d-1}:\left|\left\langle\theta, \theta^{\prime}\right\rangle\right| \geq \cos (\alpha)\right\}$ denote the union of two spherical caps centered around $\theta$ and $-\theta$, respectively, of angular radius $\alpha$ each. Note that each new point $\theta$ in the construction induces a forbidden region on $\mathbb{S}^{d-1}$ of surface area at most $|\mathbb{K}(\theta, \alpha)|=2 S(\alpha)$ in which following points cannot lie, since any point within $\mathbb{K}(\theta, \alpha)$ has distance at most $2 \sin \left(\frac{\alpha}{2}\right)=\Delta$ to either $\theta$ or $-\theta$. Also, the surface area of the forbidden region is upper bounded by $2 S(\alpha)$ since $\mathbb{K}(\theta, \alpha)$ may overlap with $\mathbb{K}\left(\theta^{\prime}, \alpha\right)$ for some previous $\theta^{\prime}$. Therefore, by the step at which the greedy construction terminates, one must have put at least $\frac{2 S\left(\frac{\pi}{2}\right)}{2 S(\alpha)}=\frac{S\left(\frac{\pi}{2}\right)}{S(\alpha)}$ many points, where the numerator is nothing but the surface area of $\mathbb{S}^{d-1}$. [Blachman and Few. 1963. Eq. (1)] upper bounds the area ratio between
two spherical caps on $\mathbb{S}^{d-1}$ as follows:

$$
\begin{equation*}
\frac{S(\alpha)}{S(\beta)}=\frac{\int_{0}^{\alpha} \sin ^{d-2}(x) \mathrm{d} x}{\int_{0}^{\beta} \sin ^{d-2}(x) \mathrm{d} x} \leq \sec (\alpha) \cdot \frac{\sin ^{d-1}(\alpha)}{\sin ^{d-1}(\beta)} \tag{226}
\end{equation*}
$$

Using this bound, we conclude that the cardinality of the constructed packing set is at least $\frac{\cos (\alpha)}{\sin ^{d-1}(\alpha)}$, which finishes the proof.

By Lemma 16 , a $\Delta$-packing set can be obtained by setting $2 \sin \left(\frac{\alpha}{2}\right)=\Delta$. Therefore,

$$
\begin{gather*}
\sin \left(\frac{\alpha}{2}\right)=\frac{\Delta}{2}, \quad \cos \left(\frac{\alpha}{2}\right)=\sqrt{1-\frac{\Delta^{2}}{4}}  \tag{227}\\
\sin (\alpha)=2 \sin \left(\frac{\alpha}{2}\right) \cos \left(\frac{\alpha}{2}\right)=\Delta \sqrt{1-\frac{\Delta^{2}}{4}}, \quad \cos (\alpha)=1-2 \sin ^{2}\left(\frac{\alpha}{2}\right)=1-\frac{\Delta^{2}}{2} \tag{228}
\end{gather*}
$$

and the cardinality of the packing set is at least

$$
\begin{equation*}
\left(1-\frac{\Delta^{2}}{2}\right)\left(\Delta \sqrt{1-\frac{\Delta^{2}}{4}}\right)^{-(d-1)} \tag{229}
\end{equation*}
$$

Setting $\Delta \sqrt{1-\frac{\Delta^{2}}{4}}=\frac{1}{8}$, we have $\Delta=\frac{1}{2} \sqrt{8-3 \sqrt{7}} \geq \frac{1}{8}$ and we get a $\frac{1}{8}$-packing set $\Theta_{M}=$ $\left\{\theta_{m}\right\}_{m \in[M]}$ of size at least $M \geq \frac{3 \sqrt{7}}{8} \cdot 8^{-(d-1)} \geq 0.99 \cdot 8^{-(d-1)}$.
The above construction can be used in place of the one described at the beginning of Appendix B. 2 The resulting lower bound on the minimax error rate follows from similar reasoning with suitably adjusted numerical constants. We sketch the rest of the proof below. By Fano's method |Wainwright. 2019 . Proposition 15.2] (see also (140)),

$$
\begin{align*}
\mathrm{M}(n, d, \delta, t) & \geq \frac{\epsilon t}{16}\left[1-\frac{I\left(J ; X_{1}^{n}\right)+\log 2}{\log M}\right]  \tag{230}\\
& \geq \frac{\epsilon t}{16}\left[1-\frac{\max _{m \in[M]} \mathrm{D}_{\mathrm{KL}}\left(P_{\theta_{m}}^{(n)} \| P_{\theta_{0}}^{(n)}\right)+\log 2}{(d-1) \log 8+\log 0.99}\right]  \tag{231}\\
& \geq \frac{\epsilon t}{16}\left[\frac{1}{2}-4 \cdot \frac{\max _{m \in[M]} \mathrm{D}_{\mathrm{KL}}\left(P_{\theta_{m}}^{(n)} \| P_{\theta_{0}}^{(n)}\right)}{d}\right] \tag{232}
\end{align*}
$$

where the last inequality follows since $d \geq 2$ implies (i) $(d-1) \log 8+\log 0.99 \geq \frac{d}{4}$, (ii) $\frac{\log 2}{(d-1) \log 8+\log 0.99} \leq \frac{\log 2}{\log 8+\log 0.99} \leq \frac{1}{2}$. The proof of the bound $\max _{m \in[M]} \mathrm{D}_{\mathrm{KL}}\left(P_{\theta_{m}}^{(n)} \| P_{\theta_{0}}^{(n)}\right) \leq$ $\frac{d}{16}$ can be completely reused which implies $\mathrm{M}(n, d, \delta, t) \geq \frac{\epsilon t}{64}$. The same choice of $\epsilon$ then yields the desired lower bound on the minimax error rate in Theorem 2 for any $d \geq 2$.

## C Proofs for Section 3: Estimation of $\delta$ for a given estimate of $\theta_{*}$

## C. 1 Proof of Theorem 4; Analysis of the estimator

Proof of Theorem 4 . We first note that since the estimator (19) does not exploit the correlation between $X_{i}$ and $X_{j}$ for $|i-j| \geq 2$, we may assume that the $n / 2$ pairs of random variables ( $S_{2 i-1}, S_{2 i}$ ) are independent by multiplying each pair of samples $\left(X_{2 i-1}, X_{2 i}\right)$ by an i.i.d. random sign. By explicitly using $X_{i}=S_{i} \theta_{*}+Z_{i}$ in the definition of $\hat{\rho}$, and using the triangle inequality, we obtain that

$$
\begin{align*}
|\hat{\rho}-\rho| \leq \left\lvert\, \frac{\left\|\theta_{*}\right\|^{2}}{\left\|\theta_{\sharp}\right\|^{2}} \frac{2}{n} \sum_{i=1}^{n / 2} S_{2 i} S_{2 i-1}\right. & -\rho\left|+\left|\frac{2}{\left\|\theta_{\sharp}\right\|^{2} n} \sum_{i=1}^{n / 2} S_{2 i} \theta_{*}^{\top} Z_{2 i-1}\right|\right. \\
& +\left|\frac{2}{\left\|\theta_{\sharp}\right\|^{2} n} \sum_{i=1}^{n / 2} S_{2 i-1} \theta_{*}^{\top} Z_{2 i}\right|+\left|\frac{2}{\left\|\theta_{\sharp}\right\|^{2} n} \sum_{i=1}^{n / 2} Z_{2 i}^{\top} Z_{2 i-1}\right| . \tag{233}
\end{align*}
$$

We begin with the analysis of the first term. We note that

$$
\begin{equation*}
\frac{2}{n} \sum_{i=1}^{n / 2} S_{2 i} S_{2 i-1} \stackrel{d}{=} \frac{2}{n} \sum_{i=1}^{n / 2} R_{i} \tag{234}
\end{equation*}
$$

where

$$
R_{i}= \begin{cases}1, & \text { w.p. } \frac{1+\rho}{2}  \tag{235}\\ -1, & \text { w.p. } \frac{1-\rho}{2}\end{cases}
$$

and $\left\{R_{i}\right\}_{i \in[n / 2]}$ are i.i.d. with $\mathbb{E}\left[R_{i}\right]=1-2 \delta=\rho$ and $\mathbb{V}\left[R_{i}\right]=4 \delta(1-\delta) \leq 4 \delta$. Now, the first term in 233 is bounded by

$$
\begin{align*}
& \left|\frac{\left\|\theta_{*}\right\|^{2}}{\left\|\theta_{\sharp}\right\|^{2}} \frac{2}{n} \sum_{i=1}^{n / 2} R_{i}-\rho\right| \\
& =\left|\left(1+\frac{\left\|\theta_{*}\right\|^{2}-\left\|\theta_{\sharp}\right\|^{2}}{\left\|\theta_{\sharp}\right\|^{2}}\right) \frac{2}{n} \sum_{i=1}^{n / 2} R_{i}-\rho\right|  \tag{236}\\
& \leq\left|\frac{2}{n} \sum_{i=1}^{n / 2} R_{i}-\rho\right|+\frac{\left|\left\|\theta_{*}\right\|^{2}-\left\|\theta_{\sharp}\right\|^{2}\right|}{\left\|\theta_{\sharp}\right\|^{2}}, \tag{237}
\end{align*}
$$

where the last inequality follows since $\left|\frac{2}{n} \sum_{i=1}^{n / 2} R_{i}\right| \leq 1$. In this last equation, the last term can be considered a bias of the estimator due to the mismatch between $\theta_{*}$ and $\theta_{\sharp}$. Now, by Bernstein's inequality for a sum of independent, zero-mean, bounded random variables $R_{i}-\mathbb{E}\left[R_{i}\right] \in[-2,2]$ (cf. (323) from [Wainwright 2019, Propostion 2.14])

$$
\begin{equation*}
\mathbb{P}\left[\left|\frac{2}{n} \sum_{i=1}^{n / 2} R_{i}-\mathbb{E}\left[R_{i}\right]\right| \geq t\right] \leq 2 \exp \left(-\frac{\frac{n}{2} t^{2}}{2\left(4 \delta+\frac{2 t}{3}\right)}\right) \tag{238}
\end{equation*}
$$

and so with probability larger than $1-\epsilon$

$$
\begin{align*}
\left|\frac{2}{n} \sum_{i=1}^{n / 2} R_{i}-\mathbb{E}\left[R_{i}\right]\right| & \leq \frac{4}{3 n} \log \left(\frac{2}{\epsilon}\right)+\frac{4}{3 n} \sqrt{9 n \delta \log \left(\frac{2}{\epsilon}\right)+\log ^{2}\left(\frac{2}{\epsilon}\right)}  \tag{239}\\
& \leq \sqrt{\frac{16 \delta \log \left(\frac{2}{\epsilon}\right)}{n}}+\frac{8 \log \left(\frac{2}{\epsilon}\right)}{3 n}  \tag{240}\\
& \leq 7 \log \left(\frac{2}{\epsilon}\right) \sqrt{\frac{\delta}{n}} \tag{241}
\end{align*}
$$

since $\delta \geq \frac{1}{n}$.
The second term in 233) (and similarly, the third term in 233) is

$$
\begin{equation*}
\frac{2}{\left\|\theta_{\sharp}\right\|^{2} n} \sum_{i=1}^{n / 2} S_{2 i} \theta_{\sharp}^{\top} Z_{2 i-1} \sim N\left(0, \frac{2}{\left\|\theta_{\sharp}\right\|^{2} n}\right) . \tag{242}
\end{equation*}
$$

Thus, by the standard Chernoff bound for Gaussian random variables

$$
\begin{equation*}
\mathbb{P}\left[\left|\frac{2}{\left\|\theta_{\sharp}\right\|^{2} n} \sum_{i=1}^{n / 2} S_{2 i} \theta_{\sharp}^{\top} Z_{2 i-1}\right|>t\right] \leq 2 e^{-\frac{\left\|\theta_{\sharp}\right\|^{2} n t^{2}}{4}}, \tag{243}
\end{equation*}
$$

and so with probability larger than $1-\epsilon$

$$
\begin{equation*}
\left|\frac{2}{\left\|\theta_{\sharp}\right\|^{2} n} \sum_{i=1}^{n / 2} S_{2 i} \theta_{\sharp}^{\top} Z_{2 i-1}\right| \leq \sqrt{\frac{4}{n\left\|\theta_{\sharp}\right\|^{2}} \log \left(\frac{2}{\epsilon}\right)} . \tag{244}
\end{equation*}
$$

The fourth term in 233) satisfies

$$
\begin{equation*}
\frac{2}{\left\|\theta_{\sharp}\right\|^{2} n} \sum_{i=1}^{n / 2} Z_{2 i}^{\top} Z_{2 i-1} \stackrel{d}{=} \frac{2}{\left\|\theta_{\sharp}\right\|^{2} n} \sum_{i=1}^{n d / 2} W_{i} \tilde{W}_{i} \tag{245}
\end{equation*}
$$

where $\left\{W_{i}\right\}_{i \in[n d / 2]}$ and $\left\{\tilde{W}_{i}\right\}_{i \in[n d / 2]}$ are i.i.d. and $W_{i}, \tilde{W}_{i} \sim N(0,1)$ are independent. It is well known [Vershynin, 2018, Lemma 2.7.7] that the product of two subGaussian random variables (even if they are not independent) is sub-exponential. Here, we have a simpler and exact characterization. Letting

$$
\begin{equation*}
W_{i} \tilde{W}_{i}=\left(\frac{W_{i}+\tilde{W}_{i}}{2}\right)^{2}-\left(\frac{W_{i}-\tilde{W}_{i}}{2}\right)^{2} \stackrel{d}{=} V_{i}^{2}-\tilde{V}_{i}^{2} \tag{246}
\end{equation*}
$$

and since $V_{i}=\frac{1}{2}\left(W_{i}+\tilde{W}_{i}\right) \sim N(0,1)$ and $\tilde{V}_{i}=\frac{1}{2}\left(W_{i}-\tilde{W}_{i}\right) \sim N(0,1)$ are uncorrelated, they are independent. Letting $\chi_{n d / 2}^{2}$ and $\tilde{\chi}_{n d / 2}^{2}$ be a pair of independent chi-square random variables with $n d / 2$ degrees of freedom, it then holds that

$$
\begin{align*}
\frac{2}{\left\|\theta_{\sharp}\right\|^{2} n} \sum_{i=1}^{n / 2} Z_{2 i}^{\top} Z_{2 i-1} & \stackrel{d}{=} \frac{2}{\left\|\theta_{\sharp}\right\|^{2} n}\left(\chi_{n d / 2}^{2}-\tilde{\chi}_{n d / 2}^{2}\right)  \tag{247}\\
& =\frac{2}{\left\|\theta_{\sharp}\right\|^{2} n}\left(\chi_{n d / 2}^{2}-\mathbb{E}\left[\chi_{n d / 2}^{2}\right]+\mathbb{E}\left[\tilde{\chi}_{n d / 2}^{2}\right]-\tilde{\chi}_{n d / 2}^{2}\right) . \tag{248}
\end{align*}
$$

From the chi-square tail bound in (330) and (331) it holds that

$$
\begin{equation*}
\mathbb{P}\left[\left|\chi_{n d / 2}^{2}-\mathbb{E}\left[\chi_{n d / 2}^{2}\right]\right| \geq 2 \sqrt{n d / 2 t}+2 t\right] \leq 2 e^{-t} \tag{249}
\end{equation*}
$$

and so it holds with probability $1-\epsilon$ that

$$
\begin{equation*}
\left|\chi_{n d / 2}^{2}-\mathbb{E}\left[\chi_{n d / 2}^{2}\right]\right| \leq \sqrt{2 n d \log \left(\frac{2}{\epsilon}\right)}+2 \log \left(\frac{2}{\epsilon}\right) \leq 4 \sqrt{n d} \log \left(\frac{2}{\epsilon}\right) \tag{250}
\end{equation*}
$$

Hence, by the union bond

$$
\begin{equation*}
\frac{2}{\left\|\theta_{\sharp}\right\|^{2} n} \sum_{i=1}^{n / 2} Z_{i}^{\top} Z_{i-1} \leq \frac{16 \log \left(\frac{2}{\epsilon}\right)}{\left\|\theta_{\sharp}\right\|^{2}} \sqrt{\frac{d}{n}} . \tag{251}
\end{equation*}
$$

with probability $1-\epsilon$. The claim (20) follows from the analysis of the terms above, the choice of $\epsilon=\frac{2}{n}$ and a union bound.

## C. 2 Proof of Proposition 6: Impossibility lower bound

The effect of knowledge of $\theta_{*} \quad$ To begin, it is apparent that if $\left\|\theta_{*}\right\|^{2}=\left\|\theta_{\sharp}\right\|^{2}$, the first term in 20) vanishes, and the loss is bounded by the remaining terms. Furthermore, if $\left\|\theta_{*}\right\| \leq 1$ then the dominant term in the brackets is $\frac{1}{\left\|\theta_{*}\right\|^{2}} \sqrt{\frac{d}{n}}$, and evidently, dominant error term suffers from a penalty of $\sqrt{d}$, even though, essentially, $\rho$ is a one-dimensional parameter. If $\theta_{*}$ is known exactly up to a sign, that is $\theta_{\sharp}= \pm \theta_{*}$ (and it is not just their norms which are equal), then the estimation error can be reduce to the case of $d=1$. This is a simple consequence of the rotational invariance of the distribution of the Gaussian noise, which implies that the projections $\left( \pm \theta_{*}^{\top} X_{i}\right)_{i=1}^{n}$ are sufficient statistics for the estimation of $\rho$. The estimator constructs the projections

$$
\begin{equation*}
U_{i}:=\frac{\theta_{\sharp}^{\top} X_{i}}{\left\|\theta_{\sharp}\right\|}= \pm\left\|\theta_{*}\right\| \cdot S_{i}+\frac{ \pm \theta_{*}^{\top} Z_{i}}{\left\|\theta_{*}\right\|} \stackrel{d}{=}\left\|\theta_{*}\right\| \cdot S_{i}+W_{i}, \tag{252}
\end{equation*}
$$

where $W_{i} \sim N(0,1)$. This is effectively a one-dimensional model with parameter given by $\left\|\theta_{*}\right\|$, and so Corollary 5 immediately follows from Theorem 4.

Tightness of the impossibility lower bound Evidently, Corollary 5 and Proposition 6 match in their dependence on the number of samples $\Theta\left(\frac{1}{\sqrt{n}}\right)$, but it is not clear what the optimal dependence on $\left\|\theta_{*}\right\|$ is. The estimator we propose is based on the moment $\rho=\mathbb{E}\left[U_{i} U_{i-1}\right]$ and can be contrasted with likelihood based methods as follows. Letting $p_{\delta}\left(s_{0}^{n}\right)$ denote the probability of the sign sequence $s_{0}^{n}$ with flip probability $\delta$, and letting $\varphi(x)$ denote the standard Gaussian density, the likelihood of $u_{1}^{n}$ is given by

$$
\begin{equation*}
P_{\delta, \theta_{*}}\left(u_{1}^{n}\right)=\sum_{s_{0}^{n} \in\{ \pm 1\}^{n+1}} p_{\delta}\left(s_{0}^{n}\right) \cdot \prod_{i \in[n]} \varphi\left(u_{i}-s_{i}\left\|\theta_{*}\right\|\right) . \tag{253}
\end{equation*}
$$

This function is a large degree polynomial in $\delta$ on the order of $n$. Even if one sums only over $s_{0}^{n}$ with the typical number of flips, then this degree is $\Theta(\delta n)$, which means this polynomial has a degree which blows up with $n$. Thus, the MLE may indeed be sensitive to empirical errors. The update of the Baum-Welch algorithm (or EM) can also be easily computed and contrasted with our proposed estimator. Letting $P_{\delta, \theta_{*}}$ denote the probability distribution of the corresponding model with flip probability $\delta$ and mean parameter $\theta_{*}$, the Baum-Welch estimate $\hat{\delta}^{(j)}$ at iteration $j$ is given by

$$
\begin{equation*}
\hat{\delta}^{(j)}=\frac{1}{n} \sum_{i=1}^{n} P_{\hat{\delta}^{(j-1)}, \theta_{*}}\left(S_{i} \neq S_{i-1} \mid u_{1}^{n}\right) \tag{254}
\end{equation*}
$$

Evidently, the inner estimate is the probability that $S_{i} \neq S_{i-1}$ conditioned on the entire sample $u_{1}^{n}$, rather than just $u_{i}, u_{i-1}$ as in our proposed estimator in $(19)$. However, deriving sharp error rate bounds for this estimator seems to be a challenging task.

In terms of minimax lower bounds, the difficulty arises since a reduction to an (almost) memoryless GMM as in lower bound for the estimation of $\theta_{*}$ in Theorem 2 does not seem fruitful. A standard application of Le-Cam's method requires bounding the total variation between models $P_{\delta_{1}, \theta_{*}}$ and $P_{\delta_{2}, \theta_{*}}$ for some "close" $\delta_{1}, \delta_{2}$ (e.g., $\delta_{1}=0$ and $\delta_{2}=\epsilon$ ). The total variation is then typically bounded by the KL divergence. However, the KL divergence does not tensorize (due to the memory), and using the chain rule requires evaluating the KL divergence for the process $U_{1}^{n}$ which is not a Markov process (but rather an HMM). Alternatively, further bounding the $n$-dimensional KL divergence with a chi-square divergence, which is a convenient choice for mixture models, leads to an excessively large bound.

The proof's main ideas The proof is based on Le-Cam's method that is applied to a genieaided model in which the estimator knows every other sign $S_{0}, S_{2}, S_{4}, \ldots, S_{n}$. The main technical challenge is then to bound the total variation for this genie-aided model. A complete proof is presented below.

Proof of Proposition 6] For a given $\delta \in[0,1]$ let $P_{\delta}^{(n)}$ denote the probability distribution of $U_{1}^{n}$ under the model 252], that is $U_{i}=t \cdot S_{i}+W_{i}, i \in[n]$, where $t \equiv\left\|\theta_{*}\right\|$ and $S_{i}$ is a binary symmetric Markov chain with flip probability $\delta$, and $\mathbb{P}\left[S_{0}=1\right]=\frac{1}{2}$. It should be noted that $P_{\delta}^{(n)}$ is a distribution on $\mathbb{R}^{n}$ which is not a product distribution. To lower bound the estimation error of an estimator $\hat{\delta}\left(U_{1}^{n}\right)$, we consider a genie-aided estimator which is informed with the values of $S_{0}, S_{2}, S_{4}, \ldots S_{n}$ (assuming for simplicity that $n$ is even). We then set $\epsilon \in\left(0, \frac{1}{2}\right)$ and use Le-Cam's two point method |Wainwright 2019. Section 15.2] with $\delta_{0}=\frac{1}{2}$ and $\delta_{1}=\frac{1}{2}-\epsilon$. Let us denote by $Q_{U_{1}^{n} S_{0}^{n}}$ (resp. $P_{U_{1}^{n}} S_{0}^{n}$ ) the joint probability distribution of $U_{1}^{n}$ and $S_{0}^{n}$ under $\delta=\frac{1}{2}$ (resp. $\delta=\frac{1}{2}-\epsilon$ ), and marginals and conditional versions by standard notation, e.g., $Q_{U_{2} U_{4} \mid S_{0} S_{2}}$. The proof of the proposition follows from Le-Cam's two point method, which states that for any estimator $\hat{\delta}\left(U_{1}^{n}, S_{0}, S_{2}, S_{4}, \ldots, S_{n}\right)$, and thus also for any less informed estimator $\hat{\delta}\left(U_{1}^{n}\right)$,

$$
\begin{equation*}
\mathbb{E}[|\hat{\delta}-\delta|] \geq \frac{\epsilon}{2} \cdot\left(1-\mathrm{d}_{\mathrm{TV}}\left(P_{U_{1} U_{2} \cdots U_{n} S_{0} S_{2} S_{4} \cdots S_{n}}, Q_{U_{1} U_{2} \cdots U_{n} S_{0} S_{2} S_{4} \cdots S_{n}}\right)\right) \tag{255}
\end{equation*}
$$

We next obtain a bound on the total variation distance in Lemma 17 of $4 \sqrt{n} \epsilon$ for $t \leq \frac{1}{\sqrt{2}}$ and choosing $\epsilon=\frac{1}{8 \sqrt{n}}$ in 255 completes the proof of the lower bound.

Lemma 17. If $t \leq 1 / \sqrt{2}$ then

$$
\begin{equation*}
\mathrm{d}_{\mathrm{TV}}\left(P_{U_{1} U_{2} \cdots U_{n}, S_{0} S_{2} S_{4} \cdots S_{n}}, Q_{U_{1} U_{2} \cdots U_{n}, S_{0} S_{2} S_{4} \cdots S_{n}}\right) \leq \sqrt{\frac{5}{2} n \epsilon}+\sqrt{8 n} t \epsilon \tag{256}
\end{equation*}
$$

Proof. By Pinsker's inequality (e.g., [Tsybakov, 2008, Lemma 2.5]), it holds for any pair of probability measures $P, Q$ that $\mathrm{d}_{\mathrm{TV}}(P, Q) \leq \sqrt{\frac{1}{2} \mathrm{D}_{\mathrm{KL}}(Q, P)}$. We next upper bound the KL divergence. To this end, recall the chain rule, that for any joint distributions $P_{X Y}$ and $Q_{X Y}$ (with conditional distributions $P_{Y \mid X}, Q_{Y \mid X}$ ) the chain rule for the KL divergence states that $\mathrm{D}_{\mathrm{KL}}\left(P_{X Y} \| Q_{X Y}\right)=\mathrm{D}_{\mathrm{KL}}\left(P_{X} \| Q_{X}\right)+\mathrm{D}_{\mathrm{KL}}\left(P_{Y \mid X} \| Q_{Y \mid X} \mid P_{X}\right)$ where $\mathrm{D}_{\mathrm{KL}}\left(P_{Y \mid X} \|\right.$ $\left.Q_{Y \mid X} \mid P_{X}\right):=\int \mathrm{D}_{\mathrm{KL}}\left(P_{Y \mid X} \| Q_{Y \mid X}\right) \mathrm{d} P_{X}$ is the conditional KL divergence [Cover and Thomas 2006, Theorem 2.5.3]. Thus,

$$
\begin{align*}
& \mathrm{D}_{\mathrm{KL}}\left(P_{U_{1} U_{2} \cdots U_{n}, S_{0} S_{2} S_{4} \cdots S_{n}} \| Q_{U_{1} U_{2} \cdots U_{n}, S_{0} S_{2} S_{4} \cdots S_{n}}\right) \\
& =\mathrm{D}_{\mathrm{KL}}\left(P_{S_{0} S_{2} S_{4} \cdots S_{n}} \| Q_{S_{0} S_{2} S_{4} \cdots S_{n}}\right)  \tag{257}\\
& \quad \quad+\mathrm{D}_{\mathrm{KL}}\left(P_{U_{1} U_{2} \cdots U_{n} \mid S_{0} S_{2} S_{4} \cdots S_{n}} \| Q_{U_{1} U_{2} \cdots U_{n} \mid S_{0} S_{2} S_{4} \cdots S_{n}} \mid P_{S_{0} S_{2} S_{4} \cdots S_{n}}\right) . \tag{258}
\end{align*}
$$

We next bound each of the two KL divergences appearing in 257) and 258. First,

$$
\begin{align*}
\mathrm{D}_{\mathrm{KL}}\left(P_{S_{0} S_{2} S_{4} \cdots S_{n}} \| Q_{S_{0} S_{2} S_{4} \cdots S_{n}}\right) & \stackrel{(a)}{\leq} \mathrm{D}_{\mathrm{KL}}\left(P_{S_{0}^{n}} \| Q_{S_{0}^{n}}\right)  \tag{259}\\
& =\sum_{s_{0}^{n} \in\{-1,1\}^{n+1}} P_{S_{0}^{n}}\left(s_{0}^{n}\right) \log \frac{P_{S_{0}^{n}}\left(s_{0}^{n}\right)}{Q_{S_{0}^{n}}\left(s_{0}^{n}\right)}  \tag{260}\\
& \stackrel{(b)}{=}(n+1) \log 2-\sum_{s_{0}^{n} \in\{-1,1\}^{n+1}} P_{S_{0}^{n}}\left(s_{0}^{n}\right) \log \frac{1}{P_{S_{0}^{n}}\left(s_{0}^{n}\right)}  \tag{261}\\
& \stackrel{(c)}{=}(n+1) \log 2-H\left(S_{0}, S_{1}, \ldots, S_{n}\right)  \tag{262}\\
& \stackrel{(d)}{=}(n+1) \log 2-H\left(S_{0}\right)-\sum_{i=1}^{n} H\left(S_{i} \mid S_{0}^{i-1}\right)  \tag{263}\\
& \stackrel{(e)}{=} n \cdot \log 2-\sum_{i=1}^{n} H\left(S_{i} \mid S_{i-1}\right)  \tag{264}\\
& \stackrel{(f)}{=} n \cdot\left[\log 2-h_{b}\left(\frac{1}{2}-\epsilon\right)\right] \tag{265}
\end{align*}
$$

where ( $a$ ) follows from the convexity of the KL divergence (recall $S_{0}^{n}=S_{0}, S_{1}, S_{2}, \ldots, S_{n}$ ), (b) since $Q_{S_{0}^{n}}\left(s_{0}^{n}\right)=2^{-(n+1)}$ for any $s_{0}^{n} \in\{ \pm 1\}^{n+1},(c)$ follows by defining the entropy $H$ of $S_{0}^{n}$ (under the probability measure $P$ ), (d) follows from the chain rule of entropy and the definition of conditional entropy [Cover and Thomas, 2006, Theorem 2.2.1], (e) follows from Markovity and $H\left(S_{0}\right)=\log 2,(f)$ follows from $H\left(S_{i} \mid S_{i-1}\right)=h_{b}\left(\delta_{1}\right)=h_{b}\left(\frac{1}{2}-\epsilon\right)$ where $h_{b}(\delta):=$ $-\delta \log \delta-(1-\delta) \log (1-\delta)$ is the binary entropy function. Now, for $\epsilon \in\left(0, \frac{1}{2}\right)$ the power series expansion of the binary entropy function results the bound

$$
\begin{align*}
h_{b}\left(\frac{1}{2}-\epsilon\right) & =\log 2-\sum_{k=1}^{\infty} \frac{(2 \epsilon)^{2 k}}{2 k(2 k-1)}  \tag{266}\\
& \geq \log 2-(2 \epsilon)^{2} \sum_{k=1}^{\infty} \frac{1}{2 k(2 k-1)}  \tag{267}\\
& \geq \log 2-4 \epsilon^{2} \sum_{k=1}^{\infty} \frac{1}{(2 k-1)^{2}}  \tag{268}\\
& =\log 2-\frac{\pi^{2}}{2} \epsilon^{2}  \tag{269}\\
& \geq \log 2-5 \epsilon^{2} \tag{270}
\end{align*}
$$

Inserting this bound into 265) results that the first term in 257) is upper bounded as

$$
\begin{equation*}
\mathrm{D}_{\mathrm{KL}}\left(P_{S_{0} S_{2} S_{4} \cdots S_{n}} \| Q_{S_{0} S_{2} S_{4} \cdots S_{n}}\right) \leq 5 n \epsilon^{2} \tag{271}
\end{equation*}
$$

We now move on to bound the second term in (258). To this end, we note that the distribution of $U_{1}^{n}$ conditioned on $S_{0}, S_{2}, S_{4}, \ldots S_{n}$ can be decomposed in a simple way. First, under $Q$, the signs $S_{0}^{n}$
are i.i.d., and so $U_{1}^{n}$ is a vector of independent samples from a Gaussian mixture model. Furthermore, $U_{i}$ depends on the sign $S_{i}$ but otherwise is independent of all other $S_{1}^{n} \backslash\left\{S_{i}\right\}$. Hence,

$$
\begin{equation*}
Q_{U_{1} U_{2} \cdots U_{n} \mid S_{0} S_{2} S_{4} \cdots S_{n}}=Q_{U_{1}} \cdot Q_{U_{2} \mid S_{2}} \cdot Q_{U_{3}} \cdot Q_{U_{4} \mid S_{4}} \cdots Q_{U_{n-1}} \cdot Q_{U_{n} \mid S_{n}} \tag{272}
\end{equation*}
$$

that is, a model of independent samples, where the odd samples are drawn from a Gaussian mixture model, and the even samples from a Gaussian location model with known sign $S_{2 i}$. Second, under $P$, an application of Bayes rule and the Markovity assumption results the decomposition to pairs of samples given by

$$
\begin{align*}
& P_{U_{1} U_{2} \cdots U_{n} \mid S_{0} S_{2} S_{4} \cdots S_{n}} \\
& =P_{U_{1} U_{2} \mid S_{0} S_{2} S_{4} \cdots S_{n}} \cdot P_{U_{3} U_{4} \mid S_{0} S_{2} S_{4} \cdots S_{n} U_{1} U_{2}} \cdots P_{U_{n-1} U_{n} \mid S_{0} S_{2} S_{4} \cdots S_{n} U_{1}^{n-2}}  \tag{273}\\
& =P_{U_{1} U_{2} \mid S_{0} S_{2}} \cdot P_{U_{3} U_{4} \mid S_{2} S_{4}} \cdots P_{U_{n-1} U_{n} \mid S_{n-2} S_{n}} . \tag{274}
\end{align*}
$$

The chain rule for the KL divergence therefore implies that

$$
\begin{align*}
& \mathrm{D}_{\mathrm{KL}}\left(P_{U_{1} U_{2} \cdots U_{n} \mid S_{0} S_{2} S_{4} \cdots S_{n}} \| Q_{U_{1} U_{2} \cdots U_{n} \mid S_{0} S_{2} S_{4} \cdots S_{n}} \mid P_{S_{0} S_{2} S_{4} \cdots S_{n}}\right) \\
& \stackrel{(a)}{=} \mathrm{D}_{\mathrm{KL}}\left(P_{U_{1} U_{2} \mid S_{0} S_{2}} \| Q_{U_{1}} \cdot Q_{U_{2} \mid S_{2}} \mid P_{S_{0} S_{2}}\right)+\mathrm{D}_{\mathrm{KL}}\left(P_{U_{3} U_{4} \mid S_{2} S_{4}} \| Q_{U_{3}} \cdot Q_{U_{4} \mid S_{4}} \mid P_{S_{2} S_{4}}\right) \\
& \quad+\cdots+\mathrm{D}_{\mathrm{KL}}\left(P_{U_{n-1} U_{n} \mid S_{n-2} S_{n}} \| Q_{U_{n-1}} \cdot Q_{U_{n} \mid S_{n}} \mid P_{S_{n-2} S_{n}}\right)  \tag{275}\\
& \stackrel{(b)}{=} \frac{n}{2} \cdot \mathrm{D}_{\mathrm{KL}}\left(P_{U_{1} U_{2} \mid S_{0} S_{2}} \| Q_{U_{1}} \cdot Q_{U_{2} \mid S_{2}} \mid P_{S_{0} S_{2}}\right)  \tag{276}\\
& \stackrel{(c)}{=} \frac{n}{2} \cdot \mathrm{D}_{\mathrm{KL}}\left(P_{U_{1} \mid S_{0} S_{2}} \| Q_{U_{1}} \mid P_{S_{0} S_{2}}\right)+\frac{n}{2} \cdot \mathrm{D}_{\mathrm{KL}}\left(P_{U_{2} \mid S_{0} S_{2}} \| Q_{U_{2} \mid S_{2}} \mid P_{S_{0} S_{2}}\right)  \tag{277}\\
& \stackrel{(d)}{=} \frac{n}{2} \cdot \mathrm{D}_{\mathrm{KL}}\left(P_{U_{1} \mid S_{0} S_{2}} \| Q_{U_{1}} \mid P_{S_{0} S_{2}}\right)  \tag{278}\\
& \stackrel{(e)}{=} \frac{n}{2} \cdot \mathrm{D}_{\mathrm{KL}}\left(P_{U_{1} \mid S_{0}=1, S_{2}} \| Q_{U_{1}} \mid P_{S_{2} \mid S_{0}=1}\right) \tag{279}
\end{align*}
$$

where $(a)$ follows from $\sqrt{274}$ and the chain rule for KL divergence, and (b) follows from the stationarity of the Markov chain, $(c)$ follows again from the chain rule, $(d)$ follows since $\mathrm{D}_{\mathrm{KL}}\left(P_{U_{2} \mid S_{0} S_{2}} \| Q_{U_{2} \mid S_{2}} \mid P_{S_{0} S_{2}}\right)=\mathrm{D}_{\mathrm{KL}}\left(P_{U_{2} \mid S_{2}} \| Q_{U_{2} \mid S_{2}} \mid P_{S_{2}}\right)=0$, and (e) follows since by symmetry, we may condition on $S_{0}=1$. Thus, the last KL divergence in 279 should be averaged over the cases $S_{2}=-1$ and $S_{2}=1$. For the first case, it holds that $P_{S_{1} \mid S_{0} S_{2}}(\cdot \mid 1,-1)$ is a uniform distribution on $\{ \pm 1\}$. Hence, conditioned on $S_{0}=1, S_{2}=-1$, under $P, U_{1}$ is a sample from a balanced Gaussian mixture, just as under $Q$. So,

$$
\begin{equation*}
\mathrm{D}_{\mathrm{KL}}\left(P_{U_{1} \mid S_{0}=1, S_{2}=-1} \| Q_{U_{1}}\right)=0 \tag{280}
\end{equation*}
$$

So, the KL divergence is only comprised of the term in the second case $S_{2}=1$. Continuing to evaluate the KL divergence in 279, we get

$$
\begin{equation*}
\mathrm{D}_{\mathrm{KL}}\left(P_{U_{1} \mid S_{0}=1, S_{2}} \| Q_{U_{1}} \mid P_{S_{2} \mid S_{0}=1}\right)=P_{S_{2} \mid S_{0}=1}(1) \cdot \mathrm{D}_{\mathrm{KL}}\left(P_{U_{1} \mid S_{0}=1, S_{2}=1} \| Q_{U_{1}}\right) \tag{281}
\end{equation*}
$$

$$
\begin{align*}
& \stackrel{(a)}{\leq} \mathrm{D}_{\mathrm{KL}}\left(P_{U_{1} \mid S_{0}=1, S_{2}=1} \| Q_{U_{1}}\right)  \tag{282}\\
& \stackrel{(b)}{\leq} \chi^{2}\left(P_{U_{1} \mid S_{0}=1, S_{2}=1} \| Q_{U_{1}}\right) \tag{283}
\end{align*}
$$

where $(a)$ follows since $P_{S_{2} \mid S_{0}}(1 \mid 1)=\left(1-\delta_{1}\right)^{2}+\delta_{1}^{2} \leq 1$, and (b) follows from the chisquare divergence bound on the KL divergence (e.g., [Tsybakov, 2008, Eq. (2.27)]). Now, recall that under $Q_{U_{1}}$, it holds that $U_{1} \sim \frac{1}{2} N(t, 1)+\frac{1}{2} N(-t, 1)$, and under $P_{U_{1} \mid S_{0}=1, S_{2}=1}$ it holds that $U_{1} \sim(1-\alpha) \cdot N(t, 1)+\alpha N(-t, 1)$ where

$$
\begin{equation*}
\alpha=P_{S_{1} \mid S_{0} S_{2}}(-1 \mid 1,1)=\frac{\delta_{1}^{2}}{\left(1-\delta_{1}\right)^{2}+\delta_{1}^{2}} . \tag{284}
\end{equation*}
$$

Thus, the chi-square divergence in 283 is between a balanced Gaussian mixture (with probability that $S_{1}=1$ being $\frac{1}{2}$ ) and an unbalanced Gaussian mixture (with probability that $S_{1}=1$ being $1-\alpha>\frac{1}{2}$ ). Let $\psi_{\beta}(u)$ denote the probability density function of $U=t \cdot S+W$ with $\mathbb{P}[S=-1]=$ $1-\mathbb{P}[S=1]=\beta$ and $W \sim N(0,1)$. Then, if $\varphi(u) \equiv \frac{1}{\sqrt{2 \pi}} e^{-u^{2} / 2}$ is the standard Gaussian density function, it holds that

$$
\begin{equation*}
\psi_{\beta}(u)=\beta \cdot \varphi(u+t)+(1-\beta) \cdot \varphi(u+t) \tag{285}
\end{equation*}
$$

$$
\begin{align*}
& =\beta \cdot \frac{1}{\sqrt{2 \pi}} e^{-(u+t)^{2} / 2}+(1-\beta) \cdot \frac{1}{\sqrt{2 \pi}} e^{-(u-t)^{2} / 2}  \tag{286}\\
& =e^{-t^{2} / 2} \cdot \varphi(u) \cdot\left[\beta \cdot e^{-u t}+(1-\beta) \cdot e^{u t}\right] \tag{287}
\end{align*}
$$

Specifically, for $\beta=\frac{1}{2}$ it holds that $\psi_{1 / 2}(u)=e^{-t^{2} / 2} \cdot \varphi(u) \cdot \cosh (u t)$. Therefore, the chi-square divergence from 283 is upper bounded as

$$
\begin{align*}
& \chi^{2}\left(P_{U_{1} \mid S_{0}=1, S_{2}=1} \| Q_{U_{1}}\right) \\
& =\int_{-\infty}^{\infty} \frac{\left[P_{U_{1} \mid S_{0} S_{2}}(u \mid 1,1)-Q_{U_{1}}(u)\right]^{2}}{Q_{U_{1}}(u)} \mathrm{d} u  \tag{288}\\
& =e^{-t^{2} / 2} \cdot \int_{-\infty}^{\infty} \varphi(u) \frac{\left[\left(\alpha-\frac{1}{2}\right) \cdot e^{-u t}+\left(1-\alpha-\frac{1}{2}\right) \cdot e^{u t}\right]^{2}}{\cosh (u t)} \mathrm{d} u  \tag{289}\\
& \stackrel{(a)}{\leq} e^{-t^{2} / 2} \cdot \int_{-\infty}^{\infty} \varphi(u)\left[\left(\alpha-\frac{1}{2}\right) \cdot e^{-u t}+\left(1-\alpha-\frac{1}{2}\right) \cdot e^{u t}\right]^{2} \mathrm{~d} u  \tag{290}\\
& =e^{-t^{2} / 2} \cdot \int_{-\infty}^{\infty} \varphi(u)\left[\left(\alpha-\frac{1}{2}\right)^{2} \cdot e^{-2 u t}+2\left(\alpha-\frac{1}{2}\right)\left(1-\alpha-\frac{1}{2}\right)+\left(1-\alpha-\frac{1}{2}\right)^{2} \cdot e^{2 u t}\right] \mathrm{d} u  \tag{291}\\
& =e^{-t^{2} / 2} \cdot\left[\left(\alpha-\frac{1}{2}\right)^{2} e^{2 t^{2}}+2\left(\alpha-\frac{1}{2}\right)\left(1-\alpha-\frac{1}{2}\right)+\left(1-\alpha-\frac{1}{2}\right)^{2} e^{2 t^{2}}\right]  \tag{292}\\
& (b)  \tag{293}\\
& \left.\leq\left(\alpha-\frac{1}{2}\right)^{2}\left(1+4 t^{2}\right)+2\left(\alpha-\frac{1}{2}\right)\left(1-\alpha-\frac{1}{2}\right)+\left(1-\alpha-\frac{1}{2}\right)^{2}\left(1+4 t^{2}\right)\right]  \tag{294}\\
& =4 t^{2} \cdot\left[\left(\alpha-\frac{1}{2}\right)^{2}+\left(1-\alpha-\frac{1}{2}\right)^{2}\right]  \tag{295}\\
& =2 t^{2} \cdot[1-4 \alpha(1-\alpha)]  \tag{296}\\
& \stackrel{(c)}{=} 2 t^{2} \cdot\left[1-\left(\frac{2-8 \epsilon^{2}}{2+8 \epsilon^{2}}\right)^{2}\right]  \tag{297}\\
& =2 t^{2} \cdot\left[1-\frac{2-8 \epsilon^{2}}{2+8 \epsilon^{2}}\right]\left[1+\frac{2-8 \epsilon^{2}}{2+8 \epsilon^{2}}\right]  \tag{298}\\
& =2 t^{2} \cdot\left[\frac{64 \epsilon^{2}}{\left(2+8 \epsilon^{2}\right)^{2}}\right]  \tag{299}\\
& \leq 32 t^{2} \epsilon^{2}
\end{align*}
$$

where $(a)$ follows since $\cosh (x) \geq 1$ for all $x \in \mathbb{R},(b)$ follows using the assumption $t \leq \frac{1}{\sqrt{2}}$, and so $e^{2 t^{2}} \leq 1+4 t^{2}$, and $(c)$ is by setting

$$
\begin{equation*}
4 \alpha(1-\alpha)=4 \frac{\delta_{1}^{2}\left(1-\delta_{1}\right)^{2}}{\left[\left(1-\delta_{1}\right)^{2}+\delta_{1}^{2}\right]^{2}}=4 \frac{\left(\frac{1}{2}-\epsilon\right)^{2}\left(\frac{1}{2}+\epsilon\right)^{2}}{\left[\left(\frac{1}{2}+\epsilon\right)^{2}+\left(\frac{1}{2}-\epsilon\right)^{2}\right]^{2}}=\left(\frac{2-8 \epsilon^{2}}{2+8 \epsilon^{2}}\right)^{2} \tag{300}
\end{equation*}
$$

Using the bound 299) in 283, We thus conclude that

$$
\begin{equation*}
\mathrm{D}_{\mathrm{KL}}\left(P_{U_{1} \mid S_{0}=1, S_{2}} \| Q_{U_{1}} \mid P_{S_{2} \mid S_{0}=1}\right) \leq 32 t^{2} \epsilon^{2} \tag{301}
\end{equation*}
$$

and then in 279) that the second KL term of 258) is upper bounded as

$$
\begin{equation*}
\mathrm{D}_{\mathrm{KL}}\left(\overline{\left.P_{U_{1} U_{2} \cdots U_{n} \mid S_{0} S_{2} S_{4} \cdots S_{n}} \| Q_{U_{1} U_{2} \cdots U_{n} \mid S_{0} S_{2} S_{4} \cdots S_{n}} \mid P_{S_{0} S_{2} S_{4} \cdots S_{n}}\right) \leq 16 n t^{2} \epsilon^{2} . . . ~}\right. \tag{302}
\end{equation*}
$$

Combining this bound with the bound on the first KL term of 258 given in 271 we obtain

$$
\begin{equation*}
\mathrm{D}_{\mathrm{KL}}\left(P_{U_{1} U_{2} \cdots U_{n}, S_{0} S_{2} S_{4} \cdots S_{n}} \| Q_{U_{1} U_{2} \cdots U_{n}, S_{0} S_{2} S_{4} \cdots S_{n}}\right) \leq 5 n \epsilon^{2}+16 n t^{2} \epsilon^{2} \tag{303}
\end{equation*}
$$

The aforementioned Pinsker bound on the total variation distance, and the relation $\sqrt{a+b} \leq \sqrt{a}+\sqrt{b}$ then completes the proof.

## D Proofs for Section 3: Analysis of Algorithm 1

Proof of Theorem 7. Before presenting the analysis of Algorithm11, we first specify the choices of the constants $\lambda_{\theta} \geq 1$ and $\lambda_{\delta} \geq 1$. Recall that from Theorem 1 for a Markov model $\left(\theta_{*}, \delta\right)$ in low dimension, $d \leq \delta n$, there exists a numerical constant $\lambda_{\theta}>0 \cdot{ }^{3}$ so that the estimator $\hat{\theta} \equiv \hat{\theta}_{\operatorname{cov}}\left(X_{1}^{n} ; k=\right.$ $\left.\frac{1}{8 \delta}\right)$, which assumes a perfect knowledge of $\delta$, achieves with probability $1-O\left(\frac{1}{n}\right)$,

$$
\operatorname{loss}\left(\hat{\theta}, \theta_{*}\right) \leq \lambda_{\theta} \cdot \log (n) \cdot \begin{cases}\left\|\theta_{*}\right\|, & \left\|\theta_{*}\right\| \leq\left(\frac{\delta d}{n}\right)^{1 / 4}  \tag{304}\\ \frac{1}{\left\|\theta_{*}\right\|} \sqrt{\frac{\delta d}{n}}, & \left(\frac{\delta d}{n}\right)^{1 / 4} \leq\left\|\theta_{*}\right\| \leq \sqrt{\delta} \\ \sqrt{\frac{d}{n}}, & \left\|\theta_{*}\right\| \geq \sqrt{\delta}\end{cases}
$$

if $d \leq \delta n$, and

$$
\operatorname{loss}\left(\hat{\theta}, \theta_{*}\right) \leq \lambda_{\theta} \cdot \log (n) \cdot \begin{cases}\left\|\theta_{*}\right\|, & \left\|\theta_{*}\right\| \leq \sqrt{\frac{d}{n}}  \tag{305}\\ \sqrt{\frac{d}{n}}, & \left\|\theta_{*}\right\| \geq \sqrt{\frac{d}{n}}\end{cases}
$$

if $\delta n \leq d \leq n$. We assume that $\lambda_{\theta} \geq 1$, and otherwise replace $\lambda_{\theta}$ by 1 .
In addition, it holds from Theorem 4 that for the estimator $\hat{\delta}$ for $\delta$, which is based on a mismatched mean $\theta_{\sharp}$, i.e.,

$$
\begin{equation*}
\hat{\delta} \equiv \hat{\delta}_{\text {corr }}\left(X_{1}^{n} ; \theta_{\sharp}\right):=\frac{1}{2}\left(1-\hat{\rho}_{\text {corr }}\left(X_{1}^{n} ; \theta_{\sharp}\right)\right), \tag{306}
\end{equation*}
$$

there exists a numerical constant $\lambda_{\delta}>0$, for which with probability larger than $1-O\left(\frac{1}{n}\right)$,

$$
\begin{equation*}
\operatorname{loss}(\hat{\delta}, \delta) \leq \lambda_{\delta}\left[\frac{\left|\left\|\theta_{*}\right\|^{2}-\left\|\theta_{\sharp}\right\|^{2}\right|}{\left\|\theta_{\sharp}\right\|^{2}}+\frac{\log (n)}{\left\|\theta_{\sharp}\right\|^{2}} \sqrt{\frac{d}{n}}\right] \tag{307}
\end{equation*}
$$

assuming that $\left\|\theta_{\sharp}\right\| \leq 2$ in order to simplify the bound to the regime of interest. We assume here too that $\lambda_{\delta} \geq 1$, and otherwise replace $\lambda_{\delta}$ by 1 .
For the analysis we assume that all three steps of the algorithm are successful estimation events. By the union bound, this occurs with probability $1-O\left(\frac{1}{n}\right)$.
Analysis of Step A:
First, suppose that

$$
\begin{equation*}
\left\|\theta_{*}\right\| \leq \lambda_{\theta} \log (n) \cdot\left(\frac{d}{n}\right)^{1 / 4} \tag{308}
\end{equation*}
$$

Then, it holds by (304) and (305) that

$$
\begin{equation*}
\left\|\hat{\theta}^{(A)}\right\| \leq\left\|\theta_{*}\right\|+\lambda_{\theta} \log (n) \cdot\left(\frac{\delta d}{n}\right)^{1 / 4} \leq 2 \lambda_{\theta} \log (n)\left(\frac{d}{n}\right)^{1 / 4} \tag{309}
\end{equation*}
$$

and thus the algorithm will stop and output $\hat{\theta}=0$, for which it holds that $\operatorname{loss}\left(\theta_{*}, \hat{\theta}^{(A)}\right)=\left\|\theta_{*}\right\|$. This agrees with (26) and (27).
Second, suppose that $\left\|\theta_{*}\right\| \geq 1$. Now, it holds by 304 and 305

$$
\begin{equation*}
\left\|\hat{\theta}^{(A)}\right\| \geq\left\|\theta_{*}\right\|-\lambda_{\theta} \log (n) \sqrt{\frac{d}{n}} \geq \frac{1}{2} \tag{310}
\end{equation*}
$$

by the assumption $d \leq \frac{1}{4 \lambda_{\theta}^{2} \log ^{2}(n)} \cdot n$. Thus the algorithm will stop and output $\hat{\theta}=\hat{\theta}^{(A)}$, for which it also holds that

$$
\begin{equation*}
\operatorname{loss}\left(\theta_{*}, \hat{\theta}^{(A)}\right) \leq \lambda_{\theta} \cdot \log (n) \cdot \sqrt{\frac{d}{n}} \tag{311}
\end{equation*}
$$

[^0](irrespective of the value of $\delta$ ). This also agrees with (26) and (27).
Thus, the algorithm will continue to Step B only if
\[

$$
\begin{equation*}
\lambda_{\theta} \log (n) \cdot\left(\frac{d}{n}\right)^{1 / 4} \leq\left\|\theta_{*}\right\| \leq 1 \tag{312}
\end{equation*}
$$

\]

which we henceforth assume. As a preparation for Step B, in which $\hat{\theta}^{(A)}$ will play the rule of $\theta_{\sharp}$ (the mismatched mean estimator of Section 3) in an estimation of $\delta$, we bound the absolute value $\left|\left\|\hat{\theta}^{(A)}\right\|-\left\|\theta_{*}\right\|\right|$ and also show that $\left\|\hat{\theta}^{(A)}\right\| \geq \frac{1}{2}\left\|\theta_{*}\right\|$. To this end, we may repeat the arguments of the proof of Theorem 1, and specifically, arguments similar to 126) to show that

$$
\begin{align*}
\left|\left\|\hat{\theta}^{(A)}\right\|-\left\|\theta_{*}\right\|\right| & \lesssim \log (n) \cdot \psi\left(n, d, \delta=\frac{1}{2}, k=1\right)  \tag{313}\\
& =\log (n) \cdot\left[2 \sqrt{\frac{\delta}{n}} \cdot\left\|\theta_{*}\right\|^{2}+2 \sqrt{\frac{d}{n}} \cdot\left\|\theta_{*}\right\|+13 \sqrt{\frac{d}{n}}+10 \frac{d}{n}\right] . \tag{314}
\end{align*}
$$

Note that the probability of this event is $1-O\left(\frac{1}{n}\right)$, and is included in the successful estimation event mentioned in the beginning of the proof. Under the assumption $\left\|\theta_{*}\right\| \leq 1$ it then holds that

$$
\begin{equation*}
\left|\left\|\hat{\theta}^{(A)}\right\|-\left\|\theta_{*}\right\|\right| \leq \lambda_{\theta} \log (n) \cdot \sqrt{\frac{d}{n}} \tag{315}
\end{equation*}
$$

Note that we take $\lambda_{\theta}>0$ to be large enough so that this holds. The assumed case $\left\|\theta_{*}\right\| \geq \lambda_{\theta} \log (n)$. $\left(\frac{d}{n}\right)^{1 / 4}$, and the assumption of the theorem $d \leq \frac{n}{16}$ then imply that

$$
\begin{equation*}
\frac{1}{2}\left\|\theta_{*}\right\| \leq\left\|\hat{\theta}^{(A)}\right\| \leq 2\left\|\theta_{*}\right\| \leq 2 \tag{316}
\end{equation*}
$$

## Analysis of Step B:

Setting $\theta_{\sharp}=\hat{\theta}^{(A)}$ and utilizing the two properties just derived in 315 and 316, 307 implies that

$$
\begin{align*}
\operatorname{loss}\left(\hat{\delta}^{(B)}, \delta\right) & \leq \lambda_{\delta}\left[\frac{\left|\left\|\theta_{*}\right\|^{2}-\left\|\theta_{\sharp}\right\|^{2}\right|}{\left\|\theta_{\sharp}\right\|^{2}}+\frac{\log (n)}{\left\|\theta_{\sharp}\right\|^{2}} \sqrt{\frac{d}{n}}\right]  \tag{317}\\
& \leq 4 \lambda_{\delta} \lambda_{\theta} \frac{\log (n)}{\left\|\theta_{*}\right\|^{2}} \sqrt{\frac{d}{n}} . \tag{318}
\end{align*}
$$

There are two cases. First suppose that $\left\|\theta_{*}\right\| \geq \sqrt{8 \lambda_{\delta} \lambda_{\theta} \log (n)}\left(\frac{d}{\delta^{2} n}\right)^{1 / 4}$. In this case,

$$
\begin{equation*}
\operatorname{loss}\left(\hat{\delta}^{(B)}, \delta\right) \leq \frac{\delta}{2} \tag{319}
\end{equation*}
$$

and so

$$
\begin{equation*}
\frac{1}{2} \delta \leq \hat{\delta}^{(B)} \leq 2 \delta \tag{320}
\end{equation*}
$$

In addition, from 316 it holds that

$$
\begin{equation*}
64 \lambda_{\delta} \lambda_{\theta} \frac{\log (n)}{\left\|\hat{\theta}^{(A)}\right\|^{2}} \sqrt{\frac{d}{n}} \geq 16 \lambda_{\delta} \lambda_{\theta} \frac{\log (n)}{\left\|\theta_{*}\right\|^{2}} \sqrt{\frac{d}{n}} \geq 2 \delta \geq \hat{\delta}^{(B)} \tag{321}
\end{equation*}
$$

and so the algorithm will continue to Step C. Now, suppose that $\left\|\theta_{*}\right\| \leq \sqrt{8 \lambda_{\delta} \lambda_{\theta} \log (n)}\left(\frac{d}{\delta^{2} n}\right)^{1 / 4}$. If the algorithm stops then $\hat{\theta}=\hat{\theta}^{(A)}$, and the estimation rates of the Gaussian mixture model are achieved, which agrees with 26 and 27 . If the algorithm does not stop and proceeds to Step C then it holds that

$$
\begin{equation*}
\hat{\delta}^{(B)} \geq 64 \lambda_{\delta} \lambda_{\theta} \frac{\log (n)}{\left\|\hat{\theta}^{(A)}\right\|^{2}} \sqrt{\frac{d}{n}} \geq 16 \lambda_{\delta} \lambda_{\theta} \frac{\log (n)}{\left\|\theta_{*}\right\|^{2}} \sqrt{\frac{d}{n}} \geq \delta \tag{322}
\end{equation*}
$$

by 316 and the assumption $\left\|\theta_{*}\right\| \leq \sqrt{8 \lambda_{\delta} \lambda_{\theta} \log (n)}\left(\frac{d}{\delta^{2} n}\right)^{1 / 4}$.

To conclude, there are cases in which the algorithm proceeds to Step C. If $\left\|\theta_{*}\right\| \geq$ $\sqrt{8 \lambda_{\delta} \lambda_{\theta} \log (n)}\left(\frac{d}{\delta^{2} n}\right)^{1 / 4}$ then it proceeds Step C and 320 holds. Otherwise, the algorithm might proceed to Step C, yet now only $\hat{\delta}^{(B)} \geq \frac{1}{2} \delta$ is assured.
Analysis of Step C:
If the algorithm has proceeded to Step $C$, then it is guaranteed that $\hat{\delta}^{(B)} \geq \frac{1}{2} \delta$ (in any event). Recall that, ideally, had $\delta$ was known, the choice of the block length $k$ for the estimator $\hat{\theta}_{\text {cov }}\left(X_{2 n+1}^{3 n} ; k\right)$ is $k_{*}:=\frac{1}{8 \delta}$ (Section 2). It can be readily verified that the analysis of Section 2 is valid when using any smaller blocklength $k$, and that the error rate improve as $k$ increases from $k=1$ to $k=k_{*}$. In accordance to the analysis of Step $B$, there are two possible cases. If $\left\|\theta_{*}\right\| \geq \sqrt{8 \lambda_{\delta} \lambda_{\theta} \log (n)}\left(\frac{d}{\delta^{2} n}\right)^{1 / 4}$ then it holds that $\frac{1}{2} \delta \leq \hat{\delta}^{(B)} \leq 2 \delta$. Using $k=\frac{1}{16 \hat{\delta}(B)}$ then assures that $k \leq k_{*}$. In addition, since the error bound scales linearly with $\psi(n, d, \delta, k)$ (see 92 ), the error can increase by a factor at most 2 . Thus, $\hat{\theta}^{(C)}$ achieves the error rates in 304) and 305), with a factor of 2. There are again two cases to consider. If $d \geq \frac{1}{64 \lambda_{\delta}^{2} \lambda_{\theta}^{2} \log ^{2}(n)} \delta^{4} n$, then the interval $\left(\sqrt{8 \lambda_{\delta} \lambda_{\theta} \log (n)}\left(\frac{d}{\delta^{2} n}\right)^{1 / 4}, \sqrt{\delta}\right)$ is empty (its right end point is smaller than its left end point), and plugging $k$ in the error rates of (304) and 305 , the resulting error rates are as in the Gaussian mixture model. If $d \leq \frac{1}{64 \lambda_{\delta}^{2} \lambda_{\theta}^{2} \log ^{2}(n)} \delta^{4} n$ then the aforementioned interval is non-empty, and the resulting error rates agree with (26) and 27. Otherwise, if $\left\|\theta_{*}\right\| \leq \sqrt{8 \lambda_{\delta} \lambda_{\theta} \log (n)}\left(\frac{d}{\delta^{2} n}\right)^{1 / 4}$, and the estimator $\hat{\theta}_{\text {cov }}\left(X_{2 n+1}^{3 n} ; k\right)$ operates with $k \leq k_{*}$, but $k$ can be as low as 1 . Thus, error rates of the Gaussian mixture are again achieved, and this agrees with 26) and 27.

## E Useful results

Bernstein's inequality Let $X_{1}, \cdots, X_{\ell}$ be independent random variables and $\left|X_{i}\right| \leq b$ almost surely for every $i \in[\ell]$. Then [Wainwright, 2019, Proposition 2.14 ${ }^{4}$ states

$$
\begin{equation*}
\mathbb{P}\left[\left|\frac{1}{\ell} \sum_{i=1}^{\ell}\left(X_{i}-\mathbb{E}\left[X_{i}\right]\right)\right| \geq \delta\right] \leq 2 \exp \left(-\frac{\ell \delta^{2} / 2}{\frac{1}{\ell} \sum_{i=1}^{\ell} \mathbb{V}\left[X_{i}\right]+b \delta / 3}\right) \tag{323}
\end{equation*}
$$

Norm of subGaussian random vectors A random vector $X \in \mathbb{R}^{d}$ is said to be $\sigma^{2}$-subGaussian if $\mathbb{E}[X]=0$ and

$$
\begin{equation*}
\mathbb{E}\left[\exp \left(t v^{\top} X\right)\right] \leq \exp \left(\frac{\sigma^{2} t^{2}}{2}\right) \tag{324}
\end{equation*}
$$

for every $t \in \mathbb{R}$ and every $v \in \mathbb{S}^{d-1}$. Let $X \in \mathbb{R}^{d}$ be a $\sigma^{2}$-subGaussian random vector. Then from |Rigollet and Hütter 2019, Theorem 1.19], we have for any $\delta>0$,

$$
\begin{equation*}
\mathbb{P}\left[\max _{v \in \mathbb{B}^{d}} v^{\top} X \leq 4 \sigma \sqrt{d}+2 \sigma \sqrt{2 \log \left(\frac{1}{\delta}\right)}\right]=\mathbb{P}\left[\|X\| \leq 4 \sigma \sqrt{d}+2 \sigma \sqrt{2 \log \left(\frac{1}{\delta}\right)}\right] \geq 1-\delta . \tag{325}
\end{equation*}
$$

Gaussian covariance estimation Let $W \in \mathbb{R}^{\ell \times d}$ be a matrix with i.i.d. $N(0,1)$ entries. Then [Wainwright, 2019, Example 6.2] implies that for any $\delta>0$,

$$
\begin{equation*}
\mathbb{P}\left[\left\|\frac{1}{\ell} W^{\top} W-I_{d}\right\|_{\mathrm{op}} \leq 2\left(\sqrt{\frac{d}{n}}+\delta\right)+\left(\sqrt{\frac{d}{n}}+\delta\right)^{2}\right] \geq 1-2 e^{-n \delta^{2} / 2} \tag{326}
\end{equation*}
$$

Davis-Kahan's perturbation bound Let $\Sigma, \hat{\Sigma}$ be symmetric matrices with the same dimensions. Let $\lambda_{i}(\Sigma)$ and $v_{i}(\Sigma)$ denote the $i$ th largest eigenvalue and the associated eigenvector (of unit norm) of $\Sigma$. Fix $i$ and assume that $\lambda_{i}(\Sigma)$ is well-separated from the rest of the spectrum of $\Sigma$ :

$$
\begin{equation*}
\min _{j \neq i}\left|\lambda_{i}(\Sigma)-\lambda_{j}(\Sigma)\right|=\delta>0 \tag{327}
\end{equation*}
$$

[^1]Then from [Vershynin, 2018, Theorem 4.5.5], we have

$$
\begin{equation*}
\operatorname{loss}\left(v_{i}(\Sigma), v_{i}(\hat{\Sigma})\right) \leq 4 \frac{\|\Sigma-\hat{\Sigma}\|_{\mathrm{op}}}{\delta} \tag{328}
\end{equation*}
$$

where loss $(\cdot, \cdot)$ was defined in (3).
Packing number Let $M\left(\delta ; \mathbb{B}^{d},\|\cdot\|\right)$ and $N\left(\delta ; \mathbb{B}^{d},\|\cdot\|\right)$ denote the $\delta$-packing number and $\delta$ covering number of $\mathbb{B}^{d}$ w.r.t. $\|\cdot\|$, respectively. From Wainwright 2019, Lemma 5.7 and Example 5.8], we have

$$
\begin{equation*}
M\left(\delta ; \mathbb{B}^{d},\|\cdot\|\right) \geq N\left(\delta ; \mathbb{B}^{d},\|\cdot\|\right) \geq\left(\frac{1}{\delta}\right)^{d} \tag{329}
\end{equation*}
$$

Chi-square tail bounds From the chi-square tail bound in [Boucheron et al., 2013, Remark 2.11]

$$
\begin{equation*}
\mathbb{P}\left[\chi_{\ell}^{2}-\ell \geq 2 \sqrt{\ell t}+2 t\right] \leq e^{-t} \tag{330}
\end{equation*}
$$

and

$$
\begin{equation*}
\mathbb{P}\left[\chi_{\ell}^{2}-\ell \leq-2 \sqrt{\ell t}\right] \leq e^{-t} \tag{331}
\end{equation*}
$$

## F Numerical validation

In this section, we provide numerical validation of the performance of the estimators $\hat{\theta}_{\text {cov }}\left(X_{1}^{n} ; k\right)$ (cf. (14)), $\hat{\delta}_{\text {corr }}\left(X_{1}^{n} ; \theta_{\sharp}\right)=\frac{1}{2}\left(1-\hat{\rho}_{\text {corr }}\left(X_{1}^{n} ; \theta_{\sharp}\right)\right)$ (cf. (19)) and Algorithm 1 proposed and analyzed in Theorem 1, Theorem 4 and Corollary 5, and Theorem 7, respectively. Numerical results for these estimators/algorithm are plotted in Figures 3,4 and 5 , respectively.


Figure 3: Numerical validation of the performance of the $\theta_{*}$-estimator $\hat{\theta}_{\text {cov }}\left(X_{1}^{n} ; k\right)$ (cf. (14)) used in the proof of Theorem 1 which assumes a known $\delta$. We take the minimum loss (cf. (3)) achieved by $\hat{\theta}_{\text {cov }}\left(X_{1}^{n} ; k=\frac{1}{8 \delta}\right)$ and the trivial estimator $\hat{\theta}_{0}\left(X_{1}^{n}\right)=0$. We plot the loss as a function of $t=\left\|\theta_{*}\right\|$ for $t \in[0,5]$ with step size 0.05 . We take $n=5000, d=250, \delta=0.05$. Therefore, $\delta>\frac{1}{n}$ and $d<\delta n$.

## G Open directions

We discuss some open directions pertaining to Remark 3

1. Suppose $Z_{i} \sim N(0, \Sigma)$ i.i.d. for some general covariance $\Sigma \succ 0$. If $\Sigma$ is unknown, in contrast to the observation above, then the problem becomes significantly more delicate and challenging. A well-known and intuitive example (see, e.g., [Ferguson, 1982]) shows that the maximum likelihood estimator (MLE) does not exist even for estimating the mean $\mu$ and variance $\sigma^{2}$ of a Gaussian mixture with two components $N(0,1)$ and $N\left(\mu, \sigma^{2}\right)$ where $\mu \in \mathbb{R}, \sigma \in \mathbb{R}_{+}$. In fact, fitting both mean and scale parameters was studied in Dwivedi


Figure 4: Numerical validation of the performance of the $\delta$-estimator $\hat{\delta}_{\text {corr }}\left(X_{1}^{n} ; \theta_{\sharp}\right)=\frac{1}{2}(1-$ $\left.\hat{\rho}\left(X_{1}^{n} ; \theta_{\sharp}\right)\right)$ where $\hat{\rho}_{\text {corr }}\left(X_{1}^{n} ; \theta_{\sharp}\right)$ is given by (19). This estimator is used in Theorem 4 (which assumes $\theta_{\sharp}$ is a mismatched estimate of $\theta_{*}$ ), Corollary 5 (which assume $\theta_{\sharp}= \pm \theta_{*}$ ) and also Theorem 7 (which concerns estimating $\theta_{*}$ without the knowledge of $\delta$ ). We take the loss $|\hat{\delta}-\delta|$ achieved by $\hat{\delta}_{\text {corr }}$ and the trivial estimators $\hat{\delta}_{0}\left(X_{1}^{n}\right)=0, \hat{\delta}_{1}\left(X_{1}^{n}\right)=1, \hat{\delta}_{1 / 2}\left(X_{1}^{n}\right)=\frac{1}{2}$. We plot the loss as a function of $t=\left\|\theta_{*}\right\|$ for $t \in[0,1]$ with step size 0.05 . We take $n=500, d=250, \delta=0.1$. In the left panel, we assume the estimator has access to $\theta_{\sharp}$ with $\left\|\theta_{\sharp}\right\|=1.2 \cdot\left\|\theta_{*}\right\|$. In the right panel, we assume $\theta_{\sharp}=\theta_{*}$ and therefore the model (2) is equivalent to the model (252).


Figure 5: Numerical validation of the performance of Algorithm 1 in Theorem 7. We take the minimum loss (cf. (3)) achieved by Algorithm 1 and the trivial estimator $\hat{\theta}_{0}\left(X_{1}^{3 n}\right)=0$. We plot the loss as a function of $t=\left\|\theta_{*}\right\|$ for $t \in[0,4]$ with step size 0.05 . We take $n=100, d=5, \delta=0.1$.
et al. 2019, Ren et al. 2022 in the context of the EM algorithm, in a rather restricted setting: The distribution of the samples is standard Gaussian $N\left(0, I_{d}\right)$, yet the estimator is allowed to (over)fit a two-component Gaussian mixture, with symmetric means $\pm \theta$ and a covariance matrix $\sigma^{2} \cdot I_{d}$ with any $\sigma>0$. Even in this restricted setting, the result is rather delicate and there are differences, e.g., between one- and multi-dimensional models. We finally remark that even method-of-moments based estimators (as the one we use in our paper) the analysis is also typically made for isotropic noise, e.g., [Hsu and Kakade, 2013, Wu and Yang, 2020]. Addressing these issues in the context of Markovian model or models with more sophisticated dependence structures is left as an important yet challenging future task.
2. Another direction beyond Gaussian noise is to look at noise with a heavy-tailed distribution. There has been some recent progress on this topic in high-dimensional statistics Hopkins 2020, Hopkins et al., 2020, Cherapanamjeri et al., 2020]. The estimation error rate is expected to depend on the decay rate of the tail. Additional ideas and techniques will most likely be needed in order to handle heavy tails. We leave it for future research.


[^0]:    ${ }^{3}$ The constants can be deduced from the proof, though they were not optimized. We reiterate that both $\lambda_{\theta}$ and $\lambda_{\delta}$ below are universal numerical constants. In particular, they do not depend on $\theta_{*}, \delta$ or any other problem parameters. The subscripts are merely to emphasize that they are obtained in the estimation procedure for $\theta_{*}$ and $\delta$, respectively.

[^1]:    ${ }^{4}$ In the original form in |Wainwright 2019 Proposition 2.14], on the right hand side of the inequality, $\mathbb{V}\left[X_{i}\right]$ is replaced with $\mathbb{E}\left[X_{i}^{2}\right]$ which leads to a seemingly worse concentration bound. However, applying this form of Bernstein's inequality to $Y_{i}=X_{i}-\mathbb{E}\left[X_{i}\right]$ allows us to conclude (323).

