

A Diffusion Process Perspective on Posterior Contraction Rates for Parameters

Wenlong Mou[◦] Nhat Ho[∗] Martin J. Wainwright^{◦,†,‡}
 Peter Bartlett^{◦,†} Michael I. Jordan^{◦,†}

Department of EECS[◦], Department of Statistics[†], UC Berkeley

Department of Statistics and Data Science, UT Austin[∗]

Department of EECS, MIT[‡]

August 18, 2022

Abstract

We analyze the posterior contraction rates of parameters in Bayesian models via the Langevin diffusion process, in particular by controlling moments of the stochastic process and taking limits. Analogous to the non-asymptotic analysis of statistical M-estimators and stochastic optimization algorithms, our contraction rates depend on the structure of the population log-likelihood function, and stochastic perturbation bounds between the population and sample log-likelihood functions. Convergence rates are determined by a non-linear equation that relates the population-level structure to stochastic perturbation terms, along with a term characterizing the diffusive behavior. Based on this technique, we also prove non-asymptotic versions of a Bernstein-von-Mises guarantee for the posterior. We illustrate this general theory by deriving posterior convergence rates for various concrete examples, as well as approximate posterior distributions computed using Langevin sampling procedures.

1 Introduction

Bayesian inference is one of the central pillars of statistics. In Bayesian analysis, we first endow the parameter space with a prior distribution chosen by modeling considerations, and then apply Bayes’ rule, combining the prior with the likelihood, so as to form the posterior distribution. From a statistical perspective, this posterior is of fundamental interest, and there are various questions associated with its behavior, including its consistency as the sample size goes to infinity, and from a more refined point of view, its contraction rate in various metrics.

The earliest work on posterior consistency dates back to the seminal work of Doob [12], who demonstrated that the posterior distribution is consistent for all parameters apart from a set of zero measure. Subsequent work by Freedman [16, 17] provided examples showing that this null set can be problematic for Bayesian consistency in non-parametric settings. In order to address this issue, Schwartz [44] proposed a general framework for establishing posterior consistency for both semiparametric and nonparametric models. Since then, a number of researchers have isolated conditions that are useful for studying posterior distributions [4, 56, 57].

Moving beyond posterior consistency, convergence rates for the posterior density function, along with associated parameters of models, remains an active area of research. For posterior densities, Ghosal et al. [20] gave a general testing framework for proving convergence rates for both finite and infinite dimensional models; it has been used by various researchers to analyze posterior

densities for Dirichlet and nonparametric Beta mixtures [21, 22, 42, 45]. Other work [5, 60, 59] established minimax optimal rates for regression functions in nonparametric regression models. Related problems include adaptive rates for the density in nonparametric Bayesian inference [11, 19], and posterior contraction rates of density under misspecified models [28]. Other popular general frameworks for analyzing the density functions of posterior distributions include those of Shen and Wasserman [46], and Walker et al. [58].

1.1 From frequentist to Bayesian analysis

The focus of this paper is on posterior convergence rates for parameters—namely, how for parametric Bayesian models, the posterior distribution assigns mass to certain regions of the parameter space. Our contributions can be put into perspective by considering known results for M -estimators. In the world of frequentist statistics, estimators based on maximizing empirically-defined objective functions—known as M -estimators—play a central role. In the parametric setting, a generic M -estimator takes the form

$$\hat{\theta}_n := \arg \max_{\theta \in \Theta} F_n(\theta) \quad \text{where } F_n(\theta) := \frac{1}{n} \sum_{i=1}^n f(\theta; X_i), \text{ with } X_i \stackrel{\text{i.i.d.}}{\sim} \mathbb{P} \text{ for } i = 1, \dots, n, \quad (1)$$

while the parameters θ range over some constraint set Θ , and the real-valued function f has domain $\Theta \times \mathcal{X}$. Maximum-likelihood is the archetypal example, obtained when f is the log likelihood.

There is now a rich and well-developed theory—one which exploits ideas from both optimization theory and empirical process theory—for deriving sharp non-asymptotic bounds on the difference between the estimate $\hat{\theta}_n$ and the maximizer θ^* of the population-level objective (e.g., see the books [53, 51, 55]). This theory leverages properties of the population-level objective $F(\theta) := \mathbb{E}[f(\theta, X)]$ where the expectation is taken with respect to $X \sim \mathbb{P}$. At a high level, there are two key steps in the analysis of an M -estimator: exploiting the structure of F , and linking the behavior of the empirical objective F_n to the population objective F . In the simplest setting, the population objective is strongly concave around its unique maximum θ^* . More generally, when F is differentiable, one can consider a condition of the following type

$$-\langle \nabla F(\theta), \theta - \theta^* \rangle \geq \psi(\|\theta - \theta^*\|_2), \quad (2a)$$

assumed to hold uniformly for all θ in a local neighborhood of θ^* . Here ψ is an increasing function on the positive real-line, with $\psi(t) = \frac{\mu}{2}t^2$ being the one obtained for a μ -strongly concave function. The second step is to relate the empirical and population objective, for instance by establishing a uniform bound on their gradients—say

$$\|\nabla F_n(\theta) - \nabla F(\theta)\|_2 \leq \zeta(\|\theta - \theta^*\|_2) \varepsilon_n, \quad (2b)$$

where the function ζ is again defined on the positive real line, and ε_n measures the magnitude of the noise.

When the functions F and F_n satisfy bounds of the form (2a) and (2b), it can be shown that the estimate $\hat{\theta}_n$ satisfies a bound of the form $\|\hat{\theta}_n - \theta^*\|_2 \lesssim r_n$, where $r_n > 0$ is the largest positive solution to the inequality¹

$$\psi(r) \leq \varepsilon_n \zeta(r). \quad (3)$$

¹This solution exists and is unique under mild regularity conditions on the pair (ψ, ζ) .

This framework is very convenient to use, since optimization theory and empirical process theory give us various tools for establishing the local growth condition (2a) and the stochastic perturbation bound (2b).

By using this framework with care, one can often obtain sharp results in terms of *problem dimension* d , in both the rate itself and sample size lower bound needed to achieve such rates. Moreover, the local growth condition (2a) is relatively flexible; for instance, it allows for models in which the Fisher information matrix is singular (so that the function ψ is *not* quadratic). There are many different instantiations of this general approach in past work, including various methods or establishing growth conditions and empirical process bounds [48, 37], analysis of iterative optimization algorithm [3, 15, 32, 24], as well as regularized and constrained M -estimators [31, 9].

1.2 Our contributions

Moving back to the Bayesian setup, it is natural to seek to a similarly flexible and user-friendly method for establishing finite-sample results for posterior contraction. The main contribution of this paper is do so by using the Langevin diffusion process—a stochastic differential equation that can encode the posterior distribution—as a lens of analysis.

There are natural parallels between our mode of analysis, and deterministic analyses of optimization algorithms via differential equations [49, 47]. To provide such intuition, recall the M -estimator defined by the objective function (1). Under the given conditions, its optimum θ^* can be characterized as the limiting point of an *ordinary differential equation* known as the gradient flow, and the rate (3) via the gradient flow dynamics for population and empirical loss functions, respectively. Now consider the analogous approach for studying *not* the M -estimator, but rather (in the Bayesian set-up) the posterior distribution. It is well-known [40] that under mild regularity conditions, the posterior distribution can be represented as the stationary distribution of a *stochastic differential equation* known as the Langevin diffusion. Consequently, just as information about the M -estimator can be recovered by studying the gradient flow, we can recover information about the posterior distribution by studying the Langevin diffusion. In particular, we do so by leveraging stochastic calculus so as to control the moments of this diffusion process. At a high-level, our main results involving showing that, under assumptions of the form (2), the posterior convergence rate is governed by the inequality $\psi(r) \leq \varepsilon_n \zeta(r) + \frac{d}{n}$. By comparison to inequality (3), relevant for M -estimation, we see that this inequality includes an additional $\frac{d}{n}$ term: it characterizes the diffusive behavior (with dimension d and sample size n) induced from sampling from the Gibbs measure e^{-F_n} as opposed to taking its maximum.

With this overview in place, we now summarize the different classes of contributions that are made in this paper:

Globally concave problems: We begin with the simplest setting, in which the population log-likelihood function is strongly concave in a global sense. Under certain regularity conditions,² we prove that the posterior contraction rate around the true parameter is $(d/n)^{1/2}$. Our technique allows us to specify precise non-asymptotic conditions on the sample size and other model properties under which a guarantee of this type holds. We then relax our assumption from strongly concave to (weakly) concave, and prove related guarantees. We illustrate these general results for three

²Briefly, we require the prior distribution to be sufficiently smooth and the perturbation error between the population and empirical log-likelihood function to be well-controlled.

concrete classes of models: Bayesian non-linear regression models, over-specified Bayesian location Gaussian mixture models, and Bayesian logistic regression models. Our theory reveals the influence of different modeling assumptions on the behavior of the posterior.

From global to local concavity: In order to extend the scope of our theory, we next relax the global nature of our conditions. We study posterior contraction when the population log-likelihood function F is only locally concave in a ball around θ^* , thereby allowing for multi-modality. In this setting, we find key properties that govern the convergence rate: the rate of growth of the population log-likelihood, and the deviations between the gradients of the sample and population log-likelihoods. In particular, consider a log-likelihood function such that

$$-\langle \nabla F(\theta), \theta - \theta^* \rangle \gtrsim \|\theta - \theta^*\|_2^{\alpha+1}, \quad \text{and} \quad \|\nabla F_n(\theta) - \nabla F(\theta)\|_2 \lesssim \|\theta - \theta^*\|_2^\beta \cdot \sqrt{d/n}$$

for some positive values of α and β with $\alpha > \beta$. Our theory guarantees that the posterior convergence rate of parameters is given by $O((d/n)^{\min\{\frac{1}{1+\alpha}, \frac{1}{2(\alpha-\beta)}\}})$. This result not only recovers the classical results when the Fisher information is non-singular—i.e., when $\alpha = 1$ and $\beta = 0$ —in a non-asymptotic way for a suitable range of n , but also applies to a broad class of models with singular Fisher information—i.e., for which $\alpha > 1$ and $\beta \geq 0$. The proof relies on the similar diffusion process considered in the globally concave settings, with a modified version of the potential function that exhibits the same local behavior as the empirical log-likelihood function.

Guarantees for approximate posteriors computed via Langevin algorithms: By adapting the continuous-time arguments to a discrete-time setting, we show contraction rate bounds for the output of the unadjusted Langevin algorithm. Working with the local strongly convex setting, we show that the output of Langevin algorithm satisfies contraction bounds that (up to logarithmic factors) match the optimal posterior contraction behavior. Compared to existing works, our result does not put stringent assumptions on the stepsize, allowing for faster convergence of the algorithm.

Non-asymptotic Bernstein-von-Mises (BvM) results: Our final contribution is to establish two non-asymptotic BvM results for models with non-degenerate Fisher information. For the first result, we derive a non-asymptotic upper bound on the Kullback-Leibler (KL) divergence between the posterior distribution and the limiting Gaussian distribution with mean given by maximum a posteriori (MAP) estimate, and covariance matrix by the inverse of Hessian matrix of the population log-likelihood function. This bound scales at the order $\mathcal{O}(1/n)$ in terms of the sample size n . Second, we prove non-asymptotic tail bounds that are satisfied by the posterior distribution; those bounds almost match the tail bounds that are satisfied by the limiting Gaussian law, up to high-order terms. In particular, we show that the posterior mass concentrates within an ellipsoid whose shape is determined by the Hessian matrix of the population log-likelihood at θ^* . We note that the diffusion process approach plays a central role in this proof: in particular, a key technical ingredient is an error estimate between the underlying diffusion process and an Ornstein-Uhlenbeck (OU) process, whose stationary distribution is the limiting Gaussian law in BvM theorems.

The remainder of the paper is organized as follows. In Section 2, we set up the basic framework for Bayesian models and introduce a diffusion process that admits posterior distribution as its stationary distribution. Section 3 is devoted to establishing the general results for posterior convergence rates of parameters under various assumptions on the global concavity of the population log-likelihood. We then study these convergence rates under the locally concave settings

of the population log-likelihood function in Section 4.1. Section 4.2 is devoted to non-asymptotic BvM results for models with non-degenerate Fisher information. We discuss an application of these general theories to Bayesian logistic regression and Gaussian mixture models in Section 5 and other statistical models in Appendix A. We conclude our work with a discussion in Section 6 while proofs of results in the paper are in the supplementary material [35].

Notation. In the paper, the expression $a_n \gtrsim b_n$ will be used to denote $a_n \geq cb_n$ for some positive universal constant c that does not change with n . Additionally, we write $a_n \asymp b_n$ if both $a_n \gtrsim b_n$ and $a_n \lesssim b_n$ hold. For any $n \in \mathbb{N}$, we denote $[n] = \{1, 2, \dots, n\}$. The notation \mathbb{S}^{d-1} stands for the unit sphere, namely, the set of vectors $u \in \mathbb{R}^d$ such that $\|u\|_2 = 1$. For any subset Θ of \mathbb{R}^d , $r \geq 1$, and $\varepsilon > 0$, we denote $\mathcal{N}(\varepsilon, \Theta, \|\cdot\|_r)$ the covering number of Θ under $\|\cdot\|_r$ norm, namely, the minimum number of ε -balls under $\|\cdot\|_r$ norm to cover the entire set Θ . Given a positive-definite matrix $M \succ 0$, we use $\lambda_{\max}(M)$ and $\lambda_{\min}(M)$ to denote its largest and smallest eigenvalue, respectively, and we use $\kappa(M) := \lambda_{\max}(M)/\lambda_{\min}(M)$ to denote its condition number. Finally, for any $x, y \in \mathbb{R}$, we denote $x \vee y = \max\{x, y\}$ and $x \wedge y = \min\{x, y\}$.

2 Background and problem formulation

This section is devoted to background material along with formulation of the problems studied in this paper. We first set up the problem of studying convergence rates for posterior distributions over parameters in Section 2.1, and provide background on its representation as the stationary distribution of a Langevin diffusion process in Section 2.2. Finally, we define the population likelihood function, and introduce various smoothness conditions in Section 2.3.

2.1 Posterior contraction rates for parameters

Consider a parametric family of distributions $\{P_\theta \mid \theta \in \Theta\}$. Throughout the paper, we assume that each distribution P_θ has density p_θ with respect to the Lebesgue measure. Let $X_1^n := (X_1, \dots, X_n)$ be a sequence of random variables drawn i.i.d. from P_{θ^*} , where $\theta^* \in \Theta$ is the true parameter, albeit unknown. Given a prior π over the parameter space, we define the the log-likelihood

$$F_n(\theta) := \frac{1}{n} \sum_{i=1}^n \log p_\theta(X_i), \quad \text{along with the posterior} \quad \Pi(\theta \mid X_1^n) := \frac{e^{nF_n(\theta)} \pi(\theta)}{\int_{\Theta} e^{nF_n(u)} \pi(u) du}. \quad (4)$$

As the sample size n increases, we expect that the posterior distribution will concentrate more of its mass over increasingly smaller neighborhoods of the true parameter θ^* . Posterior contraction rates allow us to study how quickly this concentration of mass takes place. In particular, for a given norm, we study the posterior mass of a ball of the form $\|\theta - \theta^*\| \leq \rho$ for a suitably chosen radius $\rho > 0$. For a given $\delta \in (0, 1)$, our goal is to prove statements of the form $\Pi(\|\theta - \theta^*\| \geq \rho(n, d, \delta) \mid X_1^n) \leq \delta$, with probability at least $1 - \delta$ over the randomly drawn data X_1^n . Our interest is in the scaling of the radius $\rho(n, d, \delta)$ as a function of sample size n , problem dimension d , and the error tolerance δ , as well as other problem-specific parameters.

2.2 From diffusion processes to the posterior distribution

The analysis of this paper relies on a well-known connection between the posterior distribution and a particular stochastic differential equation (SDE) known as the Langevin diffusion. For a

parameter $\beta > 0$, the Langevin diffusion can be written as

$$d\theta_t = -\nabla U(\theta_t)dt + \sqrt{\frac{2}{\beta}} dB_t, \quad (5)$$

where $(B_t, t \geq 0)$ is a standard d -dimensional Brownian motion [39], and $U : \mathbb{R}^d \rightarrow \mathbb{R}$ is known as the potential function. Suppose that we impose the following regularity conditions on the potential: (a) its gradient ∇U is locally Lipschitz, and (b) its gradient satisfies the inequality $\langle \nabla U(\theta), \theta \rangle \geq c_1 \|\theta\|_2 - c_2$ for any $\theta \in \mathbb{R}^d$, for some strictly positive constants c_1, c_2 . Under these conditions, by known results on general Langevin diffusions [2], the solution to the Langevin diffusion (5) exists and is unique in the strong sense. Furthermore, the density of θ_t converges in \mathbb{L}^2 to the stationary distribution with density proportional to $e^{-\beta U}$.

In the context of Bayesian inference, we can apply this argument to the potential function $U_n(\theta) := -nF_n(\theta) - \log \pi(\theta)$. Doing so will require us to verify that U_n satisfies the requisite regularity conditions. Assuming this validity, we are guaranteed that the posterior distribution $\Pi(\theta | X_1^n)$ is the stationary distribution of the SDE

$$d\theta_t = \frac{1}{2}\nabla F_n(\theta_t)dt + \frac{1}{2n}\nabla \log \pi(\theta_t)dt + \frac{1}{\sqrt{n}}dB_t, \quad (6)$$

with initial condition $\theta_0 = \theta^*$. Moreover, the density of θ_t converges in \mathbb{L}^2 to the posterior density.

It should be noted that this SDE-based representation of the posterior underlies various algorithms for drawing samples from the posterior distribution; we refer the reader to the papers [10, 13, 14] for some recent state-of-the-art results in this direction. In this paper, we exploit this SDE-based representation for statistical analysis (as opposed to efficient computation). In particular, by characterizing the behavior of the process $(\theta_t, t \geq 0)$ as a function of time, we can obtain bounds on the posterior distribution by taking limits. The following proposition guarantees the convergence of the moments based on a uniform-in-time moment upper bound and a convergence in total variation distance.

Proposition 1. *Consider a sequence of distributions $(\pi_t)_{t \geq 0}$ on \mathbb{R}^d such that $d_{\text{TV}}(\pi_t, \pi^*) \rightarrow 0$, and suppose that $\sup_{t \geq 0} \mathbb{E}_{\pi_t} [\|X\|_2^p] < +\infty$ and $\mathbb{E}_{\pi^*} [\|X\|_2^p] < +\infty$ for any even integer $p \geq 2$. We then have $\lim_{t \rightarrow +\infty} \mathbb{E}_{\pi_t} [\|\bar{X}\|_2^p] = \mathbb{E}_{\pi^*} [\|X\|_2^p]$.*

See Appendix E.1 in our supplementary material [35] for the proof of this proposition.

Given this limiting behavior, we can establish posterior contraction rates for the parameters by controlling the moments of the diffusion process $\{\theta_t\}_{t \geq 0}$. The main theoretical results of this paper are obtained by following this general roadmap.

2.3 From empirical to population likelihood

Before proceeding to our main results, let us introduce some additional definitions and conditions. A useful notion for our analysis is the population log-likelihood F . It corresponds to the limit of log-likelihood function F_n , as previously defined in equation (4), as the sample size n goes to infinity—viz.

$$F(\theta) := \mathbb{E} [\log p_\theta(X)], \quad (7)$$

where the expectation is taken with respect to $X \sim P_{\theta^*}$. Throughout the paper, we impose the following smoothness conditions on the population log-likelihood F and the log prior density $\log \pi$:

(A) There exist positive constants L_1 and L_2 such that for any $\theta_1, \theta_2 \in \mathbb{R}^d$, we have

$$\|\nabla F(\theta_1) - \nabla F(\theta_2)\|_2 \leq L_1 \|\theta_1 - \theta_2\|_2, \quad \text{and} \quad \|\nabla \log \pi(\theta_1) - \nabla \log \pi(\theta_2)\|_2 \leq L_2 \|\theta_1 - \theta_2\|_2.$$

(B) There exists a non-negative constant $B \geq 0$ such that

$$\langle \nabla \log \pi(\theta), \theta - \theta^* \rangle \leq B \|\theta - \theta^*\|_2 \quad \text{for all } \theta \in \mathbb{R}^d.$$

Although the constant B in Assumption (B) can depend on θ^* , we suppress this dependence so as to keep the notation streamlined. When the function $\log \pi$ is globally Lipschitz (so that $\|\nabla \log \pi(\theta)\|_2$ is uniformly bounded), Assumption (B) is automatically satisfied, but it only requires a one-sided control, allowing for important examples such as Gaussian prior.

The above conditions are relatively mild, and we provide a number of examples in the sequel for which they are satisfied.

3 Results under global conditions

We now turn to our first set of results, which provide bounds on posterior contraction rates under global concavity conditions on the population log-likelihood function. Results under milder local conditions are given in Section 4 to follow.

In Section 3.1, we present a result (theorem 1) that establishes the posterior convergence under strong concavity. Section 3.2 answers the same question when the population log-likelihood is only weakly concave; see the statement of theorem 2.

3.1 Posterior contraction under strong concavity

We begin with results under strong concavity conditions. For this part, the following assumptions underlie our analysis:

(S.1) There exists a scalar $\mu > 0$ such that

$$-\langle \nabla F(\theta), \theta^* - \theta \rangle \geq \mu \|\theta - \theta^*\|_2^2 \quad \text{for any } \theta \in \mathbb{R}^d.$$

(S.2) There exist non-negative functions ε_1 and ε_2 that map from $\mathbb{N} \times (0, 1]$ to \mathbb{R}_+ such that for any radius $r > 0$ and any $\delta \in (0, 1)$, we have

$$\sup_{\theta \in \mathbb{B}(\theta^*, r)} \|\nabla F_n(\theta) - \nabla F(\theta)\|_2 \leq \varepsilon_1(n, \delta)r + \varepsilon_2(n, \delta) \quad \text{with prob. at least } 1 - \delta.$$

Assumption (S.1) is a standard strong concavity condition of function F around θ^* , whereas Assumption (S.2) provides uniform control on the gradients of the population and sample log-likelihoods. It is important to note that these assumptions, along with other assumptions to follow, *do not* require the data-generating distribution P to belong to the specified parametric class. Indeed, the results throughout this paper apply to both well-specified and mis-specified models. In the latter case, the parameter θ^* is typically the KL-projection of the true model, i.e., $\theta^* \in \arg \min_{\theta \in \Theta} D_{\text{KL}}(\mathbb{P} \parallel \mathbb{P}_\theta)$.

Given the above assumptions, we are ready to state our first result regarding the posterior convergence rate of parameters for a strongly concave population log likelihood:

Theorem 1. Suppose that Assumptions **(A)**, **(B)**, **(S.1)**, and **(S.2)** hold. Then there is a universal constant c such that for any $\delta \in (0, 1)$ and any sample size n for which $\varepsilon_1(n, \delta) \leq \frac{\mu}{6}$, we have

$$\Pi\left(\|\theta - \theta^*\|_2 \geq c\sqrt{\frac{d}{n\mu}} + \frac{B}{n\mu} + \frac{\varepsilon_2(n, \delta)}{\mu} + c\sqrt{\frac{\log(1/\delta)}{n\mu}} \mid X_1^n\right) \leq \delta$$

with probability $1 - \delta$, taken with respect to the random observations X_1^n .

See appendix C.3 for the proof of theorem 1.

This result guarantees posterior convergence at the rate $(d/n)^{1/2}$ when the log likelihood is strongly concave. To be clear, such rate of posterior contraction for the parameters can be derived from the asymptotic behavior of the posterior distribution via the classical Bernstein-von-Mises theorem. However, the guarantee in theorem 1 is non-asymptotic, and provides explicit dependence of the rate on other model parameters, including B and μ , both of which might vary as a function of θ^* . At the moment, we do not know whether the dependence of these parameters is optimal. This guarantee is valid as long as the error term $\varepsilon_1(n, \delta)$ is less than an absolute constant; such a bound typically holds as long as $n \gtrsim d$. In theorem 5 to follow, we also provide near-optimal non-asymptotic contraction bounds on the posterior distribution that nearly match the exact shape of the posterior distribution.

Although our set-up is focused on simple sampling models, it should be noted that our method is sufficiently flexible so as to accommodate certain non-i.i.d. forms of sampling, along with misspecified models. After the first version was posted, Mazumdar et al. [33] used a variant of this result to study the posterior contraction rates for Thompson sampling in contextual bandits. In their problem, the data are adaptively collected instead of being i.i.d., and the empirical process bound **(S.2)** can be verified using martingale concentration inequalities.

Proof overview: As described in our motivating introduction, the proof of theorem 1 is based on analyzing the Langevin diffusion $(\theta_t)_{t \geq 0}$ from equation (6). The key idea—one which plays a key role in the proofs throughout the entire paper—is the use of a Lyapunov function Φ_t . In particular, we use Itô calculus to track the growth of Φ_t over time t . By taking $t \rightarrow +\infty$, the bounds on the Lyapunov function carry over to the stationary distribution.

In more detail, we prove theorem 1 using the Lyapunov function $\Phi_t := \frac{1}{2}e^{\frac{\mu t}{2}}\|\theta_t - \theta^*\|_2^2$ and bounding the moments of this stochastic process. Some calculation leads to the upper bound

$$e^{\frac{\mu t}{2}}\|\theta_t - \theta^*\|_2^2 \leq \frac{1}{\sqrt{n}} \int_0^t e^{\mu s/2} \langle \theta_s - \theta^*, dB_s \rangle + c \left(\frac{d}{n} + \frac{\varepsilon_2^2(n, \delta)}{\mu} + \frac{B}{n^2} \right) \cdot \frac{e^{\mu t/2}}{\mu}.$$

The last term is deterministic, and gives rise to the terms $\sqrt{\frac{d}{\mu n}} + \frac{\varepsilon_2(n, \delta)}{\mu} + \frac{B}{n\mu}$ in theorem 1. Taking expectations on both sides of the bound yields a non-asymptotic bound on the second moment of the posterior distribution. In order to provide a high probability bound, as stated in the claim, we control the martingale term by invoking the Burkholder-Davis-Gundy (BDG) inequality for continuous-time martingales; doing so produces the term $\sqrt{\frac{\log(1/\delta)}{n}}$ in our bound. The full proof is given in appendix C.3.

3.2 Posterior contraction under weak concavity

theorem 1 requires global strong concavity, which is relatively strong. In this section, we relax this assumption in two ways: we relax the growth condition locally around θ^* so as to allow for

weak concavity, and the global behavior need not coincide with this local behavior. Weakly concave log-likelihoods arise for singular problems, for which the Fisher information matrix at the true parameter θ^* is rank-degenerate. Examples of such singular problems include Bayesian non-linear regression models with certain choices of link functions [34], as well as over-specified mixture models [43], in which the fitted mixture model has more components than the true mixture distribution. The mismatch between local and global concavity conditions exists not only in such models, but also in non-singular problems such as Bayesian logistic regression. We discuss implications of these examples in the supplementary material [35]. Note that the results in this section still require the global maximum θ^* to be unique, so that the posterior is unimodal. This requirement is removed in the analysis of the next section.

Our analysis in the weakly concave setting is based on the following assumptions:

(W.1) There exists a convex, non-decreasing function $\psi : [0, +\infty) \rightarrow \mathbb{R}$ such that

$$-\langle \nabla F(\theta), \theta - \theta^* \rangle \geq \psi(\|\theta - \theta^*\|_2) \quad \text{for any } \theta \in \mathbb{R}^d.$$

Assumption **(W.1)** characterizes the weak concavity of the function F around the global maxima θ^* . This condition can hold when the log likelihood is locally strongly concave around θ^* but only weakly concave in a global sense, or it can hold when the log likelihood is weakly concave but nowhere strongly concave. An example of the former type is the logistic regression model analyzed in Section 5.1, whereas an example of the latter type is given by certain kinds of non-linear regression models, as analyzed in Appendix A.1.

Our next assumption controls the deviation between the gradients of the population and sample likelihoods, and involves a failure probability $\delta \in (0, 1)$:

(W.2) There exist a function $\varepsilon : \mathbb{N} \times (0, 1] \mapsto \mathbb{R}_+$ and a non-decreasing function $\zeta : \mathbb{R} \rightarrow \mathbb{R}$ with that $\zeta(0) \geq 0$ such that for any radius $r > 0$, we

$$\sup_{\theta \in \mathbb{B}(\theta^*, r)} \|\nabla F_n(\theta) - \nabla F(\theta)\|_2 \leq \varepsilon(n, \delta) \zeta(r) \quad \text{with prob. at least } 1 - \delta.$$

Note that the function ζ can depend on the sample size n and other model parameters; such dependence arises in our analysis of over-specified Bayesian mixture model given in section 5.2. In this main text, we suppress this dependence so as to keep the notation streamlined.

The previous conditions involved two functions, namely ψ and ζ . We let $\xi : \mathbb{R}_+ \rightarrow \mathbb{R}$ denote the inverse function of the strictly increasing function $r \mapsto r\zeta(r)$. Our third assumption imposes certain inequalities on these functions and their derivatives:

(W.3) The function $r \mapsto \psi(\xi(r))$ is convex, and ψ and ζ satisfy the differential inequalities

$$\begin{aligned} r\psi'(r)\zeta(r) &\stackrel{(i)}{\geq} r\psi(r)\zeta'(r) + \psi(r)\zeta(r), \quad \text{and} \\ r^2\psi''(r)\zeta(r) + r\psi'(r)\zeta(r) &\stackrel{(ii)}{\geq} 3\psi(r)\zeta(r) + r^2\psi(r)\zeta''(r) \quad \text{for all } r > 0. \end{aligned}$$

These differential inequalities are needed controlling the moments of the diffusion process $\{\theta_t\}_{t>0}$ in equation (6). In our discussion of concrete examples, we provide instances for which they are satisfied.

Our result involves a certain fixed point equation that depends on the parameters and functions in our assumptions. In particular, for any tolerance parameter $\delta \in (0, 1)$ and sample size n , consider the following fixed point equation in the variable $z > 0$:

$$\psi(z) = \varepsilon(n, \delta)\zeta(z)z + \frac{B}{n}z + \frac{d}{n} + \frac{\log(1/\delta)}{n}. \quad (8)$$

In order to ensure that this equation has a unique positive solution, our final assumption imposes certain condition on the growth of the functions ψ and ζ :

(W.4) The limit $\liminf_{z \rightarrow +\infty} \frac{\psi(z)}{z\zeta(z)}$ is strictly positive, and the sample size n and tolerance parameter $\delta \in (0, 1)$ are such that $\varepsilon(n, \delta) < \liminf_{z \rightarrow +\infty} \frac{\psi(z)}{z\zeta(z)}$.

With this set-up, we are now ready to state our second main result:

Theorem 2. *Suppose that Assumptions **(A)**, **(B)**, and **(W.1)**—**(W.3)** hold. Then for any given sample size n and $\delta \in (0, 1)$ such that Assumption **(W.4)** holds, equation (8) has a unique positive solution $z^*(n, \delta)$ such that*

$$\Pi\left(\|\theta - \theta^*\|_2 \geq z^*(n, \delta) \mid X_1^n\right) \leq \delta \quad \text{with probability } 1 - \delta \text{ w.r.t. } X_1^n. \quad (9)$$

See appendix C.5 for the proof of theorem 2.

A few comments are in order. First, the convergence guarantee (9) depends on the weak convexity function ψ and the perturbation function ζ through the non-linear equation (8). See the proof sketch below for the origins of this equation. Second, at least in general, it is not possible to compute an explicit form for the positive solution $z^*(n, \delta)$ to the non-linear equation (8). However, for certain forms of the function ψ and ζ , we can derive a relatively simple upper bound. For instance, given some positive parameters (α, β) such that $\alpha > \beta$, suppose that these functions are defined locally, in a interval above zero, as follows:

$$\psi(r) = r^{\alpha+1}, \quad \text{and} \quad \zeta(r) = r^\beta \quad \text{for all } r \text{ in some interval } [0, \bar{r}). \quad (10a)$$

Moreover, suppose that the perturbation function takes the form

$$\varepsilon(n, \delta) = \sqrt{\left(d + \log\left(\frac{1}{\delta}\right)\right)/n}. \quad (10b)$$

As shown in in section 5, these particular forms arise in several statistical models, including Bayesian logistic regression and over specified Bayesian Gaussian mixture models. Under these conditions, we have the following simple upper bound:

Corollary 1. *Assume that the functions ψ , ζ have the local behavior (10a), and the perturbation term $\varepsilon(n, \delta)$ has the form (10b). If, in addition, the global forms of ψ and ζ satisfy Assumption **(W.3)**, then the scalar $z^*(n, \delta)$ from theorem 2 satisfies the bound $z^*(n, \delta) \leq c \left(\frac{d+\log(1/\delta)}{n}\right)^{\frac{1}{2(\alpha-\beta)}} \vee \left(\frac{d+\log(1/\delta)}{n}\right)^{\frac{1}{\alpha+1}} + \left(\frac{B}{n}\right)^{\frac{1}{\alpha}}$.*

Note that Corollary 1 ensures that the posterior has the following contraction property

$$\Pi\left(\|\theta - \theta^*\|_2 \geq c \left(\frac{d+\log(1/\delta)}{n}\right)^{\frac{1}{2(\alpha-\beta)}} \wedge \left(\frac{B}{n}\right)^{\frac{1}{\alpha}} \mid X_1^n\right) \leq \delta \quad \text{with prob. } 1 - \delta \quad (11)$$

with respect to the training data. The posterior convergence rate scales as $(d/n)^{\frac{1}{2(\alpha-\beta)}}$ when $\alpha \geq 2\beta + 1$. On the other hand, this rate becomes $(d/n)^{\frac{1}{\alpha+1}}$ when $\alpha < 2\beta + 1$.

Proof overview: Similar to the proof of theorem 1, the proof of theorem 2 is based on tracking the behavior of a Lyapunov function along the trajectory of diffusion process (6). In doing so, we study the moments $\mathbb{E}[\|\theta_t - \theta^*\|_2^p]$ for $p \geq 2$. Unlike the strongly concave case, however, the negative term in the expression is no longer the p -th moment itself, but rather a quantity depending on the local geometry of the population log-likelihood F . More precisely, we adopt the Lyapunov function $\Phi_t := \mathbb{E}[\|\theta_t - \theta^*\|_2^{p-2} \psi(\|\theta_t - \theta^*\|_2)]$, where ψ is the function from Assumption **(W.1)**. Under the conditions on the functions ψ and ζ given in Assumption **(W.3)**, the time derivative of the p^{th} moment $\mathbb{E}[\|\theta_t - \theta^*\|_2^p]$ can then be controlled as a function of Φ_t . Since the moment converges to a finite quantity when $t \rightarrow +\infty$, its time derivative cannot converge to a positive number. Using the convexity of ψ , the bound on the Lyapunov function leads to the inequality

$$\lim_{t \rightarrow +\infty} (\mathbb{E}(\|\theta_t - \theta^*\|_2^p))^{\frac{1}{p}} \leq z_p^*,$$

where z_p^* is the unique positive solution to the equation $\psi(z) = \varepsilon(n, \delta)\zeta(z)z + \frac{B}{n}z + \frac{p+d}{n}$. In light of the above result and proposition 1, when p is of the order $\log(1/\delta)$, we obtain the posterior convergence rate (9). The full proof is given in appendix C.5.

4 Results under local conditions

In this section, we present results without the global conditions on the population log-likelihood function in section 3.1 and section 3.2. Our set-up allows the posterior distribution to be multi-modal in nature; only local growth conditions and empirical process bounds around θ^* are needed in our analysis. In section 4.1, we establish the posterior convergence rate of parameters under mild local conditions on the population and empirical log-likelihood functions, and also extend the results to Langevin algorithms. Finally, we provide non-asymptotic Bernstein-von-Mises results in section 4.2.

4.1 Non-asymptotic contraction rates under local assumptions

We begin with posterior concentration results. When the log-likelihood function satisfies suitable growth conditions and perturbation bounds in a local neighborhood of θ^* , we show posterior convergence rates conditionally on such a local ball. We further extend our results to contraction bounds of the last iterate of Langevin algorithm, again under such local conditions.

4.1.1 Conditional posterior contraction

For some local radius $r_0 > 0$, we make the following assumptions with the population and sample log-likelihood functions within the local region $\mathbb{B}(\theta^*, r_0)$:

(LWC.1) There exist $\alpha \geq 0$, $\mu > 0$ and $\varsigma \geq 0$ such that for $\theta \in \mathbb{B}(\theta^*, r_0)$, we have

$$\langle \nabla F(\theta), \theta - \theta^* \rangle \leq -\mu \|\theta - \theta^*\|_2^{\alpha+1} + \varsigma.$$

Assumption **(LWC.1)** characterizes the local growth of the function F around the global maximum θ^* . We note that in either the well-specified case ($\mathbb{P} = \mathbb{P}_{\theta^*}$), or the mis-specified case when \mathbb{P}_{θ^*} is the KL-projection of \mathbb{P} , it follows from the optimality condition that this assumption is satisfied with $\varsigma = 0$. Relaxing to values $\varsigma > 0$ allows us to accommodate mis-specified cases in which θ^* is not the exact projection, or situations in which variants of the log-likelihood are used. See appendix A.3 in the supplementary material for an application of this result with $\varsigma > 0$ to a Bayesian location model with singular densities on the density function in the Ibragimov-Khasminskii sense [26].

The parameters (α, μ) control the rate of local growth of the log-likelihood. When $\alpha = 1$, the function F is locally strongly concave around θ^* , so that one should expect posterior convergence at the rate given in theorem 1. On the other hand, when $\alpha > 1$, the log likelihood is only weakly concave in a local neighborhood; such behavior arises when the Fisher information matrix at θ^* is degenerate. Concrete instances of such degenerate models include over-specified mixture distributions, and certain types of non-linear regression models. See section 5.2 in the supplementary material for discussion of these specific examples.

Our next assumption concerns the deviation between the gradients of population and sample log-likelihood functions within the ball $\mathbb{B}(\theta^*, r_0)$.

(LWC.2) There exists $\beta \in (-1, \alpha)$ and $\varepsilon(n, \delta) > 0$ such that with probability $1 - \delta$, we have

$$\sup_{\theta \in \mathbb{B}(\theta^*, r_0)} \frac{\|\nabla F_n(\theta) - \nabla F(\theta)\|_2}{\|\theta - \theta^*\|_2^\beta} \leq \varepsilon(n, \delta).$$

Note that the assumption **(LWC.2)** requires that $\beta < \alpha$, which means that the variance of score functions cannot decay too quickly around a neighborhood of θ^* . This condition is needed to make the presentation simpler. On the other hand, when $\beta \geq \alpha$, exact recovery of θ^* is possible, and the Bayesian approach may lead to sub-optimal results. A detailed development for this setting is left for the future work.

Under Assumptions **(LWC.1)** and **(LWC.2)**, we have the following result on posterior convergence for the parameters. It involves the radius r_n given by

$$r_n := \left(\frac{\log(1/\vartheta) + d}{n\mu} + \frac{\varsigma}{\mu} \right)^{\frac{1}{\alpha+1}} + \left(\frac{2\varepsilon(n, \delta)}{\mu} \right)^{\frac{1}{\alpha-\beta}} + \left(\frac{B}{n\mu} \right)^{\frac{1}{\alpha}},$$

where $\vartheta \in (0, 1)$ is a pre-specified tolerance parameter.

Theorem 3. *Assume that Assumptions **(LWC.1)**, **(LWC.2)** and **(B)** hold. For any given $\nu \in (0, 1)$ and pair (n, δ) such that $\varepsilon(n, \delta) \leq \frac{\mu}{2} \left(\frac{r_0}{2} \right)^{\alpha-\beta}$, we have*

$$\Pi(\|\theta - \theta^*\|_2 \leq r_n \mid X_1^n) \geq (1 - \nu) \Pi(\mathbb{B}(\theta^*, r_0/2) \mid X_1^n). \quad (12)$$

See appendix C.1 for the proof.

Remark on $\Pi(\mathbb{B}(\theta^, r_0/2) \mid X_1^n)$:* As shown in the bound (12), the posterior convergence rate depends on the non-asymptotic behavior of the probability mass $\Pi(\mathbb{B}(\theta^*, r_0/2) \mid X_1^n)$ that the posterior assigns to a local ball around θ^* . In particular, given a (potentially non-sharp) non-asymptotic posterior contraction that ensures concentration within a *constant-radius ball* $\mathbb{B}(\theta^*, r_0/2)$, theorem 3 automatically improves it to a concentration result with the *optimal* radius. Moreover, in the non-identifiable case where the global maxima of the population log-likelihood function F is not unique, one can still apply theorem 3 to obtain concentration around the (finite) set of global maxima (see Corollary 7 in appendix B for more details).

Remarks on r_n : Let us consider the three different terms in r_n . First, consider idealized situation in which the empirical log-likelihood function replaced by the population one, and we ignore the contribution from the prior π . The “posterior” in this case takes the form e^{-nF} ; note that it satisfies the contraction bounds with radius $\left(\frac{\log(1/\nu)+d}{n\mu} + \frac{\varsigma}{\mu}\right)^{\frac{1}{\alpha+1}}$. The second term $\left(\frac{2\varepsilon(n,\delta)}{\mu}\right)^{\frac{1}{\alpha-\beta}}$ characterizes the effect of using empirical data instead of population-level functions. This term coincides with the non-asymptotic rates for the maximal likelihood estimator in a local neighborhood of θ^* . Finally, the last term $\left(\frac{B}{n\mu}\right)^{\frac{1}{\alpha}}$ characterizes the effect of the prior density π . Under mild regularity conditions on π , this term is of higher order compared to the first term $\left(\frac{\log(1/\nu)+d}{n\mu}\right)^{\frac{1}{\alpha+1}}$, as in the case of theorem 1.

Remarks on the proof: Let us provide some high-level comments on the proof. The argument involves constructing a Lyapunov function similar to that used theorem 2. However, since the condition **(LWC.1)** holds only in a small ball $\mathbb{B}(\theta^*, r_0)$, the leading term $\langle \nabla F(\theta), \theta - \theta^* \rangle$ in Itô’s formula *cannot* be uniformly upper bounded by a negative function of the distance $\|\theta - \theta^*\|_2$. In order to overcome this issue, we first study a modified version of the posterior distribution, and then transform the result back to the posterior distribution itself. In particular, we construct a probability density function $\tilde{\Pi}$ over \mathbb{R}^d such that:

- Within the local ball $\mathbb{B}(\theta^*, r_0/2)$, the shape of the function $\tilde{\Pi}$ exactly matches that of the true posterior $\Pi(\cdot \mid X_1^n)$, up to a multiplicative constant.
- Outside the larger ball $\mathbb{B}(\theta^*, r_0)$, the function $\tilde{\Pi}$ behaves as a Gaussian density—in particular, we have $\tilde{\Pi}(\theta) \propto \exp\left(-\frac{nL_1}{2} \|\theta - \theta^*\|_2^2\right)$.
- In the annulus between the two balls, we interpolate between the two regimes so as to ensure that $\log \tilde{\Pi}$ is smooth.

By applying the analysis in theorem 2 to the modified density $\tilde{\Pi}$, one can show that this modified density concentrates within $\mathbb{B}(\theta^*, r_n)$ with high probability. Since the shapes of $\tilde{\Pi}$ and $\Pi(\cdot \mid X_1^n)$ are exactly the same inside $\mathbb{B}(\theta^*, r_0/2)$, we can prove that conditionally in the ball $\mathbb{B}(\theta^*, r_0/2)$, the posterior $\Pi(\cdot \mid X_1^n)$ also contracts around θ^* with the correct radius.

4.1.2 Contraction of approximate posterior via Langevin algorithm

We have analyzed posterior contraction properties using the Langevin diffusion process (5), upon which most posterior sampling algorithms are built. It is therefore natural to extend our techniques to the discretized Langevin process, and obtain contraction rates for the approximate posterior

distribution computed via Langevin algorithm. In this section, we analyze the following forward Euler discretization of Langevin diffusion, a widely-used algorithm for computation of posterior [13, 13, 10].

$$\theta_{k+1} = \theta_k + \eta \nabla F_n(\theta_k) + \sqrt{\frac{2\eta}{n}} W_k, \quad \text{for } k = 0, 1, \dots \quad (13)$$

where $(W_k)_{k=0,1,\dots}$ are i.i.d. standard Gaussian random vectors.

As Euler discretization can be unstable when applied to functions with growth at infinity faster than quadratic (see [41]), we focus on the case where $\alpha = 1$ and $\beta = 0$. The general case, for which a more stable discretization scheme may be employed, is an important direction of future research. We also restrict our attention to algorithms with *local initialization*, satisfying $\|\theta_0 - \theta^*\| \leq r_0/2$. Finally, we require the stepsize η and the sample size n to satisfy the following conditions:

$$\eta \leq \frac{\mu}{3L^2} \quad \text{and} \quad \frac{3\varepsilon(n, \delta)}{\mu} + \frac{3B}{n\mu} + \sqrt{\frac{3cd}{n\mu} \log^3 \frac{T}{\delta}} \leq \frac{r_0}{2}. \quad (14)$$

Under such setup, we have the following theorem:

Theorem 4. *Under Assumptions (LWC.1) and (LWC.2) with $\alpha = 1$ and $\beta = 0$, for sample size n and stepsize satisfying Eq (14), given a local initialization satisfying, we have the following with probability $1 - \delta$ with respect to both the data and the randomness in the algorithm:*

$$\|\theta_T - \theta^*\|_2 \leq e^{-\frac{T\mu\eta}{12\log(1/\delta)}} \|\theta_0 - \theta^*\|_2 + c \left\{ \frac{\varepsilon(n, \delta)}{\mu} + \frac{B}{\mu n} + \log(1/\delta) \cdot \sqrt{\frac{d + \log(1/\delta)}{\mu n}} \right\}. \quad (15)$$

See appendix C.4 for the proof of this theorem. A few remarks are in order.

The contraction rate for Langevin algorithm consists of four terms: the first term depends on the initial distance $\|\theta_0 - \theta^*\|_2$, and is exponentially decaying with the number of iterations T . By taking number of iterations $T \geq \frac{c}{\mu\eta} \log^2 \left(\frac{r_0^2 \mu n}{\delta} \right)$, the first term in the bound (15) becomes dominated by other terms. The rest three terms in Eq (15) matches the optimal posterior contraction rates in theorem 1, up to extra logarithmic factors in $1/\delta$. The sample size requirement in Eq (14) is essentially the sample size needed for the bound to be smaller than a constant r_0 . Notably, unlike existing literature on contraction analysis for Langevin algorithm [33, 18] the stepsize requirement in Eq (14) does not depend on the sample size n or the problem dimension d . Indeed, we only require it to be smaller than a stability threshold $\frac{\mu}{3L^2}$. This makes it possible for Langevin algorithms to use larger stepsize and achieve faster convergence, while still preserving good posterior contraction properties. Such distinction is due to the proof technique: instead of bounding the error between the distribution of Langevin algorithm iterates and the true posterior, we directly analyze the dynamics itself following the same approach as we analyze posterior contraction.

4.2 Non-asymptotic Bernstein-von-Mises results

In this section, we develop non-asymptotic Bernstein-von-Mises results using the diffusion process (6). Under mild assumptions on the population-level and empirical-level landscapes, we establish the KL divergence between the posterior distribution and the limiting Gaussian distribution based on the posterior convergence rates of the parameters.

In order to obtain the non-asymptotic Bernstein-von-Mises results, we first need the following assumptions on the second order derivatives with respect to the parameters (or equivalently Hessian matrices) of the empirical and population log-likelihoods:

(BvM.1) There exists $A > 0$ such that the population log-likelihood function F satisfies the one-point Lipschitz condition:

$$\forall \theta \in \mathbb{R}^d, \quad \|\nabla^2 F(\theta) - \nabla^2 F(\theta^*)\|_{\text{op}} \leq A \|\theta - \theta^*\|_2.$$

(BvM.2) For any $\delta > 0$, there exist non-negative functions $\varepsilon_1^{(2)}$ and $\varepsilon_2^{(2)}$ with domain $\mathbb{N} \times (0, 1]$ such that

$$\sup_{\theta \in \mathbb{B}(\theta^*, r)} \|\nabla^2 F_n(\theta) - \nabla^2 F(\theta)\|_{\text{op}} \leq \varepsilon_1^{(2)}(n, \delta)r + \varepsilon_2^{(2)}(n, \delta),$$

for any radius $r > 0$ with probability at least $1 - \delta$.

The first condition **(BvM.1)** is a standard smoothness condition needed to prove quantitative results about asymptotic normality (e.g., the paper [38]), and satisfied by many models such as exponential family models, location density models, as well as their mixtures and hierarchical composition. The second condition **(BvM.2)** is an empirical process condition on the Hessian matrix $\nabla^2 F_n$. This condition can usually be verified using suitable concentration bounds for each θ , as well as smoothness conditions on $\nabla^2 F_n$ used in controlling metric entropies. Both assumptions are naturally needed: the limiting Gaussian law $\mathcal{N}(\widehat{\theta}^{(n)}, (nH^*)^{-1})$, which depends on the population-level Hessian at the point θ^* . The shape of posterior distribution, on the other hand, depends on the sample-level Hessian $\nabla^2 F_n$ in a local neighborhood of θ^* . These two conditions are needed to relate the shape of the sample-level posterior with the matrix H^* . As before, we note that these assumptions do not require the model to be well-specified, and our non-asymptotic Bernstein-von-Mises theorems applies to the mis-specified case, where θ^* is the KL-projection of the model to this parametric class.

Consider the MAP estimate $\widehat{\theta}^{(n)} := \arg \max_{\theta \in \mathbb{R}^d} (F_n(\theta) + \frac{1}{n} \log \pi(\theta))$. Then, we have the following upper bound on the difference between the posterior distribution of the parameters and the Gaussian distribution with mean $\widehat{\theta}^{(n)}$ and covariance matrix $(nH^*)^{-1}$, where $H^* := -\nabla^2 F(\theta^*)$.

Proposition 2. *Under Assumptions **(BvM.1)**, **(BvM.2)** and **(B)**, suppose that $H^* \succ 0$, and that $\|\widehat{\theta}^{(n)} - \theta^*\|_2 \leq \sigma \sqrt{\frac{d}{n}}$ and $\mathbb{E}_{\Pi}(\|\theta - \theta^*\|_2^4 | X_1^n)^{1/4} \leq \sigma \sqrt{\frac{d}{n}}$ with prob. $1 - \delta$. Then there exists a constant c such that the KL divergence $D_{KL}(\Pi(\cdot | X_1^n) \| \mathcal{N}(\widehat{\theta}^{(n)}, (nH^*)^{-1}))$ is at most*

$$c \cdot \frac{1}{\lambda_{\min}(H^*)} \left(\frac{A^2 d^2 \sigma^4}{n} + \frac{\varepsilon_1^{(2)}(n, \delta)^2 d^2 \sigma^4}{n} + \sigma^2 \left(\varepsilon_2^{(2)}(n, \delta)^2 + \frac{L_2^2}{n^2} \right) d \right) \quad \text{with prob. at least } 1 - 2\delta.$$

See appendix C.6 for the proof of this claim.

A few remarks are in order. First, assuming that the problem-dependent constants (A, σ, L_2) are of constant order, and that the deviation bound scales as $\varepsilon_2^{(2)}(n, \delta) = O(1/\sqrt{n})$, proposition 2 shows that the KL divergence between the posterior distribution and the Gaussian limit is of order $O(1/n)$; second, the non-asymptotic behavior of posterior distribution depends on the Hessian matrix $H^* = -\nabla^2 F(\theta^*)$. In the well-specified case where the data points X_1^n are i.i.d. samples from the

distribution \mathbb{P}_{θ^*} , the standard Fisher-information identity $H^* = \mathbb{E}_{\theta^*} [\nabla \log p_{\theta^*}(X) \nabla \log p_{\theta^*}(X)^\top]$ holds true, and the Bayesian credible set is asymptotically the same as the confidence set in the frequentist sense. On the other hand, in the mis-specified models where $\theta^* = \arg \min_{\theta \in \Theta} D_{\text{KL}}(\mathbb{P} \parallel \mathbb{P}_\theta)$, the limiting Gaussian law is $\mathcal{N}(\widehat{\theta}^{(n)}, (nH^*)^{-1})$, depending on the Hessian matrix but not the covariance of the log-likelihood. This result coincides with the asymptotic Bernstein-von-Mises theorem for mis-specified parametric models [29], providing a non-asymptotic characterization. Using Pinsker's inequality and Talagrand's T_2 -inequality [50], the KL divergence bound can also be transformed into bounds in term of total variation and Wasserstein-2 distances, yielding a non-asymptotic $O(1/\sqrt{n})$ rate of convergence.

We can also use the diffusion process approach to derive more fine-grained concentration bounds for the posterior distribution, with behavior matching the limiting Gaussian law. Doing so requires the following stronger version of the posterior contraction condition:

$$\left(\mathbb{E}_\Pi \left[\|\theta - \theta^*\|_2^{2p} \mid X_1^n \right] \right)^{1/p} \leq \frac{\sigma^2 pd}{n}, \quad \text{for all } p > 0 \text{ with probability at least } 1 - \delta. \quad (16)$$

In addition, we define the function

$$\mathcal{H}_n(t, \delta) := (A + \varepsilon_1^{(2)}(n, \delta))^2 \cdot \frac{\sigma^4 d^2 t^2}{n^2} + \frac{\sigma d}{n} \left(\varepsilon_2^{(2)}(n, \delta)^2 + \frac{L_2^2}{n^2} + (A + \varepsilon_1^{(2)}(n, \delta))^2 \frac{\sigma d}{n} \right),$$

which plays the role of a higher-order term. Equipped with this notation, we have:

Theorem 5. *Suppose that conditions **(BvM.1)** and **(BvM.2)** are in force, the Hessian H^* is strictly positive definite, and the high-probability posterior contraction condition (16) holds. Then for any $\delta \in (0, 1)$, uniformly over all $\omega \in (0, 1)$ and $t > 0$, we have*

$$\Pi \left(\left\| \theta - \widehat{\theta}^{(n)} \right\|_{H^*}^2 \geq (1 + \omega) \frac{d}{n} + c \frac{1 + \log \kappa(H^*)}{\omega} \left(\frac{t}{n} + \mathcal{H}_n(t, \delta) \right) \mid X_1^n \right) \leq e^{-t}, \quad (17)$$

with probability at least $1 - \delta$.

See appendix C.2 for the proof of the theorem.

A few remarks are in order. Note that the limiting Gaussian density $\gamma_n = \mathcal{N}(0, (nH^*)^{-1})$ satisfies a tail bound of the form $\gamma_n \left(\|\theta - \widehat{\theta}^{(n)}\|_{H^*}^2 \geq \frac{d}{n} + \frac{t}{n} \right) \leq e^{-t/2}$ for any $t > 0$. Unless the posterior is actually Gaussian in finite samples, it cannot satisfy this bound exactly. However, theorem 5 provides a bound with near-matching behavior: note that the leading-order term scales $\frac{d}{n}$, matching the asymptotics with a pre-factor $1 + \omega$ that can be made arbitrarily close to 1 (at the expense of the other term). The $\frac{t}{n}$ dependency on the tail probability comes with a mild $\log \kappa(H^*)$ factor due to technical reasons. The bound also contains a high-order term $\mathcal{H}_n(t, \delta)$, which scales as $O(n^{-2})$. It is also worth noticing that the terms in theorem 5 depend on the tail probability $\nu = e^{-t}$ only logarithmically, allowing for very small value of ν . We can therefore use equation (17) to construct non-asymptotic credible sets of ellipsoid shape, adapted to the geometry of local Hessian matrix H^* .

Proof outline: The proofs of both proposition 2 and theorem 5 rely on a first-order approximation of the gradient ∇F_n . In particular, the diffusion process (6) can be written in the form $d\theta_t = -\frac{1}{2}H^*(\theta_t - \widehat{\theta}^{(n)})dt + \frac{1}{2}e_n(\theta_t)dt + \frac{1}{2n}\log \pi(\theta_t)dt + \frac{1}{\sqrt{n}}dB_t$, where we have defined the linearization

error $e_n(\theta) := \nabla F_n(\theta) + H^*(\theta - \theta^*)$. Under the smoothness assumption **(BvM.1)** and the empirical process bound **(BvM.2)**, one can show that $\|e_n(\theta)\|_2 \leq \|\theta - \theta^*\|_2 \cdot O(\sqrt{d/n})$ with high probability. When this error term is ignored, the diffusion process is an Ornstein-Uhlenbeck process whose stationary distribution is $\mathcal{N}(\widehat{\theta}^{(n)}, (nH^*)^{-1})$. Therefore, given the non-asymptotic bounds on the error $e_n(\theta)$ stated above, we can provide a non-asymptotic characterization of the distance between the stationary distribution and the limiting Gaussian law. In order to prove proposition 2, we use the Gaussian log-Sobolev inequality [23] to control the KL divergence, whereas proving theorem 5 is based on using Itô calculus to study the growth of a Lyapunov function defined using the metric induced by H^* . Full proofs for the two results are given in appendix C.6 and appendix C.2, respectively.

5 Some illustrative examples

Having developed some general theory, we now use it to derive some concrete results for two examples of interest in statistical analysis: Bayesian logistic regression and Gaussian mixture models. Due to space constraints, we defer the treatment of additional examples to appendix A.

5.1 Bayesian logistic regression

Logistic regression is a classical way of modeling the relationship between a binary response variable $Y \in \{-1, +1\}$ and a vector $X \in \mathbb{R}^d$ of explanatory variables (e.g., see the book [34]). In the logistic regression model, the pair (X, Y) are related by the conditional distribution

$$\mathbb{P}(Y = 1 \mid X, \theta) = \frac{e^{\langle X, \theta \rangle}}{1 + e^{\langle X, \theta \rangle}}, \quad \text{where } \theta \in \mathbb{R}^d \text{ is a parameter vector.} \quad (18)$$

Suppose that we observe a collection $Z_1^n = \{Z_i\}_{i=1}^n$ of n i.i.d paired samples $Z_i = (X_i, Y_i)$, each generated in the following way. First, the covariate vector X_i is drawn from a standard Gaussian distribution $N(0, I_d)$, and then the binary response Y_i is drawn according to the conditional distribution $\mathbb{P}(\cdot \mid X_i, \theta^*)$ from equation (18), where $\theta^* \in \mathbb{R}^d$ is a fixed but unknown value of the parameter vector. Given these assumptions, the sample log-likelihood function of the samples Z_1^n takes the form $F_n^R(\theta) := \frac{1}{n} \sum_{i=1}^n \{\log \mathbb{P}(Y_i \mid X_i, \theta) + \log \phi(X_i)\}$, where ϕ denotes the density of a standard normal vector. Combining this log likelihood with a given prior π over θ yields the posterior distribution in the usual way. We assume that the prior function π satisfies Assumptions **(A)** and **(B)**, and recall the constant B defined in the latter assumption.

With this set-up, the following result establishes the posterior convergence rate of θ around θ^* , conditionally on the observations Z_1^n .

Corollary 2. *For any $\delta \in (0, 1)$, given $\frac{n}{\log n} \geq c'd\log(\frac{1}{\delta})$ i.i.d. samples from the Bayesian logistic regression model (18), we have $\Pi\left(\|\theta - \theta^*\|_2 \geq c\left\{\sqrt{\frac{d}{n}} + \sqrt{\frac{\log(1/\delta)}{n}} + \frac{B}{n}\right\} \mid Z_1^n\right) \leq \delta$ with probability $1 - \delta$ over the data Z_1^n .*

See appendix D.1 for the proof of this claim.

A few comments are in order. First, the result of Corollary 2 shows that for Bayesian logistic regression model (18), the posterior convergence rate for the parameter is of the order $(d/n)^{1/2}$. Furthermore, this result also gives a concrete dependence of the rate on B characterizing the degree to which the prior is concentrated away from the true parameter. Second, by taking the

sample size in the function F_n^R to infinity, we find that the population log-likelihood is given by $F^R(\theta) := \mathbb{E}_{(X,Y)} [-\log(1 + e^{-Y\langle X, \theta \rangle}) + \log \phi(X)]$. Here ϕ denotes the standard normal density in \mathbb{R}^d , and the outer expectation in the above display is taken with respect to X and $Y | X$ from the logistic model (18).

Let us sketch how theorem 2 can be applied so as to prove this corollary. The first step in our proof, as given in appendix D.1, is to show that there are universal constants c, c_1, c_2 such that

$$-\langle \nabla F^R(\theta), \theta - \theta^* \rangle \geq c_1 \begin{cases} \|\theta - \theta^*\|_2^2, & \text{for all } \|\theta - \theta^*\|_2 \leq 1 \\ \|\theta - \theta^*\|_2, & \text{otherwise} \end{cases}, \quad \text{and} \quad (19a)$$

$$\sup_{\theta \in \mathbb{R}^d} \|\nabla F_n^R(\theta) - \nabla F^R(\theta)\|_2 \leq c_2 \left(\sqrt{\frac{d}{n}} + \sqrt{\frac{\log(1/\delta)}{n}} + \frac{\log(1/\delta)}{n} \right), \quad (19b)$$

for any $r > 0$ with probability $1 - \delta$ as long as $\frac{n}{\log n} \geq cd \log(1/\delta)$. Using these results, we show that Assumptions **(W.1)** and **(W.2)** hold with

$$\psi(r) = c_1 \begin{cases} r^2 & \text{for all } r \in (0, 1), \text{ and} \\ r & \text{otherwise} \end{cases}, \quad \text{and} \quad \zeta(r) = c_2 \quad \text{for all } r > 0. \quad (20)$$

We can check that the functions ψ and ζ satisfy the conditions in Assumptions **(W.3)** and **(W.4)**. Therefore, applying theorem 2 to these functions yields the posterior contraction rate claimed in corollary 2. See appendix D.1 for the details.

5.2 Over-specified Bayesian Gaussian mixture models

Gaussian mixtures are widely used for modeling heterogenous datasets; clusters in the data are naturally associated with different mixture components [30]. In fitting such models, the true number of components is generally unknown, and several approaches have been proposed to deal with this challenge. One of the most popular methods is to deliberately include a large number of components, leading to what are known as overspecified Gaussian mixture models [43]. While the behavior of posterior densities in such mixture models is relatively well-understood [21], the behavior of the posterior in terms of its parametric components is not as well understand. When the covariance matrices are known and the parameter space is bounded, the location parameters have been shown to have posterior convergence rates of the order $n^{-1/4}$ in the Wasserstein-2 metric [36]. However, neither the dependence on dimension d nor on the true number of components have been established.

In this section, we consider the behavior of overspecified Gaussian mixture models in a particular setting, and provide convergence rates for the parameters with precise dependence on the dimension d , and without requiring any boundedness assumption. In order to model the simplest form of over-specification, suppose that we fit a Bayesian location mixture model to a collection of i.i.d. samples $X_1^n = (X_1, \dots, X_n)$ drawn from a Gaussian Gaussian distribution $\mathcal{N}(\theta^*, I_d)$. (For concreteness, we set $\theta^* = 0$.) We study the behavior of the Bayesian Gaussian mixture model

$$\theta \sim \pi(\cdot), \quad V_i \in \{-1, 1\} \stackrel{\text{i.i.d.}}{\sim} \text{Cat}(1/2, 1/2), \quad X_i | V_i, \theta \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(V_i \theta, I_d), \quad (21)$$

where $\text{Cat}(1/2, 1/2)$ stands for the categorical distribution with parameters $(1/2, 1/2)$. We assume that the prior π satisfies the smoothness Assumptions **(A)** and **(B)**; one example is a Gaussian

distribution (over the location parameter θ . Our goal in this section is to characterize the posterior contraction rate of the location parameter θ around θ^* .

In order to do so, we first define the sample log-likelihood function F_n^G given data X_1^n . It has the form $F_n^G(\theta) := \frac{1}{n} \sum_{i=1}^n \log \left(\frac{1}{2} \phi(X_i; -\theta, I_d) + \frac{1}{2} \phi(X_i; \theta, I_d) \right)$, where $x \mapsto \phi(x; \theta, I_d) = (2\pi)^{-d/2} e^{-\|x-\theta\|_2^2/2}$ denotes the density of multivariate Gaussian distribution $\mathcal{N}(\theta, \sigma^2 I_d)$. Similarly, the population log-likelihood function is given by $F^G(\theta) := \mathbb{E}_X \left[\log \left(\frac{1}{2} \phi(X; -\theta, I_d) + \frac{1}{2} \phi(X; \theta, I_d) \right) \right]$, where the outer expectation in the above display is taken with respect to $X \sim \mathcal{N}(\theta^*, I_d)$.

In appendix D.3, we prove that there is a universal constant $c_1 > 0$ such that

$$-\langle \nabla F_n^G(\theta), \theta - \theta^* \rangle \geq \begin{cases} c_1 \|\theta - \theta^*\|_2^4, & \text{for all } \|\theta - \theta^*\|_2 \leq \sqrt{2} \\ 4c_1 \left(\|\theta - \theta^*\|_2^2 - 1 \right), & \text{otherwise} \end{cases}, \quad (22a)$$

and moreover, there are universal constants (c, c_2) such that for any $\delta \in (0, 1)$, given a sample size $n \geq cd \log(1/\delta)$, we have

$$\sup_{\theta \in \mathbb{B}(\theta^*, r)} \|\nabla F_n^G(\theta) - \nabla F^G(\theta)\|_2 \leq c_2 \left(r + \frac{1}{\sqrt{n}} \right) \left(\sqrt{\frac{d}{n}} + \sqrt{\frac{\log(\log(n/\delta))}{n}} \right) \quad \text{with prob. } 1 - \delta. \quad (22b)$$

Given the above results, the functions ψ and ζ in Assumptions **(W.1)** and **(W.2)** take the form

$$\psi(r) = \begin{cases} c_1 r^4, & \text{for all } 0 < r \leq \sqrt{2} \\ 4c_1 (r^2 - 1), & \text{otherwise} \end{cases}, \quad \text{and} \quad \zeta(r) = r + \frac{1}{\sqrt{n}} \quad \text{for all } r > 0. \quad (23)$$

These functions satisfy the conditions of Assumptions **(W.3)** and **(W.4)**. Therefore, it leads to the following result regarding the posterior contraction rate of parameters under overspecified Bayesian location Gaussian mixtures (21):

Corollary 3. *Given the overspecified Bayesian location Gaussian mixture model (21), there are universal constants c, c' such that given any $\delta \in (0, 1)$ and a sample size $n \geq c'd \log(1/\delta)$, we have $\Pi \left(\|\theta - \theta^*\|_2 \geq c \left(\frac{d}{n} + \frac{\log(\log(n/\delta))}{n} \right)^{1/4} + \left(\frac{B}{n} \right)^{1/3} \mid X_1^n \right) \leq \delta$ with probability $1 - \delta$ over the data X_1^n . Here, B is the non-negative constant in Assumption **(B)**.*

See appendix D.3 for the proof of Corollary 3.

The dependence on n in the posterior contraction rate of θ in Corollary 3 is consistent with the previous result with location parameters in the overspecified Bayesian location Gaussian mixtures [8, 27, 36]. Novel aspects of the bound include $d^{1/4}$ -dependence on dimension d and the $B^{1/3}$ -dependence on the smoothness parameter B . Finally, our result does not require the boundedness of the parameter space, in contrast to past work [8, 27, 36].

6 Discussion

In this paper, we described an approach for analyzing the posterior contraction rates of parameters based on the diffusion processes. Our theory depends on two important features: the convex-analytic structure of the population log-likelihood function F and stochastic perturbation bounds between the gradient of F and the gradient of its sample counterpart F_n . We studied the problem

under both global and local assumptions on the log-likelihood. For log-likelihoods that are globally strongly concave around the true parameter θ^* , we established posterior convergence rates for parameter estimation of the order $(d/n)^{1/2}$, valid under appropriate smoothness conditions on the prior distribution π and mild conditions on the perturbation error between ∇F_n and ∇F . On the other hand, when the population log-likelihood function is globally weakly concave, our analysis shows that convergence rates are more delicate: they depend on an interaction between the degree of weak convexity, and the stochastic error bounds. In this setting, we proved that the posterior convergence rate of parameter is upper bounded by the unique positive solution of a non-linear equation determined by the previous interplay. We also provided results under weaker local conditions on the growth of log-likelihood, and the empirical process defined by the likelihood gradients over some neighborhood $\mathbb{B}(\theta^*, r_0)$ of the global maximum. Finally, we demonstrated the utility of the diffusion process approach by deriving non-asymptotic forms of Bernstein-von Mises results for models with non-degenerate Fisher information.

Let us now discuss a few directions that arise naturally from our work. First, in the weakly convex setting, though we have established non-asymptotic posterior contraction bounds, the current results do not provide information on the shape of the asymptotic posterior distribution. For example, when F is locally strongly concave around θ^* , it is well-known from the Bernstein-von Mises theorem that the posterior distribution of parameter converges to a multivariate normal distribution centered at the maximum likelihood estimation (MLE) with the covariance matrix is given by $1/(nI(\theta^*))$ (e.g., see the book [52]), where $I(\theta^*)$ denotes the Fisher information matrix at θ^* . When the F is only weakly concave, then the Fisher information matrix $I(\theta^*)$ is degenerate, so that the posterior distribution can no longer be approximated by a multivariate Gaussian distribution. It is interesting to consider how the diffusion approach might provide insight into the posterior behavior in this setting.

Second, the contraction rates given in this paper can give information about the over-specification of the latent variable models, thereby having potential applications for model selection. As a concrete example, for the symmetric two-component Gaussian mixture model example discussed in Section 5.2, the posterior distribution concentrates around $\theta^* = 0$ at a rate $O((d/n)^{1/4})$ in the over-specified case. On the other hand, for a non-degenerate mixture with symmetric modes at θ^* and $-\theta^*$ (with $\theta^* \neq 0$), it concentrates at the usual rate $O((d/n)^{1/2})$. Consequently, the degree of dispersion in the posterior serves as an indicator of over-specification. Furthermore, since our results are non-asymptotic, they also give guidance on how this procedure could be performed with finite sample size n . Finally, whereas this paper focused on posterior contraction for parametric models, we suspect that the diffusion process approach used here might also be fruitfully applied to non-parametric models.

Appendices

In our appendices, we provide the details of our general theory applied to various examples, along with all details for the proofs of our general results. Appendix A covers the additional examples mentioned in the main text that serve to illustrate the diffusion process approach to posterior contraction. Additional general theory for the posterior convergence rate of parameters when the population log-likelihood function is non-convex is in appendix B. The proofs of theorems and propositions are given in appendix C. The proofs of our main corollaries are in appendix D, while proofs of the remaining results in the paper are in appendix E.

A Additional examples

This appendix continues the discussion of Section 5, providing consequence of our theorems for some additional examples. Our discussion includes: Bayesian non-linear regression models with polynomial link functions, general Bayesian Gaussian mixture models, and one-dimensional location models with a singular density function. These examples feature different aspects of the diffusion process approach, covering local and global conditions, as well as strongly and weakly concave log-likelihood functions.

A.1 Bayesian non-linear regression models

We now turn to analyzing a certain type of non-linear regression model, known as a single index model, but in a simplified form in which the link function is known. These models are a natural generalization of linear regression, and have applications in econometrics, biostatistics, and computational imaging [7].

Given a collection of d -dimensional covariate vectors $\{X_i\}_{i=1}^n$, suppose that we observe responses of the form

$$Y_i = g(X_i^\top \theta^*) + \epsilon_i, \quad \text{for } i = 1, \dots, n, \quad (24)$$

for a known link function of the form g . In the analysis given here, we study the family $g(t) = t^p$ for some $p \geq 2$ given. The special case $p = 2$ leads to an idealized instance of the problem of noisy phase retrieval.

We assume moreover that the additive noise variables $\{\epsilon_i\}_{i=1}^n$ are i.i.d. and standard Gaussian, whereas the covariate vectors $\{X_i\}_{i=1}^n$ are also i.i.d., independent of the noise, and standard multivariate Gaussian. Conditioning on X_i and θ , we have

$$Y_i | X_i, \theta \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}\left(g\left(X_i^\top \theta\right), 1\right). \quad (25)$$

Moreover, we endow the parameter space \mathbb{R}^d with a prior function π that satisfies the Assumptions **(A)** and **(B)**. As in the previous example, we first study the structure of the sample log-likelihood function around the true parameter θ^* , and then we establish a uniform perturbation bound between the population and sample log-likelihood functions.

Given the Bayesian single index model (25), the sample log-likelihood function F_n^I of the samples $Z_1^n = \{Z_i\}_{i=1}^n$ admits the following form

$$F_n^I(\theta) := \frac{1}{n} \left(\sum_{i=1}^n -\frac{(Y_i - g(X_i^\top \theta))^2}{2} + \log \phi(X_i) \right), \quad (26)$$

where ϕ is the standard normal density function of X_1, \dots, X_n . Hence, the population log-likelihood function F^I has the following form

$$F^I(\theta) := \mathbb{E}_{(X, Y)} \left[-\frac{(Y - g(X^\top \theta))^2}{2} + \log \phi(X) \right], \quad (27)$$

where the outer expectation in the above display is taken with respect to $X \sim \mathcal{N}(0, I_d)$ and $Y|X = x \sim \mathcal{N}(g(x^\top \theta^*), 1)$.

The interesting case to consider is $\theta^* = 0$, in which case, for any link function of the function $g(t) = t^p$ with $p \geq 2$, the function F^I is weakly concave around θ^* . Given our choices of g and θ^* , the population log-likelihood function takes on the closed-form expression

$$F^I(\theta) = \frac{1 + (2p-1)!! \|\theta - \theta^*\|_2^{2p}}{2} \quad \text{for all } \theta \in \mathbb{R}^d.$$

Furthermore, in appendix D.2, we prove that there is a universal constant $c_1 > 0$ such that

$$\langle \nabla F^I(\theta), \theta^* - \theta \rangle \geq c_1 \|\theta - \theta^*\|_2^{2p} \quad \text{for all } \theta \in \mathbb{R}^d, \quad (28a)$$

and there are universal constants (c, c_2) such that for any $r > 0$ and $\delta \in (0, 1)$, as long as $n \geq c(d \log(d/\delta))^{2p}$, we have

$$\sup_{\theta \in \mathbb{B}(\theta^*, r)} \|\nabla F_n^I(\theta) - \nabla F^I(\theta)\|_2 \leq c_2 (r^{p-1} + r^{2p-1}) \sqrt{\frac{d + \log(1/\delta)}{n}}, \quad (28b)$$

with probability at least $1 - \delta$. Therefore, the functions ψ and ζ in Assumptions **(W.1)** and **(W.2)** take the specific forms

$$\psi(r) = c_1 r^{2p}, \quad \text{and} \quad \zeta(r) = r^{p-1} + r^{2p-1}, \quad (29)$$

for all $r > 0$. Simple algebra shows that these functions satisfy Assumptions **(W.3)** and **(W.4)**. With this set-up, applying theorem 2 yields:

Corollary 4. *Consider the Bayesian single index model (24) with true parameter $\theta^* = 0$ and link function $g(r) = r^p$ for some $p \geq 2$. Then there are universal constants c, c' such that for any $\delta \in (0, 1)$, given a sample size $n \geq c'(d + \log(d/\delta))^{2p}$, we have*

$$\Pi \left(\|\theta - \theta^*\|_2 \geq c \left(\frac{d + \log(1/\delta) + B}{n} \right)^{1/(2p)} \mid Z_1^n \right) \leq \delta$$

with probability $1 - \delta$ over the data Z_1^n . Here, B is the non-negative constant in Assumption **(B)**.

See appendix D.2 for the proof of corollary 4.

It is worth noting that the proof of corollary 4 actually leads to the following stronger uniform perturbation bound:

$$\begin{aligned} \sup_{\theta \in \mathbb{B}(\theta^*, r)} \|\nabla F_n^I(\theta) - \nabla F^I(\theta)\|_2 &\leq c r^{p-1} \left(\sqrt{\frac{d + \log \frac{1}{\delta}}{n}} + \frac{1}{n^{3/2}} \left(d + \log \frac{n}{\delta} \right)^{p+1} \right) \\ &\quad + r^{2p-1} \left(\sqrt{\frac{d + \log(1/\delta)}{n}} + \frac{1}{n^{3/2}} \left(d + \log \frac{n}{\delta} \right)^{2p+1} \right), \end{aligned}$$

valid for each $r > 0$ with probability $1 - \delta$. The condition $n \geq c(d + \log(d/\delta))^{2p}$ is required to guarantee that the RHS of the above display is upper bounded by the RHS of equation (28b); this bound permits us to apply theorem 2 to establish the posterior convergence rate of parameter under the Bayesian single index models.

A.2 Bayesian Gaussian mixture models with multiple centers

We now consider a class of well-separated location Gaussian mixture models. In particular, we consider i.i.d. data X_1^n from the mixture distribution $\frac{1}{K} \sum_{j=1}^K \mathcal{N}(u_j^*, I_d)$ for some $K \geq 2$ and $u_1^*, u_2^*, \dots, u_K^*$ are distinct parameters. We use the following Bayesian mixture model to fit the data:

$$u_1, u_2, \dots, u_K \stackrel{\text{i.i.d.}}{\sim} \pi(u), \quad c_i \in [K] \stackrel{\text{i.i.d.}}{\sim} \text{Cat}(1/K, 1/K, \dots, 1/K), \quad X_i | c_i, \theta \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\theta_{c_i}, I_d). \quad (30)$$

This model is well-specified in the sense that the true model belongs to the class of models being considered, and the number of components in the model equals the true number of (distinct) components. The model is identifiable only up to permutation of the labels, as there are $M = K!$ many global minima of the population-level log-likelihood that parametrizes the same probability distribution. Given a permutation function $\sigma : [K] \rightarrow [K]$, we denote $\theta_\sigma^* := (u_{\sigma(1)}^*, u_{\sigma(2)}^*, \dots, u_{\sigma(K)}^*)$. As an application of the results from Section 4, we establish the posterior contraction rate of parameters as well as the Bernstein-von-Mises phenomena for this model around θ_σ^* , for each permutation function σ .

To state the corollary, for each permutation function $\sigma : [K] \rightarrow [K]$, we define the Fisher information matrix:

$$H_\sigma^* := \mathbb{E}_{\theta^*} \left[\nabla_\theta \log p(X; \theta_\sigma^*) \cdot \nabla_\theta \log p(X; \theta_\sigma^*)^\top \right].$$

By symmetry, the matrices H_σ^* are permutations of each other. In particular, for σ_1, σ_2 , we have:

$$H_{\sigma_1}^* = (I_d \otimes P_\sigma) H_{\sigma_2}^* (I_d \otimes P_\sigma)^\top,$$

where P_σ is a $K \times K$ permutation matrix defined by the permutation σ , and \otimes denotes the Kronecker product.

When $(u_j^*)_{j \in [K]}$, it is known (see, e.g. [25]) that the Fisher information is positive definite. We denote its smallest eigenvalue $\mu := \lambda_{\min}(H_{Id}^*) > 0$. For notational convenience, we also introduce the notation

$$\sigma_X := \sup_{u \in \mathbb{S}^{d-1}, j \in [K]} \|\langle X - u_j^*, u \rangle\|_{\psi_2},$$

where $\|Y\|_{\psi_2}$ denotes the Orlicz ψ_2 norm for a random variable Y . We can see that $1 \leq \sigma_X \leq c \left(1 + \sup_{j, \ell \in [K]} \|u_j^* - u_\ell^*\|_2 \right) < +\infty$.

The log-likelihood of this mixture model can have multiple global maxima due to the symmetry. We use $\hat{\theta}_\sigma^{(n)}$ to denote the one corresponding to θ_σ^* :

$$\hat{\theta}_\sigma^{(n)} := \arg \min_{\theta \in \arg \max F_n} \|\theta - \theta_\sigma^*\|_2.$$

Corollary 5. *Under the mixture of K location Gaussian distributions (30), there exists $n_{\min} > 0$ depending on θ^* , such that for any $\delta, \vartheta, \omega \in (0, 1)$, given sample size $n \geq n_{\min} (\log \delta^{-1} + \log \vartheta^{-1})^2$,*

we have the following concentration bounds on the posterior distribution:

$$\begin{aligned} \Pi \left[\text{There exists permutation } \sigma : \|\theta - \theta_\sigma^*\|_2 \leq \frac{cK\sigma_X}{\mu} \sqrt{\frac{Kd \log(Kd) + \log \delta^{-1}}{n}} \right. \\ \left. + c \sqrt{\frac{\log \vartheta^{-1}}{\mu n}} \mid X_1^n \right] \geq 1 - \vartheta, \quad \text{and} \end{aligned} \quad (31a)$$

$$\begin{aligned} \Pi \left[\text{There exists permutation } \sigma : \left\| \theta - \widehat{\theta}_\sigma^{(n)} \right\|_{H_\sigma^*} \geq (1 + \omega) \frac{d}{n} \right. \\ \left. + c \frac{1 + \log \kappa(H_\sigma^*)}{\omega} \left(\frac{\log \vartheta^{-1}}{n} + \frac{a'(\log \vartheta^{-1} + \log \delta^{-1})^2}{n^2} \right) \mid X_1^n \right] \geq 1 - \vartheta, \end{aligned} \quad (31b)$$

where $a' > 0$ is a constant depending on K, d and θ^* , while $c > 0$ is a universal constant.

See appendix D.4 for the proof of this corollary.

A.3 Bayesian location families with singularities

Consider a one-dimensional location family $\{f(\cdot - \theta)\}_{\theta \in \mathbb{R}}$, where f is a density function with respect to the Lebesgue measure on \mathbb{R} . It was observed by Ibragimov and Khasminskii [26] that the discontinuities and singularities in the density function f reveal more information about the location parameter, leading to rates of parameters even faster than the usual $n^{-1/2}$ -rates of regular models. In this section, we show how theorem 3 can be used to obtain the optimal posterior concentration rates of parameters in such models. For our example, we only consider the singularity of the second type (see Chapter 6.1 in the book [26]), while a similar argument can be applied to the case of discontinuities in the densities. We leave an extension of our framework to the first and third type of singularities for the future work.

Without loss of generality, we assume the singularity happens at 0. Following [26], given $\beta \in (0, 1/2)$, we assume the following representation:

$$f(x) = h(x) \exp \left(\ell(x) |x|^\beta \right) \quad \forall x \in \mathbb{R}. \quad (32)$$

The function h is assumed to be everywhere differentiable, with the following quantitative assumption:

$$c_1 := \sup_{x \in \mathbb{R}} \left| \frac{\partial}{\partial x} \log h(x) \right| < +\infty. \quad (33)$$

The function ℓ is smooth except for possible discontinuity at 0, with additional assumption that $|\ell(0^-)| + |\ell(0^+)| > 0$. The fluctuations in ℓ can be absorbed into the pre-factor $h(\cdot)$. In such case, without loss of generality, we can assume that:

$$\ell(x) = a \mathbf{1}_{\{x < 0\}} + b \mathbf{1}_{\{x > 0\}}.$$

By the translation invariance of location families, we assume $\theta^* = 0$ without loss of generality. Though the empirical process condition **(LWC.2)** does not generally hold for the gradient of log-likelihood of singular location families, the analysis can still be done via the smoothing technique.

In particular, let $F_n^S := \frac{1}{n} \sum_{i=1}^n \log f(X_i - \theta)$, we define:

$$\tilde{F}_n^S(\theta) := \frac{1}{2a_n} \int_{\theta-a_n}^{\theta+a_n} F_n^S(z) dz, \quad \tilde{F}^S(\theta) := \mathbb{E}_{\theta^*} [\tilde{F}_n^S(\theta)]. \quad (34)$$

We can then define the smoothed posterior distribution:

$$\tilde{\Pi}(\theta) = \tilde{Z}^{-1} \pi(\theta) \cdot \exp(-n\tilde{F}_n^S(\theta)), \quad \text{where } \tilde{Z} := \int \pi(\theta) \cdot \exp(-n\tilde{F}_n^S(\theta)) d\theta.$$

For simplicity of presentation, we assume that the prior distribution π is supported on the interval $[-1, 1]$, and satisfies the smoothness condition **(B)** on its support. Then, we prove in Appendix D.5 that there exist constants $q_1, q_2, q_3, r_0 > 0$ that depend on the density function f but independent of n and a_n , such that:

$$\begin{aligned} -\langle \theta, \nabla \tilde{F}^S(\theta) \rangle &\geq q_1 |\theta|^{1+2\beta} - q_2 a_n^{1+2\beta}, \quad \text{for } \theta \in (-r_0/2, r_0/2), \\ \sup_{\theta \in [-1, 1]} |\nabla \tilde{F}^S(\theta) - \nabla \tilde{F}_n^S(\theta)| &\leq q_3 \left(a_n^{\beta-1/2} \sqrt{\frac{\log n/\delta}{n}} + a_n^{\beta-1} \frac{\log n/\delta}{n} \right), \end{aligned}$$

with probability $1 - \delta$. Based on these results, an application of Theorem 3 with local concavity assumption on the population log-likelihood function leads to the following result on the posterior convergence rates of parameters under model with density function (32).

Corollary 6. *Given a Bayesian location model with density specified in equation (32) with $\beta \in (0, 1/2)$, under above setup, there exists a pair of constants (q_0, q') depending on the function f , such that given any $\delta \in (0, 1)$, for $n \geq q_0 \log^{1+\frac{1}{2\beta}} \delta^{-1}$, we have the following bound with probability $1 - \delta$:*

$$\forall \vartheta \in (0, 1), \quad \tilde{\Pi} \left(|\theta| > q' n^{-\frac{1}{1+2\beta}} \left(\log^{\frac{1}{2\beta}} \frac{n}{\delta} + \log^{\frac{1}{1+2\beta}} \vartheta^{-1} \right) \mid X_1^n \right) \leq \vartheta.$$

See appendix D.5 for the proof of this corollary.

B Multiple global maxima setting

When the population log-likelihood is non-convex, there may be multiple global maxima. Nonetheless, given some conditions on the form of non-convexity, it is possible to establish a contraction result that allows for multiple global maxima as a consequence of theorem 3.

More concretely, suppose that there is a finite collection $\mathcal{M}^* = \{\theta_1^*, \theta_2^*, \dots, \theta_M^*\}$ of global maxima of the population log-likelihood function F , and that the following conditions are in force:

- (C.1) There exists $r_0 > 0$, such that for any $j \in [M]$, the conditions **(LWC.1)** and **(LWC.2)** hold true with parameters $(\mu_j, \alpha_j, \beta_j, \varepsilon_j(n, \delta), \varsigma_j)$.
- (C.2) The gap Δ_0 for the log-likelihood outside the radius r_0 is strictly positive, i.e.,

$$\Delta_0 := \inf \left\{ F(\theta_1^*) - F(\theta) : \theta \in \bigcap_{j=1}^M \mathbb{B}^c(\theta_j^*, r_0) \right\} > 0. \quad (36a)$$

Furthermore, for any $R > 0$ and $\delta > 0$, there exists $\bar{\varepsilon}_{n,\delta}(R) > 0$ with $\lim_{n \rightarrow +\infty} \bar{\varepsilon}_{n,\delta}(R) \rightarrow 0$, such that with probability $1 - \delta$, we have:

$$\sup_{\theta \in \mathbb{B}(0, R)} |F(\theta) - F_n(\theta)| \leq \bar{\varepsilon}_{n,\delta}(R). \quad (36b)$$

In addition, the prior density function satisfies the lower bound $\max_{j \in [M]} \pi(\theta_j^*) \geq \pi_0$.

(C.3) For any $\delta > 0$, there exists a radius $R_\delta > 0$, such that with probability $1 - \delta$, the following bound holds true:

$$\forall \theta \notin \mathbb{B}(0, R_0), \quad \langle \nabla F_n(\theta), \theta \rangle \leq 0. \quad (37a)$$

Additionally, there exists $c_\pi > 0$ such that

$$\forall \theta \notin \mathbb{B}(0, R_\delta)^c, \quad -\langle \nabla \log \pi(\theta), \theta \rangle \geq c_\pi \|\theta\|_2^2. \quad (37b)$$

The last condition requires the prior density π to have sub-Gaussian tail. This condition is satisfied, for example, by any Gaussian density. We use this condition to simplify the arguments of unbounded parameter space, with quite weak assumption (37a) required on the log-likelihood function itself. Under stronger conditions on the log-likelihood (for example, when the right-hand-side of equation (37a) is replaced by a quadratic function), this requirement on the prior density can be removed.

An unconditional posterior concentration result can then be established under this setup. Recall that the quantities Δ_0 and $\bar{\varepsilon}_{n,\delta}$ are defined in equations (36a) and (36b), and for each $j \in [M]$, the parameters $(\mu_j, \alpha_j, \beta_j, \varepsilon_j(n, \delta), \varsigma_j)$ are the parameters in Assumptions **(W.1)** and **(W.2)**.

Corollary 7. *Under Assumptions **(C.1)**, **(C.2)**, and **(C.3)**, with probability $1 - 3\delta$, denote $\tilde{r}_0 := r_0 \wedge \sqrt{\frac{\Delta_0}{8L_1}}$, for sample size satisfying the inequality*

$$\begin{aligned} n &\geq \frac{4}{\Delta_0} \left(\log(1/\vartheta) + \log \pi_0^{-1} + d \log \frac{d}{\tilde{r}_0} \right), \quad \text{and} \\ \bar{\varepsilon}_{n,\delta} \left(c(R_\delta \log(1/\vartheta) + \sqrt{(d + \log(1/\vartheta))/c_\pi}) \right) &< \frac{\Delta_0}{4}, \end{aligned}$$

we have the posterior concentration result for any $\vartheta > 0$:

$$\Pi \left(\bigcup_{j=1}^M \mathbb{B}(\theta_j^*, r_n^{(j)}) \mid X_1^n \right) \geq 1 - \vartheta,$$

where the radius $r_n^{(j)}$ is defined as:

$$r_n^{(j)} := \left(\frac{\log(1/\vartheta) + d}{n\mu_j} + \frac{\varsigma_j}{\mu_j} \right)^{\frac{1}{\alpha_j+1}} + \left(\frac{2\varepsilon_j(n, \delta)}{\mu_j} \right)^{\frac{1}{\alpha_j-\beta_j}} + \left(\frac{B}{n\mu_j} \right)^{\frac{1}{\alpha_j}}.$$

The proof of corollary 7 is in Appendix E.2.

The contraction radius $r_n^{(j)}$ around each θ_j^* corresponds to the contraction radius r_n in theorem 3 with corresponding local conditions on the log-likelihood. In many examples such as well-specified Bayesian mixture models, the global maxima of the population log-likelihood are permutations of each other (see our example in Appendix A.2), and the contraction radii for each center θ_j^* are the same. In general, however, the global maxima of the population-level log-likelihood landscape can have different geometric behaviors, leading to different contraction radii around different centers.

C Proofs

This section is devoted to the proofs of our main theorems.

C.1 Proof of theorem 3

We begin with some notation and definitions that are central to the analysis. First, we define the annulus $\mathbb{A}(\theta^*, r_0) := \mathbb{B}(\theta^*, r_0) \setminus \mathbb{B}(\theta^*, r_0/2)$. Second, we define a pair of functions with domain \mathbb{R}^d as follows:

$$\Psi(\theta) := \begin{cases} F(\theta), & \theta \in \mathbb{B}(\theta^*, r_0/2), \\ \left(2 - \frac{2\|\theta - \theta^*\|_2}{r_0}\right) F\left(\frac{r_0(\theta - \theta^*)}{2\|\theta - \theta^*\|_2}\right) + \left(\frac{2\|\theta - \theta^*\|_2}{r_0} - 1\right) (F(\theta^*) - \frac{L_1 r_0^2}{2}), & \theta \in \mathbb{A}(\theta^*, r_0), \\ F(\theta^*) - \frac{L_1}{2} \|\theta - \theta^*\|_2^2, & \theta \in \mathbb{B}^c(\theta^*, r_0), \end{cases} \quad (38a)$$

and

$$\zeta_n(\theta) := \begin{cases} (F_n(\theta) - F(\theta)) - (F_n(\theta^*) - F(\theta^*)) & \theta \in \mathbb{B}(\theta^*, r_0/2), \\ 2\frac{r_0 - \|\theta - \theta^*\|_2}{r_0} \zeta_n\left(\theta^* + \frac{r_0}{2} \cdot \frac{\theta - \theta^*}{\|\theta - \theta^*\|_2}\right) & \theta \in \mathbb{A}(\theta^*, r_0), \\ 0 & \theta \in \mathbb{B}^c(\theta^*, r_0). \end{cases} \quad (38b)$$

A few comments to provide intuition are in order. Inside the ball $\mathbb{B}(\theta^*, r_0/2)$, the function Ψ is the population log-likelihood, whereas the function ζ_n specifies a “noise” term that can be controlled using empirical process methods. On the other hand, outside of the ball $\mathbb{B}(\theta^*, r_0)$, the function Ψ corresponds to a quadratic upper bound on the population log-likelihood F , whereas the function ζ_n is identically zero. In the annulus region between the two balls, we interpolate linearly between the two behaviors.

It can be verified that both Ψ and ζ_n are almost everywhere continuously differentiable and locally Lipschitz functions. Moreover, a direct computation yields

$$\langle \nabla \Psi(\theta), \theta - \theta^* \rangle = \begin{cases} \langle \nabla F(\theta), \theta - \theta^* \rangle, & \theta \in \mathbb{B}(\theta^*, r_0/2), \\ \frac{2\|\theta - \theta^*\|_2}{r_0} \left(F(\theta^*) - \frac{L_1}{2} r_0^2 - F\left(\frac{r_0(\theta - \theta^*)}{2\|\theta - \theta^*\|_2}\right)\right), & \theta \in \mathbb{A}(\theta^*, r_0), \\ -L_1 \|\theta - \theta^*\|_2^2, & \theta \in \mathbb{B}^c(\theta^*, r_0). \end{cases}$$

By Assumption **(LWC.1)**, we have the following inequalities:

$$\langle \nabla \Psi(\theta), \theta - \theta^* \rangle \leq \begin{cases} -\mu \|\theta - \theta^*\|_2^{\alpha+1} + \varsigma, & \theta \in \mathbb{B}(\theta^*, r_0/2), \\ -\frac{3L_1 r_0}{8} \|\theta - \theta^*\|_2, & \theta \in \mathbb{A}(\theta^*, r_0), \end{cases}$$

Based on the above bounds, we define the following function:

$$\psi(r) := \begin{cases} \mu r^{\alpha+1}, & r \in [0, r_0/2], \\ 2(r_0 - r)\mu r_0^\alpha + L_1 r_0(2r - r_0), & r \in (r_0/2, r_0], \\ L_1 r^2, & r > r_0. \end{cases}$$

Since $\alpha \geq 1$, it is clear that ψ is a convex function, and we have:

$$\langle \nabla \Psi(\theta), \theta - \theta^* \rangle \leq -\psi(\|\theta - \theta^*\|_2) + \varsigma \quad \forall \theta \in \mathbb{R}^d.$$

For the function ζ_n , we have

$$|\langle \nabla \zeta_n(\theta), \theta - \theta^* \rangle| \leq \begin{cases} \|\nabla F_n(\theta)\|_2 \cdot \|\theta - \theta^*\|_2, & \theta \in \mathbb{B}(\theta^*, r_0/2), \\ \frac{2\|\theta - \theta^*\|_2}{r_0} \left| \zeta_n \left(\theta^* + \frac{r_0}{2} \cdot \frac{\theta - \theta^*}{\|\theta - \theta^*\|_2} \right) \right|, & \theta \in \mathbb{A}(\theta^*, r_0), \\ 0, & \theta \in \mathbb{B}^c(\theta^*, r_0). \end{cases}$$

Note that for $\theta \in \mathbb{A}(\theta^*, r_0)$, conditionally on the event that the inequality in Assumption **(LWC.2)** holds, we have the bound

$$\begin{aligned} \frac{2\|\theta - \theta^*\|_2}{r_0} \left| \zeta_n \left(\theta^* + \frac{r_0}{2} \cdot \frac{\theta - \theta^*}{\|\theta - \theta^*\|_2} \right) \right| \\ \leq 2 \int_0^1 \|(\nabla F_n - \nabla F_n)(\gamma\theta + (1-\gamma)\theta^*)\|_2 \cdot \|\theta - \theta^*\|_2 d\gamma \\ \leq 2\varepsilon(n, \delta) \|\theta - \theta^*\|_2^{\beta+1}. \end{aligned}$$

Therefore, on the event that Assumption **(LWC.2)** holds, we have

$$|\langle \nabla \zeta_n(\theta), \theta - \theta^* \rangle| \leq 2\varepsilon(n, \delta) \|\theta - \theta^*\|_2^{\beta+1} \mathbf{1}_{\{\theta \in \mathbb{B}(\theta^*, r_0)\}}. \quad (39)$$

Now we consider the distribution $\tilde{\Pi}_n$ given by

$$\tilde{\Pi}_n(\theta) := \tilde{Z}_n^{-1} \exp(n\Psi(\theta) + n\zeta_n(\theta) + \log \pi(\theta)), \quad \text{for all } \theta \in \mathbb{R}^d,$$

where \tilde{Z}_n is a normalizing constant. Intuitively, the density function $\tilde{\Pi}_n$ is a “localized” version of the posterior distribution: the distribution $\tilde{\Pi}_n$ inherits the local behavior of the posterior Π itself, while behavior as Gaussian outside this local neighborhood. This allows us to capture the effect of local geometry of the log-likelihood function, and apply an argument similar to the proof of theorem 2.

Within the ball $\mathbb{B}(\theta^*, r_0/2)$, we have

$$\frac{\tilde{\Pi}_n(\theta)}{\Pi(\theta|X_1^n)} = \frac{Z_n}{\tilde{Z}_n} \exp(-n(F_n(\theta^*) - F(\theta^*))),$$

which is a fixed quantity independent of θ . So for $r_n < r_0/2$, we obtain

$$\begin{aligned} \Pi(\mathbb{B}(\theta^*, r_0/2)|X_1^n)^{-1} \Pi(\mathbb{B}(\theta^*, r_n)|X_1^n) &= \tilde{\Pi}_n(\mathbb{B}(\theta^*, r_0/2))^{-1} \tilde{\Pi}_n(\mathbb{B}(\theta^*, r_n)) \\ &\geq \tilde{\Pi}_n(\mathbb{B}(\theta^*, r_n)). \end{aligned}$$

Conditionally on X_1^n , the distribution $\tilde{\Pi}_n$ can be seen as the stationary distribution for the following Itô diffusion process:

$$d\tilde{\Theta}_t^{(n)} = \nabla \left(\Psi(\tilde{\Theta}_t^{(n)}) + \zeta_n(\tilde{\Theta}_t^{(n)}) + \frac{1}{n} \log \pi \right) dt + \sqrt{\frac{2}{n}} dB_t, \quad \tilde{\Theta}_0^{(n)} = \theta^*. \quad (40)$$

On the other hand, for $t \geq 0$, we define the Lyapunov function Φ as:

$$\Phi_t := \mathbb{E} \left[\left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-2} \psi \left(\left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2 \right) \right].$$

For $q \in (0, p-1]$, we define $g_q(z) := z^{\frac{p-2}{q}} \psi(z^{\frac{1}{q}})$ for all z . Now, we claim that the one-dimensional function g_q is strictly increasing and convex. Furthermore, we have

$$\begin{aligned} \mathbb{E} \left\| \tilde{\Theta}_T^{(n)} - \theta^* \right\|_2^p &\leq -p \int_0^T \Phi_t dt + \frac{pB}{n} \int_0^T g_{p-1}^{-1}(\Phi_t) dt + 2\varepsilon(n, \delta) \mu^{-\frac{p+\beta-1}{p+\alpha-1}} \int_0^T \Phi_t^{\frac{p+\beta-1}{p+\alpha-1}} dt \\ &\quad + \left(p\varsigma + \frac{p(p+d-1)}{n} \right) \int_0^T g_{p-2}^{-1}(\Phi_t) dt, \end{aligned} \quad (41)$$

when $\alpha > \beta$.

Taking the above claim as given for the moment, let us now complete the proof of the theorem. When $\alpha > \beta$, we define the function ϕ_1 as follows:

$$\phi_1(h) := -h + \frac{B}{n} g_{p-1}^{-1}(h) + 2\varepsilon(n, \delta) \mu^{-\frac{p+\beta-1}{p+\alpha-1}} h^{\frac{p+\beta-1}{p+\alpha-1}} + \left(\varsigma + \frac{p+d-1}{n} \right) g_{p-2}^{-1}(h).$$

By lemma 2, since both $\lim_{t \rightarrow +\infty} \mathbb{E} \left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^p$ and $\lim_{t \rightarrow +\infty} \Phi_t$ exist, we have

$$\lim_{T \rightarrow +\infty} \Phi_T \leq \inf \{ h > 0 : \forall h' > h, \phi_1(h') < 0 \}.$$

Note that ϕ_1 is a concave function and the equation $\phi_1(h) = 0$ only admits two solutions. One of them is $h = 0$ and the other one is the RHS of the above bound of $\lim_{T \rightarrow +\infty} \Phi_T$. Therefore, if $\frac{B}{n\mu} \leq (r_0/2)^\alpha$, $\frac{p+d}{n\mu} \leq (r_0/2)^{\alpha+1}$ and $\frac{2\varepsilon(n, \delta)}{\mu} \leq (r_0/2)^{\alpha-\beta}$, we have

$$\lim_{T \rightarrow +\infty} \Phi_T \leq \Phi_* := \mu^{-\frac{p-1}{\alpha}} \left(\frac{B}{n} \right)^{\frac{p-1+\alpha}{\alpha}} \vee \mu^{-\frac{p-2}{\alpha+1}} \left(\frac{p+d}{n} + \varsigma \right)^{\frac{p-1+\alpha}{\alpha+1}} \vee (2\varepsilon(n, \delta))^{\frac{p+\alpha-1}{\alpha-\beta}} \mu^{-\frac{p+\beta-1}{\alpha-\beta}}.$$

An application of Jensen's inequality shows that

$$\mathbb{E}_{\tilde{\Pi}_n} \left(\|\theta - \theta^*\|_2^{p-1} \right) = \lim_{T \rightarrow +\infty} \mathbb{E} \left\| \tilde{\Theta}_T^{(n)} - \theta^* \right\|_2^{p-1} \leq g^{-1} \left(\lim_{T \rightarrow +\infty} \Phi_T \right).$$

Putting the above results together, for $\Phi_* \leq \mu(r_0/2)^{p-1+\alpha}$ we obtain that

$$\left(\mathbb{E}_{\tilde{\Pi}_n} \|\theta - \theta^*\|_2^{p-1} \right)^{\frac{1}{p-1}} \leq \left(\frac{\Phi_*}{\mu} \right)^{\frac{1}{p+\alpha-1}} \leq \left(\frac{B}{n\mu} \right)^{\frac{1}{\alpha}} \vee \left(\frac{p+d}{n\mu} + \varsigma \right)^{\frac{1}{\alpha+1}} \vee \left(\frac{2\varepsilon(n, \delta)}{\mu} \right)^{\frac{1}{\alpha-\beta}}.$$

Hence, we obtain the conclusion of the theorem when $\alpha > \beta$.

C.1.1 Proof of claim (41)

By Itô's formula, for any $p \geq 2$, we have

$$\mathbb{E} \left\| \tilde{\Theta}_T^{(n)} - \theta^* \right\|_2^p \leq pI_1 + \frac{p}{n} I_2 + pI_3 + \frac{p(p+d-1)}{n} I_4, \quad (42)$$

where

$$\begin{aligned}
I_1 &:= \mathbb{E} \int_0^T \langle \nabla \Psi(\tilde{\Theta}_t^{(n)}), \tilde{\Theta}_t^{(n)} - \theta^* \rangle \left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-2} dt, \\
I_2 &:= \mathbb{E} \int_0^T \langle \nabla \log \pi, \tilde{\Theta}_t^{(n)} - \theta^* \rangle \left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-2} dt, \\
I_3 &:= \mathbb{E} \int_0^T \left| \langle \nabla \zeta_n(\tilde{\Theta}_t^{(n)}), \tilde{\Theta}_t^{(n)} - \theta^* \rangle \right| \cdot \left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-2} dt, \quad \text{and} \\
I_4 &:= \mathbb{E} \int_0^T \left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-2} dt.
\end{aligned}$$

Beginning with the first term I_1 , by using the properties of the function Ψ , we have the bound

$$\begin{aligned}
I_1 &\leq -\mathbb{E} \int_0^T \psi \left(\left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2 \right) \left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-2} dt + \varsigma \cdot \mathbb{E} \int_0^T \left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-2} dt \\
&= - \int_0^T \Phi_t dt + \varsigma \int_0^T \mathbb{E} \left[\left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-2} \right] dt.
\end{aligned} \tag{43a}$$

Turning to the second term I_2 , applying Assumption **(B)** yields the upper bound

$$I_2 \leq B \mathbb{E} \int_0^T \left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-1} dt.$$

Applying Jensen's inequality then leads to

$$\begin{aligned}
g_{p-1} \left(\mathbb{E} \left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-1} \right) &\leq \mathbb{E} g_{p-1} \left(\left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-1} \right) \\
&= \mathbb{E} \left(\left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-2} \psi \left(\left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2 \right) \right) = \Phi_t.
\end{aligned}$$

Collecting the above results, we find that

$$I_2 \leq B_2 \int_0^T g_{p-1}^{-1}(\Phi_t) dt. \tag{43b}$$

For the third term I_3 , from the bound (39) we have

$$I_3 \leq 2\varepsilon(n, \delta) \cdot \mathbb{E} \int_0^T \left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p+\beta-1} \mathbf{1}_{\{\tilde{\Theta}_t^{(n)} \in \mathbb{B}(\theta^*, r_0)\}} dt.$$

Since $\alpha > \beta$, invoking Jensen's inequality leads to

$$\begin{aligned}
\mathbb{E} \left(\left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{\beta+p-1} \mathbf{1}_{\{\tilde{\Theta}_t^{(n)} \in \mathbb{B}(\theta^*, r_0)\}} \right) &\leq \left(\mathbb{E} \left(\left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{\alpha+p-1} \mathbf{1}_{\{\tilde{\Theta}_t^{(n)} \in \mathbb{B}(\theta^*, r_0)\}} \right) \right)^{\frac{\beta+p-1}{\alpha+p-1}} \\
&\leq \mu^{-\frac{p+\beta-1}{p+\alpha-1}} \mathbb{E} \left(\psi \left(\left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2 \right) \left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-2} \right)^{\frac{p+\beta-1}{p+\alpha-1}} \\
&= \mu^{-\frac{p+\beta-1}{p+\alpha-1}} \Phi_t^{\frac{p+\beta-1}{p+\alpha-1}}.
\end{aligned}$$

Consequently, the term I_3 is upper bounded as

$$I_3 \leq 2\varepsilon(n, \delta) \mu^{-\frac{p+\beta-1}{p+\alpha-1}} \int_0^T \Phi_t^{\frac{p+\beta-1}{p+\alpha-1}} dt. \quad (43c)$$

For the fourth term I_4 , invoking Jensen's inequality yields

$$\begin{aligned} g_{p-2} \left(\mathbb{E} \left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-2} \right) &\leq \mathbb{E} g_{p-2} \left(\left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-2} \right) \\ &= \mathbb{E} \left(\left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2^{p-2} \psi \left(\left\| \tilde{\Theta}_t^{(n)} - \theta^* \right\|_2 \right) \right) = \Phi_t. \end{aligned}$$

The above inequality shows that

$$I_4 \leq \int_0^T g_{p-2}^{-1}(\Phi_t) dt. \quad (43d)$$

Collecting the bounds for I_1 — I_4 given in equations (43a)–(43d) respectively, we find that $\mathbb{E} \left\| \tilde{\Theta}_T^{(n)} - \theta^* \right\|_2^p$ is at most

$$\begin{aligned} &- p \int_0^T \Phi_t dt + \frac{pB}{n} \int_0^T g_{p-1}^{-1}(\Phi_t) dt \\ &+ 2\varepsilon(n, \delta) \mu^{-\frac{p+\beta-1}{p+\alpha-1}} \int_0^T \Phi_t^{\frac{p+\beta-1}{p+\alpha-1}} dt + \left(p\varsigma + \frac{p(p+d-1)}{n} \right) \int_0^T g_{p-2}^{-1}(\Phi_t) dt. \end{aligned}$$

Thus, we have established the claim (41).

C.1.2 Structure of the function g_q

For $q \in (0, p-1]$, we define $g_q(z) := z^{\frac{p-2}{q}} \psi(z^{\frac{1}{q}})$ for all z . Since ψ is strictly increasing and $p \geq 2$, we can check that g_q is strictly increasing. By taking the derivative of g_q , we have

$$\frac{dg_q(z)}{dz} = \frac{p-2}{q} z^{\frac{p-q-1}{q}} \frac{\psi(z^{\frac{1}{q}})}{z^{\frac{1}{q}}} + \frac{1}{q} z^{\frac{p-1-q}{q}} \psi'(z^{\frac{1}{q}}).$$

By the construction, ψ is a convex function on \mathbb{R}_+ and $\psi(0) = 0$. Therefore, ψ' is non-decreasing, and therefore $\frac{1}{r}\psi(r) = \frac{1}{r} \int_0^r \psi'(s) ds$ is also non-decreasing. For $q \leq p-1$, the function $z^{\frac{p-q-1}{q}}$ is also non-decreasing in z , and apparently, for $r \geq 0$, both $\psi'(r)$ and $\psi(r)/r$ are non-negative. Therefore, for any $q \in (0, p-1]$, the function $\frac{dg_q}{dz}$ is non-decreasing in z . Therefore, g_q is a convex function.

C.2 Proof of theorem 5

For any fixed $T > 0$, we define the sequence of potential functions $\Phi_t : \mathbb{R}^d \rightarrow \mathbb{R}$

$$\Phi_t(\theta) := (\theta - \hat{\theta}^{(n)})^\top H^* e^{H^*(t-T)} (\theta - \hat{\theta}^{(n)}), \quad \text{for each } t \in [0, T].$$

Once again, we consider the diffusion process with the initial condition $\theta_0 = \widehat{\theta}^{(n)}$:

$$d\theta_t = -\nabla F_n(\theta_t)dt + \frac{1}{n}\nabla \log \pi(\theta_t)dt + dB_t.$$

Using Itô's formula, for $t \in [0, T]$, we have

$$\begin{aligned} \Phi_t(\theta_t) &= \int_0^t \frac{\partial \Phi_s}{\partial s}(\theta_s)ds - \int_0^t \langle \nabla \Phi_s(\theta_s), \nabla F_n(\theta_s) - \frac{\nabla \log \pi(\theta_s)}{n} \rangle ds \\ &\quad + \sqrt{\frac{2}{n}} \int_0^t \langle \nabla \Phi_s(\theta_s), dB_s \rangle + \frac{1}{n} \int_0^t \Delta \Phi_s(\theta_s)ds \\ &= \underbrace{\int_0^t \left(H^*(\theta_s - \widehat{\theta}^{(n)}) - \nabla F_n(\theta_s) + \frac{\nabla \log \pi(\theta_s)}{n} \right)^\top H^* e^{H^*(s-T)} (\theta_s - \widehat{\theta}^{(n)}) ds}_{:= I_1(t)} \\ &\quad + \underbrace{\sqrt{\frac{2}{n}} \int_0^t (\theta_s - \widehat{\theta}^{(n)})^\top H^* e^{(s-T)H^*} dB_s}_{I_2(t)} + \underbrace{\frac{1}{n} \int_0^t \text{Tr} \left(H^* e^{H^*(s-T)} \right) ds}_{I_3(t)}. \end{aligned} \quad (44)$$

Note that the matrices H^* and $e^{(s-T)H^*}$ commute, so that we may write their product in an arbitrary order.

Defining the linearization error

$$\Delta_s := (A + \varepsilon_1^{(2)}(n, \delta)) \left(\|\theta_s - \theta^*\|_2 + \|\widehat{\theta}^{(n)} - \theta^*\|_2 \right) + \varepsilon_2^{(2)}(n, \delta) + \frac{L_2}{n},$$

we claim that the following bounds hold for each $t \in [0, T]$:

$$\begin{aligned} I_1(t) &\leq \frac{2 + \log \kappa(H^*)}{a} \sup_{0 \leq s \leq t} \Phi_s(\theta_s) \\ &\quad + a \int_0^t \Delta_s^2 \left(\|\theta_s - \theta^*\|_2^2 + \|\widehat{\theta}^{(n)} - \theta^*\|_2^2 \right) e^{-\frac{\lambda_{\min}(H^*)}{2}(s-T)} ds, \end{aligned} \quad (45a)$$

$$\left(\mathbb{E} \sup_{0 \leq t \leq T} |I_2(t)|^p \right)^{1/p} \leq c \sqrt{\frac{p(1 + \log \kappa(H^*))}{n}} \left(\mathbb{E} \sup_{0 \leq t \leq T} \Phi_t(\theta_t)^{p/2} \right)^{1/p}, \quad \text{and} \quad (45b)$$

$$I_3(t) \leq \frac{d}{n}. \quad (45c)$$

Here $c > 0$ is an universal constant. We prove all of these bounds in the subsections to follow.

Taking these bounds as given for the moment, let us complete the proof of the theorem. By Jensen's inequality, for an even integer $p \geq 2$, the moments of the integral term in equation (45a) can be bounded as

$$\begin{aligned} \mathbb{E} \left(\int_0^T \Delta_s^2 \left(\|\theta_s - \theta^*\|_2^2 + \|\widehat{\theta}^{(n)} - \theta^*\|_2^2 \right) e^{-\frac{\lambda_{\min}(H^*)}{2}(s-T)} ds \right)^p \\ \leq \left(\frac{c}{\lambda_{\min}(H^*)} \right)^{p-1} \cdot \mathbb{E} \int_0^T \Delta_s^{2p} \left(\|\theta_s - \theta^*\|_2^{2p} + \|\widehat{\theta}^{(n)} - \theta^*\|_2^{2p} \right) e^{-\frac{\lambda_{\min}(H^*)}{2}(s-T)} ds, \end{aligned} \quad (46)$$

for a universal constant $c > 0$.

For any $\omega \in (0, 1)$, by taking supremum on both sides of the decomposition (44), combining with the bounds (45a) and (45c), and taking $a = c \frac{2 + \log \kappa(H^*)}{\omega}$, we arrive at the inequality

$$\begin{aligned} \sup_{0 \leq t \leq T} \Phi_t(\theta_t) &\leq (1 + \omega) \left(\frac{d}{n} + \sup_{0 \leq t \leq T} I_2(t) \right) \\ &\quad + \frac{c(2 + \log \kappa(H^*))}{\omega} \int_0^T \Delta_t^2 \left(\|\theta_t - \theta^*\|_2^2 + \|\widehat{\theta}^{(n)} - \theta^*\|_2^2 \right) e^{-\frac{\lambda_{\min}(H^*)}{2}(t-T)} dt. \end{aligned}$$

Taking p -th moment on both sides of the inequality, combining with the bounds (45b) and (46), and applying Minkowski's inequality, we arrive at the bound

$$\begin{aligned} \left(\mathbb{E} \sup_{0 \leq t \leq T} \Phi_t(\theta_t)^p \right)^{1/p} &\leq (1 + \omega) \frac{d}{n} + \sqrt{\frac{cp(1 + \log \kappa(H^*))}{n}} \cdot \left(\mathbb{E} \sup_{0 \leq t \leq T} \Phi_t(\theta_t)^p \right)^{\frac{1}{2p}} \\ &\quad + \frac{c(2 + \log \kappa(H^*))}{\omega \lambda_{\min}(H^*)} \left(\sup_{0 \leq t \leq T} \mathbb{E} \left[\Delta_t^{2p} \left(\|\theta_t - \theta^*\|_2^{2p} + \|\widehat{\theta}^{(n)} - \theta^*\|_2^{2p} \right) \right] \right)^{1/p}. \end{aligned}$$

Substituting with the definition of the last term, and applying Young's inequality, we find that

$$\left(\mathbb{E} \sup_{0 \leq t \leq T} \Phi_t(\theta_t)^p \right)^{1/p} \leq (1 + \omega) \frac{d}{n} + c \frac{1 + \log \kappa(H^*)}{\omega} \left(\frac{p}{n} + \frac{\mathcal{H}_n(p, \delta)}{\lambda_{\min}(H^*)} \right),$$

where the high-order term $\mathcal{H}_n(p, \delta)$ is defined as

$$\begin{aligned} \mathcal{H}_n(p, \delta) &:= (A + \varepsilon_1^{(2)}(n, \delta))^2 \left(\mathbb{E}_{\Pi} \|\theta - \theta^*\|_2^{4p} \right)^{1/p} \\ &\quad + \left\| \widehat{\theta}^{(n)} - \theta^* \right\|_2^2 \left(\varepsilon_2^{(2)}(n, \delta)^2 + \frac{L_2^2}{n^2} + (A + \varepsilon_1^{(2)}(n, \delta))^2 \left\| \widehat{\theta}^{(n)} - \theta^* \right\|_2^2 \right). \end{aligned}$$

Putting together the pieces yields the conclusion of the theorem.

C.2.1 Proof of claim (45a)

We first bound the term $I_1(t)$. Noting the defining identity $\nabla F_n(\widehat{\theta}^{(n)}) + \frac{1}{n} \nabla \log \pi(\widehat{\theta}^{(n)}) = 0$, we have the following bound:

$$\begin{aligned} &\left\| H^*(\theta_s - \widehat{\theta}^{(n)}) - \nabla F_n(\theta_s) + \nabla \log \pi(\theta_s)/n \right\|_2 \\ &= \left\| \int_0^1 \left(H^* - \nabla^2 F_n(\gamma \theta_s + (1 - \gamma) \widehat{\theta}^{(n)}) + \nabla^2 \log \pi(\gamma \theta_s + (1 - \gamma) \widehat{\theta}^{(n)})/n \right) (\theta_s - \widehat{\theta}^{(n)}) d\gamma \right\|_2 \\ &\leq \int_0^1 \|H^* - \nabla^2 F_n(\gamma \theta_s + (1 - \gamma) \widehat{\theta}^{(n)}) + \nabla^2 \log \pi(\gamma \theta_s + (1 - \gamma) \widehat{\theta}^{(n)})/n\|_{\text{op}} \cdot \left\| \theta_s - \widehat{\theta}^{(n)} \right\|_2 d\gamma. \end{aligned}$$

By Assumptions **(BvM.1)**, **(BvM.2)**, and **(A)**, for any $\theta \in \mathbb{R}^d$, we have the bound

$$\begin{aligned} & \|H^* - \nabla^2 F_n(\theta) + \nabla^2 \log \pi(\theta)/n\|_{\text{op}} \\ & \leq \|H^* - \nabla^2 F(\theta)\|_{\text{op}} + \|\nabla^2 F(\theta) - \nabla^2 F_n(\theta)\|_{\text{op}} + \|\nabla^2 \log \pi(\theta)/n\|_{\text{op}} \\ & \leq A \|\theta - \theta^*\|_2 + \varepsilon_1^{(2)}(n, \delta) \|\theta - \theta^*\|_2 + \varepsilon_2^{(2)}(n, \delta) + \frac{L_2}{n}. \end{aligned}$$

Substituting into the bound for $I_1(t)$, for any $a > 0$, we have that

$$\begin{aligned} I_1(t) & \leq \int_0^t \| (H^*)^{1/2} e^{H^*(s-t)/2} \|_{\text{op}} \\ & \quad \times \| H^* - \nabla^2 F_n(\theta_s) + \nabla^2 \log \pi(\theta_s)/n \|_2 \left\| \theta_s - \hat{\theta}^{(n)} \right\|_2 \sqrt{\Phi_s(\theta_s)} ds \\ & \leq a^{-1} \sup_{0 \leq s \leq t} \Phi_s(\theta_s) \cdot \int_0^t \| (H^*)^{1/2} e^{H^*(s-T)/4} \|_{\text{op}}^2 ds \\ & \quad + a \int_0^t \| H^* - \nabla^2 F_n(\theta_s)^2 + \nabla^2 \log \pi(\theta_s)/n \|_{\text{op}}^2 \cdot \left\| \theta_s - \hat{\theta}^{(n)} \right\|_2^2 \| e^{H^*(s-T)/4} \|_{\text{op}}^2 ds \\ & \leq \frac{2 + \log \kappa(H^*)}{a} \sup_{0 \leq s \leq t} \Phi_s(\theta_s) \\ & \quad + a \int_0^t \Delta_s^2 \left(\|\theta_s - \theta^*\|_2^2 + \left\| \hat{\theta}^{(n)} - \theta^* \right\|_2^2 \right) e^{-\frac{\lambda_{\min}(H^*)}{2}(s-T)} ds. \end{aligned}$$

Therefore, claim (45a) follows.

C.2.2 Proof of claim (45b)

Note that $I_2(t)$ is a martingale with respect to the Brownian filtration. Applying the Burkholder-Gundy-Davis inequality for an arbitrary $p \geq 2$ yields

$$\begin{aligned} \left(\mathbb{E} \sup_{0 \leq t \leq T} |I_2(t)|^p \right)^{1/p} & \leq c \sqrt{\frac{p}{n}} \left(\mathbb{E} \left(\int_0^T \left\| H^* e^{(t-T)H^*} (\theta_t - \hat{\theta}^{(n)}) \right\|_2^2 dt \right)^{\frac{p}{2}} \right)^{1/p} \\ & \leq C \sqrt{\frac{p}{n}} \left(\mathbb{E} \left(\int_0^T \| (H^*)^{1/2} e^{\frac{t-T}{2} H^*} \|_{\text{op}}^2 \Phi_t(\theta_t) dt \right)^{\frac{p}{2}} \right)^{1/p} \\ & \leq c \sqrt{\frac{p}{n}} \left(\mathbb{E} \sup_{0 \leq t \leq T} \Phi_t(\theta_t)^{p/2} \right)^{1/p} \cdot \sqrt{\int_0^T \| (H^*)^{1/2} e^{\frac{t-T}{2} H^*} \|_{\text{op}}^2 dt}. \end{aligned}$$

We now observe that

$$\| (H^*)^{1/2} e^{\frac{t-T}{2} H^*} \|_{\text{op}}^2 = \| H^* e^{(t-T)H^*} \|_{\text{op}} = \max_{i \in [d]} \left(\lambda_i(H^*) e^{(t-T)\lambda_i(H^*)} \right).$$

Taking the time integral leads to the bound

$$\begin{aligned} \int_0^T \| (H^*)^{1/2} e^{\frac{t-T}{2} H^*} \|_{\text{op}}^2 dt & \leq \int_0^{+\infty} \max_{i \in [d]} \left(\lambda_i(H^*) e^{-t\lambda_i(H^*)} \right) dt \\ & \leq \underbrace{\int_0^{+\infty} \max_{\lambda_{\min}(H^*) \leq \lambda \leq \lambda_{\max}(H^*)} \left(\lambda e^{-t\lambda} \right) dt}_{=: J}. \end{aligned}$$

We now split the integral J into three parts, thereby obtaining

$$\begin{aligned} J &\leq \int_0^{\lambda_{\max}(H^*)^{-1}} \lambda_{\max}(H^*) e^{-t\lambda_{\max}(H^*)} dt \\ &\quad + \int_{\lambda_{\max}(H^*)^{-1}}^{\lambda_{\min}(H^*)^{-1}} \frac{dt}{et} + \int_{\lambda_{\min}(H^*)^{-1}}^{+\infty} \lambda_{\min}(H^*) e^{-t\lambda_{\min}(H^*)} dt \\ &\leq 1 + \frac{1}{e} \log \frac{\lambda_{\max}(H^*)}{\lambda_{\min}(H^*)}. \end{aligned} \tag{47}$$

Denote $\kappa(M) := \frac{\lambda_{\max}(M)}{\lambda_{\min}(M)}$ for a positive definite matrix M . Collecting the above inequalities, we find that the term $I_2(t)$ is upper bounded as

$$\left(\mathbb{E} \sup_{0 \leq t \leq T} |I_2(t)|^p \right)^{1/p} \leq c \sqrt{\frac{p(1 + \log \kappa(H^*))}{n}} \left(\mathbb{E} \sup_{0 \leq t \leq T} \Phi_t(\theta_t)^{p/2} \right)^{1/p}$$

for a universal constant $c > 0$. This completes the proof of the claim (45b).

C.2.3 Proof of claim (45c)

Finally, the term $I_3(t)$ is straightforward to upper bound as

$$I_3(t) \leq \frac{1}{n} \text{Tr} \left(H^* \int_0^T e^{H^*(s-T)} ds \right) \leq \frac{1}{n} \text{Tr} \left(H^* \int_0^{+\infty} e^{-sH^*} ds \right) = \frac{d}{n},$$

which establishes the claim (45c).

In this Appendix, we provide proofs of remaining theorems and propositions in the main text.

C.3 Proof of theorem 1

Throughout the proof, in order to simplify notation, we omit the conditioning on the σ -field $\mathcal{F}_n := \sigma(X_1^n)$; it should be taken as given. For $\alpha = \frac{1}{2}\mu - \varepsilon_1(n, \delta) > \frac{\mu}{6}$, we claim that

$$\frac{1}{2} e^{\alpha t} \|\theta_t - \theta^*\|_2^2 \leq \frac{1}{\sqrt{n}} M_t + U_n \frac{(e^{\alpha t} - 1)}{2\alpha}, \tag{48}$$

where $U_n := \frac{3B^2}{n^2} + \frac{3\varepsilon_2^2(n, \delta)}{\mu} + \frac{d}{n}$ and $M_t := \int_0^t e^{\alpha s} \langle \theta_s - \theta^*, dB_s \rangle$, which is a martingale.

Assume that the above claim is given at the moment (the proof of that claim is deferred to the end of the proof of the proposition). In order to bound the moments of martingale M_t , for any $p \geq 4$, we invoke the Burkholder-Gundy-Davis inequality [39] to find that

$$\begin{aligned} \mathbb{E} \left[\sup_{0 \leq t \leq T} |M_t|^{\frac{p}{2}} \right] &\leq (pC)^{\frac{p}{4}} \mathbb{E} \left[[M]_T^{\frac{p}{4}} \right] = (pC)^{\frac{p}{4}} \mathbb{E} \left(\int_0^T e^{2\alpha s} \|\theta_s - \theta^*\|_2^2 ds \right)^{\frac{p}{4}} \\ &\leq (pC)^{\frac{p}{4}} \mathbb{E} \left(\sup_{0 \leq t \leq T} e^{\alpha t} \|\theta_t - \theta^*\|_2^2 \int_0^T e^{\alpha s} ds \right)^{\frac{p}{4}} \\ &\leq \left(\frac{pC e^{\alpha T}}{\alpha} \right)^{\frac{p}{4}} \mathbb{E} \left(\sup_{0 \leq t \leq T} e^{\alpha t} \|\theta_t - \theta^*\|_2^2 \right)^{\frac{p}{4}}, \end{aligned}$$

where C is a universal constant. Therefore, we arrive at the following bound:

$$\begin{aligned}\mathbb{E} \left[\left(\sup_{0 \leq t \leq T} e^{\alpha t} \|\theta_t - \theta^*\|_2 \right)^p \right] &\leq \mathbb{E} \left(\frac{2}{\sqrt{n}} M_t \right)^{\frac{p}{2}} + \left(U_n \frac{(e^{\alpha T} - 1)}{\alpha} \right)^{\frac{p}{2}} \\ &\leq \left(U_n \frac{e^{\alpha T}}{\alpha} \right)^{\frac{p}{2}} + \left(\frac{p C e^{\alpha T}}{\alpha n} \right)^{\frac{p}{4}} \mathbb{E} \left(\sup_{0 \leq s \leq T} e^{\alpha s} \|\theta_s - \theta^*\|_2^2 \right)^{\frac{p}{4}}.\end{aligned}$$

For the right hand side of the above inequality, we can relate it to the left hand side by using Young's inequality, which is given by

$$\left(\frac{p C e^{\alpha T}}{\alpha n} \right)^{\frac{p}{4}} \mathbb{E} \left(\sup_{0 \leq s \leq T} e^{\alpha s} \|\theta_s - \theta^*\|_2^2 \right)^{\frac{p}{4}} \leq \frac{1}{2} \left(\frac{p C e^{\alpha T}}{\alpha n} \right)^{\frac{p}{2}} + \frac{1}{2} \mathbb{E} \left(\sup_{0 \leq s \leq T} e^{\alpha s} \|\theta_s - \theta^*\|_2^2 \right)^{\frac{p}{2}}.$$

Putting the above results together, and let $\alpha = \frac{\mu}{2}$, we find that

$$(\mathbb{E} [\|\theta_T - \theta^*\|_2^p])^{\frac{1}{p}} \leq e^{-\alpha T} \left(\mathbb{E} \sup_{0 \leq t \leq T} (e^{\alpha t} \|\theta_t - \theta^*\|_2^p) \right)^{\frac{1}{p}} \leq C' \left(\sqrt{\frac{U_n}{\mu}} + \sqrt{\frac{2p}{n\mu}} \right),$$

for universal constant $C' > 0$. Therefore, the diffusion process defined in equation (6) satisfies the following inequality

$$\sup_{t \geq 0} (\mathbb{E} [\|\theta_t - \theta^*\|_2^p])^{\frac{1}{p}} \leq c \left(\sqrt{\frac{d}{\mu n}} + \frac{B}{\mu n} + \frac{\varepsilon_2(n, \delta)}{\mu} + \sqrt{\frac{p}{n\mu}} \right)$$

for any $p \geq 1$. Combining the above inequality with the inequality (54) yields the conclusion of the proposition.

Proof of claim (48): For the given choice $\alpha > 0$, an application of Itô's formula yields the decomposition

$$\begin{aligned}\frac{1}{2} e^{\alpha t} \|\theta_t - \theta^*\|_2^2 &= -\frac{1}{2} \int_0^t \langle \theta^* - \theta_s, \nabla F_n(\theta_s) e^{\alpha s} \rangle ds + \frac{1}{2n} \int_0^t \langle \theta_s - \theta^*, \nabla \log \pi(\theta_s) e^{\alpha s} \rangle ds \\ &\quad + \frac{d}{2n} \int_0^t e^{\alpha s} ds + \frac{1}{\sqrt{n}} \int_0^t e^{\alpha s} \langle \theta_s - \theta^*, dB_s \rangle + \frac{1}{2} \int_0^t \alpha e^{\alpha s} \|\theta_s - \theta^*\|_2^2 ds \\ &= J_1 + J_2 + J_3 + J_4 + J_5.\end{aligned}\tag{49}$$

We begin by bounding the term J_1 in equation (49). Based on Assumption **(S.2)** regarding the perturbation error between F_n and F and the strong convexity of F , we have

$$\begin{aligned}J_1 &= -\frac{1}{2} \int_0^t \langle \theta^* - \theta_s, \nabla F_n(\theta_s) e^{\alpha s} \rangle ds \\ &\leq -\frac{1}{2} \int_0^t \langle \theta^* - \theta_s, \nabla F(\theta_s) e^{\alpha s} \rangle ds + \frac{1}{2} \int_0^t \|\theta_s - \theta^*\|_2 \|\nabla F(\theta_s) - \nabla F_n(\theta_s)\|_2 e^{\alpha s} ds \\ &\leq -\frac{1}{2} \int_0^t \mu \|\theta_s - \theta^*\|_2^2 e^{\alpha s} ds + \frac{1}{2} \int_0^t \|\theta_s - \theta^*\|_2 (\varepsilon_1(n, \delta) \|\theta_s - \theta^*\|_2 + \varepsilon_2(n, \delta)) e^{\alpha s} ds \\ &\leq -\frac{1}{2} \int_0^t \mu \|\theta_s - \theta^*\|_2^2 e^{\alpha s} ds + \frac{1}{2} \int_0^t \|\theta_s - \theta^*\|_2^2 (\varepsilon_1(n, \delta) + \mu/3) e^{\alpha s} ds + \frac{3\varepsilon_2^2(n, \delta)}{2\mu} \int_0^t e^{\alpha s} ds.\end{aligned}$$

The second term J_2 involving prior π can be controlled in the following way:

$$\begin{aligned} J_2 &= \frac{1}{2n} \int_0^t \langle \theta_s - \theta^*, \nabla \log \pi(\theta_s) e^{\alpha s} \rangle ds \leq \frac{1}{2n} \int_0^t B \|\theta_s - \theta^*\|_2 e^{\alpha s} ds \\ &\leq \int_0^t \frac{\mu}{6} \|\theta_s - \theta^*\|_2^2 e^{\alpha s} ds + \frac{3B^2}{n^2 \mu} \int_0^t e^{\alpha s} ds. \end{aligned}$$

For the third term J_3 , a direct calculation leads to

$$J_3 = \frac{d(e^{\alpha t} - 1)}{2\alpha n}.$$

Moving to the fourth term J_4 , it is a martingale as $J_4 = M_t / \sqrt{n}$. Putting the above results together, as $\alpha = \frac{1}{2}\mu - \varepsilon_1(n, \delta) > \frac{\mu}{6}$, we obtain that

$$\frac{1}{2} e^{\alpha t} \|\theta_t - \theta^*\|_2^2 \leq \frac{1}{\sqrt{n}} M_t + U_n \frac{(e^{\alpha t} - 1)}{2\alpha}.$$

Putting together the pieces yields the claim (48).

C.4 Proof of theorem 4

Let the pair (Ψ, ζ_n) to be the functions defined in equations (38) in the proof of theorem 3. We denote $F_n := \Psi + \zeta_n + \frac{1}{n} \log \pi$, and consider the process generated by running Langevin algorithm on the modified posterior distribution:

$$\tilde{\theta}_{k+1} = \tilde{\theta}_k + \eta \nabla \left(\Psi + \zeta_n + \frac{1}{n} \log \pi \right) (\tilde{\theta}_k) + \sqrt{\frac{2\eta}{n}} W_k. \quad (50)$$

Note that the potential function \tilde{F} is exactly the same as F_n within the ball $\mathbb{B}(\theta^*, r_0)$. Defining the event:

$$\mathcal{E}_k := \left\{ \max_{1 \leq i \leq k} \|\theta_i - \theta^*\|_2 \leq r_0 \right\}. \quad (51)$$

On the event \mathcal{E}_k , the process $(\tilde{\theta}_i)_{1 \leq i \leq k}$ has the same law as $(\theta_i)_{1 \leq i \leq k}$. In the following, we analyze the moments of the process $(\tilde{\theta}_i)_{1 \leq i \leq k}$. As with the proof of theorem 3, condition on the random data $(X_i)_{i=1}^n$.

Defining $\Delta_k = \tilde{\theta}_k - \theta^*$, for any integer $p \geq 1$, a direct expansion of the iterates yields:

$$\begin{aligned} &\mathbb{E} \left[\|\Delta_{k+1}\|_2^{2p} \right] \\ &\leq \sum_{q=0}^p \binom{2p}{2q} \left(\sqrt{\frac{2\eta}{n}} \right)^{2q} \mathbb{E}[\|W_k\|_2^{2q}] \cdot \mathbb{E} \left[\left\| \Delta_k + \eta \nabla \tilde{F}_n(\tilde{\theta}_k) \right\|_2^{2p-2q} \right] \\ &\leq \sum_{q=0}^p \binom{p}{q} \frac{(p+1) \cdots (2p)}{(q+1) \cdots (2q) \cdot (p-q+1) \cdots (2p-2q)} \left(\sqrt{\frac{c\eta(d+q)}{n}} \right)^{2q} \cdot \mathbb{E} \left[\left\| \Delta_k + \eta \nabla \tilde{F}_n(\tilde{\theta}_k) \right\|_2^{2p-2q} \right] \\ &\leq \sum_{q=0}^p \binom{p}{q} \left(\sqrt{\frac{c\eta p^2(d+p)}{n}} \right)^{2q} \cdot \left\{ \mathbb{E} \left[\left\| \Delta_k + \eta \nabla \tilde{F}_n(\tilde{\theta}_k) \right\|_2^{2p} \right] \right\}^{\frac{p-q}{p}} \\ &\leq \left(\frac{c\eta p^2(d+p)}{n} + \left\{ \mathbb{E} \left[\left\| \Delta_k + \eta \nabla \tilde{F}_n(\tilde{\theta}_k) \right\|_2^{2p} \right] \right\}^{1/p} \right)^p. \end{aligned}$$

Using the shorthand notation $\lambda_{2p} := \left\{ \mathbb{E} \left[\left\| \Delta_k + \eta \nabla \tilde{F}_n(\tilde{\theta}_k) \right\|_2^{2p} \right] \right\}^{\frac{1}{2p}}$, we conclude that:

$$\mathbb{E} \left[\left\| \Delta_{k+1} \right\|_2^{2p} \right] \leq \left(\lambda_{2p}^2 + \frac{c \eta p^2 (d+p)}{n} \right)^p. \quad (52)$$

By the local growth conditions **(LWC.1)**, **(LWC.2)**, and the global smoothness assumptions **(A)** and **(B)**, we note that:

$$\begin{aligned} & \left\| \Delta_k + \eta \nabla \tilde{F}_n(\tilde{\theta}_k) \right\|_2^2 \\ &= \left\| \Delta_k \right\|_2^2 + \eta \langle \tilde{\theta}_k - \theta^*, \nabla \Psi(\tilde{\theta}_k) \rangle + \eta^2 \left\| \nabla \tilde{F}_n(\tilde{\theta}_k) \right\|_2^2 \\ &\leq \left\| \Delta_k \right\|_2^2 + \langle \Delta_k, \Psi(\tilde{\theta}_k) \rangle + \eta \left\| \Delta_k \right\|_2 \cdot \left(\left\| \nabla \zeta_n(\tilde{\theta}_k) \right\|_2 + n^{-1} \left\| \nabla \log \pi(\tilde{\theta}_k) \right\|_2 \right) + \eta^2 L^2 \left\| \Delta_k \right\|_2^2 \\ &\leq (1 - \eta\mu + \eta^2 L^2) \left\| \Delta_k \right\|_2^2 + \eta \left(\varepsilon(n, \delta) + \frac{B}{n} \right) \left\| \Delta_k \right\|_2 \\ &\leq (1 - 2\eta\mu/3 + \eta^2 L^2) \left\| \Delta_k \right\|_2^2 + \frac{3\eta}{\mu} \left(\varepsilon(n, \delta) + \frac{B}{n} \right)^2. \end{aligned}$$

Given the stepsize $\eta < \frac{\mu}{3L^2}$, we have that:

$$\begin{aligned} \lambda_{2p}^{2p} &= \mathbb{E} \left[\left\| \Delta_k + \eta \nabla \tilde{F}_n(\tilde{\theta}_k) \right\|_2^{2p} \right] \\ &\leq \mathbb{E} \left\{ (1 - \eta\mu/3) \left\| \Delta_k \right\|_2^2 + \frac{3\eta}{\mu} \left(\varepsilon(n, \delta) + \frac{B}{n} \right)^2 \right\}^p \\ &\leq \left\{ (1 - \mu\eta/3) \left(\mathbb{E} [\left\| \Delta_k \right\|_2^{2p}] \right)^{\frac{1}{p}} + \frac{3\eta}{\mu} \left(\varepsilon(n, \delta) + \frac{B}{n} \right)^2 \right\}^p. \end{aligned}$$

Combining with the bound (52), we conclude that:

$$\left\{ \mathbb{E} \left[\left\| \Delta_{k+1} \right\|_2^{2p} \right] \right\}^{1/p} \leq (1 - \mu\eta/3) \left(\mathbb{E} [\left\| \Delta_k \right\|_2^{2p}] \right)^{\frac{1}{p}} + \frac{3\eta}{\mu} \left(\varepsilon(n, \delta) + \frac{B}{n} \right)^2 + \frac{c \eta p^2 (d+p)}{n}.$$

Solving this recursion, we arrive at the following bound for $k = 0, 1, 2, \dots$

$$\left\{ \mathbb{E} \left[\left\| \Delta_k \right\|_2^{2p} \right] \right\}^{1/p} \leq e^{-k\mu\eta/3} \left\| \Delta_0 \right\|_2^2 + \frac{9}{\mu^2} \left(\varepsilon(n, \delta) + \frac{B}{n} \right)^2 + \frac{3cp^2(d+p)}{\mu n} \quad (53)$$

By equation (53) and a union bound over $k = 0, 1, 2, \dots, T$, with probability $1 - \vartheta$, we have that:

$$\max_{0 \leq k \leq T} \left\| \Delta_k \right\|_2 \leq \left\| \Delta_0 \right\|_2 + \frac{3\varepsilon(n, \delta)}{\mu} + \frac{3B}{n\mu} + \sqrt{\frac{3cd}{n\mu} \log^3 \frac{T}{\vartheta}}.$$

Under the condition $\left\| \theta_0 - \theta^* \right\|_2 \leq r_0/2$ and the sample size condition (14), we have the uniform bound:

$$\mathbb{P}(\mathcal{E}_T) \geq 1 - \delta/2,$$

Consequently, on the event \mathcal{E}_T , we conclude the following moment bound on the last iterate of the Langevin algorithm:

$$\left\{ \mathbb{E} \left[\|\Delta_T\|_2^{2p} \cdot \mathbf{1}_{\mathcal{E}_T} \right] \right\}^{1/p} \leq e^{-T\mu\eta/3} \|\Delta_0\|_2^2 + \frac{9}{\mu^2} \left(\varepsilon(n, \delta) + \frac{B}{n} \right)^2 + \frac{3cp^2(d+p)}{\mu n},$$

which can be readily converted into the following bound with probability $1 - \delta$:

$$\|\Delta_T\|_2 \leq e^{-\frac{T\mu\eta}{12\log(1/\delta)}} \|\Delta_0\|_2 + c \left\{ \frac{\varepsilon(n, \delta)}{\mu} + \frac{B}{\mu n} + \log(1/\delta) \cdot \sqrt{\frac{d + \log(1/\delta)}{\mu n}} \right\}.$$

C.5 Proof of theorem 2

As in the proof of theorem 1, we omit the conditioning on $\mathcal{F}_n := \sigma(X_1^n)$. For any $p \geq 2$, we define the functions on the positive real line $(0, \infty)$

$$\nu_{(p)}(r) := \psi \left(r^{\frac{1}{p-1}} \right) r^{\frac{p-2}{p-1}}, \quad \text{and} \quad \tau_{(p)}(r^{p-1}\zeta(r)) := r^{p-2}\psi(r).$$

By Assumption **(W.2)**, the function $r \mapsto r^{p-1}\zeta(r)$ is strictly increasing and surjective function that maps from $[0, +\infty)$ to $[0, +\infty)$. Therefore, it is invertible and the function $\tau_{(p)}^{-1}$ is well-defined.

Now we claim that for any $p \geq 2$, the functions $\nu_{(p)}$ and $\tau_{(p)}$ are convex and strictly increasing, and that furthermore, the expectation $\mathbb{E} [\|\theta_t - \theta^*\|_2^p]$ is upper bounded by the integral

$$\frac{p}{2} \int_0^t \left(-R_p(s) + \varepsilon(n, \delta) \tau_{(p)}^{-1}(R_p(s)) + \frac{B}{n} \nu_{(p)}^{-1}(R_p(s)) + \frac{p-1+d}{n} \nu_{(p)}^{-1}(R_p(s))^{\frac{p-2}{p-1}} \right) ds, \quad (54)$$

where $R_p(s) := \mathbb{E} \left[\|\theta_s - \theta^*\|_2^{p-2} \psi(\|\theta_s - \theta^*\|_2) \right]$.

Taking the above claims as given for the moment, let us now complete the proof of the theorem. Since for each finite $q \geq 1$, the process $(\theta_t : t \geq 0)$ converges in \mathbb{L}^q norm, the limit $\lim_{t \rightarrow +\infty} R_p(t)$ exists. Since the functions $\tau_{(p)}$ and $\nu_{(p)}$ are convex and strictly increasing, their inverse functions are concave. Moreover, simple calculation leads to

$$\nabla_r \left(\nu_{(p)}^{-1}(r)^{\frac{p-2}{p-1}} \right) = \frac{p-2}{p-1} \cdot \frac{\nu_{(p)}^{-1}(r)^{-\frac{1}{p-1}}}{\nu'_{(p)}(\nu_{(p)}^{-1}(r))}. \quad (55)$$

Since $\nu_{(p)}$ is convex and increasing, the numerator is a decreasing positive function of r . Additionally, the denominator is an increasing positive function of r . Therefore, the derivative in equation (55) is a decreasing function of r , and the function $r \mapsto \nu_{(p)}^{-1}(r)^{\frac{p-2}{p-1}}$ is concave. Define the function

$$\phi(r) := -r + \varepsilon(n, \delta) \tau_{(p)}^{-1}(r) + \frac{B}{n} \nu_{(p)}^{-1}(r) + \frac{p-1+d}{n} \nu_{(p)}^{-1}(r)^{\frac{p-2}{p-1}},$$

and observe that ϕ is concave and $\phi(0) = 0$. Let r_* be the smallest positive solution to the equation

$$r = \varepsilon(n, \delta) \tau_{(p)}^{-1}(r) + \frac{B}{n} \nu_{(p)}^{-1}(r) + \frac{p-1+d}{n} \nu_{(p)}^{-1}(r)^{\frac{p-2}{p-1}}.$$

We then have $\phi(r) < 0$ for $r > r_*$ and $\phi(r) > 0$ for $r \in (0, r_*)$. By lemma 2, we have $\lim_{t \rightarrow +\infty} R_p(t) \leq r_*$.

Since $\nu_{(p)}$ is a convex and strictly increasing function, Jensen's inequality implies that

$$R_p(t) = \mathbb{E} \left(\|\theta_t - \theta^*\|_2^{p-2} \psi(\|\theta_t - \theta^*\|_2) \right) \geq \nu_{(p)} \left(\mathbb{E} \|\theta_t - \theta^*\|_2^{p-1} \right). \quad (56)$$

Therefore, if we define $z_* := \lim_{t \rightarrow +\infty} \left(\mathbb{E} \|\theta_t - \theta^*\|_2^{p-1} \right)^{\frac{1}{p-1}}$, we have $z_*^{p-1} \leq \nu_{(p)}^{-1}(r_*)$. Hence, we arrive at the following inequality

$$\begin{aligned} z_*^{p-2} \psi(z_*) &\leq \varepsilon(n, \delta) \tau_{(p)}^{-1} (\nu_{(p)}(z_*^{p-1})) + \frac{B}{n} z_*^{p-1} + \frac{p-1+d}{n} z_*^{p-2} \\ &= \varepsilon(n, \delta) z_*^{p-1} \zeta(z_*) + \frac{B}{n} z_*^{p-1} + \frac{p-1+d}{n} z_*^{p-2}. \end{aligned}$$

As a consequence, we find that

$$\psi(z_*) \leq \varepsilon(n, \delta) \zeta(z_*) z_* + \frac{B + (p-1)d}{n}.$$

Now, we claim that there exists a unique positive solution to equation (8). Given this claim, replacing p by $(p+1)$ and putting the above results together yields

$$\lim_{t \rightarrow +\infty} (\mathbb{E} (\|\theta_t - \theta^*\|_2^p))^{\frac{1}{p}} \leq z_p^*,$$

where z_p^* is the unique positive solution to the following equation:

$$\psi(z) = \varepsilon(n, \delta) \zeta(z) z + \frac{B}{n} z + \frac{p+d}{n}.$$

Combining the above inequality with the inequality (54) yields the conclusion of the theorem.

We now return to prove our earlier claims about the behavior of the functions $\nu_{(p)}$, $\tau_{(p)}$, the moment bound (54), and the existence of unique positive solution to equation (8).

C.5.1 Structure of the function $\nu_{(p)}$

Since ψ is a convex and strictly increasing function, by taking the second derivative, we find that

$$\begin{aligned} \nu_{(p)}''(r) &= \nabla_r^2 \left(\psi \left(r^{\frac{1}{p-1}} \right) r^{\frac{p-2}{p-1}} \right) \\ &= \frac{1}{p-1} r^{\frac{1}{p-1}-1} \psi'' \left(r^{\frac{1}{p-1}} \right) + \frac{1}{p-1} r^{-1} \left(\psi' \left(r^{\frac{1}{p-1}} \right) - r^{-\frac{1}{p-1}} \psi \left(r^{\frac{1}{p-1}} \right) \right) \geq 0 \end{aligned}$$

for all $r > 0$. As a consequence, the function $\nu_{(p)}$ is convex.

C.5.2 Structure of the function $\tau_{(p)}$

This proof exploits Assumption **(W.3)** on the functions ψ and ζ . For any $p \geq 2$, we denote $\zeta_{(p)} : r \rightarrow r^{p-1}\zeta(r)$ and $\psi_{(p)} : r \rightarrow r^{p-2}\psi(r)$ two strictly increasing functions. Therefore, we can define a function $\tau_{(p)} := \psi_{(p)} \circ \zeta_{(p)}^{-1}$, namely, $\tau_{(p)}(r^{p-1}\zeta(r)) = r^{p-2}\psi(r)$, for any $r > 0$. Following some calculation, we find that

$$\begin{aligned}\nabla_r (\tau_{(p)}(r^{p-1}\zeta(r))) &= [(p-1)r^{p-2}\zeta(r) + r^{p-1}\zeta'(r)] \tau'_{(p)}(r^{p-1}\zeta(r)) \\ &= (p-2)r^{p-3}\psi(r) + r^{p-2}\psi'(r).\end{aligned}$$

Setting $z = \zeta_{(p)}(r)$ leads to

$$\nabla_z \tau_{(p)}(z) = \frac{(p-2)\psi(r) + r\psi'(r)}{(p-1)r\zeta(r) + r^2\zeta'(r)}.$$

Taking another derivative of the above term, we find that

$$\nabla_z^2 \tau_{(p)}(z) = \left(\zeta'_{(p)}(r) \right)^{-1} \frac{g(r, p)}{((p-1)r\zeta(r) + r^2\zeta'(r))^2},$$

where we denote

$$\begin{aligned}g(r, p) &:= [(p-1)r\zeta(r) + r^2\zeta'(r)] \cdot [(p-1)\psi'(r) + r\psi''(r)] \\ &\quad - [(p-1)\zeta(r) + (p+1)r\zeta'(r) + r^2\zeta''(r)] \cdot [(p-2)\psi(r) + r\psi'(r)].\end{aligned}$$

According to Assumption **(W.3)**, the function $\tau_{(2)} = \psi_{(2)} \circ \zeta_{(2)}^{-1}$ is convex. Therefore, we have $g(r, 2) \geq 0$ for any $r > 0$. Simple algebra with first order derivative of function g with respect to parameter p leads to

$$\begin{aligned}\nabla_p (g(r, p)) &= \zeta(r) \cdot [(p-1)r\psi'(r) + r^2\psi''(r) - (p-2)\psi(r) - r\psi'(r)] \\ &\quad - r\zeta'(r) [(p-2)\psi(r) + r\psi'(r)] + r\psi'(r) \cdot [(p-1)\zeta(r) + r\zeta'(r)] \\ &\quad - \psi(r) \cdot [(p-1)\zeta(r) + (p+1)r\zeta'(r) + r^2\zeta''(r)] \\ &= 2(p-2) [r\psi'(r)\zeta(r) - \psi(r)\zeta(r) - r\zeta'(r)\psi(r)] \\ &\quad + [r^2\zeta(r)\psi''(r) + r\psi'(r)\zeta(r) - 3\psi(r)\zeta(r) - r^2\psi(r)\zeta''(r)] \geq 0\end{aligned}$$

for all $r > 0$. Here the last inequality follows from Assumption **(W.3)**. Therefore, the function g is increasing function in terms of p when $p \geq 2$, so that $g(r, p) \geq g(r, 2) \geq 0$ for all $r > 0$. Given this inequality, we have $\frac{d^2}{dz^2} \tau_{(p)}(z) \geq 0$ for any $z \geq 0$, $p \geq 2$, i.e., the function $\tau_{(p)}(z)$ is a convex function for $z = \zeta_{(p)}(r)$.

C.5.3 Proof of claim (54)

For any $p \geq 2$, an application of Itô's formula yields the bound $\|\theta_t - \theta^*\|_2^p \leq \sum_{j=1}^5 T_j$, where

$$T_1 := -\frac{p}{2} \int_0^t \langle \theta^* - \theta_s, \nabla F(\theta_s) \rangle \|\theta_s - \theta^*\|_2^{p-2} ds, \quad (57a)$$

$$T_2 := \frac{p}{2} \int_0^t \langle \theta^* - \theta_s, \nabla F(\theta_s) - \nabla F_n(\theta_s) \rangle \|\theta_s - \theta^*\|_2^{p-2} ds \quad (57b)$$

$$T_3 := \frac{p}{2n} \int_0^t \langle \theta_s - \theta^*, \nabla \log \pi(\theta_s) \rangle \|\theta_s - \theta^*\|_2^{p-2} ds \quad (57c)$$

$$T_4 := p \int_0^t \|\theta_s - \theta^*\|_2^{p-2} \langle \theta_s - \theta^*, dB_s \rangle \quad (57d)$$

$$T_5 := \frac{p(p-1+d)}{2n} \int_0^t \|\theta_s - \theta^*\|_2^{p-2} ds. \quad (57e)$$

We now upper bound the terms $\{T_j\}_{j=1}^5$ in terms of functionals of the quantity R_p . From the weak convexity of F guaranteed by Assumption W.1, we have

$$\mathbb{E}[T_1] = -\frac{p}{2} \mathbb{E} \left[\int_0^t \langle \theta^* - \theta_s, \nabla F(\theta_s) \rangle \|\theta_s - \theta^*\|_2^{p-2} ds \right] \leq -\frac{p}{2} \int_0^t R_p(s) ds. \quad (58a)$$

Based on Assumption **(W.2)**, we find that

$$\begin{aligned} \mathbb{E}[T_2] &= \frac{p}{2} \mathbb{E} \left[\int_0^t \langle \theta^* - \theta_s, \nabla F(\theta_s) - \nabla F_n(\theta_s) \rangle \|\theta_s - \theta^*\|_2^{p-2} ds \right] \\ &\leq \frac{p}{2} \varepsilon(n, \delta) \int_0^t \mathbb{E} \left[\|\theta_s - \theta^*\|_2^{p-1} \zeta(\|\theta_s - \theta^*\|_2) \right] ds. \end{aligned}$$

Since the function $\tau_{(p)}$ is convex, invoking Jensen's inequality, we obtain the following inequalities:

$$\begin{aligned} \int_0^t \mathbb{E} \left[\|\theta_s - \theta^*\|_2^{p-1} \zeta(\|\theta_s - \theta^*\|_2) \right] ds &\leq \int_0^t \tau_{(p)}^{-1} \mathbb{E} \left[\tau_{(p)} \left(\|\theta_s - \theta^*\|_2^{p-1} \zeta(\|\theta_s - \theta^*\|_2) \right) \right] ds \\ &= \int_0^t \tau_{(p)}^{-1} (R_p(s)) ds. \end{aligned}$$

In light of the above inequalities, we have

$$\mathbb{E}[T_2] \leq \frac{p}{2} \varepsilon(n, \delta) \int_0^t \tau_{(p)}^{-1} (R_p(s)) ds. \quad (58b)$$

Moving to T_3 in equation (57c), given Assumption **(B)** which controls the growth of prior distribution π , its expectation is bounded as

$$\begin{aligned} \mathbb{E}[T_3] &= \frac{p}{2n} \mathbb{E} \left[\int_0^t \langle \theta_s - \theta^*, \nabla \log \pi(\theta_s) \rangle \|\theta_s - \theta^*\|_2^{p-2} ds \right] \\ &\leq \frac{pB}{2n} \int_0^t \mathbb{E} \left[\|\theta_s - \theta^*\|_2^{p-1} \right] ds. \end{aligned} \quad (58c)$$

By exploiting the bound (56) along with the fact that $\nu_{(p)}$ is strictly increasing on $[0, +\infty)$, we find that

$$\int_0^t \mathbb{E} \left(\|\theta_s - \theta^*\|_2^{p-1} \right) ds \leq \int_0^t \nu_{(p)}^{-1} (R_p(s)) ds. \quad (58d)$$

Combining the inequalities (58c) and (58d), we have

$$\mathbb{E} [T_3] \leq \frac{pB}{2n} \int_0^t \nu_{(p)}^{-1} (R_p(s)) ds. \quad (58e)$$

Moving to the fourth term T_4 from equation (57d), we have

$$\mathbb{E} [T_4] = \mathbb{E} \left[\int_0^t \|\theta_s - \theta^*\|_2^{p-2} \langle \theta_s - \theta^*, dB_s \rangle \right] = 0, \quad (58f)$$

where we have used the martingale structure.

For the last term T_5 , invoking Hölder's inequality and the bound (56), we have the moment estimate:

$$\mathbb{E} \left(\|\theta_s - \theta^*\|_2^{p-2} \right) \leq \left(\mathbb{E} \left[\|\theta_s - \theta^*\|_2^{p-1} \right] \right)^{\frac{p-2}{p-1}} \leq \nu_{(p)}^{-1} (R_p(s))^{\frac{p-2}{p-1}}.$$

Consequently, the term T_5 can be bounded in expectation as

$$\mathbb{E} [T_5] \leq \frac{p(p-1+d)}{2n} \int_0^t \nu_{(p)}^{-1} (R_p(s))^{\frac{p-2}{p-1}} ds. \quad (58g)$$

Collecting the bounds on the expectations of the terms $\{T_j\}_{j=1}^5$ from equations (58a)-(58g), respectively, yields the claim (54).

C.5.4 Unique positive solution to equation (8)

We now establish that equation (8) has a unique positive solution under the stated assumptions. Define the function

$$\vartheta(z) := \psi(z) - \left(\varepsilon(n, \delta) \zeta(z) z + \frac{B + d \log(1/\delta)}{n} \right).$$

Since $\psi(0) = 0$, we have $\vartheta(0) < 0$. On the other hand, based on Assumption **(W.4)**, $\liminf_{z \rightarrow +\infty} \vartheta(z) > 0$. Therefore, there exists a positive solution to the equation $\vartheta(z) = 0$.

Recall that $\xi : \mathbb{R}_+ \rightarrow \mathbb{R}$ is an inverse function of the strictly increasing function $z \mapsto z\zeta(z)$. Therefore, we can write the function ϑ as follows:

$$\vartheta(z) = \tilde{\vartheta}r := \psi(\xi(r)) - \varepsilon(n, \delta)r - \frac{B + d \log(1/\delta)}{n},$$

where $r = z \cdot \zeta(z)$. Given the convexity of function $r \mapsto \psi(\xi(r))$ guaranteed by Assumption **(W.3)**, the functions $\tilde{\vartheta}$ and ϑ are convex. Putting the above results together, there exists a unique positive solution to equation (8).

C.6 Proof of proposition 2

We introduce the shorthand $\mu := \mathcal{N}(\widehat{\theta}^{(n)}, (nH^*)^{-1})$ for the target density. Since $H^* \succ 0$, the Gaussian log-Sobolev inequality implies that

$$D_{\text{KL}}(\Pi(\cdot | X_1^n) \| \mu) \leq \frac{1}{n\lambda_{\min}(H^*)} \int_{\mathbb{R}^d} \|\nabla \log \Pi(\theta | X_1^n) - \nabla \log \mu(\theta)\|_2^2 \Pi(d\theta | X_1^n). \quad (59)$$

Since μ is a Gaussian density, we find that

$$\nabla \log \mu(\theta) = -nH^*(\theta - \widehat{\theta}^{(n)}).$$

For the posterior density $\Pi(\cdot | X_1^n)$, we note that

$$\begin{aligned} \nabla \log \Pi(\theta | X_1^n) &= -n\nabla F_n(\theta) + \nabla \log \pi(\theta) \\ &= \int_0^1 \left(-n\nabla^2 F_n(\gamma\theta + (1-\gamma)\widehat{\theta}^{(n)}) + \nabla^2 \log \pi(\gamma\theta + (1-\gamma)\widehat{\theta}^{(n)}) \right) \\ &\quad \times (\theta - \widehat{\theta}^{(n)}) d\gamma. \end{aligned}$$

Putting together the above equations together yields

$$\begin{aligned} &\|\nabla \log \Pi(\theta | X_1^n) - \nabla \log \mu(\theta)\|_2 \\ &\leq n \int_0^1 \|\nabla^2 F_n(\gamma\theta + (1-\gamma)\widehat{\theta}^{(n)}) - H^* + \nabla^2 \log \pi(\gamma\theta + (1-\gamma)\widehat{\theta}^{(n)})/n\|_{\text{op}} \cdot \|\theta - \widehat{\theta}^{(n)}\|_2 d\gamma. \end{aligned}$$

By Assumptions **(BvM.1)**, **(BvM.2)**, and **(A)**, we have the bounds

$$\begin{aligned} &\|\nabla^2 F_n(\gamma\theta + (1-\gamma)\widehat{\theta}^{(n)}) + \nabla^2 \log \pi(\gamma\theta + (1-\gamma)\widehat{\theta}^{(n)})/n - H^*\|_{\text{op}} \\ &\leq \|\nabla^2 F(\gamma\theta + (1-\gamma)\widehat{\theta}^{(n)}) - H^*\|_{\text{op}} \\ &\quad + \|\nabla^2 F_n(\gamma\theta + (1-\gamma)\widehat{\theta}^{(n)}) - \nabla^2 F_n(\gamma\theta + (1-\gamma)\widehat{\theta}^{(n)})\|_{\text{op}} + \frac{L_2}{n} \\ &\leq A \|\gamma\theta + (1-\gamma)\widehat{\theta}^{(n)} - \theta^*\|_2 + \varepsilon_1^{(2)}(n, \delta) \|\theta - \widehat{\theta}^{(n)}\|_2 + \varepsilon_2^{(2)}(n, \delta) + \frac{L_2}{n}. \end{aligned}$$

Substituting this bound into the bound (59) yields

$$\begin{aligned} D_{\text{KL}}(\Pi(\cdot | X_1^n) \| \mu) &\leq \frac{n}{\lambda_{\min}(H^*)} \left(A \cdot \mathbb{E}_{\Pi} \left[\|\theta - \theta^*\|_2^4 | X_1^n \right] + A \|\widehat{\theta}^{(n)} - \theta^*\|_2^4 \right) \\ &\quad + \frac{n\varepsilon_1^{(2)}(n, \delta)}{\lambda_{\min}(H^*)} \mathbb{E}_{\Pi} \left[\|\theta - \widehat{\theta}^{(n)}\|_2^3 | X_1^n \right] \\ &\quad + (\varepsilon_2^{(2)}(n, \delta) + L_2/n) \cdot \mathbb{E} \left[\|\theta - \widehat{\theta}^{(n)}\|_2^2 | X_1^n \right]. \end{aligned}$$

As a consequence, we obtain the conclusion of the proposition.

D Proofs of corollaries

In this appendix, we collect the proofs of several corollaries stated in the main text and section 5. To summarize, we make use of Theorems 2, 3, and 5 to establish the posterior contraction rates of parameters and non-asymptotic Bernstein-von Mises theorem in the examples in section 5. The crux of the proofs of these corollaries involves a verification of assumptions to invoke the respective theorems. Note that the values of universal constants may change from line-to-line.

D.1 Proof of corollary 2

We begin by verifying claim (19a) about the structure of the negative population log-likelihood function F^R and claim (19b) about the uniform perturbation error between ∇F^R and ∇F_n^R .

D.1.1 Proof of claim (19a)

Following some algebra, we find that

$$\begin{aligned} -F^R(\theta) &= \mathbb{E} \left[-Y \log \left(1 + e^{-\langle X, \theta \rangle} \right) - (1 - Y) \log \left(1 + e^{\langle X, \theta \rangle} \right) \right] \\ &= -\mathbb{E} \left[\frac{1}{1 + e^{-\langle X, \theta^* \rangle}} \log \left(1 + e^{-\langle X, \theta \rangle} \right) + \frac{1}{1 + e^{\langle X, \theta^* \rangle}} \log \left(1 + e^{\langle X, \theta \rangle} \right) \right], \end{aligned}$$

where the above expectations are taken with respect to $X \sim \mathcal{N}(0, \sigma^2 I_d)$ and $Y|X$ following probability distribution generated from logistic model (18). Taking the derivative of F^R with respect to θ yields

$$\begin{aligned} \langle \nabla F^R(\theta), \theta^* - \theta \rangle \\ = \mathbb{E} \left[\left(\frac{1 + e^{\langle X, \theta \rangle}}{1 + e^{\langle X, \theta^* \rangle}} - \frac{1 + e^{-\langle X, \theta \rangle}}{1 + e^{-\langle X, \theta^* \rangle}} \right) \frac{e^{-\langle X, \theta \rangle}}{(1 + e^{-\langle X, \theta \rangle})^2} \langle X, \theta - \theta^* \rangle \right]. \end{aligned}$$

By the mean value theorem, there exists ξ between 0 and $\langle X, \theta - \theta^* \rangle$ such that

$$\frac{1 + e^{\langle X, \theta \rangle}}{1 + e^{\langle X, \theta^* \rangle}} - \frac{1 + e^{-\langle X, \theta \rangle}}{1 + e^{-\langle X, \theta^* \rangle}} = \langle X, \theta - \theta^* \rangle \left(\frac{e^{\langle X, \theta^* \rangle} + \xi}{1 + e^{\langle X, \theta^* \rangle}} + \frac{e^{-\langle X, \theta^* \rangle} - \xi}{1 + e^{-\langle X, \theta^* \rangle}} \right).$$

In light of the above equality, we arrive at the following inequalities:

$$\begin{aligned} \langle \nabla F^R(\theta), \theta^* - \theta \rangle &\geq \mathbb{E} \left[\inf_{|\xi| \in [0, |\langle X, \theta - \theta^* \rangle|]} \left(\frac{e^{\langle X, \theta^* \rangle} + \xi}{1 + e^{\langle X, \theta^* \rangle}} + \frac{e^{-\langle X, \theta^* \rangle} - \xi}{1 + e^{-\langle X, \theta^* \rangle}} \right) \right. \\ &\quad \times \left. \frac{e^{-\langle X, \theta \rangle}}{(1 + e^{-\langle X, \theta \rangle})^2} |\langle X, \theta - \theta^* \rangle|^2 \right] \\ &\geq \mathbb{E} \left[\frac{1}{2} e^{-|\langle X, \theta - \theta^* \rangle|} \frac{e^{-\langle X, \theta \rangle}}{(1 + e^{-\langle X, \theta \rangle})^2} |\langle X, \theta - \theta^* \rangle|^2 \right] \\ &\geq \frac{1}{8} \mathbb{E} \left[e^{-|\langle X, \theta - \theta^* \rangle| - |\langle X, \theta \rangle|} |\langle X, \theta - \theta^* \rangle|^2 \right] \\ &\geq \frac{1}{8e^4} \mathbb{E} \left[\mathbf{1}_{\{|\langle X, \theta \rangle| \leq 2, |\langle X, \theta - \theta^* \rangle| \leq 2\}} |\langle X, \theta - \theta^* \rangle|^2 \right]. \end{aligned}$$

Since $X \sim \mathcal{N}(0, I_d)$, we have

$$\begin{bmatrix} \langle X, \theta \rangle \\ \langle X, \theta - \theta^* \rangle \end{bmatrix} \sim \mathcal{N} \left(0, \begin{bmatrix} \|\theta\|_2^2 & \langle \theta, \theta - \theta^* \rangle \\ \langle \theta, \theta - \theta^* \rangle & \|\theta - \theta^*\|_2^2 \end{bmatrix} \right).$$

Given that result, direct calculation leads to

$$\begin{aligned} \mathbb{E} (\mathbf{1}_{\{|\langle X, \theta \rangle| \leq 2, |\langle X, \theta - \theta^* \rangle| \leq 2\}} |\langle X, \theta - \theta^* \rangle|^2) \\ \geq \frac{c}{(1 + \|\theta\|_2)(1 + \|\theta - \theta^*\|_2)} \|\theta - \theta^*\|_2^2, \end{aligned}$$

for a universal constant $c > 0$. Collecting the above results, for all θ such that $\|\theta - \theta^*\|_2 \leq 1$, we achieve that

$$\begin{aligned} \langle \nabla F^R(\theta), \theta^* - \theta \rangle &\geq \frac{c}{(1 + \|\theta\|_2)(1 + \|\theta - \theta^*\|_2)} \|\theta - \theta^*\|_2^2 \\ &\geq c \frac{1}{1 + \|\theta^*\|_2} \|\theta - \theta^*\|_2^2. \end{aligned}$$

For θ with $\|\theta - \theta^*\|_2 > 1$, let $\tilde{\theta} = \theta^* + \frac{\theta - \theta^*}{\|\theta - \theta^*\|_2}$. Then, we find that

$$\langle \nabla F^R(\theta), \theta^* - \theta \rangle \geq \langle \nabla F^R(\tilde{\theta}), \theta^* - \theta \rangle \geq \frac{c}{2(1 + \|\theta^*\|_2)} \|\theta - \theta^*\|_2,$$

which yields the claim (19a).

D.1.2 Proof of the bound (19b)

In this appendix, we prove the uniform bound (19b) between the empirical and population likelihood gradients. It suffices to establish the following stronger result:

$$Z := \sup_{\theta \in \mathbb{R}^d} \|\nabla F_n^R(\theta) - \nabla F^R(\theta)\|_2 \leq c \left\{ \sqrt{\frac{d}{n}} + \sqrt{\frac{\log(1/\delta)}{n}} + \frac{\log(1/\delta)}{n} \right\}, \quad (60)$$

with probability at least $1 - \delta$ for any $\frac{n}{\log n} \geq c_0 d \log(1/\delta)$ where c_0 is a universal constant.

In order to prove the claim (60), we exploit a concentration inequality due to Adamczak [1]; it gives tight tail bounds for supremum of unbounded empirical processes. Throughout our derivation, we use $\|X\|_{\psi_\alpha}$ to denote the Orlicz ψ_α norm for a random variable X , for any $\alpha \in (0, 2]$. Let us state a simplified version of a theorem due to Adamczak:

Proposition 3 (Theorem 4 of [1], simplified version). *Let $(x, \theta) \mapsto f(\theta; x)$ be a function with domain $\Theta \times \mathcal{X}$, and suppose that there is a function $\bar{F} : \mathcal{X} \rightarrow \mathbb{R}$ such that $|f(\theta, x)| \leq \bar{F}(x)$ for any $\theta \in \Theta$. Let $X_1, X_2, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}_X$, and suppose that $\|\bar{F}\|_{\psi_\alpha} < +\infty$ for some $\alpha \leq 1$. Then the random variable $Z_n := \frac{1}{n} \sup_{\theta \in \Theta} |\sum_{i=1}^n f(\theta; X_i) - \mathbb{E}[f(\theta; X)]|$ satisfies the bound:*

$$\mathbb{P}(Z_n > 2\mathbb{E}[Z_n] + t) \leq \exp \left(-\frac{t^2}{2\mathbb{E}[\bar{F}(X)^2]} \right) + 3 \exp \left(- \left(\frac{t}{c \|\max_{i \in [n]} \bar{F}(X_i)\|_{\psi_\alpha}} \right)^\alpha \right),$$

for a universal constant $c > 0$.

In order to prove the claim (60), we begin by writing Z as the supremum of a stochastic process. Let \mathbb{S}^{d-1} denote the Euclidean sphere in \mathbb{R}^d , and define the stochastic process

$$Z_{u,\theta} := \left| \frac{1}{n} \sum_{i=1}^n f_{u,\theta}(X_i, Y_i) - \mathbb{E}[f_{u,\theta}(X, Y)] \right|,$$

where $f_{u,\theta}(x, y) = \frac{y\langle x, u \rangle e^{y\langle x, \theta \rangle}}{1 + e^{y\langle x, \theta \rangle}}$, indexed by vectors $u \in \mathbb{S}^{d-1}$ and $\theta \in \mathbb{B}(\theta^*; r)$. The outer expectation in the above display is taken with respect to (X, Y) drawn from the logistic model (18)

Observe that $Z = \sup_{u \in \mathbb{S}^{d-1}} \sup_{\theta \in \mathbb{R}^d} Z_{u,\theta}$. Let $\{u^1, \dots, u^N\}$ be a $1/8$ -covering of \mathbb{S}^{d-1} in the Euclidean norm; there exists such a set with $N \leq 17^d$ elements. By a standard discretization argument (see Chapter 6, [55]), we have

$$Z \leq 2 \max_{j=1, \dots, N} \sup_{\theta \in \mathbb{R}^d} Z_{u^j, \theta}.$$

Accordingly, the remainder of our argument focuses on bounding the random variable $V := \sup_{\theta \in \mathbb{R}^d} Z_{u,\theta}$, where the vector $u \in \mathbb{S}^{d-1}$ should be understood as arbitrary but fixed. For each $u \in \mathbb{S}^{d-1}$ fixed, we note that $\bar{F}(X, Y) = |\langle X, u \rangle|$ is an envelop function for the class $(f_{u,\theta}(X, Y))_{\theta \in \mathbb{R}^d}$. Additionally, by standard tail bounds for maximum of Gaussian random variables, we know that:

$$\left\| \max_{1 \leq i \leq n} \bar{F}(X_i, Y_i) \right\|_{\psi_1} \leq \sqrt{\log n}.$$

Consequently, invoking proposition 3 yields that

$$V \leq 2\mathbb{E}[V] + \sqrt{\frac{2\log(1/\delta)}{n}} + \frac{c\log(1/\delta)}{n} \sqrt{\log n} \quad (61)$$

with probability at least $1 - \delta$.

Now define the symmetrized random variable

$$V' := \sup_{\theta \in \mathbb{R}^d} \left| \frac{1}{n} \sum_{i=1}^n \varepsilon_i f_{\theta,u}(X_i, Y_i) \right|,$$

where $\{\varepsilon_i\}_{i=1}^n$ is an i.i.d. sequence of Rademacher variables. By standard symmetrization arguments, we have

$$\mathbb{E}[V] \leq 2\mathbb{E}[V'].$$

We now bound the expectation of V' , first over the Rademacher variables. Consider the function class

$$\mathcal{G} := \left\{ g_\theta : (x, y) \mapsto \langle x, u \rangle \varphi_\theta(x, y) \mid \theta \in \mathbb{R}^d \right\}.$$

It is clear that the function class \mathcal{G} has the envelope function $\bar{G}(x) := |\langle x, u \rangle|$. We claim that the L_2 -covering number of \mathcal{G} can be bounded as

$$\bar{N}(t) := \sup_Q \left| \mathcal{N} \left(\mathcal{G}, \|\cdot\|_{L^2(Q)}, t \|\bar{G}\|_{L^2(Q)} \right) \right| \leq \left(\frac{1}{t} \right)^{c(d+1)} \quad \text{for all } t > 0, \quad (62)$$

where $c > 0$ is a universal constant.

Let us take the claim (62) as given for the moment, and use it to bound the expectation of V' , first over the Rademacher variables. Define the empirical expectation $\mathbb{P}_n(\bar{G}^2) := \frac{1}{n} \sum_{i=1}^n \langle X_i, u \rangle^2$. Invoking Dudley's entropy integral bound (e.g., Theorem 5.22, [55]), we find that there are universal constants C, C' such that

$$\begin{aligned} \mathbb{E}_\varepsilon[V'] &= \mathbb{E}_\varepsilon \left[\sup_{g \in \mathcal{G}} \left| \frac{1}{n} \sum_{i=1}^n \varepsilon_i g(X_i, Y_i) \right| \right] \leq C \sqrt{\frac{\mathbb{P}_n(\bar{G}^2)}{n}} \int_0^1 \sqrt{1 + \log \bar{N}(t)} dt \\ &\leq C' \sqrt{\mathbb{P}_n(\bar{G}^2)} \sqrt{\frac{d}{n}}. \end{aligned}$$

Up to this point, we have been conditioning on the observations $\{X_i\}_{i=1}^n$. Taking expectations over them as well yields

$$\mathbb{E}_{\varepsilon, X_1^n}[V'] \leq C' \sqrt{\frac{d}{n}} \cdot \mathbb{E}_{X_1^n} \left[\sqrt{\mathbb{P}_n(\bar{G}^2)} \right] \stackrel{(i)}{\leq} C' \sqrt{\frac{d}{n}} \cdot \sqrt{\mathbb{E}_{X_1^n} [\mathbb{P}_n(\bar{G}^2)]} \stackrel{(ii)}{=} C' \sqrt{\frac{d}{n}}, \quad (63)$$

where step (i) follows from Jensen's inequality; and step (ii) uses the fact that $\mathbb{E}_{X_1^n}[\mathbb{P}_n(\bar{G}^2)] = 1$. Putting together the bounds (61) and (63) yields the following bound with probability $1 - \delta$:

$$V \leq c \sqrt{\frac{d + \log \delta^{-1}}{n}} + c \frac{\log \delta^{-1}}{n} \sqrt{\log n}.$$

This probability bound holds for each $u \in \mathbb{S}^{d-1}$. By taking the union bound over the $1/8$ -covering set $\{u^1, \dots, u^N\}$ of \mathbb{S}^{d-1} where $N \leq 17^d$ and applying above bound with $\delta' = \delta/N$, we obtain the claim (60) for sample size satisfying $\frac{n}{\log n} \geq cd \log(1/\delta)$.

D.1.3 Proof of claim (62)

We consider a fixed sequence $(x_i, y_i, t_i)_{i=1}^m$ where $y_i \in \{-1, 1\}$, $x_i \in \mathbb{R}^d$ and $t_i \in \mathbb{R}$ for $i \in [m]$. Now, we suppose that for any binary sequence $(z_i)_{i=1}^m \in \{0, 1\}^m$, there exists $\theta \in \mathbb{R}^d$ such that

$$z_i = \mathbb{I}[\langle X_i, u \rangle \varphi_\theta(X_i, Y_i) \geq t_i] \quad \text{for all } i \in [m].$$

Following some algebra, we find that

$$y_i x_i^T \theta - \log \frac{Y_i t_i}{\langle X_i, u \rangle - Y_i t_i} \begin{cases} \geq 0 & z_i = 1 \\ < 0 & z_i = 0 \end{cases}.$$

Consequently, the set $\{[y_i x_i, \log(Y_i t_i / \langle X_i, u \rangle - Y_i t_i))]\}_{i=1}^m$ of $(d+1)$ -dimensional points can be shattered by linear separators. Therefore, we have $m \leq d+2$, which leads to the VC subgraph dimension of \mathcal{G} to be at most $d+2$ (e.g., see the book [54]). As a consequence, we obtain the conclusion of the claim (62).

D.2 Proof of corollary 4

The claim (28a) of weak convexity for the negative population log-likelihood function F^I is straightforward. Therefore, we only need to establish the claim (28b) about the uniform perturbation bound between ∇F^I and ∇F_n^I .

D.2.1 Bounding the difference $\nabla F^I - \nabla F_n^I$

It is convenient to introduce the shorthand

$$p_\theta(x, y) = \left(y - (x^\top \theta)^p \right)^2 / 2 \quad \text{for all } (x, y) \in \mathbb{R}^{d+1}.$$

We then compute the gradient

$$\nabla \log p_\theta(x, y) = p \left(y - (x^\top \theta)^p \right) (x^\top \theta)^{p-1} x.$$

Fix an arbitrary $r > 0$, by applying the triangle inequality, we find that

$$\begin{aligned} & \sup_{\theta \in \mathbb{B}(\theta^*, r)} \|\nabla F_n^I(\theta) - \nabla F^I(\theta)\|_2 \\ &= \sup_{\theta \in \mathbb{B}(\theta^*, r)} \left\| \frac{1}{n} \sum_{i=1}^n \nabla \log p_\theta(X_i, Y_i) - \mathbb{E}_{(X, Y)} [\nabla \log p_\theta(X, Y)] \right\|_2 \\ &\leq p \{J_1 + J_2\}, \end{aligned}$$

where we define

$$J_1 := p \sup_{\theta \in \mathbb{B}(\theta^*, r)} \left\| \frac{1}{n} \sum_{i=1}^n Y_i X_i (X_i^\top \theta)^{p-1} \right\|_2, \quad \text{and} \quad (64a)$$

$$J_2 := p \sup_{\theta \in \mathbb{B}(\theta^*, r)} \left\| \frac{1}{n} \sum_{i=1}^n X_i (X_i^\top \theta)^{2p-1} - \mathbb{E}_X \left[X (X^\top \theta)^{2p-1} \right] \right\|_2. \quad (64b)$$

We claim that there is a universal constant c such that for any $\delta \in (0, 1)$, the quantities J_1 and J_2 can be bounded as

$$J_1 \leq c r^{p-1} \left(\sqrt{\frac{d + \log \frac{1}{\delta}}{n}} + \frac{1}{n^{3/2}} \left(d + \log \frac{n}{\delta} \right)^{p+1} \right), \quad \text{and} \quad (65a)$$

$$J_2 \leq c r^{2p-1} \left(\sqrt{\frac{d + \log \frac{1}{\delta}}{n}} + \frac{1}{n^{3/2}} \left(d + \log \frac{n}{\delta} \right)^{2p+1} \right), \quad (65b)$$

with probability at least $1 - \delta$.

Assume that the above claims are given at the moment. We proceed to finish the proof of the uniform perturbation bound between ∇F_n^I and ∇F^I in (28b). In fact, plugging the concentration bounds (65a) and (65b) into (64), we obtain that

$$\begin{aligned} & \sup_{\theta \in \mathbb{B}(\theta^*, r)} \|\nabla F_n^I(\theta) - \nabla F^I(\theta)\|_2 \\ & \leq c (r^{p-1} + r^{2p-1}) \sqrt{\frac{d + \log(1/\delta)}{n}} \\ & \quad + \frac{r^{p-1}(d + \log(1/\delta) + \log n)^{p+1} + r^{2p-1}(d + \log(1/\delta) + \log n)^{2p+1}}{n^{\frac{3}{2}}}, \end{aligned}$$

for any $r > 0$ with probability at least $1 - 2\delta$ where c is a universal constant. When $n \geq c'(d + \log(d/\delta))^{2p}$ for some universal constant c' , it is clear that the the second term is dominated by the first term in the RHS of the above inequality. As a consequence, we have proved the claim (28b).

D.2.2 Proof of claim (65a)

Following some algebra, we find that

$$\begin{aligned}
& \sup_{\theta \in \mathbb{B}(\theta^*, r)} \left\| \frac{1}{n} \sum_{i=1}^n Y_i X_i (X_i^\top \theta)^{p-1} \right\|_2 \\
& \sup_{r>0} \frac{\sup_{\theta \in \mathbb{B}(\theta^*, r)} \left\| \frac{1}{n} \sum_{i=1}^n Y_i X_i \left(X_i^\top \frac{\theta}{\|\theta\|_2} \right)^{p-1} \right\|_2}{r^{p-1}} \\
& \leq \sup_{r>0} \sup_{\theta \in \mathbb{B}(\theta^*, r)} \left\| \frac{1}{n} \sum_{i=1}^n Y_i X_i \left(X_i^\top \frac{\theta}{\|\theta\|_2} \right)^{p-1} \right\|_2 \\
& = \underbrace{\sup_{\theta \in \mathbb{S}^{d-1}} \left\| \frac{1}{n} \sum_{i=1}^n Y_i X_i (X_i^\top \theta)^{p-1} \right\|_2}_{=: Z}. \tag{66}
\end{aligned}$$

Thus, in order to establish the claim (65a), it suffices to show that there is a universal constant c such that

$$\mathbb{P} \left(Z \leq c \sqrt{\frac{d + \log(1/\delta)}{n}} + \frac{1}{n^{3/2}} (d + \log \frac{n}{\delta})^{p+1} \right) \geq 1 - \delta. \tag{67}$$

By the variational definition of the Euclidean norm, we have

$$\begin{aligned}
Z &= \sup_{\theta \in \mathbb{S}^{d-1}} \left\| \frac{1}{n} \sum_{i=1}^n Y_i X_i (X_i^\top \theta)^{p-1} \right\|_2 \\
&= \sup_{u \in \mathbb{S}^{d-1}} \sup_{\theta \in \mathbb{S}^{d-1}} \underbrace{\left| \frac{1}{n} \sum_{i=1}^n Y_i X_i^\top u (X_i^\top \theta)^{p-1} \right|}_{:= Z_u}.
\end{aligned}$$

Using a discretization argument as in appendix D.1.2, we find that

$$Z \leq 2 \sup_{u \in \mathcal{N}(\frac{1}{8}, \mathbb{S}^{d-1}, \|\cdot\|_2)} Z_u,$$

where $\mathcal{N}(\frac{1}{8}, \mathbb{S}^{d-1}, \|\cdot\|_2)$ is the $\frac{1}{8}$ -covering of \mathbb{S}^{d-1} under $\|\cdot\|_2$ norm. Therefore, it is sufficient to bound Z_u for any fixed $u \in \mathcal{N}(\frac{1}{8}, \mathbb{S}^{d-1}, \|\cdot\|_2)$.

For any even integer $q \geq 2$, a symmetrization argument (e.g., Theorem 4.10, [55]) yields

$$\begin{aligned}
& \mathbb{E} \left(\sup_{\theta \in \mathbb{S}^{d-1}} \left| \frac{1}{n} \sum_{i=1}^n Y_i X_i^\top u (X_i^\top \theta)^{p-1} \right|^q \right)^{1/q} \\
& \leq \mathbb{E} \left(\sup_{\theta \in \mathbb{S}^{d-1}} \left| \frac{2}{n} \sum_{i=1}^n \varepsilon_i Y_i X_i^\top u (X_i^\top \theta)^{p-1} \right|^q \right)^{1/q},
\end{aligned}$$

where $\{\varepsilon_i\}_{i=1}^n$ is an i.i.d. sequence of Rademacher variables. In order to facilitate the proof argument, for any $t > 0$, we introduce the shorthand $\mathcal{N}(t) := \mathcal{N}(t, \mathbb{S}^{d-1}, \|\cdot\|_2) = \{\theta_1, \dots, \theta_{\bar{N}(t)}\}$ where $\bar{N}(t) = |\mathcal{N}(t, \mathbb{S}^{d-1}, \|\cdot\|_2)|$. For any compact set $\Omega \subseteq \mathbb{R}^d$, we define the following random variable:

$$\mathcal{R}(\Omega) := \sup_{\theta \in \Omega, p' \in [1, p]} \left| \frac{2}{n} \sum_{i=1}^n \varepsilon_i Y_i X_i^\top u (X_i^\top \theta)^{p'-1} \right|.$$

By the definition of t -covering, we obtain that

$$\begin{aligned}
\mathcal{R}(\mathbb{S}^{d-1}) &= \sup_{\theta \in \mathbb{S}^{d-1}, p' \in [1, p]} \left| \frac{2}{n} \sum_{i=1}^n \varepsilon_i Y_i X_i^\top u \left(X_i^\top \theta \right)^{p'-1} \right| \\
&\leq \sup_{\theta_k \in \mathcal{N}(t), \|\eta\|_2 \leq t, p' \in [1, p]} \left| \frac{2}{n} \sum_{i=1}^n \varepsilon_i Y_i X_i^\top u \left(X_i^\top (\theta_k + \eta) \right)^{p'-1} \right| \\
&\leq \sup_{\theta_k \in \mathcal{N}(t), p' \in [1, p]} \left| \frac{4}{n} \sum_{i=1}^n \varepsilon_i Y_i X_i^\top u \left(X_i^\top \theta \right)^{p'-1} \right| \\
&\quad + \max_{p' \in [1, p]} \sum_{b=1}^{p'-1} \binom{p'-1}{b} \cdot \sup_{\|\eta\|_2 \leq t} \left| \frac{4}{n} \sum_{i=1}^n \varepsilon_i Y_i \langle X_i, u \rangle \langle X_i, \eta \rangle^b \right| \\
&\leq \mathcal{R}(\mathcal{N}(t)) + 2^{p+1} t \cdot \mathcal{R}(\mathbb{S}^{d-1}).
\end{aligned} \tag{68}$$

By choosing $t = 2^{-(p+2)}$, the above inequality leads to

$$\mathcal{R}(\mathbb{S}^{d-1}) \leq 2\mathcal{R}(\mathcal{N}(2^{-(p+2)})).$$

In order to obtain a high-probability upper bound on $\mathcal{R}(\mathcal{N}(2^{-(p+2)}))$, we bound its moments. By the union bound, for any $q \geq 1$, we have

$$\begin{aligned}
\mathbb{E} \left[\mathcal{R}^q \left(\mathcal{N}(2^{-(p+2)}) \right) \right] &\leq p \cdot \left| \mathcal{N}(2^{-(p+2)}) \right| \\
&\times \sup_{\theta \in \mathbb{S}^{d-1}, p' \in [1, p]} \underbrace{\mathbb{E} \left[\left(\left| \frac{4}{n} \sum_{i=1}^n \varepsilon_i Y_i X_i^\top u \left(X_i^\top \theta \right)^{p'-1} \right| \right)^q \right]}_{:= T_1(\theta, p')}.
\end{aligned}$$

In order to upper bound $T_1(\theta, p')$, we apply Khintchine's inequality [6]; it guarantees that there is a universal constant C such that

$$T_1(\theta, p') \leq \mathbb{E} \left[\left(\frac{Cq}{n^2} \sum_{i=1}^n Y_i^2 (X_i^\top u)^2 (X_i^\top \theta)^{2(p'-1)} \right)^{\frac{q}{2}} \right], \tag{69a}$$

for any $p' \in [1, p]$. In order to further upper bound the right hand side, we define the function $g_{\theta, u}(x, y) := y^2 (x^\top u)^2 (x^\top \theta)^{2(p'-1)}$. For any $i \in [n]$, we can verify that

$$\begin{aligned}
\mathbb{E} [g_{\theta, u}(X_i, Y_i)] &= \mathbb{E} \left[Y_i^2 \cdot \mathbb{E} \left((X_i^\top u)^2 (X_i^\top \theta)^{2(p'-1)} \right) \right] \leq (2p')^{p'}, \\
\mathbb{E} [g_{\theta, u}(X_i, Y_i)^q] &= \mathbb{E} \left[Y_i^{2q} \cdot \mathbb{E} \left((X_i^\top u)^{2q} (X_i^\top \theta)^{2(p'-1)q} \right) \right] \leq (2q)^q (2p'q)^{p'q}.
\end{aligned}$$

Given the above bounds, invoking the result of lemma 3 leads to the following probability bound

$$\begin{aligned}
\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n g_{\theta, u}(X_i, Y_i) - \mathbb{E}_{(X, Y)} [g_{\theta, u}(X, Y)] \right| \right. \\
\left. > (8p')^{p'} \sqrt{\frac{\log 4/\delta}{n}} + \frac{1}{n} \left(2p' \log \frac{n}{\delta} \right)^{p+1} \right) \leq \delta,
\end{aligned}$$

for all $\delta \in (0, 1)$. Here the outer expectation in the above display is taken with respect to (X, Y) such that $X \sim \mathcal{N}(0, I_d)$ and $Y | X = x \sim \mathcal{N}((x^\top \theta^*)^p, 1)$. Combining the previous bounds yields

$$\begin{aligned}
& \mathbb{E} \left[\left(\frac{1}{n} \sum_{i=1}^n g_{\theta, u}(X_i, Y_i) \right)^{q/2} \right] \\
& \leq 2^{q/2} (\mathbb{E}_{(X, Y)} [g_{\theta, u}(X, Y)])^{q/2} \\
& \quad + 2^{q/2} \mathbb{E} \left[\left| \frac{1}{n} \sum_{i=1}^n g_{\theta, u}(X_i, Y_i) - \mathbb{E}_{(X, Y)} [g_{\theta, u}(X, Y)] \right|^{q/2} \right] \\
& \leq (4p')^{pq} + q \int_0^{+\infty} \lambda^{q-1} \mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n g_{\theta, u}(X_i, Y_i) - \mathbb{E}_{(X, Y)} [g_{\theta, u}(X, Y)] \right| > \lambda \right) d\lambda \\
& \leq (4p')^{p'q} + q \int_0^1 (p' + 1) \left((8p')^{p'} \sqrt{\frac{\log 4/\delta}{n}} + \frac{1}{n} \left(2p' \log \frac{n}{\delta} \right)^{p'+1} \right)^q \log^{-1} \frac{4}{\delta} d\delta \\
& \leq (4p')^{p'q} + Cp'q \left(\frac{(16p')^{p'q}}{n^{\frac{q}{2}}} \Gamma(q/2) \right. \\
& \quad \left. + \frac{(2p')^{(p'+1)q}}{n^q} \left((2 \log n)^{(p'+1)q} + \Gamma((p'+1)q) \right) \right), \tag{69b}
\end{aligned}$$

where Γ denotes the Gamma function. Combining the bounds (69a) and (69b), we reach to the following upper bound for $T_1(\theta, p')$:

$$\begin{aligned}
T_1(\theta, p') & \leq \left(\frac{Cq}{n} \right)^{q/2} \left[(4p')^{p'q} + Cp'q \left(\frac{(16p')^{pq}}{n^{\frac{q}{2}}} \Gamma(q/2) \right. \right. \\
& \quad \left. \left. + \frac{(2p')^{(p'+1)q}}{n^q} \left((2 \log n)^{(p'+1)q} + \Gamma((p'+1)q) \right) \right) \right]. \tag{70}
\end{aligned}$$

Plugging the upper bounds of T_1 in equation (70) into equation (68) and taking the union bound over all $\theta_k \in \mathcal{N}(2^{-(p+2)}, \mathbb{S}^{d-1}, \|\cdot\|_2)$, we find that

$$\begin{aligned}
\mathbb{E} [\mathcal{R}^q(\mathbb{S}^{d-1})] & \leq 2^q \mathbb{E} [\mathcal{R}^q(\mathcal{N}(2^{-(p+2)}))] \\
& \leq 2^q p (2^{p+3})^d \sup_{\theta \in \mathbb{S}^{d-1}, p' \in [1, p]} T_1(\theta, p') \\
& \leq 2^q p (2^{p+3})^d \left(\frac{Cq}{n} \right)^{\frac{q}{2}} \left[(4p)^{pq} + Cpq \left(\frac{(16p)^{pq}}{n^{\frac{q}{2}}} \Gamma(q/2) \right. \right. \\
& \quad \left. \left. + \frac{(2p)^{(p+1)q}}{n^q} \left((2 \log n)^{(p+1)q} + \Gamma((p+1)q) \right) \right) \right],
\end{aligned}$$

for any given $u \in \mathcal{N}(\frac{1}{8}, \mathbb{S}^{d-1}, \|\cdot\|_2)$.

Taking the supremum over $u \in \mathcal{N}(\frac{1}{8}, \mathbb{S}^{d-1}, \|\cdot\|_2)$ of both sides in the above bound and applying

Minkowski's inequality, we obtain that

$$\begin{aligned} (\mathbb{E}|Z|^q)^{\frac{1}{q}} &\leq \left(\frac{64}{7}\right)^{d/q} \left(\mathbb{E} \left[\sup_{\theta \in \mathbb{S}^{d-1}} \left| \frac{2}{n} \sum_{i=1}^n \varepsilon_i Y_i X_i^\top u (X_i^\top \theta)^{p-1} \right|^q \right] \right)^{\frac{1}{q}} \\ &\leq 2 (10 \cdot 2^{p+3})^{d/q} \left[\sqrt{\frac{C_p q}{n}} + \frac{C_p q}{n} + \frac{C_p}{n^{3/2}} (\log n + q)^{p+1} \right], \end{aligned}$$

where C_p is a universal constant depending only on p . By choosing $q = d(p+7) + \log \frac{2}{\delta}$ and using Markov inequality, we find that

$$\mathbb{P} \left(|Z| \geq C_p \left(\sqrt{\frac{d + \log \frac{1}{\delta}}{n}} + \frac{1}{n^{3/2}} (d + \log \frac{n}{\delta})^{p+1} \right) \right) \leq \delta.$$

Thus, we have established the claim (65a).

Proof of claim (65b): In order to obtain a uniform concentration bound for J_2 , we use an argument similar to that from the proof of claim (65a). In particular, since polynomial $(x^\top \theta)^{2p-1}$ is homogeneous in terms of θ , using the same normalization as in equation (66), it suffices to demonstrate that

$$\mathbb{P} \left(W \leq c r^{2p-1} \left(\sqrt{\frac{d + \log \frac{1}{\delta}}{n}} + \frac{1}{n^{3/2}} (d + \log \frac{n}{\delta})^{2p+1} \right) \right) \geq 1 - \delta, \quad (71)$$

for any $\delta > 0$ where we define

$$W := \sup_{\theta \in \mathbb{S}^{d-1}} \left\| \frac{1}{n} \sum_{i=1}^n X_i (X_i^\top \theta)^{2p-1} - \mathbb{E}_X \left[X (X^\top \theta)^{2p-1} \right] \right\|_2.$$

For each $u \in \mathbb{R}^d$, define the random variable

$$W_u := \sup_{\theta \in \mathbb{S}^{d-1}} \left| \frac{1}{n} \sum_{i=1}^n X_i^\top u (X_i^\top \theta)^{2p-1} - \mathbb{E}_X \left[X^\top u (X^\top \theta)^{2p-1} \right] \right|.$$

It suffices to bound W_u for fixed $u \in \mathcal{N}(\frac{1}{8}, \mathbb{S}^{d-1}, \|\cdot\|_2)$. We bound W_u by controlling its moments. By a symmetrization argument, we have

$$\begin{aligned} \mathbb{E} \left[\sup_{\theta \in \mathbb{S}^{d-1}} \left| \frac{1}{n} \sum_{i=1}^n X_i^\top u (X_i^\top \theta)^{2p-1} - \mathbb{E}_X \left[X^\top u (X^\top \theta)^{2p-1} \right] \right|^q \right] \\ \leq \mathbb{E} \left[\sup_{\theta \in \mathbb{S}^{d-1}} \left| \frac{2}{n} \sum_{i=1}^n \varepsilon_i X_i^\top u (X_i^\top \theta)^{2p-1} \right|^q \right]. \end{aligned}$$

From here, we can use the same technique as that in and after inequality (68) to bound the RHS term in the above display. Therefore, we will only highlight the main differences here. For any compact set $\Omega \subseteq \mathbb{R}^d$, we define the random variable

$$\mathcal{Q}(\Omega) := \sup_{\theta \in \Omega, p' \in [1, p]} \left| \frac{2}{n} \sum_{i=1}^n \varepsilon_i X_i^\top u (X_i^\top \theta)^{2p'-1} \right|.$$

Following the similar argument as that in equation (68), we can check that $\mathcal{Q}(\mathbb{S}^{d-1}) \leq 2\mathcal{Q}(\mathcal{N}(2^{-(2p+2)}))$. A direct application of union bound leads to

$$\begin{aligned} \mathbb{E} \left[\mathcal{Q}^q \left(\mathcal{N}(2^{-(2p+2)}) \right) \right] &\leq 2p \cdot \left| \mathcal{N}(2^{-(2p+2)}) \right| \\ &\times \sup_{\theta \in \mathbb{S}^{d-1}, p' \in [1, p]} \underbrace{\mathbb{E} \left[\left(\left| \frac{4}{n} \sum_{i=1}^n \varepsilon_i X_i^\top u (X_i^\top \theta)^{2p'-1} \right| \right)^q \right]}_{:= T_2(\theta, p')} . \end{aligned}$$

We control $T_2(\theta, p')$ using the same approach as that the proof of claim (65a). For the convenience of notation, we denote $h_{\theta, u}(x) := (x^\top u)^2 (x^\theta)^{2(2p'-1)}$. Simple algebra lead to the following upper bounds:

$$\mathbb{E} [h_{\theta, u}(X_i)] \leq (4p')^{2p'}, \quad \mathbb{E} [h_{\theta, u}(X_i)^q] \leq (4p'q)^{2p'q}.$$

Invoking the result of lemma 3, the above bounds lead to the following probability bound:

$$\begin{aligned} \mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n h_{\theta, u}(X_i) - \mathbb{E}_X [h_{\theta, u}(X)] \right| \right. \\ \left. \leq (16p')^{2p'} \sqrt{\frac{\log 4/\delta}{n}} + \left(4p' \log \frac{n}{\delta} \right)^{2p'} \frac{\log 4/\delta}{n} \right) \leq \delta. \end{aligned}$$

Therefore, we further obtain that

$$\begin{aligned} \mathbb{E} \left[\left(\frac{1}{n} \sum_{i=1}^n h_{\theta, u}(X_i) \right)^{q/2} \right] &\leq (8p')^{2p'q} + Cp'q \left(\frac{(32p')^{2p'q}}{n^{\frac{q}{2}}} \Gamma(q/2) \right. \\ &\quad \left. + \frac{(4p')^{(2p'+1)q}}{n^q} \left((2 \log n)^{(2p'+1)q} + \Gamma((2p'+1)q) \right) \right). \end{aligned}$$

Combining the above bound and an upper bound of $T_2(\theta, p')$ based on Khintchine's inequality, we obtain the following inequality:

$$\begin{aligned} T_2(\theta, p') &\leq \left(\frac{Cq}{n} \right)^{q/2} \left[(8p')^{2p'q} + Cp'q \left(\frac{(32p')^{2p'q}}{n^{\frac{q}{2}}} \Gamma(q/2) \right. \right. \\ &\quad \left. \left. + \frac{(4p')^{(2p'+1)q}}{n^q} \left((2 \log n)^{(2p'+1)q} + \Gamma((2p'+1)q) \right) \right) \right]. \end{aligned}$$

Collecting the above bounds leads to

$$\begin{aligned} \mathbb{E} \left[\mathcal{Q}^q(\mathbb{S}^{d-1}) \right] &\leq 2^{q+1} p (2^{2p+3})^d \sup_{\theta \in \mathbb{S}^{d-1}, p' \in [1, p]} T_2(\theta, p') \\ &\leq 2^{q+1} p (2^{2p+3})^d \left(\frac{Cq}{n} \right)^{\frac{q}{2}} \left[(8p)^{2pq} + Cpq \left(\frac{(32p)^{2pq}}{n^{\frac{q}{2}}} \Gamma(q/2) \right. \right. \\ &\quad \left. \left. + \frac{(4p)^{(2p+1)q}}{n^q} \left((2 \log n)^{(2p+1)q} + \Gamma((2p+1)q) \right) \right) \right], \end{aligned}$$

for any fixed $u \in \mathcal{N}\left(\frac{1}{8}, \mathbb{S}^{d-1}, \|\cdot\|_2\right)$. Taking supremum over $u \in \mathcal{N}\left(\frac{1}{8}, \mathbb{S}^{d-1}, \|\cdot\|_2\right)$ of both sides in the above bound and applying Minkowski's inequality, we arrive at the following bound:

$$\begin{aligned} (\mathbb{E} [|W|^q])^{\frac{1}{q}} &\leq \left(\frac{64}{7}\right)^{\frac{d}{q}} \left(\mathbb{E} \left[\sup_{\theta \in \mathbb{S}^{d-1}} \left| \frac{2}{n} \sum_{i=1}^n \sigma_i X_i^\top u(X_i^\top \theta)^{2p-1} \right|^q \right] \right)^{\frac{1}{q}} \\ &\leq \left(\frac{10}{\varepsilon}\right)^{\frac{d}{q}} \left[\sqrt{\frac{C_p q}{n}} + \frac{C_p q}{n} + \frac{C_p}{n^{\frac{3}{2}}} (\log n + q)^{2p+1} \right], \end{aligned}$$

where C_p is a universal constant depending only upon p . With the choice of $q = d(2p+7) + \log \frac{2}{\delta}$, we obtain that

$$\mathbb{P} \left(|W| \geq C_p \left(\sqrt{\frac{d + \log \frac{1}{\delta}}{n}} + \frac{1}{n^{3/2}} \left(d + \log \frac{n}{\delta} \right)^{2p+1} \right) \right) \leq \delta.$$

Thus, we have established the claim (65b).

D.3 Proof of corollary 3

We prove corollary 3 by verifying the claims (22a) and (22b).

D.3.1 Structure of F^G

Direct algebra leads to the following equation

$$\begin{aligned} \langle \nabla F^G(\theta), \theta^* - \theta \rangle &= \left(\theta - \mathbb{E} \left[X \tanh(X^\top \theta) \right] \right)^\top (\theta - \theta^*) \\ &\geq \|\theta\|_2^2 - \|\theta\|_2 \left\| \mathbb{E} \left[X \tanh(X^\top \theta) \right] \right\|_2 \end{aligned} \tag{72}$$

where $\tanh(x) := \frac{\exp(x) - \exp(-x)}{\exp(x) + \exp(-x)}$ for all $x \in \mathbb{R}$. From Theorem 2 in Dwivedi et al. [15], we have

$$\left\| \mathbb{E} \left[X \tanh(X^\top \theta) \right] \right\|_2 \leq \left(1 - p + \frac{p}{1 + \frac{\|\theta\|_2^2}{2}} \right) \|\theta\|_2$$

for all $\theta \in \mathbb{R}^d$ where $p := \mathbb{P}(|Y| \leq 1) + \frac{1}{2}\mathbb{P}(|Y| > 1)$ where $Y \sim \mathcal{N}(0, 1)$. Plugging the above inequality into equation (72) leads to

$$\langle \nabla F^G(\theta), \theta^* - \theta \rangle \geq \frac{p \|\theta\|_2^4}{2 + \|\theta\|_2^2} \geq \begin{cases} \frac{p}{4} \|\theta\|_2^4, & \text{for } \|\theta\|_2 \leq \sqrt{2} \\ \frac{p}{2} \left(\|\theta\|_2^2 - 1 \right), & \text{otherwise} \end{cases}.$$

As a consequence, we achieve the conclusion of claim (22a).

D.3.2 Perturbation error between ∇F^G and ∇F_n^G

Direct calculation indicates the following equation:

$$\nabla F_n^G(\theta) - \nabla F^G(\theta) = \frac{1}{n} \sum_{i=1}^n X_i \tanh(X_i^\top \theta) - \mathbb{E} \left[X \tanh(X^\top \theta) \right].$$

The outer expectation in the above display is taken with respect to $X \sim \mathcal{N}(\theta^*, \sigma^2 I_d)$ where $\theta^* = 0$. Based on the proof argument of Lemma 1 from the paper [15], for each $r > 0$, we have the following concentration inequality

$$\begin{aligned} \mathbb{P} \left(\sup_{\theta \in \mathbb{B}(\theta^*, r)} \left\| \frac{1}{n} \sum_{i=1}^n X_i \tanh(X_i^\top \theta) - \mathbb{E} \left[X \tanh(X^\top \theta) \right] \right\|_2 \right. \\ \left. \leq c r \sqrt{\frac{d + \log(1/\delta)}{n}} \right) \geq 1 - \delta, \end{aligned} \quad (73)$$

for any $\delta > 0$ as long as the sample size $n \geq c'd \log(1/\delta)$ where c and c' are universal constants. For any $M \in \mathbb{N}_+$, by the concentration bound (73) and the union bound, we find that

$$\begin{aligned} \mathbb{P} \left(\forall r \in [2^{-M}, 1], \sup_{\theta \in \mathbb{B}(\theta^*, r)} \left\| \nabla F_n^G(\theta) - F^G(\theta) \right\|_2 \right. \\ \left. \leq c r \sqrt{\frac{d + \log(M/\delta)}{n}} \right) \geq 1 - \delta. \end{aligned} \quad (74)$$

On the other hand, based on the standard inequality $|\tanh(x)| \leq |x|$ for all $x \in \mathbb{R}$, we find that

$$\begin{aligned} \left\| \nabla F_n^G(\theta) - \nabla F^G(\theta) \right\|_2 &\leq \frac{1}{n} \sum_{i=1}^n \|X_i\|_2 \left| \tanh(X_i^\top \theta) \right| + \mathbb{E} \left[\|X\|_2 \left| \tanh(X^\top \theta) \right| \right] \\ &\leq \frac{1}{n} \sum_{i=1}^n \|X_i\|_2 \left| X_i^\top \theta \right| + \mathbb{E} \left[\|X\|_2 \left| X^\top \theta \right| \right] \\ &\leq \left(\frac{1}{n} \sum_{i=1}^n \|X_i\|_2^2 + \mathbb{E} \left[\|X\|_2^2 \right] \right) \|\theta\|_2. \end{aligned}$$

Therefore, we have $\left\| \nabla F_n^G(\theta) - \nabla F^G(\theta) \right\|_2 \leq 2d \|\theta\|_2 \log(1/\delta)$ with probability $1 - \delta$. By choosing $M_1 := \log(2nd)$, based on the previous bound, we obtain that

$$\mathbb{P} \left(\forall r < 2^{-M_1}, \sup_{\theta \in \mathbb{B}(\theta^*, r)} \left\| \nabla F_n^G(\theta) - \nabla F^G(\theta) \right\|_2 \leq \frac{\log(1/\delta)}{n} \right) \geq 1 - \delta. \quad (75)$$

Furthermore, for vector $\theta \in \mathbb{R}^d$ with large norm, by the concentration bound (73) combined with the union bound, for any $M' \in \mathbb{N}_+$, we find that

$$\begin{aligned} \mathbb{P} \left(\forall r \in [1, 2^{M'}], \sup_{\theta \in \mathbb{B}(\theta^*, r)} \left\| \nabla F_n^G(\theta) - F^G(\theta) \right\|_2 \right. \\ \left. \leq c r \sqrt{\frac{d + \log(M'/\delta)}{n}} \right) \geq 1 - \delta. \end{aligned}$$

When r in the above bound is too large, we can simply use the fact that \tanh is a bounded function. We thus have the upper bound

$$\|\nabla F_n^G(\theta) - \nabla F^G(\theta)\|_2 \leq \mathbb{E} [\|X\|_2] + \frac{1}{n} \sum_{i=1}^n \|X_i\|_2,$$

for any θ . Given the above bound, by choosing $M_2 := \log(2\sqrt{n})$, we obtain that

$$\begin{aligned} \mathbb{P} \left(\forall r > 2^{M_2}, \sup_{\theta \in \mathbb{B}(\theta^*, r)} \|\nabla F_n^G(\theta) - \nabla F^G(\theta)\|_2 \leq r \sqrt{\frac{d + \log(1/\delta)}{n}} \right) \\ \geq \mathbb{P} \left(\mathbb{E} [\|X\|_2] + \frac{1}{n} \sum_{i=1}^n \|X_i\|_2 \leq 2^{M_2} \sqrt{\frac{d + \log(1/\delta)}{n}} \right) \geq 1 - \delta. \end{aligned} \quad (76)$$

Putting the bounds (74), (75), and (76) together, for $n \geq cd\log(1/\delta)$, the following probability bound holds

$$\begin{aligned} \mathbb{P} \left(\forall r > 0, \sup_{\theta \in \mathbb{B}(\theta^*, r)} \|\nabla F_n^G(\theta) - \nabla F^G(\theta)\|_2 \right. \\ \left. \leq c r \sqrt{\frac{d + \log(\log n/\delta)}{n}} + \frac{\log(1/\delta)}{n} \right) \geq 1 - \delta, \end{aligned}$$

which completes the proof of the claim (22b).

D.4 Proof of corollary 5

We prove this claim by verifying the conditions in corollary 7 and theorem 5. In particular, we claim the following bounds on the population log-likelihood $F(\theta) = \mathbb{E} [\log p_\theta(X)]$ and its empirical counterpart $F_n(\theta) = \frac{1}{n} \sum_{i=1}^n \log p_\theta(X_i)$: For each permutation function σ and $R_0 > 0$, we have

$$\|\nabla^2 F(\theta) - \nabla^2 F(\theta_\sigma^*)\|_{\text{op}} \leq cK \left(\|\theta - \theta_\sigma^*\|_2^3 + \sigma_X^3 \right) \|\theta - \theta_\sigma^*\|_2 \quad \forall \theta \in \mathbb{R}^{dK}, \quad (77a)$$

$$\sup_{\theta \in \mathbb{B}(\theta_\sigma^*, R_0)} \|\nabla F_n(\theta) - \nabla F(\theta)\|_2 \leq cK(\sigma_X + R_0) \sqrt{\frac{Kd \log(Kd) + \log \delta^{-1}}{n}}, \quad \text{w.p. } 1 - \delta, \quad (77b)$$

$$\sup_{\theta \in \mathbb{B}(\theta_\sigma^*, R_0)} \|\nabla^2 F_n(\theta) - \nabla^2 F(\theta)\|_{\text{op}} \leq cK^2(\sigma_X^2 + R_0^2) \sqrt{\frac{Kd \log(Kd) + \log(1/\delta)}{n}} \quad \text{w.p. } 1 - \delta, \quad (77c)$$

$$\Pi \left(\mathbb{B}^c \left(0, 3\sqrt{K}R_X + \sqrt{K(d \log(R_X n + L_2) + \log(\pi_0(1/\vartheta)))} \mid X_1^n \right) < \vartheta, \right) \quad (77d)$$

where $R_X := \max_i \|X_i\|_2$. The proofs of these bounds are deferred to later subsections. Taking the four bounds as given, we now proceed with the proof of the corollary.

First, we define $r_0 := \frac{\mu}{4c'K(\sigma_X^3 + 1)} \wedge \left(\frac{\mu}{4c'K} \right)^{1/4} \wedge 1$, for any permutation function σ and any $\theta \in \mathbb{B}(\theta_\sigma^*, r_0)$. Then, equation (77a) guarantees the following local bound:

$$-\nabla^2 F(\theta) \succeq -\nabla^2 F(\theta_\sigma^*) - \frac{\mu}{2} I_d \succeq \frac{\mu}{2} I_d,$$

which implies the condition $-\langle \nabla F(\theta), \theta - \theta_\sigma^* \rangle \geq \frac{\mu}{2} \|\theta - \theta_\sigma^*\|_2^2$ inside the ball $\mathbb{B}(\theta_\sigma^*, r_0)$.

Combining this bound with equation (77b) by taking $R_0 = r_0$, we invoke theorem 3 and obtain the following localized posterior contraction bound with probability $1 - \delta$:

$$\Pi(\mathbb{B}(\theta_\sigma^*, r_0) \mid X_1^n)^{-1} \Pi(\mathbb{B}(\theta_\sigma^*, r_n) \mid X_1^n) \geq 1 - \vartheta, \quad (78)$$

where the contraction radius r_n is given by:

$$r_n := \frac{cK\sigma_X}{\mu} \sqrt{\frac{Kd \log(Kd) + \log \delta^{-1}}{n}} + c \sqrt{\frac{\log \vartheta^{-1}}{\mu n}}.$$

By the sub-Gaussian condition, we obtain that

$$\max_i \|X_i\|_2 \leq 2\sigma_X \sqrt{d + \log n + \log \delta^{-1}}, \quad \text{with probability } 1 - \delta.$$

Therefore, the following tail bound holds true with probability $1 - \delta$:

$$\Pi\left(\mathbb{B}^c\left(0, c\sigma_X \sqrt{Kd \log \frac{n}{\vartheta \delta \pi_0}}\right) \mid X_1^n\right) < \vartheta. \quad (79)$$

Taking $R_0 = c\sigma_X \sqrt{Kd \log \frac{n}{\vartheta \delta \pi_0}}$ and applying equation (77b), we conclude that there exists a quantity $a_0 > 0$ depending on the constants K, d, σ_X , such that

$$\sup_{\theta \in \mathbb{B}(0, R_0)} \|\nabla F_n(\theta) - \nabla F_n(\theta^*)\|_2 \leq \frac{a_0(1 + \log \delta^{-1} + \log \vartheta^{-1})}{\sqrt{n}}.$$

Note that, $F_n(0) = -\frac{1}{n} \sum_{i=1}^n \|X_i\|_2^2$. An application of sub-exponential concentration bounds [55] leads to

$$\mathbb{P}\left(|F_n(0) - F(0)| > \sigma_X^2 \frac{cd \log \delta^{-1}}{\sqrt{n}}\right) < \delta.$$

Combining the previous two bounds, there exists $a_1 > 0$, such that the following bound holds true with probability $1 - \delta$:

$$\sup_{\theta \in \mathbb{B}(0, R_0)} |F_n(\theta) - F_n(\theta^*)| \leq \frac{a_1(1 + \log \delta^{-1} + \log \vartheta^{-1})}{\sqrt{n}}.$$

On the other hand, since the Gaussian mixture model is identifiable up to permutations, there exists $\Delta_0 > 0$ depending on θ^* , such that:

$$\inf_{\theta \in (\bigcup_{\sigma: [K] \rightarrow [K]} \mathbb{B}(\theta_\sigma^*, r_0))^c} F(\theta_{\text{Id}}^*) - F(\theta) \geq \Delta_0.$$

Consequently, for $n \geq \left(\frac{3a_0}{\Delta_0} (\log \delta^{-1} + \log \vartheta^{-1})\right)^2$, with probability $1 - \delta$ we have that

$$F_n(\theta') \leq F_n(\theta) - \frac{\Delta_0}{3}$$

for all $\sigma : [K] \rightarrow [K]$, $\theta \in \mathbb{B}\left(\theta_\sigma^*, \sqrt{\frac{\Delta_0}{d+\log \delta^{-1}}}\right)$ and $\theta' \in \mathbb{B}(0, R_0) \setminus \bigcup_{\sigma' : [K] \rightarrow [K]} \mathbb{B}(\theta_{\sigma'}^*, r_0)$. Thus, we have the posterior probability bound:

$$\Pi \left[\mathbb{B}(0, R_0) \setminus \bigcup_{\sigma' : [K] \rightarrow [K]} \mathbb{B}(\theta_{\sigma'}^*, r_0) \mid X_1^n \right] \leq \exp \left(-\frac{\Delta_0}{3} n \right) \cdot \left(\frac{R_0 \sqrt{d + \log \delta^{-1}}}{\sqrt{\Delta_0}} \right)^{dK}.$$

It indicates that there exists $a_2 > 0$ depending on the problem instances θ^* , K , d , such that for $n \geq \frac{a_2}{\Delta_0} \log \vartheta^{-1}$, we have the following bound:

$$\Pi \left[\mathbb{B}(0, R_0) \setminus \bigcup_{\sigma' : [K] \rightarrow [K]} \mathbb{B}(\theta_{\sigma'}^*, r_0) \mid X_1^n \right] \leq \vartheta. \quad (80)$$

Collecting the bounds (78), (79) and (80), we conclude that for $n \geq n_{\min} \cdot \log^2 \frac{1}{\delta \vartheta}$, the following bound holds true with probability $1 - \delta$:

$$\Pi \left[\bigcup_{\sigma : [K] \rightarrow [K]} \mathbb{B}(\theta_\sigma^*, r_n) \mid X_1^n \right] \geq 1 - \vartheta,$$

for contraction radius r_n defined as:

$$r_n := \frac{cK\sigma_X}{\mu} \sqrt{\frac{Kd \log(Kd) + \log \delta^{-1}}{n}} + c \sqrt{\frac{\log \vartheta^{-1}}{\mu n}},$$

which proves the bound (31a).

Furthermore, applying theorem 5 to each local neighborhood $\mathbb{B}(\theta_\sigma^*, r_0)$, for any $\omega \in (0, 1)$, we obtain the following bound with probability $1 - \delta$:

$$\begin{aligned} \Pi \left[\left\| \theta - \hat{\theta}_\sigma^{(n)} \right\|_{H_\sigma^*}^2 \leq (1 + \omega) \frac{d}{n} + c \frac{1 + \log \kappa(H_\sigma^*)}{\omega} \left(\frac{\log \vartheta^{-1}}{n} + \frac{a'(\log \vartheta^{-1} + \log \delta^{-1})^2}{n^2} \right) \mid X_1^n \right] \\ \geq (1 - \vartheta) \Pi(\mathbb{B}(\theta_\sigma^*, r_0) \mid X_1^n), \end{aligned}$$

for a constant $a' > 0$ depending on K, d and θ^* .

Combining with the tail bounds (79) and (80), we obtain the result (31b).

D.4.1 Proof of the claim (77a)

We first verify the local conditions **(LWC.1)** and **(LWC.2)**. Given the parameters (u_1, u_2, \dots, u_K) , direct calculation yields

$$-\nabla_\theta F_n(\theta) = \left[\frac{1}{n} \sum_{i=1}^n (u_j - X_i) \frac{\exp(-\|u_j - X_i\|_2^2/2)}{\sum_{\ell=1}^K \exp(-\|u_\ell - X_i\|_2^2/2)} \right]_{j \in [K]}.$$

Given distinct centers $(u_j)_{j \in [K]}$ of each mixture component, we have that $H_\sigma^* \succ 0$ for any permutation σ . To show the local growth condition **(LWC.1)**, we study the local conditions around θ_σ^* . Denote $q_j(x; \theta) := \frac{\exp(-\|u_j - x\|_2^2/2)}{\sum_{\ell=1}^K \exp(-\|u_\ell - x\|_2^2/2)}$ for any $x \in \mathbb{R}^d$ and $j \in [K]$. Direct calculation shows that

$$\begin{aligned} -\nabla_\theta^2 \log p_\theta(X) &= \text{diag} \left((I_d + (u_j - X)(u_j - X)^\top) q_j(X; \theta) \right)_{j \in [K]} \\ &\quad - \left[(u_j - X)(u_\ell - X)^\top q_j(X; \theta) q_\ell(X; \theta) \right]_{j, \ell \in [K]}. \end{aligned}$$

For the third-order derivative, for any vector $v = [v_1 \ v_2 \ \cdots \ v_K] \in \mathbb{S}^{Kd-1}$, direct calculation leads to

$$\begin{aligned} &\|\nabla^3 F(\theta)[v]\|_{\text{op}} \\ &\leq \left\| \left[\mathbb{E} \left[\nabla_{u_\ell} q_j(X; \theta) (I_d + (u_j - X)(u_j - X)^\top) v_j \right] \right]_{j, \ell \in [K]} \right\|_{\text{op}} \\ &\quad + \max_{j \in [K]} \left| \mathbb{E} \left[q_j(X; \theta) (u_j - X)^\top v_j \right] \right| + \max_{j \in [K]} \left| \mathbb{E} \left[q_j(X; \theta) \sum_{\ell \in [K]} q_\ell(X; \theta) (u_\ell - X)^\top v_\ell \right] \right| \\ &\quad + \left\| \left[\mathbb{E} \left[q_j(X; \theta) q_\ell(X; \theta) v_\ell (u_\ell - X)^\top \right] \right]_{j, \ell \in [K]} \right\|_{\text{op}} \\ &\quad \left\| \left[\left\| \sum_{\ell \in [K]} \mathbb{E} \left[q_\ell(X; \theta) (u_j - X)(u_\ell - X)^\top v_\ell \nabla_{u_k} q_j(X; \theta) \right] \right\|_{\text{op}} \right]_{j, k \in [K]} \right\|_{\text{op}} \\ &\quad + \left\| \left[\left\| \sum_{\ell \in [K]} \mathbb{E} \left[q_j(X; \theta) (u_j - X)(u_\ell - X)^\top v_\ell \nabla_{u_k} q_\ell(X; \theta) \right] \right\|_{\text{op}} \right]_{j, k \in [K]} \right\|_{\text{op}}. \end{aligned}$$

Using Hölder inequality and the variational representation of the operator norm, we obtain that

$$\begin{aligned} &\|\nabla^3 F(\theta)[v]\|_{\text{op}} \\ &\leq cK \cdot \sup_{\substack{y, z \in \mathbb{S}^{d-1} \\ j, k, \ell \in [K]}} \mathbb{E} \left[\left| (X - u_j)^\top y \cdot (X - u_k)^\top z \cdot (X - u_\ell)^\top v_\ell \right| \right] \\ &\quad + cK \cdot \sup_{\substack{y, z \in \mathbb{S}^{d-1} \\ \ell \in [K]}} \mathbb{E} \left[\left| y^\top z (X - u_\ell)^\top v_\ell \right| \right] \\ &\leq cK \cdot \sup_{\substack{y, z \in \mathbb{S}^{d-1} \\ j, k, \ell \in [K]}} \mathbb{E} \left[\left| (X - u_j)^\top y \right|^3 \right]^{1/3} \cdot \mathbb{E} \left[\left| (X - u_k)^\top z \right|^3 \right]^{1/3} \cdot \mathbb{E} \left[\left| (X - u_\ell)^\top v_\ell \right|^3 \right]^{1/3} \\ &\quad + cK \cdot \sup_{\substack{y, z \in \mathbb{S}^{d-1} \\ \ell \in [K]}} \mathbb{E} \left[\left((X - u_\ell)^\top v_\ell \right)^2 \right]^{1/2} \\ &\leq c'K \left(\|\theta - \theta_\sigma^*\|_2^3 + \sigma_X^3 + 1 \right), \end{aligned}$$

for a universal constant $c' > 0$ and any permutation function σ . This proves the desired claim.

D.4.2 Proof of the claims (77b) and (77c)

Now we turn to the empirical process bounds for the gradient and Hessian of F_n . For $\theta \in \mathbb{R}^{dK}$ and $v, w \in \mathbb{S}^{dK-1}$, we define the following quantities

$$\begin{aligned} Y_{\theta,v}^{(1)} &:= \langle \nabla F_n(\theta), v \rangle = \frac{1}{n} \sum_{j \in [K]} \sum_{i=1}^n (u_j - X_i)^\top v_j q_j(X_i, \theta), \quad \text{and} \\ Y_{\theta,v,w}^{(2)} &:= v^\top \nabla^2 F_n(\theta) w \\ &= \frac{1}{n} \sum_{j=1}^K \sum_{i=1}^n (v^\top w + v_j^\top (u_j - X_i) w_j^\top (u_j - X_i)) q_j(X_i; \theta) \\ &\quad - \frac{1}{n} \sum_{j,\ell \in [K]} \sum_{i=1}^n v_j^\top (u_j - X_i) \cdot w_\ell^\top (u_\ell - X_i) \cdot q_j(X_i, \theta) q_\ell(X_i, \theta). \end{aligned}$$

We further define $Z_{\theta,v}^{(1)} := Y_{\theta,v}^{(1)} - \mathbb{E}[Y_{\theta,v}^{(1)}]$ and $Z_{\theta,v,w}^{(2)} := Y_{\theta,v}^{(2)} - \mathbb{E}[Y_{\theta,v}^{(2)}]$.

In the following derivation, we first regard the vectors v, w as fixed, and then use standard discretization approach to take the maximum with respect to both vectors. Similar to the proof of corollary 2, we use proposition 3 to control the concentration behavior of the above quantities. Note that, q_j is a bounded function for each $j \in [K]$. Therefore, by applying proposition 3 to each term of $Z_{\theta,v}^{(1)}$ with envelop function $\bar{G}^{(1)}(X) = 1 + R_0 + |(u_j^* - X)^\top v_j|$ for each $j \in [K]$, we obtain the following bound with probability $1 - \delta$:

$$\sup_{\theta \in \mathbb{B}(\theta_\sigma^*, R_0)} Z_{\theta,v}^{(1)} \leq 2\mathbb{E} \left[\sup_{\theta \in \mathbb{B}(\theta_\sigma^*, R_0)} Z_{\theta,v}^{(1)} \right] + K(1 + R_0 + \sigma_X) \left(\sqrt{\frac{\log \delta^{-1}}{n}} + \frac{\log \delta^{-1}}{n} \sqrt{\log n} \right).$$

Similarly, by applying proposition 3 to each term of $Z_{\theta,v}^{(2)}$ with envelop function $\bar{G}^{(2)}(X) = 1 + R_0 \cdot (|v_j^\top (u_j^* - X_i)| + |w_\ell^\top (u_\ell^* - X_i)|) + R_0^2 + |v_j^\top (u_j^* - X_i) \cdot w_\ell^\top (u_\ell^* - X_i)|$, for each $j, \ell \in [K]$, we obtain the following bound with probability $1 - \delta$:

$$\sup_{\theta \in \mathbb{B}(\theta_\sigma^*, R_0)} Z_{\theta,v,w}^{(2)} \leq 2\mathbb{E} \left[\sup_{\theta \in \mathbb{B}(\theta_\sigma^*, R_0)} Z_{\theta,v,w}^{(2)} \right] + K^2(1 + R_0^2 + \sigma_X^2) \left(\sqrt{\frac{\log \delta^{-1}}{n}} + \frac{\log \delta^{-1}}{n} \log n \right).$$

Now, we consider the function classes

$$\begin{aligned} \mathcal{G}_{v,j}^{(1)} &:= \{x \mapsto \langle \nabla \log p_\theta(x), v_j \rangle : \theta \in \mathbb{B}(\theta_\sigma^*, R_0)\}, \quad \text{and} \\ \mathcal{G}_{v,w,j,\ell}^{(2)} &:= \{x \mapsto \langle \nabla^2 \log p_\theta(X_i) v_j, w_\ell \rangle : \theta \in \mathbb{B}(\theta_\sigma^*, R_0)\}. \end{aligned}$$

Apparently, $\bar{G}^{(1)}$ and $\bar{G}^{(2)}$ are envelop functions for the corresponding classes $\mathcal{G}_{v,j}^{(1)}$ and $\mathcal{G}_{v,w,j,\ell}^{(2)}$. In order to bound the expected suprema, we define the following symmetrized random variables:

$$V_{\theta,v}^{(1)} = \frac{1}{n} \sum_{i=1}^n \varepsilon_i \langle \nabla \log p_\theta(X_i), v \rangle, \quad \text{and} \quad V_{\theta,v,w}^{(2)} = \frac{1}{n} \sum_{i=1}^n \varepsilon_i \langle \nabla^2 \log p_\theta(X_i) v, w \rangle,$$

for i.i.d. Rademacher random variables $(\varepsilon_i)_{i=1}^n$. Standard symmetrization arguments imply that $\mathbb{E} \left[\sup_{\theta \in \mathbb{B}(\theta_\sigma^*, R_0)} Z^{(i)} \right] \leq 2\mathbb{E} \left[\sup_{\theta \in \mathbb{B}(\theta_\sigma^*, R_0)} V^{(i)} \right]$ for $i \in \{1, 2\}$.

Let $P_n := \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$, we claim the following covering number bounds, conditionally on the data X_1^n :

$$\begin{aligned} \bar{N}^{(1)}(t) &:= \left| \mathcal{N} \left(\mathcal{G}_{v,j}^{(1)}, \|\cdot\|_{L^2(P_n)}, t \left\| \bar{G}^{(1)} \right\|_{L^2(P_n)} \right) \right| \\ &\leq \left(\frac{c \sum_{k'=1}^K \|u_{k'}^* - X_i\|_2^2 + cK(R_0^2 + 1)}{t} \right)^{Kd}, \end{aligned} \quad (81a)$$

$$\begin{aligned} \bar{N}^{(2)}(t) &:= \left| \mathcal{N} \left(\mathcal{G}_{v,w,j,\ell}^{(2)}, \|\cdot\|_{L^2(P_n)}, t \left\| \bar{G}^{(2)} \right\|_{L^2(P_n)} \right) \right| \\ &\leq \left(\frac{cK \sum_{k'=1}^K \|u_{k'}^* - X_i\|_2^3 + cK^2(R_0^3 + 1)}{t} \right)^{Kd}. \end{aligned} \quad (81b)$$

By Dudley's chaining integral bound, we obtain the following bounds:

$$\begin{aligned} \mathbb{E} \left[\sup_{\theta \in \mathbb{B}(\theta_\sigma^*, R_0)} V^{(1)} \right] &\leq \sqrt{\frac{1}{n} \mathbb{E} [\bar{G}^{(1)}(X)^2]} \int_0^1 \sqrt{1 + \mathbb{E} [\log \bar{N}^{(1)}(t)]} dt \\ &\leq cK(1 + R_0 + \sigma_X) \sqrt{\frac{Kd \log(Kd)}{n}}, \\ \mathbb{E} \left[\sup_{\theta \in \mathbb{B}(\theta_\sigma^*, R_0)} V^{(2)} \right] &\leq \sqrt{\frac{1}{n} \mathbb{E} [\bar{G}^{(2)}(X)^2]} \int_0^1 \sqrt{1 + \mathbb{E} [\log \bar{N}^{(2)}(t)]} dt \\ &\leq cK^2(1 + R_0^2 + \sigma_X^2) \sqrt{\frac{Kd \log(Kd)}{n}}. \end{aligned}$$

Combining with the concentration inequalities, we obtain the following bounds with probability $1 - \delta$:

$$\begin{aligned} \sup_{\theta \in \mathbb{B}(\theta_\sigma^*, R_0)} Z_{\theta,v}^{(1)} &\leq cK(1 + R_0 + \sigma_X) \left[\sqrt{\frac{Kd \log(Kd) + \log \delta^{-1}}{n}} + \frac{\log \delta^{-1}}{n} \sqrt{\log n} \right], \\ \sup_{\theta \in \mathbb{B}(\theta_\sigma^*, R_0)} Z_{\theta,v}^{(2)} &\leq cK^2(1 + R_0^2 + \sigma_X^2) \left[\sqrt{\frac{Kd \log(Kd) + \log \delta^{-1}}{n}} + \frac{\log \delta^{-1}}{n} \log n \right]. \end{aligned}$$

Finally, by taking union bound over a maximal $\frac{1}{8}$ -packing of the sphere \mathbb{S}^{d-1} , which has cardinality bounded by 17^{Kd} , for $\sigma_X \geq 1$ and $\frac{n}{\log n} \geq Kd \log \frac{Kd}{\delta}$, we conclude that

$$\begin{aligned} \sup_{\theta \in \mathbb{B}(\theta_\sigma^*, R_0)} \|\nabla F_n(\theta) - \nabla F(\theta)\|_2 &\leq cK(\sigma_X + R_0) \sqrt{\frac{Kd \log(Kd) + \log \delta^{-1}}{n}}, \\ \sup_{\theta \in \mathbb{B}(\theta_\sigma^*, R_0)} \|\nabla^2 F_n(\theta) - \nabla^2 F(\theta)\|_{\text{op}} &\leq cK^2(\sigma_X^2 + R_0^2) \sqrt{\frac{Kd \log(Kd) + \log \delta^{-1}}{n}}, \end{aligned}$$

which proves the desired bounds in claims (77b) and (77c).

Proof of equations (81a) and (81b) Given a positive number ε' to be determined later, let $\{\theta_1, \theta_2, \dots, \theta_M\}$ be a minimal ε' -covering of the parameter space $\mathbb{B}(\theta_\sigma^*, R_0)$. By standard volume arguments, we have that $M \leq \left(\frac{c}{\varepsilon'}\right)^{Kd}$.

We bound the $L^2(P_n)$ covering number by studying the Lipschitz constant for the functions in these classes. Note that for each $\theta \in \mathbb{B}(\theta_\sigma^*, R_0)$, simple derivation yields that:

$$\begin{aligned} \|\nabla_\theta \langle \nabla \log p_\theta(X_i), v_j \rangle\|_2 &\leq c \sum_{k'=1}^K \|u_{k'}^* - X_i\|_2^2 + cK(R_0^2 + 1), \quad \text{and} \\ \|\nabla_\theta \langle \nabla^2 \log p_\theta(X_i) v_j, w_\ell \rangle\|_2 &\leq cK \sum_{k'=1}^K \|u_{k'}^* - X_i\|_2^3 + cK^2(R_0^3 + 1). \end{aligned}$$

Note furthermore that $\bar{G}^{(1)}(x) \geq 1$ and $\bar{G}^{(2)}(x) \geq 1$ by definition. By taking $\varepsilon' := \frac{t}{c \sum_{k'=1}^K \|u_{k'}^* - X_i\|_2^2 + cK(R_0^2 + 1)}$, the set $\{p_{\theta_i} : i \in [M]\}$ constitutes a t -packing of the set $\mathcal{G}_{v,j}^{(1)}$. We, therefore, have the bound

$$\bar{N}^{(1)}(t) \leq \left| \mathcal{N} \left(\mathcal{G}_{v,j}^{(1)}, \|\cdot\|_{L^\infty(P_n)}, t \right) \right| \leq \left(\frac{c \sum_{k'=1}^K \|u_{k'}^* - X_i\|_2^2 + cK(R_0^2 + 1)}{t} \right)^{Kd}.$$

Similarly, we have that

$$\bar{N}^{(2)}(t) \leq \left| \mathcal{N} \left(\mathcal{G}_{v,w,j,\ell}^{(2)}, \|\cdot\|_{L^\infty(P_n)}, t \right) \right| \leq \left(\frac{cK \sum_{k'=1}^K \|u_{k'}^* - X_i\|_2^3 + cK^2(R_0^3 + 1)}{t} \right)^{Kd},$$

which proves the claim in equations (81a) and (81b).

D.4.3 Proof of the claim (77d)

For the global condition, we use an argument slightly different from the third condition in corollary 7. Note that for $\theta = [u_j]_{j \in [K]}$, the log-likelihood function takes the form

$$F_n([u_j]_{j \in [K]}) = \frac{1}{n} \sum_{i=1}^n \log \left(\sum_{j=1}^K \exp \left(- \|u_j - X_i\|_2^2 / 2 \right) \right) - \frac{\log(2\pi)d}{2} - \log(K).$$

Given X_1^n , we denote $R_X := \max_{i \in [n]} \|X_i\|_2$, and define the compact set

$$U(r) := \{[u_j]_{j \in [K]} : \|u_j\|_2 \leq r \quad \text{for all } j \in [K]\}. \quad (82a)$$

Now, we claim that for all $t > 0$

$$\Pi \left(U^c(3R_X + \sqrt{d+t}) \mid X_1^n \right) \leq \left(c(R_X + 1)\sqrt{n} + \sqrt{L_2} \right)^{d/2} \pi_0^{-1} e^{-t/2}. \quad (82b)$$

Taking this claim as given, by choosing $t = cd \log(R_X n + L_2) + c \log \frac{1}{\pi_0 \vartheta}$, we have the tail bound $\Pi \left(\mathbb{B}^c(0, 3\sqrt{K}R_X + \sqrt{K(d+t)}) \mid X_1^n \right) < \vartheta$. As a consequence, we obtain the conclusion of claim (77d).

D.4.4 Proof of the claim (82b)

Given $\theta = [u_j]_{j \in [K]}$, if $\theta \notin U(3R_X)$, there exists $j_0 \in [K]$ such that $\|u_{j_0}\|_2 > 3R_X$. We note that for each $i \in [n]$, we have:

$$\begin{aligned} \exp\left(-\frac{1}{2}\|u_{j_0} - X_i\|_2^2\right) &< \exp\left(-\frac{1}{2}(3R_X - \|X_i\|_2)^2\right) \\ &\leq \exp\left(-\frac{1}{2}(2R_X)^2\right) < \exp\left(-\frac{1}{2}\|X_1 - X_i\|_2^2\right). \end{aligned}$$

Therefore, we obtain

$$F_n([u_j]_{j \in [K]}) < F_n([u_1, \dots, u_{j_0-1}, X_1, u_{j_0+1}, \dots, u_K]).$$

Consequently, we can replace any u_j whose norm is larger than $3R_X$ with X_1 , and increase the log-likelihood function. The global maximum of the F_n is therefore attained only in the set $U(3R_X)$.

On the other hand, for any $\hat{\theta}_\sigma^{(n)} \in \arg \max_{\theta \in U(3R_X)} F_n(\theta)$, we have that

$$\|\nabla^2 F_n(\theta)\|_{\text{op}} \leq \frac{1}{n} \sum_{i=1}^n \left(\|u_j - X_i\|_2^2 + 1 \right) \leq 16R_X^2 + 1.$$

Taking the local radius $r := \frac{1}{(4R_X+1)\sqrt{n}} \wedge \frac{1}{\sqrt{L_2}}$, we have the lower bound

$$\begin{aligned} \int_{\mathbb{B}(\hat{\theta}_\sigma^{(n)}, r)} e^{nF_n(\theta)} \pi(d\theta) &\geq \pi_0 \cdot \text{Vol}(\mathbb{B}(\hat{\theta}_\sigma^{(n)}, r)) e^{nF_n(\hat{\theta}_\sigma^{(n)})} \cdot e^{-L_2 r^2/2} \cdot e^{-n(16R_X^2+1)r^2/2} \\ &\geq 4\pi_0 (cr)^d e^{nF_n(\hat{\theta}_\sigma^{(n)})}, \end{aligned}$$

for a universal constant $c > 0$.

On the other hand, for any $t > 0$, we have that:

$$\int_{\mathbb{B}^c(\hat{\theta}_\sigma^{(n)}, 3R_X + \sqrt{d+t})} e^{nF_n(\theta)} \pi(d\theta) \leq e^{nF_n(\hat{\theta}_\sigma^{(n)})} \int_{\mathbb{B}^c(\hat{\theta}_\sigma^{(n)}, t\sqrt{d})} \pi(d\theta) \leq e^{nF_n(\hat{\theta}_\sigma^{(n)})} \cdot e^{-t/2}.$$

Consequently, we have the upper bound on the posterior tail probability:

$$\begin{aligned} \Pi\left(\mathbb{B}(0, 3R_X + \sqrt{d+t}) \mid X_1^n\right) &\leq \left(\int_{\mathbb{B}(\hat{\theta}_\sigma^{(n)}, r)} e^{nF_n(\theta)} \pi(d\theta) \right)^{-1} \\ &\quad \times \int_{\mathbb{B}^c(\hat{\theta}_\sigma^{(n)}, 3R_X + \sqrt{d+t})} e^{nF_n(\theta)} \pi(d\theta) \\ &\leq \left(c(R_X + 1)\sqrt{n} + \sqrt{L_2} \right)^{d/2} \pi_0^{-1} e^{-t/2}, \end{aligned}$$

for some universal constant $c > 0$. Therefore, we obtain the conclusion of claim (82b).

D.5 Proof of corollary 6

We first invoke theorem 3 in a small local neighborhood of $\theta^* = 0$. We claim that there exist constants $q_1, q_2, q_3, R_0 - > 0$ that depend on the density function f but independent of n and a_n , such that:

$$-\langle \theta, \nabla \tilde{F}^S(\theta) \rangle \geq q_1 |\theta|^{1+2\beta} - q_2 a_n^{1+2\beta}, \quad \text{for } \theta \in (-r_0/2, r_0/2), \quad (83a)$$

$$\sup_{\theta \in [-1, 1]} \left| \nabla \tilde{F}^S(\theta) - \nabla \tilde{F}_n^S(\theta) \right| \leq q_3 \left(a_n^{\beta-1/2} \sqrt{\frac{\log n/\delta}{n}} + a_n^{\beta-1} \frac{\log n/\delta}{n} \right), \quad (83b)$$

with probability $1 - \delta$. Assume that the above claims are given at the moment (their proofs are given in Appendices D.5.2 and D.5.3). Invoking theorem 3, there exists a universal constant $c > 0$, such that:

$$\tilde{\Pi}(\mathbb{B}(\theta^*, r_0/4) \mid X_1^n)^{-1} \tilde{\Pi}(\mathbb{B}(\theta^*, c \cdot r_n) \mid X_1^n) \geq 1 - \vartheta,$$

where the scalar r_n is given by

$$\begin{aligned} r_n = q_2^{\frac{1}{1+2\beta}} a_n + \left(\frac{1 + \log(1/\vartheta)}{n} \right)^{\frac{1}{1+2\beta}} + \left(\frac{q_3}{q_1} \cdot \left(a_n^{\beta-1/2} \sqrt{\frac{\log(n/\delta)}{n}} + a_n^{\beta-1} \frac{\log(n/\delta)}{n} \right) \right)^{\frac{1}{2\beta}} \\ + \left(\frac{1}{q_3 n} \right)^{\frac{1}{2\beta}}. \end{aligned}$$

Taking $a_n = n^{-\frac{1}{1+2\beta}}$, we conclude that

$$r_n \leq q' \cdot n^{-\frac{1}{1+2\beta}} \left(\log^{\frac{1}{2\beta}}(n/\delta) + \log^{\frac{1}{1+2\beta}}(1/\vartheta) \right), \quad (84)$$

where the constant $q' > 0$ depends on q_1, q_2, q_3 and β .

It remains to lower bound the posterior probability in a small ball $\mathbb{B}(0, r_0/4)$. We claim that there exists a constant $\Delta_0 > 0$ depending on r_0 and f , such that there exists a constant $q_0 > 0$ depending on the function f and the quantities r_0, Δ_0 , when $n \geq q_0 \log^{1+\frac{1}{2\beta}} \delta^{-1}$, the following bound holds true with probability $1 - \delta$:

$$\sup_{\theta \in \mathbb{B}(0, 1) \setminus \mathbb{B}(0, r_0/4)} \tilde{F}_n^S(\theta) < \inf_{|\theta'| < a_n} \tilde{F}^S(\theta') - \frac{1}{2} \Delta_0. \quad (85)$$

Taking this bound as given, we proceed with the proof of this corollary. In order to bound the smoothed posterior probability outside the ball $\mathbb{B}(0, r_0/4)$, we note that:

$$\begin{aligned} & \tilde{\Pi}(\mathbb{B}(0, 1) \setminus \mathbb{B}(0, r_0/4) \mid X_1^n) \\ & \leq \tilde{\Pi}(\mathbb{B}(0, a_n) \mid X_1^n)^{-1} \tilde{\Pi}(\mathbb{B}(0, 1) \setminus \mathbb{B}(0, r_0/4) \mid X_1^n) \\ & \leq \frac{\sup_{\theta \in \mathbb{B}(0, 1) \setminus \mathbb{B}(0, r_0/4)} \exp(n \tilde{F}_n^S(\theta))}{2a_n \cdot \inf_{|\theta| < a_n} \exp(n \tilde{F}^S(\theta)) \cdot \inf_{|\theta| < a_n} \pi(\theta)} \\ & \leq \frac{1}{2a_n \pi(0) e^{-B}} \cdot \exp\left(-\frac{\Delta_0 n}{2}\right). \end{aligned}$$

Given $n \geq \frac{2}{\Delta_0} \left(B + c \log \frac{n}{\vartheta \pi(0)} \right)$, for a prior density π supported on the interval $[-1, 1]$, we have that

$$\tilde{\Pi}(\mathbb{B}(0, r_0/4)^C \mid X_1^n) = \tilde{\Pi}(\mathbb{B}(0, 1) \setminus \mathbb{B}(0, r_0/4) \mid X_1^n) \leq \vartheta.$$

Therefore, for $n \geq q_0 \log^{1+\frac{1}{2\beta}} \delta^{-1}$, we conclude that

$$\tilde{\Pi}(\mathbb{B}(0, cr_n) \mid X_1^n) \geq 1 - 2\vartheta$$

with probability at least $1 - \delta$, where r_n was defined in equation (84).

D.5.1 Proof of claim (85)

We first prove the result for the original population-level log-likelihood F^S , and then show that the smoothing does not affect the gap up to constant factors. Finally we show the sample-level version using the deviation bound (83b).

Denote the density function $f_\theta(x) := f(x - \theta)$. We note that

$$F^S(0) - F^S(\theta) = D_{\text{KL}}(f \parallel f_\theta) \geq 0$$

for any $\theta \in \mathbb{R}$. Furthermore, the function F^S is continuous in the interval $[-1, 1]$. On the compact set $[-1, -r_0/4] \cup [r_0/4, 1]$, the maximum point of the continuous function F^S is attainable, i.e.,

$$\exists \theta_0 \in [-1, -r_0/4] \cup [r_0/4, 1], \quad \text{s.t. } F^S(\theta_0) = \sup_{\theta \in [-1, -r_0/4] \cup [r_0/4, 1]} F^S(\theta).$$

Since $f \neq f_{\theta_0}$, we have that $D_{\text{KL}}(f \parallel f_{\theta_0}) > 0$. We define $\Delta_0 := \frac{1}{2} D_{\text{KL}}(f \parallel f_{\theta_0})$.

On the other hand, since the function F^S is continuous on the compact set $[-2, 2]$, by Heine-Cantor theorem, F^S is also uniformly continuous on $[-2, 2]$, i.e.,

$$\lim_{\delta \rightarrow 0^+} \sup_{\theta, \theta' \in [-2, 2], |\theta - \theta'| \leq \delta} |F(\theta) - F(\theta')| = 0.$$

So there exists $\delta_0 > 0$, such that when $|\theta - \theta'| < \delta_0$ for some $\theta, \theta' \in [-2, 2]$, we have that

$$|F(\theta) - F(\theta')| \leq \frac{1}{2} \Delta_0.$$

Consequently, for n large enough such that $a_n = n^{-\frac{1}{1+2\beta}} < \delta_0/2$, for any $\theta \in [-1, -r_0/4] \cup [r_0/4, 1]$, we have the following bound:

$$\begin{aligned} \tilde{F}^S(\theta) &\leq F^S(\theta) + \sup_{\theta' \in [\theta - a_n, \theta + a_n]} |F(\theta) - F(\theta')| \leq F^S(\theta_0) + \frac{1}{2} \Delta_0 \\ &\leq F^S(0) - 2\Delta_0 + \frac{1}{2} \Delta_0 \\ &\leq \inf_{|\theta'| < a_n} \tilde{F}^S(\theta') - \frac{3}{2} \Delta_0 + \sup_{\theta' \in [-a_n, a_n]} |F(0) - F(\theta')| \\ &\leq \inf_{|\theta'| < a_n} \tilde{F}^S(\theta') - \Delta_0. \end{aligned}$$

For the sample version, we note that the bound (83b) implies the following inequality with probability $1 - \delta$:

$$\begin{aligned} & \sup_{\theta, \theta' \in [-1, 1]} \left| (\tilde{F}_n^S(\theta) - \tilde{F}_n^S(\theta')) - (\tilde{F}^S(\theta) - \tilde{F}^S(\theta')) \right| \\ & \leq \sup_{\theta \in [-1, 1]} \int_{\theta'}^{\theta} \left| \nabla \tilde{F}_n^S(s) - \nabla \tilde{F}^S(s) \right| ds \\ & \leq q_3 n^{-\frac{2\beta}{1+2\beta}} \log \frac{n}{\delta}. \end{aligned}$$

Given $n \geq c \left(\frac{q_3}{\Delta_0} \log(1/\delta) \right)^{1+\frac{1}{2\beta}}$, we have that:

$$\sup_{\theta, \theta' \in [-1, 1]} \left| (\tilde{F}_n^S(\theta) - \tilde{F}_n^S(\theta')) - (\tilde{F}^S(\theta) - \tilde{F}^S(\theta')) \right| \leq \frac{1}{4} \Delta_0.$$

Combining with the population-level bound, we obtain the following bound with probability $1 - \delta$:

$$\sup_{\theta \in [-1, -r_0/4] \cup [r_0/4, 1]} \tilde{F}_n^S(\theta) \leq \inf_{|\theta'| < a_n} \tilde{F}^S(\theta') - \frac{1}{2} \Delta_0,$$

which proves the desired claim.

D.5.2 Local structure of \tilde{F}^S

Now we prove claim (83a). We first analyze the local structure of F^S , and then study the effect of smoothing. For $\theta > 0$, direct calculation yields:

$$\begin{aligned} -\nabla_{\theta} F^S(\theta) &= \int_{-\infty}^{+\infty} f(x) \nabla \log f(x - \theta) dx \\ &= \int_{-\infty}^{+\infty} (f(x + \theta) - f(x)) \nabla \log f(x) dx \\ &= \int_{-\infty}^{+\infty} \int_0^{\theta} f(x + z) (\nabla \log f(x + z) \cdot \nabla \log f(x)) dz dx = I_1 + I_2 + I_3 + I_4. \end{aligned} \quad (86)$$

where the terms I_1 , I_2 , I_3 and I_4 are defined as follows:

$$\begin{aligned} I_1(\theta) &:= \beta^2 \int_{-\infty}^{\infty} \int_0^{\theta} f(x + z) \ell(x) \ell(x + z) |x + z|^{\beta-1} |x + z|^{\beta-1} \operatorname{sgn}(x(x + z)) dz dx \\ I_2(\theta) &:= \beta^2 \int_{-\infty}^{\infty} \int_0^{\theta} f(x + z) \ell(x + z) |x + z|^{\beta-1} \nabla \log h(x) \operatorname{sgn}(x + z) dz dx, \\ I_3(\theta) &:= \beta^2 \int_{-\infty}^{\infty} \int_0^{\theta} f(x + z) \ell(x) |x|^{\beta-1} \nabla \log h(x + z) \operatorname{sgn}(x) dz dx, \\ I_4(\theta) &:= \beta^2 \int_{-\infty}^{\infty} \int_0^{\theta} f(x + z) \nabla \log h(x) \cdot \nabla \log h(x + z) dz dx. \end{aligned}$$

For the term I_1 , we note that:

$$\begin{aligned}
I_1 &= \theta^{2\beta} \ell(0^+)^2 \beta \int_1^{+\infty} y^{\beta-1} (y^\beta - (y-1)^\beta) f(\theta y) dy \\
&\quad + \theta^{2\beta} \ell(0^-)^2 \beta \int_0^{+\infty} y^{\beta-1} ((y+1)^\beta - y^\beta) f(-\theta y) dy \\
&\quad + \theta^{2\beta} \beta \ell(0^+) \int_0^1 y^{\beta-1} (\ell(0^+) y^\beta - \ell(0^-) (1-y)^\beta) f(\theta y) dy \\
&\geq \theta^{2\beta} \beta \int_0^1 \ell(0^+)^2 y^{\beta-1} y^\beta f(\theta y) dy + \int_0^1 \ell(0^-)^2 y^{\beta-1} ((y+1)^\beta - y^\beta) f(-\theta y) dy \\
&\quad - \int_0^1 \ell(0^+) \ell(0^-) y^{\beta-1} (1-y)^\beta f(\theta y) dy.
\end{aligned}$$

Since the function f is continuous at point 0, and $f(0) > 0$, there exists $r' > 0$, such that:

$$\forall s \in (-r', r'), \quad |f(s) - f(0)| \leq \frac{1}{10} f(0).$$

For $\theta \in (0, r')$, we have:

$$\begin{aligned}
\int_0^1 \ell(0^+) \ell(0^-) y^{\beta-1} (1-y)^\beta f(\theta y) dy &\leq \frac{11}{10} \ell(0^+) \ell(0^-) f(0) \int_0^1 y^{\beta-1} (1-y)^\beta dy \\
&\leq \left(\frac{4}{5} \ell(0^+)^2 + \frac{121}{320} \ell(0^-)^2 \right) f(0) \int_0^1 y^{\beta-1} (1-y)^\beta dy, \\
\int_0^1 \ell(0^-)^2 y^{\beta-1} ((y+1)^\beta - y^\beta) f(-\theta y) dy &\geq \frac{9}{10} \ell(0^-)^2 f(0) \int_0^1 y^{\beta-1} ((1+y)^\beta - y^\beta) dy, \\
\int_0^1 \ell(0^+)^2 y^{\beta-1} y^\beta f(\theta y) dy &\geq \frac{9}{10} \ell(0^+)^2 f(0) \int_0^1 y^{2\beta-1} dy.
\end{aligned}$$

Note that $y^{\beta-1} (1-y)^\beta \leq y^{\beta-1} ((1+y)^\beta - y^\beta) + y^{2\beta-1}$. Therefore, for $\theta \in (0, r')$, we have the following lower bound on I_1 :

$$I_1 \geq \frac{\beta (\ell(0^+)^2 + \ell(0^-)^2) f(0)}{10} \theta^{2\beta}$$

On the other hand, we can also deduce the following upper bound on I_1 from above expression:

$$\begin{aligned}
|I_1| &\leq \theta^{2\beta} \left(b^2 \int_1^{+\infty} y^{\beta-1} (y^\beta - (y-1)^\beta) f(\theta y) dy + a^2 \int_0^{+\infty} y^{\beta-1} ((y+1)^\beta - y^\beta) f(-\theta y) dy \right) \\
&\quad + \theta^{2\beta} b \int_0^1 y^{\beta-1} (by^\beta + a(1-y)^\beta) f(\theta y) dy \\
&\leq \theta^{2\beta} \sup_{z \in \mathbb{R}} f(z) \cdot \left(\int_1^{+\infty} (b^2 y^{\beta-1} + a^2 (y-1)^{\beta-1}) (y^\beta - (y-1)^\beta) dy \right. \\
&\quad \left. + b \int_0^1 y^{\beta-1} (by^\beta + a(1-y)^\beta) dy \right) \\
&\leq M_1 \theta^{2\beta},
\end{aligned}$$

for a constant $M_1 < +\infty$ depending on a, b and β .

Note that above arguments holds true also on the side $\theta \rightarrow 0^-$. We therefore have the following lower bound for $\theta \in (-r', r')$:

$$\frac{\beta(a^2 + b^2)f(0)}{10}|\theta|^{2\beta} \leq I_1(\theta) \leq M_1|\theta|^{2\beta}. \quad (87)$$

Now we bound each of I_2, I_3, I_4 respectively, and show that they are of order $O(\theta)$, as $\theta \rightarrow 0^+$. For the term $I_2(\theta)$, it is easy to see by definition that $I_2(0) = 0$, and by the Lebesgue differentiation theorem, we have that:

$$\begin{aligned} \left| \frac{dI_2}{d\theta}(\theta) \right| &\leq \beta^2 \int_{-\infty}^{\infty} f(x + \theta) |\ell(x + \theta)| |x + \theta|^{\beta-1} \cdot |\nabla \log h(x)| dx \\ &\leq (|\ell^+(0)| + |\ell^-(0)|) \cdot \left(\int_{-1}^1 f(y) |y|^{\beta-1} \cdot |\nabla \log h(y - \theta)| dy \right. \\ &\quad \left. + \int_{-\infty}^{\infty} f(y) |\nabla \log h(y - \theta)| dy \right). \end{aligned}$$

Invoking the assumption (33) on h , we have that:

$$\begin{aligned} \left| \frac{dI_2}{d\theta}(\theta) \right| &\leq 4(|a| + |b|)c_1 e^{|a|+|b|+c_1} f(0) \int_{-1}^1 |y|^{\beta-1} dy + c_1(|a| + |b|). \\ &\leq \frac{8(|a| + |b|)}{\beta} c_1 e^{|a|+|b|+c_1} f(0) + (|a| + |b|)c_1 =: M_2 \end{aligned}$$

So for $|\theta| \leq 1$, we have that:

$$|I_2(\theta)| \leq M_2|\theta|. \quad (88)$$

Similarly, for the term I_3 , when $|\theta| \leq 1$, we have:

$$\begin{aligned} \left| \frac{dI_3}{d\theta}(\theta) \right| &\leq (|a| + |b|) \int_{-\infty}^{+\infty} f(x + \theta) |x|^{\beta-1} |\nabla \log h(x + \theta)| dx \\ &\leq (|a| + |b|) \left(\int_{-1}^1 f(0) e^{c_1} |x|^{\beta-1} c_1 dx + c_1 \right). \\ &\leq \frac{8(|a| + |b|)}{\beta} c_1 e^{|a|+|b|+c_1} f(0) + (|a| + |b|)c_1 =: M_3, \end{aligned}$$

and consequently, for $\theta \in [-1, 1]$, we have the bound

$$|I_3(\theta)| \leq M_3|\theta|. \quad (89)$$

For the last term I_4 , simple calculation yields:

$$\begin{aligned} |I_4(\theta)| &\leq |\theta| \cdot \mathbb{E}_f \left[\sup_{z \in [0, \theta]} |\nabla \log h(X)| \cdot |\nabla \log h(X + z)| \right] \\ &\leq c_1^2 |\theta| =: M_4 |\theta|. \end{aligned} \quad (90)$$

for $\theta \in [-1, 1]$.

We define $r_0 := \min \left(r', 1, \left(\frac{\beta(a^2 + b^2)f(0)}{20(M_2 + M_3 + M_4)} \right)^{\frac{1}{1-2\beta}} \right)$. Plugging the bounds (87)-(90) to equation (86), for $|\theta| < r_0$, we have:

$$\begin{aligned} \langle \theta, \nabla F^S(\theta) \rangle &\geq \frac{\beta(a^2 + b^2)}{10} |\theta|^{1+2\beta} - (M_2 + M_3 + M_4) |\theta|^2 \\ &\geq \frac{\beta(a^2 + b^2)}{20} |\theta|^{1+2\beta}. \end{aligned}$$

Given the smoothing radius $a_n < r_0/2$, for any $\theta \in (a_n, r_0 - a_n)$, we have that:

$$\begin{aligned} \langle \theta, \nabla \tilde{F}^S(\theta) \rangle &= \frac{1}{2a_n} \int_{-a_n}^{a_n} \frac{\theta}{\theta + z} \cdot (\theta + z) \cdot \nabla F^S(\theta + z) dz \\ &\geq \frac{1}{2a_n} \int_{-a_n}^{a_n} \frac{\theta}{\theta + z} \cdot \frac{\beta(a^2 + b^2)}{20} (\theta + z)^{1+2\beta} dz \\ &\geq \frac{\beta(a^2 + b^2)}{20} \theta \cdot (\theta - a_n)^{2\beta}. \end{aligned} \tag{91}$$

For $\theta \in (0, a_n)$, we note that:

$$\begin{aligned} \langle \theta, \nabla \tilde{F}^S(\theta) \rangle &\geq \frac{1}{2a_n} \left(\int_{-\theta}^{a_n} \frac{\theta}{\theta + z} \cdot (\theta + z) \cdot \nabla F^S(\theta + z) dz - \int_{-a_n}^{-\theta} |\theta| \cdot |\nabla F^S(\theta + z)| dz \right) \\ &\geq -\frac{1}{2a_n} \cdot a_n \theta \cdot \sup_{|z| \leq a_n} (|I_1(\theta)| + |I_2(\theta)| + |I_3(\theta)| + |I_4(\theta)|) \\ &\geq -2M_1 a_n^{1+2\beta}. \end{aligned}$$

We can observe that similar bounds also hold true in the intervals $(-r_0 + a_n, -a_n)$ and $(-a_n, 0)$. Therefore, we conclude that the following bound holds true within the interval $(r_0/2, r_0/2)$:

$$\langle \theta, \nabla \tilde{F}^S(\theta) \rangle = \begin{cases} -2M_1 a_n^{1+2\beta}, & 0 \leq r \leq a_n, \\ \frac{\beta(a^2 + b^2)}{20} \cdot (r - a_n)^{1+2\beta} - 2M_1 a_n^{1+2\beta}, & a_n \leq r \leq r_0/2, \end{cases}$$

D.5.3 Bounding the difference $\nabla \tilde{F}^S - \nabla \tilde{F}_n^S$

Now, we proceed to prove claim (83b). By definition, we note that

$$\begin{aligned} \frac{d}{d\theta} \tilde{F}_n^S(\theta) &= \frac{1}{n} \sum_{i=1}^n \frac{1}{2a_n} (\log f(\theta + a_n - X_i) - \log f(\theta - a_n - X_i)) \\ &= \frac{1}{n} \sum_{i=1}^n \frac{1}{2a_n} \left(|\theta + a_n - X_i|^\beta \ell(\theta + a_n - X_i) - |\theta - a_n - X_i|^\beta \ell(\theta - a_n - X_i) \right) \\ &\quad + \frac{1}{n} \sum_{i=1}^n \frac{\log h(\theta + a_n - X_i) - \log h(\theta - a_n - X_i)}{2a_n}. \end{aligned}$$

We define the following function:

$$\begin{aligned}\eta_\theta(x) &:= \frac{\log h(\theta + a_n - x) - \log h(\theta - a_n - x)}{2a_n}, \quad \text{and} \\ \nu_\theta(x) &:= \frac{1}{2a_n} \left(|\theta + a_n - x|^\beta \ell(\theta + a_n - x) - |\theta - a_n - x|^\beta \ell(\theta + a_n - x) \right)\end{aligned}$$

We also define the following random variable:

$$Z_n^{(1)}(\theta) := \frac{1}{n} \sum_{i=1}^n \eta_\theta(X_i) - \mathbb{E}[\eta_\theta(X)], \quad \text{and} \quad Z_n^{(2)}(\theta) := \frac{1}{n} \sum_{i=1}^n \nu_\theta(X_i) - \mathbb{E}[\nu_\theta(X)].$$

D.5.4 Upper bounds for the term $Z_n^{(1)}$

By the Lipschitz assumption (33), we have

$$|\eta_\theta(x)| \leq \frac{1}{2a_n} \int_{\theta-a_n-x}^{\theta+a_n-x} |\nabla \log h(t)| dt \leq c_1.$$

By the Hoeffding bound, for any given $\theta \in [-1, 1]$ and $t > 0$, we obtain

$$\mathbb{P}(|Z_n^{(1)}(\theta)| > t) \leq 2 \exp\left(-\frac{2nt^2}{c_1^2}\right).$$

On the other hand, for $\theta_1, \theta_2 \in [-1, 1]$, we note that:

$$|\eta_{\theta_1}(x) - \eta_{\theta_2}(x)| \leq c_1 \frac{|\theta_1 - \theta_2|}{2a_n},$$

which implies that $|Z_n^{(1)}(\theta_1) - Z_n^{(1)}(\theta_2)| \leq c_1 \frac{|\theta_1 - \theta_2|}{2a_n}$ almost surely.

Let $\mathcal{M}_n := \{\theta_1, \theta_2, \dots, \theta_K\}$ be a maximal $\frac{a_n}{c_1 n}$ -packing of the interval $[-1, 1]$. By union bound, we find that

$$\mathbb{P}\left(\exists \theta \in \mathcal{M}_n, |Z_n^{(1)}(\theta)| \geq t\right) \leq 2|\mathcal{M}_n| e^{-\frac{2nt^2}{c_1^2}}.$$

Consequently, for any $\delta > 0$, we have the following uniform upper bound with probability $1 - \delta$:

$$\sup_{\theta \in [-1, 1]} |Z_n^{(1)}(\theta)| \leq \frac{2}{n} + c_1 \sqrt{\frac{1}{n} \log \frac{|\mathcal{M}_n|}{\delta}} \leq \frac{2}{n} + 3c_1 \sqrt{\frac{1}{n} \log \frac{n}{\delta}}.$$

Upper bounds for the term $Z_n^{(2)}$ We first study moment bounds of the random variable $\nu_\theta(X_i)$. For $p \geq 2$, we have

$$\mathbb{E}[|\nu_\theta(X_i)|^p] = \frac{1}{(2a_n)^p} \int_{-\infty}^{+\infty} \left| |z + a_n|^\beta \ell(z + a_n) - |z - a_n|^\beta \ell(z - a_n) \right|^p f(\theta - z) dz.$$

Define $S := \ell(0^-) + \ell(0^+)$, which is positive. To upper bound the integral, we split it into three terms:

$$\begin{aligned}\bar{I}_1(\theta) &:= \int_{-3a_n}^{3a_n} |\nu_\theta(\theta - z)|^p f(\theta - z) dz, \\ \bar{I}_2(\theta) &:= \int_{-1-3a_n}^{-3a_n} |\nu_\theta(\theta - z)|^p f(\theta - z) dz + \int_{3a_n}^{1+3a_n} |\nu_\theta(\theta - z)|^p f(\theta - z) dz, \\ \bar{I}_3(\theta) &:= \int_{-\infty}^{-1-3a_n} |\nu_\theta(\theta - z)|^p f(\theta - z) dz + \int_{1+3a_n}^{+\infty} |\nu_\theta(\theta - z)|^p f(\theta - z) dz.\end{aligned}$$

For the term $\bar{I}_1(\theta)$, we can simply take upper bounds on each term of $\nu_\theta(\theta - z)$, and obtain

$$\bar{I}_1(\theta) \leq (6a_n)^{(\beta-1)p+1} S^p \cdot \sup_{z \in [-a_n, a_n]} f(\theta - z).$$

For the term $\bar{I}_2(\theta)$, note that

$$\begin{aligned}& \int_{3a_n}^{1+3a_n} |\nu_\theta(\theta - z)|^p f(\theta - z) dz \\ &= (2a_n)^{-p} \ell(0^+)^p \int_{2a_n}^{1+2a_n} z^{p\beta} \left(\left(1 + \frac{2a_n}{z}\right)^\beta - 1 \right)^p f(\theta - z - a_n) dz \\ &\leq (2a_n)^{-p} \ell(0^+)^p \int_{2a_n}^{1+2a_n} (2\beta a_n z^{\beta-1})^p f(\theta - z - a_n) dz \\ &\leq \frac{\beta^p}{p(1-\beta)-1} \ell(0^+)^p (2a_n)^{1+p(\beta-1)} \sup_{z \in [-1-3a_n, 1+3a_n]} f(\theta - z).\end{aligned}$$

For the integral within the interval $[-1-3a_n, -3a_n]$, we have a similar upper bound. Putting them together, we obtain

$$\bar{I}_2(\theta) \leq \frac{S^p}{p(1-\beta)-1} (2a_n)^{1+p(\beta-1)} \sup_{z \in [-1-3a_n, 1+3a_n]} f(\theta - z).$$

For the last term $\bar{I}_3(\theta)$, we note that for $|z| > 1 + a_n$, there is

$$\left| |z + a_n|^\beta - |z - a_n|^\beta \right| \leq \int_{z-a_n}^{z+a_n} \beta |z - s|^{\beta-1} ds \leq 2\beta a_n \leq a_n.$$

Consequently, we have

$$\bar{I}_3(\theta) \leq (2a_n)^{-p} \left(\int_{-\infty}^{-1-3a_n} + \int_{-\infty}^{-1-3a_n} \right) a_n^p f(\theta - z) dz \leq 1.$$

Combining the above upper bounds of $\bar{I}_1(\theta)$, $\bar{I}_2(\theta)$, and $\bar{I}_3(\theta)$, for any $a_n < 1$ and $p \geq 2$, we obtain that

$$(\mathbb{E} |\nu_\theta(X_i)|^p)^{\frac{1}{p}} \leq C \left(\frac{\sup_{z \in [-4, 4]} f(\theta - z)}{p(1-\beta)-1} \right)^{\frac{1}{p}} a_n^{\beta-1+\frac{1}{p}},$$

for a universal constant $C > 0$.

Invoking Bernstein inequality, for any fixed $\theta \in [-1, 1]$ and $t > 0$, we have

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n \nu_\theta(X_i) - \mathbb{E}[\nu_\theta(X)] \right| > t \right) \leq \exp \left(\frac{-nt^2/2}{Qa_n^{2\beta-1} + a_n^{\beta-1}t/3} \right),$$

where $Q := C \frac{\sup_{z \in [-4, 4]} f(\theta - z)}{1-2\beta}$ for universal constant $C > 0$.

Now we extend the concentration inequality for a fixed θ to the uniform bound for any $\theta \in [-1, 1]$. Recall that $Z_n^{(2)}(\theta) := \frac{1}{n} \sum_{i=1}^n \nu_\theta(X_i) - \mathbb{E}[\nu_\theta(X)]$. For $\theta_1, \theta_2 \in [-1, 1]$, we have

$$\begin{aligned} & |Z_n^{(2)}(\theta_1) - Z_n^{(2)}(\theta_2)| \\ & \leq \frac{1}{na_n} \sum_{i=1}^n \left| |\theta_1 + a_n - X_i|^\beta \ell(\theta_1 + a_n - X_i) - |\theta_2 + a_n - X_i|^\beta \ell(\theta_2 + a_n - X_i) \right| \\ & + \frac{1}{na_n} \sum_{i=1}^n \left| |\theta_1 - a_n - X_i|^\beta \ell(\theta_1 - a_n - X_i) - |\theta_2 - a_n - X_i|^\beta \ell(\theta_2 - a_n - X_i) \right| \\ & \leq \frac{2S}{a_n} |\theta_1 - \theta_2|^\beta, \quad \text{a.s.} \end{aligned}$$

where we use the inequality $||x|^\beta - |y|^\beta| \leq |x - y|^\beta$ for any x, y .

Let $b_n := (\frac{a_n}{2Sn})^{\frac{1}{\beta}}$ and \mathcal{M}_{b_n} be a maximal b_n -packing of the interval $[-1, 1]$. For any $\theta \in [-1, 1]$, there exists $\theta' \in \mathcal{M}_{b_n}$, such that $|\theta - \theta'| < b_n$, which implies that $|Z_n^{(2)}(\theta) - Z_n^{(2)}(\theta')| < \frac{1}{n}$. Consequently, for any $t > 0$, we find that

$$\begin{aligned} \mathbb{P} \left(\sup_{\theta \in [-1, 1]} |Z_n^{(2)}(\theta)| > t + \frac{1}{n} \right) & \leq \mathbb{P} \left(\sup_{\theta \in \mathcal{M}_{b_n}} |Z_n^{(2)}(\theta)| > t \right) \\ & \leq |\mathcal{M}_{b_n}| \exp \left(\frac{-nt^2/2}{Qa_n^{2\beta-1} + a_n^{\beta-1}t/3} \right). \end{aligned}$$

Given $a_n > \frac{S}{n^2}$, we have $\log |\mathcal{M}_{b_n}| \leq \frac{3}{\beta} \log n$. Choosing appropriate value of t , we have

$$\sup_{\theta \in [-1, 1]} |Z_n^{(2)}(\theta)| \leq C \left(\sqrt{Q} a_n^{\beta-1/2} \sqrt{\frac{\log n/\delta}{n}} + a_n^{\beta-1} \frac{\log n/\delta}{n} + \frac{1}{n} \right),$$

with probability $1 - \delta$.

Collecting the bounds for the terms $Z_n^{(1)}$ and $Z_n^{(2)}$, for $a_n \in (\frac{S}{n^2}, 1)$, we conclude the following bound that holds true with probability $1 - \delta$:

$$\sup_{\theta \in [-1, 1]} \left| \nabla_\theta \tilde{F}_n^S(\theta) - \nabla_\theta \tilde{F}^S(\theta) \right| \leq c \cdot \left(\sqrt{Q} a_n^{\beta-1/2} \sqrt{\frac{\log n/\delta}{n}} + a_n^{\beta-1} \frac{\log n/\delta}{n} \right),$$

for a universal constant $c > 0$.

E Proofs of the remaining auxiliary results

In this appendix, we provide proofs of the remaining auxiliary results in the paper.

E.1 Proof of proposition 1

For any $p \geq 2$, we define the quantity:

$$R_p := \sup_{p \geq 0} (\mathbb{E}_{\pi_t} [\|X\|_2^p])^{1/p} \vee (\mathbb{E}_{\pi^*} [\|X\|_2^p])^{1/p}$$

For any given value $\bar{R} > 0$, we note the following decomposition:

$$\begin{aligned} & |\mathbb{E}_{\pi_t} [\|X\|_2^p] - \mathbb{E}_{\pi^*} [\|X\|_2^p]| \\ & \leq \int_{\mathbb{B}(0, \bar{R})} |\pi_t - \pi^*| \cdot \|x\|_2^p dx + \int_{\mathbb{B}(0, \bar{R})^C} \pi_t(x) \|x\|_2^p dx + \int_{\mathbb{B}(0, \bar{R})^C} \pi^*(x) \|x\|_2^p dx \\ & \leq \bar{R}^p \cdot d_{\text{TV}}(\pi_t, \pi^*) + \mathbb{E}_{\pi_t} \left[\|X\|_2^p \mathbf{1}_{\|X\|_2 > \bar{R}} \right] + \mathbb{E}_{\pi^*} \left[\|X\|_2^p \mathbf{1}_{\|X\|_2 > \bar{R}} \right] \\ & \leq \bar{R}^p \cdot d_{\text{TV}}(\pi_t, \pi^*) + \sqrt{\mathbb{E}_{\pi_t} [\|X\|_2^{2p}]} \sqrt{\pi_t (\|X\|_2 > \bar{R})} + \sqrt{\mathbb{E}_{\pi^*} [\|X\|_2^{2p}]} \sqrt{\pi^* (\|X\|_2 > \bar{R})} \\ & \leq \bar{R}^p \cdot d_{\text{TV}}(\pi_t, \pi^*) + 2R_{2p}^p \cdot R_2 / \bar{R}. \end{aligned}$$

For any $\varepsilon > 0$, take $\bar{R} := \frac{\varepsilon}{2R_{2p}^p R_2}$, we have that:

$$\lim_{t \rightarrow +\infty} |\mathbb{E}_{\pi_t} [\|X\|_2^p] - \mathbb{E}_{\pi^*} [\|X\|_2^p]| \leq \varepsilon,$$

which proves the claim.

E.2 Proof of corollary 7

It follows from theorem 3 that

$$\Pi \left(\mathbb{B}(\theta_j^*, r_n^{(j)}) \mid X_1^n \right) \leq \Pi \left(\mathbb{B}(\theta_j^*, r_0) \mid X_1^n \right) \cdot (1 - \vartheta), \quad \forall j \in [M]. \quad (92)$$

It remains to prove a lower bound on the sum $\sum_{j=1}^M \Pi(\mathbb{B}(\theta_j^*, r_0) \mid X_1^n)$. We utilize the following lemma, which controls the tail behavior of posterior distribution in an unbounded space.

Lemma 1. *Under the condition C.3 for corollary 7, for any $\vartheta > 0$, we have that:*

$$\Pi \left(\mathbb{B}(0, R(\vartheta)) \mid X_1^n \right) \geq 1 - \vartheta, \quad \text{where} \quad R(\vartheta) := 2R_\delta \log \vartheta^{-1} + \sqrt{\frac{6(d + \log \vartheta^{-1})}{c_\pi}}. \quad (93)$$

Taking this lemma as given, we proceed with the proof of the corollary. First, by the empirical process assumption within the ball $\mathbb{B}(0, R(\vartheta))$, for the sample size satisfying $\bar{\varepsilon}_{n,\delta}(R(\vartheta)) < \frac{1}{4}\Delta_0$, with probability $1 - \delta$, we have the bound

$$\sup_{\theta \in \mathbb{B}(0, R(\vartheta))} |F(\theta) - F_n(\theta)| \leq \frac{1}{4}\Delta_0.$$

Denote $F^* := F(\theta_1^*)$. By the smoothness condition (A) of the population-level log-likelihood, we denote $\tilde{r}_0 := \sqrt{\Delta_0/8L_1} \wedge r_0$. Then, we have that

$$\min_{j \in [M]} \inf_{\theta \in \mathbb{B}(\theta_j^*, \tilde{r}_0)} F(\theta) \geq F^* - \frac{1}{4}L_1(\tilde{r}_0)^2 = F^* - \frac{1}{4}\Delta_0.$$

Denote the set $\mathcal{Z} := \mathbb{B}(0, R(\vartheta)) \setminus \bigcup_{j \in [M]} \mathbb{B}(\theta_j^*, \tilde{r}_0)$. Applying above inequalities in conjunction with the gap condition on the log-likelihood, we find that

$$F_n(\theta) \geq F_n(\theta') + \frac{\Delta_0}{4}, \quad \forall \theta \in \bigcup_{j \in [M]} \mathbb{B}(\theta_j^*, \tilde{r}_0), \quad \theta' \in \mathcal{Z}.$$

Consequently, we obtain

$$\frac{\Pi\left(\theta \in \bigcup_{j \in [M]} \mathbb{B}(\theta_j^*, \tilde{r}_0) \mid X_1^n\right)}{\Pi(\theta \in \mathcal{Z} \mid X_1^n)} \geq \pi\left(\bigcup_{j \in [M]} \mathbb{B}(\theta_j^*, \tilde{r}_0)\right) \cdot e^{\frac{n\Delta_0}{4}}.$$

The prior mass can be lower bounded by the prior density condition and smoothness condition (B): let $j_0 \in [M]$ be an index such that $\pi(\theta_j^*) \geq \pi_0$, we have

$$\pi\left(\bigcup_{j \in [M]} \mathbb{B}(\theta_j^*, \tilde{r}_0)\right) \geq \pi(\mathbb{B}(\theta_{j_0}^*, \tilde{r}_0)) \geq \pi_0 \cdot e^{-\frac{L_2}{2}\tilde{r}_0^2} \cdot \text{Vol}(\mathbb{B}(\theta_{j_0}^*, \tilde{r}_0)) \geq \pi_0 \cdot e^{-\frac{L_2}{2}\tilde{r}_0^2} \cdot (\tilde{r}_0/\sqrt{d})^d.$$

Therefore, given the sample size satisfying the condition:

$$n \geq \frac{4}{\Delta_0} \left(\log \vartheta^{-1} + \log \pi_0^{-1} + L_2 \tilde{r}_0^2 + d \log \frac{d}{\tilde{r}_0} \right),$$

we have the lower bound:

$$\begin{aligned} \Pi(\mathbb{B}(0, R(\vartheta)))^{-1} \sum_{j=1}^M \Pi(\mathbb{B}(\theta_j^*, \tilde{r}_0) \mid X_1^n) &\geq 1 - \Pi(\mathbb{B}(0, R(\vartheta)))^{-1} \Pi(\mathcal{Z} \mid X_1^n) \\ &\geq 1 - \Pi\left(\bigcup_{j \in [M]} \mathbb{B}(\theta_j^*, \tilde{r}_0) \mid X_1^n\right)^{-1} \Pi(\mathcal{Z} \mid X_1^n) \geq 1 - \vartheta. \end{aligned} \quad (94)$$

Collecting the bounds (92), (93), and (94), we arrive at the lower bound

$$\Pi\left(\bigcup_{j \in [M]} \mathbb{B}(\theta_j^*, r_n^{(j)}) \mid X_1^n\right) \geq (1 - \vartheta)^3,$$

which completes the proof of this corollary.

E.2.1 Proof of lemma 1

For the simplicity of presentation, we make the following argument conditionally on X_1^n . The diffusion process (6) has $\Pi(\cdot \mid X_1^n)$ as its stationary distribution. Given $p \geq 4$, we take the potential function as:

$$\Phi(\theta) := \max(\|\theta\|_2 - R_\delta, 0)^p.$$

By Itô's formula, for $T \geq 0$, we have the expansion

$$\begin{aligned} \mathbb{E} [\Phi(\theta_T)] &= \underbrace{\frac{p}{2} \int_0^T \mathbb{E} [\langle \nabla F_n(\theta_t), \theta_t \rangle \cdot \max(\|\theta\|_2 - R_\delta, 0)^{p-2}] dt}_{:=I_1} \\ &\quad + \underbrace{\frac{p}{2n} \int_0^T \mathbb{E} [\langle \nabla \log \pi(\theta_t), \theta_t \rangle \cdot \max(\|\theta\|_2 - R_\delta, 0)^{p-2}] dt}_{:=I_2} \\ &\quad + \underbrace{\frac{p}{2n} \int_0^T \mathbb{E} [\left(\max(\|\theta_t\|_2 - R_\delta, 0)^2 d + \|\theta_t\|_2^2 (p-1) \right) \max(\|\theta_t\|_2 - R_\delta, 0)^{p-4}] dt}_{:=I_3}. \end{aligned}$$

By condition (37a) on the log-likelihood function, we have that $I_1 \leq 0$.

For the term I_2 , condition (37b) implies the following upper bound:

$$I_2 \leq -c_\pi \int_0^T \mathbb{E} [\|\theta_t\|_2^2 \cdot \max(\|\theta_t\|_2 - R_\delta, 0)^{p-2}] dt \leq -c_\pi \int_0^T \mathbb{E} [\Phi(\theta_t)] dt.$$

The term I_3 can be decomposed into two parts:

$$\begin{aligned} I_3 &\leq (p+d) \int_0^T \mathbb{E} [\max(\|\theta_t\|_2 - R_\delta, 0)^{p-2}] dt + pR_\delta^2 \int_0^T \mathbb{E} [\max(\|\theta_t\|_2 - R_\delta, 0)^{p-4}] dt \\ &= (p+d) \int_0^T \mathbb{E} [\Phi(\theta_t)^{\frac{p-2}{p}}] dt + pR_\delta^2 \int_0^T \mathbb{E} [\Phi(\theta_t)^{\frac{p-4}{p}}] dt \\ &\leq (p+d) \int_0^T \left(\mathbb{E} [\Phi(\theta_t)] \right)^{\frac{p-2}{p}} dt + pR_\delta^2 \int_0^T \left(\mathbb{E} [\Phi(\theta_t)] \right)^{\frac{p-4}{p}} dt. \end{aligned}$$

Denote $\psi_t := \mathbb{E} [\Phi(\theta_t)]$, we have the integral inequality

$$\begin{aligned} \psi_T &\leq \frac{p}{2n} \int_0^T \left(-c_\pi \psi_t + (p+d) \psi_t^{\frac{p-2}{p}} + pR_\delta^2 \psi_t^{\frac{p-4}{p}} \right) dt \\ &\leq \frac{p}{2n} \int_0^T \left(-c_\pi \psi_t + \frac{c_\pi}{3} \psi_t + \frac{(p+d)^{\frac{p}{2}}}{(c_\pi/3)^{\frac{p-2}{2}}} + \frac{c_\pi}{3} \psi_t + \frac{(pR_\delta)^{\frac{p}{4}}}{(c_\pi/3)^{\frac{p-4}{4}}} \right) dt \\ &\leq \frac{pc_\pi}{6n} \int_0^T \left(-\psi_t dt + \left(\frac{6(p+d)}{c_\pi} \right)^{\frac{p}{2}} + (pR_\delta)^p \right) dt. \end{aligned}$$

Note that $\psi_0 = 0$ by definition. Applying Grönwall inequality, we arrive at the bound

$$\psi_t \leq \left(\frac{6(p+d)}{c_\pi} \right)^{\frac{p}{2}} + (pR_\delta)^p, \quad \forall t \geq 0.$$

Consequently, we have the bound

$$\begin{aligned} (\mathbb{E} [\|\theta\|_2^p | X_1^n])^{1/p} &\leq R_\delta + \limsup_{t \rightarrow +\infty} (\mathbb{E} [\Phi(\theta_t)])^{1/p} \\ &\leq (p+1)R_\delta + \sqrt{\frac{6(p+d)}{c_\pi}}, \end{aligned}$$

which completes the proof of lemma 1.

E.3 A limit result

We begin with a lemma on the limiting behavior of a certain type of function. The lemma is used in the proof of theorem 2 in Appendix C.5.

Lemma 2. *Let ϕ be a non-increasing continuous function on the real line with $\phi(c) = 0$, and such that $\phi(t) \geq 0$ for all $t \in (c, \infty)$. Suppose that there exist two continuous functions $f, g : [0, +\infty) \rightarrow \mathbb{R}$ such that $\lim_{t \rightarrow +\infty} g(t)$ exists and $f(t) \leq \int_0^t \phi(g(s))ds$ for all $t \geq 0$. Under these conditions, we have $\lim_{t \rightarrow +\infty} g(t) \leq c$.*

Proof. Define the limit $A := \lim_{t \rightarrow +\infty} g(t)$, which exists according to the assumptions. We proceed via proof by contradiction. In particular, suppose that $A > c$. Based on the definition of A , for the positive constant $\varepsilon = (A - c)/2 > 0$, we can find a sufficiently large positive constant T such that $g(t) > A - \varepsilon$ for any $t \geq T$. According to the assumptions on ϕ , we obtain that

$$\delta := - \sup_{s \geq c+\varepsilon} \phi(s) < 0.$$

Therefore, for all $t > T$, we arrive at the following inequalities

$$0 \leq f(t) \leq \int_0^T \phi(g(s))ds + \int_T^t \phi(g(s))ds \leq \int_0^T \phi(g(s))ds - \delta(t - T).$$

By choosing $t = 1 + T + \delta^{-1} \int_0^T \phi(g(s))ds$, the above inequality cannot hold. This yields the desired contradiction, which completes the proof. \square

E.4 A tail bound based on truncation

We now state an upper deviation inequality based on a truncation argument. This lemma is used in appendix D.2 to prove the uniform concentration bound (28b). Consider a sequence of random variables $\{Y_i\}_{i=1}^n$ satisfying the moment bounds

$$\mathbb{E}[|Y_i|^q] \leq (aq)^{bq} \quad \text{for all } q = 1, 2, \dots \quad (95)$$

where a, b are universal constants.

Lemma 3. *Given an i.i.d. sequence of zero-mean random variables $\{Y_i\}_{i=1}^n$ satisfying the moment bounds (95), we have*

$$\mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n Y_i \geq (4a)^b \sqrt{\frac{\log 4/\delta}{n}} + \left(a \log \frac{n}{\delta}\right)^b \frac{\log 4/\delta}{n}\right) \leq \delta.$$

Proof. The proof of the lemma is a direct combination of truncation argument and Bernstein's inequality. In particular, for each $i \in [n]$, define the truncated random variable $\tilde{Y}_i := Y_i \mathbb{I}[|Y_i| \leq 3(a \log \frac{n}{\delta})^b]$. With this definition, we have

$$\begin{aligned} \mathbb{P}\left((Y_i)_{i=1}^n \neq (\tilde{Y}_i)_{i=1}^n\right) &= \mathbb{P}\left(\max_{1 \leq i \leq n} |Y_i| > 3\left(a \log \frac{n}{\delta}\right)^b\right) \\ &\leq n \mathbb{P}\left(|Y_i| > 3\left(a \log \frac{n}{\delta}\right)^b\right) \leq \frac{\delta}{2}. \end{aligned}$$

Therefore, it is sufficient to study a concentration behavior of the quantity $\sum_{i=1}^n \tilde{Y}_i$. Invoking Bernstein's inequality [6], we obtain that

$$\mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n \tilde{Y}_i \geq \varepsilon\right) \leq 2 \exp\left(-\frac{n\varepsilon^2}{2(2a)^{2b} + \frac{2}{3}\varepsilon \cdot 3(a \log \frac{n}{\delta})^b}\right).$$

In order to make the RHS of the above inequality less than $\frac{\delta}{2}$, it suffices to set

$$\varepsilon = (4a)^b \sqrt{\frac{\log(4/\delta)}{n}} + \left(a \log \frac{n}{\delta}\right)^b \frac{\log(4/\delta)}{n}.$$

Collecting all of the above inequalities yields the claim. \square

References

- [1] Radoslaw Adamczak. A tail inequality for suprema of unbounded empirical processes with applications to markov chains. *Electronic Journal of Probability*, 13:1000–1034, 2008.
- [2] Dominique Bakry, Franck Barthe, Patrick Cattiaux, and Arnaud Guillin. A simple proof of the Poincaré inequality for a large class of probability measures. *Electronic Communications in Probability*, 13:60–66, 2008.
- [3] S. Balakrishnan, M. J. Wainwright, and B. Yu. Statistical guarantees for the EM algorithm: From population to sample-based analysis. *Annals of Statistics*, 45:77–120, 2017.
- [4] A. Barron, M. Schervish, and L. Wasserman. The consistency of posterior distributions in nonparametric problems. *Ann. Statist.*, 27:536–561, 1999.
- [5] A. Bhattacharya, D. Pati, , and D. B. Dunson. Anisotropic function estimation using multi-bandwidth Gaussian processes. *Annals of Statistics*, 42:352–381, 2014.
- [6] S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2016.
- [7] R. J. Carroll and P. Hall. Optimal rates of convergence for deconvolving a density. *Journal of American Statistical Association*, 83:1184–1186, 1988.
- [8] J. Chen. Optimal rate of convergence for finite mixture models. *Annals of Statistics*, 23(1):221–233, 1995.
- [9] Stéphane Chrétien, Mihai Cucuringu, Guillaume Lecué, and Lucie Neirac. Learning with semi-definite programming: statistical bounds based on fixed point analysis and excess risk curvature. *Journal of Machine Learning Research*, 22(230), 2021.
- [10] A. S. Dalalyan. Theoretical guarantees for approximate sampling from smooth and log-concave densities. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(3):651–676, 2017.
- [11] R. de Jonge and J. H. van Zanten. Adaptive nonparametric Bayesian inference using location-scale mixture priors. *Annals of Statistics*, 38:3300–3320, 2010.

- [12] J. L. Doob. Application of the theory of martingales. *Actes du Colloque International Le Calcul des Probabilités et ses applications (Lyon, 28 Juin– 3 Juillet, 1948)*, pages 23–27, 1949.
- [13] Alain Durmus and Eric Moulines. Nonasymptotic convergence analysis for the unadjusted langevin algorithm. *The Annals of Applied Probability*, 27(3):1551–1587, 2017.
- [14] Alain Durmus and Eric Moulines. High-dimensional bayesian inference via the unadjusted langevin algorithm. *Bernoulli*, 25(4A):2854–2882, 2019.
- [15] R. Dwivedi, N. Ho, K. Khamaru, M. J. Wainwright, M. I. Jordan, and B. Yu. Singularity, misspecification, and the convergence rate of EM. *arXiv preprint arXiv:1810.00828*, 2018.
- [16] D. A. Freedman. On the asymptotic behavior of Bayes' estimates in the discrete case. *Annals of Statistics*, 34:1386–1403, 1963.
- [17] D. A. Freedman. On the asymptotic behavior of Bayes' estimates in the discrete case.II. *Annals of Statistics*, 36:454–456, 1965.
- [18] Sébastien Gadat, Fabien Panloup, and Clément Pellegrini. On the cost of bayesian posterior mean strategy for log-concave models. *arXiv preprint arXiv:2010.06420*, 2020.
- [19] C. Gao and H. H. Zhou. Rate exact Bayesian adaptation with modified block priors. *Annals of Statistics*, 44:318–345, 2016.
- [20] S. Ghosal, J. K. Ghosh, and A. van der Vaart. Convergence rates of posterior distributions. *Annals of Statistics*, 28:500–531, 2000.
- [21] S. Ghosal and A. van der Vaart. Entropies and rates of convergence for maximum likelihood and bayes estimation for mixtures of normal densities. *Annals of Statistics*, 29:1233–1263, 2001.
- [22] S. Ghosal and A. van der Vaart. Posterior convergence rates of Dirichlet mixtures at smooth densities. *Annals of Statistics*, 35:697–723, 2007.
- [23] Leonard Gross. Logarithmic Sobolev inequalities. *American Journal of Mathematics*, 97(4):1061–1083, 1975.
- [24] Nhat Ho, Koulik Khamaru, Raaz Dwivedi, Martin J Wainwright, Michael I Jordan, and Bin Yu. Instability, computational efficiency and statistical accuracy. *arXiv preprint arXiv:2005.11411*, 2020.
- [25] Nhat Ho and XuanLong Nguyen. Singularity structures and impacts on parameter estimation in finite mixtures of distributions. *SIAM Journal on Mathematics of Data Science*, 1(4):730–758, 2019.
- [26] Io Ao Ibragimov and RZ Khasminskii. Asymptotic theory of estimation. *Nauka, Mosow*, 1:979, 1979.
- [27] H. Ishwaran, L. F. James, and J. Sun. Bayesian model selection in finite mixtures by marginal density decompositions. *Journal of the American Statistical Association*, 96:1316–1332, 2001.

[28] B. J. K. Kleijn and A. W. van der Vaart. Misspecification in infinite-dimensional Bayesian statistics. *Annals of Statistics*, 34:837–877, 2006.

[29] Bas JK Kleijn and Aad W van der Vaart. The bernstein-von-mises theorem under misspecification. *Electronic Journal of Statistics*, 6:354–381, 2012.

[30] B. Lindsay. *Mixture Models: Theory, Geometry and Applications*. In NSF-CBMS Regional Conference Series in Probability and Statistics. IMS, Hayward, CA., 1995.

[31] Po-Ling Loh and Martin J Wainwright. Regularized M-estimators with nonconvexity: Statistical and algorithmic theory for local optima. *Advances in Neural Information Processing Systems*, 26, 2013.

[32] Cong Ma, Kaizheng Wang, Yuejie Chi, and Yuxin Chen. Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion, and blind deconvolution. *Foundations of Computational Mathematics*, 20(3):451–632, 2020.

[33] Eric Mazumdar, Aldo Pacchiano, Yian Ma, Michael Jordan, and Peter Bartlett. On approximate thompson sampling with langevin algorithms. In *International Conference on Machine Learning*, pages 6797–6807. PMLR, 2020.

[34] P. McCullagh and J. A. Nelder. *Generalized Linear Models*. Chapman and Hall/CRC, 1989.

[35] W. Mou, N. Ho, M. J. Wainwright, P. L. Bartlett, and M. I. Jordan. Supplementary material to “A Diffusion Process Perspective on Posterior Contraction Rates for Parameters”, 2019. DOI: [COMPLETED BY TYPESETTER].

[36] X. Nguyen. Convergence of latent mixing measures in finite and infinite mixture models. *Annals of Statistics*, 4(1):370–400, 2013.

[37] Dmitrii M Ostrovskii and Francis Bach. Finite-sample analysis of M-estimators using self-concordance. *Electronic Journal of Statistics*, 15(1):326–391, 2021.

[38] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4):838–855, 1992.

[39] D. Revuz and M. Yor. *Continuous Martingales and Brownian Motion*, volume 293. Springer-Verlag, third edition, 1999.

[40] Hannes Risken. *The Fokker-Planck Equation*. Springer, 1996.

[41] Gareth O Roberts and Richard L Tweedie. Exponential convergence of langevin distributions and their discrete approximations. *Bernoulli*, pages 341–363, 1996.

[42] J. Rousseau. Rates of convergence for the posterior distributions of mixtures of Beta and adaptive nonparametric estimation of the density. *Annals of Statistics*, 38:146–180, 2010.

[43] J. Rousseau and K. Mengersen. Asymptotic behaviour of the posterior distribution in overfitted mixture models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73:689–710, 2011.

[44] L. Schwartz. On Bayes procedures. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 4:10–26, 1965.

[45] W. Shen, S. R. Tokdar, , and S. Ghosal. Adaptive Bayesian multivariate density estimation with Dirichlet mixtures. *Biometrika*, 100:623–640, 2013.

[46] X. Shen and L. Wasserman. Rates of convergence of posterior distributions. *Annals of Statistics*, 29:687–714, 2001.

[47] Bin Shi, Simon S Du, Michael I Jordan, and Weijie J Su. Understanding the acceleration phenomenon via high-resolution differential equations. *Mathematical Programming*, pages 1–70, 2021.

[48] Vladimir Spokoiny. Parametric estimation. finite sample theory. *The Annals of Statistics*, 40(6):2877–2909, 2012.

[49] Weijie Su, Stephen Boyd, and Emmanuel J Candes. A differential equation for modeling nesterov’s accelerated gradient method: Theory and insights. *Journal of Machine Learning Research*, 17:1–43, 2016.

[50] Michel Talagrand. Transportation cost for gaussian and other product measures. *Geometric & Functional Analysis GAFA*, 6(3):587–600, 1996.

[51] S. van de Geer. *Empirical Processes in M-estimation*. Cambridge University Press, 2000.

[52] A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.

[53] A. W. van der Vaart and J. Wellner. *Weak Convergence and Empirical Processes*. Springer-Verlag, New York, NY, 1996.

[54] A. W. van der Vaart and J. A. Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer-Verlag, New York, NY, 2000.

[55] M. J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge University Press, 2019.

[56] S. Walker. On sufficient conditions for Bayesian consistency. *Annals of Statistics*, 90:482–488, 2003.

[57] S. Walker. New approaches to Bayesian consistency. *Annals of Statistics*, 32:2028–2043, 2004.

[58] S. G. Walker, A. Lijoi, and I. Prunster. On rates of convergence for posterior distributions in infinite-dimensional models. *Annals of Statistics*, 35:738–746, 2007.

[59] Y. Yang and D. B. Dunson. Bayesian manifold regression. *Annals of Statistics*, 44:876–905, 2016.

[60] Y. Yang and S. T. Tokdar. Minimax-optimal nonparametric regression in high dimensions. *Annals of Statistics*, 43:652–674, 2015.