

On the Approximation Accuracy of Gaussian Variational Inference

Anya Katsevich
 akatsevi@mit.edu

Philippe Rigollet
 rigollet@math.mit.edu

January 9, 2024

Abstract

The main computational challenge in Bayesian inference is to compute integrals against a high-dimensional posterior distribution. In the past decades, variational inference (VI) has emerged as a tractable approximation to these integrals, and a viable alternative to the more established paradigm of Markov Chain Monte Carlo. However, little is known about the approximation accuracy of VI. In this work, we bound the TV error and the mean and covariance approximation error of Gaussian VI in terms of dimension and sample size. Our error analysis relies on a Hermite series expansion of the log posterior whose first terms are precisely cancelled out by the first order optimality conditions associated to the Gaussian VI optimization problem.

1 Introduction

A central challenge in Bayesian inference is to compute integrals against a posterior distribution π on \mathbb{R}^d . The classical approach is to *sample* the posterior using Markov Chain Monte Carlo (MCMC), in which a Markov chain designed to converge to π is simulated for sufficiently long time. However, MCMC can be expensive, and it is notoriously difficult to identify clear-cut stopping criteria for the algorithm [6]. An alternative, often computationally cheaper, approach is *variational inference* (VI) [4]. The idea of VI is to find, among all measures in a certain parameterized family \mathcal{P} , the closest measure to π . While various notions of proximity have been proposed since the introduction of VI [7, 8], we employ here KL divergence, which is by far the most common choice. Typically, the family \mathcal{P} is selected so that integrals against measures in the family are either readily available or else easily computable. In this work, we consider the family of Gaussian distributions. We define

$$\hat{\pi} = \mathcal{N}(\hat{m}, \hat{S}) \in \operatorname{argmin}_{p \in \mathcal{P}_{\text{Gauss}}} \text{KL}(p \parallel \pi), \quad (1.1)$$

where $\mathcal{P}_{\text{Gauss}}$ denotes the family of non-degenerate Gaussian distributions on \mathbb{R}^d . With the Gaussian approximation $\hat{\pi}$ in hand, we can now approximate $\int g d\pi$ using $\int f d\hat{\pi}$. For certain g (e.g. polynomials), the latter integral can be computed in closed form in terms of \hat{m} and \hat{S} . Otherwise, Gaussian sampling can be employed, which is much cheaper than MCMC for the original measure.

A key difference between MCMC and VI is that unbiased MCMC algorithms yield arbitrarily accurate samples from π if they are run for long enough. On the other hand, the output of a perfect VI algorithm is $\hat{\pi}$, which is itself only an approximation to π . Therefore, a fundamental question in VI is to understand the quality of the approximation $\hat{\pi} \approx \pi$. In this work, we bound the approximation error $\int g d\pi - \int g d\hat{\pi}$ for a wide class of functions g , as well as several particular

metrics of interest: the TV error $\text{TV}(\pi, \hat{\pi})$ and the mean and covariance errors $\|\hat{m} - \bar{m}\|$ and $\|\hat{S} - \Sigma\|$, where \bar{m}, Σ are the true mean and covariance of π , respectively.

Of course, we cannot expect an arbitrary, potentially multimodal π to be well-approximated by a Gaussian distribution. In the setting of Bayesian inference, however, the Bernstein-von Mises theorem (BvM) guarantees that under certain regularity conditions, a posterior distribution converges to a Gaussian density in the limit of large sample size [27, Chapter 10]. To understand why this is the case, consider a generic posterior $\pi = \pi_n$ with density of the form

$$\pi_n(\theta \mid x_{1:n}) \propto \nu(\theta) \prod_{i=1}^n p_\theta(x_i) \quad (1.2)$$

Here, ν is the prior, p_θ is the model density, and $x_{1:n} = x_1, \dots, x_n$ are i.i.d observations. Provided ν and p_θ are everywhere positive, we can write π_n as

$$\pi_n(\theta) \propto e^{-nv_n(\theta)}, \quad v_n(\theta) := -\frac{1}{n} \sum_{i=1}^n \log p_\theta(x_i) - \frac{1}{n} \log \nu(\theta). \quad (1.3)$$

As n increases, we expect that v_n approaches v_∞ , the n -independent negative population log likelihood, a.k.a. the cross-entropy. Therefore, if n is large and v_∞ has a unique global minimum, then π_n will place most of its mass in a neighborhood of this point. In other words, π_n is effectively unimodal, and hence a Gaussian approximation is reasonable in this case. This reasoning drives a second strategy to approximate posterior, whereby π_n is a Gaussian centered at the global minimizer of v_n . While the accuracy of the Laplace approximation has been the subject of many works, including recent ones (see below), no comparable results exist for Gaussian VI. The main goal of this paper is to establish such approximation error bounds.

Main contributions In our above heuristic justification of a Gaussian approximation to $\pi = \pi_n$, we considered a sequence of functions v_n converging to some v_∞ as $n \rightarrow \infty$. However, our results do not require this setup. Instead, we simply consider a target measure of the form $\pi \propto e^{-nv}$ with $v \in C^4(\mathbb{R}^d)$, and we impose a requirement on the triple (v, d, n) . Namely, we assume $a_3(v)d/\sqrt{n}$ is smaller than a certain constant and $a_3(v)d/\sqrt{n} + a_4(v)d^2/n \leq 1$, where $a_3(v)$ and $a_4(v)$ bound the size of the third and fourth derivatives of v (see Assumption 2 for the precise definition). Our error bounds are in line with this condition on (v, d, n) , in the sense that the errors are small provided $a_3(v)d/\sqrt{n} + a_4(v)d^2/n$ is small. Up to this restriction, v is free to depend on both d and n .

Assuming also that v has a unique global minimizer m^* , that $\|\nabla^3 v\|$ and $\|\nabla^4 v\|$ grow at most polynomially away from m^* , and that v grows at least linearly away from m^* , we show the following error bound for a wide class of functions g :

$$\left| \int g d\pi - \int g d\hat{\pi} \right| \lesssim_v \sqrt{1 \vee \text{Var}_{\hat{\pi}}(g)} \begin{cases} d/\sqrt{n}, \\ (d/\sqrt{n})^2, & g \text{ even about } \hat{m}, \\ (d/\sqrt{n})^3, & g \text{ linear.} \end{cases} \quad (1.4)$$

We have hidden the v -dependent factors $a_3(v), a_4(v)$ within the \lesssim_v symbol for maximal clarity, though these factors may in principle affect the order of magnitude of the upper bound. See Theorem 2.1 for the precise statement of the result and observe that the mean a posteriori, i.e. Bayes estimator, is approximated at the fastest of these three rates. The second bound in (1.4) is enabled by a leading order expansion of the error. Namely, we show that

$$\int g d\pi - \int g d\hat{\pi} = \int g Q d\hat{\pi} + \text{Rem}_g, \quad |\text{Rem}_g| \lesssim_v \sqrt{1 \vee \text{Var}_{\hat{\pi}}(g)} \left(\frac{d}{\sqrt{n}} \right)^2, \quad (1.5)$$

where Q is a certain cubic polynomial which is odd about \hat{m} . Thus if g is even about \hat{m} , then the leading order term $\int gQd\hat{\pi}$ vanishes. Similarly, the third bound in (1.4) is enabled by a higher order expansion of the error. We note that having derived the leading order term not only provides valuable insight into the error, but also opens the door to *correcting* the Gaussian measure $\hat{\pi}$ to get the more accurate approximation $(1+Q)\hat{\pi}$. However, we do not delve into this possibility here and postpone it for future research. The benefit of improving a Gaussian posterior approximation via a multiplicative correction was already studied in [19], but for a different, *Laplace* approximation to the posterior (defined in (1.8)).

This main result (1.4) gives us a bound on $\int gd\pi - \int gd\hat{\pi}$ which is tailored to the observable g . We can also apply the result to derive error bounds for common Bayesian quantities of interest. For example, the following TV bound quantifies the error in using $\hat{\pi}$ to construct approximate credible sets for π :

$$\text{TV}(\pi, \hat{\pi}) \lesssim_v \frac{d}{\sqrt{n}}. \quad (1.6)$$

We also bound the mean and covariance approximation errors. The posterior mean holds particular significance in Bayesian inference, since it is the Bayes-optimal estimator of the unknown parameter with respect to the squared error risk. The posterior covariance is another measure of the uncertainty in the posterior which can be used in addition to credible sets. To account for the vanishing variance (on the order of n^{-1}) of $\pi \propto e^{-nv}$, we rescale the below approximation errors appropriately:

$$\sqrt{n}\|\hat{m} - \bar{m}\| \lesssim_v \left(\frac{d}{\sqrt{n}}\right)^3, \quad n\|\hat{S} - \Sigma\| \lesssim_v \left(\frac{d}{\sqrt{n}}\right)^2. \quad (1.7)$$

Together, the bounds (1.4), (1.6), and (1.7) give us a comprehensive understanding of the error induced by using $\hat{\pi}$ as a proxy for π in computing all of the relevant quantities of interest in Bayesian inference. In Section 3, we apply the general theory to formulate high-probability bounds on the VI approximation error for a posterior arising from logistic regression with Gaussian design.

Laplace approximation as benchmark for comparison As mentioned above, the Laplace method is another Gaussian approximation to π that is widespread in practice for its computational simplicity. We use it as a benchmark to put the above error bounds into context. The Laplace approximation to $\pi \propto e^{-nv}$ is given by

$$\pi \approx \mathcal{N}(m^*, (n\nabla^2 v(m^*))^{-1}), \quad (1.8)$$

where m^* is the global minimizer of v . This approximation simply replaces v by its second order Taylor expansion around m^* . In recent work [19], the first author shows that the TV error and covariance error of the Laplace approximation are of the same order as that of VI, in terms of both the d and n dependence. However, the VI mean error is two orders of magnitude smaller than that of Laplace: $(d/\sqrt{n})^3$ for VI vs. d/\sqrt{n} for Laplace. See Section 2.3 for a more detailed comparison of the approximation accuracy of the two methods.

We emphasize that our aim in the present work is *not* to make a value judgment about Gaussian VI in comparison to the Laplace approximation. Indeed, the accuracy of a method is not the only axis on which to judge its overall performance. An overall comparison of the two methods is a delicate undertaking that involves computational considerations and varies according to downstream statistical tasks. It is beyond the scope of the present work to produce a full comparison analysis but the approximation results should certainly be useful to consider in certain scenarios.

Instead, the main contribution of our work is to supply precise statements of Gaussian VI's approximation accuracy, an understanding of which has been virtually nonexistent, lagging far

behind algorithmic developments in the literature. We compare VI to Laplace simply because the Laplace approximation is another popular approach to ease posterior computation via a Gaussian fit.

Another way in which the Laplace approximation is relevant to VI is that the proof techniques developed by the first author in [18] to bound the Laplace approximation error also play a crucial role in the present work. Specifically, notice that the bounds (1.4), (1.6), and (1.7) are all small provided $d^2 \ll n$. This dependence on dimension is an improvement over the condition $d^3 \ll n$ required in a long line of work on the BvM and the Laplace approximation accuracy [13, 24, 23, 10, 15, 25]. Since then, the recent works [16] and [18] have developed new proof techniques which reduced the dependence to $d^2 \ll n$, though only the latter work's proof techniques are flexible enough to derive leading order asymptotics in terms of d/\sqrt{n} .

We employ these same techniques in the present work. In addition, our proof is based on crucial new insights into why minimization of the KL objective (1.1) leads to good approximation properties. We now briefly summarize these insights.

First-order optimality conditions and Hermite series expansion Our analysis of the approximation accuracy of Gaussian VI rests on a remarkable interaction between first-order optimality conditions and a Hermite series expansion of the potential v .

Hereafter, we replace θ by x and let $V = nv$, $\pi \propto e^{-V}$. The focal point of this work are the first order optimality equations for the minimization (1.1):

$$\nabla_{m,S} \text{KL}(\mathcal{N}(m, S) \parallel \pi) \big|_{(m,S)=(\hat{m},\hat{S})} = 0.$$

This is also equivalent to setting the Bures-Wasserstein gradient of $\text{KL}(p \parallel \pi)$ to zero at $p = \mathcal{N}(\hat{m}, \hat{S})$ as in [20]. Explicitly, we obtain that $(m, S) = (\hat{m}, \hat{S})$ is a solution to

$$\mathbb{E}[\nabla V(m + S^{1/2}Z)] = 0, \quad \mathbb{E}[\nabla^2 V(m + S^{1/2}Z)] = S^{-1}, \quad (1.9)$$

where $Z \sim \mathcal{N}(0, I_d)$ and $S^{1/2}$ is the positive definite symmetric square root of S ; see [20] for this calculation. In some sense, the fact that $\hat{\pi} = \mathcal{N}(\hat{m}, \hat{S})$ minimizes the KL divergence to π does not explain why $\hat{\pi}$ approximates π well. Rather, the true reason has to do with properties of solutions to the fixed point equations (1.9).

To see why, consider the function $V_0(x) = V(\hat{m} + \hat{S}^{1/2}x)$. If $\pi \propto e^{-V}$ is close to the density of $\mathcal{N}(\hat{m}, \hat{S})$, then $\rho \propto e^{-V_0}$ should be close to the density of $\mathcal{N}(0, I_d)$. In other words, we should have that $V_0(x) \approx \text{const.} + \|x\|^2/2$. This is ensured by the first order optimality equations (1.9). Indeed, note that (1.9) can be written in terms of V_0 as

$$\mathbb{E}[\nabla V_0(Z)] = 0, \quad \mathbb{E}[\nabla^2 V_0(Z)] = I_d. \quad (1.10)$$

As we explain in Section 4.2, the equations (1.10) set the first and second order coefficients in the Hermite series expansion of V_0 to 0 and I_d , respectively. As a result, $V_0(x) - \|x\|^2/2 = \text{const.} + r_3(x)$, where r_3 is a Hermite series containing only third and higher order Hermite polynomials. The accuracy of the Gaussian VI approximation for integrating linear and quadratic functions stems from the fact that the Hermite remainder r_3 is of order $r_3 \sim 1/\sqrt{n}$, and the fact that r_3 is *orthogonal to linear and quadratic functions with respect to the Gaussian measure*. See Section 4.2 for a high-level summary of this Hermite series based error analysis.

Related Work The literature on VI can be roughly divided into statistical and algorithmic works. Works on the statistical side have focused on the contraction of variational posteriors around a ground truth parameter in the large n (sample size) regime. (We use “variational posterior” as an abbreviation for variational approximation to the posterior.) For example, [28] prove an analogue of the Bernstein-von Mises theorem for the variational posterior, [30] study the contraction rate of the variational posterior around the ground truth in a nonparametric setting, and [2] study the contraction rate of variational approximations to tempered posteriors, in high dimensions.

A key difference between these works and ours is that here, we determine how well the statistics of the variational posterior match those of the posterior itself, rather than those of a limiting ($n \rightarrow \infty$) distribution. We are only aware of one other work studying the problem of quantifying posterior approximation accuracy of VI. In [14], the authors consider using the mean \hat{m} of the mean-field variational posterior as a proxy for the maximum likelihood estimate in a Bayesian latent variable model. They show that $\sqrt{n}\|\hat{m} - \text{MLE}\| \lesssim 1/n^{1/4}$.

On the algorithmic side of Gaussian VI, we refer the reader to the works [20, 12] and references therein. These works complement our analysis in that they provides rigorous convergence guarantees for several algorithms that solves the optimization problem (1.1). See also the discussion in Section 2.3.

Next, we mention a closely related approximation method called expectation propagation (EP). In Gaussian EP, one attempts to minimize the KL divergence $\text{KL}(\pi \parallel p)$ over all Gaussian measures p . (Note that the order of π and p has been switched relative to (1.1).) Minimizing this objective amounts to matching the first and second moments of p to those of π . It is not possible to minimize this objective directly without resorting to computing integrals against π . So instead, EP works by iteratively constructing hybrids between the Gaussian distribution and π , and matching the moments of the hybrid (which are easier to compute) to those of π . The works [9] and [11] studied the mean and covariance approximation accuracy of EP with respect to n . Interestingly enough, the authors showed EP has the exact same mean and covariance accuracy as shown here for traditional Gaussian VI: $\|\hat{m} - \bar{m}\| \lesssim n^{-2}$ and $\|\hat{S} - \Sigma\| \lesssim n^{-2}$ (compare to (1.7)). Whether this is a coincidence or a result of a deeper connection between the two methods is an interesting avenue to explore.

Finally, we note that [29] also uses Hermite polynomials to provide an approximation to the posterior. The approximation, given by an Edgeworth series, is effectively a perturbation of a Gaussian measure. The coefficients of the series depend on the cumulants of the posterior distribution. Thus writing down the series requires computing cumulants, which come from integrals against the posterior. Therefore, the expansion cannot itself be directly used as a tool to help approximate integrals against the posterior, which is the goal in our work.

Organization of the paper

The rest of the paper is organized as follows. In Section 2 we state our assumptions and main results on the Gaussian VI approximation accuracy. In Section 3, we apply the general theory to a logistic regression posterior. In Section 4, we outline the proof of the main theorems. Finally in Section 5, we prove the existence and uniqueness of solutions to the first-order optimality conditions (1.9). Omitted proofs can be found in the Appendix.

Notation

The notation $a \lesssim b$ means there is constant $C > 0$ such that $a \leq Cb$. Unless indicated otherwise, the suppressed constant is absolute. We let γ denote the density of the standard Gaussian distribution $\mathcal{N}(0, I_d)$ in d dimensions, and we write either $\int f d\gamma$ or $\mathbb{E}[f(Z)]$ or $\gamma(f)$ for the expectation of f

under γ . We write Z to denote a standard multivariate Gaussian random variable $Z \sim \gamma$ in \mathbb{R}^d . We let $\text{Var}_{\hat{\pi}}(f) = \int (f - \hat{\pi}(f))^2 d\hat{\pi}$, and $\|f\|_2 = (\int f^2 d\gamma)^{1/2}$.

A tensor T of order k is an array $T = (T_{i_1 i_2 \dots i_k})_{i_1, \dots, i_k=1}^d$. For two order k tensors T and S we let $\langle T, S \rangle$ be the entrywise inner product. We say T is symmetric if $T_{i_1 \dots i_k} = T_{j_1 \dots j_k}$, for all permutations $j_1 \dots j_k$ of $i_1 \dots i_k$.

Let H be a symmetric positive definite matrix. For a vector $x \in \mathbb{R}^d$, we let $\|x\|_H$ denote $\|x\|_H = \sqrt{x^T H x}$. For an order $k \geq 2$ symmetric tensor T , we define the H -weighted operator norm of T to be

$$\|T\|_H := \sup_{\|u\|_H=1} \langle T, u^{\otimes k} \rangle. \quad (1.11)$$

By Theorem 2.1 of [31], for symmetric tensors, the definition (1.11) coincides with the standard definition of operator norm:

$$\sup_{\|u_1\|_H=\dots=\|u_k\|_H=1} \langle T, u_1 \otimes \dots \otimes u_k \rangle = \|T\|_H = \sup_{\|u\|_H=1} \langle T, u^{\otimes k} \rangle.$$

When $H = I_d$, the norm $\|T\|_{I_d}$ is the regular operator norm, and in this case we omit the subscript.

2 Statement of Main Results

In light of the centrality of the fixed point equations (1.9), we begin the section by redefining (\hat{m}, \hat{S}) as solutions to (1.9) rather than minimizers of the KL divergence objective (1.1). These definitions diverge only in the case that V is not strongly convex. Indeed, if V is strongly convex then $\text{KL}(\cdot \parallel \pi)$ is strongly geodesically convex in the submanifold of Gaussian distributions; see, e.g., [20]. Therefore, in this case, $\text{KL}(\cdot \parallel \pi)$ admits a unique minimizer $\hat{\pi}$ over this submanifold, corresponding to a unique solution $(\hat{m}, \hat{S}) \in \mathbb{R}^d \times \mathbf{S}_{++}^d$ to (1.9). In general, however, there are spurious solutions (m, S) to the first order equations (1.9) that may have poor approximation properties. To see this, consider the equations in the following form, recalling that $v = V/n$:

$$\mathbb{E}[\nabla v(m + S^{1/2}Z)] = 0, \quad S \mathbb{E}[\nabla^2 v(m + S^{1/2}Z)] = \frac{1}{n} I_d. \quad (2.1)$$

Let $x \neq m_*$ be a critical point of v , that is, $\nabla v(x) = 0$, and consider the pair $(m, S) = (x, 0)$. For this (m, S) we have

$$\mathbb{E}[\nabla v(m + S^{1/2}Z)] = \nabla v(x) = 0, \quad S \mathbb{E}[\nabla^2 v(m + S^{1/2}Z)] = 0 \approx \frac{1}{n} I_d.$$

Thus $(x, 0)$ is an approximate solution to (2.1), and by continuity, we expect that there is an exact solution nearby. In other words, to each critical point x of v is associated a solution $(m, S) \approx (x, 0)$ of (2.1). The solution (m, S) of (2.1) which we are interested in, then, is the one near $(m_*, 0)$. Lemma 2.1 below formalizes this intuition; we show there is a unique solution (m, S) to (1.9) in the set

$$\mathcal{R}_V = \left\{ (m, S) \in \mathbb{R}^d \times \mathbf{S}_{++}^d : S \preceq 2H_V^{-1}, \|H_V^{1/2} S^{1/2}\|^2 + \|H_V^{1/2} (m - m_*)\|^2 \leq 8 \right\}, \quad (2.2)$$

where $H_V = \nabla^2 V(m_*) = n \nabla^2 v(m_*)$. Note that due to the scaling of H_V with n , the set \mathcal{R}_V is a small neighborhood of $(m_*, 0)$. We call this unique solution (m, S) in \mathcal{R}_V the ‘‘canonical’’ solution of (1.9). We expect the Gaussian distribution corresponding to this canonical solution to be the minimizer of (1.1), although we have not proved this. Regardless of whether it is true,

we will redefine (\hat{m}, \hat{S}) to denote the canonical solution. Indeed, whether or not $\hat{\pi} = \mathcal{N}(\hat{m}, \hat{S})$ actually minimizes the KL divergence or is only a local minimizer is immaterial for the purpose of approximating π .

In the rest of this section, we state our assumptions on v , a lemma guaranteeing a canonical solution \hat{m}, \hat{S} to (1.9), and our main results on the approximation accuracy of Gaussian VI.

2.1 Assumptions on the potential

Assumption 1 (Regularity and unique strict minimum). The potential $v \in C^4$, with unique global minimizer $x = m^*$, and $H_v = \nabla^2 v(m^*) \succ 0$.

Assumption 2 (Polynomial growth of $\|\nabla^k v\|_{H_v}$, $k = 3, 4$). There is some $q, a_3, a_4 > 0$ such that

$$\|\nabla^k v(x)\|_{H_v} \leq a_k \left(1 \vee \sqrt{n/d} \|x - m^*\|_{H_v}\right)^q \quad \forall x \in \mathbb{R}^d, \quad k = 3, 4, \quad (2.3)$$

$$a_3 \frac{d}{\sqrt{n}} \leq C(q), \quad a_3 \frac{d}{\sqrt{n}} + a_4 \frac{d^2}{n} \leq 1 \quad (2.4)$$

for a certain constant $C(q) < 1$ depending only on q .

The $\sqrt{n/d}$ factor in (2.3) is there to “zoom in” on the $\sqrt{d/n}$ scale, which is the scale at which the measure $\pi \propto e^{-nv}$ concentrates.

Assumption 3 (Lower bound on growth of v). There exists $c_0 > 0$ such that

$$c_0 \sqrt{d/n} \|x - m^*\|_{H_v} \leq v(x) - v(m^*) \quad \forall x : \|x - m^*\|_{H_v} \geq \frac{1}{2} \sqrt{d/n}. \quad (2.5)$$

See Section 2.3 for a discussion of the assumptions.

2.2 Main Results

We are now in a position to state our main results. In this section, we use the notation $\epsilon = d/\sqrt{n}$. First, we characterize the Gaussian VI parameters (\hat{m}, \hat{S}) :

Lemma 2.1. *Let Assumptions 1 and 2 be satisfied and let $H_V = \nabla^2 V(m_*) = n \nabla^2 v(m_*)$. Then there exists a unique $(m, S) = (\hat{m}, \hat{S})$ in the set*

$$\mathcal{R}_V = \left\{ (m, S) \in \mathbb{R}^d \times \mathbf{S}_{++}^d : S \preceq 2H_V^{-1}, \|H_V^{1/2} S^{1/2}\|^2 + \|H_V^{1/2} (m - m_*)\|^2 \leq 8 \right\}$$

which solves

$$\mathbb{E}[\nabla V(m + S^{1/2} Z)] = 0, \quad \mathbb{E}[\nabla^2 V(m + S^{1/2} Z)] = S^{-1}.$$

Moreover, \hat{S} satisfies

$$\frac{2}{3} H_V^{-1} \preceq \hat{S} \preceq 2H_V^{-1}. \quad (2.6)$$

For the proof of this lemma, see Section 5. We now state our bounds on the error in using $\hat{\pi} = \mathcal{N}(\hat{m}, \hat{S})$ to approximate expectations under the true measure π . The main theorem is the following.

Theorem 2.1. *Suppose v satisfies Assumptions 1, 2, and 3, and let \hat{m}, \hat{S} be as in Lemma 2.1. Let g be a function satisfying the following inequality for some $R_g > 0$:*

$$|g(x) - \int g d\hat{\pi}| \leq \exp \left(c_0 \sqrt{d} \|x - \hat{m}\|_{\hat{S}^{-1}} / 4 \right), \quad \forall x : \|x - \hat{m}\|_{\hat{S}^{-1}} \geq R_g \sqrt{d}. \quad (2.7)$$

Then

$$\left| \int g d\pi - \int g d\hat{\pi} \right| \lesssim \left(1 + \text{Var}_{\hat{\pi}}(g)^{\frac{1}{2}} \right) (a_3\epsilon + a_4\epsilon^2), \quad (2.8)$$

If additionally, g is even about \hat{m} , then

$$\left| \int g d\pi - \int g d\hat{\pi} \right| \lesssim \left(1 + \text{Var}_{\hat{\pi}}(g)^{\frac{1}{2}} \right) (a_3^2 + a_4)\epsilon^2. \quad (2.9)$$

Finally, if g is linear, then

$$\left| \int g d\pi - \int g d\hat{\pi} \right| \lesssim \left(1 + \text{Var}_{\hat{\pi}}(g)^{\frac{1}{2}} \right) ((a_3^3 + a_3 a_4) \epsilon^3 + a_4^2 \epsilon^4). \quad (2.10)$$

In all the above bounds, the suppressed constant is an increasing function of q , c_0^{-1} , and R_g .

The improved bound (2.9) actually holds for a somewhat wider class of functions g ; see Remark 2.1. For a discussion of the condition (2.7), see Section 2.3.

Theorem 2.2 (Leading order term in VI approximation error). *Suppose v satisfies Assumptions 1, 2, and 3. Define the function*

$$Q(x) = \left\langle \mathbb{E}_{X \sim \hat{\pi}} [\nabla^3 V(X)], \frac{1}{2}(x - \hat{m}) \otimes \hat{S} - \frac{1}{6}(x - \hat{m})^{\otimes 3} \right\rangle. \quad (2.11)$$

If g satisfies (2.7), then

$$\left| \int g d\pi - \int g d\hat{\pi} - \int g Q d\hat{\pi} \right| \lesssim \left(1 + \text{Var}_{\hat{\pi}}(g)^{\frac{1}{2}} \right) (a_3^2 + a_4)\epsilon^2. \quad (2.12)$$

The suppressed constant is an increasing function of q , c_0^{-1} , and R_g .

See Section 2.3 for further insight into the leading order term $\int g Q d\hat{\pi}$.

Remark 2.1. Note that the bound (2.9) from Theorem 2.1 is an immediate corollary of Theorem 2.2, since $\int g Q d\hat{\pi} = 0$ if g is even about \hat{m} .

Next, we bound the TV error. Let $\mathcal{B}(\mathbb{R}^d)$ be the set of all Borel sets of \mathbb{R}^d , and $\mathcal{S}_{\hat{m}}(\mathbb{R}^d)$ be the set of Borel sets symmetric about \hat{m} .

Corollary 2.1 (TV error). *Suppose v satisfies Assumptions 1, 2, and 3. Then*

$$\begin{aligned} \sup_{A \in \mathcal{B}(\mathbb{R}^d)} |\pi(A) - \hat{\pi}(A)| &\lesssim a_3\epsilon + a_4\epsilon^2, \\ \sup_{A \in \mathcal{S}_{\hat{m}}(\mathbb{R}^d)} |\pi(A) - \hat{\pi}(A)| &\lesssim (a_3^2 + a_4)\epsilon^2. \end{aligned} \quad (2.13)$$

The suppressed constant is an increasing function of q and c_0^{-1} .

Corollary 2.2 (Mean and covariance error). *Let $\bar{m} = \int x d\pi(x)$ be the mean of π , and $\Sigma = \int (x - \bar{m})(x - \bar{m})^T d\pi(x)$ be the covariance of π . Suppose v satisfies Assumptions 1, 2, and 3. Then*

$$\begin{aligned} \|\hat{S}^{-1/2}(\bar{m} - \hat{m})\| &\lesssim (a_3^3 + a_3 a_4) \epsilon^3 + a_4^2 \epsilon^4, \\ \|\hat{S}^{-1/2}(\Sigma - \hat{S})\hat{S}^{-1/2}\| &\lesssim (a_3^2 + a_4)\epsilon^2. \end{aligned} \quad (2.14)$$

The suppressed constant is an increasing function of q and c_0^{-1} .

Remark 2.2. Recall from (2.6) that \hat{S} is equivalent to H^{-1} with respect to the positive definite matrix ordering. Therefore, given a lower bound $n\alpha \preceq H$, we can translate the bounds in Corollary 2.2 into analogous, unweighted bounds:

$$\begin{aligned}\|\bar{m} - \hat{m}\| &\lesssim \frac{1}{\sqrt{n\alpha}} \left[(a_3^3 + a_3 a_4) \frac{d^3}{n\sqrt{n}} + a_4^2 \frac{d^4}{n^2} \right], \\ \|\Sigma - \hat{S}\| &\lesssim \frac{1}{n\alpha} \left[(a_3^2 + a_4) \frac{d^2}{n} \right]\end{aligned}$$

2.3 Discussion

Assumptions on v and g Recall that in Bayesian inference, v is the normalized negative log posterior, taking the form (1.3), where the x_i are i.i.d. In typical situations, we expect that when n is large, the prior has negligible effect and $v(\theta) \approx -\mathbb{E} \log p_\theta(X)$, in which case v depends only on d . However, our results do not only apply to this typical regime. The function v is allowed to depend on both n and d , so that the constants from the assumptions — a_3, a_4, q, c_0 — may also depend on both n and d . Our results hold as long as these quantities and d and n are constrained by (2.4). However, for the logistic regression example, we will see that a_3, a_4, q, c_0 can all be chosen to be absolute constants.

When v is convex, Assumption 3 automatically follows from (2.4). Indeed, we have the following lemma.

Lemma 2.2. *Let v be convex and suppose $a_3(1 + (1/2)^q)\sqrt{d/n} \leq 3$. Then v satisfies Assumption 3 with $c_0 = 1/8$.*

See Appendix A for the proof. If v is not convex, note that the condition (2.5) is a nontrivial requirement only for larger x . To see this, fix the largest possible r such that $a_3 r(1+r^q)\sqrt{d/n} \leq 3/2$. Then a straightforward Taylor expansion shows that

$$v(x) - v(m^*) \geq \frac{1}{4} \|x - m^*\|_{H_v}^2 \geq \frac{1}{8} \sqrt{\frac{d}{n}} \|x - m^*\|_{H_v}$$

for all x such that $(1/2)\sqrt{d/n} \leq \|x - m^*\|_{H_v} \leq r\sqrt{d/n}$. Therefore, (2.5) need only be checked for $\|x - m^*\|_{H_v} \geq r\sqrt{d/n}$.

Next, let us discuss the condition (2.7) on the function g in Theorem 2.1. The condition is easiest to check when g is of the form $g(x) = f(\hat{S}^{-1/2}(x - \hat{m}))$. In this case, we have the following sufficient condition.

Lemma 2.3. *Let $g(x) = f(\hat{S}^{-1/2}(x - \hat{m}))$ and suppose f satisfies*

$$|f(y) - f(0)| \leq e^{C_f \sqrt{d} \|y\|^\alpha} \quad \forall y \in \mathbb{R}^d, \tag{2.15}$$

for some $0 \leq \alpha < 1$ and $C_f > 0$. Then for some function $R_g = R_g(C_f, \alpha, c_0^{-1})$ which is increasing in each of the variables, the condition (2.7) is satisfied.

See Appendix G for the proof.

Leading order term: VI vs Laplace We can gain insight into the difference between the Laplace and VI approximations by considering the respective leading order terms (LOTs) of the difference $\int g d\pi - \int g d\hat{\pi}$, where $\hat{\pi}$ is either the Laplace or VI Gaussian approximation. To make this comparison, we first make note of an alternative representation of the VI leading order term $\int g Q d\hat{\pi}$. Let

$$V_0(x) = V(\hat{m} + \hat{S}^{1/2}x), \quad f(x) = g(\hat{m} + \hat{S}^{1/2}x), \quad (2.16)$$

and let \mathbf{H}_3 be the tensor of third order Hermite polynomials. See Appendix E for a primer on Hermite polynomials in dimension $d \geq 1$. Let $\mathbf{A}_3(f) = \mathbb{E}[f(Z)\mathbf{H}_3(Z)]$ and $\mathbf{A}_3(V_0) = \mathbb{E}[V_0(Z)\mathbf{H}_3(Z)] = \mathbb{E}[\nabla^3 V_0(Z)]$, where the second representation is by Gaussian integration by parts. (We don't write $\mathbf{A}_3(f)$ this way because f need not be C^3 .) It can be shown (see Theorem 4.1 and Remark 4.1) that

$$\int g Q d\hat{\pi} = -\frac{1}{3!} \langle \mathbf{A}_3(V_0), \mathbf{A}_3(f) \rangle = -\frac{1}{6} \langle \mathbb{E}[\nabla^3 V_0(Z)], \mathbb{E}[f(Z)\mathbf{H}_3(Z)] \rangle.$$

In contrast, it is shown in [19] that the leading order term of the difference $\int g d\pi - \int g d\hat{\pi}$ for the Laplace approximation $\hat{\pi} = \mathcal{N}(m^*, H_V^{-1})$ takes the form $-\frac{1}{6} \langle \nabla^3 V_0(0), \mathbb{E}[f(Z)Z^{\otimes 3}] \rangle$ where now V_0, f are defined as

$$V_0(x) = V(m^* + H_V^{-1/2}x), \quad f(x) = g(m^* + H_V^{-1/2}x). \quad (2.17)$$

To summarize, we have the following errors, informally:

$$\int g d\pi - \int g d\hat{\pi} = -\frac{1}{6} \langle \mathbb{E}[\nabla^3 V_0(Z)], \mathbb{E}[f(Z)\mathbf{H}_3(Z)] \rangle + \mathcal{O}\left(\frac{d^2}{n}\right) \quad (\text{VI})$$

$$\int g d\pi - \int g d\hat{\pi} = -\frac{1}{6} \langle \nabla^3 V_0(0), \mathbb{E}[f(Z)Z^{\otimes 3}] \rangle + \mathcal{O}\left(\frac{d^2}{n}\right) \quad (\text{Laplace})$$

Here, V_0, f are given by (2.16) in the first line, and by (2.17) in the second line. The remainder term also depends on v and g , but we have written $\mathcal{O}(d^2/n)$ for brevity. We now compare the two approximation errors for various functions f . The most natural way to do so is to compare the two LOTs for the same function f , keeping in mind that this f corresponds to two different functions g for the two methods.

Both of the two LOTs can be shown to have order d/\sqrt{n} , generically. In particular, if f is an indicator, then neither LOT vanishes. This explains why $\text{TV}(\pi, \hat{\pi}) \lesssim d/\sqrt{n}$ for both VI and Laplace. Both the VI and Laplace LOT vanish when f is even. In particular, this explains why the Laplace and VI covariance errors are both bounded by d^2/n .

Finally, the key difference between Laplace and VI is that for functions f which are orthogonal to all third order Hermite polynomials but which are *not* even, the VI LOT vanishes while the Laplace LOT does not. Probably the most important example of such a function is $f(x) = x$. This explains VI's superior mean approximation accuracy. (The fact that the error is $\mathcal{O}((d/\sqrt{n})^3)$ rather than $\mathcal{O}((d/\sqrt{n})^2)$ stems from a cancellation in the next order term.) Another example of such an f is a fifth order Hermite polynomial.

We remark that having the LOT allows one to prove lower bounds in addition to upper bounds. In [19], we showed that for a posterior stemming from logistic regression, the Laplace mean approximation error is *lower*-bounded by d/\sqrt{n} , giving definitive proof that the VI mean approximation error is better in general. See Section 2.2 of [19] for the precise statements of the Laplace error bounds summarized above.

Solution to first-order optimality conditions (1.9) vs solution to (1.1) Recall that in our results, (\hat{m}, \hat{S}) is a solution to (1.9) for which \hat{m} is in a neighborhood of m^* , the mode of the posterior π . In the case that the posterior is multimodal, (\hat{m}, \hat{S}) may not be the global minimizer of the KL objective (1.1). This may occur if the posterior landscape has a single, “correct” dominant mode as well as a “false” second mode with a larger density value, but which is nevertheless a low probability region. Such a low probability mode could arise due to outliers in the data, for example. In this setting, the solution to (1.9) would lie near the false mode, while the solution to (1.1) would lie near the correct mode. It would therefore be useful to study the global minimizer of the KL objective in this regime, rather than the critical point near the highest mode, as we do in the present work. Moreover, this is potentially a case in which VI may be preferable over the Laplace approximation. Intuitively, we expect that VI is stable to low probability perturbations, while Laplace is not.

We leave the question of VI’s performance under this kind of multimodality to future work. Note that our assumptions do allow for multimodality, as long as the negative log posterior grows linearly away from the minimum. Our results show that in this case, the solution to (1.9) near the mode is a good approximation to the posterior, whether or not this solution coincides with the global minimizer of the KL objective.

Computation Although the computational side of Gaussian VI is not our focus, for the sake of completeness we summarize the known computational guarantees on Gaussian VI algorithms. The typical algorithm iteratively updates the mean and covariance matrix according to some optimization procedure. For example, [20] studies the discretization of an ODE evolving (m_t, Σ_t) , which corresponds to a Wasserstein gradient flow on $\text{KL}(\cdot \parallel \pi)$ constrained to the submanifold of normal distributions. The authors show that if V is strongly convex, then the gradient flow (as well as its discretization) converges exponentially fast in Wasserstein distance to $\hat{\pi}$ of (1.1), in this case the unique minimizer of $\text{KL}(\cdot \parallel \pi)$ over Gaussians. In the nonconvex case, [12] obtains a stationarity point guarantee for a “forward-backward” algorithm, involving a gradient step and a proximal step. The update step in [20] involves inverting a $d \times d$ matrix, while that of [12] involves computing the square root of a $d \times d$ matrix. The update step of both algorithms also involve computing a Gaussian expectation or else drawing a single Gaussian sample, in which case the algorithm is stochastic.

A cheaper algorithm whose complexity scales linearly with dimension is *mean-field* Gaussian VI, in which the objective (1.1) is minimized only over Gaussians with diagonal covariance matrices. See [5] for a discussion of the computational complexity of a whole range of Gaussian VI variants, from diagonal to full covariance matrix. However, the accuracy of Gaussian mean-field VI is not currently known.

3 Logistic regression with Gaussian design

In logistic regression, we generate n feature vectors $X_i \in \mathbb{R}^d$ and observe their corresponding labels $Y_i \in \{0, 1\}$. The distribution of the labels given the features is modeled as

$$p(Y_i \mid X_i, \theta) = \exp(Y_i X_i^T \theta - \psi(X_i^T \theta)), \quad (3.1)$$

where $\psi(t) = \log(1 + e^t)$ and θ is the unknown coefficient vector. Note that (3.1) is just the probability distribution for $\text{Bern}(\sigma(X_i^T \theta))$, where $\sigma(t) = \psi'(t) = (1 + e^{-t})^{-1}$ is the sigmoid. We assume the model is well-specified, so that there is a true $\theta_0 \in \mathbb{R}^d$ such that

$$Y_i \sim \text{Bern}(\sigma(X_i^T \theta_0)).$$

We assume the features are generated from a standard Gaussian distribution:

$$X_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d), \quad i = 1, \dots, n.$$

We consider a Gaussian prior $\theta \sim \mathcal{N}(0, \Sigma)$ on the coefficient vector θ . The posterior distribution of θ given the labels Y_i is then $\pi(\theta) \propto e^{-nv(\theta)}$, where

$$\begin{aligned} v(\theta) &= \ell(\theta) + \frac{1}{2n} \theta^T \Sigma^{-1} \theta, \\ \ell(\theta) &:= -\frac{1}{n} \sum_{i=1}^n \log p(Y_i | X_i, \theta) = -\theta^T \left[\frac{1}{n} \sum_{i=1}^n Y_i X_i \right] + \frac{1}{n} \sum_{i=1}^n \psi(X_i^T \theta) \end{aligned} \quad (3.2)$$

Here, ℓ is the negative normalized log likelihood, and $\theta^T \Sigma^{-1} \theta / 2n$ is the contribution from the log prior. Formally, $\Sigma^{-1} = 0$ corresponds to the case of a flat prior. Note that the distribution of the X_i does not enter into the posterior.

3.1 Checking the assumptions

To verify Assumption 1, we need to prove there is a unique global minimizer θ^* of v . To do so, we first show there is a unique global minimizer θ_ℓ^* of ℓ . It is straightforward to check that ℓ is strictly convex provided the features X_i span \mathbb{R}^d , and this occurs with probability 1. Therefore, if $\nabla \ell(\theta_\ell^*) = 0$ for some θ_ℓ^* , then this point must be the unique global minimizer of ℓ . To show ℓ has a critical point, we will use the following lemma with $f = \ell$ and $\theta = \theta_0$.

Lemma 3.1 (Corollary F.6 in [17] with $A = I_d$). *Let $f \in C^2(\mathbb{R}^d)$. Suppose there is a point θ such that for some $\lambda, s > 0$, the following holds:*

$$\begin{aligned} \nabla^2 f(\theta) &\succeq \lambda I_d, \\ \|\nabla f(\theta)\| &\leq 2\lambda s, \\ \|\nabla^2 f(\theta') - \nabla^2 f(\theta)\| &\leq \frac{\lambda}{4} \quad \forall \|\theta' - \theta\| \leq s. \end{aligned} \quad (3.3)$$

Then there exists θ_f^* such that $\nabla f(\theta_f^*) = 0$ and $\|\theta_f^* - \theta\| \leq s$.

Of course, ℓ is a random function, so (3.3) cannot hold deterministically for $f = \ell$. Instead, we bound from below the probability that (3.3) holds.

Lemma 3.2. *Suppose $d/n \leq 1/10$ and $n \geq 13$. Then the event*

$$E_1 = \{\|\nabla \ell(\theta_0)\| \leq 8(d/\sqrt{n})^{1/2}\}$$

has probability at least $1 - e^{-n/4} - n^{-d/4}$ under the ground truth distribution: $X_i \sim \mathcal{N}(0, I_d)$, $Y_i | X_i \sim \text{Bern}(\sigma(X_i^T \theta_0))$, and (X_i, Y_i) , $i = 1, \dots, n$ are i.i.d.

The proof of the lemma is similar to that of Lemma 4.2 in [17], but we include the proof in Appendix B for the sake of completeness.

Lemma 3.3 (Adapted from Lemma 7, Chapter 3, [26]). *Suppose $d < n$. Then the event*

$$E_2 = \left\{ \inf_{\|\theta\| \leq c} \lambda_{\min}(\nabla^2 \ell(\theta)) \geq \tau(c) \quad \forall c \geq 0 \right\} \quad (3.4)$$

has probability at least $1 - 4e^{-Cn}$ for an absolute constant C . Here, $\tau : [0, \infty) \rightarrow (0, \infty)$ is a universal non-increasing function.

We note that the function ℓ actually considered in Lemma 7 in [26] relates to our ℓ by a \sqrt{n} rescaling. We have translated the statement of the result in terms of our function ℓ .

Lemma 3.4. *Suppose $d \leq n$. Then there exist absolute constants C_3, C_4, C such that the event*

$$E_3 = \left\{ \sup_{\theta \in \mathbb{R}^d} \|\nabla^k \ell(\theta)\| \leq C_k, \quad k = 3, 4 \right\} \quad (3.5)$$

has probability at least $1 - \exp(-C(nd)^{1/9})$.

See Appendix B for the proof, which follows almost immediately by an application of Theorem 4.2 in [1]. We can now apply Lemma 3.1 to prove ℓ has a unique global minimizer θ_ℓ^* . This then also implies v has a unique global minimizer nearby, by another application of Lemma 3.1. Lemmas 3.2, 3.3, 3.4 also give us the necessary ingredients to check the other assumptions.

Corollary 3.1. *Suppose $\|\theta_0\| \leq C$, $\|\Sigma^{-1}\| \leq C$ and that d/\sqrt{n} and n^{-1} are smaller than a sufficiently small absolute constant. Then with probability at least $1 - \exp(-C(nd)^{1/9}) - n^{-d/4} - 5e^{-Cn}$, the following statements all hold:*

1. *The functions ℓ and v have unique global minimizers θ_ℓ^* and θ^* , respectively, with $\|\theta_\ell^* - \theta_0\| \lesssim \sqrt{d/n}$ and $\|\theta^* - \theta_\ell^*\| \lesssim n^{-1}$. In particular, Assumption 1 is satisfied.*
2. *We have $\|\theta^*\| \leq C$ and $\lambda_{\min}(H_v) = \lambda_{\min}(\nabla^2 v(\theta^*)) \geq \tau(C)$.*
3. *We have $\sup_{\theta \in \mathbb{R}^d} \|\nabla^k v(\theta)\|_{H_v} \leq C$ for $k = 3, 4$. In particular, (2.3) of Assumption 2 is satisfied with $q = 0$ and $a_3 = a_4 = C$. Furthermore, (2.4) of Assumption 2 is also satisfied provided d/\sqrt{n} is smaller than some fixed deterministic absolute constant.*
4. *Assumption 3 is satisfied with $c_0 = 1/8$.*

See Appendix B for the proof.

3.2 Application of results in Section 2 to logistic regression

We first compute the function Q in the leading order term. Now, one can check that

$$\nabla^3 v(\theta) = \frac{1}{n} \sum_{i=1}^n \psi'''(\theta^T X_i) X_i^{\otimes 3},$$

and recall that $V = nv$. Therefore,

$$\mathbb{E}_{\theta \sim \hat{\pi}} [\nabla^3 V(\theta)] = \sum_{i=1}^n b_i(\theta) X_i^{\otimes 3},$$

where

$$b_i(\theta) = \mathbb{E}_{\theta \sim \hat{\pi}} [\psi'''(\theta^T X_i)]. \quad (3.6)$$

Using the definition (2.11) of Q and the above expression for $\mathbb{E}_{\theta \sim \hat{\pi}} [\nabla^3 V(\theta)]$, we get

$$Q(\theta) = \sum_{i=1}^n b_i(\theta) \left[\frac{1}{2} X_i^T (\theta - \hat{m}) X_i^T \hat{S} X_i - \frac{1}{6} (X_i^T (\theta - \hat{m}))^3 \right] \quad (3.7)$$

We now summarize the results from Section 2 applied to logistic regression. We assume the set-up from the beginning of the section: well-specified model, i.i.d. Gaussian design, bounded ground truth $\|\theta_0\| \leq C$, and Gaussian or flat prior $\mathcal{N}(0, \Sigma)$, with $\|\Sigma^{-1}\| \leq C$. In this setting, Corollary 3.1 shows that Assumptions 1, 2, 3 are satisfied, and that $q = 0, a_3 = a_4 = C$.

Theorem 3.1. Suppose d/\sqrt{n} and n^{-1} are smaller than certain absolute constants. On an event of probability at least $1 - \exp(-C(nd)^{1/9}) - 5e^{-Cn} - n^{-d/4}$, the following results hold.

There exists a unique solution \hat{m}, \hat{S} to (1.9) in the region \mathcal{R}_V . Furthermore, if g satisfies (2.7) with $c_0 = 1/8$, then

$$\begin{aligned} \left| \int g d\pi - \int g d\hat{\pi} \right| &\lesssim_{R_g} (1 + \text{Var}_{\hat{\pi}}(g)^{\frac{1}{2}}) \frac{d}{\sqrt{n}}, \\ \left| \int g d\pi - \int g d\hat{\pi} - \int g Q d\hat{\pi} \right| &\lesssim_{R_g} (1 + \text{Var}_{\hat{\pi}}(g)^{\frac{1}{2}}) \left(\frac{d}{\sqrt{n}} \right)^2, \end{aligned} \quad (3.8)$$

where Q is defined in (3.7) and (3.6). If g is even about \hat{m} , then

$$\left| \int g d\pi - \int g d\hat{\pi} \right| \lesssim_{R_g} (1 + \text{Var}_{\hat{\pi}}(g)^{\frac{1}{2}}) \left(\frac{d}{\sqrt{n}} \right)^2. \quad (3.9)$$

If g is linear, then

$$\left| \int g d\pi - \int g d\hat{\pi} \right| \lesssim_{R_g} (1 + \text{Var}_{\hat{\pi}}(g)^{\frac{1}{2}}) \left(\frac{d}{\sqrt{n}} \right)^3. \quad (3.10)$$

Next, we have

$$\begin{aligned} \sup_{A \in \mathcal{B}(\mathbb{R}^d)} |\pi(A) - \hat{\pi}(A)| &\lesssim \frac{d}{\sqrt{n}}, \\ \sup_{A \in \mathcal{B}_{s, \hat{m}}(\mathbb{R}^d)} |\pi(A) - \hat{\pi}(A)| &\lesssim \left(\frac{d}{\sqrt{n}} \right)^2. \end{aligned} \quad (3.11)$$

Finally, let $\bar{\theta} = \int \theta d\pi(\theta)$ be the mean of π and $\Sigma = \int (\theta - \bar{\theta})(\theta - \bar{\theta})^T d\pi(\theta)$ be the covariance of π . Then

$$\begin{aligned} \sqrt{n} \|\bar{m} - \hat{m}\| &\lesssim \left(\frac{d}{\sqrt{n}} \right)^3, \\ n \|\Sigma - \hat{S}\| &\lesssim \left(\frac{d}{\sqrt{n}} \right)^2 \end{aligned} \quad (3.12)$$

In (3.11) and (3.12), all suppressed constants are absolute.

The proof follows immediately from Corollary 3.1 and the results in Section 2.

3.3 Numerical Simulation

We confirm some of the theoretical results —specifically, the bounds (3.12) — in a numerical simulation. The numerical results are displayed in Figure 1. We take $d = 2$ and $n = 100, 200, \dots, 1000$, and a flat prior on θ . For each n , we draw ten sets of covariates $x_i, i = 1, \dots, n$ from $\mathcal{N}(0, I_d)$, yielding ten posterior distributions $\pi_n(\cdot | x_{1:n})$. We then compute the Laplace and VI mean and covariance approximation errors for each n and each of the ten posteriors at a given n . The solid lines in Figure 1 depict the average approximation errors over the ten distributions at each n . The shaded regions depict the spread of the middle six out of ten approximation errors. See Appendix B for details about the simulation.

In the left panel of Figure 1 depicting the mean error, the slopes of the best fit lines are -1.04 and -2.02 for Laplace and Gaussian VI, respectively. For the covariance error in the righthand panel, the slopes of the best fit lines are -2.09 and -2.12 for Laplace and Gaussian VI. This confirms that our bounds on VI, as well as the Laplace bounds, are tight in their n dependence.

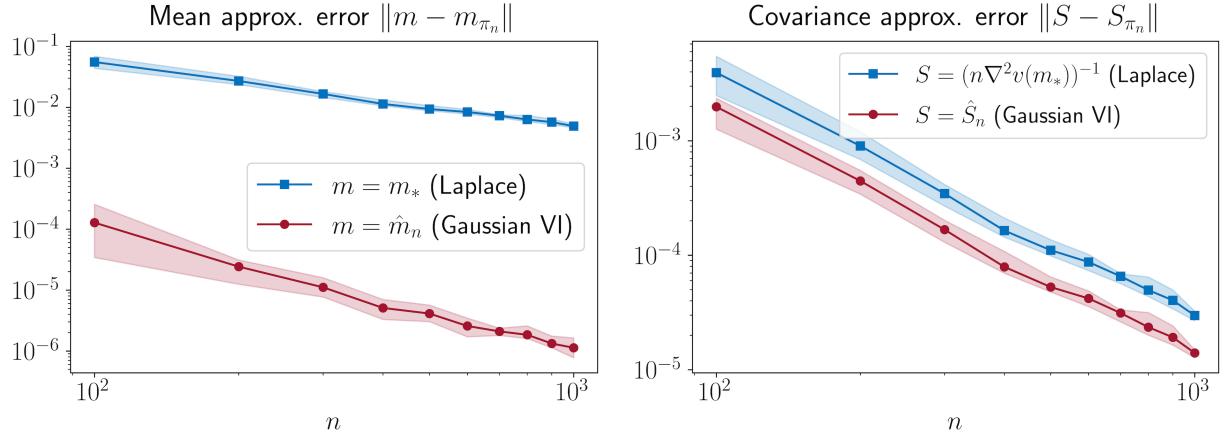


Figure 1: Gaussian VI yields a more accurate mean estimate than does Laplace, while the two covariance estimates are on the same order. Here, π_n is the likelihood of logistic regression given n observations in dimension $d = 2$. For the left-hand plot, the slopes of the best-fit lines are -1.04 for the Laplace approximation and -2.02 for Gaussian VI. For covariance: the slopes of the best-fit lines are -2.09 for Laplace, -2.12 for VI.

4 Proof of Theorems 2.1 and 2.2

In this section we explain the main steps in the proof of Theorems 2.1 and 2.2. We take Lemma 2.1 as given; its proof is outlined in Section 5.

4.1 Reduction to comparison with a standard Gaussian

Recall that $\hat{\pi} = \mathcal{N}(\hat{m}, \hat{S})$, where \hat{m} , \hat{S} are the solutions to (1.9) furnished by Lemma 2.1. We are interested in the difference $\int g d\pi - \int g d\hat{\pi}$, but we can always change variables via a bijection T , writing this difference in the form

$$\int g d\pi - \int g d\hat{\pi} = \int f dT_{\#}\pi - \int f dT_{\#}\hat{\pi}, \quad f = g \circ T^{-1}. \quad (4.1)$$

For convenience, we choose T so that $T_{\#}\hat{\pi} = \gamma$ is the standard Gaussian distribution. Specifically, we let

$$T(x) = \hat{S}^{-1/2}(x - \hat{m}), \quad (4.2)$$

and we define $\rho = T_{\#}\pi \propto e^{-V_0}$, where

$$V_0(x) := V(\hat{m} + \hat{S}^{1/2}x) = V(T^{-1}(x)). \quad (4.3)$$

Next, we reformulate Theorems 2.1 and 2.2. To do so, define $\mathbf{A}_3 = \mathbb{E}[\nabla^3 V_0(Z)]$ and

$$p_3(x) = \frac{1}{6} \langle \mathbf{A}_3, \mathbf{H}_3(x) \rangle = \left\langle \mathbf{A}_3, \frac{1}{6}x^{\otimes 3} - \frac{1}{2}x \otimes I_d \right\rangle.$$

Here, $\mathbf{H}_3(x)$ is the $d \times d \times d$ tensor of third order multivariate Hermite polynomials; see Section 4.2 and Appendix E for more details.

Theorem 4.1. Let $f = g \circ T^{-1}$, where T is as in (4.2), and let g satisfy (2.7). Also, let $\rho \propto e^{-V_0}$ as in (4.3), and $\Delta_f = \int f d\rho - \int f d\gamma$. Then under the conditions of Theorem 2.1, it holds

$$|\Delta_f| \lesssim (1 + \|f_0\|_2) (a_3\epsilon + a_4\epsilon^2), \quad (4.4)$$

$$\left| \Delta_f - \int f(-p_3) d\gamma \right| \lesssim (1 + \|f_0\|_2) (a_3^2 + a_4)\epsilon^2. \quad (4.5)$$

If f is orthogonal to all third order Hermite polynomials (in particular, if f is even), then

$$|\Delta_f| \lesssim (1 + \|f_0\|_2) (a_3^2 + a_4)\epsilon^2. \quad (4.6)$$

If f is linear, then

$$|\Delta_f| \lesssim (1 + \|f_0\|_2) ((a_3^3 + a_3 a_4) \epsilon^3 + a_4^2 \epsilon^4). \quad (4.7)$$

In all of the above bounds, the suppressed constant is an increasing function of q, c_0^{-1}, R_g .

The bounds (4.4), (4.6), and (4.7) correspond to (2.8), (2.9), and (2.10) of Theorem 2.1, respectively. This immediately follows from (4.1) and the fact that $\text{Var}_{\hat{\pi}}(g) = \text{Var}_{\gamma}(f) = \|f_0\|_2^2$. The bound (4.5) corresponds to (2.12) of Theorem 2.2. This is because $\int f(-p_3) d\gamma = \int g Q d\hat{\pi}$, as a short calculation in Appendix G shows.

4.2 Proof Outline

We now outline the proof of the equivalent Theorem 4.1. The proof is based on several key observations. First, note that the optimality equations (1.9) and the definition (4.3) of V_0 imply

$$\mathbb{E}[\nabla V_0(Z)] = 0, \quad \mathbb{E}[\nabla^2 V_0(Z)] = I_d. \quad (4.8)$$

The proof exploits (4.8) through six further observations.

1) The optimality conditions (4.8) imply that the Hermite series expansion of V_0 is of the form $V_0(x) = \text{const.} + \frac{1}{2}\|x\|^2 + r_3(x)$, where

$$r_3(x) = \sum_{k \geq 3} \frac{1}{k!} \langle \mathbf{A}_k, \mathbf{H}_k(x) \rangle. \quad (4.9)$$

In particular, $d\rho \propto e^{-V_0} \propto e^{-r_3} d\gamma$.

2) Since r_3 contains only third and higher order Hermite polynomials, r_3 is *orthogonal to all polynomials of degree 2 or lower*, with respect to the Gaussian measure. In other words, $\int f r_3 d\gamma$ for all quadratic f . As a special case (taking $f = 1$), we see that $\int r_3 d\gamma = 0$.

3) Let $p_3 = \frac{1}{3!} \langle \mathbf{A}_3, \mathbf{H}_3 \rangle$ be the third order Hermite polynomial contribution to V_0 , and write $r_3 = p_3 + r_4$, where r_4 is the Hermite series remainder of order 4. The change of variables via the transformation T ensures that $r_3 \sim p_3 \sim n^{-1/2}$ and $r_4 \sim n^{-1}$.

4) Let $f_0 = f - \gamma(f)$, so that $\Delta_f = \int f_0 d\rho$. Using that $d\rho \propto e^{-r_3} d\gamma$, that $r_3 \sim n^{-1/2}$, and that $\int f_0 d\gamma = \int r_3 d\gamma = 0$, we have

$$\begin{aligned} \Delta_f &= \int f_0 d\rho = \frac{\int f_0 e^{-r_3} d\gamma}{\int e^{-r_3} d\gamma} = \frac{\int f_0 (1 - r_3 + r_3^2/2 + \mathcal{O}(n^{-3/2})) d\gamma}{\int (1 - r_3 + \mathcal{O}(n^{-1})) d\gamma} \\ &= \frac{\int f_0 (-r_3 + r_3^2/2 + \mathcal{O}(n^{-3/2})) d\gamma}{1 + \mathcal{O}(n^{-1})} \end{aligned} \quad (4.10)$$

5) For a generic f , the order of Δ_f is that of the first remaining term in the numerator on the second line: $\Delta_f \sim \int f_0 r_3 d\gamma \sim n^{-1/2}$. But if f is linear or quadratic, then $\int f_0 r_3 d\gamma = 0$, since r_3 is orthogonal to quadratic polynomials. Therefore, Δ_f is at most of order $\int f_0 r_3^2 d\gamma \sim n^{-1}$.

6) By writing $r_3 = p_3 + r_4 = \mathcal{O}(n^{-1/2}) + \mathcal{O}(n^{-1})$ and therefore $r_3^2 = p_3^2 + \mathcal{O}(n^{-3/2})$, we can get even more refined estimates. Namely, we have

$$\Delta_f \sim \int f_0 \left(-r_3 + \frac{1}{2} r_3^2 \right) d\gamma = \underbrace{\int f_0 (-p_3) d\gamma}_{\mathcal{O}(n^{-1/2})} + \underbrace{\int f_0 \left(-r_4 + \frac{1}{2} p_3^2 \right) d\gamma}_{\mathcal{O}(n^{-1})} + \mathcal{O}(n^{-3/2}).$$

From this decomposition, we see that if f is orthogonal to third order Hermite polynomials then $\int f_0 p_3 d\gamma = 0$, and hence $\Delta_f \sim \mathcal{O}(n^{-1})$ in this case. If $f(x) = a^T x + b$ then $0 = \int f_0 p_3 d\gamma = \int f_0 r_4 d\gamma$ automatically, but we also have $\int f_0 p_3^2 d\gamma = 0$ since $f_0(x) = a^T x$ is odd and p_3^2 is even. Hence both the $\mathcal{O}(n^{-1/2})$ term and the $\mathcal{O}(n^{-1})$ term vanish, so that $\Delta_f \sim n^{-3/2}$.

This is the essence of the proof. Now that we have given this overview, let us go into a few more details about the above points. To explain 1), recall that the Hermite series expansion of V_0 is defined as

$$V_0(x) = \sum_{k=0}^{\infty} \frac{1}{k!} \langle \mathbf{A}_k(V_0), \mathbf{H}_k(x) \rangle, \quad \mathbf{A}_k(V_0) := \mathbb{E} [V_0(Z) \mathbf{H}_k(Z)] \quad (4.11)$$

Here, the $\mathbf{A}_k(V_0)$ and $\mathbf{H}_k(x)$ are tensors in $(\mathbb{R}^d)^{\otimes k}$. Specifically, $\mathbf{H}_k(x)$ is the tensor of all order k Hermite polynomials, enumerated as $H_k^{(\alpha)}$, $\alpha \in [d]^k$ with some entries repeating. For $k = 0, 1, 2$, the Hermite tensors are given by

$$\mathbf{H}_0(x) = 1, \quad \mathbf{H}_1(x) = x, \quad \mathbf{H}_2(x) = xx^T - I_d. \quad (4.12)$$

See Appendix E for further details on Hermite series and polynomials. Distinct Hermite polynomials are orthogonal to each other with respect to the Gaussian weight. Using the representation

$$\mathbf{H}_k(x) e^{-\|x\|^2/2} = (-1)^k \nabla^k e^{-\|x\|^2/2}, \quad (4.13)$$

we obtain a useful, “Gaussian integration by parts” identity for a k -times differentiable function f . Namely,

$$\mathbb{E} [f(Z) \mathbf{H}_k(Z)] = \mathbb{E} [\nabla^k f(Z)]. \quad (4.14)$$

This is a generalization of Stein’s identity, $\mathbb{E} [Z_i f(Z)] = \mathbb{E} [\partial_i f(Z)]$. Since V_0 is at four times differentiable, we can use Gaussian integration by parts to write $\mathbf{A}_1, \mathbf{A}_2$ as

$$\begin{aligned} \mathbf{A}_1(V_0) &:= \mathbb{E} [V_0(Z) \mathbf{H}_1(Z)] = \mathbb{E} [\nabla V_0(Z)] = 0, \\ \mathbf{A}_2(V_0) &:= \mathbb{E} [V_0(Z) \mathbf{H}_2(Z)] = \mathbb{E} [\nabla^2 V_0(Z)] = I_d, \end{aligned} \quad (4.15)$$

where the last equality in each line comes from the optimality conditions (4.8). Therefore the Hermite series expansion of V_0 takes the form

$$\begin{aligned} V_0(x) &= \langle A_0(V_0), H_0 \rangle + \langle 0, \mathbf{H}_1(x) \rangle + \frac{1}{2} \langle I_d, \mathbf{H}_2(x) \rangle + r_3(x) \\ &= \text{const.} + \frac{\|x\|^2}{2} + r_3(x), \end{aligned} \quad (4.16)$$

where r_3 is the third order remainder (4.9). Here, we recall from (4.12) that $\mathbf{H}_2(x) = xx^T - I_d$. This explains point 1), and point 2) is a direct application of the orthogonality of distinct Hermite polynomials.

For a heuristic proof of point 3), recall that $V_0(x) = V(\hat{m} + \hat{S}^{1/2}x)$, that $V = nv$, and that $\hat{S}^{1/2} \sim n^{-1/2}$. Therefore by the chain rule, $\nabla^3 V_0 \sim n^{-1/2}$, and $\nabla^4 V_0 \sim n^{-1}$. Hence

$$p_3 = \langle \mathbf{A}_3, \mathbf{H}_3 \rangle = \langle \mathbb{E}[\nabla^3 V_0(Z)], \mathbf{H}_3 \rangle \sim n^{-1/2}$$

and for $k \geq 4$, we can apply “partial” Gaussian integration by parts to express \mathbf{A}_k as

$$\mathbf{A}_k = \mathbb{E}[\mathbf{H}_k(Z)V_0(Z)] = \mathbb{E}[\mathbf{H}_{k-4}(Z) \otimes \nabla^4 V_0(Z)] \sim n^{-1}.$$

Therefore $r_4 \sim n^{-1}$.

Remark 4.1. Recall that in the statement of Theorem 4.1, we defined \mathbf{A}_3 as $\mathbf{A}_3 = \mathbb{E}[\nabla^3 V_0(Z)]$. We now see that $\mathbf{A}_3 = \mathbf{A}_3(V_0)$ is precisely the third order coefficient tensor in the Hermite series expansion of V_0 . Furthermore, we can now explain the formulation of the leading order term given in Section 2.3. Letting $\mathbf{A}_3(f)$ be the third order coefficient tensor in the Hermite series expansion of f , we can write

$$\int fp_3 d\gamma = \frac{1}{3!} \left\langle \mathbf{A}_3(V_0), \int f \mathbf{H}_3 d\gamma \right\rangle = \frac{1}{3!} \langle \mathbf{A}_3(V_0), \mathbf{A}_3(f) \rangle.$$

4.2.1 Dimension dependence

So far, we have only discussed the n dependence of our bounds. We now also briefly discuss the d dependence. As point 6) shows, in the typical case the leading order term is $\Delta_f \sim \int f_0 p_3 d\gamma$, which by Cauchy-Schwarz is bounded as $\|f_0\|_2 \|p_3\|_2$. In [18], we showed that for a symmetric $d \times d \times d$ tensor T , it holds

$$\mathbb{E}[\langle T, \mathbf{H}_3(Z) \rangle^p] \leq C_k \mathbb{E}[\langle T, \mathbf{H}_3(Z) \rangle^2]^{p/2} \leq C_p (d\|T\|)^p \quad (4.17)$$

for even p . Therefore, if $\|T\| \sim 1/\sqrt{n}$, then we get the desired dependence on $\epsilon = d/\sqrt{n}$.

Another difficulty that arises is in bounding $\|r_4\|_p$. Intuitively, $r_4(x)$ is dominated by the fourth order polynomial $\langle \mathbf{A}_4, \mathbf{H}_4(x) \rangle$, so we expect that

$$\mathbb{E}[|r_4(Z)|^p] \sim \mathbb{E}[|\langle \mathbf{A}_4, \mathbf{H}_4(Z) \rangle|^p] \sim (d^2/n)^p, \quad (4.18)$$

since $\mathbf{A}_4 \sim n^{-1}$ and $\mathbb{E}[\|Z\|^{4p}] \sim d^{2p}$. However, with only a series representation of $\|r_4\|_p$, it is difficult to translate this intuition into a formal bound. We therefore derive the following explicit formula for the Hermite series remainder:

Proposition 4.1. *Assume $V_0 \in C^K$. Let $V_0(x) = \sum_{j=0}^{\infty} \frac{1}{j!} \langle \mathbf{A}_j, \mathbf{H}_j(x) \rangle$ be the Hermite series expansion of V_0 , and define*

$$r_k(x) = V_0(x) - \sum_{j=0}^{k-1} \frac{1}{j!} \langle \mathbf{A}_j, \mathbf{H}_j(x) \rangle. \quad (4.19)$$

Then for all $k \leq K$, it holds

$$r_k(x) = \int_0^1 \frac{(1-t)^{k-1}}{(k-1)!} \mathbb{E} \left[\left\langle \nabla^k V_0((1-t)Z + tx), \mathbf{H}_k(x) - Z \otimes \mathbf{H}_{k-1}(x) \right\rangle \right] dt. \quad (4.20)$$

Note that (4.20) is analogous to the integral form of the remainder of a Taylor series. See Lemma E.3 in Appendix E.3 for the proof of this proposition. The formula (4.20) is known in one dimension; see Section 4.15 in [22]. However, we could not find the multidimensional version in the literature, so we have proved it here. Using this formula, we show that the heuristic argument (4.18) is essentially correct in that it gives the right scaling with d and n .

4.3 Key Lemmas

Now that we have outlined the steps of the proof informally, we turn to the rigorous proof. We formulate a sequence of lemmas from which Theorem 4.1 will follow. The omitted proofs can be found in Appendix D.

First, we define a few quantities.

Definition 4.1. Let r_3 be as defined in (4.19), and let $E := \int (e^{-r_3} - 1 + r_3) d\gamma$. Also, for a function $f \in L^2(\gamma)$, let

$$L(f) = \int f_0 \left(-r_3 + \frac{r_3^2}{2} \right) d\gamma, \quad E(f) = \int f_0 \left(e^{-r_3} - 1 + r_3 - \frac{r_3^2}{2} \right) d\gamma.$$

In the next lemma, we write $\int f d\rho - \int f d\gamma$ suggestively as $L(f)$ plus a remainder term.

Lemma 4.1. *It holds*

$$\left| \int f d\rho - \int f d\gamma - L(f) \right| \leq |E| |L(f)| + |E(f)|. \quad (4.21)$$

Proof. Recall from the proof outline that $\int f d\rho - \int f d\gamma = \int f_0 e^{-r_3} d\gamma / \int e^{-r_3} d\gamma$, and that $\int f_0 d\gamma = \int r_3 d\gamma = 0$. Therefore, we see that

$$\int f d\rho - \int f d\gamma = \frac{L(f) + E(f)}{1 + E}. \quad (4.22)$$

Furthermore, $1 + E = \int e^{-r_3} d\gamma \geq 1$ by Jensen's inequality, since $\int (-r_3) d\gamma = 0$. Subtracting $L(f)$ from both sides of (4.22) and bounding the resulting righthand side gives (4.21). \square

Consider the definition of $L(f)$, $E(f)$, and E . Since $r_3 \sim n^{-1/2}$, we expect that $L(f) \sim n^{-1/2}$, $E \sim n^{-1}$, and $E(f) \sim n^{-3/2}$. This is confirmed in the next lemma.

Lemma 4.2. *Let $f = g \circ T^{-1}$, as in Theorem 4.1. Then under the conditions of Theorem 2.1, it holds*

$$\begin{aligned} |L(f)| &\lesssim \|f_0\|_2 (a_3\epsilon + a_4\epsilon^2), \\ |E| &\lesssim (a_3\epsilon + a_4\epsilon^2)^2, \\ |E(f)| &\lesssim (1 + \|f_0\|_2) (a_3\epsilon + a_4\epsilon^2)^3, \end{aligned} \quad (4.23)$$

where the suppressed constant is an increasing function of q , c_0^{-1} , and R_g .

Lemmas 4.1 and 4.2 immediately give the following corollary, which shows that $\int f d\rho - \int f d\gamma = L(f) + \mathcal{O}(n^{-3/2})$.

Corollary 4.1. *In the same setting as Lemma 4.2, it holds*

$$\left| \int f d\rho - \int f d\gamma - L(f) \right| \lesssim (1 + \|f_0\|_2) (a_3\epsilon + a_4\epsilon^2)^3, \quad (4.24)$$

where the suppressed constant is an increasing function of q , c_0^{-1} , and R_g .

Thus it remains to study the term $L(f)$. By writing $r_3 = p_3 + r_4$, we break down $L(f)$ into a term of order $\mathcal{O}(n^{-1/2})$, a term of order $\mathcal{O}(n^{-1})$, and a remainder of order $\mathcal{O}(n^{-3/2})$.

Lemma 4.3. Let $p_3(x) = \frac{1}{3!} \langle \mathbf{A}_3, \mathbf{H}_3(x) \rangle$ and $r_4 = r_3 - p_3$. In the same setting as Lemma 4.2, it holds

$$L(f) = \int f_0(-p_3) d\gamma + \int f_0 \left(-r_4 + \frac{1}{2} p_3^2 \right) d\gamma + R,$$

where

$$\left| \int f_0(-p_3) d\gamma \right| \lesssim_q \|f_0\|_2 a_3 \epsilon, \quad (4.25)$$

$$\left| \int f_0 \left(-r_4 + \frac{1}{2} p_3^2 \right) d\gamma \right| \lesssim_q \|f_0\|_2 (a_3^2 + a_4) \epsilon^2, \quad (4.26)$$

$$|R| \lesssim_q \|f_0\|_2 (a_3 a_4 \epsilon^3 + a_4^2 \epsilon^4). \quad (4.27)$$

Combining Lemma 4.3 with Corollary 4.1, we can now prove Theorem 4.1.

Proof of Theorem 4.1. To get (4.4), we add the bound (4.24) to the bound on $|L(f)|$, which is given by the sum of (4.25), (4.26), and (4.27). To get (4.5), we add up (4.24), (4.27) and (4.26), omitting the bound (4.25) which gets incorporated into the lefthand side. Similarly, if f is orthogonal to all third order Hermite polynomials then $\int f_0 p_3 d\gamma = 0$, so the bound (4.6) stems from adding up (4.24), (4.26) and (4.27), again omitting (4.25). Finally, if $f(x) = a^T x + b$ then $\int f_0 p_3 d\gamma = \int f_0 r_4 d\gamma = 0$. Furthermore, $f_0(x) = a^T x$, which is odd, so $\int f_0 p_3^2 d\gamma = 0$. So we get (4.7) by adding (4.24) and (4.27), omitting both (4.25) and (4.26). \square

5 Existence of unique solution to stationarity conditions

As in the proof of Theorem 2.1, our first step in proving Lemma 2.1 will be to reformulate the lemma in a scale-free way. We will use the notation

$$\begin{aligned} B_r(0, 0) &= \{(m, \sigma) \in \mathbb{R}^d \times \mathbb{R}^{d \times d} : \|\sigma\|^2 + \|m\|^2 \leq r^2\}, \\ B_r &= \{\sigma \in \mathbb{R}^{d \times d} : \|\sigma\| \leq r\}, \\ S_{c_1, c_2} &= \{\sigma \in \mathbf{S}_+^d : c_1 I_d \preceq \sigma \preceq c_2 I_d\}. \end{aligned} \quad (5.1)$$

In particular, note that $S_{0, r} \subset B_r$.

5.1 Change of coordinates

Recall that m^* is the unique minimizer of $V = nv$, and that $H_V = \nabla^2 V(m^*)$. Define

$$W(x) = V(m^* + H_V^{-1/2} x). \quad (5.2)$$

Remark 5.1. Note that this change of variables is different from the one in Section 4. The change of variables (4.3) in that section already presumes existence and uniqueness of \hat{m} and \hat{S} , which is precisely what we will prove in the present section. The transformation $V \mapsto W$ in (5.2) simply recenters the minimum to be at zero, and rescales the Hessian at the minimum to be the identity matrix.

Lemma 5.1. Let W be as in (5.2) for $V = nv$, where v satisfies Assumptions 1 and 2. Then there exists a unique pair $(m, \sigma) \in B_{2\sqrt{2}}(0, 0) \cap \mathbb{R}^d \times S_{0, \sqrt{2}}$ satisfying

$$\mathbb{E}[\nabla W(m + \sigma Z)] = 0, \quad (5.3)$$

$$\mathbb{E}[\nabla^2 W(m + \sigma Z)] = (\sigma \sigma^T)^{-1}. \quad (5.4)$$

and this pair is such that $\sigma \in S_{\sqrt{2/3}, \sqrt{2}}$.

Using this lemma, a rescaling argument easily proves Lemma 2.1; see Appendix G.

Remark 5.2. The relationship between (m, σ) and (\hat{m}, \hat{S}) is as follows:

$$\hat{m} = m^* + H_V^{-1/2}m, \quad \hat{S} = H_V^{-1/2}\sigma\sigma^T H_V^{-1/2}. \quad (5.5)$$

5.2 Proof of Lemma 5.1

We first sketch the proof of Lemma 5.1, and then formally state the lemmas from which the result will follow. The proofs of the lemmas can be found in Appendix F.

Let $f : \mathbb{R}^d \times \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^d$ be given by $f(m, \sigma) = \mathbb{E}[\nabla W(m + \sigma Z)]$. Note that $f(0, 0) = 0$, so by the Implicit Function Theorem, there exists a map $m(\sigma)$ defined in a neighborhood of $\sigma = 0$ such that $f(m(\sigma), \sigma) = 0$. In Lemma 5.3, we make this statement quantitative, showing that for $r = 2\sqrt{2}$ we have the following result: for every $\sigma \in B_{r/2}$ there is a unique $m = m(\sigma)$ such that $(m, \sigma) \in B_r(0, 0)$ and $f(m, \sigma) = 0$. Since $S_{0,r/2} \subset B_{r/2}$, we have in particular that any solution (m, σ) to (5.3) in the region $B_r(0, 0) \cap \mathbb{R}^d \times S_{0,r/2}$ is of the form $(m(\sigma), \sigma)$. Thus it remains to prove there exists a unique solution $\sigma \in S_{0,r/2}$ to the equation $\mathbb{E}[\nabla^2 W(m(\sigma) + \sigma Z)] = (\sigma\sigma^T)^{-1}$.

To do so, we rewrite this equation as $F(\sigma) = \sigma$, where

$$F(\sigma) = \mathbb{E}[\nabla^2 W(m(\sigma) + \sigma Z)]^{-1/2}.$$

We show in Lemma 5.4 that F is well-defined on $S_{0,r/2}$, a contraction, and satisfies $F(S_{0,r/2}) \subset S_{c_1, c_2} \subset S_{0,r/2}$. Thus by the Contraction Mapping Theorem, there is a unique $\sigma \in S_{0,r/2}$ satisfying $F(\sigma) = \sigma$. But since F maps $S_{0,r/2}$ to S_{c_1, c_2} , the fixed point σ necessarily lies in S_{c_1, c_2} . This finishes the proof.

Using a quantitative statement of the Inverse Function Theorem given in [21], the following lemma determines the size of the neighborhood in which the map $m(\sigma)$ is defined.

Lemma 5.2. *Let $f = (f_1, \dots, f_d) : \mathbb{R}^d \times \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^d$ be C^3 , where $\mathbb{R}^{d \times d}$ is the set of $d \times d$ matrices, endowed with the standard matrix operator norm. Suppose $f(0, 0) = 0$, $\nabla_\sigma f(0, 0) = 0$, $\nabla_m f(m, \sigma)$ is symmetric for all m, σ , and $\nabla_m f(0, 0) = I_d$. Let $r > 0$ be such that*

$$\sup_{(m, \sigma) \in B_r(0, 0)} \|\nabla f(m, \sigma) - \nabla f(0, m^*)\|_{op} \leq \frac{1}{4}. \quad (5.6)$$

Then for each $\sigma \in \mathbb{R}^{d \times d}$ such that $\|\sigma\| \leq r/2$ there exists a unique $m = m(\sigma) \in \mathbb{R}^d$ such that $f(m(\sigma), \sigma) = 0$ and $(m(\sigma), \sigma) \in B_r(0, 0)$. Furthermore, the map $\sigma \mapsto m(\sigma)$ is C^2 , with

$$\begin{aligned} \frac{1}{2}I_d &\preceq \nabla_m f(m, \sigma) \Big|_{m=m(\sigma)} \preceq \frac{3}{2}I_d, \\ \|\nabla_\sigma m(\sigma)\|_{op} &\leq 1. \end{aligned} \quad (5.7)$$

See Appendix F for careful definitions of the norms appearing above, as well as the proof of the lemma.

Lemma 5.3. *Let $f : \mathbb{R}^d \times \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^d$ be given by $f(m, \sigma) = \mathbb{E}[\nabla W(\sigma Z + m)]$. Then all the conditions of Lemma 5.2 are satisfied; in particular, (5.6) is satisfied with $r = 2\sqrt{2}$. Thus the conclusions of Lemma 5.2 hold with this choice of r .*

Lemma 5.4. *Let $r = 2\sqrt{2}$ and $\sigma \in S_{0,r/2} \mapsto m(\sigma) \in \mathbb{R}^d$ be the restriction of the map furnished by Lemmas 5.2 and 5.3 to symmetric nonnegative matrices. Then the function F given by*

$$F(\sigma) = \mathbb{E} [\nabla^2 W(m(\sigma) + \sigma Z)]^{-1/2}$$

is well-defined and a strict contraction on $S_{0,r/2}$. Moreover,

$$F(S_{0,r/2}) \subseteq S_{c_1, c_2} \subseteq S_{0,r/2}, \quad \text{where } c_1 = \sqrt{2/3}, \quad c_2 = \sqrt{2} = r/2.$$

This lemma concludes the proof of Lemma 5.1 since by the Contraction Mapping Theorem there is a unique fixed point $\sigma \in S_{0,r/2}$ of F , and $F(\sigma) = \sigma$ is simply a reformulation of the second optimality equation (5.4). We know σ must lie in S_{c_1, c_2} since F maps $S_{0,r/2}$ to this set.

Acknowledgments

A. Katsevich is supported by NSF grant DMS-2202963. P. Rigollet is supported by NSF grants IIS-1838071, DMS-2022448, and CCF-2106377.

We review some notation used throughout the appendix. For a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ such that $\int |f|^p d\gamma < \infty$, we define

$$\|f\|_p = \left(\int |f|^p d\gamma \right)^{\frac{1}{p}}.$$

For any other measure λ , we indicate the dependence of the p -norm on λ explicitly. In particular, we let

$$\|f\|_{L^p(\lambda)} = \left(\int |f|^p d\lambda \right)^{\frac{1}{p}}, \quad \text{Var}_\lambda(f) = \int (f - \lambda(f))^2 d\lambda.$$

A Proofs from Section 2

Proof of Corollary 2.1. Let $g = \mathbb{1}_A$. Note that $|g(x) - \hat{\pi}(g)| \leq 1$ for all x , so g satisfies (2.7) with $R_g = 0$. Also, $\text{Var}_{\hat{\pi}}(g) \leq 1$. We now apply (2.8) and (2.9) of Theorem 2.1 to conclude. \square

Proof of Corollary 2.2. Fix $\|u\| = 1$, and let $g_u(x) = u^T \hat{S}^{-1/2}(x - \hat{m})$. We have

$$\|\hat{S}^{-1/2}(\bar{m} - \hat{m})\| = \sup_{\|u\|=1} \int g_u d\pi - \int g_u d\hat{\pi}.$$

Now, $g_u(x) = f_u(\hat{S}^{-1/2}(x - \hat{m}))$, where $f_u(y) = u^T y$. We have

$$|f_u(y) - f_u(0)| \leq \|y\| \leq e^{2\|y\|^{1/2}} \quad \forall y \in \mathbb{R}^d,$$

so f satisfies (2.15) with $\alpha = 1/2$ and $C_{f_u} = 2$. Therefore by Lemma 2.3, we have that g_u satisfies (2.7) with $R_g = C(c_0^{-1})$ (note C_{f_u} and α are absolute constants, so we don't include them). Finally, note that $\text{Var}_{\hat{\pi}}(g_u) = \text{Var}(u^T Z) = 1$ for all $\|u\| = 1$. An application of (2.10) from Theorem 2.1 concludes the proof of the bound on the mean error. Since R_g depends only on c_0^{-1} , the suppressed constant depends only on c_0^{-1} and q .

To bound the covariance error, we consider the function g_u^2 , which is even about \hat{m} . Now,

$$\begin{aligned}
\|\hat{S}^{-1/2}(\Sigma - \hat{S})\hat{S}^{-1/2}\| &= \sup_{\|u\|=1} |\text{Var}_{X \sim \pi}(g_u(X)) - 1| \\
&\leq \sup_{\|u\|=1} |\mathbb{E}_{X \sim \pi}[g_u(X)^2] - 1| + \sup_{\|u\|=1} \mathbb{E}[g_u(X)]^2 \\
&= \sup_{\|u\|=1} \left| \int g_u^2 d\pi - \int g_u^2 d\hat{\pi} \right| + \|\hat{S}^{-1/2}(\bar{m} - \hat{m})\|^2
\end{aligned} \tag{A.1}$$

For the bound on $\|\hat{S}^{-1/2}(\bar{m} - \hat{m})\|^2$, we use the mean error we have just proved. For the bound on $\int g_u^2 d\pi - \int g_u^2 d\hat{\pi}$, we use that $g_u^2(x) = f_u(\hat{S}^{-1/2}(x - \hat{m}))$, where $f_u(y) = (u^T y)^2$. We have

$$|f_u(y) - f(0)| \leq \|y\|^2 \leq e^{4\|y\|^{1/4}} \quad \forall y \in \mathbb{R}^d,$$

so f satisfies (2.15) with $\alpha = 1/4$ and $C_{f_u} = 4$. Therefore by Lemma 2.3, we have that g_u^2 satisfies (2.7) with $C_{g_u^2} = C(c_0^{-1})$ (note C_{f_u} and α are absolute constants, so we don't include them). Finally, note that $\text{Var}_{\hat{\pi}}(g_u^2) = \text{Var}((u^T Z)^2) \leq 3$ for all $\|u\| = 1$. The bound on $\int g_u^2 d\pi - \int g_u^2 d\hat{\pi}$ now follows from (2.9) of Theorem 2.1. Since $C_{g_u^2}$ depends only on c_0^{-1} , the suppressed constant depends only on c_0^{-1} and q . \square

Proof of Lemma 2.2. Let x be such that $\|x - m^*\|_{H_v} = (1/2)\sqrt{d/n}$ for some $r > 0$. A Taylor expansion of v around m^* gives

$$\begin{aligned}
v(x) - v(m^*) &= \frac{1}{2}\|x - m^*\|_{H_v}^2 + \frac{1}{3!}\langle \nabla^3 v(\xi), (x - m^*)^{\otimes 3} \rangle \\
&\geq \frac{1}{2}\|x - m^*\|_{H_v}^2 \left(1 - \frac{1}{3}\|\nabla^3 v(\xi)\|_{H_v} \|x - m^*\|_{H_v} \right) \\
&\geq \frac{1}{2}\|x - m^*\|_{H_v}^2 \left(1 - \frac{a_3}{6}\sqrt{d/n}(1 + (1/2)^q) \right) \geq \frac{1}{4}\|x - m^*\|_{H_v}^2
\end{aligned} \tag{A.2}$$

using the assumption $a_3(1 + (1/2)^q)\sqrt{d/n} \leq 3$ to get the last inequality. Therefore,

$$\inf_{\|x - m^*\|_{H_v} = (1/2)\sqrt{d/n}} \frac{v(x) - v(m^*)}{\|x - m^*\|_{H_v}} \geq \frac{1}{8}\sqrt{d/n} \tag{A.3}$$

Now, since v is convex, we have

$$\frac{v(y) - v(m^*)}{\|y - m^*\|_{H_v}} \geq \inf_{\|x - m^*\|_{H_v} = (1/2)\sqrt{d/n}} \frac{v(y - x) - v(m^*)}{\|x - m^*\|_{H_v}} \geq \frac{1}{8}\sqrt{d/n}$$

for all $\|y - m^*\|_{H_v} \geq (1/2)\sqrt{d/n}$. \square

For the proof of Lemma 2.3, see Section G.

B Logistic Regression Example

Details of Numerical Simulation

For the numerical simulation displayed in Figure 1, we take $d = 2$ and $n = 100, 200, \dots, 1000$. For each n , we draw ten sets of covariates $x_i, i = 1, \dots, n$ from $\mathcal{N}(0, \lambda^2 I_d)$ with $\lambda = \sqrt{5}$, yielding ten

posterior distributions $\pi_n(\cdot | x_{1:n})$. For each π_n we compute the ground truth mean and covariance by directly evaluating the integrals, using a regularly spaced grid (this is feasible in two dimensions). The mode m_* of π_n is found by a standard optimization procedure, and the Gaussian VI estimates \hat{m}, \hat{S} are computed using the procedure described in [20]. We used the authors' implementation of this algorithm, found at <https://github.com/marc-h-lambert/W-VI>. We then compute the Laplace and VI mean and covariance approximation errors for each n and each of the ten posteriors at a given n . The solid lines in Figure 1 depict the average approximation errors over the ten distributions at each n . The shaded regions depict the spread of the middle eight out of ten approximation errors.

Verifying the Assumptions

Proof of Lemma 3.2. Let $S = \frac{1}{n} \sum_{i=1}^n (Y_i - \mathbb{E}[Y_i | X_i])$ and \mathcal{N} be a $1/2$ -net of the sphere in \mathbb{R}^d , so that $\inf_{u \in \mathcal{N}} \|u - w\| \leq 1/2$ for all $\|w\| = 1$. Standard arguments show we can take \mathcal{N} to have at most 5^d elements. For some $s \geq 0$, we have

$$\mathbb{P}(\|S\| \geq 2s\sqrt{d/n}) \leq \mathbb{P}(\sup_{u \in \mathcal{N}} u^T S \geq s\sqrt{d/n}) \leq 5^d \sup_{\|u\|=1} \mathbb{P}(u^T S \geq s\sqrt{d/n}) \quad (\text{B.1})$$

using a union bound and the fact that $|\mathcal{N}| \leq 5^d$. For a fixed unit vector u , we have

$$\begin{aligned} \mathbb{P}(u^T S \geq s\sqrt{d/n}) &= \mathbb{E} \left[\mathbb{P} \left(u^T S \geq s\sqrt{d/n} \mid \{X_i\}_{i=1}^n \right) \right] \\ &= \mathbb{E} \left[\mathbb{P} \left(\sum_{i=1}^n (Y_i - \mathbb{E}[Y_i | X_i]) u^T X_i \geq s\sqrt{nd} \mid \{X_i\}_{i=1}^n \right) \right] \\ &\leq \mathbb{E} \left[\exp \left(-\frac{2s^2nd}{\sum_{i=1}^n (u^T X_i)^2} \right) \right] = \mathbb{E} \left[\exp \left(-\frac{2s^2nd}{\|Z\|^2} \right) \right]. \end{aligned} \quad (\text{B.2})$$

In the last line, $Z \sim \mathcal{N}(0, I_n)$. To get the third line, we used Hoeffding's inequality. Now, we have

$$\mathbb{E} \left[\exp \left(-\frac{2s^2nd}{\|Z\|^2} \right) \right] \leq \mathbb{P}(\|Z\| \geq 2\sqrt{n}) + \exp(-s^2d/2) \leq e^{-n/2} + e^{-s^2d/2}. \quad (\text{B.3})$$

Combining (B.1), (B.2), (B.3) gives

$$\mathbb{P}(\|S\| \geq 2s\sqrt{d/n}) \leq e^{d\log 5-n/2} + e^{d(\log 5-s^2/2)} \quad (\text{B.4})$$

Taking $s = \log n$ gives

$$\begin{aligned} \mathbb{P}(\|S\| \geq 2\log n\sqrt{d/n}) &\leq e^{d\log 5-n/2} + e^{d(\log 5-\frac{1}{2}\log^2 n)} \\ &\leq e^{-n/4} + e^{-d(\log n)^2/4} \leq e^{-n/4} + n^{-d/4} \end{aligned}$$

using that $\log 5 - \frac{1}{2}\log^2 n \leq -\frac{1}{4}\log^2 n$ when $n \geq 13$, and that $d\log 5 - n/2 < -n/4$ when $d/n \leq 0.1$. Finally, note that

$$2\log n\sqrt{d/n} = \frac{2\log n}{n^{1/4}} \left(\frac{d}{\sqrt{n}} \right)^{\frac{1}{2}} = \frac{8\log(n^{1/4})}{n^{1/4}} \left(\frac{d}{\sqrt{n}} \right)^{\frac{1}{2}} \leq 8 \left(\frac{d}{\sqrt{n}} \right)^{\frac{1}{2}}.$$

□

Proof of Lemma 3.4. First, note that $\nabla^k v(\theta) = \frac{1}{n} \sum_{i=1}^n \psi^{(k)}(X_i^T \theta) X_i^{\otimes k}$ for $k = 3, 4$. Therefore,

$$\|\nabla^k v(\theta)\| \lesssim 1 + \sup_{\|u\|=1} \frac{1}{n} \sum_{i=1}^n |u^T X_i|^4, \quad k = 3, 4. \quad (\text{B.5})$$

using also that $\|\psi^{(k)}\|_\infty \lesssim 1$. Now, we apply the following result.

Theorem B.1 (Adapted from Theorem 4.2 in [1]). *Let X_1, \dots, X_n be i.i.d. standard Gaussian. There exist absolute constants C, C' such that if $Ct^8[\log(2t^2)]^6 \leq n/d^2$ then*

$$\sup_{\|u\|=1} \frac{1}{n} \sum_{i=1}^n |u^T X_i|^4 \leq 4 \quad (\text{B.6})$$

with probability at least $1 - e^{-C't\sqrt{d}}$.

We apply the theorem with $t = c(n/d^2)^{1/9}$. Let $M = n/d^2$ and suppose $M \geq 1$. The condition $Ct^8[\log(2t^2)]^6 \leq n/d^2$ can then be written as $Cc^8 \log(2c^2 M^{2/9}) \leq M^{1/9}$ and this can be satisfied by taking c a small enough absolute constant. Thus the probability that (B.6) holds is $1 - \exp(-C''(nd)^{1/9})$ for some other absolute constant $C'' > 0$. Combining (B.6) with (B.5) concludes the proof. \square

Proof of Corollary 3.1. We claim the statements hold on the event $E = E_1 \cap E_2 \cap E_3$, where E_1, E_2, E_3 are defined in Lemmas 3.2, 3.3, 3.4, respectively. Note that $\mathbb{P}(E)$ satisfies the desired lower bound. Now, on E , we have $\|\nabla \ell(\theta_0)\| \leq 8(d/\sqrt{n})^{1/2}$ and $\nabla^2 \ell(\theta_0) \succeq \tau(\|\theta_0\|) I_d$. Furthermore, if $\|\theta' - \theta_0\| \leq s$, then $\|\nabla^2 \ell(\theta') - \nabla^2 \ell(\theta_0)\| \leq C_3 s$. Let $\lambda = \tau(\|\theta_0\|)$ and $s = 4(d/\sqrt{n})^{1/2}/\lambda$. If $(d/\sqrt{n})^{1/2} \leq \lambda^2/(16C_3) = \tau(\|\theta_0\|)^2/(16C_3)$, then the inequalities (3.3) are all satisfied for $f = \ell$ and $\theta = \theta_0$. We therefore conclude that there exists a unique global minimizer θ_ℓ^* of ℓ , which satisfies

$$\|\theta_\ell^* - \theta_0\| \leq (4/\tau(\|\theta_0\|))(d/\sqrt{n})^{1/2} \leq C(d/\sqrt{n})^{1/2}$$

Next, note that on E , we have

$$\|\nabla v(\theta_\ell^*)\| = \|\nabla \ell(\theta_\ell^*) + n^{-1} \Sigma^{-1} \theta_\ell^*\| = n^{-1} \|\Sigma^{-1} \theta_\ell^*\| \lesssim n^{-1} \|\theta_\ell^*\|$$

and

$$\nabla^2 v(\theta_\ell^*) \succeq \nabla^2 \ell(\theta_\ell^*) \succeq \tau(C) I_d,$$

since $\|\theta_\ell^*\| \leq C(d/\sqrt{n})^{1/2} + \|\theta_0\| \leq C$. Furthermore, $\sup_\theta \|\nabla^3 v(\theta)\| \leq C_3$, since $\nabla^3 v = \nabla^3 \ell$. Therefore, since n^{-1} is assumed to be sufficiently small, we can use Lemma 3.1 to conclude v has a critical point (and therefore a unique global minimizer) θ^* , and $\|\theta^* - \theta_\ell^*\| \lesssim n^{-1}$. That $\|\theta^*\| \leq C$ is bounded follows from the bounds on $\|\theta_0\|, \|\theta_\ell^* - \theta_0\|, \|\theta^* - \theta_\ell^*\|$. Furthermore, we have $\lambda_{\min}(H_v) \geq \lambda_{\min}(\nabla^2 \ell(\theta^*)) \geq \tau(C)$ using the definition of event E_2 . To prove point 3, we use the lower bound on $\lambda_{\min}(H_v)$ from point 2, combined with Lemma 3.4 and the fact that $\nabla^k v = \nabla^k \ell$, $k \geq 3$, to get

$$\|\nabla^k v(\theta)\|_{H_v} \leq \lambda_{\min}(H_v)^{-k/2} \|\nabla^k v(\theta)\| \leq C \quad (\text{B.7})$$

on event E . Finally, point 4 is satisfied by convexity of v , using Lemma 2.2 and the boundedness of a_3 (needed to verify the assumption of the lemma). \square

C Properties of W and V_0

In this section, we prove several key properties of the rescaled functions $W(x) = V(m^* + H_V^{-1/2}x)$ and $V_0(x) = V(\hat{m} + \hat{S}^{1/2}x) = V(T^{-1}(x))$. These properties are used in the proofs in Sections 4 and 5. First, note that for a measure λ on \mathbb{R}^d , we have

$$\left\| \|\nabla^k W\| \right\|_{L^p(\lambda)} = \left(\int \|\nabla^k W(x)\|^p d\lambda(x) \right)^{1/p},$$

where the norm inside the integral is the standard tensor operator norm

Lemma C.1. *If v satisfies Assumptions 1, 2, 3 then 0 is the unique global minimizer of W , with $\nabla^2 W(0) = I_d$, and*

$$\begin{aligned} \|\nabla^3 W(x)\| &\leq \frac{a_3}{\sqrt{n}} (1 + \|x/\sqrt{d}\|^q), \quad \forall x \in \mathbb{R}^d \\ \|\nabla^4 W(x)\| &\leq \frac{a_4}{n} (1 + \|x/\sqrt{d}\|^q), \quad \forall x \in \mathbb{R}^d. \end{aligned} \quad (\text{C.1})$$

Furthermore,

$$W(x) - W(0) \lesssim 1 + \|x\|^{3+q} \quad \forall x \in \mathbb{R}^d, \quad (\text{C.2})$$

$$W(x) - W(0) \geq c_0 \sqrt{d} \|x\| \quad \forall \|x\| \geq \sqrt{\frac{2}{3}} \sqrt{d}, \quad (\text{C.3})$$

Finally, for $m \in \mathbb{R}^d$, $\sigma \in \mathbb{R}^{d \times d}$, we have

$$\begin{aligned} \left\| \|\nabla^3 W\| \right\|_{L^p(\mathcal{N}(m, \sigma\sigma^T))} &\lesssim_{p,q} \frac{a_3}{\sqrt{n}} \left(1 + \|m/\sqrt{d}\|^q + \|\sigma\|^q \right), \\ \left\| \|\nabla^4 W\| \right\|_{L^p(\mathcal{N}(m, \sigma\sigma^T))} &\lesssim_{p,q} \frac{a_4}{n} \left(1 + \|m/\sqrt{d}\|^q + \|\sigma\|^q \right). \end{aligned} \quad (\text{C.4})$$

See the end of the section for the proof. Next, we collect properties of the function $V_0(x) = V(\hat{m} + \hat{S}^{1/2}x) = V(T^{-1}(x))$ used in the proof of Theorem 4.1. Here, \hat{m}, \hat{S} are the solutions to the first-order optimality conditions from Lemma 2.1. We first rewrite V_0 in terms of W . To do so, let $\hat{m}_0 = m, \hat{\sigma}_0 = \sigma$, where (m, σ) is the unique solution to (5.3), (5.4) furnished by Lemma 5.1 (these are just the rescaled optimality conditions). Using the relationship (5.5) between $(m, \sigma) = (\hat{m}_0, \hat{\sigma}_0)$ and $(\hat{m}, \hat{S}^{1/2})$ given in Remark 5.2, it is straightforward to show that

$$V_0(x) = V(\hat{m} + \hat{S}^{1/2}x) = W(\hat{m}_0 + \hat{\sigma}_0 x). \quad (\text{C.5})$$

Also, recall from Lemma 5.1 that

$$\|\hat{m}_0\| \leq 2\sqrt{2}, \quad \sqrt{\frac{2}{3}} I_d \preceq \hat{\sigma}_0 \preceq \sqrt{2} I_d. \quad (\text{C.6})$$

We now have the following properties of V_0 .

Lemma C.2. *It holds*

$$|V_0(x) - V_0(0)| \lesssim_q 1 + \|x\|^{3+q} \quad \forall x \in \mathbb{R}^d, \quad (\text{C.7})$$

$$V_0(x) - V_0(0) \geq \frac{c_0 \sqrt{d}}{2} \|x\| - C(q) \quad \forall \|x\| \geq \sqrt{d} + 8\sqrt{2}. \quad (\text{C.8})$$

Furthermore, if $0 < s \leq 1$, then

$$\|\|\nabla^3 V_0\|\|_{L^p(\mathcal{N}(\mu, s^2 I_d))} \lesssim_q \frac{a_3}{\sqrt{n}} \left(1 + \|\mu/\sqrt{d}\|^q\right), \quad (\text{C.9})$$

$$\|\|\nabla^4 V_0\|\|_{L^p(\mathcal{N}(\mu, s^2 I_d))} \lesssim_q \frac{a_4}{n} \left(1 + \|\mu/\sqrt{d}\|^q\right). \quad (\text{C.10})$$

In particular,

$$d \|\|\nabla^3 V_0\|\|_{L^p(\mathcal{N}(0, s^2 I_d))} \lesssim_q a_3 \epsilon, \quad (\text{C.11})$$

$$d^2 \|\|\nabla^4 V_0\|\|_{L^p(\mathcal{N}(0, s^2 I_d))} \lesssim_q a_4 \epsilon^2, \quad (\text{C.12})$$

where $\epsilon = d/\sqrt{n}$.

We now prove Lemmas C.1 and C.2.

Proof of Lemma C.1. We first note that

$$\|\nabla^k W(x)\| = n^{1-k/2} \|\nabla^k v(m^* + H_v^{-1/2} x/\sqrt{n})\|_{H_v}. \quad (\text{C.13})$$

Now, using Assumption 2, we have

$$\begin{aligned} \|\nabla^k W(x)\| &= n^{1-k/2} \|\nabla^k v(m^* + H_v^{-1/2} x)\|_{H_v} \\ &\leq \frac{a_k}{n^{\frac{k}{2}-1}} \left[1 + \left(\sqrt{n/d} \|m^* + H_v^{-1/2} x/\sqrt{n} - m^*\|_{H_v} \right)^q \right] \\ &= \frac{a_k}{n^{\frac{k}{2}-1}} \left[1 + \|x/\sqrt{d}\|^q \right], \end{aligned} \quad (\text{C.14})$$

as desired. To prove (C.2) we Taylor expand W about 0 and apply (C.1). We get

$$\begin{aligned} W(x) - W(0) &= \frac{\|x\|^2}{2} + \frac{1}{3!} \langle \nabla^3 W(tx), x^{\otimes 3} \rangle \\ &\leq \frac{\|x\|^2}{2} + \frac{\|x\|^3}{6} (a_3/\sqrt{n}) (1 + \|x/\sqrt{d}\|^q) \\ &\lesssim 1 + \|x\|^{3+q}, \end{aligned} \quad (\text{C.15})$$

recalling $a_3/\sqrt{n} \leq 1$ by (2.4). The inequality (C.3) follows from the definition of W and (C.13). To prove (C.4), we use (C.1) to get

$$\begin{aligned} \mathbb{E} [\|\nabla^3 W(m + \sigma Z)\|^p] &\leq (a_3/\sqrt{n})^p \mathbb{E} \left[1 + (\|m + \sigma Z\|/\sqrt{d})^{pq} \right] \\ &\lesssim_{p,q} (a_3/\sqrt{n})^p \left(1 + \|m/\sqrt{d}\|^{pq} + \|\sigma\|^{pq} \right). \end{aligned} \quad (\text{C.16})$$

We take the p th root to conclude. The bound on $\mathbb{E} [\|\nabla^4 W(m + \sigma Z)\|^p]$ is shown analogously. \square

Proof Lemma C.2. Recall the expression (C.5) for V_0 in terms of W . We start by bounding $W(\hat{m}_0) - W(0)$. A Taylor expansion, the fact that $\|\hat{m}_0\| \leq 2\sqrt{2}$, and (2.4), gives

$$0 \leq W(\hat{m}_0) - W(0) \leq \frac{\|\hat{m}_0\|^2}{2} + \frac{\|\hat{m}_0\|^3}{3!} \frac{a_3(1 + \|\hat{m}_0/\sqrt{d}\|^q)}{\sqrt{n}} \leq C(q).$$

Now we prove (C.7). Using (C.2) and (C.6), we have

$$\begin{aligned} |V_0(x) - V_0(0)| &= |W(\hat{m}_0 + \hat{\sigma}_0 x) - W(\hat{m}_0)| \leq |W(\hat{m}_0 + \hat{\sigma}_0 x) - W(0)| + C(q) \\ &\lesssim 1 + \|\hat{m}_0 + \hat{\sigma}_0 x\|^{3+q} + C(q) \\ &\lesssim_q 1 + \|x\|^{3+q}. \end{aligned} \quad (\text{C.17})$$

To prove (C.8), fix $\|x\| \geq \sqrt{d} + 9$. Then

$$\|\hat{m}_0 + \hat{\sigma}_0 x\| \geq \|\hat{\sigma}_0 x\| - \|\hat{m}_0\| \geq \lambda_{\min}(\hat{\sigma}_0) \|x\| - \|\hat{m}_0\| \geq \sqrt{\frac{2}{3}} \|x\| - 2\sqrt{2} \geq \frac{1}{2} \sqrt{d}.$$

Hence

$$\begin{aligned} V_0(x) - V_0(0) &= W(\hat{m}_0 + \hat{\sigma}_0 x) - W(\hat{m}) \\ &= W(\hat{m}_0 + \hat{\sigma}_0 x) - W(0) + W(0) - W(\hat{m}_0) \\ &\geq c_0 \sqrt{d} \|\hat{m}_0 + \hat{\sigma}_0 x\| + W(0) - W(\hat{m}_0) \\ &\geq c_0 \sqrt{d} \left(\sqrt{\frac{2}{3}} \|x\| - 2\sqrt{2} \right) - C(q) \\ &\geq \frac{c_0}{2} \sqrt{d} \|x\| - C(q). \end{aligned} \quad (\text{C.18})$$

To get the last line, we used that $\sqrt{2/3} \|x\| - 2\sqrt{2} \geq \|x\|/2$ when $\|x\| \geq 9$. This concludes the proof of (C.8). To prove (C.9), we use the boundedness of $\|\hat{\sigma}_0\|$ to get that

$$\|\nabla^k V_0(y)\| \lesssim \|\nabla^k W(\hat{m}_0 + \hat{\sigma}_0 y)\|, \quad k = 3, 4. \quad (\text{C.19})$$

Therefore,

$$\mathbb{E} [\|\nabla^3 V_0(Y)\|^p] \lesssim \mathbb{E} [\|\nabla^3 W(\hat{m}_0 + \hat{\sigma}_0 Y)\|^p]. \quad (\text{C.20})$$

Now, $\hat{m}_0 + \hat{\sigma}_0 Y = \hat{m}_0 + \hat{\sigma}_0(\mu + S^{1/2}Z) = m + \sigma Z$, where $m = \hat{m}_0 + \hat{\sigma}_0 \mu$ and $\sigma = \hat{\sigma}_0 S^{1/2}$. Note that $\|m\| \leq 2\sqrt{2} + \sqrt{2}\|\mu\|$ and $\|\sigma\| \leq \sqrt{2}s \leq \sqrt{2}$. Therefore, applying (C.4) from Lemma C.1, we get

$$\|\|\nabla^3 V_0\|\|_{L^p(\mathcal{N}(\mu, s^2 I_d))} \lesssim_q \frac{a_3}{\sqrt{n}} \left(1 + \|\mu/\sqrt{d}\|^q \right).$$

The bound (C.10) on $\|\|\nabla^4 V_0\|\|_{L^p(\mathcal{N}(\mu, S))}$ is analogous. The bounds (C.11), (C.12) follow immediately from (C.9) and (C.10). \square

D Proofs from Section 4

For the equivalence of (4.5) and (2.12), see Section G. Here, we prove the key lemmas from Section 4.3. We start by making two observations: first, if g satisfies (2.7) and $f(x) = g(T^{-1}(x)) = g(\hat{m} + \hat{S}^{1/2}x)$, then f satisfies

$$|f_0(x)| = |f(x) - \gamma(f)| \leq \exp\left(\frac{c_0}{4} \sqrt{d} \|x\|\right) \quad \forall \|x\| \geq R_g \sqrt{d}. \quad (\text{D.1})$$

Second, using (4.16) from the proof outline, note that

$$V_0(x) - V_0(0) = \delta + r_3(x) - \|x\|^2/2$$

for some constant δ . Taking $x = 0$ in this equation, we get that $0 = \delta + r_3(0)$, i.e. $\delta = -r_3(0)$. But $r_3(0) = \frac{1}{3!} \langle \mathbf{A}_3, \mathbf{H}_3(0) \rangle + r_4(0) = r_4(0)$. Therefore, $\delta = -r_4(0)$, i.e. we have

$$V_0(x) - V_0(0) = -r_4(0) + r_3(x) + \frac{\|x\|^2}{2}. \quad (\text{D.2})$$

Now, to prove Lemma 4.2, we prove a number of supplementary lemmas. The overall proof structure is similar to that of [18]. The outline given in Section 4.3 of that work may be useful to follow the below proofs.

Lemma D.1. *Let g, R_g be as in Theorem 2.1, and $f = g \circ T$. Let $R \geq \max(R_g, 1 + 8\sqrt{2/d})$ and $\mathcal{U}(R) = \{\|x\| \leq R\sqrt{d}\}$. Then*

$$\begin{aligned} |E(f)| &\lesssim_q \left(1 + \|e^{-r_3} \mathbb{1}_{\mathcal{U}(R)}\|_4\right) \|f_0\|_2 \|r_3\|_{12}^3 + I(R, 9 + 3q, c_0/4), \\ |E| &\lesssim_q \left(1 + \|e^{-r_3} \mathbb{1}_{\mathcal{U}(R)}\|_2\right) \|r_3\|_4^2 + I(R, 6 + 2q, c_0/2), \end{aligned} \quad (\text{D.3})$$

where

$$I(R, p, b) = \frac{1}{(2\pi)^{d/2}} \int_{\|x\| \geq R\sqrt{d}}^{\infty} \|x\|^p e^{-b\sqrt{d}\|x\|} dx. \quad (\text{D.4})$$

See the end of this section for the proof. We see that we need to bound $I(R, p, b)$, $\|e^{-r_3} \mathbb{1}_{\mathcal{U}(R)}\|_4$, and $\|r_3\|_p$. The bounds on the first two quantities are stated in the next two lemmas. The bound on $\|r_3\|_p$ is given in Lemma E.6 in Section E.4.

Lemma D.2 (Bound on tail integral). *There is some C_p depending only on p such that if $R \geq C_p(1 \vee b^{-1})^2$ then $I(R, p, b) \lesssim_p e^{-Rbd/2}$. In particular,*

$$I(R, 9 + 3q, c_0/4) \vee I(R, 6 + 2q, c_0/2) \lesssim_q e^{-Rc_0d/8} \quad \text{if } R \geq C(q)(1 \vee c_0^{-1})^2$$

for some $C(q)$ depending only on q .

See Section G for the proof.

Lemma D.3 (Bound on exponential integral in $\mathcal{U}(R)$). *Let $R \geq 1$. Then*

$$\|e^{-r_3} \mathbb{1}_{\mathcal{U}(R)}\|_4 \leq \exp(C(q)R^{4+q}\omega),$$

where $\omega = a_3\epsilon + a_4\epsilon^2$.

See the end of this section for the proof. With these lemmas in hand, as well as a bound on $\|r_3\|_p$ in Lemma E.6 below, we can now prove Lemma 4.2.

Proof of Lemma 4.2. We have

$$|L(f)| \leq \|f_0\|_2 (\|r_3\|_2 + \|r_3\|_4^2) \lesssim_q \|f_0\|_2 \omega, \quad (\text{D.5})$$

using Lemma E.6 and the fact that $\omega \leq 1$ by (2.4). Next, let $R \geq \bar{C} := \max(R_g, 1 + 8\sqrt{2/d}, C(q)(1 \vee c_0^{-1})^2)$. We can then use (D.3) from Lemma D.1 to bound $|E|$ and $|E(f)|$. Furthermore, for this choice of R we can use Lemma D.2 to get that the tail integrals $I(R, 9 + 3q, c_0/4)$ and $I(R, 6 +$

$2q, c_0/2$) are both bounded by $e^{-Rc_0d/8}$. Finally, Lemma D.3 bounds $\|e^{-r_3} \mathbb{1}_{\mathcal{U}(R)}\|_4$, and Lemma E.6 bounds $\|r_3\|_p$. Combining all of these bounds, (D.3) reduces to

$$\begin{aligned} |E| &\lesssim_q \omega^2 \exp(C(q)R^{4+q}\omega) + e^{-\frac{Rc_0}{8}d} \\ |E(f)| &\lesssim_q \|f_0\|_2 \omega^3 \exp(C(q)R^{4+q}\omega) + e^{-\frac{Rc_0}{8}d}. \end{aligned} \quad (\text{D.6})$$

Now, if $\omega = 0$ then we get $|E|, |E(f)| \lesssim e^{-\frac{Rc_0}{8}d}$, and we can take $R \rightarrow \infty$ to conclude $|E|, |E(f)| = 0$. If $\omega > 0$ then by (2.4), we know $0 < \omega \leq 1$. Let

$$R = \max\left(\bar{C}, \frac{8}{c_0 d} \log(\omega^{-3})\right), \quad (\text{D.7})$$

so that $\exp(-(Rc_0/8)d) \leq \omega^3$. Furthermore, using that $\log(1/\omega)^{4+q}\omega \leq C(q)$ since $\omega \leq 1$, we have that for R as in (D.7), the exponent $C(q)R^{4+q}\omega$ in (D.6) is bounded above by some fixed function $C(R_g, c_0^{-1}, q)$. Using these bounds, we get

$$|E| \leq C(q, c_0^{-1}, R_g)\omega^2, \quad |E(f)| \leq C(q, c_0^{-1}, R_g)(1 + \|f_0\|_2)\omega^3 \quad (\text{D.8})$$

Note that this bound is still valid if $\omega = 0$. Finally, recalling the definition of ω concludes the proof. \square

Next, we prove Lemma 4.3, which follows immediately from Lemma E.6 and a few algebraic manipulations.

Proof of Lemma 4.3. We have

$$\begin{aligned} L(f) + \int f_0 p_3 d\gamma + \int f_0 \left(r_4 - \frac{1}{2}p_3^2\right) d\gamma &= \int f_0 \left(\left[-r_3 + \frac{1}{2}r_3^2\right] + r_3 - \frac{1}{2}p_3^2\right) d\gamma \\ &= \frac{1}{2} \int f_0 ((p_3 + r_4)^2 - p_3^2) d\gamma = \int f_0 p_3 r_4 d\gamma + \frac{1}{2} \int f_0 r_4^2 d\gamma. \end{aligned}$$

Next,

$$\begin{aligned} \left| \int f_0 p_3 r_4 d\gamma + \frac{1}{2} \int f_0 r_4^2 d\gamma \right| &\leq \|f_0\|_2 (\|p_3\|_4 \|r_4\|_4 + \|r_4\|_4^2) \\ &\lesssim \|f_0\|_2 (a_3 \epsilon a_4 \epsilon^2 + (a_4 \epsilon^2)^2) \\ &\lesssim \|f_0\|_2 (a_3 a_4 \epsilon^3 + a_4^2 \epsilon^4). \end{aligned}$$

Next, we have

$$\left| \int f_0 \left(r_4 - \frac{1}{2}p_3^2\right) d\gamma \right| \leq \|f_0\|_2 (\|r_4\|_2 + \|p_3\|_4^2) \lesssim \|f_0\|_2 (a_3^2 + a_4) \epsilon^2.$$

Finally,

$$\left| \int f p_3 d\gamma \right| \leq \|f_0\|_2 \|p_3\|_2 \lesssim \|f_0\|_2 a_3 \epsilon.$$

\square

Proof of Lemma D.1. A Taylor remainder argument shows that

$$|e^{-r_3} - 1 + r_3 - r_3^2/2| \leq |r_3|^3 + |r_3|^3 e^{-r_3}.$$

Therefore,

$$\begin{aligned} |E(f)| &\leq \int |f_0| |r_3|^3 d\gamma + \int |f_0| |r_3|^3 e^{-r_3} d\gamma \\ &= \int |f_0| |r_3|^3 d\gamma + \int |f_0| r_3^3 (e^{-r_3} \mathbb{1}_{\mathcal{U}(R)}) d\gamma + \int_{\mathcal{U}(R)^c} |f_0| |r_3|^3 e^{-r_3} d\gamma \\ &\leq \|f_0\|_2 \|r_3\|_{12}^3 (1 + \|e^{-r_3} \mathbb{1}_{\mathcal{U}(R)}\|_4) + \int_{\mathcal{U}(R)^c} |f_0| |r_3|^3 e^{-r_3} d\gamma. \end{aligned} \quad (\text{D.9})$$

In the third line we used generalized Hölder with powers 2, 4, 4. Next, by (D.2) we have

$$e^{-r_3(x)} d\gamma(x) = e^{-r_4(0)} (2\pi)^{-d/2} e^{V_0(0) - V_0(x)} dx.$$

We now apply the bound (C.8) on $V_0(0) - V_0(x)$ in the region $\|x\| \geq R\sqrt{d} \geq \sqrt{d} + 8\sqrt{2}$, to get the following bound on the tail integral:

$$\begin{aligned} \int_{\mathcal{U}(R)^c} |f_0| |r_3|^3 e^{-r_3} d\gamma(x) &= \frac{e^{-r_4(0)}}{(2\pi)^{\frac{d}{2}}} \int_{\mathcal{U}(R)^c} |f_0| |r_3|^3 e^{V_0(0) - V_0(x)} dx \\ &\lesssim_q \frac{e^{-r_4(0)}}{(2\pi)^{\frac{d}{2}}} \int_{\mathcal{U}(R)^c} |f_0| |r_3|^3 e^{-\frac{c_0}{2}\sqrt{d}\|x\|} dx. \end{aligned} \quad (\text{D.10})$$

Next, we have by (D.2) that $|r_3(x)| \leq |V_0(x) - V_0(0)| + |r_4(0)| + \|x\|^2$. Therefore, by the first bound in Lemma C.2 we have

$$\begin{aligned} |r_3(x)|^3 &\lesssim |V_0(x) - V_0(0)|^3 + |r_4(0)|^3 + \|x\|^6 \\ &\lesssim_q 1 + \|x\|^{9+3q} + \|x\|^6 + |r_4(0)|^3 \\ &\lesssim_q e^{|r_4(0)|} \|x\|^{9+3q}. \end{aligned} \quad (\text{D.11})$$

We have used that $\|x\| \geq R\sqrt{d} \geq 1$. Finally, note that $|f_0(x)| \leq \exp(c_0\sqrt{d}\|x\|/4)$ for $x \in \mathcal{U}(R)$, since $\|x\| \geq R\sqrt{d} \geq R_g\sqrt{d}$. Substituting this bound and (D.11) into (D.10), and then substituting the resulting bound into (D.9) gives

$$|E(f)| \lesssim_q \|f_0\|_2 \|r_3\|_{12}^3 (1 + \|e^{-r_3} \mathbb{1}_{\mathcal{U}(R)}\|_4) + e^{2|r_4(0)|} (2\pi)^{-\frac{d}{2}} \int_{\mathcal{U}(R)^c} \|x\|^{9+3q} e^{-\frac{c_0}{4}\sqrt{d}\|x\|} dx. \quad (\text{D.12})$$

Next, note that using Lemma E.5 and (2.4), we have $|r_4(0)| \leq C(q)a_4\epsilon^2 \leq C(q)$. Using this bound in (D.12), as well as the definition of $I(R, p, b)$ with $p = 9+3q$ and $b = c_0/4$ finishes the proof. The bound on $|E|$ is shown analogously. \square

Proof of Lemma D.3. Recall that $\mathbf{A}_3 = \mathbb{E}[\nabla^3 V_0(Z)]$, and that $r_3 = \frac{1}{3!} \langle \mathbf{A}_3, \mathbf{H}_3(x) \rangle + r_4(x)$. Hence

$$\begin{aligned} \|e^{-r_3} \mathbb{1}_{\mathcal{U}(R)}\|_4 &\leq \sup_{\|x\| \leq R\sqrt{d}} e^{|r_4(x)|} \|e^{-\frac{1}{3!} \langle \mathbf{A}_3, \mathbf{H}_3 \rangle} \mathbb{1}_{\mathcal{U}(R)}\|_4 \\ &\leq e^{C(q)R^{4+q}a_4\epsilon^2} \|e^{-\frac{1}{3!} \langle \mathbf{A}_3, \mathbf{H}_3 \rangle} \mathbb{1}_{\mathcal{U}(R)}\|_4, \end{aligned} \quad (\text{D.13})$$

using the pointwise bound on $|r_4|$ from Lemma E.5 to get the second line. To bound $\|e^{-\frac{1}{3!}\langle \mathbf{A}_3, \mathbf{H}_3 \rangle} \mathbb{1}_{\mathcal{U}(R)}\|_4$, we use that if the restriction of f to an open set \mathcal{U} is Lipschitz has Lipschitz constant L , then

$$\left| \int_{\mathcal{U}} e^f d\gamma \right| \leq \exp(CL^2 + L\|f\|_1)$$

for some absolute constant C . This follows from Proposition 5.4.1 in [3] and the fact that γ satisfies a log Sobolev inequality with an absolute constant. We apply this result with $f = -\frac{1}{3!}\langle \mathbf{A}_3, \mathbf{H}_3 \rangle$. We have

$$\|\nabla f(x)\| = \frac{1}{2}\|\langle \mathbf{A}_3, xx^T - I_d \rangle\| \lesssim R^2 d\|\mathbf{A}_3\| \lesssim_q R^2 a_3 \epsilon, \quad \|x\| \leq R\sqrt{d}.$$

The bound on $d\|\mathbf{A}_3\|$ is by (C.11) of Lemma C.2. Therefore, f has Lipschitz constant $L = C(q)R^2 a_3 \epsilon$ when restricted to $\mathcal{U}(R)$. Furthermore, $\|\langle \mathbf{A}_3, \mathbf{H}_3 \rangle\|_1 \lesssim d\|\mathbf{A}_3\| \lesssim a_3 \epsilon$ by (E.32). Therefore,

$$\|e^{-\frac{1}{3!}\langle \mathbf{A}_3, \mathbf{H}_3 \rangle} \mathbb{1}_{\mathcal{U}(R)}\|_4 \leq \exp(C(q)R^4(a_3 \epsilon)^2).$$

Substituting this bound into (D.13), we get

$$\|e^{-r_3} \mathbb{1}_{\mathcal{U}(R)}\|_4 \leq \exp(C(q)[R^{4+q} a_4 \epsilon^2 + R^4(a_3 \epsilon)^2]) \leq \exp(C(q)R^{4+q} \omega), \quad (\text{D.14})$$

where $\omega = a_3 \epsilon + a_4 \epsilon^2$. Here we used that $R^4 \leq R^{4+q}$ and $(a_3 \epsilon)^2 \leq a_3 \epsilon$ by (2.4). \square

E Hermite Series Remainder

E.1 Brief Primer

Here is a brief primer on Hermite polynomials, multinomials, and series expansions. We let $H_k : \mathbb{R} \rightarrow \mathbb{R}$, $k = 0, 1, 2, \dots$ be the k th order probabilist's Hermite polynomial. We have $H_0(x) = 1$, $H_1(x) = x$, $H_2(x) = x^2 - 1$, $H_3(x) = x^3 - 3x$. For all $k \geq 1$, we can generate H_{k+1} from the recurrence relation

$$H_{k+1}(x) = xH_k(x) - kH_{k-1}(x), \quad k \geq 1. \quad (\text{E.1})$$

In particular, $H_k(x)$ is an order k polynomial given by a sum of monomials of the same parity as k . The H_k are orthogonal with respect to the Gaussian measure; namely, we have $\mathbb{E}[H_k(Z)H_j(Z)] = k!\delta_{jk}$. We also note for future reference that

$$\mathbb{E}[ZH_k(Z)H_{k+1}(Z)] = \mathbb{E}[(H_{k+1}(Z) + kH_{k-1}(Z))H_{k+1}(Z)] = (k+1)!, \quad (\text{E.2})$$

using the recurrence relation (E.1).

The Hermite multinomials are given by products of Hermite polynomials, and are indexed by $\gamma \in \{0, 1, 2, \dots\}^d$. Let $\gamma = (\gamma_1, \dots, \gamma_d)$, with $\gamma_j \in \{0, 1, 2, \dots\}$. Then

$$H_\gamma(x_1, \dots, x_d) = \prod_{j=1}^d H_{\gamma_j}(x_j),$$

which has order $|\gamma| := \sum_{j=1}^d \gamma_j$. Note that if $|\gamma| = k$ then $H_\gamma(x)$ is given by a sum of monomials of the same parity as k . Indeed, each $H_{\gamma_j}(x_j)$ is a linear combination of $x_j^{\gamma_j - 2p}$, $p \leq \lfloor \gamma_j/2 \rfloor$.

Thus $H_\gamma(x)$ is a linear combination of monomials of the form $\prod_{j=1}^d x_j^{\gamma_j - 2p_j}$, which has total order $k - 2 \sum_j p_j$. Using the independence of the entries of $Z = (Z_1, \dots, Z_d)$, we have

$$\mathbb{E}[H_\gamma(Z)H_{\gamma'}(Z)] = \gamma! \prod_{j=1}^d \delta_{\gamma_j, \gamma'_j},$$

where $\gamma! := \prod_{j=1}^d \gamma_j!$. The H_γ can also be defined explicitly as follows:

$$e^{-\|x\|^2/2} H_\gamma(x) = (-1)^{|\gamma|} \partial^\gamma \left(e^{-\|x\|^2/2} \right), \quad (\text{E.3})$$

where $\partial^\gamma f(x) = \partial_{x_1}^{\gamma_1} \dots \partial_{x_d}^{\gamma_d} f(x)$. This leads to the useful Gaussian integration by parts identity,

$$\mathbb{E}[f(Z)H_\gamma(Z)] = \mathbb{E}[\partial^\gamma f(Z)], \quad \text{if } f \in C^{|\gamma|}(\mathbb{R}^d).$$

The Hermite polynomials H_γ , $\gamma \in \{0, 1, \dots\}^d$, form a complete orthogonal basis of the Hilbert space of functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ with inner product $\langle f, g \rangle = \mathbb{E}[f(Z)g(Z)]$. In particular, if $f : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfies $\mathbb{E}[f(Z)^2] < \infty$, then f has the following Hermite expansion:

$$f(x) = \sum_{\gamma \in \{0, 1, \dots\}^d} \frac{1}{\gamma!} a_\gamma(f) H_\gamma(x), \quad a_\gamma(f) := \mathbb{E}[f(Z)H_\gamma(Z)]. \quad (\text{E.4})$$

Let

$$r_k(x) = f(x) - \sum_{|\gamma| \leq k-1} \frac{1}{\gamma!} a_\gamma(f) H_\gamma(x) \quad (\text{E.5})$$

be the remainder of the Hermite series expansion of f after taking out the order $\leq k-1$ polynomials. We can write r_k as an integral of f against a kernel. Namely, define

$$K(x, y) = \sum_{|\gamma| \leq k-1} \frac{1}{\gamma!} H_\gamma(x) H_\gamma(y). \quad (\text{E.6})$$

Note that

$$\mathbb{E}[f(Z)K(x, Z)] = \sum_{|\gamma| \leq k-1} \frac{1}{\gamma!} a_\gamma(f) H_\gamma(x)$$

is the truncated Hermite series expansion of f . Therefore, the remainder r_k can be written as

$$r_k(x) = f(x) - \mathbb{E}[f(Z)K(x, Z)] = \mathbb{E}[(f(x) - f(Z))K(x, Z)],$$

using that $\mathbb{E}[K(x, Z)] = 1$.

E.2 Exact Expression for the Remainder

Lemma E.1. *Let $k \geq 1$ and r_k , K , be as in (E.5), (E.6), respectively. Assume that $f \in C^1$, and that $\|\nabla f(x)\| \lesssim e^{c\|x\|^2}$ for some $0 \leq c < 1/2$. Then*

$$\begin{aligned} r_k(x) &= \mathbb{E}[(f(x) - f(Z))K(x, Z)] \\ &= \int_0^1 \sum_{i=1}^d \sum_{|\gamma|=k-1} \frac{1}{\gamma!} \mathbb{E}[\partial_i f((1-t)Z + tx) (H_{\gamma+e_i}(x) H_\gamma(Z) - H_{\gamma+e_i}(Z) H_\gamma(x))] dt. \end{aligned} \quad (\text{E.7})$$

The proof relies on the following identity:

Lemma E.2. *For each $i = 1, \dots, d$, it holds that*

$$K(x, y) = \frac{1}{x_i - y_i} \sum_{|\gamma|=k-1} \frac{1}{\gamma!} (H_{\gamma+e_i}(x)H_\gamma(y) - H_{\gamma+e_i}(y)H_\gamma(x)). \quad (\text{E.8})$$

The proof of this identity is given at the end of the section.

Proof of Lemma E.1. Write

$$\begin{aligned} f(x) - f(Z) &= \int_0^1 (x - Z)^T \nabla f((1-t)Z + tx) dt \\ &= \sum_{i=1}^d \int_0^1 (x_i - Z_i) \partial_i f((1-t)Z + tx) dt, \end{aligned} \quad (\text{E.9})$$

so that, using (E.8), we have

$$\begin{aligned} \mathbb{E}[(f(x) - f(Z))K(x, Z)] &= \sum_{i=1}^d \sum_{|\gamma|=k-1} \frac{1}{\gamma!} \mathbb{E} \left[\int_0^1 \partial_i f((1-t)Z + tx) (H_{\gamma+e_i}(x)H_\gamma(Z) - H_{\gamma+e_i}(Z)H_\gamma(x)) dt \right]. \end{aligned} \quad (\text{E.10})$$

By assumption,

$$\begin{aligned} \sup_{t \in [0,1]} |\partial_i f((1-t)Z + tx)| (|H_\gamma(Z)| + |H_{\gamma+e_i}(Z)|) \\ \lesssim \exp(c\|Z\|^2 + 2c\|Z\|\|x\|) (|H_\gamma(Z)| + |H_{\gamma+e_i}(Z)|) \end{aligned} \quad (\text{E.11})$$

for some $0 \leq c < 1/2$. The right-hand side is integrable with respect to the Gaussian measure, and therefore we can interchange the expectation and the integral in (E.10). Therefore,

$$\begin{aligned} \mathbb{E}[(f(x) - f(Z))K(x, Z)] &= \int_0^1 \sum_{i=1}^d \sum_{|\gamma|=k-1} \frac{1}{\gamma!} \mathbb{E} [\partial_i f((1-t)Z + tx) (H_{\gamma+e_i}(x)H_\gamma(Z) - H_{\gamma+e_i}(Z)H_\gamma(x))] dt. \end{aligned} \quad (\text{E.12})$$

□

Proof of Lemma E.2. Without loss of generality, assume $i = 1$. To simplify the proof, we will also assume $d = 2$. The reader can check that the proof goes through in the same way for general d . By the recursion relation (E.1) for 1d Hermite polynomials, we get that

$$H_{\gamma_1+1, \gamma_2}(x) = x_1 H_{\gamma_1, \gamma_2}(x) - \gamma_1 H_{\gamma_1-1, \gamma_2}(x),$$

where $x = (x_1, x_2)$. Multiply this equation by $H_\gamma(y)$ (where $y = (y_1, y_2)$) and swap x and y , to get the two equations

$$\begin{aligned} H_{\gamma_1+1, \gamma_2}(x)H_{\gamma_1, \gamma_2}(y) &= x_1 H_{\gamma_1, \gamma_2}(x)H_{\gamma_1, \gamma_2}(y) - \gamma_1 H_{\gamma_1-1, \gamma_2}(x)H_{\gamma_1, \gamma_2}(y), \\ H_{\gamma_1+1, \gamma_2}(y)H_{\gamma_1, \gamma_2}(x) &= y_1 H_{\gamma_1, \gamma_2}(x)H_{\gamma_1, \gamma_2}(y) - \gamma_1 H_{\gamma_1-1, \gamma_2}(y)H_{\gamma_1, \gamma_2}(x). \end{aligned} \quad (\text{E.13})$$

Let

$$S_{\gamma_1, \gamma_2} = H_{\gamma_1+1, \gamma_2}(x)H_{\gamma_1, \gamma_2}(y) - H_{\gamma_1+1, \gamma_2}(y)H_{\gamma_1, \gamma_2}(x).$$

Subtracting the second equation of (E.13) from the first one, and using the S_{γ_1, γ_2} notation, gives

$$S_{\gamma_1, \gamma_2} = (x_1 - y_1)H_{\gamma_1, \gamma_2}(x)H_{\gamma_1, \gamma_2}(y) + \gamma_1 S_{\gamma_1-1, \gamma_2} \quad (\text{E.14})$$

and hence

$$\frac{S_{\gamma_1, \gamma_2}}{\gamma_1! \gamma_2!} = (x_1 - y_1) \frac{H_{\gamma_1, \gamma_2}(x)H_{\gamma_1, \gamma_2}(y)}{\gamma_1! \gamma_2!} + \frac{S_{\gamma_1-1, \gamma_2}}{(\gamma_1 - 1)! \gamma_2!}.$$

Iterating this recursive relationship $\gamma_1 - 1$ times, we get

$$\frac{S_{\gamma_1, \gamma_2}}{\gamma_1! \gamma_2!} = (x_1 - y_1) \sum_{j=0}^{\gamma_1-1} \frac{H_{\gamma_1-j, \gamma_2}(x)H_{\gamma_1-j, \gamma_2}(y)}{(\gamma_1 - j)! \gamma_2!} + \frac{S_{0, \gamma_2}}{0! \gamma_2!}. \quad (\text{E.15})$$

Now, we have

$$\begin{aligned} S_{0, \gamma_2} &= H_{1, \gamma_2}(x)H_{0, \gamma_2}(y) - H_{1, \gamma_2}(y)H_{0, \gamma_2}(x) \\ &= H_1(x_1)H_{\gamma_2}(x_2)H_{\gamma_2}(y_2) - H_1(y_1)H_{\gamma_2}(y_2)H_{\gamma_2}(x_2) \\ &= (x_1 - y_1)H_{\gamma_2}(x_2)H_{\gamma_2}(y_2) \\ &= (x_1 - y_1)H_{0, \gamma_2}(x)H_{0, \gamma_2}(y). \end{aligned} \quad (\text{E.16})$$

Therefore,

$$\frac{S_{0, \gamma_2}}{0! \gamma_2!} = (x_1 - y_1) \frac{H_{\gamma_1-j, \gamma_2}(x)H_{\gamma_1-j, \gamma_2}(y)}{(\gamma_1 - j)! \gamma_2!}, \quad j = \gamma_1$$

so (E.15) can be written as

$$\frac{S_{\gamma_1, \gamma_2}}{\gamma_1! \gamma_2!} = (x_1 - y_1) \sum_{j=0}^{\gamma_1} \frac{H_{\gamma_1-j, \gamma_2}(x)H_{\gamma_1-j, \gamma_2}(y)}{(\gamma_1 - j)! \gamma_2!}$$

and hence

$$\begin{aligned} \frac{1}{x_1 - y_1} \sum_{\gamma_1 + \gamma_2 = k-1} \frac{S_{\gamma_1, \gamma_2}}{\gamma_1! \gamma_2!} &= \sum_{\gamma_1 + \gamma_2 = k-1} \sum_{j=0}^{\gamma_1} \frac{H_{\gamma_1-j, \gamma_2}(x)H_{\gamma_1-j, \gamma_2}(y)}{(\gamma_1 - j)! \gamma_2!} \\ &= \sum_{\gamma_1 + \gamma_2 \leq k-1} \frac{H_{\gamma_1, \gamma_2}(x)H_{\gamma_1, \gamma_2}(y)}{\gamma_1! \gamma_2!} = K(x, y), \end{aligned} \quad (\text{E.17})$$

using the observation that

$$\begin{aligned} &\{(\gamma_1 - j, \gamma_2) : \gamma_1 + \gamma_2 = k-1, 0 \leq j \leq \gamma_1\} \\ &= \{(\tilde{\gamma}_1, \gamma_2) : \tilde{\gamma}_1 + \gamma_2 \leq k-1\}. \end{aligned} \quad (\text{E.18})$$

Substituting back in the definition of S_{γ_1, γ_2} gives the desired result. \square

E.3 Hermite Series Remainder in Tensor Form

Using (E.7), it is difficult to obtain an upper bound on $|r_k(x)|$, since we need to sum over all γ of order $k-1$. In this section, we obtain a more compact representation of r_k in terms of a scalar product of k -tensors. We then take advantage of a very useful representation of the tensor of order- k Hermite polynomials, as an expectation of a vector outer product. This allows us to bound the scalar product in the r_k formula in terms of an operator norm rather than a Frobenius norm (the latter would incur larger d dependence).

First, let us put all the unique k th order Hermite polynomials into a tensor of d^k entries, some of which are repeating, enumerated by multi-indices $\alpha = (\alpha_1, \dots, \alpha_k) \in [d]^k$. Here, $[d] = \{1, \dots, d\}$. We do so as follows: given $\alpha \in [d]^k$, define $\gamma(\alpha) = (\gamma_1(\alpha), \dots, \gamma_d(\alpha))$ by

$$\gamma_j(\alpha) = \sum_{\ell=1}^k \mathbb{1}\{\alpha_\ell = j\},$$

i.e. $\gamma_j(\alpha)$ counts how many times index j appears in α . For this reason, we use the term *counting index* to denote indices of the form $\gamma = (\gamma_1, \dots, \gamma_d) \in \{0, 1, 2, \dots\}^d$, whereas we use the standard term “multi-index” to refer to the α ’s. Note that we automatically have $|\gamma(\alpha)| = k$ if $\alpha \in [d]^k$. Now, for $x \in \mathbb{R}^d$, define $\mathbf{H}_0(x) = 1$ and $\mathbf{H}_k(x)$, $k \geq 1$ as the tensor

$$\mathbf{H}_k(x) = \{H_{\gamma(\alpha)}(x)\}_{\alpha \in [d]^k}, \quad x \in \mathbb{R}^d.$$

When enumerating the entries of \mathbf{H}_k , we write $H_k^{(\alpha)}$ to denote $H_{\gamma(\alpha)}$. Note that for each γ with $|\gamma| = k$, there are $\binom{k}{\gamma}$ α ’s such that $\gamma(\alpha) = \gamma$.

Example E.1. Consider the $\alpha = (i, j, j, k, k, k)$ entry of the tensor $\mathbf{H}_6(x)$, where $i, j, k \in [d]$ are all distinct. We count that i occurs once, j occurs twice, and k occurs thrice. Thus

$$H_6^{(i,j,j,k,k,k)}(x) = H_1(x_i)H_2(x_j)H_3(x_k) = x_i(x_j^2 - 1)(x_k^3 - 3x_k).$$

The first two tensors $\mathbf{H}_1, \mathbf{H}_2$ can be written down explicitly. For the entries of \mathbf{H}_1 , we simply have $H_1^{(i)}(x) = H_1(x_i) = x_i$, i.e. $\mathbf{H}_1(x) = x$. For the entries of \mathbf{H}_2 , we have $H_2^{(i,i)}(x) = H_2(x_i) = x_i^2 - 1$ and $H_2^{(i,j)}(x) = H_1(x_i)H_1(x_j) = x_i x_j$, $i \neq j$. Thus $\mathbf{H}_2(x) = xx^T - I_d$.

We now group the terms in the Hermite series expansion (E.4) based on the order $|\gamma|$. Consider all γ in the sum such that $|\gamma| = k$. We claim that

$$\sum_{|\gamma|=k} \frac{1}{\gamma!} a_\gamma(f) H_\gamma(x) = \frac{1}{k!} \sum_{\alpha \in [d]^k} a_{\gamma(\alpha)}(f) H_k^{(\alpha)}(x). \quad (\text{E.19})$$

Indeed, for a fixed γ such that $|\gamma| = k$, there are $\binom{k}{\gamma}$ α ’s in $[d]^k$ for which $\gamma(\alpha) = \gamma$, and the summands in the right-hand sum corresponding to these α ’s are all identical, equalling $a_\gamma(f) H_\gamma(x)$. Thus we obtain $\binom{k}{\gamma}$ copies of $a_\gamma(f) H_\gamma(x)$, and it remains to note that $\binom{k}{\gamma}/k! = 1/\gamma!$.

Analogously to $\mathbf{H}_k(x)$, define the tensor $\mathbf{A}_k \in (\mathbb{R}^d)^{\otimes k}$, whose α ’th entry is

$$\mathbf{A}_k^{(\alpha)} = a_{\gamma(\alpha)} = \mathbb{E}[f(Z)H_{\gamma(\alpha)}(Z)] = \mathbb{E}[f(Z)H_k^{(\alpha)}(Z)].$$

We then see that the sum (E.19) can be written as $\frac{1}{k!} \langle \mathbf{A}_k, \mathbf{H}_k(x) \rangle$, and hence the series expansion of f can be written as

$$f(x) = \sum_{k=0}^{\infty} \frac{1}{k!} \langle \mathbf{A}_k(f), \mathbf{H}_k(x) \rangle, \quad \mathbf{A}_k(f) := \mathbb{E}[f(Z)\mathbf{H}_k(Z)]. \quad (\text{E.20})$$

Lemma E.3. Let f satisfy the assumptions of Lemma E.1, and additionally, assume $f \in C^k$. Then the remainder r_k , given as (E.7) in Lemma E.1, can also be written in the form

$$r_k(x) = \int_0^1 \frac{(1-t)^{k-1}}{(k-1)!} \mathbb{E} \left[\left\langle \nabla^k f((1-t)Z + tx), \mathbf{H}_k(x) - Z \otimes \mathbf{H}_{k-1}(x) \right\rangle \right] \quad (\text{E.21})$$

Proof. Recall that $\partial^\gamma := \partial_1^{\gamma_1} \dots \partial_d^{\gamma_d}$, and that

$$H_\gamma(z) e^{-\|z\|^2/2} = (-1)^{|\gamma|} \partial^\gamma (e^{-\|z\|^2/2}).$$

We then have for $|\gamma| = k-1$,

$$\begin{aligned} \mathbb{E} [\partial_i f((1-t)Z + tx) H_\gamma(Z)] &= (1-t)^{k-1} \mathbb{E} [\partial^{\gamma+e_i} f], \\ \mathbb{E} [\partial_i f((1-t)Z + tx) H_{\gamma+e_i}(Z)] &= (1-t)^{k-1} \mathbb{E} [\partial^{\gamma+e_i} f Z_i], \end{aligned} \quad (\text{E.22})$$

using the fact that $f \in C^k$. We omitted the argument $(1-t)Z + tx$ from the right-hand side for brevity. To get the second equation, we moved only γ of the $\gamma + e_i$ derivatives from $e^{-\|z\|^2/2}$ onto $\partial_i f$, leaving $-\partial_i(e^{-\|z\|^2/2}) = z_i$. Substituting these two equations into (E.7), we get

$$\begin{aligned} &\mathbb{E} [(f(x) - f(Z)) K(x, Z)] \\ &= \int_0^1 (1-t)^{k-1} \sum_{i=1}^d \sum_{|\gamma|=k-1} \frac{1}{\gamma!} \mathbb{E} [(\partial^{\gamma+e_i} f) (H_{\gamma+e_i}(x) - Z_i H_\gamma(x))] dt \\ &= \frac{1}{(k-1)!} \int_0^1 (1-t)^{k-1} \sum_{i=1}^d \sum_{|\gamma|=k-1} \binom{k-1}{\gamma} \mathbb{E} [(\partial^{\gamma+e_i} f) (H_{\gamma+e_i}(x) - Z_i H_\gamma(x))] dt \end{aligned} \quad (\text{E.23})$$

Now, define the sets

$$\begin{aligned} A &= \{(i, \gamma + e_i) : i = 1, \dots, d, \gamma \in \{0, 1, \dots\}^d, |\gamma| = k-1\}, \\ B &= \{(i, \tilde{\gamma}) : \tilde{\gamma} \in \{0, 1, \dots\}^d, |\tilde{\gamma}| = k, \tilde{\gamma}_i \geq 1\}. \end{aligned} \quad (\text{E.24})$$

It is straightforward to see that $A = B$. Therefore,

$$\begin{aligned} &\sum_{i=1}^d \sum_{|\gamma|=k-1} \binom{k-1}{\gamma} \mathbb{E} [\partial^{\gamma+e_i} f] H_{\gamma+e_i}(x) \\ &= \sum_{|\tilde{\gamma}|=k} \sum_{i: \tilde{\gamma}_i \geq 1} \binom{k-1}{\tilde{\gamma} - e_i} \mathbb{E} [\partial^{\tilde{\gamma}} f] H_{\tilde{\gamma}}(x) \\ &= \sum_{|\tilde{\gamma}|=k} \sum_{i: \tilde{\gamma}_i \geq 1} \binom{k}{\tilde{\gamma}} \frac{\tilde{\gamma}_i}{k} \mathbb{E} [\partial^{\tilde{\gamma}} f] H_{\tilde{\gamma}}(x) \\ &= \sum_{|\tilde{\gamma}|=k} \binom{k}{\tilde{\gamma}} \mathbb{E} [\partial^{\tilde{\gamma}} f] H_{\tilde{\gamma}}(x) = \langle \mathbb{E} [\nabla^k f], \mathbf{H}_k(x) \rangle \end{aligned} \quad (\text{E.25})$$

Next, note that

$$\sum_{|\gamma|=k-1} \binom{k-1}{\gamma} \partial^\gamma \partial_i f H_\gamma(x) = \langle \nabla^{k-1} \partial_i f, \mathbf{H}_{k-1}(x) \rangle,$$

and therefore

$$\sum_{i=1}^d \sum_{|\gamma|=k-1} \binom{k-1}{\gamma} \mathbb{E}[(\partial^{\gamma+e_i} f) Z_i] H_\gamma(x) = \mathbb{E}[\langle \nabla^k f, Z \otimes \mathbf{H}_{k-1}(x) \rangle]. \quad (\text{E.26})$$

Substituting (E.25) and (E.26) into (E.23) gives

$$\begin{aligned} r_k(x) &= \mathbb{E}[(f(x) - f(Z)) K(x, Z)] \\ &= \int_0^1 \frac{(1-t)^{k-1}}{(k-1)!} \mathbb{E} \left[\left\langle \nabla^k f((1-t)Z + tx), \mathbf{H}_k(x) - Z \otimes \mathbf{H}_{k-1}(x) \right\rangle \right] \end{aligned} \quad (\text{E.27})$$

□

In the next section, we obtain a pointwise upper bound on $|r_k(x)|$ in the case $f = V_0$. In order for this bound to be tight in its dependence on d , we need a supplementary result on inner products with Hermite tensors. To motivate this supplementary result, consider bounding the inner product in (E.21) by the product of the Frobenius norms of the tensors on either side. As a rough heuristic, $\|\nabla^k f\|_F \sim d^{k/2} \|\nabla^k f\|$, where recall that $\|\nabla^k f\|$ is the *operator norm* of $\nabla^k f$. Therefore, we would prefer to bound the inner product in terms of $\|\nabla^k f\|$ to get a tighter dependence on d . Apriori, however, this seems impossible, since $\mathbf{H}_k(x)$ is not given by an outer product of k vectors. But the following representation of the order k Hermite polynomials will make this possible.

$$\mathbf{H}_k(x) = \mathbb{E}[(x + iZ)^{\otimes k}], \quad (\text{E.28})$$

where $Z \sim \mathcal{N}(0, I_d)$. Using (E.28), we can bound scalar products of the form $\langle \nabla^k f, \mathbf{H}_k(x) \rangle$ and $\langle \nabla^k f, Z \otimes \mathbf{H}_{k-1}(x) \rangle$ in terms of the operator norm of $\nabla^k f$. More generally, we have the following lemma.

Lemma E.4. *Let $T \in (\mathbb{R}^d)^{\otimes k}$ be a k -tensor, and $v \in \mathbb{R}^d$. Then for all $0 \leq \ell \leq k$, we have*

$$|\langle T, v^{\otimes \ell} \otimes \mathbf{H}_{k-\ell}(x) \rangle| \lesssim_k \|T\| \|v\|^\ell (\|x\|^{k-\ell} + d^{\frac{k-\ell}{2}}).$$

Proof. Using (E.28), we have

$$\langle T, v^{\otimes \ell} \otimes \mathbf{H}_{k-\ell}(x) \rangle = \mathbb{E}[\langle T, v^{\otimes \ell} \otimes (x + iZ)^{\otimes (k-\ell)} \rangle]$$

and hence

$$\begin{aligned} |\langle T, v^{\otimes \ell} \otimes \mathbf{H}_{k-\ell}(x) \rangle| &\leq \mathbb{E} |\langle T, v^{\otimes \ell} \otimes (x + iZ)^{\otimes (k-\ell)} \rangle| \\ &\leq \|T\| \|v\|^\ell \mathbb{E} \left[\|x + iZ\|^{k-\ell} \right] \\ &\lesssim_k \|T\| \|v\|^\ell (\|x\|^{k-\ell} + \sqrt{d}^{k-\ell}). \end{aligned} \quad (\text{E.29})$$

□

E.4 Hermite Remainder Formula Applied to V_0

We now consider the case $f = V_0$, where V_0 is defined in (C.5) and satisfies the properties stated in Lemma C.2. Recall that

$$r_k(x) = V_0(x) - \sum_{j=0}^{k-1} \frac{1}{j!} \langle \mathbf{A}_j, \mathbf{H}_j(x) \rangle, \quad (\text{E.30})$$

where $\mathbf{A}_j = \mathbb{E}[V_0(Z) \mathbf{H}_j(Z)] = \mathbb{E}[\nabla^j V_0(Z)]$, and the second equality holds provided $V_0 \in C^j$. We first prove a pointwise bound on r_4 , and then bound the $L^p(\gamma)$ norm of r_3 and r_4 .

Lemma E.5. *It holds*

$$|r_4(x)| \lesssim_q \frac{a_4}{n} \left(\|x\|^4 + d^2 + \sqrt{d}\|x\|^3 \right) (1 + \|x/\sqrt{d}\|^q).$$

Proof. Let $\nabla^k V_0$ be shorthand for $\nabla^k V_0((1-t)Z + tx)$. Using (E.21) for $f = V_0$, we have

$$\begin{aligned} |r_4(x)| &\leq \int_0^1 \mathbb{E} [|\langle \nabla^4 V_0, \mathbf{H}_4(x) - Z \otimes \mathbf{H}_3(x) \rangle|] dt \\ &\lesssim \int_0^1 \mathbb{E} \left[\|\nabla^4 V_0\| \left(\|x\|^4 + d^2 + \|Z\|(\|x\|^3 + d^{3/2}) \right) \right] dt \\ &\lesssim \left(\|x\|^4 + d^2 + \sqrt{d}\|x\|^3 \right) \int_0^1 \sqrt{\mathbb{E} [\|\nabla^4 V_0((1-t)Z + tx)\|^2]} dt \\ &= \left(\|x\|^4 + d^2 + \sqrt{d}\|x\|^3 \right) \int_0^1 \|\nabla^4 V_0\|_{L^2(\mathcal{N}(x, (1-t)^2 I_d))} dt \\ &\lesssim_q \frac{a_4}{n} \left(\|x\|^4 + d^2 + \sqrt{d}\|x\|^3 \right) (1 + \|x/\sqrt{d}\|^q) \end{aligned} \quad (\text{E.31})$$

In the second line we used Lemma E.4, and in the third line we used Cauchy-Schwarz and the fact that $\mathbb{E} [\|Z\|^k] \lesssim_k d^{k/2}$. To get the fifth line, we used (C.10) of Lemma C.2 \square

The d -dependence in the bound on $\|r_3\|_p$ in the next lemma relies on Lemmas 2.1, C.2 and C.1 of [18]. Combined, the three lemmas give that for a symmetric $d \times d \times d$ tensor T , it holds

$$\|\langle T, \mathbf{H}_3 \rangle\|_{2k} \lesssim_k \|\langle T, \mathbf{H}_3 \rangle\|_2 \lesssim \|T\|_F \lesssim d\|T\| \quad (\text{E.32})$$

Lemma E.6. *Let $p_3(x) = \frac{1}{3!} \langle \mathbf{A}_3, \mathbf{H}_3 \rangle(x)$. It holds*

$$\begin{aligned} \|r_4\|_p &\lesssim_{p,q} a_4 \epsilon^2 \\ \|p_3\|_p &\lesssim_{p,q} a_3 \epsilon, \\ \|r_3\|_p &\lesssim_{p,q} a_3 \epsilon + a_4 \epsilon^2 =: \omega, \end{aligned}$$

where $\epsilon = d/\sqrt{n}$.

Proof. For the bound on $\|r_4\|_p$, we use Lemma E.5 which gives

$$\mathbb{E} [|r_4(Z)|^p] \lesssim_{p,q} \left(\frac{a_4}{n} \right)^p \mathbb{E} \left[\left(\|Z\|^4 + d^2 + \sqrt{d}\|Z\|^3 \right)^p (1 + \|Z/\sqrt{d}\|^q)^p \right] \lesssim_{p,q} \left(\frac{a_4}{n} \right)^p d^{2p}.$$

Taking the p th root gives $\|r_4\|_p \lesssim_{p,q} a_4 d^2/n = a_4 \epsilon^2$. To bound $\|p_3\|_p$, we use (E.32) and (C.11) of Lemma C.2:

$$\|p_3\|_p \leq \|\langle \mathbf{A}_3, \mathbf{H}_3 \rangle\|_p \lesssim_p d\|\mathbf{A}_3\| \leq d\mathbb{E} [\|\nabla^3 V_0(Z)\|] \lesssim_q a_3 \epsilon. \quad (\text{E.33})$$

Finally, note that $r_3 = p_3 + r_4$, and the bound on $\|r_3\|_p$ follows. \square

F Proofs from Section 5: existence and uniqueness of solutions to stationarity conditions

For the proof of Lemma 2.1 using Lemma 5.1, see Section G. In this section, we prove Lemma 5.1 via the auxiliary lemmas outlined in Section 5.2. The proofs rely on tensor-matrix and tensor-vector scalar products. Let us review the rules of such scalar products, and how to bound the

operator norms of these quantities. Let $v \in \mathbb{R}^d$, $A \in \mathbb{R}^{d \times d}$, and $T \in \mathbb{R}^{d \times d \times d}$. We define the vector $\langle T, A \rangle \in \mathbb{R}^d$ and the matrix $\langle T, v \rangle \in \mathbb{R}^{d \times d}$ by

$$\begin{aligned}\langle T, A \rangle_i &= \sum_{j,k=1}^d T_{ijk} A_{jk}, \quad i = 1, \dots, d, \\ \langle T, v \rangle_{ij} &= \sum_{k=1}^d T_{ijk} v_k, \quad i, j = 1, \dots, d.\end{aligned}\tag{F.1}$$

We will always sum over the last two or last one indices of the tensor. Note that the norm of the matrix $\langle T, v \rangle$ is given by $\|\langle T, v \rangle\| = \sup_{\|u\|=\|w\|=1} u^T \langle T, v \rangle w$, and we have

$$u^T \langle T, v \rangle w = \sum_{i,j=1}^d u_i w_j \sum_{k=1}^d T_{ijk} v_k = \langle T, u \otimes w \otimes v \rangle \leq \|T\| \|v\|.$$

Therefore, $\|\langle T, v \rangle\| \leq \|T\| \|v\|$.

We also review the notion of operator norm for derivatives of a function, and note the distinction between this kind of operator norm and the standard tensor operator norm. Specifically, consider a C^2 function $f = (f_1, \dots, f_d) : \mathbb{R}^d \times \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^d$, where $\mathbb{R}^{d \times d}$ is endowed with the standard matrix norm. Then $\nabla_\sigma f(m, \sigma)$ is a linear functional from $\mathbb{R}^{d \times d}$ to \mathbb{R}^d , and we let $\langle \nabla_\sigma f(m, \sigma), A \rangle \in \mathbb{R}^d$ denote the application of $\nabla_\sigma f(m, \sigma)$ to A . Note that we can represent $\nabla_\sigma f$ by the $d \times d \times d$ tensor $(\nabla_{\sigma_{jk}} f_i)_{i,j,k=1}^d$, so that $\langle \nabla_\sigma f(m, \sigma), A \rangle$ coincides with the definition given above of tensor-matrix scalar products. However, $\|\nabla_\sigma f\|_{\text{op}}$ is *not* the standard tensor operator norm. Rather,

$$\|\nabla_\sigma f\|_{\text{op}} = \sup_{A \in \mathbb{R}^{d \times d}, \|A\|=1} \|\langle \nabla_\sigma f, A \rangle\| = \sup_{\substack{A \in \mathbb{R}^{d \times d}, \|A\|=1, \\ u \in \mathbb{R}^d, \|u\|=1}} \langle \nabla_\sigma f, A \otimes u \rangle.$$

We continue to write $\|\nabla_\sigma f\|$ to denote the *standard* tensor operator norm, i.e.

$$\|\nabla_\sigma f\| = \sup_{\substack{u, v, w \in \mathbb{R}^d, \\ \|u\|=\|v\|=\|w\|=1}} \langle \nabla_\sigma f, u \otimes v \otimes w \rangle.$$

Note also that $\nabla_m f \in \mathbb{R}^{d \times d}$ is a matrix, and that

$$\begin{aligned}\max &\left(\|\nabla_\sigma f(m, \sigma)\|_{\text{op}}, \|\nabla_m f(m, \sigma)\|_{\text{op}} \right) \\ &\leq \|\nabla f(m, \sigma)\|_{\text{op}} \leq \|\nabla_\sigma f(m, \sigma)\|_{\text{op}} + \|\nabla_m f(m, \sigma)\|_{\text{op}}.\end{aligned}\tag{F.2}$$

Finally, recall the notation

$$\begin{aligned}B_r(0, 0) &= \{(m, \sigma) \in \mathbb{R}^d \times \mathbb{R}^{d \times d} : \|\sigma\|^2 + \|m\|^2 \leq r^2\}, \\ B_r &= \{\sigma \in \mathbb{R}^{d \times d} : \|\sigma\| \leq r\}, \\ S_{c_1, c_2} &= \{\sigma \in \mathbb{S}_+^d : c_1 I_d \preceq \sigma \preceq c_2 I_d\}.\end{aligned}\tag{F.3}$$

The proof of Lemma 5.2 uses the following lemma.

Lemma F.1 (Lemma 1.3 in Chapter XIV of [21]). *Let U be open in a Banach space E , and let $f : U \rightarrow E$ be of class C^1 . Assume that $f(0) = 0$ and $f'(0) = I$. Let $r > 0$ be such that $\bar{B}_r(0) \subset U$. If*

$$|f'(z) - f'(x)| \leq s, \quad \forall z, x \in \bar{B}_r(0)$$

for some $s \in (0, 1)$, then f maps $\bar{B}_r(0)$ bijectively onto $\bar{B}_{(1-s)r}(0)$.

Proof of Lemma 5.2. Let $\phi : \mathbb{R}^d \times \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^d \times \mathbb{R}^{d \times d}$ be given by $\phi(m, \sigma) = (f(m, \sigma), \sigma)$, so that $\phi(0, 0) = (0, 0)$, and

$$\nabla \phi(m, \sigma) = \begin{pmatrix} \nabla_m f(m, \sigma) & \nabla_\sigma f(m, \sigma) \\ 0 & I_d \end{pmatrix}. \quad (\text{F.4})$$

For each $(m, \sigma), (m', \sigma') \in B_r(0, 0)$, we have

$$\begin{aligned} \|\nabla \phi(m, \sigma) - \nabla \phi(m', \sigma')\|_{\text{op}} &= \|\nabla f(m, \sigma) - \nabla f(m', \sigma')\|_{\text{op}} \\ &\leq 2 \sup_{(m, \sigma) \in B_r(0, 0)} \|\nabla f(m, \sigma) - \nabla f(0, 0)\|_{\text{op}} \leq \frac{1}{2}. \end{aligned} \quad (\text{F.5})$$

Note also that $\nabla \phi(0, 0)$ is the identity. Thus by Lemma F.1, we have that ϕ is a bijection from $B_r(0, 0)$ to $B_{r/2}(\phi(0, 0)) = B_{r/2}(0, 0)$. Now, fix any $\sigma \in \mathbb{R}^{d \times d}$ such that $\|\sigma\| \leq r/2$. Then $(0, \sigma) \in B_{r/2}(0, 0)$, and hence there exists a unique $(m, \sigma') \in B_r(0, 0)$ such that $(0, \sigma) = \phi(m, \sigma') = (f(m, \sigma'), \sigma')$. Thus $\sigma = \sigma'$ and $f(m, \sigma) = 0$. In other words, for each σ such that $\|\sigma\| \leq r/2$ there exists a unique $m = m(\sigma)$ such that $(m(\sigma), \sigma) \in B_r(0, 0)$ and such that $0 = f(m, \sigma)$.

The map $\sigma \mapsto m(\sigma)$ is C^2 by standard Implicit Function Theorem arguments. To show that the first inequality of (5.7) holds, note that we have

$$\|\nabla_m f(m(\sigma), \sigma) - \nabla_m f(0, 0)\|_{\text{op}} \leq \|\nabla f(m(\sigma), \sigma) - \nabla f(0, 0)\|_{\text{op}} \leq 1/4 \leq 1/2$$

by (5.6) since we know that $(m(\sigma), \sigma) \in B_r(0, 0)$. Thus,

$$\begin{aligned} I_d &= \nabla^2 W(0) = \nabla_m f(0, 0) \\ &\implies \frac{1}{2} I_d \preceq \nabla_m f(m(\sigma), \sigma) \preceq \frac{3}{2} I_d. \end{aligned} \quad (\text{F.6})$$

To show the second inequality of (5.7), we first need the supplementary bound

$$\|\nabla_\sigma f(m(\sigma), \sigma)\|_{\text{op}} = \|\nabla_\sigma f(m(\sigma), \sigma) - \nabla_\sigma f(0, 0)\|_{\text{op}} \leq 1/2 \quad (\text{F.7})$$

which holds by the same reasoning as above. Now,

$$\partial_{\sigma_{jk}} m = -\nabla_m f(m, \sigma)^{-1} \partial_{\sigma_{jk}} f(m, \sigma) \in \mathbb{R}^d$$

by standard Implicit Function Theorem arguments, where $\nabla_m f(m, \sigma)$ is a matrix, $\partial_{\sigma_{jk}} f(m, \sigma)$ is a vector, and $\nabla_\sigma m$, $\nabla_\sigma f$ are linear maps from $\mathbb{R}^{d \times d}$ to \mathbb{R}^d . Hence by the first inequality in (5.7) combined with (F.7) we have

$$\begin{aligned} \|\nabla_\sigma m(\sigma)\|_{\text{op}} &= \sup_{\|A\|=1} \|\langle \nabla_\sigma m(\sigma), A \rangle\| \\ &= \sup_{\|A\|=1} \|\nabla_m f(m, \sigma)^{-1} \sum_{j,k=1}^d \partial_{\sigma_{jk}} f(m, \sigma) A_{jk}\| \\ &= \sup_{\|A\|=1} \|\nabla_m f(m, \sigma)^{-1} \langle \nabla_\sigma f, A \rangle\| \\ &\leq \|\nabla_m f(m, \sigma)^{-1}\| \|\nabla_\sigma f\|_{\text{op}} \leq 2 \times \frac{1}{2} = 1. \end{aligned} \quad (\text{F.8})$$

□

Proof of Lemma 5.3. Note that f is C^2 thanks to the fact that W is C^3 and ∇W grows polynomially by Assumption 2 (since the third derivative tensor grows polynomially). We then immediately have $f(0, 0) = \nabla W(0) = 0$, $\nabla_m f(m, \sigma) = \mathbb{E}[\nabla^2 W(m + \sigma Z)]$ is symmetric for all m, σ , and $\nabla_m f(0, 0) = \nabla^2 W(0) = I_d$. To show $\nabla_\sigma f(0, 0) = 0$, we compute the i, j, k term of this tensor:

$$\partial_{\sigma_{jk}} f_i = \partial_{\sigma_{jk}} \mathbb{E}[\partial_i W(m + \sigma Z)] = \mathbb{E}[\partial_{ij}^2 W(m + \sigma Z) Z_k],$$

so that $\partial_{\sigma_{jk}} f_i(0, 0) = \mathbb{E}[\partial_{ij}^2 W(0) Z_k] = 0$. It remains to show that for $r = 2\sqrt{2}$ we have

$$\sup_{(m, \sigma) \in B_r(0, 0)} \|\nabla f(m, \sigma) - \nabla f(0, 0)\|_{\text{op}} \leq \frac{1}{4}. \quad (\text{F.9})$$

First, note that

$$\sup_{(m, \sigma) \in B_r(0, 0)} \|\nabla f(m, \sigma) - \nabla f(0, 0)\|_{\text{op}} \leq r \sup_{(m, \sigma) \in B_r(0, 0)} \|\nabla^2 f(m, \sigma)\|_{\text{op}}, \quad (\text{F.10})$$

where $\nabla^2 f(m, \sigma)$ is a bilinear form on $(\mathbb{R}^d \times \mathbb{R}^{d \times d})^2$, and we have

$$\|\nabla^2 f(m, \sigma)\|_{\text{op}} \leq \|\nabla_\sigma^2 f(m, \sigma)\|_{\text{op}} + 2\|\nabla_\sigma \nabla_m f(m, \sigma)\|_{\text{op}} + \|\nabla_m^2 f(m, \sigma)\|_{\text{op}}. \quad (\text{F.11})$$

For $f(m, \sigma) = \mathbb{E}[\nabla W(\sigma Z + m)]$, these second order derivatives are given by

$$\begin{aligned} \partial_{m_i, m_j}^2 f(m, \sigma) &= \mathbb{E}[\partial_{i,j}^2 \nabla W(m + \sigma Z)], \\ \partial_{m_i} \partial_{\sigma_{jk}} f(m, \sigma) &= \mathbb{E}[\partial_{i,j}^2 \nabla W(m + \sigma Z) Z_k], \\ \partial_{\sigma_{jk}, \sigma_{\ell p}}^2 f(m, \sigma) &= \mathbb{E}[\partial_{j,\ell}^2 \nabla W(m + \sigma Z) Z_k Z_p], \end{aligned} \quad (\text{F.12})$$

each a vector in \mathbb{R}^d . From the first line, we get that

$$\|\nabla_m^2 f(m, \sigma)\|_{\text{op}} \leq \mathbb{E} \|\nabla^3 W(m + \sigma Z)\| = \|\|\nabla^3 W\|\|_{L^1(\mathcal{N}(m, \sigma \sigma^T))} \quad (\text{F.13})$$

where $\|\nabla^3 W\|$ is the standard tensor norm. From the second line, we get

$$\begin{aligned} \|\nabla_m \nabla_\sigma f(m, \sigma)\|_{\text{op}} &= \sup_{\|A\|=1, \|x\|=1} \left\| \mathbb{E} \left[\sum_{i,j,k=1}^d \partial_{i,j}^2 \nabla W(m + \sigma Z) Z_k x_i A_{jk} \right] \right\| \\ &= \sup_{\|A\|=1, \|x\|=1} \left\| \mathbb{E} \left[\sum_{i,j=1}^d \partial_{i,j}^2 \nabla W(m + \sigma Z) x_i (AZ)_j \right] \right\| \\ &= \sup_{\|A\|=1, \|x\|=1} \left\| \mathbb{E} \left[\langle \nabla^3 W(m + \sigma Z), x \otimes AZ \rangle \right] \right\| \\ &\leq \sup_{\|A\|=1, \|x\|=1} \mathbb{E} \left[\|x\| \|AZ\| \|\nabla^3 W(m + \sigma Z)\| \right] \\ &\leq \sqrt{d} \|\|\nabla^3 W\|\|_{L^2(\mathcal{N}(m, \sigma \sigma^T))} \end{aligned} \quad (\text{F.14})$$

A similar computation gives

$$\begin{aligned} \|\nabla_\sigma^2 f(m, \sigma)\|_{\text{op}} &\leq \sup_{\|A\|=1, \|B\|=1} \mathbb{E} [\|AZ\| \|BZ\| \|\nabla^3 W(m + \sigma Z)\|] \\ &\leq 2d \|\|\nabla^3 W\|\|_{L^2(\mathcal{N}(m, \sigma \sigma^T))} \end{aligned} \quad (\text{F.15})$$

Using (F.13), (F.14), and (F.15) in (F.11) gives

$$\begin{aligned}\|\nabla^2 f(m, \sigma)\|_{\text{op}} &\leq (2d + 2\sqrt{d} + 1) \|\|\nabla^3 W\|\|_{L^2(\mathcal{N}(m, \sigma\sigma^T))} \\ &\leq 5d \sup_{(m, \sigma) \in B_r(0, 0)} \|\|\nabla^3 W\|\|_{L^2(\mathcal{N}(m, \sigma\sigma^T))}.\end{aligned}\quad (\text{F.16})$$

Substituting (F.16) into (F.10) with $r = 2\sqrt{2}$ and then using (C.4) of Lemma C.1, we get

$$\begin{aligned}\sup_{(m, \sigma) \in B_{2\sqrt{2}}(0, 0)} \|\nabla f(m, \sigma) - \nabla f(0, 0)\|_{\text{op}} \\ &\leq 10\sqrt{2}d \sup_{(m, \sigma) \in B_{2\sqrt{2}}(0, 0)} \|\|\nabla^3 W\|\|_{L^2(\mathcal{N}(m, \sigma\sigma^T))} \\ &\leq C(q)a_3 \frac{d}{\sqrt{n}} \leq \frac{1}{4}.\end{aligned}\quad (\text{F.17})$$

The last line holds provided $a_3d/\sqrt{n} \leq 1/4C(q)$, which is ensured by (2.4). \square

Proof of Lemma 5.4. First, let $G(\sigma) = \mathbb{E}[\nabla^2 W(m(\sigma) + \sigma Z)]$ and $f(m, \sigma) = \mathbb{E}[\nabla W(m + \sigma Z)]$ as in Lemma 5.3. Note that $\nabla_m f(m, \sigma) = \mathbb{E}[\nabla^2 W(\sigma Z + m)]$, so that $G(\sigma) = \nabla_m f(m, \sigma)|_{m=m(\sigma)}$ and hence by (5.7) of Lemma 5.2 we have

$$\frac{1}{2}I_d \preceq G(\sigma) \preceq \frac{3}{2}I_d, \quad \forall \sigma \in S_{0, r/2}. \quad (\text{F.18})$$

But then $G(\sigma)$ has a unique invertible symmetric positive definite square root, and we define $F(\sigma) = G(\sigma)^{-1/2}$ to be the inverse of this square root. Moreover, using (F.18), it follows that

$$c_1 I_d \preceq F(\sigma) \preceq c_2 I_d, \quad \forall \sigma \in S_{0, r/2},$$

where $c_1 = \sqrt{2/3}$ and $c_2 = \sqrt{2} = r/2$. In other words, $F(S_{0, r/2}) \subseteq S_{c_1, c_2} \subseteq S_{0, r/2}$. It remains to show F is a contraction on $S_{0, r/2}$. Let $\sigma_1, \sigma_2 \in S_{0, \sqrt{2}}$. We will first bound $\|G(\sigma_1) - G(\sigma_2)\|$. We have

$$\|G(\sigma_1) - G(\sigma_2)\| \leq \|\sigma_1 - \sigma_2\| \sup_{\sigma \in S_{0, \sqrt{2}}} \|\nabla_\sigma G(\sigma)\|_{\text{op}}, \quad (\text{F.19})$$

and

$$\begin{aligned}\|\nabla_\sigma G(\sigma)\|_{\text{op}} &= \sup_{\|A\|=1} \|\langle \nabla_\sigma G(\sigma), A \rangle\| \\ &= \sup_{\|A\|=1} \left\| \mathbb{E} \left\langle \nabla^3 W, \langle A, \nabla_\sigma(m(\sigma) + \sigma Z) \rangle \right\rangle \right\|.\end{aligned}\quad (\text{F.20})$$

Here, the quantities inside of the $\|\cdot\|$ on the right are matrices. Indeed, $\langle \nabla_\sigma G, A \rangle$ denotes the application of $\nabla_\sigma G$ to A . Since G sends matrices to matrices, $\nabla_\sigma G$ is a linear functional which also sends matrices to matrices. In the third line, $\nabla_\sigma(m(\sigma) + \sigma Z)$ should be interpreted as a linear functional from $\mathbb{R}^{d \times d}$ to \mathbb{R}^d , so $\langle A, \nabla_\sigma(m(\sigma) + \sigma Z) \rangle$ is a vector in \mathbb{R}^d , and the inner product of this vector with the $d \times d \times d$ tensor $\nabla^3 W$ is a matrix. Using that $\|\langle T, x \rangle\| \leq \|T\| \|x\|$, as explained at the beginning of this section, we have

$$\begin{aligned}\left\| \left\langle \nabla^3 W, \langle A, \nabla_\sigma(m(\sigma) + \sigma Z) \rangle \right\rangle \right\| &\leq \|\nabla^3 W\| \|\langle A, \nabla_\sigma(m(\sigma) + \sigma Z) \rangle\| \\ &\leq \|\nabla^3 W\| \|\nabla_\sigma(m(\sigma) + \sigma Z)\|_{\text{op}} \\ &= \|\nabla^3 W\| \|\nabla_\sigma m(\sigma) + Z \otimes I_d\|_{\text{op}} \leq \|\nabla^3 W\| (1 + \|Z\|).\end{aligned}\quad (\text{F.21})$$

To get the last bound, we used that $\|\nabla_\sigma m(\sigma)\|_{\text{op}} \leq 1$, shown in Lemma 5.3. We also use the fact that $\|Z \otimes I_d\|_{\text{op}} = \sup_{\|A\|=1} \|\langle A, Z \otimes I_d \rangle\| = \sup_{\|A\|=1} \|AZ\| = \|Z\|$. (Recall that since $Z \otimes I_d$ is part of $\nabla_\sigma m$, we are considering $Z \otimes I_d$ as an operator on matrices rather than as a $d \times d \times d$ tensor, and this is why we take the supremum over matrices A .)

Substituting (F.21) back into (F.19), we get

$$\begin{aligned} \|G(\sigma_1) - G(\sigma_2)\| &\leq \|\sigma_1 - \sigma_2\| \sup_{\sigma \in S_0, \sqrt{2}} \mathbb{E} [\|\nabla^3 W(m(\sigma) + \sigma Z)\| (1 + \|Z\|)] \\ &\leq \|\sigma_1 - \sigma_2\| C(q) a_3 \sqrt{\frac{d}{n}} \leq \frac{1}{2\sqrt{2}} \|\sigma_1 - \sigma_2\| \end{aligned} \quad (\text{F.22})$$

The second inequality is by Cauchy-Schwarz and (C.4) of Lemma C.1. The third inequality holds if $a_3 \sqrt{d/n} \leq 1/(2\sqrt{2}C(q))$, which is ensured by (2.4) of Assumption 2. Now, note that thanks to Lemma 5.3, both $\lambda_{\min}(G(\sigma_1))$ and $\lambda_{\min}(G(\sigma_2))$ are bounded below by 1/2. Using Lemma F.2 and (F.22), we therefore have

$$\begin{aligned} \|F(\sigma_1) - F(\sigma_2)\| &\leq \sqrt{2} \|G(\sigma_1) - G(\sigma_2)\| \\ &\leq \frac{1}{2} \|\sigma_1 - \sigma_2\|. \end{aligned} \quad (\text{F.23})$$

Hence F is a strict contraction. \square

Lemma F.2. *Let A_0 and A_1 be psd, and $A_0^{1/2}, A_1^{1/2}$ their unique psd square roots. Assume without loss of generality that $\lambda_{\min}(A_0) \leq \lambda_{\min}(A_1)$. Then*

$$\|A_1^{-1/2} - A_0^{-1/2}\| \leq \frac{\|A_1 - A_0\|}{2\lambda_{\min}(A_0)^{3/2}}.$$

Proof. First note that

$$A_1^{-1/2} - A_0^{-1/2} = A_1^{-1/2} (A_0^{1/2} - A_1^{1/2}) A_0^{-1/2}$$

and hence

$$\|A_1^{-1/2} - A_0^{-1/2}\| \leq \|A_1^{-1/2}\| \|A_1^{1/2} - A_0^{1/2}\| \leq \frac{\|A_1^{1/2} - A_0^{1/2}\|}{\lambda_{\min}(A_0)}.$$

Now, define $A_t = A_0 + t(A_1 - A_0)$ and let $B_t = A_t^{1/2}$, where B_t is the unique psd square root of A_t . We then have $\|A_1^{1/2} - A_0^{1/2}\| \leq \sup_{t \in [0,1]} \|\dot{B}_t\|$. We will now express \dot{B}_t in terms of \dot{A}_t and B_t . Differentiating $B_t^2 = A_t$, we get

$$B_t \dot{B}_t + \dot{B}_t B_t = \dot{A}_t = A_1 - A_0. \quad (\text{F.24})$$

Now, one can check that the solution \dot{B}_t to this equation is given by

$$\dot{B}_t = \int_0^\infty e^{-sB_t} (A_1 - A_0) e^{-sB_t} ds$$

and hence

$$\|\dot{B}_t\| \leq \|A_1 - A_0\| \int_0^\infty \|e^{-sB_t}\|^2 dt = \frac{\|A_1 - A_0\|}{2\lambda_{\min}(B_t)} = \frac{\|A_1 - A_0\|}{2\sqrt{\lambda_{\min}(A_t)}}.$$

Now note that $\lambda_{\min}(A_t) \geq \lambda_{\min}(A_0)$, since A_t is just a convex combination of A_0 and A_1 . Hence $\|\dot{B}_t\| \leq \|A_1 - A_0\|/2\sqrt{\lambda_{\min}(A_0)}$ for all $t \in [0, 1]$. Combining all of the above estimates gives

$$\|A_1^{-1/2} - A_0^{-1/2}\| \leq \frac{\|A_1 - A_0\|}{2\lambda_{\min}(A_0)^{3/2}}.$$

\square

G Auxiliary proofs

We first prove the two change-of-variables results from the paper and then two other postponed lemmas, the proofs of which involve some technical calculations which are not particularly inspiring.

First, we show the equivalence of (4.5) and (2.12). It suffices to show $\int f(-p_3)d\gamma = \int gQd\hat{\pi}$. Recall that $\hat{\pi} = T_{\#}\gamma$ and $g = f \circ T$. Therefore, $\int f(-p_3)d\gamma = -\int gp_3 \circ Td\hat{\pi}$, so it suffices to show $p_3 \circ T = -Q$. Now,

$$\begin{aligned} p_3(x) &= \left\langle \mathbb{E} [\nabla^3 V_0(Z)], \frac{1}{6}x^{\otimes 3} - \frac{1}{2}x \otimes I_d \right\rangle \\ &= \left\langle \mathbb{E} [\nabla^3 V(\hat{m} + \hat{S}^{1/2}Z)], \frac{1}{6}(\hat{S}^{1/2}x)^{\otimes 3} - \frac{1}{2}(\hat{S}^{1/2}x) \otimes \hat{S} \right\rangle \\ &= \left\langle \mathbb{E}_{X \sim \hat{\pi}} [\nabla^3 V(X)], \frac{1}{6}(\hat{S}^{1/2}x)^{\otimes 3} - \frac{1}{2}(\hat{S}^{1/2}x) \otimes \hat{S} \right\rangle. \end{aligned} \quad (\text{G.1})$$

Therefore,

$$\begin{aligned} (p_3 \circ T)(x) &= p_3(\hat{S}^{-1/2}(x - \hat{m})) \\ &= \left\langle \mathbb{E}_{X \sim \hat{\pi}} [\nabla^3 V(X)], \frac{1}{6}(x - \hat{m})^{\otimes 3} - \frac{1}{2}(x - \hat{m}) \otimes \hat{S} \right\rangle = -Q(x), \end{aligned}$$

as desired.

Next, we assume the conclusion of Lemma 5.1 to prove Lemma 2.1.

Proof of Lemma 2.1. Let m, σ be the solution from Lemma 5.1. Then

$$\begin{aligned} 0 &= \mathbb{E} [\nabla W(m + \sigma Z)] = H_V^{-1/2} \mathbb{E} [\nabla V(m^* + H_V^{-1/2}(m + \sigma Z))], \\ (\sigma\sigma^T)^{-1} &= \mathbb{E} [\nabla^2 W(m + \sigma Z)] = H_V^{-1/2} \mathbb{E} [\nabla V(m^* + H_V^{-1/2}(m + \sigma Z))] H_V^{-1/2} \end{aligned}$$

Let

$$\hat{m} = m^* + H_V^{-1/2}m, \quad \hat{S} = H_V^{-1/2}\sigma\sigma^T H_V^{-1/2}.$$

We conclude that

$$\begin{aligned} 0 &= \mathbb{E} [\nabla V(\hat{m} + H_V^{-1/2}\sigma Z)], \\ \hat{S}^{-1} &= \mathbb{E} [\nabla^2 V(\hat{m} + H_V^{-1/2}\sigma Z)] \end{aligned} \quad (\text{G.2})$$

Now, we can replace $H_V^{-1/2}\sigma$ (an asymmetric square root of \hat{S}) in the above equations with $\hat{S}^{1/2}$ (the symmetric square root of \hat{S}) and the equations will still hold. Thus \hat{m}, \hat{S} satisfies (1.9). Furthermore, since $(m, \sigma) \in B_{2\sqrt{2}}(0, 0)$ and $\sigma \in S_{\sqrt{2/3}, \sqrt{2}}$, we get that

$$\|H_V^{1/2}\hat{S}H_V^{1/2}\| + \|H_V^{1/2}(\hat{m} - m^*)\|^2 = \|\sigma\sigma^T\| + \|m\|^2 = \|\sigma\|^2 + \|m\|^2 \leq 8, \quad (\text{G.3})$$

and

$$\frac{2}{3}H^{-1} \preceq \hat{S} = H_V^{-1/2}\sigma\sigma^T H_V^{-1/2} \preceq 2H^{-1}.$$

It remains to show that \hat{m}, \hat{S} is the unique solution satisfying (1.9) in the region \mathcal{R}_V . To show this is the case, note that if (m, S) is a solution to (1.9) in \mathcal{R}_V , then $(H^{1/2}(m - m^*), (H^{1/2}SH^{1/2})^{1/2}) \in B_{2\sqrt{2}}(0, 0) \cap \mathbb{R}^d \times S_{0, \sqrt{2}}$ is a solution to (5.3), (5.4). Furthermore, the map

$$(m, S) \mapsto (H^{1/2}(m - m^*), (H^{1/2}SH^{1/2})^{1/2})$$

is injective. Since we know solutions to (5.3), (5.4) in $B_{2\sqrt{2}}(0, 0) \cap \mathbb{R}^d \times S_{0, \sqrt{2}}$ are unique by Lemma 5.1, it follows that solutions to (1.9) in \mathcal{R}_V must also be unique. \square

We now turn to the technical calculations.

Proof of Lemma 2.3. If $g(x) = f(\hat{S}^{-1/2}(x - \hat{m}))$ then (2.7) is satisfied provided $|f(y) - \gamma(f)| \leq e^{\frac{c_0}{4}\sqrt{d}\|y\|}$ for all $\|y\| \geq R_g\sqrt{d}$. We now find R_g to ensure this inequality on f is satisfied. Suppose WLOG that $f(0) = 0$, and $|f(y)| \leq \exp(C\sqrt{d}\|y\|^\alpha)$ for all $y \in \mathbb{R}^d$. We first bound $\gamma(|f|)$. We have

$$\gamma(|f|) \leq \max_{\|y\| \leq (4C_f + 2\sqrt{2})\sqrt{d}} |f(y)| + (2\pi)^{-d/2} \int_{\|y\| \geq (4C_f + 2\sqrt{2})\sqrt{d}} \exp(C_f\sqrt{d}\|y\| - \|y\|^2/2) dy. \quad (\text{G.4})$$

Now, we bound the two summands separately. For the first summand, we have

$$\begin{aligned} K_1 &:= \max_{\|y\| \leq (4C_f + 2\sqrt{2})\sqrt{d}} |f(y)| \\ &\leq \exp(C_f(4C_f + 3)^\alpha(\sqrt{d})^{1+\alpha}) \\ &\leq \exp(C_f(4C_f + 3)^\alpha d) =: e^{\bar{C}d}, \end{aligned} \quad (\text{G.5})$$

where

$$\bar{C} := C_f(4C_f + 3)^\alpha.$$

For the second summand in (G.4), we have

$$\begin{aligned} K_2 &:= (2\pi)^{-d/2} \int_{\|y\| \geq (4C_f + 2\sqrt{2})\sqrt{d}} e^{-\|y\|^2/4} dy \\ &\leq 2^{\frac{d}{2}} (2\pi)^{-d/2} \int_{\|y\| \geq 2\sqrt{d}} e^{-\|y\|^2/2} dy \\ &\leq 2^{\frac{d}{2}} e^{-\frac{d}{2}} \leq 1. \end{aligned} \quad (\text{G.6})$$

In the second line, we used that $C_f\sqrt{d}\|y\| - \|y\|^2/2 \leq \|y\|^2/4 - \|y\|^2/2 = -\|y\|^2/4$ when $\|y\| \geq (4C_f + 2\sqrt{2})\sqrt{d}$. Using (G.5) and (G.6) in (G.4), we get

$$\gamma(|f|) \leq K_1 + K_2 \leq 2e^{\bar{C}d}. \quad (\text{G.7})$$

Now, if $R_g \geq \frac{4}{c_0}(\bar{C} + \log 4)$, then

$$|\gamma(f)| \leq 2e^{\bar{C}d} \leq \frac{1}{2}e^{\frac{c_0}{4}R_g d} \leq \frac{1}{2}e^{\frac{c_0}{4}\sqrt{d}\|x\|} \quad \forall \|x\| \geq R_g\sqrt{d}.$$

Next, we bound $|f(y)|$. Note that if $\|y\| \geq R_g\sqrt{d}$ and $R_g \geq 1$, then

$$\begin{aligned} C_f\sqrt{d}\|y\|^\alpha + \log 2 &= \left(\frac{C_f\sqrt{d}}{\|y\|^{1-\alpha}} + \frac{\log 2}{\|y\|} \right) \|y\| \\ &\leq \left(\frac{C_f\sqrt{d}}{R_g^{1-\alpha}\sqrt{d}^{1-\alpha}} + \frac{\log 2}{R_g\sqrt{d}} \right) \|y\| \\ &\leq \frac{(C_f + \log 2)\sqrt{d}}{R_g^{1-\alpha}} \|y\| \leq \frac{c_0}{4}\sqrt{d}\|y\| \end{aligned} \quad (\text{G.8})$$

as long as $R_g \geq \left[\frac{4}{c_0} (C_f + \log 2) \right]^{\frac{1}{1-\alpha}}$. We then have

$$|f(y)| \leq \exp(C_f \sqrt{d} \|y\|^\alpha) \leq \frac{1}{2} \exp\left(\frac{c_0}{4} \sqrt{d} \|y\|\right) \quad \forall \|y\| \geq R_g \sqrt{d}.$$

Combining all of the above estimates, we get that

$$|f(y) - \gamma(f)| \leq |f(y)| + \gamma(|f|) \leq \exp\left(\frac{c_0}{4} \sqrt{d} \|y\|\right) \quad \forall \|y\| \geq R_g \sqrt{d}$$

as long as

$$R_g \geq \max\left(\left[1 \vee \frac{4}{c_0} (C_f + \log 2)\right]^{\frac{1}{1-\alpha}}, \frac{4}{c_0} (C_f (4C_f + 3)^\alpha + \log 4)\right).$$

This is a function of C_f, α, c_0 which is increasing with C_f, α , and c_0^{-1} , as desired. \square

Proof of Lemma D.2. Lemma E.3 of [18] shows that if $Rb > 1 + p/d$, then

$$I(R, p, b) \leq (eR)^p d^{\frac{p}{2}+1} \exp\left(\left[\frac{3}{2} + \log R - Rb\right] d\right) \quad (\text{G.9})$$

Massaging this upper bound, we get

$$I(R, p, b) \lesssim_p R^p d^{\frac{p}{2}+1} \exp\left(\left[\frac{3}{2} + \log R - Rb\right] d\right) \leq d^{\frac{p}{2}+1} \exp\left(\left[\frac{3}{2} + (p+1) \log R - Rb\right] d\right).$$

Now, we have

$$(p+1) \log R = (2p+2) \log \sqrt{R} \leq (2p+2) \sqrt{R} \leq \frac{Rb}{4}$$

provided $R \geq (8p+8)^2 b^{-2}$. Therefore,

$$\frac{3}{2} + (p+1) \log R - Rb \leq \frac{3}{2} - \frac{3}{4} Rb \leq -1 - \frac{1}{2} Rb$$

for $R \geq 10b^{-1}$. Thus $I(R, p, b) \lesssim_p d^{\frac{p}{2}+1} e^{-d-Rbd/2} \lesssim_p e^{-Rbd/2}$. The lower bounds on R are met provided

$$R \geq \max((1+p)b^{-1}, (8p+8)^2 b^{-2}, 10b^{-1}),$$

which can be achieved by taking $R = C_p(1 \vee b^{-1})^2$ for some C_p depending only on p . \square

References

- [1] Radoslaw Adamczak, Alexander Litvak, Alain Pajor, and Nicole Tomczak-Jaegermann. Quantitative estimates of the convergence of the empirical covariance matrix in log-concave ensembles. *Journal of the American Mathematical Society*, 23(2):535–561, 2010.
- [2] Pierre Alquier and James Ridgway. Concentration of tempered posteriors and of their variational approximations. *The Annals of Statistics*, 48(3):1475–1497, 2020.
- [3] Dominique Bakry, Ivan Gentil, Michel Ledoux, et al. *Analysis and geometry of Markov diffusion operators*, volume 103. Springer, 2014.

- [4] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- [5] Edward Challis and David Barber. Gaussian kullback-leibler approximate inference. *Journal of Machine Learning Research*, 14(68):2239–2286, 2013.
- [6] Mary Kathryn Cowles and Bradley P. Carlin. Markov chain Monte Carlo convergence diagnostics: A comparative review. *Journal of the American Statistical Association*, 91(434):883–904, 1996.
- [7] Kamélia Daudel and Randal Douc. Mixture weights optimisation for alpha-divergence variational inference. In *Advances in Neural Information Processing Systems*, volume 34, pages 4397–4408, 2021.
- [8] Kamélia Daudel, Randal Douc, and Francois Portier. Infinite-dimensional gradient-based descent for alpha-divergence minimisation. *The Annals of Statistics*, 49(4):2250–2270, 2021.
- [9] Guillaume Dehaene and Simon Barthelmé. Expectation propagation in the large data limit. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 80(1):199–217, 2018.
- [10] Guillaume P Dehaene. A deterministic and computable Bernstein-von Mises theorem. *arXiv preprint arXiv:1904.02505*, 2019.
- [11] Guillaume P Dehaene and Simon Barthelmé. Bounding errors of expectation-propagation. *Advances in Neural Information Processing Systems*, 28, 2015.
- [12] Michael Diao, Krishnakumar Balasubramanian, Sinho Chewi, and Adil Salim. Forward-backward gaussian variational inference via jko in the bures-wasserstein space. *arXiv preprint arXiv:2304.05398*, 2023.
- [13] Subhashis Ghosal. Asymptotic normality of posterior distributions for exponential families when the number of parameters tends to infinity. *Journal of Multivariate Analysis*, 74(1):49–68, 2000.
- [14] Wei Han and Yun Yang. Statistical inference in mean-field variational Bayes. *arXiv preprint arXiv:1911.01525*, 2019.
- [15] Tapani Helin and Remo Kretschmann. Non-asymptotic error estimates for the Laplace approximation in Bayesian inverse problems. *Numerische Mathematik*, 150(2):521–549, 2022.
- [16] Mikolaj J Kasprzak, Ryan Giordano, and Tamara Broderick. How good is your Gaussian approximation of the posterior? Finite-sample computable error bounds for a variety of useful divergences. *arXiv preprint arXiv:2209.14992*, 2022.
- [17] Anya Katsevich. Improved dimension dependence in the Bernstein von Mises Theorem via a new Laplace approximation bound. *arXiv preprint arXiv:2308.06899*, 2023.
- [18] Anya Katsevich. Tight dimension dependence of the Laplace approximation. *arXiv preprint arXiv:2305.17604*, 2023.
- [19] Anya Katsevich. Tight skew adjustment to the Laplace approximation in high dimensions. *arXiv preprint arXiv:2306.07262*, 2023.

- [20] Marc Lambert, Sinho Chewi, Francis Bach, Silvère Bonnabel, and Philippe Rigollet. Variational inference via Wasserstein gradient flows. *arXiv preprint arXiv:2205.15902*, 2022.
- [21] Serge Lang. *Real and Functional Analysis*. Graduate Texts in Mathematics. Springer New York, NY, 3 edition, 1993.
- [22] N Lebedev. *Special Functions and Their Applications*,. Dover Publications, 1972.
- [23] Yulong Lu. On the Bernstein-von Mises theorem for high dimensional nonlinear Bayesian inverse problems. *arXiv preprint arXiv:1706.00289*, 2017.
- [24] Vladimir Spokoiny. Bernstein-von Mises theorem for growing parameter dimension. *arXiv preprint arXiv:1302.3430*, 2013.
- [25] Vladimir Spokoiny. Dimension free nonasymptotic bounds on the accuracy of high-dimensional laplace approximation. *SIAM/ASA Journal on Uncertainty Quantification*, 11(3):1044–1068, 2023.
- [26] Pragya Sur. *A modern maximum-likelihood theory for high-dimensional logistic regression*. PhD thesis, Stanford University, 2019.
- [27] Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.
- [28] Yixin Wang and David M Blei. Frequentist consistency of variational Bayes. *Journal of the American Statistical Association*, 114(527):1147–1161, 2019.
- [29] Ruby C Weng. A Bayesian Edgeworth expansion by stein’s identity. *Bayesian Analysis*, 5(4):741–764, 2010.
- [30] Fengshuo Zhang and Chao Gao. Convergence rates of variational posterior distributions. *The Annals of Statistics*, 48(4):2180–2207, 2020.
- [31] Xinzhen Zhang, Chen Ling, and Liqun Qi. The best rank-1 approximation of a symmetric tensor and related spherical optimization problems. *SIAM Journal on Matrix Analysis and Applications*, 33(3):806–821, 2012.