
A Central Limit Theorem for Differentially Private Query Answering

Jinshuo Dong

Department of Computer Science
Northwestern University and IDEAL*
jinshuo@northwestern.edu

Weijie J. Su

Department of Statistics and Data Science
University of Pennsylvania
suw@wharton.upenn.edu

Linjun Zhang

Department of Statistics
Rutgers University
linjun.zhang@rutgers.edu

Abstract

Perhaps the single most important use case for differential privacy is to privately answer numerical queries, which is usually achieved by adding noise to the answer vector. The central question is, therefore, to understand which noise distribution optimizes the privacy-accuracy trade-off, especially when the dimension of the answer vector is high. Accordingly, an extensive literature has been dedicated to the question and the upper and lower bounds have been successfully matched up to constant factors [BUV18, SU17]. In this paper, we take a novel approach to address this important optimality question. We first demonstrate an intriguing central limit theorem phenomenon in the high-dimensional regime. More precisely, we prove that a mechanism is approximately *Gaussian Differentially Private* [DRS21] if the added noise satisfies certain conditions. In particular, densities proportional to $e^{-\|x\|_p^\alpha}$, where $\|x\|_p$ is the standard ℓ_p -norm, satisfies the conditions. Taking this perspective, we make use of the Cramer–Rao inequality and show a “uncertainty principle”-style result: the product of privacy parameter and the ℓ_2 -loss of the mechanism is lower bounded by the dimension. Furthermore, the Gaussian mechanism achieves the constant-sharp optimal privacy-accuracy trade-off among all such noises. Our findings are corroborated by numerical experiments.

1 Introduction

Introduced in [DMNS06], to date *differential privacy* (DP) is perhaps the most popular privacy definition. One of the most important applications of differential privacy is to answer numeric queries. Given a function f of interest, which is also termed a query, our goal is to evaluate this (potentially vector-valued) query f on the sensitive data. To preserve privacy, a DP mechanism M working on a dataset D , in its simplest form, is defined as

$$M(D) = f(D) + tX. \tag{1}$$

Above, X denotes the noise term and t is a scalar, which together are selected depending on the properties of the query f and the desired privacy level. Among these, perhaps the most popular examples are the Laplace mechanism and the Gaussian mechanism where the noise X follows the Laplace distribution and the Gaussian distribution, respectively.

*the Institute for Data, Econometrics, Algorithms, and Learning

Aside from privacy considerations, the most important criterion of an algorithm is arguably the estimation accuracy in the face of choosing, for example, between the Laplace mechanism or its Gaussian counterpart for a given problem. To be concrete, consider a real-valued query f with sensitivity 1—that is, $\Delta f = \sup_{D, D'} |f(D) - f(D')| = 1$, where the supremum is over all neighboring datasets D and D' . Assuming $(\epsilon, 0)$ -DP for the mechanism M , we are interested in minimizing its ℓ_2 loss defined as

$$\text{err}(M) := \mathbb{E}(M(D) - f(D))^2 = \mathbb{E}(tX)^2 = t^2 \mathbb{E}X^2.$$

This question is commonly² addressed by setting X to a standard Laplace random variable and $t = \epsilon^{-1}$ [DMNS06]. This gives $\text{err}(M) = 2\epsilon^{-2}$. Moving forward, we *relax* the privacy constraint from $(\epsilon, 0)$ -DP to (ϵ, δ) -DP for some small δ . The canonical way, which was born together with the notion of (ϵ, δ) -DP, is to add Gaussian noise [DKM⁺06]. A well-known result demonstrates that Gaussian mechanism with X being the standard normal and $t = \frac{1}{\epsilon} \sqrt{2 \log(1.25\delta^{-1})}$ is (ϵ, δ) -DP (see, e.g., [DR14]). The ℓ_2 -loss is $\text{err}(M) = t^2 = 2\epsilon^{-2} \cdot \log(1.25\delta^{-1})$.

A quick comparison between the two errors reveals a surprising message. The latter error $2\epsilon^{-2} \cdot \log(1.25\delta^{-1})$ is larger than the former $2\epsilon^{-2}$. In fact, the extra factor $\log(1.25\delta^{-1})$ is already greater than 10 when $\delta = 10^{-5}$. At least on the surface, this observation contradicts the fact that (ϵ, δ) -DP is a relaxation of $(\epsilon, 0)$ -DP. Put differently, moving from Laplace to Gaussian, both privacy and accuracy get worse. Nevertheless, this contradiction suggests that we need a better alternative to the Gaussian mechanism instead of giving up the notion of (ϵ, δ) -DP. Indeed, the truncated Laplace mechanism has been proposed as a better alternative to achieve (ϵ, δ) -DP [GDGK20], which outperforms the Laplace mechanism in terms of estimation accuracy³.

Motivated by these facts concerning the Laplace, Gaussian, and truncated Laplace mechanisms, one cannot help asking:

- (Q1) Why was the truncated Laplace mechanism not considered in the first place? Are there any insights behind the design of such mechanisms?
- (Q2) More importantly, are these insights inherent for answering one-dimensional queries, or can we extend them to high-dimensional setting?

In this paper, we tackle these fundamental questions, beginning with explaining (Q1) in Section 2 from the decision-theoretic perspective of DP [WZ10, KOV17, DRS21]. However, our main focus is (Q2). In addressing this question, we uncover a seemingly surprising phenomenon — it is impossible to utilize the (ϵ, δ) privacy budget in high-dimensional problems the same way as the truncated Laplace mechanism utilizes it in the one-dimensional problem. More specifically, we show a central limit behavior of the noise-addition mechanism in high dimensions, which, roughly speaking, says that for general noise distributions, the corresponding mechanisms *all* behave like a Gaussian mechanism. The formal language of “a mechanism behaves like the Gaussian mechanism” has been set up in [DRS21], where a notion called *Gaussian Differential Privacy* (GDP) was proposed. Roughly speaking, a mechanism is μ -GDP if it offers as much privacy as adding $N(0, \mu^{-2})$ noise to a sensitivity-1 query. As in the (ϵ, δ) -DP case, the smaller μ is, the stronger privacy is offered.

To state our first main contribution, let f be an n -dimensional query and assume that its ℓ_2 -sensitivity is 1. Consider the noise addition mechanism $M(D) = f(D) + tX$ where X has a log-concave density $\propto e^{-\varphi(x)}$ on \mathbb{R}^n . Let $\mathcal{I}_X := \mathbb{E}[\nabla\varphi(X)\nabla\varphi(X)^T]$ be the $n \times n$ Fisher information matrix and $\|\mathcal{I}_X\|_2$ be its operator norm.

Theorem 1.1 (Central Limit Theorem (Informal version of Theorem 3.1)). *Under certain conditions on φ , for $t = \mu^{-1} \cdot \sqrt{\|\mathcal{I}_X\|_2}$, the corresponding noise addition mechanism M defined in Eq.(1) is asymptotically μ -GDP as the dimension $n \rightarrow \infty$ except for an $o(1)$ fraction of directions of $f(D) - f(D')$.*

In particular, the norm power functions $\varphi(x) = \|x\|_p^\alpha$ ($p, \alpha \geq 1$) satisfy these technical conditions. Note that this class already contains correlated noise, so the results in [DRS21] do not apply here. Numerical results in Figure 1 shows that the convergence occurs for a dimension as small as 30.

²If f is integer-valued, then the doubly geometric distribution is a better choice and yields an ℓ_2 -loss of $\frac{1}{2} \sinh^{-2} \frac{\epsilon}{2} < 2\epsilon^{-2}$. In the so-called high privacy regime, i.e. $\epsilon \rightarrow 0$, the two ℓ_2 -losses have the same order in the sense that their ratio goes to 1.

³One may blame the sub-optimality of the choice of t , but the problem remains even if the smallest possible t from [BW18] is applied.




Figure 1: Fast convergence to GDP as claimed in Theorem 1.1. Blue solid curves indicate the true privacy (i.e. ROC functions, see Section 2 for details) of the noise addition mechanism considered in Theorem 1.1. Red dashed curves are GDP limit predicted by our CLT. In all three panels the dimension $n = 30$. Numerical details can be found in the appendix.

We then elaborate on the condition “ $o(1)$ fraction of $f(D) - f(D')$ ”. Following the original definition, DP or GDP is a condition that needs to hold for arbitrary neighboring datasets D and D' . This worst case perspective is exactly what prevents us to observe the central limit behavior. For example, consider a certain pair of datasets with $f(D) = (0, 0, \dots, 0)$ and $f(D') = (1, 0, \dots, 0)$, then privacy is completely determined by the first marginal distribution of X , and the dimension n plays no role here. The “ $o(1)$ fraction of $f(D) - f(D')$ ” rules out the essentially low-dimension cases and reveals the truly high-dimensional behavior.

In summary, Theorem 1.1 suggests that when the dimension is high, a large class of noise addition mechanisms behave like the Gaussian mechanism, and hence are doomed to a poor use of the given (ϵ, δ) privacy budget, in the same fashion as we have seen in the one-dimensional example.

However, admitting the central limit phenomenon, our second theorem turns the table and characterizes the optimal privacy-accuracy trade-off and justifies the Gaussian mechanism. To see this, recall that the noise addition mechanism defined in Equation (1) is determined by the pair (t, X) . Both privacy and accuracy are jointly determined by t and X . Adopting the central limit theorem 1.1, it is convenient to take an equivalent parametrization, which is (μ, X) , where μ is the desired (asymptotic) GDP parameter. Given X , the two parametrizations are related by $t = \mu^{-1} \cdot \sqrt{\|\mathcal{I}_X\|_2}$. Using parameters (μ, X) , the corresponding mechanism $M_{\mu, X}$ is given by

$$M_{\mu, X}(D) = f(D) + \mu^{-1} \cdot \sqrt{\|\mathcal{I}_X\|_2} \cdot X$$

By Theorem 1.1, it is asymptotically μ -GDP. The following theorem states in an “uncertainty principle” fashion that the privacy parameter and the error cannot be small at the same time.

Theorem 1.2. *As long as the Fisher information of X is defined, we have*

$$\mu^2 \cdot \text{err}(M_{\mu, X}) \geq n.$$

The equality holds if X is n -dimensional standard Gaussian.

Combining Theorems 1.1 and 1.2, among all the noise that satisfies the conditions of Theorem 1.1, Gaussian yields the constant-sharp optimal privacy-accuracy trade-off. As far as we know, this is the first result characterizing optimality with the sharp constant when the dimension is high.

The privacy conclusion of Theorem 1.1 does not work for every pair of neighboring datasets, so it is worth noting that we do NOT intend to suggest this as a valid privacy guarantee. Instead, we present it as an interesting phenomenon that has been largely overlooked in the literature. Furthermore, this central limit theorem admits an elegant characterization of privacy-accuracy trade-off that is sharp in constant. From a theoretical point of view, the proof of Theorem 1.1, as we shall see in later sections, involves *non-linear* functionals of high dimensional distributions. This type of results are, to the best of our knowledge, quite underexplored compared to linear functionals, so our results may serve as an additional motivation to study this type of questions.

Related work There is a large body of literature on the characterization of privacy-accuracy trade-off for query answering mechanisms. For the one-dimensional case, the constant-sharp optimal noise for $(\varepsilon, 0)$ -DP was shown to have a piece-wise constant density by [GV16]. This complements our discussion in Figure 2. When the dimension is high, only up-to-constant-factor optimality was known. In particular, [BUV18, SU17] confirm that Gaussian mechanism is minimax rate optimal under (ε, δ) -DP by a novel lower bound technique. In addition, [ENU20] also confirms the minimax optimality of Gaussian mechanism for linear queries with a refined notion of sensitivity. Our work extends this direction by taking the CLT perspective and providing an elegant constant-sharp optimality result. There are also works studying the up-to-constant-factor minimax optimality in other models, such as the (sparse) linear regression [CWZ21], generalized linear models [CWZ20, SSTT21], Gaussian mixtures [KSSU20, ZZ21] and so on. In our work, we initialize the investigation in the (simpler) mean estimation problem, and leave the constant-sharp optimality in other problems for future work.

2 GDP and the ROC Functions

The decision theoretic interpretation of DP was first proposed in [WZ10] and then extended by [KOV17]. More recently, [DRS21] systematically studied this perspective and developed various tools. In this section we take this perspective and introduce the basics of [DRS21]. This will allow us to give an intuitive answer to (Q1).

Suppose each individual’s sensitive information is an element in the abstract set \mathcal{X} . A dataset D of k people is then an element in \mathcal{X}^k . Let a randomized algorithm M take a dataset as input and let D and D' be two neighboring datasets, i.e. they differ by one individual. Differential privacy seeks to limit the power of an adversary identifying the presence of an arbitrary individual in the dataset. That is, with the output as the observation, telling apart D and D' must be hard for the adversary. Decision theoretically, the quality of an attack is measured by the errors it makes. The more error it is forced to make, the more privacy M provides.

To breach the privacy, the adversary performs the following hypothesis testing attack:

$$H_0 : \text{output} \sim M(D) \text{ vs } H_1 : \text{output} \sim M(D').$$

By the random nature of M , $M(D)$ and $M(D')$ are two distributions. We emphasize this point by denoting them by P and Q . The errors mentioned above are simply the probabilities confusing D and D' , which are commonly known as false positive and false negative rates. Because of the symmetry of the neighboring relation, there is no need to worry about which is which.

ROC function. For simplicity assume M outputs a vector in \mathbb{R}^n . A general decision rule for testing H_0 against H_1 has the form $\phi : \mathbb{R}^n \rightarrow \{0, 1\}$. Observing $v \in \mathbb{R}^n$, hypothesis H_i is accepted if $\phi(v) = i$, for $i \in \{0, 1\}$. The false positive rate (type I error) of ϕ , i.e. mistakenly accepting $H_1 : v \sim M(D') = Q$ while actually $v \sim M(D) = P$, is $\alpha_\phi := \mathbb{P}_{v \sim P}(\phi(v) = 1) = \mathbb{E}_P(\phi)$. Similarly, the false negative rate (type II error) of ϕ is $\beta_\phi := 1 - \mathbb{E}_Q(\phi)$. Note that both errors are in $[0, 1]$. Consider the function $f_{P,Q} : [0, 1] \rightarrow [0, 1]$ defined as follows:

$$f_{P,Q}(\alpha) := \inf\{1 - \mathbb{E}_Q(\phi) : \phi \text{ satisfies } \mathbb{E}_P(\phi) \leq \alpha\}. \quad (2)$$

That is, $f_{P,Q}(\alpha)$ equals the minimum false negative rate that one can achieve when false negative is at most α . The graph of $f_{P,Q}$ is exactly the flipped ROC curve of the family of optimal tests (which, by Neyman–Pearson lemma, are the likelihood ratio tests). We call it the *ROC function of the test P vs Q* . The same notion is called *trade-off function of P and Q* in [DRS21] and is denoted by $T[P, Q]$. We avoid this name because in our paper “trade-off” mainly refers to the privacy-accuracy trade-off, but we will keep their notation.

DP and ROC function Plugging in the privacy context where $P = M(D)$, $Q = M(D')$, from the discussion above, we see that $T[M(D), M(D')]$ measures the optimal error distinguishing $M(D)$ and $M(D')$. Therefore, a lower bound on $T[M(D), M(D')]$ implies privacy of M . Indeed, [WZ10, KOV17] showed that M is (ε, δ) -DP if and only if $T[M(D), M(D')] \geq f_{\varepsilon, \delta}$ pointwise in $[0, 1]$ for any neighboring dataset D, D' . The graph of $f_{\varepsilon, \delta}$ is plotted in the left panel of Figure 2. Compared to a single (ε, δ) bound, the ROC function $T[M(D), M(D')]$ provides a more refined picture of the privacy of M . In fact, [DRS21] shows that the ROC function is equivalent to an infinite family of (ε, δ) bounds, which is called privacy profile in [BBG20].

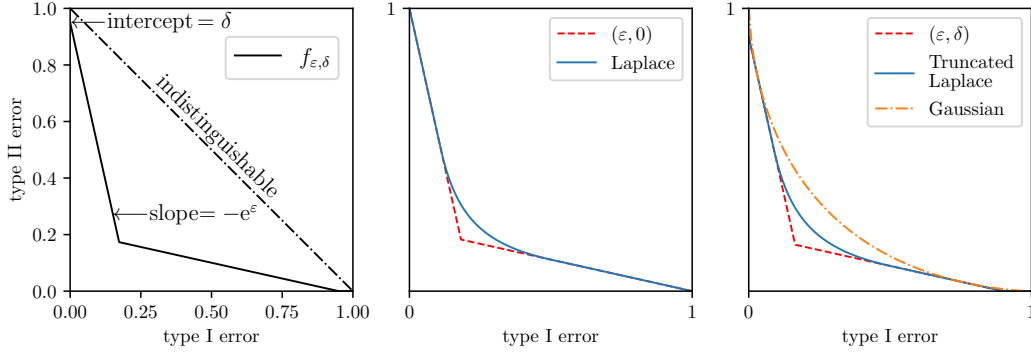


Figure 2: Left: $f_{\epsilon, \delta}$ which recovers the classical (ϵ, δ) -DP definition. Middle: Laplace mechanism is $(\epsilon, 0)$ -DP Right: Gaussian mechanism and truncated Laplace mechanism are both (ϵ, δ) -DP

Truncated Laplace vs Gaussian through the lense of ROC function Now we use ROC function to answer (Q1) in the introduction. Namely, we want to explain the embarrassing situation of the Gaussian mechanism that the privacy budget is not fully used, and the success of the truncated Laplace mechanism.

When M is the Laplace mechanism which is designed to be $(\epsilon, 0)$ -DP, it is not hard to determine $T[M(D), M(D')]$ via Neyman–Pearson lemma and verify that it is indeed lower bounded by $f_{\epsilon, 0}$ (see the middle panel of Figure 2). In fact, $T[M(D), M(D')]$ mostly agrees with $f_{\epsilon, 0}$. In other words, the $(\epsilon, 0)$ privacy budget is almost⁴ fully utilized.

When M is the Gaussian mechanism with (ϵ, δ) -DP guarantee, $T[M(D), M(D')]$ is naturally lower bounded by $f_{\epsilon, \delta}$, however, there is a large gap between the two curves (see the right panel of Figure 2). The (ϵ, δ) privacy budget is poorly utilized by the Gaussian mechanism. This explains why the l^2 -loss of Gaussian mechanism is not satisfactory.

For noise addition mechanism, if the noise is bounded, say a uniform $[-1, 1]$ distribution, then $T[M(D), M(D')] = f_{0, \delta}$ for some $\delta \in (0, 1)$. This suggests us to consider bounded noise if we want to add a δ slack in privacy to the Laplace mechanism. The obvious attempt is then to truncate Laplace noise. Indeed, the corresponding ROC function is as close to $f_{\epsilon, \delta}$ as that of the Laplace mechanism to $f_{\epsilon, 0}$ (also see the right panel of Figure 2). This not only explains the success of the truncated Laplace mechanism, but also points us to the right direction in searching for such a mechanism.

In hindsight, this achievement for one-dimensional mechanisms is due to the following fact: as we change the noise distribution, the corresponding ROC functions are significantly different. Hence we can pick the one that best utilizes our privacy budget. However, in the next section we will argue that this no longer works when the dimension is high — many (if not all) choices of noise distribution yield the same ROC function, which is the ROC of Gaussian mechanism.

ROC function of the Gaussian mechanism For $\mu \geq 0$, let $G_\mu := T[\mathcal{N}(0, 1), \mathcal{N}(\mu, 1)]$ where Φ denotes the cumulative distribution function (CDF) of the standard normal distribution. Consider a query f with sensitivity 1 and let $\text{Lap}(0, 1)$ be the standard Laplace noise. Just like ϵ -DP captures the privacy of the mechanism $M(D) = f(D) + \epsilon^{-1} \cdot \text{Lap}(0, 1)$, the function G_μ captures the privacy of $M(D) = f(D) + \mu^{-1} \cdot N(0, 1)$. In fact, if $f(D') - f(D) = 1$, then $M(D) = N(f(D), \mu^{-2})$ and $M(D') = N(f(D'), \mu^{-2})$. By its hypothesis testing construction, $T[P, Q]$ remains invariant when an invertible transformation is simultaneously applied to P and Q , resulting in

$$T[M(D), M(D')] = T[N(f(D), \mu^{-2}), N(f(D'), \mu^{-2})] = T[N(0, 1), N(\mu, 1)] = G_\mu$$

Therefore, the privacy of a Gaussian mechanism is precisely captured by the ROC function G_μ . A general mechanism M is said to be *Gaussian differentially private* (GDP) if it offers more privacy than a Gaussian mechanism. More specifically,

Definition 2.1 (GDP). *An algorithm M is μ -GDP if $T[M(D), M(D')] \geq G_\mu$ for any pair of neighboring datasets D and D' .*

⁴If the query is integer-valued, then $(\epsilon, 0)$ privacy budget can be saturated by adding doubly geometric noise.

Alternatively, M is μ -GDP if and only if $\inf_{D, D'} T[M(D), M(D')] \geq G_\mu$ where the infimum of ROC functions is interpreted pointwise, and the infimum is taken over all neighboring datasets D and D' . This inequality says M offers more privacy than the corresponding Gaussian mechanism. If the equality holds, i.e.

$$\inf_{D, D'} T[M(D), M(D')] = G_\mu \quad (3)$$

then it means the mechanism M offers exact the same amount of privacy as the corresponding Gaussian mechanism. In fact, the CLT to be presented in the next section has this flavor of conclusion.

3 Central Limit Theorem

In the following two sections we turn to addressing (Q2). This section is dedicated to the rigorous form of the CLT and the discussion.

The experience with the CLT for i.i.d. random variables suggests that the statement for the normalized special case is usually the most comprehensible. Therefore, we will state the normalized version as Theorem 3.1 and derive the general case as Corollary 3.2, which is also the rigorous version of our informal theorem 1.1 mentioned in the introduction.

Consider an n -dimensional query $f : \mathcal{X}^k \rightarrow \mathbb{R}^n$. We assume it has ℓ_2 -sensitivity 1, i.e. $\sup_{D, D'} \|f(D) - f(D')\|_2 = 1$. Suppose $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$ is convex and $e^{-\varphi}$ is integrable on \mathbb{R}^n . The log-concave random vector with density $\propto e^{-\varphi(x)}$ will be denoted by X_φ . Define the function class

$$\mathfrak{F}_n := \{\varphi : \mathbb{R}^n \rightarrow \mathbb{R} \text{ convex} \mid \varphi(x) = \varphi(-x), e^{-\varphi} \in L^1(\mathbb{R}^n), \mathbb{E}\|X_\varphi\|_2^2 < +\infty, \mathbb{E}\|\nabla\varphi(X_\varphi)\|_2^2 < +\infty\}.$$

The regularity conditions guarantee that X_φ has finite second moments and Fisher information matrix defined as $\mathcal{I}_\varphi = \mathbb{E}[\nabla\varphi(X_\varphi)\nabla\varphi(X_\varphi)^T]$. Furthermore, we also have $\mathbb{E}X_\varphi = 0$ by symmetry and $\mathbb{E}\nabla\varphi(X_\varphi) = 0$ by standard theory of Fisher information (e.g. [VdV00]). We will focus on this class of functions for the rest of paper.

The n -dimensional noise addition mechanism of interest takes the form $M(D) = f(D) + tX_\varphi$. The parameter t is only for the convenience of tuning and can be absorbed into φ . In fact, tX_φ has log-concave density $\propto e^{-\varphi(x/t)}$, so it is distributed as $X_{\tilde{\varphi}}$ where $\tilde{\varphi}(x) = \varphi(x/t)$. For the normalized CLT, we set $t = 1$ and assume \mathcal{I}_φ is the $n \times n$ identity matrix $I_{n \times n}$.

Since what we are going to present is an asymptotic result where the dimension $n \rightarrow \infty$, the above objects necessarily appear with an index n , i.e. we have $f_n, \varphi_n, X_{\varphi_n}$ and \mathcal{I}_{φ_n} . The latter two are often denoted by X_n and \mathcal{I}_n for brevity. With normalization, the n -dimensional mechanism of interest is $M_n(D) = f_n(D) + X_n$. For clarity, we choose to state the theorem first, and then present the details of the technical conditions.

Theorem 3.1. *If the function sequence φ_n satisfies conditions (D1) and (D2), then there is a sequence of positive numbers c_n with $c_n \rightarrow 0$ as $n \rightarrow \infty$ and a subset $E_n \subseteq S^{n-1}$ with $\mathbb{P}_{v \sim S^{n-1}}(v \in E_n) > 1 - c_n$ such that*

$$\left\| \inf_{D, D'} T[M_n(D), M_n(D')] - G_1 \right\|_\infty \leq c_n$$

where the infimum is taken over D, D' such that $\frac{f_n(D') - f_n(D)}{\|f_n(D') - f_n(D)\|_2} \in E_n$.

Here $v \sim S^{n-1}$ means v comes from a uniform distribution of the unit sphere $S^{n-1} \subseteq \mathbb{R}^n$. The conclusion is basically that $\inf_{D, D'} T[M_n(D), M_n(D')] \rightarrow G_1$, i.e. M_n is asymptotically GDP. Similar to the interpretation of (3), it means the mechanism M_n provides the same amount of privacy as a Gaussian mechanism in the limit of $n \rightarrow \infty$. However, a fraction of neighboring datasets has to be excluded. More specifically, the limit holds if the direction of the difference $f_n(D') - f_n(D)$ falls in E_n , an ‘‘almost sure’’ event as the dimension $n \rightarrow \infty$. As we remarked in the introduction, directions in E_n can exhibit low dimensional behavior and hence must be ruled out for any high-dimensional observation.

For a vector $v \in \mathbb{R}^n$ and $\varphi \in \mathfrak{F}_n$, let $P_v^\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$ be defined as $P_v^\varphi(x) = \varphi(x + v) - \varphi(x) - \frac{1}{2}v^T \mathcal{I}_\varphi v$. For two random variables X and Y , their Kolmogorov–Smirnov distance $\text{KS}(X, Y)$ is defined as the ℓ_∞ distance of their CDFs. A sequence of random variables is denoted by $o_P(1)$ if they converge in probability to 0. The technical conditions for the CLT are as follows. Note that each of them are conditions on the function sequence φ_n .

(D1) $\text{KS}(P_v^{\varphi_n}(X_n), v^T \nabla \varphi_n(X_n)) = o(1)$ with probability at least $1 - o(1)$ over $v \sim S^{n-1}$

(D2) $\|\nabla \varphi_n(X_n)\|_2 = \sqrt{n} \cdot (1 + o_P(1))$

Remark 1. Dropping the cumbersome subscripts n , (D1) roughly asks that

$$P_v^\varphi(X) = \varphi(X + v) - \varphi(X) - \frac{1}{2} v^T \mathcal{I}_\varphi v \approx v^T \nabla \varphi(X)$$

Since \mathcal{I}_φ is the expectation of the Hessian of φ , we see that (D1) is basically a regularity condition stating that the Taylor expansion of φ holds on average up to the second order.

Remark 2. Condition (D2) basically says that $\nabla \varphi(X)$ mostly falls on a spherical shell of radius \sqrt{n} (as it should since $\mathcal{I}_\varphi = \mathbb{E}[\nabla \varphi(X) \nabla \varphi(X)^T]$ is assumed to be identity). A deeper understanding is provided by an alternative interpretation of condition (D1), using a new notion we propose called “likelihood projection”.

Likelihood Projection. The function P_v^φ defined above is called the “likelihood projection” along direction v . It is (up to an additive constant) the log likelihood ratio of X_φ and its translation $X_\varphi - v$. In fact, X_φ has density $\frac{1}{Z_\varphi} e^{-\varphi(x)}$ and $X_\varphi - v$ has density $\frac{1}{Z_\varphi} e^{-\varphi(x+v)}$ where Z_φ is the common normalizing constant. The log likelihood ratio is $\varphi(x+v) - \varphi(x)$. This explains the word “likelihood”. To observe its nature as a “projection”, consider the special case $\varphi(x) = \frac{1}{2} \|x\|_2^2$. Straightforward calculation suggests that \mathcal{I}_φ is identity and $P_v^\varphi(x) = v^T x$. So it is indeed a generalization of the linear projection along direction v .

The alternative interpretation of condition (D1) is that when the dimension is high, the “likelihood projection” $P_v^\varphi(X)$ is roughly a linear projection to the direction v . Condition (D2) is then the “thin-shell” condition proposed in Sudakov’s theorem [Sud78] which we state in the appendix as a necessary tool for the proof of our CLT.

For the general case, consider $M_n(D) = f_n(D) + t_n X_{\varphi_n}$ where $t_n = \mu^{-1} \cdot \sqrt{\|\mathcal{I}_n\|_2}$. The factor $\sqrt{\|\mathcal{I}_n\|_2}$ normalizes the Fisher information to the identity, and the factor μ^{-1} controls the final privacy level. For this mechanism, we have

Corollary 3.2. *If the function sequence $\tilde{\varphi}_n(x) = \varphi_n(\|\mathcal{I}_n\|_2^{-\frac{1}{2}} x)$ satisfies conditions (D1) and (D2) and that $\mathcal{I}_n = \|\mathcal{I}_n\|_2 \cdot (1 + o(1)) \cdot I_{n \times n}$, then there is a sequence of positive numbers $c_n \rightarrow 0$ and a subset $E_n \subseteq S^{n-1}$ for each n with $\mathbb{P}_{v \sim S^{n-1}}(v \in E_n) > 1 - c_n$ such that*

$$\| \inf_{D, D'} T[M_n(D), M_n(D')] - G_\mu \|_\infty \leq c_n$$

where the infimum is taken over D, D' such that $\frac{f_n(D') - f_n(D)}{\|f_n(D') - f_n(D)\|_2} \in E_n$.

In particular, when p and α belong to $[1, +\infty)$, norm powers $\|x\|_p^\alpha$ satisfy the above conditions.

Lemma 3.3. *For $p \in [1, +\infty), \alpha \in [1, +\infty)$, let $c_{p,\alpha} = \alpha^{-1} \cdot p^{-\alpha + \frac{\alpha}{p}} \cdot \left(\frac{\Gamma(2 - \frac{1}{p})}{\Gamma(\frac{1}{p})}\right)^{-\frac{\alpha}{2}}$, the sequence of functions $\varphi_n(x) = n^{1 - \frac{\alpha}{p}} \cdot c_{p,\alpha} \|x\|_p^\alpha$ satisfies conditions (D1) and (D2) and that $\mathcal{I}_n = \|\mathcal{I}_n\|_2 \cdot (1 + o(1)) \cdot I_{n \times n}$.*

The parameter $c_{p,\alpha}$ and the power of n are determined by the Fisher information, which can be found in Lemma 4.2. More generally, we conjecture that

Conjecture 3.4. *All functions in \mathfrak{F}_n satisfy (D1) and (D2).*

Recall that $\varphi \in \mathfrak{F}_n$ lead to log-concave distributions. We limit the scope of our conjecture to log-concave distributions because of an interesting lemma involved in the proof of the central limit theorem 3.1. Consider the mechanism $M^t(D) = f(D) + tX$, with the emphasis on the scaling parameter t . As t increases, M^t obviously loses accuracy regardless of log-concavity of X . On the other hand, when it comes to privacy, we have

Lemma 3.5. *When X has log-concave distribution and $t \geq 0$, the ROC function $T[M^t(D), M^t(D')]$ is (pointwise) monotone increasing in t for any D, D' .*

Since larger ROC function means more privacy, this lemma confirms that M^t gains privacy as t increases. In other words, it confirms the existence of “privacy-accuracy trade-off” given the log-concavity of X . Note that without log-concavity, monotonicity in the lemma need not hold. For a

one-dimensional example, consider an X that supports on even numbers and $f(D) = 0, f(D') = 2$. When $t = 2$, $T[M^t(D), M^t(D')] = T[2X, 2X + 2] = T[X, X + 1]$. There is no privacy in this case as X and $X + 1$ has completely disjoint support. On the other hand, when $t = 1$, $T[M^t(D), M^t(D')] = T[X, X + 2]$ and incurs some privacy. That is, more noise does not imply more privacy, hence violating the conclusion of Lemma 3.5.

In summary, results in this section show that mechanisms adding noise that satisfies (D1) and (D2) (e.g. densities $\propto e^{-\|x\|_p^\alpha}$) behave like a Gaussian mechanism. Changing the noise in this class does not change the ROC function by much. Hence we cannot repeat the success at fully utilizing the (ϵ, δ) privacy budget as in Section 2.

On the other hand, our CLT involves Fisher information, and hence gives us the opportunity to relate to the (arguably) most successful tool for constant-sharp lower bound — the Cramer–Rao inequality. This will be the content of the next section.

4 Privacy-Accuracy Trade-off via Cramer–Rao Inequality

The central limit theorem in the previous section suggests that we use GDP parameter μ to measure privacy. Adopting this, we will show that the privacy-accuracy trade-off is naturally characterized by the Cramer–Rao lower bound. The conclusion has a similar flavor to the uncertainty principles.

Recall that the mechanism $M(D) = f(D) + tX_\varphi$ is determined by two “parameters”: the shape parameter $\varphi \in \mathfrak{F}_n$ which determines the distribution of X_φ , and the scale parameter t . If φ also satisfies the conditions of Theorem 3.1, then we can use the desired (asymptotic) GDP parameter μ to determine the scale parameter, i.e. $t = \mu^{-1} \cdot \sqrt{\|\mathcal{I}_\varphi\|_2}$. Using the equivalent parametrization (μ, φ) , the corresponding mechanism $M_{\mu, \varphi}$ is given by

$$M_{\mu, \varphi}(D) = f(D) + \mu^{-1} \cdot \sqrt{\|\mathcal{I}_\varphi\|_2} \cdot X_\varphi. \quad (4)$$

As we have explained in the introduction, one way to measure the accuracy of the mechanism is the mean squared error of the noise

$$\text{err}(M_{\mu, \varphi}) = \mathbb{E}\|tX_\varphi\|_2^2 = \mu^{-2} \cdot \|\mathcal{I}_\varphi\|_2 \cdot \mathbb{E}\|X_\varphi\|_2^2. \quad (5)$$

The following theorem characterizes the privacy-accuracy trade-off as the product of the mean squared error $\text{err}(M_{\mu, \varphi})$ and privacy parameter μ^2 .

Theorem 4.1 (Restating Theorem 1.2). *For any $\varphi \in \mathfrak{F}_n$ and $M_{\mu, \varphi}$ defined as in (4), we have*

$$\mu^2 \cdot \text{err}(M_{\mu, \varphi}) \geq n.$$

In addition, the equality holds if the added noise X is n -dimensional standard Gaussian.

Proof of Theorem 4.1. To simplify notations we will drop the subscript φ in X . We first claim that it suffices to show the following uncertainty-principle-like result

$$\text{Var}[X] \cdot \text{Var}[\nabla\varphi(X)] \geq n^2. \quad (6)$$

where the notation $\text{Var}[\cdot]$ is slightly abused to denote the mean squared distance of a random vector from its expectation, i.e. $\text{Var}[X] = \mathbb{E}[\|X - \mathbb{E}X\|_2^2]$.

To see why (6) suffices, notice that by (5), the interested quantity can be simplified as

$$\mu^2 \cdot \text{err}(M_{\mu, \varphi}) = \mathbb{E}\|X\|_2^2 \cdot \|\mathcal{I}_\varphi\|_2. \quad (7)$$

Recall that we have $\mathbb{E}X = 0$ by symmetry of φ and $\mathbb{E}\nabla\varphi(X) = 0$ by basic Fisher information theory. So $\text{Var}[\nabla\varphi(X)] = \mathbb{E}\|\nabla\varphi(X)\|_2^2 = \text{Tr} \mathbb{E}\nabla\varphi(X)\nabla\varphi(X)^T = \text{Tr} \mathcal{I}_\varphi$. That is, eq. (6) implies

$$\mathbb{E}\|X\|_2^2 \cdot \text{Tr} \mathcal{I}_\varphi \geq n^2. \quad (8)$$

Since \mathcal{I}_φ is positive semi-definite, by (7) and (8) we have

$$\mu^2 \cdot \text{err}(M_{\mu, \varphi}) = \mathbb{E}\|X\|_2^2 \cdot \|\mathcal{I}_\varphi\|_2 \geq \mathbb{E}\|X\|_2^2 \cdot \frac{1}{n} \text{Tr} \mathcal{I}_\varphi \geq n.$$

Table 1: Explicit expressions of Fisher information and mean squared error.

Density	$\ \mathcal{I}_\varphi\ _2$	$\mathbb{E}\ X\ _2^2$	$\mathbb{E}\ X\ _2^2 \cdot \ \mathcal{I}_\varphi\ _2$	$\mathbb{E}\ X\ _\infty^2$	$\mathbb{E}\ X\ _\infty^2 \cdot \ \mathcal{I}_\varphi\ _2$
$\propto e^{-\ x\ _1}$	1	$2n$	$2n$	$\sim (\log n)^2$	$\sim (\log n)^2$
$\propto e^{-\ x\ _2}$	$\frac{1}{n}$	$n(n+1)$	$n+1$	$\sim 2n \log n$	$\sim 2 \log n$
$\propto e^{-\ x\ _2^2}$	2	$\frac{1}{2}n$	n	$\sim \log n$	$\sim 2 \log n$
$\propto e^{-\ x\ _p^\alpha}$	Lemma 4.2	Lemma 4.2	$\sim C_p \cdot n$	Appendix	$\leq C'_p \cdot (\log n)^{\frac{2}{p}}$

Next we focus on the proof of (6). Consider the location family $\{X + \theta : \theta \in \mathbb{R}^n\}$. The Fisher information of this family is \mathcal{I}_φ at all θ . The random vector itself is an unbiased estimator of the location. Therefore, by the Cramer–Rao inequality (c.f. [VdV00]), we have that $\text{Cov}(X) - \mathcal{I}_\varphi^{-1}$ is positive semi-definite. As a consequence,

$$\text{Var}[X] = \text{Tr Cov}(X) \geq \text{Tr } \mathcal{I}_\varphi^{-1} = \lambda_1^{-1} + \dots + \lambda_n^{-1}.$$

where $\lambda_1 \geq \dots \geq \lambda_n > 0$ are the eigenvalues of \mathcal{I}_φ . We already see that $\text{Var}[\nabla\varphi(X)] = \text{Tr } \mathcal{I}_\varphi$, so by Cauchy–Schwarz inequality,

$$\text{Var}[X] \cdot \text{Var}[\nabla\varphi(X)] \geq (\lambda_1^{-1} + \dots + \lambda_n^{-1})(\lambda_1 + \dots + \lambda_n) \geq n^2$$

The proof of the inequality is complete. For standard Gaussian, we have $\text{Cov}(X) = \mathcal{I}_\varphi = I_{n \times n}$, and we have $\mu^2 \cdot \text{err}(M_{\mu,\varphi}) = \mathbb{E}\|X\|_2^2 \cdot \|\mathcal{I}_\varphi\|_2 = \text{Tr } I_{n \times n} \cdot 1 = n$. \square

Note that although Theorem 4.1 holds true for very general φ (only integrability conditions are imposed in \mathfrak{F}_n), the interpretation that μ is the asymptotic privacy parameter only holds for distributions that satisfy (D1) and (D2). Therefore, let us consider the special case where $\varphi(x) = \|x\|_p^\alpha$. The corresponding X_φ will be denoted by $X_{p,\alpha}$ and \mathcal{I}_φ by $\mathcal{I}_{p,\alpha}$. In this special case, we can compute the quantities in (8) exactly. In the following lemma, we write $a_n \sim b_n$ for the two sequences a_n and b_n if $\frac{a_n}{b_n} \rightarrow 1$ as $n \rightarrow \infty$.

Lemma 4.2. *For $1 \leq p < \infty$ and $1 \leq \alpha < \infty$, as $n \rightarrow \infty$, we have*

$$\begin{aligned} \mathbb{E}\|X_{p,\alpha}\|_2^2 &\sim n^{\frac{2}{\alpha} - \frac{2}{p} + 1} \cdot \alpha^{-\frac{2}{\alpha}} \cdot p^{\frac{2}{p}} \cdot \Gamma\left(\frac{3}{p}\right) / \Gamma\left(\frac{1}{p}\right); \\ \mathcal{I}_{p,\alpha} &\sim n^{\frac{2}{p} - \frac{2}{\alpha}} \cdot \alpha^{\frac{2}{\alpha}} \cdot p^{2 - \frac{2}{p}} \cdot \Gamma\left(2 - \frac{1}{p}\right) / \Gamma\left(\frac{1}{p}\right) \cdot I_{n \times n}. \end{aligned}$$

This result put Theorem 4.1 into a more concrete context. Some important cases with specific values of p and α are worked out in Table 1. Remarkably, in the last row, the products that characterize privacy-accuracy trade-off are asymptotically independent of α . As a by-product of this calculation, we also derive the expression for the isotropic constant of the n -dimensional ℓ_p ball, which is an important concept in convex geometry (c.f. [BGVV14]). See the appendix for more results and discussion.

Alternatively, we may want to measure the accuracy by the expected squared ℓ_∞ -norm of the noise. A similar argument suggests to consider the following quantity $\mathbb{E}\|X_\varphi\|_\infty^2 \cdot \|\mathcal{I}_\varphi\|_2$. By Theorem 4.1 and the fact that $\|x\|_\infty \geq \frac{1}{\sqrt{n}}\|x\|_2$, we have

$$\mathbb{E}\|X_\varphi\|_\infty^2 \cdot \|\mathcal{I}_\varphi\|_2 \geq \frac{1}{n} \mathbb{E}\|X_\varphi\|_2^2 \cdot \|\mathcal{I}_\varphi\|_2 \geq 1. \quad (9)$$

We would like to point out a connection to a recently resolved open problem proposed in [SU17], asking if there is a DP algorithm that answers a high-dimensional query with ℓ_2 -sensitivity 1 with $O(1)$ error in ℓ_∞ norm. In particular, the recent solution [DK20, GKM20] provides strong evidence that the lower bound in (9) is tight up to a constant factor.

An Analogy with Uncertainty Principles There are various mathematical manifestations of the uncertainty principle. The one behind Hesenberg uncertainty principle is that a function and its Fourier transform cannot both be localized simultaneously. Specifically, for a function $f \in L^2(\mathbb{R}^n)$, its Fourier transform is defined as $\hat{f}(\xi) = \int e^{-2\pi i \langle \xi, x \rangle} f(x) dx$. Fourier transform is unitary, i.e. $\|f\|_{L^2} = \|\hat{f}\|_{L^2}$. In particular, if $|f|^2$ is a probability density, then so is $|\hat{f}|^2$. Our previous

abuse of notation also applies here, for example, $\text{Var}[|f|^2] = \int (x - a)^T (x - a) |f(x)|^2 dx$ where $a = \int x |f(x)|^2 dx$. For $\|f\|_{L^2} = 1$, we have the following result⁵ (c.f. Corollary 2.8 of [FS97])

$$\text{Var}[|f|^2] \cdot \text{Var}[\hat{f}^2] \geq \frac{n^2}{16\pi^2}. \quad (10)$$

The similarity between (6) and (10) suggests that Theorem 4.1 can be considered as yet another manifestation of the uncertainty principle.

5 Conclusions and Future Works

In this work, we study constant-sharp optimality of noise addition algorithms for high-dimensional query answering with differential privacy. We demonstrate that the ROC function offers good insight in comparing the “actual spend vs budget” of differential privacy and hence in the design of one-dimensional algorithms. However, when the dimension is high, a CLT shows that (ϵ, δ) privacy budget cannot be fully spent for a large class of noise addition mechanisms as they all behave like a Gaussian mechanism. On the other hand, Fisher information naturally arises in these high-dimensional mechanisms, and the simple and fundamental quantity “privacy parameter \times error” automatically manifests itself as the quantity “information \times error” in the Cramer–Rao lower bound. Using this, we are able to show an elegant characterization of the precise privacy-accuracy trade-off, and justify the constant-sharp optimality of the Gaussian mechanism. We believe the insights offer a novel perspective to the long-lived privacy-accuracy trade-off question.

Various extensions are possible. An immediate one is to extend the CLT to a broader class of noise distributions, such as log-concave distributions as specified in Conjecture 3.4. Another condition imposed by \mathfrak{F}_n (implicitly) is that the noise must be supported on the whole space. The difficulty in removing the condition lies in the lack of a definition of Fisher information for noise with bounded support. In particular, one may consider extending the theories to cover the noise used in [DK20] and prove a corresponding lower bound like (9). For non-log-concave noise, Lemma 3.5 suggests us to believe that a corresponding log-concave noise with no less privacy and accuracy exists. For algorithms beyond noise addition or problems beyond query answering, we believe that they still exhibit some universal behavior as long as the dimension is high. As a circumstantial evidence, [BDKT12] shows that generic algorithms for query answering can be reduced to a noise addition one with better accuracy and slightly worse privacy.

Acknowledgements

We thank Jason Hartline, Yin-Tat Lee, Haotian Jiang, Qiyang Han, Sasho Nikolov, Aravindan Vijayaraghavan and Yuansi Chen for helpful comments on earlier versions of the manuscript. J. D. was supported by the NSF HDR TRIPODS award CCF-1934931. W. J. S. was supported in part by NSF through CCF-1763314 and CAREER DMS-1847415, an Alfred Sloan Research Fellowship, and a Facebook Faculty Research Award. L. Z. was supported in part by NSF through DMS-2015378.

Societal Impact

Private data analysis has positive societal impacts. The major negative concern is that too much utility is sacrificed for privacy. This work is intended to improve our theoretical understanding of such trade-off between privacy and utility.

References

- [BBG20] Borja Balle, Gilles Barthe, and Marco Gaboardi. Privacy profiles and amplification by subsampling. *Journal of Privacy and Confidentiality*, 10(1), 2020.
- [BDKT12] Aditya Bhaskara, Daniel Dadush, Ravishankar Krishnaswamy, and Kunal Talwar. Unconditional differentially private mechanisms for linear queries. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 1269–1284, 2012.

⁵Hessenberg uncertainty principle is a direct consequence of Equation (10) and the fact that the position operator and momentum operator are conjugate of each other via Fourier transform.

- [BGVV14] Silouanos Brazitikos, Apostolos Giannopoulos, Petros Valettas, and Beatrice-Helen Vritsiou. *Geometry of isotropic convex bodies*, volume 196. American Mathematical Soc., 2014.
- [BUV18] Mark Bun, Jonathan Ullman, and Salil Vadhan. Fingerprinting codes and the price of approximate differential privacy. *SIAM Journal on Computing*, 47(5):1888–1938, 2018.
- [BW18] Borja Balle and Yu-Xiang Wang. Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. In *International Conference on Machine Learning*, pages 403–412, 2018.
- [CDT98] G Calafiore, F Dabbene, and R Tempo. Uniform sample generation in l_p -balls for probabilistic robustness analysis. In *Proceedings of the 37th IEEE Conference on Decision and Control (Cat. No. 98CH36171)*, volume 3, pages 3335–3340. IEEE, 1998.
- [CWZ20] T Tony Cai, Yichen Wang, and Linjun Zhang. The cost of privacy in generalized linear models: Algorithms and minimax lower bounds. *arXiv preprint arXiv:2011.03900*, 2020.
- [CWZ21] T Tony Cai, Yichen Wang, and Linjun Zhang. The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *The Annals of Statistics*, 2021.
- [DK20] Yuval Dagan and Gil Kur. A bounded-noise mechanism for differential privacy. *arXiv preprint arXiv:2012.03817*, 2020.
- [DKM⁺06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 486–503. Springer, 2006.
- [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006.
- [DR14] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3 & 4):211–407, 2014.
- [DRS21] Jinshuo Dong, Aaron Roth, and Weijie J Su. Gaussian differential privacy. *Journal of the Royal Statistical Society, Series B*, 2021. to appear.
- [ENU20] Alexander Edmonds, Aleksandar Nikolov, and Jonathan Ullman. The power of factorization mechanisms in local and central differential privacy. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 425–438, 2020.
- [FS97] Gerald B Folland and Alladi Sitaram. The uncertainty principle: a mathematical survey. *Journal of Fourier analysis and applications*, 3(3):207–238, 1997.
- [GDGK20] Quan Geng, Wei Ding, Ruiqi Guo, and Sanjiv Kumar. Tight analysis of privacy and utility tradeoff in approximate differential privacy. In *International Conference on Artificial Intelligence and Statistics*, pages 89–99. PMLR, 2020.
- [Gia] Apostolos Giannopoulos. Geometry of isotropic convex bodies and the slicing problem.
- [GKM20] Badih Ghazi, Ravi Kumar, and Pasin Manurangsi. On avoiding the union bound when answering multiple differentially private queries. *arXiv preprint arXiv:2012.09116*, 2020.
- [GV16] Quong Geng and Pramod Viswanath. The optimal noise-adding mechanism in differential privacy. *IEEE Transactions on Information Theory*, 62(2):925–951, 2016.
- [Kla10] Bo’az Klartag. High-dimensional distributions with convexity properties. In *European Congress of Mathematics Amsterdam, 14–18 July, 2008*, pages 401–417, 2010.

- [KOV17] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. *IEEE Transactions on Information Theory*, 63(6):4037–4049, 2017.
- [KSSU20] Gautam Kamath, Or Sheffet, Vikrant Singhal, and Jonathan Ullman. Differentially private algorithms for learning mixtures of separated gaussians. In *2020 Information Theory and Applications Workshop (ITA)*, pages 1–62. IEEE, 2020.
- [SSTT21] Shuang Song, Thomas Steinke, Om Thakkar, and Abhradeep Thakurta. Evading the curse of dimensionality in unconstrained private glms. In *International Conference on Artificial Intelligence and Statistics*, pages 2638–2646. PMLR, 2021.
- [SU17] Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *Journal of Privacy and Confidentiality*, 7(2):3–22, 2017.
- [Sud78] Vladimir Nikolaevich Sudakov. Typical distributions of linear functionals in finite-dimensional spaces of higher dimension. In *Doklady Akademii Nauk*, volume 243, pages 1402–1405. Russian Academy of Sciences, 1978.
- [VdV00] Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.
- [Wan05] Xianfu Wang. Volumes of generalized unit balls. *Mathematics Magazine*, 78(5):390–395, 2005.
- [WZ10] Larry Wasserman and Shuheng Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489):375–389, 2010.
- [ZZ21] Zhe Zhang and Linjun Zhang. High-dimensional differentially-private em algorithm: Methods and near-optimal statistical guarantees. *arXiv preprint arXiv:2104.00245*, 2021.

Supplemental Materials

In Appendix A we provide the detail of the numerical experiments. The central limit theorem 3.1 (normalized) and 3.2 (general case) are proved in Appendix B. Appendix C proves Lemma 4.2 and provide some additional results that apply beyond norm powers. The proof of Lemma 3.3, which verifies that norm powers satisfy the technical conditions (D1) and (D2), requires results in Appendix C and takes significant effort, so we dedicate the entire Appendix D to it.

Appendix A Numerical Verification of the Central Limit Theorem

This section discusses the details of the numerical experiments shown in Figure 1 (repeated below) that verifies our central limit theorem.




Figure 1: Fast convergence to GDP as claimed in Theorem 1.1. Blue solid curves indicate the true privacy (i.e. ROC functions, see Section 2 for details) of the noise addition mechanism considered in Theorem 1.1. Red dashed curves are GDP limit predicted by our CLT. In all three panels the dimension $n = 30$. In order to show that our theory works for general ℓ_p -norm to the power α , we pick them to be famous mathematical constants, namely $p = \pi$, $\alpha = e$ and the coefficient being the Euler–Mascheroni constant γ .

The mechanism in consideration is $M(D) = f(D) + tX$ where the n -dimensional random vector X has density $\propto e^{-\|x\|_p^\alpha}$. We want to demonstrate that when $t = \mu^{-1} \cdot \sqrt{\|\mathcal{I}_{p,\alpha}\|_2} = \mu^{-1} \cdot n^{\frac{1}{p}-\frac{1}{\alpha}} \cdot \alpha^{\frac{1}{\alpha}} \cdot p^{1-\frac{1}{p}} \cdot \sqrt{\frac{\Gamma(2-\frac{1}{p})}{\Gamma(\frac{1}{p})}} \cdot 1 + (o(1))$ we have

$$\inf_{D, D'} T[M(D), M(D')] \approx G_\mu.$$

The infimum is taken over D, D' such that the direction of $f_n(D) - f_n(D')$ is in a large subset E_n of the unit sphere. It is hard to evaluate the infimum even numerically, but it turns out that the infimum is equal to $\inf_{v \in E_n} T[tX, tX + v]$. This is the first part of the proof of Theorem 1.1.

Therefore, it suffices to evaluate $T(tX, tX + v)$ and compare with the GDP function G_μ , but the high-dimensional nature of X prevents exact evaluation, so we will introduce a Monte Carlo approach.

Empirical ROC Function X has density $\propto e^{-\varphi(x)}$ and $X + v$ has density $\propto e^{-\varphi(x-v)}$. The log likelihood ratio is $\varphi(x) - \varphi(x - v)$. Thresholding it at h yields the following type I and type II errors

$$\begin{aligned} \alpha(h) &= \mathbb{P}_{x \sim X}(\varphi(x) - \varphi(x - v) \geq h) = \mathbb{P}(\varphi(X) - \varphi(X - v) \geq h) \\ \beta(h) &= \mathbb{P}_{x \sim X+v}(\varphi(x) - \varphi(x - v) < h) = \mathbb{P}(\varphi(X + v) - \varphi(X) < h) \end{aligned}$$

Once we have these, $T(X, X + v)$ can be obtained by eliminating h and express β as a function of α . These two probabilities can be computed by a simple Monte Carlo approach. First We can sample




Figure 3: The effect of sample size N . We see that $N = 10000$ provides a good estimate. The red dashed curve corresponds to Gaussian noise $\varphi(x) = \frac{1}{2}\|x\|_2^2$, which is why we can perform exact computation. The dimension $n = 30$.

$\{x_1, \dots, x_N\}$ as i.i.d. copies of X . Let

$$\begin{aligned} a_i &= \varphi(x_i - v) - \varphi(x_i) \\ b_i &= \varphi(x_i + v) - \varphi(x_i) \\ \hat{\alpha}(h) &= \frac{1}{N} \cdot \#\{a_i \leq -h\} \\ \hat{\beta}(h) &= \frac{1}{N} \cdot \#\{b_i < h\} \end{aligned}$$

We only evaluate on a discrete set of h such that the corresponding α forms a uniform grid $\{\frac{1}{N}, \dots, \frac{N}{N}\}$. Let $h_j = -a_{(j)}$ where $a_{(1)} \leq \dots \leq a_{(N)}$ are order statistics of $a_i, i = 1, 2, \dots, N$. Then for $j = 0, 1, 2, \dots, N$,

$$\begin{aligned} \hat{\alpha}_j &= \hat{\alpha}(h_j) = \frac{1}{N} \cdot \#\{a_i \leq a_{(j)}\} = \frac{j}{N} \\ \hat{\beta}_j &= \hat{\beta}(h_j) = \frac{1}{N} \cdot \#\{b_i < a_{(j)}\} \end{aligned}$$

Let $\hat{T}_N(X, X + v)$ be the function that linearly interpolates the values $\hat{\beta}_0, \dots, \hat{\beta}_N$ at $\frac{0}{N}, \dots, \frac{N}{N}$. As a direct consequence of the well-known Glivenko–Cantelli theorem, we have

$$\|\hat{T}_N(X, X + v) - T(X, X + v)\|_\infty \rightarrow 0 \text{ almost surely.}$$

We evaluate the effect of the sample size N in Figure 3 and observe that $N = 10000$ works quite well when the $T[X, X + v]$ is close to the 1-GDP function G_1 (which is true for all experiments in the paper).

Evaluating the Fisher Information Note that it is numerically infeasible to use the exact expression in Lemma 4.2, since Gamma function grows extremely fast (of course it does, as an interpolation of the factorial). In practice, we find the asymptotic expression in Lemma 4.2 works extremely well.

Next we present the algorithm that samples an n -dimensional random vector whose density is $\propto e^{-\|x\|_p^\alpha}$.

It is easy to see its correctness from Lemma C.5 and [CDT98].

Appendix B Proof of Theorem 3.1 and 3.2

In this section we first prove the normalized central limit theorem 3.1 and then the general case theorem 3.2. Recall that in normalized CLT, the mechanism in consideration is $M_n(D) = f_n(D) + X_n$, where X_n has density $\propto e^{-\varphi_n}$ with Fisher information \mathcal{I}_n being the identity matrix $I_{n \times n}$.

Theorem 3.1. *If the function sequence φ_n satisfies conditions (D1) and (D2), then there is a sequence of positive numbers c_n with $c_n \rightarrow 0$ as $n \rightarrow \infty$ and a subset $E_n \subseteq S^{n-1}$ with $\mathbb{P}_{v \sim S^{n-1}}(v \in E_n) >$*

Algorithm 1 Sample $\propto e^{-\|x\|_p^\alpha}$

- 1: **Input:** p, α and dimension n
 Generate $t \sim \Gamma(\frac{n}{\alpha} + 1, 1)$
 Generate $\xi_i, i = 1, 2, \dots, n$ i.i.d. from $\Gamma(\frac{1}{p}, 1)$
 Generate random vector $x \in \mathbb{R}^n$ where $x_i = \varepsilon_i \cdot \xi_i^{\frac{1}{p}}$, ε_i are Rademacher random variables (unbiased coin flips) independent from everything else.
 Generate $r \sim U[0, 1]$
 Let $V = r^{\frac{1}{n}} \cdot \frac{x}{\|x\|_p}$
 - 2: **Output** $t \cdot V$
-

$1 - c_n$ such that

$$\| \inf_{D, D'} T[M_n(D), M_n(D')] - G_1 \|_\infty \leq c_n$$

where the infimum is taken over D, D' such that $\frac{f_n(D') - f_n(D)}{\|f_n(D') - f_n(D)\|_2} \in E_n$.

(D1) $\text{KS}(P_v^{\varphi_n}(X_n), v^T \nabla \varphi_n(X_n)) = o(1)$ with probability at least $1 - o(1)$ over $v \sim S^{n-1}$

(D2) $\|\nabla \varphi_n(X_n)\|_2 = \sqrt{n} \cdot (1 + o_P(1))$

B.1 The Main Proof

Proof of Theorem 3.1. For clarity, in the proof we drop the subscript n unless the limit $n \rightarrow \infty$ is taken. First we show that

$$\inf_{D, D'} T[M(D), M(D')] = \inf_{v \in E_n} T[X, X + v].$$

Notice that

$$T[M(D), M(D')] = T[f(D) + X, f(D') + X] = T[X, X + f(D') - f(D)]$$

Consider the vector $f(D') - f(D)$. Let $v = \frac{f(D') - f(D)}{\|f(D') - f(D)\|_2}$ be its direction and r be its length. We have $f(D') - f(D) = rv$. Since f has ℓ_2 -sensitivity 1, we have $r \in [0, 1]$. The infimum over D, D' can be taken in two steps: first over v and then over r . That is,

$$\inf_{D, D'} T[M(D), M(D')] = \inf_{v \in E_n} \inf_{r \in [0, 1]} T[M(D), M(D')] = \inf_{v \in E_n} \inf_{r \in [0, 1]} T[X, X + rv]$$

By Lemma 3.5, $T[X, X + rv]$ is pointwise monotone decreasing in r , so for the inner infimum we have $\inf_{r \in [0, 1]} T[X, X + rv] = T[X, X + v]$. Back to the limiting conclusion, it suffices to show that for all $v \in E_n$,

$$\|T[X_n, X_n + v] - G_1\|_\infty \leq c_n.$$

To prove this, we use the following lemma

Lemma B.1. *Suppose random vector X has density $\propto e^{-\varphi}$ where $\varphi \in \mathfrak{F}_n$. Let F_v be the CDF of the likelihood projection $P_v^\varphi(X) = \varphi(X + v) - \varphi(X) - \frac{1}{2}v^T \mathcal{I}_\varphi v$, then for any $v \in \mathbb{R}^n$,*

$$T[X, X + v](\alpha) = F_v(-F_v^{-1}(\alpha) - v^T \mathcal{I}_\varphi v).$$

Let H_v^n and F_v^n be the CDFs of the linear projection $v^T \nabla \varphi_n(X_n)$ and the likelihood projection $P_v^{\varphi_n}(X_n) = \varphi_n(X_n + v) - \varphi_n(X_n) - \frac{1}{2}v^T \mathcal{I}_n v$. When Lemma B.1 is applied to X_n and unit vector v , we have

$$T[X_n, X_n + v](\alpha) = F_v^n(- (F_v^n)^{-1}(\alpha) - 1). \quad (11)$$

Recall that $G_1(\alpha) = \Phi(-\Phi^{-1}(\alpha) - 1)$ where Φ is the CDF of standard normal. Comparing with (11), it suffices to show F_v^n is close to Φ . To prove this and to take care of the set $E_n \subseteq S^{n-1}$, we use the two conditions (D1) and (D2), which allow us to apply Sudakov's Theorem (c.f. [Kla10]) stated below as Lemma B.2. In the following, let σ_n be the uniform measure (with total measure 1) on the unit sphere S^{n-1} .

Lemma B.2. Let Y_n be an isotropic random vector in \mathbb{R}^n . Assume that there is $a_n \rightarrow 0$ and

$$\mathbb{P}\left(\left|\frac{\|Y_n\|_2}{\sqrt{n}} - 1\right| \geq a_n\right) \leq a_n. \quad (12)$$

Then, there exists $b_n \rightarrow 0$ and $\Theta_n \subseteq S^{n-1}$ with $\sigma_n(\Theta_n) \geq 1 - b_n$, such that for any $v \in \Theta_n$,

$$\sup_{t \in \mathbb{R}} |\mathbb{P}(v^T Y_n \leq t) - \Phi(t)| \leq b_n.$$

Let $Y_n = \varphi_n(X_n)$. $v^T \nabla \varphi_n(X_n) = v^T Y_n$. We know $\mathbb{E}Y_n = \mathbb{E}\nabla \varphi_n(X_n) = 0$, and by the normalization of the Fisher information,

$$\mathbb{E}Y_n Y_n^T = \mathcal{I}_n = I_{n \times n}.$$

Therefore, Y_n is isotropic. Condition (D2) says $\|Y_n\|_2 = \sqrt{n} \cdot (1 + o_P(1))$. That is, $\left|\frac{\|Y_n\|_2}{\sqrt{n}} - 1\right| = o_P(1)$. This implies the existence of a_n in (12). Therefore, we can apply Lemma B.2 to $Y_n = \varphi_n(X_n)$ and conclude that there is $b_n \rightarrow 0$ and $\Theta_n \subseteq S^{n-1}$ with $\sigma_n(\Theta_n) \geq 1 - b_n$, such that for any $v \in \Theta_n$,

$$\|H_v^n - \Phi\|_\infty \leq b_n.$$

Condition (D1) says, there is $b'_n \rightarrow 0$ and $\Omega_n \subseteq S^{n-1} \geq b'_n$ with $\sigma_n(\Omega_n)$ such that for all $v \in \Omega_n$,

$$\|F_v^n - H_v^n\|_\infty \leq b'_n$$

So $v \in \Theta_n \cap \Omega_n$ implies $\|H_v^n - \Phi\|_\infty \leq b_n$ and $\|F_v^n - H_v^n\|_\infty \leq b'_n$. Therefore, $\|F_v^n - \Phi\|_\infty \leq b_n + b'_n$. Set $E_n = \Theta_n \cap (-\Theta_n) \cap \Omega_n \cap (-\Omega_n)$. Then $v \in E_n$ implies

$$\|F_v^n - \Phi\|_\infty \leq b_n + b'_n \text{ and } \|F_{-v}^n - \Phi\|_\infty \leq b_n + b'_n$$

That is, for any $x \in \mathbb{R}$,

$$\begin{aligned} \Phi(x) - b_n - b'_n &\leq F_v^n(x) \leq \Phi(x) + b_n + b'_n \\ \Phi(x) - b_n - b'_n &\leq F_{-v}^n(x) \leq \Phi(x) + b_n + b'_n \end{aligned}$$

As a consequence of the second inequality,

$$\Phi^{-1}(\alpha - b_n - b'_n) \leq (F_{-v}^n)^{-1}(\alpha) \leq \Phi^{-1}(\alpha + b_n + b'_n)$$

Therefore, when $v \in E_n$, by (11) and the inequalities above,

$$\begin{aligned} T[X_n, X_n + v](\alpha) &= F_v^n(- (F_{-v}^n)^{-1}(\alpha) - 1) \\ &\leq F_v^n(-\Phi^{-1}(\alpha - b_n - b'_n) - 1) \\ &\leq \Phi(-\Phi^{-1}(\alpha - b_n - b'_n) - 1) + b_n + b'_n \\ &= G_1(\alpha - b_n - b'_n) + b_n + b'_n \\ &\leq G_1(\alpha) + C\sqrt{b_n + b'_n} + b_n + b'_n \end{aligned}$$

The final step used the Hölder continuity of G_μ which we state as Lemma B.3 and prove afterwards.

Lemma B.3. Let $G_\mu = T[N(0, 1), N(\mu, 1)]$ for $\mu \geq 0$. Then G_μ is α -Hölder continuous for any $\alpha < 1$.

It is worth noting that G_μ is not 1-Hölder continuous (i.e. Lipschitz continuous) as long as $\mu > 0$.

Without loss of generality assume $C > 1$. Let $c_n = C\sqrt{b_n + b'_n} + b_n + b'_n$. Then $c_n \geq 2b_n + 2b'_n$. The above argument shows $T[X_n, X_n + v](\alpha) \leq G_1(\alpha) + c_n$. The lower bound can be obtained similarly, so for $v \in E_n$, we have

$$\|T[X_n, X_n + v] - G_1\|_\infty \leq c_n.$$

Since all four sets $\Theta_n, -\Theta_n, \Omega_n$ and $-\Omega_n$ are large, we have

$$\sigma_n(E_n) = \sigma_n(\Theta_n \cap (-\Theta_n) \cap \Omega_n \cap (-\Omega_n)) \geq 1 - (2b_n + 2b'_n) \geq 1 - c_n$$

This is the conclusion stated in the theorem. The proof is complete. \square

B.2 Proof of Lemmas

Next we provide the proofs of the lemmas used, namely Lemmas 3.5, B.1 and B.3.

Let $M^t(D) = f(D) + tX$.

Lemma 3.5. *When X has log-concave distribution and $t \geq 0$, the ROC function $T[M^t(D), M^t(D')]$ is (pointwise) monotone increasing in t for any D, D' .*

Proof of Lemma 3.5. Let $0 \leq t_1 \leq t_2$ and $f_i = T[X, X + t_i v]$, $i = 1, 2$. We need to show $f_1 \geq f_2$. Fix $\alpha \in [0, 1]$, let $A_t \subseteq \mathbb{R}^n$ be the optimal rejection region for the testing of X vs $X + tv$. That is,

$$\mathbb{P}[X \in A_t] = \alpha \quad \text{and} \quad \mathbb{P}[X + tv \notin A_t] = T[X, X + tv](\alpha)$$

In order to show $f_1(\alpha) \geq f_2(\alpha)$, consider a translated set $A_{t_1} + (t_2 - t_1)v$. This set is at the best suboptimal for the testing of X vs $X + t_2v$. If we denote $\mathbb{P}[X \in A_{t_1} + (t_2 - t_1)v]$ by α' , suboptimality means

$$\mathbb{P}[X + t_2v \notin A_{t_1} + (t_2 - t_1)v] \geq f_2(\alpha').$$

If we can show $\alpha' \leq \alpha$, then by the monotonicity of ROC functions, we have

$$\begin{aligned} f_2(\alpha) &\leq f_2(\alpha') \\ &= \mathbb{P}[X + t_2v \notin A_{t_1} + (t_2 - t_1)v] \\ &= \mathbb{P}[X + t_1v \notin A_{t_1}] \\ &= f_1(\alpha) \end{aligned}$$

So the only thing left is to show $\alpha' \leq \alpha$, or equivalently,

$$\mathbb{P}[X \in A_{t_1} + (t_2 - t_1)v] \leq \mathbb{P}[X \in A_{t_1}].$$

In fact, we will show $A_{t_1} + (t_2 - t_1)v \subseteq A_{t_1}$. This is where log-concavity kicks in. To phrase it more generally, we are going to show that $A_t + sv \subseteq A_t$ for general $t, s \geq 0$. Suppose X has density $e^{-\varphi(x)}$ where $\varphi : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ is a (potentially extended) convex function. By Neyman–Pearson lemma, $A_t = \{x : \varphi(x) - \varphi(x - tv) > h\}$ for some threshold h . We would like to show that $x \in A_t$ implies $x + sv \in A_t$. It suffices to show

$$\varphi(x + sv) - \varphi(x + sv - tv) \geq \varphi(x) - \varphi(x - tv).$$

In fact, $\varphi(x + sv) - \varphi(x + sv - tv)$ is monotone increasing as a function of s . This is a direct consequence of the convexity of $\varphi(x + sv)$ as a function of s . More specifically, let $g(s) = \varphi(x + sv)$. Its convexity follows from the convexity of φ . For $t \geq 0$, $\varphi(x + sv) - \varphi(x + sv - tv) = g(s) - g(s - t)$ is easily seen to be monotone by taking a derivative, or from a more rigorous approach following just the definition of convex functions. \square

Recall that $\tilde{\varphi}_n(x) = \varphi_n(\frac{x}{t_n})$. The corresponding random vector with density $\propto e^{-\tilde{\varphi}_n}$ is $\tilde{X}_n = X_{\tilde{\varphi}_n}$ and has the same distribution as $t_n X_{\varphi_n}$. The scaling factor normalizes its Fisher information to be an $n \times n$ scalar matrix. In fact,

$$\tilde{\mathcal{I}}_n := \mathcal{I}_{\tilde{\varphi}_n} = t_n^{-2} \mathcal{I}_{\varphi_n} = \mu^2 \cdot I_{n \times n}.$$

Proof of Lemma B.1. We are interested in the hypothesis testing $H_0 : X$ vs $H_1 : X + v$. By definition of ROC function in Equation (2), we need to find out the optimal type II error at a given level α . By Neymann–Pearson lemma, it suffices to consider likelihood ratio tests. The log density of the null is (up to an additive constant) $-\varphi(x)$, while that of the alternative is $-\varphi(x - v)$. So the log likelihood ratio is $\varphi(x) - \varphi(x - v)$. Under null it is distributed as $\varphi(X) - \varphi(X - v)$ and thresholding at h yields type I error

$$\begin{aligned} \alpha &= \mathbb{P}(\varphi(X) - \varphi(X - v) > h) \\ &= \mathbb{P}(\varphi(X - v) - \varphi(X) < -h) \\ &= \mathbb{P}(\varphi(X - v) - \varphi(X) - \frac{1}{2}v^T \mathcal{I}_\varphi v < -h - \frac{1}{2}v^T \mathcal{I}_\varphi v) \\ &= F_{-v}(-h - \frac{1}{2}v^T \mathcal{I}_\varphi v) \end{aligned}$$

Under the alternative, the log likelihood ratio is distributed as $\varphi(X+v) - \varphi(X)$, so the corresponding type II error is

$$\begin{aligned}\beta &= \mathbb{P}(\varphi(X+v) - \varphi(X) < h) \\ &= \mathbb{P}(\varphi(X+v) - \varphi(X) - \frac{1}{2}v^T \mathcal{I}_\varphi v < hh - \frac{1}{2}v^T \mathcal{I}_\varphi v) \\ &= F_v(h - \frac{1}{2}v^T \mathcal{I}_\varphi v)\end{aligned}$$

From the expression of α we can solve for h :

$$h = -F_{-v}^{-1}(\alpha) - \frac{1}{2}v^T \mathcal{I}v$$

Plugging this into the expression of β yields

$$\begin{aligned}\beta &= F_v(h - \frac{1}{2}v^T \mathcal{I}_\varphi v) \\ &= F_v(-F_{-v}^{-1}(\alpha) - v^T \mathcal{I}v)\end{aligned}$$

The ROC function maps α to the minimal β , so this is exactly the expression of $T[X, X+v]$. \square

Proof of Lemma B.3. We know that $G_\mu(x) = \Phi(-\Phi^{-1}(x) - \mu)$. It suffices to show that $1 - G_\mu(x) = \Phi(\Phi^{-1}(x) + \mu)$ is α -Hölder continuous for any $\alpha < 1$.

Lemma B.4. Consider $f, g : [0, 1] \rightarrow \mathbb{R}$. Suppose f, g and $g - f$ is monotone increasing and g is α -Hölder continuous, then f is also α -Hölder continuous.

To see this, notice that by monotonicity of $g - f$, we have $g(x) - f(x) \leq g(y) - f(y)$ for $x < y$. Hence

$$|f(y) - f(x)| = f(y) - f(x) \leq g(y) - g(x) \leq Cx^\alpha.$$

Lemma B.5. For each $\alpha < 1$, there is an $\varepsilon = \varepsilon(\alpha) > 0$ such that $x^\alpha - (1 - G_\mu(x))$ is monotone increasing in $[0, \varepsilon]$.

Let $\alpha = 1 - \delta$. $h(x) = x^\alpha - (1 - G_\mu(x)) = x^\alpha - \Phi(\Phi^{-1}(x) + \mu)$. Let $y = \Phi^{-1}(x)$. Then $x = \Phi(y)$

$$\begin{aligned}h'(x) &= \alpha x^{\alpha-1} - \frac{d}{dx} \Phi(y + \mu) \\ &= \alpha x^{-\delta} - \phi(y + \mu) \cdot \frac{dy}{dx} \\ &= \alpha x^{-\delta} - \phi(y + \mu) / \frac{dx}{dy} \\ &= \alpha x^{-\delta} - \frac{\phi(y + \mu)}{\phi(y)} \\ &= \alpha x^{-\delta} - e^{-\mu y - \frac{1}{2}\mu^2} \\ &= e^{\delta \log x^{-1} + \log \alpha} - e^{-\mu y - \frac{1}{2}\mu^2}\end{aligned}$$

It is known that $|\Phi^{-1}(x)| = -\Phi^{-1}(x) \leq \sqrt{2 \log x^{-1}}$. So for fixed α, δ and μ , there is an ε such that when $x \in [0, \varepsilon]$, we have

$$\delta \log x^{-1} + \log \alpha \geq \mu \Phi^{-1}(x) - \frac{1}{2}\mu^2 = -\mu y - \frac{1}{2}\mu^2$$

Hence,

$$h'(x) = e^{\delta \log x^{-1} + \log \alpha} - e^{-\mu y - \frac{1}{2}\mu^2} \geq 0$$

\square

Interestingly, this implies the following result:

Proposition B.6. For each $\alpha \in [0, 1)$, there is a $C > 0$ such that

$$\int_{a+1}^{b+1} e^{-x^2} dx \leq C \left(\int_a^b e^{-x^2} dx \right)^\alpha.$$

For a convex φ such that $e^{-\varphi}$ is integrable, let F_v^φ be the cdf of $P_v^\varphi(X_\varphi)$. Dropping the unnecessary subscripts and superscripts of φ , we have

Appendix C Proof of Lemma 4.2

The major goal of this section is the following extended version of Lemma 4.2. We proceed by first presenting some general results in Appendix C.1, followed by calculation for norm powers in Appendix C.2.

Lemma C.1. For $1 \leq p < \infty$ and $1 \leq \alpha < \infty$, as $n \rightarrow \infty$, we have

$$\begin{aligned} \mathbb{E} \|X_{p,\alpha}\|_2^2 &= \frac{\Gamma(\frac{n}{\alpha} + 1 + \frac{2}{\alpha})}{\Gamma(\frac{n}{p} + 1 + \frac{2}{p})} \cdot \frac{\Gamma(\frac{n}{p} + 1)}{\Gamma(\frac{n}{\alpha} + 1)} \cdot n \cdot \frac{\Gamma(\frac{3}{p})}{\Gamma(\frac{1}{p})} \\ &\sim n^{\frac{2}{\alpha} - \frac{2}{p} + 1} \cdot \alpha^{-\frac{2}{\alpha}} \cdot p^{\frac{2}{p}} \cdot \frac{\Gamma(\frac{3}{p})}{\Gamma(\frac{1}{p})} \\ \mathcal{I}_{p,\alpha} &= \alpha^2 \cdot \frac{\Gamma(\frac{n+2\alpha-2}{\alpha})}{\Gamma(\frac{n+2p-2}{p})} \cdot \frac{\Gamma(\frac{n}{p})}{\Gamma(\frac{n}{\alpha})} \cdot \frac{\Gamma(2 - \frac{1}{p})}{\Gamma(\frac{1}{p})} \cdot I_{n \times n} \\ &\sim n^{\frac{2}{p} - \frac{2}{\alpha}} \cdot \alpha^{\frac{2}{\alpha}} \cdot p^{2 - \frac{2}{p}} \cdot \frac{\Gamma(2 - \frac{1}{p})}{\Gamma(\frac{1}{p})} \cdot I_{n \times n}. \end{aligned}$$

C.1 Regarding Homogeneous φ

In this section, in addition to that $\varphi \in \mathfrak{F}_n$, we further assume that $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$ is positively homogeneous. Recall that φ is (positively) homogeneous of degree $\alpha > 0$ if $\varphi(tx) = |t|^\alpha \varphi(x)$ for $t \in \mathbb{R}, x \in \mathbb{R}^n$. This implies $\varphi(0) = 0$.

The first result takes care of the normalizer Z_φ defined as $\int e^{-\varphi(x)} dx$.

Lemma C.2. $Z_\varphi = \Gamma(\frac{n}{\alpha} + 1) \cdot \text{vol}(K_\varphi)$ where $K_\varphi = \{x : \varphi(x) \leq 1\}$.

Proof of Lemma C.2. We use polar coordinate. For any function $f : \mathbb{R}^n \rightarrow \mathbb{R}$,

$$\int_{\mathbb{R}^n} f(x) dx = \int_0^\infty \int_{S^{n-1}} f(r\theta) r^{n-1} d\theta dr$$

So

$$\begin{aligned} Z &= \int_{\mathbb{R}^n} e^{-\varphi(x)} dx = \int_{S^{n-1}} \int_0^\infty e^{-\varphi(r\theta)} r^{n-1} dr d\theta \\ &= \int_{S^{n-1}} \int_0^\infty e^{-r^\alpha \varphi(\theta)} r^{n-1} dr d\theta \end{aligned}$$

Let $t = r^\alpha, r = t^{\frac{1}{\alpha}}, dr = \frac{1}{\alpha} t^{\frac{1}{\alpha}-1} dt$.

$$\begin{aligned} \int_0^\infty e^{-r^\alpha \varphi(\theta)} r^{n-1} dr &= \int_0^\infty e^{-t\varphi(\theta)} \cdot t^{\frac{n-1}{\alpha}} \cdot \frac{1}{\alpha} t^{\frac{1}{\alpha}-1} dt \\ &= \frac{1}{\alpha} \int_0^\infty e^{-t\varphi(\theta)} \cdot t^{\frac{n}{\alpha}-1} dt \\ &= \frac{1}{\alpha} \cdot \frac{\Gamma(\frac{n}{\alpha})}{\varphi(\theta)^{\frac{n}{\alpha}}} \end{aligned}$$

So

$$Z = \int_{S^{n-1}} \frac{1}{\alpha} \cdot \frac{\Gamma(\frac{n}{\alpha})}{\varphi(\theta)^{\frac{n}{\alpha}}} d\theta$$

On the other hand, consider a set defined with polar coordinate:

$$K := \{(r, \theta) : r \leq \rho(\theta)\}.$$

Its volume is

$$\begin{aligned} \text{vol}(K) &= \int_{\mathbb{R}^n} 1_K(x) dx \\ &= \int_{S^{n-1}} \int_0^{\rho(\theta)} r^{n-1} dr d\theta \\ &= \frac{1}{n} \int_{S^{n-1}} \rho(\theta)^n d\theta \end{aligned}$$

We see that

$$\begin{aligned} Z &= \int_{S^{n-1}} \frac{1}{\alpha} \cdot \frac{\Gamma(\frac{n}{\alpha})}{\varphi(\theta)^{\frac{n}{\alpha}}} d\theta \\ &= \frac{\Gamma(\frac{n}{\alpha})}{\alpha} \cdot n \cdot \frac{1}{n} \int_{S^{n-1}} [\varphi(\theta)^{-\frac{1}{\alpha}}]^n d\theta \\ &= \frac{n}{\alpha} \cdot \Gamma(\frac{n}{\alpha}) \cdot \text{vol}(K_\varphi) \end{aligned}$$

where

$K_\varphi = \{(r, \theta) : r \leq \varphi(\theta)^{-\frac{1}{\alpha}}\} = \{(r, \theta) : r^\alpha \varphi(\theta) \leq 1\} = \{(r, \theta) : \varphi(r\theta) \leq 1\} = \{x : \varphi(x) \leq 1\}$.
Noticing $\Gamma(z+1) = z\Gamma(z)$, we have

$$Z = \Gamma(\frac{n}{\alpha} + 1) \cdot \text{vol}(K_\varphi)$$

□

Lemma C.3. The m -th moment of a $\Gamma(k, 1)$ distribution is $\frac{\Gamma(m+k)}{\Gamma(k)}$.

Proof of Lemma C.3. The m -th moment of a $\Gamma(k, 1)$ distribution is

$$\frac{1}{\Gamma(k)} \int_0^\infty x^m \cdot x^{k-1} e^{-x} dx = \frac{\Gamma(m+k)}{\Gamma(k)}.$$

□

The following result also appears in [Wan05].

Lemma C.4.

$$\text{vol}(K_p) = 2^n \cdot \frac{\Gamma(\frac{1}{p} + 1)^n}{\Gamma(\frac{n}{p} + 1)}.$$

Proof of Lemma C.4. By Lemma C.2, we have

$$\begin{aligned} \text{vol}(K_p) &= \frac{1}{\Gamma(\frac{n}{p} + 1)} \cdot \int e^{-\sum |x_i|^p} dx \\ &= \frac{1}{\Gamma(\frac{n}{p} + 1)} \cdot \left(\int_{-\infty}^{+\infty} e^{-|x|^p} dx \right)^n \\ \int_{-\infty}^{+\infty} e^{-|x|^p} dx &= 2 \int_0^{+\infty} e^{-x^p} dx \\ &= 2 \int_0^{+\infty} e^{-y} dy^{\frac{1}{p}} \\ &= \frac{2}{p} \int_0^{+\infty} e^{-y} y^{\frac{1}{p}-1} dy \\ &= \frac{2}{p} \Gamma(\frac{1}{p}) = 2\Gamma(\frac{1}{p} + 1) \end{aligned}$$

□

Lemma C.5. Let $t \sim \Gamma(\frac{n}{\alpha} + 1, 1)$. Let V_φ has uniform distribution over K_φ where $K_\varphi = \{x : \varphi(x) \leq 1\}$ independently from t . Then $t^{\frac{1}{\alpha}} \cdot V$ has density $\frac{1}{Z} e^{-\varphi}$.

Proof of Lemma C.5. We use a more principled way: assume r has density $p(r)$ over $(0, +\infty)$ and rV has density $\frac{1}{Z} e^{-\varphi}$, find $p(r)$. Let B be a small ball.

$$\begin{aligned} \mathbb{P}(rV \in x + B) &= \int_0^\infty \mathbb{P}(V \in \frac{x+B}{r}) \cdot p(r) \, dr \\ \mathbb{P}(V \in \frac{x+B}{r}) &= \begin{cases} 0, & \text{if } \frac{x}{r} \notin K_\varphi \\ \frac{\text{vol}(\frac{B}{r})}{\text{vol}(K_\varphi)}, & \text{if } \frac{x}{r} \in K_\varphi. \end{cases} \end{aligned}$$

$\frac{x}{r} \in K_\varphi \Leftrightarrow \varphi(\frac{x}{r}) \leq 1 \Leftrightarrow \varphi(x) \leq r^\alpha \Leftrightarrow r \geq \varphi(x)^{1/\alpha}$. So

$$\begin{aligned} \mathbb{P}(rV \in x + B) &= \int_0^\infty \mathbb{P}(V \in \frac{x+B}{r}) \cdot p(r) \, dr \\ &= \int_{\varphi(x)^{1/\alpha}}^\infty \frac{\text{vol}(\frac{B}{r})}{\text{vol}(K_\varphi)} \cdot p(r) \, dr \\ &= \text{vol}(B) \cdot \int_{\varphi(x)^{1/\alpha}}^\infty \frac{p(r)r^{-n}}{\text{vol}(K_\varphi)} \, dr. \end{aligned}$$

So the density of rV at x is $\int_{\varphi(x)^{1/\alpha}}^\infty \frac{p(r)r^{-n}}{\text{vol}(K_\varphi)} \, dr$. In order to match it with $\frac{1}{Z} e^{-\varphi}$, we have

$$\begin{aligned} \int_{\varphi(x)^{1/\alpha}}^\infty \frac{p(r)r^{-n}}{\text{vol}(K_\varphi)} \, dr &= \frac{e^{-\varphi(x)}}{\Gamma(\frac{n}{\alpha} + 1) \cdot \text{vol}(K_\varphi)} \\ \int_{\varphi(x)^{1/\alpha}}^\infty p(r)r^{-n} \, dr &= \frac{e^{-\varphi(x)}}{\Gamma(\frac{n}{\alpha} + 1)} \end{aligned}$$

Let $\varphi(x)^{1/\alpha} = u$, we have

$$\int_u^\infty p(r)r^{-n} \, dr = \frac{1}{\Gamma(\frac{n}{\alpha} + 1)} \cdot e^{-u^\alpha}.$$

Taking derivative with respect to u , we have

$$p(u)u^{-n} = \frac{1}{\Gamma(\frac{n}{\alpha} + 1)} \cdot e^{-u^\alpha} \cdot \alpha u^{\alpha-1}.$$

It's straightforward to show that if $t \sim \Gamma(\frac{n}{\alpha} + 1, 1)$, then $t^{\frac{1}{\alpha}}$ has the above density $p(u)$. \square

A simple but useful corollary of Lemma C.5 is

Corollary C.6. Let $t \sim \Gamma(\frac{n}{\alpha} + 1, 1)$. Let U has uniform distribution over ∂K_φ , independently from t and r with density nx^{n-1} over $[0, 1]$, independent from t and U . Then $t^{\frac{1}{\alpha}} \cdot r \cdot U$ has density $\frac{1}{Z} e^{-\varphi}$.

We commented that we can compute the isotropic constants for ℓ_p balls. The rest of the section is dedicated to this kind of results.

Let μ be a log-concave probability measure on \mathbb{R}^n with density $f_\mu : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$. The isotropic constant of μ is defined by (see e.g. [Gia])

$$L_\mu = \left(\sup_{x \in \mathbb{R}^n} f_\mu(x) \right)^{\frac{1}{n}} \cdot \left(\det \text{Cov}(\mu) \right)^{\frac{1}{2n}}.$$

As a special case, when μ is the uniform distribution over the convex body K , the corresponding isotropic constant is denoted by L_K and has expression

$$L_K = \text{vol}(K)^{-\frac{1}{n}} \cdot \left(\det \text{Cov}(\mu) \right)^{\frac{1}{2n}}.$$

For homogeneous and convex $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$, we use L_φ to denote the isotropic constant of its associated probability distribution, i.e. the one with density $\frac{1}{Z_\varphi} e^{-\varphi(x)}$. With the help of Lemma C.5, we can relate L_φ to the isotropic constant of its unit ball L_{K_φ} .

Lemma C.7.

$$L_\varphi = \frac{[\Gamma(\frac{n}{\alpha} + 1 + \frac{2}{\alpha})]^{\frac{1}{2}}}{[\Gamma(\frac{n}{\alpha} + 1)]^{\frac{1}{2} + \frac{1}{n}}} \cdot L_{K_\varphi}$$

Proof.

$$\begin{aligned} \text{Cov}(X_\varphi) &= \mathbb{E}[X_\varphi X_\varphi^T] = \mathbb{E}[t^{\frac{2}{\alpha}} \cdot V_\varphi V_\varphi^T] = \mathbb{E}t^{\frac{2}{\alpha}} \cdot \mathbb{E}[V_\varphi V_\varphi^T] = \mathbb{E}t^{\frac{2}{\alpha}} \cdot \text{Cov}(V_\varphi) \\ \det \text{Cov}(X_\varphi) &= (\mathbb{E}t^{\frac{2}{\alpha}})^n \cdot \det \text{Cov}(V_\varphi) \end{aligned}$$

$$\begin{aligned} L_\varphi &= Z_\varphi^{-\frac{1}{n}} \cdot (\det \text{Cov}(X_\varphi))^{\frac{1}{2n}} \\ &= Z_\varphi^{-\frac{1}{n}} \cdot (\mathbb{E}t^{\frac{2}{\alpha}})^{\frac{1}{2}} \cdot (\det \text{Cov}(V_\varphi))^{\frac{1}{2n}} \\ &= \Gamma(\frac{n}{\alpha} + 1)^{-\frac{1}{n}} \cdot \text{vol}(K_\varphi)^{-\frac{1}{n}} \cdot (\mathbb{E}t^{\frac{2}{\alpha}})^{\frac{1}{2}} \cdot (\det \text{Cov}(V_\varphi))^{\frac{1}{2n}} \quad (\text{Lemma C.2}) \\ &= (\mathbb{E}t^{\frac{2}{\alpha}})^{\frac{1}{2}} \cdot \Gamma(\frac{n}{\alpha} + 1)^{-\frac{1}{n}} \cdot L_{K_\varphi} \end{aligned}$$

By Lemma C.3, $\mathbb{E}t^{\frac{2}{\alpha}} = \frac{\Gamma(\frac{n}{\alpha} + 1 + \frac{2}{\alpha})}{\Gamma(\frac{n}{\alpha} + 1)}$. So

$$\begin{aligned} L_\varphi &= \left(\frac{\Gamma(\frac{n}{\alpha} + 1 + \frac{2}{\alpha})}{\Gamma(\frac{n}{\alpha} + 1)} \right)^{\frac{1}{2}} \cdot \Gamma(\frac{n}{\alpha} + 1)^{-\frac{1}{n}} \cdot L_{K_\varphi} \\ &= \frac{[\Gamma(\frac{n}{\alpha} + 1 + \frac{2}{\alpha})]^{\frac{1}{2}}}{[\Gamma(\frac{n}{\alpha} + 1)]^{\frac{1}{2} + \frac{1}{n}}} \cdot L_{K_\varphi} \end{aligned}$$

□

The last result is a sufficient condition of the Fisher information being a scalar matrix.

Lemma C.8. *If φ is invariant under the action of $\{\pm 1\}^n$ and cyclic group of size n , i.e.*

1. $\varphi(\pm x_1, \pm x_2, \dots, \pm x_n) = \varphi(x_1, x_2, \dots, x_n)$
2. $\varphi(x_1, x_2, \dots, x_{n-1}, x_n) = \varphi(x_2, x_3, \dots, x_n, x_1)$

then $\mathcal{I}_\varphi = \frac{1}{n} \mathbb{E} \|\nabla \varphi\|_2^2 \cdot I$.

Proof of Lemma C.8. First we use the symmetry to show $\mathcal{I}_\varphi = cI$ for some c .

$$\begin{aligned} \varphi(x_1, x_2, \dots, x_n) &= \varphi(-x_1, x_2, \dots, x_n) \\ \partial_1 \varphi(x_1, x_2, \dots, x_n) &= -\partial_1 \varphi(-x_1, x_2, \dots, x_n) \\ \partial_2 \varphi(x_1, x_2, \dots, x_n) &= \partial_2 \varphi(-x_1, x_2, \dots, x_n) \end{aligned}$$

This shows $\partial_1 \varphi \cdot \partial_2 \varphi$ is an odd function of x_1 . On the other hand, we know the density $e^{-\varphi}$ is an even function of x_1 . So we conclude that $\mathbb{E}[\partial_1 \varphi \cdot \partial_2 \varphi] = 0$. Similarly, we can show that for any $i \neq j$, $\mathbb{E}[\partial_i \varphi \cdot \partial_j \varphi] = 0$. This shows that \mathcal{I}_φ is a diagonal matrix.

By cyclic symmetry, we have

$$\begin{aligned} \varphi(x_1, x_2, \dots, x_{n-1}, x_n) &= \varphi(x_2, x_3, \dots, x_n, x_1) \\ \partial_1 \varphi(x_1, x_2, \dots, x_{n-1}, x_n) &= \partial_n \varphi(x_2, x_3, \dots, x_n, x_1) \\ \partial_1 \varphi(x_1, x_2, \dots, x_{n-1}, x_n)^2 e^{-\varphi(x_1, x_2, \dots, x_{n-1}, x_n)} &= \partial_n \varphi(x_2, x_3, \dots, x_n, x_1)^2 e^{-\varphi(x_1, x_2, \dots, x_{n-1}, x_n)} \\ &= \partial_n \varphi(y_1, y_2, \dots, y_{n-1}, y_n)^2 e^{-\varphi(y_1, y_2, \dots, y_{n-1}, y_n)} \\ &= \partial_n \varphi(y_1, y_2, \dots, y_{n-1}, y_n)^2 e^{-\varphi(y_1, y_2, \dots, y_n)} \end{aligned}$$

This shows $\mathbb{E} \partial_1 \varphi^2 = \dots = \mathbb{E} \partial_n \varphi^2$. Hence $\mathcal{I}_\varphi = cI$ for some c .

$$\text{Tr } \mathcal{I}_\varphi = \text{Tr } \mathbb{E}[\nabla \varphi \nabla \varphi^T] = \mathbb{E}[\text{Tr } \nabla \varphi \nabla \varphi^T] = \mathbb{E}[\text{Tr } \nabla \varphi^T \nabla \varphi] = \mathbb{E} \|\nabla \varphi\|_2^2.$$

On the other hand, $\text{Tr } \mathcal{I}_\varphi = \text{Tr } cI = cn$, so $c = \frac{1}{n} \mathbb{E} \|\nabla \varphi\|_2^2$. □

C.2 Calculation for Norm Powers

Lemma 4.2. For $1 \leq p < \infty$ and $1 \leq \alpha < \infty$, as $n \rightarrow \infty$, we have

$$\begin{aligned}\mathbb{E}\|X_{p,\alpha}\|_2^2 &\sim n^{\frac{2}{\alpha}-\frac{2}{p}+1} \cdot \alpha^{-\frac{2}{\alpha}} \cdot p^{\frac{2}{p}} \cdot \Gamma(\frac{3}{p})/\Gamma(\frac{1}{p}); \\ \mathcal{I}_{p,\alpha} &\sim n^{\frac{2}{p}-\frac{2}{\alpha}} \cdot \alpha^{\frac{2}{\alpha}} \cdot p^{2-\frac{2}{p}} \cdot \Gamma(2-\frac{1}{p})/\Gamma(\frac{1}{p}) \cdot I_{n \times n}.\end{aligned}$$

We divide the proof into two parts, one for each of the equations.

Proof of Lemma 4.2 (variance part). By Lemma C.5, let $t \sim \Gamma(\frac{n}{\alpha} + 1, 1)$ random variable and V_p has uniform distribution over the ℓ_p unit ball K_p .

$$\mathbb{E}\|X_{p,\alpha}\|_2^2 = \mathbb{E}t^{\frac{2}{\alpha}} \cdot \mathbb{E}\|V_p\|_2^2 = \frac{\Gamma(\frac{n}{\alpha} + 1 + \frac{2}{\alpha})}{\Gamma(\frac{n}{\alpha} + 1)} \cdot \mathbb{E}\|V_p\|_2^2 \quad (13)$$

Setting $\alpha = p$ yields

$$\mathbb{E}\|X_{p,p}\|_2^2 = \frac{\Gamma(\frac{n}{p} + 1 + \frac{2}{p})}{\Gamma(\frac{n}{p} + 1)} \cdot \mathbb{E}\|V_p\|_2^2 \quad (14)$$

The reason we do this is that $\mathbb{E}\|X_{p,p}\|_2^2$ can be computed explicitly. In fact,

$$\begin{aligned}\mathbb{E}\|X_{p,p}\|_2^2 &= \frac{1}{Z_n} \int \sum x_i^2 \cdot e^{-\sum |x_i|^p} dx \\ &= \frac{n}{Z_n} \int x_1^2 \cdot e^{-\sum |x_i|^p} dx \\ &= \frac{n}{Z_n} \int x_1^2 \cdot e^{-|x_1|^p} dx_1 \cdot Z_{n-1}\end{aligned}$$

where $Z_n = \int e^{-\sum |x_i|^p} dx$. We know by Lemma C.4 that

$$Z_n = \Gamma(\frac{n}{p} + 1) \cdot \text{vol}(K_p) = 2^n \cdot \Gamma(\frac{1}{p} + 1)^n$$

and

$$\int_{-\infty}^{+\infty} x^2 \cdot e^{-|x|^p} dx = 2 \int_0^{\infty} x^2 \cdot e^{-x^p} dx = 2 \int_0^{\infty} y^{\frac{2}{p}} \cdot e^{-y} \cdot \frac{1}{p} y^{\frac{1}{p}-1} dy = \frac{2}{p} \int_0^{\infty} y^{\frac{3}{p}-1} \cdot e^{-y} dy = \frac{2}{p} \cdot \Gamma(\frac{3}{p})$$

So

$$\begin{aligned}\mathbb{E}\|X_{p,p}\|_2^2 &= \frac{n}{Z_n} \int x_1^2 \cdot e^{-|x_1|^p} dx_1 \cdot Z_{n-1} \\ &= n \cdot \frac{2}{p} \cdot \Gamma(\frac{3}{p}) \cdot \frac{Z_{n-1}}{Z_n} \\ &= \frac{n \cdot \frac{2}{p} \cdot \Gamma(\frac{3}{p})}{2\Gamma(\frac{1}{p} + 1)} \\ &= \frac{n \cdot \frac{2}{p} \cdot \Gamma(\frac{3}{p})}{\frac{2}{p}\Gamma(\frac{1}{p})} = n \cdot \frac{\Gamma(\frac{3}{p})}{\Gamma(\frac{1}{p})}\end{aligned}$$

Plugging this into (14), we have

$$\begin{aligned}\mathbb{E}\|V_p\|_2^2 &= \frac{\Gamma(\frac{n}{p} + 1)}{\Gamma(\frac{n}{p} + 1 + \frac{2}{p})} \cdot \mathbb{E}\|X_{p,p}\|_2^2 \\ &= \frac{\Gamma(\frac{n}{p} + 1)}{\Gamma(\frac{n}{p} + 1 + \frac{2}{p})} \cdot n \cdot \frac{\Gamma(\frac{3}{p})}{\Gamma(\frac{1}{p})}.\end{aligned}$$

Using this in (13),

$$\begin{aligned}\mathbb{E}\|X_{p,\alpha}\|_2^2 &= \frac{\Gamma(\frac{n}{\alpha} + 1 + \frac{2}{\alpha})}{\Gamma(\frac{n}{\alpha} + 1)} \cdot \mathbb{E}\|V_p\|_2^2 \\ &= \frac{\Gamma(\frac{n}{\alpha} + 1 + \frac{2}{\alpha})}{\Gamma(\frac{n}{\alpha} + 1)} \cdot \frac{\Gamma(\frac{n}{p} + 1)}{\Gamma(\frac{n}{p} + 1 + \frac{2}{p})} \cdot n \cdot \frac{\Gamma(\frac{3}{p})}{\Gamma(\frac{1}{p})}\end{aligned}$$

In order to study the asymptotics of $\mathbb{E}\|X_{p,\alpha}\|_2^2$ as $n \rightarrow \infty$, recall Stirling's formula

$$\Gamma(z + 1) \sim \sqrt{2\pi z} \left(\frac{z}{e}\right)^z.$$

So we have

$$\begin{aligned}\frac{\Gamma(\frac{n}{\alpha} + 1 + \frac{2}{\alpha})}{\Gamma(\frac{n}{\alpha} + 1)} &\sim \frac{\sqrt{\frac{n+2}{\alpha}} \cdot (\frac{n+2}{\alpha e})^{\frac{n+2}{\alpha}}}{\sqrt{\frac{n}{\alpha}} \cdot (\frac{n}{\alpha e})^{\frac{n}{\alpha}}} \\ &\sim \left(\frac{n+2}{\alpha e} \cdot \frac{\alpha e}{n}\right)^{\frac{n}{\alpha}} \cdot \left(\frac{n+2}{\alpha e}\right)^{\frac{2}{\alpha}} \\ &= \left(1 + \frac{2}{n}\right)^{\frac{n}{\alpha}} \cdot \left(\frac{n+2}{\alpha e}\right)^{\frac{2}{\alpha}} \\ &\sim \left(\frac{n+2}{\alpha}\right)^{\frac{2}{\alpha}}\end{aligned}$$

Hence

$$\mathbb{E}\|X_{p,\alpha}\|_2^2 \sim \left(\frac{n+2}{\alpha}\right)^{\frac{2}{\alpha}} / \left(\frac{n+2}{p}\right)^{\frac{2}{p}} \cdot n \cdot \frac{\Gamma(\frac{3}{p})}{\Gamma(\frac{1}{p})} \sim n^{\frac{2}{\alpha} - \frac{2}{p} + 1} \cdot \alpha^{-\frac{2}{\alpha}} \cdot p^{\frac{2}{p}} \cdot \frac{\Gamma(\frac{3}{p})}{\Gamma(\frac{1}{p})}$$

□

Before we proceed to the proof of the Fisher information part of Lemma 4.2, we derive the isotropic constants results as promised, using Lemma C.7 and the variance part of Lemma 4.2.

Corollary C.9. *The isotropic constant of n -dimensional ℓ_p ball is*

$$L_{K_p}^2 = \frac{p^2}{4} \cdot \frac{\Gamma(\frac{3}{p})}{[\Gamma(\frac{1}{p})]^3} \cdot \frac{[\Gamma(\frac{n}{p} + 1)]^{1 + \frac{2}{n}}}{\Gamma(\frac{n}{p} + 1 + \frac{2}{p})}$$

Proof of Corollary C.9. For a general convex body K , let V_K be a random vector with the uniform distribution over K . Recall that the isotropic constant of K is

$$L_K = \text{vol}(K)^{-\frac{1}{n}} \cdot (\det \text{Cov}(V_K))^{\frac{1}{2n}}.$$

Now we focus on the unit ball of ℓ_p norm K_p . The corresponding random vector is denoted by V_p . By a symmetry argument similar to Lemma C.8, we have that

$$\text{Cov}(V_p) = \frac{1}{n} \cdot \mathbb{E}\|V_p\|_2^2 \cdot I_{n \times n}.$$

Combining this and Lemma C.4, we have

$$\begin{aligned}
L_{K_p}^2 &= \text{vol}(K_p)^{-\frac{2}{n}} \cdot \frac{1}{n} \mathbb{E} \|V_p\|_2^2 \\
&= \left(2^n \cdot \frac{\Gamma(\frac{1}{p} + 1)^n}{\Gamma(\frac{n}{p} + 1)} \right)^{-\frac{2}{n}} \cdot \frac{\Gamma(\frac{n}{p} + 1)}{\Gamma(\frac{n}{p} + 1 + \frac{2}{p})} \cdot \frac{\Gamma(\frac{3}{p})}{\Gamma(\frac{1}{p})} \\
&= \frac{1}{4} \cdot \frac{\Gamma(\frac{n}{p} + 1)^{\frac{2}{n}}}{\Gamma(\frac{1}{p} + 1)^2} \cdot \frac{\Gamma(\frac{n}{p} + 1)}{\Gamma(\frac{n}{p} + 1 + \frac{2}{p})} \cdot \frac{\Gamma(\frac{3}{p})}{\Gamma(\frac{1}{p})} \\
&= \frac{1}{4} \cdot \frac{\Gamma(\frac{n}{p} + 1)^{\frac{2}{n}}}{\frac{1}{p^2} \Gamma(\frac{1}{p})^2} \cdot \frac{\Gamma(\frac{n}{p} + 1)}{\Gamma(\frac{n}{p} + 1 + \frac{2}{p})} \cdot \frac{\Gamma(\frac{3}{p})}{\Gamma(\frac{1}{p})} \\
&= \frac{p^2}{4} \cdot \frac{\Gamma(\frac{3}{p})}{[\Gamma(\frac{1}{p})]^3} \cdot \frac{[\Gamma(\frac{n}{p} + 1)]^{1 + \frac{2}{n}}}{\Gamma(\frac{n}{p} + 1 + \frac{2}{p})}
\end{aligned}$$

□

Corollary C.10. When $\varphi(x) = \|x\|_p^\alpha$,

$$L_{p,\alpha}^2 = \frac{p^2}{4} \cdot \frac{\Gamma(\frac{3}{p})}{[\Gamma(\frac{1}{p})]^3} \cdot \frac{\Gamma(\frac{n}{\alpha} + 1 + \frac{2}{\alpha})}{\Gamma(\frac{n}{p} + 1 + \frac{2}{p})} \cdot \left(\frac{\Gamma(\frac{n}{p} + 1)}{\Gamma(\frac{n}{\alpha} + 1)} \right)^{1 + \frac{2}{n}}$$

Proof of Corollary C.10. Directly follows from the above result and Lemma C.7. □

Now we turn our attention back to the proof of Lemma 4.2.

Proof of Lemma 4.2 (variance part). By Lemma C.8, $\mathcal{I}_{p,\alpha} = \|\mathcal{I}_{p,\alpha}\|_2 \cdot I_{n \times n}$ and $\|\mathcal{I}_{p,\alpha}\|_2 = \frac{1}{n} \mathbb{E} \|\nabla \varphi\|_2^2$ where $\varphi(x) = \|x\|_p^\alpha$. In this case the gradient has an explicit expression:

$$\begin{aligned}
[\nabla \varphi(x)]_i &= \alpha \cdot \left(\sum |x_i|^p \right)^{\frac{\alpha}{p} - 1} \cdot |x_i|^{p-1} \cdot \text{sgn}(x_i) = \alpha \|x\|_p^{\alpha-p} \cdot |x_i|^{p-1} \cdot \text{sgn}(x_i) \\
\|\nabla \varphi(x)\|_2^2 &= \alpha^2 \sum |x_i|^{2p-2} \cdot \left(\sum |x_i|^p \right)^{\frac{2\alpha}{p} - 2} = \alpha^2 \|x\|_{2p-2}^{2p-2} \cdot \|x\|_p^{2\alpha-2p}
\end{aligned}$$

By Corollary C.6, $X_{p,\alpha} \stackrel{d}{=} t^{\frac{1}{\alpha}} \cdot r \cdot U_p$, where $t \sim \Gamma(\frac{n}{\alpha} + 1, 1)$, U has uniform distribution over ∂K_p and r has density $n x^{n-1}$ over $[0, 1]$.

$$\begin{aligned}
\|\nabla \varphi\|_2^2 &= \alpha^2 \|t^{\frac{1}{\alpha}} r U\|_{2p-2}^{2p-2} \cdot \|t^{\frac{1}{\alpha}} r U\|_p^{2\alpha-2p} \\
&= \alpha^2 t^{\frac{2\alpha-2}{\alpha}} r^{2\alpha-2} \cdot \|U\|_{2p-2}^{2p-2} \cdot \|U\|_p^{2\alpha-2p}
\end{aligned}$$

Since t, r, U_p are independent and $\|U_p\|_p = 1$, we have

$$\|\mathcal{I}_{p,\alpha}\|_2 = \frac{1}{n} \mathbb{E} \|\nabla \varphi\|_2^2 = \frac{\alpha^2}{n} \cdot \mathbb{E} t^{\frac{2\alpha-2}{\alpha}} \cdot \mathbb{E} r^{2\alpha-2} \cdot \mathbb{E} \|U\|_{2p-2}^{2p-2} \quad (15)$$

By Lemma C.3,

$$\mathbb{E} t^{\frac{2\alpha-2}{\alpha}} = \frac{\Gamma(\frac{n}{\alpha} + 1 + \frac{2\alpha-2}{\alpha})}{\Gamma(\frac{n}{\alpha} + 1)} = \frac{\Gamma(\frac{n+2\alpha-2}{\alpha} + 1)}{\Gamma(\frac{n}{\alpha} + 1)}.$$

The moment of r can be computed directly

$$\mathbb{E} r^{2\alpha-2} = \int_0^1 r^{2\alpha-2} \cdot n r^{n-1} dr = \frac{n}{n + 2\alpha - 2}.$$

Plugging into (15), we have

$$\begin{aligned}
\|\mathcal{I}_{p,\alpha}\|_2 &= \frac{\alpha^2}{n} \cdot \mathbb{E}t^{\frac{2\alpha-2}{\alpha}} \cdot \mathbb{E}r^{2\alpha-2} \cdot \mathbb{E}\|U_p\|_{2p-2}^{2p-2} \\
&= \frac{\alpha^2}{n} \cdot \frac{\Gamma(\frac{n+2\alpha-2}{\alpha} + 1)}{\Gamma(\frac{n}{\alpha} + 1)} \cdot \frac{n}{n+2\alpha-2} \cdot \mathbb{E}\|U_p\|_{2p-2}^{2p-2} \\
&= \frac{\alpha^2}{n} \cdot \frac{\frac{n+2\alpha-2}{\alpha} \cdot \Gamma(\frac{n+2\alpha-2}{\alpha})}{\frac{n}{\alpha} \cdot \Gamma(\frac{n}{\alpha})} \cdot \frac{n}{n+2\alpha-2} \cdot \mathbb{E}\|U_p\|_{2p-2}^{2p-2} \\
&= \frac{\alpha^2}{n} \cdot \frac{\Gamma(\frac{n+2\alpha-2}{\alpha})}{\Gamma(\frac{n}{\alpha})} \cdot \mathbb{E}\|U_p\|_{2p-2}^{2p-2} \tag{*}
\end{aligned}$$

On the other hand, when $p = \alpha$, we have

$$\|\nabla\varphi(x)\|_2^2 = p^2 \sum |x_i|^{2p-2}.$$

In this case, x has joint density $\propto e^{-\|x\|_p^p} = e^{-\sum |x_i|^p}$. Let Y_1, Y_2, \dots, Y_n be i.i.d. random variables with density $\propto e^{-|y|^p}$. Then we have

$$\|\mathcal{I}_{p,p}\|_2 = \frac{1}{n} \mathbb{E}\|\nabla\varphi\|_2^2 = \frac{1}{n} \cdot p^2 \cdot \mathbb{E} \sum |Y_i|^{2p-2} = p^2 \mathbb{E}[|Y_i|^{2p-2}]$$

Let $z = \int_{-\infty}^{+\infty} e^{-|y|^p} dy$.

$$\begin{aligned}
\mathbb{E}[|Y_i|^{2p-2}] &= \frac{1}{z} \int_{-\infty}^{+\infty} |y|^{2p-2} \cdot e^{-|y|^p} dy \\
&= \frac{\int_0^{+\infty} y^{2p-2} \cdot e^{-y^p} dy}{\int_0^{+\infty} e^{-y^p} dy} \\
&= \frac{\int_0^{+\infty} x^{\frac{2p-2}{p}} \cdot x^{\frac{1}{p}-1} \cdot e^{-x} dx}{\int_0^{+\infty} x^{\frac{1}{p}-1} \cdot e^{-x} dx} \\
&= \frac{\Gamma(2 - \frac{1}{p})}{\Gamma(\frac{1}{p})}
\end{aligned}$$

Relating to (*) in the special case of $\alpha = p$,

$$\frac{p^2}{n} \cdot \frac{\Gamma(\frac{n+2p-2}{p})}{\Gamma(\frac{n}{p})} \cdot \mathbb{E}\|U_p\|_{2p-2}^{2p-2} = \|\mathcal{I}_{p,p}\|_2 = p^2 \cdot \frac{\Gamma(2 - \frac{1}{p})}{\Gamma(\frac{1}{p})}$$

Hence

$$\mathbb{E}\|U_p\|_{2p-2}^{2p-2} = \frac{n}{p^2} \cdot \frac{\Gamma(\frac{n}{p})}{\Gamma(\frac{n+2p-2}{p})} \cdot \|\mathcal{I}_{p,p}\|_2 = n \cdot \frac{\Gamma(\frac{n}{p})}{\Gamma(\frac{n+2p-2}{p})} \cdot \frac{\Gamma(2 - \frac{1}{p})}{\Gamma(\frac{1}{p})}$$

Using (*) again,

$$\begin{aligned}
\|\mathcal{I}_{p,\alpha}\|_2 &= \frac{\alpha^2}{n} \cdot \frac{\Gamma(\frac{n+2\alpha-2}{\alpha})}{\Gamma(\frac{n}{\alpha})} \cdot \mathbb{E}\|U_p\|_{2p-2}^{2p-2} \\
&= \frac{\alpha^2}{n} \cdot \frac{\Gamma(\frac{n+2\alpha-2}{\alpha})}{\Gamma(\frac{n}{\alpha})} \cdot n \cdot \frac{\Gamma(\frac{n}{p})}{\Gamma(\frac{n+2p-2}{p})} \cdot \frac{\Gamma(2 - \frac{1}{p})}{\Gamma(\frac{1}{p})} \\
&= \alpha^2 \cdot \frac{\Gamma(\frac{n+2\alpha-2}{\alpha})}{\Gamma(\frac{n}{\alpha})} \cdot \frac{\Gamma(\frac{n}{p})}{\Gamma(\frac{n+2p-2}{p})} \cdot \frac{\Gamma(2 - \frac{1}{p})}{\Gamma(\frac{1}{p})}
\end{aligned}$$

In order to study the asymptotics, using Stirling's formula again,

$$\begin{aligned}
\frac{\Gamma\left(\frac{2\alpha-2+n}{\alpha}\right)}{\Gamma\left(\frac{n}{\alpha}+1\right)} &\sim \frac{\sqrt{\frac{n+\alpha-2}{\alpha}} \cdot \left(\frac{n+\alpha-2}{\alpha e}\right)^{\frac{n+\alpha-2}{\alpha}}}{\sqrt{\frac{n}{\alpha}} \cdot \left(\frac{n}{\alpha e}\right)^{\frac{n}{\alpha}}} \\
&\sim \left(\frac{n+\alpha-2}{\alpha e} \cdot \frac{\alpha e}{n}\right)^{\frac{n}{\alpha}} \cdot \left(\frac{n+\alpha-2}{\alpha e}\right)^{\frac{\alpha-2}{\alpha}} \\
&= \left(1 + \frac{\alpha-2}{n}\right)^{\frac{n}{\alpha}} \cdot \left(\frac{n+\alpha-2}{\alpha e}\right)^{\frac{\alpha-2}{\alpha}} \\
&\sim e^{\frac{\alpha-2}{\alpha}} \cdot \left(\frac{n+\alpha-2}{\alpha e}\right)^{\frac{\alpha-2}{\alpha}} \\
&\sim \left(\frac{n+\alpha-2}{\alpha}\right)^{\frac{\alpha-2}{\alpha}} \\
&\sim n^{1-\frac{2}{\alpha}} \cdot \alpha^{-1+\frac{2}{\alpha}} \\
\|\mathcal{I}_{p,\alpha}\|_2 &= \alpha^2 \cdot \frac{\Gamma\left(\frac{n+2\alpha-2}{\alpha}\right)}{\Gamma\left(\frac{n}{\alpha}\right)} \cdot \frac{\Gamma\left(\frac{n}{p}\right)}{\Gamma\left(\frac{n+2p-2}{p}\right)} \cdot \frac{\Gamma\left(2-\frac{1}{p}\right)}{\Gamma\left(\frac{1}{p}\right)} \\
&= \alpha^2 \cdot \frac{\Gamma\left(\frac{n+2\alpha-2}{\alpha}\right)}{\Gamma\left(\frac{n}{\alpha}+1\right)} \cdot \frac{n}{\alpha} \cdot \frac{\Gamma\left(\frac{n}{p}\right)}{\Gamma\left(\frac{n+2p-2}{p}\right)} \cdot \frac{p}{n} \cdot \frac{\Gamma\left(2-\frac{1}{p}\right)}{\Gamma\left(\frac{1}{p}\right)} \\
&\sim \alpha p \cdot n^{1-\frac{2}{\alpha}} \cdot \alpha^{-1+\frac{2}{\alpha}} \cdot n^{-1+\frac{2}{p}} \cdot p^{1-\frac{2}{p}} \cdot \frac{\Gamma\left(2-\frac{1}{p}\right)}{\Gamma\left(\frac{1}{p}\right)} \\
&\sim \alpha^{\frac{2}{\alpha}} \cdot p^{2-\frac{2}{p}} \cdot \frac{\Gamma\left(2-\frac{1}{p}\right)}{\Gamma\left(\frac{1}{p}\right)} \cdot n^{\frac{2}{p}-\frac{2}{\alpha}}
\end{aligned}$$

This finishes the entire proof of Lemma 4.2. \square

Appendix D Proof of Lemma 3.3

Without the loss of generality, we assume $\mu = 1$ in the proof.

By Lemma 4.2, we then have $t_n = \frac{1}{\mu} \cdot \sqrt{\|\mathcal{I}_{\varphi_n}\|_2} \asymp n^{\frac{1}{\alpha}-\frac{1}{p}}$, so the rescaled $\tilde{\varphi}_n : \mathbb{R}^n \rightarrow \mathbb{R}$ has the form

$$\tilde{\varphi}_n(x) = c_{p,\alpha} n^{1-\frac{\alpha}{p}} \|x\|_p^\alpha,$$

where $c_{p,\alpha} = \alpha^{-1} \cdot p^{-\alpha+\frac{\alpha}{p}} \cdot \left(\frac{\Gamma\left(2-\frac{1}{p}\right)}{\Gamma\left(\frac{1}{p}\right)}\right)^{-\frac{\alpha}{2}}$.

D.1 Lemmas

We first state a few auxiliary lemmas.

Lemma D.1.

$$\|X\|_p = \left(\left(\frac{1}{\alpha}\right)^{1/\alpha} + o_P(1)\right) \cdot n^{\frac{1}{p}}.$$

Lemma D.2. *If $p > 1, \alpha > 0$, then*

$$\begin{aligned}
\sum_{i:|X_i|\leq 2|v_i|} (|X_i+v_i|^p - |X_i|^p) &= o_P(1). \\
\sum_{i:|X_i|\leq 2|v_i|} p \operatorname{sgn}(X_i) |X_i|^{p-1} v_i &= o_P(1). \\
\sum_{i:|X_i|\leq 2|v_i|} \frac{p(p-1)}{2} |X_i|^{p-2} v_i^2 &= o_P(1)
\end{aligned}$$

D.2 Main proof

We will first prove the case where $p > 1$. The proof of the corner case where $p = 1$ will be given in Section D.5

Verification of Condition D1.

Recall that

$$\tilde{\varphi}_n(x) = c_{p,\alpha} n^{1-\frac{\alpha}{p}} \|x\|_p^\alpha,$$

where $c_{p,\alpha} = \alpha^{-1} \cdot p^{-\alpha+\frac{\alpha}{p}} \cdot \left(\frac{\Gamma(2-\frac{1}{p})}{\Gamma(\frac{1}{p})}\right)^{-\frac{\alpha}{2}}$.

We then have

$$\begin{aligned} P_v^\varphi(X) &= \varphi(X+v) - \varphi(X) - \frac{1}{2} v^T \mathcal{I}_\varphi v \\ &= c_{p,\alpha} n^{1-\frac{\alpha}{p}} (\|X+v\|_p^\alpha - \|X\|_p^\alpha - \frac{1}{2}) \\ \nabla \phi(X) &= c_{p,\alpha} n^{1-\frac{\alpha}{p}} \cdot \alpha \|X\|_p^{\alpha-p} \cdot \text{sgn}(X) \odot X^{\odot(p-1)} \end{aligned}$$

Now let us consider

$$P_v^\varphi(X) - v^T \nabla \varphi(X) = c_{p,\alpha} n^{1-\frac{\alpha}{p}} (\|X+v\|_p^\alpha - \|X\|_p^\alpha - \frac{1}{2} - \langle \alpha \|X\|_p^{\alpha-p} \cdot \text{sgn}(X) \odot X^{\odot(p-1)}, v \rangle).$$

Then

$$\begin{aligned} & \|X+v\|_p^\alpha - \|X\|_p^\alpha - \langle \alpha \|X\|_p^{\alpha-p} \cdot \text{sgn}(X) \odot X^{\odot(p-1)}, v \rangle \\ & \asymp (1 + o_P(1)) \frac{\alpha \|X\|_p^{\alpha-p}}{p} \left(\|X+v\|_p^p - \|X\|_p^p - p \text{sgn}(X) \odot X^{\odot(p-1)} \odot v \right) \\ & \asymp \frac{1}{p} \cdot \alpha^{\frac{p}{\alpha}} \cdot n^{\frac{\alpha-p}{p}} \left(\|X+v\|_p^p - \|X\|_p^p - p \text{sgn}(X) \odot X^{\odot(p-1)} \odot v \right) \\ & \asymp \frac{1}{p} \cdot \alpha^{\frac{p}{\alpha}} \cdot n^{\frac{\alpha-p}{p}} \left(\sum_{i=1}^n (|X_i+v_i|^p - |X_i|^p - p \text{sgn}(X_i) \cdot X_i^{p-1} v_i) \right) \end{aligned}$$

To prove (1), it suffices to show $\left(\sum_{i=1}^n (|X_i+v_i|^p - |X_i|^p) - p \text{sgn}(X_i) \cdot X_i^{p-1} v_i - \frac{1}{2} \right) = o_P(1)$.

We expand this expression as

$$\begin{aligned} \sum_{i=1}^n (|X_i+v_i|^p - |X_i|^p - p \text{sgn}(X_i) \cdot X_i^{p-1} v_i) &= \sum_{i:|X_i|>2|v_i|} (|X_i+v_i|^p - |X_i|^p - p \text{sgn}(X_i) \cdot X_i^{p-1} v_i) \\ &+ \sum_{i:|X_i|\leq 2|v_i|} (|X_i+v_i|^p - |X_i|^p) \\ &- \sum_{i:|X_i|\leq 2|v_i|} p \text{sgn}(X_i) |X_i|^{p-1} v_i \end{aligned}$$

When $|X_i| > 2|v_i|$, note that

$$|X_i+v_i|^p = |X_i|^p (1 + v_i/X_i)^p = |X_i|^p \left(1 + p \frac{v_i}{X_i} + \frac{p(p-1)}{2} \frac{v_i^2}{X_i^2} + O\left(\frac{v_i^3}{X_i^3}\right) \right).$$

Combing with Lemma D.2, we then have

$$\begin{aligned} & \sum_{i=1}^n (|X_i+v_i|^p - |X_i|^p - p \text{sgn}(X_i) \cdot X_i^{p-1} v_i) \\ &= \sum_{i:|X_i|>2|v_i|} |X_i|^p \left(\frac{p(p-1)}{2} \frac{v_i^2}{X_i^2} + O\left(\frac{v_i^3}{X_i^3}\right) \right) + o_P(1) \end{aligned}$$

In order to verify Condition D1, we need to show: 1). $\sum_{i:|X_i|>2|v_i|} \frac{p(p-1)}{2} |X_i|^{p-2} v_i^2$ converges to a constant; 2). Show the third order term is vanishing, that is, $\sum_i |X_i|^{p-3} |v_i|^3 \rightarrow 0$.

We will prove the proposition in the following two steps:

Step 1.

We need to prove that

$$\sum_{i:|X_i|>2|v_i|} \frac{p(p-1)}{2} |X_i|^{p-2} v_i^2$$

converges to a constant.

This can be computed similarly as before. By Lemma D.2

$$\sum_{i:|X_i|<2|v_i|} \frac{p(p-1)}{2} |X_i|^{p-2} v_i^2 = o_P(1)$$

Therefore, we have

$$\sum_{i:|X_i|>2|v_i|} \frac{p(p-1)}{2} |X_i|^{p-2} v_i^2 = \sum_i \frac{p(p-1)}{2} |X_i|^{p-2} v_i^2 + o_P(1).$$

We further obtain

$$\sum_i \frac{p(p-1)}{2} |X_i|^{p-2} v_i^2 = \frac{p(p-1)}{2} \sum_i v_i^2 \cdot \left(\frac{n}{\alpha}\right)^{\frac{p-2}{\alpha}} \cdot |U_i|^{p-2}.$$

Again, we also have

$$\begin{aligned} \sum_{i=1}^n v_i^2 |\tilde{X}_i|^{p-2} &\sim \left(\frac{n}{p}\right)^{\frac{p-2}{p}} \cdot \sum_{i=1}^n v_i^2 |\tilde{U}_i|^{p-2} \\ &= \left(\frac{n}{p}\right)^{\frac{p-2}{p}} \cdot \sum_{i=1}^n v_i^2 |U_i|^{p-2} \cdot n^{\frac{p-2}{\alpha} - \frac{p-2}{p}} \\ &= n^{\frac{p-2}{\alpha}} \cdot p^{\frac{2}{p}-1} \cdot \sum_{i=1}^n v_i^2 |U_i|^{p-2}. \end{aligned}$$

This implies

$$\sum_{i=1}^n v_i^2 |U_i|^{p-2} = n^{-\frac{p-2}{\alpha}} \cdot p^{1-\frac{2}{p}} \sum_i v_i^2 |\tilde{X}_i|^{p-2} \xrightarrow{p} n^{-\frac{p-2}{\alpha}} \cdot p^{1-\frac{2}{p}} \cdot \frac{\Gamma(1-\frac{1}{p})}{\Gamma(\frac{1}{p})}.$$

Then we have

$$\begin{aligned} \sum_i \frac{p(p-1)}{2} |X_i|^{p-2} v_i^2 &= \frac{p(p-1)}{2} \cdot \left(\frac{n}{\alpha}\right)^{\frac{p-2}{\alpha}} \sum_i v_i^2 \cdot |U_i|^{p-2} \\ &\xrightarrow{p} \frac{1}{2} \cdot \alpha^{-\frac{p-2}{\alpha}} \cdot p^{3-\frac{2}{p}} \cdot \frac{\Gamma(2-\frac{1}{p})}{\Gamma(\frac{1}{p})}. \end{aligned}$$

Step 2. Show the third order term is vanishing

Prove that

$$\sum_i |X_i|^{p-3} |v_i|^3 \rightarrow 0.$$

$$\sum_i |X_i|^{p-3} |v_i|^3 \leq n \cdot \left(\sqrt{\frac{\log n}{n}}\right)^3 \cdot \max_i |X_i|^{p-3}$$

According to Theorem C.5, we have

$$X \stackrel{d}{=} t^{1/\alpha} \cdot r \cdot U,$$

where $t \sim \Gamma(\frac{n}{\alpha} + 1, 1)$, U has uniform distribution over $\partial K_\varphi = \{x : \|x\|_p = n^{1/p-1/\alpha}\}$, r has density nx^{n-1} over $[0, 1]$, and U, r, t are independent.

As a result, by Corollary C.6

$$\max_i |X_i| \sim \left(\frac{n}{\alpha}\right)^{\frac{1}{\alpha}} \cdot \max_i |U_i|.$$

Consider $\tilde{X} \sim e^{-\|x\|_p^p}$, we then have

$$\tilde{X} \stackrel{d}{=} \tilde{t} \cdot r \cdot \tilde{U},$$

where $\tilde{t} \sim \Gamma(\frac{n}{p} + 1, 1)$ and $\tilde{U} = n^{1/\alpha-1/p}U$.

Therefore,

$$\begin{aligned} \max_i |\tilde{X}_i| &\sim \left(\frac{n}{p}\right)^{\frac{1}{p}} \cdot \max_i |\tilde{U}_i| \\ &= \left(\frac{n}{p}\right)^{\frac{1}{p}} \cdot \max_i |U_i| \cdot n^{\frac{1}{\alpha}-\frac{1}{p}} \\ &= n^{\frac{1}{\alpha}} \cdot p^{-\frac{1}{p}} \cdot \max_i |U_i|. \end{aligned}$$

Therefore, we have

$$\max_i |U_i| \sim n^{-\frac{1}{\alpha}} \cdot p^{\frac{1}{p}} \max_i |\tilde{X}_i| \lesssim n^{-\frac{1}{\alpha}} \cdot p^{\frac{1}{p}} (\log n)^{1/p},$$

and we have

$$\max_i |X_i| \sim \left(\frac{n}{\alpha}\right)^{\frac{1}{\alpha}} \cdot \max_i |U_i| \lesssim (\log n)^{1/p}$$

As a result,

$$\sum_i |X_i|^{p-3} |v_i|^3 \leq n \cdot \left(\sqrt{\frac{\log n}{n}}\right)^3 \cdot \max_i |X_i|^{p-3} = (\log n)^{2.5-3/p} \cdot n^{-1/2} \rightarrow 0.$$

Verification of Condition (D2). Use Sudakov's theorem to prove asymptotic normality.

Note that by Lemma D.2

$$\sum_{i: |X_i| > 2|v_i|} p \operatorname{sgn}(X_i) |X_i|^{p-1} v_i = o_P(1) + \sum_{i=1}^n p \operatorname{sgn}(X_i) |X_i|^{p-1} v_i.$$

It suffices to show $\sum_{i=1}^n p \operatorname{sgn}(X_i) |X_i|^{p-1} v_i$ is asymptotically a normal random variable.

Firstly, we have

$$\sum_{i=1}^n (p \operatorname{sgn}(X_i) |X_i|^{p-1})^2 = p^2 \sum_{i=1}^n |X_i|^{2(p-1)}$$

According to Theorem C.5, we have

$$X \stackrel{d}{=} t^{1/\alpha} \cdot r \cdot U,$$

where $t \sim \Gamma(\frac{n}{\alpha} + 1, 1)$, U has uniform distribution over $\partial K_\varphi = \{x : \|x\|_p = n^{1/p-1/\alpha}\}$, r has density nx^{n-1} over $[0, 1]$, and U, r, t are independent.

As a result, by Corollary C.6

$$\sum_{i=1}^n |X_i|^{2(p-1)} \sim \left(\frac{n}{\alpha}\right)^{\frac{2p-2}{\alpha}} \cdot \sum_{i=1}^n |U_i|^{2p-2}.$$

Consider $\tilde{X} \sim e^{-\|x\|_p^p}$, we then have

$$\tilde{X} \stackrel{d}{=} \tilde{t} \cdot r \cdot \tilde{U},$$

where $\tilde{t} \sim \Gamma(\frac{n}{p} + 1, 1)$ and $\tilde{U} = n^{1/\alpha-1/p}U$.

Therefore,

$$\begin{aligned} \sum_{i=1}^n |\tilde{X}_i|^{2(p-1)} &\sim \left(\frac{n}{p}\right)^{\frac{2p-2}{p}} \cdot \sum_{i=1}^n |\tilde{U}_i|^{2p-2} \\ &= \left(\frac{n}{p}\right)^{\frac{2p-2}{p}} \cdot \sum_{i=1}^n |U_i|^{2p-2} \cdot n^{\frac{2p-2}{\alpha} - \frac{2p-2}{p}} \\ &= n^{\frac{2p-2}{\alpha}} \cdot p^{\frac{2}{p}-2} \cdot \sum_{i=1}^n |U_i|^{2p-2}. \end{aligned}$$

This implies

$$\sum_{i=1}^n |U_i|^{2p-2} = n^{-\frac{2p-2}{\alpha}} \cdot p^{2-2/p} \sum_{i=1}^n |\tilde{X}_i|^{2(p-1)} \xrightarrow{p} p^{2-2/p} n^{1-\frac{2p-2}{\alpha}} \cdot \mathbb{E}[|\tilde{X}_i|^{2p-2}],$$

where

$$\mathbb{E}[|\tilde{X}_i|^{2p-2}] = \frac{1}{\frac{2}{p}\Gamma(\frac{1}{p})} \int_{-\infty}^{\infty} |x|^{2p-2} e^{-|x|^p} dx = 2 \int_0^{\infty} x^{2p-2} e^{-x^p} dx = \frac{\Gamma(2 - \frac{1}{p})}{\Gamma(\frac{1}{p})}$$

is a constant when p is of constant order.

As a result, let $C = p^{4-2/p} \cdot \alpha^{-\frac{2p-2}{\alpha}} \cdot \frac{\Gamma(2-\frac{1}{p})}{\Gamma(\frac{1}{p})}$, we have

$$\sum_{i=1}^n (p \operatorname{sgn}(X_i) |X_i|^{p-1})^2 = p^2 \sum_{i=1}^n |X_i|^{2(p-1)} \sim p^2 \left(\frac{n}{\alpha}\right)^{\frac{2p-2}{\alpha}} \cdot \sum_{i=1}^n |U_i|^{2p-2} \xrightarrow{p} Cn,$$

satisfying the thin-shell condition of Sudakov's theorem and therefore $\sum_{i=1}^n p \operatorname{sgn}(X_i) |X_i|^{p-1} v_i$ is asymptotically normal with variance Cn .

D.3 Proof of Lemma 10.1

According to Theorem C.5, we have

$$X \stackrel{d}{=} t^{1/\alpha} \cdot r \cdot U,$$

where $t \sim \Gamma(\frac{n}{\alpha} + 1, 1)$, U has uniform distribution over $\partial K_\varphi = \{x : \|x\|_p = n^{1/p-1/\alpha}\}$, r has density nx^{n-1} over $[0, 1]$, and U, r, t are independent.

As a result,

$$\|X\|_p = |t^{1/\alpha}| \cdot |r| \cdot n^{1/p-1/\alpha} \sim \left(\frac{1}{\alpha}\right)^{1/\alpha} \cdot n^{1/p}$$

D.4 Proof of Lemma 10.2

Denote $z_i = I(|X_i| \leq 2|v_i|) = I(|t^{1/\alpha} \cdot r \cdot U_i| \leq 2|v_i|) \leq I(|t^{1/\alpha} \cdot r \cdot U_i| \leq 2\sqrt{\frac{2 \log n}{n}})$.

Consider $\tilde{X} \sim e^{-\|x\|_p^p}$, we then have

$$\tilde{X} \stackrel{d}{=} \tilde{t}^{1/p} \cdot r \cdot \tilde{U},$$

where $\tilde{t} \sim \Gamma(\frac{n}{p} + 1, 1)$ and $\tilde{U} = n^{1/\alpha-1/p}U$.

Since $p/n = o(1)$, we then have $\tilde{t} \sim \frac{\alpha}{p} \cdot t$ and $\tilde{t}^{1/p} \sim \frac{n^{1/p-1/\alpha} \cdot \alpha^{1/\alpha}}{p^{1/p}} \cdot t^{1/\alpha}$, then

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n I(|t^{1/\alpha} \cdot r \cdot U_i| \leq 2\sqrt{\frac{2 \log n}{n}}) &= \frac{1}{n} \sum_{i=1}^n I(|\tilde{t}^{1/p} \cdot r \cdot \tilde{U}_i| \leq 2\frac{\alpha^{1/\alpha}}{p^{1/p}} \sqrt{\frac{2 \log n}{n}}) \\ &= \frac{1}{n} \sum_{i=1}^n I(|\tilde{X}_i| \leq 2\frac{n^{1/p-1/\alpha} \cdot \alpha^{1/\alpha}}{p^{1/p}} \sqrt{\frac{2 \log n}{n}}) \\ &\sim \mathbb{P}(|\tilde{X}_i| \leq 2\frac{\alpha^{1/\alpha}}{p^{1/p}} \sqrt{\frac{2 \log n}{n}}), \end{aligned}$$

where \tilde{X}_i are *i.i.d.* drawn from the population with density $\propto e^{-|x|^p}$.

When $|x| \leq 2\frac{\alpha^{1/\alpha}}{p^{1/p}} \sqrt{\frac{2 \log n}{n}}$, and $2\frac{\alpha^{1/\alpha}}{p^{1/p}} \sqrt{\frac{2 \log n}{n}} = o(1)$, we then have

$$\mathbb{P}(|\tilde{X}_i| \leq 2\frac{\alpha^{1/\alpha}}{p^{1/p}} \sqrt{\frac{2 \log n}{n}}) \asymp \frac{\alpha^{1/\alpha}}{p^{1/p}} \sqrt{\frac{2 \log n}{n}},$$

which implies that

$$\frac{1}{n} \sum_{i=1}^n I(|X_i| \leq 2|v_i|) \asymp \frac{\alpha^{1/\alpha}}{p^{1/p}} \sqrt{\frac{\log n}{n}}.$$

As a result, we have that when $p > 1$,

$$\sum_{i:|X_i| \leq 2|v_i|} (|X_i + v_i|^p - |X_i|^p) \lesssim \frac{\alpha^{1/\alpha}}{p^{1/p}} \sqrt{n \log n} \cdot \left(\frac{\log n}{\sqrt{n}}\right)^p = \frac{\alpha^{1/\alpha}}{p^{1/p}} n^{(1-p)/2} (\log n)^{p+1/2} = o(1).$$

Similarly, for Lemma 5.4, we can use the same idea to show

$$\sum_{|X_i| < 2|v_i|} p \operatorname{sgn}(X_i) |X_i|^{p-1} v_i = o_P(1).$$

In fact, by using the same derivation, we have

$$\sum_{|X_i| < 2|v_i|} p \operatorname{sgn}(X_i) |X_i|^{p-1} v_i \lesssim \frac{\alpha^{1/\alpha}}{p^{1/p}} \sqrt{n \log n} \cdot \left(\frac{\log n}{\sqrt{n}}\right)^p = \frac{\alpha^{1/\alpha}}{p^{1/p}} n^{(1-p)/2} (\log n)^{p+1/2} = o(1).$$

D.5 $p = 1$

We now study the case where $p = 1$. Since, for general p , we have

$$\tilde{\varphi}_n(x) = c_{p,\alpha} n^{1-\frac{\alpha}{p}} \|x\|_p^\alpha,$$

where $c_{p,\alpha} = \alpha^{-1} \cdot p^{-\alpha+\frac{\alpha}{p}} \cdot \left(\frac{\Gamma(2-\frac{1}{p})}{\Gamma(\frac{1}{p})}\right)^{-\frac{\alpha}{2}}$.

Letting $p = 1$ we get

$$\tilde{\varphi}_n(x) = \frac{1}{\alpha} n^{1-\alpha} \|x\|_1^\alpha.$$

We first study the limit of $\|X\|_1$ when the density of X is given by $\frac{1}{2} e^{-\tilde{\varphi}_n(x)}$.

According to Theorem C.5, we have

$$X \stackrel{d}{=} t^{1/\alpha} \cdot r \cdot U,$$

where $t \sim \Gamma(\frac{n}{\alpha} + 1, 1)$, U has uniform distribution over $\partial K_\varphi = \{x : \|x\|_1 = (\alpha n^{\alpha-1})^{1/\alpha}\}$, r has density $n x^{n-1}$ over $[0, 1]$, and U, r, t are independent.

As a result,

$$\|X\|_1 = |t^{1/\alpha}| \cdot |r| \cdot (\alpha n^{\alpha-1})^{1/\alpha} \sim \left(\frac{n}{\alpha}\right)^{1/\alpha} \cdot 1 \cdot (\alpha n^{\alpha-1})^{1/\alpha} = n.$$

Now let us consider $\nabla \tilde{\varphi}_n(x)$, and we have

$$\nabla \tilde{\varphi}_n(X) = \frac{1}{\alpha} n^{1-\alpha} \cdot \alpha \cdot \|x\|_1^{\alpha-1} \cdot \text{sgn}(X) = \text{sgn}(X).$$

Since X is symmetric, the above expression implies the **Condition D3**, that is, $\|\nabla \tilde{\varphi}_n(X)\|_2 \sim \sqrt{n}$ and therefore $v^\top \nabla \tilde{\varphi}_n(X) \rightarrow N(0, 1)$.

To prove **Condition D1**, since $\frac{1}{2} v^\top I_\varphi v = \frac{1}{2}$, it suffices to show when $\|v\|_2 = 1, \alpha = 1$,

$$\|X + v\|_1 - \|X\|_1 \rightarrow N\left(\frac{1}{2}, 1\right).$$

The case where $\alpha > 1$ can be reduced to this setting by using the following technique. Let us write

$$X \stackrel{d}{=} t^{1/\alpha} \cdot r \cdot U,$$

where $t \sim \Gamma\left(\frac{n}{\alpha} + 1, 1\right)$, U has uniform distribution over $\partial K_\varphi = \{x : \|x\|_p = n^{1-1/\alpha}\}$, r has density $n x^{n-1}$ over $[0, 1]$, and U, r, t are independent.

$$\tilde{X} \stackrel{d}{=} \tilde{t} \cdot r \cdot \tilde{U},$$

where $\tilde{t} \sim \Gamma(n + 1, 1)$ and $\tilde{U} = n^{1/\alpha-1} U$.

Then

$$\begin{aligned} \sum_i (|X_i + v_i| - |X_i|) &\sim t^{1/\alpha} \cdot r \cdot \sum_i (|U_i + \left(\frac{n}{\alpha}\right)^{-1/\alpha} v_i| - |U_i|) \\ &\sim t^{1/\alpha} \cdot r \cdot n^{1-1/\alpha} \sum_i (|\tilde{U}_i + n^{-1} \cdot \alpha^{1/\alpha} v_i| - |\tilde{U}_i|) \\ &\sim \left(\frac{1}{\alpha}\right)^{1/\alpha} \cdot \tilde{t}^{1/\alpha} \cdot r \cdot n^{1-1/\alpha} \sum_i (|\tilde{U}_i + n^{-1} \cdot \alpha^{1/\alpha} v_i| - |\tilde{U}_i|) \\ &\sim \left(\frac{1}{\alpha}\right)^{1/\alpha} \cdot n^{1-1/\alpha} \sum_i (|\tilde{X}_i + n^{1/\alpha-1} \cdot \alpha^{1/\alpha} v_i| - |\tilde{X}_i|), \end{aligned}$$

which reduces to the $\alpha = 1$ setting up to some scaling.

Therefore, it suffices to show the asymptotic normality of $\|X + v\|_1 - \|X\|_1$ when X has density $\propto e^{-\|X\|_1}$. We are going to use the Berry-Esseen theorem. Suppose we have n independent random variables X_1, \dots, X_n with $\mathbb{E}X_i = \mu_i, \text{Var}X_i = \sigma_i^2, \mathbb{E}|X_i - \mu_i|^3 = \rho_i^3$. Consider the normalized random variable

$$S_n := \frac{\sum_{i=1}^n X_i - \mu_i}{\sqrt{\sum_{i=1}^n \sigma_i^2}}.$$

Denote its cdf by F_n . Then

Theorem D.3 (Berry-Esseen). *There exists a universal constant $C > 0$ such that*

$$\sup_{x \in \mathbb{R}} |F_n(x) - \Phi(x)| \leq C \cdot \frac{\sum_{i=1}^n \rho_i^3}{\left(\sum_{i=1}^n \sigma_i^2\right)^{3/2}}.$$

In the following, we proceed to calculating these three moments for the random variable $X_1, \dots, X_n \sim p$ with density $p(x) = \frac{1}{2} e^{-|x|}$ for $x \in \mathbb{R}$.

D.5.1 Expectation

Without loss of generality we assume $v > 0$.

$$\begin{aligned}
\mathbb{E}[|X + v|] - \mathbb{E}[|X|] &= \frac{1}{2} \int_{-\infty}^{\infty} |x + v| e^{-|x|} dx - \frac{1}{2} \int_{-\infty}^{\infty} |x| e^{-|x|} dx \\
&= \frac{1}{2} \left(\int_{-\infty}^{-v} |x + v| e^{-|x|} dx + \int_{-v}^0 |x + v| e^{-|x|} dx + \int_0^{\infty} |x + v| e^{-|x|} dx - \int_{-\infty}^0 |x| e^{-|x|} dx - \int_0^{\infty} |x| e^{-|x|} dx \right) \\
&= \frac{1}{2} \left(\int_{-\infty}^{-v} -(x + v) e^x dx + \int_{-v}^0 (x + v) e^x dx + \int_0^{\infty} (x + v) e^{-x} dx + \int_{-\infty}^0 x e^x dx - \int_0^{\infty} x e^{-x} dx \right) \\
&= \frac{1}{2} \left(\int_v^{\infty} x e^{-x} dx - \int_0^v x e^{-x} dx - \int_0^{\infty} x e^{-x} dx + v \left(- \int_v^{\infty} e^{-x} dx + \int_0^v e^{-x} dx + \int_0^{\infty} e^{-x} dx \right) \right) \\
&= - \int_0^v x e^{-x} dx + v \int_0^v e^{-x} dx \\
&= - (1 - e^{-v} - v e^{-v}) + v(1 - e^{-v}) \\
&= (v - 1 + e^{-v}) = \frac{1}{2} v^2 + o(v^2)
\end{aligned}$$

D.5.2 Variance

$$\begin{aligned}
2\mathbb{E}[|X + v| \cdot |X|] &= \int_{-\infty}^{\infty} |x + v| \cdot |x| e^{-|x|} dx \\
&= \int_{-\infty}^{-v} |x + v| \cdot |x| e^{-|x|} dx + \int_{-v}^0 |x + v| \cdot |x| e^{-|x|} dx + \int_0^{\infty} |x + v| \cdot |x| e^{-|x|} dx \\
&= \int_{-\infty}^{-v} (x + v) \cdot x e^x dx - \int_{-v}^0 (x + v) x e^x dx + \int_0^{\infty} (x + v) \cdot x e^{-x} dx \\
&= \int_{-\infty}^{-v} x^2 e^x dx - \int_{-v}^0 x^2 e^x dx + \int_0^{\infty} x^2 e^{-x} dx + v \left(\int_{-\infty}^{-v} x e^x dx - \int_{-v}^0 x e^x dx + \int_0^{\infty} x e^{-x} dx \right) \\
&= 2 \int_v^{\infty} x^2 e^{-x} dx + 2v \cdot \int_0^v x e^{-x} dx \\
&= 2v^2 e^{-v} + 4v e^{-v} + 4e^{-v} + 2v \cdot (1 - v e^{-v} - e^{-v}) \\
&= 2v + 2v e^{-v} + 4e^{-v} \sim 4 + o(v^2)
\end{aligned}$$

$$\begin{aligned}
2\mathbb{E}[|X + v|^2] &= \int_{-\infty}^{\infty} (x + v)^2 e^{-|x|} dx \\
&= \int_{-\infty}^{\infty} x^2 e^{-|x|} dx + v^2 \int_{-\infty}^{\infty} e^{-|x|} dx \\
&= 2 \int_0^{\infty} x^2 e^{-x} dx + 2v^2 \int_{-\infty}^{\infty} e^{-|x|} dx \\
&= 4 + 2v^2
\end{aligned}$$

$$\begin{aligned}
2\mathbb{E}[|X|^2] &= \int_{-\infty}^{\infty} |x|^2 e^{-|x|} dx \\
&= 2 \int_0^{\infty} x^2 e^{-x} dx \\
&= 4
\end{aligned}$$

$$\begin{aligned}
2\mathbb{E}[(|X+v| - |X|)^2] &= \mathbb{E}[|X+v|^2] + \mathbb{E}[|X|^2] - 2\mathbb{E}[|X+v| \cdot |X|] \\
&= \frac{1}{2}(8 + 2v^2 - 2(4 + o(v^2))) \\
&= v^2 + o(v^2)
\end{aligned}$$

D.5.3 Third moment

$$\begin{aligned}
\mathbb{E}[|X|^3] &= \int_{-\infty}^{\infty} |x|^3 e^{-|x|} dx \\
&= 2 \int_0^{\infty} x^3 e^{-x} dx \\
&= 12
\end{aligned}$$

$$\begin{aligned}
2\mathbb{E}[|X+v|^3] &= \int_{-\infty}^{\infty} |x+v|^3 e^{-|x|} dx \\
&= \int_{-\infty}^{-v} |x+v|^3 e^{-|x|} dx + \int_{-v}^0 |x+v|^3 e^{-|x|} dx + \int_0^{\infty} |x+v|^3 e^{-|x|} dx \\
&= \int_{-\infty}^{-v} -(x+v)^3 e^x dx + \int_{-v}^0 (x+v)^3 e^x dx + \int_0^{\infty} (x+v)^3 e^{-x} dx \\
&= 6e^{-v} + (-6 + 6e^{-v} + 6v - 3v^2 + v^3) + (v^3 + 3v^2 + 6v + 6) \\
&= 12e^{-v} + 12v + 2v^3
\end{aligned}$$

$$\begin{aligned}
2\mathbb{E}[|X+v|^2 \cdot |X|] &= \int_{-\infty}^{\infty} (x+v)^2 \cdot |x| e^{-|x|} dx \\
&= - \int_{-\infty}^0 (x+v)^2 x e^x dx + \int_0^{\infty} (x+v)^2 \cdot x e^{-x} dx \\
&= (v^2 - 4v + 6) + (v^2 + 4v + 6) \\
&= 2v^2 + 12
\end{aligned}$$

$$\begin{aligned}
2\mathbb{E}[|X+v| \cdot |X|^2] &= \int_{-\infty}^{\infty} |x+v| \cdot x^2 e^{-|x|} dx \\
&= - \int_{-\infty}^{-v} (x+v) \cdot x^2 e^x dx + \int_{-v}^0 (x+v)x^2 e^x dx + \int_0^{\infty} (x+v)x^2 e^{-x} dx \\
&= e^{-v}(v^2 + 4v + 6) + [2v - 6 + e^{-v}(v^2 + 4v + 6)] + 2(v + 3) \\
&= 4v + 2e^{-v}(v^2 + 4v + 6)
\end{aligned}$$

$$\begin{aligned}
\mathbb{E}[(|X+v| - |X|)^3] &= \frac{1}{2}(-12 + 12e^{-v} + 12v + 2v^3 - 3(2v^2 + 12) + 3(4v + 2e^{-v}(v^2 + 4v + 6))) \\
&= \frac{1}{2}(-48 + 48e^{-v} + 24v + 2v^3 - 6v^2 + 6v^2 e^{-v} + 24ve^{-v}) \\
&= \frac{1}{2}(48(-v + v^2/2 + o(v^2)) + 24v + 2v^3 - 6v^2 + 6v^2(1 - v + o(v)) + 24v(1 - v + o(v))) \\
&= o(v^2)
\end{aligned}$$

Combining the pieces, we get

$$\begin{aligned}\sum_{i=1}^n \mathbb{E}[|X_i + v_i| - |X_i|] &= \frac{1}{2} \sum_{i=1}^n v_i^2 + o(1) = \frac{1}{2} + o(1) \\ \sum_{i=1}^n \text{Var}[|X_i + v_i| - |X_i|] &= \sum_{i=1}^n v_i^2 + o(1) = 1 + o(1) \\ \sum_{i=1}^n \mathbb{E}[(|X_i + v_i| - |X_i| - \mathbb{E}[|X_i + v_i| - |X_i|])^3] &= o\left(\sum_{i=1}^n v_i^2\right) = o(1).\end{aligned}$$

Therefore, by using Theorem D.3, we get the desired result.