
One-bit Submission for Locally Private Quasi-MLE: Its Asymptotic Normality and Limitation

Hajime Ono¹

Kazuhiro Minami^{1,2}

Hideitsu Hino^{1,2,3}

¹ The Graduate University for Advanced Studies (SOKENDAI),

² The Institute of Statistical Mathematics, ³ RIKEN AIP

Abstract

Local differential privacy (LDP) is an information-theoretic privacy definition suitable for statistical surveys that involve an untrusted data curator. An LDP version of quasi-maximum likelihood estimator (QMLE) has been developed, but the existing method to build LDP QMLE is difficult to implement for a large-scale survey system in the real world due to long waiting time, expensive communication cost, and the boundedness assumption of derivative of a log-likelihood function. We provided an alternative LDP protocol without those issues, which is potentially much easily deployable to a large-scale survey. We also provided sufficient conditions for the consistency and asymptotic normality and limitations of our protocol. Our protocol is less burdensome for the users, and the theoretical guarantees cover more realistic cases than those for the existing method.

1 INTRODUCTION

The collection and use of data related to individuals continue at an unprecedented pace, raising a critical question: How do we balance the benefits of data use with the inherent privacy risks involved? One option is privacy protection based on differential privacy (DP) (Dwork et al., 2006; Dwork and Roth, 2014), whose information-theoretic definition requires data curators such as IT companies to stochastically perturb the results of research before making them available to third parties or the public. DP statistical

data processing has been widely studied both theoretically (Bassily et al., 2014) and empirically (Abadi et al., 2016). However, protection with DP does not work when the curator is adversarial. In fact, IT companies sometimes betray their users (Day et al., 2019).

To ensure that user privacy is protected even if the company is adversarial, local differential privacy (LDP) (Kasiviswanathan et al., 2011; Duchi et al., 2013) can be employed. By definition, LDP requires that the users themselves stochastically perturb their sensitive records before providing the records to a company. This perturbation ensures that no one can deterministically know the records. Notably, Google and Apple have conducted statistical surveys that guarantee user privacy based on this definition (Erlingsson et al., 2014; Apple Differential Privacy Team, 2017).

LDP versions of many statistical tools have been developed, including heavy-hitter estimation (Erlingsson et al., 2014; Fantì et al., 2016; Bassily and Smith, 2015; Qin et al., 2016), discrete distribution estimation (Kairouz et al., 2016; Ding et al., 2017), t-tests (Ding et al., 2018), chi-squared tests (Gaboardi and Rogers, 2018) and sparse linear regression (Wang and Xu, 2019). An LDP quasi-maximum likelihood estimator (QMLE) can also be included among these tools. QMLE is an estimator of a parameter likely approximating a distribution F generating a set of observations $D_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, from model family $\{F_\theta : \theta \in \Theta\}$. The likelihood of parameter θ is evaluated using the log-likelihood function $\ell(\theta; D_n) = \sum_{i=1}^n \log f_\theta(\mathbf{x}_i)/n$, and QMLE $\hat{\theta}_n$ is defined as the maximizer of $\ell(\theta; D_n)$. MLE is a special case in which there is a correct model: $F \in \{F_\theta : \theta \in \Theta\}$. Since no one observes the raw data under the LDP constraint, it is too optimistic to assume that we can specify a family including the true distribution. In this paper, we mainly consider QMLE rather than MLE. Under regularity conditions, QMLE has asymptotic normality. By understanding its normality, the curator is able to determine how likely and by how much the estimator is to deviate from the optimal point. Moreover, with

the asymptotic normality, we can perform the Wald test, which is an important application (Vaart, 2000).

Bhowmick et al. (2018) provided a framework for LDP M-estimators, which is a superclass of LDP QMLEs. It approximates the maximizer of an objective function with stochastic gradient descent. They showed that the covariance matrix of the normal distribution on which the estimator converges agrees with minimax optimal ones up to a constant.

However, the existing protocol may be difficult to deploy for a large-scale system in the real world due to the following three problems: (i) it requires a long waiting time for users, (ii) it is communication inefficient, and (iii) it requires finiteness of the derivative of the objective function. The existing protocol is interactive wherein the communication of the i th user depends on those of the previous $i - 1$ users. Though this interactivity gives more accurate statistics (Smith et al., 2017), it causes a long waiting time for users when millions of users are involved in the protocol. Communication efficiency is a non-ignorable problem for large-scale implementation, especially on Edge or IoT devices. When the parameter is d -dimensional and each component of the parameter uses float as a data type, each user submits $32d$ bits. It is also of great practical importance to be able to apply to unbounded domain data. The LDP constraints require a user to perturb her record so as to be indistinguishable from the other candidate records in the domain. An unbounded domain makes it difficult to satisfy this requirement since no one knows how many candidate values exist in the domain.

We provide low-user-side-cost protocols that involve no waiting time, require no boundedness assumption, and avoid high communication costs for QMLEs of regression. In this paper, we focus on regression which is a wide and important class. To eliminate waiting time, we abandon interactivity. Although less accurate than interactive methods, our protocol has a significant advantage in that the execution time on the user side is constant regardless of the number of users. To remove the boundedness assumption, we incorporate truncation into the protocol. This simple technique makes it possible for the protocol to perform safely even when the record domain is unbounded. For communication efficiency, we adopt the one-bit submission strategy whereby a record is stochastically quantized into a binary value (McGregor et al., 2010; Seide et al., 2014; Bassily and Smith, 2015; Ding et al., 2018; Wang et al., 2018). This strategy significantly reduces the communication cost. See Table 1 for a quick comparison of the communication costs and waiting time.

As the main contributions of this paper, (i) we give

consistency and asymptotic normality theorem with their sufficient conditions for our QMLEs, and (ii) we make explicit the limitations of the scope of our theoretical analysis. The asymptotic normality is useful for curators to adequately decide sample size n and privacy parameter ϵ . The sufficient conditions for our consistent and normality theorems are conditions on the model family and the true distribution. The curator should check the conditions for the model family when selecting the family. On the other hand, no one can evaluate the conditions on the true distribution. We recommend that the curator should carefully consider these conditions with the help of experts.

To discuss the sufficient conditions for our theorems on a concrete problem, we consider α -quantile linear regression (Davino et al., 2013). With this example, we can see that it is not so difficult to make a model family satisfying the conditions. Given $\alpha \in (0, 1)$, coefficients estimation for α -quantile regression is one of the standard statistical data analyses and QMLE is one of the solutions. For explanatory variables \mathbf{X} on \mathbb{R}^d and objective variable Y on \mathbb{R} , the goal of the α -quantile regression is to find coefficient $\beta \in \mathcal{B} \subset \mathbb{R}^d$ such that the inner product $\beta^\top \mathbf{X}$ well approximates the α -quantile of the distribution of Y , i.e., $\inf\{y | \Pr(Y \leq y | \mathbf{X}) > \alpha\}$. If we consider asymmetric Laplace distributions as the model family, this problem is a likelihood-maximizing problem. With this example, we are able to confirm that the conditions regarding the model family are easily satisfied. In addition, using real data, we observe the asymptotic behavior of our QMLE. The observations imply that the Frobenius norm of empirical covariance matrix shrinks in proportion to $1/n$ as expected in the asymptotic normality theorem.

We mention some related works. LDP regression by non-interactive algorithms has been studied in the context of LDP empirical risk minimization e.g., (Smith et al., 2017; Zheng et al., 2017; Wang et al., 2018, 2019a, 2021). Their targets are not analyses of asymp-

Table 1: Comparison of communication costs in number of submitting bits and waiting time of the protocols of the existing protocol (Bhowmick et al., 2018) and our protocol in two scenarios where explanatory variables \mathbf{X} are public and private, d is the dimension of parameter, k is number of explanatory variables, and n is the number of users.

Id	Scenario	Server	User	Wait
Bhowmick2018	\mathbf{X} pub	$32k$	$32d$	$O(n)$
	\mathbf{X} pri	0	$32k$	
Ours	\mathbf{X} pub	$32k$	1	$O(1)$
	\mathbf{X} pri	0	$d + 1$	

otic normality but seeking smaller risk. The studies for non-local differentially privately M-estimators took different ways from us (Smith, 2011; Chaudhuri and Hsu, 2012; Avella-Medina, 2020). Due to the difference in the privacy models, we do not compare our results with theirs. Bhowmick et al. (2018) showed asymptotic normality of their estimator relying on Polyak and Juditsky (1992) ’s asymptotic-normality proof for the estimators obtained by stochastic gradient descent. Since we do not use stochastic gradient descent, we prove our theorem by a different method.

The remainder of the paper is organized as follows: In Section 2, we introduce the notation used in this paper and some of the basic concepts. In Section 3, we describe our protocols for building QMLEs. In Section 4, we discuss QMLE for α -quantile regression as an illustrative application of the protocol. In Section 5, we report the results of a numerical experiment with real data. In Section 6, we offer concluding remarks.

2 PRELIMINARIES

We begin by defining some of the notation used in the paper. We denote by 0_d the d -length zero vector. When we take expectation while emphasizing the distribution F , we use $Fg = \mathbb{E}_{X \sim F}[g(X)]$ where g is a function. A comprehensive summary of our notation is provided in Appendix A.

2.1 Local Differential Privacy

Local differential privacy is a rigorous privacy definition for distributed statistical analyses. The definition requires each user to protect her sensitive record individually by stochastic perturbation. In particular, we consider the case in which users receive no feedback from the curator. LDP in such a situation is called non-interactive LDP; in this paper, we refer to non-interactive LDP simply as LDP.

We can now formally define LDP. Assume there are n users, each of whom possesses a sensitive record R_i for $i = 1, \dots, n$. Let \mathcal{R} be the domain of the records. Assume that there is also a curator who will perform a statistical analysis on the users’ records and that each user will submit her perturbed record to the curator. We can define the perturbation as a conditional distribution $Q(\cdot|R = r)$ and LDP as a property of Q .

Definition 1 (ϵ -LDP). *Given $\epsilon > 0$, distribution Q is ϵ -locally differentially private if, for any $r, r' \in \mathcal{R}$,*

$$\sup_{S \in \sigma(\mathcal{X})} Q(S|R = r) \leq e^\epsilon Q(S|R = r'),$$

where $\sigma(\mathcal{X})$ is a σ -algebra on \mathcal{X} .

This definition requires that the conditional distributions $Q(\cdot|r)$ and $Q(\cdot|r')$ are not so different from each other for any pair r, r' of records in \mathcal{R} . The ϵ represents the similarity of the conditional distributions. A smaller ϵ implies stricter privacy protection but less information of the outputs. ϵ thus controls the trade-off between privacy protection and utility.

This paper uses the bit flip (Ding et al., 2018) for the concrete implementation of conditional distribution Q satisfying ϵ -LDP. The bit flip stochastically maps a finite continuous interval $[c_l, c_u]$, where c_l and c_u are some real constants such that $c_l < c_u$, into discrete binary values $\{z_-, z_+\}$. Then, for any input $v \in [c_l, c_u]$ and with $C_\epsilon = \frac{e^\epsilon + 1}{e^\epsilon - 1}$, the bit flip is defined as

$$Q_{\text{bf}}(Z = z|v) = \begin{cases} \frac{1}{2} - \frac{v - \frac{c_l + c_u}{2}}{(c_u - c_l)C_\epsilon} & \text{if } z = z_-, \\ \frac{1}{2} + \frac{v - \frac{c_l + c_u}{2}}{(c_u - c_l)C_\epsilon} & \text{if } z = z_+. \end{cases}$$

When the input is close to c_u , the output is likely to be z_+ ; conversely, when the input is close to c_l , the output is likely to be z_- .

2.2 Quasi-Maximum Likelihood Estimator

Given observations $D_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ generated by distribution F , the likelihood of parameter θ of a model F_θ is evaluated by the log-likelihood function

$$\ell(\theta; D_n) = \frac{1}{n} \sum_{i=1}^n \log f_\theta(\mathbf{x}_i),$$

where f_θ is the density function of F_θ . Roughly speaking, the log-likelihood is the log of the probability that the observations are obtained assuming they are sampled from F_θ . For the likelihood function, QMLE $\hat{\theta}_n$ is defined as $\hat{\theta}_n = \operatorname{argmax}_{\theta \in \Theta} \ell(\theta; D_n)$. Not only D_n but also $\hat{\theta}_n$ itself is a random variable.

In this subsection, we review the consistency and asymptotic normality theorems of QMLEs by White (1982). To define the log-likelihood function well, we first need to make some assumptions. The first is that the observations are independently generated from a distribution F and that F has a regular Radon–Nikodym density function f . The second condition requires that the model family also has regular density functions.

Assumption 1. *Let ν be an appropriate measure on \mathcal{X} . For a constant k , the independent $1 \times k$ random vectors $\mathbf{X}_i, i = 1, \dots, n$, have common joint distribution function F on \mathcal{X} , a measurable Euclidean space, with measurable Radon–Nikodym density $f = dF/d\nu$.*

Assumption 2. *The family of distribution functions $F_\theta(\mathbf{x})$ has Radon–Nikodym densities $f_\theta(\mathbf{x}) = dF_\theta(\mathbf{x})/d\nu$ which are measurable in x for every $\theta \in \Theta$,*

a compact subset of a d -dimensional Euclidean space, and continuous in θ for every $\mathbf{x} \in \mathcal{X}$.

To guarantee consistency, we introduce an additional technical assumption.

Assumption 3. (a) $F \log f$ exists, and $|\log f_\theta(\mathbf{x})| \leq h(\mathbf{x})$ for all $\theta \in \Theta$, where h is integrable with respect to F ; (b) $F \log f_\theta$ has a unique maximum at $\theta^* \in \Theta$.

Under these regularity conditions, the QMLE converges to $\theta^* = \operatorname{argmax}_{\theta \in \Theta} F \log f_\theta$.

Theorem 1 (Theorem 2.2 in (White, 1982)). *Given Assumptions 1 to 3, $\hat{\theta}_n \rightarrow \theta^*$ as $n \rightarrow \infty$ for almost every sequence $\{\mathbf{X}_i\}_{i=1}^n$.*

We also have asymptotic normality under some additional assumptions regarding the existence of scores $\partial \log f_\theta(\mathbf{x})/\partial \theta$ and related quantities.

Assumption 4. $\partial \log f_\theta(\mathbf{x})/\partial \theta_j, j = 1, \dots, d$, are measurable of \mathbf{x} for each $\theta \in \Theta$ and continuously differentiable functions of θ for each $\mathbf{x} \in \mathcal{X}$.

Assumption 5. $|\partial \log f_\theta(\mathbf{x})/\partial \theta_{j_1} \cdot \partial \log f_\theta(\mathbf{x})/\partial \theta_{j_2}|$ and $|\partial^2 \log f_\theta(\mathbf{x})/\partial \theta_{j_1} \partial \theta_{j_2}|$, for $j_1, j_2 = 1, \dots, d$ are dominated by functions integrable with respect to F for all \mathbf{x} in \mathcal{X} and θ in Θ .

Assumption 6. (a) θ^* is interior to Θ ; (b) $B(\theta) = (F(\partial \log f_\theta/\partial \theta)(\partial \log f_\theta/\partial \theta)^\top)$ is nonsingular at $\theta = \theta^*$; (c) θ^* is a regular point of $A(\theta) = F \partial^2 \log f_\theta/\partial \theta^2$.

The following shows the asymptotic normality.

Theorem 2 (Theorem 3.2 in (White, 1982)). *Given Assumptions 1 to 6,*

$$\sqrt{n}(\hat{\theta}_n - \theta^*) \rightarrow \mathcal{N}(0, C(\theta^*))$$

where $C(\theta) = A(\theta)^{-1} B(\theta) A(\theta)^{-1}$.

When $F_{\theta^*} = F$, $C(\theta^*)$ is called the Fisher information matrix.

2.3 Quantile Regression

Linear quantile regression deals with the statistical problem of finding coefficients $\beta \in \mathcal{B} \subset \mathbb{R}^d$ such that, given \mathbf{x} , the inner product $\beta^\top \mathbf{x}$ well approximates the α -quantile $\inf\{y | F(Y \leq y | \mathbf{x}) > \alpha\}$ of $Y | \mathbf{x}$. The problem is often formulated as an optimization problem finding $\beta \in \mathcal{B}$ that minimizes the following objective function: Given observations $\{\mathbf{x}_i, y_i\}_{i=1}^n$,

$$\sum_{i=1}^n \rho_\alpha(y_i - \beta^\top \mathbf{x}_i) \text{ where } \rho_\alpha(\tau) = \begin{cases} (\alpha - 1)\tau & \text{if } \tau \leq 0, \\ \alpha\tau & \text{if } \tau > 0. \end{cases} \quad (1)$$

ρ_α is a convex function, which is called the check loss.

If we assume that objective variable Y is sampled from the asymmetric Laplace distribution defined below, the minimization of (1) is equivalent to the likelihood maximization for the parameter of the distributions: With $\sigma > 0$,

$$f_Y(y; \alpha, \mu, \sigma) = \frac{\alpha(1-\alpha)}{\sigma} \exp\left(-\rho_\alpha\left(\frac{y-\mu}{\sigma}\right)\right). \quad (2)$$

Hence the log-likelihood function is written as

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n \log f_Y(y_i; \alpha, \beta^\top \mathbf{x}_i, \sigma) \\ &= \log \frac{\alpha(1-\alpha)}{\sigma} - \frac{1}{n\sigma} \sum_{i=1}^n \rho_\alpha(y_i - \beta^\top \mathbf{x}_i). \end{aligned} \quad (3)$$

Finally, we revisit the classical result of the asymptotic normality of the MLE. Let $\hat{\beta}_n \in \mathcal{B}$ be the MLE that minimizes (3), and let β^* be the coefficient such that $F(Y \leq y | \mathbf{X} = \mathbf{x}) = F_Y(y; \alpha, \beta^{*\top} \mathbf{x}, \sigma)$ for almost every \mathbf{x} and y with appropriate α and σ . Then, the sequence of MLEs $\{\hat{\beta}_n\}_n$ converges as

$$\sqrt{n}(\hat{\beta}_n - \beta^*) \rightarrow \mathcal{N}(0_d, I^{-1}), \quad (4)$$

where $\mathcal{N}(0_d, I^{-1})$ is the normal distribution whose mean and covariance are 0_d and I^{-1} , respectively (Davino et al., 2013). Assuming that $\mathbb{E}[\mathbf{X}\mathbf{X}^\top]$ is non-singular, I is the Fisher information matrix defined as

$$I = \frac{\alpha(1-\alpha)}{\sigma^2} \mathbb{E}[\mathbf{X}\mathbf{X}^\top]. \quad (5)$$

3 PROPOSED PROTOCOL

We provide two protocols for building QMLEs of regression in two different privacy scenarios and give their asymptotic normality theorem. Then, we remark on their advantages, limitations, and possible future works.

3.1 Regression with Public \mathbf{X}

In this subsection, we consider regression with sensitive objective variable Y and public explanatory variables \mathbf{X} . This situation may seem strange, but we will give a practical use case. Consider a situation in which a company is planning to conduct a customer opinion survey on a new product. The company can control its features set \mathbf{X} and gives a new product with certain features $\mathbf{X} = \mathbf{x}$ to each customer. The customer gives an evaluation Y for $\mathbf{X} = \mathbf{x}$. The target of the company is to understand the conditional distribution of Y . In the survey, the company knows the \mathbf{X} s and their distribution, and they are public.

The system model is as follows: There are a single curator and n users. The curator selects distribution $F_{\mathbf{X}}$ on $\mathcal{X} \subset \mathbb{R}^k$, a measurable Euclidean space, generates \mathbf{X}_i for each user $i = 1, \dots, n$ following $F_{\mathbf{X}}$, and passes them to each user. Given $\mathbf{X}_i = \mathbf{x}_i$, user i independently generates Y_i following unknown conditional distribution $F(\cdot|\mathbf{x}_i)$ on $\mathcal{Y} \subset \mathbb{R}$, a measurable space, and truncates it into interval $[c_l, c_u]$. Let \bar{Y}_i be the truncated version of Y_i :

$$\bar{Y}_i = t(Y_i) \equiv \begin{cases} c_l & \text{if } Y_i \leq c_l, \\ Y_i & \text{if } c_l < Y_i < c_u, \\ c_u & \text{if } Y_i \geq c_u. \end{cases} \quad (6)$$

We let \bar{y}_i be a realization of \bar{Y}_i . Then, the user perturbs \bar{y}_i by the bit flip. Z_i that is perturbed \bar{Y}_i distributes as

$$p(Z_i = z|\mathbf{X}_i = \mathbf{x}) = \int Q_{\text{bf}}(z|t(y))dF(y|\mathbf{x}). \quad (7)$$

User i submits z_i which is a realization of Z_i to the curator. The user submission is always only one bit.

The curator considers model family $\{F_\beta(\cdot|\mathbf{x}) : \beta \in \mathcal{B}, \mathbf{x} \in \mathcal{X}\}$ that consists of conditional distributions parameterized by \mathcal{B} , a compact subset of a d -dimensional Euclidean space. For each $\beta \in \mathcal{B}$, we define conditional density function $p_\beta(z|\mathbf{x})$ by replacing F by F_β in (7). The target of the curator is to find β such that P_β well approximates P . In this subsection, we write P and P_β to designate joint distributions $P(\mathbf{x}, z)$ and $P_\beta(\mathbf{x}, z)$ rather than conditional distributions $P(z|\mathbf{x})$ and $P_\beta(z|\mathbf{x})$.

Given observations $D_n = \{(z_i, \mathbf{x}_i)\}_{i=1}^n$, the log-likelihood function is defined as

$$\begin{aligned} \ell(\beta; D_n) &\equiv \frac{1}{n} \sum_{i=1}^n \log p_\beta(\mathbf{x}_i, z_i) \\ &= \frac{1}{n} \sum_{i=1}^n (z_i \log \Lambda_\epsilon(\beta, \mathbf{x}_i) + (1 - z_i) \log(1 - \Lambda_\epsilon(\beta, \mathbf{x}_i))) \\ &\quad + \log F_{\mathbf{X}}(\mathbf{x}_i) \end{aligned}$$

where $\Lambda_\epsilon(\beta, \mathbf{x}) = p_\beta(z = 1|\mathbf{x})$. We define $\hat{\beta}_n = \arg \max_{\beta \in \mathcal{B}} \ell(\beta; D_n)$ and $\beta^* = \arg \max_{\beta \in \mathcal{B}} P \log p_\beta$. The model selection and optimization are performed by the curator, and the users do not have to care about them. The curator can change hyperparameters excepting c_u, c_l and ϵ and can try multiple model families without any additional cost for the users. The pseudo-code is included in Appendix B.

Now, we analyze the behavior of $\hat{\beta}_n$. To derive the consistency of our QMLE, we replace F and F_θ in Theorem 1 with P and P_β , respectively. We find the conditions under which Assumptions 1 to 3 are satisfied while replacing F and F_θ with P and P_β . To

satisfy Assumptions 1 and 2, we introduce the following assumptions.

Assumption 7. *Conditional distribution $F(\cdot|\mathbf{x})$ has a Radon–Nikodym density function $f(y|\mathbf{x}) = dF(y|\mathbf{x})/d\nu$ which is measurable in y for every $\mathbf{x} \in \mathcal{X}$.*

Assumption 8. *$F_{\mathbf{X}}$ has a measurable Radon–Nikodym density $f_{\mathbf{X}} = dF_{\mathbf{X}}/d\mu$ with some appropriate measure μ .*

Assumption 9. *The family of distribution functions $F_\beta(y|\mathbf{x})$ has Radon–Nikodym densities $f_\beta(y|\mathbf{x}) = dF_\beta(y|\mathbf{x})/d\nu$ which are measurable in y for every $\mathbf{x} \in \mathcal{X}$ and $\beta \in \mathcal{B}$, and continuous in β for every $\mathbf{x} \in \mathcal{X}$ and $y \in \mathcal{Y}$.*

These assumptions are satisfied with many distributions e.g., Gaussian and Bernoulli distributions. From these assumptions, it is obvious that $P(\mathbf{x}, z), P_\beta(\mathbf{x}, z)$ are measurable and that the density functions $p(\mathbf{x}, z) = p(z|\mathbf{x})f(\mathbf{x})$ and $p_\beta(\mathbf{x}, z) = p_\beta(z|\mathbf{x})f(\mathbf{x})$ exist.

In order for the QMLE for regression parameter to satisfy Assumption 3, we consider the following two conditions. The first one is the existence of $P \log p$ and integrable function $h(\mathbf{x}, z)$ such that $|\log p_\beta(\mathbf{x}, z)| \leq h(\mathbf{x}, z)$ for all $\beta \in \mathcal{B}$. $P \log p$ can be extended as

$$P \log p = F_{\mathbf{X}}(P_{|\mathbf{X}} \log p(\cdot|\mathbf{X}) + \log f_{\mathbf{X}}(\mathbf{X})).$$

Since $\log p(\cdot|\mathbf{X})$ is always bounded away from $-\infty$ and $+\infty$ by the following lemma, $\log p(\cdot|\mathbf{X})$ is always integrable with respect to P .

Lemma 1. *The value of $\Lambda_\epsilon(\beta, \mathbf{x})$ is bounded away from 0 and 1, for all $\beta \in \mathcal{B}$ and $\mathbf{x} \in \mathcal{X}$.*

See Appendix C.1.1 for the proof. Thus, if $F_{\mathbf{X}} \log F_{\mathbf{X}}$ exists, $P \log p$ also exists. Similarly, the existence of $P \log p_\beta$ depends on the existence of $F_{\mathbf{X}} \log f_{\mathbf{X}}$.

Assumption 10. *$F_{\mathbf{X}} \log f_{\mathbf{X}}$ exists.*

The second condition relates to the uniqueness of the maximum of the log-likelihood function. Because the maxima are not always unique, we adopt the following assumption.

Assumption 11. *$P \log p_\beta$ has a unique maximum.*

We now have consistency.

Theorem 3. *Suppose Assumptions 7 to 11 hold. Then, $\hat{\beta}_n \rightarrow \beta^*$ as $n \rightarrow \infty$ surely.*

Next, we derive the asymptotic normality. We find the conditions under which Assumptions 4 to 6 are satisfied. Assumption 4 specifies the continuous differentiability of $\partial \log p_\beta / \partial \beta$. The partial derivative is extended as

$$\frac{\partial}{\partial \beta} \log(p_\beta(\mathbf{x}, z)) = \frac{(2z - 1)\Lambda'_\epsilon(\beta, \mathbf{x})}{\Lambda_\epsilon(\beta, \mathbf{x})^z(1 - \Lambda_\epsilon(\beta, \mathbf{x}))^{1-z}}$$

where $\Lambda'_\epsilon(\beta, \mathbf{x}) = \partial \Lambda_\epsilon(\beta, \mathbf{x}) / \partial \beta$. By Lemma 1, the following is sufficient to satisfy the requirement.

Assumption 12. *Each element of $\Lambda'_\epsilon(\beta, \mathbf{x})$ is measurable of \mathbf{x} for each $\beta \in \mathcal{B}$ and continuously differentiable functions of β for each $\mathbf{x} \in \mathcal{X}$.*

Assumption 5 states that $|\partial^2 \log p_\beta / \partial \beta_{j_1} \partial \beta_{j_2}|$ and $|\partial \log p_\beta / \partial \beta_{j_1} \cdot \partial \log p_\beta / \partial \beta_{j_2}|$ for $j_1, j_2 = 1, \dots, d$ are bounded by functions integrable with respect to P . To verify this, we extend these values.

$$\begin{aligned} \frac{\partial^2 \log p_\beta(\mathbf{x}, z)}{\partial \beta^2} &= (2z - 1) \frac{\Lambda''_\epsilon(\beta, \mathbf{x})}{\Lambda_\epsilon(\beta, \mathbf{x})^{2z} (1 - \Lambda_\epsilon(\beta, \mathbf{x}))^{1-z}} \\ &\quad - \frac{\Lambda'_\epsilon(\beta, \mathbf{x}) \Lambda'_\epsilon(\beta, \mathbf{x})^\top}{\Lambda_\epsilon(\beta, \mathbf{x})^{2z} (1 - \Lambda_\epsilon(\beta, \mathbf{x}))^{2(1-z)}} \end{aligned}$$

where $\Lambda''_\epsilon(\beta, \mathbf{x}) = \partial^2 \Lambda_\epsilon(\beta, \mathbf{x}) / \partial \beta^2$, and

$$\begin{aligned} &\left(\frac{\partial}{\partial \beta} \log p_\beta(\mathbf{x}, z) \right) \left(\frac{\partial}{\partial \beta} \log p_\beta(\mathbf{x}, z) \right)^\top \\ &= \frac{\Lambda'_\epsilon(\beta, \mathbf{x}) \Lambda'_\epsilon(\beta, \mathbf{x})^\top}{\Lambda_\epsilon(\beta, \mathbf{x})^{2z} (1 - \Lambda_\epsilon(\beta, \mathbf{x}))^{2(1-z)}}. \end{aligned}$$

The denominators are always non-zero by Lemma 1. Thus, the following assumption is sufficient to satisfy the requirement.

Assumption 13. *The absolute values of each element of $\Lambda'_\epsilon(\beta, \mathbf{x})$ and $\Lambda''_\epsilon(\beta, \mathbf{x})$ are bounded by integrable functions with respect to P .*

Assumption 6 consists of three parts. The first part is that β^* is interior to \mathcal{B} . We assume this.

Assumption 14. *β^* is interior to \mathcal{B} .*

The second part is the non-singularity of $P((\partial \log p_\beta / \partial \beta)(\partial \log p_\beta / \partial \beta)^\top)$ at $\beta = \beta^*$.

$$\begin{aligned} &P \left(\frac{\partial}{\partial \beta} \log p_\beta \right) \left(\frac{\partial}{\partial \beta} \log p_\beta \right)^\top = \\ &F_{\mathbf{X}} \left(\frac{p(Z=1|\mathbf{X})}{\Lambda_\epsilon(\beta, \mathbf{X})^2} + \frac{p(Z=0|\mathbf{X})}{(1 - \Lambda_\epsilon(\beta, \mathbf{X}))^2} \right) \Lambda'_\epsilon(\beta, \mathbf{X}) \Lambda'_\epsilon(\beta, \mathbf{X})^\top. \end{aligned}$$

Thus, the following assumption is a sufficient condition of the requirement.

Assumption 15. *$F_{\mathbf{X}} \Lambda'_\epsilon(\beta^*, \mathbf{X}) \Lambda'_\epsilon(\beta^*, \mathbf{X})^\top$ is non-singular.*

The third part is non-singularity of $P \partial^2 \log p_\beta / \partial \beta^2$ at $\beta = \beta^*$. We obtain this from Assumption 11. If $P \log p_\beta$ has a second partial derivative along β and β^* is interior to \mathcal{B} , then $\partial^2 P \log p_\beta / \partial \beta^2$ must be negative-definite. If not, there exists β' such that $P \log p_{\beta'} = P \log p_{\beta^*}$ and $\beta' \neq \beta^*$. Finally, we obtain asymptotic normality.

Theorem 4. *Suppose Assumptions 7 to 15 hold. Then, $\sqrt{n}(\hat{\beta}_n - \beta^*) \rightarrow \mathcal{N}(0_d, C(\beta^*))$ where $C(\beta) = A^{-1}(\beta)B(\beta)A^{-1}(\beta)$ with $A(\beta) = P \partial^2 \log p_\beta / \partial \beta^2$ and $B(\beta) = P(\partial \log p_\beta / \partial \beta)(\partial \log p_\beta / \partial \beta)^\top$.*

3.2 Regression with Private X

Next, we consider regression when both objective variables and explanatory variables are sensitive and are submitted with perturbation. The system model is that each user i generates \mathbf{X}_i following unknown distribution $F_{\mathbf{X}}$ and then generates Y_i following unknown conditional distribution $F(\cdot | \mathbf{X}_i)$.

The communication protocol is as follows. User i stochastically perturbs \mathbf{X}_i and Y_i by LDP mechanism Q . We denote the perturbed ones by $\mathbf{Z}^{(\mathbf{X})}$ and $Z^{(Y)}$, respectively. Q consists of $Q_{Z^{(Y)}}$ and $Q_{\mathbf{Z}^{(\mathbf{X})}}$ perturbing Y_i and \mathbf{X}_i , respectively. The privatized objective variable $Z^{(Y)}$ is the same as Z in the previous subsection without the privacy budget consumed by the LDP mechanisms. On the other hand, since $\mathbf{Z}^{(\mathbf{X})}$ was not defined in the previous section, we need to define $Q_{\mathbf{Z}^{(\mathbf{X})}}$. We use the bit flip as $Q_{\mathbf{Z}^{(\mathbf{X})}}$ in an element-wise manner. Each element is randomized with privacy budget $\epsilon/(k+1)$. The total consumption of the privacy budget per user does not exceed ϵ by the sequential composition theorem (McSherry, 2009). We set the domain of $Q_{\mathbf{Z}^{(\mathbf{X})}}$ to $\{-1, +1\}^k$. For each $\mathbf{z}^{(\mathbf{X})} \in \{-1, +1\}^k$,

$$Q_{\mathbf{Z}^{(\mathbf{X})}}(\mathbf{z}^{(\mathbf{X})} | \mathbf{x}) = \prod_{j=1}^k \left(\frac{1}{2} + \frac{t(x_j) z_j^{(\mathbf{X})}}{2C_{\epsilon/(k+1)}} \right). \quad (8)$$

The generated privatized variables $(\mathbf{z}_i^{(\mathbf{X})}, z_i^{(Y)})$ are submitted to the curator.

In the communication protocol, each user submits $(k+1)$ bits to the curator, and the curator sends no information to the users. This privacy scenario is nearly the same as the Bhowmich's one, and our communication protocol is more efficient than theirs. In their protocol, each user receives and submits d float or double values, either $64d$ bits or $144d$ bits. Thus, our protocol results in communication costs that are roughly 64 or 144 times smaller than their protocol when $k \leq d$.

The curator defines model family $\{F_\beta(y|\mathbf{x}) : \beta \in \mathcal{B}, \mathbf{x} \in \mathcal{X}\}$ and provisional distribution $\hat{F}_{\mathbf{X}}$. Though the true $F_{\mathbf{X}}$ is unknown, the curator must assume some distribution of \mathbf{X} to compute the log-likelihood function, as we will see later. $\hat{F}_{\mathbf{X}}$ is a kind of prior distribution.

Since the discussion of consistency and asymptotic normality has much in common with the previous subsection, here we describe only the differences. See the appendix for details. Given observations $D_n = \{(\mathbf{z}_i^{(\mathbf{X})}, z_i^{(Y)})\}_{i=1}^n$, the likelihood function is

$$\ell(\beta; D_n) = \frac{1}{n} \sum_{i=1}^n (\log \hat{p}_{\mathbf{Z}}^{(\mathbf{X})}(\mathbf{z}_i^{(\mathbf{X})}) + z_i^{(Y)} \log \Phi(\beta, \mathbf{z}_i^{(\mathbf{X})}))$$

$$\begin{aligned}
 & + (1 - z_i^{(Y)}) \log(1 - \Phi(\beta, \mathbf{z}_i^{(\mathbf{X})})) \\
 \text{where } \hat{p}_{\mathbf{Z}}^{(\mathbf{X})}(\mathbf{z}^{(\mathbf{X})}) & \equiv \int Q_{\mathbf{Z}^{(\mathbf{X})}}(\mathbf{z}^{(\mathbf{X})} | \mathbf{x}) d\hat{F}_{\mathbf{X}}(\mathbf{x}), \\
 \Phi(\beta, \mathbf{z}^{(\mathbf{X})}) & \equiv \frac{\hat{F}_{\mathbf{X}}(\Lambda_{\epsilon/(d+1)}(\beta, \mathbf{X}) Q_{\mathbf{Z}^{(\mathbf{X})}}(\mathbf{Z}^{(\mathbf{X})} | \mathbf{X}))}{\hat{p}_{\mathbf{Z}}^{(\mathbf{X})}(\mathbf{z}^{(\mathbf{X})})}
 \end{aligned}$$

QMLE $\hat{\beta}_n$ is defined as $\hat{\beta}_n \equiv \operatorname{argmin}_{\beta \in \mathcal{B}} \ell(\beta; D_n)$.

We can show consistency based on Theorem 1 under the assumption that the curator chooses a regular distribution as $\hat{F}_{\mathbf{X}}$.

Assumption 16. $\hat{F}_{\mathbf{X}}$ has a measurable Radon-Nikodym density $\hat{f}_{\mathbf{X}} = d\hat{F}_{\mathbf{X}}/d\mu$.

Theorem 5. Suppose Assumptions 7 to 9, 11 and 16 hold. Then, $\hat{\beta}_n \rightarrow \beta^*$ as $n \rightarrow \infty$ for almost every sequence $\{(\mathbf{Z}_i^{(\mathbf{X})}, Z_i^{(Y)})\}_i$.

For details, see Appendix C.2. This consistent theorem does not require the existence of $F_{\mathbf{X}} \log f_{\mathbf{X}}$ unlike Theorem 3. We can obtain the existence from Assumption 16 and the properties of $\hat{p}_{\mathbf{Z}^{(\mathbf{X})}}$. The discretization by the bit flip relaxes the integrable condition.

To show asymptotic normality, we adopt several additional assumptions.

Assumption 17. $\Phi'(\beta, \mathbf{z}^{(\mathbf{X})})$ is continuous differentiable function of β .

Assumption 18. Each component of $\Phi''(\beta, \mathbf{z}^{(\mathbf{X})})$ and $(\Phi'(\beta, \mathbf{z}^{(\mathbf{X})}))(\Phi'(\beta, \mathbf{z}^{(\mathbf{X})}))^\top$ is bounded by integrable functions with respect to P .

Assumption 19. $\mathbb{E}_{\mathbf{Z}^{(\mathbf{X})}} [(\Phi'(\beta, \mathbf{Z}^{(\mathbf{X})}))(\Phi'(\beta, \mathbf{Z}^{(\mathbf{X})}))^\top]$ is non-singular at $\beta = \beta^*$.

Assumption 17 is used to prove the requirement corresponding to Assumption 4. The requirement corresponding to Assumption 5 is satisfied with Assumption 18, which requires that the curator should design Φ such that its first and second derivatives almost surely take finite values. The requirement corresponding to Assumption 6 is satisfied with Assumptions 11, 14 and 19. We now have asymptotic normality.

Theorem 6. Suppose Assumptions 7 to 9, 11, 14 and 16 to 19 hold. Then, $\sqrt{n}(\hat{\beta}_n - \beta^*) \rightarrow \mathcal{N}(0_d, C(\beta^*))$ where $C(\beta) = A^{-1}(\beta)B(\beta)A^{-1}(\beta)$ with $A(\beta) = P\partial^2 \log p_\beta / \partial \beta^2$ and $B(\beta) = P(\partial \log p_\beta / \partial \beta)(\partial \log p_\beta / \partial \beta)^\top$.

3.3 Remark and Limitation

The assumptions for proving consistency and asymptotic normality in Theorems 3 to 6 are not relevant to privacy preservation. Even if those assumptions do not hold, users' privacy is still protected as long as the ϵ -LDP mechanisms correctly work. The users who

supply data do not need to worry about these assumptions at all.

The requirements of our theorems clarify the properties of the model that the curator should check. The curator is free to choose any linear or non-linear model as long as it satisfies these properties. In addition, those requirements place few restrictions on model selection since the curator can modify the model after data collection.

As we see in Section 4, it is not so difficult to craft a model satisfying the requirements. We thus expect that most standard regression models satisfy them.

The first limitation relates to the problem of choosing \hat{F} . Although any \hat{F} satisfying Assumptions 10 and 16 can be acceptable, a poor choice of \hat{F} may make it difficult to satisfy the other assumptions. The theorem provides no method for choosing a better \hat{F} , which remains an open problem.

The second limitation relates to the true distribution, which is a common problem in most statistical theories. We have no method to evaluate Assumptions 7, 11 and 14. The curator never know the exact value of β^* and $C(\beta^*)$. The curator should carefully consider these assumptions with the help of experts.

The exploration of better mechanisms is our future work. There may exist $Q_{\mathbf{Z}^{(\mathbf{X})}}$ giving us a more sharp covariance matrix. In the context of LDP, vector submission is studied by many researchers e.g., (Duchi et al., 2013; Erlingsson et al., 2014; Bassily and Smith, 2015; Wang et al., 2019b).

Better selection of c_l and c_u is another future work. Whether certain c_l and c_u are good or bad strongly depends on F , and we have no general strategy to select better c_l and c_u .

One of the potential applications of our algorithms is bootstrapping. In the above subsections, we described that our algorithms output only one estimator in each protocol. However, without additional privacy loss, the curator can compute many estimators using the subsets of the submitted data. The post-processing invariant enables us to perform such an operation. This is one of the advantages of a non-interactive algorithm.

Another potential application is a misspecification test to determine whether the model family contains the true distribution (White, 1982). In the LDP setting, since the raw data are distributed, no single entity has knowledge on the statistical properties of the raw data. It is difficult to evaluate whether a model family is appropriate. A curator performs the test as a preliminary experiment. The results of the test would help the curator to quantitatively assess the confidence

level of the main survey.

4 EXAMPLE: QUANTILE REGRESSION

In this section, we show the QMLEs for quantile regression as a concrete example of our QMLEs. One of the main goals of this section is to show that it is possible to replace some of the assumptions noted in the previous section with a concrete implementation of the model. We note that the notation used in Section 4.1 and Section 4.2 is the same as that used in Section 3.1 and Section 3.2. Here, $k = d$.

4.1 With Public \mathbf{X}

As described in Section 2.3, we can formulate the α -quantile regression as a quasi-maximum likelihood estimation problem. For some $\sigma > 0$, we set f_β as

$$f_\beta(y|\mathbf{x}) = \frac{\alpha(1-\alpha)}{\sigma} \exp\left(-\rho_\alpha\left(\frac{y-\beta^\top\mathbf{x}}{\sigma}\right)\right)$$

for each y and \mathbf{x} , where ρ_α is defined in (1). This construction satisfies Assumption 9: measurable and continuous.

When we choose the product of independent d uniform distributions on interval $[-1, +1]$ as $F_{\mathbf{X}}$, Assumption 10 is satisfied.

Let Ψ_ϵ be the function such that $\Lambda_\epsilon(\beta, \mathbf{x}) = \Psi_\epsilon(\beta^\top\mathbf{x})$. Then, $\Lambda'_\epsilon(\beta, \mathbf{x}) = \Psi'_\epsilon(\beta^\top\mathbf{x})\mathbf{x}$ and $\Lambda''_\epsilon(\beta, \mathbf{x}) = \Psi''_\epsilon(\beta^\top\mathbf{x})\mathbf{x}\mathbf{x}^\top$ where $\Psi'_\epsilon(\theta) = \partial\Psi_\epsilon(\theta)/\partial\theta$ and $\Psi''_\epsilon(\theta) = \partial^2\Psi_\epsilon(\theta)/\partial\theta^2$. It has the following property.

Lemma 2. *$\Psi_\epsilon(\theta)$ is a strictly monotonically increasing function and is bounded away from 0 and 1. $\Psi'_\epsilon(\theta)$ and $\Psi''_\epsilon(\theta)$ exist and for any $\theta \in \mathbb{R}$, and their absolute values are bounded.*

See Appendix C.3 for the proof. From the second part of Lemma 2, Assumptions 12 and 13 are satisfied. Mover, $F_{\mathbf{X}}(\mathbf{X}\mathbf{X}^\top)$ is a non-singular matrix since

$$F_{\mathbf{X}}X_{j_1}X_{j_2} = \begin{cases} 0 & \text{if } j_1 \neq j_2, \\ \frac{1}{3} & \text{if } j_1 = j_2. \end{cases}$$

Thus, Assumption 15 is satisfied.

As a consequence of Theorems 3 and 6, we have the following corollaries.

Corollary 1. *Suppose Assumptions 7, 11 and 14 hold. Then, $\hat{\beta}_n \rightarrow \beta^*$ almost surely and $\sqrt{n}(\hat{\beta}_n - \beta^*) \rightarrow \mathcal{N}(0_d, C(\beta^*))$.*

To prove this corollary, we need only three assumptions. The concrete constructions of the model remove some of the assumptions used in Theorems 3 and 6.

Although the accuracy of our QMLEs is not a focus of this paper, we did conduct a rough comparison of accuracy with existing works. As a result, we found with $\epsilon \downarrow 0$, the Fisher information of our MLEs is $\sigma^2/\alpha(1-\alpha)$ times smaller than the upper bound shown in (Barnes et al., 2020). For details, see Appendix D.

4.2 With Private \mathbf{X}

In this setting, the curator does not know $F_{\mathbf{X}}$. Instead of $F_{\mathbf{X}}$, we adopt the product distribution of d symmetric binary distributions on $\{-1, +1\}$. Then, Assumption 16 is satisfied.

With $\epsilon' = \epsilon/(d+1)$, Φ is extended as

$$\Phi(\beta, \mathbf{z}^{(\mathbf{X})}) = \frac{\sum_{\mathbf{x} \in \{\pm 1\}^d} \Psi_{\epsilon'}(\beta^\top\mathbf{x}) \exp\left(\epsilon'\mathbf{1}\left[z_j^{(\mathbf{X})} = x_j\right]\right)}{p_{\mathbf{Z}(\mathbf{X})}(\mathbf{z}^{(\mathbf{X})})(e^{\epsilon'} + 1)^{d2^d}},$$

where $\Psi_{\epsilon'}$ is defined in the previous subsection. Due to the properties of $\Psi_{\epsilon'}$, which we evaluated in the previous subsection, Assumptions 17 and 18 are obviously satisfied. By the monotonicity of $\Psi_{\epsilon'}$, Assumption 19 is also satisfied. Now, as a corollary of Theorems 5 and 6, we obtain the following result.

Corollary 2. *Suppose Assumptions 7, 11 and 14 hold. Then, $\sqrt{n}(\hat{\beta}_n - \beta^*) \rightarrow \mathcal{N}(0_d, C(\beta^*))$ where $C(\beta) = A^{-1}(\beta)B(\beta)A^{-1}(\beta)$.*

5 NUMERICAL EVALUATION

In this section, we observe the behavior of our QMLE for real data. We consider the QMLE for quantile regression in the public \mathbf{X} case. Since we do not know the true distribution generating the real data, we cannot perform exact comparisons with the theoretical result, Corollary 3. Here, we observe the empirical covariance of the QMLEs to evaluate the convergence of the distribution of the QMLE. For additional numerical evaluations, see Appendix E.

We numerically compare the covariance matrices with varying n and ϵ . We use CO and NOx emission data set (Kaya et al., 2019), which consists of 36,733 records of 11 sensors attached to a turbine of a power plant. Although this data is not sensitive, we chose this data because of its large number of records and its format. We treat the 11th column as y and treat the columns from the first to 9th as \mathbf{x} . We set $c_u = 110, c_l = 40, \sigma = 1.0$ and $\alpha = 0.3$. These specific values of hyperparameters do not have a particular meaning. We vary n from 5,000 to 35,000 in increments of 5,000 for $\epsilon \in \{1, 2.5, 5, 10\}$. For each combination of n and ϵ , we sub-sample n records 1,000 times without replacement from the 36,733 records. For each sub-data, we perturb ys and compute a QMLE as described in




Figure 1: Frobenius norm of covariance matrices. The norms decrease in proportion to $1/n$ for each ϵ .

Section 4.1. With the 1,000 QMLEs, we obtain the empirical covariance matrix and its Frobenius norm. We implemented the simulations with Python 3.9.2, NumPy 1.19.2, and SciPy 1.6.1. The Python code is contained in the supplementary material.

Figure 1 shows the result. The horizontal and vertical axes show n and the value of each Frobenius norm in log-scale, respectively. For each ϵ , with large n , the norm of the covariance matrix is smaller. The decreasing speed is $O(1/n)$, and this result is compatible with the theoretical result. Greater ϵ also gives smaller covariance. In this case, the QMLE is concentrated in one point, and, as n increases, the distribution becomes more concentrated at that point.

6 CONCLUSION

We developed the simple protocols for building QMLEs from distributed data while guaranteeing ϵ -LDP for the users. They address the two different privacy scenarios. In the protocols, users submit only one or a few bits to the curator and do not need to wait for one another. Moreover, the users do not need to perform complex computations such as integration or derivation. Thus, the protocols are highly user-friendly and suitable for low-priced devices. We clarified the sufficient conditions for the QMLEs to be consistent and asymptotically normal, and showed their limitations. We showed that the sufficient conditions are relaxed with a concrete implementation.

Acknowledgements

H.H. is partly supported by JSPS KAKENHI No. JP19H04113 and CREST JST CREST Grant No. JP-MJCR2015, and K.M. is partly supported by JSPS KAKENHI No. JP21H04403.

References

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in

private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography*, pages 265–284, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3–4):211–407, August 2014.

R. Bassily, A. Smith, and A. Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 464–473, 2014.

Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS '16*, pages 308–318, New York, NY, USA, 2016. Association for Computing Machinery.

Matt Day, Giles Turner, and Natalia Drozdak. Amazon workers are listening to what you tell alexa, Apr 2019. URL <https://www.bloomberg.com/news/articles/2019-04-10/is-anyone-listening-to-you-on-alexa-a-global-team-reviews-audio>.

Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.

J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Local privacy and statistical minimax rates. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 429–438, Oct 2013.

Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, CCS '14*, page 1054–1067, New York, NY, USA, 2014. Association for Computing Machinery.

Apple Differential Privacy Team. Learning with privacy at scale, 2017. URL <https://machinelearning.apple.com/2017/12/06/learning-with-privacy-at-scale.html>.

Giulia Fanti, Vasyl Pihur, and Úlfar Erlingsson. Building a rappor with the unknown: Privacy-preserving learning of associations and data dictionaries. *Proceedings on Privacy Enhancing Technologies*, 3:41–61, 2016.

Raef Bassily and Adam Smith. Local, private, efficient protocols for succinct histograms. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, STOC '15*, page 127–135, New

- York, NY, USA, 2015. Association for Computing Machinery.
- Zhan Qin, Yin Yang, Ting Yu, Issa Khalil, Xiaokui Xiao, and Kui Ren. Heavy hitter estimation over set-valued data with local differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS '16*, page 192–203, New York, NY, USA, 2016. Association for Computing Machinery.
- Peter Kairouz, Keith Bonawitz, and Daniel Ramage. Discrete distribution estimation under local privacy. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2436–2444, New York, New York, USA, 20–22 Jun 2016. PMLR.
- Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In *NIPS*, pages 3574–3583, 2017.
- Bolin Ding, Harsha Nori, Paul Li, and Joshua Allen. Comparing population means under local differential privacy: With significance and power. In *AAAI*, pages 26–33, 2018.
- Marco Gaboardi and Ryan Rogers. Local private hypothesis testing: Chi-square tests. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1626–1635. PMLR, 10–15 Jul 2018.
- Di Wang and Jinhui Xu. On sparse linear regression in the local differential privacy model. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6628–6637. PMLR, 09–15 Jun 2019.
- A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 2000.
- Abhishek Bhowmick, John Duchi, Julien Freidinger, Gaurav Kapoor, and Ryan Rogers. Protection against reconstruction and its applications in private federated learning. *arXiv preprint arXiv:1812.00984*, 2018.
- A. Smith, A. Thakurta, and J. Upadhyay. Is interaction necessary for distributed private learning? In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 58–77, 2017.
- Andrew McGregor, Ilya Mironov, Toniann Pitassi, Omer Reingold, Kunal Talwar, and Salil Vadhan. The limits of two-party differential privacy. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 81–90, 2010. doi: 10.1109/FOCS.2010.14.
- Frank Seide, Hao Fu, Jasha Droppo, Gang Li, and Dong Yu. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns. In *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- Di Wang, Marco Gaboardi, and Jinhui Xu. Empirical risk minimization in non-interactive local differential privacy revisited. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS'18*, page 973–982, Red Hook, NY, USA, 2018. Curran Associates Inc.
- Cristina Davino, Marilena Furno, and Domenico Vistocco. *Quantile regression: theory and applications*. Wiley series in probability and statistics. Wiley, Hoboken, NJ, 2013.
- Kai Zheng, Wenlong Mou, and Liwei Wang. Collect at once, use effectively: Making non-interactive locally private learning possible. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 4130–4139. PMLR, 06–11 Aug 2017.
- Di Wang, Adam Smith, and Jinhui Xu. Noninteractive locally private learning of linear models via polynomial approximations. In Aurélien Garivier and Satyen Kale, editors, *Proceedings of the 30th International Conference on Algorithmic Learning Theory*, volume 98 of *Proceedings of Machine Learning Research*, pages 898–903. PMLR, 22–24 Mar 2019a.
- Di Wang, Huangyu Zhang, Marco Gaboardi, and Jinhui Xu. Estimating smooth GLM in non-interactive local differential privacy model with public unlabeled data. In Vitaly Feldman, Katrina Ligett, and Sivan Sabato, editors, *Proceedings of the 32nd International Conference on Algorithmic Learning Theory*, volume 132 of *Proceedings of Machine Learning Research*, pages 1207–1213. PMLR, 16–19 Mar 2021.
- Adam Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing, STOC '11*, page 813–822, New York, NY, USA, 2011. Association for Computing Machinery.
- Kamalika Chaudhuri and Daniel Hsu. Convergence rates for differentially private statistical estimation. In *Proceedings of the 29th International Conference on Machine Learning, ICML'12*, page 1715–1722, Madison, WI, USA, 2012. Omnipress.

- Marco Avella-Medina. Privacy-preserving parametric inference: A case for robust statistics. *Journal of the American Statistical Association*, 0(0):1–15, 2020.
- B. T. Polyak and A. B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
- Halbert White. Maximum likelihood estimation of misspecified models. *Econometrica*, 50(1):1–25, 1982.
- Frank D. McSherry. Privacy integrated queries: An extensible platform for privacy-preserving data analysis. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of Data*, SIGMOD '09, page 19–30, New York, NY, USA, 2009. Association for Computing Machinery.
- Ning Wang, Xiaokui Xiao, Yin Yang, Jun Zhao, Siu Cheung Hui, Hyejin Shin, Junbum Shin, and Ge Yu. Collecting and analyzing multidimensional data with local differential privacy. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, pages 638–649, 2019b.
- L. P. Barnes, W. N. Chen, and A. Özgür. Fisher information under local differential privacy. *IEEE Journal on Selected Areas in Information Theory*, 1(3): 645–659, 2020.
- Heysem Kaya, Pinar Tüfekci, and Erdinç Uzun. Predicting co and nox emissions from gas turbines: novel data and a benchmark pems. *Turkish Journal of Electrical Engineering & Computer Sciences*, 27(6):4783–4796, 2019.

Supplementary Material: One-bit Submission for Locally Private Quasi-MLE: Its Asymptotic Normality and Limitation

A SUMMARY OF NOTATION

A.1 Defined in Section 2

Q_{bf} is the bit flip, and $C_\epsilon = \frac{\epsilon^\epsilon + 1}{\epsilon^\epsilon - 1}$ is a value used to define the bit flip.

With $0 < \alpha < 1$, check loss for α -quantile is defined as

$$\rho_\alpha(\tau) = \begin{cases} (\alpha - 1)\tau & \text{if } \tau \leq 0, \\ \alpha\tau & \text{if } \tau > 0. \end{cases}$$

A.2 Defined in Section 3.1

Y and \mathbf{X} are objective and explanatory variables. $F_{\mathbf{X}}$ is the distribution of \mathbf{X} , and $F(\cdot|\mathbf{X} = \mathbf{x})$ is the distribution of Y conditioned by $\mathbf{X} = \mathbf{x}$. \mathcal{X} and \mathcal{Y} are the domains of \mathbf{X} and \mathbf{Y} . For each $i = 1, \dots, n$, (\mathbf{X}_i, Y_i) is an independent copy of (\mathbf{X}, Y) , which is possessed by the i th user. \bar{Y}_i is truncated version of Y_i , and $t(\cdot)$ is the truncating function mapping \mathcal{Y} into $[c_l, c_u]$, where c_l and c_u are real values such that $c_l < c_u$. Z_i is the perturbed version of Y_i , and its distribution is

$$p(Z_i = z|\mathbf{X}_i = \mathbf{x}) = \int Q_{\text{bf}}(z|t(y))dF(y|\mathbf{x}).$$

$\{F_\beta(\cdot|\mathbf{x}) : \beta \in \mathcal{B}, \mathbf{x} \in \mathcal{X}\}$ is the model family and \mathcal{B} is the parameter set. For each $\beta \in \mathcal{B}$, $p_\beta(z|\mathbf{x})$ is the density function which is obtained by replacing F by F_β in (7). In this Section 3.1, we write P and P_β to designate joint distributions $P(\mathbf{x}, z)$ and $P_\beta(\mathbf{x}, z)$ rather than conditional distributions $P(z|\mathbf{x})$ and $P_\beta(z|\mathbf{x})$.

$\Lambda_\epsilon(\beta, \mathbf{x}) = p_\beta(z = 1|\mathbf{x})$. Its first and second derivatives along β are denoted by $\Lambda'_\epsilon(\beta, \mathbf{x})$ and $\Lambda''_\epsilon(\beta, \mathbf{x})$.

A.3 Defined in Section 3.2

$F_{\mathbf{X}}$ is the model of \mathbf{X} . Q consists of $Q_{Z^{(Y)}}$ and $Q_{Z^{(X)}}$ perturbing Y_i and \mathbf{X}_i , respectively.

$\hat{p}(\mathbf{z}^{(\mathbf{X})})$ is the model of $\mathbf{z}^{(\mathbf{X})}$ written as

$$\hat{p}_X(\mathbf{z}^{(\mathbf{X})}) = \int Q_{Z^{(X)}}(\mathbf{z}^{(\mathbf{X})}|\mathbf{x})d\hat{F}(\mathbf{x}).$$

$\Phi(\beta, \mathbf{z}^{(\mathbf{X})})$ is the probability that $Z^{(Y)}|\mathbf{X} = \mathbf{x}$ is 1 when model F_β is correct:

$$\Phi(\beta, \mathbf{z}^{(\mathbf{X})}) = \frac{\hat{F}_{\mathbf{X}}(\Lambda_{\epsilon/(d+1)}(\beta, \mathbf{X})Q_{Z^{(X)}}(\mathbf{Z}^{(\mathbf{X})}|\mathbf{X}))}{\hat{p}_{\mathbf{Z}^{(\mathbf{X})}}(\mathbf{z}^{(\mathbf{X})})}.$$

Its first and second derivatives along β are denoted by $\Phi'(\beta, \mathbf{z}^{(\mathbf{X})})$ and $\Phi''(\beta, \mathbf{z}^{(\mathbf{X})})$.

B PSEUDO-CODE

Algorithm 1 and Algorithm 2 are the pseudo-codes of the protocols described in Section 3.1 and Section 3.2, respectively. In the for loops, the processing of each user does not need to be synchronized.

Algorithm 1: Protocol with Public \mathbf{X}

Input: Unknown distribution F , privacy parameter ϵ and \mathcal{B}

Curator set F_X ;

for $i = 1$ *to* n **do**

Curator generates $\mathbf{x}_i \sim F_X$;

Send \mathbf{x}_i to user i ;

User i generates $y_i \sim F(\cdot|\mathbf{x}_i)$;

User i computes \bar{y}_i as (6);

User i generates $z_i \sim Q_{\text{bf}}(\cdot|\bar{y}_i)$;

Send z_i to curator;

end

Let $D_n = \{(\mathbf{x}_i, z_i)\}_{i=1}^n$;

Curator computes $\ell(\beta; D)$;

Computes $\hat{\beta}_n = \arg \max_{\beta \in \mathcal{B}} \ell(\beta; D)$;

Output: $\hat{\beta}_n$

C MATHEMATICAL NOTES

C.1 for Section 3.1

C.1.1 Proof of Lemma 1

By the definition of $\Lambda_\epsilon(\beta, \mathbf{x})$, it is written as

$$\Lambda_\epsilon(\beta, \mathbf{x}) = p_\beta(Z = 1|\mathbf{X} = \mathbf{x}) = \int Q_{\text{bf}}(1|t(y))dF_\beta(y|\mathbf{x}).$$

From the definition of Q_{bf} , we have

$$\frac{1}{e^\epsilon + 1} \leq Q_{\text{bf}}(1|t(y)) \leq \frac{e^\epsilon}{e^\epsilon + 1}$$

for any $y \in \mathcal{Y}$. Thus, the following relation holds.

$$\Lambda_\epsilon(\beta, \mathbf{x}) \leq \int \frac{e^\epsilon}{e^\epsilon + 1} dF_\beta(y|\mathbf{x}) = \frac{e^\epsilon}{e^\epsilon + 1}.$$

The last equation is by the fact that F_β is a probability distribution. Similarly, we have

$$\Lambda_\epsilon(\beta, \mathbf{x}) \geq \frac{1}{e^\epsilon + 1}.$$

C.2 for Section 3.2

For each $\mathbf{z}^{(\mathbf{X})} \in \{-1, +1\}^d$, the curator considers the probability distribution of $\mathbf{Z}^{(\mathbf{X})}$ at $\mathbf{z}^{(\mathbf{X})}$ as

$$\hat{p}_{\mathbf{Z}}^{(\mathbf{X})}(\mathbf{z}^{(\mathbf{X})}) = \int Q_{\mathbf{Z}^{(\mathbf{X})}}(\mathbf{z}^{(\mathbf{X})}|\mathbf{x})d\hat{F}(\mathbf{x}).$$

The joint density is

$$p_\beta(\mathbf{z}^{(\mathbf{X})}, z^{(Y)}) = \int Q_{\mathbf{Z}^{(\mathbf{X})}}(\mathbf{z}^{(\mathbf{X})}|\mathbf{x})Q_{Z^{(Y)}}(z^{(Y)}|t(y))dF_\beta(y|\mathbf{x})d\hat{F}_X(\mathbf{x}).$$

The conditional distribution of $Z^{(Y)}$ is written as

$$p_\beta(z^{(Y)}|\mathbf{z}^{(\mathbf{X})}) = \frac{\hat{F}_X(p_\beta(z^{(Y)}|\mathbf{X})Q_{\mathbf{Z}^{(\mathbf{X})}}(\mathbf{z}^{(\mathbf{X})}|\mathbf{X}))}{\hat{p}_{\mathbf{Z}}^{(\mathbf{X})}(\mathbf{z}^{(\mathbf{X})})} = \Phi(\beta, \mathbf{z}^{(\mathbf{X})})^{z^{(Y)}}(1 - \Phi(\beta, \mathbf{z}^{(\mathbf{X})}))^{1-z^{(Y)}}.$$

Algorithm 2: Protocol with Private \mathbf{X}

Input: Unknown distribution F, F_X , privacy parameter ϵ and \mathcal{B}

Curator set \hat{F}_X ;

for $i = 1$ to n **do**

User i generates $\mathbf{x}_i \sim F_X$;

User i generates $y_i \sim F(\cdot | \mathbf{x}_i)$;

User i computes \bar{y}_i as (6);

User i generates $z_i^{(Y)} \sim Q_{\text{bf}}(\cdot | \bar{y}_i)$;

for $j = 1$ to d **do**

User i compute \bar{x}_{ij} as (6);

Generate $z_{ij}^{(\mathbf{X})} \sim Q_{\text{bf}}(\cdot | \bar{x}_{ij})$;

end

Let $\mathbf{z}_i^{(\mathbf{X})} = (z_{ij}^{(\mathbf{X})})_i$;

Send $(\mathbf{z}_i^{(\mathbf{X})}, z_i^{(Y)})$ to curator;

end

Let $D_n = \{(\mathbf{z}_i^{(\mathbf{X})}, z_i^{(Y)})\}_{i=1}^n$;

Curator computes $\ell(\beta; D)$;

Computes $\hat{\beta}_n = \arg \max_{\beta \in \mathcal{B}} \ell(\beta; D)$;

Output: $\hat{\beta}_n$

With Φ , the joint density is written as

$$p_\beta(\mathbf{z}^{(\mathbf{X})}, z^{(Y)}) = \Phi(\beta, \mathbf{z}^{(\mathbf{X})})^{z^{(Y)}} (1 - \Phi(\beta, \mathbf{z}^{(\mathbf{X})}))^{1-z^{(Y)}} \hat{p}_{\mathbf{Z}}^{(\mathbf{X})}(\mathbf{z}^{(\mathbf{X})}).$$

We analyze the sufficient conditions under which Assumptions 1 to 3 are satisfied while replacing F and F_θ in Theorem 1 with P and P_β . We adopt Assumptions 7 to 9 and 16. From these assumptions, it is obvious that $P(\mathbf{z}^{(\mathbf{X})}, z^{(Y)})$ and $P_\beta(\mathbf{z}^{(\mathbf{X})}, z^{(Y)})$ are measurable, and that density functions $p(\mathbf{z}^{(\mathbf{X})}, z^{(Y)})$ and $p_\beta(\mathbf{z}^{(\mathbf{X})}, z^{(Y)})$ exist.

The condition corresponding to Assumption 3 consists of two parts. The first part is the existence of $P \log p$ integrable function $h(\mathbf{z}^{(\mathbf{X})}, z^{(Y)})$ such that $|\log p_\beta(\mathbf{z}^{(\mathbf{X})}, z^{(Y)})| \leq h(\mathbf{z}^{(\mathbf{X})}, z^{(Y)})$ for all β . $P \log p$ is expanded as follows:

$$P \log p = P_{\mathbf{Z}^{(\mathbf{X})}}(P_{\cdot | \mathbf{Z}^{(\mathbf{X})}} \log p(\cdot | \mathbf{Z}^{(\mathbf{X})}) + \log p_{\mathbf{Z}^{(\mathbf{X})}}).$$

To evaluate the bound condition, it is necessary to analyze $p_{\mathbf{Z}^{(\mathbf{X})}}$ and $p(\cdot | \mathbf{Z}^{(\mathbf{X})})$.

Lemma 3. For any $\mathbf{z}^{(\mathbf{X})} \in \{-1, +1\}^d$,

$$\left(\frac{1}{e^{\epsilon/(d+1)} + 1} \right)^d \leq \hat{p}_{\mathbf{Z}^{(\mathbf{X})}}(\mathbf{z}^{(\mathbf{X})}) \leq \left(\frac{e^{\epsilon/(d+1)}}{e^{\epsilon/(d+1)} + 1} \right)^d.$$

$p_{\mathbf{Z}^{(\mathbf{X})}}(\mathbf{z}^{(\mathbf{X})})$ has the same bounds.

Lemma 4. For any $\beta \in \mathcal{B}$ and $\mathbf{z}^{(\mathbf{X})} \in \{-1, +1\}^d$,

$$\frac{1}{e^{\epsilon/(d+1)} + 1} \leq \Phi(\beta, \mathbf{z}^{(\mathbf{X})}) \leq \frac{e^{\epsilon/(d+1)}}{e^{\epsilon/(d+1)} + 1}.$$

With $0 < \epsilon < +\infty$ and $1 \leq d < +\infty$, $\Phi(\beta, \mathbf{z}^{(\mathbf{X})})$ are always bounded away from 0 and 1. Also,

$$\frac{1}{e^{\epsilon/(d+1)} + 1} \leq 1 - \Phi(\beta, \mathbf{z}^{(\mathbf{X})}) \leq \frac{e^{\epsilon/(d+1)}}{e^{\epsilon/(d+1)} + 1}.$$

By the above lemmas, $\log p(\cdot | \mathbf{z}^{(\mathbf{X})})$ and $\log p_{\mathbf{Z}^{(\mathbf{X})}}$ are always bounded away from $\pm\infty$, and $\log p(\cdot | \mathbf{z}^{(\mathbf{X})})$ and $\log p_{\mathbf{Z}^{(\mathbf{X})}}$ are always integrable with respect to P . The existence of integral function h is also obtained.

The second part is the uniqueness of the log-likelihood function. To guarantee that this property holds, we again adopt Assumption 11. Then, we have Theorem 5.

We next analyze the conditions under which Assumptions 4 to 6 are satisfied. The condition corresponding to Assumption 4 is the continuous differentiability of $\partial \log p_\beta / \partial \beta$. The partial derivative is expanded as

$$\frac{\partial}{\partial \beta} \log p_\beta(\mathbf{z}(\mathbf{X}), z^{(Y)}) = z^{(Y)} \frac{\Phi'(\beta, \mathbf{z}(\mathbf{X}))}{\Phi(\beta, \mathbf{z}(\mathbf{X}))} - (1 - z^{(Y)}) \frac{\Phi'(\beta, \mathbf{z}(\mathbf{X}))}{1 - \Phi(\beta, \mathbf{z}(\mathbf{X}))} = \frac{\Phi'(\beta, \mathbf{z}(\mathbf{X}))(z^{(Y)} - \Phi(\beta, \mathbf{z}(\mathbf{X})))}{\Phi(\beta, \mathbf{z}(\mathbf{X}))(1 - \Phi(\beta, \mathbf{z}(\mathbf{X})))}$$

where

$$\Phi'(\beta, \mathbf{z}(\mathbf{X})) \equiv \frac{\partial}{\partial \beta} \Phi(\beta, \mathbf{z}(\mathbf{X})).$$

By Lemma 4, $\Phi(\beta, \mathbf{z}(\mathbf{X}))$ always takes values greater than 0 and less than 1. So, if Assumption 17 holds, Assumption 4 is satisfied.

The condition corresponding to Assumption 5 is that there exist integrable functions with respect to P that upper bound the absolute values of each component of $\partial^2 \log p_\beta(\mathbf{z}(\mathbf{X}), z^{(Y)}) / \partial \beta^2$ and $(\partial \log p_\beta(\mathbf{z}(\mathbf{X}), z^{(Y)}) / \partial \beta) (\partial \log p_\beta(\mathbf{z}(\mathbf{X}), z^{(Y)}) / \partial \beta)^\top$. The second-order derivative is

$$\frac{\partial^2}{\partial \beta^2} \log p_\beta(\mathbf{z}(\mathbf{X}), z^{(Y)}) = (2z^{(Y)} - 1) \frac{\Phi''(\beta, \mathbf{z}(\mathbf{X}))}{\Psi_\epsilon(\beta^\top x)^z (1 - \Psi_\epsilon(\beta^\top x))^{1-z^{(Y)}}} - \frac{\Phi'(\beta, \mathbf{z}(\mathbf{X})) (\Phi'(\beta, \mathbf{z}(\mathbf{X})))^\top}{\Phi(\beta, \mathbf{z}(\mathbf{X}))^{2z^{(Y)}} (1 - \Phi(\beta, \mathbf{z}(\mathbf{X})))^{2(1-z^{(Y)})}}$$

where we define $\Phi''(\beta, \mathbf{z}(\mathbf{X})) \equiv \frac{\partial^2}{\partial \beta^2} \Phi(\beta, \mathbf{z}(\mathbf{X}))$. $(\partial \log p_\beta(\mathbf{z}(\mathbf{X}), z^{(Y)}) / \partial \beta) (\partial \log p_\beta(\mathbf{z}(\mathbf{X}), z^{(Y)}) / \partial \beta)^\top$ is

$$\left(\frac{\partial}{\partial \beta} \log p_\beta(\mathbf{z}(\mathbf{X}), z^{(Y)}) \right) \left(\frac{\partial}{\partial \beta} \log p_\beta(\mathbf{z}(\mathbf{X}), z^{(Y)}) \right)^\top = \left(\frac{z - \Phi(\beta, \mathbf{z}(\mathbf{X}))}{\Phi(\beta, \mathbf{z}(\mathbf{X}))(1 - \Phi(\beta, \mathbf{z}(\mathbf{X})))} \right)^2 \Phi'(\beta, \mathbf{z}(\mathbf{X})) \Phi'(\beta, \mathbf{z}(\mathbf{X}))^\top$$

By Lemma 4, Assumption 18 is sufficient to make the requirement hold.

The requirement corresponding Assumption 6 consists of three parts. The first part is that β is interior of \mathcal{B} . We assume this as Assumption 14. For enough large \mathcal{B} , this assumption is not particularly strong. Letting

$$A(\beta) = P \frac{\partial^2}{\partial \beta^2} \log p_\beta \quad \text{and} \quad B(\beta) = P \left(\frac{\partial}{\partial \beta} \log p_\beta \right) \left(\frac{\partial}{\partial \beta} \log p_\beta \right)^\top,$$

the second and third parts are the regularity of $A(\beta^*)$ and $B(\beta^*)$. We have already assumed that $A(\beta^*)$ is regular in Assumption 11. We consider the regularity of $B(\beta^*)$ here. $B(\beta)$ is

$$B(\beta) = \mathbb{E}_{\mathbf{Z}(\mathbf{x})} \left[\left(\frac{p(1|\mathbf{Z}(\mathbf{X}))}{\Phi(\beta, \mathbf{Z}(\mathbf{X}))^2} + \frac{p(0|\mathbf{Z}(\mathbf{X}))}{(1 - \Phi(\beta, \mathbf{Z}(\mathbf{X})))^2} \right) (\Phi'(\beta, \mathbf{Z}(\mathbf{X})) (\Phi'(\beta, \mathbf{Z}(\mathbf{X})))^\top) \right]$$

Since the scalar part is always finite and positive, Assumption 19 is a sufficient condition of the regularity of $B(\beta^*)$.

Summarizing the above discussions, we obtain Theorem 6.

C.2.1 Proof of Lemma 3

Proof. By definition, for any $\mathbf{z}(\mathbf{X}) \in \{-1, +1\}^d$, we have

$$\hat{p}_X(\mathbf{z}(\mathbf{X})) = \int Q_{\mathbf{Z}(\mathbf{x})}(\mathbf{z}(\mathbf{X})|\mathbf{x}) d\hat{F}(\mathbf{x}) \leq \int \left(\frac{e^{\epsilon/(d+1)}}{e^{\epsilon/(d+1)} + 1} \right)^d d\hat{F}(\mathbf{x}) = \left(\frac{e^{\epsilon/(d+1)}}{e^{\epsilon/(d+1)} + 1} \right)^d.$$

Similarly, we have

$$\hat{p}_X(\mathbf{z}(\mathbf{X})) \geq \left(\frac{1}{e^{\epsilon/(d+1)} + 1} \right)^d.$$

□

C.2.2 Proof of Lemma 4

Proof. By Lemma 1 and (8), we have

$$\Phi(\beta, \mathbf{z}^{(\mathbf{X})}) = \frac{\hat{F}_X \Lambda(\beta, \mathbf{X}) Q_{\mathbf{Z}^{(\mathbf{X})}}(\mathbf{z}^{(\mathbf{X})} | \mathbf{X})}{p_{\mathbf{Z}^{(\mathbf{X})}}(\mathbf{z}^{(\mathbf{X})})} \leq \frac{\hat{F}_X \frac{e^{\epsilon/(d+1)}}{e^{\epsilon/(d+1)} + 1} Q_{\mathbf{Z}^{(\mathbf{X})}}(\mathbf{z}^{(\mathbf{X})} | \mathbf{X})}{p_{\mathbf{Z}^{(\mathbf{X})}}(\mathbf{z}^{(\mathbf{X})})} = \frac{e^{\epsilon/(d+1)}}{e^{\epsilon/(d+1)} + 1}.$$

Similarly, we have

$$\Phi(\beta, \mathbf{z}^{(\mathbf{X})}) \geq \frac{1}{e^{\epsilon/(d+1)} + 1}.$$

Replacing \hat{F}_X by F_X , we obtain the arguments with regard to $p(z^{(Y)} | \mathbf{z}^{(\mathbf{X})})$. \square

C.3 for Section 4

In this section, we derive $\Psi_\epsilon(\theta)$ used in Section 4. As a consequence of the analysis, we obtain Lemma 2. We analyze the function in different three cases. The first case is the case where $c_1 < \theta < c_u$. For the sake of simplicity of notation, we let $G = \exp\left(-\frac{\alpha-1}{\sigma}(c_1 - \theta)\right)$ and $H = \exp\left(-\frac{\alpha}{\sigma}(c_u - \theta)\right)$. These values appear many times throughout the remainder of this section. First, we extend the probability $F_\theta(Y_i \leq c_1)$.

$$\begin{aligned} F_\theta(Y_i \leq c_1) &= \int_{-\infty}^{c_1} \frac{\alpha(1-\alpha)}{\sigma} \exp\left(-\rho\left(\frac{y_i - \theta}{\sigma}\right)\right) dy_i \\ &= \int_{-\infty}^{c_1} \frac{\alpha(1-\alpha)}{\sigma} \exp\left(-\frac{\alpha-1}{\sigma}(y_i - \theta)\right) dy_i \\ &= \left[\frac{\alpha(1-\alpha)}{\sigma} \left(-\frac{\sigma}{\alpha-1}\right) \exp\left(-\frac{\alpha-1}{\sigma}(y_i - \theta)\right) \right]_{-\infty}^{c_1} \\ &= \alpha \exp\left(-\frac{\alpha-1}{\sigma}(c_1 - \theta)\right) - \alpha \times 0 \\ &= \alpha \exp\left(-\frac{\alpha-1}{\sigma}(c_1 - \theta)\right) = \alpha G. \end{aligned}$$

Similarly, the probability $F_\theta(Y_i \geq c_u)$ is expanded as:

$$\begin{aligned} F_\theta(Y_i \geq c_u) &= \int_{c_u}^{+\infty} \frac{\alpha(1-\alpha)}{\sigma} \exp\left(-\rho\left(\frac{y_i - \theta}{\sigma}\right)\right) dy_i \\ &= \int_{c_u}^{+\infty} \frac{\alpha(1-\alpha)}{\sigma} \exp\left(-\frac{\alpha}{\sigma}(y_i - \theta)\right) dy_i \\ &= \left[\frac{\alpha(1-\alpha)}{\sigma} \left(-\frac{\sigma}{\alpha}\right) \exp\left(-\frac{\alpha}{\sigma}(y_i - \theta)\right) \right]_{c_u}^{+\infty} \\ &= -(1-\alpha) \times 0 + (1-\alpha) \exp\left(-\frac{\alpha}{\sigma}(c_u - \theta)\right) \\ &= (1-\alpha) \exp\left(-\frac{\alpha}{\sigma}(c_u - \theta)\right) = (1-\alpha)H. \end{aligned}$$

The probability $P_\theta(Z_i = 1)$ is written as follows:

$$\begin{aligned} P_\theta(Z_i = 1) &= \alpha G \left(\frac{1}{2} - \frac{1}{2C_\epsilon} \right) \\ &\quad + \frac{\alpha(1-\alpha)}{\sigma} \int_{c_1}^{\theta} \left(\frac{1}{2} + \frac{y_i - \frac{c_u+c_1}{2}}{C_\epsilon(c_u - c_1)} \right) \exp\left(-(\alpha-1)\frac{y_i - \theta}{\sigma}\right) dy_i \\ &\quad + \frac{\alpha(1-\alpha)}{\sigma} \int_{\theta}^{c_u} \left(\frac{1}{2} + \frac{y_i - \frac{c_u+c_1}{2}}{C_\epsilon(c_u - c_1)} \right) \exp\left(-\alpha\frac{y_i - \theta}{\sigma}\right) dy_i \\ &\quad + (1-\alpha)H \left(\frac{1}{2} + \frac{1}{2C_\epsilon} \right) \end{aligned} \tag{9}$$

Now, we extend each term.

$$\begin{aligned}
 & \frac{\alpha(1-\alpha)}{\sigma} \int_{c_1}^{\theta} \left(\frac{1}{2} + \frac{y_i - \frac{c_u+c_1}{2}}{C_\epsilon(c_u - c_1)} \right) \exp\left(-(\alpha-1)\frac{y_i - \theta}{\sigma}\right) dy_i \\
 &= \frac{\alpha(1-\alpha)}{\sigma} \left(\frac{1}{2} - \frac{c_u + c_1}{2C_\epsilon(c_u - c_1)} + \frac{\theta}{C_\epsilon(c_u - c_1)} \right) \int_{c_1}^{\theta} \exp\left(-(\alpha-1)\frac{y_i - \theta}{\sigma}\right) dy_i \\
 & \quad + \frac{\alpha(1-\alpha)}{\sigma} \frac{1}{C_\epsilon(c_u - c_1)} \int_{c_1}^{\theta} (y_i - \theta) \exp\left(-(\alpha-1)\frac{y_i - \theta}{\sigma}\right) dy_i \\
 &= \frac{\alpha(1-\alpha)}{\sigma} \left(\frac{1}{2} - \frac{c_u + c_1}{2C_\epsilon(c_u - c_1)} + \frac{\theta}{C_\epsilon(c_u - c_1)} \right) \frac{\sigma}{1-\alpha} (1-G) \\
 & \quad + \frac{\alpha(1-\alpha)}{2\sigma} \frac{1}{C_\epsilon(c_u - c_1)} \left(-\frac{\sigma^2}{(1-\alpha)^2} - \frac{\sigma}{1-\alpha} (c_1 - \theta)G + \frac{\sigma^2}{(1-\alpha)^2} G \right) \\
 &= \alpha \left(\frac{1}{2} - \frac{c_u + c_1}{2C_\epsilon(c_u - c_1)} + \frac{\theta}{C_\epsilon(c_u - c_1)} \right) (1-G) + \frac{\alpha}{C_\epsilon(c_u - c_1)} \left(-\frac{\sigma}{1-\alpha} - (c_1 - \theta)G + \frac{\sigma}{1-\alpha} G \right). \tag{10}
 \end{aligned}$$

Similarly,

$$\begin{aligned}
 & \frac{\alpha(1-\alpha)}{\sigma} \int_{\theta}^{c_u} \left(\frac{1}{2} + \frac{y_i - \frac{c_u+c_1}{2}}{C_\epsilon(c_u - c_1)} \right) \exp\left(-\alpha\frac{y_i - \theta}{\sigma}\right) dy_i \\
 &= -(1-\alpha) \left(\frac{1}{2} - \frac{c_u + c_1}{2C_\epsilon(c_u - c_1)} + \frac{\theta}{C_\epsilon(c_u - c_1)} \right) (H-1) + \frac{\alpha(1-\alpha)}{\sigma C_\epsilon(c_u - c_1)} \left(-\frac{\sigma}{\alpha} (c_u - \theta)H - \frac{\sigma^2}{\alpha^2} H + \frac{\sigma^2}{\alpha^2} \right) \\
 &= -(1-\alpha) \left(\frac{1}{2} - \frac{c_u + c_1}{2C_\epsilon(c_u - c_1)} + \frac{\theta}{C_\epsilon(c_u - c_1)} \right) (H-1) \\
 & \quad + \frac{1-\alpha}{C_\epsilon(c_u - c_1)} \left(-(c_u - \theta)H - \frac{\sigma}{\alpha} H + \frac{\sigma}{\alpha} \right). \tag{11}
 \end{aligned}$$

Substituting (10) and (11) into (9), we have

$$\begin{aligned}
 \Psi_\epsilon(\theta) &= P_\theta(Z_i = 1) \\
 &= \frac{\theta}{C_\epsilon(c_u - c_1)} + \frac{\alpha}{1-\alpha} \frac{\sigma}{C_\epsilon(c_u - c_1)} \exp\left(-\frac{\alpha-1}{\sigma}(c_1 - \theta)\right) - \frac{1-\alpha}{\alpha} \frac{\sigma}{C_\epsilon(c_u - c_1)} \exp\left(-\frac{\alpha}{\sigma}(c_u - \theta)\right) \\
 & \quad + \frac{1}{2} + \left(-\frac{\alpha}{1-\alpha} + \frac{1-\alpha}{\alpha} \right) \frac{\sigma}{C_\epsilon(c_u - c_1)} - \frac{c_u + c_1}{2C_\epsilon(c_u - c_1)}.
 \end{aligned}$$

The first and second derivatives are

$$\begin{aligned}
 \Psi'_\epsilon(\theta) &= \frac{1}{C_\epsilon(c_u - c_1)} - \frac{\alpha}{C_\epsilon(c_u - c_1)} \exp\left(\frac{1-\alpha}{\sigma}(c_1 - \theta)\right) - \frac{1-\alpha}{C_\epsilon(c_u - c_1)} \exp\left(-\frac{\alpha}{\sigma}(c_u - \theta)\right), \\
 \Psi''_\epsilon(\theta) &= \frac{\alpha(1-\alpha)}{\sigma C_\epsilon(c_u - c_1)} \left(e^{\frac{1-\alpha}{\sigma}(c_1 - \theta)} - e^{-\frac{\alpha}{\sigma}(c_u - \theta)} \right).
 \end{aligned}$$

By $c_1 < \theta < c_u$, $\frac{1-\alpha}{\sigma}(c_1 - \beta^\top \mathbf{x})$ and $-\frac{\alpha}{\sigma}(c_u - \beta^\top \mathbf{x})$ are always negative.

$$|\Psi'_\epsilon(\theta)| < \frac{1}{C_\epsilon(c_u - c_1)} \quad \text{and} \quad |\Psi''_\epsilon(\theta)| < \frac{\alpha(1-\alpha)}{\sigma C_\epsilon(c_u - c_1)}.$$

The second case is the case where $\theta \leq c_1$. $\Psi_\epsilon(\theta)$ is computed as

$$\begin{aligned}
 \Psi_\epsilon(\theta) &= \left(-(1-\alpha) \exp\left(-\alpha\frac{c_1 - \theta}{\sigma}\right) + 1 \right) \left(\frac{1}{2} - \frac{1}{2C_\epsilon} \right) \\
 & \quad + \frac{\alpha(1-\alpha)}{\sigma} \frac{1}{C_\epsilon(c_u - c_1)} \left(\frac{\sigma}{\alpha}(c_1 - \theta) + \frac{\sigma^2}{\alpha^2} \right) \exp\left(-\frac{\alpha}{\sigma}(c_1 - \theta)\right)
 \end{aligned}$$

$$\begin{aligned}
 & -\frac{\alpha(1-\alpha)}{\sigma} \frac{1}{C_\epsilon(c_u - c_l)} \left(\frac{\sigma}{\alpha}(c_u - \theta) + \frac{\sigma^2}{\alpha^2} \right) \exp\left(-\frac{\alpha}{\sigma}(c_u - \theta)\right) \\
 & + \frac{\alpha(1-\alpha)}{\sigma} \left(\frac{1}{2} + \frac{\theta}{C_\epsilon(c_u - c_l)} - \frac{c_u + c_l}{2C_\epsilon(c_u - c_l)} \right) \frac{\sigma}{\alpha} \left(\exp\left(-\frac{\alpha}{\sigma}(c_l - \theta)\right) - \exp\left(-\frac{\alpha}{\sigma}(c_u - \theta)\right) \right) \\
 & + (1-\alpha) \exp\left(-\frac{\alpha}{\sigma}(c_u - \theta)\right) \left(\frac{1}{2} + \frac{1}{2C_\epsilon} \right) \\
 = & \left(-(1-\alpha) \exp\left(-\alpha \frac{c_l - \theta}{\sigma}\right) + 1 \right) \left(\frac{1}{2} - \frac{1}{2C_\epsilon} \right) \\
 & + \frac{\alpha(1-\alpha)}{\sigma} \frac{1}{C_\epsilon(c_u - c_l)} \left(\frac{\sigma}{\alpha}c_l + \frac{\sigma^2}{\alpha^2} \right) \exp\left(-\frac{\alpha}{\sigma}(c_l - \theta)\right) \\
 & - \frac{\alpha(1-\alpha)}{\sigma} \frac{1}{C_\epsilon(c_u - c_l)} \left(\frac{\sigma}{\alpha}c_u + \frac{\sigma^2}{\alpha^2} \right) \exp\left(-\frac{\alpha}{\sigma}(c_u - \theta)\right) \\
 & + \frac{\alpha(1-\alpha)}{\sigma} \left(\frac{1}{2} - \frac{c_u + c_l}{2C_\epsilon(c_u - c_l)} \right) \frac{\sigma}{\alpha} \left(\exp\left(-\frac{\alpha}{\sigma}(c_l - \theta)\right) - \exp\left(-\frac{\alpha}{\sigma}(c_u - \theta)\right) \right) \\
 & + (1-\alpha) \exp\left(-\frac{\alpha}{\sigma}(c_u - \theta)\right) \left(\frac{1}{2} + \frac{1}{2C_\epsilon} \right) \\
 = & \frac{1}{2} - \frac{1}{2C_\epsilon} + \frac{(1-\alpha)\sigma}{\alpha} \frac{1}{C_\epsilon(c_u - c_l)} \exp\left(-\frac{\alpha}{\sigma}(c_l - \theta)\right) - \frac{(1-\alpha)\sigma}{\alpha} \frac{1}{C_\epsilon(c_u - c_l)} \exp\left(-\frac{\alpha}{\sigma}(c_u - \theta)\right).
 \end{aligned}$$

Its first and second derivatives are

$$\begin{aligned}
 \Psi'_\epsilon(\theta) &= \frac{1-\alpha}{C_\epsilon(c_u - c_l)} \left(\exp\left(-\frac{\alpha}{\sigma}(c_l - \theta)\right) - \exp\left(-\frac{\alpha}{\sigma}(c_u - \theta)\right) \right), \\
 \Psi''_\epsilon(\theta) &= \frac{\alpha(1-\alpha)}{\sigma C_\epsilon(c_u - c_l)} \left(\exp\left(-\frac{\alpha}{\sigma}(c_l - \theta)\right) - \exp\left(-\frac{\alpha}{\sigma}(c_u - \theta)\right) \right)
 \end{aligned}$$

Since $\theta \leq c_l$ and $c_u > c_l$, $\Psi'_\epsilon(\theta)$ is positive, and $\Psi''_\epsilon(\theta)$ is positive. Moreover, we have

$$|\Psi'_\epsilon(\theta)| < \frac{1-\alpha}{C_\epsilon(c_u - c_l)} \quad \text{and} \quad |\Psi''_\epsilon(\theta)| < \frac{\alpha(1-\alpha)}{\sigma C_\epsilon(c_u - c_l)}.$$

The last case is the case where $\theta \geq c_u$.

$$\begin{aligned}
 \Psi_\epsilon(\theta) &= \alpha \exp\left(\frac{1-\alpha}{\sigma}(c_l - \theta)\right) \left(\frac{1}{2} - \frac{1}{2C_\epsilon} \right) \\
 & - \frac{\alpha(1-\alpha)}{\sigma} \frac{1}{C_\epsilon(c_u - c_l)} \left(\frac{\sigma}{1-\alpha}(c_l - \theta) - \frac{\sigma^2}{(1-\alpha)^2} \right) \exp\left(\frac{1-\alpha}{\sigma}(c_l - \theta)\right) \\
 & + \frac{\alpha(1-\alpha)}{\sigma} \frac{1}{C_\epsilon(c_u - c_l)} \left(\frac{\sigma}{1-\alpha}(c_u - \theta) - \frac{\sigma^2}{(1-\alpha)^2} \right) \exp\left(-\frac{1-\alpha}{\sigma}(c_u - \theta)\right) \\
 & + \frac{\alpha(1-\alpha)}{\sigma} \left(\frac{1}{2} + \frac{\theta}{C_\epsilon(c_u - c_l)} - \frac{c_u + c_l}{2C_\epsilon(c_u - c_l)} \right) \frac{\sigma}{1-\alpha} \left(-\exp\left(\frac{1-\alpha}{\sigma}(c_l - \theta)\right) + \exp\left(\frac{1-\alpha}{\sigma}(c_u - \theta)\right) \right) \\
 & + \left(1 - \alpha \exp\left(\frac{1-\alpha}{\sigma}(c_u - \theta)\right) \right) \left(\frac{1}{2} + \frac{1}{2C_\epsilon} \right) \\
 = & \alpha \exp\left(\frac{1-\alpha}{\sigma}(c_l - \theta)\right) \left(\frac{1}{2} - \frac{1}{2C_\epsilon} \right) \\
 & - \frac{\alpha(1-\alpha)}{\sigma} \frac{1}{C_\epsilon(c_u - c_l)} \left(\frac{\sigma}{1-\alpha}c_l - \frac{\sigma^2}{(1-\alpha)^2} \right) \exp\left(\frac{1-\alpha}{\sigma}(c_l - \theta)\right) \\
 & + \frac{\alpha(1-\alpha)}{\sigma} \frac{1}{C_\epsilon(c_u - c_l)} \left(\frac{\sigma}{1-\alpha}c_u - \frac{\sigma^2}{(1-\alpha)^2} \right) \exp\left(-\frac{1-\alpha}{\sigma}(c_u - \theta)\right) \\
 & + \frac{\alpha(1-\alpha)}{\sigma} \left(\frac{1}{2} - \frac{c_u + c_l}{2C_\epsilon(c_u - c_l)} \right) \frac{\sigma}{1-\alpha} \left(-\exp\left(\frac{1-\alpha}{\sigma}(c_l - \theta)\right) + \exp\left(\frac{1-\alpha}{\sigma}(c_u - \theta)\right) \right)
 \end{aligned}$$

$$\begin{aligned}
 & + \left(1 - \alpha \exp\left(\frac{1-\alpha}{\sigma}(c_u - \theta)\right)\right) \left(\frac{1}{2} + \frac{1}{2C_\epsilon}\right) \\
 & = \frac{1}{2} + \frac{1}{2C_\epsilon} + \frac{\alpha\sigma}{1-\alpha} \frac{1}{C_\epsilon(c_u - c_1)} \exp\left(\frac{1-\alpha}{\sigma}(c_1 - \theta)\right) - \frac{\alpha\sigma}{1-\alpha} \frac{1}{C_\epsilon(c_u - c_1)} \exp\left(-\frac{1-\alpha}{\sigma}(c_u - \theta)\right).
 \end{aligned}$$

Its first and second derivatives are

$$\begin{aligned}
 \Psi'_\epsilon(\theta) & = \frac{\alpha}{C_\epsilon(c_u - c_1)} \left(-\exp\left(\frac{1-\alpha}{\sigma}(c_1 - \theta)\right) + \exp\left(\frac{1-\alpha}{\sigma}(c_u - \theta)\right)\right), \\
 \Psi''_\epsilon(\theta) & = \frac{\alpha(1-\alpha)}{\sigma C_\epsilon(c_u - c_1)} \left(\exp\left(\frac{1-\alpha}{\sigma}(c_1 - \theta)\right) - \exp\left(\frac{1-\alpha}{\sigma}(c_u - \theta)\right)\right).
 \end{aligned}$$

Since $\theta \geq c_u$ and $c_u > c_1$, $\Psi'_\epsilon(\theta)$ is positive, and $\Psi''_\epsilon(\theta)$ is negative. Moreover, since $(1-\alpha)(c_1 - \theta)/\sigma < (1-\alpha)(c_u - \theta)/\sigma \leq 0$, we have

$$|\Psi'_\epsilon(\theta)| < \frac{\alpha}{C_\epsilon(c_u - c_1)} \quad \text{and} \quad |\Psi''_\epsilon(\theta)| < \frac{\alpha(1-\alpha)}{\sigma C_\epsilon(c_u - c_1)}.$$

We also analyze their behavior on the boundaries. $\Psi_\epsilon(\theta)$ is continuous at $\theta = c_1$ and c_u if and only if $\lim_{\theta \downarrow c_u} \Psi_\epsilon(\theta) = \lim_{\theta \uparrow c_u} \Psi_\epsilon(\theta)$ and $\lim_{\theta \downarrow c_1} \Psi_\epsilon(\theta) = \lim_{\theta \uparrow c_1} \Psi_\epsilon(\theta)$. As we see below, these equations hold.

$$\begin{aligned}
 \lim_{\theta \downarrow c_u} \Psi_\epsilon(\theta) & = \lim_{\theta \uparrow c_u} \Psi_\epsilon(\theta) = \frac{1}{2} + \frac{1}{2C_\epsilon} + \frac{\alpha\sigma}{1-\alpha} \frac{1}{C_\epsilon(c_u - c_1)} \exp\left(\frac{1-\alpha}{\sigma}(c_1 - c_u)\right) - \frac{\alpha\sigma}{1-\alpha} \frac{1}{C_\epsilon(c_u - c_1)}, \\
 \lim_{\theta \downarrow c_1} \Psi_\epsilon(\theta) & = \lim_{\theta \uparrow c_1} \Psi_\epsilon(\theta) = \frac{1}{2} - \frac{1}{2C_\epsilon} + \frac{(1-\alpha)\sigma}{\alpha} \frac{1}{C_\epsilon(c_u - c_1)} - \frac{(1-\alpha)\sigma}{\alpha} \frac{1}{C_\epsilon(c_u - c_1)} \exp\left(-\frac{\alpha}{\sigma}(c_u - c_1)\right).
 \end{aligned}$$

We next evaluate the existence of first and second derivatives at $\theta = c_1$ and c_u .

$$\begin{aligned}
 \lim_{\theta \downarrow c_u} \Psi'_\epsilon(\theta) & = \lim_{\theta \uparrow c_u} \Psi'_\epsilon(\theta) = \frac{\alpha}{C_\epsilon(c_u - c_1)} \left(-\exp\left(\frac{1-\alpha}{\sigma}(c_1 - c_u)\right) + 1\right), \\
 \lim_{\theta \downarrow c_1} \Psi'_\epsilon(\theta) & = \lim_{\theta \uparrow c_1} \Psi'_\epsilon(\theta) = \frac{1-\alpha}{C_\epsilon(c_u - c_1)} \left(1 - \exp\left(-\frac{\alpha}{\sigma}(c_u - c_1)\right)\right), \\
 \lim_{\theta \downarrow c_u} \Psi''_\epsilon(\theta) & = \lim_{\theta \uparrow c_u} \Psi''_\epsilon(\theta) = \frac{\alpha(1-\alpha)}{C_\epsilon(c_u - c_1)} \left(\exp\left(\frac{1-\alpha}{\sigma}(c_1 - c_u)\right) - 1\right), \\
 \lim_{\theta \downarrow c_1} \Psi''_\epsilon(\theta) & = \lim_{\theta \uparrow c_1} \Psi''_\epsilon(\theta) = \frac{\alpha(1-\alpha)}{\sigma C_\epsilon(c_u - c_1)} \left(1 - \exp\left(-\frac{\alpha}{\sigma}(c_u - c_1)\right)\right).
 \end{aligned}$$

D COMPARISON WITH NON-PRIVATE ESTIMATOR

For comparison with existing work, we also consider the correct model case.

Assumption 20. Given $\mathbf{x} \in \mathbb{R}^d$, Y is a random variable sampled from the asymmetric Laplace distribution $f(\cdot; \alpha, \beta^\top \mathbf{x}, \sigma)$, which is defined in (2). For each $i \in [n]$, y_i is a realization of random variable Y_i that is a copy of Y .

Under this condition, Corollary 1 is more specified.

Corollary 3. Suppose Assumptions 7, 10, 11 and 20 hold. The MLE $\hat{\beta}_n$ is distributed asymptotically normally as $\sqrt{n}(\hat{\beta}_n - \beta^*) \rightarrow \mathcal{N}(0_d, I_{\beta^*}^{-1})$ where $I_{\beta^*} = F_{\mathbf{X}} \frac{\Psi'_\epsilon(\beta^{*\top} \mathbf{X})^2}{\Psi_\epsilon(\beta^{*\top} \mathbf{X})(1 - \Psi_\epsilon(\beta^{*\top} \mathbf{X}))} \mathbf{X} \mathbf{X}^\top$.

To obtain an intuitive understanding of the result, we roughly compare the Fisher information matrix derived in Corollary 3 and the non-private Fisher matrix (5), and analyze some extreme cases. First, we consider the concentrated case in which the scale parameter σ is extremely small. For a σ sufficiently small that $\sigma \ll |(1-\alpha)(c_1 - \beta^{*\top} \mathbf{x})|$ and $\sigma \ll |\alpha(c_u - \beta^{*\top} \mathbf{x})|$ for most \mathbf{x} ,

$$\Psi(\beta^{*\top} \mathbf{x}) \approx \frac{1}{2} + \left(-\frac{\alpha}{1-\alpha} + \frac{1-\alpha}{\alpha}\right) \frac{\sigma}{2C_\epsilon} + \frac{\beta^{*\top} \mathbf{x}}{2C_\epsilon} \quad \text{and} \quad \Psi'(\beta^{*\top} \mathbf{x}) \approx \frac{1}{2C_\epsilon}.$$

Thus,

$$\frac{\Psi'(\beta^{*\top} \mathbf{x})^2}{\Psi(\beta^{*\top} \mathbf{x})(1 - \Psi(\beta^{*\top} \mathbf{x}))} \approx \frac{1}{C_\epsilon^2 - \left(\frac{1-2\alpha}{\alpha(1-\alpha)}\sigma + \beta^{*\top} \mathbf{x}\right)^2} \geq \frac{1}{C_\epsilon^2 - \left(\frac{1-2\alpha}{\alpha(1-\alpha)}\sigma + c_1\right)^2}.$$

In comparing this with (4), we can see that the Fisher information matrix of our LDP estimator is $\Omega\left(\epsilon^2 \frac{\sigma^2}{\alpha(1-\alpha)}\right)$ times smaller than that of the non-private estimator as $\epsilon \downarrow 0$. This lower bound agrees with the complexity of ϵ but is $\sigma^2/\alpha(1-\alpha)$ times lower. Since we assumed that σ is small, this gap can be large. Although our MLE tends to lose more information regarding the structure of f_{β^*} than an optimal MLE, it experiences minimum information loss due to perturbation for privacy.

We omit the comparison of the MLE of the regression coefficient with the private \mathbf{X} . The Fisher information matrix strongly depends on the structure of the distribution of \mathbf{X} . We have no informative comparison in this case.

E ADDITIONAL NUMERICAL EVALUATION

In this section, we perform some additional numerical evaluations with the real data, which is the same data used in Section 5.

We implemented our simulation in Python. The necessary packages are written in requirements.txt. The main part is written in experiment.py. We made the Jupyter notebook files corresponding to each numerical evaluation. Visualization of the results is also in the Jupyter notebook files. You can open these files and run the simulations on your Jupyter notebook or Jupyter Lab.

Our supplemental material does not contain the real data used in the numerical evaluation. Before running our program, please download the data from "https://archive.ics.uci.edu/ml/datasets/Gas+Turbine+CO+and+NOx+Emission+Data+Set". Then, put them into the folder "data/emission/".

E.1 Evaluation of Private \mathbf{X}

Here, we observe the behavior of our QMLE for the private \mathbf{X} scenario, which is described in Section 4.2.

Due to implementation needs, we have made some modifications to the description in the main part. First, we made some changes to $\Phi(\beta, \mathbf{z}^{(\mathbf{X})})$. Theoretically, Φ and $1 - \Phi$ never take negative values. However, we found that the value of Φ can exceed 1 by a small amount due to rounding error. Then, $1 - \Phi$ is negative, and the computation corrupts since the log function is inputted a negative value. To avoid this undesirable situation, we multiplied Φ by $e^{-0.000001}$.

Second, we changed the domain of $\hat{F}_{\mathbf{X}}$ because no element of each \mathbf{x}_i is in the interval $[-1, 1]$. In the simulation, each user truncates the components of \mathbf{X}_i into the intervals $[5, 10]$, $[1000, 1030]$, $[70, 100]$, $[4, 6]$, $[20, 30]$, $[1000, 1100]$, $[530, 570]$, $[130, 170]$, and $[10, 15]$. We recommend that the curators should set the intervals with the help of experts when they use our algorithm in reality.

We observe the covariance matrices for $n \in \{100, 1000, 10000\}$ and $\epsilon \in \{5.0, 10, 25\}$ with $\alpha = 0.3$ and $\sigma = 1.0$. For each combination of n and ϵ , we sub-sample n records 1,000 times without replacement from the 36,733 records. For each sub-data, we simulate the protocol described in Section 4.2 and obtain a QMLE. Then, we compute the Frobenius norm of covariance matrices of the 1,000 QMLEs,

Fig. 2 shows the result. The horizontal and vertical axes show n and the value of each Frobenius norm in log-scale, respectively. For each ϵ , with large n , the norm of the covariance matrix is smaller. The decreasing speed is $O(1/n)$. These properties are similar to those in the public \mathbf{X} scenario, which is described in Section 5.

E.2 Evaluation of Effect of Truncation

In this subsection, we evaluate the effect of the truncation in the public \mathbf{X} scenario.




Figure 2: Frobenius norm of covariance matrices in private \mathbf{X} scenario. The norm decrease in proportion to $1/n$ for each ϵ .




Figure 3: Frobenius norm of covariance matrices with various $[c_l, c_u]$ s. A smaller interval makes the norm smaller.

With $\epsilon = 2.5$ and $n = 10,000$, we try intervals $[50, 100]$, $[40, 110]$, $[30, 120]$ and $[20, 130]$ for the truncation. The other setting is the same as Section 5.

Fig. 3 shows the result. A shorter interval makes the estimators more concentrated. We remark that the concentration does not necessarily imply a good approximation of the true distribution. In general, there is a trade-off between bias and variance.

E.3 Comparison with Non-private Estimator

In this subsection, we evaluate the difference between the centers of the distributions of our QMLEs and the non-private QMLEs which is described in Section 2.3. Our theoretical result does not say that those QMLEs converge to the same point. Thus, we consider it with numerical simulations.

First, we observe the behavior of the non-private QMLE. Fig. 4 shows the Frobenius norm of covariance matrices. It is seen that the non-private QMLEs converge to one point. We treat the average vector of the non-private QMLEs with $n = 30,000$ as the grand truth in the main observation as described below. We remark that the "grand truth" can be biased.

We use the same simulation result used in Section 5. We compute the difference of the average vector of our QMLE and the grand truth and observe the norm for each n and ϵ .

Fig. 5 shows the main result. The horizontal and vertical axes show n and the value of the norm of the covariance matrices, respectively. The bias is not zero for all ϵ . Smaller ϵ tends to give smaller bias. It is seen that n does not affect the bias. This result implies that the non-private QMLE and our QMLE can converge to different points.




Figure 4: Frobenius norm of covariance matrices of non-private QMLE. It seems that the non-private QMLEs converge to one point.




Figure 5: Norm of difference between the centers of non-private QMLEs and our QMLEs with various n and ϵ . The difference does not depend on n .