Identification and Formal Privacy Guarantees

Tatiana Komarova[a] and **Denis Nekipelov**[b*]

This version: April 2021

Abstract

Abstract: Empirical economic research crucially relies on highly sensitive individual datasets. At the same time, increasing availability of public individual-level data that comes from social networks, public government records and directories makes it possible for adversaries to potentially de-identify anonymized records in sensitive research datasets. Most commonly accepted formal definition of an individual non-disclosure guarantee is referred to as differential privacy. With differential privacy in place the researcher interacts with the data by issuing queries that evaluate the functions of the data. Differential privacy guarantee is achieved by replacing the actual outcome of the query with a randomized outcome with the amount of randomness determined by the sensitivity of the outcome to individual observations in the data.

While differential privacy does provide formal non-disclosure guarantees, its impact on the identification of empirical economic models as well as its impact on the performance of estimators in nonlinear empirical Econometric models has not been sufficiently studied. Since privacy protection mechanisms are inherently finite-sample procedures, we define the notion of identifiability of the parameter of interest under differential privacy as a property of the limit of experiments. It is naturally characterized by the concepts from the random sets theory and is linked to the asymptotic behavior in measure of differentially private estimators.

We demonstrate that particular instances of regression discontinuity design may be problematic for inference with differential privacy. Those parameters turn out to be neither point nor partially identified. The set of differentially private estimators converges weakly to a random set. This result is clearly supported by our simulation evidence. Our analysis suggests that many other estimators that rely on nuisance parameters may have similar properties with the requirement of differential privacy. Identification becomes possible if the target parameter can be deterministically localized within the random set. In that case, a full exploration of the random set of the weak limits of differentially private estimators can allow the data curator to select a sequence of instances of differentially private estimators that is guaranteed to converge to the target parameter in probability. We provide a decision-theoretic approach to this selection.

JEL Classification: C35, C14, C25, C13.

Keywords: Differential privacy; average treatment effect; regression discontinuity; random sets; identification

---

*[a] Department of Economics, London School of Economics and Political Science and [b] Department of Economics and Computer Science, University of Virginia.

Support from STICERD and the NSF is gratefully acknowledged.

1 Introduction

A large portion of empirical work in Social Sciences and most notably in Economics relies on *highly sensitive* data. Sensitivity of the data may be well understood from the personal experiential perspective and is largely associated with risks of potential exposure of individuals in the data when the data may reveal their personal or financial information which consequently would make them sensitive to the adversaries or embarrass them. At the same time, there has been a long struggle in attempts to formalize the concept of sensitivity of the data and the related concept of "privacy guarantee" that would measure how and to what extent sensitive attributes of the data are protected.

The most significant progress in the efforts to formalize privacy protection has been made in the Computer Science literature. The mainstream approach there would consider any attribute of the data as "sensitive." At the same time, the security risks are considered in the worst-case scenario setting where public components of the protected dataset or its summaries are accessed by an adversary. The adversary is assumed to have an arbitrary amount of auxiliary information that he or she can use to expose individuals in the data.

In order to be able to talk about privacy-preserving approaches and their properties, one has to be able to characterize formally the level of exposure induced by the adversary. This can done, for instance, in accordance with the US Census Bureau's analysis in [34] which distinguishes between *identity disclosure*, where an adversary is able to identify if a specific data entry belongs to a certain individual, and *attribute disclosure* where an adversary is able to find out if a particular individual has a particular characteristic (e.g. belongs to a certain group). Formal legal approaches to protection of individual data, for instance, behind HIPAA and FERPA are based on clear prioritization of identity disclosure over attribute disclosure and mandate the removal of specific demographic and personal identifiers to make the data "less sensitive." Computer Science literature, at the same time, clearly demonstrates that the reduction in "sensitivity" of the data based on potential exposure level by removing individual characteristics is highly ineffective.[1] The examples of successful attacks on the "anonymized" data lead to the early work on the formal definition of privacy guarantees that resulted, in particular, in the development and the implementation (see [50], [51], [52], [36], [6], [37], [13], among others) of the so-called k-anonymity approach. A database instance is said to provide k-anonymity, for some number k, if every way of singling an individual out of the database returns records for at least k individuals. In other words, anyone whose information is stored in the

[1] For instance, [52] identified the medical records of William Weld, then governor of Massachusetts, by linking voter registration records to "anonymized" Massachusetts Group Insurance Commission (GIC) medical encounter data, which retained the birthdate, sex, and zip code of the patient. In another example in [46] the risk of disclosure was identified in the so-called "Netflix prize dataset." In 2009 Netflix announced a competition with a grand prize of $1M for developers of the prediction algorithm that would use the information on past viewership history for given consumer and how this consumer rated those movies and be be able to predict how that consumer would rate the movies that he or she has not seen yet. For this competition Netflix released an "anonymized" dataset containing 100,480,507 ratings produced by 480,189 consumers. [46] used the public movie review data of users from imdb.com and were able to link a significant fraction of "anonymous" consumers on Netflix to imdb.com users based on the uniqueness of watch histories.

database can be "confused" with k others. It is important to note that the k-anonymity approach was primarily targeted to prevent the identity disclosure and, as is well known, does not prevent the attribute disclosure. It also implicitly requires that the data curator responsible for protecting the data is aware of all possible auxiliary information that could be available to the adversary. These features make k-anonymity and its variants rather impractical.

The concept of *differential privacy* developed in [14] provides a measure of privacy guarantees without these complications and addresses both identity and attribute disclosure concerns. Differential privacy formalizes the interaction of a user with a "sensitive" database via queries that are submitted through a secure server with a privacy protection mechanism. These queries are functions that need to be computed on the data representing summaries or tabulations of the data. The assumed risk to the database is that the queries to the database can be issued by an adversary. The privacy protection mechanism has two main elements that allow it to be effective against arbitrary adversaries. First, it takes into account the "sensitivity" of the query to the data, that measures the maximum change in the output of the function computed on the data if any data entry is deleted or altered. Second, it independently randomizes the outcome of the query which ensures that the produced outcome is not correlated with any auxiliary information that an adversary may have. Differential privacy ensures that no single individual/observation in the data can have a significant impact on the distribution of the randomized query and, thus, it gives a broad set of guarantees of data protection from adversarial attacks.

Given the universal theoretical appeal of differential privacy, its adoption as a security standard has been considered by a variety of private enterprises and government bodies. For instance, Facebook recently announced ([4]) that it will provide public access to the data on shared links within the social network via a differentially private protocol. In 2019 the US Census bureau has announced ([2]) that it will use differential privacy as the baseline for privacy protection in 2020 Census. The transition to the differential privacy standards has caused an outcry in Social Science community due to obvious concerns about the impact of the differential privacy technology on the quality of the data-driven research and resulted in an open letter signed by a large number of academic researchers ([3]).

In the evaluation of privacy technologies the Computer Science literature focuses on the "privacy-utility" trade-off. In this literature, e.g. discussed in [38], [28], [9], [56]), the main idea is that privacy protection makes the data "noisier" which reduces its "utility" but at the same time does not preclude one from recovering the parameter of interest from the data, albeit possibly less accurately. In our previous work [32] we showed that this thought process is flawed for techniques in the spirit of k-anonymity, which are aimed at the prevention of the identity disclosure. In particular, [32] argues that since privacy protection is fundamentally a *finite sample* paradigm, its impact on the properties of estimators has to be studied using the concept of limits of experiments. From this perspective, the asymptotic behavior of the estimator of interest is viewed as the limit of experiments where an estimator with privacy constraints is produced from samples of an increasing size. As it will be

clear, this is a general approach we will also take in this paper albeit with some very important differences in technical and practical implementation. Also, our earlier work [32] (as well as in [30], [31]) does not analyze differential privacy as it is focused on the techniques in the spirit of k-anonymity in a framework where the data needed for estimation has to be combined from split datasets. [32] in particularly shows that the prevention of the *identity disclosure* when combining split datasets using quasi-identifiers is incompatible with the parameter identification in conditional moment models, which comprise many commonly used econometric models. Namely, in [32] with any degree of identity disclosure prevention imposed in the data one can only recover a *pseudo-identified set* for the parameter of interest. This set is generally a non-singleton subset of the parameter space and it does not include the true parameter. In other words, the "privacy-utility" trade-off is meaningless if one wants to enforce guarantees in the spirit of k-anonymity.

In this paper we develop a notion of identification under differential privacy. Differential privacy guarantee is provided for finite samples whereas, as is well known, identification in Economics and Econometrics is a population property. Thus, to bring these two worlds – the population and finite sample – together, we propose to look at the identification (or lack of such) as the property of the limit of experiments.[2] We first consider a set of estimators that can obtained in a finite sample by applying differentially private mechanisms for a particular privacy guarantee. We then analyze the limit of this set as the sample size increases and associate parameter identification with limit being a singleton.

In this paper we demonstrate that certain important classes of estimators may not be identified with differential privacy in place. Moreover, they are also not partially identified, i.e. contained in a non-trivial subset of the parameter space with probability 1. Finite sample differentially private estimators turn out to have limits within a random set and thus cannot be analyzed via traditional machinery used for inference in point or partially identified models. Instead, the analysis has to rely on the well-developed theory of random sets of [44] and [45].

This result stems from the fundamental asymmetry of information between the data curator and the researcher. Since the researcher cannot "explore" the data, she does not know the properties of its distribution. As a result, some statistical functions that the researcher would like to evaluate may not be accurately tuned to that distribution. Given this simple intuition, we emphasize that our findings do not directly parallel existing impossibility results for differential privacy (e.g. [17] and [26]) since those largely stem from the statistical lower bounds for inference of functions from data. They instead arise from the asymptotic limits of finite sample data distribution and the noise-infused differentially private output.

More importantly, point identification in our scenario is not impossible and can be achieved if the data curator is willing to collaborate with the researcher. If the data curator is able to locate the

[2] We used a similar idea in our previous work [32] for privacy under data combination. However, the identification approach we suggest under differential privacy is different from that in [32], which is not surprising given that differentially privacy is a fundamentally distinct approach for privacy protection.

parameter of interest relative to the limit random set of differentially private outputs, then she can select an element of that limit random set which is the closest to the parameter of interest. This allows her to construct a mapping selecting an element of the random set that results for any finite data sample. We show that such a mapping accommodates the point identification of the target parameter.

In the setting where the data curator cannot construct such a mapping but is able to convey the procedure she used to produce the differentially private output to the data curator, point identification may not be possible. However, the researcher can construct bounds for the target parameter via Bayesian credible regions. These can serve as a simple to produce characterization of the differentially private output without requiring the researcher to analyze the entire probability distribution of the resulting random set.

In Section 2 we formalize the concept of identification with differential privacy guarantees and provide formal characterization of the limit of the set of differentially private estimators applied to finite samples.

In Section 3 we show how parameter identification can be restored with collaboration between the researcher and the data curator and provide a decision-theoretic framework that leads to a consistent estimator for the parameter of interest from the set of differentially private outputs. We also show how a researcher can construct probabilistic bounds for the parameter of interest when the data curator is not willing to collaborate with her to point identify the target parameter.

In Section 4 we conduct the identification analysis for differentially private estimators in regression discontinuity design models. We establish the lack of identification for differentially private versions of nonparametric regression at the boundary and local linear estimators. We also discuss implications of differential privacy for specification tests in the RDD framework and present some simulation evidence.

Section 5 concludes.

The proofs of all the results as well as additional examples are collected in the Appendix.

2 Limit Properties of Estimators with Formal Privacy Guarantees

2.1 Preliminaries

We start be reviewing a formal definition of differential privacy.

When the query to the data is an estimator $\widehat{\theta}$, differential privacy requires to replace the point estimator of interest $\widehat{\theta}$ with a randomized function $\theta(\mathbb{P}_N, \nu_N)$, where $\mathbb{P}_N$ is the empirical distribution of the dataset and $\nu_N \in \mathcal{V}$ is an independent random element which is assumed to be belong to a Banach space $\mathcal{V}$. In other words, it is required that the estimator is randomized using independent noise. The independence of the random element ν from the data sample is essential, since any

correlation may allow an adversary to recover potentially sensitive attributes of observations in the data.

We now provide a formal definition of differential privacy in application to randomized estimator $\theta(\mathbb{P}_N, \nu_N)$.

DEFINITION 1 ((ϵ, δ)-differential privacy [14]). *A randomized estimator $\theta(\mathbb{P}_N, \nu_N)$ is (ϵ, δ)-differentially private if for any two empirical measures $\mathbb{P}_N$ and $\mathbb{P}'_N$ over N support points and differing in exactly two support points, we have that for all measurable sets A of possible outputs the following holds:*

$$\mathbb{P}_{\nu_N}\left[\theta(\mathbb{P}_N, \nu_N) \in A\right] \leq e^{\epsilon}\mathbb{P}_{\nu_N}\left[\theta(\mathbb{P}'_N, \nu_N) \in A\right] + \delta, \tag{2.1}$$

where $\epsilon > 0$, $\delta \in [0, 1)$ are privacy parameters (sometimes we will refer to them as privacy constants and probabilities are taken over randomness in ν_N.

In addition, if $\delta = 0$, then the estimator $\theta(\mathbb{P}_N, \nu_N)$ is referred to as ϵ-differentially private.

In this definition we use notation $\mathbb{P}_{\nu}(\cdot)$ to emphasize that differentially private estimator $\theta(\mathbb{P}_N, \nu_N)$ is based on the distribution of random element ν_N while the distributions of two adjacent datasets $\mathbb{P}_N$ and $\mathbb{P}'_N$ are fixed. The randomized estimator $\theta(\mathbb{P}_N, \nu_N)$ in Definition 1 ensures that information regarding individual data entry cannot be reverse-engineered from its values. We treat ν_N as a "seed" for randomness represented by a fixed standardized random variable that is then transformed by the estimator into the random variable used in a particular mechanism for differential privacy.

There are two important features of the privacy preserving methodology we want to emphasize. The first one is that even though the distribution of ν_N can be adjusted to the general properties of the *population* data distribution in $\mathbb{P}_N$ (such as the the number of variable and the support variables) and can depend on the sample size N, it may not depend on the specific values of observations producing $\mathbb{P}_N$. The second feature is that the privacy protection has to be guaranteed for *every* possible realizations of the data $\mathbb{P}_N$ and $\mathbb{P}'_N$ different in one support point no matter how small the probabilities of realizations of these datasets are. These two features contribute to the powerful privacy-preserving framework delivered by differential privacy. However, as we will see later, these two properties will contribute to the possible lack of identification of parameters in differentially private versions of some important econometric models (examples of those are given in section 4).

One seemingly effective way to mitigate the randomness introduced by ν_N is by considering repeatedly sampled versions of the estimator $\theta(\mathbb{P}_N, \nu_N)$. This, however, is impossible since differential privacy would require one to consider *the entire set* of those queries *as a single query.* This is often implemented by allocating a *privacy budget* ϵ^* to all statistical functions that would ever be evaluated from a given dataset. Each function would be allocated a fraction of that ϵ^* for the corresponding (ϵ, δ) - differentially private output (desirably, with $\epsilon \ll \epsilon^*$) and the more functions are evaluated from the data, the less of the the privacy budget ϵ^* will be allocated to the remaining queries and, consequently, the more noise will have to be added to those queries. A possible approach to avoid

the unnecessary erosion of the privacy budget is to allow each unique function of the data to be evaluated exactly once. Each subsequent evaluation of that function would use exactly the same noise that would be added to ensure differential privacy and output the value identical to that in its first ever evaluation. We will discuss the consequences of this design in the next section.

If $\delta = 0$, the inequality (2.1) can be rewritten in terms of a more familial log-likelihood ratio:

$$\log \frac{\mathbb{P}_\nu \left[\theta(\mathbb{P}_N, \nu) \in A\right]}{\mathbb{P}_\nu \left[\theta(\mathbb{P}'_N, \nu) \in A\right]} \leq \epsilon,$$

with the constant ϵ capturing the maximum privacy loss. The differential privacy requirement with $\delta = 0$ is a rather strong one as it puts a bound on the privacy loss even for very unlikely events (when $\mathbb{P}_{\nu_N} \left[\theta(\mathbb{P}_N, \nu_N) \in A\right]$ or $\mathbb{P}_{\nu_N} \left[\theta(\mathbb{P}'_N, \nu_N) \in A\right]$ is small). With $\delta > 0$, "bad" events from the perspective of the log-likelihood ratio are allowed to happen but the probability of such is bounded by δ.

It is clear from the inequality in the definition of differential privacy that the knowledge that a certain estimator is (ϵ, δ)-differentially private is not very precise in the sense that this estimator can also be (ϵ', δ')-differentially private for some $\epsilon' < \epsilon$ or $\delta' < \delta$, and smaller values of privacy constants ε and δ obviously correspond to stricter privacy restrictions.[3] This leads us to introducing the notion of a *sharp* (ϵ, δ)*-differentially private estimator*.

DEFINITION 2 (sharp (ϵ, δ)-differential privacy)**.** *A randomized estimator $\theta(\mathbb{P}_N, \nu_N)$ is sharp (ϵ, δ)-differentially private if there is exists a measurable set A such that (2.1) holds with equality with a strictly positive probability $\mathbb{P}_{\nu_N}$.*

A sharp (ϵ, δ)-differentially private estimator by definition cannot be (ϵ', δ')-differentially private for some $\epsilon' < \epsilon$ or $\delta' < \delta$, and, thus, in this case ϵ and δ describe the exact privacy guarantees.

In practice privacy constants ϵ and δ are chosen by a data curator. As noted in [15], it is advisable that parameters in the definition of (ε, δ)-differential privacy are calibrated such that both of them are allowed to approach zero as the sample size increases. We therefore will write them as $(\varepsilon_N, \delta_N)$. Our coverage and the identification notion will allow for a variety of situations – we can$(\varepsilon_N, \delta_N)$ to be constant as well as decreasing with N and this is formally given in Section 2.2.

We now provide two examples that demonstrate different asymptotic "regimes" that arise when the estimator of interest needs to be (ϵ, δ)- differentially private. In the first regime, the resulting estimator is consistent and, moreover, additional statistical noise that ensures differential privacy is asymptotically negligible relative to the sampling noise in the estimator (additional similar examples can be found in [16], among others). In the second regime, the resulting estimator may not converge in probability and, thus, the statistical noise added due to requirement of differential privacy asymptotically dominates the sampling noise in the estimator.

[3] It is especially important for a smaller ε as it measures the range of the likelihood ratio of distributions of randomized estimators with two adjacent datasets while δ measures the (lack of) overlap between the sets of these randomized estimators).

EXAMPLE 1 (Sample mean of a random variable with a bounded support)**.** *Suppose that our goal is the estimation of the mean of a single-dimensional random variable X with support on $[0, 1]$ from the sample of i.i.d. observations $\{x_i\}_{i=1}^N$. We consider the following two mechanisms to produce (ϵ, δ)-differentially private estimator for $\mathbb{E}[X]$.*

Laplace mechanism: *Let $a_N(\nu_N) \sim Lap\left(0, \frac{1}{\epsilon N}\right)$, where $Lap(0, \lambda)$ is a Laplace random variable with density $\frac{1}{2\lambda}\exp(-|z|/\lambda)$. The output is*

$$\theta(\mathbb{P}_N, \nu_N) = \bar{x} + a_N(\nu_N)$$

with $\bar{x} = \frac{1}{N}\sum_{i=1}^N x_i$.

Exponential mechanism [4]*: Consider the distribution density $p(z; \mathbb{P}_N) \propto \exp(-\frac{1}{2}\epsilon^2 N^{1-2\gamma} \sum_{i=1}^N (z - x_i)^2)$ up to an appropriate scaling constant with some small $\gamma > 0$. In other words, density $p(z; \mathbb{P}_N)$ is the scaled and exponentiated quadratic loss whose minimizer is the sample mean (i.e. the estimator for the target $\mathbb{E}[X]$). The output is*

$$\theta(\mathbb{P}_N, \nu_N) \sim p(z; \mathbb{P}_N).$$

We note that both outputs of the Laplace mechanism and the exponential mechanism $\theta(\mathbb{P}_N, \nu_N)$ converge in probability to the target $\mathbb{E}[X]$.

For the Laplace mechanism, the variance of added Laplace noise is $1/(N^2\epsilon^2)$, meaning that $a_N(\nu_N) \xrightarrow{p} 0$ and $\theta(\mathbb{P}_N, \nu_N) \xrightarrow{p} \mathbb{E}[X]$ by the law of large numbers.

For the exponential mechanism, we note that its density can be written as

$$p(z; \mathbb{P}_N) = \frac{\epsilon N^{1-\gamma}}{\sqrt{2\pi}} \exp(-\frac{1}{2}\epsilon^2 N^{2-2\gamma} (z - \bar{x})^2),$$

which is the density of normal random variable with mean $\bar{x}$ and variance $1/(\epsilon^2 N^{2-2\gamma})$. In other words, we can decompose $\theta(\mathbb{P}_N, \nu_N)$ produced by exponential mechanism as a sum of $\bar{x}$ and $a_N(\nu_N)$ which is a mean zero normal random variable with variance $1/(\epsilon^2 N^{2-2\gamma})$. Consequently, just as the output of the Laplace mechanism, the output of the exponential mechanism converges in probability to $\mathbb{E}[X]$.

To analyze privacy consider empirical distributions $\mathbb{P}_N$ and $\mathbb{P}'_N$ where in $\mathbb{P}'_N$ a single observation x_i is replaced with an arbitrary value x'_i. Next, consider the likelihood ratio of distributions of $\theta(\mathbb{P}_N, \nu_N)$ and $\theta(\mathbb{P}'_N, \nu_N)$ for each of the two mechanisms.

[4] We note that our mechanism does not match exactly the definition of the exponential mechanism used in the differential privacy literature. It does, however, capture the spirit of that mechanism where the minimization of a criterion function is replaced with sampling from a distribution constructed from that criterion function.

For the Laplace mechanism the density of its output is

$$f_{\theta(\mathbb{P}_N,\nu_N)}(z) = \frac{N\epsilon}{2}\exp\left(-N\,\epsilon\,|z-\bar{x}|\right).$$

Then

$$f_{\theta(\mathbb{P}'_N,\nu_N)}(z) = \frac{N\epsilon}{2}\exp\left(-N\,\epsilon\,|z-\bar{x}'|\right) = \frac{N\epsilon}{2}\exp\left(-N\,\epsilon\,|z-\bar{x}-\frac{1}{N}(x'_i-x_i)|\right).$$

Then the log-likelihood ratio

$$\log\left(\frac{f_{\theta(\mathbb{P}_N,\nu_N)}(z)}{f_{\theta(\mathbb{P}'_N,\nu_N)}(z)}\right) = -N\epsilon\left(|z-\bar{x}| - |z-\bar{x}-\frac{1}{N}(x'_i-x_i)|\right) \le \epsilon|x'_i-x_i| \le \epsilon,$$

where bound is sharp with $x_i = 1$ and $x'_i = 0$. According to Definition 2, the output $\theta(\mathbb{P}_N,\nu_N)$ of the Laplace mechanism is sharply $(\epsilon,0)$-differentially private.

For the exponential mechanism the density of its output is

$$f_{\theta(\mathbb{P}_N,\nu_N)}(z) = p(z;\mathbb{P}_N).$$

We note that the likelihood ratio in this case is unbounded due to the fast decay of normal distribution tails. Instead, it is possible to bound it with a positive probability. To do that we first explicitly express the log-likelihood ratio

$$\log\left(\frac{f_{\theta(\mathbb{P}_N,\nu_N)}(z)}{f_{\theta(\mathbb{P}'_N,\nu_N)}(z)}\right) = \frac{1}{2}N^{-2\gamma}\epsilon^2(x'_i-x_i)^2 - \epsilon^2\,N^{1-2\gamma}\,(x'_i-x_i)(z-\bar{x}).$$

The log-likelihood ratio evaluated at the output of the exponential mechanism $\theta(\mathbb{P}_N,\nu_N)$ is a random variable. Since we established that $\theta(\mathbb{P}_N,\nu_N)$ is normal with mean $\bar{x}$ and variance $1/(\epsilon^2\,N^{2-2\gamma})$, then $\log\left(\frac{f_{\theta(\mathbb{P}_N,\nu_N)}(\theta(\mathbb{P}_N,\nu_N))}{f_{\theta(\mathbb{P}'_N,\nu_N)}(\theta(\mathbb{P}_N,\nu_N))}\right)$ is normal with mean $\frac{\epsilon^2}{2N^{2\gamma}}\,(x'_i-x_i)^2$ and variance $\frac{\epsilon^2}{N^{2\gamma}}(x'_i-x_i)^2$. As a result,

$$\mathbb{P}\left(\log\left(\frac{f_{\theta(\mathbb{P}_N,\nu_N)}(\theta(\mathbb{P}_N,\nu_N))}{f_{\theta(\mathbb{P}'_N,\nu_N)}(\theta(\mathbb{P}_N,\nu_N))}\right) \ge \epsilon\right) = \Phi\left(\frac{\epsilon}{2N^{\gamma}}|x'_i-x_i| - \frac{N^{\gamma}}{|x'_i-x_i|}\right) \le \Phi\left(-N^{\gamma}+.5\epsilon\right),$$

given that function $a\,z - b/z$ for $a,\,b>0$ is strictly increasing and the sharp bound is attained at $x_i = 1$ and $x'_i = 0$. As a result, the output of the exponential mechanism is $(\epsilon, \Phi\left(-N^{\gamma}+.5\epsilon\right))$-differentially private.

As we can see in Example 1, it is possible to produce differentially private estimators for sample means of the random variable with a bounded range with inference being pretty much the same as for the original estimator. Note that the standard deviation of the noise associated with the Laplace mechanism is $1/(\epsilon\,N)$ and the standard deviation of the noise associated with the exponential mechanism is $1/(\epsilon\,N^{1-\gamma})$. Thus, if random variable X has finite variance, then the standard deviation of the sample mean $\bar{x}$ is $O(1/\sqrt{N})$. As a result, the noise associated with either mechanisms used to

guarantee differential privacy ($(\epsilon, 0)$ and $(\epsilon, \Phi(-N^{\gamma} + .5\epsilon))$ respectively) is asymptotically negligible (provided that $\gamma < .5$) and the statistical analysis can proceed simply ignoring additional noise. The situation may change if we allow the differential privacy parameters ϵ and δ to depend on N. In particular, if ϵ_N is taken to converge to zero too quickly ($\epsilon_N = O(1/N)$), then the described differentially private estimator is no longer consistent as the variance of the noise dominates the variance of the original estimator.

In Example 2 below we illustrate a contrasting scenario where differential privacy guarantee substantially impacts the estimator of the target parameter making it inconsistent. This example involves a weighted mean of a random variable of interest. The reason why we want to highlight this specific example is because, as we discuss further in this paper, many relevant estimators in Economics can be viewed as being constructed from weighted means. In fact, in the RDD models estimation in Section 4 it will be clear that the lack of identifiability under differential privacy and poor statistical performance of differentially private estimators (even with relatively weak privacy requirements) will stem, in particular, from the weights of observations not being fixed but decided by the data. These weights may vary in a certain fixed interval even when only one observation in the dataset becomes different. Example 2 below considers an asymptotic regime where with the increase of the sample size, the weights are more likely to take extreme values. In this case an increasing size of the data is accompanied by an increase in its diversity with more individuals in the data having unique (low probability) combinations of features. Example 2 is not directly related to the RDD setup and does not provide any immediate result for the RDD. However, the choice of thresholding in this example and the impact of the "effective number of observations" in the RDD case have a similar flavor.

EXAMPLE 2 (Weighted sample mean of a random variable with bounded support)**.** *Consider a different target to that in Example 1. In particular, suppose that X and W are single-dimensional random variables independent from each other, each with support on $[0, 1]$. However, the object of interest is the weighted mean $\mathbb{E}[X/W] \in [0,\ 1]$ from the sample of i.i.d. observations $\{x_i,\ w_i\}_{i=1}^{N}$. Our goal in this example is to explore application of two mechanisms considered in Example 1: Laplace mechanism and exponential mechanism to guarantee (ϵ, δ)-differential privacy.*

In Example 1 we used scaling by ϵ in the noise addition for each of the two mechanisms to achieve differential privacy guarantee. The direct analysis of the likelihood ratios for both Laplace and exponential mechanisms shows that they are unbounded regardless of how these mechanisms are scaled. Therefore, for any scaling of those two mechanisms cannot provide differential privacy guarantee without a further modification of the estimator.[5]

Problems with both considered mechanisms arise due to range of the weight w_i being between 0 and 1 which means that the range of individual terms in the weighted sum x_i/w_i is unbounded.

One possible approach to attempt to address this issue is to exclude observations i where w_i is close to zero by using adaptive truncation by choosing a sequence of thresholds δ_N such that only observations

[5]This is formally evaluated in the appendix.

with $w_i \geq \delta_N$ are included for consideration. The threshold δ_N is a free parameter that controls the severity of truncation.

Suppose that $F_N^x(\cdot)$ and $F_N^w(\cdot)$ are finite-sample distributions of x_i and w_i respectively and suppose that limits of these distributions exist as $N \to \infty$. Let $\pi_N = 1 - F_N^w(\delta_N)$. $N\,\pi_N$ can be interpreted as the "effective sample size" evaluating the expected number of observations where w_i exceeds the threshold. We will normalize the estimators using this effective sample size.[6]

Laplace mechanism *in this case takes $a_N(\nu_N) \sim Lap\,(0, 1/(\epsilon N \pi_N \delta_N))$ and outputs*

$$\theta(\mathbb{P}_N, \nu_N) = \frac{1}{N\pi_N} \sum_{i=1}^{N} \frac{x_i}{w_i} \mathbf{1}\{w_i \geq \delta_N\} + a_N(\nu_N)$$

projected on $[0,1]$ (given that $E[X/W] \in [0,1]$). The first term corresponds to the mean in the truncated sample normalized by the effective sample size $N\pi_N$. Following the same steps as in Example 1, we can verify that the output of this Laplace mechanism with truncation is $(\epsilon, 0)$- differentially private.

Exponential mechanism *takes density*

$$p(z; \mathbb{P}_N) \propto \exp\left(-\frac{1}{2}\epsilon^2 \delta_N^2 (N\pi_N)^{2-2\gamma} \left(z - \frac{1}{N\pi_N} \sum_{i=1}^{N} x_i/w_i \mathbf{1}\{w_i \geq \delta_N\} \right)^2 \right)$$

and samples $\theta(\mathbb{P}_N, \nu_N) \sim p(\cdot; \mathbb{P}_N)$ projected on $[0,1]$. As for the Laplace mechanism, using the analysis from Example 1 we can verify that the output of exponential mechanism is $(\epsilon, \Phi\,(-(N\pi_N)^\gamma + .5\epsilon))$-differentially private.

In other words, from the perspective of differential privacy guarantees, both mechanisms for weighted means with truncation behave identically to those in case of an unweighted mean of a random variable with a bounded range. Statistical convergence of the outputs $\theta(\mathbb{P}_N, \nu_N)$, however, will vary in three different "regimes."

Regime 1: $N\,\pi_N \delta_N \to 0$**.** *The variance of the noise in the Laplace mechanism $(\epsilon N \pi_N \delta_N)^{-2} \to \infty$ and the variance of the noise in the exponential mechanism $(\epsilon\,\delta_N)^{-2}(N\pi_N)^{-2+2\gamma} \to \infty$. Due to projection on $[0,1]$, the outputs $\theta(\mathbb{P}_N, \nu_N)$ of both Laplace and exponential mechanism converge in distribution to a Bernoulli random variable with parameter $1/2$.*

Regime 2: $N\,\pi_N \delta_N \to \infty$ **and** $\delta_N \to \underline{\delta} > 0$**.** *In this case the variance of the noise in both the Laplace mechanism and the exponential mechanism converges to 0 as $N \to \infty$. At the same time the truncated mean $\frac{1}{N\,\pi_N} \sum_{i=1}^{N} \frac{x_i}{w_i} \mathbf{1}\{w_i \geq \delta_N\}$ converges in probability to $\underline{b} = \mathbb{E}[X/W\,\mathbf{1}\{W \geq \underline{\delta}\}]$. Due to projection on $[0,1]$, the outputs $\theta(\mathbb{P}_N, \nu_N)$ of both Laplace and exponential mechanism converge in probability to $\min\{1, \max\{0, \underline{b}\}\}$. Depending on the value $\underline{\delta}$, the limit can be any point in $[0,1]$.*

[6] Our results can be immediately to incorporate the actual sample size after truncation at the cost of additional analysis that does not, however, change the results qualitatively.

Regime 3: $N\pi_N\delta_N \to \infty$ **and** $\delta_N \to 0$**.** *As in Regime 2 the variance of the noise in both the Laplace mechanism and the exponential mechanism converges to* 0 *as* $N \to \infty$*. Given that* $\delta_N \to 0$ *the truncated mean* $\frac{1}{N\pi_N}\sum_{i=1}^{N}\frac{x_i}{w_i}\mathbf{1}\{w_i \geq \delta_N\}$ *converges in probability to* $\mathbb{E}[X/W]$*, the parameter of interest. Thus the outputs* $\theta(\mathbb{P}_N, \nu_N)$ *of both Laplace and exponential mechanism converge to the target parameter.*

To summarize Example 2, depending on the asymptotic behavior of the weights in the weighted sample mean, differentially private output can be (i) consistent and converge to the target parameter $\mathbb{E}[X/W]$; (ii) inconsistent and converge to any point in the parameter range $[0, 1]$; (iii) not converge in probability and instead converge to a random limit.

One reason we have focused on the consistency property so far is because it is a minimum desirable property of a good estimator. The second reason is directly related to what we do in the remainder of this section – our notion of identification is related to weak limits of differentially private estimators and the property of consistency describe special cases of those weak limits.

In our previous work [32] we considered identification in models where a researcher has access to a dataset that was obtained by combining split datasets subject to constraints on the prevention of the identity disclosure. There, both the data combination procedure and the impositions of privacy constraints such as k-anonymity) were intrinsically finite sample procedures. To reconcile the nature of these procedures with the population nature of the identification, we argued that the identification notion for econometric models from combined and privacy-protected data can only be defined as a limit of a combined output of the privacy preserving procedure and finite sample distribution of the data. In this paper, we aim to develop an approach to analyze identification of econometric models when the data curator requires the output of the econometric procedure to be differentially private. For that, we can use ideas related to those in [32].

2.2 Formal framework

In this section, we present our formal approach to identification for models with differentially private outcomes.

We consider a sequence of statistical experiments indexed by the sample size N ($N \to \infty$ along this sequence), where for each N we generate an i.i.d. sample $\{z_i\}_{i=1}^{N}$ from the joint distribution of d-dimensional random vector Z leading to empirical distribution $\mathbb{P}_N$. We assume that the parameter of interest θ_0 is in the interior of p-dimensional convex compact parameter space $\Theta \subset \mathbb{R}^p$

We then consider randomized estimators $\theta(\mathbb{P}_N, \nu_N) \in \Theta$ where random element ν_N ensures that these randomized estimators are $(\varepsilon_N, \delta_N)$-differentially private for some sequences ε_N and δ_N. To ensure protection from adversarial attacks on data $\mathbb{P}_N$ added random element ν_N has to be statistically independent from ν_N.

In our analysis in this section we use the techniques from the theory of random sets which reflect the

spirit of data analysis with differential privacy: the random element ν_N or the technique that is used to produce the randomized estimator $\theta(\mathbb{P}_N, \nu_N)$ aiming to represent parameter of interest θ_0 may not be available to the researcher. Instead, data curator controlling the dataset inducing the empirical distribution $\mathbb{P}_N$ reports parameters $(\varepsilon_N, \delta_N)$ which yield the upper bound guarantee for differential privacy of a given estimated output (and, possibly, the algorithm used for implementation of $\theta(\cdot, \cdot)$). This means that there can be an entire class of estimators for parameter θ_0 that satisfy differential privacy with these parameters.

While we provided examples of mechanisms that can be used to achieve differential privacy in the previous section, we need to formally define the structure of a general differentially private estimator. A differentially private estimator takes as an input data sample that produced empirical distribution $\mathbb{P}_N$ and a random element ν_N and outputs a point in Θ. We now give a formal description of the class of our considered estimators.

ASSUMPTION 1. *The class of differentially estimators $\mathcal{T}$ is associated with operators $\theta(\cdot, \cdot)$ such that:*

(i) For each $\nu \in \mathcal{V}$, $\theta(\cdot, \nu) : D(\mathbb{R}^d; [0,1]) \mapsto \Theta$ is a compact and Lipschitz-continuous operator in the $\mathbf{L}_\infty$ norm (where $D(\mathbb{R}^d; [0,1])$ is the Skorohod space of functions);

(ii) For each $F \in D(\mathbb{R}^d; [0,1])$ $\theta(F, \cdot)$ is a measurable Lipschitz-continuous function supported on $\mathcal{V}$;

(iii) For any empirical distribution $\mathbb{P}_N$, random variable $\theta(\mathbb{P}_N, \nu_N)$ induced by random variable ν_N satisfies (2.1).

We can illustrate this assumption for differentially private estimate of a sample mean using Laplace mechanism for the sample i.i.d. draws $\{X_i\}_{i=1}^N$ of random varible with bounded support and random element $\nu_N \sim U[0,1]$. In this case $\theta(F, \nu) = \int_{-\infty}^{+\infty} z \, dF(z) + \frac{\text{diam}(\Theta)}{N\varepsilon_N} F_\Lambda^{-1}(\nu)$ (where $F_\Lambda^{-1}(\cdot)$ is the inverse cdf of the standard Laplace distribution). We note that the empirical moment induced by this operator produces $(\varepsilon_N, 0)$-differentially private estimator as discussed in Example 1.

Once we imposed the structure on the class of estimators that we consider (that inherit the structure of the underlying functional space), we want to focus our attention on the classes of differentially private estimators that have certain desirable properties. These properties are formulated in Definition 3 below.

Here and thereafter we will consider differential privacy parameters $(\varepsilon_N, \delta_N)$ such that $\varepsilon_N \leq \bar{\varepsilon}$ and $\delta_N \leq \bar{\delta}$ for all N for some universal constants $\bar{\varepsilon}$ and $\bar{\delta}$.

DEFINITION 3. *For a given sequence of $(\varepsilon_N, \delta_N)$, we say that an $(\varepsilon_N, \delta_N)$-differentially private estimator $\theta(\cdot, \cdot) : \mathcal{Z}^N \times \mathcal{V} \to \Theta$ satisfying Assumption 1 is regular for the parameter of interest θ_0 if the following conditions hold:*

(i) $\theta(\mathbb{P}_N, \nu_N)$ *is a continuous random variable with respect to the Lebesgue measure;*

(ii) in the absence of the mechanism noise – that is, when the estimator is $\theta(\mathbb{P}_N, 0)$*, – there exists a function* $\bar{R}(N\kappa)$*, such that* $\lim_{N\to\infty} \bar{R}(N,\kappa) = 0$ *that for all* N *and for all* $\kappa > 0$*:*

$$\mathbb{P}\left(\|\theta(\mathbb{P}_N, 0) - \theta_0\| > \kappa\right) \leq \bar{R}(N,\kappa); \tag{2.2}$$

(iii) $\theta(\mathbb{P}_N, \nu_N)$ *has a weak limit if the sequence* $(\varepsilon_N, \delta_N)$ *is convergent.*

Condition (i) states that the distribution of $\theta(\cdot,\cdot)$ has a density. Condition (ii) implies that in the absence of any mechanism noise are informative for the parameter of interest θ_0 – in particular, the estimator is consistent and has a guaranteed rate of convergence (in most practical scenarios $\bar{R}(N,\kappa)$ would be required to be exponentially decreasing in N and κ). The condition that function $\theta(\cdot,\cdot)$ takes values in Θ only, ensures that in cases when the mechanism noise would drive the estimator outside of the parameter space Θ, such an estimator would be projected on the boundary of Θ. Finally, condition (iii) requires the differentially private estimator to converge in distribution.

Now, having defined a class of regular differentially private estimators, we make the next step towards the notion of identification.

Our next notion will depend on the set of sequences of $(\varepsilon_N, \delta_N)$ a data curator is willing to consider. E.g. it could be the set of sequences where ε_N and δ_N do not change with N, or it could a set of sequences converging to zero at a certain rate. We will refer to a fixed set of sequences as $\mathcal{E}$. We will suppose that this set of sequences is a join-semilattice in the coordinate-wise partial order for $(\varepsilon_N, \delta_N)$ – that is, the join of any two sequences from $\mathcal{E}$ is also in $\mathcal{E}$.

Following [44], we use the concept of measurable selection to define the set of all regular differentially private estimators.

DEFINITION 4. *Consider set* $\mathbb{T}^*_{N,\mathcal{E}}$ *of all random variables* $\theta(\mathbb{P}_N, \nu_N) : \mathcal{Z}^N \times \mathcal{V} \mapsto \Theta$ *satisfying Definition 3 and corresponding to sequences* $(\varepsilon_N, \delta_N)$ *from* $\mathcal{E}$*. We define random set* $\mathbb{T}_{N,\mathcal{E}}$ *referred to as the set of regular differentially private estimators for* θ_0 *for a given* $\mathcal{E}$ *as the completion of the set* $\mathbb{T}^*_{N,\mathcal{E}} \cap \mathbf{L}_1(\mathbb{P})$ *with respect to* $\mathbf{L}_1(\mathbb{P})$*-norm.*

$\mathbf{L}_1(\mathbb{P})$ is a space of measurable functions that map the elements of the σ-algebra on $\mathcal{Z}^N$ (with the product measure defined by the probability measure on $\mathcal{Z}$) and the σ-algebra associated with random elements ν_N mapping into $\mathbb{R}^p$, and for which the Euclidean norm is integrable. By definition of $\theta(\cdot,\cdot)$ in 3 and compactness of Θ, all elements in $\mathbb{T}^*_{N,\mathcal{E}}$ and, thus, in $\mathbb{T}_{N,\mathcal{E}}$ as well are bounded in the $\mathbf{L}_1$-norm.

Random set $\mathbb{T}_{N,\mathcal{E}}$ is compact and convex in the sense of Definitions 1.30 and 4.32 in [44], as shown in Lemma 1.

LEMMA 1. $\mathbb{T}_{N,\mathcal{E}}$ *is convex and compact random set.*

Our next step in the formulation of the identification notion is to suggest which notion of limit to apply to random sets. At the first glance, the use of the probability limit might seem like a natural approach if we want to somehow relate the identifiability of the parameter to the fact that random sets $\mathbb{T}_{N,\mathcal{E}}$ overlap with an arbitrary neighborhood of θ_0 with probability approaching 1. However, the notion of the probability limit is too strong as it will not allow us to talk about the limit in the following simple instance of regular differentially private estimators:

$$\theta(\mathbb{P}_N, \nu_N) = \theta(\mathbb{P}_N, 0) + a_N(\nu_N)$$

where the variance of $a_N(\nu_N)$ remains constant or increases as $N \to \infty$ (the estimator $\theta(\mathbb{P}_N, 0)$ in the absence of the mechanism noise is, of course, consistent by the condition (ii) in Definition 3). As we will see later, such situations will be prevalent in the estimation of RDD models. In fact, we already conveyed our intention to consider weak limits in condition (iii) in Definition 3. The next lemma demonstrates that weak convergence is the strongest plausible convergence concept to consider unless weak limits of regular differentially private estimators are constant.

LEMMA 2. *Suppose that $\theta(\mathbb{P}_N, \nu_N) \xrightarrow{W} \tau$ as $N \to \infty$. Then if τ is not constant with a positive probability, then there exists $\bar{\kappa} > 0$ and $\gamma > 0$ such that for all $\kappa \leq \bar{\kappa}$*

$$\limsup_{N\to\infty} \mathbb{P}\left(|\theta(\mathbb{P}_N, \nu_N) - \tau| > \kappa\right) > \gamma.$$

Theorem 1 establishes weak convergence of the convex compact random set $\mathbb{T}_{N,\mathcal{E}}$ when all the sequences of $(\varepsilon_N, \delta_N)$ in $\mathcal{E}$ are convergent.

THEOREM 1. *For a $\mathcal{E}$ be the join-semilattice that consists of only convergent sequences of $(\varepsilon_N, \delta_N)$. The random set $\mathbb{T}_{N,\mathcal{E}}$ as defined in Definition 4 weakly converges to a random set $\mathbf{T}_{\mathcal{E}}$, which is the closure of all weak limits of estimators in the respective random set $\mathbb{T}^*_{N,\mathcal{E}}$.*

The convergence result of Theorem 1 is essentially a result of what happens in the limit of statistical experiments. Naturally, we will base our notion of identifiability on some characteristics of the random set $\mathbf{T}_{\mathcal{E}}$.

2.2.1 Characterization of the limit random set $\mathbf{T}_{\mathcal{E}}$

Our next step will be to produce a tangible characterization of the random set $\mathbf{T}_{\mathcal{E}}$ and define the information content delivered by this set $\mathbf{T}_{\mathcal{E}}$ with regard to the parameter of interest. The best case scenario from the information content point of view is the case when $\mathbf{T}_{\mathcal{E}}$ is the degenerate distribution concentrated at θ_0, which essentially means that the sequence of the statistical experiments delivers the true parameter values in the limit. Generally, however, this may not be the case. One of the main difficulties in characterizing this set and its information content is that it may not contain the target parameter θ_0 (i.e. it can be "biased") and that the distribution of $\mathbf{T}_{\mathcal{E}}$ may not be degenerate (i.e. it is not "consistently" estimating θ_0). Some important work in the random sets literature, such

as [8] and [7], defined the information content (or, in other words, identified set) as the selection expectation of a random set. In our previous work [32] we used the concept related to selection expectation to analyze the impact of privacy guarantees (in particular, k-anonymity) under data combination.[7] Thus, selection expectation might seem like a promising approach to explore in our framework as well, especially given that such characterization is deterministic, but for the reasons explained earlier and again in more detail below (privacy budget for differentially private mechanisms and the impossibility of repeated experiments), in the differentially private setting we do not see this approach as a fruitful one.

Having given the gist of the content of this section, we now turn to a more detailed discussion.

First, as mentioned above, in the context of converging random sets, as for converging random variables, the notion of the expectation or the median (or some other quantile) might appear a natural way to characterize the limit (and consequently provide the framework for identification). Recall that the *selection (or Aumann) expectation* of random set $\mathbb{T}_{N,\mathcal{E}}$ denoted $\mathbb{E}\,\mathbb{T}_{N,\mathcal{E}}$ is the closure of the set $\{\mathbb{E}\,\xi : \xi \in \mathbb{T}_{N,\mathcal{E}}\}$. However, the selection expectation (along with any first-order statistic, such as the Vorob'ev's expectation) fails to be "representative" for the limiting random set $\mathbf{T}_{\mathcal{E}}$. This lack of representativeness stems from the inherent impossibility to replicate a statistical experiment whose outcome is the regular differentially private estimator driven by the required structure of differentially private systems as discussed in [16]. Indeed, in the context of differentially private systems no function evaluated on the data can be considered in isolation. Differential privacy is the property of the *entirety of all functions that have ever been or will ever be evaluated from a given dataset.* By the composition property of differential privacy, two different functions that, for instance, are $(\varepsilon/2, 0)$-differentially private each constitute one combined (at most) $(\varepsilon, 0)$-differentially private data query. As we mentioned it above, the way to address the multitude of statistical function that can be evaluated from the data is to define the "privacy budget" that will determine differential privacy parameter ϵ determining the privacy of all those functions in totality. The more functions need to be evaluated from the data, the more noise will need to be added to each function to ensure the entire set of functions is within the "privacy budget." The evaluation of K functions to ensure $(\varepsilon, 0)$-differential privacy requires the output of each function to be $(\varepsilon/K, 0)$-differentially private, i.e. each output should be K times less sensitive to arbitrary changes in individual observations in the dataset.

One natural way to extend the "privacy budget" of a given dataset is to ensure that the output corresponding to the evaluation of the exactly same function of the data always remains the same. In this case, the random element ν_N to produce the regular differentially private estimator $\theta(\mathbb{P}_N, \nu_N)$ is generated only once and then any data user who wants to estimate a given parameter θ_0 will observe exactly the same value of the randomized estimator $\theta(\mathbb{P}_N, \nu_N)$. In other words, repeated identical queries to the data always result in the same (randomized only once) output.

From this perspective, the concept of the selection expectation or other related statistics is clearly

[7] As discussed in the introduction, k-anonymity is a formal privacy guarantee that predates differential privacy).

misleading for a characterization of the limiting random set $\mathbf{T}_{\mathcal{E}}$. While it may be the case that $\mathbb{E}_{\nu_N}\,\theta(\mathbb{P}_N,\nu_N)$ is close to θ_0 with high probability (as $N\to\infty$) for all measurable selections $\theta(\mathbb{P}_N,\nu_N)$ of $\mathbb{T}_{N,\mathcal{E}}$, there is no guarantee that $\theta(\mathbb{P}_N,\nu_N)$ is also close to θ_0 with high probability and under "privacy budget" considerations there is no way a researcher can access repeated samples from the distribution $\theta(\mathbb{P}_N,\nu_N)$ corresponding to a given empirical distribution $\mathbb{P}_N$ to "average out" the added noise ν_N. To put in other words, given the "privacy budget" considerations, the selection expectation $\mathbb{E}_{\nu_N}\,\theta(\mathbb{P}_N,\nu_N)$ is not "feasible" in the differentially private framework. Moreover, as we illustrate in continuation of Example 2 below, differential privacy may be in conflict even with the identification of expectations. It is clear that we need to use a different approach to characterizing the limiting random set $\mathbf{T}_{\mathcal{E}}$.

To provide a comprehensive characterization of random sets $\mathbb{T}_{N,\mathcal{E}}$ and $\mathbf{T}_{\mathcal{E}}$ we use the notion of the containment functional adopted from [44]:

DEFINITION 5. *Functional $C_{\mathbf{X}}(K)=\mathbb{P}\left(\mathbf{X}\subset K\right)$ for convex compact subset of Θ is referred to as the containment functional of random set $\mathbf{X}$.*

By Theorem 1.7.8 in [44] containment functional provides a complete characterization of convex compact random set. Moreover, it is sufficient to choose the "test sets" K to be convex polytopes. Coverage functional preserves the property of the weak convergence of the sequence of random sets.

We summarize this in the following theorem

THEOREM 2. *Under conditions of Theorem 1 for any convex polytope $K\subset\Theta$*

$$C_{\mathbb{T}_{N,\mathcal{E}}}(K)\to C_{\mathbf{T}_{\mathcal{E}}}(K),\ \text{as } N\to\infty.$$

This theorem is a simple corollary of Theorem 1.6.5 in [44] and it ensures that the coverage functional preserves the properties of the converging sequence of random sets $\mathbb{T}_{N,\mathcal{E}}$ and, more importantly, its limit $\mathbb{T}_{N,\mathcal{E}}$. The characterization of the limiting coverage functional would equivalently characterize the limiting random set. In other words, the analysis of weak convergence of random sets can be replaced with the analysis of pointwise convergence of the coverage function on the set of convex polytopes contained in Θ.

EXAMPLE 2 (continued). *We now return to our Example 2 that focused on the differentially private estimation of the weighted mean $\theta_0=\mathbb{E}[X/W]$. We considered this estimation in the asymptotic regime where finite sample distributions of (independent) random variables X and W, $F_N^x(\cdot)$ and $F_N^w(\cdot)$ converge to their limits as $N\to\infty$. In order to ensure that differential privacy guarantees can be achieved with a finite noise addition, we trimmed the sample with the adaptive trimming threshold δ_N such that all observations with $w_i<\delta_N$ are discarded.*

We note that for this example the operator $M_{\theta,\nu}(\cdot)$ discussed in Assumption 1 can be written explicitly as

$$\theta(\mathbb{P}_N,\nu_N)=\frac{1}{N\pi_N}\sum_{i=1}^{N}\frac{x_i}{w_i}\mathbf{1}\{w_i\geq\delta_N\}+F_{\Lambda}^{-1}(\nu_N)$$

for the Laplace mechanism (where $F_\Lambda(\cdot)$ is the cdf of the Laplace distribution with parameter $1/(\epsilon\, N\, \pi_N\, \delta_N)$ and ν_N is the uniform random variable on $[0,1]$). For exponential mechanism we can write it as a solution of

$$\int_{-\infty}^{\theta(\mathbb{P}_N,\nu_N)} p(t;\mathbb{P}_N)\,dt - \nu_N = 0.$$

where ν_N is the uniform random variable on $[0,1]$.

For the sake of the simplicity of the exposition, we ignore Regime 2 of Example 2 where $\delta_N \to \underline{\delta} > 0$ and instead focus on Regimes 1 and 3 where $\delta_N \to 0$. Let $\theta^{(1)}(\mathbb{P}_N,\nu_N)$ be a single instance of the output of either the Laplace or the exponential mechanism in Example 2 in Regime 1. Also, let $\theta^{(3)}(\mathbb{P}_N,\nu_N)$ be a single instance of the output of either mechanism in Regime 3.

Define $\mathbb{T}^*_{N,\mathcal{E}} = co\{\theta^{(1)}(\mathbb{P}_N,\nu_N),\ \theta^{(3)}(\mathbb{P}_N,\nu_N)\}$. Provided that random set $\mathbb{T}^*_{N,\mathcal{E}}$ is a convex combination of two elements, then $\mathbb{T}_{N,\mathcal{E}} \equiv \mathbb{T}^*_{N,\mathcal{E}}$. Since $\theta^{(1)}(\mathbb{P}_N,\nu_N)$ and $\theta^{(3)}(\mathbb{P}_N,\nu_N)$ are $(\epsilon,0)$-differentially private if Laplace mechanism was used (or $(\epsilon,\Phi(-(N\pi_N)^\gamma+.5\epsilon))$-differentially private if exponential mechanism was used) then each element of $\mathbb{T}_{N,\mathcal{E}}$ is also $(\epsilon,0)$- differentially private ($(\epsilon,\Phi(-(N\pi_N)^\gamma+.5\epsilon))$-respectively for the exponential mechanism). Moreover, as $N\to\infty$, regardless of the mechanism the output will be $(\epsilon,0)$-differentially private (given that $\lim_{N\to\infty}\Phi(-(N\pi_N)^\gamma+.5\epsilon)=0$).

From Example 2 it follows that $\theta^{(1)}(\mathbb{P}_N,\nu_N)$ converges in distribution to Bernoulli random variable with parameter $1/2$ and $\theta^{(3)}(\mathbb{P}_N,\nu_N)$ converges in probability to $\theta_0 = \mathbb{E}[X/W] \in [0,1]$. This also means that

$$\Big(\theta^{(1)}(\mathbb{P}_N,\nu_N),\theta^{(3)}(\mathbb{P}_N,\nu_N)\Big) \xrightarrow{W} (B(.5),\,\theta_0).$$

As a result, the limiting random set $\mathbf{T}_\mathcal{E}$ will be a convex hull of this vector. More explicitly, we can write

$$\mathbf{T}_\mathcal{E} = \begin{cases} [0,\,\theta_0]\,, & \text{with probability } 1/2,\\ [\theta_0,\,1]\,, & \text{with probability } 1/2. \end{cases}$$

The containment functional also takes a simple form: $C_{\mathbf{T}_\mathcal{E}}(K) = \frac{1}{2}\,\mathbf{1}\{[0,\,\theta_0]\subset K\} + \frac{1}{2}\,\mathbf{1}\{[\theta_0,\,1]\subset K\}$.

2.3 Smoothness and separability of differentially private estimators

While our previous discussion considers a general, possibly non-separable form for regular differentially private estimator all existing approaches to inducing differential privacy lead to much simpler (approximately) separable estimators. We now further narrow down the class of regular differentially private estimators to reflect this property.

DEFINITION 6. *We say that regular differentially private estimator $\theta(\mathbb{P}_N,\nu_N)$ is smooth, if there exists a functional $\psi(\cdot)$ and function a_N with support on $\mathcal{V}$ such that*

$$\theta(\mathbb{P}_N,\nu_N) = \psi(\mathbb{P}_N) + a_N(\nu_N) + \Delta_N,$$

with $\mathbb{E}\left[\left(\sqrt{\log \frac{1}{\bar{R}(N,\kappa)}}\Delta_N\right)^2\right] \to 0$ *as* $N \to \infty$ *for all* $\kappa > 0$, *where* $\bar{R}(N,\kappa)$ *is provided in Definition 3.*

Definition 6 focuses on the regular differentially private estimators that are approximately separable in the way they depend on the data sample and on the added noise to achieve differential privacy. The residual Δ_N is required to be "small" relative to the rate of convergence of the version of the estimator $\theta(\mathbb{P}, \nu_N)$ that is infused with the "trivial" noise that does not perturb it.

We note that for two early mechanisms that were proposed to provide differential privacy: Laplace and Gaussian mechanisms, Definition 6 applies trivially. In Laplace mechanisms differential privacy is achieved by the addition of the double-exponential noise ν_N to the original estimator and in the Gaussian mechanism ν_N is the additive normal noise. In other words, by construction for both of these mechanisms $\Delta_N \equiv 0$.

We now consider additional classes of mechanisms for differential privacy that are not explicitly additive.

One interesting example of a popular non-separable mechanism for differential privacy is the exponential mechanism developed in [42]. In application to extremum estimators, the mechanism replaces extremum estimator $\widehat{\theta}$ that maximizes sample objective function $Q(\theta\,;\mathbb{P}_N)$ over $\theta \in \Theta$ with a draw from *quasi-posterior* distribution implied by $Q(\cdot\,;\mathbb{P}_N)$. This class of estimators is directly related to randomized estimators developed in [12]. [12] consider the case where the population analog of the objective function $Q(\theta\,;\mathbb{P}_N)$ satisfies the information matrix equality. To form the estimator they propose to consider a prior distribution $\pi(\cdot)$ over θ and a *quasi-likelihood function* $\exp\left(Q(\theta\,;\mathbb{P}_N)\right)$ (so that the original objective function is quasi-log-likelihood function). Then the estimator is a mean of quasi-posterior distribution $\propto \exp\left(Q(\theta\,;\mathbb{P}_N)\right)\pi(\theta)$ which consistently estimates the maximizer of the population objective function regardless of the shape of the prior distribution $\pi(\cdot)$ (under mild regularity conditions). Moreover, the variance of this quasi-posterior distribution accurately estimates the asymptotic variance of the original extremum estimator. A significant advantage of this estimator over the original extremum estimator $\widehat{\theta}$ is that it does not require maximization of potentially non-smooth or hard to optimize function $Q(\theta\,;\mathbb{P}_N)$.

In the follow-up work in [33] for the cases where $Q(\theta\,;\mathbb{P}_N)$ may be steep in the vicinity of the maximum, which may lead to slow convergence of the simulations required to sample from the quasi-posterior, it is proposed to scale the exponent in the pseudo-likelihood function as $\exp\left(\lambda\, Q(\theta\,;\mathbb{P}_N)\right)$ using a constant λ which is selected based on the speed of mixing of the simulated Markov Chain (produced using the new pseudo-posterior). The corresponding posterior mean remains a consistent estimator for the maximizer of the population objective function while its asymptotic variance can be estimated by scaling the variance of the quasi-posterior using λ.

The exponential mechanism for differential privacy considered in [42] is a simple implementation of the idea in [12]: the estimator is a single draw from the quasi-posterior $\propto \exp(\lambda\, Q(\theta\,;\mathbb{P}_N))\,\pi(\theta)$.

The resulting estimator turns out to be $(\lambda, \Delta Q\,, 0)$-differentially private, where

$$\Delta Q = \sup_{\theta \in \Theta,\, \mathbb{P}_N, \mathbb{P}'_N} |Q(\theta\,;\, \mathbb{P}'_N) - Q(\theta\,;\, \mathbb{P}_N)|$$

is the *global sensitivity* of the objective function $Q(\theta\,;\, \mathbb{P}_N)$ evaluated over all empirical distributions $\mathbb{P}'_N$ that are different from $\mathbb{P}_N$ in any one single support point if ΔQ is bounded. If ΔQ is unbounded, the estimator is (ϵ, δ) -differentially private with vanishing δ for an appropriately calibrated scaling constant λ.

[12] and later [33] focus on the cases where $Q(\theta\,;\, \mathbb{P}_N)$ is stochastically equicontinuous and the quasi-posterior is asymptotically equivalent to

$$\propto \exp\left(-\frac{1}{2}\lambda(\theta - \widehat{\theta})' H\,(\theta - \widehat{\theta}) + o_p(\|\theta - \widehat{\theta}\|^2)\right),$$

where H is the Hessian of the population objective function. This means that a single draw from this quasi-posterior, corresponding to the exponential mechanism for differential privacy can be represented as

$$\tilde{\theta} = \widehat{\theta} + \lambda\,\xi + o_p(1),$$

where ξ is a multivariate normal random vector with mean zero and covariance matrix H^{-1}. The extremum estimator $\widehat{\theta}$ is only depends on the data distribution $\mathbb{P}_N$ and is not affected by the noise. Therefore, the exponential mechanism is smooth in the sense of Definition 6.

We now present a simple lemma that outlines the additive representation of smooth differentially private estimators which we will use in our applications.

LEMMA 3. *Consider additive randomized estimator* $\theta^*(\mathbb{P}_N, \nu_N) = \psi(\mathbb{P}_N) + a_N(\nu_N)$, *where*

- *For each* $\kappa > 0$, $\lim_{N\to\infty} \mathbb{P}\left(|\psi(\mathbb{P}_N) - \theta_0| > \kappa\right) = 0$
- $a_N(0) = 0$ *for all* N.

Then estimator $\theta^*(\mathbb{P}_N, \nu_N)$ *is regular in the sense of Definition 3. Moreover for any converging sequence* $(\varepsilon_N, \delta_N)$ *and any smooth* $(\varepsilon_N, \delta_N)$*- differentially private estimator* $\theta(\mathbb{P}_N, \nu_N)$ *(i.e. as in Definition 6) there exists an asymptotically equivalent regular estimator* $\theta^*(\mathbb{P}_N, \nu_N)$.

This lemma, which follows from our smoothness definition, allows us to focus our further analysis on such separable estimators. Given the asymptotic equivalence, the set of weak limits of these estimators coinsides with the target random set $\mathbf{T}_{\mathcal{E}}$.

3 Point and Set Identification with Formal Privacy Guarantees

3.1 Preliminarities

Given that, as established in the previous section, the set of regular differentially private estimators weakly converges to a limiting random set, we can define the pseudo-identified set (we use the prefix "pseudo-" because it does not necessarily contain the true parameter value) as the random set $\mathbf{T}_{\mathcal{E}}$ itself. The notion of a random identified or a random pseudo-identified set is not traditional in econometrics as usually researchers are able to extract deterministic consensus about which parameter values can be driving observables. We argue that in the differentially private setting this is the preferable approach. Some other work, such as [29] employed the notion of random identified set. In [29], the source of probability that induces the random identified set is the posterior uncertainty for the identifiable parameters. In our case, this source is a combination of the sampling uncertainty of the observations and the mechanism noise, and it is not possible to separate these two sources.

We now formulate the notion of identifiability of the parameter of interest.

DEFINITION 7 (Identifiability of parameter under differential privacy)**.** *Let $\mathcal{E}$ include only some converging sequences of $(\varepsilon_N, \delta_N)$. We will say that the parameter θ_0 is identified in the regular $(\varepsilon_N, \delta_N)$-differentially private framework, where the sequences of $(\varepsilon_N, \delta_N)$ belong to $\mathcal{E}$, if and only if for any $\alpha \in (0,1)$ and any convex polytope $K \ni \theta_0$*

$$C_{\mathbf{T}_{\mathcal{E}}}(K) \geq 1 - \alpha.$$

Theorem 3 below gives necessary and sufficient conditions for the identification of parameter θ_0.

THEOREM 3. *Suppose the conditions of Theorem 1 hold. For any sequence of $(\varepsilon_N, \delta_N)$ from $\mathcal{E}$ it holds that any regular $(\varepsilon_N, \delta_N)$-differentially private estimator $\theta(\mathbb{P}_N, \nu_N)$ is such that*

$$\theta(\mathbb{P}_N, \nu_N) \xrightarrow{p} \theta_0,$$

if and only if for any $\alpha \in (0,1)$ and any convex polytope $K \ni \theta_0$ we have $C_{\mathbf{T}_{\mathcal{E}}}(K) \geq 1-\alpha$ and, thus, the parameter θ_0 is identifiable even under differential privacy.

Theorem 3 provides our characterization of identifiability which corresponds to the convergence of the sequence of random sets to a singleton. In other words, this parallels consistency for sequences of ordinary random variables.

Based on the same principles we can characterize the case of non-identifiability .

DEFINITION 8 (non-identifiability of parameter under differential privacy)**.** *Let $\mathcal{E}$ consist of converging sequences of $(\varepsilon_N, \delta_N)$. We will say that the parameter θ_0 is non-identified in the regular $(\varepsilon_N, \delta_N)$-differentially private framework, where the sequences of $(\varepsilon_N, \delta_N)$ belong to $\mathcal{E}$, if and only if there exists $\beta \in (0,1)$ and a convex polytope $K_\beta \ni \theta_0$ such that*

$$C_{\mathbf{T}_{\mathcal{E}}}(K_\beta) \leq 1 - \beta.$$

Non-identifiability implies that the limiting random set is not degenerate. Therefore, it becomes impossible to pinpoint the true parameter θ_0 by tracing a "mass point" of the coverage function of that limiting random set $\mathbf{T}_{\mathcal{E}}$. This makes the analysis of partial identification in our case different from the traditional approach where partial identification aims to construct a deterministic set that contains the parameter of interest θ_0. In our case where the coverage function is non-degenerate in the limit it is impossible to construct such a deterministic set. At the same time, coverage function itself may be difficult to work with in practice. In the next section we focus on developing a more practical approach to identification as well as construction of coverage sets for the parameter of interest.

3.2 Identification

3.2.1 Identification from data curator's decisions

As we show in the applications below, the setting where the limiting random set $\mathbf{T}_{\mathcal{E}}$ is not a singleton is not an exception but instead may be a commonplace scenario. This means, that even when N is large a smooth regular differentially private estimator $\theta(\mathbb{P}_N, \nu_N)$ is not guaranteed to be close to the target parameter θ_0 with a high probability. Even more importantly, given that limit $\mathbf{T}_{\mathcal{E}}$ is a random set it itself cannot be the target of inference since events $\{\mathbf{T}_{\mathcal{E}} \subset K\}$ for convex polytopes K may occur with non-singular probability (i.e. strictly between 0 and 1). This means that general settings of differential privacy provide the framework where the in the absence of point identification one cannot use the classic tools of partial identification (see [40] and [8]) where target set containing parameter θ_0 can be contained with probability approaching 1 as $N \to \infty$. The parameter of interest is then neither point nor partially identified.

One important key component in the proposed "practical" implementation of differentially private data analysis (e.g. [1]) is the adherence to a specific protocol for noise infusion to ensure differential privacy of the target output. This means that the differential privacy guarantee becomes "verifiable" where one can re-trace all the steps of the algorithm for differential privacy starting from raw data to the final output. In the context of our result characterizing the limits of regular differentially private outputs as a random set, we can formalize this idea by introducing a decision functional that takes random set $\mathbf{T}_{N,\mathcal{E}}$ and selects its specific element $\theta(\mathbb{P}_N, \nu_N)$.

DEFINITION 9. *Mapping τ_f from the elements of Fell topology on Θ into Θ indexed by a strongly convex function f on Θ such that for a fixed realization of $\mathbb{T}_{N,\mathcal{E}}$*

$$\tau_f\left(\mathbb{T}_{N,\mathcal{E}}\right) = \text{Argmin}_{z \in \mathbb{T}_{N,\mathcal{E}}} f(z)$$

is referred to as the decision mapping of the data curator.

Definition 9 provides two key elements of "transparent behavior" of the data curator who wants to impose differential privacy constraints. First, the data curator learns *the entire random set* $\mathbb{T}_{N,\mathcal{E}}$ containing regular differentially private estimators for the target parameter. Second, she uses

a principled way of selecting a point within that set that depends on the shape of that set by minimizing a convex function $f(\cdot)$ that she communicates to the public.

We note that this definition is conceptually different from the idea of "transparent reproducible algorithm" for differential privacy, which simply corresponds to an element of the random set $\mathbb{T}_{N,\mathcal{E}}$ without the guarantee that it is some special element of that set (such as its center of gravity). As we discuss this formally below, without the knowledge of a relative position of given differentially private output in $\mathbb{T}_{N,\mathcal{E}}$ one would not be able to identify the parameter of interest θ_0.

LEMMA 4. *Decision mapping $\tau_f(\mathbb{T}_{N,\mathcal{E}})$ is a random variable.*

Proof: Given that $\mathbb{T}_{N,\mathcal{E}}$ is convex and compact then $\text{Argmin}_{z\in\mathbb{T}_{N,\mathcal{E}}} f(z) \subset \mathbb{T}_{N,\mathcal{E}}$. Moreover, it is a singleton by strong convexity of $f(\cdot)$. As a result, $\tau_f(\mathbb{T}_{N,\mathcal{E}})$ is a measurable selection of $\mathbb{T}_{N,\mathcal{E}}$ and, therefore, a random variable. *Q.E.D.*

Next we analyze how the convergence of random sets $\mathbb{T}_{N,\mathcal{E}}$ to the limit $\mathbf{T}_{\mathcal{E}}$ translates into the covergence of the realizations of the decision mapping. To do that note that $\tau_f(\mathbb{T}_{N,\mathcal{E}}) \leq t$ means that sets $\mathbb{T}_{N,\mathcal{E}}$ and convex compact set $\{f(z) \leq t\}$ overlap. In other words

$$\mathbb{P}(\tau_f(\mathbb{T}_{N,\mathcal{E}}) \leq t) = \mathbb{P}(\mathbb{T}_{N,\mathcal{E}} \cap \{f(z) \leq t\} \neq \emptyset) = 1 - \mathcal{C}_{\mathbb{T}_{N,\mathcal{E}}}(\{f(z) > t\}).$$

Combining this result with Theorem 1 we formulate the following theorem.

THEOREM 4. *The sequence of random variables $\tau_f(\mathbb{T}_{N,\mathcal{E}})$ weakly converges as $N \to \infty$ with the limit $\tau_f(\mathbb{T}_{\mathcal{E}})$.*

Since $\tau_f(\mathbb{T}_{N,\mathcal{E}})$ is a measurable selection of $\mathbb{T}_{N,\mathcal{E}}$, by our construction it is an absolutely continuous random variable. Thus, we can write its likelihood function $\ell_{f,\mathbb{T}_{N,\mathcal{E}}}(t)$. Note that generically this likelihood function cannot be used directly since one gets a single draw of the random variable $\tau_f(\mathbb{T}_{N,\mathcal{E}})$ (due to the privacy budget considerations discussed above). The only instance where it could be used directly is when function $f(\cdot)$ determining the decision mapping ensures that the limiting distribution of $\tau_f(\mathbb{T}_{\mathcal{E}})$ is degenerate at θ_0.

Finding an appropriate $f(\cdot)$ (adapted to the structure of the random set $\mathbf{T}_{\mathcal{E}}$ for a given problem) is the most important task of the data curator that would ensure that the differential privacy framework renders the dataset "useful" for estimation of the target parameter θ_0.

We note that the distribution of the random set characterized by the containment functional $\mathcal{C}_{\mathbf{T}_{\mathcal{E}}}(\cdot)$ is ultimately determined by the data generating process indexed by parameter θ_0 and, thus we can write $\mathcal{C}_{\mathbf{T}_{\mathcal{E}}}(\cdot;\theta_0)$ indexing the containment functional by θ_0. The distributional knowledge of the data curator regarding the structure of the of the limiting random set $\mathbf{T}_{\mathcal{E}}$ implies that for parameter of the data generating process θ, $\mathcal{C}_{\mathbf{T}_{\mathcal{E}}}(\cdot;\theta)$ is known. Moreover, for each realization of the random set $\mathbf{T}_{\mathcal{E}}$, parameter θ can be geometrically located within $\mathbf{T}_{\mathcal{E}}$ (e.g. relative to the extreme points of this realization).

This means that the data curator with full distributional knowledge of $\mathbf{T}_{\mathcal{E}}$ can select $f(\cdot)$ such that

$$\mathbb{P}\left(\tau_f\left(\mathbf{T}_{\mathcal{E}}\right)=\theta\,;\,\theta\right)=1.$$

The decision mapping with this property ensures that regular differentially private output selected by the data curator is "meaningful" by approximating θ_0 with high probability.

THEOREM 5. *Suppose that the decision rule of the data curator is such that* $\mathbb{P}\left(\tau_f\left(\mathbf{T}_{\mathcal{E}}\right)=\theta\,;\,\theta\right)=1$ *for almost all* $\theta\in\Theta$*, then*

$$\tau_f\left(\mathbb{T}_{N,\mathcal{E}}\right)\xrightarrow{p}\theta_0.$$

The proof of this theorem follows directly from the proof of Theorem 4 and equivalency of the weak convergence and the convergence in probability when both have degenerate limits.

We now illustrate this point using an extremely simplified example, that has an advantage of the closed-form expression for the limiting random set $\mathbf{T}_{\mathcal{E}}$ and a suitable decision function of the data curator.

EXAMPLE 2 (continued). *We continue discussion of Example 2 focusing on estimation of a weighed mean where we demonstrated that the limit random set of considered differentially private outputs are such that all of its elements are* $(\epsilon,0)$ *-differentially private. The limit random set takes the form*

$$\mathbf{T}_{\mathcal{E}}=\begin{cases}[0,\,\theta_0], & \text{with probability } \frac{1}{2},\\ [\theta_0,\,1]\,, & \text{with probability } \frac{1}{2}.\end{cases}$$

Note that this structure reflects the point that we made in this section: if the data curator fully explores the distributional structure of the limit random set then she knows the structure of the containment functional $\mathcal{C}_{\mathbf{T}_{\mathcal{E}}}(K)=\frac{1}{2}\mathbf{1}\{[0,\theta_0]\subset K\}+\frac{1}{2}\mathbf{1}\{[\theta_0,1]\subset K\}$*. This means that if she constructs the decision mapping that "picks out" the extremal point of the random set that is not on (or close to) the boundary of the parameter space, then she can estimate the target parameter* θ_0*. For instance, we can construct the decision mapping by minimizing*

$$f(z)=\left(z-\min_{z\in\mathbf{T}_{\mathcal{E}}}z\;\mathbf{1}\{|\max_{z\in\mathbf{T}_{\mathcal{E}}}z-1|<\delta\}-\max_{z\in\mathbf{T}_{\mathcal{E}}}z\;\mathbf{1}\{|\min_{z\in\mathbf{T}_{\mathcal{E}}}z|<\delta\}\right)^2 \tag{3.3}$$

over $\mathbf{T}_{\mathcal{E}}$ *for some small* δ*. This function attains its minimum at* θ_0 *for any realization of* $\mathbf{T}_{\mathcal{E}}$*.*

As a result, the decision mapping that minimizes (3.3) for finite sample realization $\mathbb{T}_{N,\mathcal{E}}$ *selects its element which is a consistent regular differentially private estimator for* θ_0*. Therefore,* θ_0 *is identified in terms of Definition 7.*

We note that "non-strategic" choice of the point within $\mathbf{T}_{\mathcal{E}}$ *leads to the loss of identification. For instance, suppose that the data curator chooses a point within* $\mathbf{T}_{\mathcal{E}}$ *uniformly at random. In this case output* $\theta(\mathbb{P}_N,\nu_N)$ *is a continuous random variable with density*

$$\varphi_{\theta(\mathbb{P}_N,\nu_N)}(t)=\begin{cases}\frac{.5}{\theta_0}, & \text{if } t<\theta_0,\\ \frac{.5}{1-\theta_0}, & \text{if } t>\theta_0.\end{cases}$$

This distribution is not degenerate at θ_0 and by Definition 8, θ_0 is not identified.

We note that uniformly at random choice of a point within $\mathbf{T}_{\mathcal{E}}$ in Example 2 above satisfies the "transparency" principle of differential privacy (e.g. considered in [20]). In fact, an external researcher can re-trace all steps undertaken by the data curator from the raw data to the final output. However, the resulting regular (ϵ_N, δ_N) - differentially private output is not a consistent estimator for the target parameter. Consistent estimation in these settings requires a complete knowledge of the distribution of the random set by the data curator and a carefully constructed decision mapping that selects the point in each realization of the limiting random set.

3.2.2 Computational approach

Example 1 above gives a parsimonious closed form expression for the limiting random set of regular differentially private estimators and the decision mapping of the data curator. We envision the data curator to engage in a case-by-case analysis of the distribution structure of the random set that would produce a suitable decision mapping. When such an analysis is difficult or impossible we can consider computational approach which produces a numerical approximation to the decision mapping under an additional condition.

Parameter location within the random set: *Denote by $\mathbf{T}_{\mathcal{E}}^{\theta}$ the limiting random set of regular differentially private estimators corresponding to the data generating process with parameter $\theta \in \Theta$. There exists fixed continuous mapping $M(\cdot)$ defined by Lipschitz-continuous function $m(\cdot)$ such that for each realization of the random set $\mathbf{T}_{\mathcal{E}}^{\theta}$,*

$$\theta = M\left(\mathbf{T}_{\mathcal{E}}^{\theta}\right) = \int_{\mathbf{T}_{\mathcal{E}}^{\theta}} z\, m(z)\, dz \tag{3.4}$$

Condition (3.4) implies that the location of the parameter of the data generating process is encoded in the shape of the realization of the random set as an integral over the fixed function. Once a particular realization is fully known, the parameter of the data generating process can be recovered.

This condition can be given a different interpretation by invoking the divergence theorem. In particular, if $\mu(\cdot)$ is such that

$$m(z) = \sum_{i=1}^{p} \frac{\partial(z_i \mu(z))}{\partial z_i},$$

then by the divergence theorem

$$\int_{\mathbf{T}_{\mathcal{E}}^{\theta}} m(z)\, dz = \oint_{\partial \mathbf{T}_{\mathcal{E}}^{\theta}} z\, \mu(z)\, dS, \tag{3.5}$$

where $\oint_{\partial \mathbf{T}_{\mathcal{E}}^{\theta}}$ is the surface integral over the boundary $\partial \mathbf{T}_{\mathcal{E}}^{\theta}$ of the realization of the random set. As a result, with application of the divergence theorem, Condition (3.4) implies that the parameter θ of

the generating process is encoded in the boundary of all realizations of the random set. We can then interpret this as formalization of the result in Example 1 where the parameter of the data generating process always belongs to the boundary of the realization of the random set and, thus, it can be located as the point in the random set that does not belong to the boundary of the parameter space.

Expression (3.5) expresses the parameter of the data generating process as the center of mass distributed over the boundary of the realization of the random set $\mathbf{T}_{\mathcal{E}}^{\theta}$ with density $\mu(\cdot)$.

The decision mapping of the data curator can constructed in this case by minimizing

$$f(z) = \left\| z - \oint_{\partial \mathbf{T}_{\mathcal{E}}^{\theta}} z\, \mu(z)\, dS \right\|^2 .$$

Our computational strategy relies on two requirements for computing abilities of the data curator. First, the data curator is able to construct realizations of random sets $\mathbf{T}_{\mathcal{E}}$ for data generating process indexed by $\theta \in \Theta$. Second, the data curator represents random sets $\mathbf{T}_{\mathcal{E}}$ by encoding their boundaries (e.g. via vertices of an approximating polyhedron).

Then the goal of the data curator will be to find a numerical approximation to $\mu(\cdot)$ in (3.5). To operationalize this approximation, we define a grid $G = \{\theta^{(k)}\}_{k=1}^{K}$ over the parameter space Θ. For each grid point $\theta^{(k)}$ we simulate a single realization $T_{\theta^{(k)}}$ of the random set $\mathbf{T}_{\mathcal{E}}^{\theta^{(k)}}$.

We represent "density" $\mu(\cdot)$ as

$$\widehat{\mu}(z) = \sum_{r=0}^{R} \alpha_r\, h_r(z)$$

using orthogonal polynomials $h_r(\cdot)$. This representation allows us to reduce the problem of finding function $\mu(\cdot)$ to the problem of finding $R+1$ coefficients of its orthogonal representation.

Then for each simulated realization $T_{\theta^{(k)}}$, we compute $2(R+1)$ surface integrals

$$H_{rk} = \oint_{\partial T_{\theta^{(k)}}} z\, h_r(z)\, dS, \quad \text{and} \quad Q_{rk} = \oint_{\partial T_{\theta^{(k)}}} h_r(z)\, dS,\ r = 0, \ldots, R.$$

Coefficients $\widehat{\alpha}_r$ of the orthogonal representation are found by solving a constrained quadratic optimization problem:

$$\begin{aligned} &\min_{\alpha_0,\ldots,\alpha_R} \sum_{k=1}^{K} \left(\theta^{(k)} - \sum_{r=0}^{R} \alpha_r\, H_r \right)^2, \\ &\text{s.t. } \tfrac{1}{K} \sum_{k=1}^{K} \alpha_r Q_{rk} = 1. \end{aligned} \tag{3.6}$$

The constraint ensures that numerical approximation for the density $\mu(\cdot)$ approximately integrates to 1 over the instances of $T_{\theta^{(k)}}$.

Then the numerical approximation for the convex function used by the data curator for the realization of the random set $\mathbf{T}_{N,\mathcal{E}}$

$$\widehat{f}(z) = \left\| z - \sum_{r=0}^{R} \widehat{\alpha}_r \oint\limits_{\partial \mathbf{T}_{N,\mathcal{E}}} z\, h_r(z)\, dS \right\|^2 .$$

LEMMA 5. *Mapping $\widehat{M}(\cdot)$ constructed as*

$$\widehat{M}(T) = \sum_{r=0}^{R} \widehat{\alpha}_r \oint\limits_{\partial T} z\, h_r(z)\, dS$$

over elements T of Fell topology on Θ where $\widehat{\alpha}_r$ solve (3.6) uniformly converges to mapping $M(\cdot)$ in (3.4).

Proof: By Theorem III.3 in [55] linear partial differential equation

$$\sum_{i=1}^{p} \frac{\partial(z_i \mu(z))}{\partial z_i} = m(z)$$

has solution $\mu(\cdot)$ which is unique for each Lipschitz-continuous function $m(\cdot)$. Moreover, the solution is also Lipschitz-continuous. From the divergence theorem (3.5) mapping $M(\cdot)$ in (3.4) can be represented via surface integral. By uniqueness of representation of $m(\cdot)$ via $\mu(\cdot)$, this surface integral representation is also unique.

Finally, by Lipschitz-continuity of $\mu(\cdot)$ by Weirstrass' theorem it can be approximated by high-degree polynomial with approximation error decreasing as the degree of the polynomial grows *Q.E.D.*

3.2.3 Bayesian credible regions

In some circumstances the data curator may be unable or unwilling to use a decision mapping that produces the target parameter θ_0 for the limiting random set (e.g. due to institutional constraints). The corresponding differentially private output of the decision mapping $\tau_f\left(\mathbb{T}_{N,\mathcal{E}}\right)$ weakly converges to a random limit.

In this case the goal of the researcher would be to use all available information including possible structure of the random set and the rule used by the data curator to produce a differentially private output to make the best possible inference regarding the target parameter.

Our proposed approach of researcher for inference in this setting parallels that in either classic point estimation or set estimation for partially identified models uses a Bayesian approach. Attractive features of the Bayesian approach include, first of all, to incorporate prior information regarding the target parameter. This becomes convenient when preliminary inference can be done using a public dataset that does not require differential privacy guarantees. Then inference using curated dataset that requires differential privacy can be viewed as improving the accuracy of the estimate produced

from the publicly available data. Second, Bayesian approach allows one to produce credible regions for target parameter θ_0. Although, unlike traditional confidence sets, when the limiting random set $\mathbf{T}_{\mathcal{E}}$ is not singleton the credible region does not shrink as sample size increases, operationally it is very similar to the confidence set. This makes it a convenient tool for practitioners to use when differential privacy requirement is imposed by the data curator but the differential private output produced by the data curator is not tailored for the parameter of interest θ_0.

The Bayesian model utilizes the likelihood function $\ell_{f,\mathbb{T}_{N,\mathcal{E}}}(t)$ of the output of the decision mapping constructed from its distribution $\mathbb{P}\left(\tau_f\left(\mathbb{T}_{N,\mathcal{E}}\,;\,\theta\right)\leq t\right)$ induced by the distribution of the random set $\mathbb{T}_{N,\mathcal{E}}$. We explicitly use index $\theta\in\Theta$ for the probability to indicate that $\mathbb{T}_{N,\mathcal{E}}$ is the random set of regular differentially private estimators corresponding to the data generating process indexed by θ.

The Bayesian model then is determined by the tuple $(\mathbb{T}_{N,\mathcal{E}}, f(\cdot), \pi(\cdot))$ where $\pi(\cdot)$ is the prior of the researcher over the values of the parameter of the data generating process.

The researcher combines the likelihood function $\ell_{f,\mathbb{T}_{N,\mathcal{E}}}(t;\theta)$ with the information from the prior to form the posterior distribution for the target parameter given the realization $t=\tau_f\left(\mathbb{T}_{N,\mathcal{E}}\right)$ of the decision mapping of the data curator :

$$\Pi_N(B\,;\,t)\propto\int_B \ell_{f,\mathbb{T}_{N,\mathcal{E}}}(t;\theta)\,\pi(\theta)\,d\theta,$$

where prior $\pi(\cdot)$ plays the role of weighting the likelihood functions corresponding to data generating processes indexed by a given θ.

By weak convergence of random sets $\mathbb{T}_{N,\mathcal{E}}$ to $\mathbf{T}_{\mathcal{E}}$, the posterior $\Pi_N(B\,;\,t)$ converges pointwise to the limiting distribution

$$\Pi(B\,;\,t)\propto\int_B \ell_{f,\mathbf{T}_{\mathcal{E}}}(t;\theta)\,\pi(\theta)\,d\theta,$$

as $N\to\infty$. We then view inference drawn from finite sample prior $\Pi_N(B\,;\,t)$ as an approximation for inference that would be drawn from the limiting distribution of the random set $\mathbf{T}_{\mathcal{E}}$ and the corresponding posterior $\Pi(B\,;\,t)$.

DEFINITION 10. *Let $\Pi_N(B\,;\,t)$ be posterior distribution constructed as above from a single realization t of the differentially private output $\theta(\mathbb{P}_N,\nu_N)$ which converges weakly to the population counterpart $\Pi(B\,;\,t)$. Then set $\mathcal{B}(t)$ with the property*

$$\Pi\left(\mathcal{B}(t)\,;\,t\right)\geq 1-\alpha,\ \alpha\in(0,1)$$

is referred to as the α-level credible region for target parameter θ_0.

α-level credible region is our proposed method to draw informative inference whenever the data curator does not tailor the decision mapping to the target parameter.

EXAMPLE 2 (continued) Suppose that the data curator picks the point in the random set of regular differentially private estimators uniformly at random and the researcher uses uninformative

prior. Based on our previous discussion the likelihood for the difrerentially private output t can then be expressed as

$$\ell_{f,\mathbf{T}_\mathcal{E}}(t;\theta) = \begin{cases} \frac{.5}{\theta}, \text{ if } t < \theta, \\ \frac{.5}{1-\theta}, \text{ if } t \geq \theta. \end{cases}$$

Then the posterior density can be expressed as

$$p(\theta\,;\, t) = \left(-\log\left(t(1-t)\right)\right)^{-1} \left((1-\theta)^{-1}\,\mathbf{1}\{\theta < t\} + \theta^{-1}\,\mathbf{1}\{\theta \geq t\}\right).$$

Thus, we can define credible region $\mathcal{B}(t) = [z(t), 1-z(t)]$ such that

$$\Pi(\mathcal{B}(t); t) = \int_{z(t)}^{1-z(t)} p(\theta\,;\, t)\, d\theta = 1-\alpha.$$

Solving this equation for $z(t)$ yields $z(t) = 1-(t(1-t))^\alpha$ and $\mathcal{B}(t) = [1-(t(1-t))^\alpha,\, (t(1-t))^\alpha]$ whenever data curator produces a differentially private output t.

4 Regression discontinuity design

Regression discontinuity design is an important empirical tool for estimation of treatment effects in a variety of disciplines. The literature on RDD goes back to the work by [53]. A lot of important theoretical and empirical work on RDD has emerged in Economics the last two decades, with too many papers to list here. For a general review of this literature, see [22], [25], [35], [11].

In the usual setting for the RD design, the object of interest is the causal effect of binary treatment on the outcome and units are either exposed or not exposed to a treatment. The effect of the treatment can be heterogenous across units but we will focus on the case of when this effect is homogeneous as our main goal is to highlight the loss of identification of of the treatment effect in these models when estimator is subject to differential privacy guarantees.

Following the tradition of the treatment effect literature, we let Y_0 and Y_1 denote the pair of potential outcomes (without the treatment with exposure to the treatment, respectively) and the actual treatment denoted by Y and defined as

$$Y = W \cdot Y_1 + (1-W) \cdot Y_0,$$

where W is the treatment indicator. The goal is to evaluate the average treatment effect (ATE) of the treatment. The observables are (W, Y, X), where X is a pre-treatment covariate (the so-called *forcing* or *running* variable).

We first give a review of two main designs used in this literature. We then formulate conditions under which differentially private mechanism applied to traditional RDD methods leads to the lack of identification of the treatment effect. We then show that these situations of non-identifiability will be generic due to the global sensitivity of RDD estimators being bounded away from zero even as the sample size increases. These findings no doubt will be of interest to researchers as usually RD design is considered to be one of the most credible identification strategies for causal inference and it loses this powerful feature under differential privacy.

4.1 Overview

4.1.1 Sharp regression discontinuity design

In the sharp design there is a deterministic relation between the running variable and the treatment indicator:

$$W = 1(X \geq c),$$

and the average causal effect β is given by the discontinuity in the conditional expectation of the outcome given the covariate effect of the treatment:

$$\beta = \lim_{X \downarrow c} E[W|X] - \lim_{X \uparrow c} E[W|X].$$

Even though sometimes researchers rely on parametric methods by estimating, for instance, the linear model

$$Y = \alpha + \gamma W + \beta \cdot W + \delta \cdot DW + \varepsilon$$

(or analogous models with more polynomial terms), these approaches may work poorly in practice due to their reliance on a functional form.

The state-of-the art methods are local and rely on learning β from a small neighborhood of c, the size of which becomes increasingly smaller as the sample size increases. In other words, these approaches rely on employing only observations in a small neighborhood $(c-h, c+h)$ of c, where the size of the neighborhood is described by bandwidth $h = h(N)$ that depends on sample size N, where $h(N) \to 0$ as $N \to \infty$. Such estimators can be roughly classified into two categories: *nonparametric regression at the boundary* and *local linear regression.*

Nonparametric regression at the boundary This method selects a kernel $K(\cdot)$ and takes the estimator

$$\hat{\tau}_{S,NR} = \hat{\tau}_r(c) - \hat{\tau}_l(c), \tag{4.1}$$

where

$$\hat{\tau}_r(c) = \frac{\sum_{X_i \geq c} Y_i \cdot K\left(\frac{X_i - c}{h}\right)}{\sum_{X_i \geq c} K\left(\frac{X_i - c}{h}\right)}, \quad \hat{\tau}_l(c) = \frac{\sum_{X_i < c} Y_i \cdot K\left(\frac{X_i - c}{h}\right)}{\sum_{X_i < c} K\left(\frac{X_i - c}{h}\right)}$$

for a chosen bandwidth $h = h(N)$. Even though this estimator is quite intuitive, it has well-known drawbacks particularly with respect to the bias term being linear in bandwidth, as discussed in [22], [48] and [25], among others. Nevertheless, it is instructive for us to analyze a differentially private version of this estimator.

Local linear regression This method conducts two optimization problems by fitting linear regression functions to the observations within an h-neighborhood on either side of the discontinuity

point:

$$(\widehat{\alpha}_L, \widehat{\beta}_L) = \arg\min_{\alpha_L, \beta_L} \sum_{i:X_i<c} K\left(\frac{X_i - c}{h}\right)(Y_i - \alpha_L - \beta_L(X_i - c))^2,$$
$$(\widehat{\alpha}_R, \widehat{\beta}_R) = \arg\min_{\alpha_R, \beta_R} \sum_{i:c\leq X_i} K\left(\frac{X_i - c}{h}\right)(Y_i - \alpha_R - \beta_R(X_i - c))^2,$$

and then estimating the average treatment effect as

$$\widehat{\tau}_{S,LocLin} = \widehat{\alpha}_R - \widehat{\alpha}_L. \tag{4.2}$$

The asymptotic properties of this estimator can be found e.g. in [22], among others, and they are based on the theory in [18] and [19].

There is a significant literature that analyzes the properties and performance of these estimators as well as the choice of the bandwidth and the construction of confidence sets. These issues, even though important, are not immediately relevant to our research question which is to illustrate that a generic differential privacy algorithm will prevent researchers to be able to to identify the average treatment effect in RD design in the limit of statistical experiments. We aim to keep the discussion as general as possible without focusing on just one specific local estimation approach.

4.1.2 Fuzzy regression discontinuity design

In the fuzzy RD design, there is a jump in the *probability* of assignment to the treatment at the threshold:

$$\lim_{x\uparrow c} P\left(W_i = 1 | X_i = x\right) \neq \lim_{x\downarrow c} P\left(W_i = 1 | X_i = x\right),$$

and average treatment effect in defined as

$$\tau_{FRD} = \frac{\lim_{x\downarrow c} E[Y|X = x] - \lim_{x\uparrow c} E[Y|X = x]}{\lim_{x\downarrow c} P(W = 1|X = x) - \lim_{x\uparrow c} P(W = 1|X = x)}.$$

One way to estimate τ_{FRD} would be to use the estimator in (4.1). As first pointed out by [22], in the context of the fuzzy design this estimator is a Wald estimator or, in other words, the IV estimator that treats W_i as an endogenous regressor and the indicator $1(X_i \geq c)$ as an instrument.

The local linear regression estimator is defined as

$$\widehat{\tau}_{F,LocLin} = \frac{\widehat{\alpha}_{y,R} - \widehat{\alpha}_{y,L}}{\widehat{\alpha}_{w,R} - \widehat{\alpha}_{w,L}}, \tag{4.3}$$

where the objects on the right-hand side are obtained from the following estimations:

$$\left(\widehat{\alpha}_{y,L},\widehat{\beta}_{y,L}\right)=\arg\min_{\alpha_{y,L},\beta_{y,L}}\sum_{i:c-h_N\leq X_i<c}(Y_i-\alpha_{y,L}-\beta_{y,L}(X_i-c))^2,$$
$$\left(\widehat{\alpha}_{y,R},\widehat{\beta}_{y,R}\right)=\arg\min_{\alpha_{y,R},\beta_{y,R}}\sum_{i:c\leq X_i\leq c+h_N}(Y_i-\alpha_{y,R}-\beta_{y,R}(X_i-c))^2,$$
$$\left(\widehat{\alpha}_{w,L},\widehat{\beta}_{w,L}\right)=\arg\min_{\alpha_{w,L},\beta_{w,L}}\sum_{i:c-h_N\leq X_i<c}(W_i-\alpha_{w,L}-\beta_{w,L}(X_i-c))^2,$$
$$\left(\widehat{\alpha}_{w,R},\widehat{\beta}_{w,R}\right)=\arg\min_{\alpha_{w,R},\beta_{w,R}}\sum_{i:c\leq X_i\leq c+h_N}(W_i-\alpha_{w,R}-\beta_{w,R}(X_i-c))^2.$$

As is well known, this estimator can be interpreted as the IV estimator with the main equation being the regression of Y_i on exogenous constant, $1\left(X_i-c<0\right)\left(X_i-c\right)$, $1\left(X_i-c\geq 0\right)\left(X_i-c\right)$ and endogenous W_i, while using the indicator $1(X_i\geq c)$ as the excluded instrument. This estimator can, of course, be easily generalized to include more polynomial terms in each estimation. In the definition of $\widehat{\alpha}_{y,L}$, $\widehat{\alpha}_{y,R}$ and $\widehat{\alpha}_{w,L}$, $\widehat{\alpha}_{w,R}$ we, for simplicity, used the uniform kernel. However, one could use other kernels.

4.2 Asymptotic behavior of global sensitivity and parameter non-identifiability

In this section, we establish general results that connect the asymptotic behavior of the global sensitivity with the inconsistency of regular differentially private estimators, allowing us, in light of our notion of identification in Section 2, to make conclusions about the non-identifiability of the parameter of interest under differentially private mechanisms.

We consider smooth estimators as given in Definition 6. The smoothness property allows us to essentially consider estimators with an additive mechanism noise ξ_N:

$$\widehat{\theta}=\psi(\mathbb{P}_N)+\xi_N \tag{4.4}$$

(as discussed in the introduction, ν_N in the Definitions 3 and 6, among others, play the role of the "seed" and, thus, the actual independent from the data additive noise ξ_N is a transformation of ν_N). For this estimator to satisfy regularity requirements in Definition 3, it has to, among other things, be consistent in the absence of any mechanism noise, thus giving the identification of the parameter in the limit of statistical experiments. This immediately leads us to the condition that

$$\psi(\mathbb{P}_N)\xrightarrow{p}\theta_0, \tag{4.5}$$

where θ_0 denotes the true parameter value.

Suppose the family from which the distribution of ξ_N is drawn is described by the density $f_{N;\sigma^2}$, where σ^2 denotes the variance (the actual value of σ^2 in practice depend on the sample size). The notation $f_{N;\sigma^2}$ is not meant to say that the distributional family for the additive noise is fully described by the variance parameter. This parameter is introduced explicitly in the notation as

usually the situations when differentially private mechanisms do not prevent the identification of θ_0 in the limit are characterized the behavior of this parameter as the sample size increases.[8] Even though for now we consider just one distributional family, one has to keep in mind that potentially several different distributions could be used, in which one has to consider all of them. To keep this exposition simple, we will focus on one family f_{σ^2}, as this already will give us a rather comprehensive analysis.

For any $\Delta > 0$, σ^2 and $a < b$, define

$$\mathcal{D}_N(\Delta, \sigma^2, a, b) \equiv \sup_{z \in [a,b]} |\log f_{N;\sigma^2}(z) - \log f_{N;\sigma^2}(z + \Delta)|, \tag{4.6}$$

which is the discrepancy between the logarithms of two densities with a fixed variance, from which one is calculated with a shifts Δ, on some interval $[a, b]$. As can be seen from the Definition 1 of differential privacy, the behaviour of this object across different Δ and σ^2 is directly related to the parameters of the differential privacy. The properties of a density imply that if $[a, b]$ is large enough, then for a fixed sample size N we have

$$\mathcal{D}_N(\Delta, \sigma^2, a, b) \to +\infty \quad \text{as } \sigma^2 \to 0, \quad \Delta \neq 0. \tag{4.7}$$

In the majority of applications the family of distributions $f_{N;\sigma^2}$ is not indexed by N and the change in the distribution of the estimators with the sample size may be driven by different variances. In fact, this is the case for all commonly used mechanisms – Laplace, Gaussian, exponential mechanisms and their variations. To make our discussion more general, we can allow for different distributional families across different N, in which case we will require the following

Algorithm condition 1 (AC1). *For any Δ, σ^2 and $a < b$*

$$\sup_N \mathcal{D}_N(\Delta, \sigma^2, a, b) \leq \mathcal{D}(\Delta, \sigma^2, a, b) \tag{4.8}$$

$$\mathcal{D}(\Delta, \sigma^2, a, b) \to +\infty \quad \textit{as } \sigma^2 \to 0, \quad \Delta \neq 0. \tag{4.9}$$

As is clear from the discussion later, in the differential privacy implementation Δ will be associated with with the global sensitivity of $\psi(\mathbb{P}_N)$. Condition AC1 tells us that if the global sensitivity remains bounded away from 0 as the sample size increases, then the diminishing variance of the noise and (ϵ_N, δ_N), $\epsilon_N \leq \bar{\epsilon}$, $\delta_N \leq \bar{\delta}$ differential privacy guarantees are incompatible with each other. This will allow us to immediately draw conclusions about the situation of non-identifiability in the limit of statistical experiments of the parameter of interest.

We give another algorithmic condition on the family of the distribution of the noise variable, which holds generally for differentially private mechanisms and which also help to establish further results related to the identifiability (or lack of such) of the parameter.

[8] For very popular mean-zero Laplace and Gaussian mechanisms, the distribution of ξ_N is fully characterized by the variance parameter.

Algorithm condition 2 (AC2). *For any $\Delta_N \to 0$ as $N \to \infty$, it is possible to indicate $\sigma_N^2 \to 0$ and $[a_N, b_N]$ such that*

$$\mathcal{H}\left([a_N, b_N], \mathbb{R}\right) \to 0$$

and

$$\mathcal{D}(\Delta_N, \sigma_N^2, a_N, b_N) \to 0, \tag{4.10}$$

where $\mathcal{D}(\Delta, \sigma^2, a, b)$ is as defined in (4.10).

To give an example, consider the Laplace mechanism (discussed in Example 1) and note that $\log f_{2\lambda^2}(z + \Delta) = -\frac{|z+\Delta|}{\lambda}$, and thus,

$$\mathcal{D}(\Delta, \sigma^2, a, b) \leq \frac{|\Delta|}{\lambda},$$

which in particular implies that we can even take $a = -\infty$ and $b = +\infty$ in AC2. We can see that this case trivially satisfies AC2 and, of course, condition AC1 if one chooses the Laplace mechanism for any N.

It is exactly condition AC2 that would give a hope for the identifiability of the parameter of interest in cases when global sensitivity of $\psi(\mathbb{P}_N)$ goes to 0. Indeed, the convergence $\Delta_N \to 0$ is meant to capture the case when the global sensitivity of the estimator converges to zero.

When this happens, condition AC2 says that it is possible to take $\sigma_N^2 \to 0$ and $[a_N, b_N]$ converging to the whole real line at the right rates (which clearly depend on the rate of a decreasing global sensitivity) in such a way that $\mathcal{D}(\mu_{1,N}, \mu_{2,N}, \sigma_N^2, a_N, b_N)$ remains bounded by a small $\varepsilon_N \to 0$ with the probability of at least $1 - \delta_N$. This will ensure that the differentially privacy criteria will be satisfied for some sequences $\epsilon_N, \delta_N)$ converging to $(0, 0)$ while delivering consistent differentially private estimator. This is given in Proposition 1 below.

Proposition 1. *Consider a smooth (ϵ_N, δ_N)-differentially private estimator – without a loss of generality represented as (4.4), – and suppose that in the absence of the mechanism noise this estimator (denoted as $\psi(\mathbb{P}_N)$) is consistent, i.e. (4.5) holds. Suppose that the global sensitivity of $\psi(\mathbb{P}_N)$ denoted as $G(N)$ converges to 0 with the sample size and the mean on ξ_N converges to 0. If AC2 holds, then the differentially private estimator is consistent if ϵ_N and δ_N both converging to 0 have slow enough rates.*

Note that under the conditions of Proposition 1 implies that an (ϵ, δ)-differentially private estimator is consistent for fixed (ϵ, δ) as well since these requirements are weaker than requiring that (ϵ_N, δ_N) converges to 0. The result of Proposition 1 gives a hope that the parameter of interest can be identified in the limit of experiments with a suitable choice of the set $\mathcal{E}$ of sequences (ϵ_N, δ_N).

Our next step is to establish that a generic $(\varepsilon_N, \delta_N)$-differentially private algorithm will give an inconsistent estimator if the global sensitivity remains bounded away from zero as the sample size increases even if the mean of the mechanism noise ξ_N converges to 0.

THEOREM 6. *Consider a smooth (ϵ_N, δ_N)-differentially private estimator $\widehat{\theta}$ – without a loss of generality represented as (4.4), – and suppose that in the absence of the mechanism noise this estimator (denoted as $\psi(\mathbb{P}_N)$) is consistent, i.e. (4.5) holds. Suppose that AC1 holds and the global sensitivity of $\psi(\mathbb{P}_N)$ denoted as $G(N)$ does not converge to 0 with the sample size whereas the mean on ξ_N converges to 0.*

Then this (ϵ_N, δ_N)-differentially private estimator is inconsistent even if ϵ_N does not change with N.

We end this section by giving sufficient condition for when the parameter of interest is not identified from the differential privacy estimation in the limit of statistical experiments.

COROLLARY 1. *Consider a class of smooth (ϵ_N, δ_N)-differentially private estimators $\widehat{\theta}$ – without a loss of generality represented as (4.4), – and suppose that in the absence of the mechanism noise these estimators (correspond to $\psi(\mathbb{P}_N)$) are consistent, i.e. (4.5) holds. Suppose that AC1 holds and the global sensitivity of $\psi(\mathbb{P}_N)$ denoted as $G(N)$ does not converge to 0 with the sample size whereas the mean on ξ_N converges to 0.*

For any join-semilattice $\mathcal{E}$ of sequences of (ϵ_N, δ_N), with $\epsilon_N \leq \bar{\epsilon}$, $\delta_N \leq \bar{\delta}$, parameter θ_0 is not identified in the limit of experiments.

This corollary directly follows from Theorem 6.

4.3 Regression discontinuity under differential privacy

In this section, we use results of Section 4.2 to analyze the identifiability of the average treatment effect of differentially private regression discontinuity design estimators. Even though this section will be focused on the issue of identifiability, there are other important issues one would want to explore in the RDD framework. One of these issues is the question of how differential privacy requirements would affect the visual analysis of the data, which is on e of the fundamental steps in the practice of RDD. Another issue is the the question of the credibility of specification tests (continuity of the density of the running variable at the cut-off, placebo tests with pre-treatment covariates), under differential privacy. Even though our main focus is on identifiability of the average treatment effect under differential privacy, we do discuss these other related issues in Section 4.4 albeit in less detail.

Even though sometimes there are some parametric RDD estimation methods which are global in nature, the state-of-the-art RDD techniques are local in nature and employ some elements of non-parametric methods. These latter methods focus on a neighborhood around the switch point with the size of this neighborhood being determined by a kernel $K(\cdot)$ and a respective bandwidth $h = h(N)$. We will suppose that $h(N)$ is chosen by a certain differential privacy algorithm according to some rule in such a way that $h(N) = o(1)$ as $N \to \infty$. Then the expected number of observations from a sample of size N in a right-hand side neighborhood of c is $N \cdot Pr\left(c \leq X < c + h(N)\right)$, and in the

left-hand side neighborhood is $N \cdot Pr\,(c - h(N) < X < c)$. There are, of course, some well known approaches for selecting a bandwidth, such as [24], [10], among others. Our analysis will apply to general bandwidth choices subject.

4.3.1 Sharp RD under differential privacy

We start with the analysis of differentially private RDD estimators for the sharp design. We begin our series of formal results with establishing the results on the global sensitivity of nonparametric regression at the boundary and local linear (polynomial) estimation. Propositions 2-4 look at nonparametric regression at the boundary and various properties of kernels that affect the global sensitivity result. Proposition 5 looks at the local linear estimation. In light of the results in Section 4.2, the knowledge of the asymptotic behavior of the global sensitivity of these estimators will allow us to analyze wether smooth differentially private approaches are compatible with the identifiability of the ATE of interest. As we show, the exact results on the global sensitivity even depend on the type of kernel used in the above-mentioned estimation techniques.

Before formulate to Proposition 2, we formulate what we mean by kernels with a bounded support.

DEFINITION 11. *We say that the kernel function $K(\cdot) : \mathbb{R} \to \mathbb{R}^+$ has a bounded support if there is a value $u_0 > 0$ such that $K(u) = 0$ when $|u| > u_0$. If this condition is not satisfied, then we will say that the kernel has an unbounded support.*

Uniform, Epanechnikov, triangular kernels are examples of kernel functions with bounded supports. Gaussian and logistic kernels are examples of kernel functions with unbounded supports. Even if one considers kernels with a bounded support, we will see that it will make a difference whether the kernel is continuous (like the triangular kernel) or has discontinuities (like the uniform kernel).

DEFINITION 12. *For a given kernel function $K(\cdot)$ with a bounded support and a given bandwidth h, we define a K-h-neighborhood to the right of c as a set $[c, c + \Delta_{K,r}(h))$, where $\Delta_{K,r}(h) > 0$, such that $K(\frac{u-c}{h}) > 0$ for $u > c$ if and only if $u \in [c, c + \Delta_r(h))$. In other words, this is the set of point to the right of c that will be used in the nonparametric regression.*

Analogously, we define a K-h-neighborhood to the left of c as a set $(c-\Delta_{K,l}(h), c)$, where $\Delta_{K,l}(h) > 0$, such that $K(\frac{u-c}{h}) > 0$ for $u < c$ if and only if $u \in (c - \Delta_{K,l}(h), c)$.

A differentially private algorithm takes the support of the observed variables as given and usually depends on this support, and thus, uses the supports of Y in both right-hand side and left-hand side K-h-neighborhoods as inputs. For a kernel with a bounded support, these supports can be denoted as $\mathcal{Y}^r(h)$ and $\mathcal{Y}^l(h)$, respectively, and may generally depend on h. However, they are naturally approximated by

$$\mathcal{Y}^r = \lim_{h \downarrow 0} \mathcal{Y}^r(h) \text{ and } \mathcal{Y}^l = \lim_{h \downarrow 0} \mathcal{Y}^l(h) \tag{4.11}$$

that no longer depend on the bandwidth choice (these limits are well defined as $\{\mathcal{Y}^r(h)\}$ and $\{\mathcal{Y}^l(h)\}$ are sequences of monotonically decreasing events.)

We will suppose that $\mathcal{Y}^r$ and $\mathcal{Y}^l$ are convex non-singleton sets. As further notations, we will use

$$\overline{Y}^r = \sup \mathcal{Y}^r, \quad \underline{Y}^r = \inf \mathcal{Y}^r,$$
$$\overline{Y}^l = \sup \mathcal{Y}^l, \quad \underline{Y}^l = \inf \mathcal{Y}^l.$$

We now present a series of results on the global sensitivity.

Proposition 2. *Consider a nonparametric regression at the boundary estimator that uses a* ***continuous kernel with a bounded support****.*

Suppose that for a data-driven choice of bandwidth $h = h(N)$, for any sample size N it is possible to have realizations of the data $\{(Y_i, W_i, X_i)\}_{i=1}^N$ that will deliver the minimum number of observations $m^r(N) \geq 1$ in the K-h-neighborhood to the right of c and the minimum number of observations $m^l(N) \geq 1$ in the K-h-neighborhood to the left of c.

(a) If the supports $\mathcal{Y}^r$ and $\mathcal{Y}^l$ are bounded, then the global sensitivity of the nonparametric regression at the boundary estimator is

$$\overline{Y}^r - \underline{Y}^r + \overline{Y}^l - \underline{Y}^l$$

and, hence, it does not depend on the sample size.

(b) If at least one of the supports $\mathcal{Y}^r$ and $\mathcal{Y}^l$ is unbounded, the global sensitivity of of the nonparametric regression at the boundary estimator is $+\infty$.

We next consider the case of a kernel function with a bounded support and discontinuities at the support boundaries $-u_0$ and u_0 with the main example here being the uniform kernel. We define

$$\underline{K} \equiv \inf_{u \in (-u_0, u_0)} K(u). \tag{4.12}$$

The discontinuities of $K(\cdot)$ at $-u_0$ and u_0 imply that $\underline{K} > 0$. For expositional simplicity, in the formulation of Proposition 3 we only indicate the rate of the global sensitivity. However, the proof of this proposition in the Appendix gives the exact expression for this sensitivity.

Proposition 3. *Consider a nonparametric regression at the boundary that uses a* ***kernel with a bounded support and*** $\underline{K} > 0$*, where $\underline{K}$ is as defined in 4.12.*

Suppose that for a data-driven choice of bandwidth h, for any sample size N it is somehow possible to guarantee the minimum number of observations $m^r(N) \geq 1$ in the K-h-neighborhood to the right of c and the minimum number of observations $m^l(N) \geq 1$ in the K-h-neighborhood to the left of c.

(a) If the supports $\mathcal{Y}^r$ and $\mathcal{Y}^l$ are bounded are bounded, then the global sensitivity of the nonparametric regression at the boundary estimator is proportional to $\frac{1}{\min\{m^r(N), m^l(N)\}}$.

(b) If at least one of the supports $\mathcal{Y}^r$ or $\mathcal{Y}^l$ is unbounded, the global sensitivity of of the nonparametric regression at the boundary estimator is $+\infty$.

Part (a) in Proposition 3 seemingly gives some hope of achieving the situation when the global sensitivity may be going to zero as $N \to \infty$ if it can be ensured that $\min\{m^r(N), m^l(N)\} \to \infty$. This hope, however, is a false one as the situation of being able to guarantee a minimal number (growing to ∞) of observations in each neighborhood for any sample $\{X_i\}_{i=1}^N$ with a given support of X is a rather hypothetical scenario as the probabilities that the number of observations in the right- and left-hand side neighborhoods is strictly less than $m^r(N)$ and $m^l(N)$, respectively:

$$\sum_{k=0}^{m^r(N)} \binom{N}{k} F_X(c)^{N-k} \left(1 - F_X(c)\right)^k$$

is the probability of fewer than $m^r(N)$ observations to the right of c,

$$\sum_{k=0}^{m^l(N)} \binom{N}{k} F_X(c)^{k} \left(1 - F_X(c)\right)^{N-k}$$

is the probability of fewer than $m^l(N)$ observations to the left of c.

Both these probabilities are clearly strictly positive when c is an interior point of the support of X. The non-stochastic nature of the global sensitivity concept effectively leads to the situation of the global sensitivity always being bounded away from 0 as $N \to \infty$ in the case of the kernel with a bounded support and $\underline{K} > 0$.

Our final case is that of a kernel with an unbounded support, that is, the kernel is strictly positive on the whole real line and approaches 0 at $\pm\infty$ (such as the Gaussian kernel). In this case we will take that a differentially private algorithm uses the supports of $Y|X \geq c$ and $Y|X < c$ since the kernel weights are technically never equal to zero. Let's denote these supports as $\mathcal{Y}^r_{all}$ and $\mathcal{Y}^l_{all}$, respectively.[9] Also denote

$$\overline{Y}^r_{all} = \sup \mathcal{Y}^r_{all}, \quad \underline{Y}^r_{all} = \inf \mathcal{Y}^r_{all},$$
$$\overline{Y}^l_{all} = \sup \mathcal{Y}^l_{all}, \quad \underline{Y}^l_{all} = \inf \mathcal{Y}^l_{all}.$$

Since the kernel approaches zero arbitrarily closely, the global sensitivity results for this case will be similar to those in Proposition 4 where the infimum of the values of the kernel on its support bounded support is 0.

Proposition 4. *Consider a nonparametric regression at the boundary that uses a kernel function with a unbounded support.*

[9] Potentially differentially private algorithms may use more a complicated support for $Y|X$ that could depend on X. This is not going to change our qualitative findings on the global sensitivity being bounded away from 0, even though the exact numerical values for global sensitivities may be different.

(a) *If the supports $\mathcal{Y}^r_{all}$ and $\mathcal{Y}^l_{all}$ are bounded, then the global sensitivity of a nonparametric regression at the boundary estimator is*

$$\overline{Y}^r_{all} - \underline{Y}^r_{all} + \overline{Y}^l_{all} - \underline{Y}^l_{all}$$

– that is, it does not depend on the sample size.

(b) *If at least one of the supports of $\mathcal{Y}^r_{all}$ and $\mathcal{Y}^l_{all}$ is unbounded, then the global sensitivity of a nonparametric regression at the boundary estimator is $+\infty$.*

Thus, the results of Propositions 2-4 and the discussion following Proposition 3 lead us to conclude that the global sensitivity of a nonparametric at boundary estimator is always bounded away from zero. The implications of this for asymptotic properties of differentially private estimators is given in Theorem 6 and Corollary 1 allow us to conclude that the *ATE is not identified in the limit of statistical experiments of smooth differentially private nonparametric regression at boundary estimators.*

Our next step is to analyze whether things get better with a local linear (and more generally, polynomial) estimator as defined in (4.2). Proposition 5 below establishes that this is not the case and the global sensitivity of this estimator is in fact infinite even if the support of the outcome variable is bounded.

Proposition 5. *Consider the local linear estimator as defined in (4.2). The global sensitivity of this estimator is bounded away from zero as $N \to \infty$.*

Theorem 6 and Corollary 1 allow us to conclude that the *ATE is not identified in the limit of statistical experiments of smooth differentially private local linear estimators.*

Our findings for the sharp design are summarized in Theorem 7.

THEOREM 7. *In the sharp regression discontinuity design case, any smooth $(\varepsilon_N, \delta_N)$-differentially private nonparametric regression at the boundary estimator and any smooth $(\varepsilon_N, \delta_N)$-differentially private local linear estimator is inconsistent for any bounded sequences of positive $\{\varepsilon_N\}$ and non-negative $\{\delta_N\}$.*

If we add other covariates to our estimation or use more terms in the local polynomial estimation, the conclusion of Theorem 7 remains exactly the same, even though in Propositions 2-5 quantitatively the global sensitivities of the estimators could be different. Indeed, from the proofs in the Appendix one could easily see that the global sensitivity for the local polynomial estimator would once again rely on the formula for the OLS estimator for the intercept whereas in the nonparametric regression at the boundary estimator the lower bounds on global sensitivities could be obtained in the same way as the proofs in Propositions 2-4.

4.3.2 Fuzzy RD under differential privacy

In the case of the fuzzy design the results for the global sensitivities of the estimators are analogous to the sharp design case. Naturally, we are also able the conclude that differentially pirate versions of traditional estimators in this framework are inconsistent.

Indeed, as discussed in Section 4.1.2, the estimator (4.1) may be used in the fuzzy design case as well, with the asymptotic properties analogous to the sharp design scenario. The global sensitivity of this estimator remains the same and, therefore, its properties are described by Propositions 2-4. As for the local linear estimator, the result is the same, once again, even though the proof is slightly more elaborate than in the sharp design case. For the sake of completeness, we establish this result formally in Proposition 6 below.

Proposition 6. *Consider the local linear estimator as defined in (4.3). The global sensitivity of this estimator is bounded away from 0 as* $N \to \infty$.

The proof of Proposition 6 in the Appendix for simplicity does not employ other covariates. It is worth mentioning, however, that the situation with other covariates (let's call them S_i with the support $\mathcal{S}$) may be even worse as in the case we may have that

$$\inf_{S_i \mathcal{S}} \left| \lim_{x \uparrow c} P\left(W_i = 1 | X_i = x\right) \neq \lim_{x \downarrow c} P\left(W_i = 1 | X_i = x\right) \right| = 0.$$

If this situation occurs, then even in the nonparametric regression at the boundary type of estimators the global sensitivities may be converging to ∞ as $N \to \infty$, giving even more severe implications for the asymptotic properties of smooth differentially private estimators.

In the definition of the local linear estimator in Section 4.1.2 we for simplicity used the uniform kernel. However, the result of Proposition 6 remains true of other kernels are used.

Then, relying on the results in Propositions 2-5, 6 and Theorem 6, we can immediately obtain the result of Theorem 8 below.

THEOREM 8. *In fuzzy regression discontinuity design, any smooth* $(\varepsilon_N, \delta_N)$*-differentially private nonparametric regression at the boundary estimator and any smooth* $(\varepsilon_N, \delta_N)$*-differentially private local linear estimator is inconsistent for any bounded sequences of positive* $\{\varepsilon_N\}$ *and non-negative* $\{\delta_N\}$.

To summarize the results of Theorems 7 and 8 in Sections 4.3.1 and 4.3.2, in the sharp and fuzzy regression discontinuity design *the requirements of* (ε, δ)*-differential privacy either with fixed* ε, δ *or with these parameters decreasing with the sample size are incompatible with the consistent estimation of the average treatment effect. Therefore, given our notion of identifiability in Section 2, they are also incompatible with the identifiability of the average treatment effect in the limit of statistical experiments.*

4.4 Specification testing

So far we mostly have focused on the non-identifiability of the average treatment effect under smooth differentially private mechanisms. However, every regression discontinuity design analysis is traditionally accompanied by specification testing. This includes checking for the possibility of other changes at the cutoff value c of the forcing variable X_i and also checking for the manipulation of X_i.

The first type of checks include testing the null hypothesis of a zero average effect on pseudo outcomes known not to be affected by the treatment. The outcomes of such placebo tests would also have to be $(\varepsilon_N, \delta_N)$-differentially private. A traditional differentially private literature approach in this case would add noise to the true test statistic and then adjust the asymptotic distribution to compute correct p-values (see e.g. [54]).[10] It is not surprising that, once again, this brings a range of issues in the context of placebo tests in regression discontinuity designs. Indeed, let $\widehat{\tau}_{pl}$ denote the true regression discontinuity design estimator in the analysis of treatment of pseudo-outcomes, and let τ_{pl} denote the true parameter. The testing of the null $H_0 : \tau_{pl} = 0$ is based on the t-ratio $\frac{\widehat{\tau}_{pl}}{se(\widehat{\tau}_{pl})}$. Without giving formal results on this, we nevertheless want to point out that utilizing our techniques in the proofs of Propositions 2-5 and 6, we can establish that the global sensitivity of this ratio either increases to ∞ with the sample size or is already ∞ in a finite sample. This implies that the noise added to this ratio would asymptotically dominate this ratio. Even if the critical values are corrected to account for the added noise, it is clear that the conclusion of the $(\varepsilon_N, \delta_N)$-differentially private test based on such a procedure are not credible and, in particular, result in a lower power of the test (in the limit, the power of this test is trivial). To make our point more transparent, let us focus on a stylized version of the test when the asymptotic variance of $\sqrt{Nh}\,(\widehat{\tau}_{pl} - \tau_{pl})$ is known. We will denote is as $Avar\,(\widehat{\tau}_{pl})$. In this stylized version we want to create a differentially private version of the test statistic $t_N = \frac{\sqrt{Nh}\widehat{\tau}_{pl}}{\sqrt{Avar(\widehat{\tau}_{pl})}}$.

As we have shown in Sections 4.3.1 and 4.3.2, the global sensitivity of $\widehat{\tau}_{pl}$ may be constant and bounded away from zero or may even be infinite for every N (in situations when we add other covariates, it may be increasing to ∞ wth the sample size). Given that $h = h(N)$ is chosen in a way to give $Nh \to \infty$, this will imply immediately that the global sensitivity of t_N is either increasing to infinity with the sample size or is infinite. This means that the variance of the independent noise added in the differentially private algorithm will increasingly dominate the asymptotically constant variance of t_N. Instead of using the standard normal distribution critical values, one would take the critical values from the distribution that suitably combines the standard normal distribution and the distribution of noise. However, as N increases, the testing essentially becomes inference about the mechanism noise, leading to a decreasing power of the test, which asymptotically diminishes to the trivial power. [11]

[10] There are also approaches to hypotheses testing in the differential privacy literature that are based on adding noise to the inputs – see e.g. [27]. They. may, however, be considered less reliable than the approaches based on the output perturbation.

[11] In contrast, in a simpler case of the estimation of mean, as in Example 1, the global sensitivity of t-ratio would be finite and bounded away from zero as N increases. Upon the correction of the critical values, in this situation the

A second type of tests are for the manipulation of the forcing variable. We will illustrate issues associated with (ε, δ)-differentially private versions of these tests by considering the test of the continuity of the density at cutoff by [41]. The test is based on the ratio $\frac{\widehat{\theta}}{\widehat{\sigma}_\theta}$, where

$$\widehat{\theta} = \ln \widehat{f}^+ - \ln \widehat{f}^-, \quad \widehat{\sigma}_\theta = \sqrt{\frac{1}{Nh}\frac{24}{5}} \left(\frac{1}{\widehat{f}^+} + \frac{1}{\widehat{f}^-} \right),$$

where h is the bandwidth and

$$\widehat{f}^+ = \sum_{X_i > c} K\left(\frac{X_i - c}{h}\right) \frac{S^+_{N,2} - S^+_{N,1}(X_i - c)}{S^+_{N,2} S^+_{N,0} - \left(S^+_{N,1}\right)^2} Y_i = \sum_{X_i > c} \frac{K\left(\frac{X_i - c}{h}\right)}{M^+_N} \frac{\frac{S^+_{N,2}}{M^+_N} - \frac{S^+_{N,1}}{M^+_N}(X_i - c)}{\frac{S^+_{N,2}}{M^+_N}\frac{S^+_{N,0}}{M^+_N} - \left(\frac{S^+_{N,1}}{M^+_N}\right)^2} Y_i, \quad (4.13)$$

$$\widehat{f}^- = \sum_{X_i < c} K\left(\frac{X_i - c}{h}\right) \frac{S^-_{N,2} - S^-_{N,1}(X_i - c)}{S^-_{N,2} S^-_{N,0} - \left(S^-_{N,1}\right)^2} Y_i = \sum_{X_i < c} \frac{K\left(\frac{X_i - c}{h}\right)}{M^-_N} \frac{\frac{S^-_{N,2}}{M^-_N} - \frac{S^-_{N,1}}{M^-_N}(X_i - c)}{\frac{S^-_{N,2}}{M^-_N}\frac{S^-_{N,0}}{M^-_N} - \left(\frac{S^-_{N,1}}{M^-_N}\right)^2} Y_i, \quad (4.14)$$

where $M^+_N = \sum_{X_i > c} K\left(\frac{X_i - c}{h}\right)$, $S^+_{N,k} = \sum_{X_i > c} K\left((X_i - c)/h\right)(X_i - c)^2$ and analogous definition for the objects with the minus superscript. In the rewritten expression for $\widehat{f}^+$ and $\widehat{f}^-$ on the right-hand side of (4.13) and (4.14), we can see that these definitions are weighted averages with the weights $\frac{K((X_i - c)/h)}{M^+_N}$ and $\frac{K((X_i - c)/h)}{M^-_N}$, respectively. This brings us to the situations analogous to the ones in Propositions 2-4 and, thus, allows us to analyze the behavior of the global sensitivities of $\widehat{\theta}$ and the ratio $\frac{\widehat{\theta}}{\widehat{\sigma}_\theta}$ using similar tools. Even if only one observation in a sample changes, these weights may vary from 0 to 1 for continuous kernels with abounded support (like in Proposition 2) and for kernels with an unbounded support (like in Proposition 4), and they may vary within a range that is bounded away from 0 for kernels with a bounded support and $\underline{K}$ described in Proposition 3. These are exactly the features that would allow us to establish that the global sensitivity of $\widehat{\theta}$ is bounded away from zero (may even be infinite in some situations) as $N \to \infty$ and the global sensitivity of $\frac{\widehat{\theta}}{\widehat{\sigma}_\theta}$ goes to ∞ (or may even be infinite for a finite N) as $N \to \infty$. This implies that for smooth estimators the independent mechanism noise combined with the test statistic $\frac{\widehat{\theta}}{\widehat{\sigma}_\theta}$ to make such a ratio $(\varepsilon_N, \delta_N)$-differentially private, will dominate the asymptotic behavior of the differentially private ratio. Even if the critical values are corrected by taking into account the distribution of the mechanism noise, the conclusions of this test are not credible and the power of the test is very low (in the limit, the power of this test is trivial).

4.5 Graphical analyses

Graphical analyses have become an essential part of regression discontinuity design applications as they give a powerful way to visualize the identification strategy of the RD design. [25], among others, summarizes three type of analyses particularly useful.

significance testing would still have non-trivial power even though the power of the test would be lower.

The first one is the histogram-type estimate of the average value of the outcome for different values of the forcing variable. Usually, the support of the forcing variable is split into a number of equal-sized bins, with first bins being defined at the cut-point and then taken further to the right and to the left. It is essential that no bin mounts over the cut-off point, as this graph is meant to provide a visual guidance about whether there is in fact a discontinuity at the cut-off. In addition, it allows to analyze whether there is evidence of jumps at points other than the cut-off, thus helping to analyze whether or not the jump at the cut-off can be solely attributed to the treatment of interest (see the discussion in [35]). In an nutshell, in this analysis one would want to plot the average bin values

$$\mathcal{Y}_k = \frac{\sum_{i=1}^{N} Y_i \cdot 1\left(b_k < X_i \leq b_{k+1}\right)}{\sum_{i=1}^{N} 1\left(b_k < X_i \leq b_{k+1}\right)}$$

against the average points $\frac{b_k+b_{k+1}}{2}$ of bins $[b_k, b_{k+1}]$.

The differentially private literature has developed a variety of methods for outputting differentially private histograms (for a review see e.g. [43]). Starting with first papers on differentially private histograms (such as [39]), it has been recognized that the partitioning can leak information about the data and for that reason usually the data are split into bins using a privacy-preserving clustering algorithm that is compatible with differential privacy, such as the k-means algorithm or similar. This means that usually the bins would be chosen by an algorithm rather than a data curator or a researcher. in such a situation in general the cut-off point will not be the separator of two of the bins, and when it is not, then the main purpose of this visual analysis is defeated. In case a researcher wants to impose restrictions that the cut-off is a separator of the two bins, then the issues encountered under these restrictions will be similar to the ones in the nonparametric regression at the boundary estimation with the uniform kernel (see Proposition 3). In particular, this means that generally the global sensitivity of means in the bins next to the threshold does not converge to zero with the sample size, implying that in any $(\varepsilon_N, \delta_N)$-differentially private histogram-like output for conditional means of the outcome the impact of the mechanism noise will be persistent for any sample size, leading to the lack of credibility of such an analysis.

The second type of analyses is similar to the first one but plots average bin values of other covariates against the average points of bin values. These types of graphs are useful in detecting potential specification problems. The issues with delivering differentially private graphs are completely analogous to those described in the first type of analyses above.

Finally, the third type of analyses plots the histogram of the distribution of the forcing variable as this helps to inspect whether there is a discontinuity in the distribution of the forcing variable at the cut-off and, thus, to analyze whether there is any manipulation of the forcing variable. As argued in [35], it is preferable to use histograms rather than smoothed density estimates as histograms can provide a sense in which any jump at the threshold is unusual. Once again, the issues with making these type of histograms differentially private are the same as before: a) usually the bins would be chosen by a mechanisms adaptively and therefore generally one the bins will contain the cut-off point as an interior point thus defeating the purpose of this visualization; b) if a researcher asks for

separate histograms to the right and to the left of the cut-off, this will result in either having global sensitivity that is large and does not diminish to 0 with the sample size (similar to the issues we had with the nonparametric regression at the boundary estimator) or in the bins that are too wide. In either case this leads to a significant amount of information about the distribution of the forcing variable.

4.6 Monte Carlo illustrations

In this section, we want to illustrate our findings of the generally poor performance of the smooth differentially private RDD estimators. We consider the sharp design and illustrate paths of the differentially private local linear estimator with a triangular kernel for increasing sample sizes with different degrees of the privacy protection. These paths are constructed for increasing samples from the size of 300 till the size of 4000. For visual simplicity, we give paths for 20 independent realizations of datasets.

Scenario 1.

The forcing variable X has a uniform distribution on $[-1, 1]$. The regression function is a fifth-order polynomial, with separate coefficients for $X_i < 0$ and $X_i > 0$:

$$m(x) = \begin{cases} 0.35 + 1.27x + 7.18x^2 + 20.21x^3 + 21.54x^4 + 7.33x^5, & \text{if } x < 0, \\ 0.65 + 0.84x - 3x^2 + 7.99x^3 - 9.01x^4 + 3.56x^5, & \text{if } x \geq 0, \end{cases}$$

and the error u having a symmetric uniform distribution on $[-0.12952 \cdot \sqrt{3}, 0.12952 \cdot \sqrt{3}]$. The bandwidth in the local linear estimation is chosen using the approach in [24].

Differentially private estimators are obtained by using the Laplace mechanism, which, as discussed above, draws a mechanism nose from the Laplace distribution with mean zero and variance depending on the global sensitivity of the local linear estimator a noise to the estimator.

In Panel 1 in Figure 1 we show the paths of the estimator in the absence of the mechanism noise when the mechanism noise variance equal to 0.002 for any sample size (for a conservative lower bound of $4 \cdot 0.12952 \cdot \sqrt{3}$ on the global sensitivity of the estimator[12], this would correspond to ϵ_N being 10 times of this $\delta_N = 0$). Panel 2 in Figure 1 depicts the paths of the estimator when the mechanism noise variance equal to 0.002 for any sample size (for a conservative lower bound of $4 \cdot 0.12952 \cdot \sqrt{3}$ on the global sensitivity of the estimator, this would correspond to ϵ_N being equal to 10 times of this and $\delta_N = 0$). Panel 3 in Figure 1 shows the paths of the estimator when the mechanism noise variance equal to 2 for any sample size (for a conservative lower bound of $4 \cdot 0.12952 \cdot \sqrt{3}$ on the global sensitivity of the estimator, this would correspond to ϵ_N being equal to this and $\delta_N = 0$). Finally, Panel 4 in Figure 1 illustrates the paths of the estimator when the mechanism noise variance equal to 200 for any sample size (for a conservative lower bound of $4 \cdot 0.12952 \cdot \sqrt{3}$ on the global

[12] See our discussion in the proof of Proposition 5 in the Appendix.

Setting	Var(mech noise =0)		Var(mech noise =0.002)		Var(mech noise =2)		Var(mech noise =200)	
	5%	1%	5%	1%	5%	1%	5%	1%
$N = 500$	1	1	0.6846	0.3706	0.0666	0.0286	0.056	0.023
$N = 2000$	1	1	0.7252	0.3880	0.0664	0.0290	0.0666	0.0272
$N = 5000$	1	1	0.7260	0.3844	0.0694	0.0284	0.0594	0.0250

Table 1: Rejection rates in 5000 simulations of the false null hypothesis $H_0 : \tau = 0$ in Scenario 1. N denotes the number of observations.

sensitivity of the estimator, this would correspond to ϵ_N being equal to the one-tenth of this and $\delta_N = 0$). Note the different range of the values on the vertical axis in these panels.

In Table 1 we focus on the rejection of the null $H_0 : \tau = 0$ against $H_1 : \tau \neq 0$ when a researcher uses differentially private estimates and their standard errors (note that this is different from our discussion of the differentially private release of t-tests in Section 4.4).

Scenario 2. The only difference here from Scenario 1 is that u is normally distributed with mean zero and variance 0.12952. The support of the outcome variable is unbounded and the global sensitivity is clearly infinite (see the proof of Proposition 5 in the Appendix) which means that regardless of ϵ_N, $\epsilon_N \leq \bar{\epsilon}$, in the Laplace differentially private mechanism the noise has to be drawn from the distribution with an infinite variance. Figure 2 shows that paths of differentially private local linear estimators for when the variance is equal to 10^6.

Here we could have conducted similar power analysis based on a large number of simulations, like in Scenario 1, and we would have obtained that power of the test $H_0 : \tau = 0$ vs $H_1 : \tau \neq 0$ based on differentially private estimates is very low.

5 Conclusion

Differential privacy is a powerful data security concept that precludes a potential adversary from linking sensitive data with outside information, inferring data attributes or determining if particular individual is included in the dataset. The implementation of the differentially private data analysis is based on consideration of randomized estimator where independent randomness is the key instrument that provides the differential privacy guarantee.

In this paper we focused on identification of Econometric models under differential privacy. We concluded that even with relatively simple models identification in this context requires the concepts and methods from the random set theory. We consider identification from the perspective of the

limit of statistical experiments where differentially private implementation of the estimator is applied to the datasets of an increasing size. Identification in this case is the property of the set of weak limits of such estimators. Under our mild regularity conditions this limiting set is a convex compact random set and, thus, it needs to be characterized in probabilistic terms, for instance, using the containment functional.

We apply our theory to the analysos of the regression discontinuity design (RDD) where we consider both sharp and fuzzy design. We show that for both models the random set of weak limits of differentially private estimators contains non-degenerate random elements which precludes point identification of the parameters of interest. We illustrate this finding in a series of Monte Carlo simulations.

Our result, in part, is driven by the structure of the estimators which have to rely on local properties of the underlying distribution. This may indicate that under differential privacy a similar behavior is to be expected for other Econometric models that rely on nuisance parameters.

6 Appendix

6.1 Proofs of theorems and propositions in section 2

Proof of Lemma 1. Let ω_ν be the element of the σ-algebra $\mathcal{F}_\nu$ associated with random element ν_N and ω_S be the element of the σ-algebra of the subsets of $\mathcal{Z}^n$. Since $\mathbb{T}_N$ is the closure of the set of measurable selections that form a closed and bounded space, then for almost all elements (ω_N, ω_S) the set of values $\theta(\mathbb{P}_N(\cdot;\omega_S), \nu_N(\omega_\nu))$ is closed and bounded and, therefore, compact. Thus, random set $\mathbb{T}_N$ is compact.

To prove convexity, it is enough to consider $\theta(\mathbb{P}_N, \nu_N)$ and $\theta'(\mathbb{P}_N, \nu_N)$ that are realizations of two regular (ϵ_N, δ_N)-differential private and (ϵ'_N, δ'_N)-differential private estimators, respectively, where sequences of (ϵ_N, δ_N) and (ϵ'_N, δ'_N) are in $\mathcal{E}$. Then by union bound their convex combination satisfies (2.2) with the right-hand side bound of at most $2\bar{R}(n, \kappa)$. Also, any convex combination $\tau\theta(S_N, \nu_N) + (1-\tau)\theta'(S_N, \nu_N)$ is a realization of the estimator $\tau\theta(\cdot,\cdot)+(1-\tau)\theta'(\cdot,\cdot)$. This estimator is differentially private for the sequence of $(\max\{\epsilon_N, \epsilon'_N\}, \max\{\delta_N, \delta'_N\})$ which belongs to $\mathcal{E}$ by our assumption of $\mathcal{E}$ being a join-semilattice. Finally not that the estimator $\tau\theta(\cdot,\cdot) + (1-\tau)\theta'(\cdot,\cdot)$ has a weak limit from the continuous mapping theorem as it is straightforward to show that $(\theta, \theta')^T(\cdot,\cdot)$ has a joint weak limit (of course, we would use the fact that θ, θ' do not depend on N). Thus, set $\mathbb{T}_N$ is convex random set. ■

Proof of Lemma 2: Assume, contrary to the statement of the Lemma that $\Delta_N = \theta(S_N, \nu_N) - \tau \xrightarrow{p} 0$. Then $\theta(S_N, \nu_N) = \tau + \Delta_N$, and because τ is not constant, then conditional on S_N and S_{N+1}, estimator $\theta(S_N, \nu_N)$ and $\theta(S_{N+1}, \nu_{N+1})$ cannot be independent. This, in its turn, will contradict the independence of elements ν_N and ν_{N+1}, which is a fundamental requirement for differential privacy. ■

Proof of Theorem 1. Provided that $\theta(\cdot,\cdot)$ belong to a compact subset of a separable space in $\mathbf{L}_1$, for each χ there exists K such that for a finite set of elements $\{\theta^{(1)}(\cdot,\cdot), \ldots, \theta^{(K)}(\cdot,\cdot)\}$, their convex hull Θ_N^K is within χ-Hausdorf distance from set $\mathbb{T}_{N,\mathcal{E}}$. Since each $\theta^{(k)}(\cdot,\cdot)$ is a function of the same elements (ω_N, ω_S), weak convergence of $\theta^{(k)}(\mathbb{P}_N, \nu_N)$ implies joint weak convergence of the set $\{\theta^{(1)}(\cdot,\cdot), \ldots, \theta^{(K)}(\cdot,\cdot)\}$ and, therefore, weak convergence of their convex hull. We then choose sequence χ_N, which induces sequence K_N such that $\sup_A \left| C_{\Theta^K_{K_N}}(A) - C_{\mathbb{T}_{N,\mathcal{E}}}(A) \right|$ is a decreasing function of N. Then by Theorem 6.26 in [44] sequence of random sets $\mathbb{T}_{N,\mathcal{E}}$ converges weakly. ■

Proof of Theorem 3. a) Suppose that for any sequence of (ϵ_N, δ_N) from $\mathcal{E}$ it holds that any regular (ϵ_N, δ_N)-differentially private estimator $\theta(\mathbb{P}_N, \nu_N)$ is such that $\theta(\mathbb{P}_N, \nu_N) \xrightarrow{p} \theta_0$. Then $\mathbf{T}_\mathcal{E} = \{\theta_0\}$ (degenerate distribution at θ_0). Then, clearly, for any convex polytope $K \ni \theta_0$ we have $C_{\mathbf{T}_\mathcal{E}}(K) = 1 \geq 1 - \alpha$ for any $\alpha \in (0, 1)$.
b) Suppose for any $\alpha \in (0, 1)$ and any convex polytope $K \ni \theta_0$

$$C_{\mathbf{T}_\mathcal{E}}(K) \geq 1 - \alpha.$$

Since α can be taken to be arbitrarily close to 0, this means that $C_{\mathbf{T}_\mathcal{E}}(K) = 1$. Since convex polytope $K \ni \theta_0$ can be taken to have arbitrarily small volume, this means that $\mathbf{T}_\mathcal{E} = \{\theta_0\}$ (degenerate distribution at θ_0). Indeed, take a decreasing sequence $K_m \ni \theta_0$ of convex polytopes such that $\cap_{m=1}^{\infty} K_m = \{\theta_0\}$. By the continuity theorem for monotone sequences of events $C_{\mathbf{T}_\mathcal{E}}(\cap_{m=1}^{\infty} K_m) = \lim_{m\to\infty} C_{\mathbf{T}_\mathcal{E}}(K_m) = 1$, which immediately implies that $\mathbf{T}_\mathcal{E} = \{\theta_0\}$ meaning that every $\theta(S_N, \eta_N)$ converges weakly to θ_0, and thus, $\theta(S_N, \eta_N) \xrightarrow{p} \theta_0$. ■

6.1.1 On the selection expectation of the set of regular differentially private estimators for θ_0 in Section ??

Here we briefly discuss the properties of the selection expectation of the limiting random set $\mathbf{T}_\mathcal{E}$ of regular differentially private estimators for θ_0 in the context of the example in Section **??**.

Let $\Theta_{N,\mathcal{E}}$ denote the selection expectation of the random set $\mathbb{T}_{N,\mathcal{E}}$ of all regular (ϵ_N, δ_N)-differentially private estimators for θ_0 with sequences (ϵ_N, δ_N) from $\mathcal{E}$. The Hausdorff limit of $\Theta_{N,\mathcal{E}}$ as $N \to \infty$ is denoted as $\Theta_{\infty,\mathcal{E}}$.

$\Theta_{\infty,\mathcal{E}}$, loosely speaking, contains all limits of expectations of regular differentially private estimators. Ideally, if differentially private estimators are compatible with consistency, for a broad range of sequences ϵ_N and δ_N converging to zero the set $\Theta_{\infty,\mathcal{E}}$ should be a singleton $\{\theta_0\}$.

Theorem 1.45 in [44] links $\Theta_{\infty,\mathcal{E}}$ to the selection expectation of the limit random set $\mathbf{T}_\mathcal{E}$. Indeed, that theorem immediately implies that under conditions of Theorem 1, $\mathbb{E}\mathbb{T}_{N,\mathcal{E}}$ converges to $\mathbb{E}\mathbf{T}_\mathcal{E}$ in the Hausdorff metric and the Lebesgue measure of $\mathbb{E}\mathbb{T}_{N,\mathcal{E}}$ converges to the Lebesgue measure of $\mathbb{E}\mathbf{T}_\mathcal{E}$ as $N \to \infty$. In other words, weak convergence of a sequence of random sets implies the convergence of the selection expectation.

Going back to the discussion of our example in Section **??**, we note that even though in that example differential privacy-inducing mechanisms perturb the estimator with random noise symmetric at zero, there is no guarantee that the limiting $\Theta_{\infty,\mathcal{E}}$ is a singleton at θ_0. In If we collect differentially private estimators across all three regimes in that example, we find that the corresponding limiting set of selection expectations will include

$$\mathbb{E}\, U\left(\{\mathrm{Argmin}_{\theta\in\Theta}\,\theta, \mathrm{Argmax}_{\theta\in\Theta}\,\theta\}\right),\ \{\mathbb{E}\,\Lambda(c),\, c \in [0,+\infty)\},\ 0,\ \mathbb{E}[X].$$

We note that in this case the target expectation $\mathbb{E}[X]$ belongs to $\Theta_{\infty,\mathcal{E}}$. At the same time, the set $\Theta_{\infty,\mathcal{E}}$ itself is clearly large.

We also note that if we exclude the elements of the selection expectation that result from Regime 3 where the scale of double exponential noise asymptotically increases, the selection expectation of our considered family of estimator will be a linear segment in Θ that connects points 0 and $\mathbb{E}[X]$ since the set $\{\mathbb{E}\,\Lambda(c),\, c \in [0,+\infty)\}$ is a line in Θ that connects 0 and $\mathbb{E}[X]$.

6.2 Proofs of theorems and propositions in section 4

Proof of Proposition 1. Treat $G(N)$ as Δ_N in AC2 and choose respective σ_N^2 and $[a_N, b_N]$ as in condition AC2. Take the definition of the differentially private estimator and note that for two estimators $\psi(S_N)$ and $\psi(S'_N)$ based on two datasets that differ in one observation only, we have

$$\begin{aligned}
P_{\nu_N \sim f_{N;\sigma_N^2}}\left(\nu_N + \psi(S_N) \in B\right) = P_{\nu_N \sim f_{N;\sigma_N^2}}\left(\nu_N + \psi(S_N) \in B | \nu_N \in [a_N, b_N]\right) + \\
+ P_{\nu_N \sim f_{N;\sigma_N^2}}\left(\nu_N + \psi(S_N) \in B | \nu_N \notin [a_N, b_N]\right) \leq P_{\nu_N \sim f_{N;\sigma_N^2}}\left(\nu_N + \psi(S_N) \in B | \nu_N \in [a_N, b_N]\right) \\
+ P_{\nu_N \sim f_{N;\sigma_N^2}}\left(\nu_N + \psi(S_N) \in B | \nu_N \notin [a_N, b_N]\right) \\
\leq e^{\mathcal{D}(G(N),\sigma_N^2,a_N,b_N)} P_{\nu_N \sim f_{N;\sigma_N^2}}\left(\nu_N + \psi(S'_N) \in B | \nu_N \in [a_N, b_N]\right) \\
+ P_{\nu_N \sim f_{N;\sigma_N^2}}\left(\nu_N + \psi(S_N) \in B | \nu_N \notin [a_N, b_N]\right),
\end{aligned}$$

where B is any subset of $\mathbb{R}$. Thus, if ϵ_N is greater os equal $e^{\mathcal{D}(G(N),\sigma_N^2,a_N,b_N)}$ (which also gives acceptable rates of convergence for ϵ_N to 0) and $P_{\nu_N \sim f_{N;\sigma_N^2}}\left(\nu_N + \psi(S_N) \in B | \nu_N \notin [a_N, b_N]\right) \delta_N$ (which also gives acceptable rates of convergence for δ_N to 0), then (ϵ_N, δ_N)-differential privacy criterion is satisfied.

At the same time, because both the mean and the variance of ξ_N converge to 0, then ξ_N converges in probability to 0 and, therefore, in light of (4.5) the estimator $\widehat{\theta}$ in (4.4) is consistent. ■

Proof of Theorem 6.

As we will show, the inconsistency of the estimator $\widehat{\theta}$ in (4.4) stems from the fact that the variance of the mechanism noise ξ_N does not go to 0 with the sample size, which in its turn is explained by the fact that the global sensitivity of $\psi(S_N)$ does not go to 0 with the sample size.

Indeed, suppose that $G(N)$ is bounded away from zero as $N \to \infty$ and is also bounded from above. Then for a fixed large enough interval $[a_N, b_N]$, the value of $\mathcal{D}(G(N), \sigma_N^2, a_N, b_N)$ has to be bounded from above by ε_N. Since $G(N)$ is bounded away from 0, from AC1 we have that

$$\mathcal{D}(G(N), \sigma_N^2, a, b) \to +\infty$$

if $\sigma_N^2 \to 0$ for a fixed interval $[a, b]$. The definition of $\mathcal{D}()$ as a supremum implies that the same property will hold if instead of the fixed interval $[a, b]$ we take $[a_N, b_N]$ converging to $\mathbb{R}$. This implies, of course, that σ_N^2 has to be bounded away from zero. This, in turn, implies that ξ_N does not converge in probability to zero even if the mean of ξ_N converges to 0. Hence, $\widehat{\theta}$ dose not converge in probability to the true parameter value.

If $G(N) = +\infty$, then to guarantee $(\varepsilon_N, \delta_N)$-differential privacy, one would have to take $\sigma_N^2 = +\infty$, clearly leading to the inconsistency of $\widehat{\theta}$.

Note that this inconsistency result applies even to (ϵ_N, δ_N) not changing with N. It will also be true under stronger requirements of differential privacy when both parameters converge to 0. ■

Lemmas 6 and 7 will help to establish results in Propositions 2-5.

LEMMA 6. *Consider two weighted averages*

$$q_1 = \sum_{i=1}^{T} w_i a_i + w_{T+1} a_{T+1}, \qquad \text{where } w_i = \frac{b_i}{\sum_{i=1}^{T+1} b_i}, \quad i = 1, \ldots, T+1,$$

$$q_2 = \sum_{i=1}^{T} \tilde{w}_i a_i + \tilde{w}_T \tilde{a}_T, \qquad \text{where } \tilde{w}_i = \frac{b_i}{\sum_{i=1}^{T} b_i + \tilde{b}_{T+1}}, \quad i = 1, \ldots, T, \quad \tilde{w}_{T+1} = \frac{\tilde{b}_{T+1}}{\sum_{i=1}^{T} b_i + \tilde{b}_{T+1}},$$

and

$$0 \le c_1 \le (\text{ or } <)\, b_i,\ \tilde{b}_i \le c_2, \tag{6.1}$$

$$d_1 \le a_i \le d_2, \quad i = 1, \ldots, T+1, \qquad \text{and } d_1 < d_2. \tag{6.2}$$

Then

(a) if $c_1 = 0$ *and* $|d_1|, |d_2| < \infty$, *then*

$$\max_{a_1,\ldots,a_T,a_{T+1},\tilde{a}_{T+1},b_1,\ldots,b_T,b_{T+1},\tilde{b}_{T+1} \text{ s.t. } (6.1),(6.2)} |q_1 - q_2| = d_2 - d_1.$$

(b) if $c_1 > 0$ *and* $|d_1|, |d_2| < \infty$, *then*

$$\max_{a_1,\ldots,a_T,a_{T+1},\tilde{a}_{T+1},b_1,\ldots,b_T,b_{T+1},\tilde{b}_{T+1} \text{ s.t. } (6.1),(6.2)} |q_1 - q_2| = \frac{c_2(d_2 - d_1)}{T \cdot c_1 + c_2}.$$

(c) if $d_1 = -\infty$ *or* $d_2 = +\infty$, *then*

$$\max_{a_1,\ldots,a_T,a_{T+1},\tilde{a}_{T+1},b_1,\ldots,b_T,b_{T+1},\tilde{b}_{T+1} \text{ s.t. } (6.1),(6.2)} |q_1 - q_2| = +\infty.$$

In cases (a)-(c), $\max |q_1 - q_2|$ *can be attained by a positive change as well as by a negative change – that is, there are values of* a_t*'s,* b_t*'s and* $\tilde{a}_{T+1}$, $\tilde{b}_{T+1}$ *such that* $q_1 - q_2 = \max |q_1 - q_2|$, *and there are values of* a_t*'s,* b_t*'s and* $\tilde{a}_{T+1}$, $\tilde{b}_{T+1}$ *such as* $q_1 - q_2 = -\max |q_1 - q_2|$.

Proof of Lemma 6.

(a) In this case, we can take

- $b_1 = \ldots = b_T \approx 0$; $b_{T+1} = \tilde{b}_{T+1} = c_2$;
- $a_1, \ldots, a_T$ can be arbitrary values that satisfy (6.2); $a_{T+1} = d_1$, $\tilde{a}_{T+1} = d_2$.

This gives us $q_2 - q_1 = d_2 - d_1$. Therefore, we should have $\max |q_2 - q_1| \ge d_2 - d_1$. At the same time each weighted average q_1 and q_2 has to belong to $[d_1, d_2]$, which is the range for a_i's. Therefore,

necessarily $\max|q_2 - q_1| \le d_2 - d_1$. This implies that $\max|q_2 - q_1| = d_2 - d_1$. Note that if above we take $a_{T+1} = d_2$, $\tilde{a}_{T+1} = d_1$, then $q_2 - q_1 = d_1 - d_2 = -|d_2 - d_1|$.

(b) In this case, to evaluate the largest change in the weighted average we have to consider extreme situations. The first extreme situation is when $q_1 = d_1$ and the $(T+1)$-th component in this average has the largest weight and changes to the the other extreme d_2 in the new average q_2.

This situation can be described as

- $b_1 = \ldots = b_T = c_1$; $b_{T+1} = \tilde{b}_{T+1} = c_2$;
- $a_1, \ldots, a_T = d_1$; $a_{T+1} = d_1$, $\tilde{a}_{T+1} = d_2$.

This will give us $q_2 - q_1 = \frac{c_2(d_2-d_1)}{T\cdot c_1 + c_2} > 0$.

In the second extreme scenario where b_t's and $\tilde{b}_{T+1}$ are the same as above but $q_1 = d_2$ and the $(T+1)$-th component in this average has the largest weight and changes to the the other extreme d_1 in the new average q_2, we obtain that $q_2 - q_1 = -\frac{c_2(d_2-d_1)}{T\cdot c_1 + c_2} < 0$. These two extreme scenarios give us exactly the same $|q_2 - q_1|$.

Thus, $\max|q_2 - q_1| = \frac{c_2(d_2-d_1)}{T\cdot c_1 + c_2}$.

(c) In this case, consider the case when b_i, $i = 1, \ldots, T+1$, and $\tilde{b}_{T+1}$ are any values that satisfy (6.1). Suppose $d_2 = +\infty$. Let a_i, $i = 1, \ldots, T+1$, take any finite values while $\tilde{a}_{T+1}$ is very (arbitrarily) large. This gives $q_1 - q_2 = -\infty$ and, thus, $|q_1 - q_2| = +\infty$. Therefore, in this case $\max|q_1 - q_2| = +\infty$. If, of course, $\tilde{a}_{T+1}$ is taking a finite value while a_T is very (arbitrarily) large, then $q_1 - q_2 = +\infty$.

The case of when d_2 is finite but $d_1 = -\infty$ is analyzed analogously.

■

LEMMA 7. *Consider two weighted averages*

$$q_1 = \sum_{i=1}^{T} w_i a_i + w_{T+1} a_{T+1}, \qquad \textit{where } w_i = \frac{b_i}{\sum_{j=1}^{T+1} b_j}, \quad i = 1, \ldots, T+1,$$

$$q_2 = \sum_{i=1}^{T} \tilde{w}_i a_i, \qquad \textit{where } \tilde{w}_i = \frac{b_i}{\sum_{j=1}^{T} b_j}, \quad i = 1, \ldots, T,$$

where b_i and a_i satisfy conditions (6.1) and (6.2), respectively. Then

(a) if $c_1 = 0$ and $|d_1|, |d_2| < \infty$, then

$$\max_{a_1,\ldots,a_T,a_{T+1},\tilde{a}_{T+1},b_1,\ldots,b_T,b_{T+1} \ s.t. \ (6.1),(6.2)} |q_1 - q_2| = d_2 - d_1.$$

(b) if $c_1 > 0$ and $|d_1|, |d_2| < \infty$, then

$$\max_{a_1,\ldots,a_T,a_{T+1},\tilde{a}_{T+1},b_1,\ldots,b_T,b_{T+1} \ s.t. \ (6.1),(6.2)} |q_1 - q_2| = \frac{c_2(d_2 - d_1)}{T\cdot c_1 + c_2}.$$

(c) if $d_1 = -\infty$ *or* $d_2 = +\infty$*, then*

$$\max_{a_1,\ldots,a_T,a_{T+1},\tilde{a}_{T+1},b_1,\ldots,b_T,b_{T+1} \ s.t. \ (6.1),(6.2)} |q_1 - q_2| = +\infty.$$

In cases (a)-(c), $\max|q_1 - q_2|$ *can be attained by a positive change as well as by a negative change – that is, there are values of* a_t*'s,* b_t*'s such that* $q_1 - q_2 = \max|q_1 - q_2|$*, and there are values of* a_t*'s,* b_t*'s such that* $q_1 - q_2 = -\max|q_1 - q_2|$*.*

Proof of Lemma 7.

(a) In this case, we can take

- $b_1 = \ldots = b_T \approx 0$ (arbitrarily close values to 0); $b_{T+1} = c_2$;
- $a_1 = \ldots = a_T = d_1$; $a_{T+1} = d_2$.

This will give $q_1 \approx d_1$, $q_2 \approx d_2$ and, thus, $q_2 - q_1 \approx d_2 - d_1 > 0$.

Therefore, we should have $\max|q_2 - q_1| \geq d_2 - d_1$. At the same time each weighted average q_1 and q_2 has to belong to $[d_1, d_2]$, which is the range for a_i's. Therefore, necessarily $\max|q_2 - q_1| \leq d_2 - d_1$. This implies that $\max|q_2 - q_1| = d_2 - d_1$.

(b) In this case, to evaluate the largest change in the weighted average we have to consider extreme situations. An extreme situation is when in q_1 the $(T+1)$-th component (which is later dropped in when defining q_2) has the largest weight and the value that is maximally different from the values of the first T components.

The first extreme situation can be described as

- $b_1 = \ldots = b_T = c_1$; $b_{T+1} = c_2$;
- $a_1, \ldots, a_T = d_1$; $a_{T+1} = d_2$.

This gives us $|q_2 - q_1| = \frac{c_2(d_2 - d_1)}{T \cdot c_1 + c_2}$.

The second extreme situation, where b_t's are the same as above but $a_1, \ldots, a_T = d_2$ and $a_{T+1} = d_1$, gives us the exactly same value of $|q_2 - q_1|$. Thus, $\max|q_2 - q_1| = \frac{c_2(d_2 - d_1)}{T \cdot c_1 + c_2}$.

(c) In this case, consider the case when b_i, $i = 1, \ldots, T+1$, and $\tilde{b}_{T+1}$ are any values that satisfy (6.1). Also, let a_i, $i = 1, \ldots, T+1$, take any finite values while $\tilde{a}_{T+1}$ is very large (arbitrarily large) in the absolute value. This gives $|q_1 - q_2| = +\infty$. therefore, in this case $\max|q_1 - q_2| = +\infty$. ■

Proof of Proposition 2.

(a) The global sensitivity of the estimator is calculated by comparing the results of the estimation for two datasets that differ only in on data point. In order to calculate the global sensitivity, we need to keep in mind the following things:

(i) the new data point can enter a K-h-neighborhood of c (thus, the old data point was outside of both K-h-neighborhoods of c)

(ii) the new data point can fall outside of both K-h-neighborhoods of c (thus, the old data point was inside one of K-h-neighborhoods of c)

(iii) the new data point remains in the same neighborhood

(iv) the new data point can switch neighborhoods

In order to find global sensitivity, it is enough for us to find maximum absolute changes in the estimate in these four situations and then takes their maximum. Let us consider these four different situation listed above. In this proof we use Lemmas 6 and 7 with $c_2 = \bar{K}$, where $\bar{K}$ denotes the maximum value the kernel $K(\cdot)$, and $c_1 = 0$ since $K(\cdot)$ is continuous and, therefore, $\underline{K} = 0$.

(i) Suppose the new data point enters the K-h-neighborhood to the left of c while the old data point was outside of both K-h-neighborhoods of c. Then by part (a) of Lemma 7, the maximum absolute change G_{0L} in the estimate in this case is $G_{0L} = \overline{Y}^l - \underline{Y}^l$. Analogously we can consider the case when a new data point enters the K-h-neighborhood to the right of c. Then the maximum absolute change G_{0R} in the estimate in this case is $G_{0R} = \overline{Y}^r - \underline{Y}^r$.

(ii) In this case we have two situations – in one situation the old data point was in the left K-h-neighborhood and in the other situation the old data point was in the right K-h-neighborhood. In both situations the new data point falls outside of both neighborhoods. In the former case the maximum absolute change in the estimate coincides with G_{0L}, and in the latter case the maximum absolute change in the estimate coincides with G_{0R}.

(iii) When the observation remains in the left K-h-neighborhood, we apply part (a) of Lemma 6 to obtain that the maximum absolute change G_{LL} in the estimate in this case is $G_{LL} = \overline{Y}^l - \underline{Y}^l$.

When the observation remains in the right K-h-neighborhood, we consider the maximum absolute change G_{RR} in estimate and analogously to above show that $G_{RR} = \overline{Y}^r - \underline{Y}^r$.

(iv) Suppose an observation moves from the left K-h-neighborhood to the the right K-h-neighborhood. Our estimator of interest is the *difference* between the weighted means in the right and the left K-h neighborhoods of c. Therefore, the move of the observation from one neighborhood to the other affects both parts of the estimator.

As we know from part (a) of Lemma 6, the maximum absolute change in the weighted average for the right-hand side is $G_{0R} = \overline{Y}^r - \underline{Y}^r$ and that this degree of change can be attained as a positive change (increase). Similarly, the maximum absolute change in the weighted average for the left-hand side is $G_{LO} = \overline{Y}^l - \underline{Y}^l$ and that this degree of change can be attained as a negative change (decrease). In order to obtain the maximum absolute changes for the difference in weighted means we have to look at the cases when these two weighted means change in opposite directions, which leads to the maximum change being $G_{LR} = \overline{Y}^r - \underline{Y}^r + \overline{Y}^l - \underline{Y}^l$.

Analogously, we can consider an observation moves from the right K-h-neighborhood to the the left K-h-neighborhood and show that in this case the maximum absolute change is $G_{RL} = \overline{Y}^r - \underline{Y}^r + \overline{Y}^l - \underline{Y}^l$.

To sum up the results of part (a), the global sensitivity is

$$G(N) = \max\{G_{L0}, G_{R0}, G_{0L}, G_{0R}, G_{LR}, G_{RL}, G_{LL}, G_{RR}\} = \overline{Y}^r - \underline{Y}^r + \overline{Y}^l - \underline{Y}^l.$$

(b) Suppose for instance that the support $\mathcal{Y}^r$ is unbounded. Then part (c) of Lemmas 6 and 7 will immediately give us that for $G_{R0}, G_{0R}, G_{RL}, G_{RR}$ defined above,

$$G_{R0} = G_{0R} = G_{RL} = G_{RR} = +\infty,$$

which implies this part of the proposition. Other cases in this part of the proposition lead the same conclusion.

■

Proof of Proposition 3.

Just like in Proposition 2, the global sensitivity is

$$G(N) = \max\{G_{L0}, G_{R0}, G_{0L}, G_{0R}, G_{LR}, G_{RL}, G_{LL}, G_{RR}\},$$

where $G_{L0}, G_{R0}, G_{0L}, G_{0R}, G_{LL}, G_{RR}, G_{LR}, G_{RL}$ are defined as in the proof of Proposition 2.

Once again, we will rely on the results in Lemmas 6 and 7 but this time in part (b) in both lemmas as we will take $c_1 = \underline{K}$ and $c_2 = \overline{K}$, where $\overline{K}$ is the maximum value of kernel K.

(a) Applying results of part (b) of Lemma 7 and noting that the minimum number of observations in the left and the right K-h-neighborhoods of c is $m^l(N)$ and $m^r(N)$ respectively, we obtain that

$$G_{L0} = G_{0L} = \frac{\overline{K}(\overline{Y}^l - \underline{Y}^l)}{m^l(N) \cdot \underline{K} + \overline{K}},$$
$$G_{R0} = G_{0R} = \frac{\overline{K}(\overline{Y}^r - \underline{Y}^r)}{m^r(N) \cdot \underline{K} + \overline{K}}.$$

Applying results of part (b) of Lemma 6, we have

$$G_{LL} = \frac{\overline{K}(\overline{Y}^l - \underline{Y}^l)}{(m^l(N) - 1) \cdot \underline{K} + \overline{K}}, \quad G_{RR} = \frac{\overline{K}(\overline{Y}^r - \underline{Y}^r)}{(m^r(N) - 1) \cdot \underline{K} + \overline{K}}.$$

We next consider G_{LR} which quantifies the case when an observation from the left K-h neighborhood of c moved in to the right-hand side neighborhood. Suppose we started with $T + 1$ observations in the left K-h-neighborhood, where $m^l(N) \leq T \leq N - m^r(N)$. We need to evaluate the biggest

change that happened in the left-hand side neighborhood, the biggest change in the right-hand side neighborhood and evaluate their directions (whether these changes are acting in the same or opposite directions). Relying on the results of part (b) of Lemma 7 we can establish that given T the largest absolute change in the weighted mean in the left K-h neighborhood of c is $\frac{\overline{K}(\overline{Y}^l - \underline{Y}^l)}{T \cdot \underline{K} + \overline{K}}$, and the largest absolute change in the weighted mean from acquiring an extra point in that neighborhood is $\frac{\overline{K}(\overline{Y}^r - \underline{Y}^r)}{(N-T-1) \cdot \underline{K} + \overline{K}}$.

As shown in Lemma 7, these changes can be either positive or negative. Since our estimator of interest is the *difference* between the weighted means in the right and the left K-h neighborhoods of c, to get maximum absolute changes for a given T we have to look at the cases when these two weighted means change in opposite directions. For a given T this gives us the maximum absolute change

$$\frac{\overline{K} \cdot (\overline{Y}^l - \underline{Y}^l)}{T \cdot \underline{K} + \overline{K}} + \frac{\overline{K} \cdot (\overline{Y}^r - \underline{Y}^r)}{(N - T - 1) \cdot \underline{K} + \overline{K}}.$$

Now we just need to find the maximum of this over T such that $m^l(N) \leq T \leq N - m^r(N)$. If $\overline{Y}^r - \underline{Y}^r > \overline{Y}^l - \underline{Y}^l$, then the maximum is attained at $T = N - m^r(N)$, otherwise it is attained at $T = m^l(N)$. To summarize,

$$G_{LR} = \max \left\{ \frac{\overline{K} \cdot (\overline{Y}^l - \underline{Y}^l)}{m^l(N) \cdot \underline{K} + \overline{K}} + \frac{\overline{K} \cdot (\overline{Y}^r - \underline{Y}^r)}{(N - m^l(N) - 1) \cdot \underline{K} + \overline{K}}, \ \frac{\overline{K} \cdot (\overline{Y}^l - \underline{Y}^l)}{(N - m^r(N)) \cdot \underline{K} + \overline{K}} + \frac{\overline{K} \cdot (\overline{Y}^r - \underline{Y}^r)}{(m^r(N) - 1) \cdot \underline{K} + \overline{K}} \right\}.$$

The case of G_{RL} which quantifies the case when an observation from the right K-h neighborhood of c moved in to the left-hand side neighborhood is considered analogously. In this case,

$$G_{RL} = \max \left\{ \frac{\overline{K} \cdot (\overline{Y}^l - \underline{Y}^l)}{(N - m^r(N) - 1) \cdot \underline{K} + \overline{K}} + \frac{\overline{K} \cdot (\overline{Y}^r - \underline{Y}^r)}{m^r(N) \cdot \underline{K} + \overline{K}}, \ \frac{\overline{K} \cdot (\overline{Y}^l - \underline{Y}^l)}{(m^l(N) - 1) \cdot \underline{K} + \overline{K}} + \frac{\overline{K} \cdot (\overline{Y}^r - \underline{Y}^r)}{(N - m^l(N)) \cdot \underline{K} + \overline{K}} \right\}.$$

This gives the result that $G(N)$ is of the rate $\frac{1}{\min\{m^l(N), m^r(N)\}}$.

(b) Suppose e.g. that the support $\mathcal{Y}^r$ is unbounded. Then part (c) of Lemmas 6 and 7 give us that

$$G_{0R} = G_{R0} = G_{RR} = +\infty,$$

which implies that $G(N) = +\infty$.

■

Proof of Proposition 4. The proof in this case is analogous to the proof of Proposition 2. Since the kernel has an unbounded support, there are no longer case of observations falling outside of either neighborhood or entering a neighborhood. Therefore, the global sensitivity is

$$G(N) = \max\{G_{LR}, G_{RL}, G_{LL}, G_{RR}\},$$

where $G_{LL}, G_{RR}, G_{LR}, G_{RL}$ are defined as in the proof of Proposition 2. Throughout the proof we apply Lemmas 6 and 7 with the strict inequality version ($0 = c_1 < b_i$) in (6.1).

(a) When the observation remains to the left of c, we apply part (a) of Lemma 6 to obtain that the maximum absolute change G_{LL} in the estimate in this case is $G_{LL} = \overline{Y}^l_{all} - \underline{Y}^l_{all}$.

When the observation remains to the right of c, we consider the maximum absolute change G_{RR} in estimate and analogously to above show that $G_{RR} = \overline{Y}^r_{all} - \underline{Y}^r_{all}$.

Suppose an observation moves from the left of c to the right of c, or from the right of c to the left of c. Analogously to the proof of of Proposition 2, we can establish that

$$G_{LR} = G_{RL} = \overline{Y}^r_{all} - \underline{Y}^r_{all} + \overline{Y}^l_{all} - \underline{Y}^l_{all}.$$

(b) Analogous to the proof in Proposition 2.

■

Proof of Proposition 5.

Just like in Propositions 2 and 3, we want to find

$$G(N) = \max\{G_{L0}, G_{R0}, G_{0L}, G_{0R}, G_{LR}, G_{RL}, G_{LL}, G_{RR}\}$$

where G_{L0} and G_{R0} are sensitivities in situations of a new observation leaving the left or the right h-neighborhood, respectively; and G_{0L} and G_{0R} are sensitivities in situations of a new observation entering the left or the right h-neighborhood, respectively; G_{LR} and G_{RL} are sensitivities in cases of an observation switching the neighborhoods; G_{LL} and G_{RR} are sensitivities in cases when an observation changes within the same neighborhood.

Since the local linear estimator effectively considers observations whose running variable values are in a small neighborhood around c, we employ (4.11) as approximations of the support for the outcome in one-sided neighborhoods of c.

As we know,

$$\widehat{\alpha}_R = \overline{y}_R - (\overline{x}_R - c)\frac{\sum_{i=1}^N (q_i x_i - \overline{x}_R) q_i y_i \cdot 1(c \le x_i)}{\sum_{i=1}^N (x_i q_i - \overline{x}_R)^2 \cdot 1(c \le x_i)}$$
$$\widehat{\alpha}_L = \overline{y}_L - (\overline{x}_L - c)\frac{\sum_{i=1}^N (x_i q_i - \overline{x}_L) y_i q_i \cdot 1(x_i < c)}{\sum_{i=1}^N (x_i q_i - \overline{x}_R)^2 \cdot 1(x_i < c)},$$

where $q_i = K\left(\frac{x_i - c}{h_N}\right)$, $\overline{y}_R = \frac{\sum_{i=1}^N q_i y_i 1(c \le x_i)}{\sum_{i=1}^N q_i 1(c \le x_i)}$, $\overline{x}_R = \frac{\sum_{i=1}^N q_i x_i 1(c \le x_i)}{\sum_{i=1}^N q_i 1(c \le x_i)}$, $\overline{y}_L = \frac{\sum_{i=1}^N q_i y_i 1(x_i < c)}{\sum_{i=1}^N q_i 1(x_i < c)}$, $\overline{x}_L = \frac{\sum_{i=1}^N q_i x_i 1(x_i < c)}{\sum_{i=1}^N q_i 1(x_i < c)}$.

If at least one of $\overline{Y}^l$, $\underline{Y}^l$, $\overline{Y}^r$ or $\underline{Y}^r$ is unbounded, then it is obvious that the global sensitivity is infinite as in this case we can consider a change in the value of the outcome variable for one

observation and we can imagine that it changes from some finite number to the one arbitrary large one in the absolute value.

However, we want to show that $G(N) = \infty$ even if the support of the outcome is bounded. Let us show that $G_{RR} = \infty$. Consider a situation when the first $T \leq N$ observations in our data are in the right-hand side neighborhood. Consider, for example, realizations of the datasets when only T-th observation in the right-hand side neighborhood changes its value x_T. Suppose we have the following realized data in the right-hand side neighborhood:

$$x_i = c + \Delta_N, \quad i = 1, \ldots, m^r(N) - 1 \tag{6.3}$$

$$x_T = c + u_0 h_N - \Delta_N, \quad x'_T = c + \Delta_N - \Delta_N \delta, \tag{6.4}$$

for some $0 < \Delta_N << u_0 h_N$ and $0 < \delta < 1$. For a kernel $K(\cdot)$ with a bounded support $u_0 > 0$ is the value such that $(-u_0, u_0)$ is the support of this kernel. If $K(\cdot)$ has an unbounded support, then we can take u_0 to be a very large positive number. Whatever the situation is, we can take

$$q_i \approx K(0), \quad i = 1, \ldots, m^r(N) - 1$$
$$q_T \approx \underline{K}, \quad q'_T = K(0),$$

where $\underline{K} = \inf_{u \in (-u_0, u_0)} K(u)$. Suppose that $y_T = y'_T$. Then

$$\widehat{\alpha}_R \approx \overline{y}_R - \left(\frac{(c + \Delta_N)(T-1)K(0)}{(T-1)K(0) + \underline{K}} - c \right) \times$$
$$\times \frac{K(0)((c + \Delta_N)K(0) - \frac{(c+\Delta_N)(T-1)K(0)}{(T-1)K(0)+\underline{K}}) \sum_{i=1}^{T-1} y_i + \underline{K} y_T \left((c + u_0 h_N - \Delta_N)\underline{K} - \frac{(c+\Delta_N)(T-1)K(0)}{(T-1)K(0)+\underline{K}} \right)}{((c + \Delta_N)K(0) - \frac{(c+\Delta_N)(T-1)K(0)}{(T-1)K(0)+\underline{K}})^2 + \left((c + u_0 h_N - \Delta_N)\underline{K} - \frac{(c+\Delta_N)(T-1)K(0)}{(T-1)K(0)+\underline{K}} \right)^2}$$

$$\widehat{\alpha}_R \approx \overline{y}_R - \Delta_N \left(1 - \frac{\delta}{T} \right) \times \frac{-\frac{\Delta_N \delta}{T} \sum_{i=1}^{T-1} y_i - y_T \frac{T-1}{T} \Delta_N \delta}{\frac{(T-1)\Delta_N^2 \delta^2}{T^2} + (\frac{T-1}{T} \Delta_N \delta)^2}$$

For fixed T, h_N, Δ_N, it is possible to have $\delta \downarrow 0$, in which case we have that

$$|\widehat{\alpha}'_R - \widehat{\alpha}_R| \to \infty.$$

Since there are no changes in $\widehat{\alpha}_L$, we conclude that $G_{RR} = \infty$, and thus, $G = \infty$.

Note that when the kernel either has an unbounded support or has a bounded support with $\underline{K} = 0$, then even without using $\delta \downarrow 0$, we can establish that the global sensitivity is bounded away from zero for any N, using techniques similar to those in Propositions 2 and 4. The proof above is based on the ability to have realizations of the data such that the minimum eigenvalue of the matrix $\frac{1}{T} \tilde{X}_r^T \tilde{X}_r$ can be arbitrarily close to zero, where $\tilde{X}_r$ is the $T \times 2$ matrix of regressors $(1, x_i - c)$ for when $x_i \geq c$. If in the implementation of a differentially private a data curator wants to establish a strictly positive lower bound on the minimum eigenvalue of this matrix, then in the case of a kernel with a bounded support and $\underline{K} > 0$ (like the uniform kernel), the global sensitivity in this case can be shown to

have the rate $\frac{1}{\min\{m^r(N),m^l(N)\}}$. However, an issue with this is given in discussion after Proposition 3 and is related to the fact that there is a always a strictly positive probability of the number of observations being strictly less $m^r(N)$ in the K-h-neighborhood to the right or strictly less than $m^l(N)$ in the K-h-neighborhood to the left.

It is obvious that when the support of $Y|X$ in the neighborhood to the right of c or to the left of c is unbounded, then $G(N) = +\infty$, which can be shown by just changing one value of Y_i only.

■

Proof of Proposition 6.

Just like in Proposition 5, we can formulate the problem as that of finding

$$G(N) = \max\{G_{L0}, G_{R0}, G_{0L}, G_{0R}, G_{LR}, G_{RL}, G_{LL}, G_{RR}\},$$

where G_{L0} and G_{R0} are sensitivities in situations of a new observation leaving the left or the right h-neighborhood, respectively; and G_{0L} and G_{0R} are sensitivities in situations of a new observation entering the left or the right h-neighborhood, respectively; G_{LR} and G_{RL} are sensitivities in cases of an observation switching the neighborhoods; G_{LL} and G_{RR} are sensitivities in cases when an observation changes within the same neighborhood. When we say "leaves a neighborhood" or "enters a neighborhood", we mean that with respect the value of X_i.

If we follow the textbook definition of differential privacy and, thus, consider all possible realizations with data no matter how small the probability of these realizations is as long as it is strictly positive, then we can show that $G(N) = +\infty$. Indeed, let us show e.g. that following the textbook definition of differential privacy, we have that $G_{RR} = \infty$.

Consider a situation when the first $T \leq N$ observations in our data are in the right-hand side neighborhood in the values of X_i. Consider, for example, realizations of the datasets when only T-th observation in the right-hand side neighborhood changes its value x_T while its values W_T and Y_T do not change. Then, of course,

$$\widehat{\alpha}_{y,L} = \widehat{\alpha}'_{y,L}, \quad \widehat{\alpha}_{w,L} = \widehat{\alpha}'_{w,L}.$$

Now we will use the same realizations for $X_1, \ldots, X_T$, X'_T as given in (6.3)-(6.4) in the proof of Proposition 5. Also, we will take the realization of the dataset when $W_i = 1$, $i = 1, \ldots, T$ or $W_i = 0$, $i = 1, \ldots, T$ (again, due to the fuzzy scenario, the probability of this scenario may be perceived as low but it is strictly positive and, thus, has to be taken into account by a differentially private mechanism). Then we can take, of course, that

$$\widehat{\alpha}_{w,R} = \widehat{\alpha}'_{w,R}$$

as the values of indicators $1(X_i \geq c)$, $i = 1, \ldots, T$, do not add any explanatory power in the local linear regression of W_i on constant and $1(X_i \geq c)$ in the right-hand side neighborhood (one can

think of this situation as the situation of the perfect fit in the reduced form in the IV regression even though technically $\widehat{\alpha}_{w,R}$ and $\widehat{\beta}_{w,R}$ may not be separately estimated in a sample like the one we suggested). Thus, changes in the value of $\widehat{\tau}_{F,LocLin}$ in (4.3) happen only because of the changes in the numerator. The changes in the numerator are, of course, the same as the changes in $\widehat{\alpha}_R$ described in the proof of Proposition 5 and, thus, manipulating δ, we can make this change arbitrarily large in the absolute value, leading us to conclusion that $G_{RR} = +\infty$.

Even if one wanted to deviate from the textbook definition of differential privacy and restrict W_i, $i = 1, \ldots, T$, to have some variation in each neighborhood – e.g. by requiring a minimum number of zero's and one's in each neighborhood or a fixed proportion – in this case, it would be straightforward to show that the global sensitivity would be bounded away from zero as $N \to \infty$.

Indeed, in this case even using the same example with T realizations of $X_1, \ldots, X_T$, X_T' in the right-hand side neighborhood with the values given in (6.3)-(6.4) in the proof of Proposition 5, we would obtain that manipulating δ approaching 0, change in both the numerator $\widehat{\alpha}_{y,R} - \widehat{\alpha}_{y,L}$ and the denominator $\widehat{\alpha}_{w,R} - \widehat{\alpha}_{w,L}$ are arbitrarily large in the absolute value but they become arbitrarily large with the same rate in δ, thus allowing us to conclude that the change is constant and show that this constant change may not diminish to 0 with the sample size.

It is obvious that when the support of $Y|X$ in the neighborhood to the right of c or to the left of c is unbounded, then $G(N) = +\infty$, which can be shown by just changing one value of Y_i only.

Note that for simplicity we used the uniform kernel to define (4.3). In the case if one were using a kernel with a bounded support but $\underline{K} > 0$ or a kernel with an unbounded support, the proof of the global sensitivity bounded away from zero would even be more straightforward.

■

Proof of Proposition 7 Since the support of Y is bounded by Assumption 3, the global sensitvity is determined by variation of the empirical weight $1/\widehat{P}(x)$ over $\mathcal{X}$. Then

$$\sup_{x,x' \in \mathcal{X}} \left| 1/\widehat{P}(x) - 1/\widehat{P}(x') \right| \geq \bar{K}/(h_N K(\text{diam}(\mathcal{X})/h_N)).$$

Note whenever $\text{diam}(\mathcal{X})$ is infinite, an infinte lower bound applies and (ii) immediately follows.

When $\text{diam}(\mathcal{X})$ is finite, global sensitivity of $\psi(\mathbb{P}_N)$ is bounded from below by $\bar{K}/(N\, h_N K(\text{diam}(\mathcal{X})/h_N))$. For $h_N = o(N^{-1/4})$, since $\lim_{|z| \to \infty} |z|^k K(|z|) = 0$ for $k \leq d$ with $d > 3$, then $N\, h_N K(\text{diam}(\mathcal{X})/h_N) = o(1)$ and, thus global sensitivity of $\psi(\mathbb{P}_N)$ does not decrease as $N \to \infty$.

■

Analysis of Example 2 without truncation

Laplace mechanism: Let $a_N(\nu_N) \sim Lap\,(0, 1/(C\,N))$ for a given C. Output

$$\theta(\mathbb{P}_N, \nu_N) = \bar{\xi} + a_N(\nu_N)$$

with $\bar{\xi} = \frac{1}{N}\sum_{i=1}^{N} \frac{x_i}{w_i}$.

Exponential mechanism: Take distribution density $p(z; \mathbb{P}_N) \propto \exp(-\frac{1}{2}C^2\,N^{1-2\gamma}\sum_{i=1}^{N}(z-\xi_i)^2)$ for a given C with $\xi_i = x_i/w_i$. Output

$$\theta(\mathbb{P}_N, \nu_N) \sim p(z; \mathbb{P}_N).$$

For each of the mechanisms we consider likelihood ratio of the distribution of the output of each of the two mechanisms with empirical distributions $\mathbb{P}_N$ and $\mathbb{P}'_N$ where in $\mathbb{P}'_N$ observation (x_i, w_i) is replaced with (x'_i, w'_i) while all other observations are identical to those in $\mathbb{P}_N$.

The log-likelihood ratio for the Laplace mechanism is

$$\log\left(\frac{f_{\theta(\mathbb{P}_N,\nu_N)}(z)}{f_{\theta(\mathbb{P}'_N,\nu_N)}(z)}\right) = -C\,N\left(|z-\bar{\xi}| - \left|z-\bar{\xi}-\frac{1}{N}\left(\frac{x'_i}{w'_i}-\frac{x_i}{w_i}\right)\right|\right).$$

Denote $\Delta = \frac{x'_i}{w'_i} - \frac{x_i}{w_i}$ and, without loss of generality, let $\Delta > 0$. Then the likelihood ratio is a decreasing function of z attaining its maximum $C\Delta$ for any $z \le \bar{\xi} - \Delta/N$ and minimum $-C\Delta$ for any $z \ge \bar{\xi}$. Then the likelihood ratio evaluated at the output of the Laplace mechanism $\theta(\mathbb{P}_N, \nu_N)$ is distributed as $Lap(C\,\Delta,\, 2)$ truncated at $[-C\,\Delta,\, C\Delta]$. Thus

$$\mathbb{P}\left(\log\left(\frac{f_{\theta(\mathbb{P}_N,\nu_N)}(\theta(\mathbb{P}_N,\nu_N))}{f_{\theta(\mathbb{P}'_N,\nu_N)}(\theta(\mathbb{P}_N,\nu_N))}\right) \ge \epsilon\right) = 1 - \exp(\epsilon - C\Delta).$$

When w'_i approaches 0, then Δ approaches $+\infty$. This means that for any ϵ parameter of differential privacy and any finite value of parameter C of the exponential mechanism, the log-likelihood ratio is greater than ϵ with probability 1. As a result, the for any finite scale parameter C Laplace mechanism cannot guarantee (ϵ, δ) -differential privacy for any ϵ and δ.

Similarly, we can express the log-likelihood ratio of the output of the exponential mechanism

$$\log\left(\frac{f_{\theta(\mathbb{P}_N,\nu_N)}(z)}{f_{\theta(\mathbb{P}'_N,\nu_N)}(z)}\right) = \frac{1}{2}N^{-2\gamma}C^2\Delta^2 - C^2\,N^{1-2\gamma}\,\Delta(z-\bar{\xi}).$$

Therefore,

$$\mathbb{P}\left(\log\left(\frac{f_{\theta(\mathbb{P}_N,\nu_N)}(\theta(\mathbb{P}_N,\nu_N))}{f_{\theta(\mathbb{P}'_N,\nu_N)}(\theta(\mathbb{P}_N,\nu_N))}\right) \ge \epsilon\right) = \Phi\left(\frac{C\Delta}{2\,N^{\gamma}} - \frac{\epsilon}{C\Delta}N^{\gamma}\right).$$

Similarly to the Laplace mechanism, we note that when w'_i approaches zero, then Δ approaches infinity. Given that function $a\,z - b/z$ is strictly increasing, then for any finite value of parameter C, the likelihood ratio exceeds ϵ with probability 1. As a result, any scaling parameter C of the exponential mechanism cannot guarantee (ϵ, δ)-differential privacy for any ϵ and $\delta < 1$.

6.3 Additional applications

6.3.1 Average treatment effect

A central problem in evaluation studies is that potential outcomes that program participants would have received in the absence of the program is not observed. Letting D_i denote a binary variable taking the value 1 if treatment was given to agent i, and 0 otherwise, and letting Y_{0i},Y_{1i} denote potential outcome variables, we refer to $Y_{1i} - Y_{0i}$ as the *treatment effect* for the i'th individual. A parameter of interest for identification and estimation is the *average treatment effect*, defined as:

$$\theta = E[Y_{1i} - Y_{0i}] \tag{6.5}$$

As in the previous section our notation will be to denote realizations of random variables by lower case letters and the random variables themselves by capital letters. One identification strategy for θ was proposed in [49], under the following assumption:

ASSUMPTION 2 (ATE under Conditional Independence)**.** *Let the following hold:*

(i) *There exists an observed variable X_i s.t.*

$$D_i \perp (Y_{0i}, Y_{1i})|X_i$$

(ii) $0 < P(D_i = 1|X_i) < 1 \quad \forall X_i$

See also [21], [23], [5]. The above assumption can be used to identify α as

$$\theta = E[E[Y|D = 1, X] - E[Y|D = 0, X]]. \tag{6.6}$$

The above parameter can be written as:

$$\theta_0 = E\left[\frac{Y(D - p(X))}{p(X)(1 - p(X))}\right], \tag{6.7}$$

where $p(X) = P(D = 1|X)$ is the propensity score. This parameter is a weighted moment condition where the denominator gets small if the propensity score approaches 0 or 1. Also, identification is lost when we remove any region in the support of X (so, fixed trimming will not identify θ above).

Consider the general setting of the treatment effect model under unconfoundeness with two potential continuous outcomes Y_1 and Y_0 and treatment D along with the vector of (continuous and discrete) covariates X. We assume that $(Y_1, Y_0) \perp D \,|\, X$. The observed outcome is

$$Y = Y_1 D + Y_0(1 - D).$$

In our setup the propensity score needs to be estimated as a function of X. In the further discussion without loss of generality we assume that X is single-dimensional. Our theory will be based on the following structure:

ASSUMPTION 3. *(i) X has a support $\mathcal{X}$ is a closed and continuous (but possible unbounded) set.*

(ii) $(Y_1, Y_0) \mid X = x$ has an absolutely continuous density for each $x \in \mathcal{X}$. Moreover the support of Y_k for $k = 0, 1$ is bounded.

(iii) The propensity score is strictly positive $P(\cdot) > 0$ on its support.

We consider the following procedure to implement an estimator for θ_0. First, a non-parametric estimator is used to estimate the propensity score

$$\widehat{P}(x) = \frac{\frac{1}{N}\sum_{i=1}^{N} D_i \, K\left(\frac{x - X_i}{h_N}\right)}{\frac{1}{N}\sum_{i=1}^{N} K\left(\frac{x - X_i}{h_N}\right)}, \tag{6.8}$$

where $K(\cdot)$ is a symmetric kernel and h_N is the bandwidth. Then the average treatment effect θ_0 is estimated as:

$$\widehat{\theta} = \psi(\mathbb{P}_N) \equiv \frac{1}{N}\sum_{i=1}^{N}\left(\frac{Y_i \, D_i}{\widehat{P}(X_i)} - \frac{Y_i \, (1 - D_i)}{1 - \widehat{P}(X_i)}\right). \tag{6.9}$$

In our analysis we focus on the kernel-based estimator for the propensity score without the loss of generality. One can use a different approach such as the series estimator where the number of terms used to approximate the function would play the role of the tuning parameter equivalent to the bandwidth parameter.

We consider the kernel functions $K(\cdot)$ with sub-polynomial tail behavior. In particular, we assume that there exists natural number $d > 3$ such that for all $k \leq d$, $\lim_{|z| \to \infty} |z|^k K(|z|) = 0$. This ensures existence of moments of kernel-weighted statistics over the distribution of X which is particularly helpful when the support of X is unbounded. We note that all "standard" kernel functions such as the bounded support uniform, quadratic and Epanechnikov kernel as well as the most commonly used Gaussian kernel satisfy this condition.

The bandwidth is required to satisfy $h_N \gg \frac{\log N}{N}$ (see [47]) to ensure uniform convergence of the propensity score estimator and is typically chosen so that $h_N = o(N^{-1/4})$ to avoid the propagation of the non-parametric bias to the estimator of the average treatment effect.

We now consider the impact of diffrential privacy on estimation of θ_0. As in Section 4 we rely on our smoothness assuption that allows us to focus on additive mechanisms to induce differential privacy. Also, like in our previous analysis of the regression discontinuity design we start with the analyis of the global sensitivity of $\psi(\mathbb{P}_N)$.

Proposition 7. *Suppose that average treatment effect estimator (6.9) uses propensity score estimator (6.8) and kernel function $K(\cdot)$ is such that $|K(\cdot)| \leq \bar{K}$ and there exists natural number $d > 3$ such that for all $k \leq d$, $\lim_{|z| \to \infty} |z|^k K(|z|) = 0$. then*

(i) If the support $\mathcal{X}$ is bounded and $h_N = o(N^{-1/4})$ the global sensitivity of functional $\psi(\mathbb{P}_N)$ in (6.9)) is bounded away from zero as $N \to \infty$.

(ii) If the support $\mathcal{X}$ is unbounded then the global sensitivity of functional $\psi(\mathbb{P}_N)$ in (6.9)) is $+\infty$.

This result allows us to formulate the following theorem.

THEOREM 9. *For estimation of average treatment effect any smooth $(\varepsilon_N, \delta_N)$-differentially private propensity score-weighted estimator is inconsistent for any bounded sequence $\{\varepsilon_N\}$ and non-negative $\{\delta_N\}$. As a result, the limiting random set $\mathbf{T}_{\mathcal{E}}$ contains at least one non-degenerate element different from $\{\theta_0\}$.*

References

[1] https:www.census.gov/about/policies/privacy/statistical$_s afeguards/disclosure - avoidance - 2020 - census.html$.

[2] Disclosure avoidance and the 2020 census. ***https***: ***//www.census.gov/about/policies/privacy/statistical_safeguards/disclos avoidance-2020-census.html***, 2019.

[3] Open letter to census bureau leadership. ***https***: ***//ipums.org/changes-to-census-bureau-data-products/open-letter-to-census-bureau-leadership***, 2019.

[4] Social science one announces access to facebook dataset of publicly shared urls for research. ***https***: ***//socialscience.one/blog/social-science-one-announces-access-facebook-dataset-publicly-shared-urls***, 2019.

[5] Alberto Abadie and Guido W Imbens. Large sample properties of matching estimators for average treatment effects. *Econometrica*, 74(1):235–267, 2006.

[6] G. Aggarwal, T. Feder, K. Kenthapadi, R. Motwani, R. Panigrahy, D. Thomas, and A. Zhu. Approximation algorithms for k-anonymity. *Journal of Privacy Technology*, 2005112001, 2005.

[7] A. Beresteanu, I. Molchanov, and F. Molinari. Partial identification using random set theory. *Journal of Econometrics*, 166(1):17–32, 2012.

[8] Arie Beresteanu and Francesca Molinari. Asymptotic properties for a class of partially identified models. *Econometrica*, 76(4):763–814, 2008.

[9] J. Brickell and V. Shmatikov. The cost of privacy: destruction of data-mining utility in anonymized data publishing. In *Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 70–78. ACM, 2008.

[10] Sebastian Calonico, Matias D. Cattaneo, and Rocio Titiunik. Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326, 2014.

[11] Matias D. Cattaneo, NicolÃ¡s Idrobo, and RocÃo Titiunik. *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Elements in Quantitative and Computational Methods for the Social Sciences. Cambridge University Press, 2020.

[12] Victor Chernozhukov and Han Hong. An mcmc approach to classical estimation. *Journal of Econometrics*, 115(2):293–346, 2003.

[13] V. Ciriani, S.D.C. di Vimercati, S. Foresti, and P. Samarati. k-anonymity. *Secure Data Management in Decentralized Systems. Springer-Verlag*, 2007.

[14] C. Dwork. Differential privacy. *Automata, languages and programming*, pages 1–12, 2006.

[15] 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.

[16] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.

[17] Cynthia Dwork, Adam Smith, Thomas Steinke, and Jonathan Ullman. Exposed! a survey of attacks on private data. *Annual Review of Statistics and Its Application*, 4:61–84, 2017.

[18] Jianqing Fan. Design-adaptive nonparametric regression. *Journal of the American Statistical Association*, 87(420):998–1004, 1992.

[19] Jianqing Fan and Irene Gijbels. *Local polynomial modelling and its applications*. Number 66 in Monographs on statistics and applied probability series. Chapman Hall.

[20] Ruobin Gong. Transparent privacy is principled privacy. *arXiv preprint arXiv:2006.08522*, 2020.

[21] Jinyong Hahn. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, pages 315–331, 1998.

[22] Jinyong Hahn, Petra Todd, and Wilbert Van der Klaauw. Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*, 69(1):201–209, 2001.

[23] Keisuke Hirano, Guido W Imbens, and Geert Ridder. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4):1161–1189, 2003.

[24] Guido Imbens and Karthik Kalyanaraman. Optimal bandwidth choice for the regression discontinuity estimator. *Review of Economic Studies*, 79:933–959, 2012.

[25] Guido W. Imbens and Thomas Lemieux. Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, 142:615–635, 2008.

[26] Matthew Jagielski, Jonathan Ullman, and Alina Oprea. Auditing differentially private machine learning: How private is private sgd? *arXiv preprint arXiv:2006.07709*, 2020.

[27] Aaron Johnson and Vitaly Shmatikov. Privacy-preserving data exploration in genome-wide association studies. In Inderjit S. Dhillon, Yehuda Koren, Rayid Ghani, Ted E. Senator, Paul Bradley, Rajesh Parekh, Jingrui He, Robert L. Grossman, and Ramasamy Uthurusamy, editors, *KDD*, pages 1079–1087. ACM, 2013.

[28] A. F. Karr, C. N. Kohnen, A. Oganian, J. P. Reiter, and A. P. Sanil. A framework for evaluating the utility of data altered to protect confidentiality. *The American Statistician*, 60(3):224–232, 2006.

[29] Toru Kitagawa. Estimation and inference for set-identified parameters using posterior lower probability. *working paper*, 2012.

[30] T. Komarova, D. Nekipelov, and E. Yakovlev. Estimation of treatment effects from combined data: Identification versus data security. In A. Goldfarb, S.M. Greenstein, and C.E. Tucker, editors, *Economic Analysis of the Digital Economy*. The University of Chicago Press, Chicago, 2015.

[31] Tatiana Komarova, Denis Nekipelov, Ahnaf Al Rafi, and Evgeny Yakovlev. K-anonymity: A note on the trade-off between data utility and data security. *Applied Econometrics*, 48:44â–62, 2017.

[32] Tatiana Komarova, Denis Nekipelov, and Evgeny Yakovlev. Identification, data combination, and the risk of disclosure. *Quantitative Economics*, 9(1):395–440, 2018.

[33] Anna Kormilitsina and Denis Nekipelov. Consistent variance of the laplace-type estimators: Application to dsge models. *International Economic Review*, 57(2):603–622, 2016.

[34] D. Lambert. Measures of disclosure risk and harm. *Journal of Official Statistics*, 9:313–313, 1993.

[35] David S. Lee and Thomas Lemieux. Regression discontinuity designs in economics. *Journal of Economic Literature*, 48:281â355, 2008.

[36] K. LeFevre, D.J. DeWitt, and R. Ramakrishnan. Incognito: Efficient full-domain k-anonymity. In *Proceedings of the 2005 ACM SIGMOD international conference on Management of data*, pages 49–60. ACM, 2005.

[37] K. LeFevre, D.J. DeWitt, and R. Ramakrishnan. Mondrian multidimensional k-anonymity. In *Data Engineering, 2006. ICDE'06. Proceedings of the 22nd International Conference*, pages 25–25. IEEE, 2006.

[38] Y. Lindell and B. Pinkas. Privacy preserving data mining. In *Advances in Cryptologyï¿œCRYPTO 2000*, pages 36–54. Springer, 2000.

[39] Ashwin Machanavajjhala, Daniel Kifer, John Abowd, Johannes Gehrke, and Lars Vilhuber. Privacy: Theory meets practice on the map. In *Proceedings of the 2008 IEEE 24th International Conference on Data Engineering*, ICDE â08, page 277â286, USA, 2008. IEEE Computer Society.

[40] C.F. Manski and E.T. Tamer. Inference on regressions with interval data on a regressor or outcome. *Econometrica*, 70(2):519–546, 2002.

[41] Justin McCrary. Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2):698–714, February 2008.

[42] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 94–103. IEEE, 2007.

[43] Xue Meng, Hui Li, and Jiangtao Cui. Different strategies for differentially private histogram publication. *J. Comm. Inform. Networks*, 2(3):68–77, 2017.

[44] I. Molchanov. *Theory of random sets*. Springer, 2005.

[45] Ilya Molchanov and Francesca Molinari. *Random Sets in Econometrics*. Econometric Society monographs. Cambridge University Press.

[46] A. Narayanan and V. Shmatikov. Robust de-anonymization of large sparse datasets. In *Security and Privacy, 2008. SP 2008. IEEE Symposium on*, pages 111–125. IEEE, 2008.

[47] David Pollard. Empirical processes: theory and applications. In *NSF-CBMS regional conference series in probability and statistics*, pages i–86. JSTOR, 1990.

[48] Jack Porter. Estimation in the regression discontinuity model. *working paper*, 2003.

[49] Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.

[50] P. Samarati and L. Sweeney. Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression. Technical report, Citeseer, 1998.

[51] L. Sweeney. Achieving k-anonymity privacy protection using generalization and suppression. *International Journal of Uncertainty Fuzziness and Knowledge-Based Systems*, 10(5):571–588, 2002.

[52] L. Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty Fuzziness and Knowledge Based Systems*, 10(5):557–570, 2002.

[53] Donald L. Thistlethwaite and Donald T. Campbell. Regression-discontinuity analysis: An alternative to the ex post facto experiment. *Journal of Educational Psychology*, 51(6):309–317, 1960.

[54] Caroline Uhler, Aleksandra B. Slavkovic, and Stephen E. Fienberg. Privacy-preserving data sharing for genome-wide association studies. *J. Priv. Confidentiality*, 5(1), 2013.

[55] publisher=Nauka, Moscow year=1976 Vladimirov, Vasilij. *Equations of mathematical physics*.

[56] M. Woo, J. P. Reiter, A. Oganian, and A. F. Karr. Global measures of data utility for microdata masked for disclosure limitation. *Journal of Privacy and Confidentiality*, 1(1):111–124, 2009.

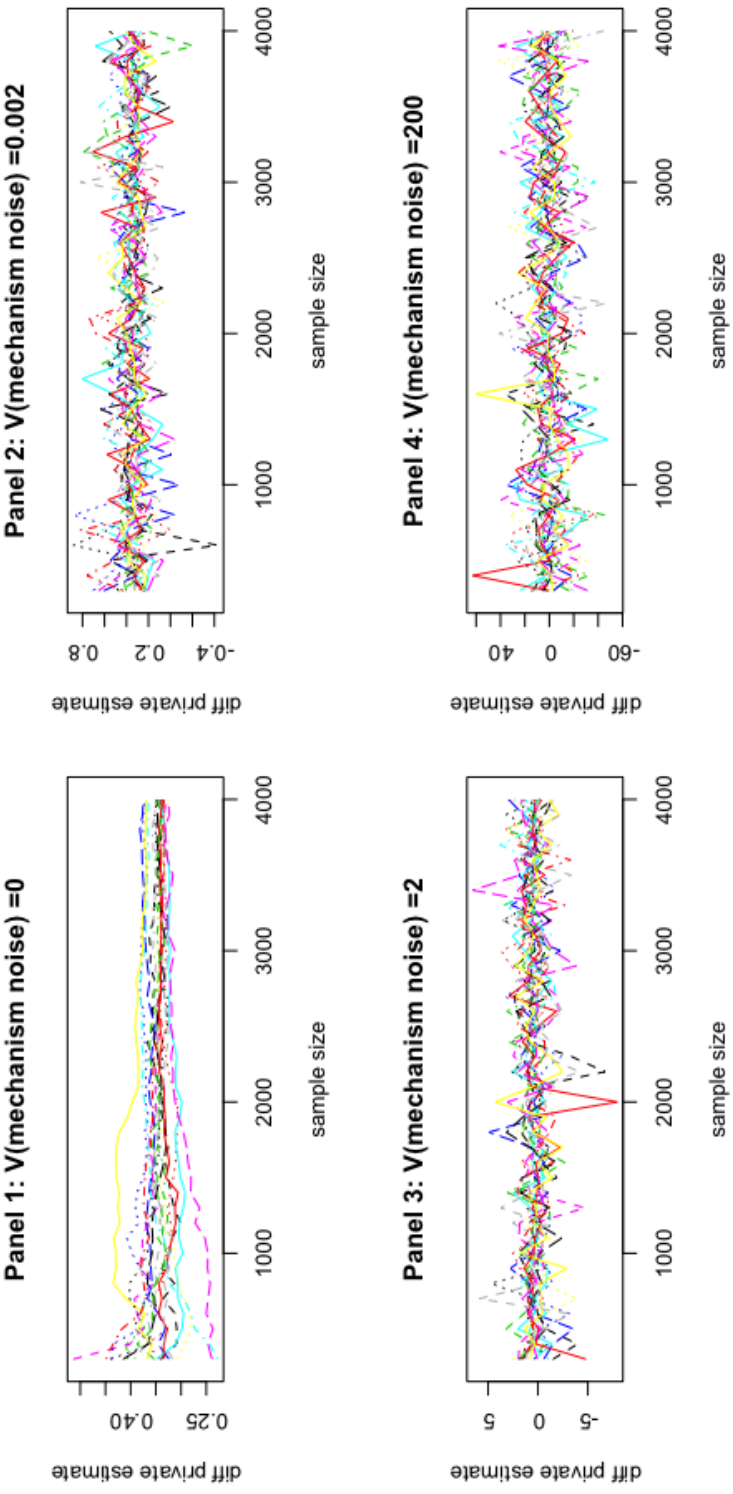


Figure 1: Illustration to Scenario 1. Twenty independent paths of differentially private estimators local linear estimators for increasing sample sizes for various degrees of differential privacy protection.

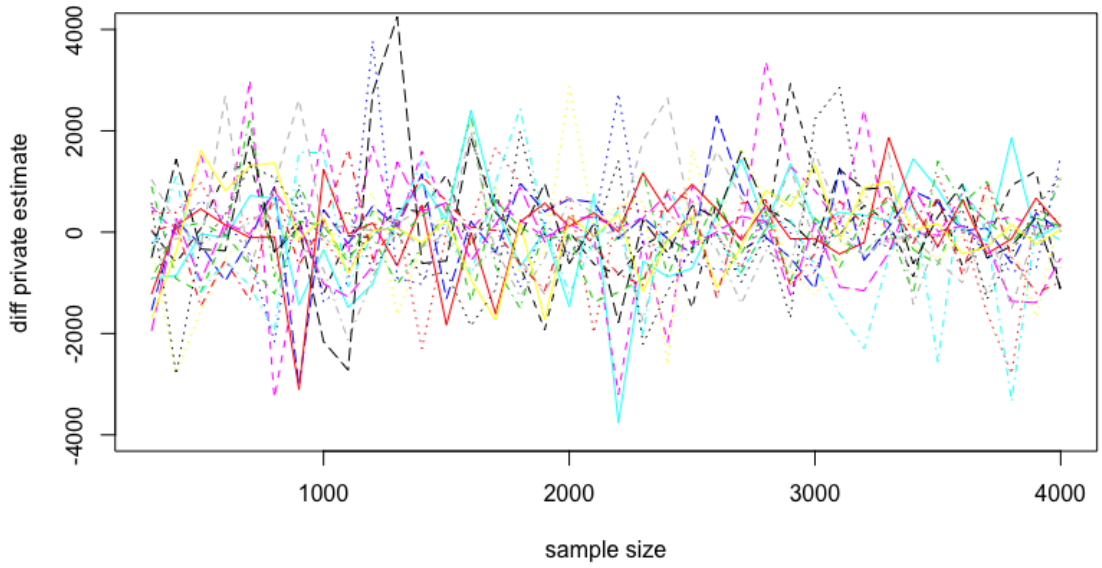


Figure 2: Illustration to Scenario 2. Twenty independent paths of differentially private estimators local linear estimators for increasing sample sizes when the Laplace mechanism noise has the variance equal to 10^6.