COMPARING APPROXIMATE AND PROBABILISTIC DIFFERENTIAL PRIVACY PARAMETERS

VINCENT GUINGONA, ALEXEI KOLESNIKOV, JULIANNE NIERWINSKI, AND AVERY SCHWEITZER

Abstract. This paper compares two notions of differential privacy: approximate differential privacy (ADP) and probabilistic differential privacy (PrDP). It is well-known that the PrDP implies the ADP; and it was established in [7] that the ADP implies the PrDP, after a penalty on the parameters ε and δ that are used in the definitions of both properties. We investigate how large do these penalties have to be. We show, in particular, that if the penalty on δ is $O(\delta)$, then the penalty on ε has to be bounded away from zero.

1. Introduction

The notion of differential privacy provides a rigorous framework for studying privacy properties of data analysis algorithms. The property was developed in a series of works, culminating in [4], [2], and [3]. Since its introduction in 2006, the notion of differential privacy was modified and extended in a variety of ways. A recent survey [1] has identified "approximately 225" different notions that build on the original definition. The survey provides a convenient taxonomy of the notions, identifying several "dimensions" along which the original definition is modified and describing the known connections between various definitions. Our paper focuses on the *Quantification of privacy loss* dimension, namely, on the connection between *approximate differential privacy* (ADP) and the *probabilistic differential privacy* (PrDP).

Both ADP and PrDP are properties of randomized algorithms; each of them depends on two parameters: ε, a bound on the privacy loss due to the algorithm, and δ, an allowance on the failure of the privacy

Date: May 9, 2022.

This work was partially supported by National Security Agency (NSA) grant H98230-20-1-0005, which funds a Research Experiences for Undergraduates program at Towson University. This work also received financial support from the Towson University Jess and Mildred Fisher College of Science and Mathematics.

loss bound. The fact that (ε, δ)-PrDP is a sufficient condition for (ε, δ)-ADP is well known. It is also known that ADP does not imply PrDP (for instance, [6] provides explanations and examples). Zhao et al. [7] established that (ε, δ)-ADP implies (ε', δ')-PrDP for $\varepsilon' > \varepsilon$ and $\delta' > \delta$, where the additive penalty $\varepsilon' - \varepsilon$ on the first parameter can be arbitrarily small and the multiplicative penalty on δ can be expressed as a function of ε and ε' (essentially, as a function of the additive penalty on ε). It follows from the formula in [7] that the multiplicative penalty on δ is unbounded as $\varepsilon' - \varepsilon \to 0$.

We note that the results of [7] have not yet appeared in a peer-reviewed journal. We provide an independent proof of the relevant result and investigate to what extent the result can be improved. Does (ε, δ)-ADP imply (ε', δ')-PrDP if $\varepsilon = \varepsilon'$ or if $\delta = \delta'$? If $\delta' = g(\delta)$, is it possible for $g(\delta)$ to have a bounded growth rate at 0?

We show that if either ε or δ is fixed, no reasonable penalty on the other parameter will suffice. Namely:

(1) For every $\varepsilon > 0$ and any function g such that $\lim\limits_{x \to 0^+} g(x) = 0$, there is an arbitrarily small δ and a randomized algorithm which is (ε, δ)-ADP, but not $(\varepsilon, g(\delta))$-PrDP.

(2) For all positive ε, ε', and δ, there is a randomized algorithm that is (ε, δ)-ADP, but not (ε', δ)-PrDP.

We establish a lower bound on the additive adjustment to ε in terms of the penalty function g on δ. We show, in particular, that if $g(x) = O(x)$ as $x \to 0$, then there is a constant ε^* that depends only on the choice of g such that for every $\varepsilon > 0$ there is an arbitrarily small δ' and a randomized algorithm which is (ε, δ')-ADP, but is not $(\varepsilon^*, g(\delta'))$-PrDP.

2. Background

We begin by defining the main notions. Our presentation here largely follows [5] and [1].

Definition 2.1. If B is a countable set, we denote by $\Delta(B)$ the set of all probability mass functions on B. That is,

$$\Delta(B) = \Big\{ f : B \to [0,1] \mid \sum_{b \in B} f(b) = 1 \Big\}.$$

A *randomized algorithm* $\mathcal{M}$ with domain A and range B is an algorithm associated with a function $M : A \to \Delta(B)$. On an input $a \in A$, the algorithm $\mathcal{M}$ outputs $\mathcal{M}(a) = b$ with probability $[M(a)](b)$ for every $b \in B$.

We will use short-hand notation $\mathbb{P}(\mathcal{M}(a) = b)$ and $\mathbb{P}(\mathcal{M}(a) \in S)$ to denote $[M(a)](b)$ and $\sum_{b \in \mathcal{S}} [M(a)](b)$, respectively. If $a \in A$ and P is a property on B, we use $\mathbb{P}_{b\sim\mathcal{M}(a)}(P(b))$ to denote the probability that $\mathcal{M}(a)$ satisfies P; i.e.,

$$\mathbb{P}_{b\sim\mathcal{M}(a)}(P(b)) = \mathbb{P}\big(\mathcal{M}(a) \in \{b \in B \mid P(b)\}\big).$$

The domain of a randomized algorithm is typically a collection of datasets, drawn from a large universe of records. The notions of differential privacy capture the intuitive idea that the probability distributions $M(a)$ and $M(a')$ should be "similar" if the datasets a and a' differ by a single record. Formally, one equips the domain A of the algorithm with a discrete integer-valued metric d. Elements $a, a' \in A$ are called *adjacent* if $d(a, a') \leq 1$. There is a variety of ways to capture similarity of two distributions; this leads to slightly different notions of differential privacy.

Definition 2.2. A randomized algorithm $\mathcal{M}$ with domain A and range B is (ε, δ)-*approximately differentially private* (ADP) if for all $S \subseteq B$ and for all adjacent $a, a' \in A$

$$\mathbb{P}(\mathcal{M}(a) \in S) \leq e^{\varepsilon}\mathbb{P}(\mathcal{M}(a') \in S) + \delta.$$

A randomized algorithm $\mathcal{M}$ is (ε, δ)-*probabilistically differentially private* (PrDP) if for all adjacent $a, a' \in A$

$$\mathbb{P}_{c\sim\mathcal{M}(a)}\left[\frac{\mathbb{P}(\mathcal{M}(a) = c)}{\mathbb{P}(\mathcal{M}(a') = c)} > e^{\varepsilon}\right] \leq \delta.$$

In practice, the (ε, δ)-PrDP is often easier to verify than (ε, δ)-ADP: the condition states that the set of outputs of $\mathcal{M}$ that witness large privacy loss has small measure. It is a well-known fact that for all $\varepsilon, \delta > 0$, if $\mathcal{M}$ is (ε, δ)-PrDP, then $\mathcal{M}$ is (ε, δ)-ADP, a complete proof can be found, for example, in [7, Appendix I]. It is also immediate from the definitions that $(\varepsilon, 0)$-ADP implies $(\varepsilon, 0)$-PrDP.

However, one can show that, in general, (ε, δ)-ADP does not imply (ε, δ)-PrDP, one example can be found in [6]. A number of papers mention that ADP implies PrDP up to a small loss in parameters without quantifying the loss. The paper [7] provides a precise statement and a proof. Zhao et al. state the definition of PrDP in a "two-tailed" form, so the parameters for PrDP listed in Lemma 4 of [7] are different from what we state here. We use the definition of PrDP from [1] (we found this form to be more common in the literature). For the definition of PrDP we use, the change in parameters to obtain PrDP from ADP can be derived from Lemma 12 of [7]. In the terminology of our paper, the lemma states that for positive ε and δ and $\hat{\varepsilon} > \varepsilon$, if $\mathcal{M}$ is a (ε, δ)-ADP randomized algorithm, then $\mathcal{M}$ is $\left(\hat{\varepsilon}, \frac{\delta}{1-\exp(\varepsilon-\hat{\varepsilon})}\right)$-PrDP.

Note that $\hat{\varepsilon}$ can be arbitrarily close to ε, but then the multiplicative penalty $\frac{1}{1-\exp(\varepsilon-\hat{\varepsilon})}$ on δ becomes unbounded. Our goal is to gain a better understanding of the trade-off between the penalties on ε and δ.

In the context of machine learning, randomized algorithms are often defined for datasets of varying sizes. Formulas describing properties of the algorithm typically connect the size m of the datasets with the privacy parameters ε and δ as well as the accuracy parameters. This, in turn, allows to estimate what level of privacy one can expect from the algorithm that inputs a dataset of size m; i.e., ε and δ can be treated as functions of m. To achieve meaningful privacy, one needs to have $\delta(m) = o(1/m^k)$ for all k (see [5] for a discussion); the parameter ε is more forgiving, but one generally expects that $\varepsilon(m) \to 0$ as $m \to \infty$.

We find it convenient to treat the penalty on δ as a function g of δ, and then write the additive penalty on ε in terms of δ and $g(\delta)$. Restating Lemma 12 of [7] in these terms, we get the following result.

Lemma 2.3. *Let ε and δ be positive real numbers and suppose $g : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is a continuous function with $g(0) = 0$. If*

$$\hat{\varepsilon} \geq \varepsilon - \ln\left(1 - \frac{\delta}{g(\delta)}\right)$$

and $\mathcal{M} : A \to B$ is a (ε, δ)-ADP randomized algorithm, then $\mathcal{M}$ is $(\hat{\varepsilon}, g(\delta))$-PrDP.

We include a short self-contained proof of the lemma in Section 4. The main goal of our paper is to explore to what extent can the bound in the lemma be improved.

3. Results

We now summarize the main results of this paper. In this section, we will provide informal outlines of the proofs of the statements. The formal proofs will appear in Section 4.

A useful tool in the proof of Theorem 3.2 and Theorem 3.3 is the following lemma.

Lemma 3.1. *Let ε, $\hat{\varepsilon}$, δ, $\hat{\delta}$, p and q be positive real numbers and let $n \geq 2$ be an integer such that the following conditions hold:*

(1) $p, q \in (0, 1)$;
(2) $p + (n-1)q = 1$;
(3) $e^{\hat{\varepsilon}} q < p \leq e^{\varepsilon} q + \delta$; *and*
(4) $p > \hat{\delta}$.

Then there is a randomized algorithm $\mathcal{M} : [n] \to [n]$ that is (ε, δ)-ADP, but is not $(\hat{\varepsilon}, \hat{\delta})$-PrDP.

The algorithm $\mathcal{M}$ in Lemma 3.1 takes as input an element $a \in [n]$ and outputs a with probability p and outputs any other element of $[n]$ with probability q. The first two conditions ensure that this is a valid probability distribution. The conditions (3) and (4) ensure that the privacy loss $\ln \dfrac{\mathbb{P}(\mathcal{M}(a) = a)}{\mathbb{P}(\mathcal{M}(a') = a)}$ from observing the output a is more than $\hat{\varepsilon}$, the probability of observing the output a is at least $\hat{\delta}$, and that the (ε, δ)-ADP holds. The randomized algorithm constructed in the lemma inputs datasets of size 1; it is possible to modify the construction to define $\mathcal{M}$ on datasets of size $m > 1$.

Using the lemma, we establish that, when we fix δ, for any choice of ε and $\hat{\varepsilon}$, (ε, δ)-ADP does not imply $(\hat{\varepsilon}, \delta)$-PrDP.

Theorem 3.2. *For all $0 < \varepsilon < 1$, $0 < \delta < \frac{1}{2}$, and $\hat{\varepsilon} > 0$, there exists a (ε, δ)-ADP randomized algorithm $\mathcal{M}$ that is not $(\hat{\varepsilon}, \delta)$-PrDP.*

To prove Theorem 3.2, we appeal to Lemma 3.1. We need to choose p and q satisfying the following linear system

$$\begin{cases} p - \frac{e^{\hat{\varepsilon}} + e^{\varepsilon}}{2} q = \frac{\delta}{2}, \\ p + (n-1)q = 1. \end{cases}$$

The first equation gives us condition (3) and the second equation gives us condition (1) of Lemma 3.1. A sufficiently large choice of n yields conditions (2) and (4).

The following theorem shows, in particular, that for any reasonable penalty function g on δ, no matter how fast growing, there is an arbitrarily small δ' such that (ε, δ')-ADP does not imply $(\varepsilon, g(\delta'))$-PrDP.

Theorem 3.3. *Let $\hat{\varepsilon}, \varepsilon > 0$ and $0 < \delta < 1/2$. Let $g : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be a continuous function such that $g(0) = 0$ and*

$$\hat{\varepsilon} < \varepsilon + \ln\left(1 + \frac{x}{g(x)}\right)$$

for all $0 < x < \delta$. Then there exists $\delta' \in (0, \delta)$ and a randomized function $\mathcal{M}$ that is (ε, δ')-ADP but not $(\hat{\varepsilon}, g(\delta'))$-PrDP.

To prove Theorem 3.3, we need to construct δ', n, p, and q satisfying the conditions of Lemma 3.1 (with $\hat{\delta} = g(\delta')$). This is similar to the proof of Theorem 3.2, with the added complication due to the function g. First we choose n sufficiently large to be able to use the Intermediate Value Theorem to select $\delta' \in (0, \delta)$ such that

$$\frac{e^{\hat{\varepsilon}} q + e^{\varepsilon} q + \delta'}{2} = g(\delta') + \delta'.$$

Letting p be the common value above and $q = \frac{1-p}{n-1}$, we are able to satisfy all of the conditions of Lemma 3.1.

From Theorem 3.3, we can immediately conclude the following.

Corollary 3.4. *(1) For every $\varepsilon > 0$, every $0 < \delta < 1/2$, and every continuous $g : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ such that $g(0) = 0$, there exists $\delta' \in (0, \delta)$ and a randomized function $\mathcal{M}$ that is (ε, δ')-ADP but not $(\varepsilon, g(\delta'))$-PrDP.*

(2) If $g : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is a continuous function such that $g(0) = 0$ and $g(x) \leq Kx$ for all sufficiently small x, then for every $\varepsilon > 0$ and every $\hat{\varepsilon} \in (0, \ln(1 + 1/K))$, there is $\delta' > 0$ and a randomized function $\mathcal{M}$ that is (ε, δ')-ADP but not $(\hat{\varepsilon}, g(\delta'))$-PrDP.

The second statement in the corollary shows that, for practical purposes, the penalty function on δ cannot be bounded by a constant multiple of δ. Otherwise, the parameter $\hat{\varepsilon}$ would be bounded away from zero, for arbitrarily small ε.

4. Proofs

Proof of Lemma 2.3. Let $\hat{\varepsilon} \geq \varepsilon - \ln\left(1 - \frac{\delta}{g(\delta)}\right)$ and let $a, a' \in A$ be adjacent.

Let $S = \{c \in B \mid \mathbb{P}(\mathcal{M}(a) = c) > e^{\hat{\varepsilon}}\mathbb{P}(\mathcal{M}(a') = c)\}$. Note that $\mathbb{P}(\mathcal{M}(a') \in S) < e^{-\hat{\varepsilon}}\mathbb{P}(\mathcal{M}(a) \in S)$. Since $\mathcal{M}$ is (ε, δ)-ADP, we have

$$\begin{aligned}\mathbb{P}(\mathcal{M}(a) \in S) \leq e^{\varepsilon}\mathbb{P}(\mathcal{M}(a') \in S) + \delta &< e^{\varepsilon - \hat{\varepsilon}}\mathbb{P}(\mathcal{M}(a) \in S) + \delta \\ &= \left(1 - \frac{\delta}{g(\delta)}\right)\mathbb{P}(\mathcal{M}(a) \in S) + \delta.\end{aligned}$$

It follows that $\mathbb{P}(\mathcal{M}(a) \in S) < g(\delta)$, as needed. □

Proof of Lemma 3.1. Consider the randomized algorithm $\mathcal{M}$ associated with the function $M : [n] \to \Delta([n])$ defined by

$$[M(a)](c) = \begin{cases} p, & c = a \\ q, & \text{otherwise.} \end{cases}$$

for $a, c \in [n]$. We show this algorithm has (ε, δ)-ADP.

Let $a, b \in [n]$ be distinct elements, they will play the role of adjacent datasets.

For $c \in [n]$, if $c = a$, then by property (3)

$$[M(a)](c) \leq e^{\varepsilon}[M(b)](c) + \delta.$$

If $c \neq a$, then by the definition of $\mathcal{M}$ and property (3), we get

$$[M(a)](c) \leq e^{\varepsilon}[M(b)](c).$$

If S is a subset of $[n]$, adding the inequalities for $c \in S$, we see that

$$\mathbb{P}(\mathcal{M}(a) \in S) \leq e^{\varepsilon}\mathbb{P}(\mathcal{M}(b) \in S) + \delta,$$

so $\mathcal{M}$ is (ε, δ)-ADP.

We now verify that

$$\mathbb{P}_{c \sim \mathcal{M}(a)}\left[\frac{\mathbb{P}(\mathcal{M}(a) = c)}{\mathbb{P}(\mathcal{M}(b) = c)} > e^{\hat{\varepsilon}}\right] = \mathbb{P}_{c \sim \mathcal{M}(a)}(c = a) = p.$$

Since $p > \hat{\delta}$ by property (4), it would follow that $\mathcal{M}$ is not $(\hat{\varepsilon}, \hat{\delta})$-PrDP.

If $c = a$, then by property (3)

$$\frac{\mathbb{P}(\mathcal{M}(a) = c)}{\mathbb{P}(\mathcal{M}(b) = c)} = \frac{p}{q} > e^{\hat{\varepsilon}}.$$

If $c \neq a$, then by the definition of $\mathcal{M}$ we have

$$\frac{\mathbb{P}(\mathcal{M}(a) = c)}{\mathbb{P}(\mathcal{M}(b) = c)} \leq 1 \leq e^{\hat{\varepsilon}}.$$

This completes the proof. □

Proof of Theorem 3.2. Without loss of generality, we may assume that $\hat{\varepsilon} > \varepsilon - \ln(1 - \delta)$, since failure of PrDP for a large value of $\hat{\varepsilon}$ implies failure of PrDP for smaller values of $\hat{\varepsilon}$. By Lemma 3.1 it is enough to find $n \geq 2$, p, and q satisfying

(1) $p + (n - 1)q = 1$,
(2) $p, q \in (0, 1)$,
(3) $e^{\hat{\varepsilon}}q < p \leq e^{\varepsilon}q + \delta$, and
(4) $p > \delta$.

The number n needs to be chosen carefully: it must be large enough to ensure (3) but small enough to ensure (4). Let $n = \left\lceil \frac{(e^{\hat{\varepsilon}} + e^{\varepsilon})(1-\delta)}{\delta} \right\rceil$. Then

$$n - 1 < \frac{(e^{\hat{\varepsilon}} + e^{\varepsilon})(1 - \delta)}{\delta}. \tag{$*$}$$

Since $\delta < \frac{1}{2}$ it can verified that

$$\frac{(e^{\hat{\varepsilon}} + e^{\varepsilon})(1 - \delta)}{\delta} - \frac{e^{\hat{\varepsilon}} - e^{\varepsilon} - \delta e^{\hat{\varepsilon}}}{\delta} > 1,$$

we also have

$$n - 1 > \frac{e^{\hat{\varepsilon}} - e^{\varepsilon} - \delta e^{\hat{\varepsilon}}}{\delta} \tag{$**$}$$

The choice of $\hat{\varepsilon}$ guarantees that $\frac{e^{\hat{\varepsilon}}-e^{\varepsilon}-\delta e^{\hat{\varepsilon}}}{\delta} > 0$. Now, we let p be the average of $e^{\hat{\varepsilon}}q$ and $e^{\varepsilon}q+\delta$, so p and q are solutions to the system

$$\begin{cases} p - \frac{e^{\hat{\varepsilon}}+e^{\varepsilon}}{2}q = \frac{\delta}{2}, \\ p+(n-1)q = 1. \end{cases}$$

Namely, $q = \frac{2-\delta}{2(n-1)+e^{\hat{\varepsilon}}+e^{\varepsilon}}$, $p = \frac{e^{\hat{\varepsilon}}q+e^{\varepsilon}q+\delta}{2}$. By $(*)$ and $(**)$ we get

$$\frac{\delta}{e^{\hat{\varepsilon}}+e^{\varepsilon}} < q < \frac{\delta}{e^{\hat{\varepsilon}}-e^{\varepsilon}}.$$

The first inequality and our choice of p implies (4) and the second inequality and our choice of p gives us (3). Since the four properties are verified, the claim holds by Lemma 3.1.

□

Proof of Theorem 3.3. By Lemma 3.1, it suffices to show that for some integer n there exist values of p, q, and $\delta' \in (0,\delta)$ such that the following properties hold:

(1) $p, q \in (0,1)$,
(2) $p+(n-1)q = 1$,
(3) $e^{\hat{\varepsilon}}q < p \le e^{\varepsilon}q+\delta'$,
(4) $p > g(\delta')$.

For a suitably chosen n, we find δ', p, and q such that

$$(*) \qquad \begin{cases} p+(n-1)q = 1 \\ p = \frac{1}{2}\left(e^{\hat{\varepsilon}}q+e^{\varepsilon}q+\delta'\right) \\ p = g(\delta')+\delta'. \end{cases}$$

From the first two equations of $(*)$ we have $p = \dfrac{\delta'(n-1)+e^{\hat{\varepsilon}}+e^{\varepsilon}}{2(n-1)+e^{\hat{\varepsilon}}+e^{\varepsilon}}$. We want to find $\delta' \in (0,\delta)$ such that $p = g(\delta')+\delta'$. Choose n so that

$$n > \frac{(1-\delta)(e^{\hat{\varepsilon}}+e^{\varepsilon})}{\delta}+1.$$

Note that $n > 3$. Let $h(x) := \dfrac{x(n-1)+e^{\hat{\varepsilon}}+e^{\varepsilon}}{2(n-1)+e^{\hat{\varepsilon}}+e^{\varepsilon}}$. By our choice of n, we have that $e^{\hat{\varepsilon}}+e^{\varepsilon} < \delta[(n-1)+e^{\hat{\varepsilon}}+e^{\varepsilon}]$. So

$$h(\delta) = \frac{\delta(n-1)+e^{\hat{\varepsilon}}+e^{\varepsilon}}{2(n-1)+e^{\hat{\varepsilon}}+e^{\varepsilon}} < \frac{\delta[2(n-1)+e^{\hat{\varepsilon}}+e^{\varepsilon}]}{2(n-1)+e^{\hat{\varepsilon}}+e^{\varepsilon}} = \delta.$$

Since $h(0) > g(0) + 0$ and $h(\delta) < g(\delta) + \delta$, the Intermediate Value Theorem implies that there exists some $\delta' \in (0, \delta)$ such that

$$(**) \qquad h(\delta') = \frac{\delta'(n-1) + e^{\hat{\varepsilon}} + e^{\varepsilon}}{2(n-1) + e^{\hat{\varepsilon}} + e^{\varepsilon}} = g(\delta') + \delta'.$$

Property (2) holds since p, q, and δ' satisfy $(*)$. Since $n \geq 2$ and $\delta' < 1/2$, we have that $p \in (0,1)$. Thus $q \in (0,1)$ by property (2). So property (1) holds.

Since $p = \dfrac{e^{\hat{\varepsilon}}q + e^{\varepsilon}q + \delta'}{2}$, to show property (3) it suffices to show $e^{\hat{\varepsilon}}q < e^{\varepsilon}q + \delta'$. By our assumption that $\hat{\varepsilon} < \varepsilon + \ln\left(1 + \dfrac{x}{g(x)}\right)$ for all $x \in (0, \delta)$, we have $e^{\varepsilon}\delta' > g(\delta')(e^{\hat{\varepsilon}} - e^{\varepsilon})$. So

$$\begin{aligned}\delta'(e^{\varepsilon} + e^{\hat{\varepsilon}}) &= 2e^{\varepsilon}\delta' + (e^{\hat{\varepsilon}}\delta' - e^{\varepsilon}\delta') > 2g(\delta')(e^{\hat{\varepsilon}} - e^{\varepsilon}) + (e^{\hat{\varepsilon}}\delta' - e^{\varepsilon}\delta') \\ &= (2g(\delta') + \delta')(e^{\hat{\varepsilon}} - e^{\varepsilon}).\end{aligned}$$

Therefore, we have that $\dfrac{e^{\hat{\varepsilon}} + e^{\varepsilon}}{2g(\delta') + \delta'} > \dfrac{e^{\hat{\varepsilon}} - e^{\varepsilon}}{\delta'}$. By Equation $(**)$, we have that

$$\begin{aligned}n - 1 &= \frac{e^{\hat{\varepsilon}} + e^{\varepsilon}}{2g(\delta') + \delta'}(1 - \delta' - g(\delta')) \\ &= \frac{e^{\hat{\varepsilon}} + e^{\varepsilon}}{2g(\delta') + \delta'}(1 - p) > \frac{e^{\hat{\varepsilon}} - e^{\varepsilon}}{\delta'}(1 - p).\end{aligned}$$

Then by property (2), we have $e^{\hat{\varepsilon}}q < e^{\varepsilon}q + \delta'$.

For property (4), since $\delta' > 0$ we have that $p = g(\delta') + \delta' > g(\delta')$. As the four properties are verified, the claim holds by Lemma 3.1.

□

References

[1] D. Desfontaines and B. Pejó, *Sok: differential privacies*, Proceedings on privacy enhancing technologies **2020** (2020), no. 2, 288–313.

[2] C. Dwork, *Differential privacy*, International colloquium on automata, languages, and programming, 2006, pp. 1–12.

[3] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, *Our data, ourselves: Privacy via distributed noise generation*, Annual international conference on the theory and applications of cryptographic techniques, 2006, pp. 486–503.

[4] C. Dwork, F. McSherry, K. Nissim, and A. Smith, *Calibrating noise to sensitivity in private data analysis*, Theory of cryptography conference, 2006, pp. 265–284.

[5] C. Dwork, A. Roth, et al., *The algorithmic foundations of differential privacy.*, Found. Trends Theor. Comput. Sci. **9** (2014), no. 3-4, 211–407.

[6] S. Meiser, *Approximate and probabilistic differential privacy definitions.*, IACR Cryptol. ePrint Arch. **2018** (2018), 277.

[7] J. Zhao, T. Wang, T. Bai, K.-Y. Lam, Z. Xu, S. Shi, X. Ren, X. Yang, Y. Liu, and H. Yu, *Reviewing and improving the gaussian mechanism for differential privacy*, arXiv preprint arXiv:1911.12060 (2019).

Towson University, 7800 York Rd., Towson, MD, 21252
Email address: `vguingona@towson.edu`
Email address: `akolesnikov@towson.edu`