Offset-Symmetric Gaussians for Differential Privacy

Parastoo Sadeghi and Mehdi Korki

August 31, 2022

P. Sadeghi's research is supported by the Future Fellowship Scheme, Australian Research Council.

Summary of contributions

- A new "hybrid" differential privacy mechanism is proposed, which is somewhat between Laplace and Gaussian mechanisms.

- It is rooted in the Gaussian mechanism → analytical privacy performance derivations.

- It can have sub-Gaussian tail → desirable for reducing outliers or post-processing bias.

- At the same utility (measured by variance), it has better (ϵ, δ) and better Rényi differential privacy performance than the Gaussian.

- To achieve the same (ϵ, δ) differential privacy levels, it adds less noise to query (measured by noise variance).

Reconstruction attacks

Facts

- Between 50% to 80% of people in the US are uniquely identified by their full DoB, ZIP code and sex.
- Simple “anonymization” of datasets and publishing them or using them in training algorithms is a privacy risk.
- In mid 1990s, Latanya Sweeney managed to hack into “anonymized” health records of the Governor of MA by linking it with publicly available voter data.

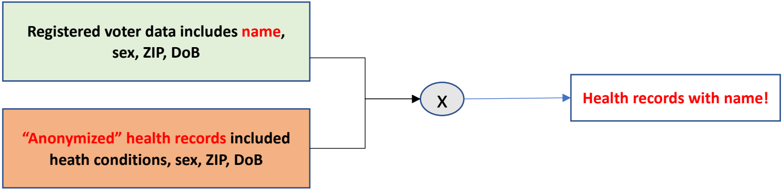

Registered voter data includes **name**, sex, ZIP, DoB

“Anonymized” health records included heath conditions, sex, ZIP, DoB

X

Health records with name!

Reverse engineering aggregated/swapped datacells has become easy

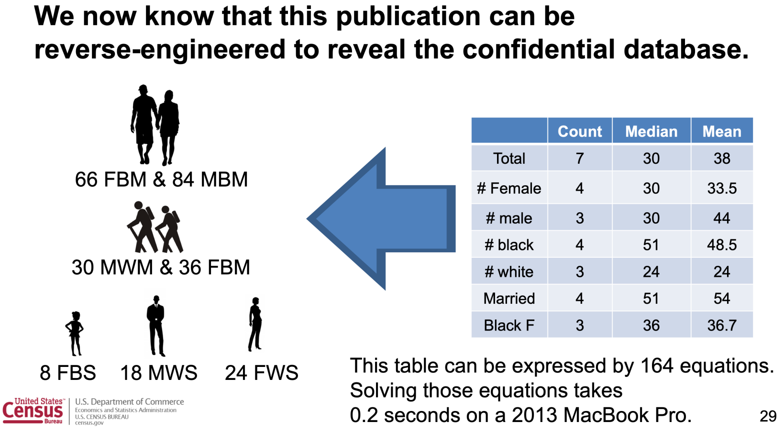

We now know that this publication can be reverse-engineered to reveal the confidential database.

66 FBM & 84 MBM

30 MWM & 36 FBM

8 FBS 18 MWS 24 FWS

	Count	Median	Mean
Total	7	30	38
# Female	4	30	33.5
# male	3	30	44
# black	4	51	48.5
# white	3	24	24
Married	4	51	54
Black F	3	36	36.7

This table can be expressed by 164 equations.
Solving those equations takes
0.2 seconds on a 2013 MacBook Pro.

United States Census Bureau | U.S. Department of Commerce
Economics and Statistics Administration
U.S. CENSUS BUREAU
census.gov

29

Source: `https://simson.net/ref/2019/2019-07-16%20Deploying%20Differential%20Privacy%20for%20the%202020%20Census.pdf`

Reconstructions attacks are possible

Facts

- Reconstructed over 300 million records from the 2010 census publicly released data
- Used 4 commercial databases which included real people's name, address, age, sex
- Linked reconstructed records with commercial databases including name, address, age, sex, ethnicity and race
- Comparing with confidential US census data could get all variables (including race & ethnicity) right for 17% of the US population

Source: `https://simson.net/ref/2019/2019-07-16%20Deploying%20Differential%20Privacy%20for%20the%202020%20Census.pdf`

Use of differential privacy in the 2020 US Census was announced in 2018

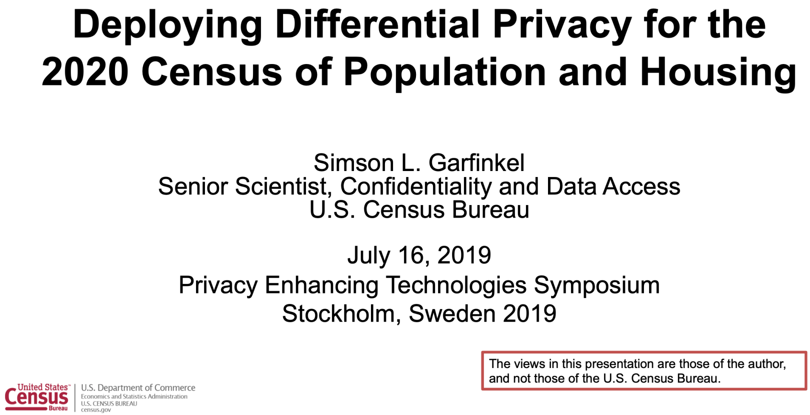

Deploying Differential Privacy for the 2020 Census of Population and Housing

Simson L. Garfinkel
Senior Scientist, Confidentiality and Data Access
U.S. Census Bureau

July 16, 2019
Privacy Enhancing Technologies Symposium
Stockholm, Sweden 2019

The views in this presentation are those of the author, and not those of the U.S. Census Bureau.

United States Census Bureau
U.S. Department of Commerce
Economics and Statistics Administration
U.S. CENSUS BUREAU
census.gov

Source: `https://simson.net/ref/2019/2019-07-16%20Deploying%20Differential%20Privacy%20for%20the%202020%20Census.pdf`

NY Times Article, Feb 2020

US Census in 2020

- “Every person matters for federal funding.”
- To preserve privacy: “Imaginary people will be added to some locations and real people will be removed from others.”
- “Minorities and rural areas at most risk.”

Opinion

Changes to the Census Could Make Small Towns Disappear

By Gus Wezerek and David Van Riper

FEB. 6, 2020

Source: `https://www.nytimes.com/interactive/2020/02/06/opinion/census-algorithm-privacy.html#commentsContainer`

NY Times Article, Feb 2020

Outliers Matter

- Y axis: true population count (log scale)
- X axis: reported count relative to the true population count (log scale)

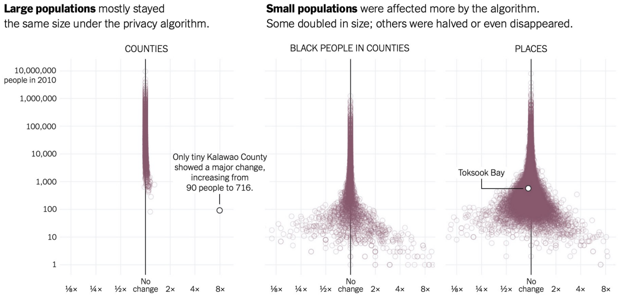

Large populations mostly stayed the same size under the privacy algorithm.

Small populations were affected more by the algorithm. Some doubled in size; others were halved or even disappeared.

Only tiny Kalawao County showed a major change, increasing from 90 people to 716.

Toksook Bay

Source: `https://www.nytimes.com/interactive/2020/02/06/opinion/census-algorithm-privacy.html#commentsContainer`

NY Times Article, Feb 2020

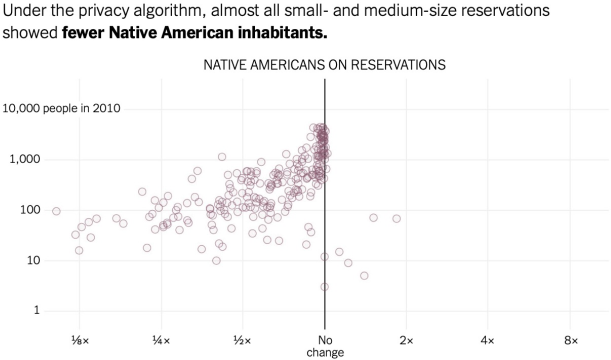

Under the privacy algorithm, almost all small- and medium-size reservations showed **fewer Native American inhabitants.**

Source: `https://www.nytimes.com/interactive/2020/02/06/opinion/census-algorithm-privacy.html#commentsContainer`

ABS is also interested in the potential of DP in its release methodologies

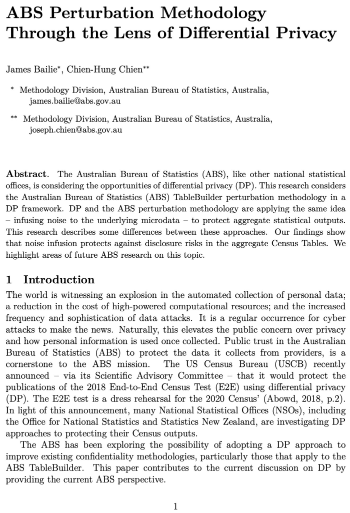

ABS Perturbation Methodology Through the Lens of Differential Privacy

James Bailie*, Chien-Hung Chien**

* Methodology Division, Australian Bureau of Statistics, Australia, james.bailie@abs.gov.au

** Methodology Division, Australian Bureau of Statistics, Australia, joseph.chien@abs.gov.au

Abstract. The Australian Bureau of Statistics (ABS), like other national statistical offices, is considering the opportunities of differential privacy (DP). This research considers the Australian Bureau of Statistics (ABS) TableBuilder perturbation methodology in a DP framework. DP and the ABS perturbation methodology are applying the same idea – infusing noise to the underlying microdata – to protect aggregate statistical outputs. This research describes some differences between these approaches. Our findings show that noise infusion protects against disclosure risks in the aggregate Census Tables. We highlight areas of future ABS research on this topic.

1 Introduction

The world is witnessing an explosion in the automated collection of personal data; a reduction in the cost of high-powered computational resources; and the increased frequency and sophistication of data attacks. It is a regular occurrence for cyber attacks to make the news. Naturally, this elevates the public concern over privacy and how personal information is used once collected. Public trust in the Australian Bureau of Statistics (ABS) to protect the data it collects from providers, is a cornerstone to the ABS mission. The US Census Bureau (USCB) recently announced – via its Scientific Advisory Committee – that it would protect the publications of the 2018 End-to-End Census Test (E2E) using differential privacy (DP). The E2E test is a dress rehearsal for the 2020 Census' (Abowd, 2018, p.2). In light of this announcement, many National Statistical Offices (NSOs), including the Office for National Statistics and Statistics New Zealand, are investigating DP approaches to protecting their Census outputs.

The ABS has been exploring the possibility of adopting a DP approach to improve existing confidentiality methodologies, particularly those that apply to the ABS TableBuilder. This paper contributes to the current discussion on DP by providing the current ABS perspective.

1

Source: `http://www.unece.org/fileadmin/DAM/stats/documents/ece/ces/ge.46/2019/mtg1/SDC2019_S2_ABS_Bailie_D.pdf`

The crux of the problem

How to educate researchers, the public and politicians about such tradeoff?

Managing the Tradeoff

United States Census Bureau

U.S. Department of Commerce
Economics and Statistics Administration
U.S. CENSUS BUREAU
census.gov

Noise adding mechanism for counting queries

Basic notation

- x is a database, counting query $q(x)$ returns true count n:

$$\underbrace{M(x)}_{\text{mechanism noisy output}} = \underbrace{n}_{q(x) \text{ is true count}} + \underbrace{Y}_{\text{noise}}$$

- Data-independent noise added to the true count.

$$\Pr(M(x) = i|n) = \Pr(Y = i - n|n) \underbrace{=}_{\text{independence}} \Pr(Y = i - n)$$

Differential privacy for counting queries

- For every two neighbouring datasets $x \sim x'$ differing on the attribute of one individual (e.g., smokes or does not smoke), the output of the mechanism should be almost indistinguishable.
- ϵ-DP:

$$e^{-\epsilon} \leq \frac{\Pr(M(x) = i)}{\Pr(M(x') = i)} = \frac{\Pr(M = i|n)}{\Pr(M = i|n+1)} \leq e^{\epsilon}, \forall i, n.$$

(ϵ, δ)-differential privacy

Formal definition

- A randomized mechanism is $M : \mathcal{X}^n \to \mathcal{Y}$ (n elements in dataset).
- If for all neighboring datasets $x \sim x' \in \mathcal{X}^n$ and all events $E \subset \mathcal{Y}$, we have

  $$\mathbb{P}[M(x) \in E] \leq e^{\varepsilon}\mathbb{P}[M(x') \in E] + \delta,$$

  then we say M satisfies (ϵ, δ)-DP.
- Pure differential privacy: $(\epsilon, 0)$-DP.
- Approximate differential privacy: (ϵ, δ)-DP, $0 < \delta < 1$.

Laplace mechanism

Basic notation

- Assume two neighboring datasets $x \sim x'$ differing on one element.
- Query function q is called; Laplace noise $Y \leftarrow \mathcal{L}(0, \lambda)$ is added to q.
- $q(x) = 0 \qquad \Rightarrow \qquad M(x) = 0 + Y.$
- $q(x') = \Delta \qquad \Rightarrow \qquad M(x') = \Delta + Y.$

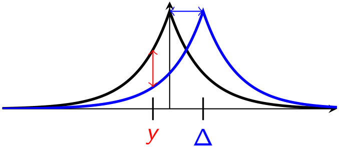

- The worst-case ratio of the two probability density functions will determine the differential privacy (DP) performance.
- $\lambda = \frac{\Delta}{\varepsilon} \qquad \Rightarrow \qquad \varepsilon$-DP.

Gaussian mechanism

Basic notation

- $M(x) = q(x) + Y, \qquad Y \leftarrow \mathcal{N}(0, \sigma_{\text{gauss}}^2).$

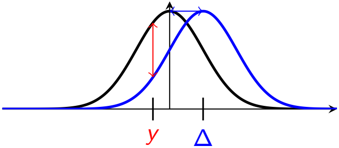

- There is no finite ε for which the ratio of the two probability density functions is bounded by e^{ε} $\Rightarrow$ We need to use either (ε, δ)-DP or Rényi DP frameworks. In other words:

$$E = \{y : \frac{f(y)}{f(y-\Delta)} > e^{\varepsilon}\} \neq \emptyset$$

Pros and cons of Laplace DP mechanism

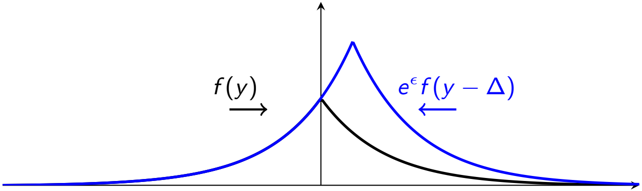

- Pro: Can achieve pure differential privacy: $(\epsilon, 0)$-DP:
  The blue curve never goes below the black curve.
- Easy-to-design: e.g., $\varepsilon = 0.5$, $\Delta = 1$, Laplace scale $\lambda = \frac{\Delta}{\varepsilon} = 2$.
- Utility optimal in some scenarios.
- Con: Heavy tail → outliers and bias→ challenges in post processing.

Pros and cons of Gaussian DP mechanism

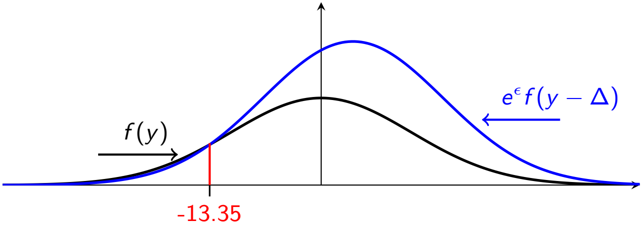

- Pro: Sub-Gaussian tail → less outliers or bias after post-processing.
- Con: Only approximate (ε, δ)-DP is possible → δ is a measure of impossibility of pure ε-DP.
- For $\epsilon = 0.5, \Delta = 1$ and $\sigma^2_{\text{gauss}} \approx 27.7$, values $y < -13.35$ contribute to $\delta \approx 3.2 \times 10^{-4}$ (figure not to scale).

Computation of δ

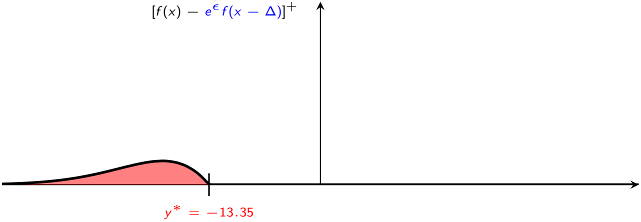

- Critical value of y, below which all values contribute to δ:
  $y^{*} = \frac{\Delta}{2} - \frac{\epsilon\sigma^2_{\text{gauss}}}{\Delta} = \frac{1}{2} - \frac{0.5*27.7}{1} = -13.35$
- δ is given by the area under the curve:

$$\delta(\epsilon) = \int_{-\infty}^{y^{*}} (f(x) - e^{\epsilon} f(x - \Delta))dx,$$

  where f denotes the pdf of the Gaussian $\mathcal{N}(0, \sigma^2_{\text{gauss}})$.
- Figure not to scale.

Motivations

- Discrete Gaussian is used for the release of 2020 US Census.
- One main reason for switching from Laplace to Gaussian was the "heavy tail" problem of Laplace leading to outliers and bias during processing.
- δ should be much smaller than $1/n$, where n is the size of population.

Motivations

- Discrete Gaussian is used for the release of 2020 US Census.
- One main reason for switching from Laplace to Gaussian was the "heavy tail" problem of Laplace leading to outliers and bias during processing.
- δ should be much smaller than $1/n$, where n is the size of population.

Research Question: Can we keep light tail of Gaussian but improve on its DP-variance tuple (joint) performance?

Motivations

- Discrete Gaussian is used for the release of 2020 US Census.
- One main reason for switching from Laplace to Gaussian was the “heavy tail” problem of Laplace leading to outliers and bias during processing.
- δ should be much smaller than $1/n$, where n is the size of population.

Research Question: Can we keep light tail of Gaussian but improve on its DP-variance tuple (joint) performance?

The answer is YES!

Intuition behind the proposed distribution - 1

Definition and properties

$$f(y) = \frac{1}{\underbrace{S'}_{\text{normalization}}} \quad \underbrace{e^{-\frac{y^2}{2\sigma^2}}}_{\text{Gaussian}} \quad \underbrace{e^{-\frac{|y|m}{\sigma^2}}}_{\text{Laplace factor } \mathcal{L}(0,\, \frac{m}{\sigma^2})} \tag{1}$$

- Point-wise product of a Gaussian pdf and a Laplace pdf.
- Original idea: moderate or slow down the Gaussian decay to reduce δ through factoring in a Laplace pdf.
- Two parameters in pdf: σ^2 and m (more DoF in privacy-utility tradeoff).
- Can be interpreted in different ways as we will see next (including why it is called Offset Symmetric Gaussian Tails (OSGT)).

Alternative Description of pdf: The Reason behind OSGT Naming - 1

Towards the OSGT pdf

Our original definition can be rewritten as:

$$f(y) = \frac{1}{S'} \quad e^{-\frac{y^2}{2\sigma^2} - \frac{|y|m}{\sigma^2}} \qquad \overset{\text{Expand } |y| \text{ cases}}{\Longrightarrow}$$

$$= \begin{cases} \frac{1}{S'} e^{-\frac{y^2}{2\sigma^2} + \frac{ym}{\sigma^2}}, & y \leq 0, \\ \frac{1}{S'} e^{-\frac{y^2}{2\sigma^2} - \frac{ym}{\sigma^2}}, & y > 0, \end{cases} \qquad \overset{\text{Multiply by 1 or by } e^{+\frac{m^2}{2\sigma^2}} e^{-\frac{m^2}{2\sigma^2}}}{\Longrightarrow}$$

$$= \begin{cases} \frac{e^{+\frac{m^2}{2\sigma^2}}}{S'} e^{-\frac{(y-m)^2}{2\sigma^2}}, & y \leq 0, \\ \underbrace{\frac{e^{+\frac{m^2}{2\sigma^2}}}{S'}}_{\text{New constant } 1/S} \quad e^{-\frac{(y+m)^2}{2\sigma^2}}, & y > 0, \end{cases}$$

Alternative Description of pdf: The Reason behind OSGT Naming - 2

OSGT pdf

$$f(y) = \begin{cases} \frac{1}{S} e^{-\frac{(y-m)^2}{2\sigma^2}}, & y \leq 0, \\ \frac{1}{S} e^{-\frac{(y+m)^2}{2\sigma^2}}, & y > 0, \end{cases}$$

Normalization $S = \sqrt{2\pi\sigma^2} 2Q(m/\sigma)$ takes care of the area under the Gaussian tails. Note the resulting distribution has zero mean.

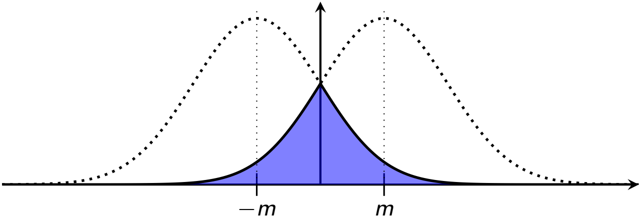

The Gaussian Q-function:

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-t^2/2} dt$$

Sampling: rejecting the non-tail values of $\mathcal{N}(0, \sigma^2)$

Rejection, followed by a shift

- This view will be useful in sampling from the Gaussian, rejecting values between $[-m, m)$, and shifting the acceptable tail values to obtain a sample from the OSGT.

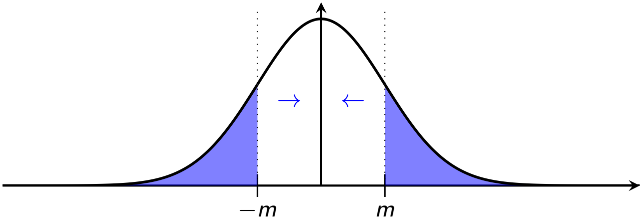

Intuition behind the proposed distribution - 2

Approximation at small $|y| \approx$ scaled Laplace

$$f(y) = \frac{1}{\underbrace{S'}_{\text{normalization}}} \underbrace{e^{-\frac{y^2}{2\sigma^2}}}_{\text{Gaussian}} \underbrace{e^{-\frac{|y|m}{\sigma^2}}}_{\text{Laplace}}$$

$$= \frac{1}{S'} e^{-\frac{|y|(|y|+2m)}{2\sigma^2}}$$

$$\approx \frac{1}{S'} e^{-\frac{|y|(\cancel{|y|}+2m)}{2\sigma^2}}, \qquad |y| \ll 2m$$

$$= \frac{1}{S'} e^{-\frac{m|y|}{\sigma^2}}, \qquad |y| \ll 2m.$$

The shape of the proposed pdf

- All three distributions are compared at the same variance.
- The “Laplace-like” behaviour of the new distribution at small $|y|$ can be observed.

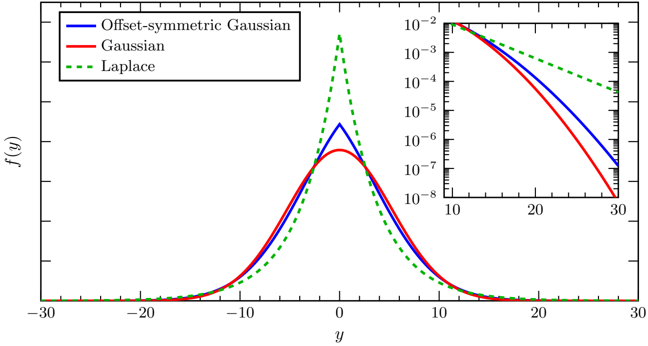

Figure: $\sigma^2 = 40, m = 3,$ but **actual variance is** $V \approx 27.7$.

Variance of OSGT random variable

Variance of OSGT is smaller than its input parameter σ^2

Let $Y \leftarrow \mathcal{T}(0, V(m, \sigma^2))$ be a zero-mean OSGT distributed random variable. Its variance is given by

$$\mathbb{E}[Y^2] = V(m, \sigma^2) = \sigma^2 + m^2 - \frac{m\sigma \exp\left(-\frac{m^2}{2\sigma^2}\right)}{\sqrt{2\pi}Q\left(\frac{m}{\sigma}\right)} < \sigma^2. \qquad (2)$$

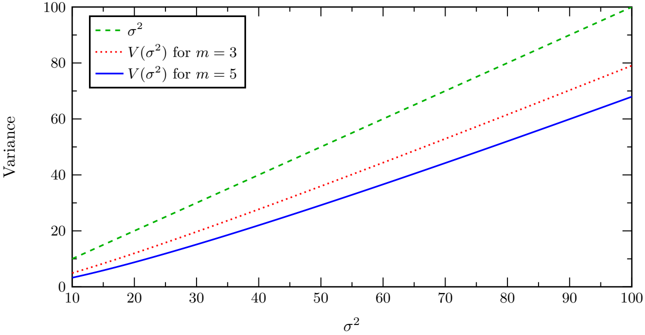

Sub-Gaussian tail

Proposition

- The OSGT distribution $\mathcal{T}(0, V(m, \sigma^2))$ is σ^2-sub-Gaussian if

$$\frac{m}{\sigma} \leq Q^{-1}(0.25) \approx 0.6745.$$

- In this case, we have:

$$\underset{Y \leftarrow \mathcal{T}(0, V(m,\sigma^2))}{\mathbb{P}} [Y \geq y] \leq \exp\left(-\frac{y^2}{2\sigma^2}\right).$$

Computing δ for the OSGT mechanism - 1

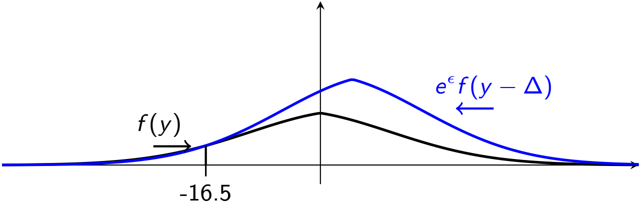

- At $\epsilon = 0.5, \Delta = 1$, $m = 3$, and $\sigma^2 \approx 40$, values $y < -16.5$ contribute to $\delta \approx 6.8 \times 10^{-5}$.
- The actual variance is $V = 27.7$.
- Recall, for the same variance $\sigma^2_{\text{gauss}} = 27.7$, same $\epsilon = 0.5$ and $\Delta = 1$, the Gaussian gives $y < -13.35$ and $\delta \approx 3.2 \times 10^{-4}$.

Numerical evaluation of (ε, δ) for OSGT versus Gaussian

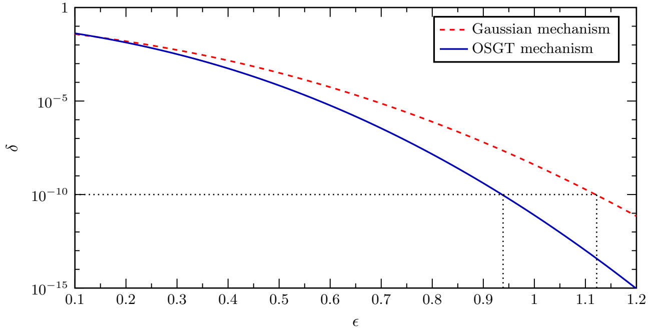

Figure: $\Delta = 1,\ m = 3,\ \sigma^2 = 40,\ V(m, \sigma^2) = \sigma^2_{\text{gauss}} \approx 27.7.$

- The OSGT mechanism achieves $(\varepsilon, \delta) \approx (0.94, 10^{-10})$.
- The Gaussian mechanism achieves $(\varepsilon, \delta) \approx (1.12, 10^{-10})$.

Numerical evaluation of (ε, δ) for OSGT versus Gaussian - 2

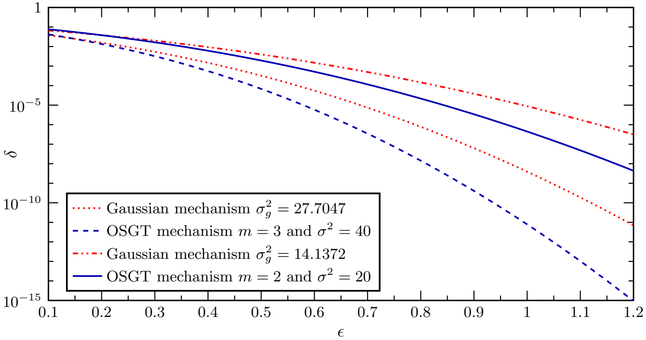

Figure: (ε, δ)-DP performance of OSGT and Gaussian mechanisms at the same variance.

- The OSGT mechanism achieves better (ε, δ) at the same variance.

Numerical evaluation of (ε, δ) for OSGT versus Gaussian - 3

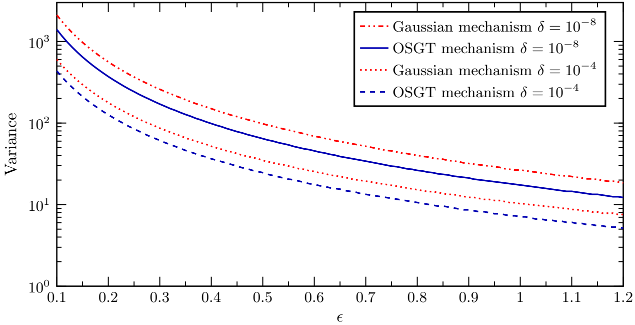

Figure: Variance performance of OSGT and Gaussian mechanisms at the same (ε, δ)-DP level.

- The OSGT mechanism needs smaller variance (adds less noise; proxy for utility) at the same (ε, δ).

Computing δ for the OSGT Mechanism

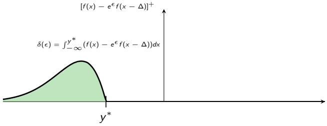

Analytical formula for δ, single-variate OSGT mechanism (1/2 cases shown)

$$\delta^{\mathcal{T}}(\varepsilon) = \frac{1}{2Q\left(\frac{m}{\sigma}\right)} \left(\underbrace{Q\left(\frac{\sigma\varepsilon}{\Delta} - \frac{\Delta}{2\sigma}\right) - e^{\varepsilon} Q\left(\frac{\sigma\varepsilon}{\Delta} + \frac{\Delta}{2\sigma}\right)}_{\text{decreasing function of } \sigma} \right), \quad \varepsilon > \frac{\Delta^2}{2\sigma^2} + \frac{\Delta m}{\sigma^2}.$$

Computing δ for the OSGT Mechanism

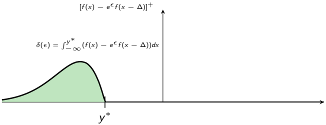

Analytical formula for δ, single-variate OSGT mechanism (1/2 cases shown)

$$\delta^{\mathcal{T}}(\varepsilon) = \frac{1}{2Q\left(\frac{m}{\sigma}\right)} \left(\underbrace{Q\left(\frac{\sigma\varepsilon}{\Delta} - \frac{\Delta}{2\sigma}\right) - e^{\varepsilon} Q\left(\frac{\sigma\varepsilon}{\Delta} + \frac{\Delta}{2\sigma}\right)}_{\text{decreasing function of } \sigma} \right), \quad \varepsilon > \frac{\Delta^2}{2\sigma^2} + \frac{\Delta m}{\sigma^2}.$$

Recall that for a fair utility comparison: $\sigma_{\text{gauss}}^2 = V(m, \sigma^2) < \sigma^2$.

Comparison with δ of the Gaussian mechanism

$$\delta^{\mathcal{N}}(\varepsilon) = Q\left(\frac{\sigma_{\text{gauss}}\varepsilon}{\Delta} - \frac{\Delta}{2\sigma_{\text{gauss}}}\right) - e^{\varepsilon} Q\left(\frac{\sigma_{\text{gauss}}\varepsilon}{\Delta} + \frac{\Delta}{2\sigma_{\text{gauss}}}\right), \quad \varepsilon \geq \frac{\Delta^2}{2\sigma_{\text{gauss}}^2}.$$

Multidimensional OSGT mechanism

Definition

- The k-dimensional query function $q : \mathcal{X}^n \to \mathbb{R}^k$.
- The OSGT mechanism is obtained as $M(x) = q(x) + Y$.
- Each noise component Y_i in $Y = (Y_1, \cdots, Y_k)$ is drawn from an i.i.d. scalar OSGT distribution $\mathcal{T}(0, V(m, \sigma^2))$, leading to:

$$f^{\mathcal{T}}_{M(x)}(y) = \frac{1}{S'^k} \exp\left(\underbrace{-\frac{\|y - q\|_2^2}{2\sigma^2}}_{k\text{-dim Gaussian}} \underbrace{-\frac{m\,\|y - q\|_1}{\sigma^2}}_{k\text{-dim Laplace}} \right).$$

Multidimensional OSGT mechanism

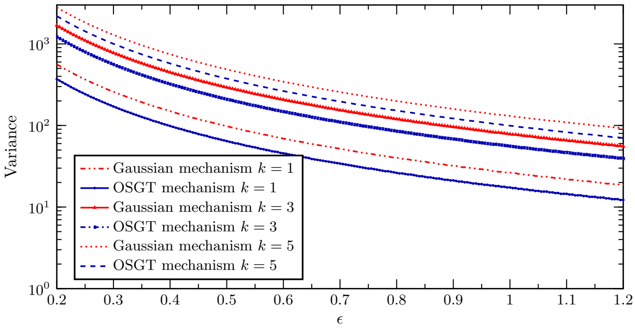

Figure: Variance performance of OSGT and Gaussian mechanisms at the same (ε, δ)-.

- The OSGT mechanism needs smaller variance (adds less noise; proxy for utility) at the same better (ε, δ) for k-dimensional query.

- However, the improvement seems to diminish for large k.

- The method for computation of δ is numerical.

Rényi DP (RDP) of the OSGT mechanism

Single-letter RDP upper bound

The Rényi DP of the OSGT mechanism is upper bounded by

$$\underbrace{D_\alpha(M(x)\|M(x'))}_{\text{Rényi DP}} \le \alpha\frac{\Delta_2^2}{2\sigma^2} + \zeta, \qquad \zeta = \frac{k}{\alpha-1}\log\left(\frac{1-Q(\frac{m}{\sigma})}{Q(\frac{m}{\sigma})}\right).$$

where $D_\alpha(M(x)\|M(x'))$ is the Rényi divergence of order α.

Gaussian achieves RDP as $\alpha\frac{\Delta_2^2}{2\sigma_{\text{gauss}}^2} > \alpha\frac{\Delta_2^2}{2\sigma^2}$.

Exact computation of Rényi divergence for OSGT mechanism - 1

Key equations

- The Rényi Divergence for k-dimensional OSGT mechanism can be analytically written as

$$D_\alpha(M(x)\|M(x')) \leq \frac{\alpha\Delta_2^2}{2\sigma^2} + \frac{k}{\alpha-1}\log\left(\frac{\sqrt{2\pi\sigma^2}}{S}\overline{B}\right)$$

- where:

$$\overline{B} = \Phi((-m+(\alpha-1)\Delta)/\sigma) + \Phi((-m-\alpha\Delta)/\sigma) +$$
$$e^{\alpha(\alpha-1)(4m\Delta+4m^2)/2\sigma^2} \times \left(\Phi\left(\frac{\alpha\Delta - m(1-2\alpha)}{\sigma}\right) - \Phi\left(\frac{(\alpha-1)\Delta - m(1-2\alpha)}{\sigma}\right)\right).$$

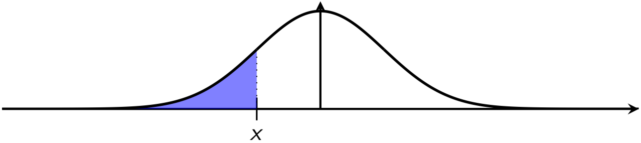

Φ is the Gaussian cumulative distribution function (cdf):

$$\Phi(x) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{x} e^{-t^2/2}dt = 1 - Q(x).$$

Components of Rényi divergence for OSGT mechanism

Key equation

$$D_\alpha(M(x)\|M(x')) \leq \underbrace{\alpha\frac{\Delta_2^2}{2\sigma^2}}_{\alpha\rho} + \underbrace{\frac{k}{\alpha-1}\log\left(\frac{\sqrt{2\pi\sigma^2}}{S}\overline{B}\right)}_{\zeta}$$

We numerically evaluate the scale of ζ and $\alpha\rho$ as a function of α.

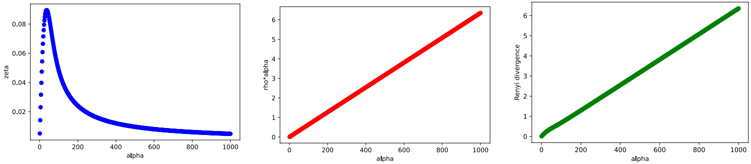

Figure: $k = 8, m = 15, \sigma^2 = 630, \Delta = 1, \Delta_2^2 = k\Delta = 8$. It can be observed the effect of ζ in the overall Rényi divergence is rather small.

Rényi DP $\to \delta$

Analytical computation of Rényi DP to get a bound on δ

$$D_\alpha(M(x)\|M(x')) \leq \tau \qquad \Rightarrow \qquad \delta(\varepsilon) = \frac{\exp((\alpha-1)(\tau-\varepsilon))}{\alpha-1}\left(1-\frac{1}{\alpha}\right)^{\alpha}$$

$$\tau = \frac{\alpha\Delta_2^2}{2\sigma^2} + \frac{k}{\alpha-1}\log\left(\frac{\sqrt{2\pi\sigma^2}}{S}\overline{B}\right).$$

For more details and applications (such as composition and noise filtering), please see our paper published in 2022 in IEEE Trans. Information Forensics and Security (TIFS) with the same title.