Privacy Loss Classes: The Central Limit Theorem in Differential Privacy

David Sommer, Sebastian Meiser, Esfandiar Mohammadi

Privacy: The Question of Who

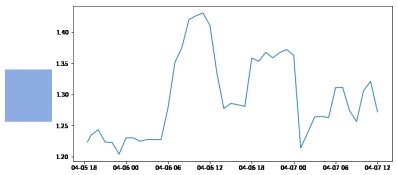

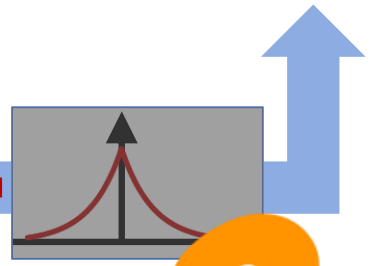

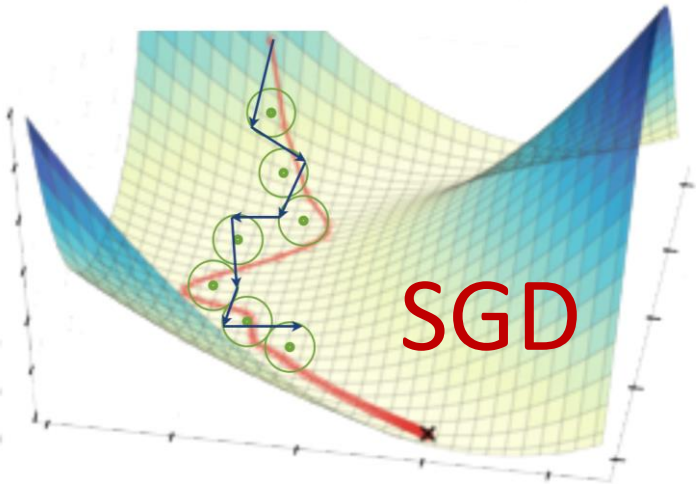

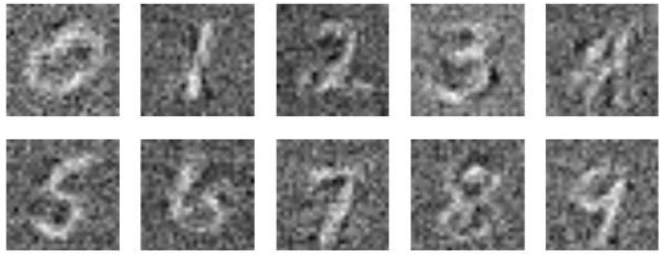

3. Synthesized MNIST images from the RON-Gauss model out the DP component using half of the full dimensions.

(ε, δ)-Approximate DP

(ε, δ)-Probabilistic DP

Rényi-Differential Privacy

Concentrated-DP

ε-Differential Privacy

Our Contribution / Outline

- Connect Differential Privacy Definitions
- Privacy Leakage under Independent Composition
  - Exact Formula for Gaussian Noise
- Comparison of mechanisms (Privacy Loss Classes)
  - e.g., Gaussian better than Laplacian Noise

Basics: Pure ε-Differential Privacy (ε- DP)

Definition (informal)

$$\Pr\left[M(D_0) \in S\right] \leq e^{\varepsilon} \cdot \Pr\left[M(D_1) \in S\right]$$

- View on individual atomic events $o \in [M]$

$$\Pr\left[M(D_0) \in S\right] = \sum_{o \in S} \Pr\left[M(D_0) = o\right]$$

$$\log \frac{\Pr\left[M(D_0) \in S\right]}{\Pr\left[M(D_1) \in S\right]} \leq \varepsilon$$

Privacy Loss Variable

Definition (informal)

for $o \in [M]$, if both probabilities are non-zero

$$\mathcal{L}_{M(D_0),M(D_1)}(o) = \log \frac{\Pr[M(D_0) = o]}{\Pr[M(D_1) = o]}$$

If only $\Pr[M(D_1) = o] = 0$, then

$$\mathcal{L}_{M(D_0),M(D_1)}(o) = +\infty$$

Else, $\mathcal{L}_{M(D_0),M(D_1)}(o) = -\infty$

- For specific inputs D_0, D_1
- Example: Additive noise

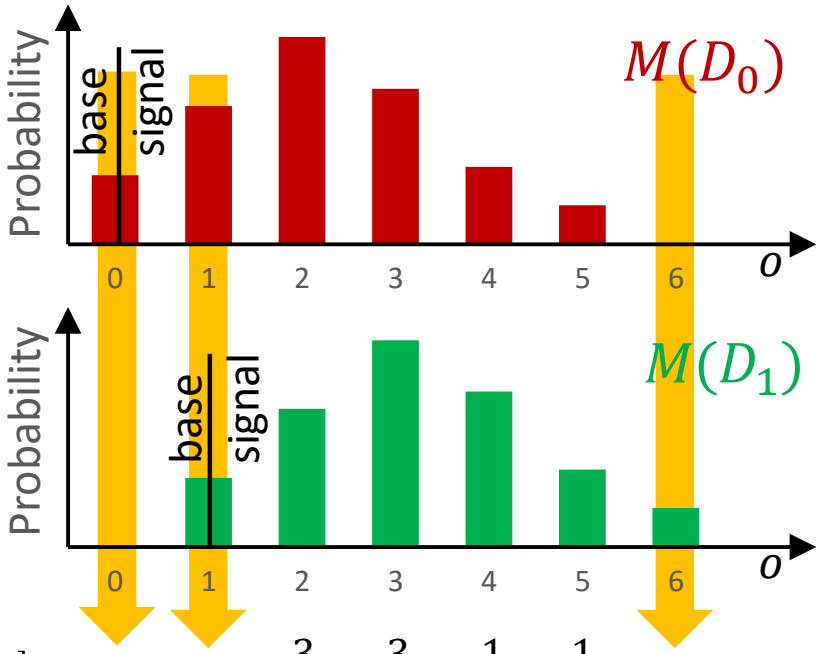

$\frac{\Pr[M(D_0) = o]}{\Pr[M(D_1) = o]}$:	$+\infty$	2	$\frac{3}{2}$	$\frac{3}{4}$	$\frac{1}{2}$	$\frac{1}{2}$	0
$\mathcal{L}_{M(D_0),M(D_1)}(o)$:	$+\infty$	0.69	0.41	-0.29	-0.69	-0.69	$-\infty$

Privacy Loss Distribution (PLD)

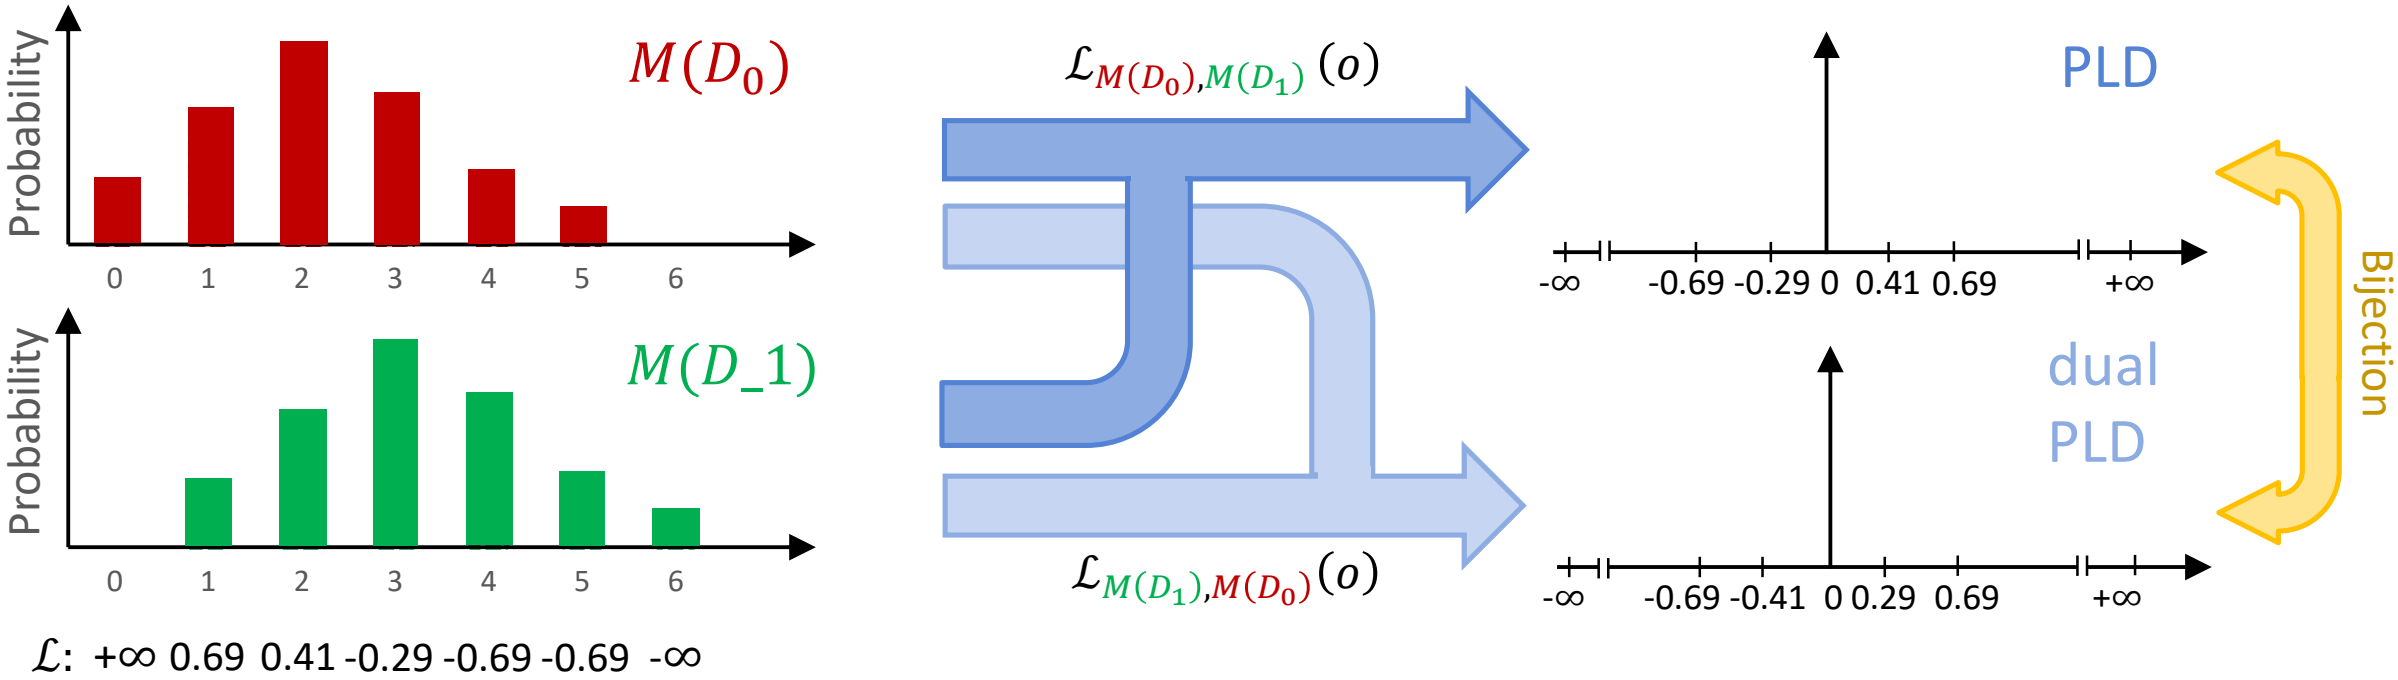

$\mathcal{L}$: +∞ 0.69 0.41 -0.29 -0.69 -0.69 -∞

- Different input-pairs D_0, D_1 can lead to different PLDs
  - In practice, often one pair of PLDs suffices

(ε, δ)-Probabilistic Differential Privacy (PDP)

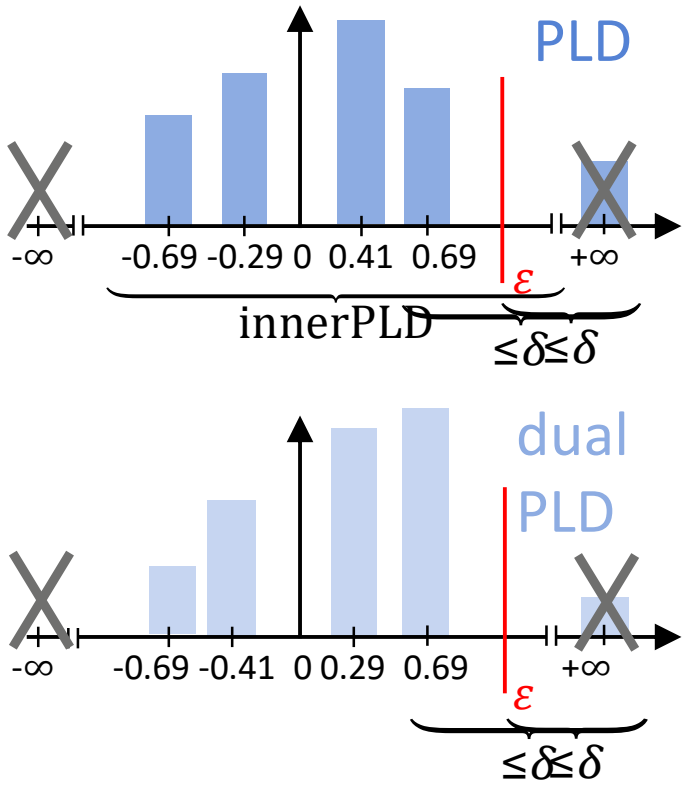

ε-Differential Privacy

Definition (informal)

With $\varepsilon, \delta \geq 0$, M is $(\boldsymbol{\varepsilon}, \boldsymbol{\delta})$-**PDP** if

$$\Pr\left[\left\{ o \mid \mathcal{L}_{M(D_0),M(D_1)}(o) > \varepsilon \right\} \in M(D_0)\right] \leq \delta$$

$$\Pr\left[\left\{ o \mid \mathcal{L}_{M(D_1),M(D_0)}(o) > \varepsilon \right\} \in M(D_1)\right] \leq \delta$$

(ε, δ)-Approximate Differential Privacy (ADP)

Definition

A mechanism M [illegible] with $D_0 \approx^1 D_1$, and for all sets S [illegible] tion holds:

ADP: same as ε-DP, plus δ

$$\Pr[M(D_0) \in S] \leq e^{\varepsilon} \cdot \Pr[M(D_1) \in S] + \delta$$

$$\delta(\varepsilon) = \mathrm{PLD}(+\infty) + \sum_{y \geq \varepsilon} \left(1 - e^{\varepsilon - y}\right) \mathrm{PLD}(y)$$

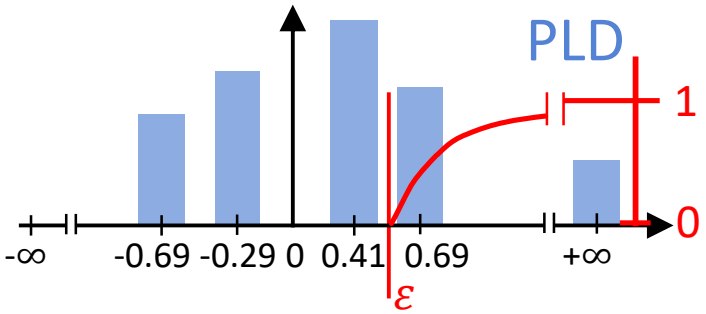

Equivalence

[if finite number of atomic events]

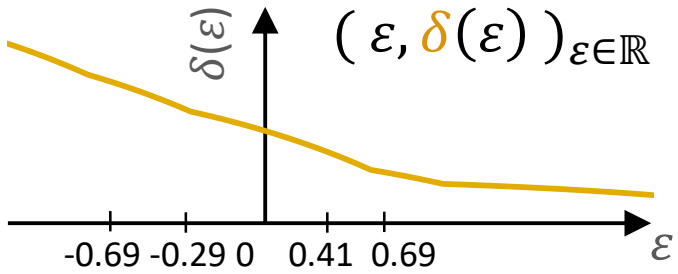

(α, ε)-Rényi Differential Privacy (RDP)

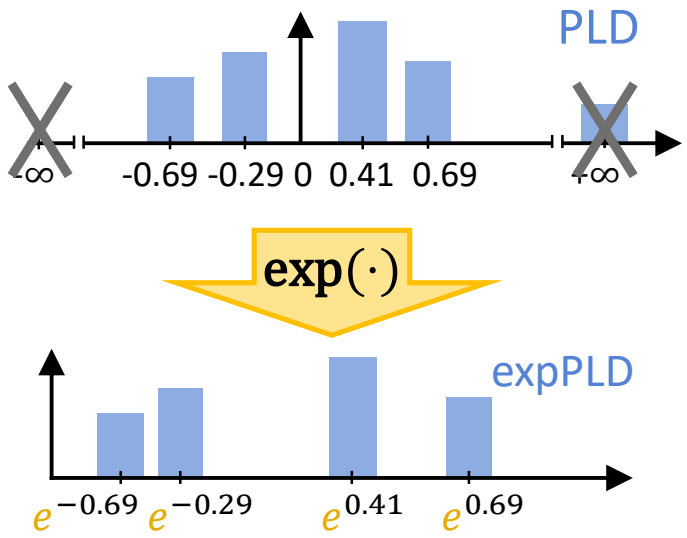

$$\log\left(\mathop{\mathbb{E}}_{y\sim PLD}\left[(e^{y})^{\lambda}\right]\right)^{\frac{1}{\lambda}} = \frac{1}{\lambda}\ln \mathop{\mathbb{E}}_{o\sim M(D_1)}\left[\left(\frac{\Pr\left[M(D_0)=o\right]}{\Pr\left[M(D_1)=o\right]}\right)^{\lambda+1}\right]$$

$$= \mathcal{D}_{\underbrace{\lambda+1}_{\alpha}}\left(M(D_0)\middle|M(D_1)\right)$$

Definition (informal)

(α, ε)-RDP: $\forall\, D_0, D_1\colon\ \mathcal{D}_{\alpha}\left(M(D_0)\middle|M(D_1)\right) \leq \varepsilon$.

- PLD ⇒ RDP: $\left(\alpha,\ \mathcal{D}_{\alpha}\left(M(D_0)\middle|M(D_1)\right)\right)_{\alpha\in\mathbb{N}}$ ✓
- PLD ⇐ RDP: (see paper) (✓)

$RDP \iff \text{PLD} \iff ADP$

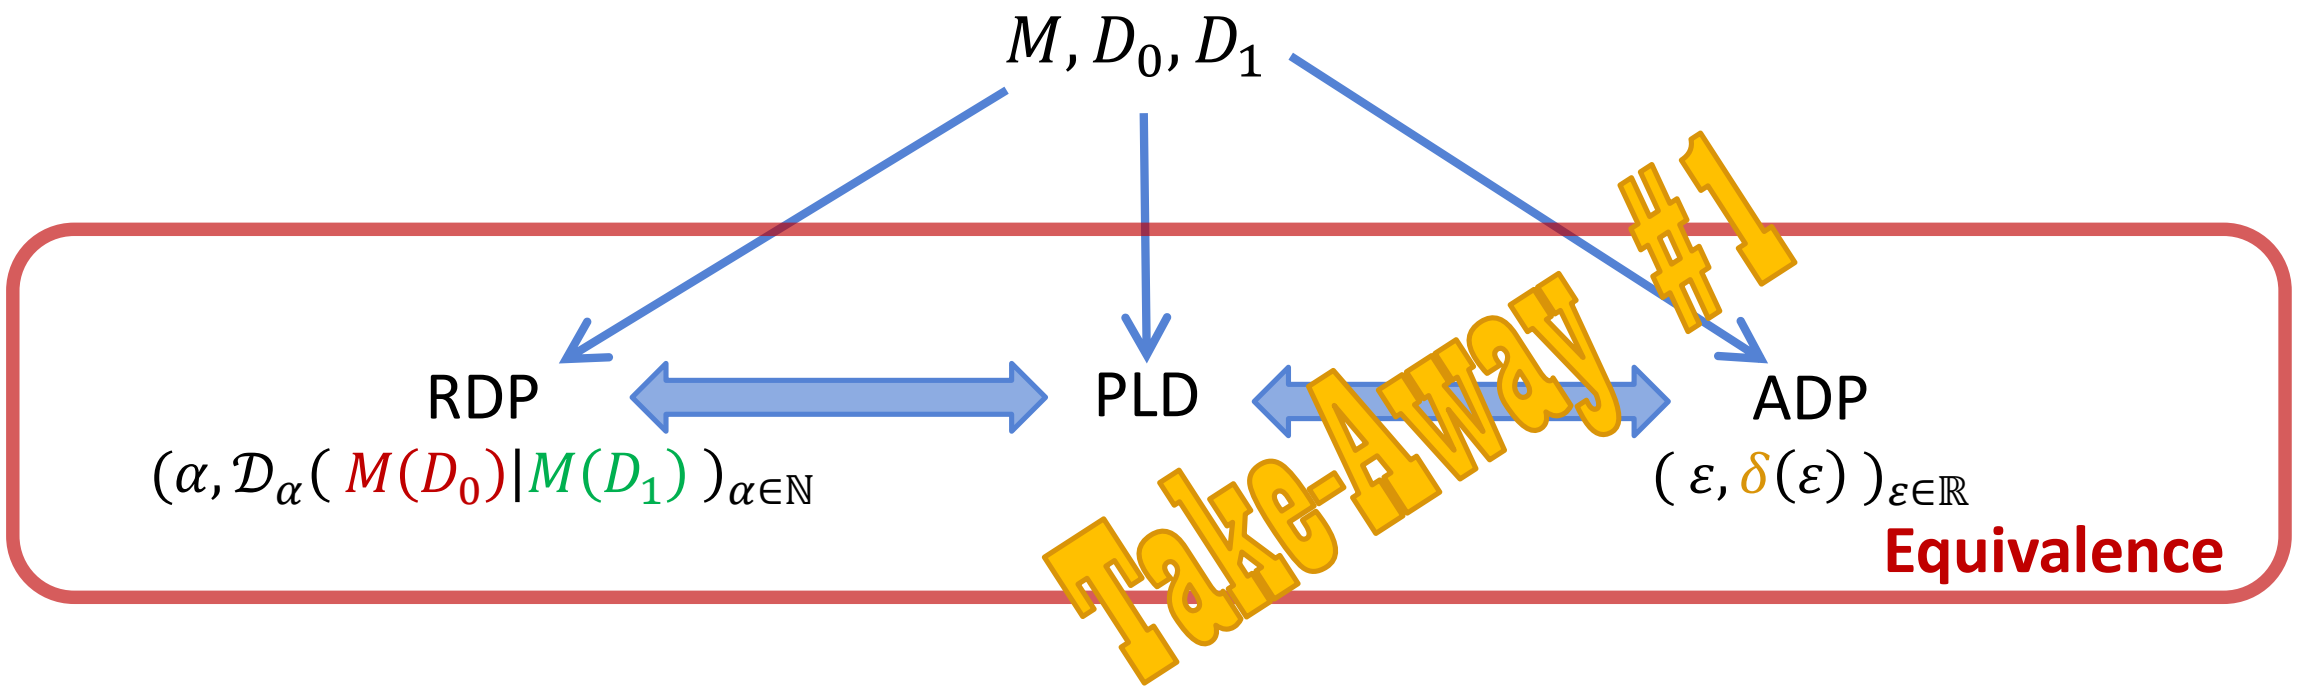

- [technical side conditions (see paper)]

Independent Composition

$$\mathcal{L}_{\left(M(D_0),\mathrm{M}'\left(\mathrm{D}_0'\right)\right),\left(M(D_1),\mathrm{M}'\left(\mathrm{D}_1'\right)\right)}(\mathrm{o},\mathrm{o}') = \ln\frac{\Pr\left[o\in M(D_0),o'\in M'\left(D_0'\right)\right]}{\Pr\left[o\in M(D_1),o'\in M'\left(D_1'\right)\right]}$$

$$= \ln\frac{\Pr[o\in M(D_0)]\cdot\Pr\left[o'\in M'\left(D_0'\right)\right]}{\Pr[o\in M(D_1)]\cdot\Pr\left[o'\in M'\left(D_1'\right)\right]}$$

$$= \ln\frac{\Pr[o\in M(D_0)]}{\Pr[o\in M(D_1)]} + \ln\frac{\Pr\left[o'\in M'\left(D_0'\right)\right]}{\Pr\left[o'\in M'\left(D_1'\right)\right]}$$

$$= \mathcal{L}_{M(D_0),M(D_1)}(\mathrm{o}) + \mathcal{L}_{\mathrm{M}'\left(\mathrm{D}_0'\right),\mathrm{M}'\left(\mathrm{D}_1'\right)}(\mathrm{o}')$$

- Composition is Convolution of **inner** PLDs

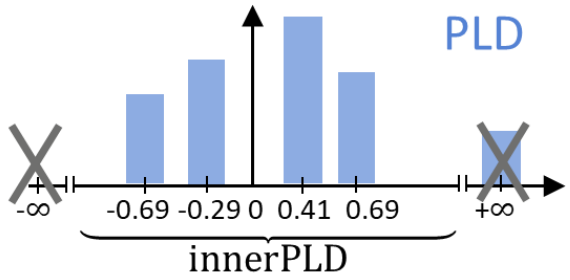

Privacy Loss Classes

- Convolution of **inner**PLD
  - ➡ Central Limit Theorem

μ = mean(**inner**PLD)

σ^2 = variance(**inner**PLD)

PLD(+∞)

n-invocations

$$\left.\begin{matrix}\mu \\ \sigma^2 \\ \text{PLD}(+\infty)\end{matrix}\right\} \quad \mathcal{N}(n\cdot\mu, n\cdot\sigma^2)$$

$$\text{PLD}^n(+\infty) = 1 - [1-\text{PLD}(+\infty)]^n$$

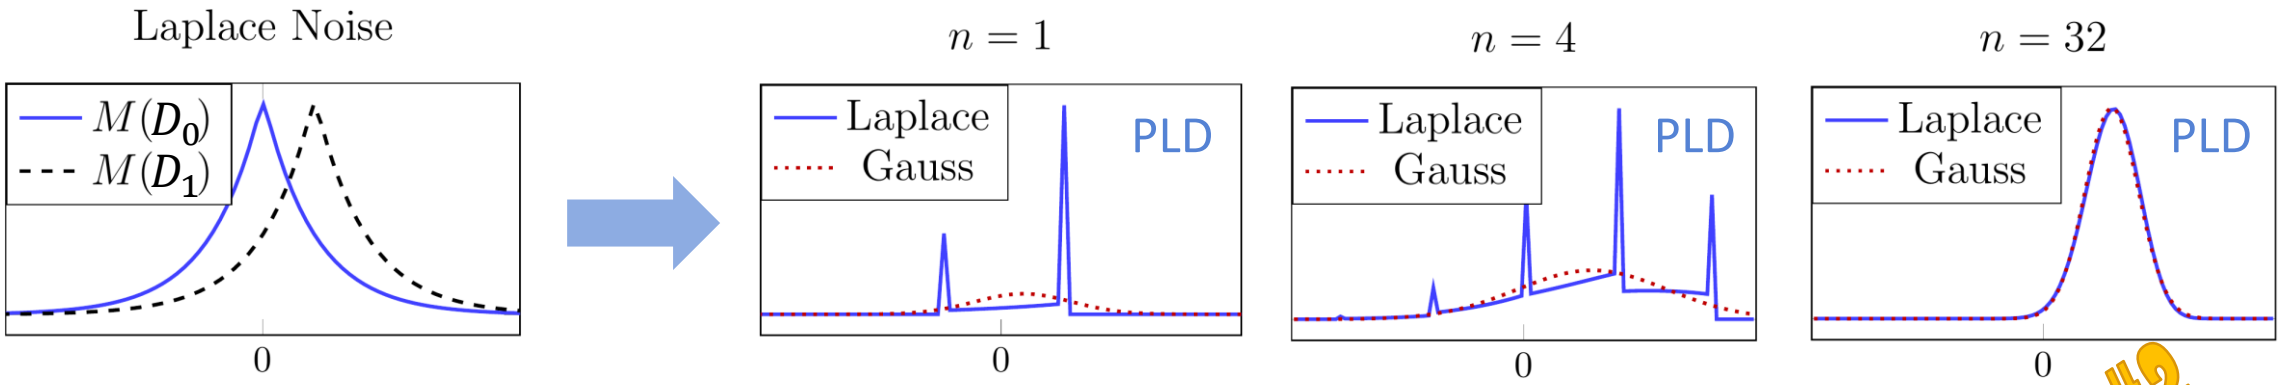

- Privacy Loss Class $\left(\mu, \sigma^2, \text{PLD}(+\infty)\right)$ characterises convergence
- Applicable to non-equal PLDs

Take-Away #2

TL;DR: Use Gaussian Noise!

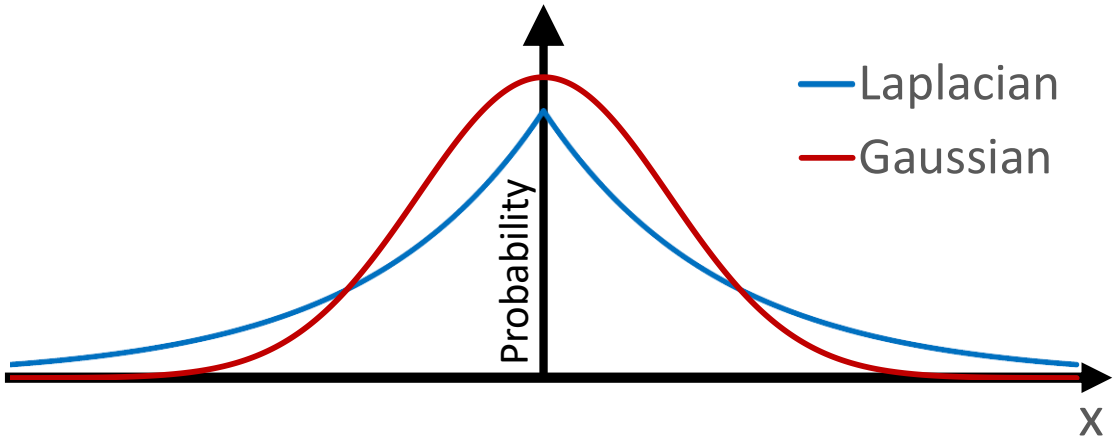

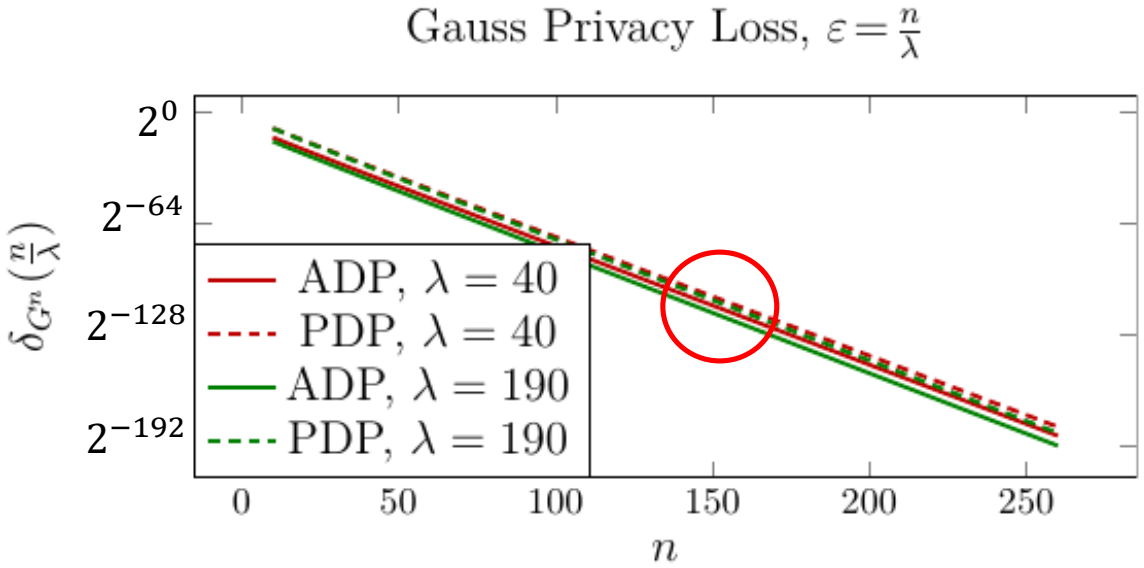

- Gaussian half the variance
- Similar Privacy Loss Class $(\mu, \sigma^2, \mathrm{PLD}(+\infty))$

  ⇒ higher utility than Laplacian at similar privacy guarantees

- After ~100 compositions: use Gaussian noise

Take-Away #3

Example: 1D-Gaussian Mechanism

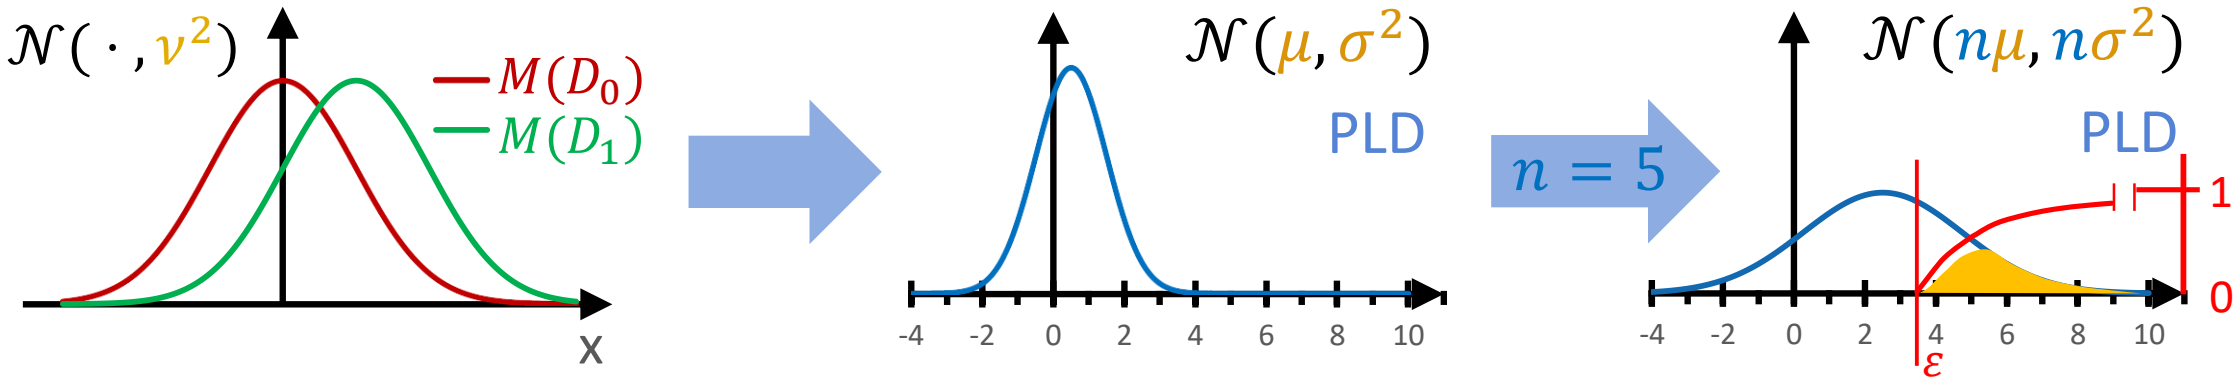

- Exact Formula:

$$\sigma^2 = \frac{1}{\nu^2}$$

$$\mu = \frac{\sigma^2}{2}$$

$$\delta_{ADP}(\varepsilon) = \frac{1}{2}\left[\operatorname{erfc}\left(\frac{\varepsilon - n\mu}{\sqrt{2n}\sigma}\right) - e^{\varepsilon} \cdot \operatorname{erfc}\left(\frac{\varepsilon + n\mu}{\sqrt{2n}\sigma}\right)\right]^{1}$$

Take-Away #4

$$\operatorname{erfc}(z) = \frac{2}{\sqrt{\pi}} \int_z^{\infty} e^{-t^2}\, dt$$

- efficiently computable
  - **gsl_sf_log_erfc** from *GNU Scientific Library*

[1] Balle at al. *Improving the Gaussian Mechanism for Differential Privacy: Analytical Calibration and Optimal Denoising*. ICML, 2018

Summary

- Introduced Privacy Loss Distribution (PLD)
- #1 (conditioned) Equivalence RDP, PLD, ADP
- #2 Composition for arbitary mechanisms
- #3 Gaussian is better than Laplace
- #4 CLT and the exact result for Gaussian
- ADP bounds based on CLT (see paper/ask)
- ADP Markov-bound (see paper)

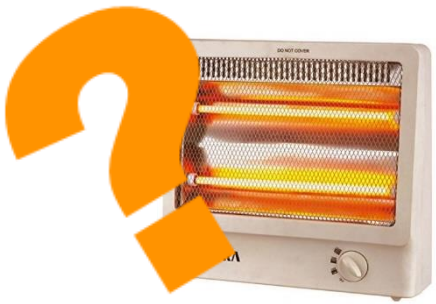

ADP under Composition

- Distance bounds $PLD^n \Leftrightarrow \mathcal{N}(n \cdot \mu, n \cdot \sigma^2)$
- Berry-Esseen (absolute error)
  - $|F_n(x) - \phi(x)| < C \cdot \frac{\rho}{\sigma^3 \sqrt{n}}$
  - Good for high number of compositions
- Nagaev (tail estimation)

Gaussian vs. Gaussian Mixture

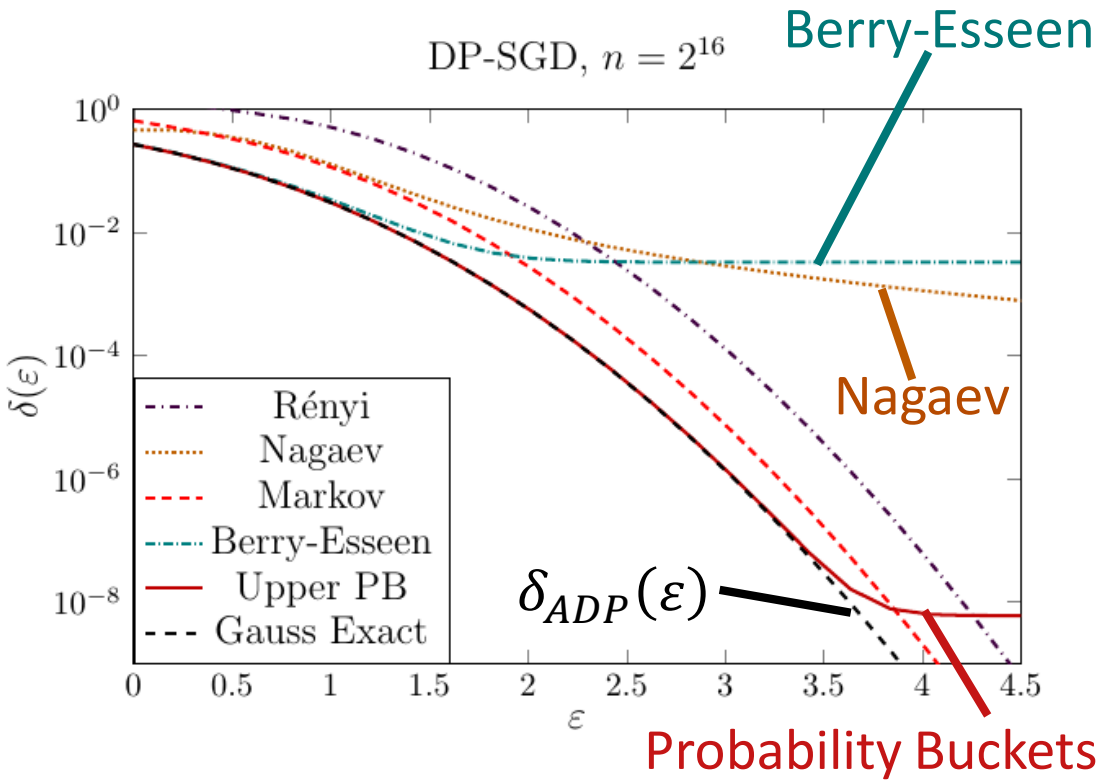

PB: Meiser S, Mohammadi E. *Tight on budget?: Tight bounds for r-fold approximate differential privacy*. CCS'18.

Backup Slides

Refernces Pictures

- https://www.worldtravelguide.net/wp-content/uploads/2018/06/shu-Europe-Sweden-Stockholm-613199033-Andrey-Shcherbukhin-2500x1045.jpg
- https://de.wiktionary.org/wiki/Datei:Snow_flake.svg
- https://i2.wp.com/www.bhakari.com/wp-content/uploads/2016/02/Electric-heater-500x500.jpg
- https://i1.wp.com/www.fixofix.se/wp-content/uploads/2017/11/Radiator_Standard-2.jpg?fit=500%2C500&ssl=1
- https://www.bathroomradiatorsuk.com/images/detailed/1/RA410.jpg
- https://www.britishgas.co.uk/digital/business/assets/smarter-working/smart-metering/how-your-electricity-smart-meter-works/_jcr_content/par/panel_container_0/par/bootstrap_accordion/par/image.img.jpg/1502201367102.jpg
- https://blog.upad.co.uk/hs-fs/hub/292131/file-3913227893-jpg/blog-files/istock_000008512727xsmall-213x300.jpg?width=213&height=300&name=istock_000008512727xsmall-213x300.jpg
- https://www.pinclipart.com/maxpin/iRJxobR/
- https://www.researchgate.net/publication/333000550/figure/fig1/AS:757018751156224@1557498743030/Virtual-memory-usage-graph.ppm
- https://vignette.wikia.nocookie.net/rickandmorty/images/4/41/Morty_Smith.jpg/revision/latest/scale-to-width-down/310?cb=20170217193441
- Chanyaswad T, Liu C, Mittal P. RON-Gauss: Enhancing utility in non-interactive private data release. Proceedings on Privacy Enhancing Technologies. 2019 Jan 1;2019(1):26-46.

Basics: Pure ε-Differential Privacy (ε- DP)

Definition

A mechanism M is ***ε-differentially private***, where $\varepsilon \geq 0$, if for all databases D_0 and D_1 with $D_0 \overset{1}{\approx} D_1$, and for all sets $S \subseteq [M]$, where $[M]$ is the range of M, the following equation holds:

$$\Pr[M(D_0) \in S] \leq e^{\varepsilon} \cdot \Pr[M(D_1) \in S]$$

- View on individual atomic events $o \in [M]$

$$\Pr[M(D_0) \in S] = \sum_{o \in S} \Pr[M(D_0) = o]$$

$$\log \frac{\Pr[M(D_0) \in S]}{\Pr[M(D_1) \in S]} \leq \varepsilon$$

Privacy Loss Distribution (PLD)

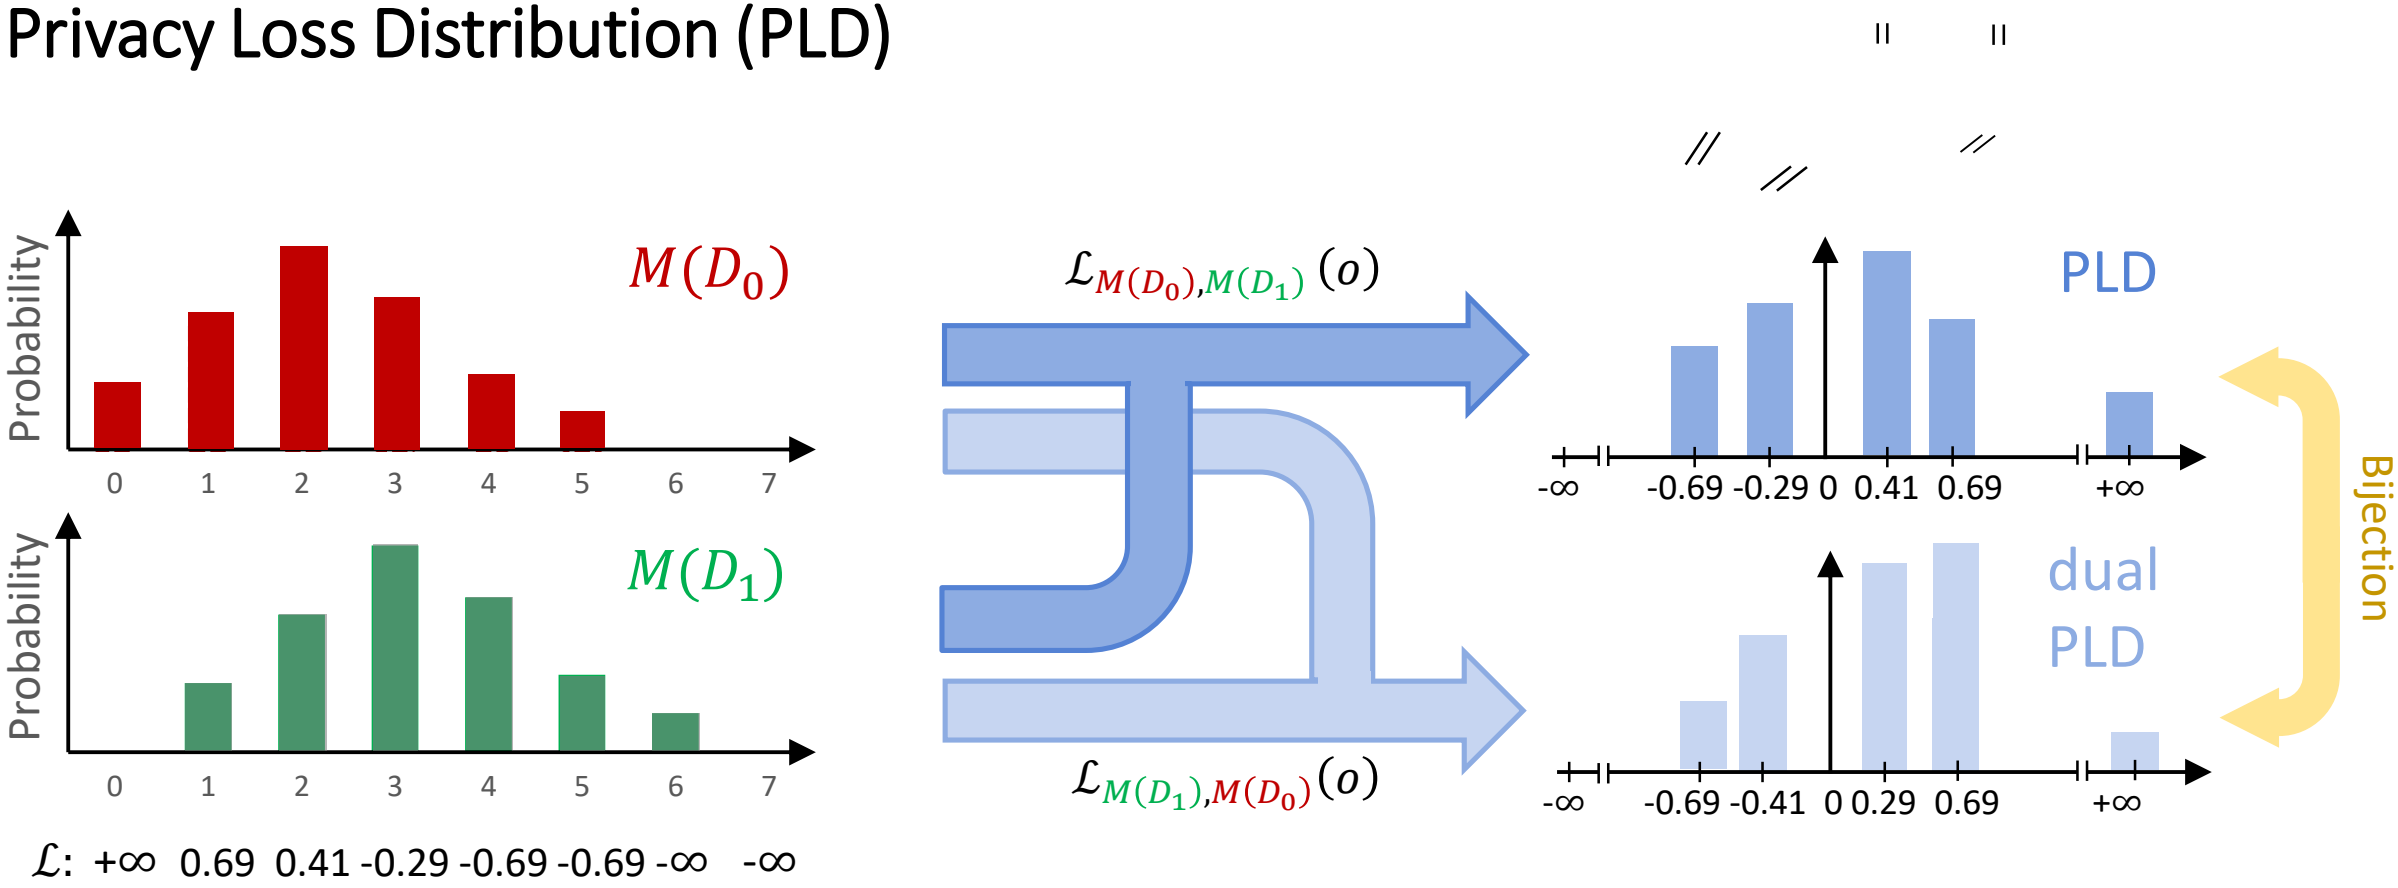

$\mathcal{L}$: +∞ 0.69 0.41 -0.29 -0.69 -0.69 -∞ -∞

- Different input-pairs D_0, D_1 can lead to different PLDs
  - In practice, often equal PLDs (sensitivity)

(α, ε)-Rényi Differential Privacy (RDP)

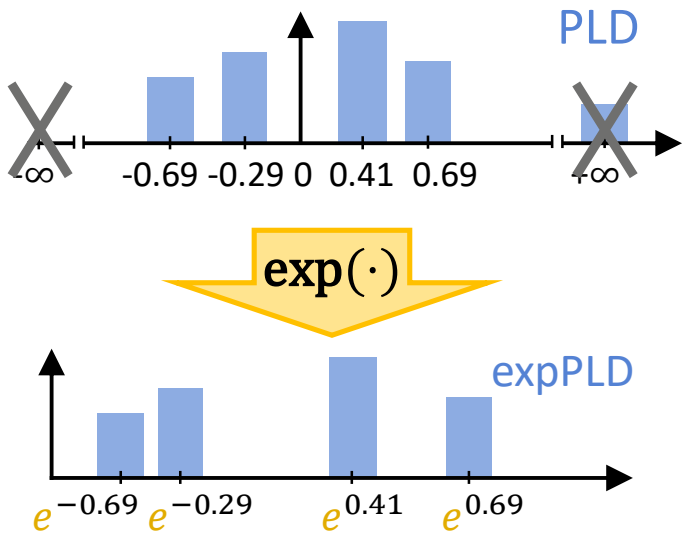

$$\log\left(\mathop{\mathbb{E}}_{y\sim PLD}\left[(e^{y})^{\lambda}\right]\right)^{\frac{1}{\lambda}} = \frac{1}{\lambda}\ln \mathop{\mathbb{E}}_{o\sim M(D_1)}\left[\left(\frac{\Pr\left[M(D_0)=o\right]}{\Pr\left[M(D_1)=o\right]}\right)^{\lambda+1}\right]$$

$$= \mathcal{D}_{\underbrace{\lambda+1}_{\alpha}}\left(M(D_0)\middle|M(D_1)\right)$$

Definition (informal)

(α, ε)-RDP: $\forall\, D_0, D_1\colon\ \mathcal{D}_{\alpha}\left(M(D_0)\middle|M(D_1)\right) \leq \varepsilon$.

- PLD ⇒ RDP: $\left(\alpha,\ \mathcal{D}_{\alpha}\left(M(D_0)\middle|M(D_1)\right)\right)_{\alpha\in\mathbb{N}}$ ✓
- PLD ⇐ RDP: If $\mathcal{D}_{\lambda+1}\left(M(D_0)\middle|M(D_1)\right) < \frac{1}{\lambda}\ln\left(c\, d^{\lambda}\, \lambda!\right)$ (with $c, d > 0$) (✓)
- With +∞: Approximate-RDP

$\lambda + 1 \rightarrow \alpha$

Rényi-Differential Privacy (RDP)

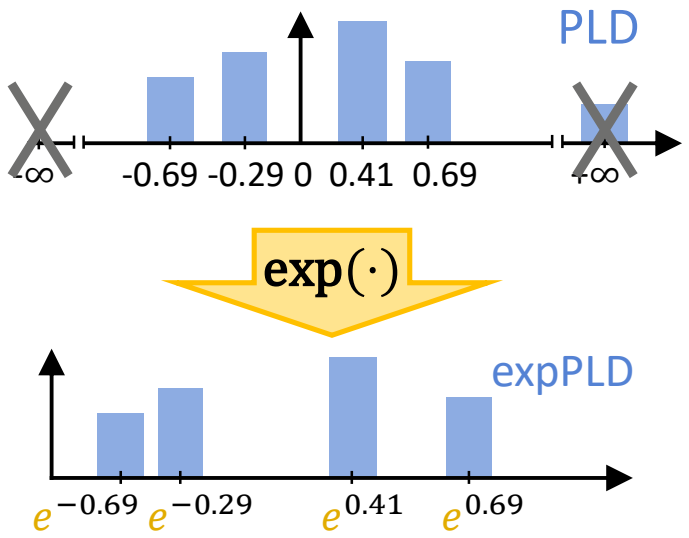

$$\ln\left(\mathop{\mathbb{E}}_{y\sim PLD}\left[(e^y)^\lambda\right]\right)^{\frac{1}{\lambda}} = \frac{1}{\lambda}\ln \mathop{\mathbb{E}}_{o\sim M(D_o)}\left[\left(\frac{\Pr\left[M(D_0)=o\right]}{\Pr\left[M(D_1)=o\right]}\right)^{\lambda}\right]$$

$$= \frac{1}{\lambda}\ln \mathop{\mathbb{E}}_{o\sim M(D_1)}\left[\left(\frac{\Pr\left[M(D_1)=o\right]}{\Pr\left[M(D_0)=o\right]}\right)^{\lambda+1}\right]$$

$$— = \mathcal{D}_{\lambda+1}\left(M(D_0)\middle|M(D_1)\right)$$

Definition (informal)

(α, ε)**-RDP:** $\forall\, D_0, D_1\colon\ \mathcal{D}_\alpha\left(M(D_0)\middle|M(D_1)\right) \leq \varepsilon$.

- PLD -> RDP: $\left(\alpha,\ \mathcal{D}_\alpha\left(M(D_0)\middle|M(D_1)\right)\right)_{\alpha\in\mathbb{N}}$ ✓
- PLD <- RDP: If $\mathcal{D}_{\lambda+1}\left(M(D_0)\middle|M(D_1)\right) < \frac{1}{\lambda}\ln\left(c\, d^\lambda\, \lambda!\right)$ (with $c, d > 0$) (✓)
- With +∞: Approximate-RDP

(ε, δ)-Probabilistic Differential Privacy (PDP)

- Given an ε_0, how to handle leftovers in PLD?
- We get a delta.
- This is PDP:

Definition (informal)

With $\varepsilon, \delta \geq 0$, M is $(\boldsymbol{\varepsilon}, \boldsymbol{\delta})$-**PDP**, if the total probability mass of all atomic events o is

$$\Pr[\{o \mid \mathcal{L}_{M(D_0),M(D_1)}(o) > \varepsilon\} \in M(D_0)] \leq \delta$$
$$\Pr[\{o \mid \mathcal{L}_{M(D_1),M(D_0)}(o) > \varepsilon\} \in M(D_1)] \leq \delta$$

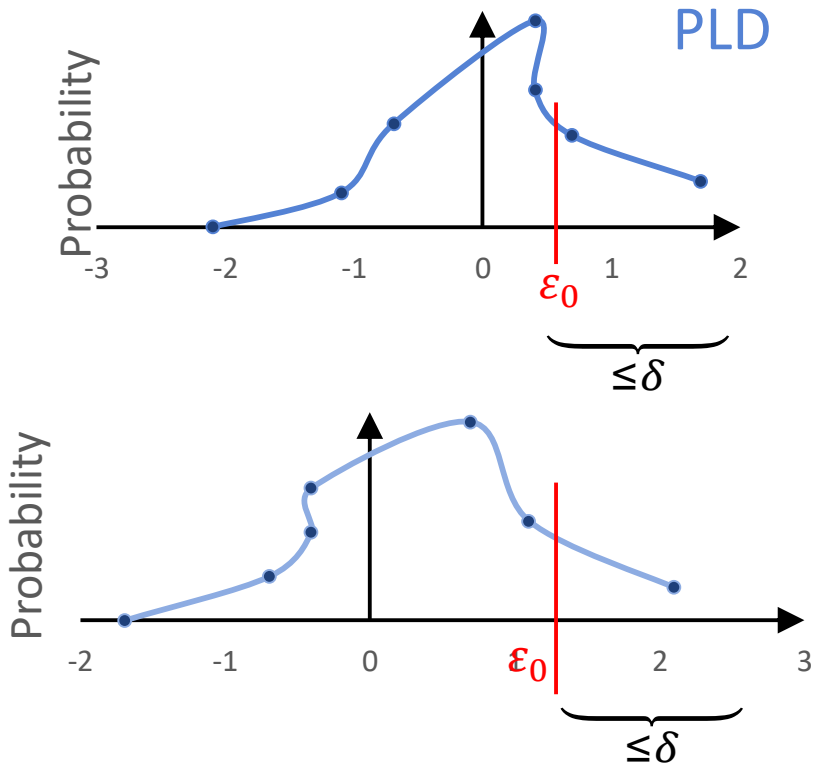

- Maybe explain (eps,delta) graphs here.

Gauss vs. Laplace Mechanism

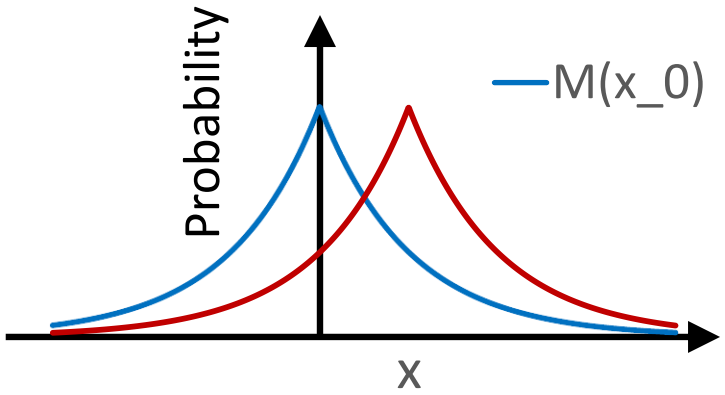

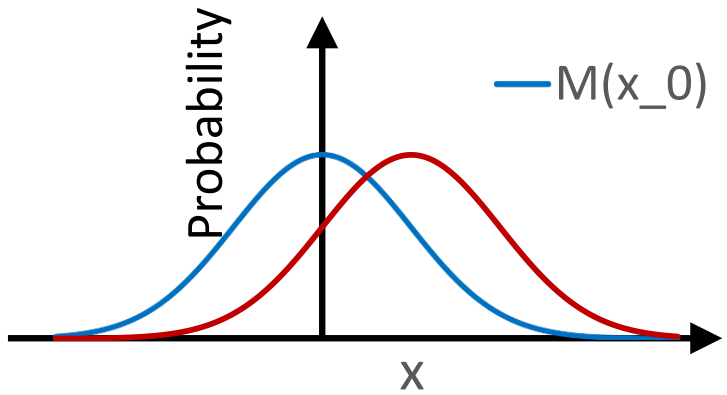

- Show PLD
- Show variances
- Give intuition why Gauss has lower var (heavier tail)
- Show graph illustating negligibility.
- Make graph for same variance , that shows smaler eps.

Independent Composition

- Draw twice (joint distribution)
  - $\left(M(D_0), \mathrm{M}'(\mathrm{D}_0')\right)$ vs. $\left(M(D_1), \mathrm{M}'(\mathrm{D}_1')\right)$
  - $\Pr[o \in M(D_0), o' \in M(D_0')] = \Pr[o \in M(D_0)] \cdot \Pr[o' \in M(D_0')]$

- $$\mathcal{L}_{\left(M(D_0),\mathrm{M}'(\mathrm{D}_0')\right),\left(M(D_1),\mathrm{M}'(\mathrm{D}_1')\right)}(\mathrm{o}, \mathrm{o}') = \ln \frac{\Pr\left[o \in M(D_0), o' \in M\left(D_0'\right)\right]}{\Pr\left[o \in M(D_1), o' \in M\left(D_1'\right)\right]}$$
  $$= \ln \frac{\Pr[o \in M(D_0)]}{\Pr[o \in M(D_1)]} + \ln \frac{\Pr\left[o' \in M\left(D_0'\right)\right]}{\Pr\left[o' \in M\left(D_1'\right)\right]}$$
  $$= \mathcal{L}_{M(D_0),M(D_1)}(\mathrm{o}) + \mathcal{L}_{\mathrm{M}'(\mathrm{D}_0'),\mathrm{M}'(\mathrm{D}_1')}(\mathrm{o}')$$

- Composition of Mechanisms is <u>Convolution</u> of *inner*PLD (pic of inner PLD)

TL;DR: Use Gaussian Noise!

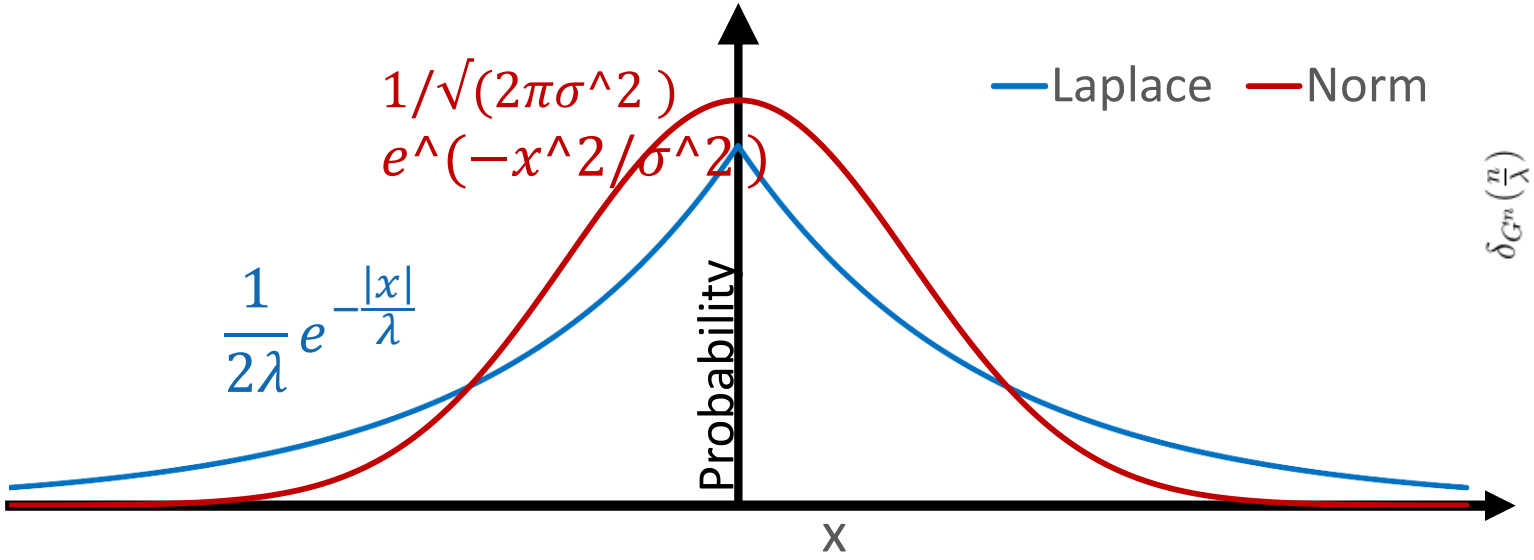

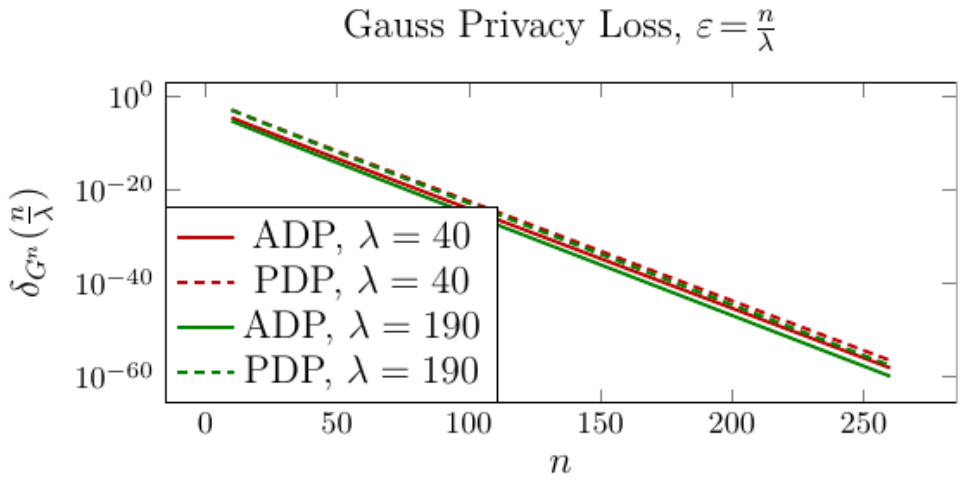

- After ~100 compositons: use Gaussian noise
- Gauss half the variance
- Same privacy loss class as Laplace.
- -> Gauss has higher utility than Laplace at similar privacy guarantees

Markov-ADP Bound

- Mironov: PDP bound based on Markov-bound
- Made better ADP bound by not over-approximating buckets
- Graph for illustration.
- Graph for showing superiority.

Dual PLD

- Illustrate

Worst Case distributions

- Exist for certain problems
- (forward) our mechanism allows different outputs compute composition leakage exactly

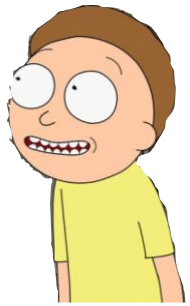

Definition:

Definition

A mechanism M is ε-*differentially private*, where $\varepsilon \geq 0$, if for all databases D_0 and D_1 with $D_0 \overset{1}{\approx} D_1$, and for all sets $S \subseteq [M]$, where $[M]$ is the range of M, the following equation holds:

$$\Pr[M(D_0) \in S] \leq e^{\varepsilon} \cdot \Pr[M(D_1) \in S]$$