Confidentialité différentielle à risque : Relier les sources d'aléa et un budget de confidentialité

Ashish Dandekar[1], Debabrota Basu[2],
Pierre Senellart[1], Stéphane Bressan[3]

[1]DI, École normale supérieure, Paris
[2]SequeL, Inria Lille Nord Europe; Chalmers University of Technology
[3]School of Computing, National University of Singapore

October 29, 2020
BDA, 2020

Data Privacy: ϵ-Differential Privacy [Dwork et al., 2014]

Information in input/database becomes private if it is indistinguishable from the output of a query/algorithm.

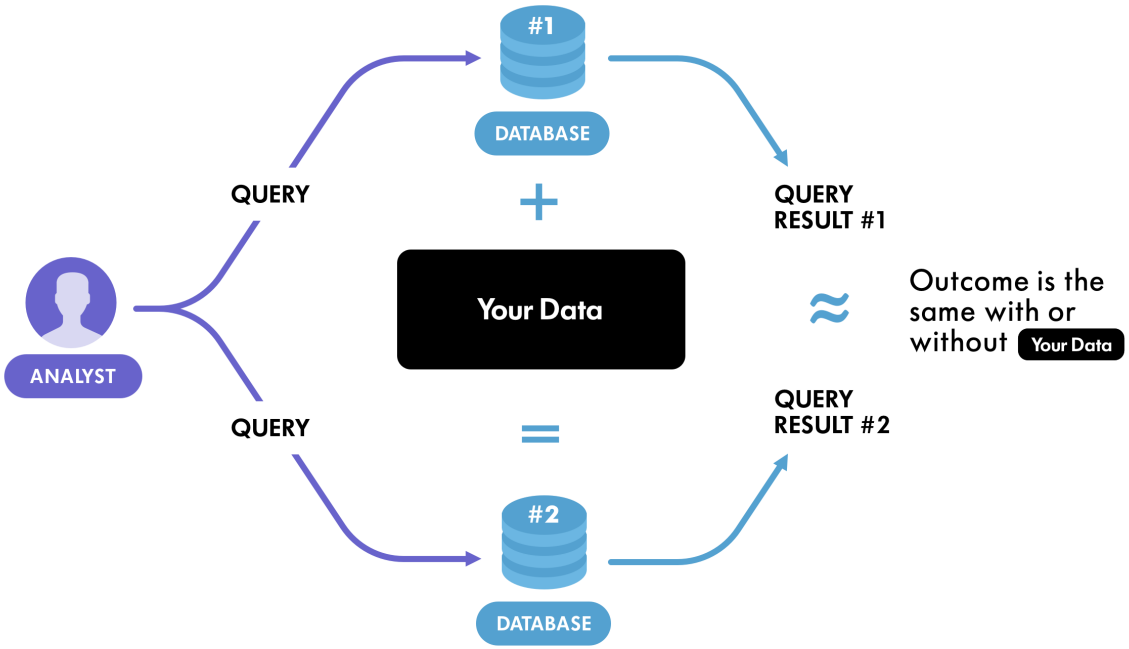

$$\frac{\mathbb{P}(\mathcal{A}(\mathrm{DB} + \text{my data}) = O)}{\mathbb{P}(\mathcal{A}(\mathrm{DB}) = O)} \leq \mathrm{e}^{\epsilon} \quad \longrightarrow \quad \epsilon - \mathrm{DP}$$

Image Courtesy: www.winton.com

What are We Proposing?

ϵ-Differential privacy is the worst-case privacy guarantee.

We propose **Privacy at Risk** that

1. provides *probabilistic bounds on the privacy guarantee* of differential privacy, and
2. the probabilistic bound *quantifies various sources of randomness*.

The privacy guarantee ϵ is too abstract to be actionable.

We propose a **cost model** that *translates* the abstract *privacy guarantee to a compensation budget* estimated by a GDPR compliant business entity.

Privacy at risk: Definition

Privacy at Risk (PaR)[1]

$$\mathbb{P}\left[\text{Algorithm } \mathcal{A} \text{ satisfies } \epsilon\text{-DP}\right] \geq 1 - \mathrm{PaR}$$

$$\mathbb{P}\left[\log\left|\frac{\mathbb{P}(\mathcal{A}(\mathrm{DB} + \text{my data}) = O)}{\mathbb{P}(\mathcal{A}(\mathrm{DB}) = O)}\right| \leq \epsilon\right] \geq 1 - \gamma$$

Computing Privacy at Risk (γ) requires quantification of:

1. *Implicit randomness* of data generating distribution or limited access to it,
2. *Explicit randomness* induced by the privacy preserving algorithm.

[1] Risk analysts use *Value at Risk* [Jorion, 2000] to quantify the loss in investments for a given portfolio and an acceptable confidence bound.

Privacy at risk: Properties

Post-processing

Privacy at risk guarantee does not stay the same after the privatised output is further processed.

Weak Post-processing

An algorithm satisfying (ϵ, γ)-privacy at risk will satisfy (ϵ, γ)-approximate DP after processing of its output.

Convexity

An algorithm satisfies $(\lambda\epsilon_1 + (1-\lambda)\epsilon_2, p\gamma_1 + (1-p)\gamma_2)$ PaR, if it chooses the outputs of $\mathcal{A}_1$ with (ϵ_1, γ_1) PaR and $\mathcal{A}_2$ with (ϵ_2, γ_2) PaR with probability p and $1-p$.

Composition of Privacy at Risk and Efficient Moment Accounting

Theorem: Applying an algorithm satisfying ϵ_0-DP and (ϵ, γ)-PaR for n times will satisfy (ϵ', δ)-approximate DP, such that

$$\epsilon' = \epsilon_0 \sqrt{2n \ln \frac{1}{\delta}} + n\mu(\gamma, \epsilon, \epsilon_0).$$

Application: We verify efficiency of composing PaR by using it for moment accounting in PATE framework.

δ	#Queries	Privacy level for moment accountant(ϵ)	
		with differential privacy	with privacy at risk
10^{-5}	100	2.04	**1.81**
10^{-5}	1000	8.03	**5.95**

Computation of privacy at risk

Explicit randomness

Source: the noise distribution of the privacy-preserving mechanism.

Probabilistic differential privacy [Machanavajjhala et al., 2008]

Implicit randomness

Source: the data-generation distribution and limited access to it.

Random differential privacy [Hall et al., 2012]

Privacy at risk for Laplace mechanism

Laplace mechanism [Dwork et al., 2014]

$$\text{noisy output} \leftarrow f(DB) + Lap\left(0, 2\left(\frac{\Delta_f}{\epsilon_0}\right)^2\right)$$

Source(s) of randomness	**Analytical result**	**Contribution**
Laplace distribution	Closed form solution (Theorem 4.1)	Overlap of output distribution under the sensitivity constraint.
Data-generation distribution	Upper bound on the confidence level (Theorem 4.2)	Sensitivity estimation using samples from data-generation distribution.
Laplace distribution + data-generation distribution	Upper bound on the confidence level (Theorem 4.4)	Overlap of output distribution under the estimated sensitivity constraint.

A sample result: Explicit Randomness

Privacy at risk level $\gamma_1 \in [0, 1]$ with which a Laplace Mechanism $\mathcal{L}_{\epsilon_0}^{\Delta_f}$ satisfies a privacy level $\epsilon \geq 0$ for a query f is

$$\gamma_1 = \frac{\mathbb{P}(T \leq \epsilon | T \sim \text{BesselK}\left(k, \frac{\Delta_f}{\epsilon_0}\right))}{\mathbb{P}(T \leq \epsilon_0 | T \sim \text{BesselK}\left(k, \frac{\Delta_f}{\epsilon_0}\right))}.$$

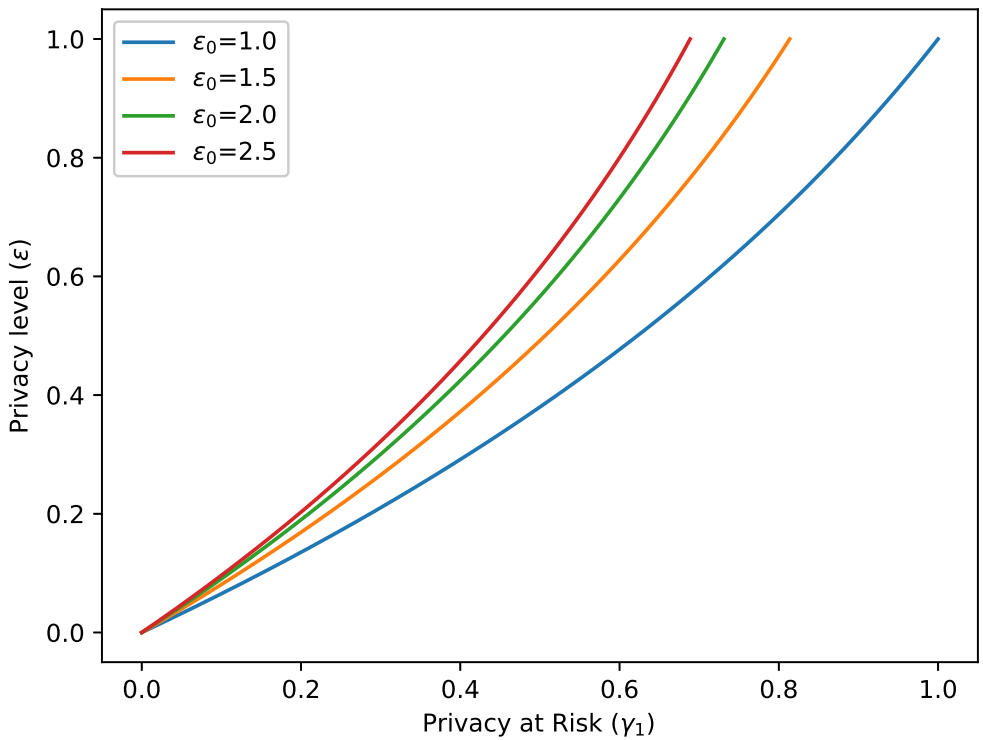

Choosing an ϵ: Privacy-utility trade-off

We run differentially private ridge regression on 2000 US Census dataset.

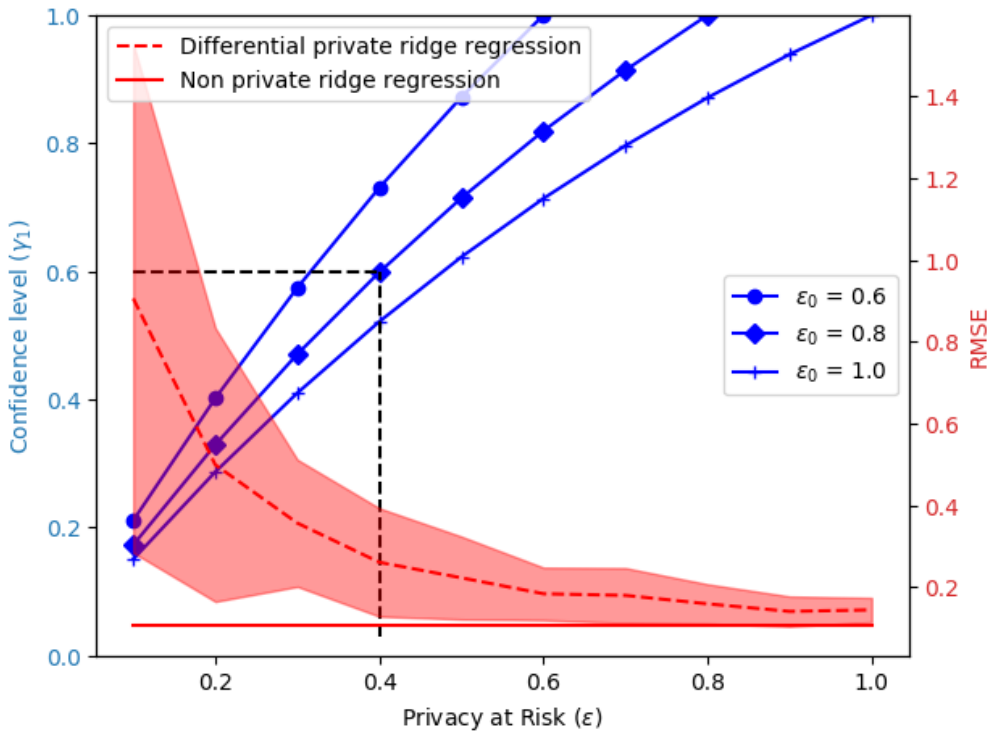

The privacy-utility plot navigates a data steward to choose a privacy level ϵ based on the utility requirements.

Can we make ϵ any less abstract?

Art. 82 GDPR

Right to compensation and liability

(1) Any person who has suffered material or non-material damage as a result of an infringement of this Regulation shall have the right to receive compensation from the controller or processor for the damage suffered.

We assume that the compensation budget secured by a GDPR compliant business entity is directly dependent on the DP guarantee provided by the business entity while processing the data.

An economic translation of privacy guarantee

Compensation budgets per stakeholder

E $\leftarrow$ in absence of privacy
E_ϵ^{dp} $\leftarrow$ under ϵ-DP
$E_{\epsilon_0}^{par}(\epsilon, \gamma)$ $\leftarrow$ under (ϵ, γ)-PaR satisfied by an ϵ_0-DP mechanism

Properties of a cost model

- For all $\epsilon \in \mathbb{R}^{\geq 0}$, $E_\epsilon^{dp} \leq E$.
- As $\epsilon \to 0$, $E_\epsilon^{dp} \to 0$.
- As $\epsilon \to \infty$, $E_\epsilon^{dp} \to E$.
- E_ϵ^{dp} is a monotonically increasing function of ϵ.

An economic translation of privacy guarantee

Cost model for ϵ-differential privacy

$$E_{\epsilon}^{dp} \triangleq Ee^{-\frac{c}{\epsilon}}$$

Cost model for (ϵ, γ)-privacy at risk

$$E_{\epsilon_0}^{par}(\epsilon, \gamma) \triangleq \gamma E_{\epsilon}^{dp} + (1-\gamma)E_{\epsilon_0}^{dp} \quad \leftarrow \text{Convex function!}$$

Minimal privacy budget

There exists a privacy at risk level (ϵ^*, γ^*) for a specified ϵ_0-differentially private mechanism that yields the smallest compensation budget!

An illustration of the economic model

Case study: Obesity related data breach in an organisation

$E \leftarrow$ \$5500, which is average increment in the premiums for health insurances with morbid obesity [Moriarty et al., 2012].

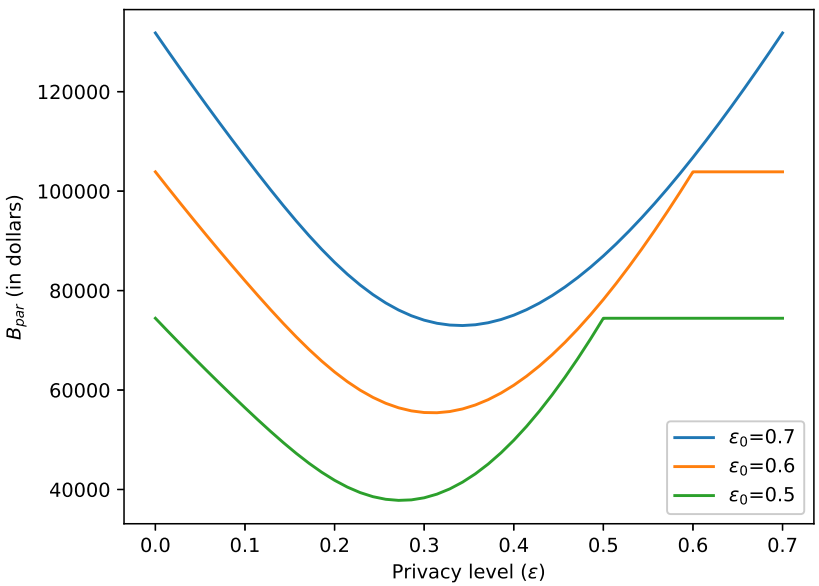

$$E_{0.5}^{dp} = \$74434.40$$

$$E_{0.5}^{par}(0.29, 0.64) = \$37805.86$$

$$\text{Savings} = \$36628.54$$

Explicit Randomness

noisy output $\leftarrow \mathcal{A}(DB) + Lap\left(0, 2\left(\frac{\Delta_f}{\epsilon_0}\right)^2\right)$

Implicit Randomness

estimated sensitivity $\leftarrow$ upper bound on sensitivity from n data samples

Privacy at Risk: Probabilistic Privacy Guarantee

$\mathbb{P}\left[\text{Algorithm } \pi \text{ satisfies } \epsilon\text{-DP}\right] \geq 1 - \mathrm{PaR}$

A Convex Cost Model

$$COST(\epsilon|\mathrm{PaR}) \triangleq \mathrm{PaR} \times COST(\epsilon_0|0) + (1 - \mathrm{PaR}) \times COST(\epsilon|0)$$

Minimum Privacy Budget

$(COST^*, \epsilon^*, PaR^*)$

[0] Longer paper: `https://arxiv.org/abs/2003.00973`

References I

Dwork, C., Roth, A., et al. (2014).
The algorithmic foundations of differential privacy.
Foundations and Trends® in Theoretical Computer Science, 9(3–4):211–407.

Hall, R., Rinaldo, A., and Wasserman, L. (2012).
Random differential privacy.
Journal of Privacy and Confidentiality, 4(2):43–59.

Jorion, P. (2000).
Value at risk: The new benchmark for managing financial risk.

Machanavajjhala, A., Kifer, D., Abowd, J., Gehrke, J., and Vilhuber, L. (2008).
Privacy: Theory meets practice on the map.
In *Data Engineering, 2008. ICDE 2008. IEEE 24th International Conference on*, pages 277–286. IEEE.

Moriarty, J. P., Branda, M. E., Olsen, K. D., Shah, N. D., Borah, B. J., Wagie, A. E., Egginton, J. S., and Naessens, J. M. (2012).
The effects of incremental costs of smoking and obesity on health care costs among adults: a 7-year longitudinal study.
Journal of Occupational and Environmental Medicine, 54(3):286–291.

A means to bound the epsilon!

Suppose that we have a maximum permissible expected mean absolute error T.

$$\frac{1}{T} \leq \epsilon \leq \left[\ln \left(\frac{\gamma E}{B - (1-\gamma) E_{\epsilon_0}^{dp}} \right) \right]^{-1} \tag{1}$$