The Composition Theorem for Differential Privacy

Sewoong Oh

Department of ISE
University of Illinois at Urbana-Champaign

Joint work with Peter Kairouz (UIUC) and Pramod Viswanath (UIUC)

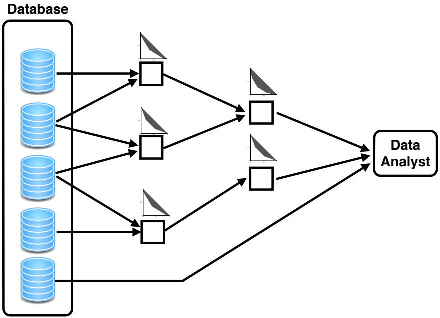

Privacy calculus

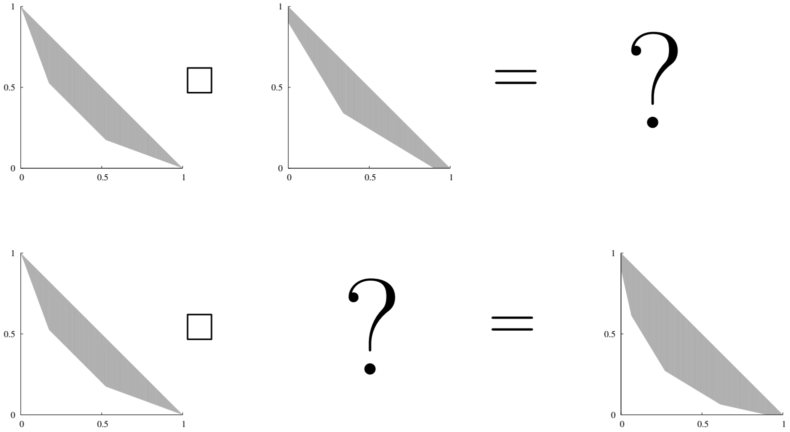

Privacy via plausible deniability [Warner 1965]

Have you ever used illegal drugs?

say yes

answer truthfully

ε-differential privacy

D_0

Alice	22
Bob	45
⋮	⋮
Me	23

D_1

Alice	22
Bob	45
⋮	⋮

$$\mathbb{P}(q(D_0) \in S) \quad \leq \quad e^{\varepsilon}\, \mathbb{P}(q(D_1) \in S)$$

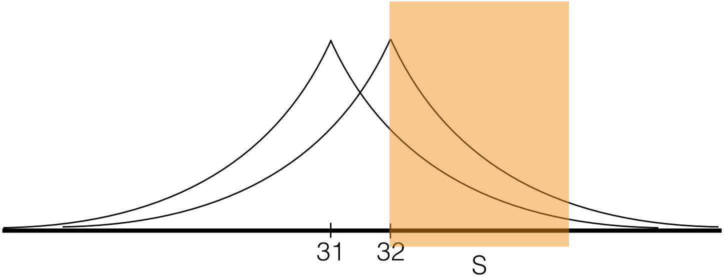

(ε, δ)-differential privacy

D_0

Alice	22
Bob	45
⋮	⋮
Me	23

D_1

Alice	22
Bob	45
⋮	⋮

$$\mathbb{P}(q(D_0) \in S) \quad \leq \quad e^{\varepsilon}\,\mathbb{P}(q(D_1) \in S) + \delta$$

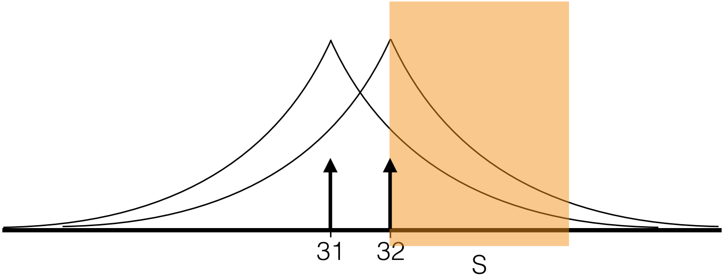

Composition of Differentially Private Mechanisms

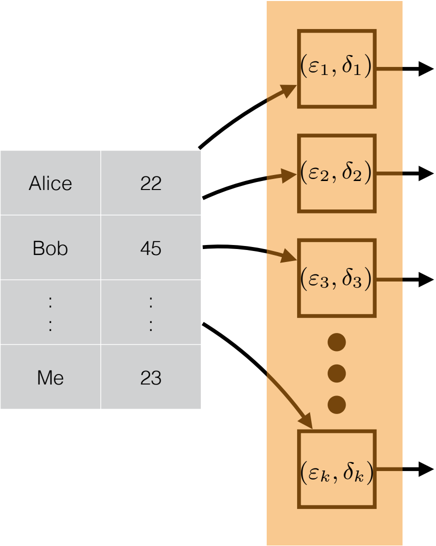

How much privacy is lost in the end?

$\left(\sum_{i=1}^{k} \varepsilon_i, \sum_{i=1}^{k} \delta_i \right)$-differentially private

Connections to binary hypothesis testing

H_0: database is D_0

H_1: database is D_1

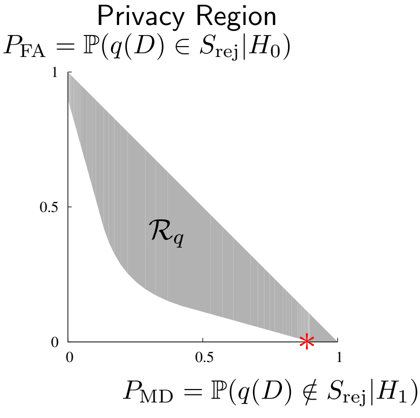

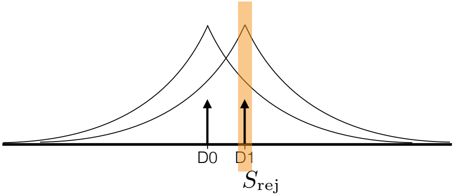

Connections to binary hypothesis testing

H_0: database is D_0

H_1: database is D_1

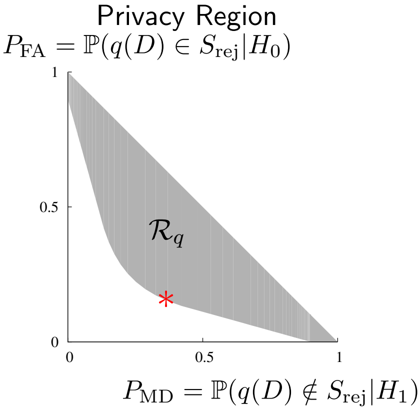

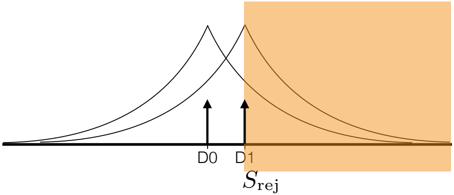

Connections to binary hypothesis testing

H_0: database is D_0

H_1: database is D_1

Privacy Region

$P_{\mathrm{FA}} = \mathbb{P}(q(D) \in S_{\mathrm{rej}} | H_0)$

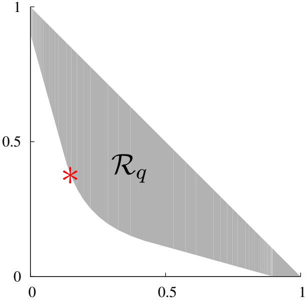

$P_{\mathrm{MD}} = \mathbb{P}(q(D) \notin S_{\mathrm{rej}} | H_1)$

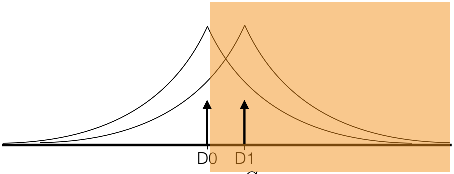

S_{rej}

Connections to binary hypothesis testing

H_0: database is D_0

H_1: database is D_1

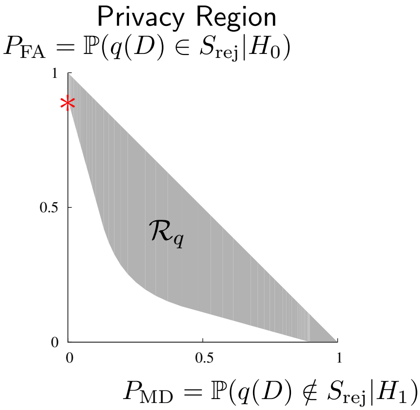

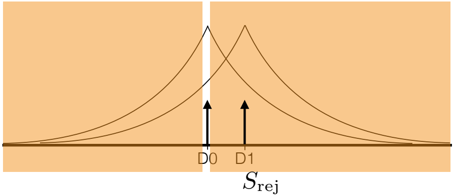

Differential privacy and privacy region are equivalent

$$\mathbb{P}(q(D_0) \in S) \le e^{\varepsilon}\mathbb{P}(q(D_1) \in S) + \delta$$

$$\begin{aligned} P_{\text{FA}} + e^{\varepsilon} P_{\text{MD}} &\ge 1 - \delta \\ e^{\varepsilon} P_{\text{FA}} + P_{\text{MD}} &\ge 1 - \delta \end{aligned}$$

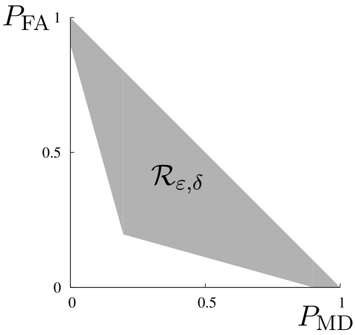

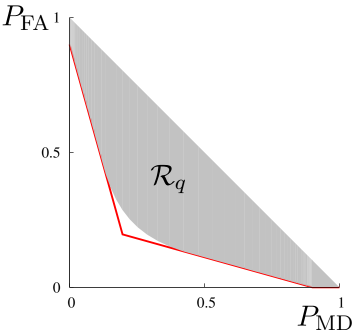

q is (ε, δ)-differentially private $\iff$ $\mathcal{R}_q \subseteq \mathcal{R}_{\varepsilon,\delta}$

Data Processing Inequality (DPI)

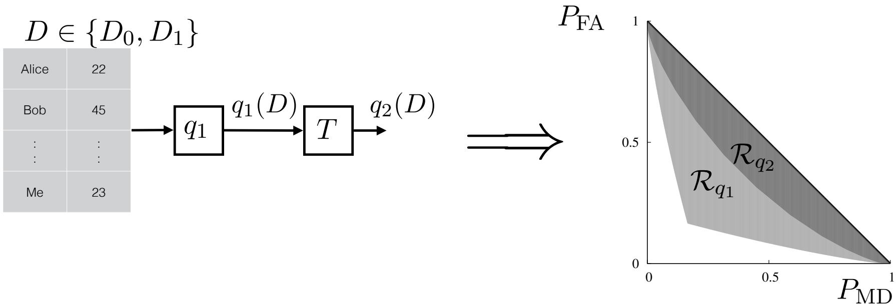

Data Processing Inequality (DPI)

$$D\text{–}q_1(D)\text{–}q_2(D) \quad \Longrightarrow \quad \mathcal{R}_{q_2} \subseteq \mathcal{R}_{q_1}$$

Converse to DPI

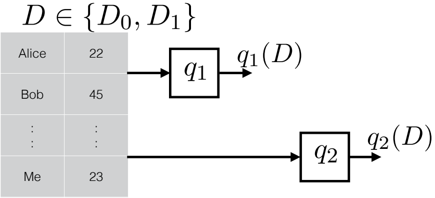

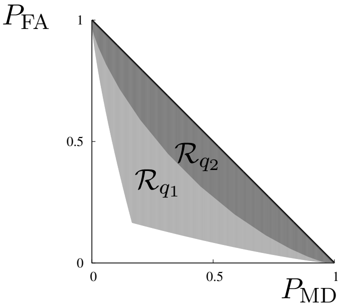

Converse to the Data Processing Inequality [KOV '15]

$$D–q_1(D)–q_2(D) \quad \Longleftarrow \quad \mathcal{R}_{q_2} \subseteq \mathcal{R}_{q_1}$$

precisely, there exists a coupling of $q_1(D)$ and $q_2(D)$ such that

(a) $D–q_1(D)–q_2(D)$; or equivalently

(b) $q_2(D) = T\big(q_1(D)\big)$. [Blackwell 1953]

Converse to DPI

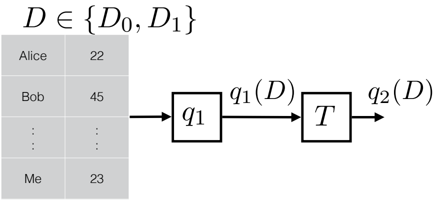

$\Longleftarrow$

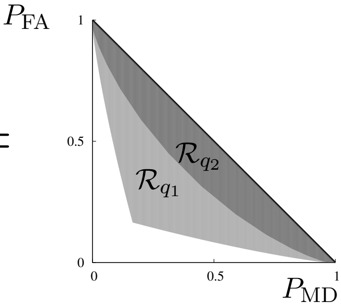

Converse to the Data Processing Inequality [KOV '15]

$$D – q_1(D) – q_2(D) \quad \Longleftarrow \quad \mathcal{R}_{q_2} \subseteq \mathcal{R}_{q_1}$$

precisely, there exists a coupling of $q_1(D)$ and $q_2(D)$ such that

(a) $D – q_1(D) – q_2(D)$; or equivalently

(b) $q_2(D) = T\big(q_1(D)\big)$. [Blackwell 1953]

Dominant mechanisms for (ε, δ)-differential privacy

the converse DPI implies that the following randomized response $q_{\varepsilon,\delta}$ dominates over all (ε, δ)-differentially private mechanisms

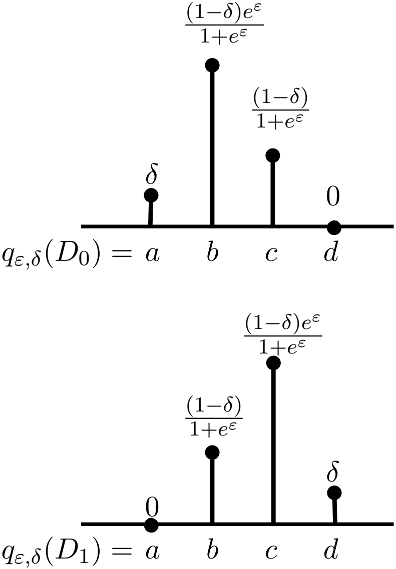

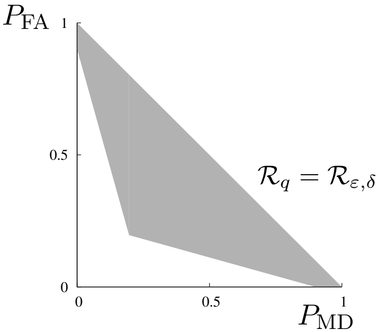

Dominant mechanism under composition

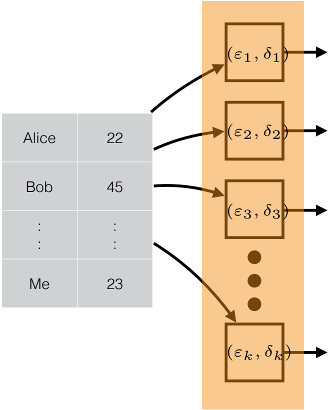

How much privacy is lost in the end?

For what values of (ε, δ), is the resulting composition still differentially private?

How does privacy region evolve under composition?

Dominant mechanism under composition

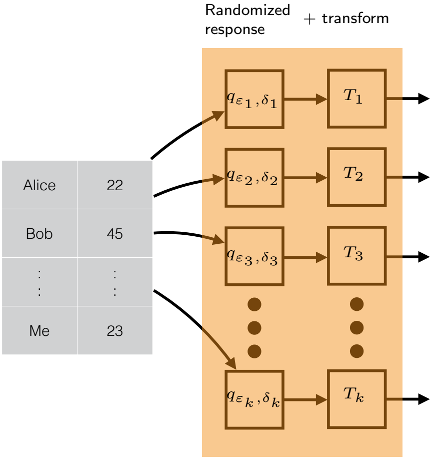

How much privacy is lost in the end?

For what values of (ε, δ), *is the resulting composition still differentially private?*

How does privacy region evolve under composition?

Dominant mechanism under composition

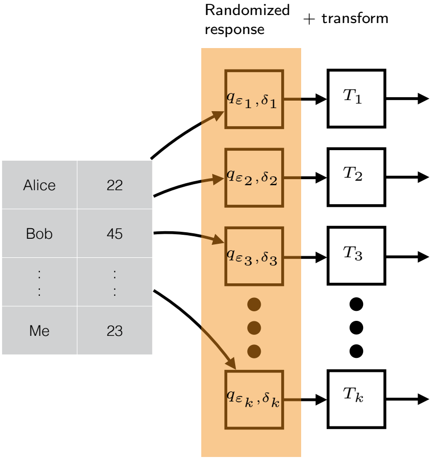

How much privacy is lost in the end?

For what values of (ε, δ), is the resulting composition still differentially private?

How does privacy region evolve under composition?

Composition of dominant mechanisms

k composition of $(0.4, 0.1)$-differential private mechanisms

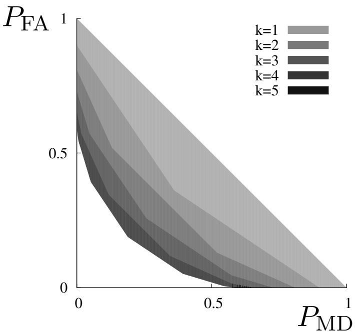

this gives the exact evolution of privacy, such that any known results on composition are corollaries.

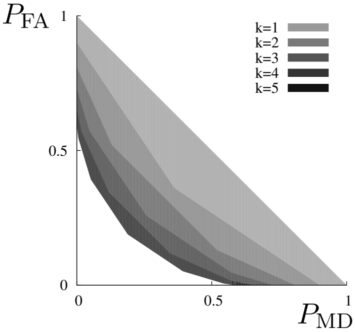

The composition theorem [Dwork, et al '10], [KOV '15]

k-fold composition of (ε, δ)-differentially private mechanisms satisfy $(\tilde{\varepsilon}, k\delta + \tilde{\delta})$-differential privacy with

$$\tilde{\varepsilon}_{\tilde{\delta}} = k\varepsilon^2 + \varepsilon\sqrt{2k\log(1/\tilde{\delta})}$$

significant improvement over $\big(k\varepsilon, k\delta\big)$-guarantee when $\varepsilon \to 0$

Comparisons with the state-of-the-art

30-fold composition of $(0.1, 0.001)$-differentially private mechanisms

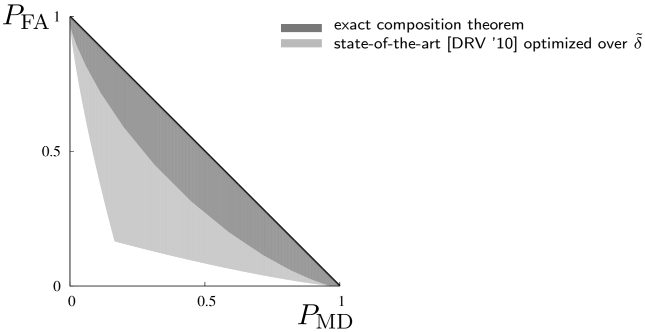

Recap

- Computational tool for exact composition

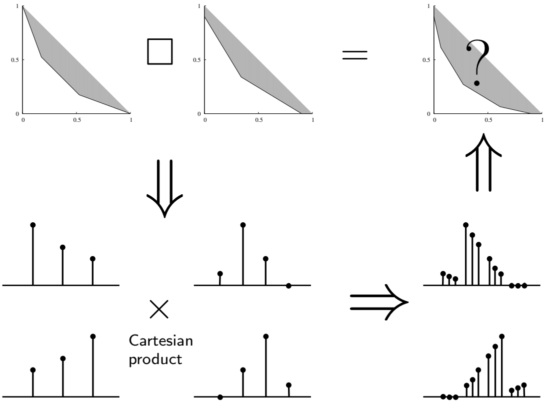

- Improved "cut-and-paste" composition theorem
  $(\tilde{\varepsilon}, k\delta + \tilde{\delta})$-differential privacy with

$$\tilde{\varepsilon}_{\tilde{\delta}} = k\varepsilon^2 + \varepsilon\sqrt{2k\log(1/\tilde{\delta})}$$

Going forward

- Computational Complexity [Vadhan, Murtagh '15]
- *"Optimality of non-interactive randomized response"*, arXiv:1407.1546

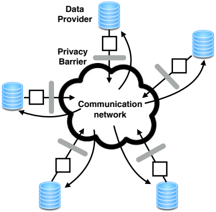

- Dominant Mechanisms for Large Alphabets
  "Extremal mechanisms for local differential privacy", NIPS 2014

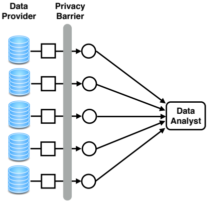

- *"The composition theorem for differential privacy"*, ICML 2015