Maximal α-Leakage for Quantum Privacy Mechanisms and Operational Meaning of Measured Rényi Capacity

(arXiv: 2403.14450)

Bo-Yu Yang, Hsuan Yu, and Hao-Chung Cheng

National Taiwan University

ISIT 2024, Athens

Information Leakage

In the era Internet of Things (IoT), data may be published or shared via **privacy mechanisms** everywhere.

- social network, online data sharing, mobile computing

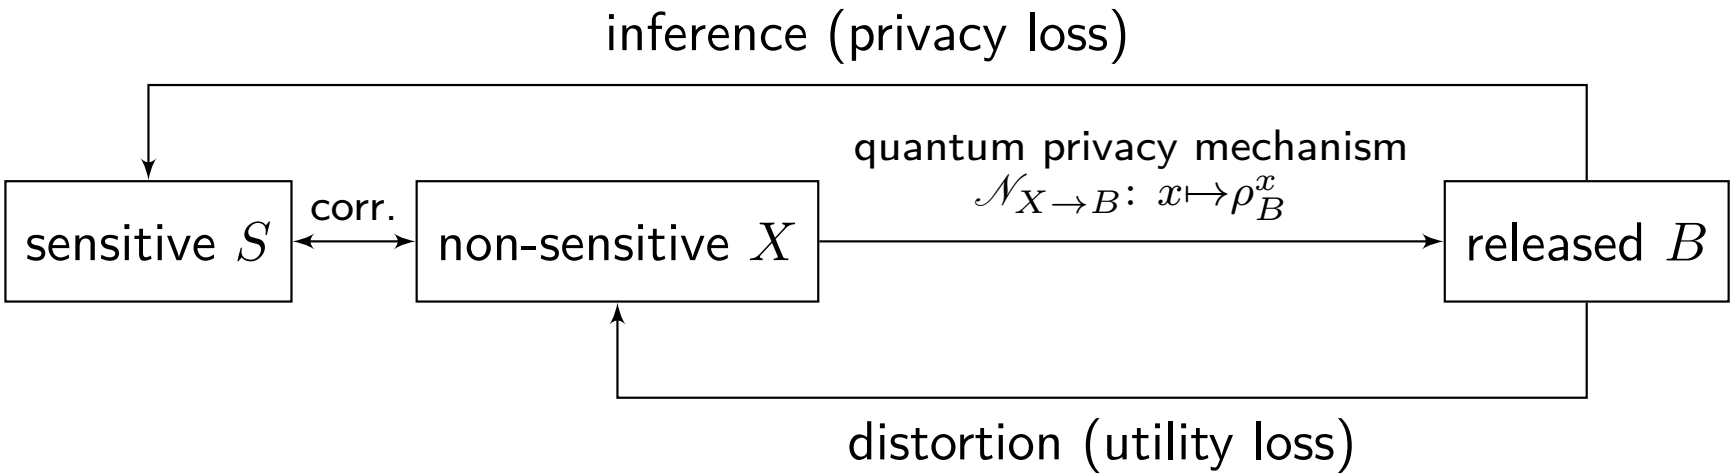

How much information does a **quantum system** B leak about S?

Quantum Systems

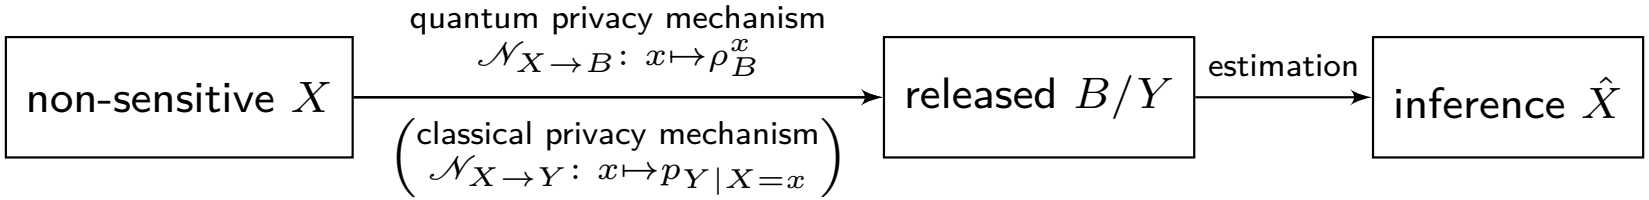

The adversary's goal: to **infer** data X from observation B (classical Y) with estimation $\hat{X}$

$\Rightarrow$ How much information does the adversary gain/lose?

- Classical system: prob. distr. $\Rightarrow p_{Y|X=x}(y)\geq 0$, $\sum_y p_{Y|X=x}(y)=1$
- Quantum system: density matrix $\Rightarrow \rho_B^x\geq 0, \mathsf{Tr}[\rho_B^x]=1$
  - $\rho_B^x=\sum_y \lambda_y^{X=x}\,|\psi_y\rangle\langle\psi_y|$, where $\lambda_y^{X=x}\geq 0$, $\sum_y \lambda_y^{X=x}=1$

Quantum Privacy Mechanisms

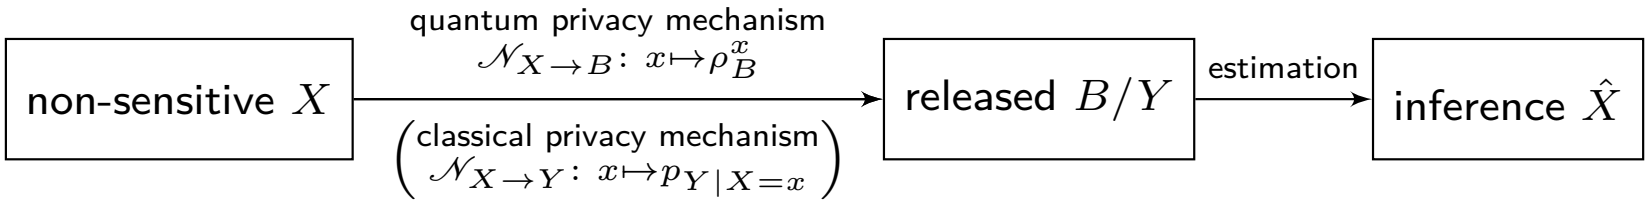

	classical Y	quantum B
privacy mechanism	$\mathcal{N}_{X\to Y}: x\mapsto p_{Y\|X=x}$	$\mathcal{N}_{X\to B}: x\mapsto \rho_B^x$
adversary	$p_{\hat{X}\|Y}$	Π_{XB} (POVM)

- **Positive Operator-Valued Measure** (POVM)
  $\Pi_{XB} := \sum_{x\in\mathcal{X}} |x\rangle\langle x| \otimes \Pi_B^x$ s.t. $\Pi_B^x \geq 0$ and $\sum_{x\in\mathcal{X}} \Pi_B^x = \mathbb{1}_B$
  - Born's rule: $\Pr\left\{\hat{X}=X \mid X=x, B\right\} = \mathrm{Tr}[\rho_B^x \Pi_B^x]$
  - can be viewed as quantum generalization of classical strategies

Parameterized Gain/Loss Functions (1/2)

We propose

- Minimal expected α-loss

$$\varepsilon_\alpha(X|B)_\rho := \inf_{\text{POVM } \Pi_{XB}} \frac{\alpha}{\alpha-1}\left(1-\mathsf{Tr}\left[\rho_{XB}\Pi_{XB}^{\frac{\alpha-1}{\alpha}}\right]\right)$$

- Maximal expected α-gain

$$\mathsf{P}_\alpha(X|B)_\rho := \sup_{\text{POVM } \{\Pi_B^x\}_{x\in\mathcal{X}}} \sum_{x\in\mathcal{X}} p_X(x)\mathsf{Tr}\left[\rho_B^x(\Pi_B^x)^{\frac{\alpha-1}{\alpha}}\right]$$

- To minimize loss $=$ To maximize gain

$$\mathsf{P}_\alpha(X|B)_\rho = 1-\frac{\alpha-1}{\alpha}\varepsilon_\alpha(X|B)_\rho$$

Parameterized Gain/Loss Functions (2/2)

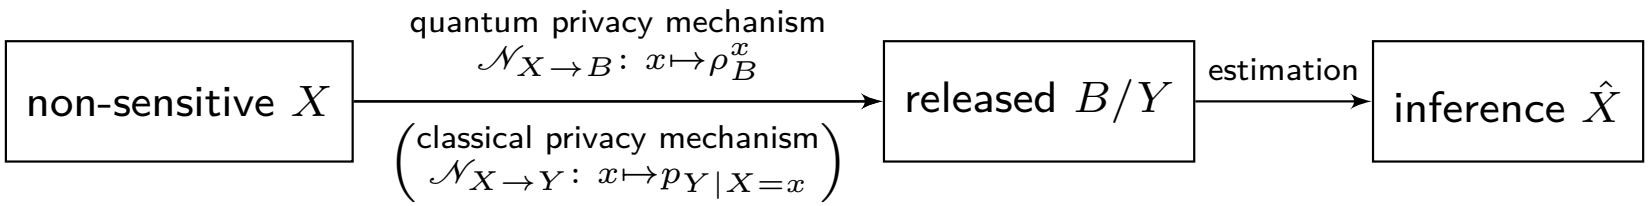

Expected [...]	classical Y [Liao'19][1]	**quantum** B (our work)
0-1 loss ($\alpha=\infty$)	$1-\mathbb{E}_{X,Y}\left[p_{\hat{X}\mid Y}\right]$	$1-\mathsf{Tr}\left[\rho_{XB}\Pi_{XB}\right]$
log loss ($\alpha=1$)	$-\mathbb{E}_{X,Y}\left[\log p_{\hat{X}\mid Y}\right]$	$-\mathsf{Tr}\left[\rho_{XB}\log\Pi_{XB}\right]$
α-loss ($\alpha\in[1,\infty]$)	$\frac{\alpha}{\alpha-1}\left(1-\mathbb{E}_{X,Y}\left[p_{\hat{X}\mid Y}^{\frac{\alpha-1}{\alpha}}\right]\right)$	$\frac{\alpha}{\alpha-1}\left(1-\mathsf{Tr}\left[\rho_{XB}\Pi_{XB}^{\frac{\alpha-1}{\alpha}}\right]\right)$

- 0-1 loss: probability of incorrectly guessing
- log loss: uncertainty of the strategy
- α-loss: tunable function for various adversarial strategies

[1] J. Liao, O. Kosut, L. Sankar, and F. du Pin Calmon. Tunable measures for information leakage and applications to privacy-utility tradeoffs. IEEE Transactions on Information Theory, 2019.

Existing Results

IEEE TRANSACTIONS ON INFORMATION THEORY, VOL. 55, NO. 9, SEPTEMBER 2009 4337

The Operational Meaning of Min- and Max-Entropy

Robert König, Renato Renner, and Christian Schaffner

For $\alpha = \infty$, the error exponent of maximal expected ∞-gain is reduced to the so-called **min-entropy** (guessing entropy) introduced by [König-Renner-Schaffner'09]

$$-\log \mathsf{P}_\infty(X|B)_\rho = H^*_\infty(X|B)_\rho := -\inf_{\sigma_B \in \mathcal{S}(\mathcal{H}_B)} D^*_\infty(\rho_{XB} \,\|\, \mathbb{1}_X \otimes \sigma_B).$$

- Maximal probability of correctly guessing:
  $\mathsf{P}_\infty(X|B)_\rho := \sup_{\text{POVM}\,\{\Pi^x_B\}_{x\in\mathcal{X}}} \sum_{x\in\mathcal{X}} p_X(x) \mathsf{Tr}\left[\rho^x_B(\Pi^x_B)\right]$

Characterization 1: Maximal Expected α-Gain

Theorem

For any classical-quantum state ρ_{XB} and $\alpha \in [1, \infty]$, the maximal expected α-gain is given by

$$\mathsf{P}_\alpha(X|B)_\rho = \mathrm{e}^{\frac{1-\alpha}{\alpha} H_\alpha^{\mathbb{M}}(X|B)_\rho}.$$

- Measured conditional Rényi entropy [Hanson et al.'20][2]: $H_\alpha^{\mathbb{M}}(X|B)_\rho := -\inf_{\sigma_B} D_\alpha^{\mathbb{M}}(\rho_{XB} \| \mathbb{1}_X \otimes \sigma_B)$
- Variational formula of measured Rényi divergence: $D_\alpha^{\mathbb{M}}(\rho\|\sigma) := \sup_{\omega>0} \frac{1}{\alpha-1} \log \left(\left(\mathsf{Tr}\left[\rho \omega^{1-\frac{1}{\alpha}}\right] \right)^\alpha (\mathsf{Tr}\left[\sigma\omega\right])^{1-\alpha} \right)$
- Proof sketch: transform the condition of POVM in $\mathsf{P}_\alpha(X|B)_\rho$ into Lagrange multiplier (penalty terms), and leverage duality

[2] E. P. Hanson, V. Katariya, N. Datta and M. M. Wilde. Guesswork with Quantum Side Information: Optimal Strategies and Aspects of Security. IEEE International Symposium on Information Theory. 2020.

Information Leakage Measure: α-Leakage

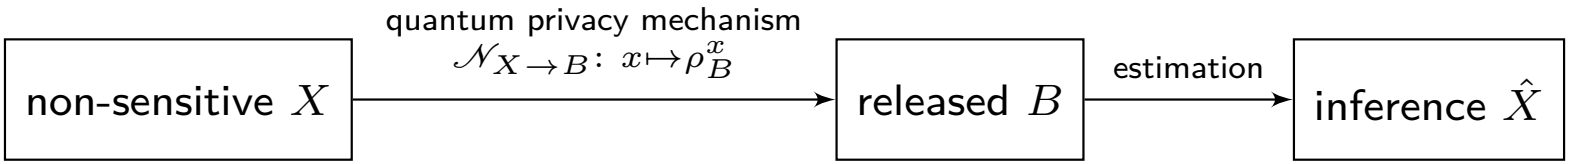

The α-leakage is defined as

$$\mathcal{L}_\alpha(X\to B)_\rho := \frac{\alpha}{\alpha-1}\log\frac{\mathsf{P}_\alpha(X|B)_\rho}{\mathsf{P}_\alpha(X)_\rho}.$$

- maximal expected α-gain after observing B
  $\mathsf{P}_\alpha(X|B)_\rho := \sup_{\text{POVM }\{\Pi_B^x\}_{x\in\mathcal{X}}} \sum_{x\in\mathcal{X}} p_X(x)\mathsf{Tr}\left[\rho_B^x(\Pi_B^x)^{\frac{\alpha-1}{\alpha}}\right]$
- max-expected α-gain for an inference $\hat{X}$ before observing B:
  $\mathsf{P}_\alpha(X)_\rho := \sup_{p_{\hat{X}}:\hat{X}\perp X} \mathbb{E}_{x\sim p_X}\left[\mathsf{Pr}(\hat{X}=X|X=x)^{\frac{\alpha-1}{\alpha}}\right]$
- depict the **multiplicative increase** of the adversary's knowledge

Characterization 2: α-Leakage

Theorem

For $\alpha \in [1, \infty]$, α-leakage can be expressed as

$$\mathcal{L}_\alpha(X \to B)_\rho = I_\alpha^{\mathrm{A},\mathbb{M}}(X:B)_\rho.$$

Proof sketch:

- α-Leakage:

$$\mathcal{L}_\alpha(X \to B)_\rho := \frac{\alpha}{\alpha-1} \log \mathsf{P}_\alpha(X|B)_\rho - \frac{\alpha}{\alpha-1} \log \mathsf{P}_\alpha(X)_\rho$$

- Measured Arimoto information:

$$I_\alpha^{\mathrm{A},\mathbb{M}}(X:B)_\rho := \left(-H_\alpha^{\mathbb{M}}(X|B)_\rho\right) - \left(-H_\alpha(X)_p\right)$$

Information Leakage Measure: Maximal α-Leakage

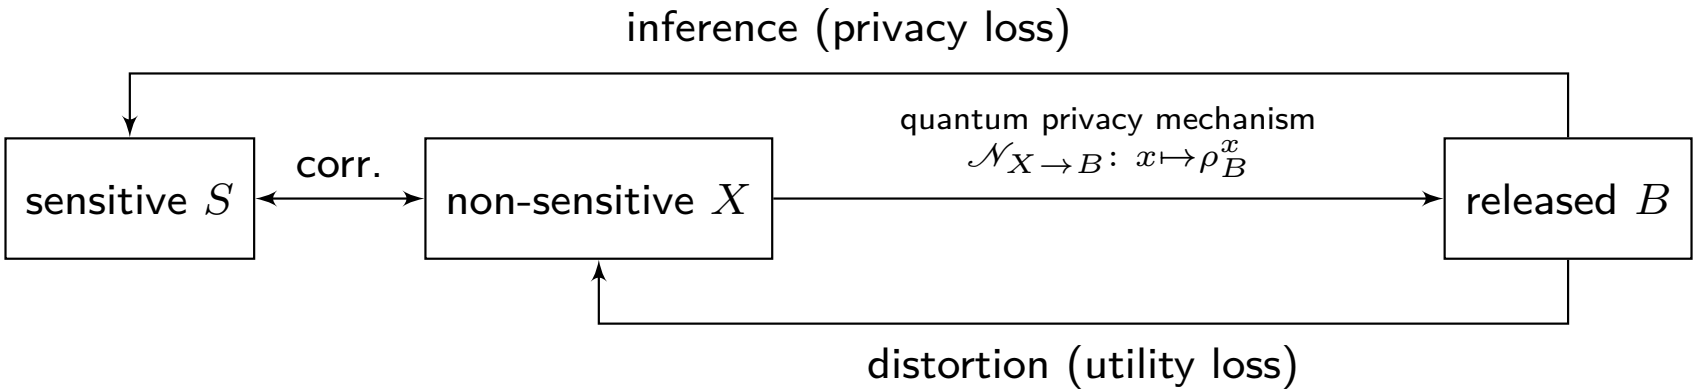

Let S be any function of X. The maximal α-leakage is defined as

$$\mathcal{L}_\alpha^{\max}(X\to B)_\rho := \sup_{p_{S|X}:S-X-B} \mathcal{L}_\alpha(S\to B)_\rho.$$

Characterization 3: Maximal α-Leakage

Theorem

For $\alpha \in [1, \infty]$, the maximal α-leakage can be expressed as

$$\mathcal{L}_\alpha^{\max}(X \to B)_\rho = \begin{cases} C_\alpha^{\mathbb{M}}(\mathcal{N}_{\text{supp}(p_X)\to B}) = R_\alpha^{\mathbb{M}}(\mathcal{N}_{\text{supp}(p_X)\to B}), & \alpha > 1; \\ I_1^{\mathrm{A},\mathbb{M}}(X:B)_\rho = \mathcal{L}_1(X \to B)_\rho, & \alpha = 1. \end{cases}$$

- **Measured Rényi capacity**: $C_\alpha^{\mathbb{M}}(\mathcal{N}_{\mathcal{X}\to B}) := \sup_{p_X} I_\alpha^{\mathbb{M}}(X:B)_\rho$
- Measured Rényi divergence radius
  $R_\alpha^{\mathbb{M}}(\mathcal{N}_{\mathcal{X}\to B}) := \inf_{\sigma_B} \sup_{x\in\mathcal{X}} D_\alpha^{\mathbb{M}}(\mathcal{N}(x)\|\sigma_B) \equiv \inf_{\sigma_B} \sup_{x\in\mathcal{X}} D_\alpha^{\mathbb{M}}(\rho_B^x\|\sigma_B)$
- Measured Rényi information: $I_\alpha^{\mathbb{M}}(X:B)_\rho := \inf_{\sigma_B} D_\alpha^{\mathbb{M}}(\rho_{XB}\|\rho_X \otimes \sigma_B)$
- For $\alpha > 1$, $\mathcal{L}_\alpha^{\max}(X \to B)_\rho$ only depends on X through its support.

Properties of Maximal α-Leakage (1/2)

$$\mathcal{L}_\alpha^{\max}(X \to B)_\rho = \begin{cases} \sup\limits_{\bar{p}_X \in \mathcal{P}(\mathrm{supp}(p_X))} \inf\limits_{\sigma_B \in \mathcal{S}(\mathcal{H}_B)} D_\alpha^{\mathbb{M}}(\bar{\rho}_{XB} \| \bar{\rho}_X \otimes \sigma_B), \ \alpha > 1; \\ I_1^{\mathrm{A},\mathbb{M}}(X:B)_\rho, \ \alpha = 1. \end{cases}$$

- concave program of $\bar{p}_X$; concave function of input p_X
- quasi-convex in ρ_B^x given $\bar{p}_X$ or p_X
- non-decreasing in α
- satisfies data-processing inequality (DPI)
- upper bound and lower bound:

$$0 \le \mathcal{L}_\alpha^{\max}(X \to B)_\rho \le \begin{cases} \log(|\mathrm{supp}(p_X)|), \ \alpha > 1; \\ H_1(X)_p, \ \alpha = 1. \end{cases}$$

Properties of Maximal α-Leakage: Composition Property (2/2)

- $B_1 - X - B_2$ form a Markov chain
- Quantum privacy mechanisms: $\mathcal{N}_{\mathcal{X}\to B_1} : x \mapsto \rho^x_{B_1}$, $\mathcal{N}_{\mathcal{X}\to B_2} : x \mapsto \rho^x_{B_2}$
- For $\alpha \in [1, \infty]$, if the adversary can access more than one released version B_1, B_2 of the non-sensitive data X

Theorem

$$\mathcal{L}_\alpha^{\max}(X \to B_1, B_2)_\rho \leq \mathcal{L}_\alpha^{\max}(X \to B_1)_\rho + \mathcal{L}_\alpha^{\max}(X \to B_2)_\rho$$

Asymptotic Behavior of Information Leakage

- So far, we discussed about $\mathcal{L}_\alpha$ and $\mathcal{L}_\alpha^{\max}$ under $\mathscr{N}_{\mathcal{X}\to B}$ in the **one-shot** setting $\Rightarrow$ what are their asymptotic behaviors?
- Adversary: cannot be restricted to attack in the i.i.d. sense
- System designer: arguably easier to employ i.i.d. quantum privacy mechanisms

$$\mathscr{N}_{\mathcal{X}\to B}^{\otimes n} : x_1 x_2 \cdots x_n \mapsto \rho_{B_1}^{x_1} \otimes \rho_{B_2}^{x_2} \otimes \cdots \otimes \rho_{B_n}^{x_n} =: \rho_{B^n}^{x^n}.$$

- **Regularization**: average information leakage as $n \to \infty$

Characterization 4: Regularized α-Leakage

Theorem

Let $\rho_{XB} = \sum_{x\in\mathcal{X}} p_X(x)|x\rangle\langle x| \otimes \rho_B^x$. For $\alpha \geq 1$, then regularized α-leakage from i.i.d. X^n to B^n under $\mathcal{N}_{\mathcal{X}\to B}^{\otimes n}$ is given by

$$\lim_{n\to\infty} \frac{1}{n} \mathcal{L}_\alpha(X^n \to B^n)_{\rho^{\otimes n}} = I_\alpha^* (X:B)_{\rho^{(\alpha)}} .$$

- Sandwiched Rényi information: $I_\alpha^* (X:B)_\rho := \inf_{\sigma_B} D_\alpha^*(\rho_{XB} \| \rho_X \otimes \sigma_B)$
- Denote $\rho_{XB}^{(\alpha)} \equiv \sum_{x\in\mathcal{X}} p_X^{(\alpha)}(x) \, |x\rangle\langle x| \otimes \rho_B^x$
- α-tilted distribution: $p_X^{(\alpha)}(x) := \frac{p_X(x)^\alpha}{\sum_{x\in\mathcal{X}} p_X(x)^\alpha}$
- (Char. 2) $\mathcal{L}_\alpha(X \to B)_\rho = I_\alpha^{\mathrm{A},\mathbb{M}}(X:B)_\rho \neq I_\alpha^{\mathbb{M}}(X:B)_\rho$

Characterization 5: Regularized Maximal α-Leakage

Theorem

For any $\alpha > 1$, the regularized maximal α-leakage from X^n with any arbitrary distribution $p_{X^n} \in \mathcal{P}(\mathcal{X}^n)$ that has full support to B^n under $\mathcal{N}_{\mathcal{X}\to B}^{\otimes n}$ is given by

$$\lim_{n\to\infty} \frac{1}{n} \mathcal{L}_\alpha^{\max}(X^n \to B^n)_{\rho^n} = C_\alpha^*(\mathcal{N}_{\mathcal{X}\to B}) = R_\alpha^*(\mathcal{N}_{\mathcal{X}\to B}).$$

- Sandwiched Rényi capacity: $C_\alpha^*(\mathcal{N}_{\mathcal{X}\to B}) := \sup_{p_X \in \mathcal{P}(\mathcal{X})} I_\alpha^*(X:B)_\rho$
- Sandwiched Rényi divergence radius:
  $R_\alpha^*(\mathcal{N}_{\mathcal{X}\to B}) := \inf_{\sigma_B} \sup_{x\in\mathcal{X}} D_\alpha^*(\mathcal{N}(x)\|\sigma_B) \equiv \inf_{\sigma_B} \sup_{x\in\mathcal{X}} D_\alpha^*(\rho_B^x\|\sigma_B)$
- Denote $\rho_{X^nB^n}^n := \sum_{x^n\in\mathcal{X}^n} p_{X^n}(x^n)|x^n\rangle\langle x^n| \otimes \rho_{B^n}^{x^n}$

Takeaways

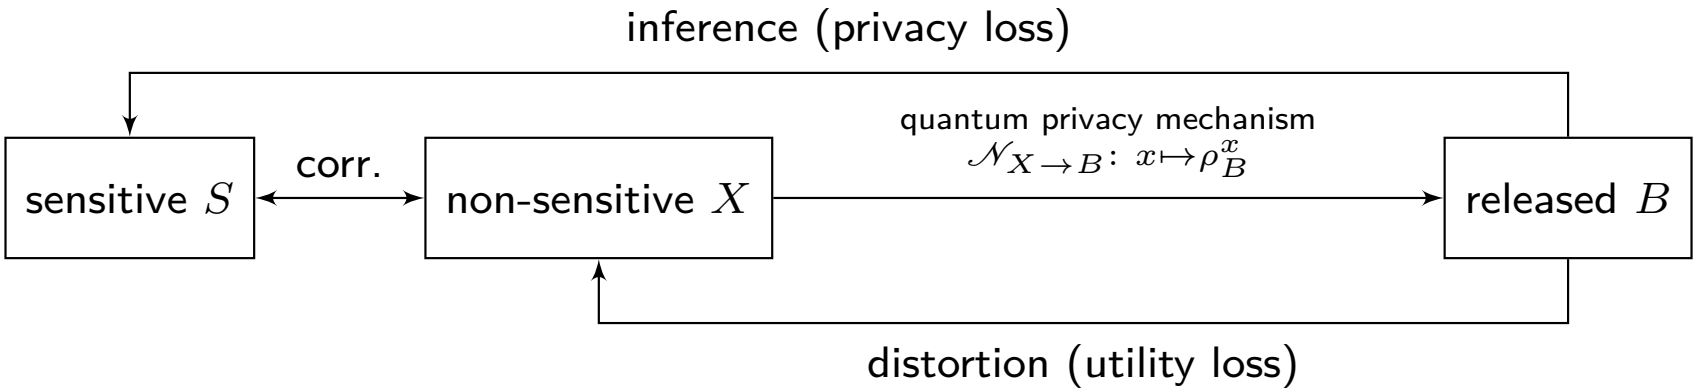

- $\mathcal{L}_\alpha^{\max}$ provided operational interpretation for measured Rényi capacity when $\alpha > 1$
- We characterized
  - one-shot: $\mathsf{P}_\alpha(X|B)_\rho$, $\mathcal{L}_\alpha$, $\mathcal{L}_\alpha^{\max}$
  - asymptotic: $\lim_{n\to\infty}\frac{1}{n}\mathcal{L}_\alpha$, $\lim_{n\to\infty}\frac{1}{n}\mathcal{L}_\alpha^{\max}$
- **Open problem**:
  - fully quantum case $S-A-B$?
  - Privacy-utility tradeoff (PUT)