# QUANTUM PUFFERFISH PRIVACY: A FLEXIBLE PRIVACY FRAMEWORK FOR QUANTUM SYSTEMS

THESHANI NURADHA

CORNELL INFORMATION THEORY DAY 2023

Joint work with Ziv Goldfeld and Mark M. Wilde

arXiv:2306.13054

# PROTECTION FOR QUANTUM DATA

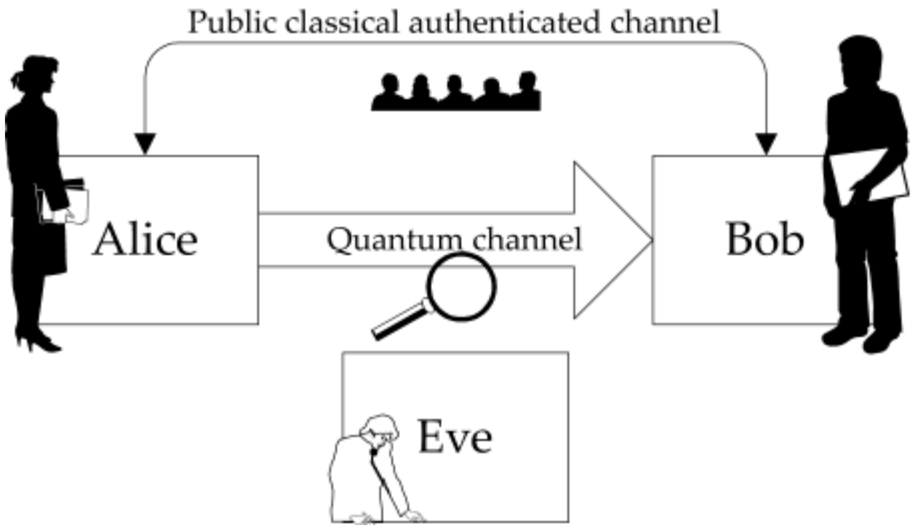


- Quantum cryptography is well studied
- What happens if we want Bob to know only certain aspects of Alice's data?
- This is captured by "statistical privacy frameworks"

# STATISTICAL PRIVACY FRAMEWORKS (CLASSICAL)

- Differential privacy:
  - Answering aggregate queries about a database while keeping individual records private
- Its limitations:
  - Accounts for one type of private information only—records of individual users
  - Does not allow encoding domain knowledge into the framework
- Pufferfish Privacy:
  - Customizing which information is regarded as private
  - Explicitly integrates distributional assumptions

# MATHEMATICAL OBJECTS

- Quantum state: PSD operator with unit trace
- Quantum channel: Completely positive trace preserving map
- Quantum measurement:
  - Measurement operator: $0 \leq M \leq I$
  - Positive operator valued measure (POVM): collection of PSD operators $\{M_y\}_{y \in \mathcal{Y}}$ such that $\sum_{y \in \mathcal{Y}} M_y = I$
- Born rule: Probability of observing outcome $y = \text{Tr}[M_y \rho]$

# QUANTUM DIFFERENTIAL PRIVACY (QDP)

$\rho \sim \sigma$ Neighboring relation: e.g., closeness in trace distance

QDP –[ZY'17], [HRF'22]

$\mathcal{A}$ is $(\varepsilon, \delta)$-QDP if

$$\mathrm{Tr}\left[M\mathcal{A}(\rho)\right] \leq e^{\varepsilon}\,\mathrm{Tr}\left[M\mathcal{A}(\sigma)\right] + \delta$$

for every measurement operator $M$ and all $\rho \sim \sigma$.

QDP guarantees that all pairs of states that are classified as neighbors are approximately indistinguishable, i.e., cannot be identified under all possible measurements.

# Beyond QDP: Need for More

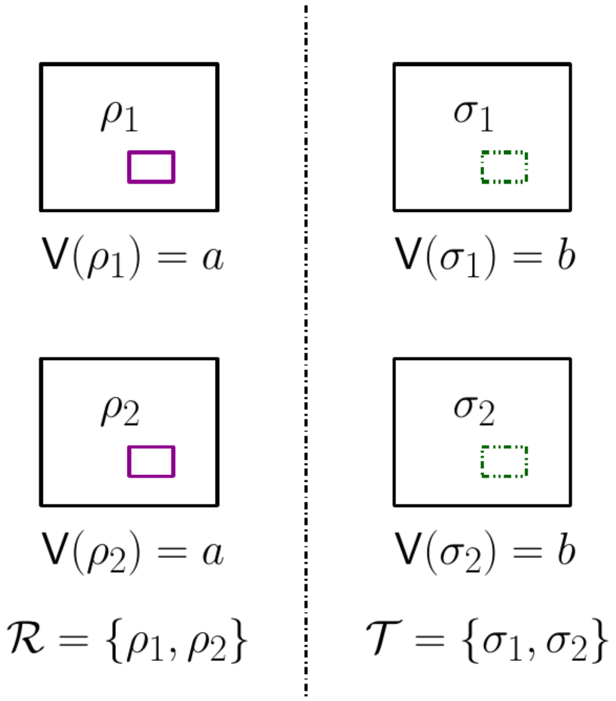


Flexible Secrets: Secrets containing collection of states

Domain Knowledge: Likelihood of observing different states

Relaxing Worst-case Measurements: Physical limitations of LOCC vs joint measurements

Goal: A Flexible Privacy Framework for Quantum Systems

# QUANTUM PUFFERFISH PRIVACY FRAMEWORK

Ingredients:

$\mathcal{S}$ Set of potential secrets $\mathcal{S} = \bigcup_{i=1}^{n} \mathcal{T}_i \quad \mathcal{T}_i = \{\rho \in \mathcal{D}(\mathcal{H}) : \mathsf{V}(\rho) = a_i\}$

Values the secret function can take

$\mathcal{Q}$ Set of discriminative pairs — Which pairs of secrets to be indistinguishable

Symmetric: $(\mathcal{R}, \mathcal{T}) \in \mathcal{Q}$ iff $(\mathcal{T}, \mathcal{R}) \in \mathcal{Q}$

Hiding different values

$\mathcal{Q} = \bigcup_{i \neq j} \{(\mathcal{T}_i, \mathcal{T}_j)\}$

$\Theta$ Set of data distributions $X \sim P_X \in \Theta$ $\quad \rho^X$ models a density operator that is randomly chosen according to $P_X$

$\mathcal{M}$ Set of possible measurements — Subset of measurements possible under physical, legal, or ethical constraints

# QPP DEFINITION

$\mathcal{A}$ is $(\varepsilon, \delta)$- QPP in the framework $(\mathcal{S}, \mathcal{Q}, \Theta, \mathcal{M})$ if for all $P_X \in \Theta$, $(\mathcal{R}, \mathcal{T}) \in \mathcal{Q}$ with $P_X(\mathcal{R}), P_X(\mathcal{T}) > 0$, and all $M \in \mathcal{M}$,

$$\mathrm{Tr}\big[M\mathcal{A}(\rho^{\mathcal{R}})\big] \leq e^{\varepsilon}\mathrm{Tr}\big[M\mathcal{A}(\rho^{\mathcal{T}})\big] + \delta$$

Conditional average states

$$\rho^{\mathcal{R}} := \sum_{\{x:\rho^x \in \mathcal{R}\}} q_{\mathcal{R}}(x)\rho^x$$

$$q_{\mathcal{R}}(x) := \frac{P_X(x)}{P_X(\mathcal{R})}$$

$$P_X(\mathcal{R}) := \sum_{\{x:\rho^x \in \mathcal{R}\}} P_X(x)$$

Semantic meaning:

For a state $\rho^X$ chosen according to $X \sim P_X \in \Theta$ and input to the quantum channel $\mathcal{A}$, an adversary applying measurement $M \in \mathcal{M}$ on the channel output $\mathcal{A}(\rho^X)$ draws the same conclusions regardless of whether $\rho^X$ belongs to $\mathcal{R}$ or $\mathcal{T}$

# OTHER PRIVACY FRAMEWORKS WHICH ARE SPECIAL CASES

By choosing specific ingredients to QPP

- Quantum differential privacy
- Classical pufferfish privacy
- Utility optimized privacy models

# DATTA-LEDITZKY (DL) DIVERGENCE

$$\overline{D}^{\delta}(\rho\|\sigma) = \ln\inf\left\{\lambda \geq 0 : \mathrm{Tr}[(\rho - \lambda\sigma)_+] \leq \delta\right\}$$

Positive eigenspace

$$(A)_+ := \sum_{i:a_i \geq 0} a_i |i\rangle\langle i|$$

Equivalent formulation of QPP (for all possible measurements)

$$\sup_{\Theta,(\mathcal{R},\mathcal{T})\in\mathcal{Q}} \overline{\mathrm{D}}^{\delta}\left(\mathcal{A}(\rho^{\mathcal{R}})\middle\|\mathcal{A}(\rho^{\mathcal{T}})\right) \leq \varepsilon$$

$$\mathrm{Tr}\left[M\mathcal{A}(\rho^{\mathcal{R}})\right] \leq e^{\varepsilon}\mathrm{Tr}\left[M\mathcal{A}(\rho^{\mathcal{T}})\right] + \delta$$

Operational Interpretation:

minimal $\varepsilon$ that can be achieved for fixed $\delta$ via the indistinguishability condition of the QPP framework

# DL DIVERGENCE

As a Semi-Definite Program:

$$\overline{D}^{\delta}(\rho\|\sigma) = \ln \inf_{\lambda, Z \geq 0} \{\lambda : \mathrm{Tr}[Z] \leq \delta,\ Z \geq \rho - \lambda\sigma\}$$

$$= \ln \sup_{\mu, W \geq 0} \{\mathrm{Tr}[W\rho] - \mu\delta : \mathrm{Tr}[W\sigma] \leq 1,\ W \leq \mu I\}$$

Properties:

Data processing: For every positive trace preserving map $\overline{D}^{\delta}(\rho\|\sigma) \geq \overline{D}^{\delta}\big(\mathcal{N}(\rho)\|\mathcal{N}(\sigma)\big)$

Joint-quasi convexity:
$$\overline{D}^{\delta}\left(\sum_{i=1}^{k} p_i\rho_i \middle\| \sum_{i=1}^{k} p_i\sigma_i\right) \leq \max_i \overline{D}^{\delta}(\rho_i\|\sigma_i)$$

Quasi subadditivity:

$$\overline{D}^{\delta_1'+\delta_2'}(\rho_1 \otimes \rho_2\|\sigma_1 \otimes \sigma_2) \leq \overline{D}^{\delta_1}(\rho_1\|\sigma_1) + \overline{D}^{\delta_2}(\rho_2\|\sigma_2) - \ln\big((1-\delta_1)(1-\delta_2)\big)$$

with $\delta_i' := \sqrt{\delta_i(2-\delta_i)} \in (0,1)$

# PROPERTIES OF QPP

- Post-processing:
  - Passing the output of a QPP mechanism through a channel still preserves QPP

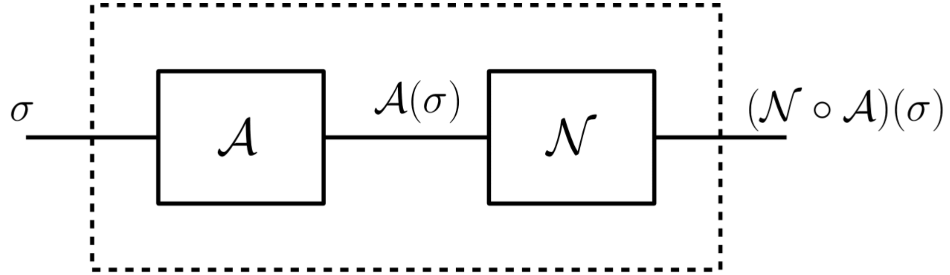


- Convexity:
  - Applying a QPP mechanism that is randomly chosen from a given set of such mechanisms still satisfies QPP
- Composability:
  - Parallel- QPP holds after applying composed mechanism to the input $\rho^{X_1} \otimes \rho^{X_2} \otimes \cdots \otimes \rho^{X_k}$
  - Adaptive- Each subsequently composed mechanism is chosen based on the outputs of the preceding ones
  - Correlated input states

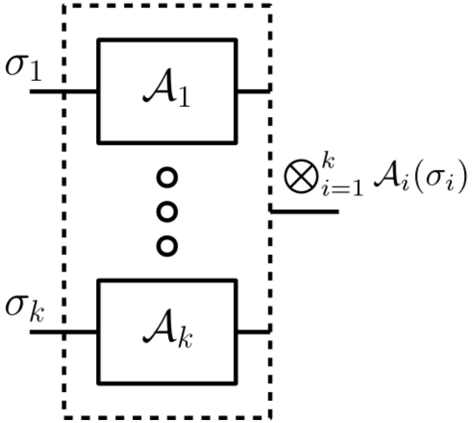


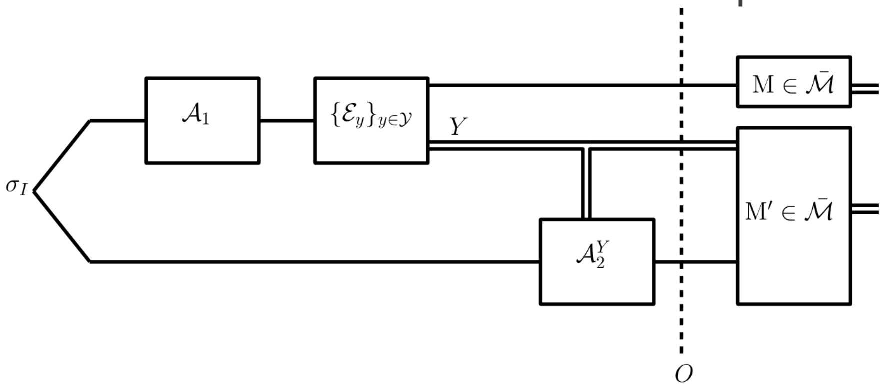


★ For all possible measurements: proof follows from properties of DL divergence

## Parallel Composability: (K=2)

With product measurements (semi-classical) $(\varepsilon_1 + \varepsilon_2, \delta_1 + \delta_2)$- QPP

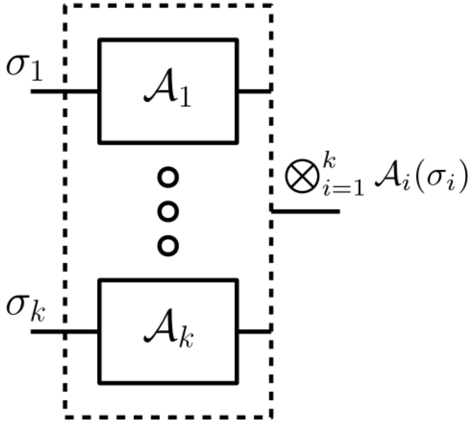


With all measurements including joint measurements $(\varepsilon', \delta')$- QPP

$$\varepsilon' := \varepsilon_1 + \varepsilon_2 + \ln\left(\frac{1}{(1-\delta_1)(1-\delta_2)}\right)$$
$$\delta' := \sqrt{\delta_1(2-\delta_1)} + \sqrt{\delta_2(2-\delta_2)}$$

★ Distinction between classical and quantum cases: joint measurements can infer more information and thus privacy degrades

# MECHANISMS

Depolarization mechanism

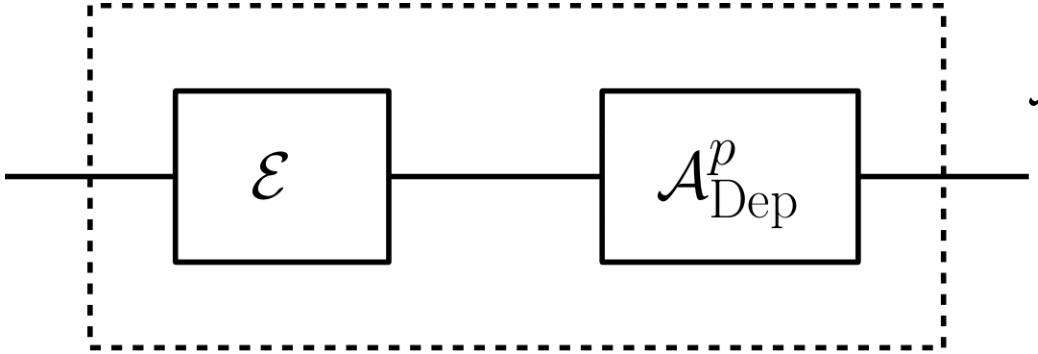

$$\mathcal{A}^p_{\mathrm{Dep}}(\rho) := (1-p)\rho + \frac{p}{d} I$$

$\mathcal{A}^p_{\mathrm{Dep}}(\mathcal{E}(\cdot))$ is $\varepsilon$-QPP if

$$p \geq \frac{dK}{dK + e^{\varepsilon} - 1}$$

$$K := \sup_{M \in \mathcal{M}} \frac{\|M\|_\infty}{\mathrm{Tr}[M]} \times \sup_{\Theta,(\mathcal{R},\mathcal{T})\in\mathcal{Q}} \frac{\left\|\mathcal{E}(\rho^{\mathcal{R}}) - \mathcal{E}(\rho^{\mathcal{T}})\right\|_1}{2}$$

# AUDITING PRIVACY

- Aims to detect violations in privacy guarantees and reject incorrect algorithms
- In classical settings: translate the privacy requirement to a weaker privacy notion that is efficiently computable
  - Not satisfying relaxed notion implies that original requirement is violated
- The pitfall of this approach is the impossibility of quantifying the gap between the original and relaxed privacy notions

Goal: Auditing without translating to a relaxed privacy notion

# AUDITING QPP

- Using SDPs for DL divergence and equivalent form: Runtime polynomial in dimension, but exponential in number of qubits

- Trace distance estimation techniques and equivalent formulation via hockey-stick divergence:
  - Equivalent form for QDP: $\sup_{\rho\sim\sigma} \mathsf{E}_{e^{\varepsilon}}(\mathcal{A}(\rho)\|\mathcal{A}(\sigma)) \leq \delta$ $\quad$ $\mathsf{E}_{\gamma}(\rho\|\sigma) := \mathrm{Tr}[(\rho - \gamma\sigma)_{+}]$
  - Hockey stick divergence $\mathsf{E}_{\gamma}(\rho\|\sigma) = \frac{1}{2}\|\rho - \gamma\sigma\|_1 + \frac{(1-\gamma)}{2}$
  - Use of quantum algorithms to estimate trace distance

- Hypothesis testing based auditing pipeline: Formal Guarantees on Type-I error

# SUMMARY

Contributions:

- Proposed notion of QPP provides a flexible privacy framework for quantum systems
- An operational interpretation of DL divergence
  - Study properties of QPP mechanisms
  - Characterize privacy-utility tradeoffs
- Mechanisms via depolarization channel
- Methodology to audit quantum privacy
- Variants of QPP
- Connections to information-theoretic tools and quantum fairness

Thank you!