2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

Quantum Privacy and Hypothesis-Testing

Farhad Farokhi

Abstract— **A novel definition for quantum privacy based on hypothesis testing is presented. This privacy notion possesses an operational interpretation based on the false positive and negative rates of adversaries attempting to distinguish among private categories to which quantum states belong using arbitrary measurements. Important properties of post processing and composition are established. The relationship between privacy against hypothesis-testing adversaries and quantum differential privacy are then examined. This enables analysis of hypothesis testing under differential privacy and in noisy quantum systems. It also provides an operational interpretation for quantum differential privacy.**

I. Introduction

Quantum computing has garnered sizeable attention due to speedups in several classically-difficult problems, such as factorising [1]. These breakthroughs and the added attention have paved the way for development of new algorithms in big-data processing and quantum machine learning [2]–[4]. However, data processing can result in unintended information leakage [5]. This is an important issue because, as quantum hardware becomes more commercially available, quantum algorithms can be implemented on real-world sensitive, private, or proprietary datasets. Therefore, there is a need to better understand private information leakage in quantum computing and to construct privacy-preserving algorithms.

In classical computing literature, differential privacy has become the gold standard of privacy analysis and private algorithm design [6]–[8]. This is often attributed to the fact that differential privacy makes minimal assumptions on data (e.g., range rather than distribution) and adversary (i.e., computationally unbounded attacker). It also meets important properties of post processing and compositions [9]. Although possessing powerful guarantees, differential privacy has been polarizing [10]–[12]. Criticisms surroundings conservativeness of differential privacy have motivated studies on information-theoretic privacy to better handle privacy-utility trade-off [13]–[17]. Adoption of hypothesis-testing and estimation-based adversaries have been proposed as less conservative alternatives to differential privacy by social scientists following implementation of differential privacy in the 2020 US Census [10]. Nonetheless, differential privacy has been recently extended to quantum computing [18]–[20]. Little attention has been paid to other forms of privacy in quantum systems. In this paper, we investigate privacy against hypothesis-testing adversaries. This is of particular interest to us due to multiple reasons. First, this notion of privacy provides an operational, intuitive measure of privacy risk. Second, by investigating its relationship with quantum differential privacy, we provide an operational meaning to quantum differential privacy and study hypothesis testing under quantum differential privacy. Finally, by establishing the effect of depolarizing channels on privacy against hypothesis testing adversaries, we shed light on hypothesis testing in noisy quantum systems.

We particularly propose a novel definition for quantum data privacy based on hypothesis testing. This notion of privacy possesses an operational interpretation, specifically for general lay-users, based on false positive and negative rates of a computationally-unbounded adversary in distinguishing private classes to which quantum states belong (e.g., diagnosis of a disease in health datasets or belonging to training dataset in membership inference attacks in quantum machine learning) based on arbitrary measurements. We prove two important properties for the new notion of privacy: post processing and composition. These properties are highly sought-after in privacy definitions [18] and information leakage metrics [13]. Subsequently, we investigate the relationship between privacy against hypothesis-testing adversaries and quantum differential privacy. This enables us to provide an interpretation for parameters of differential privacy based on its relationship with privacy against hypothesis-testing adversaries. This also enables analysis of hypothesis testing under quantum differential privacy and in noisy quantum systems (by establishing privacy of depolarizing channels).

The remainder of this paper is organized as follows. We provide a review of basic concepts in quantum computing and information in Section II. The definition and results on privacy against hypothesis-testing adversaries is presented in Section III. Section IV presents quantum differential privacy and its relationship with privacy against hypothesis-testing adversaries. Finally, we present some concluding remarks and future directions for research in Section V.

II. Quantum states and channels

The definitions and preliminary results in this section are borrowed from [21]. When the results or definitions are from outside this source, appropriate citations are presented.

A quantum system is modelled by a Hilbert space $\mathcal{H}$, i.e., a complex vector space, equipped with an inner product, that is complete with respect to the norm defined by the inner product. Dirac's notation is used to denote quantum states. That is, a *pure quantum state*, which is a vector in Hilbert space $\mathcal{H}$ with unit norm, is denoted by 'ket' $|\cdot\rangle$, e.g., $|\psi\rangle \in \mathcal{H}$. The inner product of two states $|\phi\rangle$ and $|\psi\rangle$ is denoted by $\langle\phi|\psi\rangle$. Here, 'bra' $\langle\psi|$ is used to refer to conjugate transpose of $|\psi\rangle$ and $\langle\phi|\psi\rangle := \langle\phi|\ |\psi\rangle \in \mathbb{C}$.

The basic element of interest in quantum information theory is a quantum bit, which is referred to as *qubit*. A qubit

is a vector in a 2-dimensional Hilbert space. Any qubit can be written in terms of the so-called computational basis $|0\rangle$ and $|1\rangle$ that form an orthonormal basis for the 2-dimensional Hilbert space, i.e., $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ with $\alpha, \beta \in \mathbb{C}$ such that $|\alpha|^2 + |\beta|^2 = 1$. Combination of two qubits $|\phi\rangle$ and $|\psi\rangle$ is denoted by their tensor product $|\phi\rangle \otimes |\psi\rangle$, where $\otimes$ is the Kronecker or tensor product. For the sake of brevity, we sometimes refer to $|\phi\rangle \otimes |\psi\rangle$ as $|\phi\rangle|\psi\rangle$ or $|\phi, \psi\rangle$. When $|\phi\rangle$ and $|\psi\rangle$ belong to or assigned to two distinct registers (users) A and B (e.g., used by two separate parties), and this is either unclear from the context or must be emphasized, we write $|\phi\rangle_A \otimes |\psi\rangle_B$ or $|\phi\rangle_A |\psi\rangle_B$. A quantum gate is a unitary operator, e.g., U such that $U^\dagger U = I$ with $U^\dagger$ denoting conjugate transpose of U, that acts on quantum states.

A *mixed quantum state* is represented by ensemble $\{(p_1, |\psi_1\rangle), \ldots, (p_k, |\psi_k\rangle)\}$ such that $p_i \geq 0, \forall i \in [k] := \{1, \ldots, k\}$ and $\sum_{i\in[k]} p_i = 1$, i.e., the quantum system is in pure state $|\psi_i\rangle$ with probability p_i. The density operator corresponding to ensemble $\{(p_1, |\psi_1\rangle), \ldots, (p_k, |\psi_k\rangle)\}$ is $\rho := \sum_{i\in[k]} p_i |\psi_i\rangle\langle\psi_k|$. By construction, $\mathrm{tr}(\rho) = 1$. A pure quantum states $|\phi\rangle$ can be modelled using rank-one density operator $\rho = |\phi\rangle\langle\phi|$. Combination of two density operators ρ and σ is denoted by their tensor product $\rho \otimes \sigma$.

A basic operation in quantum systems is measurement, which enables extraction of information about the states. A measurement is modelled by a set of operators $M = \{K_i\}_{i\in[m]}$ with normalization constraint that $\sum_{i\in[m]} K_i^\dagger K_i = I$. By performing measurement M on a quantum system with state ρ, we observe output $i \in [m]$ with probability $\mathrm{tr}(K_i \rho K_i^\dagger)$ in which case the state of the quantum system collapses to $K_i \rho K_i^\dagger / \mathrm{tr}(K_i \rho K_i^\dagger)$ after measurement. When the post-measurement state of the quantum system is of no interest, we can use the positive operator-valued measure (POVM) framework, which is a set of positive semi-definite Hermitian matrices $F = \{F_i\}_{i\in[m]}$ such that $\sum_{i\in[m]} F_i = I$. In this case, the probability of obtaining output $i \in [m]$ when taking a measurement on a system with quantum state ρ is $\mathrm{tr}(\rho F_i) = \mathrm{tr}(F_i \rho)$.

A quantum channel is a mapping from the space of density operators to potentially another space of density operators that is both completely positive and trace preserving. Quantum channels model open quantum systems, i.e., quantum systems that interact with environment, and thus can model noisy quantum behaviours. According to Choi-Kraus theorem [21, Theorem 4.4.1], for each quantum channel $\mathcal{E}$, there exists a family of linear operators $\{E_j\}_{j\in[n]}$ for some $n \in \mathbb{N}$ such that $\sum_j E_j^\dagger E_j = I$ and $\mathcal{E}(\rho) = \sum_{j\in[n]} E_j \rho E_j^\dagger$ for all density operators ρ. This is referred to as the Kraus representation of quantum channels. For instance, a quantum gate with unitary operator U can be represented by $\mathcal{E}(\rho) = U\rho U^\dagger$. Also, if we discard or delete the outcome of measurement $M = \{K_i\}_{i\in[m]}$, the quantum state transition can be modelled by quantum channel $\mathcal{E}(\rho) = \sum_{i\in[k]} K_i \rho K_i^\dagger$. We define the tensor product of quantum channels $\mathcal{E}_1$ and $\mathcal{E}_2$ as $\mathcal{E}_1 \otimes \mathcal{E}_2(\rho_1 \otimes \rho_2) := \mathcal{E}_1(\rho_1) \otimes \mathcal{E}_2(\rho_2)$ for all density operators ρ_1 and ρ_2.

The trace norm or Schatten 1-norm of any linear operator M is defined as $\|M\|_1 := \mathrm{tr}(|M|) = \mathrm{tr}(\sqrt{M^\dagger M})$. Based on this, we can define the trace distance between any two density operators ρ and σ with $\mathcal{T}(\rho, \sigma) := \frac{1}{2}\|\rho - \sigma\|_1 \in [0, 1]$. Recall that density operators belong to the set of linear operators (i.e., matrices). The distance is equal to zero when two quantum states are equal. However, the distance attains its maximum value when two quantum states have support on orthogonal subspaces. For $\upsilon \in [0, 1]$, the υ-relative entropy between two quantum states ρ and σ is defined as $D^\upsilon(\rho\|\sigma) = -\log\left(\min\{\mathrm{tr}(Q\sigma) | 0 \preceq Q \preceq I, \mathrm{tr}(Q\rho) \geq 1-\upsilon\}\right)$. The υ-relative entropy satisfies a few important properties that we will use in this paper. These properties are borrowed from [22]. First, $D^\upsilon(\rho\|\sigma) \geq 0$ with equality if $\rho = \sigma$ and $\upsilon = 0$. Second, υ-relative entropy enjoys data processing inequality, i.e., $D^\upsilon(\mathcal{E}(\rho)\|\mathcal{E}(\sigma)) \leq D^\upsilon(\rho\|\sigma)$ for all density operators ρ, σ and all quantum channels $\mathcal{E}$. Also, $D^\upsilon(\rho\|\sigma) \leq (S(\rho\|\sigma) + H_b(\upsilon))/(1-\upsilon)$, where $H_b(\upsilon) = -\upsilon\log(\upsilon) - (1-\upsilon)\log(1-\upsilon)$ is the binary entropy function and $S(\rho\|\sigma) := \mathrm{tr}(\rho(\log(\rho) - \log(\sigma)))$ is the usual relative entropy in quantum information theory. The υ-relative entropy and the trace distance also satisfy the following relationship $\nu/(1-\nu)\|\rho-\sigma\|_1 \leq D^\upsilon(\rho\|\sigma)$ [23]. The smooth max-relative entropy is defined as $D^\upsilon_{\max}(\rho\|\sigma) = \inf_{\tau\in\mathcal{B}^\upsilon(\rho)} D_{\max}(\tau\|\sigma)$, where $D_{\max}(\tau\|\sigma) = \inf\{\lambda \geq 0 \,|\, \rho \preceq \exp(\lambda)\sigma\}$ and $\mathcal{B}^\upsilon(\rho) := \{\tau \,|\, \tau^\dagger = \tau \succeq 0, \|\rho - \tau\|_1 \leq 2\upsilon\}$.

Global depolarizing channel is an important type of quantum noise that is represented by

$$\mathcal{E}_{\mathrm{Dep}}(\rho) := \frac{p}{D} I + (1-p)\rho, \tag{1}$$

where D is the dimension of the Hilbert space to which the system belongs and $p \in [0, 1]$ is a probability parameter.

III. Quantum Hypothesis Testing and Privacy

Consider a quantum hypothesis testing scenario where a decision maker aims to distinguish between two quantum states ρ (null hypothesis) and σ (alternative hypothesis). This is done by performing POVM $M := \{M_1, M_2\}$ with $M_1 + M_2 = I$ and $0 \preceq M_i \preceq I$ for $i = 1, 2$. If measurement outcome corresponding to the operator M_1 is realized, the decision maker guesses that the state is ρ while, if measurement outcome corresponding to the operator M_2 is realized, the decision maker guesses that the state is σ. The probability of a type-I error (false positive) is equal

$$\alpha(M_2) := \mathrm{tr}(M_2\rho). \tag{2}$$

The probability of a type-II error (false negative) is given by

$$\beta(M_1) := \mathrm{tr}(M_1\sigma). \tag{3}$$

The optimal test, which seeks to minimize the false negative probability subject to a constraint on maintaining the false positive probability below $\eta \in [0, 1]$, is given by

$$\beta_\eta(\rho, \sigma) := \min_{M_1, M_2 \succeq 0} \beta(M_1), \tag{4a}$$

$$\text{s.t.} \quad M_1 + M_2 = I, \tag{4b}$$

$$\alpha(M_2) \leq \eta. \tag{4c}$$

This is referred to as *asymmetric quantum hypothesis testing* [24]. The following well-known result (see, e.g., [22]) can be easily derived based on the definition of $\beta_\eta(\rho,\sigma)$ and η-relative entropy $D^\eta(\rho\|\sigma)$.

Proposition 1: $\beta_\eta(\rho,\sigma)=2^{-D^\eta(\rho\|\sigma)}$.

Proof: Note that

$$\begin{aligned}\beta_\eta(\rho,\sigma)&=\min_{M_1,M_2\succeq 0}\{\mathrm{tr}(M_1\sigma)|M_1+M_2=I,\mathrm{tr}(M_2\rho)\leq\eta\}\\&=\min_{I\succeq M_1\succeq 0}\{\mathrm{tr}(M_1\sigma)|\,\mathrm{tr}((I-M_1)\rho)\leq\eta\}\\&=\min_{I\succeq M_1\succeq 0}\{\mathrm{tr}(M_1\sigma)|1-\eta\leq\mathrm{tr}(M_1\rho)\}\\&=2^{-D^\eta(\rho\|\sigma)}.\end{aligned}$$

This concludes the proof. ∎

Alternatively, a combination of false positive and false negative probabilities can be minimized:

$$\begin{aligned}p_{\mathrm{err}}(\rho,\sigma):=\min_{M_1,M_2\succeq 0}\ &p_\rho\alpha(M_2)+p_\sigma\beta(M_1),\\ \text{s.t.}\quad &M_1+M_2=I,\end{aligned}$$

where $p_\rho\in[0,1]$ and $p_\sigma\in[0,1]$, respectively, denote the prior probability that quantum state ρ and the prior probability that quantum state σ are prepared. By construct, $p_\rho+p_\sigma=1$. This is referred to as *symmetric quantum hypothesis testing* [24].

Theorem 1 (Helstrom-Holevo theorem [21, p. 254-255]) $p_{\mathrm{err}}(\rho,\sigma)=\frac{1}{2}\left(1-\|p_\rho\rho-p_\sigma\sigma\|_1\right)$.

The most indistinguishable quantum states are $\rho=\sigma$. In this case, a decision maker would not be able to identify the quantum states *because their observables are equivalent*. Therefore, we can define

$$p_{\max}:=p_{\mathrm{err}}(\rho,\rho)=\frac{1}{2}\left(1-|p_\rho-p_\sigma|\right).$$

Therefore, for general density operators, we have

$$p_{\mathrm{err}}(\rho,\sigma)=p_{\max}+\frac{1}{2}\left(|p_\rho-p_\sigma|-\|p_\rho\rho-p_\sigma\sigma\|_1\right).$$

In quantum data privacy, it is desired to protect the quantum state of a system (which is being used for quantum computation) from being accurately estimated. Particularly, given a quantum state ρ, we want to make sure that no decision maker can identify whether the quantum state of the system is ρ or another *similar* quantum state σ. Similarity is modelled or captured using the neighbourhood relationship, c.f., differential privacy [20].

Definition 1 (Neighbouring Relationship) *A neighbouring or similarity relationship over the set of density operators is a mathematical relation that is both reflective and symmetric. The notation $\rho\sim\sigma$ signifies that two quantum states ρ and σ are neighbouring or similar. Note that, by definition, $\rho\sim\rho$ (reflectivity) and $\rho\sim\sigma$ implies $\sigma\sim\rho$ (symmetry).*

An example of neighbouring or similarity relationship is the notion defined using trace distance in [18]. In this case, we say $\rho\sim\sigma$ if and only if $\mathcal{T}(\rho,\sigma)\leq d$ for some constant $d>0$. However, we may select another notion of similarity that ensures that two quantum states are neighbouring if they are constructed based on two private datasets that differ in the data of one individual. Such a definition is well-suited for quantum machine learning with privacy guarantees [25].

Definition 2 ((ε,η)-Privacy) *For any $\varepsilon\geq 0$ and $\eta\in[0,1]$, a quantum channel $\mathcal{E}$ is (ε,η)-private (against hypothesis-testing adversary) if $D^\eta(\mathcal{E}(\rho)\|\mathcal{E}(\sigma))\leq\varepsilon$ for all neighbouring states $\rho\sim\sigma$.*

This definition implies that a channel $\mathcal{E}$ is private if distinguishing output states $\mathcal{E}(\rho)$ and $\mathcal{E}(\sigma)$, for any two neighbouring states $\rho\sim\sigma$, is hard by *any* decision maker. In fact, Proposition 1 shows that probability of false negatives $\beta(M_1)$ for any detection mechanism $M=\{M_1,M_2\}$ is lower bounded by $2^{-\varepsilon}$ if the probability of false positives bounded by $\alpha(M_2)\leq\eta$. Therefore, as privacy budget ε tends to zero (i.e., privacy guarantee is strengthened), the probability of false negatives tends to one (i.e., the decision maker would become overwhelmed by false negatives).

Proposition 2: Assume that a quantum channel $\mathcal{E}$ is (ε,η)-private. Then, the quantum channel $\mathcal{E}$ is (ε',η')-private if $\eta'\geq\eta$ and $\varepsilon\leq\varepsilon'$.

Proof: If $\eta'\geq\eta$, we have

$$\begin{aligned}2^{-D^\eta(\rho,\sigma)}&=\min\{\mathrm{tr}(Q\sigma)|0\preceq Q\preceq I,\mathrm{tr}(Q\rho)\geq 1-\eta\}\\&\leq\min\{\mathrm{tr}(Q\sigma)|0\preceq Q\preceq I,\mathrm{tr}(Q\rho)\geq 1-\eta'\}\\&=2^{-D^{\eta'}(\rho,\sigma)},\end{aligned}$$

where the inequality follows from that $\{Q|0\preceq Q\preceq I,\mathrm{tr}(Q\rho)\geq 1-\eta\}\subseteq\{Q|0\preceq Q\preceq I,\mathrm{tr}(Q\rho)\geq 1-\eta'\}$. Therefore, for all $\sigma\sim\rho$, we get $D^{\eta'}(\rho,\sigma)\leq D^\eta(\rho,\sigma)\leq\epsilon\leq\epsilon'$. ∎

The following corollary, building on Proposition 2, shows that $(\varepsilon,0)$-privacy against hypothesis testing adversary is the strongest notion of privacy and thus, (ε,η)-privacy can be thought of as relaxations of $(\varepsilon,0)$-privacy.

Corollary 1: Assume that a quantum channel $\mathcal{E}$ is $(\varepsilon,0)$-private. Then, the quantum channel $\mathcal{E}$ is (ε,η)-private for all $\eta\in[0,1]$.

Although privacy in Definition 2 is defined in terms of asymmetric quantum hypothesis testing, we prove the following bound on symmetric quantum hypothesis testing.

Theorem 2: For any (ε,η)-private quantum channel $\mathcal{E}$,

$$p_{\mathrm{err}}(\mathcal{E}(\rho),\mathcal{E}(\sigma))\geq\Gamma_{p_\rho,p_\sigma}(\varepsilon,\eta)\tag{5}$$

where

$$\Gamma_{p_\rho,p_\sigma}(\varepsilon,\eta):=\max\left\{p_{\max}-\frac{\varepsilon\min\{p_\rho,p_\sigma\}(1-\eta)}{2\eta},0\right\}.$$

Proof: Using [23], we have

$$\frac{\eta}{1-\eta}\|\mathcal{E}(\rho)-\mathcal{E}(\sigma)\|_1\leq D^\eta(\mathcal{E}(\rho)\|\mathcal{E}(\sigma)).$$

Therefore, if $\mathcal{E}$ is (ε,η)-private, we get

$$\|\mathcal{E}(\rho)-\mathcal{E}(\sigma)\|_1\leq\frac{1-\eta}{\eta}\varepsilon.$$

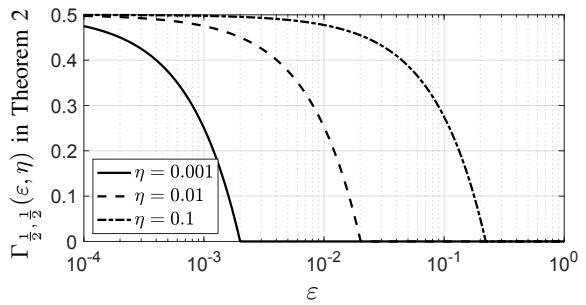

Fig. 1. The lower bound $\Gamma_{\frac{1}{2},\frac{1}{2}}(\varepsilon,\eta)$ on $p_{\rm err}(\mathcal{E}(\rho),\mathcal{E}(\sigma))$ in Theorem 2 versus the privacy budget ε for various choices of η. As expected, reducing the privacy budget ε strengthens the privacy guarantees.

We have

$$\begin{aligned}\|p_\rho\mathcal{E}(\rho)-p_\sigma\mathcal{E}(\sigma)\|_1 =& p_\sigma\left\|\frac{p_\rho}{p_\sigma}\mathcal{E}(\rho)-\mathcal{E}(\sigma)\right\|_1\\ =& p_\sigma\left\|\frac{p_\rho-p_\sigma}{p_\sigma}\mathcal{E}(\rho)+\mathcal{E}(\rho)-\mathcal{E}(\sigma)\right\|_1\\ \le& |p_\rho-p_\sigma|\|\mathcal{E}(\rho)\|_1+p_\sigma\|\mathcal{E}(\rho)-\mathcal{E}(\sigma)\|_1\\ \le& |p_\rho-p_\sigma|+\frac{1-\eta}{\eta}\varepsilon p_\sigma. \end{aligned} \tag{6}$$

Following the same line of reasoning, we can also show that

$$\|p_\rho\mathcal{E}(\rho) - p_\sigma\mathcal{E}(\sigma)\|_1 \le |p_\rho - p_\sigma| + \frac{1-\eta}{\eta}\varepsilon p_\rho. \tag{7}$$

Combining (6) and (7), we get

$$\|p_\rho\mathcal{E}(\rho) - p_\sigma\mathcal{E}(\sigma)\|_1 \le |p_\rho - p_\sigma| + \frac{1-\eta}{\eta}\varepsilon \min\{p_\rho, p_\sigma\}.$$

Therefore,

$$\begin{aligned} p_{\rm err}(\mathcal{E}(\rho),\mathcal{E}(\sigma)) =& p_{\max} + \frac{1}{2}\Big(|p_\rho - p_\sigma| \\ & - \|p_\rho\mathcal{E}(\rho) - p_\sigma\mathcal{E}(\sigma)\|_1\Big)\\ \ge& p_{\max} - \frac{\varepsilon\min\{p_\rho,p_\sigma\}(1-\eta)}{2\eta}. \end{aligned}$$

This concludes the proof. ∎

Theorem 2 shows that, by decreasing privacy budget ε, the combined probabilities of false positive and false negative denoted by $p_{\rm err}(\mathcal{E}(\rho),\mathcal{E}(\sigma))$ increases towards its maximum value $p_{\max}$. Figure 1 illustrates the lower bound $\Gamma_{p_\rho,p_\sigma}(\varepsilon,\eta)$ on $p_{\rm err}(\mathcal{E}(\rho),\mathcal{E}(\sigma))$ versus the privacy budget ε for various choices of η for the case that $p_\rho = p_\sigma = \frac{1}{2}$. As expected, reducing the privacy budget ε strengthens the privacy guarantees.

It is stipulated that any useful notion of privacy should admit two important properties of post processing and composition [20]. In the remainder of this section, we discuss these properties and their application to privacy against hypothesis-testing adversaries.

Theorem 3 (Post Processing) *Let $\mathcal{E}$ be any (ε,η)-private and $\mathcal{N}$ be an arbitrary quantum channel, then $\mathcal{N}\circ\mathcal{E}$ is (ε,η)-private.*

Proof: The proof follows from that, for all ρ and σ, $D^\eta(\mathcal{N}(\mathcal{E}(\rho))\|\mathcal{N}(\mathcal{E}(\sigma))) \le D^\eta(\mathcal{E}((\rho)\|\mathcal{E}(\sigma))$ [22]. ∎

Theorem 3 shows that an adversary cannot weaken the privacy guarantees by processing the received quantum information in any way. This is also a useful property for developing privacy-preserving algorithms as it ensures that it is enough to guarantee privacy at the beginning of the data analysis chain. This is the motivation behind the so-called local differential privacy, see, e.g., [26].

Theorem 4 (Composition) *Let $\mathcal{E}_1$ be any $(\varepsilon_1,0)$-private and $\mathcal{E}_2$ be any $(\varepsilon_2,0)$-private. Assume that $\rho_1\otimes\rho_2 \sim \sigma_1\otimes\sigma_2$ if $\rho_1\sim\sigma_1$ and $\rho_2\sim\sigma_2$. Then, $\mathcal{E}_1\otimes\mathcal{E}_2$ is $(\varepsilon_1+\varepsilon_2,0)$-private.*

Proof: Using the additivity results in [27, Appendix A], we get $D^0(\rho_1\otimes\rho_2\|\sigma_1\otimes\sigma_2) = D^0(\rho_1\|\sigma_1) + D^0(\rho_2\|\sigma_2)$. Therefore, if $D^0(\rho_i\|\sigma_i)\le\varepsilon_i$ for $i=1,2$, then $D^0(\rho_1\otimes\rho_2\|\sigma_1\otimes\sigma_2)\le\varepsilon_1+\varepsilon_2$. ∎

In practical data processing applications, there is often a need to deal with complicated algorithms in which responses from several queries based on private user data are fused together to extract useful statistical information from the data. For instance, when training machine learning models, iterative gradient descent is used and the gradient at each epoch can be modelled as a query on the private data used for training [28]. In this case, it is desirable to establish composition rules for combination of several privacy-preserving quantum operations. Theorem 4 provides such a result for privacy against hypothesis-testing adversaries.

IV. Quantum Differential Privacy

The gold standard of privacy analysis in computer science literature is differential privacy, which has been recently extended to quantum computing [18]. In this section, we establish a relationship between differential privacy and privacy against hypothesis-testing adversaries.

Definition 3: For any $\epsilon,\delta\ge 0$, a quantum channel $\mathcal{E}$ is (ϵ,δ)-differentially private if

$$\operatorname{tr}(M\mathcal{E}(\rho)) \le \exp(\epsilon)\operatorname{tr}(M\mathcal{E}(\sigma)) + \delta, \tag{8}$$

for all measurements $0\preceq M\preceq I$ and neighbouring density operators $\rho\sim\sigma$.

We can prove the following result regarding the relationship between quantum differential privacy and privacy against hypothesis testing adversaries.

Theorem 5: The following two statements hold:

- If $\mathcal{E}$ be (ε,η)-private, then $\mathcal{E}$ is $(\varepsilon,\sqrt{2\eta})$-differentially private.
- If $\mathcal{E}$ be $(\epsilon,0)$-differentially private, then $\mathcal{E}$ is (ϵ,η)-private for all $\eta\in[0,1]$.

Proof: First, $D^{\sqrt{2\nu}}_{\max}(\mathcal{E}(\rho),\mathcal{E}(\sigma)) \le D^\nu(\mathcal{E}(\rho),\mathcal{E}(\sigma))$ [23, Proposition 4.1]. Therefore, if $\mathcal{E}$ is (ε,η)-private, we get $D^{\sqrt{2\eta}}_{\max}(\mathcal{E}(\rho),\mathcal{E}(\sigma)) \le D^\eta(\mathcal{E}(\rho),\mathcal{E}(\sigma)) \le \varepsilon$. From Lemma III.2 in [20], a quantum channel $\mathcal{E}$ is (ϵ,δ)-differentially private if and only if $D^\delta_{\max}(\mathcal{E}(\rho),\mathcal{E}(\sigma))\le\epsilon$. This proves that $\mathcal{E}$ is $(\varepsilon,\sqrt{2\eta})$-differentially private.

For the second part, note that $D^\eta(\mathcal{E}(\rho),\mathcal{E}(\sigma)) \le D^0_{\max}(\mathcal{E}(\rho),\mathcal{E}(\sigma))$ [23, Proposition 4.1]. Therefore, if $\mathcal{E}$

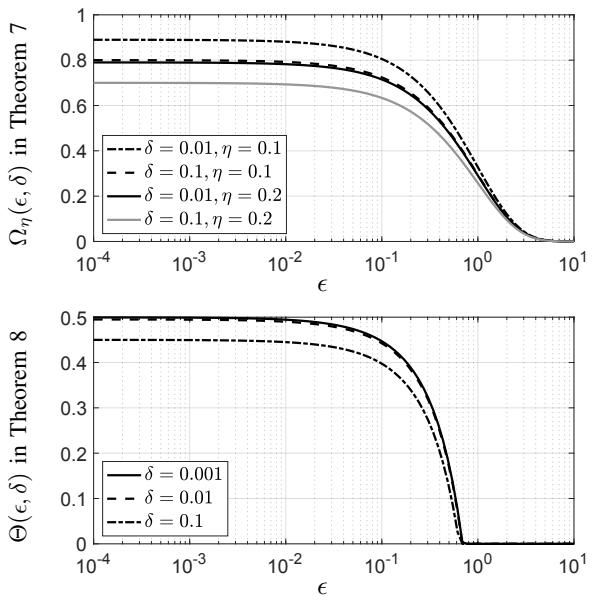

Fig. 2. Lower bound on $\beta_\eta(\mathcal{E}(\rho),\mathcal{E}(\sigma))$ in Theorem 7 [top] and lower bound on $p_{\rm err}(\mathcal{E}(\rho),\mathcal{E}(\sigma))$ in Theorem 8 [bottom] versus privacy budget ϵ for various choices of δ.

is $(\epsilon,0)$-differentially private, we have $D^\eta(\mathcal{E}(\rho),\mathcal{E}(\sigma)) \leq D^0_{\max}(\mathcal{E}(\rho),\mathcal{E}(\sigma)) \leq \epsilon$. This implies that $\mathcal{E}$ is (ϵ,η)-private for all $\eta \in [0,1]$. ∎

Theorem 6 (Lemma IV.2 [20]) *Assume $\rho \sim \sigma$ if $\mathcal{T}(\rho,\sigma) \leq \kappa$. Then, global depolarizing channel $\mathcal{E}_{\rm Dep}$ is (ϵ,δ)-differentially private with $\delta = \max\{0,(1-\exp(\epsilon))p/D + (1-p)\kappa\}$.*

Corollary 2: Assume $\rho \sim \sigma$ if $\mathcal{T}(\rho,\sigma) \leq \kappa$. Then, global depolarizing channel $\mathcal{E}_{\rm Dep}$ is (ε,η)-private with $\varepsilon = \log(1+(1-p)D\kappa/p)$ and all $\eta \in [0,1]$.

Proof: First, note that Theorem 6 shows that the depolarizing channel $\mathcal{E}_{\rm Dep}(\rho)$ is (ϵ,δ)-differentially private with $\delta = \max\{0,(1-\exp(\epsilon))p/D + (1-p)\kappa\}$. If we select $\epsilon = \log(1+(1-p)D\kappa/p)$, we get $\delta = 0$. Using Theorem 5, then $\mathcal{E}$ is (ϵ,η)-private for all $\eta \in [0,1]$. ∎

We finish this section with analysing the performance of hypothesis-testing adversaries for differentially-private quantum channels. This enables to establish bounds on hypothesis testing under quantum differential privacy constraint. Furthermore, we can establish an operational interpretation for guarantees of quantum differential privacy.

Theorem 7: For any (ϵ,δ)-differentially private quantum channel $\mathcal{E}$,

$$\beta_\eta(\mathcal{E}(\rho),\mathcal{E}(\sigma)) \geq \Omega_\eta(\epsilon,\delta), \tag{9}$$

where $\Omega_\eta(\epsilon,\delta) := \exp(-\epsilon)(1-\eta-\delta)$.

Proof: Assume that $\rho \sim \sigma$. Because of (ϵ,δ)-differential privacy, $\mathrm{tr}(M\mathcal{E}(\sigma)) \geq \exp(-\epsilon)(\mathrm{tr}(M\mathcal{E}(\rho)) - \delta)$ for all measurements $0 \preceq M \preceq I$. Therefore, $\beta_\eta(\mathcal{E}(\rho),\mathcal{E}(\sigma)) = \min_{I\succeq M\succeq 0}\{\mathrm{tr}(M\mathcal{E}(\sigma))|1-\eta \leq \mathrm{tr}(M\mathcal{E}(\rho))\} \geq \exp(-\epsilon)(1-\eta-\delta)$. ∎

Theorem 7 provides a lower bound for the false negative rate of any hypothesis testing mechanism under quantum differential privacy constraint. The lower bound grows and the decision maker would get overwhelmed by false negatives when decreasing ϵ and δ. Therefore, the privacy guarantees strengthens as the privacy budget reduces in quantum differential privacy. This is illustrated in Figure 2 [top].

Theorem 8: For any (ϵ,δ)-differentially private quantum channel $\mathcal{E}$,

$$p_{\rm err}(\mathcal{E}(\rho),\mathcal{E}(\sigma)) \geq \Theta(\epsilon,\delta), \tag{10}$$

where $\Theta(\epsilon,\delta) := \max\{p_{\max} + \max\{p_\rho,p_\sigma\}(1-\exp(\epsilon)-\delta),0\}$.

Proof: First, assume that $p_\rho \geq p_\sigma$. The definition of differential privacy implies that $\mathrm{tr}(\Lambda\mathcal{E}(\rho)) \leq \exp(\epsilon)\,\mathrm{tr}(\Lambda\mathcal{E}(\sigma)) + \delta$ for all $0 \preceq \Lambda \preceq I$. As a result,

$$\begin{aligned}\mathrm{tr}(\Lambda(p_\rho\mathcal{E}(\rho) - p_\sigma\mathcal{E}(\sigma))) &\leq (p_\rho\exp(\epsilon) - p_\sigma)\,\mathrm{tr}(\Lambda\mathcal{E}(\sigma)) + p_\rho\delta \\ &\leq p_\rho\exp(\epsilon) - p_\sigma + p_\rho\delta,\end{aligned}$$

where the last inequality follows from that $p_\rho\exp(\epsilon) - p_\sigma \geq p_\rho - p_\sigma \geq 0$ and that $\mathrm{tr}(\Lambda\mathcal{E}(\sigma)) \leq 1$ because $0 \preceq \Lambda \preceq I$. Therefore, using Lemma 1 in the appendix, we have

$$\begin{aligned}\frac{1}{2}\|p_\rho\mathcal{E}(\rho) - p_\sigma\mathcal{E}(\sigma)\|_1 &\leq p_\rho\exp(\epsilon) - p_\sigma + p_\rho\delta + \frac{p_\sigma - p_\rho}{2} \\ &\leq p_\rho(\exp(\epsilon)+\delta) - \frac{p_\sigma + p_\rho}{2}.\end{aligned} \tag{11}$$

Alternatively, assume that $p_\sigma \geq p_\rho$. Following the same line of reasoning, we get

$$\frac{1}{2}\|p_\sigma\mathcal{E}(\sigma) - p_\rho\mathcal{E}(\rho)\|_1 \leq p_\sigma(\exp(\epsilon)+\delta) - \frac{p_\sigma + p_\rho}{2}. \tag{12}$$

Combining (11) and (12) gives

$$\begin{aligned}\frac{1}{2}\|p_\sigma\mathcal{E}(\sigma) - p_\rho\mathcal{E}(\rho)\|_1 \leq \max\{p_\rho,p_\sigma\}(\exp(\epsilon)+\delta) \\ - \frac{p_\sigma + p_\rho}{2}.\end{aligned}$$

Therefore,

$$\begin{aligned}&|p_\rho - p_\sigma| - \|p_\rho\mathcal{E}(\rho) - p_\sigma\mathcal{E}(\sigma)\|_1 \\ &\quad\geq |p_\rho - p_\sigma| + (p_\sigma + p_\rho) - 2\max\{p_\rho,p_\sigma\}(\exp(\epsilon)+\delta) \\ &\quad= 2\max\{p_\rho,p_\sigma\}(1-\exp(\epsilon)-\delta).\end{aligned}$$

This concludes the proof. ∎

Theorem 8 provides a lower bound for the combined false positive and negative rates of any hypothesis testing mechanism. The lower bound grows towards $p_{\max}$ as ϵ and δ become smaller, which demonstrates that the privacy guarantees strengthen as the privacy budget reduces in quantum differential privacy. This is illustrated in Figure 2 [bottom].

V. Conclusions and Future Work

We presented a definition for privacy based on quantum hypothesis testing. Important properties of post processing and composition were proved. We then examined the relationship between privacy against hypothesis-testing adversaries and quantum differential privacy. In the composition rules for privacy against hypothesis adversaries, we only considered the case of $\eta = 0$. Future work can expand

these results for general case of $\eta \in [0,1]$. Furthermore, we only showed that $(\epsilon, 0)$-differential privacy can be translated to privacy against hypothesis testing adversaries (the inverse results are more general in this paper). Therefore, another avenue for future research is to expand these results to general (ϵ, δ)-differential privacy. Finally, an important direction for future research is to use the proposed framework in numerical setups based on real-world data.

REFERENCES

[1] P. W. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," in *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pp. 124–134, Ieee, 1994.

[2] S. Lloyd, M. Mohseni, and P. Rebentrost, "Quantum principal component analysis," *Nature Physics*, vol. 10, no. 9, pp. 631–633, 2014.

[3] H.-Y. Huang, R. Kueng, and J. Preskill, "Information-theoretic bounds on quantum advantage in machine learning," *Physical Review Letters*, vol. 126, no. 19, p. 190505, 2021.

[4] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.

[5] M. Kearns and A. Roth, *The ethical algorithm: The science of socially aware algorithm design*. Oxford University Press, 2019.

[6] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*, pp. 265–284, Springer, 2006.

[7] C. Dwork, "Differential privacy: A survey of results," in *Theory and Applications of Models of Computation: 5th International Conference, TAMC 2008, Xi'an, China, April 25-29, 2008. Proceedings 5*, pp. 1–19, Springer, 2008.

[8] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 308–318, 2016.

[9] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations and Trends in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

[10] V. J. Hotz, C. R. Bollinger, T. Komarova, C. F. Manski, R. A. Moffitt, D. Nekipelov, A. Sojourner, and B. D. Spencer, "Balancing data privacy and usability in the federal statistical system," *Proceedings of the National Academy of Sciences*, vol. 119, no. 31, p. e2104906119, 2022.

[11] R. Bhaskar, A. Bhowmick, V. Goyal, S. Laxman, and A. Thakurta, "Noiseless database privacy," in *International Conference on the Theory and Application of Cryptology and Information Security*, pp. 215–232, Springer, 2011.

[12] S. U. Nabar, B. Marthi, K. Kenthapadi, N. Mishra, and R. Motwani, "Towards robustness in query auditing," in *Proceedings of the 32nd International Conference on Very Large Data Bases*, pp. 151–162, VLDB Endowment, 2006.

[13] I. Issa, A. B. Wagner, and S. Kamath, "An operational approach to information leakage," *IEEE Transactions on Information Theory*, vol. 66, no. 3, pp. 1625–1657, 2019.

[14] F. Farokhi and N. Ding, "Measuring information leakage in non-stochastic brute-force guessing," in *2020 IEEE Information Theory Workshop (ITW)*, pp. 1–5, IEEE, 2021.

[15] 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*, vol. 65, no. 12, pp. 8043–8066, 2019.

[16] Z. Li, T. J. Oechtering, and D. Gündüz, "Privacy against a hypothesis testing adversary," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 6, pp. 1567–1581, 2018.

[17] F. Farokhi and H. Sandberg, "Fisher information as a measure of privacy: Preserving privacy of households with smart meters using batteries," *IEEE Transactions on Smart Grid*, vol. 9, no. 5, pp. 4726–4734, 2017.

[18] L. Zhou and M. Ying, "Differential privacy in quantum computation," in *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pp. 249–262, IEEE, 2017.

[19] S. Aaronson and G. N. Rothblum, "Gentle measurement of quantum states and differential privacy," in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pp. 322–333, 2019.

[20] C. Hirche, C. Rouzé, and D. S. França, "Quantum differential privacy: An information theory perspective," *arXiv preprint arXiv:2202.10717*, 2022.

[21] M. Wilde, *Quantum Information Theory*. Quantum Information Theory, Cambridge University Press, 2013.

[22] L. Wang and R. Renner, "One-shot classical-quantum capacity and hypothesis testing," *Physical Review Letters*, vol. 108, no. 20, p. 200501, 2012.

[23] F. Dupuis, L. Kraemer, P. Faist, J. M. Renes, and R. Renner, "Generalized entropies," in *XVIIth International Congress on Mathematical Physics*, pp. 134–153, 2014.

[24] B. Regula, L. Lami, and M. M. Wilde, "Postselected quantum hypothesis testing," *arXiv preprint arXiv:2209.10550*, 2022.

[25] W. M. Watkins, S. Y.-C. Chen, and S. Yoo, "Quantum machine learning with differential privacy," *arXiv preprint arXiv:2103.06232*, 2021.

[26] G. Cormode, S. Jha, T. Kulkarni, N. Li, D. Srivastava, and T. Wang, "Privacy at scale: Local differential privacy in practice," in *Proceedings of the 2018 International Conference on Management of Data*, pp. 1655–1658, 2018.

[27] X. Yuan, "Hypothesis testing and entropies of quantum channels," *Physical Review A*, vol. 99, no. 3, p. 032317, 2019.

[28] N. Wu, F. Farokhi, D. Smith, and M. A. Kaafar, "The value of collaboration in convex machine learning with differential privacy," in *2020 IEEE Symposium on Security and Privacy (SP)*, pp. 304–317, IEEE, 2020.

APPENDIX

Lemma 1: The following identity holds:

$$\frac{1}{2}\|p_\rho\rho - p_\sigma\sigma\|_1 = \max_{0\preceq\Lambda\preceq I} \operatorname{tr}(\Lambda(p_\rho\rho - p_\sigma\sigma)) + \frac{p_\sigma - p_\rho}{2}.$$

Proof: The proof is similar to the standard argument for trace distance. The difference operator $p_\rho\rho - p_\sigma\sigma$ is Hermitian. So we can diagonalize it as $p_\rho\rho - p_\sigma\sigma = \sum_i \lambda_i |i\rangle\langle i|$, where $\{|i\rangle\}_i$ is an orthonormal basis of eigenvectors and $\{\lambda_i\}_i$ is a set of real eigenvalues. Define matrices $P := \sum_{i:\lambda_i>0} \lambda_i |i\rangle\langle i| \succeq 0$ and $Q := \sum_{i:\lambda_i<0}(-\lambda_i)|i\rangle\langle i| \succeq 0$. Evidently, by construction, $p_\rho\rho - p_\sigma\sigma = P - Q$. Note that,

$$\begin{aligned}\|p_\rho\rho - p_\sigma\sigma\|_1 = \operatorname{tr}(|p_\rho\rho - p_\sigma\sigma|) &= \operatorname{tr}(|P-Q|)\\ &= \operatorname{tr}(P+Q)\\ &= 2\operatorname{tr}(P) + (p_\sigma - p_\rho),\end{aligned}$$

where the last equality follows from

$$\begin{aligned}\operatorname{tr}(P) - \operatorname{tr}(Q) = \operatorname{tr}(P-Q) &= \operatorname{tr}(p_\rho\rho - p_\sigma\sigma)\\ &= p_\rho\operatorname{tr}(\rho) - p_\sigma\operatorname{tr}(\sigma)\\ &= p_\rho - p_\sigma.\end{aligned}$$

For all $0 \preceq \Lambda \preceq I$, we have

$$\begin{aligned}\operatorname{tr}(\Lambda(p_\rho\rho - p_\sigma\sigma)) &= \operatorname{tr}(\Lambda(P-Q))\\ &\le \operatorname{tr}(\Lambda P)\\ &\le \operatorname{tr}(P)\\ &= \frac{1}{2}\|p_\rho\rho - p_\sigma\sigma\|_1 + \frac{p_\rho - p_\sigma}{2},\end{aligned}$$

with equality achieved if $\Lambda = \sum_{i:\lambda_i>0}|i\rangle\langle i|$. This implies that

$$\frac{1}{2}\|p_\rho\rho - p_\sigma\sigma\|_1 = \max_{0\preceq\Lambda\preceq I} \operatorname{tr}(\Lambda(p_\rho\rho - p_\sigma\sigma)) + \frac{p_\sigma - p_\rho}{2}.$$

■