Differential Privacy for Hybrid Quantum-Classical Algorithms: Frameworks, Mechanisms, and Applications in Quantum Machine Learning

by
Jingtong Ge

A dissertation submitted in fulfillment
of the requirements for the degree of
Doctor of Philosophy

University of Technology Sydney
2025

Certificate of Original Authorship

I, Jingtong Ge, declare that this thesis is submitted in fulfilment of the requirements for the award of Doctor of Philosophy, in the Faculty of Engineering and Information Technology at the University of Technology Sydney.

This thesis is wholly my own work unless otherwise referenced or acknowledged. In addition, I certify that all information sources and literature used are indicated in the thesis.

This document has not been submitted for qualifications at any other academic institution.

This research was supported by an Australian Government Research Training Program (RTP) Scholarship doi.org/10.82133/C42F-K220.

Production Note:
Signature removed prior to publication.

18-09-2025

Acknowledgements

During my Ph.D. study, I have greatly enjoyed my life and research experience in Sydney. This beautiful city, with its stunning scenery, excellent environment, and friendly people, has left me with unforgettable memories. Studying and living here not only broadened my horizons but also gave me the opportunity to grow as a person and as a researcher.

I would like to express my deepest gratitude to my supervisors, Prof. Yuan Feng and Prof. Sanjiang Li, for their continuous guidance and support throughout my Ph.D. study. Their patient answers to my questions, constructive suggestions in every group meeting, and thoughtful advice on both research and life matters have been invaluable to me. I am especially grateful for their encouragement in times of difficulty, for never giving up on me, and for constantly inspiring me to pursue rigorous and meaningful research.

I am also sincerely thankful to my collaborator and co-supervisor, Dr. Ji Guan, for his invaluable help during our collaboration. Without his guidance, it would have been impossible for me to complete the revisions of my papers. From the early drafts to the final submissions and responses to reviewers, his insightful suggestions and careful guidance have shaped the development of my work step by step. I owe him a deep debt of gratitude for his dedication and support.

I would like to express my heartfelt appreciation to Prof. Mingsheng Ying, whose rigorous attitude toward scientific research has always inspired me and influenced my academic path. His dedication and vision have become a lifelong example for me to look up to and follow. I am equally deeply grateful to Prof. Lihong Zhi, my Master's supervisor, who introduced me to the University of Technology Sydney and thereby opened the door to my Ph.D. journey. From her, I have learned not only research skills but also many admirable qualities in life. She has been both a mentor in science and a role model in life, and I regard her as a lifelong source of inspiration.

I would further like to thank Prof. Youming Qiao and Prof. Xianzhi Wang for serving on my CA3 panel and for their valuable feedback and suggestions, which significantly improved the quality of my thesis.

My thanks also go to the members of the CQSI for creating a stimulating research environment and for the many enlightening discussions. I am grateful as well to the administrative and support staff for their assistance with various academic and logistical matters, which has helped me focus on my research.

I would also like to acknowledge my fellow students and colleagues, including Xin Hong, Gang Tang, and Yunqi Huang, for their friendship, collaboration, and the enjoyable moments we shared during this journey.

Finally, I would like to express my heartfelt gratitude to my parents and my sister, for their unconditional love, encouragement, and trust, which have always been a source

of strength throughout my life. Most importantly, I am deeply thankful to my wife, Qiyi Wang, for her unwavering support, understanding, and companionship. Her love and encouragement have been my greatest motivation and comfort during the most challenging moments of my Ph.D. journey. Sincerely thank you all.

Abstract

Differential privacy is a mathematical framework for protecting individual data in algorithmic outputs, ensuring that small changes to the input do not significantly alter the output distribution. It has become a cornerstone of privacy-preserving data analysis and machine learning in classical settings, and has recently been extended to quantum computing to safeguard information encoded in quantum states. However, in the Noisy Intermediate-Scale Quantum (NISQ) era, where practical quantum applications rely heavily on hybrid quantum-classical algorithms due to inherent quantum noise, the integration of differential privacy in such algorithms has been largely overlooked.

This thesis fills that gap by proposing a hybrid quantum-classical differential privacy (HDP) framework tailored for hybrid algorithms with fixed quantum measurements, which serve as the primary interface between quantum and classical systems. We focus on designing differentially private quantum measurements using both quantum depolarizing noise and a measurement-based exponential mechanism (MBEM), which enables privacy guarantees while preserving utility in practical hybrid settings. To ensure robustness under post-processing and composability under repeated measurements, we establish post-processing and composition theorems within the HDP framework. We validate these theoretical results through extensive numerical experiments on various quantum circuits, demonstrating the practical effectiveness of our approach in preserving privacy with manageable utility trade-offs.

Furthermore, we investigate Rényi Differential Privacy (RDP), a refinement of standard differential privacy with strong composability and tunable privacy-utility trade-offs, in the context of quantum machine learning (QML). QML algorithms, typically implemented via hybrid quantum-classical procedures, use fixed quantum measurements to extract classical outputs from quantum states. Leveraging the Heisenberg picture, we formulate RDP in terms of Rényi divergence between measurement-induced output distributions and derive analytically tractable upper bounds for certifying privacy guarantees. Our results demonstrate that QML provides a principled and tractable pathway for privacy analysis, bridging classical RDP techniques with practical quantum learning systems.

Together, these contributions form a unified and practical approach for achieving differential privacy in hybrid quantum-classical algorithms, offering rigorous tools for privacy-preserving quantum machine learning on real-world NISQ systems.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

A quantum algorithm comprises a series of computational instructions for manipulating quantum circuits on a quantum computer to solve specific computational problems. Numerous quantum algorithms, such as Shor's algorithm and Grover's algorithm [Sho94; Gro96], have been shown to possess remarkable capabilities that speed up classical algorithms. The enhanced speed primarily stems from the distinctive benefits of quantum states, including superposition and entanglement [NC01]. These algorithms could be executed on forthcoming fault-tolerant quantum computers to showcase the practical superiority of quantum computing.

In the present era of *Noisy Intermediate-Scale Quantum (NISQ)* computing [Pre18], characterized by quantum computers containing hundreds of noisy quantum bits (qubits), noise inevitably impacts the quantum computing process. Consequently, there has been a surge in the development of *hybrid quantum-classical algorithms*. These algorithms leverage the complementary strengths of classical and quantum computing to address the limited computational capacity resulting from medium scalability and noise in NISQ computers. Hybrid quantum-classical algorithms operate akin to classical machine learning algorithms, utilizing parameterized quantum circuits (analogous to classical neural networks) alongside classical optimizers to adjust parameters based on the output distributions of quantum data obtained from quantum measurements, thereby enabling the

resolution of intricate computational tasks. Notable examples of such algorithms include the Variational Quantum Eigensolver (VQE) [Per+14] and the Quantum Approximate Optimization Algorithm (QAOA) [FGG14]. VQE is tailored for determining the ground state energy of molecules and materials. It accomplishes this by generating potential states using parameterized quantum circuits and subsequently employing classical optimizers to minimize the expected value. QAOA, on the other hand, concentrates on addressing combinatorial optimization problems by producing feasible solutions through parameterized quantum circuits and utilizing classical optimizers to refine parameters to optimize the objective function. Consequently, hybrid quantum-classical algorithms present a promising approach for integrating quantum computing into practical applications within the current NISQ era. They offer advantages in terms of performance and resilience to noise [Jon+19; End+21].

In algorithms like these, the handling of privacy-sensitive classical data (e.g., personal financial records and drug information) stored in quantum data is crucial, as highlighted in several studies [CRA18; De +21; OML19]. This heightened awareness emphasizes the importance of protecting users' privacy within these algorithms. In classical algorithms, addressing personal privacy concerns often involves employing differential privacy [DR+14], which aims to reduce the impact of individual data differences in neighboring datasets on algorithm outcomes. The concept of differential privacy has also been utilized to improve quantum data privacy in quantum information processing by establishing meaningful relationships between neighboring quantum states. These relationships are mainly assessed using informative metrics, such as local operation [AR19] and trace distance [ZY17], to describe the similarity between quantum states qualitatively and quantitatively, respectively. The local operation method involves a classical similarity achieved through a local operation to transition one state to another, while trace distance measures the difference between quantum states on a scale from 0 to 1. As a result, there is a growing body of literature exploring quantum differential privacy [SMT17; WCY23; HRF23; ADK22; AK22; QAS21; Du+21; NGW24] in various scenarios. However, these studies predominantly fo-

cus on the quantum realm and often overlook privacy concerns in hybrid quantum-classical algorithms where quantum and classical information are exchanged through quantum measurements.

Although both classical and quantum differential privacy approaches have proven effective in their respective domains, they are not directly applicable to hybrid quantum-classical settings.The limitations of these methods are outlined below. *Classical differential privacy (CDP)* [DR+14] has been extensively studied across fields such as data analysis, machine learning, and dataset queries [Aba+16; MK19; Hu+21; Gad+22; Ave+17; XX15; Li+17]. These methods ensure robust privacy guarantees for deterministic algorithms by introducing randomized noise mechanisms into their outputs. Key approaches include the Laplace, Gaussian, and exponential mechanisms, each offering distinct strategies for achieving CDP. The selection between these mechanisms depends on the specific data characteristics and privacy objectives. The Laplace and Gaussian mechanisms enforce differential privacy by adding noise sampled from their respective distributions to the output of a deterministic function. However, these mechanisms are unsuitable for developing differentially private quantum measurements due to the inherent randomness in the measurements. Instead, we consider adapting the exponential mechanism into the quantum domain for implementing differentially private quantum measurements as it works for categorical outputs that can be regarded as measurement outcomes. The exponential mechanism upholds differential privacy by selecting outputs based on their quality scores, with probabilities determined by an exponential function of the score. To apply this mechanism to quantum measurements, we leverage the original measurement outcome probability distribution as quality scores for measurement outcomes, redistributing outcomes based on the exponential function of the scores. This extension broadens the utility of the exponential mechanism to the quantum realm within the probabilistic domain. Through the innovative measurement-based exponential mechanism, the creation of differentially private quantum measurements is simplified, reinforcing the privacy of hybrid quantum-classical algorithms.

Researchers have been exploring how to incorporate differential privacy into quantum algorithms [HRF23; ADK22; AK22; QAS21; Du+21]. They have investigated the effects of both inherent quantum noise in quantum algorithms and additional noise on the overall differential privacy guarantees of these algorithms. Specifically, *quantum differential privacy (QDP)* is designed to protect the output states of quantum algorithms against privacy breaches under arbitrary quantum measurements. In the context of hybrid quantum-classical algorithms, however, only one specific measurement is revealed, and thus the focus is on mitigating privacy risks associated with this particular measurement rather than all possible measurements. To address this issue, a depolarizing mechanism has been proposed to introduce quantum noise directly into the quantum measurement. This approach simplifies the process of achieving differential privacy for the measurement and ensures the desired level of privacy protection for the hybrid quantum-classical algorithm. Moreover, this measurement-dependent notion of protection provides a stronger guarantee (i.e., a tighter lower bound) than traditional QDP, as demonstrated in Theorem 6 and its corollary.

In this thesis, we present a novel hybrid quantum-classical differential privacy framework that aims to protect the privacy of hybrid quantum-classical algorithms, thereby filling a gap in the current literature. Our framework focuses on developing differentially private quantum measurements to enable the implementation of differentially private hybrid quantum-classical algorithms. To maintain privacy guarantees even after subsequent data processing, we introduce a post-processing theorem. Although quantum measurements typically act as randomized functions mapping quantum states to a finite set of outcomes, they often fall short of providing the desired level of differential privacy. To enhance privacy protection by leveraging the hybrid nature of quantum measurements, we incorporate classical noise mechanisms after the measurements, similar to classical approaches, or introduce a quantum noise mechanism before the measurements.

To operationalize this concept effectively, we propose utilizing quantum depolarizing noise (similar to the classical randomized response mechanism [Gua24]) together with an

analytically derived privacy budget, and employing the measurement-based exponential mechanism (an adaptation of the classical exponential mechanism [DR+14]) during the measurement phase. The measurement-based exponential mechanism uses the original outcome distributions of quantum measurements as a utility function to release outcomes while ensuring differential privacy. This method offers advantages beyond its classical counterpart, such as not relying on the sensitivity of utility functions, since a universal bound can be established. Additionally, we introduce a composition theorem to guarantee the differential privacy of complex hybrid quantum-classical algorithms.

Finally, through a series of numerical experiments, we validate the efficacy of our framework, particularly with the inclusion of quantum depolarizing noise and the measurement-based exponential mechanism. Last but not least, our hybrid differential privacy framework not only achieves classical differential privacy through neighboring-preserving encoding from classical data to quantum data, but also ensures quantum differential privacy through quantum noise that is independent of quantum measurements.

A particularly promising application of quantum computing is *Quantum Machine Learning (QML)*, which aims to enhance the expressivity, optimization, and generalization of machine learning models using quantum resources. In the current *Noisy Intermediate-Scale Quantum (NISQ)* era [Pre18], QML algorithms are typically implemented via *hybrid quantum-classical algorithms*. These algorithms alternate between parameterized quantum circuits and classical optimization routines, with classical feedback used to iteratively update quantum parameters. Representative examples include the Variational Quantum Eigensolver (VQE) [Per+14], the Quantum Approximate Optimization Algorithm (QAOA) [FGG14], and the Variational Quantum Classifier (VQC). A common feature of such algorithms is the use of *fixed quantum measurements*, usually in the computational basis, to extract classical information from quantum states.

As QML begins to be applied in domains involving sensitive input data (e.g., genomic sequences, financial records, and proprietary simulations), privacy concerns become critical. Although quantum states are inaccessible during their evolution, the correspond-

ing measurement outcomes, being classical, can still potentially leak private information. Fortunately, the fixed-measurement structure of QML provides a well-defined interface between quantum computation and classical output, which enables rigorous privacy analysis based solely on measurement statistics.

To address privacy in classical computation, *Differential Privacy (DP)* [DR+14] has become the standard for limiting the influence of any single input on algorithmic output. This framework has been extended to the quantum setting via *Quantum Differential Privacy (QDP)* [ZY17; AR19], where the notion of neighboring datasets is generalized to quantum states. Various studies [SMT17; WCY23; HRF23; ADK22; AK22; QAS21; Du+21; NGW24] have investigated QDP from different perspectives. However, most of these works focus on general-purpose quantum channels or arbitrary measurements. In contrast, the QML setting, characterized by a fixed quantum measurement repeated across inputs, remains underexplored in privacy research.

In this thesis, we propose to analyze privacy in QML through the lens of *Rényi Differential Privacy (RDP)* [Mir17], a relaxation of DP based on Rényi divergence. RDP introduces a tunable order parameter α that allows more flexible privacy-utility tradeoffs and supports tighter composition bounds. It has been widely used in classical machine learning, including in DP-SGD [PS21], posterior sampling [GSC17], and generalization control [WLF16].

We show that in QML, where a fixed measurement is applied to a parameterized quantum state, the entire process can be viewed, through the Heisenberg picture, as an effective quantum measurement acting on the input. This insight enables us to define and analyze the Rényi differential privacy of the resulting output distribution.

Summary of Contributions. This thesis makes the following main contributions:

1. *Formulating* a hybrid quantum-classical differential privacy framework that ensures both classical and quantum privacy guarantees by leveraging differentially private quantum measurements. Within this framework, we *introduce* a post-processing

theorem and a composition theorem to formalize privacy preservation across multiple stages of hybrid algorithms.

2. *Designing* effective noise mechanisms for differentially private quantum measurements, including *quantum depolarizing noise* and the *measurement-based exponential mechanism*. These mechanisms allow the privatization of quantum measurement outcomes under a specified privacy budget, thereby strengthening privacy protection while maintaining practical utility.

3. *Validating* the HDP framework through a series of numerical experiments on benchmark circuits. The results highlight the effectiveness of integrating differentially private quantum measurements with quantum depolarizing noise and measurement-based exponential mechanisms in balancing privacy and utility.

4. *Extending* Rényi Differential Privacy to fixed quantum measurements, particularly tailored for quantum machine learning circuits. This extension bridges the theoretical development of RDP with practical QML scenarios, enabling fine-grained, scalable, and efficient privacy analysis.

5. *Developing* an efficient eigenvalue-based algorithm to compute accurate and meaningful privacy guarantees under the RDP framework. This algorithm provides tighter and more scalable guarantees compared to classical ϵ-DP approaches, making it suitable for realistic QML applications.

6. *Conducting* extensive experiments on realistic QML circuits and open-source datasets. The experiments demonstrate that our analytical RDP method consistently provides tight and scalable guarantees, outperforming ϵ-DP baselines and validating its effectiveness under diverse noise settings and circuit configurations.

1.1 Publications Related to This Thesis

- [GG25a] Jingtong Ge and Ji Guan, "Differential Privacy of Hybrid Quantum-Classical Algorithms," submitted to NeurIPS 2025.
- [GG25b] Jingtong Ge and Ji Guan, "Rényi Differential Privacy in Quantum Machine Learning," submitted to AAAI 2025.

1.1.1 Thesis Organization

This thesis addresses the problem of privacy protection in hybrid quantum-classical algorithms with fixed quantum measurements. It presents two central research threads: (i) the development of a hybrid differential privacy framework, and (ii) the extension and verification of Rényi differential privacy in quantum machine learning. The thesis is organized into five chapters as follows:

Chapter 1 – Introduction: This chapter outlines the motivation and background of the research, emphasizing the privacy risks in hybrid quantum-classical algorithms such as Variational Quantum Eigensolvers and Variational Quantum Classifiers. It identifies the limitations of classical differential privacy and quantum DP under fixed measurement settings. The chapter introduces the key objectives of the thesis: to define a practically applicable HDP model and to extend Rényi differential privacy to quantum learning settings. It concludes with a summary of contributions and the overall structure of the thesis.

Chapter 2 – Preliminaries: This chapter provides essential theoretical background and notation, divided into five parts:

- Basics of quantum computing: quantum states, gates, channels, and measurement operators;
- Modeling hybrid quantum-classical algorithms, including circuit structures and Heisenberg-picture analysis;
- Classical differential privacy and its variants: (ϵ, δ)-DP and Rényi DP;

- Quantum DP: definitions of neighboring quantum states, privacy sensitivity, and encoding strategies;
- Classical and quantum noise mechanisms, including Laplace, Gaussian, exponential, and depolarizing noise, and their applicability to discrete vs. continuous settings.

Chapter 3 – Hybrid Differential Privacy Framework: This chapter presents a new framework tailored to hybrid quantum-classical algorithms with fixed quantum measurements:

- It introduces two definitions of neighboring quantum states (trace-distance-based and measurement-based), and compares HDP to classical DP and quantum DP in terms of assumptions and applicability;
- It establishes two key theoretical properties: the post-processing theorem and the composition theorem for HDP;
- It proposes two practical mechanisms to realize HDP: (i) quantum depolarizing noise and (ii) a measurement-based exponential mechanism (MBEM) adapted to the discrete outcome space;
- It provides numerical experiments analyzing the privacy-utility trade-offs across different circuit sizes, noise strengths, and measurement structures.

Chapter 4 – Rényi Differential Privacy in Quantum Machine Learning: This chapter extends Rényi differential privacy to quantum learning algorithms with fixed measurements:

- It defines quantum RDP based on classical Rényi divergence over output distributions induced by fixed measurements;
- It proves post-processing and composition properties for quantum RDP;
- It establishes an equivalent condition and an approximate converse for verifying quantum Rényi differential privacy guarantees under fixed measurements, and in-

troduces both analytical and numerical methods for computing optimal privacy parameters;

- It validates the proposed methods through experiments on variational quantum circuits, comparing different noise types and RDP bounds with classical DP results.

Chapter 5 – Conclusion and Future Directions: This chapter concludes the thesis by summarizing its main contributions. The proposed HDP framework provides a practical solution for achieving differential privacy in hybrid quantum-classical algorithms with fixed measurements. The extension of Rényi differential privacy to quantum learning offers a flexible tool for analyzing privacy-utility trade-offs. Together, these results lay the groundwork for privacy-preserving quantum machine learning.

Future research may extend these frameworks to multi-measurement or adaptive scenarios, explore optimal hybrid noise mechanisms, and investigate the connection between quantum RDP and generalization error. Experimental validation on real quantum hardware and applications in federated quantum learning also present promising directions.

Chapter 2

Preliminaries

This chapter introduces the foundational concepts and technical background relevant to our study of differential privacy in hybrid quantum-classical algorithms. We consider two complementary approaches: one based on extending classical ϵ-differential privacy to hybrid settings, and another based on applying Rényi differential privacy to quantum machine learning circuits. In the hybrid setting, privacy can be preserved using either classical or quantum noise. To support these developments, we present necessary notations, review classical and quantum Rényi divergence, and introduce key concepts from quantum information and parameterized quantum circuits.

First, we introduce some notations and terminologies used in this thesis.

Symbol	**Meaning**
$\|\cdot\rangle$ ($\langle\cdot\|$)	Column (row) vectors in Hilbert space $\mathcal{H}$
$\|i\rangle$	i^{th} computational basis vector
ρ, σ	Density matrices representing quantum states
$\mathcal{L}(V)$	Set of all linear operators acting on vector space V
$\mathbb{1}$ ($\mathbb{1}_V$)	Identity operator on $\mathbb{C}^2$ (or on space V)
$U(\theta)$	Parameterized quantum circuit with parameters θ
$\mathcal{E}$	Quantum channel (CPTP map)
$\{M_k\}_{k\in\mathcal{O}}$	POVM measurement operators with classical outcome set $\mathcal{O}$
$\text{Tr}(A)$	Trace of operator A
$\|\cdot\|_{\text{tr}}$	Trace norm: $\|A\|_{\text{tr}} = \sum_i \omega_i$ where ω_i are singular values
$\|\cdot\|_\infty$	Operator norm (largest singular value)
$\rho \sim \sigma$	Neighboring quantum states for DP definitions
$D_\alpha(P\|Q)$	Rényi divergence of order α between distributions P and Q
ϵ, δ	Privacy parameters for differential privacy
$\mathbb{N}, \mathbb{R}, \mathbb{C}$	Sets of natural, real, and complex numbers
α	Rényi divergence order parameter
η	Privacy radius
$\mathcal{E}_{\text{Dep}}$	Depolarizing channel

2.1 Quantum Basics

In this subsection, we briefly review basic quantum concepts essential for understanding quantum machine learning and quantum differential privacy. We cover the definitions of *quantum states*, *quantum circuits*, *quantum measurements*, and *quantum encoding* techniques. These fundamental components lay the groundwork for the formalization of quantum privacy mechanisms under fixed measurement settings. For a comprehensive and authoritative treatment of these topics, we refer the reader to Nielsen and Chuang [NC01].

2.1.1 Quantum States

Input Quantum State: Ideally, a quantum algorithm takes a *pure quantum state* as input, which is mathematically represented as a complex unit column vector in a 2^n-dimensional *Hilbert (linear) space* $\mathcal{H}$, where n represents the number of *quantum bits (qubits)* in the space. Such states are commonly denoted by the Dirac notation like $|\psi\rangle$ or $|\phi\rangle$. The inner product of two quantum states $|\psi\rangle$ and $|\phi\rangle$ is denoted as $\langle\psi|\phi\rangle$. The magnitude of a pure quantum state is normalized to 1, which can be expressed as $\|\psi\| = \sqrt{\langle\psi|\psi\rangle} = 1$.

The state space of a qubit is specifically a 2-dimensional Hilbert space. When the computational basis (a set of mutually orthogonal pure quantum states) of $\mathcal{H}$ is defined as $|0\rangle = (1,0)^\top$ and $|1\rangle = (0,1)^\top$, a qubit denoted as $|\psi_q\rangle$ can be expressed as

$$|\psi_q\rangle = a\,|0\rangle + b\,|1\rangle = (a,b)^\top,$$

where a and b represent complex numbers and serve as the amplitudes of $|\psi_q\rangle$ satisfying the normalization condition $|a|^2 + |b|^2 = 1$, and $\top$ is the transpose operation. Moreover, the 1-qubit basis pure states, $|0\rangle$ and $|1\rangle$, can be merged to form the computational basis of more extensive quantum systems using their tensor product: $|k_0k_1\rangle = |k_0\rangle \otimes |k_1\rangle$ for

$k_0, k_1 \in \{0, 1\}$. In the case of a 2-qubit system, there are four basis states:

$$|00\rangle = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, |01\rangle = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix}, |10\rangle = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}, |11\rangle = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}.$$

The tensor product of the basis states $|0\rangle$ and $|1\rangle$ represents the binary string formed by the classical bits 0 and 1.

In current Noisy Intermediate-Scale Quantum (NISQ) computers, the input state for a hybrid quantum-classical algorithm becomes a *mixed quantum state* due to noise affecting it. This mixed state is typically represented by an ensemble $\{(p_k, |\psi_k\rangle)\}_k$, where each $|\psi_k\rangle$ is the state with probability p_k. An n-qubit mixed state can be expressed using a 2^n-by-2^n density matrix ρ as $\rho = \sum_k p_k |\psi_k\rangle \langle\psi_k|$, where $\langle\psi_k| = |\psi_k\rangle^\dagger$ is the Hermitian adjoint (complex entry-wise conjugate transpose) of $|\psi_k\rangle$ and $|\psi_k\rangle \langle\psi_k|$ represents the outer product of $|\psi_k\rangle$ with itself. For example, for a qubit state $|\psi_q\rangle$ defined before, we have

$$\langle\psi_q| = a^* \langle 0| + b^* \langle 1| = a^* \cdot (1, 0) + b^* \cdot (0, 1) = (a^*, b^*).$$

In particular, a pure quantum state $|\psi\rangle$ can be represented as an degenerated ensemble $\{(1, |\psi\rangle)\}$ with a density matrix form $\psi = |\psi\rangle\langle\psi|$. For the 1-qubit case, we have

$$\psi_q = |\psi_q\rangle\langle\psi_q| = \begin{pmatrix} a \\ b \end{pmatrix} (a^*, b^*) = \begin{pmatrix} aa^* & ab^* \\ ba^* & bb^* \end{pmatrix}.$$

Mathematically, a density matrix ρ is a positive semi-definite matrix with a unit trace, meaning that the sum of its diagonal elements equals one, denoted as $\mathrm{tr}(\rho) = 1$. In this paper, we study differential privacy for hybrid quantum-classical algorithms and thus investigate the set of mixed quantum states in the Hilbert space $\mathcal{H}$, which is denoted as $\mathcal{D}(\mathcal{H})$. For ease of discussion, we will henceforth refer to quantum states as mixed quantum

states.

2.1.2 Quantum Circuits

A quantum circuit denoted as $\mathcal{E}$ comprises a series of quantum gates U_i represented as

$$U = U_d \cdots U_1,$$

where each U_i stands for a quantum gate and d indicates the circuit's depth. When acting on an input state ρ, the resulting state of $\mathcal{E}$ is determined by the equation:

$$\mathcal{E}(\rho) = U\rho U^\dagger = U_d(\cdots(U_1\rho U_1^\dagger)\cdots)U_d^\dagger, \tag{2.1}$$

where $U_i^\dagger$ represents the Hermitian adjoint of U_i. In mathematical terms, a quantum gate U is a unitary matrix ($UU^\dagger = U^\dagger U = I$, the identity matrix) of size 2^n-by-2^n operating on an n-qubit Hilbert space $\mathcal{H}$.

We now introduces standard quantum gates frequently used in quantum computation and hybrid quantum-classical algorithms, particularly in parameterized quantum circuits for quantum machine learning.

Single-qubit logic gates. Among the commonly employed one-qubit logic gates are the Pauli gates X, Y, and Z, which correspond to the basic quantum bit-flip and phase-flip operations:

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

In addition, the Hadamard gate H plays a key role in creating superposition:

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

Other frequently used single-qubit gates include the phase gate S and the $\pi/8$ gate T:

$$S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}, \quad T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}.$$

1-qubit rotation gates. Single-qubit rotation gates apply unitary rotations along different axes of the Bloch sphere. These are especially important for parameterized quantum circuits and data encoding:

$$R_x(\theta) = e^{-i\theta X/2} = \begin{pmatrix} \cos\frac{\theta}{2} & -i\sin\frac{\theta}{2} \\ -i\sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{pmatrix},$$
$$R_y(\theta) = e^{-i\theta Y/2} = \begin{pmatrix} \cos\frac{\theta}{2} & -\sin\frac{\theta}{2} \\ \sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{pmatrix},$$
$$R_z(\theta) = e^{-i\theta Z/2} = \begin{pmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{pmatrix}.$$

These rotation gates are widely used both to encode classical data into quantum states and to introduce trainable parameters in quantum machine learning circuits.

Controlled gates. Two-qubit gates such as the controlled-U gates apply a one-qubit unitary U to the target qubit conditioned on the control qubit being in state $|1\rangle$. The most notable example is the controlled-NOT (CNOT or CX) gate:

$$\mathrm{CX} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

Another example is the controlled-Z gate:

$$\mathrm{CZ} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}.$$

Parameterized controlled rotation gates, such as $\mathrm{CR}_X(\theta)$, are also used in quantum circuits:

$$\mathrm{CR}_X(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cos\frac{\theta}{2} & -i\sin\frac{\theta}{2} \\ 0 & 0 & -i\sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{pmatrix}.$$

These controlled operations enable entanglement and conditional transformations, essential for universal quantum computation and the expressive power of variational quantum models.

2.1.3 Quantum Channel

Hybrid quantum-classical algorithms on NISQ computers encounter noise effects that cause quantum gates to deviate from behaving unitarily, leading to the introduction of uncertainty in the quantum circuit. A quantum circuit typically consists of a sequence of quantum gates (unitary operations) followed by quantum measurements. Mathematically, such a circuit can be interpreted as a *quantum operation* represented by a super-operator $\mathcal{E}$.

In particular, the unitary gates within the circuit transform the quantum state according to $U\rho U^\dagger$. However, due to the presence of noise, decoherence, and measurements, the evolution of the quantum state in realistic circuits becomes non-unitary. This evolution can be described using the framework of completely positive and trace-preserving (CPTP) maps, under which the quantum circuit as a whole can be viewed as a super-operator $\mathcal{E}$

acting on the quantum input state.

The map $\mathcal{E}$ is a physically valid quantum operation if and only if it satisfies the following *completely positive and trace-preserving* conditions:

- **Completely Positive (CP):** A linear map $\mathcal{E} : \mathcal{D}(\mathcal{H}) \to \mathcal{D}(\mathcal{H})$ is completely positive if, for any auxiliary Hilbert space $\mathcal{H}'$ of arbitrary dimension, the extended map $\mathcal{E} \otimes \mathcal{I}_{\mathcal{H}'}$ maps positive semidefinite operators to positive semidefinite operators:

$$\rho \geq 0 \;\Rightarrow\; (\mathcal{E} \otimes \mathcal{I}_{\mathcal{H}'})(\rho) \geq 0.$$

  This ensures that $\mathcal{E}$ preserves positivity even when applied to part of an entangled system.

- **Trace-Preserving (TP):** The map $\mathcal{E}$ is trace-preserving if it preserves the trace of every input state:

$$\mathrm{Tr}[\mathcal{E}(\rho)] = \mathrm{Tr}[\rho], \quad \forall \rho \in \mathcal{D}(\mathcal{H}).$$

  This guarantees that the total probability of measurement outcomes remains normalized after the quantum operation.

When a map $\mathcal{E}$ satisfies both CP and TP conditions, it is referred to as a *quantum channel* or a *quantum operation*. Importantly, any sequence of unitary gates followed by measurements in a quantum circuit can be equivalently described by such a CPTP map, where measurements are incorporated via the measurement postulate into the evolution, resulting in a probabilistic map over the post-measurement states or their classical outcomes.

Mathematically, the super-operator $\mathcal{E}$ can be expressed in the *Kraus operator-sum representation*, where a finite set $\{E_k\}_{k \in K}$ of matrices exists, satisfying the normalization condition $\sum_{k \in K} E_k^\dagger E_k = I$, such that

$$\mathcal{E}(\rho) = \sum_{k \in K} E_k \rho E_k^\dagger. \tag{2.2}$$

Here, $\{E_k\}_{k\in K}$ are the Kraus operators associated with the quantum circuit, encompassing the effects of unitary gates, noise, and measurements.

In the ideal case without noise and measurements, the Kraus representation reduces to a single unitary evolution $\mathcal{E}(\rho) = U\rho U^\dagger$, which aligns with Eq. 2.1.

2.1.4 Quantum Measurements

In the realm of hybrid quantum-classical algorithms, a vital concluding stage involves a quantum measurement denoted as $\mathcal{M}$. This measurement extracts classical information from the quantum system, with the process itself being randomized. Quantum measurement $\mathcal{M}$ is essentially a randomized function from the set of quantum states $\mathcal{D}(\mathcal{H})$ to the finite set of potential measurement outcomes $\mathcal{O}$, characterized by a set of positive matrices $\{M_i\}_{i\in\mathcal{O}}$ defined on its state space (Hilbert space) $\mathcal{H}$. In particular, if the quantum state $\rho' = \mathcal{E}(\rho)$ represents the output of the quantum circuit $\mathcal{E}$ before the measurement $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$, then the probability p_i of observing a specific measurement outcome $i \in \mathcal{O}$ is determined by

$$p_i = \Pr[\mathcal{M}(\rho') = i] = \mathrm{Tr}(M_i\rho').$$

To ensure that $\{p_i\}_{i\in\mathcal{O}}$ forms a valid probability distribution across $\mathcal{O}$, the sum of all M_i must equate to the identity matrix I (i.e., $\sum_{i\in\mathcal{O}} M_i = I$), so that

$$\sum_{i\in\mathcal{O}} p_i = \sum_{i\in\mathcal{O}} \mathrm{Tr}(M_i\rho') = \mathrm{Tr}\left[\left(\sum_{i\in\mathcal{O}} M_i\right)\rho'\right] = \mathrm{Tr}(\rho') = 1.$$

This type of quantum measurement is commonly known as the *Positive Operator-Valued Measure* (POVM), which focuses solely on the measurement outcomes without considering the post-measurement state, distinct from ρ, which is influenced by the observed outcome. This differs from classical systems where observation does not alter the state of the system.

2.1.5 Hybrid Quantum-Classical Algorithms

In this study, we explore the differential privacy aspect of hybrid quantum-classical algorithms. Essentially, these algorithms involve an *input quantum state* $\rho_{\vec{v}}$ (which encodes an input classical vector $\vec{v}$), a *quantum circuit* $\mathcal{E}$ used for executing computational tasks, and a *quantum measurement* $\mathcal{M}$ employed to reveal classical results through measurement outcomes following a specific probability distribution. The operational sequence of hybrid algorithms is illustrated in Fig. 2.1 adopting the *Schrödinger picture* that emphasizes the evolution of quantum states.

Now we can give a formal definition for hybrid quantum-classical algorithms.

Definition 1. *A hybrid quantum-classical algorithm* $\mathcal{A} = (\mathcal{E}, \mathcal{M} = \{M_i\}_{i\in\mathcal{O}})$ *on a Hilbert space* $\mathcal{H}$ *is a randomized function* $\mathcal{A} : \mathcal{D}(\mathcal{H}) \to \mathcal{O}$ *from quantum state set* $\mathcal{D}(\mathcal{H})$ *to classical measurement outcome set* $\mathcal{O}$ *satisfying the measurement outcome distribution:*

$$\Pr[\mathcal{A}(\rho) = i] = \mathrm{Tr}(M_i\mathcal{E}(\rho)) \quad \forall i \in \mathcal{O},\ \rho \in \mathcal{D}(\mathcal{H}).$$

Additionally, a contrasting perspective known as the *Heisenberg picture* will be introduced later, focusing on the evolution of quantum measurements to facilitate the design of differentially private quantum measurements. Hereafter, we will discuss each of these four components of hybrid quantum-classical algorithms individually.

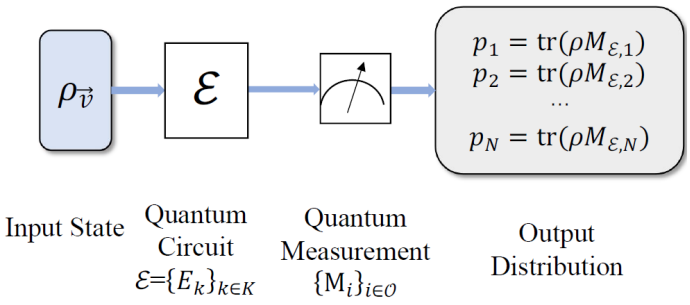

Figure 2.1: Hybrid quantum-classical algorithm

2.1.6 Quantum Encoding

Studying classical problems using hybrid quantum-classical algorithms often involves employing quantum encoding as a crucial method. Quantum encoding facilitates the conversion of classical information into quantum states. Various quantum encoding techniques are commonly utilized, including basis encoding, amplitude encoding, and others [ADK22; LC20; Kor+22].

1. *Basis encoding:* In basis encoding, a classical bit string is directly linked to one of the computational bases of the quantum system. For instance, a classical 3-bit string "101" would be encoded as a 3-qubit pure state represented as $|101\rangle$ and further as a mixed state $|101\rangle\langle 101|$.

2. *Amplitude encoding:* With amplitude encoding, a classical normalized vector $\vec{\nu} = (\nu_0, ..., \nu_{n-1})$ is transformed into pure state $|\vec{\nu}\rangle$ by assigning amplitudes and then $\rho_{\vec{v}}$:

$$|\vec{\nu}\rangle = \sum_{i=0}^{n-1} \nu_i |i\rangle \text{ and } \rho_{\vec{\nu}} = |\vec{\nu}\rangle\langle\vec{\nu}| . \tag{2.3}$$

   For example, encoding normalized vector $\vec{\mu} = \left(\frac{3}{5}, \frac{4}{5}\right)$ results in pure state $|\vec{\mu}\rangle = \frac{3}{5}|0\rangle + \frac{4}{5}|1\rangle$ and further mixed state $\rho_{\vec{m}u} = |\vec{\mu}\rangle\langle\vec{\mu}|$. Additional unnormalized vectors can be normalized and subsequently encoded into quantum states using this method.

2.1.7 Heisenberg Picture

In the following discussion, it is vital to introduce the Heisenberg picture, which examines the progression of quantum measurements (from right to left in the sequence depicted in Fig. 2.1), for designing differentially private quantum measurements, in contrast to the Schrödinger picture that focuses on the development of quantum states as previously mentioned (from left to right in the sequence shown in Fig. 2.1). To clarify this concept, when analyzing a quantum measurement denoted by $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$ and a quantum circuit represented by $\mathcal{E}$ with Kraus matrices $\{E_j\}$, given a quantum state ρ, one can perform the following calculations to determine the probability $\Pr[\mathcal{M}(\mathcal{E}(\rho)) = i] = \mathrm{tr}(M_i\mathcal{E}(\rho))$ of

detecting measurement outcome i:

$$\mathrm{tr}(M_i \sum_j E_j \rho E_j^\dagger) = \mathrm{tr}(\sum_j E_j^\dagger M_i E_j \rho) = \mathrm{tr}(\mathcal{E}^\dagger(M_i)\rho). \tag{2.4}$$

Here, the initial equation arises from the commutative nature of the trace operation ($\mathrm{tr}(AB) = \mathrm{tr}(BA)$ for matrices A and B). The symbol $\mathcal{E}^\dagger$ denotes the adjoint function of $\mathcal{E}$, which is defined as $\mathcal{E}^\dagger(\rho) = \sum_j E_j^\dagger \rho E_j$. In essence, the quantum circuit $\mathcal{E}$ is utilized to represent the transformation of the input quantum state ρ under the measurement $\mathcal{M}$. On the other hand, $\mathcal{E}^\dagger$ is employed to elucidate the transformation of the quantum measurement $\mathcal{M}$ into a new measurement $\mathcal{M}_\mathcal{E} = \{\mathcal{E}^\dagger(M_i)\}_{i\in\mathcal{O}}$, allowing the measurement outcome probability to be interpreted as a quantum state ρ measured by $\mathcal{M}_\mathcal{E}$. This concept is further illustrated below.

$$\Pr[\mathcal{M}(\mathcal{E}(\rho)) = i] = \Pr[\mathcal{M}_\mathcal{E}(\rho) = i] = \mathrm{tr}(\mathcal{E}^\dagger(M_i)\rho).$$

Thus, the statistical development of hybrid quantum-classical algorithm $\mathcal{A} = (\mathcal{E}, \mathcal{M})$ can be replicated through the evolution of a quantum measurement $\mathcal{M}_\mathcal{E}$.

To help readers better understand the Heisenberg picture, let us illustrate it with the following example.

Example 2.1.1. *We have selected a 3-qubit quantum circuit known as the GHZ circuit, which is designed to produce a GHZ state. These states find important applications in quantum computing and communication, serving purposes like quantum error correction and secure quantum key distribution [GHZ89; Mer90; CFS02]. The specific configuration of the circuit is*

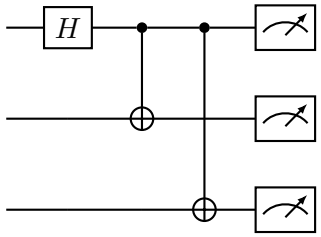

After the quantum circuit, we employ the quantum measurement $\mathcal{M} = \{M_i\}_{0\leq i\leq 7}$ *to retrieve the classical information:*

$$\begin{aligned}
M_0 &= |000\rangle\langle 000|, \quad M_1 = |001\rangle\langle 001|, \\
M_2 &= |010\rangle\langle 010|, \quad M_3 = |011\rangle\langle 011|, \\
M_4 &= |100\rangle\langle 100|, \quad M_5 = |101\rangle\langle 101|, \\
M_6 &= |110\rangle\langle 110|, \quad M_7 = |111\rangle\langle 111|.
\end{aligned}$$

As stated earlier in the Heisenberg picture, the GHZ circuit and the quantum measurement can be considered together as a new quantum measurement. Utilizing Eq. 2.4, we can derive the measurement as $\mathcal{M}_{GHZ} = \{M_{GHZ,i}\}_{0\leq i\leq 7}$ *with*

$$\begin{aligned}
M_{GHZ,0} &= \tfrac{1}{2}|000\rangle\langle 000| + \tfrac{1}{2}|100\rangle\langle 100|, \\
M_{GHZ,1} &= \tfrac{1}{2}|001\rangle\langle 001| + \tfrac{1}{2}|101\rangle\langle 101|, \\
M_{GHZ,2} &= \tfrac{1}{2}|010\rangle\langle 010| + \tfrac{1}{2}|110\rangle\langle 110|, \\
M_{GHZ,3} &= \tfrac{1}{2}|011\rangle\langle 011| + \tfrac{1}{2}|111\rangle\langle 111|, \\
M_{GHZ,4} &= \tfrac{1}{2}|011\rangle\langle 011| + \tfrac{1}{2}|111\rangle\langle 111|, \\
M_{GHZ,5} &= \tfrac{1}{2}|010\rangle\langle 010| + \tfrac{1}{2}|110\rangle\langle 110|, \\
M_{GHZ,6} &= \tfrac{1}{2}|001\rangle\langle 001| + \tfrac{1}{2}|101\rangle\langle 101|, \\
M_{GHZ,7} &= \tfrac{1}{2}|000\rangle\langle 000| + \tfrac{1}{2}|100\rangle\langle 100|.
\end{aligned}$$

Remark 1. *In the above derivation, the operators* $\{M_{\mathrm{GHZ},i}\}$ *represent the effective POVM obtained by absorbing the GHZ circuit into the measurement in the Heisenberg picture. These operators are not intended to be implemented as a physical measurement, but instead provide a mathematically equivalent representation of the measurement statistics induced by the GHZ circuit followed by the computational-basis measurement. Due to the symmetry of the GHZ transformation, several of these effective operators coincide (e.g.,* $M_{\mathrm{GHZ},3} = M_{\mathrm{GHZ},4}$*), reflecting that certain computational-basis outcomes become indistinguishable after the action of the circuit. This degeneracy does not affect the validity of the*

original measurement $\mathcal{M}$ nor the analysis in this section, as the POVM used in practice remains the standard projective measurement in the computational basis.

2.2 Differential Privacy

In the classical world, differential privacy aims to safeguard privacy leakages from neighboring relations of datasets. This concept has been adapted to quantum scenarios by defining different types of neighboring relationships among quantum states. In this section, we review the classical differential privacy framework and its extension to quantum differential privacy.

2.2.1 Classical Differential Privacy

In 1977, statistician Tore Dalenius introduced a privacy objective for statistical datasets: anything that can be learned about a member in the statistical dataset, should also be learnable without access to the dataset [Dal77]. Subsequently, it was demonstrated that achieving this goal was not feasible [Dwo06]. In the work presented in [Dwo+06a], a technique called (classical) differential privacy was proposed. This method entails introducing noise or perturbations into the data, thereby making it challenging to accurately pinpoint the details of individual entities during data analysis and thus preserving their privacy. The classical differential privacy framework requires that modifying a single entry in a dataset induces only a small change in the output distribution of queries, as observed by an adversary.

In the classical differential privacy framework, the input dataset is typically represented as a vector $\vec{\nu}$ of n records (or individuals), where each record is drawn from a domain Ω such as $\{0,1\}^d$ or a bounded subset of $\mathbb{R}^d$. Within this context, a classical computational process is represented by a randomized function denoted as $\mathcal{K}$. The dataset $\vec{\nu}$ serves as an input to the computation through $\mathcal{K}$, resulting in $\mathcal{K}(\vec{\nu})$. This output belongs to Range($\mathcal{K}$), representing the range of function $\mathcal{K}$. Two datasets $\vec{\nu}$ and $\vec{\omega}$ are said to be *neighboring*, denoted $\vec{\nu} \sim \vec{\omega}$, if they differ in exactly one individual's record (i.e., they differ in only one

coordinate).

The foundational notion is the so-called ϵ-*differential privacy*, which requires that for any two neighboring datasets that differ in only one individual's data, the output distributions remain nearly indistinguishable.

Definition 2 (ϵ-Differential Privacy)**.** *Let $\epsilon \geq 0$. A randomized function $\mathcal{K}$ satisfies ϵ-differential privacy if for all neighboring datasets $\vec{\nu}, \vec{\omega} \in \Omega$ with $\vec{\nu} \sim \vec{\omega}$, and for any subset $S \subseteq Range(\mathcal{K})$, it holds that*

$$\Pr[\mathcal{K}(\vec{\nu}) \in S] \leq \exp(\epsilon) \cdot \Pr[\mathcal{K}(\vec{\omega}) \in S].$$

While this strong guarantee offers robust privacy, it can be too strict in practical scenarios, especially when the output distribution may exhibit rare but significant deviations for a small subset of the output space. To accommodate this, a relaxed variant known as (ϵ, δ)-*differential privacy* is often used in practice.

Definition 3 ((ϵ, δ)-Differential Privacy)**.** *Let $\epsilon \geq 0$ and $0 \leq \delta < 1$. A randomized function $\mathcal{K}$ satisfies* (ϵ, δ)-differential privacy *if for all neighboring datasets $\vec{\nu}, \vec{\omega} \in \Omega$ with $\vec{\nu} \sim \vec{\omega}$, and for any subset $S \subseteq Range(\mathcal{K})$, it holds that*

$$\Pr[\mathcal{K}(\vec{\nu}) \in S] \leq \exp(\epsilon) \cdot \Pr[\mathcal{K}(\vec{\omega}) \in S] + \delta.$$

Here, the parameter δ allows for a small probability of the privacy bound being violated, which provides flexibility in the presence of unlikely but high-impact outputs. In essence, (ϵ, δ)-differential privacy is a natural relaxation of ϵ-differential privacy. It permits significantly improved utility since it avoids the need to inject excessive noise in order to satisfy the strict ϵ-DP guarantee globally.

In real-world applications, particularly when the output space contains outliers or highly sensitive regions, the difference $\mathcal{K}(\vec{\nu})$ vs. $\mathcal{K}(\vec{\omega})$ can be large with very small probability. Enforcing strict ϵ-DP in such cases would require a large amount of noise, leading to severe utility degradation. The (ϵ, δ) formulation instead ensures that the stronger

ϵ-DP guarantee holds with high probability at least $1 - \delta$, which is typically sufficient for practical privacy needs.

Differential privacy has therefore become a widely accepted standard for measuring how well an algorithm protects individual privacy, with its theoretical foundations in classical machine learning and data analytics now well established [DR+14; JLE14]. However, designing algorithms with rigorous differential privacy guarantees is subtle and error-prone, and many published algorithms have been shown to violate differential privacy in practice. This challenge has led to the development of formal frameworks for verifying the differential privacy of classical algorithms [BO13; Bar+16a; Bar+14; Bar+13; Bar+12; Bar+16b], with a wide range of verification techniques extended into this domain to enable systematic analysis and debugging of privacy guarantees in real-world systems.

In practice, mechanisms such as the Laplace and Gaussian mechanisms [DR+14] are used to achieve these privacy guarantees, depending on the choice of privacy parameters. A more detailed introduction to these mechanisms will be presented later in this section.

Weakened Variants of Differential Privacy

While (ϵ, δ)-differential privacy provides a practical relaxation of the strict ϵ-DP definition, it remains a worst-case guarantee, which can still be overly conservative when analyzing privacy loss under multiple invocations or in learning applications. To address this, a line of research has proposed *weakened* or *relaxed* variants of differential privacy, which offer more refined tools for analyzing privacy-utility trade-offs, particularly in machine learning.

A prominent class of such relaxations is based on viewing the *privacy loss* as a random variable and controlling its concentration. One representative formalism is *concentrated differential privacy (CDP)* [DR16], which upper bounds the moment generating function of the privacy loss. CDP enables tighter analysis of composition and postprocessing, and has been widely adopted in private empirical risk minimization and federated learning scenarios. Variants of CDP include zero-concentrated differential privacy (zCDP) and Rényi differential privacy (RDP), which use different divergences to define privacy loss.

In particular, *Rényi differential privacy (RDP)* [Mir17] formalizes privacy in terms of Rényi divergence, a parameterized family of divergences that generalizes the Kullback-Leibler divergence. RDP is analytically convenient and enjoys strong composition properties. It can be viewed as a specific instantiation of the concentrated DP framework using Rényi divergence to quantify privacy leakage. We defer a detailed discussion of RDP to a later section, where it serves as the central tool for analyzing privacy in hybrid quantum-classical algorithms.

Beyond CDP and RDP, many other weakenings of differential privacy have been proposed, each capturing different notions of average-case or relaxed guarantees. Some notable examples include:

- **Approximate DP** [Dwo+06b]: Is based on the (ϵ, δ) definition, bounding the *approximate max divergence* $D_\infty^\delta(P\|Q)$.
- **Personalized DP** [ESS15; LWS15]: Allows per-user bounds such as $D_\infty(P\|Q)$ or $D_\infty^\delta(P\|Q)$ for each individual $z \in Z$.
- **KL-Privacy** [WLF16; BD14; Bas+16]: Is based on bounding the Kullback-Leibler divergence $D_{\mathrm{KL}}(P\|Q)$.
- **TV-Privacy** [BD14; Bas+16]: Controls total variation distance $\|P - Q\|_{\mathrm{TV}}$.
- **Randomized Privacy (Rand-Privacy)** [HRW11]: Allows the privacy guarantee to hold with high probability over a random choice of data universe.
- **On-Average KL-Privacy** [WLF16]: Controls the expected privacy loss over a distribution of data points.

These relaxations often trade off some level of interpretability or worst-case protection for improved utility, analytical tractability, or compatibility with probabilistic learning frameworks. The choice of relaxation depends on the application domain, threat model, and composition needs.

2.2.2 Quantum Differential Privacy

The concept of differential privacy, originally introduced for classical datasets, can be extended to quantum settings to protect sensitive quantum data. In this section, we formalize the notion of quantum differential privacy (QDP).

Definition 4 (Quantum Differential Privacy)**.** *[ZY17] Let $\epsilon \geq 0$ and $0 \leq \delta < 1$. A quantum circuit $\mathcal{E}$ is said to be* (ϵ, δ)-quantum differentially private *if for all neighboring quantum states $\rho \sim \sigma$, any measurement $\mathcal{M} = \{M_i\}_{i \in \mathcal{O}}$, and any subset $S \subseteq \mathcal{O}$, it holds that*

$$\sum_{i \in S} tr(M_i \mathcal{E}(\rho)) \leq \exp(\epsilon) \cdot \sum_{i \in S} tr(M_i \mathcal{E}(\sigma)) + \delta.$$

Here, just as in the classical differential privacy framework, setting $\delta = 0$ yields the pure ϵ-QDP variant.

It is important to recall that in classical differential privacy, neighboring datasets refer to two datasets differing in only one individual entry. In the quantum setting, however, the objects of interest are quantum states, which are represented as density matrices as introduced in Section 2.1. These quantum states possess more intricate mathematical properties than classical vectors, and encode classical information through mechanisms such as basis encoding or amplitude encoding. As a result, defining a notion of "neighboring quantum states" becomes more subtle, and multiple definitions have been proposed in the literature. Two commonly adopted definitions are as follows:

1. *Trace distance-based definition* [ZY17]: Two quantum states ρ and σ are considered η-neighboring for a fixed constant $0 \leq \eta \leq 1$ if their trace distance satisfies

$$\tau(\rho, \sigma) = \frac{1}{2} \mathrm{tr}(|\rho - \sigma|) \leq \eta,$$

where $|\cdot|$ denotes the trace norm of a matrix A, defined as $|A| = \sqrt{A^\dagger A}$ [NC01]. The trace distance is a fundamental metric in quantum information theory and quantifies the statistical distinguishability of two quantum states, making it a natural choice

for defining neighboring states under QDP.

2. *Local operation-based definition* [AR19]: Two quantum states ρ and σ are neighboring if there exists a local quantum operation (e.g., a single-qubit unitary gate) that acts on at most one qubit and transforms ρ into σ or vice versa. For instance, the 3-qubit pure states $|000\rangle\langle 000|$ and $|100\rangle\langle 100|$ are neighboring under this definition, as they can be transformed into each other by applying an X-gate (bit-flip) on the first qubit. This approach resembles the classical notion of neighboring datasets, where binary vectors such as "000" and "100" differ in only one bit.

A hybrid definition that combines the trace distance and local operation approaches was also proposed in [ADK23], aiming to bridge the gap between mathematical rigor and operational interpretability.

Neighboring-preserving Quantum encoding

As discussed in Section 2.1, various quantum encoding techniques exist for transforming classical datasets into quantum states. When a quantum encoding method can maintain the neighboring relationships present in the original classical datasets within the resulting quantum states, it is referred to as a *neighboring-preserving quantum encoding* technique [ADK23; Gua+23]. In other words, if two classical datasets $\vec{\nu}, \vec{\omega}$ are encoded using a neighboring-preserving quantum encoding Λ to produce quantum states $\Lambda(\vec{\nu})$ and $\Lambda(\vec{\omega})$, the neighboring relationship between the original classical datasets should be preserved in the quantum domain, implying that

$$\vec{\nu}, \vec{\omega} \text{ are neighboring} \Rightarrow \Lambda(\vec{\nu}), \Lambda(\vec{\omega}) \text{ are neighboring.}$$

Let us demonstrate with an example how amplitude encoding, as defined in Eq. 2.3, can be used to encode classical data into quantum states in a way that preserves the neighboring relationship. Specifically, this encoding ensures that if two classical inputs are close, then the corresponding quantum states are also η-neighboring under the trace

distance criterion.

Example 2.2.1. *Consider two neighboring classical vector datasets* $\vec{\nu} = (\nu_0, ..., \nu_{n-1}), \vec{\omega} = (\omega_0, ..., \omega_{n-1}) \in \Omega$*, differing only in a single element, let's say the* $(k+1)$*-th element, i.e.* $\nu_k \neq \omega_k$*. Initially, we normalize these vectors to obtain* $\frac{\vec{\nu}}{\|\vec{\nu}\|}$ *and* $\frac{\vec{\omega}}{\|\vec{\omega}\|}$*, where* $\|\vec{\nu}\|$ *and* $\|\vec{\omega}\|$ *represent the norms of* $\vec{\nu}$ *and* $\vec{\omega}$*, respectively. Subsequently, employing the amplitude encoding as depicted in Eq. 2.3, these vectors are transformed into quantum states* $\Lambda(\frac{\vec{\nu}}{\|\vec{\nu}\|})$ *and* $\Lambda(\frac{\vec{\omega}}{\|\vec{\omega}\|})$*, correspondingly.*

The trace distance between these states is given by:

$$\tau\left(\Lambda(\frac{\vec{\nu}}{\|\vec{\nu}\|}), \Lambda(\frac{\vec{\omega}}{\|\vec{\omega}\|})\right) = \sqrt{1 - \left|\frac{\vec{\nu}^\dagger}{\|\vec{\nu}\|} \cdot \frac{\vec{\omega}}{\|\vec{\omega}\|}\right|^2}.$$

Let $M = \max_{\vec{\nu}\in\Omega} \max_i \frac{|\nu_i|^2}{\|\vec{\nu}\|^2}$ *denote the maximum square norm of an element among all normalized vectors. Then assuming* $|\nu_k| \geq |\omega_k|$*, we obtain:*

$$\begin{aligned}
\left|\frac{\vec{\nu}^\dagger}{\|\vec{\nu}\|} \cdot \frac{\vec{\omega}}{\|\vec{\omega}\|}\right|^2 &= \left|\frac{\nu_0\omega_0 + ... + \nu_{n-1}\omega_{n-1}}{\|\vec{\nu}\|\|\vec{\omega}\|}\right|^2 \\
&= \left|\frac{\|\vec{\nu}\|^2 - |\nu_k|^2 + \nu_k\omega_k}{\|\vec{\nu}\|\|\vec{\omega}\|}\right|^2 \\
&\geq \left|\frac{\|\vec{\nu}\|^2 - 2|\nu_k|^2}{\|\vec{\nu}\|^2}\right|^2 \\
&\geq (1 - 2M)^2.
\end{aligned}$$

Consequently, we derive:

$$\tau(\Lambda(\frac{\vec{\nu}}{\|\vec{\nu}\|}), \Lambda(\frac{\vec{\omega}}{\|\vec{\omega}\|})) \leq \sqrt{4M - 4M^2}.$$

Given the arbitrariness of $\vec{\nu}$ *and* $\vec{\omega}$*, we can set* $\eta = \sqrt{4M - 4M^2}$ *to maintain the classical neighboring relationship in the trace distance-based* η*-neighboring relationship within the encoded quantum states through amplitude encoding. Thus, we have successfully implemented a neighboring-preserving quantum encoding utilizing amplitude encoding.*

This example illustrates that neighboring-preserving quantum encodings, such as am-

plitude encoding, allow classical data to be analyzed under quantum differential privacy frameworks. Nevertheless, the correctness of such an approach critically depends on the appropriate definition of neighboring relations and the careful selection of quantum encoding schemes, as these directly influence the privacy guarantees of the resulting quantum algorithms.

2.3 Rényi Differential Privacy

Rényi Differential Privacy (RDP) is a natural and analytically convenient relaxation of differential privacy, introduced to provide tighter and more flexible privacy accounting, particularly under composition. It generalizes classical (ϵ, δ)-differential privacy by replacing worst-case probability ratios with a measure based on Rényi divergence, which captures the statistical distance between output distributions of neighboring datasets more precisely.

Definition 5 (Rényi Differential Privacy [Mir17])**.** *Let $\alpha > 1$ and $\epsilon \geq 0$. A randomized mechanism $\mathcal{K}$ is said to satisfy* (α, ϵ)-Rényi differential privacy *if for all neighboring datasets $D \sim D'$, the Rényi divergence of order α between the output distributions of $\mathcal{K}$ on D and D' is at most ϵ, that is,*

$$D_\alpha(\mathcal{K}(D) \,\|\, \mathcal{K}(D')) \leq \epsilon,$$

where the Rényi divergence of order α is defined as

$$D_\alpha(P \,\|\, Q) = \frac{1}{\alpha - 1} \log \mathbb{E}_{x \sim Q} \left[\left(\frac{P(x)}{Q(x)} \right)^\alpha \right].$$

RDP has several important properties that make it suitable for practical use:

- **Post-processing invariance**: RDP is preserved under arbitrary data-independent post-processing.

- **Additive composition**: If two mechanisms satisfy (α, ϵ_1)- and (α, ϵ_2)-RDP respectively, their composition satisfies $(\alpha, \epsilon_1 + \epsilon_2)$-RDP.

- **Conversion to (ϵ, δ)-DP**: RDP guarantees can be transformed into (ϵ, δ)-DP via a standard bound:

$$\text{If } \mathcal{K} \text{ is } (\alpha, \epsilon)\text{-RDP, then it is also } (\epsilon + \frac{\log(1/\delta)}{\alpha - 1}, \delta)\text{-DP for any } 0 < \delta < 1.$$

Unlike (ϵ, δ)-DP, which provides a binary bound that may include worst-case disclosure with probability δ, RDP tracks higher moments of the privacy loss variable, offering a more nuanced quantification of privacy leakage. This is particularly advantageous when analyzing privacy loss under repeated application of mechanisms (e.g., in iterative learning algorithms), as the additive nature of RDP leads to more accurate bounds with fewer composition losses.

Moreover, RDP facilitates the analysis of common mechanisms:

- The Gaussian mechanism with variance σ^2 satisfies $(\alpha, \frac{\alpha}{2\sigma^2})$-RDP.

- The Laplace mechanism admits a closed-form RDP expression involving exponential functions of α and the scale parameter.

- The randomized response mechanism has RDP bounds expressible via a convex combination of probabilities.

RDP can be visualized through a privacy budget curve parameterized by α, where each point represents the privacy cost under a specific moment of the privacy loss distribution. In practical systems, precomputing and reporting RDP curves at a fixed set of α values enables efficient tracking of cumulative privacy loss over complex workflows.

In the quantum context, the Rényi divergence also serves as a fundamental tool in defining quantum RDP, providing a promising direction for privacy-preserving quantum computations.

2.3.1 Existing Extensions of Rényi Differential Privacy to Quantum Settings.

Several studies have explored how RDP can be generalized to the quantum setting. One such extension appears in [HRF23], where the authors define quantum Rényi differential privacy (QRDP) by replacing the classical Rényi divergence with its quantum analogues. In their formulation, a quantum channel $\mathcal{A}$ is said to satisfy (α, ϵ)-QRDP if for all neighboring quantum states $\rho \sim \sigma$, the quantum Rényi divergence of order α between the output states is bounded as

$$D_\alpha(\mathcal{A}(\rho) \,\|\, \mathcal{A}(\sigma)) \leq \epsilon.$$

Here, D_α can be instantiated using different definitions of quantum Rényi divergence, such as:

- **Petz Rényi divergence**:

$$D_\alpha^{\mathrm{Petz}}(\rho\|\sigma) = \frac{1}{\alpha - 1} \log \mathrm{tr}[\rho^\alpha \sigma^{1-\alpha}],$$

  defined when $\mathrm{supp}(\rho) \subseteq \mathrm{supp}(\sigma)$;

- **Sandwiched Rényi divergence**:

$$\widetilde{D}_\alpha(\rho\|\sigma) = \frac{1}{\alpha - 1} \log \mathrm{tr}\left[\left(\sigma^{\frac{1-\alpha}{2\alpha}} \rho \sigma^{\frac{1-\alpha}{2\alpha}}\right)^\alpha\right],$$

  which satisfies the data processing inequality for $\alpha \geq \frac{1}{2}$.

These definitions preserve many desirable properties of classical RDP, such as post-processing invariance and composability, when appropriate divergence forms are used.

Further, [NGW24] presents a generalized privacy framework called quantum Pufferfish privacy, which incorporates quantum Rényi differential privacy as a special case. In this framework, privacy guarantees are parameterized by a divergence measure (e.g., Rényi) and a set of distinguishable state pairs. A channel $\mathcal{A}$ satisfies RDP under this framework

if

$$D_\alpha(\mathcal{A}(\rho)\|\mathcal{A}(\sigma)) \leq \epsilon$$

for all (ρ, σ) in a specified secret-pair set. This model accommodates both classical and quantum adversaries and generalizes several differential privacy definitions via a common structure.

These foundational works provide formal tools for reasoning about privacy in quantum and hybrid quantum-classical systems, particularly when using information-theoretic divergences like Rényi divergence as a measure of distinguishability.

2.4 Noise Mechanisms for Privacy

Noise mechanisms play a central role in differential privacy, providing a principled means to protect sensitive data while retaining utility. The core idea of differential privacy is to ensure that the inclusion or exclusion of any single individual's data does not significantly affect the output distribution of an algorithm, thereby preventing adversaries from inferring the presence or specific value of any individual data point.

To achieve this, carefully calibrated randomness is systematically introduced into computations. This added noise "masks" the precise contributions of individual data points, making it statistically difficult for an observer to determine whether a particular individual's data was included or to infer its exact value. At the same time, the noise is designed to be as small as possible while still providing rigorous privacy guarantees, preserving the utility of the output for downstream analysis and decision-making.

In essence, noise mechanisms translate theoretical privacy guarantees into practical privacy protection, balancing the trade-off between privacy and utility. Without noise, outputs would precisely reflect the input data, leading to privacy violations. With appropriately calibrated noise, sensitive information remains protected while enabling the extraction of meaningful insights from the data.

Thus, introducing randomness through noise mechanisms is a fundamental technique

in differential privacy, limiting the information that can be inferred about any single input while ensuring that outputs remain statistically similar even when individual entries in the input change. This prevents adversaries from confidently inferring the participation or value of any individual data point based on the output of the computation.

This section provides an overview of both *classical* and *quantum* noise mechanisms used to enforce privacy guarantees in computation.

On the classical side, we review the widely used Laplace and Gaussian mechanisms, which add calibrated noise directly to numeric query outputs, effectively blurring the exact value of the output to protect individual contributions. We also discuss the exponential mechanism, which introduces randomness by selecting outputs according to a utility-based probability distribution, preserving differential privacy while ensuring that higher-utility outputs are still more likely.

On the quantum side, privacy can be achieved by introducing noise via quantum channels, which reduce the distinguishability between quantum states and thus protect against privacy breaches. These noise channels can be applied either before measurement (as part of the quantum computation) or as a post-processing step after measurement. We introduce three representative quantum noise mechanisms: the depolarizing channel (DEP), which uniformly mixes the quantum state and randomizes information; the generalized amplitude damping (GAD) channel, which models energy dissipation to a thermal bath and introduces decoherence; and the phase damping channel, which dephases quantum states while preserving populations, reducing coherence that can leak information.

Together, these mechanisms illustrate how both classical and quantum noise can be harnessed to implement differentially private computations under various settings. By systematically introducing noise, these mechanisms control the sensitivity of output distributions to individual inputs, making it difficult for adversaries to infer specific data contributions while preserving the overall utility of computations.

2.4.1 Classical Noise Mechanisms for Privacy-Preserving Computation.

Recent studies have proposed achieving quantum differential privacy by leveraging the addition of classical noise to the outputs of quantum measurements. This method provides privacy guarantees by randomizing the results of measurements on quantum states, thereby making the output distributions for neighboring quantum states statistically indistinguishable. Specifically, the method employs two widely-used classical noise mechanisms:

- **Laplace Mechanism:** By adding noise drawn from a Laplace distribution with a scale parameter b, where $b \geq \Delta f/\epsilon$ (with Δf representing the sensitivity of the measurement and ϵ the privacy budget), this mechanism ensures that the measurement output remains ϵ-differentially private.

- **Gaussian Mechanism:** Gaussian noise with variance $\sigma^2 \geq 2\ln(1.25/\delta)\Delta^2/\epsilon^2$ is added to the measurement results, achieving (ϵ, δ)-differential privacy. Here, Δ denotes the sensitivity, ϵ is the privacy budget, and δ controls the probability of privacy failure.

Although the above two mechanisms are widely used in classical differential privacy, they are originally designed for applications with continuous output domains. In contrast, the *exponential mechanism* provides a noise-adding strategy tailored for finite output sets. As we will see later, this mechanism is particularly well-suited for quantum differential privacy, since the outputs of quantum measurements typically lie in a finite, discrete set.

Exponential Mechanism [DR+14]

The exponential mechanism is a method used to select the "best" output while ensuring differential privacy, particularly when adding noise directly to the output would significantly degrade its utility. In this mechanism, a utility function is defined to quantify how desirable each possible output is for a given dataset. Using this utility function, the exponential mechanism constructs a probability distribution over all possible outputs, assigning higher probabilities to outputs with higher utility scores while maintaining differential privacy guarantees.

Steps of the Exponential Mechanism

1. **Define Utility Function**: First, define a utility function $u : \mathbb{N}^{|X|} \times R \to \mathbb{R}$, which maps a dataset $\vec{x}$ and output r to a utility score. Here, X denotes the data domain (the set of all possible values of individual data entries), and the dataset $\vec{x}$ is represented as a histogram vector in $\mathbb{N}^{|X|}$, where each component counts the occurrences of the corresponding value in the dataset. The utility function represents the quality or benefit of each possible output.

2. **Compute Utility Scores**: For a given dataset $\vec{x}$ and all possible outputs $r \in R$, compute the utility score $u(\vec{x}, r)$ for each r.

3. **Compute Probability Distribution**: Assign a probability to each possible output based on its utility score. Specifically, the probability of the output r being chosen is proportional to $\exp\left(\frac{\epsilon u(\vec{x},r)}{2\Delta u}\right)$, where Δu is the sensitivity of the utility function, defined as:

$$\Delta u = \max_{r \in R, \vec{x}, \vec{x}'} |u(\vec{x}, r) - u(\vec{x}', r)|$$

   Here, $\vec{x}$ and $\vec{x}'$ are two neighboring datasets.

4. **Normalization**: Normalize these probabilities so that they sum to 1. Specifically, the probability of choosing output r is:

$$P(r|\vec{x}) = \frac{\exp\left(\frac{\epsilon u(\vec{x},r)}{2\Delta u}\right)}{\sum_{r' \in R} \exp\left(\frac{\epsilon u(\vec{x},r')}{2\Delta u}\right)}$$

5. **Sampling**: Sample an output from the set of possible outputs R according to the computed probability distribution.

Intuitive Explanation The core idea of the exponential mechanism is to choose outputs with higher utility scores more frequently by assigning them higher probabilities. This ensures that even with the addition of noise to protect privacy, an output close to the optimal can still be selected. As the utility score decreases, the probability of selection

decreases exponentially, ensuring that high-utility outputs have a greater chance of being chosen than low-utility ones.

We show an example to illustrate this mechanism. Suppose there are five bidders, each bidding a different amount for an item: \$1.00, \$2.00, \$3.00, \$4.00, and \$5.00. The objective is to choose a price that maximizes total revenue. The utility function $u(\vec{x}, p)$ is defined as the total revenue at price p, where $\vec{x}$ denotes the set of all bids. Using the exponential mechanism, we select a price based on this utility function.

For instance, consider the dataset of bids $\vec{x} = [1, 2, 3, 4, 5]$. The utility function $u(\vec{x}, p)$ computes the total revenue for a given price p. If the price is \$1, all five bidders will purchase the item, resulting in a total revenue of \$5. For higher prices, the revenues are as follows: at \$2, the revenue is \$8; at \$3, the revenue is \$9; at \$4, the revenue is \$8; and at \$5, the revenue is \$5.

Next, we apply the exponential mechanism to this example. For a given privacy parameter $\epsilon = 1.0$ and sensitivity $\Delta u = 1$, the selection probability of each price is proportional to $\exp(\epsilon u(\vec{x}, p)/2\Delta u)$. Let

$$Z = \exp(2.5) + \exp(4) + \exp(4.5) + \exp(4) + \exp(2.5)$$

be the normalization constant. After normalizing so that probabilities sum to one, we obtain

$$P(1) = \frac{\exp(2.5)}{Z}, \quad P(2) = \frac{\exp(4)}{Z}, \quad P(3) = \frac{\exp(4.5)}{Z},$$
$$P(4) = \frac{\exp(4)}{Z}, \quad P(5) = \frac{\exp(2.5)}{Z}.$$

Finally, we sample a price from these normalized probabilities. This guarantees that prices with higher utility scores (total revenue) are more likely to be chosen, while still preserving differential privacy.

2.4.2 Quantum Noise Mechanisms for Privacy-Preserving Computation.

In privacy-preserving quantum computation, one approach to achieving differential privacy is to inject quantum noise after the execution of a quantum algorithm. This can be

modeled by composing a quantum channel $\mathcal{E}$, representing the original (ideal) quantum computation, with a noise channel $\mathcal{E}_N$ serving as a privacy mechanism. The resulting channel is described by the composition:

$$\mathcal{E}_N \circ \mathcal{E} : \rho \mapsto \mathcal{E}_N(\mathcal{E}(\rho)),$$

where ρ is the input quantum state. The noise channel $\mathcal{E}_N$ acts as a quantum operation designed to reduce distinguishability between neighboring inputs, thereby protecting privacy.

Several families of quantum noise mechanisms have been proposed for this purpose. The following three are among the most commonly analyzed [ZY17]:

- **Depolarizing noise:** The depolarizing channel represents a uniform randomization over all possible Pauli errors, effectively replacing a quantum state with the maximally mixed state with some probability. For a single-qubit input state ρ, the depolarizing channel with noise rate $\lambda \in [0, 1]$ is defined as:

  $$\mathcal{D}_\lambda(\rho) = (1 - \lambda)\rho + \lambda\frac{I}{2}.$$

  This channel is unital and basis-independent, meaning it affects all quantum states equally regardless of their initial orientation. It transforms ρ toward the maximally mixed state $I/2$, thereby reducing the trace distance between any pair of input states. Its ability to uniformly "flatten" the state makes it a canonical tool in enforcing differential privacy in a quantum context.

- **Generalized Amplitude Damping:** The GAD channel models energy dissipation in a quantum system interacting with a thermal environment at non-zero temperature. It generalizes the amplitude damping channel by allowing relaxation toward a thermal equilibrium state. The channel is parameterized by a damping probability

γ and a thermal excitation probability p, with Kraus operators:

$$E_0 = \sqrt{p}\begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{pmatrix}, \quad E_1 = \sqrt{p}\begin{pmatrix} 0 & \sqrt{\gamma} \\ 0 & 0 \end{pmatrix},$$
$$E_2 = \sqrt{1-p}\begin{pmatrix} \sqrt{1-\gamma} & 0 \\ 0 & 1 \end{pmatrix}, \quad E_3 = \sqrt{1-p}\begin{pmatrix} 0 & 0 \\ \sqrt{\gamma} & 0 \end{pmatrix}.$$

The GAD channel is non-unital and introduces asymmetry in state evolution, driving states toward a fixed point that depends on the thermal environment.

- **Phase Damping (PAD):** Also known as dephasing noise, the PAD channel models the loss of quantum coherence without energy loss. It preserves the diagonal elements of the density matrix but suppresses off-diagonal (coherence) terms. The Kraus representation for phase damping with parameter λ is:

$$\mathcal{Z}_\lambda(\rho) = E_0\rho E_0^\dagger + E_1\rho E_1^\dagger, \quad \text{where } E_0 = \sqrt{1-\lambda}I, \quad E_1 = \sqrt{\lambda}Z.$$

Here Z is the Pauli-Z matrix. This type of noise is especially relevant when the quantum information is stored in populations (i.e., diagonal entries), and is resilient to phase errors.

Among these, the depolarizing mechanism is particularly useful for privacy protection due to its symmetric nature. It guarantees that any two input states are driven closer in trace distance by an amount determined solely by λ. In contrast, GAD and PAD introduce basis-dependent or state-dependent effects, which may offer more targeted noise but require careful calibration based on the specific structure of the data or algorithm.

These noise mechanisms provide concrete implementations of the privacy-preserving map $\mathcal{E}_N$, and their quantitative effects on privacy can be analyzed via the reduction in distinguishability between neighboring inputs.

2.5 Motivation and Problem Statement

Differential privacy (DP) has become a central paradigm for ensuring data confidentiality in classical settings [DR+14], while recent studies have extended its principles to the quantum domain under the name of quantum differential privacy (QDP) [ZY17; AR19]. However, both directions exhibit limitations when applied to hybrid quantum-classical algorithms—particularly those used in variational quantum algorithms (VQAs) and quantum machine learning (QML)—where classical and quantum components are closely interwoven.

On one hand, classical DP techniques typically assume deterministic or continuous outputs, which do not directly apply to inherently probabilistic and discrete quantum measurement results. On the other hand, most existing QDP frameworks treat quantum algorithms as end-to-end quantum processes and seek privacy guarantees against arbitrary quantum measurements. However, such QDP frameworks can be overly conservative or inapplicable in practical hybrid workflows, where only one specific (fixed) quantum measurement is exposed.

Hybrid Differential Privacy (HDP). To address these challenges, our first contribution introduces a hybrid differential privacy framework tailored for hybrid quantum-classical algorithms. In contrast to classical DP, which adds noise to deterministic functions, or traditional QDP, which assumes arbitrary measurement leakage, our HDP model focuses on privatizing the *specific quantum measurement* used in the hybrid workflow.

Our framework supports two complementary strategies: (1) applying *classical noise* to the measurement outcome (e.g., via a measurement-based exponential mechanism), and (2) inserting *quantum noise* before measurement (e.g., depolarizing noise), modifying the quantum state to reduce distinguishability. Unlike existing works that commit to only one type of noise or require fully quantum treatments, our HDP framework flexibly integrates classical and quantum techniques within a unified post-processing and composition analysis.

Rényi DP under Fixed Measurement. Our second contribution focuses on the analysis of Rényi differential privacy (RDP) in the specific setting of fixed quantum measurements, which naturally arise in QML algorithms. While existing quantum RDP definitions [HRF23] are designed for general quantum channels, they do not exploit the structure of fixed measurements, often leading to loose bounds or intractable divergence computation. In contrast, we reframe the problem using the Heisenberg picture and treat the quantum-classical hybrid process as an effective measurement on input states. This allows us to analyze the output distributions via classical Rényi divergence, yielding tighter and tractable privacy guarantees. We also derive subset-based conditions and analytical upper bounds that facilitate practical privacy verification in variational circuits.

Chapter 3

Hybrid Differential Privacy

As discussed in the preceding chapters, hybrid quantum-classical algorithms introduce unique privacy challenges: quantum states are manipulated through quantum circuits, but only classical information, obtained through fixed quantum measurements, is ultimately accessible. This constraint implies that standard quantum differential privacy (which considers arbitrary measurements) can be overly conservative or inapplicable, while classical differential privacy techniques cannot directly handle the quantum component of the workflow.

To address this gap, we introduce a hybrid differential privacy framework that explicitly models the structure of hybrid algorithms. The key idea is to enforce privacy guarantees at the measurement level, by ensuring that the output distribution of measurement results is insensitive to small changes in the quantum input states. Our framework accommodates both classical and quantum noise mechanisms and supports post-processing and composition analysis.

To apply differential privacy in practical hybrid quantum-classical algorithms, it is necessary to incorporate a clear threat model to specify what adversarial behaviors the privacy guarantees are designed to defend against. This makes the DP framework actionable, helps tailor noise mechanisms precisely to realistic scenarios, and enables rigorous quantification of the privacy-utility trade-offs under practical attack settings.

We begin by defining the threat model under which HDP is formulated.

3.1 Threat Model

In this section, we outline the adversary's capabilities and objectives in order to define the threat model within our hybrid quantum-classical differential privacy framework.

In the classical scenario, the adversary has the ability to input queries into classical algorithms, and examine and analyze the results to deduce personal privacy. Considering the hybrid nature of quantum-classical algorithms with quantum input and classical output, it is assumed that the adversary has both quantum and classical capabilities in our hybrid differential privacy framework:

- *Quantum computing capability:* The adversary can input any quantum state into the hybrid algorithm.
- *Classical computing capability:* The adversary can access the measurement results of the hybrid algorithm and analyze the obtained results.

The primary objective of the adversary addressed in this paper is to infer quantum state information (encoded classical information) by submitting quantum input and analyzing the classical output of the hybrid quantum-classical algorithm. Our hybrid differential privacy framework utilizes quantum or classical noise mechanisms to create differentially private quantum measurements, aiming to safeguard against privacy breaches by the adversary.

Representative reference	DP frameworks	Adversary access capabilities	Defense strategy
Dwork et al. [Dwo+06a]	classical	classical input and classical output	classical noise-adding mechanisms
Zhou and Ying [ZY17]	quantum	quantum input and quantum output	quantum noise-adding circuits
Our work	hybrid	quantum input and classical output	quantum or classical noise-adding measurement

Table 3.1: The comparison of various differential privacy (DP) frameworks.

To evaluate our hybrid differential privacy framework in comparison to quantum and classical differential privacy frameworks against potential threats, we outline the capabilities of adversaries and the corresponding defense strategies in Table 3.1. In the QDP framework, the adversary has access to quantum input and output states [ZY17] and only

quantum noise can be applied to protect the privacy. However, within our framework, we can utilize quantum or classical noise-adding quantum measurements to generate differentially private quantum measurements to ensure privacy. Our approach integrates defense strategies from both classical and quantum differential privacy frameworks, offering the advantage of providing both classical and quantum differential privacy. Further elaboration on this will be provided subsequently.

3.2 Hybrid Differential Privacy

In this section, we introduce the concept of *hybrid differential privacy* as a framework to safeguard the privacy of hybrid quantum-classical algorithms. This approach relies on quantum measurements to link classical and quantum differential privacy through the hybrid nature of these measurements.

To support this framework, we develop corresponding post-processing and composition theorems that facilitate the practical application of HDP in algorithm design. These theorems formalize two fundamental properties of differential privacy under the HDP framework:

- *Post-processing theorem*: This theorem ensures that once a quantum measurement is designed to satisfy differential privacy, any subsequent classical or quantum processing applied to the measurement results does not degrade its privacy guarantees. This allows differentially private outputs to be safely used in further data analysis or integrated into algorithmic pipelines without requiring additional privacy protection.

- *Composition theorem*: This theorem quantifies the cumulative privacy loss when multiple differentially private quantum measurements are performed on the same dataset. By providing explicit upper bounds on the accumulated privacy budget, it enables practical privacy budget management in the design of multi-stage or iterative hybrid quantum-classical algorithms.

Together, these theorems ensure that the privacy guarantees under the HDP framework

remain robust under post-processing and scalable under repeated computations, making HDP not only theoretically rigorous but also practically applicable for real-world hybrid quantum-classical systems.

A key component of our approach is the construction of *differentially private quantum measurements*, which serve as the primary mechanism for enforcing privacy in hybrid settings.

A quantum measurement maps an input quantum state to a probability distribution over a finite set of measurement outcomes. Accordingly, we model the measurement $\mathcal{M}$ as a randomized function on quantum states, denoted as $\mathcal{M}(\rho)$ and defined as follows.

Given a quantum measurement $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$ and a quantum state $\rho \in \mathcal{D}(\mathcal{H})$, the measurement induces a probability distribution over the outcome set $\mathcal{O}$ such that, for any subset $S \subset \mathcal{O}$,

$$\Pr(\mathcal{M}(\rho) \in S) = \sum_{i\in S} \mathrm{tr}(M_i\rho).$$

We now proceed to formally establish the definition and properties of hybrid differential privacy within this framework.

Definition 6 (Hybrid Differential Privacy)**.** *For constants $\epsilon \geq 0$ and $1 > \delta \geq 0$, a quantum measurement $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$ is considered to be* (ϵ, δ)-hybrid differentially private ((ϵ, δ)-HDP) *if it satisfies the condition that for any neighboring quantum states $\rho \sim \sigma$ and any subset $S \subseteq \mathcal{O}$, the inequality below holds:*

$$\Pr[\mathcal{M}(\rho) \in S] \leq \exp(\epsilon) \cdot \Pr[\mathcal{M}(\sigma) \in S] + \delta.$$

Here

$$\Pr[\mathcal{M}(\rho) \in S] = \sum_{i\in S} \Pr[\mathcal{M}(\rho) = i] = \sum_{i\in S} tr(M_i\rho).$$

When $\delta = 0$, we achieve the ϵ-HDP. Analogous to QDP and classical differential privacy, the (ϵ, δ)-HDP guarantees that the absolute value of privacy loss between any neighboring quantum states ρ and σ will be limited by ϵ with a probability of at least

$1 - \delta$.

We have not given a precise explanation for neighboring states in our definition. Any reasonable definition of neighboring quantum states can be utilized in this context and

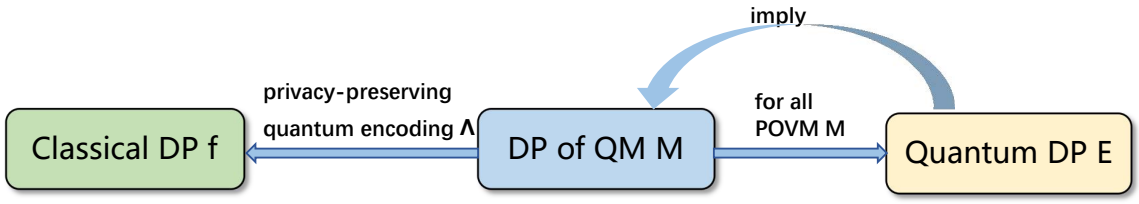

Figure 3.1: The relationship among various DP frameworks.

Our HDP definition combines the definitions of classical differential privacy and QDP from Definitions 2 and 4, respectively. Consequently, we can connect them with our HDP as visualized in Fig. 3.1. To this end, we first show how HDP provides classical differential privacy by a *neighboring-preserving quantum encoding method*. A more detailed comparison of these privacy notions is presented in Table 3.2.

Table 3.2: Comparison of classical DP, QDP and HDP frameworks.

Privacy Type	**Input**	**Mechanism**	**Output**
Classical DP	Classical Dataset	Randomized Function	Classical Distribution
QDP	Quantum State*	Quantum Circuit	Quantum State
HDP	Quantum State*	Quantum Measurement	Classical Distribution

* Quantum state can also be derived from classical data through quantum encoding.

As reviewed in Section 2.2.2, a neighboring-preserving quantum encoding ensures that classical neighboring datasets are mapped to quantum states that also satisfy a well-defined quantum notion of neighboring—such as being within a bounded trace distance. This allows the privacy guarantees applied to quantum measurement outputs to be meaningfully translated back to the classical domain. Therefore, by applying differential privacy mechanisms to the output of quantum measurements, we can effectively enforce classical

differential privacy on the original classical data through the encoding pipeline.

In addition to utilizing amplitude encoding for implementing a neighboring-preserving quantum encoding approach, basis encoding can also be employed for classical bit string datasets, maintaining local operation-based neighboring relationships introduced in Section 2.2.2. Basis encoding inherently converts classical bit strings to quantum bit strings; for instance, "0001" is represented as $|0001\rangle\langle 0001|$ through basis encoding. The local operation-based neighboring relationship is essentially an extension of its classical counterpart.

However, trying to create a neighboring-preserving quantum encoding through basis encoding and a trace-distance-based neighboring relationship would be unproductive and inconsequential. This is because, in basis encoding, setting $\eta = 1$ is required to achieve this, since any two quantum states that encode neighboring bit strings always have a trace distance of 1. Given that the trace distance ranges from 0 to 1, all quantum states are considered neighboring states, making the concept less significant. Therefore, it is essential to precisely define neighboring relationships of quantum states, especially when exploring various quantum encoding techniques, particularly in the context of utilizing quantum computing for addressing classical problems with differential privacy guarantees.

Now, by leveraging neighboring-preserving quantum encoding and the Heisenberg picture as outlined in Section 2.1, we can connect classical differential privacy and QDP through our HDP, as demonstrated in the following theorem and depicted in Fig. 3.1.

Theorem 1. *Let $\mathcal{E}$ represent a quantum circuit and let Λ be a neighboring-preserving quantum encoding. It follows that:*

- *$\mathcal{E}$ is (ϵ, δ)-QDP if and only if for any quantum measurement $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$, the transformed measurement $\mathcal{M}_\mathcal{E} = \{\mathcal{E}^\dagger(M_i)\}_{i\in\mathcal{O}}$ in the Heisenberg picture is (ϵ, δ)-HDP .*

- *If a quantum measurement $\mathcal{M}$ is (ϵ, δ)-HDP, then the classical randomized function $\mathcal{M} \circ \Lambda$ is (ϵ, δ)-classical differential privacy.*

Here, $\mathcal{M} \circ \Lambda$ denotes the functional composition of $\mathcal{M}$ and Λ.

Proof. The first claim of the theorem can be deduced from the definitions of hybrid differential privacy and quantum differential privacy.

Proving the second assertion of the theorem: It is easy to find that $\text{Range}(\mathcal{M} \circ \Lambda) = \mathcal{O}$. For any two neighboring classical input states $\vec{\nu} \sim \vec{\omega}$, since Λ is neighboring-preserving, $\Lambda(\vec{\nu}) \sim \Lambda(\vec{\omega})$. As $\mathcal{M}$ is (ϵ, δ)-HDP, we have

$$\text{tr}(M_S(\Lambda(\vec{\nu})) \leq e^{\epsilon}\text{tr}(M_S(\Lambda(\vec{\omega})) + \delta,$$

for all $S \subseteq \mathcal{O}$. Then we have

$$\begin{aligned} \Pr[\mathcal{M} \circ \Lambda(\vec{\nu}) \in S] &= \text{tr}(M_S(\Lambda(\vec{\nu})) \\ &\leq e^{\epsilon}\text{tr}(M_S(\Lambda(\vec{\omega})) + \delta \\ &= e^{\epsilon} \Pr[\mathcal{M} \circ \Lambda(\vec{\omega}) \in S] + \delta, \end{aligned}$$

for all $S \subseteq \mathcal{O}$. That means $\mathcal{M} \circ \Lambda$ is (ϵ, δ)-classical differential privacy. □

From the above theorem, our HDP framework can provide classical differential privacy by neighboring-preserving quantum encoding. Additionally, it can also provide QDP through quantum noise that makes any quantum measurement differentially private.

Recall that incorporating a randomized mechanism can protect the privacy of classical algorithms. However, while quantum measurement inherently involves randomness due to the probabilistic aspect of quantum computing, its primary purpose is not to ensure privacy protection. Hence, quantum measurement typically does not offer an effective guarantee of differential privacy. To illustrate this point more clearly, we offer the following example. This highlights the need for the development of differentially private quantum measurements.

Example 3.2.1 (Continuing Example 2.1.1)**.** *We investigate the HDP budgets offered by the quantum measurement $\mathcal{M}_{GHZ}$ as illustrated in Example 2.1.1, employing the lo-*

cal operation-based neighboring relationship for quantum states $\rho \sim \sigma$. Consider two neighboring quantum states $\rho = |000\rangle\langle 000|$ and $\sigma = |001\rangle\langle 001|$. The probability distributions of measuring the two states by $\mathcal{M}_{GHZ}$ over the outcome set $\{0, 1, \ldots, 7\}$ are $(\frac{1}{2}, 0, 0, 0, 0, 0, 0, \frac{1}{2})$ and $(0, \frac{1}{2}, 0, 0, 0, 0, \frac{1}{2}, 0)$.

As per the definition of HDP in Definition 6, the probabilities for each outcome should be relatively similar. When considering ϵ-HDP, it is observed that for the outcome 0, the measurement $\mathcal{M}_{GHZ}$ fails to meet the differential privacy requirement regardless of the chosen magnitude of ϵ. Moreover, in the context of (ϵ, δ)-DP, it necessitates δ to be 1 violating the constrait $\delta < 1$. Irrespective of the type of HDP considered, it is evident that this quantum measurement lacks differential privacy protection attributes.

As demonstrated above, it is necessary to develop differentially private quantum measurements to effectively safeguard hybrid quantum-classical algorithms. Given the mixed nature of quantum measurements, incorporating both quantum and classical noise is essential. Similar to classical and quantum differential privacy frameworks, it is crucial to establish post-processing and composition theorems initially to ensure the efficiency and scalability of differentially private quantum measurements, respectively, in subsequent sections.

3.2.1 Post-Processing Theorem

Post-processing is a crucial aspect of differential privacy as it prevents any additional computational analysis by an adversary from divulging more information about an individual's privacy. In parallel to the classical case, we also establish a post-processing theorem for HDP.

Theorem 2. *Let $\mathcal{M} = \{M_i\}_{i \in \mathcal{O}}$ be an (ϵ, δ)-HDP quantum measurement on a Hilbert space $\mathcal{H}$. For any randomized function $\mathcal{K} : \mathcal{O} \to \mathcal{O}'$, $\mathcal{K} \circ \mathcal{M} : \mathcal{D}(\mathcal{H}) \to \mathcal{O}'$ is (ϵ, δ)-HDP.*

This follows from the post-process theorem in classical differential privacy [DR+14, Proposition 2.1, Remark 3.1] by noting that $\mathcal{M}$ is a randomized function over a classical

set. Specifically, consider a randomized function $\mathcal{K} : \mathcal{O} \to \mathcal{O}'$. Observe that

$$D_\infty^\delta(\mathcal{K}(P)\|\mathcal{K}(Q)) \leq D_\infty^\delta(P\|Q).$$

It means that if $\mathcal{M}$ is (ϵ, δ)-differentially privacy, so is $\mathcal{K} \circ \mathcal{M}$. Here $D_\infty^\delta(P\|Q)$ stands for the δ-max divergence of random variables P and Q, the specific definition is as follows:

$$D_\infty^\delta(P\|Q) = \max_{S\subseteq\mathcal{O};\Pr[P\in S]\geq\delta} \ln \frac{\Pr[P \in S] - \delta}{\Pr[Q \in S]}.$$

This theorem demonstrates that classical post-processing does not affect overall privacy guarantees. This is critical, as it assures that adversarial manipulations do not compromise privacy. In the context of quantum differential privacy, distinguishing between the inherent noise introduced by the quantum machine and the noise added for privacy preservation is crucial. The inherent quantum noise arises from factors such as imperfections in quantum gates and decoherence, which are unavoidable in practical quantum computations. On the other hand, the noise added for privacy preservation is deliberately introduced to ensure differential privacy while performing quantum measurements or computations.

Our theorem addresses these challenges. Specifically, it guarantees that even when noise from the quantum machine and the noise we introduce for privacy preservation interact, the overall differential privacy property remains intact. This means that the added privacy noise does not interfere with the inherent quantum noise in a way that would compromise the accuracy of the computations. The theorem ensures that the differential privacy of the system is maintained, regardless of the interplay between these two types of noise.

HDP as a Bridge for Classical Differential Privacy in Quantum Settings. As we mentioned in the introduction, with the rapid development of quantum computing and the superior computational capabilities of quantum computers, an increasing number of classical problems are now being addressed using hybrid quantum-classical algorithms

denoted as $\mathcal{A} = (\mathcal{E}, \mathcal{M})$. The process involves converting a classical dataset $\vec{\nu}$ into a quantum input state $\rho_{\vec{\nu}}$ through quantum encoding, as depicted in Fig. 3.2 of Section 2.1. Subsequently, the quantum circuit $\mathcal{E}$ and quantum measurement $\mathcal{M}$ are applied, leading to the retrieval of classical data (measurement outcome set $\mathcal{O}$) based on a probability distribution determined by the measurement $\mathcal{M}$. This approach serves as a model for leveraging hybrid quantum-classical algorithms to solve classical problems. In the context of privacy considerations, a pertinent question arises: Does the post-process theorem still apply when employing quantum computing for classical problem-solving? Encouragingly, favorable outcomes can be observed with the use of neighboring-preserving quantum encoding techniques and noting that we can use measurement $\mathcal{M}_{\mathcal{E}}$ to describe the above evolution of hybrid quantum-classical algorithm $\mathcal{A} = (\mathcal{E}, \mathcal{M})$ in the Heisenberg picture in Section 2.1.

Corollary 1. *Suppose Λ represents a neighboring-preserving quantum encoding. If $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$ is a quantum measurement that is (ϵ, δ)-HDP, then $\mathcal{K} \circ \mathcal{M} \circ \Lambda$ is (ϵ, δ)-classical differential privacy for any randomized function $\mathcal{K} : \mathcal{O} \to \mathcal{O}'$.*

In the result above, $\mathcal{M} \circ \Lambda$ represents a classical randomized function with classical input, classical output, and quantum processing. This indicates that integrating a differentially private quantum measurement into a classical problem can also provide privacy protection against any subsequent post-processing analysis. Moreover, when combined with the second assertion in Theorem 1, our HDP can deliver efficient privacy safeguarding in classical differential privacy that can withstand post-process attacks. This improves the privacy protection when utilizing quantum computers for solving classical problems.

3.2.2 Composition Theorem

Having established the post-processing guarantees of our HDP framework, we now turn to another essential property of differential privacy—**composition**. In practical applications, hybrid quantum-classical algorithms are often invoked multiple times or involve multiple stages of computation. Therefore, it is crucial to understand how the overall privacy

guarantee degrades under such sequential uses. The composition theorem provides formal bounds on cumulative privacy loss, ensuring that privacy remains controlled even after repeated application of differentially private mechanisms.

In classical computing, it is crucial to explore how two algorithms can be combined, as this approach naturally leads to the creation of more advanced algorithms. Researchers have introduced composition theorems for classical and quantum differential privacy to safeguard the privacy of such combined algorithms [DR+14; Gua+23]. Similarly, in our HDP framework, we aim to investigate whether the fusion of two differentially private quantum measurements maintains its differentially private nature. To accomplish this goal, we must initially introduce how two quantum measurements can be combined and subsequently formulate the relevant composition theorem by extending the relationships between neighboring states from a single system to those within composed systems.

Given two quantum measurements, $\mathcal{M}_1 = \{M_i\}_{i\in\mathcal{O}_1}$ and $\mathcal{M}_2 = \{M_j\}_{j\in\mathcal{O}_2}$, we define their joint measurement over the composite system as

$$\mathcal{M}_{1,2} = \{M_{i,j} = M_i \otimes M_j \mid i \in \mathcal{O}_1,\, j \in \mathcal{O}_2\}.$$

This joint measurement $\mathcal{M}_{1,2}$ can be viewed as a randomized function mapping quantum states from the tensor product space $\mathcal{D}(\mathcal{H}_1)\otimes\mathcal{D}(\mathcal{H}_2)$ to a probability distribution over the product outcome space $\mathcal{O}_1 \times \mathcal{O}_2$.

To ensure that this definition is well-formed, we verify that $\mathcal{M}_{1,2}$ constitutes a valid quantum measurement:

Lemma 1. *$\mathcal{M}_{1,2}$ as defined above is a POVM.*

Proof. For all $i \in \mathcal{O}_1$ and $j \in \mathcal{O}_2$, define $M_{i,j} = M_i \otimes M_j$. Since M_i and M_j are POVM elements, they are positive semidefinite operators. Hence, there exist operators A_i and B_j such that

$$M_i = A_i^\dagger A_i, \qquad M_j = B_j^\dagger B_j.$$

It follows that

$$M_{i,j} = M_i \otimes M_j = (A_i^\dagger A_i) \otimes (B_j^\dagger B_j) = (A_i \otimes B_j)^\dagger (A_i \otimes B_j),$$

so each $M_{i,j}$ is also positive semidefinite.

Moreover, the completeness condition of a POVM holds:

$$\begin{aligned}\sum_{i,j} M_{i,j} = \sum_{i,j} M_i \otimes M_j &= \sum_i \left(M_i \otimes \sum_j M_j \right) \\ &= \sum_i (M_i \otimes I) = \left(\sum_i M_i \right) \otimes I = I \otimes I = I,\end{aligned}$$

where we used $\sum_j M_j = I$ and $\sum_i M_i = I$. Therefore, $\mathcal{M}_{1,2}$ satisfies the properties of a POVM. □

In the following paragraph, we define the notion of neighboring relationships in composite (product) quantum state spaces, which is essential for establishing composition theorems in our HDP framework.

In the composed state space $\mathcal{D}(\mathcal{H}_1) \otimes \mathcal{D}(\mathcal{H}_2)$, quantum states are in the form of tensor products, such as $\rho_{1,2} = \rho_1 \otimes \rho_2$, where $\rho_i \in \mathcal{D}(\mathcal{H}_i)$ for $i \in \{0, 1\}$. This type of state $\rho_{1,2}$ is referred to as a *product state*. Two product states, $\rho_{1,2} = \rho_1 \otimes \rho_2$ and $\sigma_{1,2} = \sigma_1 \otimes \sigma_2$, are considered neighboring, denoted as $\rho_{1,2} \sim \sigma_{1,2}$, when the respective subsystem states are also neighboring, meaning $\rho_i \sim \sigma_i$ for $i \in \{1, 2\}$ [ZY17]. This notion of neighboring relationships can be straightforwardly generalized to any finite composite state space $\mathcal{D}(\mathcal{H}_1) \otimes \mathcal{D}(\mathcal{H}_2) \otimes \cdots \otimes \mathcal{D}(\mathcal{H}_n)$.

Now, we present a composition theorem for differentially private quantum measurements within our HDP framework.

Theorem 3. *If quantum measurements $\mathcal{M}_1 = \{M_i\}_{i \in \mathcal{O}_1}$ and $\mathcal{M}_2 = \{M_j\}_{j \in \mathcal{O}_2}$ are (ϵ_1, δ_1)-HDP and (ϵ_2, δ_2)-HDP on Hilbert spaces $\mathcal{H}_1$ and $\mathcal{H}_2$ respectively, then the combined measurement $\mathcal{M}_{1,2}$ is $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$-HDP.*

To prove the composition theorem, we need the following lemma.

Lemma 2. *Let μ_1, ν_1 be distributions over a finite set K_1, and let μ_2, ν_2 be distributions over a finite set K_2, such that:*
For any $S_1 \subseteq K_1$,

$$\mu_1(S_1) \leq e^{\epsilon_1} \cdot \nu_1(S_1) + \delta_1.$$

For any $S_2 \subseteq K_2$,

$$\mu_2(S_2) \leq e^{\epsilon_2} \cdot \nu_2(S_2) + \delta_2.$$

Here $\mu_i(S_i) = \sum_{s \in S_i} \mu_i(s)$ and $\nu_i(S_i) = \sum_{s \in S_i} \nu_i(s)$, for $i \in \{1, 2\}$.

Define $\mu_{1,2}$ and $\nu_{1,2}$ be the distributions over set $K_1 \times K_2$, when $s = (k_1, k_2) \in K_1 \times K_2$,

$$\mu_{1,2}(s) \triangleq \mu_1(k_1) \cdot \mu_2(k_2), \ \nu_{1,2}(s) \triangleq \nu_1(k_1) \cdot \nu_2(k_2).$$

Then, for any $S \subseteq K_1 \times K_2$,

$$\mu_{1,2}(S) \leq e^{\epsilon_1 + \epsilon_2} \nu_{1,2}(S) + \delta_1 + \delta_2.$$

Here $\mu_{1,2}(S) = \sum_{s \in S} \mu_{1,2}(s)$ and $\nu_{1,2}(S) = \sum_{s \in S} \nu_{1,2}(s)$.

Proof. First, we separate K_1 into two parts A_1, A_2, satisfying $A_1 \cup A_2 = K_1, A_1 \cap A_2 = \emptyset$, for any $a \in A_1$, $\mu_1(a) > e^{\epsilon_1} \nu_1(a)$, and for any $a \in A_2$, $\mu_1(a) \leq e^{\epsilon_1} \nu_1(a)$. We separate K_2 into two parts like this: $B_1 \cup B_2 = K_2, B_1 \cap B_2 = \emptyset$, for any $b \in B_1$, $\mu_2(b) > e^{\epsilon_2} \nu_2(b)$, and for any $b \in B_2$, $\mu_2(b) \leq e^{\epsilon_2} \nu_2(b)$. Then we separate S into three parts $C_1 = S \cap (A_1 \times K_2), C_2 = (S - C_1) \cap (K_1 \times B_1), C_3 = S - C_1 - C_2$.

Let $C_a = (a \times K_2) \cap S, a \in K_1, C_b = (K_1 \times b) \cap (S - C_1), b \in K_2$. For all $a \in A_1$, it is easy to prove that $\mu_{1,2}(C_a) - e^{\epsilon_1 + \epsilon_2} \nu_{1,2}(C_a) \leq \mu_1(a) \delta_2$. Furthermore, for all $b \in B_1$, it is easy to prove that $\mu_{1,2}(C_b) - e^{\epsilon_1 + \epsilon_2} \nu_{1,2}(C_b) \leq \mu_2(b) \delta_1$. For any $c \in C_3$, it is easy to prove

that $\mu_{1,2}(c) - e^{\epsilon_1+\epsilon_2}\nu_{1,2}(c) \leq 0$. Then we have

$$\begin{aligned}
& \mu_{1,2}(S) - e^{\epsilon_1+\epsilon_2}\nu_{1,2}(S) \\
= & \bigcup_{a\in A_1} \mu_{1,2}(C_a) + \bigcup_{b\in B_1} \mu_{1,2}(C_b) + \bigcup_{c\in C_3} \mu_{1,2}(c) \\
& -e^{\epsilon_1+\epsilon_2}(\bigcup_{a\in A_1} \nu_{1,2}(C_a) + \bigcup_{b\in B_1} \nu_{1,2}(C_b) + \bigcup_{c\in C_3} \nu_{1,2}(c)) \\
\leq & \bigcup_{a\in A_1} \mu_1(a)\delta_2 + \bigcup_{b\in B_2} \mu_2(b)\delta_1 \\
\leq & \delta_1 + \delta_2.
\end{aligned}$$

□

Now we can prove Theorem 3.

Proof. When $i \in \mathcal{O}_1, j \in \mathcal{O}_2$,

$$\begin{aligned}
\mathrm{tr}(M_{i,j} \cdot (\rho \otimes \sigma)) &= \mathrm{tr}((M_i \otimes M_j)(\rho \otimes \sigma)) \\
&= \mathrm{tr}(M_i\rho) \times \mathrm{tr}(M_j\sigma).
\end{aligned}$$

According to lemma 2, for any $S \subseteq \mathcal{O}_1 \times \mathcal{O}_2$, we have

$$\begin{aligned}
& \sum_{(i,j)\in S} \mathrm{tr}(M_{i,j} \cdot (\rho \otimes \sigma)) \\
\leq & e^{(\epsilon_1+\epsilon_2)} \cdot \sum_{(i,j)\in S} \mathrm{tr}(M_{i,j} \cdot (\rho' \otimes \sigma')) + \delta_1 + \delta_2.
\end{aligned}$$

So the combination $\mathcal{M}_{1,2}$ is $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$-differentially private.

□

Remark 2. *First, we compare the above composition theorem with the composition theorem presented in [Gua+23]. Due to the differing definitions of neighboring employed in the two approaches, the composition theorems are distinct and cannot be derived from each other.*

In our framework, when defining the combination of two quantum measurements, we do not impose any restrictions on the definition of the neighboring relation between the input quantum states in $\mathcal{D}(\mathcal{H}_1)$ and $\mathcal{D}(\mathcal{H}_2)$. In other words, our composition theorem holds regardless of the specific neighboring relation chosen for the state spaces $\mathcal{D}(\mathcal{H}_1)$ and $\mathcal{D}(\mathcal{H}_2)$.

On the other hand, in [Gua+23], neighboring between two input states is defined via the trace distance, where states in $\mathcal{D}(\mathcal{H}_1)$ are considered neighboring if their trace distance is less than η_1, and similarly in $\mathcal{D}(\mathcal{H}_2)$ if it is less than η_2. After composition, the requirement for neighboring between two quantum states in the combined space $\mathcal{D}(\mathcal{H}_1 \otimes \mathcal{H}_2)$ is that the trace distance should be less than $\eta_1\eta_2$.

Hence, the composition theorems under the two definitions are entirely different due to the differing notions of neighboring. Even if we adopt the trace distance to define neighboring in our setting, the definition of neighboring after composition would still be completely different from that in [Gua+23].

Second, compared to the composition theorem in QDP [HRF23], which states that the composition satisfies $(\epsilon_1 + \epsilon_2, \bar{\delta})$-differential privacy with $\bar{\delta} = \min\{\delta_1 + e^{\epsilon_1}\delta_2, e^{\epsilon_2}\delta_1 + \delta_2\}$, our composition theorem in the HDP framework achieves a tighter bound. Specifically, we show that the composition satisfies $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$-HDP. This result aligns the composition bound in the HDP framework with the well-established classical differential privacy composition bound, significantly simplifying the analysis and improving the privacy guarantee compared to QDP. Thus, our composition theorem highlights the advantage of the HDP framework in bridging classical and quantum differential privacy while offering a more efficient composition bound.

This result can be naturally generalized to the case of finite compositions.

Corollary 2. *For each $k \in [n] = \{1, 2, \ldots, n\}$, suppose $\mathcal{M}_k = \{M_{i_k}\}_{i_k \in \mathcal{O}_k}$ is an (ϵ_k, δ_k)-HDP quantum measurement on the Hilbert space $\mathcal{H}_k$. Then their joint measurement $\mathcal{M}_{[n]}$*

is $(\sum_{i=1}^{n} \epsilon_i, \sum_{i=1}^{n} \delta_i)$*-HDP, where*

$$\mathcal{M}_{[n]} = \left\{ M_{i_1,\ldots,i_n} = \bigotimes_{k=1}^{n} M_{i_k} : (i_1, \ldots, i_n) \in \mathcal{O}_1 \times \cdots \times \mathcal{O}_n \right\}.$$

Given the above result along with the post-processing theorem, our method for differentially private quantum measurements proves to be successful. As a result, we can now focus on developing such measurements to protect the privacy of hybrid quantum-classical algorithms. This shift is motivated by the limitations of conventional quantum measurements as illustrated in Example 3.2.1. The following section will detail the effective implementation approaches involving quantum and classical noise mechanisms for this objective.

3.3 Differentially Private Quantum Measurements

In this section, we present how to design a differentially private quantum measurement by incorporating classical or quantum noise to make it practical. Fig. 3.2 depicts the overall control flow of our HDP framework.

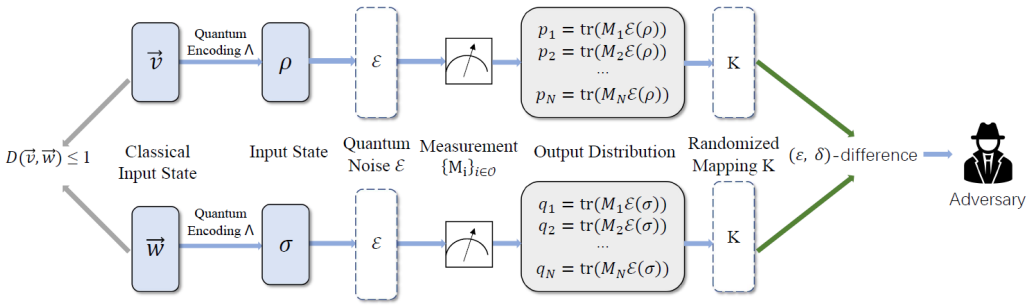

Figure 3.2: Our hybrid quantum-classical differential privacy framework

3.3.1 Classical Noise Strategy

In the realm of classical differential privacy, methods like Laplace noise and Gaussian noise have proven effective in rendering classical algorithms differentially private. These noise mechanisms are particularly useful in situations where the data domain is continuous and real-valued. However, in the context of quantum computing, quantum measurement outcomes are typically discrete and finite, making these traditional methods unsuitable. While some studies [ADK23; ADK22] have explored the use of Laplace noise and Gaussian noise to safeguard the privacy of quantum measurements, their applicability is limited due to this mismatch in suitability.

Therefore, we propose an alternative approach known as the *measurement-based exponential mechanism (MBEM)* for introducing classical noise that aligns well with quantum measurements. This method tackles the discrete nature of quantum measurement outcomes and strives to effectively balance privacy and utility. Our MBEM integrates the classical exponential mechanism [DR+14] into the quantum domain to enable the implementation of differentially private quantum measurements. The exponential mechanism is adept at handling categorical outputs akin to measurement results, ensuring differential privacy by selecting outputs based on utility scores determined by a utility function, with probabilities determined through an exponential function of these scores. The classical exponential mechanism has been briefly introduced in the chapter 2.

Measurement-based exponential mechanism. Building upon the classical exponential mechanism, we now introduce the measurement-based exponential mechanism tailored to quantum settings. Adapting this mechanism to quantum measurements involves utilizing the original probability distribution of measurement outcomes as utility scores, thereby redistributing outcomes based on the exponential function of these scores. This extension expands the utility of the exponential mechanism into the quantum domain, specifically within the realm of probabilistic outcomes.

To operationalize this idea and rigorously implement MBEM, we follow a four-step

process:

1. *Defining measurement-based utility function:* Let's start by revisiting that the outcomes of a quantum measurement $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$ create a probability distribution over a finite set $\mathcal{O}$. The probability of observing output i given input quantum state ρ is

$$P_{\mathcal{M}}(i|\rho) = \Pr[\mathcal{M}(\rho) = i] = \mathrm{tr}(M_i\rho).$$

   This distribution can be utilized as a utility function u where $u(\rho, i) = P_{\mathcal{M}}(i|\rho)$.

2. *Computing sensitivity:* Sensitivity Δu represents the maximum impact of the utility function u concerning changes in the neighboring quantum states and outcome set $\mathcal{O}$.

$$\Delta u = \max_{i\in\mathcal{O},\rho\sim\rho'} |u(\rho, i) - u(\rho', i)|. \tag{3.1}$$

   As u is a distribution, we can set a global maximum of $\Delta u = 1$ regardless of quantum state neighboring relationships. This is advantageous compared to the classical scenario, as it avoids the need to calculate the sensitivity Δu, which can be challenging in practice.

3. *Calculating probabilities:* Following the classical exponential mechanism, using an exponential function, we determine the probability of selecting each measurement outcome i based on the utility function $u(\rho, i)$ and the specified privacy parameter ϵ:

$$P_{EM}(i|\rho) = \frac{\exp\left(\frac{\epsilon u(\rho,i)}{2\Delta u}\right)}{\sum_{j\in\mathcal{O}} \exp\left(\frac{\epsilon u(\rho,j)}{2\Delta u}\right)} \tag{3.2}$$

4. *Selecting measurement outcome:* By using the calculated probability distribution $P_{EM}(i|\rho)$, a measurement outcome is randomly chosen from the set $\mathcal{O}$.

The overall procedure of MBEM is summarized in Algorithm 1.

Algorithm 1 Measurement-Based Exponential Mechanism (MBEM)

Input: Quantum state ρ, measurement $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$, privacy parameter $\epsilon > 0$
Output: Noisy measurement outcome $\tilde{i} \in \mathcal{O}$

▷ Evaluate utility values

for all $i \in \mathcal{O}$ **do**
    $u(\rho, i) \leftarrow \mathrm{tr}(M_i\rho)$
end for

▷ Set global sensitivity bound

$\Delta u \leftarrow 1$

▷ Compute exponential scores

for all $i \in \mathcal{O}$ **do**
    $s_i \leftarrow \exp\left(\frac{\epsilon\, u(\rho,i)}{2\Delta u}\right)$
end for
$Z \leftarrow \sum_{j\in\mathcal{O}} s_j$

▷ Normalize into a probability distribution

for all $i \in \mathcal{O}$ **do**
    $P_{\mathrm{EM}}(i \mid \rho) \leftarrow s_i/Z$
end for

▷ Sample noisy measurement outcome

Sample $\tilde{i} \in \mathcal{O}$ according to $P_{\mathrm{EM}}(\cdot \mid \rho)$
return $\tilde{i}$

This construction transforms any quantum measurement into a randomized function with tunable privacy guarantees. The following theorem formalizes this property.

Theorem 4. *The MBEM defined through the above procedure can enable any quantum measurement to be ϵ-HDP.*

Proof. For all $\rho \sim \sigma$ and $i \in \mathcal{O}$, we have

$$\begin{aligned}
\frac{P_{EM}(i|\rho)}{P_{EM}(i|\sigma)} &= \frac{\left(\frac{\exp\left(\frac{\epsilon u(\rho,i)}{2\Delta u}\right)}{\sum_{j\in\mathcal{O}}\exp\left(\frac{\epsilon u(\rho,j)}{2\Delta u}\right)}\right)}{\left(\frac{\exp\left(\frac{\epsilon u(\sigma,i)}{2\Delta u}\right)}{\sum_{j\in\mathcal{O}}\exp\left(\frac{\epsilon u(\sigma,j)}{2\Delta u}\right)}\right)} \\
&= \left(\frac{\exp\left(\frac{\epsilon u(\rho,i)}{2\Delta u}\right)}{\exp\left(\frac{\epsilon u(\sigma,i)}{2\Delta u}\right)}\right)\cdot\left(\frac{\sum_{j\in\mathcal{O}}\exp\left(\frac{\epsilon u(\sigma,j)}{2\Delta u}\right)}{\sum_{j\in\mathcal{O}}\exp\left(\frac{\epsilon u(\rho,j)}{2\Delta u}\right)}\right) \\
&= \exp\left(\frac{\epsilon(u(\rho,i)-u(\sigma,i))}{2\Delta u}\right)\cdot\left(\frac{\sum_{j\in\mathcal{O}}\exp\left(\frac{\epsilon u(\sigma,j)}{2\Delta u}\right)}{\sum_{j\in\mathcal{O}}\exp\left(\frac{\epsilon u(\rho,j)}{2\Delta u}\right)}\right) \\
&\leq \exp(\frac{\epsilon}{2})\exp(\frac{\epsilon}{2})\cdot\left(\frac{\sum_{j\in\mathcal{O}}\exp\left(\frac{\epsilon u(\rho,j)}{2\Delta u}\right)}{\sum_{j\in\mathcal{O}}\exp\left(\frac{\epsilon u(\rho,j)}{2\Delta u}\right)}\right) \\
&= \exp(\epsilon).
\end{aligned}$$

Similarly, by the symmetry, we obtain that

$$\frac{P_{EM}(i|\rho)}{P_{EM}(i|\sigma)} \geq \exp(-\epsilon).$$

□

To demonstrate the effectiveness of our MBEM, we utilize it to tackle the privacy issues of $\mathcal{M}_{\mathrm{GHZ}}$ (introduced in Example 2.1.1), as depicted in Example 3.2.1.

Example 3.3.1 (Continuing Example 2.1.1)**.** *We examine quantum measurement $\mathcal{M}_{GHZ}$ as detailed in Example 2.1.1. For quantum state $\rho_{\vec{v}} = |000\rangle\langle 000|$, $\mathcal{M}_{GHZ}$ results in outcomes $0, 1, ..., 7$ with corresponding probabilities of $(\frac{1}{2}, 0, 0, 0, 0, 0, 0, \frac{1}{2})$. In the subsequent steps, our objective is to employ the MBEM to compute a new outcome distribution and randomly choose the measurement result to ensure privacy preservation.*

We define the utility function for a general input state ρ and outcome i as

$$u(\rho, i) = \mathrm{tr}\big(M_{GHZ,i}\,\rho\big),$$

which captures the likelihood of observing outcome i under the effective measurement $\mathcal{M}_{GHZ}$. For the specific state $\rho_{\vec{\nu}} = |000\rangle\langle 000|$, this yields

$$u(\rho_{\vec{\nu}}, 0) = u(\rho_{\vec{\nu}}, 7) = \tfrac{1}{2}, \qquad u(\rho_{\vec{\nu}}, i) = 0 \text{ for } 1 \leq i \leq 6.$$

We set the sensitivity Δu to 1. Then, using Eq. 3.2, the distribution of selecting outcomes is

$$P_{EM}(0|\rho_{\vec{\nu}}) = P_{EM}(7|\rho_{\vec{\nu}}) = \frac{\exp(\frac{\epsilon}{4})}{2\exp(\frac{\epsilon}{4}) + 6},$$
$$P_{EM}(1|\rho_{\vec{\nu}}) = \cdots = P_{EM}(6|\rho_{\vec{\nu}}) = \frac{1}{2\exp(\frac{\epsilon}{4}) + 6}.$$

Here, we have illustrated how to add noise to quantum state $\rho_{\vec{\nu}}$. This same process can be used for other input states as well. In conclusion, the quantum measurement $\mathcal{M}_{GHZ}$ using our MBEM adheres to ϵ-HDP as proven in Theorem 4.

Our MBEM serves as a practical instance of a classical exponential mechanism. Using this approach, we can protect the differential privacy of measurement outcomes in hybrid quantum-classical algorithms. This method allows us to maintain high utility, ensuring both the accuracy and privacy of the algorithms. In our experiments presented in section 3.4.3, we compare this method with the Laplace and Gaussian mechanisms and find that, under the examples we provide, our MBEM approach achieves significantly higher utility while maintaining the desired privacy guarantees.

3.3.2 Quantum Noise Strategy

We turn to introduce quantum noise to enable differentially private quantum measurements. We opt for quantum depolarizing noise as the method for achieving this, given its ability to operate independently of the target quantum measurement and the quantum state neighboring relationships. By utilizing this noise, we can enhance the assurance of achieving better HDP.

Let's delve into the concept of quantum noise. Quantum noise can be represented

mathematically using Kraus operators, in the same formalism as the quantum circuits described in Eq.2.2 of Section2.1. A quantum noise, denoted as $\mathcal{E}_N$, consists of a set of Kraus matrices $\{E_{N,k}\}_{k\in K}$ and operates as follows:

$$\mathcal{E}_N(\rho) = \sum_{k\in K} E_{N,k}\rho E_{N,k}^{\dagger} \text{ for all } \rho \in \mathcal{D}(\mathcal{H}).$$

For instance, a 1-qubit depolarizing noise, denoted as $\mathcal{E}_{\text{Dep}}$, can be represented using Pauli gate matrices $\{X, Y, Z\}$ presented in Section 2.1 and the identity matrix I, incorporating a parameter $0 \leq p \leq 1$ in its Kraus matrices:

$$\mathcal{E}_{\text{Dep}}(\rho) = \frac{4-3p}{4}\rho + \frac{p}{4}(X\rho X + Y\rho Y + Z\rho Z),$$

with Kraus matrices as $\{\frac{\sqrt{4-3p}}{2}I, \frac{\sqrt{p}}{2}X, \frac{\sqrt{p}}{2}Y, \frac{\sqrt{p}}{2}Z\}$.

Another useful representation of $\mathcal{E}_{\text{Dep}}$ in [NC01] is given by:

$$\mathcal{E}_{\text{Dep}}(\rho) = (1-p)\rho + p\frac{I}{2}.$$

This representation indicates that quantum state ρ remains unaffected by depolarizing noise $\mathcal{E}_{\text{Dep}}$ with a probability of $1-p$, while with a probability p, ρ transitions to the quantum state $\frac{I}{2}$. Thus p is a *noisy probability*. Quantum depolarizing noise was considered a quantum adaptation of classical randomized response, aimed at safeguarding the privacy of quantum algorithms within the framework of quantum local differential privacy [Gua24]. This approach ensures the protection of all quantum states, not just the neighboring quantum states.

In a more general scenario, for an n-qubit Hilbert space $\mathcal{H}$, the depolarizing noise $\mathcal{E}_{\text{Dep}}$ can be expressed as:

$$\mathcal{E}_{\text{Dep}}(\rho) = (1-p)\rho + \frac{pI}{\text{Dim}},$$

where $\text{Dim} = 2^n$ represents the dimension of the state Hilbert space, and p is the noisy probability.

We will now delve into incorporating quantum noise into quantum measurements to enhance HDP in any dimensional Hilbert space $\mathcal{H}$. To assess this enhancement quantitatively, we will utilize trace distance-based neighboring relationships between quantum states in the following discussion of this section. As a quick reminder, as discussed in chapter 2, two quantum states ρ and σ are considered η-neighboring if their trace distance is within η, that is, $\rho \sim \sigma$ if and only if $\tau(\rho, \sigma) \leq \eta$.

Building upon this concept, we first prove that introducing any quantum noise $\mathcal{E}_N$ will not degrade the differential privacy of any quantum measurement $\mathcal{M}$.

Theorem 5. *If a quantum measurement $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$ is (ϵ, δ)-HDP, then $\mathcal{M}_{\mathcal{E}_N} = \{\mathcal{E}_N^\dagger(M_i)\}_{i\in\mathcal{O}}$ is also (ϵ, δ)-HDP for any quantum noise $\mathcal{E}_N$.*

This theorem demonstrates that incorporating quantum noise does not compromise privacy protection, even if the quantum noise is not generated by a differential privacy mechanism (a designated noise), but rather arises from the intrinsic random noise in quantum devices or a combination thereof. This substantiates the efficacy of our HDP framework in the existing NISQ era where random noise is an inevitable factor of hybrid quantum-classical algorithms.

Theorem 5 tells us that adding noise does not degrade overall privacy guarantees, which naturally leads to the question: can we leverage noise addition to actively enhance privacy protection in quantum measurements? The affirmative answer to this lies in the utilization of depolarizing noise, as supported by insights from hybrid quantum-classical algorithms [Gua+23, Theorem 4.1]. For completeness, we restate this known characterization within the HDP framework for quantum measurements, which will be used in our subsequent analysis.

Proposition 1 (Known Result [Gua+23, Theorem 4.1])**.** *Let $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$ represent a quantum measurement. The following conditions hold:*

- *$\mathcal{M}$ is (ϵ, δ)-HDP if and only if*

$$\delta \geq \max_{S\subseteq\mathcal{O}} \delta_S,$$

where

$$\delta_S = \eta\lambda_{\max}(M_S) - (e^\epsilon + \eta - 1)\lambda_{\min}(M_S),$$

and $M_S = \sum_{k\in S} M_k$. *Here,* $\lambda_{\max}(M_S)$ *and* $\lambda_{\min}(M_S)$ *denote the maximum and minimum eigenvalues of* M_S*, respectively.*

- $\mathcal{M}$ *is* ϵ*-HDP if and only if* $\epsilon \geq \epsilon^*$*, where*

$$\epsilon^* = \ln[(\kappa^* - 1)\eta + 1] \quad \textit{and} \quad \kappa^* = \max_{S\subseteq\mathcal{O}} \kappa(M_S),$$

with $\kappa(M_S) = \frac{\lambda_{\max}(M_S)}{\lambda_{\min}(M_S)}$ *being the condition number of* M_S *(taken as* $+\infty$ *if* $\lambda_{\min}(M_S) = 0$*).*

This result provides a practical tool for evaluating (noisy) quantum measurements under the HDP framework, enabling the determination of optimal privacy budgets even when noise is present. We will utilize this characterization in the following sections to quantify privacy guarantees for quantum measurements under depolarizing noise.

Theorem 6. *Suppose we have a quantum measurement denoted by* $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$ *and a depolarizing noise denoted by* $\mathcal{E}_{Dep}$ *with a noisy probability* p*. Then the obtained noisy measurement* $\mathcal{M}_{\mathcal{E}_{Dep}}$ *is* $(\epsilon, \bar{\delta})$*-HDP for any* $\epsilon \geq 0$*, where*

$$\bar{\delta} = \max_{S\subseteq\mathcal{O}} \left[(1-p)\gamma_S - (e^\epsilon - 1)\frac{ptr(M_S)}{Dim}\right]$$

with $\gamma_S = \eta\lambda_{\max}(M_S) - (e^\epsilon + \eta - 1)\lambda_{\min}(M_S)$ *and* $M_S = \sum_{k\in S} M_k$*. If we focus on* ϵ*-HDP, then* $\mathcal{M}_{\mathcal{E}_{Dep}}$ *is* $\bar{\epsilon}$*-HDP, where* $\bar{\epsilon} = \ln[(\bar{\theta} - 1)\eta + 1]$ *and* $\bar{\theta} = \max_{S\subseteq\mathcal{O}} \theta_S$ *with*

$$\theta_S = \frac{Dim\cdot(1-p)\lambda_{\max}(M_S) + tr(M_S)p}{Dim\cdot(1-p)\lambda_{\min}(M_S) + tr(M_S)p}. \tag{3.3}$$

Proof. According to the definition of the depolarizing noise, we have that, for any matrix A,

$$\mathcal{E}^\dagger_{\mathrm{Dep}}(A) = (1-p)A + \mathrm{tr}(A)\frac{p}{\mathrm{Dim}}I.$$

So we can get that

$$\begin{aligned}\sum_{k\in S}\mathcal{E}^{\dagger}_{\mathrm{Dep}}(M_k) &= \mathcal{E}^{\dagger}_{\mathrm{Dep}}(\sum_{k\in S}M_k) = \mathcal{E}^{\dagger}_{Dep}(M_S)\\ &= (1-p)M_S + \mathrm{tr}(M_S)\frac{p}{\mathrm{Dim}}I.\end{aligned}$$

Susequently, we have

$$\begin{aligned}\lambda_{\max}(\mathcal{E}^{\dagger}_{\mathrm{Dep}}(M_S)) &= (1-p)\lambda_{\max}(M_S) + \mathrm{tr}(M_S)\tfrac{p}{\mathrm{Dim}},\\ \lambda_{\min}(\mathcal{E}^{\dagger}_{\mathrm{Dep}}(M_S)) &= (1-p)\lambda_{\min}(M_S) + \mathrm{tr}(M_S)\tfrac{p}{\mathrm{Dim}}.\end{aligned}$$

Then by using Proposition 1, we can get the conclusion.

□

Intuition. Depolarizing noise transforms any operator A into a convex combination of A and the identity operator, namely

$$\mathcal{E}^{\dagger}_{\mathrm{dep}}(A) = (1-p)A + \frac{p\,\mathrm{tr}(A)}{\mathrm{Dim}}I.$$

This operation shrinks the eigenvalue spread of A: the largest eigenvalue decreases toward the average value $\mathrm{tr}(A)/\mathrm{Dim}$, and the smallest eigenvalue increases toward it. As a result, the condition number of A becomes smaller after depolarizing noise, meaning that the measurement outcomes induced by neighboring states become less distinguishable.

Under the HDP framework, privacy guarantees are controlled by the extremal eigenvalues of $M_S = \sum_{k\in S} M_k$. Depolarizing noise makes these eigenvalues closer to each other, which directly reduces the HDP privacy cost. Theorem 6 formalizes this intuition by applying Proposition 1 to the depolarized operator $\mathcal{E}^{\dagger}_{\mathrm{dep}}(M_S)$, yielding explicit expressions for (ϵ,δ)-HDP and ϵ-HDP.

By the above result, we can adjust the noisy probability p to achieve the intended

privacy budget $\bar{\epsilon}$ and $\bar{\delta}$ in the differentially private noisy measurement $\mathcal{M}_{\mathcal{E}_{\text{Dep}}}$.

The result of Theorem 6 relies on the quantum measurement $\mathcal{M}$. If the details of the measurement are unknown or if we aim to ensure that all quantum measurements are differentially private, we can still employ depolarizing noise to accomplish this with a worse (higher) ϵ-HDP.

Corollary 3. *Let $\mathcal{E}_{Dep}$ represent a depolarizing noise characterized by a noisy probability p. For any quantum measurement $\mathcal{M}$, the obtained noisy measurement $\mathcal{M}_{\mathcal{E}_{Dep}}$ is ϵ-HDP, where $\epsilon = \ln\left(\frac{Dim\cdot(1-p)}{p}\eta + 1\right)$.*

Proof. For any quantum measurement $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$ and for any $S \subseteq \mathcal{O}$, we have $0 \leq \lambda_{\min}(M_s) \leq \lambda_{\max}(M_S) \leq \text{tr}(M_S)$. Then according to Theorem 6

$$\begin{aligned}
\theta_S &= \frac{\text{Dim}\cdot(1-p)\lambda_{\max}(M_S) + \text{tr}(M_S)p}{\text{Dim}\cdot(1-p)\lambda_{\min}(M_S) + \text{tr}(M_S)p} \\
&\leq \frac{\text{Dim}\cdot(1-p)\text{tr}(M_S) + \text{tr}(M_S)p}{\text{tr}(M_S)p} \\
&= \frac{\text{Dim}\cdot(1-p)+p}{p}.
\end{aligned}$$

So $\bar{\theta} = \max_{S\subseteq\mathcal{O}}\theta_S \leq \frac{\text{Dim}\cdot(1-p)+p}{p}$ and

$$\begin{aligned}
\bar{\epsilon} &\leq \ln[(\frac{\text{Dim}\cdot(1-p)+p}{p} - 1)\eta + 1] \\
&= \ln\left(\frac{\text{Dim}\cdot(1-p)}{p}\eta + 1\right).
\end{aligned}$$

Which means $\mathcal{M}_{\mathcal{E}_{\text{Dep}}}$ is ϵ-HDP, where

$$\epsilon = \ln\left(\frac{\text{Dim}\cdot(1-p)}{p}\eta + 1\right).$$

□

Here, we present an upper limit for HDP budgets in Theorem 6 that is applicable to all measurements. Notably, this finding aligns with the outcome achieved in the QDP framework with the inclusion of depolarizing noise on quantum circuits [ZY17]. Consequently,

our HDP framework can offer QDP protection through the utilization of depolarizing noise mechanisms.

Moreover, if we aim to ensure privacy protection within different neighboring relationships beyond the trace distance-based metric discussed earlier, such as the local operation-based framework, we can still achieve this through the utilization of depolarizing noise mechanisms. This is feasible because the maximum trace distance between quantum states is 1, implying that if we consider 1-neighboring relationships based on the trace distance, then all quantum states are considered neighbors. By setting $\eta = 1$ in Theorem 6, we can derive the HDP budget facilitated by the depolarizing noise. Therefore, quantum depolarizing noise $\mathcal{E}_{\text{Dep}}$ can enable any quantum measurement to offer $\ln\left(\frac{\text{Dim}\cdot(1-p)}{p}+1\right)$-hybrid differential privacy, irrespective of the proximity relationships among quantum states.

Let's use an example to demonstrate how incorporating depolarizing noise can change the measurement $\mathcal{M}_{\text{GHZ}}$ in Example 2.1.1 from lacking HDP to ensuring HDP.

Example 3.3.2 (Continuing Example 2.1.1)**.** *We employ the quantum measurement $\mathcal{M}_{GHZ}$ as outlined in Example 2.1.1, maintaining the same local operation-based neighboring relationship and input states $\rho = |000\rangle\langle 000|$ and $\sigma = |001\rangle\langle 001|$ from Example 3.2.1 on a 3-qubit system ($Dim = 8$).*

To find the optimal $\bar{\epsilon}$ in HDP given by noisy measurement $\mathcal{M}_{\mathcal{E}_{GHZ}}$ with a noisy probability $p = 1/3$ and $\eta = 1$, according to Theorem 6, we first determine the maximum and minimum eigenvalues of $M_S = \sum_{k\in S} M_{GHZ,k}$, for all $S \subseteq \{0, ..., 7\}$. Let $|S|$ indicate the number of elements in the set.

For $|S| = 1$ and $i \in \{0, \ldots, 7\}$, we have $\lambda_{\max}(M_{GHZ,i}) = \frac{1}{2}$, $\lambda_{\min}(M_{GHZ,i}) = 0$, $tr(M_{GHZ,i}) = 1$, and $\theta_S = 9$ by Eq. 3.3.

For $|S| \geq 2$ and any $S \subseteq \{0, ..., 7\}$, we have $tr(M_S) = |S|$. Combining the fact $0 \leq \lambda_{\min}(M_S) \leq \lambda_{\max}(M_S) \leq 1$, we obtain

$$\theta_S \leq \frac{8 \cdot (1 - \frac{1}{3}) \cdot 1 + \frac{1}{3} \cdot |S|}{\frac{1}{3} \cdot |S|} \leq 9.$$

Therefore, $\max_{S\subseteq\{0,\ldots,7\}} \theta_S = 9$, *and* $\bar{\epsilon} = \ln 9$.

3.4 Evaluation of HDP

In above, we have introduced an HDP framework that utilizes both quantum and classical methods to add noise, ensuring privacy. Theorems 4 and 6 demonstrate that incorporating either an MBEM or a depolarizing noise mechanism can improve the level of differential privacy protection for quantum measurements. We utilize a running example, starting from Example 2.1.1 to 3.3.2, to demonstrate the application of our framework. To further assess the efficacy and implications of these methods within our framework, we perform a series of numerical experiments. Moreover, we carry out experiments to emphasize the benefits of our framework compared to the existing QDP framework in achieving the same privacy level through noise addition when specifying a quantum measurement. Additionally, we empirically examine the privacy-utility trade-off of our MBEM and depolarizing noise mechanisms across various aspects, including their specific parameters and comparison both internally and with other mechanisms.

Utility Loss: To analyze the trade-off between utility and privacy, we employ Kullback-Leibler (KL) divergence to evaluate the effectiveness of differentially private mechanisms. KL divergence measures the discrepancy between a model's probability distribution and the actual distribution. This metric is commonly used to measure the distinguishability between the probability distributions of private and original data in the context of classical differential privacy [DJW13; GI24; KOV14; AZ24]. In our evaluation of noise impact on utility within the HDP framework, we calculate the maximum KL divergence as the utility loss between the outcome distributions of measurements before and after noise introduction across all quantum states. Let $\mathcal{M}$ represent the initial quantum measurement and $\mathcal{M}_N$ denote the measurement after applying a differentially private noise mechanism N (e.g., MBEM and depolarizing noise). The utility loss (UL) introduced by N on a quantum measurement $\mathcal{M}$ is defined as

$$\mathrm{UL}(N) = \max_{\rho \in \mathcal{D}(\mathcal{H})} D_{\mathrm{KL}}(\mathcal{M}(\rho) \| \mathcal{M}_N(\rho))$$

In this context, the maximum is computed in all quantum states subjected to quantum measurements. A lower UL(N) signifies minimal noise impact, indicating preserved utility, while a higher value suggests a more significant utility loss. This approach offers a quantitative assessment of noise effects on utility.

Platform: Our experiments were carried out on a MacBook Pro equipped with an Apple M1 Pro chip, 10-core CPU, 16 GB of unified memory, and integrated GPU, operating on macOS Monterey 12.0.1.

3.4.1 Hybrid Differential Privacy Mechanisms

We evaluate the efficacy of the MBEM and the depolarizing noise mechanism in achieving ϵ-HDP across various ϵ values.

Measurement-based exponential mechanism. To assess the suitability of the MBEM, we investigate the correlation between utility loss and the sensitivity parameter Δu, displayed in Fig. 3.3. The study depicts four curves corresponding to different privacy budget parameters ϵ, showcasing the variation in utility loss across various Δu values. Since Δu gauges the sensitivity of probability distributions as defined in Eq. 3.1, it is constrained to a maximum value of 1. A more precise upper limit can be computed using the same equation, yielding $\Delta u = 1/2$. Hence, in Fig. 3.3, the range of Δu spans from 1/2 to 1.

The results shown in Fig. 3.3 indicate that using a smaller value of Δu (i.e., a tighter upper bound on sensitivity) reduces utility loss. This demonstrates that carefully selecting a tighter Δu not only improves data utility but also maintains the required privacy guarantees with minimal noise. These findings confirm the effectiveness of MBEM in balancing privacy and utility.

Furthermore, detailed experimental data and analyses on Δu and its relationship with the privacy level ϵ are provided in the Section 3.4.4. These analyses further reveal the trade-off between Δu and utility loss, and demonstrate how optimizing Δu can achieve more accurate privacy protection.

Quantum depolarizing noise mechanism. To verify the possibility of enhancing the

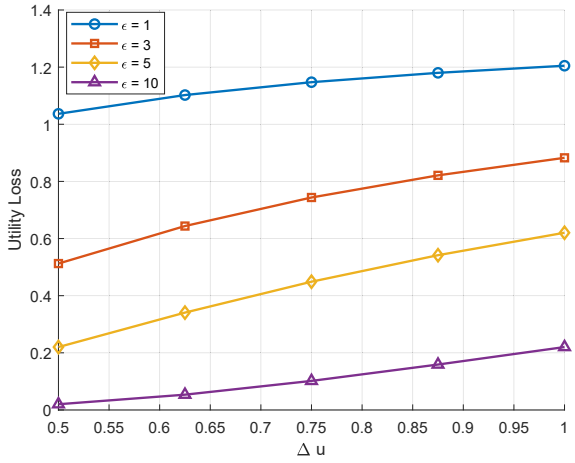

Figure 3.3: Trade-off between utility loss and MBEM sensitivity Δu for different ϵ-HDP.

probability p of quantum depolarizing noise to achieve the desired level of HDP, we deploy the relationship between p and the differential privacy parameter ϵ as per the equation stated in Corollary 3 and get Fig. 3.4. Our analysis is conducted on a 3-qubit system (Dim $= 8$) focusing on the trace distance-based η-neighboring relationships for varying η values. The depicted graph illustrates that with an increase in the added noise (value of p), the level of HDP protection improves consistently regardless of the η value. This observation suggests that augmenting the noisy probability p contributes to enhancing HDP.

Furthermore, we investigate the relationship between the neighboring parameter η and the privacy budget ϵ, as illustrated in Fig. 3.5, where we evaluate different depolarizing noise levels $p \in \{0.2, 0.3, 0.4\}$. The figure shows that as η increases—indicating a broader set of considered neighboring quantum states—it becomes more difficult to ensure privacy protection against all such neighbors, leading to a higher value of ϵ. This result highlights the inherent trade-off between the granularity of η-neighboring relationships and the achievable HDP privacy budget ϵ.

3.4.2 Comparison to Quantum Differential Privacy

The primary difference between HDP and QDP lies in their focus: the former guarantees privacy for a specific quantum measurement, whereas the latter ensures privacy across all

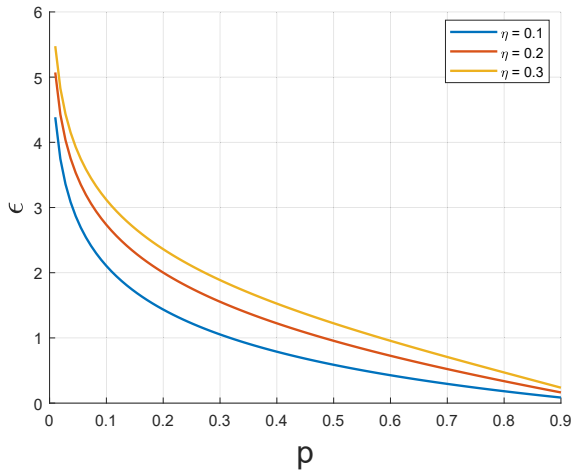

Figure 3.4: Trade-off between ϵ-HDP and noisy probability p of depolarizing noise for different η-neighboring relationships.

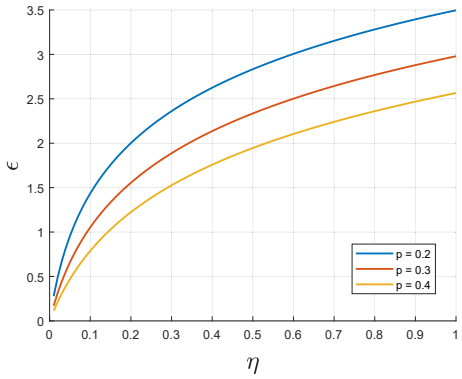

Figure 3.5: Trade-off between ϵ-HDP and η-neighboring relationship under depolarizing noise with noisy probability p.

quantum measurements, as depicted in Fig. 3.1. Moreover, when utilizing quantum depolarizing noise, the QDP framework proves to offer inferior privacy protection compared to our hybrid framework. Fig. 3.6 demonstrates this contrast within the context of Example 3.3.2. The disparity between the estimated differential privacy value from QDP and the actual value from our HDP is evident in the figure. Essentially, the HDP framework proposed in this research provides significant advantages in privacy protection for known quantum measurements of hybrid quantum-classical algorithms.

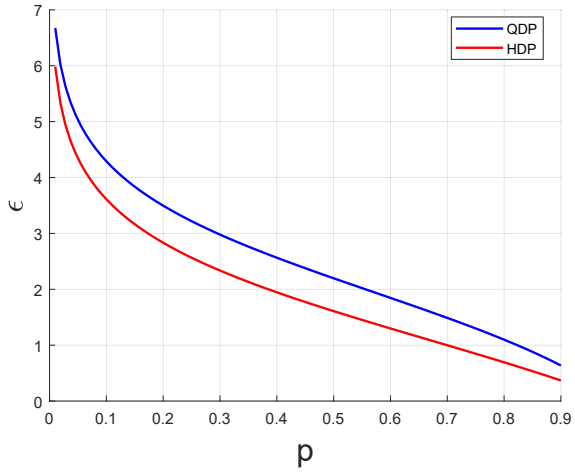

Figure 3.6: Comparison of ϵ-QDP and ϵ-HDP in terms of depolarizing noise probability p.

3.4.3 Privacy-Utility Trade-off Analysis

In this study, we conducted experiments using three types of quantum circuits: a **variational quantum circuit (VQC)**, a **GHZ circuit**, and a **randomly generated circuit**, to analyze and compare their performance.

- **VQC**: The variational quantum circuit is widely used in quantum machine learning as it serves as a quantum analog of classical neural networks, enabling parameter optimization for specific tasks. Our VQC consists of 3 qubits and 3 layers. Each layer includes parameterized $R_y(\theta)$ rotation gates with randomly initialized parameters applied to each qubit, followed by CNOT gates that entangle neighboring qubits.

- **GHZ Circuit**: The GHZ circuit is used to prepare highly entangled quantum states, as described in earlier sections.

- **Random Circuit**: The random circuit consists of 3 qubits with a depth randomly chosen between 3 and 5 layers. Single-qubit gates (e.g., H, X, Z, $R_x(\theta)$, $R_z(\theta)$) and CX gates are randomly applied, where the rotation angles θ are uniformly sampled from $[0, 2M]$. This circuit serves as a baseline to evaluate the performance of structured circuits like VQC and GHZ.

For each circuit, the input states were initialized to classical binary strings, and the

output states were simulated using the `Statevector` method to compute exact probability distributions. These circuits were evaluated under specific noise models, and the relationship between privacy parameters (ϵ) and utility loss (measured via KL divergence) was analyzed for comparison.

The circuits include the GHZ circuit described earlier and two additional randomly generated quantum circuits. For each circuit, the theoretical output distributions are modified separately using MBEM and the depolarization mechanism. The KL divergence between the original and modified distributions is computed as a measure of utility loss, while the privacy guarantee is determined by varying the privacy parameter ϵ. The figure illustrates the relationship between privacy (ϵ) and utility loss (KL divergence) across the three circuits for both noise addition methods, highlighting the effect of classical and quantum noise on the balance between privacy and utility.

This section provides supplementary experiments to further evaluate the effectiveness and utility of hybrid differential privacy mechanisms. Specifically, we present detailed analyses and experimental results for two key mechanisms: the Measurement-Based Exponential Mechanism and the Depolarizing Noise mechanism. The goal is to explore how these mechanisms perform under different configurations, with a focus on balancing privacy and utility.

Internal Comparison: We initially assess the balance between privacy and utility in our MBEM utilizing classical noise and the depolarizing noise mechanism employing quantum noise. The trade-offs between privacy and utility for both approaches on three 3-qubit quantum circuits are illustrated in Fig. 3.7. The graph clearly illustrates that both MBEM and the depolarizing noise mechanism result in a reduction in utility as the level of privacy protection increases. However, the extent of this reduction varies among the different circuits. Additionally, we notice that the depolarizing noise mechanism, incorporating quantum noise, exhibits better utility compared to the MBEM method using classical noise, at the same level of privacy protection. Nonetheless, the correlation between utility and privacy for these methods does not exhibit a significant distinction.

This suggests that both methods are valid choices for privacy protection, and the selection between them should be based on the specific requirements of the practical issue at hand.

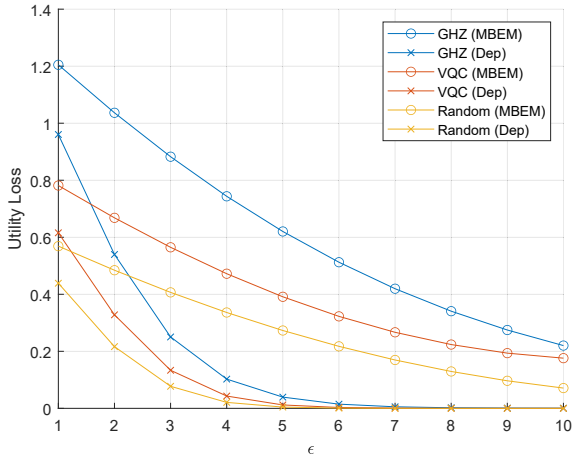

Figure 3.7: Privacy-utility trade-off for MBEM and depolarizing noise mechanisms (Dep) on three 3-qubit quantum circuits.

MBEM v.s. Other Classical Noise: We evaluate the effectiveness of our MBEM approach in comparison to the commonly employed Laplace mechanism and the Gaussian noise mechanism within the HDP framework. The specific implementation details can be found in Section 2.4.1 By plotting privacy-utility curves, illustrated in Fig. 3.8, we examine the trade-offs between privacy and utility. The results indicate that the MBEM method offers notably improved utility without compromising the level of privacy provided. Additionally, it is apparent that the utility of MBEM increases as the level of privacy protection lessens.

Qubit Number: At last, the privacy-utility trade-offs for MBEM and depolarizing noise mechanisms on quantum random circuits and variational quantum circuits (VQC) with varied qubit numbers are conducted. The result of MBEM on VQC is presented in Fig. 3.9 , and other similar results are presented in Section 3.4.5. In these experiments, VQC and RC allow for a range of 3 to 12 qubits. Although our approach is applicable to any qubit number, the computation of utility loss becomes more resource-intensive with increasing qubit numbers due to the exponential growth in the dimension of the measurement outcome distribution. To streamline operations, we present the privacy-

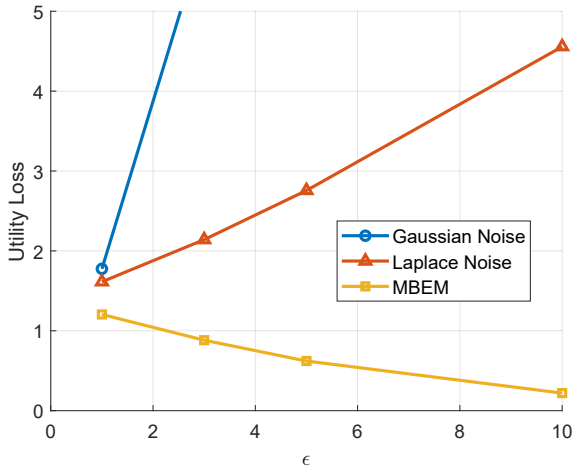

Figure 3.8: Comparison of the privacy-utility trade-off of different classical noise mechanisms on a GHZ circuit.

utility trade-off analysis within the confines of 12 qubits. The results highlight that our methods can be applied to any quantum circuits with different qubit number.

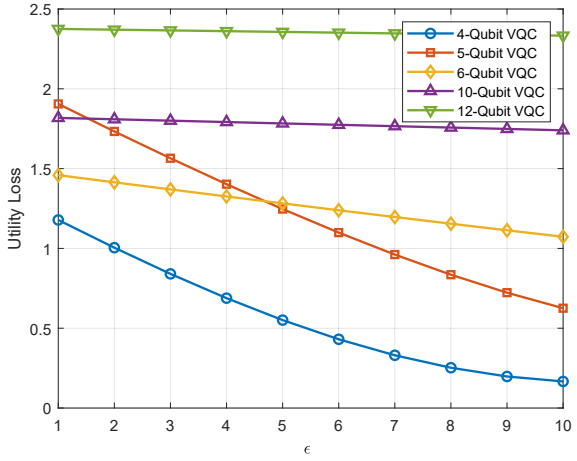

Figure 3.9: Privacy-utility curves for MBEM on variational quantum circuits (VQC) with varied qubit numbers.

3.4.4 Detailed Analysis of MBEM

In this subsection, we extend the analysis of the MBEM introduced in Example 3.3.1. We explore the application of the MBEM to safeguard the ϵ-HDP of a quantum state $\rho_{\vec{v}} = |000\rangle\langle 000|$ against any other state when measured by the quantum measurement

$\mathcal{M}_{\text{GHZ}}$. Table 3.3 illustrates the resulting probability distributions with varying levels of noise error introduced by the MBEM to achieve different levels of privacy as defined by ϵ in the HDP framework. The second row in the table represents the distribution of measurement outcomes without the mechanism, which does not ensure privacy protection ($\epsilon = \infty$). The first column in the table denotes the expected level of differential privacy provided by the MBEM in Theorem 4, while columns two through nine display the adjusted measurement outcome distributions due to the mechanism. The final column indicates the more actual level, denoted by ϵ_D, of differential privacy achieved by the mechanism according to the HDP definition in Definition 2.4, as a baseline.

The table illustrates that as the value of ϵ decreases, the MBEM introduces more noise, thereby enhancing privacy protection. This highlights the mechanism's effectiveness in ensuring privacy. However, the level of privacy guaranteed by ϵ deviates significantly from the true privacy level ϵ_D. *To address this discrepancy and obtain a more accurate estimation of the privacy protection offered by the MBEM, selecting a more suitable value for Δu, which represents the sensitivity of the mechanism, is crucial.* In Example 3.3.1 and Table 3.3, we assume $\Delta u = 1$, but a more precise value can be calculated using the formula in Eq. 3.1, resulting in $\Delta u = 1/2$. By substituting this refined value into the example, we derive the outcomes presented in Table 3.4.

A comparison between Tables 3.3 and 3.4 reveals that to achieve the same expected differential privacy effect ϵ, setting Δu to $1/2$ reduces the added noise, as evidenced by the probability distribution post-noise addition closely resembling the original distribution. Furthermore, a more accurate Δu leads to ϵ approaching ϵ_D, indicating that a precise Δu can deliver the desired privacy protection with minimal interference from noise, thus balancing privacy and utility. This trade-off is demonstrated in Fig. 3.3.

ϵ	$P_{EM}(0\|\rho_{\vec{v}})$	$P_{EM}(1\|\rho_{\vec{v}})$	$P_{EM}(2\|\rho_{\vec{v}})$	$P_{EM}(3\|\rho_{\vec{v}})$	$P_{EM}(4\|\rho_{\vec{v}})$	$P_{EM}(5\|\rho_{\vec{v}})$	$P_{EM}(6\|\rho_{\vec{v}})$	$P_{EM}(7\|\rho_{\vec{v}})$	ϵ_D
∞	0.5	0	0	0	0	0	0	0.5	∞
1	≈ 0.1499	≈ 0.1167	≈ 0.1167	≈ 0.1167	≈ 0.1167	≈ 0.1167	≈ 0.1167	≈ 0.1499	0.25
3	≈ 0.2069	≈ 0.0977	≈ 0.0977	≈ 0.0977	≈ 0.0977	≈ 0.0977	≈ 0.0977	≈ 0.2069	0.75
5	≈ 0.2680	≈ 0.0770	≈ 0.0770	≈ 0.0770	≈ 0.0770	≈ 0.0770	≈ 0.0770	≈ 0.2680	1.25
10	≈ 0.4006	≈ 0.0329	≈ 0.0329	≈ 0.0329	≈ 0.0329	≈ 0.0329	≈ 0.0329	≈ 0.4006	2.50

Table 3.3: The distributions for different values of ϵ by the measurement-based exponential mechanism with $\Delta u = 1$.

ϵ	$P_{EM}(0\|\rho_{\vec{v}})$	$P_{EM}(1\|\rho_{\vec{v}})$	$P_{EM}(2\|\rho_{\vec{v}})$	$P_{EM}(3\|\rho_{\vec{v}})$	$P_{EM}(4\|\rho_{\vec{v}})$	$P_{EM}(5\|\rho_{\vec{v}})$	$P_{EM}(6\|\rho_{\vec{v}})$	$P_{EM}(7\|\rho_{\vec{v}})$	ϵ_D
∞	0.5	0	0	0	0	0	0	0.5	∞
1	≈ 0.1770	≈ 0.1076	≈ 0.1076	≈ 0.1076	≈ 0.1076	≈ 0.1076	≈ 0.1076	≈ 0.1770	0.50
3	≈ 0.2990	≈ 0.0668	≈ 0.0668	≈ 0.0668	≈ 0.0668	≈ 0.0668	≈ 0.0668	≈ 0.2990	1.50
5	≈ 0.4010	≈ 0.0329	≈ 0.0329	≈ 0.0329	≈ 0.0329	≈ 0.0329	≈ 0.0329	≈ 0.4010	2.50
10	≈ 0.4901	≈ 0.0033	≈ 0.0033	≈ 0.0033	≈ 0.0033	≈ 0.0033	≈ 0.0033	≈ 0.4901	5.00

Table 3.4: The distributions for different values of ϵ by the measurement-based exponential mechanism with $\Delta u = 1/2$.

3.4.5 Extended Privacy-Utility Trade-off Evaluation on Multi-Qubit Circuits

While the preceding sections established the fundamental behavior of HDP mechanisms using small-scale quantum circuits, it remains essential to examine how these mechanisms perform when applied to larger and more practical quantum circuits. Small-scale experiments provide a clear and controlled environment for understanding the interaction between differential privacy mechanisms and quantum measurements; however, real-world applications in hybrid quantum-classical algorithms typically involve circuits with multiple qubits and deeper structures, where scalability and robustness under varying qubit configurations become critical.

To address this, we extend our evaluation to multi-qubit quantum circuits to assess the scalability of our HDP methods while ensuring that privacy guarantees are preserved without incurring excessive utility loss. Specifically, we present privacy-utility trade-offs for both the MBEM and depolarizing noise mechanisms applied to two types of quantum circuits: random circuits and variational quantum circuits (VQC). The experiments are conducted with varied qubit numbers to evaluate the performance of these mechanisms in multi-qubit scenarios. Our goal is to explore how the hybrid differential privacy (HDP)

framework behaves under classical and quantum noise in multi-qubit settings.

Experiment Setup and Circuit Configurations In these experiments, we utilize two types of quantum circuits: **random circuits** and **variational quantum circuits (VQC)**. Their configurations are described as follows:

- **Random Circuit**: The random circuit consists of 3 to 12 qubits, with a depth randomly chosen between 3 and 5 layers. Single-qubit gates (e.g., H, X, Z, $R_x(\theta)$, $R_z(\theta)$) and CX gates are randomly applied, where the rotation angles θ are uniformly sampled from $[0, 2M]$. This circuit serves as a baseline to evaluate the performance of structured circuits like VQC.

- **VQC**: The variational quantum circuit (VQC) is widely used in quantum machine learning as it serves as a quantum analog of classical neural networks, enabling parameter optimization for specific tasks. In our experiments, the VQC consists of 3 to 12 qubits and multiple layers proportional to the number of qubits. Each layer includes parameterized $R_y(\theta)$ rotation gates with randomly initialized parameters applied to each qubit, followed by CNOT gates that entangle neighboring qubits.

For both types of circuits, the input states were initialized as classical binary strings, and the output states were simulated using the `Statevector` method to compute exact probability distributions. Errors were introduced into the circuits using the MBEM and depolarizing noise mechanisms, resulting in noisy output distributions. The utility loss was quantified by computing the KL divergence between the original (noise-free) distributions and the noisy output distributions, while the privacy guarantee was assessed by varying the privacy parameter ϵ.

The results of these experiments are shown in Figures 3.10[1] to 3.13, which illustrate the privacy-utility trade-offs for the two mechanisms applied to random circuits and VQC, respectively. Specifically:

[1]Fig.3.10 is a copy of Fig. 3.9.

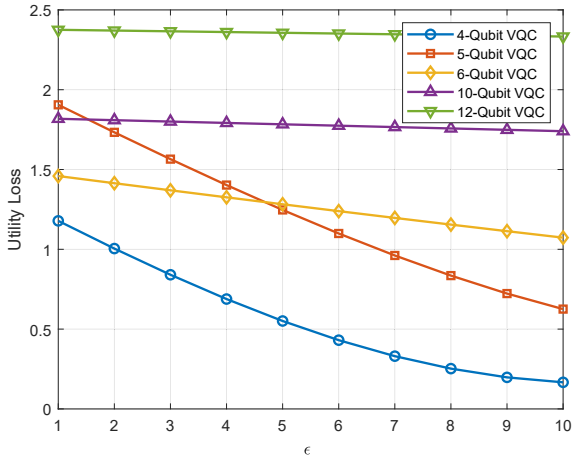

Figure 3.10: Privacy-utility curves for MBEM on variational quantum circuits (VQC) with varied qubit numbers.

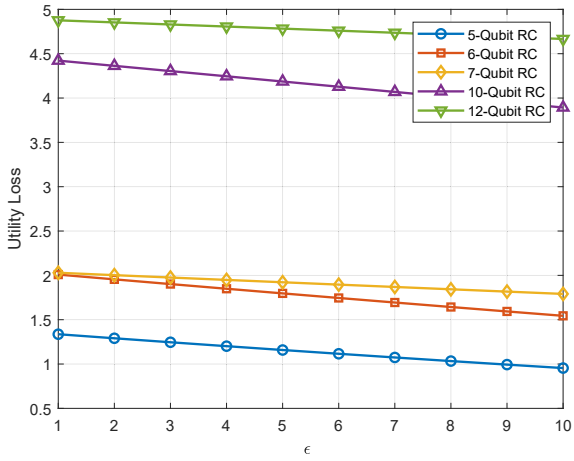

Figure 3.11: Privacy-utility curves for MBEM on quantum random circuits with varied qubit numbers

- **Figures 3.10 and 3.11** present the privacy-utility trade-offs for MBEM applied to random circuits and VQC, respectively. The results demonstrate that as privacy protection weakens (ϵ increases), utility loss decreases, indicating that reduced privacy leads to improved usability. Furthermore, the utility of the algorithm decreases as the number of qubits increases, with the degradation more evident in MBEM compared to the depolarizing noise mechanism.

- **Figures 3.12 and 3.13** display the corresponding privacy-utility trade-offs for the depolarizing noise mechanism. Similar to MBEM, weaker privacy protection im-

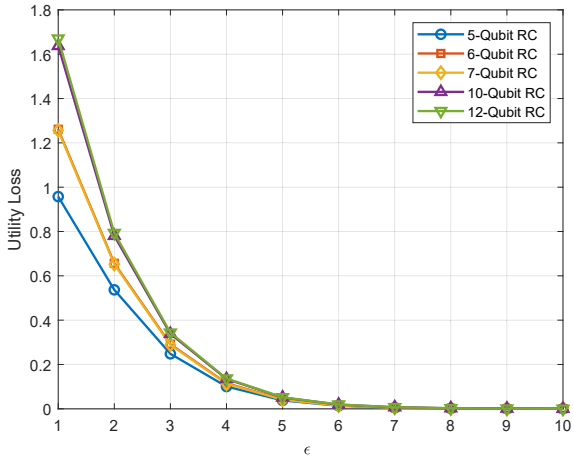

Figure 3.12: Privacy-utility curves for depolarizing noise mechanism on quantum random circuits with varied qubit numbers

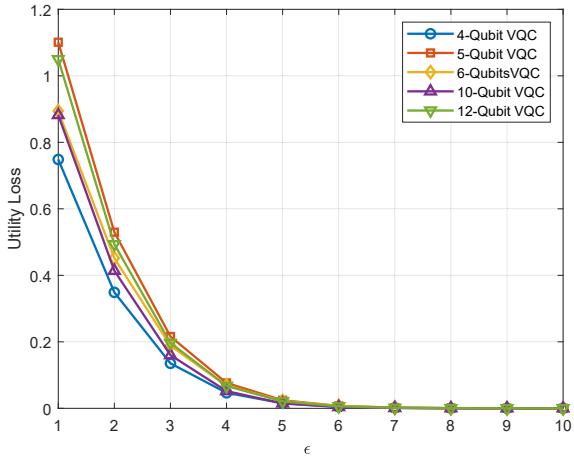

Figure 3.13: Privacy-utility curves for depolarizing noise mechanism on variational quantum circuits with varied qubit numbers

proves utility. However, the impact of qubit numbers on utility is less significant for the depolarizing mechanism compared to MBEM.

3.4.6 Summary and Implications

From these results, we observe that the depolarizing noise mechanism exhibits a privacy-utility trade-off that is less sensitive to the number of qubits, making it more predictable in multi-qubit scenarios. On the other hand, MBEM is more sensitive to qubit numbers, indicating that it requires detailed privacy-utility trade-off analyses tailored to specific

qubit configurations when applied in practice.

Circuit Type	Mechanism	10 Qubits	12 Qubits
Random Circuits (RC)	MBEM	$\approx$ 0s	1s
	Depolarizing noise	27s	7min 7s
Variational Quantum Circuits (VQC)	MBEM	$\approx$ 0s	3s
	Depolarizing noise	34s	10min 15s

Table 3.5: Execution times for privacy-utility trade-off analysis of MBEM and depolarizing noise mechanisms on RCs and VQCs.

Table 3.5 compares the execution times of the two mechanisms for different qubit numbers. MBEM demonstrates significantly faster execution, completing analyses within seconds, even for 12 qubits. In contrast, the depolarizing mechanism requires much longer computational times as the qubit number increases. For example:

- On random circuits with 10 qubits, MBEM completes in almost 0 seconds, whereas the depolarizing noise mechanism requires 27 seconds. For 12 qubits, MBEM takes 1 second, but the depolarizing noise mechanism requires 7 minutes and 7 seconds.
- On VQC circuits with 10 qubits, MBEM completes in almost 0 seconds, while the depolarizing noise mechanism requires 34 seconds. For 12 qubits, MBEM takes 3 seconds, whereas the depolarizing noise mechanism requires 10 minutes and 15 seconds.

In summary, these findings highlight two key takeaways:

1. The depolarizing mechanism's utility-privacy trade-off is less affected by the qubit number of quantum systems, making it a robust choice in multi-qubit scenarios.
2. MBEM, while more sensitive to qubit number, offers extremely fast execution times, allowing for efficient privacy-utility trade-off analyses without significantly impacting performance. This makes MBEM suitable for balancing privacy and utility in time-sensitive applications where computational efficiency is critical.

Chapter 4

Rényi Differential Privacy in Quantum Machine Learning

In classical machine learning, Rényi Differential Privacy (RDP) has become a widely used tool for analyzing privacy guarantees, particularly in iterative algorithms such as stochastic gradient descent, where its strong composition properties enable tighter bounds than standard (ϵ, δ)-differential privacy. Motivated by these successes, in this section we extend the use of RDP to quantum machine learning (QML) circuits. By focusing on fixed measurement settings that naturally arise in QML, we show how RDP can provide a more accurate and tractable characterization of privacy loss in the quantum setting.

4.1 Rényi Differential Privacy of Quantum Measurements

In quantum machine learning, quantum states often encode sensitive classical data. While the internal quantum computations may not be directly observable, QML algorithms produce classical outputs via fixed quantum measurements. Therefore, to analyze the differential privacy of QML algorithms, it is natural to focus on the privacy properties of the measurement process.

As previously discussed, the behavior of a QML algorithm can be characterized by a quantum channel $\mathcal{E}$ followed by a fixed quantum measurement $\mathcal{M} = \{M_i\}_{i \in \mathcal{O}}$, which

together map input quantum states to classical distributions over the outcome set $\mathcal{O}$. Under the Heisenberg picture, the combined effect of the channel and measurement can be represented as an equivalent measurement $\mathcal{M}_{\mathcal{E}} = \{\mathcal{E}^{\dagger}(M_i)\}_{i\in\mathcal{O}}$ directly applied to the input state.

This perspective allows us to study the privacy properties of the entire QML process by analyzing the distinguishability between the output distributions of this effective measurement applied to neighboring quantum states.

A quantum measurement maps a quantum state to a classical probability distribution over a set of possible outcomes. Specifically, let $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$ be a quantum measurement and $\rho \in \mathcal{D}(\mathcal{H})$ a quantum state. We define the measurement-induced distribution $\mathcal{M}(\rho)$ over the outcome set $\mathcal{O}$ as follows:

$$\Pr[\mathcal{M}(\rho) \in S] = \sum_{i\in S} \mathrm{tr}(M_i\rho), \quad \forall S \subseteq \mathcal{O}.$$

In this chapter, we apply Rényi differential privacy (RDP) to the classical distributions induced by quantum measurements. This formulation captures the observable privacy loss in QML and allows for efficient analytical evaluation.

Definition 7 ((α,ϵ)-RDP of a Quantum Measurement)**.** *Let $\epsilon \geq 0$ and $\alpha \in (1,\infty)$. A quantum measurement $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$ is said to satisfy (α,ϵ)-Rényi differential privacy if, for all neighboring quantum states $\rho \sim \sigma$, it holds that*

$$D_\alpha(\mathcal{M}(\rho)\|\mathcal{M}(\sigma)) \leq \epsilon.$$

This definition directly extends classical RDP to the output distributions of quantum measurements. It is particularly relevant in QML, where measurements are typically fixed and repeated, and thus serve as the only classical interface through which information about the input state may be revealed.

Compared to (ϵ,δ)-differential privacy, which bounds the worst-case multiplicative deviation in outcome probabilities, RDP provides a tunable relaxation that can yield

tighter privacy bounds and more flexible utility-privacy trade-offs. The order parameter α controls the sensitivity of the divergence to rare events: as $\alpha \to 1$, it approximates average-case behavior; as $\alpha \to \infty$, it converges to worst-case ϵ-DP.

We now present two standard connections between RDP and (ϵ, δ)-DP.

Proposition 2 (From RDP to (ϵ, δ)-DP)**.** *If a measurement $\mathcal{M}$ satisfies (α, ϵ)-RDP, then it also satisfies $(\epsilon_\delta, \delta)$-DP for any $\delta \in (0, 1)$, where*

$$\epsilon_\delta = \epsilon + \frac{\log(1/\delta)}{\alpha - 1}.$$

Proof. The proof follows the same argument as Proposition 3 in Mironov [Mir17]. In our setting, the inputs are neighboring quantum states ρ and σ, and the mechanism $\mathcal{M}$ is a fixed quantum measurement that maps quantum states to classical output distributions. The Rényi divergence $D_\alpha(\mathcal{M}(\rho)\|\mathcal{M}(\sigma))$ is therefore computed between these classical distributions.

Since the output of $\mathcal{M}$ is classical, the same reasoning from the classical case applies directly. By applying the standard result that (α, ϵ)-RDP implies $(\epsilon_\delta, \delta)$-DP for classical distributions, we obtain the same guarantee in this hybrid quantum-classical context. No modification is needed beyond interpreting the mechanism as the measurement-induced classical distribution. □

Proposition 3 (From ϵ-DP to RDP)**.** *If a measurement $\mathcal{M}$ satisfies ϵ-differential privacy, then for any $\alpha > 1$, it also satisfies $(\alpha, \frac{\alpha}{\alpha-1} \cdot \epsilon)$-RDP:*

$$D_\alpha(\mathcal{M}(\rho)\|\mathcal{M}(\sigma)) \leq \frac{\alpha}{\alpha - 1} \cdot \epsilon.$$

Proof. Let $p_i = \mathrm{tr}(M_i\rho)$ and $q_i = \mathrm{tr}(M_i\sigma)$ denote the output probabilities of the measurement outcomes. The definition of Rényi divergence of order $\alpha > 1$ is:

$$D_\alpha(p\|q) = \frac{1}{\alpha - 1} \log \left(\sum_{i \in \mathcal{O}} p_i^\alpha q_i^{1-\alpha} \right).$$

Observe that:

$$p_i^\alpha q_i^{1-\alpha} = \left(\frac{p_i}{q_i}\right)^\alpha q_i \leq e^{\alpha\epsilon} q_i,$$

where the inequality follows from the assumption $\frac{p_i}{q_i} \leq e^\epsilon$ for all i.

Summing over i, we obtain:

$$\sum_i p_i^\alpha q_i^{1-\alpha} \leq e^{\alpha\epsilon} \sum_i q_i = e^{\alpha\epsilon}.$$

Therefore,

$$D_\alpha(p\|q) \leq \frac{1}{\alpha - 1} \log\left(e^{\alpha\epsilon}\right) = \frac{\alpha}{\alpha - 1} \cdot \epsilon,$$

which completes the proof. □

In summary, analyzing the Rényi differential privacy of quantum measurements provides a practical and principled method for quantifying the privacy of QML algorithms. Since all observable outputs in QML arise from measurements, this perspective allows us to directly study the privacy impact of quantum processing pipelines using established divergence-based tools.

Next, we observe that several desirable properties satisfied by quantum (ϵ, δ) -DP, such as post-processing and composition, also hold under our framework.

4.1.1 Post-Processing Properties of RDP for Quantum Measurements

The Rényi differential privacy guarantees for quantum measurements are preserved under both classical post-processing of measurement outcomes and quantum pre-processing of input states. These properties, analogous to those in classical RDP, ensure that privacy remains intact even when additional operations are composed with the measurement procedure. We state both results formally below.

Theorem 7 (Classical Post-Processing)**.** *Let $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$ be a quantum measurement satisfying (α, ϵ)-RDP. Let $g : \mathcal{O} \to \mathcal{O}'$ be a (possibly randomized) classical function applied*

to the measurement outcomes. Then the composed mechanism $g \circ \mathcal{M}$, defined as

$$g \circ \mathcal{M}(\rho) := g(\mathcal{M}(\rho)),$$

also satisfies (α, ϵ)-RDP.

Proof. This result follows directly from the data processing inequality for Rényi divergence [VH14]. Consider a randomized mapping $g : \mathcal{O} \to \mathcal{O}'$. By the data processing inequality, we have

$$D_\alpha(\mathcal{M}(\rho) \| \mathcal{M}(\sigma)) \geq D_\alpha(g(\mathcal{M}(\rho)) \| g(\mathcal{M}(\sigma))).$$

Therefore, if $\mathcal{M}$ satisfies (α, ϵ)-RDP, so does $g \circ \mathcal{M}$. □

This result states that applying any classical transformation to the output distribution—such as grouping, labeling, or perturbing outcomes—does not increase the privacy loss beyond that of the original measurement.

Theorem 8 (Quantum Pre-Processing)**.** *Let $\mathcal{M} = \{M_i\}_{i \in \mathcal{O}}$ be a quantum measurement satisfying (α, ϵ)-RDP. Let $\mathcal{E}_N$ be a quantum channel. Define the modified measurement as*

$$\mathcal{M}_{\mathcal{E}_N} := \{\mathcal{E}_N^\dagger(M_i)\}_{i \in \mathcal{O}},$$

where $\mathcal{E}_N^\dagger$ is the dual (Heisenberg-picture) channel of $\mathcal{E}_N$. Then $\mathcal{M}_{\mathcal{E}_N}$ also satisfies (α, ϵ)-RDP.

Proof. By Proposition 2 in Zhou et al. [ZY17], quantum channels contract trace distance. That is, for any neighboring quantum states ρ and σ, we have

$$D(\mathcal{E}_N(\rho), \mathcal{E}_N(\sigma)) \leq D(\rho, \sigma).$$

This implies that if ρ and σ are neighboring, then $\mathcal{E}_N(\rho)$ and $\mathcal{E}_N(\sigma)$ are also neighboring. Since $\mathcal{M}$ satisfies (α, ϵ)-RDP, it follows that the composed process $\mathcal{M} \circ \mathcal{E}_N$ also satisfies

(α, ϵ)-RDP. In the Heisenberg picture, this is equivalent to saying that the effective measurement $\mathcal{M}_{\mathcal{E}_N} = \{\mathcal{E}_N^\dagger(M_i)\}_{i\in\mathcal{O}}$ also satisfies (α, ϵ)-RDP. Here $D(\rho, \sigma)$ denotes the trace distance between ρ and σ, consistent with the notation used throughout this paper. □

This theorem shows that applying quantum noise (or any CPTP map) before measurement preserves the RDP guarantee. This includes practical noise processes such as depolarization, amplitude damping, or encoding layers.

Together, these results demonstrate that RDP guarantees are robust under both classical and quantum operations that precede or follow the measurement. In the context of QML, this means privacy guarantees at the measurement interface are preserved even in the presence of noise or classical post-processing, making RDP a modular and reliable privacy model for analyzing quantum measurement outcomes.

4.1.2 Composition Theorem for RDP of Quantum Measurements

To support the modular analysis of QML algorithms, it is important to understand how Rényi differential privacy composes across multiple measurements. In this section, we formalize the composition rule when two or more measurements are applied independently to different subsystems of a composite quantum state.

Let $\mathcal{M}_1 = \{M_i\}_{i\in\mathcal{O}_1}$ and $\mathcal{M}_2 = \{M_j\}_{j\in\mathcal{O}_2}$ be two quantum measurements defined on Hilbert spaces $\mathcal{H}_1$ and $\mathcal{H}_2$, respectively. Their joint measurement is defined as

$$\mathcal{M}_{1,2} = \{M_{i,j} := M_i \otimes M_j \mid i \in \mathcal{O}_1,\, j \in \mathcal{O}_2\},$$

which maps product states $\rho_{1,2} = \rho_1 \otimes \rho_2 \in \mathcal{D}(\mathcal{H}_1 \otimes \mathcal{H}_2)$ to classical distributions over $\mathcal{O}_1 \times \mathcal{O}_2$.

We say that two product states $\rho_{1,2} = \rho_1 \otimes \rho_2$ and $\sigma_{1,2} = \sigma_1 \otimes \sigma_2$ are neighboring, denoted $\rho_{1,2} \sim \sigma_{1,2}$, if $\rho_1 \sim \sigma_1$ and $\rho_2 \sim \sigma_2$, where neighboring is defined via trace distance (see Section 3.2.2). This definition extends naturally to n-partite tensor products.

We now state the composition theorem.

Theorem 9 (Composition of RDP for Quantum Measurements)**.** *Let $\mathcal{M}_1$ be an (α, ϵ_1)-RDP quantum measurement on $\mathcal{H}_1$, and let $\mathcal{M}_2$ be an (α, ϵ_2)-RDP quantum measurement on $\mathcal{H}_2$. Then the joint measurement $\mathcal{M}_{1,2}$ satisfies $(\alpha, \epsilon_1 + \epsilon_2)$-RDP on $\mathcal{H}_1 \otimes \mathcal{H}_2$.*

Proof. As $\mathcal{M}_1$ is $(\alpha, \epsilon_i$1-RDP. We can get that for all neighboring state ρ_1, ρ_1' over H_1, $D_\alpha(\mathcal{M}_1(\rho_1) \| \mathcal{M}_1(\rho_1')) \leq \epsilon_1$. That is $\frac{1}{\alpha-1} \ln(\sum_i [\mathrm{tr}(M_i\rho_1)^\alpha \cdot \mathrm{tr}(M_i\rho_1')^{1-\alpha}]) \leq \epsilon_1$, similarly, $\frac{1}{\alpha-1} \ln(\sum_j [\mathrm{tr}(M_j\rho_2)^\alpha \cdot \mathrm{tr}(M_j\rho_2')^{1-\alpha}]) \leq \epsilon_2$. So

$$
\begin{aligned}
& D_\alpha(\mathcal{M}_{1,2}(\rho_1 \otimes \rho_2) || \mathcal{M}_{1,2}(\rho_1' \otimes \rho_2')) \\
= \;& \frac{1}{\alpha-1} \ln(\sum_{i,j} (M_{i,j}(\rho_1 \otimes \rho_2))^\alpha \cdot (M_{i,j}(\rho_1' \otimes \rho_2'))^{1-\alpha}]) \\
= \;& \frac{1}{\alpha-1} \ln(\sum_{i,j} [\mathrm{tr}(M_i\rho_1) \cdot \mathrm{tr}(M_j\rho_2)]^\alpha \cdot [\mathrm{tr}(M_i\rho_1') \cdot \mathrm{tr}(M_j\rho_2')]^{1-\alpha}) \\
= \;& \frac{1}{\alpha-1} \ln(\sum_i [\mathrm{tr}(M_i\rho_1)^\alpha \cdot \mathrm{tr}(M_i\rho_1')^{1-\alpha}] \cdot \sum_j [\mathrm{tr}(M_j\rho_2)^\alpha \cdot \mathrm{tr}(M_j\rho_2')^{1-\alpha}]) \\
= \;& \frac{1}{\alpha-1} \{\ln(\sum_i [\mathrm{tr}(M_i\rho_1)^\alpha \cdot \mathrm{tr}(M_i\rho_1')^{1-\alpha}]) + \ln(\sum_j [\mathrm{tr}(M_j\rho_2)^\alpha \cdot \mathrm{tr}(M_j\rho_2')^{1-\alpha}])\} \\
\leq \;& \epsilon_1 + \epsilon_2.
\end{aligned}
$$

□

This result demonstrates that the RDP guarantees of quantum measurements are additive under composition, similar to the behavior of classical and quantum (ϵ, δ)-DP. It enables the privacy analysis of complex QML workflows by composing the privacy parameters of independently measured quantum components.

If the individual measurements are defined using different Rényi orders α_1 and α_2, one can use the monotonicity of Rényi divergence in α to upper-bound both measurements under a common $\alpha := \max\{\alpha_1, \alpha_2\}$ before applying the theorem.

4.2 Rényi Differential Privacy Verification

In this section, we present a core theoretical result that enables tractable verification of Rényi differential privacy for quantum measurements. Our formulation leverages a subset-

based inequality that arises naturally due to the structure of quantum measurements, and allows us to analytically upper bound the RDP guarantee.

4.2.1 Equivalent Condition and Approximate Converse for Rényi Differential Privacy

Let $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$ be a quantum measurement and $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ be neighboring quantum states. We consider the output distributions $\mathcal{M}(\rho)$ and $\mathcal{M}(\sigma)$ over measurement outcomes. The following theorem provides an equivalent condition and an approximate converse for Rényi differential privacy.

Theorem 10 (Equivalent Condition and Approximate Converse for RDP)**.** *Let $\mathcal{M} = \{M_i\}_{i\in\mathcal{O}}$ be a quantum measurement and $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ be neighboring quantum states. Then:*

1. *If $\mathcal{M}$ satisfies (α, ϵ)-Rényi differential privacy, then for all subsets $S \subseteq \mathcal{O}$,*

$$\sum_{i\in S} tr(M_i\rho) \leq \left(e^{\epsilon} \cdot \sum_{i\in S} tr(M_i\sigma)\right)^{\frac{\alpha-1}{\alpha}}.$$

2. *Conversely, if the above inequality holds for all $S \subseteq \mathcal{O}$, then $\mathcal{M}$ satisfies $(\alpha, \epsilon + \frac{\log|\mathcal{O}|}{\alpha-1})$-Rényi differential privacy.*

Proof. **Part 1.** This follows directly from Proposition 10 in Mironov [Mir17].

Part 2. Assume the subset inequality holds for all subsets $S \subseteq \mathcal{O}$:

$$\sum_{i\in S} \mathrm{tr}(M_i\rho) \leq \left(e^{\epsilon} \sum_{i\in S} \mathrm{tr}(M_i\sigma)\right)^{\frac{\alpha-1}{\alpha}}. \tag{1}$$

Applying this to singleton sets $S = \{i\}$ gives

$$\mathrm{tr}(M_i\rho) \leq \left(e^{\epsilon}\mathrm{tr}(M_i\sigma)\right)^{\frac{\alpha-1}{\alpha}}. \tag{2}$$

Raising both sides of (2) to the power α yields

$$\mathrm{tr}(M_i\rho)^\alpha \leq e^{\epsilon(\alpha-1)}\mathrm{tr}(M_i\sigma)^{\alpha-1}. \tag{3}$$

Multiplying (3) by $\mathrm{tr}(M_i\sigma)^{1-\alpha}$ and summing over all $i \in \mathcal{O}$, we obtain

$$\sum_{i\in\mathcal{O}} \mathrm{tr}(M_i\rho)^\alpha \mathrm{tr}(M_i\sigma)^{1-\alpha} \leq e^{\epsilon(\alpha-1)}|\mathcal{O}|. \tag{4}$$

By the definition of Rényi divergence,

$$D_\alpha(\mathcal{M}(\rho)\,\|\,\mathcal{M}(\sigma)) = \frac{1}{\alpha-1}\log\sum_i \mathrm{tr}(M_i\rho)^\alpha \mathrm{tr}(M_i\sigma)^{1-\alpha}. \tag{5}$$

Combining (4) with (5) gives

$$D_\alpha(\mathcal{M}(\rho)\,\|\,\mathcal{M}(\sigma)) \leq \epsilon + \frac{\log|\mathcal{O}|}{\alpha-1},$$

which completes the proof. □

This result enables a conservative but practical approach to RDP verification. While the subset inequality is strictly weaker than full RDP, it yields a meaningful upper bound on the privacy parameter that depends only on the number of measurement outcomes $|\mathcal{O}|$. This dependence reflects the information leakage potential of the measurement and is a direct consequence of the classical nature of its outcome statistics.

4.2.2 Analytical Computation of Rényi Differential Privacy Bounds

Based on Theorem 10, we define an upper bound ϵ^* on the Rényi differential privacy of a quantum measurement as follows:

$$\epsilon^* := \inf\left\{ \epsilon \geq 0 \;\middle|\; \forall S \subseteq \mathcal{O},\ \rho \sim \sigma,\ \sum_{i\in S}\mathrm{tr}(M_i\rho) \leq \left(e^\epsilon \sum_{i\in S}\mathrm{tr}(M_i\sigma)\right)^{\frac{\alpha-1}{\alpha}} \right\}.$$

Then, the measurement $\mathcal{M}$ satisfies $(\alpha, \epsilon^* + \frac{\log|\mathcal{O}|}{\alpha-1})$-RDP. This gives a tractable way to estimate RDP guarantees based on subset probabilities, without directly computing the full Rényi divergence between output distributions.

To compute ϵ^*, we aim to determine the smallest ϵ such that the following supremum is non-positive for all $S \subseteq \mathcal{O}$:

$$\sup_{\rho,\sigma,D(\rho,\sigma)\leq\eta} \left(\sum_{i\in S} \mathrm{tr}(M_i\rho) - \left(e^{\epsilon} \cdot \sum_{i\in S} \mathrm{tr}(M_i\sigma) \right)^{\frac{\alpha-1}{\alpha}} \right) \leq 0, \tag{4.1}$$

Let $M_S := \sum_{i\in S} M_i$. Then, the constraint simplifies to:

$$\sup_{\rho,\sigma,D(\rho,\sigma)\leq\eta} \left(\mathrm{tr}(M_S\rho) - (e^{\epsilon} \cdot \mathrm{tr}(M_S\sigma))^{\frac{\alpha-1}{\alpha}} \right) \leq 0. \tag{4.2}$$

Lower Bound. To establish a lower bound for ϵ, we evaluate the supremum by selecting a specific choice of states:

$$\gamma = \eta|\psi\rangle\langle\psi| + (1-\eta)|\phi\rangle\langle\phi|, \quad \phi = |\phi\rangle\langle\phi|, \tag{4.3}$$

where:

- $|\psi\rangle$ is the eigenvector of M_S corresponding to the largest eigenvalue $\lambda_{\max}(M_S)$,
- $|\phi\rangle$ is the eigenvector of M_S corresponding to the smallest eigenvalue $\lambda_{\min}(M_S)$,
- $|\psi\rangle \perp |\phi\rangle$ (orthogonal).

Since the trace distance between γ and ϕ is:

$$D(\gamma,\phi) = \frac{1}{2}\|\gamma-\phi\|_1 = \eta\sqrt{1-|\langle\psi|\phi\rangle|^2} = \eta,$$

this choice respects the trace distance constraint.

Hence:

$$
\begin{aligned}
& \sup_{\rho,\sigma,D(\rho,\sigma)\leq\eta} \left(\mathrm{tr}(M_S\rho) - (e^{\epsilon}\cdot \mathrm{tr}(M_S\sigma))^{\frac{\alpha-1}{\alpha}}\right) \\
\geq\ & \mathrm{tr}(M_S\gamma) - (e^{\epsilon}\cdot \mathrm{tr}(M_S\phi))^{\frac{\alpha-1}{\alpha}} \\
=\ & \eta\lambda_{\max}(M_S) + (1-\eta)\lambda_{\min}(M_S) - (e^{\epsilon}\cdot\lambda_{\min}(M_S))^{\frac{\alpha-1}{\alpha}}\ .
\end{aligned}
$$

Upper Bound. On the other hand, for any quantum states $\rho,\sigma \in \mathcal{D}(\mathcal{H})$ with $D(\rho,\sigma)\leq\eta$, let $\rho-\sigma = \Delta_+ - \Delta_-$ be a decomposition into orthogonal positive and negative parts (i.e., $\Delta_\pm \geq 0$ and $\Delta_+\Delta_- = 0$). Then, we have $\mathrm{tr}(\Delta_+) = \mathrm{tr}(\Delta_-)$, since $0 = \mathrm{tr}(\rho-\sigma) = \mathrm{tr}(\Delta_+ - \Delta_-)$. Thus, the trace norm condition gives:$D(\rho,\sigma) = \mathrm{tr}(\Delta_+) \leq \eta$. Furthermore, we compute:

$$
\begin{aligned}
& \sup_{\rho,\sigma,D(\rho,\sigma)\leq\eta} \left[\mathrm{tr}(M_S\rho) - (e^{\epsilon}\mathrm{tr}(M_S\sigma))^{\frac{\alpha-1}{\alpha}}\right] \\
=\ & \sup_{\mathrm{tr}(\Delta_-)=\mathrm{tr}(\Delta_+)\leq\eta} \left[\mathrm{tr}(M_S(\Delta_+ - \Delta_- + \sigma)) - (e^{\epsilon}\mathrm{tr}(M_S\sigma))^{\frac{\alpha-1}{\alpha}}\right] \\
=\ & \sup_{\mathrm{tr}(\Delta_-)=\mathrm{tr}(\Delta_+)\leq\eta} [\mathrm{tr}(M_S\Delta_+) - \mathrm{tr}(M_S\Delta_-) \\
& + \mathrm{tr}(M_S\sigma) - (e^{\epsilon}\mathrm{tr}(M_S\sigma))^{\frac{\alpha-1}{\alpha}}\Big] \\
=\ & \sup_{\mathrm{tr}(\Delta_-)=\mathrm{tr}(\Delta_+)\leq\eta} \left[\mathrm{tr}(\Delta_+)\mathrm{tr}(M_S\frac{\Delta_+}{\mathrm{tr}(\Delta_+)}) \right. \\
& \left. -\mathrm{tr}(\Delta_-)\mathrm{tr}(M_S\frac{\Delta_-}{\mathrm{tr}(\Delta_-)}) + \mathrm{tr}(M_S\sigma) - (e^{\epsilon}\mathrm{tr}(M_S\sigma))^{\frac{\alpha-1}{\alpha}}\right] \\
\leq\ & \sup_{\mathrm{tr}(\Delta_-)=\mathrm{tr}(\Delta_+)\leq\eta} \left[\mathrm{tr}(\Delta_+) \max_{\rho_1\in\mathcal{D}(\mathcal{H})} \mathrm{tr}(M_S\rho_1) \right. \\
& \left. -\mathrm{tr}(\Delta_-) \min_{\rho_2\in\mathcal{D}(\mathcal{H})} \mathrm{tr}(M_S\rho_2) + \mathrm{tr}(M_S\sigma) - (e^{\epsilon}\mathrm{tr}(M_S\sigma))^{\frac{\alpha-1}{\alpha}}\right] \\
=\ & \sup_{\mathrm{tr}(\Delta_-)=\mathrm{tr}(\Delta_+)\leq\eta} [\mathrm{tr}(\Delta_+)(\lambda_{\max}(M_S) - \lambda_{\min}(M_S)) \\
& + \mathrm{tr}(M_S\sigma) - (e^{\epsilon}\mathrm{tr}(M_S\sigma))^{\frac{\alpha-1}{\alpha}}\Big] \\
\leq\ & \eta\lambda_{\max}(M_S) - \eta\lambda_{\min}(M_S) + \mathrm{tr}(M_S\sigma) - (e^{\epsilon}\mathrm{tr}(M_S\sigma))^{\frac{\alpha-1}{\alpha}}
\end{aligned}
$$

If we consider the expression

$$\mathrm{tr}(M_S\sigma) - (e^{\epsilon}\mathrm{tr}(M_S\sigma))^{\frac{\alpha-1}{\alpha}}$$

as a function of $\mathrm{tr}(M_S\sigma)$, we observe that it is a convex function. This implies that its maximum value is attained at the boundary points of its domain.

Since we have the constraint:

$$\lambda_{\min}(M_S) \leq \mathrm{tr}(M_S\sigma) \leq \lambda_{\max}(M_S),$$

the supremum of this function is achieved at either $\mathrm{tr}(M_S\sigma) = \lambda_{\min}(M_S)$ or $\mathrm{tr}(M_S\sigma) = \lambda_{\max}(M_S)$. Thus, its optimal value is given by:

$$\max\{\lambda_{\min}(M_S) - (e^{\epsilon}\lambda_{\min}(M_S))^{\frac{\alpha-1}{\alpha}}, \lambda_{\max}(M_S) - (e^{\epsilon}\lambda_{\max}(M_S))^{\frac{\alpha-1}{\alpha}}\}$$

Therefore, we conclude the following:

$$\begin{aligned} & \sup_{\rho,\sigma,D(\rho,\sigma)\leq\eta} \left[\mathrm{tr}(M_S\rho) - (e^{\epsilon}\mathrm{tr}(M_S\sigma))^{\frac{\alpha-1}{\alpha}}\right] \\ \leq \quad & \max\{\eta\lambda_{\max}(M_S) + (1-\eta)\lambda_{\min}(M_S) - (e^{\epsilon}\lambda_{\min}(M_S))^{\frac{\alpha-1}{\alpha}}, \\ & (1+\eta)\lambda_{\max}(M_S) - \eta\lambda_{\min}(M_S) - (e^{\epsilon}\lambda_{\max}(M_S))^{\frac{\alpha-1}{\alpha}}\} \end{aligned}$$

Two cases. We now consider two possible cases depending on which term achieves the maximum value in the following expression:

$$\max\left\{\begin{array}{l} \eta\lambda_{\max}(M_S) + (1-\eta)\lambda_{\min}(M_S) - (e^{\epsilon}\lambda_{\min}(M_S))^{\frac{\alpha-1}{\alpha}}, \\ (1+\eta)\lambda_{\max}(M_S) - \eta\lambda_{\min}(M_S) - (e^{\epsilon}\lambda_{\max}(M_S))^{\frac{\alpha-1}{\alpha}} \end{array}\right\}.$$

Case 1. If the first term achieves the maximum, i.e.,

$$\begin{aligned} & \eta\lambda_{\max}(M_S) + (1-\eta)\lambda_{\min}(M_S) - (e^{\epsilon}\lambda_{\min}(M_S))^{\frac{\alpha-1}{\alpha}} \\ \geq \quad & (1+\eta)\lambda_{\max}(M_S) - \eta\lambda_{\min}(M_S) - (e^{\epsilon}\lambda_{\max}(M_S))^{\frac{\alpha-1}{\alpha}}, \end{aligned}$$

then this upper bound matches exactly with the lower bound we constructed earlier. In this case, we can conclude that the supremum is tight, and the optimal Rényi DP parameter ϵ^* is given by:

$$\epsilon^* = \ln\left(\frac{[\eta\lambda_{\max}(M_S) + (1-\eta)\lambda_{\min}(M_S)]^{\frac{\alpha}{\alpha-1}}}{\lambda_{\min}(M_S)}\right). \tag{4.4}$$

Case 2. If instead the second term is larger, i.e.,

$$\begin{aligned} & \eta\lambda_{\max}(M_S) + (1-\eta)\lambda_{\min}(M_S) - (e^{\epsilon}\lambda_{\min}(M_S))^{\frac{\alpha-1}{\alpha}} \\ < \quad & (1+\eta)\lambda_{\max}(M_S) - \eta\lambda_{\min}(M_S) - (e^{\epsilon}\lambda_{\max}(M_S))^{\frac{\alpha-1}{\alpha}}, \end{aligned}$$

then the supremum is dominated by the second expression, and our lower-bound construction no longer achieves the maximum. In this case, we cannot solve for the exact value of ϵ^*, but we can still guarantee Rényi differential privacy by ensuring that this larger upper bound is less than or equal to zero. That is, we require:

$$\eta\lambda_{\max}(M_S) - \eta\lambda_{\min}(M_S) + \lambda_{\max}(M_S) - (e^{\epsilon}\lambda_{\max}(M_S))^{\frac{\alpha-1}{\alpha}} \leq 0, \tag{4.5}$$

which gives an *upper bound* on ϵ^* that guarantees the supremum remains non-positive. Specifically, we have:

$$\epsilon^* \leq \ln\left(\frac{[(1+\eta)\lambda_{\max}(M_S) - \eta\lambda_{\min}(M_S)]^{\frac{\alpha}{\alpha-1}}}{\lambda_{\max}(M_S)}\right).$$

Following the analytical construction discussed above, we formalize the procedure for computing the optimal privacy parameter ϵ^* into a practical algorithmic form. The algorithm iterates over all subsets $S \subseteq \mathcal{O}$ of measurement outcomes and computes two candidate values for each subset: one based on the analytical lower bound (referred to as the *tight bound*) and the other based on a general upper bound. The larger of the two is selected as the candidate privacy parameter ϵ_S for the subset S. Among all subsets, the algorithm identifies the maximum such value and returns it as the final result ϵ^*.

This procedure, presented as Algorithm 3, not only provides a direct computation of Rényi differential privacy bounds but also indicates whether the bound is analytically tight or conservatively estimated. It serves as a bridge between theoretical analysis and practical verification of privacy guarantees.

The verification of quantum Rényi differential privacy is divided into two steps: first, transforming the quantum channel and measurement into an effective measurement in the Heisenberg picture (Algorithm 2); second, computing the optimal Rényi privacy parameter based on the resulting measurement (Algorithm 3). In the following, we analyze the correctness and complexity of these two algorithms.

Correctness. Algorithm 2 correctly computes the Heisenberg dual measurement operators $\{W_k\}_{k\in\mathcal{O}}$ from the quantum channel $\mathcal{E}$ and measurement $\{M_k\}$ via the adjoint map $\mathcal{E}^\dagger$, satisfying the relation

$$\operatorname{tr}\left(M_k\,\mathcal{E}(\rho)\right) = \operatorname{tr}\left(\mathcal{E}^\dagger(M_k)\,\rho\right),$$

for all quantum states ρ and measurement indices $k \in \mathcal{O}$. This transformation ensures that the resulting effective measurement preserves the statistical behavior of the original quantum process, thereby maintaining the correctness of subsequent privacy evaluations.

Algorithm 3 implements the analytical bound derived in Theorem 10, computing an upper bound $\hat{\epsilon} \geq \epsilon^*$ for the optimal Rényi differential privacy parameter. For each subset $S \subseteq \mathcal{O}$, the algorithm evaluates two candidate expressions corresponding to tight and upper bounds and selects the dominant one. The result $\hat{\epsilon} = \max_S \epsilon_S$ matches the analytical form, and the algorithm correctly records whether the bound is achieved tightly (i.e., when the tight expression dominates). The quantities A_S and B_S used in the computation follow the definitions in Theorem 10.

Complexity. Let d be the dimension of the Hilbert space, r the number of Kraus operators in the channel $\mathcal{E}$, and $m = |\mathcal{O}|$ the number of measurement outcomes. The

computational complexity of Algorithm 1 is $O(mrd^3)$, dominated by the application of each Kraus operator $E_j^\dagger M_k E_j$ and the accumulation over r operators for each of the m measurements.

For Algorithm 3, all subsets $S \subseteq \mathcal{O}$ are enumerated, leading to 2^m evaluations. For each subset, the computation includes matrix summation $M_S = \sum_{k \in S} W_k$ with cost $O(md^2)$, and eigenvalue computation of the Hermitian matrix M_S, costing $O(d^3)$. Thus, the total complexity is $O(2^m \cdot d^3)$. This is tractable for small m, as typically encountered in fixed quantum measurements for near-term quantum machine learning circuits.

Algorithm 2 Heisenberg Dual Measurement Transformation

Input: Quantum channel $\mathcal{E} = \{E_j\}_{j \in J}$, quantum measurement operators $\{M_k\}_{k \in \mathcal{O}}$
Output: Heisenberg dual measurement operators $\{W_k\}_{k \in \mathcal{O}}$

1: **for all** $k \in \mathcal{O}$ **do**
2: $\quad W_k \leftarrow \mathcal{E}^\dagger(M_k) = \sum_{j \in J} E_j^\dagger M_k E_j$
3: **end for**
4: **return** $\{W_k\}_{k \in \mathcal{O}}$

Algorithm 3 Computation of an Upper Bound $\hat{\epsilon}$ for the Optimal Parameter ϵ^*

Input: Quantum measurement $\{W_k\}_{k\in\mathcal{O}}$, Rényi order $\alpha > 1$, privacy radius $\eta \geq 0$
Output: Estimated upper bound $\hat{\epsilon}$ and whether it corresponds to a tight bound

$\hat{\epsilon} \leftarrow 0$, $S^* \leftarrow \emptyset$, $tightFlag \leftarrow$ `false`
for all subset $S \subseteq \mathcal{O}$ **do**
    $M_S \leftarrow \sum_{k\in S} W_k$
    Compute $\lambda_{\min}(M_S)$ and $\lambda_{\max}(M_S)$
    $\epsilon_{\text{tight}} \leftarrow \ln\left(\frac{A_S^{\frac{\alpha}{\alpha-1}}}{\lambda_{\min}(M_S)}\right)$
    $\epsilon_{\text{upper}} \leftarrow \ln\left(\frac{B_S^{\frac{\alpha}{\alpha-1}}}{\lambda_{\max}(M_S)}\right)$
    if $\epsilon_{\text{tight}} \geq \epsilon_{\text{upper}}$ **then**
        $\epsilon_S \leftarrow \epsilon_{\text{tight}}$, flag $\leftarrow$ `true`
    else
        $\epsilon_S \leftarrow \epsilon_{\text{upper}}$, flag $\leftarrow$ `false`
    end if
    if $\epsilon_S > \hat{\epsilon}$ **then**
        $\hat{\epsilon} \leftarrow \epsilon_S$, $S^* \leftarrow S$, $tightFlag \leftarrow$ flag
    end if
end for
if *tightFlag* is `true` **then**
    return $\hat{\epsilon}$ is tight; Rényi DP holds with exact bound
else
    return $\hat{\epsilon}$ is an upper bound
end if

4.3 Evaluation of RDP

We evaluate the Rényi differential privacy guarantees of quantum machine learning circuits with fixed quantum measurements. All simulations were conducted in Python using Qiskit and PennyLane, with state evolution performed via Qiskit's statevector simulator. The RDP parameter ϵ, representing an upper bound on the privacy cost, is computed analytically via eigenvalue-based analysis using Algorithm 2 and Algorithm 3.

Platform. Experiments were run on a desktop machine with Intel64 Family 6 Model 167 CPU, 17 GB RAM, running Windows 10 and no GPU acceleration.

4.3.1 Quantum Circuits and Datasets

We evaluate four QML models trained on open-source datasets [DG+17; MST20; LCB+10; Has+25]. All classical features are encoded into quantum states using angle encoding, a standard and hardware-efficient strategy in QML implementations.

- **QML classifiers on standard datasets:**
  - **MNIST_10:** Trained on the MNIST dataset to distinguish handwritten digits "0" and "1".
  - **Iris_4:** Trained on a binary version of the Iris dataset to classify flower species using four numerical features.
- **QML models trained on privacy-sensitive data:**
  - **GC_6:** Trained on the German Credit dataset to classify whether an individual has a good credit rating.
  - **AI_9:** Trained on the UCI Adult Income dataset to predict whether a person earns more than $50,000.

4.3.2 Privacy Evaluation Procedure

We use Algorithm 3 to compute an upper bound ϵ on the Rényi differential privacy guarantee for each QML circuit. This method enables a fine-grained analysis of how noise strength and mechanism impact the distinguishability between neighboring quantum states, and thereby the resulting privacy guarantee.

To simulate noisy quantum execution in the NISQ regime, we inject two types of quantum noise, bit-flip and depolarizing, into each circuit after unitary evolution but before measurement. Noise is applied globally to all qubits. For each noise model, we vary the noise probability over a logarithmic scale $p \in \{10^{-1}, 10^{-2}, \ldots, 10^{-6}\}$. For each configuration, we compute ϵ over a range of privacy radii $\eta \in [0.001, 1.0]$, with the Rényi order fixed at $\alpha = 5$.

The resulting ϵ-versus-η curves under different noise probabilities provide insights into the privacy-enhancing effects of quantum noise in hybrid quantum-classical algorithms.

Due to computational limits of classical simulation, our experiments focus on circuits involving up to 10 qubits. However, the proposed analytical framework is designed to scale to larger quantum circuits, as it only relies on the eigenvalue structure of measurement operators rather than full quantum state simulation. Moreover, by leveraging composition theorems for Rényi differential privacy, our method can be extended to analyze the cumulative privacy cost of multi-step or modular quantum algorithms operating over higher-dimensional quantum systems.

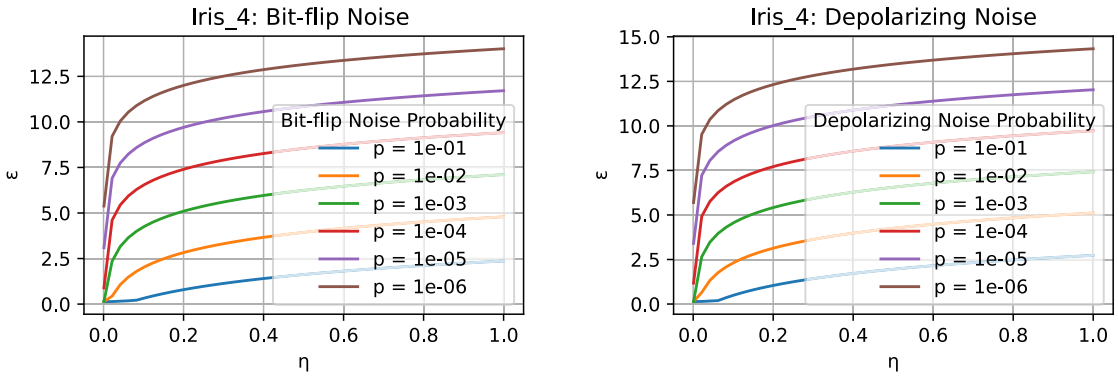

(a) Iris_4 circuit under bit-flip noise.

(b) Iris_4 circuit under depolarizing noise.

Figure 4.1: RDP upper bound ϵ results of Iris_4 circuit under bit-flip and depolarizing noise.

4.3.3 Comparison with Classical ϵ-DP Bounds

To evaluate the effectiveness of our proposed Rényi differential privacy analysis, we compare it against the classical ϵ-differential privacy bound derived from condition number analysis. We conduct experiments under two common quantum noise models: bit-flip and depolarizing noise.

The results are shown in Figure 4.5a and Figure 4.5b. Our RDP method consistently yields tighter and more stable bounds than ϵ-DP. This confirms that the analytical RDP approach offers a more precise characterization of differential privacy guarantees in noisy

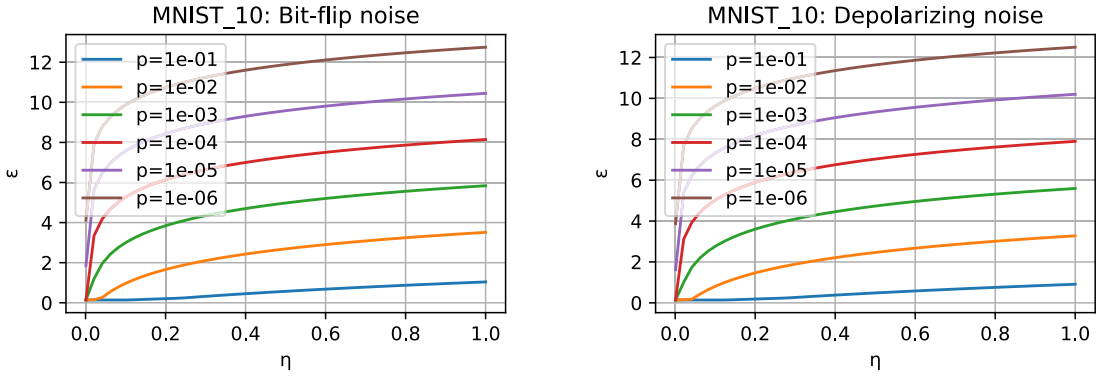

(a) MNIST_10 circuit under bit-flip noise.

(b) MNIST_10 circuit under depolarizing noise.

Figure 4.2: RDP upper bound ϵ results of MNIST_10 circuit under bit-flip and depolarizing noise.

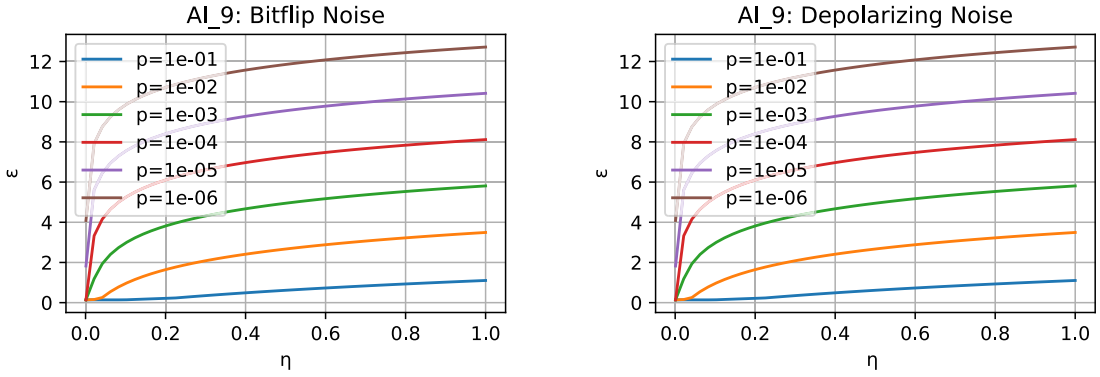

(a) AI_9 circuit under bit-flip noise.

(b) AI_9 circuit under depolarizing noise.

Figure 4.3: RDP upper bound ϵ results of AI_9 circuit under bit-flip and depolarizing noise.

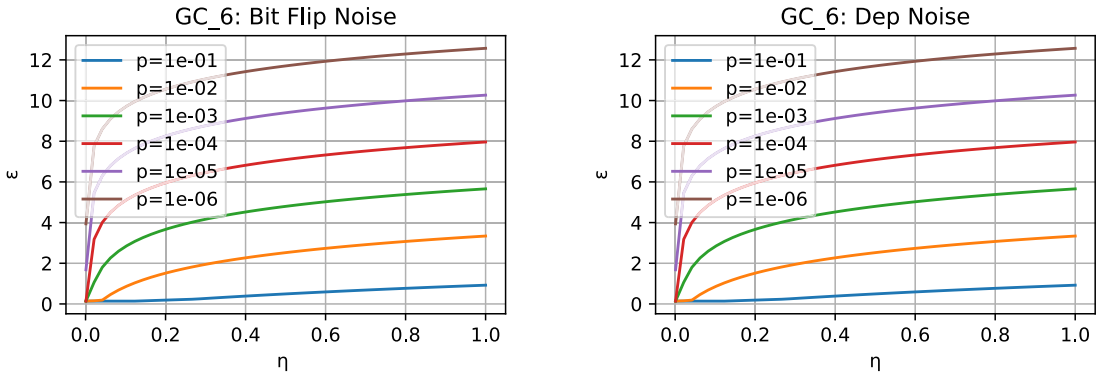

(a) GC_6 circuit under bit-flip noise.

(b) GC_6 circuit under depolarizing noise.

Figure 4.4: RDP upper bound ϵ results of GC_6 circuit under bit-flip and depolarizing noise.

quantum machine learning circuits.

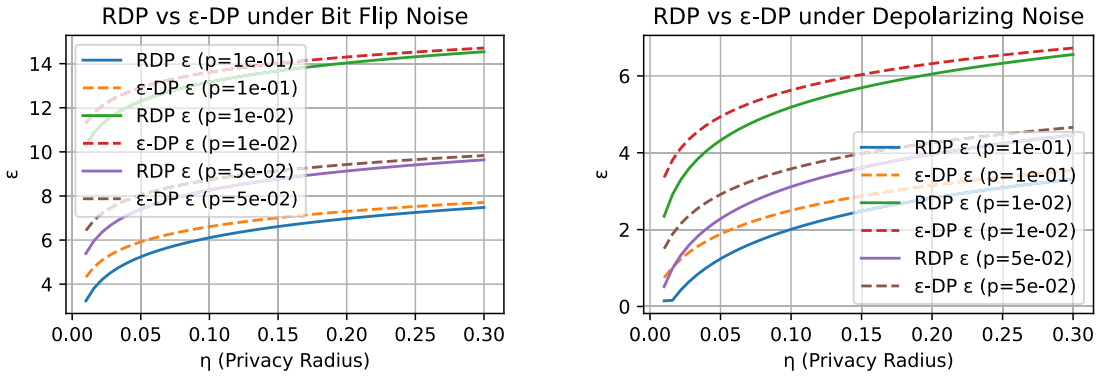

(a) Bit-flip noise

(b) Depolarizing noise

Figure 4.5: Comparison of Rényi DP and classical ϵ-DP bounds under two types of quantum noise.

Chapter 5

Conclusion and Future Direction

This thesis presents a systematic study on differential privacy in hybrid quantum-classical algorithms, focusing on fixed quantum measurements prevalent in practical quantum machine learning.

In Chapter 3, we proposed and established the Hybrid Differential Privacy (HDP) framework. This framework bridges classical and quantum differential privacy by focusing on protecting measurement outcomes in hybrid workflows. We developed post-processing and composition theorems ensuring that privacy guarantees are preserved under realistic algorithmic pipelines. We introduced two mechanisms:

- A Measurement-Based Exponential Mechanism (MBEM), adapting the classical exponential mechanism to quantum measurements to provide flexible privacy protection with tunable parameters.
- A quantum depolarizing noise strategy, allowing privacy guarantees independent of measurement details and neighboring definitions.

We validated these mechanisms through extensive experiments on single and multi-qubit circuits, analyzing privacy-utility trade-offs under varied noise strengths and circuit depths. The results demonstrate that both mechanisms effectively enforce privacy while maintaining high utility in practical quantum learning scenarios.

In Chapter 4, we extended Rényi Differential Privacy to the context of quantum machine learning with fixed measurements. We formulated RDP using Rényi divergence over measurement-induced distributions and established post-processing and composition properties in this setting, providing a framework to quantify and analyze privacy guarantees in quantum machine learning tasks under fixed measurements.

Together, these results provide a practical and unified framework for achieving and verifying differential privacy in hybrid quantum-classical algorithms, using both ϵ-DP and RDP perspectives. By focusing on fixed measurement structures, this work connects classical privacy theory with quantum implementations, facilitating secure and privacy-preserving quantum machine learning applicable to current and future quantum computing systems.

Future Work

There are several promising directions for future research:

- **Privacy in Adaptive and Multi-Measurement Settings:** Extending our framework to handle adaptive measurements and scenarios involving multiple incompatible measurements is a natural next step. Such settings frequently arise in quantum tomography, quantum kernel learning, and adaptive variational algorithms, where the interplay between measurement adaptivity and privacy guarantees remains largely unexplored.

- **Optimal Mechanism Design:** While the noise mechanisms we proposed are effective, designing optimal noise channels with provable privacy–utility trade-offs remains an open challenge. Future work could investigate broader classes of quantum channels and employ optimization-based approaches to develop privacy-preserving measurements that balance theoretical guarantees with practical performance.

- **Experimental Validation on Quantum Hardware:** As quantum devices continue to advance, validating our mechanisms on real hardware (e.g., IBM Quantum

superconducting processors and IonQ trapped-ion devices) under realistic noise will be essential. Planned evaluations include comparing empirical measurement distributions with theoretical predictions, estimating empirical Rényi divergence, and assessing the privacy–utility trade-off in practical settings. These experiments would provide insight into the robustness of our mechanisms and their feasibility for deployment in real quantum machine-learning workflows.

Through this work, we hope to contribute not only to the theoretical foundation of quantum differential privacy but also to its practical adoption in hybrid quantum-classical algorithms. As quantum computing technologies evolve, the importance of secure and privacy-preserving quantum data processing will only continue to grow.

Bibliography

[AR19] Scott Aaronson and Guy N Rothblum. "Gentle measurement of quantum states and differential privacy". In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. 2019, pp. 322–333.

[Aba+16] Martin Abadi et al. "Deep learning with differential privacy". In: *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*. 2016, pp. 308–318.

[ADK22] Armando Angrisani, Mina Doosti, and Elham Kashefi. "Differential privacy amplification in quantum and quantum-inspired algorithms". In: *arXiv preprint arXiv:2203.03604* (2022).

[ADK23] Armando Angrisani, Mina Doosti, and Elham Kashefi. "A unifying framework for differentially private quantum algorithms". In: *arXiv preprint arXiv:2307.04733* (2023).

[AK22] Armando Angrisani and Elham Kashefi. "Quantum local differential privacy and quantum statistical query model". In: *arXiv preprint arXiv:2203.03591* (2022).

[AZ24] Shahab Asoodeh and Huanyu Zhang. "Contraction of locally differentially private mechanisms". In: *IEEE Journal on Selected Areas in Information Theory* (2024).

[Ave+17] Brendan Avent et al. "{BLENDER}: Enabling local search with a hybrid differential privacy model". In: *26th USENIX Security Symposium (USENIX Security 17)*. 2017, pp. 747–764.

[BD14] Rina Foygel Barber and John C Duchi. "Privacy and statistical risk: Formalisms and minimax bounds". In: *arXiv preprint arXiv:1412.4451* (2014).

[BO13] Gilles Barthe and Federico Olmedo. "Beyond differential privacy: Composition theorems and relational logic for f-divergences between probabilistic programs". In: *International Colloquium on Automata, Languages, and Programming*. Springer. 2013, pp. 49–60.

[Bar+12] Gilles Barthe et al. "Probabilistic relational reasoning for differential privacy". In: *Proceedings of the 39th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. 2012, pp. 97–110.

[Bar+13] Gilles Barthe et al. "Verified computational differential privacy with applications to smart metering". In: *2013 IEEE 26th Computer Security Foundations Symposium*. IEEE. 2013, pp. 287–301.

[Bar+14] Gilles Barthe et al. “Proving differential privacy in Hoare logic”. In: *2014 IEEE 27th Computer Security Foundations Symposium*. IEEE. 2014, pp. 411–424.

[Bar+16a] Gilles Barthe et al. “Advanced probabilistic couplings for differential privacy”. In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. 2016, pp. 55–67.

[Bar+16b] Gilles Barthe et al. “Proving differential privacy via probabilistic couplings”. In: *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science*. 2016, pp. 749–758.

[Bas+16] Raef Bassily et al. “Algorithmic stability for adaptive data analysis”. In: *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. 2016, pp. 1046–1059.

[CRA18] Yudong Cao, Jhonathan Romero, and Alán Aspuru-Guzik. “Potential of quantum computing for drug discovery”. In: *IBM Journal of Research and Development* 62.6 (2018), pp. 6–1.

[CFS02] Carlton M Caves, Christopher A Fuchs, and Rüdiger Schack. “Unknown quantum states: the quantum de Finetti representation”. In: *Journal of Mathematical Physics* 43.9 (2002), pp. 4537–4559.

[Dal77] Tore Dalenius. “Towards a methodology for statistical disclosure control”. In: (1977).

[De +21] Nathalie P De Leon et al. “Materials challenges and opportunities for quantum computing hardware”. In: *Science* 372.6539 (2021), eabb2823.

[Du+21] Yuxuan Du et al. “Quantum noise protects quantum classifiers against adversaries”. In: *Physical Review Research* 3.2 (2021), p. 023153.

[DG+17] Dheeru Dua, Casey Graff, et al. “UCI machine learning repository”. In: (2017).

[DJW13] John C Duchi, Michael I Jordan, and Martin J Wainwright. “Local privacy and statistical minimax rates”. In: *2013 IEEE 54th annual symposium on foundations of computer science*. IEEE. 2013, pp. 429–438.

[Dwo06] Cynthia Dwork. “Differential privacy”. In: *International colloquium on automata, languages, and programming*. Springer. 2006, pp. 1–12.

[DR+14] Cynthia Dwork, Aaron Roth, et al. “The algorithmic foundations of differential privacy.” In: *Found. Trends Theor. Comput. Sci.* 9.3-4 (2014), pp. 211–407.

[DR16] Cynthia Dwork and Guy N Rothblum. “Concentrated differential privacy”. In: *arXiv preprint arXiv:1603.01887* (2016).

[Dwo+06a] Cynthia Dwork et al. “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*. Springer. 2006, pp. 265–284.

[Dwo+06b] Cynthia Dwork et al. “Our data, ourselves: Privacy via distributed noise generation”. In: *Advances in cryptology-EUROCRYPT 2006: 24th annual international conference on the theory and applications of cryptographic techniques, st. Petersburg, Russia, May 28-June 1, 2006. proceedings 25*. Springer. 2006, pp. 486–503.

[ESS15] Hamid Ebadi, David Sands, and Gerardo Schneider. “Differential privacy: Now it’s getting personal”. In: *Acm Sigplan Notices* 50.1 (2015), pp. 69–81.

[End+21] Suguru Endo et al. “Hybrid quantum-classical algorithms and quantum error mitigation”. In: *Journal of the Physical Society of Japan* 90.3 (2021), p. 032001.

[FGG14] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. “A quantum approximate optimization algorithm”. In: *arXiv preprint arXiv:1411.4028* (2014).

[Gad+22] Andrea Gadotti et al. “Pool Inference Attacks on Local Differential Privacy: Quantifying the Privacy Guarantees of Apple’s Count Mean Sketch in Practice”. In: *31st USENIX Security Symposium (USENIX Security 22)*. 2022, pp. 501–518.

[GG25a] Jingtong Ge and Ji Guan. “Differential Privacy of Hybrid Quantum-Classical Algorithms”. Submitted to NeurIPS 2025. 2025.

[GG25b] Jingtong Ge and Ji Guan. “Rényi Differential Privacy in Quantum Machine Learning”. Submitted to AAAI 2025. 2025.

[GSC17] Joseph Geumlek, Shuang Song, and Kamalika Chaudhuri. “Renyi differential privacy mechanisms for posterior sampling”. In: *Advances in Neural Information Processing Systems* 30 (2017).

[GI24] Elena Ghazi and Ibrahim Issa. “Total Variation Meets Differential Privacy”. In: *IEEE Journal on Selected Areas in Information Theory* (2024).

[GHZ89] Daniel M Greenberger, Michael A Horne, and Anton Zeilinger. “Going beyond Bell’s theorem”. In: *Bell’s theorem, quantum theory and conceptions of the universe*. Springer, 1989, pp. 69–72.

[Gro96] Lov K Grover. “A fast quantum mechanical algorithm for database search”. In: *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*. 1996, pp. 212–219.

[Gua24] Ji Guan. “Optimal Mechanisms for Quantum Local Differential Privacy”. In: *arXiv preprint arXiv:2407.13516* (2024).

[Gua+23] Ji Guan et al. “Detecting violations of differential privacy for quantum algorithms”. In: *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. 2023, pp. 2277–2291.

[HRW11] Rob Hall, Alessandro Rinaldo, and Larry Wasserman. “Random differential privacy”. In: *arXiv preprint arXiv:1112.2680* (2011).

[Has+25] Widi Hastomo et al. “Quantum Iris Classifier: A Pennylane-Based Approach”. In: *International Conference on Engineering, Construction, Renewable Energy, and Advanced Materials*. 2025.

[HRF23] Christoph Hirche, Cambyse Rouzé, and Daniel Stilck França. "Quantum differential privacy: An information theory perspective". In: *IEEE Transactions on Information Theory* (2023).

[Hu+21] Changhui Hu et al. "How to make private distributed cardinality estimation practical, and get differential privacy for free". In: *30th USENIX security symposium (USENIX Security 21)*. 2021, pp. 965–982.

[JLE14] Zhanglong Ji, Zachary C Lipton, and Charles Elkan. "Differential privacy and machine learning: a survey and review". In: *arXiv preprint arXiv:1412.7584* (2014).

[Jon+19] Tyson Jones et al. "Variational quantum algorithms for discovering Hamiltonian spectra". In: *Physical Review A* 99.6 (2019), p. 062304.

[KOV14] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. "Extremal mechanisms for local differential privacy". In: *Advances in neural information processing systems* 27 (2014).

[Kor+22] Kamil Korzekwa et al. "Encoding classical information into quantum resources". In: *IEEE Transactions on Information Theory* 68.7 (2022), pp. 4518–4530.

[LC20] Ryan LaRose and Brian Coyle. "Robust data encodings for quantum classifiers". In: *Physical Review A* 102.3 (2020), p. 032420.

[LCB+10] Yann LeCun, Corinna Cortes, Chris Burges, et al. "MNIST handwritten digit database". In: (2010).

[Li+17] Ninghui Li et al. *Differential privacy: From theory to practice*. Springer, 2017.

[LWS15] Ziqi Liu, Yu-Xiang Wang, and Alexander Smola. "Fast differentially private matrix factorization". In: *Proceedings of the 9th ACM Conference on Recommender Systems*. 2015, pp. 171–178.

[Mer90] N David Mermin. "Quantum mysteries revisited". In: *Am. J. Phys* 58.8 (1990), pp. 731–734.

[Mir17] Ilya Mironov. "Rényi differential privacy". In: *2017 IEEE 30th computer security foundations symposium (CSF)*. IEEE. 2017, pp. 263–275.

[MST20] Ramaravind K Mothilal, Amit Sharma, and Chenhao Tan. "Explaining machine learning classifiers through diverse counterfactual explanations". In: *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 2020, pp. 607–617.

[MK19] Takao Murakami and Yusuke Kawamoto. "{Utility-Optimized} local differential privacy mechanisms for distribution estimation". In: *28th USENIX Security Symposium (USENIX Security 19)*. 2019, pp. 1877–1894.

[NC01] Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*. Vol. 2. Cambridge university press Cambridge, 2001.

[NGW24] Theshani Nuradha, Ziv Goldfeld, and Mark M Wilde. "Quantum pufferfish privacy: A flexible privacy framework for quantum systems". In: *IEEE Transactions on Information Theory* (2024).

[OML19] Román Orús, Samuel Mugel, and Enrique Lizaso. "Quantum computing for finance: Overview and prospects". In: *Reviews in Physics* 4 (2019), p. 100028.

[PS21] Nicolas Papernot and Thomas Steinke. "Hyperparameter tuning with renyi differential privacy". In: *arXiv preprint arXiv:2110.03620* (2021).

[Per+14] Alberto Peruzzo et al. "A variational eigenvalue solver on a photonic quantum processor". In: *Nature communications* 5.1 (2014), p. 4213.

[Pre18] John Preskill. "Quantum computing in the NISQ era and beyond". In: *Quantum* 2 (2018), p. 79.

[QAS21] Yihui Quek, Srinivasan Arunachalam, and John A Smolin. "Private learning implies quantum stability". In: *Advances in Neural Information Processing Systems* 34 (2021), pp. 20503–20515.

[SMT17] Makhamisa Senekane, Mhlambululi Mafu, and Benedict Molibeli Taele. "Privacy-preserving quantum machine learning using differential privacy". In: *2017 IEEE AFRICON*. IEEE. 2017, pp. 1432–1435.

[Sho94] Peter W Shor. "Algorithms for quantum computation: discrete logarithms and factoring". In: *Proceedings 35th annual symposium on foundations of computer science*. Ieee. 1994, pp. 124–134.

[VH14] Tim Van Erven and Peter Harremos. "Rényi divergence and Kullback-Leibler divergence". In: *IEEE Transactions on Information Theory* 60.7 (2014), pp. 3797–3820.

[WLF16] Yu-Xiang Wang, Jing Lei, and Stephen E Fienberg. "On-average kl-privacy and its equivalence to generalization for max-entropy mechanisms". In: *Privacy in Statistical Databases: UNESCO Chair in Data Privacy, International Conference, PSD 2016, Dubrovnik, Croatia, September 14–16, 2016, Proceedings*. Springer. 2016, pp. 121–134.

[WCY23] William M Watkins, Samuel Yen-Chi Chen, and Shinjae Yoo. "Quantum machine learning with differential privacy". In: *Scientific Reports* 13.1 (2023), p. 2453.

[XX15] Yonghui Xiao and Li Xiong. "Protecting locations with differential privacy under temporal correlations". In: *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*. 2015, pp. 1298–1309.

[ZY17] Li Zhou and Mingsheng Ying. "Differential privacy in quantum computation". In: *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*. IEEE. 2017, pp. 249–262.