When Does Quantum Differential Privacy Compose?

Daniel Alabi
University of Illinois at Urbana-Champaign, IL, USA

Theshani Nuradha
University of Illinois at Urbana-Champaign, IL, USA

Abstract

Composition is a cornerstone of classical differential privacy, enabling strong end-to-end guarantees for complex algorithms through composition theorems (e.g., basic and advanced). In the quantum setting, however, privacy is defined operationally against arbitrary measurements, and classical composition arguments based on scalar privacy-loss random variables no longer apply. As a result, it has remained unclear when meaningful composition guarantees can be obtained for quantum differential privacy (QDP).

In this work, we clarify both the limitations and possibilities of composition in the quantum setting. We first show that classical-style composition fails in full generality for POVM-based approximate QDP: even quantum channels that are individually perfectly private can completely lose privacy when combined through correlated joint implementations.

We then identify a setting in which clean composition guarantees can be restored. For tensor-product channels acting on product neighboring inputs, we introduce a *quantum moments accountant* based on an operator-valued notion of privacy loss and a matrix moment-generating function. Although the resulting Rényi-type divergence does not satisfy a data-processing inequality, we prove that controlling its moments suffices to bound measured Rényi divergence, yielding operational privacy guarantees against arbitrary measurements. This leads to advanced-composition-style bounds with the same leading-order behavior as in the classical theory.

Our results demonstrate that meaningful composition theorems for quantum differential privacy require carefully articulated structural assumptions on channels, inputs, and adversarial measurements, and provide a principled framework for understanding which classical ideas do and do not extend to the quantum setting.

2012 ACM Subject Classification Security and privacy → Information-theoretic techniques; Theory of computation → Quantum information theory

Keywords and phrases Quantum Information, Differential Privacy

Digital Object Identifier 10.4230/LIPIcs.ITC.2026.10

Related Version *Full Version*: https://arxiv.org/abs/2601.00337

Funding *Daniel Alabi*: Research supported by the Simons Foundation (965342, D.A.) as part of the Junior Fellowship from the Simons Society of Fellows.
Theshani Nuradha: Research supported by the Department of Mathematics and the IQUIST Postdoctoral Fellowship from the Illinois Quantum Information Science and Technology Center at the University of Illinois Urbana-Champaign.

Acknowledgements DA acknowledges helpful discussions with Guy Rothblum and Salil Vadhan regarding the composition of classical and quantum private mechanisms. TN acknowledges helpful discussions with Sujeet Bhalerao, Felix Leditzky, Vishal Singh, and Mark M. Wilde on the composition of quantum private mechanisms.

1 Introduction

Differential privacy (DP) provides a rigorous framework for limiting information leakage about sensitive inputs under randomized data analysis [11, 10]. With the growth of quantum technologies and the generation of quantum data, it is crucial to protect sensitive information related to the data generated and the data that is fed into these developing hybrid classical–quantum systems. While the theory of classical differential privacy is by now mature, extending its guarantees to quantum information processing raises both conceptual and technical challenges [31]. In fully or hybrid (i.e., classical–quantum) quantum processing pipelines, we need a definition that is native to that setting. The outputs of the algorithms and the processing techniques have to obey the laws of quantum mechanics. Classical DP does not capture this because the released object (if any) from quantum algorithms should be a quantum state, and an adversary may apply arbitrary measurements, including collective measurements across multiple outputs, to that state.

With that, quantum DP is defined in a way to ensure privacy for quantum states such that any two quantum states that are classified as neighbors are indistinguishable under all measurements up to a certain limit demanded by privacy parameters [31, 16] (see also Definition 14). In the quantum setting, outputs of quantum systems are quantum states rather than classical samples, adversaries may perform arbitrary measurements, and correlations or entanglement across multiple outputs can fundamentally alter distinguishability properties [22, 19]. These features complicate even basic questions about composition, a cornerstone of the classical DP theory.

In classical DP, composition guarantees are derived by tracking a scalar *privacy-loss random variable* whose moment generating function (MGF) is additive under independent composition [12]. This structure underlies both basic composition and advanced composition results, including the moments accountant framework of Abadi et al. [2]. In contrast, for quantum channels privacy loss is not a scalar quantity prior to measurement, and naively taking a supremum over all measurements destroys additivity since measured quantum divergences are not always additive or at least sub-additive. As a result, classical composition arguments do not directly extend to quantum differential privacy (QDP).

Much of the existing literature on composition in quantum differential privacy has taken an explicitly adversary-centric viewpoint [16, 25]. In this line of work, composition is analyzed primarily through the lens of what an adversary may do to the outputs of multiple private mechanisms, most notably by quantifying over increasingly powerful classes of measurements [25, 26]. While this perspective is natural from an operational security standpoint, it largely abstracts away how quantum systems are actually composed in practice. In quantum information processing, composition is implemented by specific channel constructions (such as tensor-product composition, factorized releases, or correlated joint implementations) whose structural properties can be as consequential for privacy as the adversary's measurement capabilities [1]. By explicitly separating the model of channel composition from the model of adversarial measurement, our work shifts part of the focus from "what measurements are allowed" to "how mechanisms are combined," and shows that many composition phenomena in quantum differential privacy are driven as much by the structure of the joint channel as by the power of the adversary.

This work develops a principled framework for understanding when and how composition guarantees can be recovered in the quantum setting. Our analysis makes two key points. First, classical-style composition fails in full generality for POVM-based approximate QDP when correlated joint channels or entangled neighboring inputs are allowed. Second, under

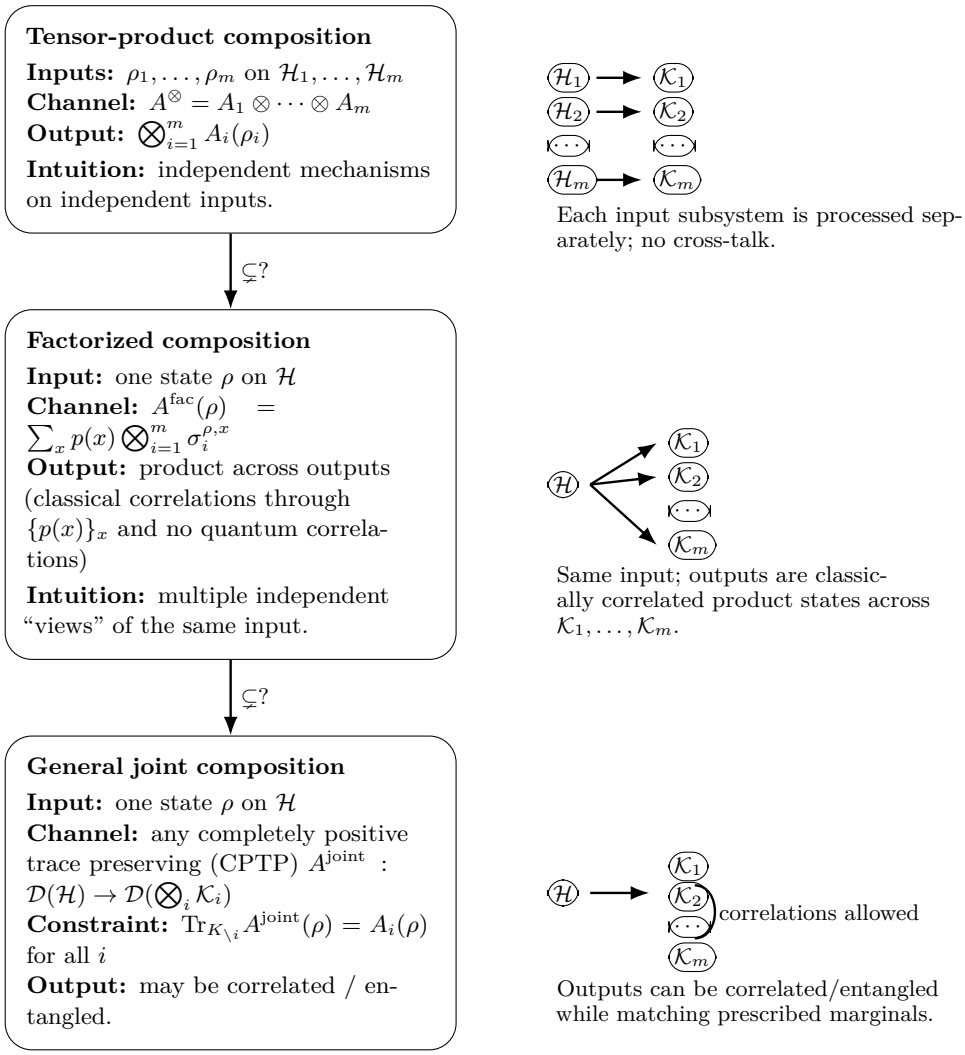

Figure 1 A hierarchy for multi-output composition models.

carefully articulated structural assumptions (most notably tensor-product channels acting on product neighboring inputs) one can recover clean and quantitatively sharp composition guarantees via an operator-level analogue of the classical moments accountant.

To achieve this, we introduce a *quantum moments accountant* (QMA) based on the privacy-loss operator and a matrix moment-generating function. While the resulting Rényi-type divergence does not satisfy a data-processing inequality in general, we show that controlling its moments suffices to bound *measured Rényi divergence*, which by definition captures worst-case distinguishability over all measurements. This yields a direct route from operator-level moment bounds to operational (ε, δ)-QDP guarantees.

Throughout the paper, we emphasize the importance of distinguishing between different models of composition. We propose a strict hierarchy (i.e., tensor-product, factorized $\subsetneq$ general joint composition) and show that failures of classical composition arise precisely when moving beyond tensor-product structure. Our results therefore clarify not only what is possible in quantum composition, but also why additional assumptions are unavoidable.

Contributions

This paper makes the following contributions.

- **A hierarchy of quantum composition models.** We formalize and distinguish tensor-product, factorized, and general joint composition of quantum channels, and show that these form a strict hierarchy. This framework, illustrated in Figure 1, isolates exactly where classical composition guarantees break down in the quantum setting. See Section 2.3.
- **Impossibility of basic composition for general joint channels.** We prove, in Theorem 15, that even pure $(0, 0)$-QDP mechanisms can fail to compose under correlated joint implementations. In particular, we show that, for $\varepsilon_1, \varepsilon_2, \delta_1, \delta_2 \geq 0$ no general $(\varepsilon_1 + \varepsilon_2, \delta_1 + \delta_2)$-type composition theorem can hold for POVM-based approximate QDP without further restrictions. See Section 4 for more details.
- **Safe basic composition under restricted adversaries.** For tensor-product channels on product neighboring inputs, we prove (in Theorem 18) basic composition guarantees against one-way classical communication and local-operations-based (one-way LOCC) adversaries, providing a clean quantum analogue of classical sequential composition under explicitly stated assumptions. See Section 5 for more details.
- **A quantum moments accountant.** In Section 6, we introduce a non-commutative analogue of the classical moments accountant based on the quantum privacy-loss operator and a matrix moment-generating function. We show, via Theorem 27, that this accountant composes additively under tensor-product channels and product neighbors.
- **From operator moments to operational privacy.** We prove (Theorem 30) that bounds on the quantum moments accountant imply bounds on measured Rényi divergence, yielding operational (ε, δ)-QDP guarantees against arbitrary POVMs. This provides an advanced-composition-style bound with the familiar $\sqrt{\sum_i \varepsilon_i^2 \log(1/\delta)}$ scaling over k $(\varepsilon_i, 0)$-QDP mechanisms. See Proposition 28 and Corollary 31.
- **Advanced composition for QDP.** We prove an advanced composition result for tensor-product composition (Theorem 33) of k $(\varepsilon_i, \delta_i)$-QDP mechanisms under all possible adversaries for $\delta_i = 0$, with the composed mechanism satisfying the privacy parameter $\sqrt{\sum_i \varepsilon_i^2 \log(1/\delta)}$ for $\delta \in (0, 1)$.
- **Clarifying the quantum–classical gap.** Our results identify measurement incompatibility and the absence of a joint classical probability space as the fundamental obstacles to classical composition arguments, rather than a failure of tensor-product additivity at the channel level. See Section 2.3.

2 Preliminaries and Definitions

We collect notation, definitions, and structural distinctions that will be used throughout the paper. Because composition behavior in the quantum setting depends sensitively on how channels are combined, how neighboring inputs are defined, and what adversarial measurements are permitted, we make these modeling choices explicit at the outset. In particular, we distinguish several notions of multi-output composition that coincide classically but diverge sharply in the quantum setting.

2.1 Notation and Conventions

All Hilbert spaces are finite-dimensional. For a Hilbert space $\mathcal{H}$, we write $\mathcal{D}(\mathcal{H})$ for the set of density operators on $\mathcal{H}$, i.e., positive semidefinite operators with unit trace.

▶ **Definition 1** (Density operators). *Let $\mathcal{H}$ be a finite-dimensional Hilbert space. We write $\mathcal{D}(\mathcal{H}) := \{\rho \in \mathsf{L}(\mathcal{H}) \mid \rho \succeq 0,\ \mathrm{Tr}(\rho) = 1\}$, where $\mathsf{L}(\mathcal{H})$ denotes the set of linear operators on $\mathcal{H}$. Elements of $\mathcal{D}(\mathcal{H})$ are called* density operators *or* quantum states.

Given a composite system $\mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_m$, we write $\mathrm{Tr}_{\mathcal{H}\setminus i}$ for the partial trace over all subsystems except $\mathcal{H}_i$. For operators X, Y we write $X \preceq Y$ to denote that $Y - X$ is positive semidefinite.

Throughout, logarithms are natural unless otherwise stated. When defining operator expressions involving inverses or logarithms, we adopt the convention that the expression is $+\infty$ whenever the required support conditions fail.

2.2 Channels

Quantum channels are completely positive, trace-preserving (CPTP) maps between spaces of density operators.

▶ **Definition 2** (Tensor-product channels). *Let $\mathcal{H}_1, \mathcal{H}_2, \ldots, \mathcal{H}_m$ and $\mathcal{K}_1, \mathcal{K}_2, \ldots, \mathcal{K}_m$ be finite-dimensional Hilbert spaces. Let $A_i : \mathcal{D}(\mathcal{H}_i) \to \mathcal{D}(\mathcal{K}_i)$, $i \in \{1, 2, \ldots, m\}$, be quantum channels (i.e., completely positive, trace-preserving maps). The* tensor-product channel *associated with $(A_1, \ldots, A_m)$ is the channel $A_1 \otimes \cdots \otimes A_m : \mathcal{D}(\mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_m) \longrightarrow \mathcal{D}(\mathcal{K}_1 \otimes \cdots \otimes \mathcal{K}_m)$ defined by $(A_1 \otimes \cdots \otimes A_m)(\rho_1 \otimes \cdots \otimes \rho_m) := A_1(\rho_1) \otimes \cdots \otimes A_m(\rho_m)$, and extended linearly to all inputs $\rho \in \mathcal{D}(\mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_m)$. Equivalently, $A_1 \otimes \cdots \otimes A_m$ is a CPTP map whose Kraus operators are all tensor products of Kraus operators of $A_1, \ldots, A_m$.*

▶ **Definition 3** (Factorized channels). *Let $\mathcal{H}$ be a finite-dimensional Hilbert space and let $\mathcal{K}_1, \ldots, \mathcal{K}_m$ be finite-dimensional Hilbert spaces. A quantum channel $A : \mathcal{D}(\mathcal{H}) \to \mathcal{D}(\mathcal{K}_1 \otimes \cdots \otimes \mathcal{K}_m)$ is called* factorized *if $A(\rho) = \sum_x p(x) \bigotimes_{i=1}^m \sigma_i^{\rho,x}$, where $\sigma_i^{\rho,x}$ is a quantum state for all i and x, and $\{p(x)\}_x$ is a probability distribution with $p(x) \geq 0$ and $\sum_x p(x) = 1$. In this case, the m outputs are not entangled and independent given the input state.*

Note that a factorized channel can measure the input, copy the classical data, and then prepare quantum states conditioned on the measurement outcome. This has been done in previous work [25].

▶ **Definition 4** (General joint channels). *Let $\mathcal{H}$ be a finite-dimensional Hilbert space and let $\mathcal{K}_1, \ldots, \mathcal{K}_m$ be finite-dimensional Hilbert spaces. A* general joint channel *is any quantum channel $A : \mathcal{D}(\mathcal{H}) \to \mathcal{D}(\mathcal{K}_1 \otimes \cdots \otimes \mathcal{K}_m)$ that is completely positive and trace-preserving, with no further structural restrictions. Equivalently, a general joint channel may produce arbitrary correlations or entanglement among the output subsystems $\mathcal{K}_1, \ldots, \mathcal{K}_m$, and need not admit any factorization or tensor-product decomposition.*

2.3 Models of Composition

Classically, releasing multiple outputs of differentially private mechanisms implicitly defines a joint distribution over all outputs. In the quantum setting, however, there is no unique joint state consistent with a collection of marginals, and different joint implementations can exhibit drastically different privacy behavior. We therefore distinguish three composition models:

1. **Tensor-product composition**, where independent channels act on independent input subsystems.
2. **Factorized composition**, where multiple channels act on the input but the joint output is constrained to be a product state (not entangled).
3. **General joint composition**, where only the marginal behavior of each output is fixed, and the joint output may be arbitrarily correlated or entangled.

While these notions might coincide in the classical setting, they form a hierarchy in the quantum setting. Much of the subtlety of quantum composition arises from the gap between factorized and general joint composition.

▶ **Definition 5** (Composition into tensor-product channels). *Let* $A_i : \mathcal{D}(\mathcal{H}_i) \to \mathcal{D}(\mathcal{K}_i),\ i = 1, \ldots, m$, *be quantum channels acting on (possibly distinct) input systems. The* composition into a tensor-product channel *is the joint channel* $A^{\otimes} := A_1 \otimes \cdots \otimes A_m : \mathcal{D}(\mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_m) \to \mathcal{D}(\mathcal{K}_1 \otimes \cdots \otimes \mathcal{K}_m)$, *defined by* $A^{\otimes}(\rho_1 \otimes \cdots \otimes \rho_m) = A_1(\rho_1) \otimes \cdots \otimes A_m(\rho_m)$, *and extended linearly to all joint inputs. This models the composition of* independent mechanisms acting on independent inputs*.*

Note that tensor-product composition is the quantum analogue of classical parallel composition.

▶ **Definition 6** (Composition into factorized channels). *The* composition into a factorized channel *is the joint channel* $A^{\mathrm{fac}} : \mathcal{D}(\mathcal{H}) \to \mathcal{D}(\mathcal{K}_1 \otimes \cdots \otimes \mathcal{K}_m)$ *defined by* $A^{\mathrm{fac}}(\rho) = \sum_x p(x) \bigotimes_{i=1}^m \sigma_i^{\rho,x}$ *In this composition, all outputs are conditionally independent given the input state (not entangled), but each output may reveal different information about that input.*

▶ **Definition 7** (Composition into general joint channels). *Let* $A_i : \mathcal{D}(\mathcal{H}) \to \mathcal{D}(\mathcal{K}_i),\ i = 1, \ldots, m$, *be quantum channels with a common input space. A* composition into a general joint channel *is any quantum channel* $A^{\mathrm{joint}} : \mathcal{D}(\mathcal{H}) \to \mathcal{D}(\mathcal{K}_1 \otimes \cdots \otimes \mathcal{K}_m)$ *whose marginals coincide with the individual channels, i.e.,* $\mathrm{Tr}_{\mathcal{K}_{\setminus i}}\big(A^{\mathrm{joint}}(\rho)\big) = A_i(\rho)$ *for all* $\rho \in \mathcal{D}(\mathcal{H})$ *and all* i*, where* $\mathcal{K}_{\setminus i} := \bigotimes_{j \neq i} \mathcal{K}_j$*. Such a composition may introduce arbitrary classical or quantum correlations (including entanglement) among the outputs and strictly generalizes both tensor-product and factorized compositions.*

▶ Remark 8 (Tensor-product vs. factorized composition). Tensor-product and factorized composition impose constraints along different axes. Tensor-product composition restricts how a joint channel is *implemented*, requiring independent local mechanisms, but allows correlated or entangled outputs on entangled inputs. In contrast, factorized composition restricts the *output structure*, requiring a product state for every system (possibly with classical correlations), but does not require that the channel decompose as a tensor product of local channels.

On product inputs $\rho = \rho_1 \otimes \cdots \otimes \rho_m$, the output of a tensor-product channel is product: $(A_1 \otimes \cdots \otimes A_m)(\rho) = A_1(\rho_1) \otimes \cdots \otimes A_m(\rho_m)$, but this coincidence does not extend to entangled inputs.

Examples

▶ **Example 9** (Example of tensor-product channel). Let $\mathcal{H}_1 = \mathcal{H}_2 = \mathbb{C}^2$ and $\mathcal{K}_1 = \mathcal{K}_2 = \mathbb{C}^2$. Let A_1 and A_2 be single-qubit depolarizing channels, $A_i(\rho) = (1 - p_i)\rho + p_i \frac{I}{2},\ i \in \{1, 2\}$. Then the tensor-product channel $A_1 \otimes A_2$ acts on two-qubit inputs as $(A_1 \otimes A_2)(\rho_1 \otimes \rho_2) = A_1(\rho_1) \otimes A_2(\rho_2)$, and introduces no correlations between the two output qubits beyond those already present in the input. If the input is a product state, the output is also a product state.

▶ **Example 10** (Example of factorized channel). Let $\mathcal{H} = \mathbb{C}^2$ and let $\mathcal{K}_i = \mathbb{C}^2$ for $i \in \{1, \ldots, m\}$. Consider measure and prepare channel A that measures the state ρ with the POVM $\{M_x\}_x$ and then depending on the classical outcome, prepare state $\otimes_{i=1}^m \sigma_i^x$. Then the joint channel $\mathcal{A}(\rho) = \sum_x \mathrm{Tr}(M_x \rho)\ \sigma_1^x \otimes \cdots \otimes \sigma_m^x$ is a valid factorized channel.

▶ **Example 11** (Example of general joint channels). Let $\mathcal{H} = \mathbb{C}^2$ and $\mathcal{K}_1 = \mathcal{K}_2 = \mathbb{C}^2$. Define a channel A that appends an ancilla qubit initialized to $|0\rangle$ and applies a CNOT gate:

$$A(\rho) = \mathrm{CNOT}\,(\rho \otimes |0\rangle\langle 0|)\,\mathrm{CNOT}^\dagger.$$

The resulting joint output state on $\mathcal{K}_1 \otimes \mathcal{K}_2$ is generally entangled. This channel is a valid joint channel but is neither tensor-product nor factorized, since the two output subsystems are correlated even when the input is pure.

▶ Remark 12 (Distinguishing the three notions). Tensor-product channels describe independent mechanisms acting on independent inputs. Factorized channels describe multiple independent outputs derived from the same input. General joint channels allow arbitrary correlations or entanglement across outputs. These distinctions are crucial when reasoning about composition and privacy guarantees in the quantum setting.

▶ **Proposition 13** (Containment relations between composition models). *Let $m \geq 2$ and let $\mathcal{H}, \mathcal{H}_1, \ldots, \mathcal{H}_m$ and $\mathcal{K}_1, \ldots, \mathcal{K}_m$ be finite-dimensional Hilbert spaces.*

1. ***Tensor-product implies factorization on product inputs.*** *Let $A^{\otimes} := A_1 \otimes \cdots \otimes A_m$ be a tensor-product composition of channels $A_i : \mathcal{D}(H_i) \to \mathcal{D}(K_i)$. Then for every product input state$\rho = \rho_1 \otimes \cdots \otimes \rho_m$, the output of $A^{\otimes}$ is factorized: $A^{\otimes}(\rho) = A_1(\rho_1) \otimes \cdots \otimes A_m(\rho_m)$. Equivalently, when restricted to product inputs, a tensor-product channel coincides with a factorized channel (notice that in this case we set the probability distribution to a singleton).*
2. ***Factorized $\subseteq$ general joint.*** *Every factorized composition is a special case of general joint composition. Indeed, if $A^{\mathrm{fac}}(\rho) = \sum_x p(x) \bigotimes_{i=1}^m \sigma_i^{\rho,x}$, then A^{fac} is a joint channel whose marginals are exactly the $A_i(\rho) = \sum_x p(x)\sigma_i^{x,\rho}$.*
3. ***Strictness of inclusions.*** *Both inclusions above are strict:*
   - *There exist factorized channels that are not tensor-product channels (e.g., multiple outputs that are independent and possibly only classically correlated for a given input state).*
   - *There exist general joint channels that are not factorized (e.g., channels that produce entangled outputs while preserving the same marginals).*

Proof. Follows from Definition 5, Definition 6, Definition 7 and Example 9,Example 10,Example 11. ◀

2.4 Operational Quantum Differential Privacy

In the classical setting, differential privacy is defined in terms of the output distributions of randomized algorithms. In the quantum setting, the output of a mechanism is a quantum state, and privacy must therefore be defined operationally in terms of the statistics produced by measurements. This leads naturally to a definition that quantifies over all POVM effects applied to the channel output [31, 16]:

▶ **Definition 14** ((Approximate) Quantum Differential Privacy). *A quantum channel (CPTP map) A is (ε, δ)-*QDP *if for all neighboring input states $\rho \sim \sigma$ and for every measurement operator M (i.e., satisfying $0 \preceq M \preceq I$),*

$$\mathrm{Tr}\big(MA(\rho)\big) \;\leq\; e^{\varepsilon}\mathrm{Tr}\big(MA(\sigma)\big) \;+\; \delta. \tag{1}$$

Importantly, this operational definition captures adversaries with unrestricted measurement power. While alternative notions of quantum privacy can be defined (e.g., by restricting measurements or comparing states directly via trace distance [25]), the general quantum differential privacy definition most directly mirrors the classical adversarial model and is the strongest notion considered in this work.

Neighboring Inputs

We use the standard notion of neighboring inputs adapted to quantum states. Two density operators $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ are said to be *neighbors*, denoted $\rho \sim \sigma$, if they differ in the data of a single individual or record. The precise definition is application-dependent and left abstract in this work; all results hold for any fixed neighboring relation. When considering multi-system inputs, we distinguish between:

- *Product neighbors*, where $\rho = \bigotimes_i \rho_i$ and $\sigma = \bigotimes_i \sigma_i$ with $\rho_i \sim \sigma_i$ for each subsystem; and
- *General neighbors*, which may be entangled or classically correlated across subsystems.

This distinction plays a critical role in our composition results and impossibility theorems.

2.5 Roadmap

The distinctions introduced above are not merely definitional. In Section 4, we show that basic composition fails for approximate QDP under general joint composition. In contrast, Sections 5 and 6 show that clean composition guarantees can be recovered for tensor-product channels acting on product neighboring inputs, culminating in an advanced-composition-style bound via the quantum moments accountant. First, we discuss some additional related work.

3 Related Work

Classical differential privacy and composition

Composition theorems are central to classical differential privacy [17]. Basic composition and advanced composition results were established early in the literature, culminating in the moments accountant framework of Abadi et al. [2] for tracking privacy loss in iterative algorithms such as DP-SGD. Rényi differential privacy (RDP) and its variants further systematized composition by working directly with Rényi divergences [20].

Quantum differential privacy

Quantum differential privacy (QDP) was first introduced in [31], and several variants of QDP considering distinguishability under all possible measurements (adversaries) have been studied in [1, 16, 3, 25, 4, 15, 14]. Furthermore, a variant that generalizes QDP by encoding domain knowledge and measurement capabilities (practical possibilities of the adversary) of the systems, known as quantum pufferfish privacy (QPP), was introduced in [25] and further studied in [26] with the inspiration from its classical variants [18, 24]. This variant highlights, in some cases, how certain mechanisms are private when we consider the limitations of the measurements that can be performed, in contrast to those mechanisms being non-private when all measurements are allowed. Prior works have explored relationships between QDP mechanisms, quantum hypothesis testing, and quantum Rényi divergences, as well as connections to classical DP under commuting states [16, 25, 13, 4, 27, 6, 7].

In terms of composition of private mechanisms, for the variant of quantum local differential privacy (Definition 14 for $\delta = 0$ and neighbors declared as $\rho \sim \sigma$ for all pairs of distinct states), (basic) composition of several such private mechanisms when allowed measurements

are having locally motivated structures has been studied in [15], and in [16, 25] for QDP and QPP, some basic composition results were provided. In the setting where one applies private channels one after the other sequentially, sequential composition results for QDP have been studied using strong data-processing inequalities [16, 23]. However, more generally, composition guarantees in the quantum setting remain far less understood in a unified way, particularly in the presence of entanglement and correlated outputs.

Rényi divergences in quantum information

Quantum Rényi divergences, including the Petz and sandwiched variants [21, 5, 30], play a central role in quantum information theory [9]. The sandwiched Rényi divergence satisfies a data-processing inequality and has been used to derive operational guarantees in a variety of settings. In contrast, the MMGF-induced divergence we consider is tailored to moment accounting and exact additivity, rather than monotonicity under channels.

Limits of composition in non-classical settings

Impossibility results related to composition have appeared in other non-classical or non-i.i.d. settings, where joint distributions or correlated releases invalidate union-bound-based arguments [17]. Our no-go results for general joint quantum channels fit squarely into this theme and provide a concrete quantum-mechanical explanation rooted in measurement incompatibility.

Our contribution in context

Relative to prior work, this paper provides the first systematic treatment of when classical-style composition can and cannot be recovered for quantum differential privacy. By explicitly separating structural assumptions on channels, neighboring inputs, and adversaries, we reconcile negative results with positive composition theorems and introduce a moments-accountant-style framework that is both quantum-native and operationally meaningful.

4 Why Basic Composition for Approximate DP Fails for General Joint Channels

Basic composition (classical)

In the classical setting, if M_1 is $(\varepsilon_1, \delta_1)$-DP and M_2 is $(\varepsilon_2, \delta_2)$-DP, then the joint release (M_1, M_2) is $(\varepsilon_1 + \varepsilon_2, \delta_1 + \delta_2)$-DP. The proof relies on representing privacy loss as a scalar random variable and applying a union bound to the "δ-bad" events.

4.1 Failure of Basic Composition for Quantum DP

We now show that the classical basic composition theorem does not extend to Definition 14.

▶ **Theorem 15** (No-go for basic composition under POVM-QDP). *There exist quantum channels A_1, A_2 and neighboring inputs $\rho \sim \sigma$ such that:*

1. *each of A_1 and A_2 is $(0, 0)$-QDP, yet*
2. *there exists a composition of A_1, A_2 into a joint channel (Definition 7) that is not (ε, δ)-QDP for any $\delta < 1$.*

Consequently, no general rule of the form

$$(\varepsilon_1, \delta_1)\text{-}QDP + (\varepsilon_2, \delta_2)\text{-}QDP \Longrightarrow (\varepsilon_1 + \varepsilon_2, \delta_1 + \delta_2)\text{-}QDP$$

can hold for POVM-based quantum differential privacy.

Proof. Let the database be a single bit $b \in \{0, 1\}$ encoded by orthogonal states $\rho_0 = |0\rangle\langle 0|$ and $\rho_1 = |1\rangle\langle 1|$, which we declare neighboring. Define a joint channel A_{12} with two-qubit output systems A and B by

$$A_{12}(\rho_0) = |\Phi^+\rangle\langle\Phi^+|, \qquad A_{12}(\rho_1) = |\Phi^-\rangle\langle\Phi^-|,$$

where $|\Phi^\pm\rangle := \frac{1}{\sqrt{2}}(|00\rangle \pm |11\rangle)$ are Bell states.

Define the individual mechanisms by taking marginals: $A_1 := \mathrm{Tr}_B \circ A_{12},\ A_2 := \mathrm{Tr}_A \circ A_{12}$. Since $\mathrm{Tr}_B(|\Phi^+\rangle\langle\Phi^+|) = \mathrm{Tr}_B(|\Phi^-\rangle\langle\Phi^-|) = \frac{I}{2}$, we have $A_1(\rho_0) = A_1(\rho_1)$, and similarly $A_2(\rho_0) = A_2(\rho_1)$. Hence for every POVM effect M, $\mathrm{Tr}(MA_i(\rho_0)) = \mathrm{Tr}(MA_i(\rho_1))$, $i \in \{1, 2\}$, so both A_1 and A_2 are $(0, 0)$-QDP. However, the joint outputs $A_{12}(\rho_0)$ and $A_{12}(\rho_1)$ are orthogonal. Let $M = |\Phi^+\rangle\langle\Phi^+|$. Then $\mathrm{Tr}(MA_{12}(\rho_0)) = 1,\ \mathrm{Tr}(MA_{12}(\rho_1)) = 0$. If A_{12} were (ε, δ)-QDP with $\delta < 1$, Definition 14 would imply $1 \leq \delta$, a contradiction. ◀

▶ Remark 16 (Clarification about Theorem 15). The joint channel refers to the correlated implementation (see Definition 7) given by the single joint channel A_{12} outputting both subsystems, not to the independent tensor product channel $A_1 \otimes A_2$ (see Definition 5).

Source of the composition failure

The failure of composition in Theorem 15 does *not* arise from entangled neighboring inputs. The neighboring states $\rho_0 = |0\rangle\langle 0|$ and $\rho_1 = |1\rangle\langle 1|$ are single-qubit, orthogonal, and unentangled. Rather, the pathology arises from allowing a general joint channel whose marginals are fixed but whose joint action introduces correlations between outputs. Although each marginal channel is perfectly private, the joint channel can amplify distinguishability through correlated outputs.

4.2 Why This Phenomenon Is Quantum

Classical approximate DP relies on two properties:

1. *Existence of a single joint probability space* supporting all outcomes and all privacy-loss events.
2. *Scalar privacy loss*: the likelihood ratio is a random variable, and δ controls the probability of large deviations via a union bound.

In contrast, Definition 14 quantifies over all POVMs *after* the channel. Different POVMs are generally incompatible and cannot be realized jointly. As a result, there is no single outcome space on which "bad events" can be union-bounded. This lack of joint measurability is the same structural feature responsible for Bell/CHSH violations in quantum mechanics [22, 19].

Basic composition for approximate DP fundamentally relies on classical probabilistic structure. Under POVM-based quantum differential privacy, this structure is absent: local indistinguishability does not imply joint indistinguishability. Consequently, classical $(\varepsilon_1 + \varepsilon_2, \delta_1 + \delta_2)$ composition has no general quantum analogue without additional assumptions.

The impossibility of basic $(\varepsilon_1 + \varepsilon_2, \delta_1 + \delta_2)$ composition for approximate QDP stems from measurement incompatibility and the absence of a joint classical probability space, rather than from a failure of tensor-product additivity at the channel level.

5 Basic Composition for Product Neighbors and Tensor-Product Channels

5.1 Setup

Let $\mathcal{H}_1, \mathcal{H}_2$ be finite-dimensional Hilbert spaces. For $i \in \{1, 2\}$, let $A_i : \mathcal{D}(\mathcal{H}_i) \to \mathcal{D}(\mathcal{K}_i)$ be quantum channels (CPTP maps), and define the tensor-product channel composition (as in Definition 5) $A_{12} := A_1 \otimes A_2 : \mathcal{D}(\mathcal{H}_1 \otimes \mathcal{H}_2) \to \mathcal{D}(\mathcal{K}_1 \otimes \mathcal{K}_2)$.

We use the POVM-based approximate quantum differential privacy definition (Definition 14).

Product-neighbor model

We say that (ρ, σ) are *product neighbors* if $\rho = \rho_1 \otimes \rho_2$, $\sigma = \sigma_1 \otimes \sigma_2$, and $\rho_i \sim \sigma_i$ for each $i \in \{1, 2\}$.

5.2 A Composition Theorem

The clean "δ-adds" basic composition proof in the classical setting relies on the fact that a joint event can be decomposed into conditional events on each release (equivalently, one can implement any test sequentially without changing the underlying probability space). In the quantum setting, this exact argument goes through *verbatim* provided the adversary's test on the joint output is *separable*. In particular, any local operations and classical communications (LOCC) test is separable. For *arbitrary* global POVMs, the corresponding statement is not generally known to hold under the POVM-based definition, and the classical proof technique does not directly apply.

▶ **Definition 17** (One-way LOCC two-outcome tests). *A two-outcome test on $K_1 \otimes K_2$ is one-way LOCC ($K_1 \to K_2$) if it can be implemented as:*

- *measure K_1 with a POVM $\{E_t\}_t$,*
- *on outcome t, measure K_2 with the POVM $\{M_{2,t}, I - M_{2,t}\}$ such that $0 \preceq M_{2,t} \preceq I$ and accept iff it accepts.*

Equivalently, its acceptance probability on product states satisfies $\mathrm{Tr}(M\, \xi_1 \otimes \xi_2) = \sum_t \mathrm{Tr}(E_t \xi_1)\mathrm{Tr}(M_{2,t} \xi_2)$.

▶ **Theorem 18** (Tensor-product composition on product neighbors against one-way LOCC tests). *Let $A_i : \mathcal{D}(H_i) \to \mathcal{D}(K_i)$ be quantum channels for $i \in \{1, 2\}$. Assume that A_i is $(\varepsilon_i, \delta_i)$-QDP: for all neighboring $\rho_i \sim \sigma_i$ and all measurement operators $0 \preceq M_i \preceq I$ on K_i,*

$$\mathrm{Tr}(M_i A_i(\rho_i)) \le e^{\varepsilon_i} \mathrm{Tr}(M_i A_i(\sigma_i)) + \delta_i.$$

Fix any product *neighboring pair $\rho = \rho_1 \otimes \rho_2$, $\sigma = \sigma_1 \otimes \sigma_2$, $\rho_i \sim \sigma_i$. Let $A_{12} := A_1 \otimes A_2$. Then for every one-way LOCC ($K_1 \to K_2$) two-outcome test on $K_1 \otimes K_2$ in the sense of Definition 17 (equivalently, for every measurement operator $0 \preceq M \preceq I$ admitting the factorization in (3) below), we have*

$$\mathrm{Tr}\big(M\, A_{12}(\rho)\big) \;\le\; e^{\varepsilon_1+\varepsilon_2} \mathrm{Tr}\big(M\, A_{12}(\sigma)\big) \;+\; e^{\varepsilon_2}\delta_1 \;+\; \delta_2. \tag{2}$$

By swapping the roles of the two subsystems (i.e., using one-way LOCC tests $K_2 \to K_1$), we also obtain the symmetric bound $\mathrm{Tr}\big(M\, A_{12}(\rho)\big) \;\le\; e^{\varepsilon_1+\varepsilon_2} \mathrm{Tr}\big(M\, A_{12}(\sigma)\big) \;+\; \delta_1 \;+\; e^{\varepsilon_1}\delta_2$. *Consequently, against one-way LOCC adversaries (in either direction) the tensor-product channel $A_1 \otimes A_2$ is $(\varepsilon_1 + \varepsilon_2, \delta_{\mathrm{comp}})$-QDP on product neighbors for $\delta_{\mathrm{comp}} := \min\{\, e^{\varepsilon_2}\delta_1 + \delta_2,\; \delta_1 + e^{\varepsilon_1}\delta_2 \,\}$.*

Proof. Let $\omega_i := A_i(\rho_i)$ and $\eta_i := A_i(\sigma_i)$ for $i \in \{1,2\}$. Then $A_{12}(\rho) = \omega_1 \otimes \omega_2$ and $A_{12}(\sigma) = \eta_1 \otimes \eta_2$.

Fix a one-way LOCC $(K_1 \to K_2)$ two-outcome test with acceptance operator $0 \preceq M \preceq I$. By Definition 17, there exist a POVM $\{E_t\}_t$ on K_1 (so $E_t \succeq 0$ and $\sum_t E_t = I$) and measurement operator $0 \preceq M_{2,t} \preceq I$ on K_2 such that for all product states $\xi_1 \otimes \xi_2$,

$$\mathrm{Tr}\big(M(\xi_1 \otimes \xi_2)\big) \;=\; \sum_t \mathrm{Tr}(E_t \xi_1)\,\mathrm{Tr}(M_{2,t}\xi_2). \tag{3}$$

Applying (3) to $\omega_1 \otimes \omega_2$ gives $\mathrm{Tr}\big(M(\omega_1 \otimes \omega_2)\big) = \sum_t \mathrm{Tr}(E_t\omega_1)\,\mathrm{Tr}(M_{2,t}\omega_2)$. For each t, since $0 \preceq M_{2,t} \preceq I$ and A_2 is $(\varepsilon_2, \delta_2)$-QDP, we have $\mathrm{Tr}(M_{2,t}\omega_2) \le e^{\varepsilon_2}\mathrm{Tr}(M_{2,t}\eta_2) + \delta_2$. Multiplying by $\mathrm{Tr}(E_t\omega_1) \ge 0$ and summing over t yields

$$\begin{aligned}\mathrm{Tr}\big(M(\omega_1 \otimes \omega_2)\big) &\le \sum_t \mathrm{Tr}(E_t\omega_1)\Big(e^{\varepsilon_2}\mathrm{Tr}(M_{2,t}\eta_2) + \delta_2\Big) \\ &= e^{\varepsilon_2}\sum_t \mathrm{Tr}(E_t\omega_1)\mathrm{Tr}(M_{2,t}\eta_2) \;+\; \delta_2 \sum_t \mathrm{Tr}(E_t\omega_1).\end{aligned} \tag{4}$$

Since $\sum_t E_t = I$ and $\mathrm{Tr}(\omega_1) = 1$, we have $\sum_t \mathrm{Tr}(E_t\omega_1) = 1$, so

$$\mathrm{Tr}\big(M(\omega_1 \otimes \omega_2)\big) \le e^{\varepsilon_2}\sum_t \mathrm{Tr}(E_t\omega_1)\mathrm{Tr}(M_{2,t}\eta_2) \;+\; \delta_2. \tag{5}$$

Define the scalars $s_t := \mathrm{Tr}(M_{2,t}\eta_2) \in [0,1]$ and define the operator $N := \sum_t s_t E_t$. Because $0 \le s_t \le 1$ and $\sum_t E_t = I$, it follows that $0 \le N \le I$. Moreover,
$\sum_t \mathrm{Tr}(E_t\omega_1)\mathrm{Tr}(M_{2,t}\eta_2) = \sum_t \mathrm{Tr}(E_t\omega_1)\, s_t = \mathrm{Tr}(N\omega_1)$. Substituting into (5) gives

$$\mathrm{Tr}\big(M(\omega_1 \otimes \omega_2)\big) \le e^{\varepsilon_2}\mathrm{Tr}(N\omega_1) + \delta_2. \tag{6}$$

Now apply $(\varepsilon_1, \delta_1)$-QDP for A_1 to the measurement operator N (valid since $0 \le N \le I$): $\mathrm{Tr}(N\omega_1) \le e^{\varepsilon_1}\mathrm{Tr}(N\eta_1) + \delta_1$. Plugging into (6) yields

$$\mathrm{Tr}\big(M(\omega_1 \otimes \omega_2)\big) \le e^{\varepsilon_2}\Big(e^{\varepsilon_1}\mathrm{Tr}(N\eta_1) + \delta_1\Big) + \delta_2 = e^{\varepsilon_1+\varepsilon_2}\mathrm{Tr}(N\eta_1) + e^{\varepsilon_2}\delta_1 + \delta_2. \tag{7}$$

Finally, apply (3) to $\eta_1 \otimes \eta_2$:
$\mathrm{Tr}\big(M(\eta_1 \otimes \eta_2)\big) = \sum_t \mathrm{Tr}(E_t\eta_1)\mathrm{Tr}(M_{2,t}\eta_2) = \sum_t \mathrm{Tr}(E_t\eta_1)\, s_t = \mathrm{Tr}(N\eta_1)$. Substituting $\mathrm{Tr}(N\eta_1) = \mathrm{Tr}(M(\eta_1 \otimes \eta_2))$ into (7) gives $\mathrm{Tr}\big(M(\omega_1 \otimes \omega_2)\big) \le e^{\varepsilon_1+\varepsilon_2}\mathrm{Tr}\big(M(\eta_1 \otimes \eta_2)\big) + e^{\varepsilon_2}\delta_1 + \delta_2$, which is exactly (2). The symmetric claim follows by repeating the same argument with the roles of the subsystems swapped. ◀

▶ Remark 19. Theorem 18 avoids any appeal to separable decompositions of M. It uses only the *defining* sequential factorization of one-way LOCC tests and the one-shot POVM-based QDP inequalities for each channel. This is the direct quantum analogue of the classical sequential/interactive proof of basic composition.

▶ **Corollary 20** (Small-ε simplification). *In the setting of Theorem 18, let* $\varepsilon_{\max} := \max\{\varepsilon_1, \varepsilon_2\}$. *Then*

$$\delta_{\mathrm{comp}} = \min\{\, e^{\varepsilon_2}\delta_1 + \delta_2,\; \delta_1 + e^{\varepsilon_1}\delta_2 \,\} \le e^{\varepsilon_{\max}}(\delta_1 + \delta_2).$$

In particular, if $\varepsilon_{\max} \le 1$, *then using* $e^x \le 1 + 2x$ *for all* $x \in [0,1]$, $\delta_{\mathrm{comp}} \;\le\; (1 + 2\varepsilon_{\max})(\delta_1 + \delta_2)$.

Proof. Since $e^{\varepsilon_2} \le e^{\varepsilon_{\max}}$ and $e^{\varepsilon_1} \le e^{\varepsilon_{\max}}$,

$$\delta_{\text{comp}} = \min\{\, e^{\varepsilon_2}\delta_1 + \delta_2,\ \delta_1 + e^{\varepsilon_1}\delta_2 \,\} \le e^{\varepsilon_{\max}}\delta_1 + \delta_2 \le e^{\varepsilon_{\max}}(\delta_1 + \delta_2),$$

and similarly using the other order. This proves the first inequality.

If $\varepsilon_{\max} \le 1$, the elementary bound $e^x \le 1 + 2x$ on $[0, 1]$ gives $e^{\varepsilon_{\max}} \le 1 + 2\varepsilon_{\max}$, hence $\delta_{\text{comp}} \le e^{\varepsilon_{\max}}(\delta_1 + \delta_2) \le (1 + 2\varepsilon_{\max})(\delta_1 + \delta_2)$. ◀

▶ Remark 21 (Relation to previous works). One can also arrive at Theorem 18 by using the general result for all joint measurements for the QDP setting in [16, Corollary III.3] using syb-additivity properties of hockey-stick divergence. Theorem 18 can also be obtained by invoking results in the Quantum Pufferfish Framework (i.e., defining potential secrets, discriminative pairs, data distributions, and measurements). Specifically, by using quasi-subadditivity of the Datta-Leditzky information spectrum divergence [8] in Proposition 2 of [25], the same result in Theorem 18 can be obtained. In this work, we provide an alternative proof that works for one-way LOCC measurements.

In this section, we obtained composition guarantees when the adversary is allowed to perform one-way LOCC two-outcome tests as in Definition 17. This can be understood as the composed mechanism satisfying a flexible private variant of QDP (also a special case of quantum pufferfish privacy in [25]). In particular, Theorem 18 shows that the composed mechanism satisfies quantum differential privacy with the neighbors defined as $\rho_1 \otimes \rho_2 \sim \sigma_1 \otimes \sigma_2,\ \rho_i \sim \sigma_i$ and $\mathcal{M}$ having measurement operators corresponding to two-outcome tests belonging to the category of one-way LOCC in Definition 17.

6 Quantum Moments Accountant

In the classical setting, the *moments accountant* tracks the log-moment generating function (MGF) of the privacy-loss random variable and exploits its additivity under independent composition. In the quantum setting, the main obstruction is that privacy loss cannot be represented as a scalar random variable prior to measurement, and taking a supremum over all POVMs destroys additivity. Nevertheless, we show that a natural non-commutative analogue of the MGF (i.e., defined using the privacy-loss operator and a Petz/exponential Rényi divergence induced by the MMGF) restores exact additivity. This yields a clean and composable quantum version of the classical moments accountant over pure DP channels.

The results we provide here are over tensor-product channels.

6.1 Quantum Privacy-Loss Operator

For neighboring states $\rho \sim \sigma$, define the *quantum privacy-loss operator*

$$L(\rho, \sigma) \coloneqq \log\big(\sigma^{-1/2}\rho\,\sigma^{-1/2}\big), \tag{8}$$

which is well defined on the support of σ. For a quantum channel A, we write $L_A(\rho, \sigma) \coloneqq L(A(\rho), A(\sigma))$. This operator is a natural non-commutative analogue of the classical privacy loss $\log \frac{P(o)}{Q(o)}$ for classical probability measures P, Q and outcome/event o.

Support convention

Throughout, when defining $L(\rho, \sigma) = \log\big(\sigma^{-1/2}\rho\,\sigma^{-1/2}\big)$, we implicitly restrict to pairs (ρ, σ) satisfying $\mathsf{supp}(\rho) \subseteq \mathsf{supp}(\sigma)$. If this condition fails, we define the corresponding privacy-loss operator and MMGF to be $+\infty$. Consequently, $\alpha_A(\lambda)$ may take the value $+\infty$ for some channels.

6.2 Matrix Moment-Generating Function

For $\lambda > 0$, quantum channel A, and states $\rho \sim \sigma$, we define the *matrix moment-generating function* (MMGF) of the privacy-loss operator by

$$\mathrm{MMGF}_A(\lambda; \rho, \sigma) \coloneqq \mathrm{Tr}\Big[A(\sigma)^{1/2}\, \exp\big(\lambda L_A(\rho,\sigma)\big)\, A(\sigma)^{1/2}\Big]. \tag{9}$$

Since $e^{\lambda L_A(\rho,\sigma)} = (A(\sigma)^{-1/2}A(\rho)A(\sigma)^{-1/2})^{\lambda}$, we can rewrite this as

$$\mathrm{MMGF}_A(\lambda; \rho, \sigma) = \mathrm{Tr}\Big[A(\sigma)\,\big(A(\sigma)^{-1/2}A(\rho)A(\sigma)^{-1/2}\big)^{\lambda}\Big]. \tag{10}$$

6.3 Quantum Moments Accountant

We now define the quantum analogue of the classical moments accountant.

▶ **Definition 22** (Quantum Moments Accountant)**.** *For a quantum channel A and $\lambda > 0$, define the quantum moments accountant (QMA) by $\alpha_A(\lambda) \coloneqq \sup_{\rho\sim\sigma} \log \mathrm{MMGF}_A(\lambda;\rho,\sigma)$, where $\mathrm{MMGF}_A(\lambda;\rho,\sigma) \coloneqq \mathrm{Tr}\big[A(\sigma)^{1/2}\exp\big(\lambda L_A(\rho,\sigma)\big)A(\sigma)^{1/2}\big]$, $L_A(\rho,\sigma) \coloneqq \log\big(A(\sigma)^{-1/2}A(\rho)A(\sigma)^{-1/2}\big)$.*

As we will discuss, Definition 22 induces an exponential Rényi-type divergence that composes additively under tensor products:

▶ **Definition 23** (Petz–Rényi divergence [28, 29])**.** *Let ρ and σ be density operators on a finite-dimensional Hilbert space $\mathcal{H}$, and let $\alpha \in (0,1)\cup(1,\infty)$. The* Petz–Rényi divergence *of order α is defined as $D_\alpha^{\mathrm{Petz}}(\rho\|\sigma) \coloneqq \frac{1}{\alpha-1}\log \mathrm{Tr}\big[\rho^\alpha\,\sigma^{1-\alpha}\big]$, with the convention that $D_\alpha^{\mathrm{Petz}}(\rho\|\sigma) = +\infty$ if $\mathrm{supp}(\rho) \not\subseteq \mathrm{supp}(\sigma)$.*

Relation to Rényi-type divergences induced by the MMGF

The quantum moments accountant is defined via the log moment generating function

$$\log \mathrm{MMGF}_A(\lambda;\rho,\sigma) = \log \mathrm{Tr}\left[A(\sigma)\left(A(\sigma)^{-1/2}A(\rho)A(\sigma)^{-1/2}\right)^{\lambda}\right].$$

It is convenient to associate to this quantity an *MMGF-induced* (exponential) Rényi-type divergence:

▶ **Definition 24** (MMGF-induced divergence)**.** *Let ρ and σ be density operators on a finite-dimensional Hilbert space $\mathcal{H}$, and let $\alpha \in (0,1)\cup(1,\infty)$. For $\alpha > 1$, the* MMGF-induced divergence *of order α is defined as*

$$D_\alpha^{\mathrm{MMGF}}(\rho\|\sigma) \;\coloneqq\; \frac{1}{\alpha-1}\log \mathrm{Tr}\left[\sigma\left(\sigma^{-1/2}\rho\sigma^{-1/2}\right)^{\alpha-1}\right], \tag{11}$$

with the convention $D_\alpha^{\mathrm{MMGF}}(\rho\|\sigma) = +\infty$ if $\mathsf{supp}(\rho) \not\subseteq \mathsf{supp}(\sigma)$.

With this notation, for $\alpha = 1+\lambda$ we have the identity
$\frac{1}{\lambda}\log \mathrm{MMGF}_A(\lambda;\rho,\sigma) = D_{1+\lambda}^{\mathrm{MMGF}}\big(A(\rho)\|A(\sigma)\big)$.

While D_α^{MMGF} composes additively under tensor products for product inputs (Section 6.4), it need not satisfy a data-processing inequality in general. To obtain an operational privacy guarantee against *arbitrary* POVMs, we instead show in Theorem 30 that bounds on the MMGF moments imply a bound on *measured* Rényi divergence, which directly captures worst-case distinguishability over all measurements:

▶ **Definition 25** (Measured Rényi divergence). *Let $\alpha > 1$. The* measured Rényi divergence *of order α is $D_\alpha^{\mathrm{meas}}(\rho\|\sigma) := \sup_{M\in\mathrm{POVM}} D_\alpha\big(P_M(\rho)\,\|\,P_M(\sigma)\big)$, where $P_M(\rho)$ denotes the classical outcome distribution induced by measuring ρ with POVM M, and $D_\alpha(\cdot\|\cdot)$ is the classical Rényi divergence of order α.*

6.4 Additivity Under Tensor-Product Composition

Let $A_1, \ldots, A_k$ act on disjoint subsystems, and let $A^{(k)} := \bigotimes_{i=1}^k A_i$. For neighboring product states $\rho = \bigotimes_{i=1}^k \rho_i$ and $\sigma = \bigotimes_{i=1}^k \sigma_i$, we have $L_{A^{(k)}}(\rho, \sigma) = \sum_{i=1}^k L_{A_i}(\rho_i, \sigma_i)$, where each summand acts on its own tensor factor.

Because exponentials of tensor-factor sums factorize, and because $\mathrm{Tr}(X \otimes Y) = \mathrm{Tr}(X)\mathrm{Tr}(Y)$, we obtain the following:

▶ **Lemma 26** (Additivity of the Matrix MGF). *For all $\lambda > 0$,* $\mathrm{MMGF}_{A^{(k)}}(\lambda; \rho, \sigma) = \prod_{i=1}^k \mathrm{MMGF}_{A_i}(\lambda; \rho_i, \sigma_i)$.

Taking logarithms and suprema yields:

▶ **Theorem 27** (Additivity of the quantum moments accountant under product neighbors). *Fix $\lambda > 0$. For a channel $\mathcal{A}$ define the* product-neighbor accountant

$$\alpha_{\mathcal{A}}^{\mathrm{prod}}(\lambda) := \sup_{\substack{\rho=\otimes_{i=1}^k \rho_i,\ \sigma=\otimes_{i=1}^k \sigma_i \\ \rho_i \sim \sigma_i\ \forall i}} \log \mathrm{MMGF}_{\mathcal{A}}(\lambda; \rho, \sigma).$$

Let $\mathcal{A}^{(k)} := \bigotimes_{i=1}^k \mathcal{A}_i$ be a tensor-product channel on disjoint subsystems. Then for all $\lambda > 0$, $\alpha_{\mathcal{A}^{(k)}}^{\mathrm{prod}}(\lambda) = \sum_{i=1}^k \alpha_{\mathcal{A}_i}(\lambda)$. Moreover, for the unrestricted *accountant $\alpha_{\mathcal{A}^{(k)}}(\lambda)$ of Definition 22, with neighbors having the same dimension as $\otimes_{i=1}^k \rho_i$, we always have the lower bound $\alpha_{\mathcal{A}^{(k)}}(\lambda) \geq \alpha_{\mathcal{A}^{(k)}}^{\mathrm{prod}}(\lambda) = \sum_{i=1}^k \alpha_{\mathcal{A}_i}(\lambda)$.*

Proof. Let $\rho = \otimes_i \rho_i$ and $\sigma = \otimes_i \sigma_i$ be product neighbors. Then $\mathcal{A}^{(k)}(\rho) = \otimes_i \mathcal{A}_i(\rho_i)$ and $\mathcal{A}^{(k)}(\sigma) = \otimes_i \mathcal{A}_i(\sigma_i)$. The privacy-loss operators add across tensor factors, hence their exponentials factorize. Using $\mathrm{Tr}(X \otimes Y) = \mathrm{Tr}(X)\mathrm{Tr}(Y)$ yields $\mathrm{MMGF}_{\mathcal{A}^{(k)}}(\lambda; \rho, \sigma) = \prod_{i=1}^k \mathrm{MMGF}_{\mathcal{A}_i}(\lambda; \rho_i, \sigma_i)$, so taking log gives additivity for each fixed product neighbor pair. Taking the supremum over product neighbors yields $\alpha_{\mathcal{A}^{(k)}}^{\mathrm{prod}}(\lambda) = \sum_i \alpha_{\mathcal{A}_i}(\lambda)$. Finally, since the unrestricted supremum is over a larger set, $\alpha_{\mathcal{A}^{(k)}}(\lambda) \geq \alpha_{\mathcal{A}^{(k)}}^{\mathrm{prod}}(\lambda)$. ◀

This mirrors the classical moments accountant exactly, but crucially *without requiring any measurement.*

▶ **Proposition 28** (Advanced composition via the quantum moments accountant (measured Rényi route)). *Let $A_1, \ldots, A_k$ be quantum channels and let $A^{(k)} := A_1 \otimes \cdots \otimes A_k$ denote their tensor-product composition. Consider product neighboring inputs $\rho = \bigotimes_{i=1}^k \rho_i$ and $\sigma = \bigotimes_{i=1}^k \sigma_i$ with $\rho_i \sim \sigma_i$. Fix $\alpha > 1$ and write $\lambda := \alpha - 1$.*

Assume that for each $i \in \{1, \ldots, k\}$ there exist parameters $\varepsilon_i \in (0, 1]$, $c_i \geq 0$, and $\alpha_i > 1$ such that for all $\alpha \in (1, \alpha_i]$ and all $\rho_i \sim \sigma_i$,

$$\log \mathrm{Tr}\Big[A_i(\sigma_i)\Big(A_i(\sigma_i)^{-1/2} A_i(\rho_i) A_i(\sigma_i)^{-1/2}\Big)^{\alpha}\Big] \leq \frac{1}{2}\varepsilon_i^2(\alpha-1)^2 + c_i(\alpha-1)^3. \tag{12}$$

Let

$$S := \sum_{i=1}^k \varepsilon_i^2, \qquad C := \sum_{i=1}^k c_i, \qquad \bar{\alpha} := \min_{i\in[k]} \alpha_i, \qquad \bar{\lambda} := \bar{\alpha} - 1.$$

Then for every $\delta \in (0,1)$ and every $\lambda \in (0, \bar{\lambda}]$ the channel $A^{(k)}$ satisfies $(\varepsilon(\lambda), \delta)$-QDP against arbitrary *POVMs, where*

$$\varepsilon(\lambda) \;\coloneqq\; \frac{S}{2}\lambda \;+\; C\lambda^2 \;+\; \frac{\log(1/\delta)}{\lambda}. \tag{13}$$

In particular, letting

$$\lambda^\star \coloneqq \sqrt{\frac{2\log(1/\delta)}{S}} \quad \textit{and} \quad \hat{\lambda} \coloneqq \min\{\bar{\lambda}, \lambda^\star\},$$

we obtain the explicit bound

$$\varepsilon(\hat{\lambda}) \;\le\; \sqrt{2S\log(1/\delta)} \;+\; \frac{2C\log(1/\delta)}{S} \;+\; \frac{S}{2}\,(\bar{\lambda} - \lambda^\star)_+, \tag{14}$$

where $(x)_+ \coloneqq \max\{x, 0\}$.

Proof. Fix $\delta \in (0,1)$ and $\lambda \in (0, \bar{\lambda}]$, and set $\alpha \coloneqq 1 + \lambda$.

Step 1: Moment additivity under tensor products (product neighbors). For each i define

$$X_i \;\coloneqq\; A_i(\sigma_i)^{-1/2} A_i(\rho_i) A_i(\sigma_i)^{-1/2} \succeq 0.$$

Since $A^{(k)} = \bigotimes_{i=1}^k A_i$ and the neighbors are products, we have $A^{(k)}(\rho) = \bigotimes_{i=1}^k A_i(\rho_i)$, $A^{(k)}(\sigma) = \bigotimes_{i=1}^k A_i(\sigma_i)$, and thus $X \;\coloneqq\; A^{(k)}(\sigma)^{-1/2} A^{(k)}(\rho) A^{(k)}(\sigma)^{-1/2} = \bigotimes_{i=1}^k X_i$. Using $(\bigotimes_i X_i)^\alpha = \bigotimes_i X_i^\alpha$ and multiplicativity of the trace,

$$\operatorname{Tr}\big(A^{(k)}(\sigma)\, X^\alpha\big) = \operatorname{Tr}\left(\bigotimes_{i=1}^k A_i(\sigma_i)\; \bigotimes_{i=1}^k X_i^\alpha\right) = \prod_{i=1}^k \operatorname{Tr}\big(A_i(\sigma_i)\, X_i^\alpha\big). \tag{15}$$

Taking logarithms and applying the bound (12) yields

$$\log \operatorname{Tr}\big(A^{(k)}(\sigma)\, X^\alpha\big) \;\le\; \sum_{i=1}^k \left(\frac{1}{2}\varepsilon_i^2\lambda^2 + c_i\lambda^3\right) = \frac{S}{2}\lambda^2 + C\lambda^3. \tag{16}$$

Step 2: Convert the moment bound to measured Rényi DP. By Theorem 30 applied to $A^{(k)}$ at order $\alpha = 1 + \lambda$, for all product neighbors $\rho \sim \sigma$,

$$D_\alpha^{\mathrm{meas}}\big(A^{(k)}(\rho) \,\|\, A^{(k)}(\sigma)\big) \le \frac{1}{\alpha - 1} \log \operatorname{Tr}\big(A^{(k)}(\sigma)\, X^\alpha\big) = \frac{1}{\lambda} \log \operatorname{Tr}\big(A^{(k)}(\sigma)\, X^\alpha\big).$$

Combining with (16) gives

$$D_{1+\lambda}^{\mathrm{meas}}\big(A^{(k)}(\rho) \,\|\, A^{(k)}(\sigma)\big) \;\le\; \frac{1}{\lambda}\left(\frac{S}{2}\lambda^2 + C\lambda^3\right) = \frac{S}{2}\lambda + C\lambda^2. \tag{17}$$

Equivalently, $A^{(k)}$ satisfies $(\alpha, \varepsilon_\alpha)$-measured Rényi DP with $\varepsilon_\alpha = \frac{S}{2}\lambda + C\lambda^2$.

Step 3: Convert measured Rényi DP to (ε, δ)-QDP. Fix an arbitrary POVM M. By definition of measured Rényi divergence, (17) implies that the induced classical distributions $P_M(A^{(k)}(\rho))$ and $P_M(A^{(k)}(\sigma))$ satisfy
$D_\alpha\big(P_M(A^{(k)}(\rho)) \,\|\, P_M(A^{(k)}(\sigma))\big) \le \varepsilon_\alpha$. Applying the standard (classical) conversion from Rényi DP to approximate DP yields that for all $\delta \in (0,1)$,

$$\Pr[M(A^{(k)}(\rho)) = 1] \le \exp\left(\varepsilon_\alpha + \frac{\log(1/\delta)}{\alpha - 1}\right) \Pr[M(A^{(k)}(\sigma)) = 1] + \delta.$$

Since $\alpha - 1 = \lambda$, we obtain $(\varepsilon(\lambda), \delta)$-QDP with $\varepsilon(\lambda) = \varepsilon_\alpha + \frac{\log(1/\delta)}{\lambda} = \big(\frac{S}{2}\lambda + C\lambda^2\big) + \frac{\log(1/\delta)}{\lambda}$, which is exactly (13).

Step 4 (Optimization). For fixed $\delta \in (0,1)$, the conversion yields the bound $\varepsilon(\lambda) := \frac{\lambda}{2} S + C\lambda^2 + \frac{\log(1/\delta)}{\lambda}$, valid for $\lambda \in (0, \bar{\lambda}]$, where $S := \sum_{i=1}^{k} \varepsilon_i^2$. Ignoring the constraint $\lambda \le \bar{\lambda}$ for a moment, the function $g(\lambda) := \frac{\lambda}{2} S + \frac{\log(1/\delta)}{\lambda}$ is minimized over $\lambda > 0$ at $\lambda^\star := \sqrt{\frac{2\log(1/\delta)}{S}}$. We therefore choose the feasible parameter $\hat{\lambda} := \min\{\lambda^\star, \bar{\lambda}\}$. Substituting $\lambda = \hat{\lambda}$ gives the valid bound

$$\varepsilon \le \frac{\hat{\lambda}}{2} S + C\hat{\lambda}^2 + \frac{\log(1/\delta)}{\hat{\lambda}}. \tag{18}$$

In particular, if $\lambda^\star \le \bar{\lambda}$ then $\hat{\lambda} = \lambda^\star$ and $\varepsilon(\lambda^\star) = \sqrt{2S\log(1/\delta)} + \frac{2C\log(1/\delta)}{S}$. If instead $\lambda^\star > \bar{\lambda}$, then $\hat{\lambda} = \bar{\lambda}$ and we simply obtain the explicit feasible bound $\varepsilon(\bar{\lambda}) = \frac{\bar{\lambda}}{2} S + C\bar{\lambda}^2 + \frac{\log(1/\delta)}{\bar{\lambda}}$. Combining these cases yields (18) with $\hat{\lambda} = \min\{\lambda^\star, \bar{\lambda}\}$. ◀

▶ **Lemma 29** (Scalar Jensen bound for operator moments). *Let $X \succeq 0$ be a positive semidefinite operator on a finite-dimensional Hilbert space and let $\tau \in \mathcal{D}(\mathcal{H})$ be a density operator. Then for every $t \ge 1$, $\big(\mathrm{Tr}(\tau X)\big)^t \le \mathrm{Tr}\big(\tau X^t\big)$.*

Proof. Let $X = \sum_j x_j \Pi_j$ be the spectral decomposition of X with eigenvalues $x_j \ge 0$ and orthogonal projectors Π_j. Define a probability distribution r on eigen-indices by $r_j := \mathrm{Tr}(\tau\Pi_j)$, so that $r_j \ge 0$ and $\sum_j r_j = \mathrm{Tr}(\tau) = 1$. Then $\mathrm{Tr}(\tau X) = \sum_j r_j x_j$, $\mathrm{Tr}(\tau X^t) = \sum_j r_j x_j^t$. Since $f(x) = x^t$ is convex on $\mathbb{R}_+$ for $t \ge 1$, Jensen's inequality gives $\left(\sum_j r_j x_j\right)^t \le \sum_j r_j x_j^t$, which is exactly the desired inequality. ◀

▶ **Theorem 30** (Quantum moments accountant ⇒ measured Rényi DP). *Fix $\alpha > 1$ and let A be a quantum channel. Define, for neighbors $\rho \sim \sigma$,*

$$X(\rho,\sigma) := A(\sigma)^{-1/2} A(\rho) A(\sigma)^{-1/2},$$

with the convention that if $\mathsf{supp}(A(\rho)) \not\subseteq \mathsf{supp}(A(\sigma))$, then the expressions below are $+\infty$. Suppose there exists a function $\alpha_A(\alpha) \ge 0$ such that for all neighbors $\rho \sim \sigma$,

$$\log \mathrm{Tr}\big(A(\sigma)\, X(\rho,\sigma)^\alpha\big) \le (\alpha - 1)\,\varepsilon_\alpha \qquad \text{for some } \varepsilon_\alpha \ge 0. \tag{19}$$

Then A satisfies $(\alpha, \varepsilon_\alpha)$-measured Rényi DP, i.e., for all $\rho \sim \sigma$,

$$D_\alpha^{\mathrm{meas}}\big(A(\rho) \,\|\, A(\sigma)\big) \le \varepsilon_\alpha.$$

Proof. Fix neighbors $\rho \sim \sigma$ and an arbitrary POVM $M = \{M_z\}_z$ on the output space. Let the induced classical distributions be $p_z := \mathrm{Tr}(M_z A(\rho))$, $q_z := \mathrm{Tr}(M_z A(\sigma))$. If $\mathsf{supp}(A(\rho)) \not\subseteq \mathsf{supp}(A(\sigma))$, then $\mathrm{Tr}(A(\sigma)X^\alpha) = +\infty$ and the claim is trivial, so assume $\mathsf{supp}(A(\rho)) \subseteq \mathsf{supp}(A(\sigma))$.

Define the positive operators $B_z := A(\sigma)^{1/2} M_z A(\sigma)^{1/2} \succeq 0$. Then $\sum_z B_z = A(\sigma)$, and moreover $q_z = \mathrm{Tr}(B_z)$, $\quad p_z = \mathrm{Tr}(M_z A(\rho)) = \mathrm{Tr}\big(B_z X(\rho,\sigma)\big)$. For each z with $q_z > 0$, define the normalized state $\tau_z := B_z/q_z \in \mathcal{D}$ so that $\frac{p_z}{q_z} = \mathrm{Tr}(\tau_z X)$. Now consider the classical Rényi divergence of order $\alpha > 1$:

$$\exp\big((\alpha-1)D_\alpha(p\|q)\big) = \sum_z p_z^\alpha q_z^{1-\alpha} = \sum_z q_z \left(\frac{p_z}{q_z}\right)^\alpha = \sum_z q_z \big(\mathrm{Tr}(\tau_z X)\big)^\alpha.$$

Apply Lemma 29 with $t = \alpha$ to each term: $\big(\mathrm{Tr}(\tau_z X)\big)^\alpha \le \mathrm{Tr}(\tau_z X^\alpha)$. Hence

$$\sum_z q_z \big(\mathrm{Tr}(\tau_z X)\big)^\alpha \le \sum_z q_z\, \mathrm{Tr}(\tau_z X^\alpha) = \sum_z \mathrm{Tr}(B_z X^\alpha) = \mathrm{Tr}\left(\sum_z B_z\, X^\alpha\right) = \mathrm{Tr}\big(A(\sigma)\, X^\alpha\big).$$

Therefore, $\exp\big((\alpha-1)D_\alpha(p\|q)\big) \le \mathrm{Tr}\big(A(\sigma)\,X^\alpha\big)$. Taking logs and dividing by $\alpha-1$ gives $D_\alpha(p\|q) \le \frac{1}{\alpha-1}\log \mathrm{Tr}\big(A(\sigma)\,X^\alpha\big)$. By the assumed accountant bound (19), the RHS is at most ε_α. Since M was arbitrary, taking the supremum over all POVMs yields $D_\alpha^{\mathrm{meas}}\big(A(\rho)\,\|\,A(\sigma)\big) = \sup_M D_\alpha\big(P_M(A(\rho))\,\|\,P_M(A(\sigma))\big) \le \varepsilon_\alpha$, which is the desired $(\alpha,\varepsilon_\alpha)$ measured Rényi DP guarantee. ◀

▶ **Corollary 31** (Composition via the quantum moments accountant (measured Rényi route))**.** *Let $A_1,\dots,A_k$ be quantum channels and let $A^{(k)} \coloneqq A_1\otimes\cdots\otimes A_k$ be their tensor-product composition. Consider product neighboring inputs $\rho=\bigotimes_{i=1}^k\rho_i$ and $\sigma=\bigotimes_{i=1}^k\sigma_i$ with $\rho_i\sim\sigma_i$. Fix $\alpha>1$ and set $\lambda\coloneqq\alpha-1$.*

Assume each A_i admits a quantum moments accountant at order α, i.e., there exists $\alpha_{A_i}(\alpha)\ge 0$ such that for all $\rho_i\sim\sigma_i$,

$$\log \mathrm{Tr}\Big[A_i(\sigma_i)\,\Big(A_i(\sigma_i)^{-1/2}A_i(\rho_i)A_i(\sigma_i)^{-1/2}\Big)^{\alpha}\Big] \;\le\; (\alpha-1)\,\alpha_{A_i}(\alpha). \tag{20}$$

*Then the composed channel $A^{(k)}$ satisfies $(\alpha,\varepsilon_\alpha)$-*measured Rényi DP *against arbitrary POVMs with*

$$\varepsilon_\alpha = \sum_{i=1}^k \alpha_{A_i}(\alpha). \tag{21}$$

Consequently, for every $\delta\in(0,1)$, the composition satisfies (ε',δ)-QDP with

$$\varepsilon' = \sum_{i=1}^k \alpha_{A_i}(\alpha) \;+\; \frac{\log(1/\delta)}{\alpha-1}. \tag{22}$$

Proof. Fix $\alpha>1$ and write $\lambda\coloneqq\alpha-1$.

Step 1: Additivity of the α-order moment bound under tensor products. Let $\rho=\bigotimes_{i=1}^k\rho_i$ and $\sigma=\bigotimes_{i=1}^k\sigma_i$ be product neighbors. Define

$$X_i \coloneqq A_i(\sigma_i)^{-1/2}A_i(\rho_i)A_i(\sigma_i)^{-1/2}, \qquad X \coloneqq A^{(k)}(\sigma)^{-1/2}A^{(k)}(\rho)A^{(k)}(\sigma)^{-1/2}.$$

Because $A^{(k)}=\bigotimes_{i=1}^k A_i$ and the inputs are products, we have $A^{(k)}(\rho)=\bigotimes_{i=1}^k A_i(\rho_i)$, $A^{(k)}(\sigma)=\bigotimes_{i=1}^k A_i(\sigma_i)$, $X=\bigotimes_{i=1}^k X_i$. Therefore, using $(\bigotimes_i X_i)^\alpha=\bigotimes_i X_i^\alpha$ and multiplicativity of the trace,

$$\mathrm{Tr}\big(A^{(k)}(\sigma)\,X^\alpha\big) = \mathrm{Tr}\left(\bigotimes_{i=1}^k A_i(\sigma_i)\,\bigotimes_{i=1}^k X_i^\alpha\right) = \prod_{i=1}^k \mathrm{Tr}\big(A_i(\sigma_i)\,X_i^\alpha\big),$$

and hence $\log \mathrm{Tr}\big(A^{(k)}(\sigma)\,X^\alpha\big) = \sum_{i=1}^k \log \mathrm{Tr}\big(A_i(\sigma_i)\,X_i^\alpha\big)$. Applying (20) term-by-term yields $\log \mathrm{Tr}\big(A^{(k)}(\sigma)\,X^\alpha\big) \le (\alpha-1)\sum_{i=1}^k \alpha_{A_i}(\alpha)$.

Step 2: Convert the moment bound to measured Rényi DP. Applying Theorem 30 at order α gives, for all product neighbors $\rho\sim\sigma$, $D_\alpha^{\mathrm{meas}}\big(A^{(k)}(\rho)\,\|\,A^{(k)}(\sigma)\big) \le \frac{1}{\alpha-1}\log \mathrm{Tr}\big(A^{(k)}(\sigma)\,X^\alpha\big) \le \sum_{i=1}^k \alpha_{A_i}(\alpha)$, which proves (21).

Step 3: Convert measured Rényi DP to (ε',δ)-QDP. Fix any POVM M. By definition of measured Rényi divergence,

$$D_\alpha\big(P_M(A^{(k)}(\rho))\,\|\,P_M(A^{(k)}(\sigma))\big) \le D_\alpha^{\mathrm{meas}}\big(A^{(k)}(\rho)\,\|\,A^{(k)}(\sigma)\big) \le \varepsilon_\alpha.$$

Applying the standard classical conversion from $(\alpha, \varepsilon_\alpha)$-Rényi DP to (ε', δ)-DP yields

$$\Pr[M(A^{(k)}(\rho)) = 1] \leq e^{\varepsilon'} \Pr[M(A^{(k)}(\sigma)) = 1] + \delta \quad \text{with} \quad \varepsilon' = \varepsilon_\alpha + \frac{\log(1/\delta)}{\alpha - 1}.$$

Substituting (21) gives (22). ◀

▶ Remark 32 (Comparison to the classical moments accountant of Abadi et al.). Proposition 28 and Corollary 31 are most naturally compared to the classical moments accountant analysis of Abadi et al. for DP-SGD.

Classical moments accountant (Abadi et al.)

In the classical setting one considers a sequence of (randomized) mechanisms and tracks the log moment generating function of the *privacy loss random variable*. Under adaptive composition, these log-moments add, and one converts the resulting moment bound to an (ε, δ)-DP guarantee via Markov's inequality, followed by an optimization over the moment order.

Quantum analogue in this work

In the quantum setting, privacy is defined operationally against *all* measurements (POVMs). A direct quantum analogue of the classical privacy loss random variable is not available prior to measurement. Instead, we work with an operator-level moment quantity: for neighbors $\rho \sim \sigma$ we form the positive operator $X(\rho, \sigma) = A(\sigma)^{-1/2} A(\rho) A(\sigma)^{-1/2}$ and track the moment functional $\mathrm{Tr}\big(A(\sigma) X(\rho, \sigma)^\alpha\big)$. Theorem 30 shows that controlling these moments implies an *operational* Rényi-type privacy guarantee, namely measured Rényi DP, which by definition quantifies worst-case distinguishability over all POVMs.

Additivity vs. post-processing

As in the classical analysis, additivity under composition is obtained at the level of moments: for tensor-product channels and product neighboring inputs, the moments factorize exactly by tensor-product identities and multiplicativity of the trace, yielding the additive accountant in Corollary 31. Unlike the classical case, post-processing by measurements is handled *inside* the privacy notion via measured Rényi divergence: the supremum over POVMs is built into D_α^{meas}, so no additional data-processing inequality is needed at this stage.

Resulting advanced-composition behavior

After converting measured Rényi DP to (ε, δ)-QDP using the standard classical Rényi-DP-to-approximate-DP conversion, one recovers an "advanced composition" tradeoff with dominant term $\sqrt{\sum_i \varepsilon_i^2 \log(1/\delta)}$ (cf. Abadi et al.). The conceptual difference is that the quantum analysis separates (i) an operator-level accounting step (moments add under tensor products) from (ii) an operational step (worst-case over POVMs), whereas in the classical setting the privacy loss is classical from the outset.

7 Advanced Composition for Quantum Differential Privacy

In this section, we provide advanced composition results for QDP mechanisms under the tensor product composition (aka parallel composition) and neighboring notion based on product neighbors (i.e., $\otimes_{i=1}^k \rho_i \sim \otimes_{i=1}^k \sigma_i \iff \rho_i \sim \sigma_i \ \forall i \in \{1, \ldots, k\}$).

For this setting, if one uses basic composition of k QDP mechanisms, we would get a composed mechanism satisfying $(k\varepsilon, 0)$-QDP, if each of them satisfies $(\varepsilon, 0)$-QDP [31, 16]. Next, we show that one can even obtain strong privacy guarantees that scale as $\sqrt{k}\varepsilon$ for sufficiently small ε, which may find use in high privacy regimes with $\varepsilon \leq 1$ and iterative protocols where the impact of k is significant.

▶ **Theorem 33.** *Let $\varepsilon_i \in [0,1]$ for all $1 \leq i \leq k$ and $\delta \in (0,1]$. Let the channel A_i satisfy $(\varepsilon_i, 0)$-QDP. Then, the parallel composition (Definition 5) of $A_1, \ldots, A_k$ (i.e., $A_1 \otimes \cdots \otimes A_k$) satisfies (ε', δ)-QDP with $\varepsilon' \coloneqq \frac{1}{2}\sum_{i=1}^{k}\varepsilon_i^2 + \sqrt{2\log\left(\frac{1}{\delta}\right)\sum_{i=1}^{k}\varepsilon_i^2}$. Furthermore, for $\varepsilon_i = \varepsilon \leq 1$ for all i, we have that $\varepsilon' = \frac{k\varepsilon^2}{2} + \sqrt{k}\varepsilon\sqrt{2\log\left(\frac{1}{\delta}\right)}$.*

Proof. Choose D_α satisfying data-processing and additivity for $\alpha > 1$ (For example, Sandwiched Rényi divergence) We also have $i \in \{1, \ldots, k\}$. A_i satisfying $(\varepsilon_i, 0)$-QDP [1] implies that we have from [25, Proposition 9, Proposition 13]

$$\sup_{\rho\sim\sigma} D_\alpha(A_i(\rho)\|A_i(\sigma)) \leq \min\left\{\frac{\varepsilon_i^2\alpha}{2}, \varepsilon_i\right\}. \tag{23}$$

Then, consider the following:

$$\sup_{\rho_i\sim\sigma_i\ \forall i} D_\alpha\left(A_1(\rho_1)\otimes\cdots\otimes A_k(\rho_k)\|A_1(\sigma_1)\otimes\cdots\otimes A_k(\sigma_k)\right) = \sup_{\rho_i\sim\sigma_i\ \forall i}\sum_{i=1}^{k} D_\alpha\left(A_i(\rho_i)\|A_i(\sigma_i)\right) \tag{24}$$

$$\leq \sum_{i=1}^{k}\min\left\{\frac{\varepsilon_i^2\alpha}{2}, \varepsilon_i\right\}, \tag{25}$$

where the equality follows by the additivity of Rényi divergence and the inequality follows from (23).

Let $0 \preceq M_x \preceq I$ (note that M_x is a measurement operator on k sub-systems), define the following: $p_x(\rho) \coloneqq \mathrm{Tr}[M_x\left(A_1(\rho_1)\otimes\cdots\otimes A_k(\rho_k)\right)]$, $q_x(\sigma) \coloneqq \mathrm{Tr}[M_x\left(A_1(\sigma_1)\otimes\cdots\otimes A_k(\sigma_k)\right)]$. Then, by the data-processing of Rényi divergence, we have that

$$\sup_{\rho_i\sim\sigma_i\ \forall i} D_\alpha\left(A_1(\rho_1)\otimes\cdots\otimes A_k(\rho_k)\|A_1(\sigma_1)\otimes\cdots\otimes A_k(\sigma_k)\right) \geq \sup_{\rho_i\sim\sigma_i\ \forall i}\sup_{\{M_x\}_x} D_\alpha(p_x(\rho)\|q_x(\sigma)), \tag{26}$$

where the second supremization is over all POVMs over k-subsystems. So we have that for all $\{M_x\}_x$ and neighboring pairs, we have

$$D_\alpha(p_x(\rho)\|q_x(\sigma)) \leq \sum_{i=1}^{k}\min\left\{\frac{\varepsilon_i^2\alpha}{2}, \varepsilon_i\right\} \leq \sum_{i=1}^{k}\frac{\varepsilon_i^2}{2}\alpha = \zeta\alpha, \tag{27}$$

where $\zeta \coloneqq \sum_{i=1}^{k}\varepsilon_i^2/2$. This is the point where we can utilize various classical procedures to obtain the desired composition result.

By following the reasoning of the proof of [20, Proposition 3] by utilizing [20, Proposition 10], we have that $p_x(\rho) \leq e^{\varepsilon_\alpha}q_x(\sigma) + \delta$ for $\delta \in (0,1)$ and $\varepsilon_\alpha \coloneqq \zeta\alpha + \frac{\log(1/\delta)}{\alpha-1}$ This holds for all $0 \leq M_x \leq I$ and all neighboring pairs, so we have the guarantees of $(\varepsilon_\alpha, \delta)$-QDP.

Note that the above analysis is valid for all $\alpha > 1$. With that the privacy parameter can be optimized by obtaining $\varepsilon' = \min_{\alpha>1}\left(\zeta\alpha + \frac{\log(1/\delta)}{\alpha-1}\right)$. Let $f(\alpha' \equiv \alpha - 1) = \zeta(\alpha'+1) + b/\alpha'$ by denoting $b \equiv \log(1/\delta)$. Then, $f'(\alpha') = \zeta - b/(\alpha')^2$. With that, the minimum is achieved at $\alpha' = \sqrt{b/\zeta}$. So that, we have $\varepsilon' = \zeta + 2\sqrt{\zeta\log\left(\frac{1}{\delta}\right)}$, and plugging that $\zeta = \sum_{i=1}^{k}\varepsilon_i^2/2$, we conclude the proof. ◀

[1] In [25], the result is shown for the Quantum Pufferfish Privacy (QPP) framework, which generalizes the QDP definition.

8 Conclusion

This work clarifies the landscape of composition guarantees for quantum differential privacy. We showed that classical composition theorems fail in full generality for POVM-based approximate QDP, due to measurement incompatibility and correlated joint channels unique to the quantum setting. At the same time, we demonstrated that these failures are not inherent to quantum channels per se, but arise from specific structural features absent in classical analysis. By restricting attention to tensor-product channels acting on product neighboring inputs, we recovered clean composition guarantees using a quantum moments accountant. Our framework separates operator-level accounting from operational privacy guarantees, enabling advanced-composition-style bounds against arbitrary measurements.

Several directions remain open. It would be valuable to understand whether variants of the quantum moments accountant can handle broader classes of channels, such as factorized but non-tensor-product compositions, or whether approximate DP guarantees can be incorporated without losing additivity. More generally, our results suggest that progress on quantum differential privacy will require careful alignment between mathematical structure and operational threat models, rather than direct transplantation of classical proofs. This work provides both technical tools and conceptual clarity for future investigations into privacy in quantum information processing.

References

1 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, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 322–333. ACM, 2019. doi:10.1145/3313276.3316378.

2 Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, pages 308–318, New York, NY, USA, 2016. ACM. doi:10.1145/2976749.2978318.

3 Armando Angrisani, Mina Doosti, and Elham Kashefi. A unifying framework for differentially private quantum algorithms, 2023. doi:10.48550/arXiv.2307.04733.

4 Armando Angrisani and Elham Kashefi. Quantum differential privacy in the local model. *IEEE Transactions on Information Theory*, 71(5):3675–3692, 2025. doi:10.1109/TIT.2025.3552671.

5 Mario Berta, Kaushik P. Seshadreesan, and Mark M. Wilde. Rényi generalizations of the conditional quantum mutual information. *Journal of Mathematical Physics*, 56(2):022205, February 2015.

6 Hao-Chung Cheng, Christoph Hirche, and Cambyse Rouzé. Sample complexity of locally differentially private quantum hypothesis testing, 2024. doi:10.48550/arXiv.2406.18658.

7 Ayanava Dasgupta, Naqueeb Ahmad Warsi, and Masahito Hayashi. Quantum information ordering and differential privacy, 2025. arXiv:2511.01467.

8 Nilanjana Datta and Felix Leditzky. Second-order asymptotics for source coding, dense coding, and pure-state entanglement conversions. *IEEE Trans. Inf. Theory*, 61(1):582–608, 2015. doi:10.1109/TIT.2014.2366994.

9 D. Ding and M. M. Wilde. Strong converse for the feedback-assisted classical capacity of entanglement-breaking channels. *Probl. Inf. Transm.*, January 2018.

10 Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *EUROCRYPT*, 2006.

11 Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, 2006.

12 Cynthia Dwork, Guy Rothblum, and Salil Vadhan. Boosting and differential privacy. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS '10)*, pages 51–60. IEEE, 23–26 october 2010. doi:10.1109/FOCS.2010.12.

13 Farhad Farokhi. Quantum privacy and hypothesis-testing. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 2841–2846, 2023. doi:10.1109/CDC49753.2023.10384192.

14 Theshani Nuradha Piliththuwasam Gallage. *Theory of Privacy and Testing in a Quantum World*. PhD thesis, Cornell University, 2025. URL: https://www.proquest.com/openview/2cb580e718241481b84185ba6d289ce7/.

15 Ji Guan. Optimal mechanisms for quantum local differential privacy, 2024. arXiv:2407.13516.

16 Christoph Hirche, Cambyse Rouzé, and Daniel Stilck França. Quantum differential privacy: An information theory perspective. *IEEE Transactions on Information Theory*, 69(9):5771–5787, 2023. doi:10.1109/TIT.2023.3272904.

17 P. Kairouz, S. Oh, and P. Viswanath. The composition theorem for differential privacy. *IEEE Transactions on Information Theory*, 63(6):4037–4049, June 2017. doi:10.1109/TIT.2017.2685505.

18 D. Kifer and A. Machanavajjhala. Pufferfish: A framework for mathematical privacy definitions. *ACM Transactions on Database Systems*, 39(1):1–36, 2014. doi:10.1145/2514689.

19 A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*. American Mathematical Society, USA, 2002.

20 Ilya Mironov. Rényi differential privacy. In *30th IEEE Computer Security Foundations Symposium*, 2017.

21 Martin Müller-Lennert, Frédéric Dupuis, Oleg Szehr, Serge Fehr, and Marco Tomamichel. On quantum rényi entropies: A new generalization and some properties. *Journal of Mathematical Physics*, 54(12):122203, December 2013. doi:10.1063/1.4838856.

22 Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

23 Theshani Nuradha, Ian George, and Christoph Hirche. Non-linear strong data-processing for quantum hockey-stick divergences, 2025. doi:10.48550/arXiv.2512.16778.

24 Theshani Nuradha and Ziv Goldfeld. Pufferfish privacy: An information-theoretic study. *IEEE Transactions on Information Theory*, 69(11):7336–7356, 2023. doi:10.1109/TIT.2023.3296288.

25 Theshani Nuradha, Ziv Goldfeld, and Mark M. Wilde. Quantum pufferfish privacy: A flexible privacy framework for quantum systems. *IEEE Trans. Inf. Theory*, 70(8):5731–5762, 2024. doi:10.1109/TIT.2024.3404927.

26 Theshani Nuradha, Vishal Singh, and Mark M. Wilde. Measured hockey-stick divergence and its applications to quantum pufferfish privacy. In *2025 IEEE International Symposium on Information Theory (ISIT)*, pages 1–6, 2025. doi:10.1109/ISIT63088.2025.11195501.

27 Theshani Nuradha and Mark M. Wilde. Contraction of private quantum channels and private quantum hypothesis testing. *IEEE Transactions on Information Theory*, 71(3):1851–1873, March 2025. arXiv:2406.18651v2. doi:10.1109/TIT.2025.3527859.

28 Dénes Petz. Quasi-entropies for States of a von Neumann Algebra. *Publications of the Research Institute for Mathematical Sciences*, 21:787–800, 1985. doi:10.2977/prims/1195178929.

29 Dénes Petz. Quasi-entropies for finite quantum systems. *Reports in Mathematical Physics*, 23:57–65, 1986.

30 Tobias Rippchen, Sreejith Sreekumar, and Mario Berta. Locally-measured rényi divergences. In *2024 IEEE International Symposium on Information Theory (ISIT)*, pages 351–356, 2024. doi:10.1109/ISIT57864.2024.10619141.

31 Li Zhou and Mingsheng Ying. Differential privacy in quantum computation. In *Proceedings of IEEE Computer Security Foundations Symposium (CSF)*, pages 249–262. IEEE, 2017. doi:10.1109/CSF.2017.23.