Information Geometry
https://doi.org/10.1007/s41884-026-00207-5

PERSPECTIVE

Foundations of quantum local asymptotic normality via information geometry

Akio Fujiwara[1,2]

Received: 20 August 2025 / Revised: 26 February 2026 / Accepted: 27 April 2026

Abstract
Local asymptotic normality is a cornerstone of classical asymptotic statistics. Extending it to quantum statistical models entails formidable conceptual and technical challenges arising from their noncommutative nature, above all, the definition of a quantum likelihood ratio, whose very choice dictates the subsequent theoretical development. Employing quantum information geometry based on the symmetric logarithmic derivative, we formulate a quantum Lebesgue decomposition and the square-root likelihood ratio. These constructions provide a natural framework for a quantum analogue of Le Cam's third lemma, the formulation of quantum local asymptotic normality, and the asymptotic representation results developed in our earlier work linking regular quantum statistical models to quantum Gaussian shift models. In this way, quantum information geometry emerges as a natural and unifying foundation for asymptotic quantum statistics.

Keywords Quantum statistics · Local asymptotic normality · Lebesgue decomposition · Likelihood ratio · Quantum information geometry

1 Introduction

The theory of local asymptotic normality (LAN) is a cornerstone of classical asymptotic statistics [1]. It provides a rigorous framework for approximating complex statistical models, in shrinking neighbourhoods of fixed parameter points, by Gaussian shift models. Through this approximation, optimality results, lower bounds, and efficient estimation procedures established in the Gaussian setting can be transferred

Communicated by Nihat Ay.

✉ Akio Fujiwara
fujiwara@math.sci.osaka-u.ac.jp

1 Department of Mathematics, The University of Osaka, Toyonaka, Osaka 560-0043, Japan

2 Center for Quantum Information and Quantum Biology, The University of Osaka, Toyonaka, Osaka 560-0043, Japan

Published online: 19 May 2026

to a broad class of regular statistical models. Since its inception, LAN has become a unifying principle in the asymptotic analysis of statistical inference.

In quantum statistics, analogous questions naturally arise: given a sequence of quantum statistical models indexed by the number of systems n, what is the optimal procedure for estimating the parameter in the asymptotic regime as $n \to \infty$? Moreover, can these models be approximated, in a mathematically precise sense, by simpler quantum Gaussian models?

Earlier approaches to quantum local asymptotic normality (qLAN) [2–5] have yielded important results, but they also exhibit notable limitations: applicability confined to specific models, strong regularity assumptions on the reference state, and a lack of direct connection to parameter estimation in certain contexts. These limitations highlight the need for a more flexible and operationally meaningful quantum analogue of the Radon–Nikodym derivative, one that extends naturally to non-faithful quantum states and directly supports statistical inference tasks.

This paper reconstructs our series of studies on qLAN [6–8] from the perspective of quantum information geometry [9, 10]. At the heart of this framework is a new noncommutative Radon–Nikodym derivative, obtained via a quantum Lebesgue decomposition formulated using the symmetric logarithmic derivative. This construction extends the classical notions of absolute continuity and singularity to arbitrary quantum states, with a clear operational meaning in quantum parameter estimation. Building on this foundation, we give a coherent conceptual account of the quantum analogue of Le Cam's third lemma, the formulation of qLAN, and related asymptotic representation results established in our earlier work, thereby clarifying how quantum asymptotic theory can be placed on a conceptual footing parallel to that of the classical theory.

The paper is organised as follows. Section 2 recalls the necessary measure-theoretic preliminaries and reviews the notion of sufficiency, highlighting parallels between classical and quantum statistics. Section 3 is devoted to extending local asymptotic normality to the quantum domain. After brief introductory remarks, Subsection 3.1 reviews classical LAN and its connection to Gaussian shift models, while Subsection 3.2 outlines the difficulties in extending LAN to the quantum domain. Subsection 3.3 provides a concise review of quantum information geometry, followed in Subsection 3.4 by the introduction of a new notion of absolute continuity for quantum states and the quantum Lebesgue decomposition, both based on the symmetric logarithmic derivative (SLD) quantum information geometry and fundamentally different from the conventional operator-algebraic formulation. Equipped with these new notions, Subsection 3.5 outlines the resulting framework of quantum LAN, including the quantum analogue of Le Cam's third lemma and references to the asymptotic representation theorem proved in our earlier work. Section 4 summarises the main conceptual message and discusses implications for quantum asymptotic theory. For the reader's convenience, an elementary proof of the sufficiency theorem based on quantum information geometry, together with brief accounts of quantum Gaussian states and quantum weak convergence, is presented in Appendices A and B.

2 Likelihood ratio revisited: from classical to quantum

Let $(\Omega, \mathcal{F}, \mu)$ be a σ-finite measure space, and let P be a probability measure that is absolutely continuous with respect to μ. Then, by the Radon–Nikodym theorem [11, Section 14.13], there exists an $\mathcal{F}$-measurable function p such that, for every $A \in \mathcal{F}$, the probability $P(A)$ is given by

$$P(A) = \int_A p\, d\mu, \qquad (\forall A \in \mathcal{F}).$$

Such a measurable function p is referred to as the probability density function (pdf) of P with respect to μ. Throughout this paper, uppercase letters will denote probability measures, and the corresponding lowercase letters will denote their pdfs, making their relationship visually clear.

Now let P and Q be probability measures with $Q \ll P$; that is, $\operatorname{supp} q \subset \operatorname{supp} p$ up to μ-null sets. By the Radon–Nikodym theorem,

$$dQ = \frac{dQ}{dP}\, dP = \frac{q}{p}\, dP,$$

so that Q can be recovered from P via the Radon–Nikodym derivative dQ/dP, or equivalently, the likelihood ratio q/p.

In some contexts, especially in statistics [1, Section 6.1], the Radon–Nikodym derivative is interpreted in a generalised sense. Given two probability measures P and Q, consider the Lebesgue decomposition of Q relative to P:

$$Q = Q^{ac} + Q^{\perp} \qquad (Q^{ac} \ll P,\ \ Q^{\perp} \perp P),$$

where Q^{ac} and $Q^{\perp}$ denote, respectively, the absolutely continuous and singular parts of Q. In this case, dQ/dP refers to dQ^{ac}/dP, which is also commonly written as q/p. Thus, the likelihood ratio carries only the information about the absolutely continuous part of Q relative to P.

The likelihood ratio plays a central role in many areas of classical statistics [1], including hypothesis testing, f-divergences, model comparison, local asymptotic normality, among others. A natural question is whether analogous concepts exist in quantum statistics. Indeed, certain versions of the quantum likelihood ratio have played essential roles in quantum statistics.

Let ρ and σ be quantum states on a finite dimensional Hilbert space $\mathcal{H}$. A quantum state is a noncommutative analogue of the probability density and is represented by a density operator: a selfadjoint, positive operator of unit trace. Assume $\operatorname{supp} \sigma \subset \operatorname{supp} \rho$, where the support of a selfadjoint operator is the orthogonal complement of its kernel. This support condition ensures that the Umegaki quantum relative entropy

$$D(\sigma, \rho) := \operatorname{Tr} \sigma (\log \sigma - \log \rho),$$

a quantum analogue of the Kullback–Leibler (KL) divergence, is finite.

The Umegaki quantum relative entropy satisfies the monotonicity property [12, Theorem 1.5]:

$$D(\sigma, \rho) \geq D(\sigma|_{\mathcal{A}}, \rho|_{\mathcal{A}}) \tag{1}$$

for any $*$-subalgebra $\mathcal{A}$ of $B(\mathcal{H})$, the algebra of all linear operators on $\mathcal{H}$. Here, $\sigma|_{\mathcal{A}}$ denotes the state on $\mathcal{A}$ obtained by restricting σ to observables in $\mathcal{A}$, that is,

$$\mathrm{Tr}\,(\sigma|_{\mathcal{A}})A = \mathrm{Tr}\,\sigma A, \qquad (\forall A \in \mathcal{A}).$$

Consider first the case when $\mathcal{A}$ is commutative and generated by a projective measurement consisting of a complete set of mutually orthogonal projections. In this setting, the right-hand side of (1) reduces to the classical KL divergence between the corresponding outcome distributions under σ and ρ. The inequality thus expresses the general principle that information is inevitably lost through measurement.

This naturally leads to the following question: under what conditions does equality hold in (1)? The sufficiency theorem answers this: for commutative $\mathcal{A}$, equality holds if and only if there exists a selfadjoint $X \in \mathcal{A}$ such that $\sigma = \rho X$. In this situation, σ, ρ, and X commute. Thus, this result can be intuitively restated as: $\mathcal{A}$ is sufficient if and only if the "quantum" likelihood ratio σ/ρ belongs to $\mathcal{A}$. This theorem is well known when both σ and ρ are strictly positive [12, Proposition 1.16]. However, for general states satisfying only the support condition, the result does not seem to have appeared explicitly in the literature. Therefore, for the reader's convenience, a simple proof is provided in Appendix A as an elementary application of quantum information geometry.

The sufficiency theorem extends to arbitrary (possibly noncommutative) $*$-subalgebras $\mathcal{A}$. Given a quantum parametric model $\{\rho_\theta : \theta \in \Theta\}$, quantum sufficiency is defined as the existence of a completely positive trace preserving (CPTP) map $\Gamma : \mathcal{A} \to B(\mathcal{H})$ such that

$$\Gamma(\rho_\theta|_{\mathcal{A}}) = \rho_\theta, \qquad (\forall \theta \in \Theta).$$

In other words, the sufficiency of $\mathcal{A}$ means that all the information about the parameter θ is already contained in the restricted model $\rho_\theta|_{\mathcal{A}}$. For technical convenience, we assume that there exists a strictly positive (faithful) state ω in the closed convex hull of $\{\rho_\theta\}$. Then the quantum sufficiency theorem [13, Theorem 9.7] asserts that $\mathcal{A}$ is sufficient for $\{\rho_\theta\}$ if and only if

$$D(\rho_\theta, \omega) = D(\rho_\theta|_{\mathcal{A}}, \omega|_{\mathcal{A}}), \qquad (\forall \theta \in \Theta).$$

Equivalently, sufficiency holds if and only if the Connes cocycle derivative [12, 13], defined by the contraction operator $\rho_\theta^{\sqrt{-1}t}\omega^{-\sqrt{-1}t}$, remains unchanged under restriction to $\mathcal{A}$ for all $t \in \mathbb{R}$ and $\theta \in \Theta$:

$$\rho_\theta^{\sqrt{-1}t}\omega^{-\sqrt{-1}t} = (\rho_\theta|_{\mathcal{A}})^{\sqrt{-1}t}(\omega|_{\mathcal{A}})^{-\sqrt{-1}t} \in \mathcal{A} \qquad (\forall t \in \mathbb{R})$$

Note that if ρ_θ and ω commute, the Connes cocycle derivative reduces to a complex power of the likelihood ratio ρ_θ/ω.

Recall that in classical statistics [14, Section 1.5], a σ-subalgebra $\mathcal{G}$ of $\mathcal{F}$ is sufficient for a classical statistical model $\{P_\theta\}$ if and only if there exists a Q in the closed convex hull of the model such that

$$D(P_\theta, Q) = D(P_\theta|_{\mathcal{G}}, Q|_{\mathcal{G}}),$$

or equivalently,

$$\frac{dP_\theta}{dQ} = \frac{d(P_\theta|_{\mathcal{G}})}{d(Q|_{\mathcal{G}})} \in m\mathcal{G},$$

where $m\mathcal{G}$ denotes the set of $\mathcal{G}$-measurable functions. The analogy between the classical and quantum sufficiency theorems is thus clear.

Due to this strong parallelism, some authors argue that the Connes cocycle derivative is the proper quantum counterpart of the Radon–Nikodym derivative. But we take a different view: just as quantum information theory provides a variety of quantum extensions of KL divergence, the Radon–Nikodym derivative should also admit multiple quantum analogues. Relying on a single concept for quantum extensions is, in our opinion, a risky oversimplification. To illustrate this perspective, we now turn to the main theme of this paper, namely, the theory of local asymptotic normality in the quantum domain.

3 Local asymptotic normality in the quantum domain

To clarify our motivation, let us begin with a typical example. Suppose we have n copies of a quantum system, each in the same state depending on an unknown parameter θ. Our task is to estimate θ by making some measurement on the n systems. A key feature of the quantum setting is that we are allowed to use collective measurement, rather than a separable measurement. In a separable measurement scheme, each system is measured individually. In contrast, a collective measurement treats the n systems as a single quantum object, allowing entangled measurements that act jointly across all n subsystems. This approach generally enables us to extract more information about θ than when we measure the systems individually.

This raises a natural question: as $n \to \infty$, what is the ultimate precision achievable when both the measurement and the subsequent data processing are optimised? The purpose of this study is to address this question by extending the theory of local asymptotic normality (LAN), a cornerstone of classical asymptotic statistics, to quantum statistical models. In doing so, we will see that the choice of a suitable quantum analogue of the classical likelihood ratio, as discussed in the preceding section, plays a crucial role in formulating and proving the quantum LAN framework.

3.1 Review of classical local asymptotic normality

In classical statistics, a sequence $\{P_\theta^{(n)} : \theta \in \Theta \subset \mathbb{R}^d\}$ of statistical models on measurable spaces $(\Omega^{(n)}, \mathcal{F}^{(n)})$ is said to be *locally asymptotically normal* (LAN) at θ_0 [1, Definition 7.14] if the log-likelihood ratio $\log(dP_\theta^{(n)}/dP_{\theta_0}^{(n)})$ admits an expansion in the local parameter $h := \sqrt{n}(\theta - \theta_0)$ of the form

$$\log \frac{dP^{(n)}_{\theta_0 + h/\sqrt{n}}}{dP^{(n)}_{\theta_0}} = h^i \Delta_i^{(n)} - \frac{1}{2} h^i h^j J_{ij} + o_{P^{(n)}_{\theta_0}}(1). \tag{2}$$

Here, $\Delta^{(n)} = (\Delta_1^{(n)}, \ldots, \Delta_d^{(n)})$ is a d-dimensional random vectors satisfying

$$\Delta^{(n)} \overset{P^{(n)}_{\theta_0}}{\rightsquigarrow} N(0, J),$$

J is a positive definite $d \times d$ real symmetric matrix, the squiggling arrow $\rightsquigarrow$ denotes weak convergence under the measure specified above the arrow, the remainder term $o_{P^{(n)}_{\theta_0}}(1)$ converges in probability to zero under $P^{(n)}_{\theta_0}$, and Einstein's summation convention is used.

A prototypical example of a LAN sequence is the i.i.d. extension of a base model $\{p_\theta\}$. In this case, the log-likelihood ratio expands as

$$\begin{aligned}
&\log \frac{p^{\otimes n}_{\theta_0 + h/\sqrt{n}}}{p^{\otimes n}_{\theta_0}}(X_1, \ldots, X_n) \\
&\quad = h^i \left\{ \frac{1}{\sqrt{n}} \sum_{k=1}^{n} \partial_i \log p_{\theta_0}(X_k) \right\} - \frac{1}{2} h^i h^j \left\{ -\frac{1}{n} \sum_{k=1}^{n} \partial_i \partial_j \log p_{\theta_0}(X_k) \right\} + o\left(\frac{1}{n}\right)
\end{aligned}$$

By the central limit theorem, the first term converges weakly to the normal distribution $N(0, J)$ under p_{θ_0}, where J is the Fisher information matrix of the base model p_θ. By the law of large numbers, the second term converges in probability to $-\frac{1}{2} h^\top J h$. In this way, the LAN property naturally arises from the asymptotic behaviour of i.i.d. models, and the definition (2) may be viewed as formalising and extending this phenomenon.

An important observation is that LAN expansion (2) mirrors the likelihood ratio in a Gaussian shift model:

$$\log \frac{dN(Jh, J)}{dN(0, J)}(X_1, \ldots, X_d) = h^i X_i - \frac{1}{2} h^i h^j J_{ij}.$$

This striking similarity suggests that LAN models are, in some sense, asymptotically equivalent to Gaussian shift models. But, to translate this similarity into a rigorous statistical methodology, one requires an asymptotic analogue of absolute continuity. This is provided by the notion of contiguity, introduced by Le Cam.

A sequence $Q^{(n)}$ of probability measures is called *contiguous* to another sequence $P^{(n)}$ of probability measures [1, Section 6.2], denoted $Q^{(n)} \lhd P^{(n)}$, if $P^{(n)}(A^{(n)}) \to 0$ implies $Q^{(n)}(A^{(n)}) \to 0$. We also refer to this property as *asymptotic absolute continuity*.

Since the concept of absolute continuity is closely related to the Radon–Nikodym theorem, it is natural to seek its weak convergence counterpart. This is precisely the content of Le Cam's third lemma [1, Theorem 6.6]: if $Q^{(n)} \lhd P^{(n)}$ and the pair $\left(X^{(n)}, dQ^{(n)}/dP^{(n)}\right)$ jointly converges weakly to (X, V) under $P^{(n)}$, then under $Q^{(n)}$ the random variable $X^{(n)}$ converges weakly to the law:

$$\mathcal{L}(B) := E[\mathbb{1}_B(X)V]$$

for every Borel set B. Here, the random variable V plays the role of a limiting Radon–Nikodym derivative, allowing the asymptotic reconstruction of $Q^{(n)}$ from $P^{(n)}$.

Le Cam's third lemma integrates seamlessly with the LAN framework [1, Example 6.7]. Suppose $\{P_\theta^{(n)}\}$ is LAN at θ_0, and that the statistic $\Delta^{(n)}$ in (2) and another statistic $X^{(n)}$ jointly satisfy

$$\begin{pmatrix} X^{(n)} \\ \Delta^{(n)} \end{pmatrix} \overset{P_{\theta_0}^{(n)}}{\rightsquigarrow} N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \Sigma & \tau \\ \tau^\top & J \end{pmatrix}\right).$$

In this situation, $P_{\theta_0+h/\sqrt{n}}^{(n)}$ and $P_{\theta_0}^{(n)}$ are mutually contiguous for all h, and Le Cam's third lemma concludes that, under shifted measures $P_{\theta_0+h/\sqrt{n}}^{(n)}$,

$$X^{(n)} \overset{P_{\theta_0+h/\sqrt{n}}^{(n)}}{\rightsquigarrow} N(\tau h, \Sigma).$$

This establishes rigorously that a LAN model is locally asymptotically equivalent to a Gaussian shift model, confirming the intuition from the formal similarity of their likelihood ratios.

3.2 Challenges in extending LAN to the quantum domain

In moving from classical to quantum statistics, it is natural to seek a quantum counterpart of the theory of LAN, which would provide a unified framework for analysing the asymptotic behaviour of quantum statistical models. However, formulating such a theory entails several fundamental challenges. First, what should be adopted as the quantum analogue of the Radon–Nikodym derivative? Second, how should the notions of contiguity and Le Cam's third lemma be extended to the quantum domain? Third, how can one formulate a meaningful notion of weak convergence for quantum statistical models? Resolving these issues is essential for establishing a general theory of quantum LAN (qLAN). In this paper, we address the first of these problems: identifying a quantum analogue of the Radon–Nikodym derivative that is well suited to quantum parameter estimation.

Before presenting our result, let us briefly outline the history of qLAN. The first formulation was pioneered by Guţă and Kahn [2, 3], who considered the i.i.d. extension of a local model in a shrinking neighbourhood of θ_0. They showed that such a model can be embedded into an infinite dimensional Hilbert space by a CPTP map that is uniformly close to a quantum Gaussian shift model σ_h:

$$\lim_{n\to\infty} \sup_{h\in K^{(n)}} \left\| \sigma_h - \Gamma^{(n)}(\rho_{\theta_0+h/\sqrt{n}}^{\otimes n}) \right\|_1 = 0$$

and, conversely:

$$\lim_{n\to\infty} \sup_{h\in K^{(n)}} \left\| \Lambda^{(n)}(\sigma_h) - \rho_{\theta_0+h/\sqrt{n}}^{\otimes n} \right\|_1 = 0.$$

Here, the convergence is in the trace-class norm rather than in weak convergence. While powerful, this formulation has inherent limitations: the proof applies only to i.i.d. sequences and imposes restrictions on the eigenvalues of the reference state ρ_{θ_0}, which are natural for isolating the generic nondegenerate regime but do not appear to be essential from the broader viewpoint of qLAN.

Subsequently, Guţă and Jenčová [5] sought to develop a weak convergence version of qLAN by adopting the Connes cocycle derivative as the quantum analogue of the Radon–Nikodym derivative. However, they were unable to establish the desired asymptotic expansion formula that would allow the recovery of information about the alternative quantum state from the reference state. This shortcoming suggests that the Connes cocycle derivative may not be an appropriate tool for qLAN.

Motivated by this, we adopt a completely different approach, seeking a proper counterpart of the Radon–Nikodym derivative adapted to the qLAN setting, drawing on methods from quantum information geometry, specifically the symmetric logarithmic derivative (SLD) geometry.

3.3 Quantum information geometry

We begin with a brief overview of quantum information geometry [9, 10]. Let $\mathcal{H}$ be a finite-dimensional Hilbert space, and consider the manifold of strictly positive density operators with unit trace:

$$\mathcal{S} = \mathcal{S}(\mathcal{H}) = \{\, \rho \in B(\mathcal{H}) : \rho > 0,\ \mathrm{Tr}\,\rho = 1 \,\}.$$

On this manifold, quantum Fisher metrics are introduced through the pairing of the m- and e-representations of tangent vectors:

$$g_\rho(X, Y) = \mathrm{Tr}\left\{ X^{(m)} Y^{(e)} \right\}, \qquad (X, Y \in T_\rho \mathcal{S}).$$

Here, a tangent vector X at ρ is understood as a directional differential operator acting on functions on $\mathcal{S}$. Hence, the m-representation of a tangent vector $X \in T_\rho\mathcal{S}$, defined by

$$X^{(m)} := X\rho,$$

is the directional derivative of the state ρ in the direction X, and thus a traceless selfadjoint element of $B(\mathcal{H})$.

The key difference from the classical setting lies in the e-representation $X^{(e)}$ of a tangent vector X at ρ. In classical information geometry, the e-representation is just the derivative of $\log p$. In the quantum case, it is defined by

$$X^{(e)} := X_f^{(e)} := f(\Delta_\rho)\left\{ \rho^{-1}(X\rho) \right\},$$

where $\Delta_\rho : B(\mathcal{H}) \to B(\mathcal{H})$ is the modular operator

$$\Delta_\rho A := \rho A \rho^{-1},$$

and f is a function on the positive reals satisfying certain conditions.

For example, when $f(t) = t \log t/(t-1)$, the e-representation $X_f^{(e)}$ reduces to the operator $X \log \rho$, closely paralleling the classical case, and the corresponding metric is known as the Bogoliubov metric. Alternatively, when $f(t) = 2t/(1+t)$, the e-representation $X_f^{(e)}$ reduces to the symmetric logarithmic derivative (SLD), the selfadjoint operator L_X^S satisfying

$$X\rho = \frac{1}{2}\left(L_X^S \rho + \rho L_X^S\right),$$

and the corresponding metric is called the SLD metric.

Among various quantum Fisher metrics, those of particular interest are the *monotone metrics*, which, like the quantum relative entropy, are contractive under CPTP maps $\Gamma : \mathcal{S}(\mathcal{H}) \to \mathcal{S}(\mathcal{H}')$ as

$$g_\rho(X, X) \geq g_{\Gamma\rho}(\Gamma_* X, \Gamma_* X).$$

By Petz's theorem [15], a metric g defined via f is monotone if and only if f is operator monotone. Building on these insights, Nagaoka proposed a dualistic structure on the quantum state space as follows. The m-connection $\nabla^{(m)}$ is defined so that the parallel translation

$$\Pi^{(m)} : T_\rho \mathcal{S} \to T_\rho \mathcal{S} : X\rho \mapsto X\rho$$

is induced by the convex structure of the quantum state space $\mathcal{S}$. The e-connection $\nabla^{(e)}$ is then defined by the duality relation:

$$Xg(Y, Z) = g(\nabla_X^{(m)} Y, Z) + g(Y, \nabla_X^{(e)} Z).$$

This construction yields a family of dualistic structures parameterised by operator monotone functions f. Nagaoka proved that the triple $(g, \nabla^{(m)}, \nabla^{(e)})$ is dually flat if and only if g is the Bogoliubov metric associated with $f(t) = t \log t/(t-1)$; see [9, Theorem 7.1] for a proof. However, this dually flat structure appears to have no direct applications to quantum estimation problems.

Instead, the structure induced by the SLD metric, corresponding to $f(t) = 2t/(1+t)$, plays a far more central role in quantum estimation problems. In fact, it is this structure that led us to a new noncommutative Radon–Nikodym derivative suited to the theory of weak qLAN.

For a one-dimensional parametric family ρ_θ, the SLD L_θ^S is the selfadjoint operator satisfying

$$\frac{d\rho_\theta}{d\theta} = \frac{1}{2}(L_\theta^S \rho_\theta + \rho_\theta L_\theta^S).$$

If the SLDs are independent of θ up to an additive constant ensuring zero expectation, this differential equation can be integrated to obtain

$$\rho_\theta = e^{\frac{1}{2}(\theta L - \psi(\theta))} \rho_0 \, e^{\frac{1}{2}(\theta L - \psi(\theta))},$$

where the operator L is related to the SLD by $L_\theta^S = L - \psi'(\theta)$. This defines the *SLD $\nabla^{(e)}$-geodesic* or the *SLD quantum exponential family*. Such families enjoy desirable properties. First, the quantum Cramér–Rao lower bound is uniformly achievable; hence θ admits an efficient estimator, as in the classical exponential family. Furthermore, this family naturally extends to non-faithful states, allowing ρ_0 to be degenerate, including pure states. In particular, the associated SLD $\nabla^{(e)}$-geodesic is rank-preserving: if $\operatorname{rank} \rho_0 = k$, then $\operatorname{rank} \rho_\theta = k$ for all θ.

Motivated by these properties, we propose a new definition of absolute continuity for quantum states, in which two quantum states connected by an SLD $\nabla^{(e)}$-geodesic are regarded as mutually absolutely continuous. In the next subsection, we harness this SLD-based notion of absolute continuity to define a new noncommutative Radon–Nikodym derivative.

3.4 Quantum lebesgue decomposition

Let σ and ρ be quantum states on a finite dimensional Hilbert space $\mathcal{H}$, possibly non-faithful, meaning that they may have zero eigenvalues. In what follows, our scope extends to the full quantum state space:

$$\overline{\mathcal{S}} = \overline{\mathcal{S}(\mathcal{H})} = \{\rho \in B(\mathcal{H}) : \rho \geq 0, \ \operatorname{Tr} \rho = 1\}.$$

We say σ is *absolutely continuous* to ρ, written $\sigma \ll \rho$, if there exists a positive operator R such that

$$\sigma = R\rho R.$$

If R is strictly positive, then multiplying both sides by R^{-1} shows that ρ is also absolutely continuous to σ. This is precisely the case when σ and ρ are connected by an SLD $\nabla^{(e)}$-geodesic. Conversely, we say that σ is *singular* to ρ, denoted $\sigma \perp \rho$, if their supports are orthogonal.

Our definition of absolute continuity is substantially more flexible than the conventional requirement

$$\operatorname{supp} \sigma \subset \operatorname{supp} \rho$$

used in the operator-algebraic framework, which is modelled directly on the classical measure-theoretic notion. This "classical" condition is overly restrictive, as demonstrated most clearly in the case of pure states. For pure states $\rho = |\xi\rangle\langle\xi|$ and $\sigma = |\eta\rangle\langle\eta|$ with unit vectors $\xi, \eta \in \mathcal{H}$, the inclusion $\operatorname{supp} \sigma \subset \operatorname{supp} \rho$ holds if and only if $\sigma = \rho$. By contrast, under our definition, $\sigma \ll \rho$ if and only if ξ and η are not orthogonal. In fact, the "only if" part is obvious. For the "if" part, let

$$R := I + \frac{1}{|\langle\eta|\xi\rangle|} |\eta\rangle\langle\eta| - |\xi\rangle\langle\xi|.$$

Then $R > 0$ and

$$R|\xi\rangle = \frac{\langle\eta|\xi\rangle}{|\langle\eta|\xi\rangle|}\,|\eta\rangle,$$

showing that $\sigma = R\rho R$. This enhanced flexibility constitutes not merely a technical refinement but a fundamental improvement, ensuring the broad applicability of the qLAN theory.

A first major outcome of this definition is a natural noncommutative analogue of the Lebesgue decomposition. Given two quantum states ρ and σ, there exists a unique decomposition

$$\sigma = R\rho R + \tau \qquad (R \geq 0,\ \tau \geq 0,\ \tau \perp \rho), \tag{3}$$

where the first term is the absolutely continuous part of σ with respect to ρ, and the second term is the singular part.[1] This result generalises the classical Lebesgue decomposition theorem to the quantum setting, which we therefore refer to it as the *quantum Lebesgue decomposition* [7]. In what follows, we denote the absolutely continuous and singular parts by σ^{ac} and $\sigma^{\perp}$, respectively.

With the quantum Lebesgue decomposition in place, we can define a quantum analogue of the Radon–Nikodym derivative. We call the positive operator R appearing in the quantum Lebesgue decomposition (3) the *square-root likelihood ratio*, denoted $\mathcal{R}(\sigma \mid \rho)$. This operator provides a noncommutative analogue of the Radon–Nikodym derivative adapted to quantum estimation theory, extending the classical concept so that it applies to all pairs of quantum states without the restrictive support-inclusion assumption. It plays a central role in extending the notions of contiguity, weak convergence, and related asymptotic concepts to the quantum domain, and thereby in formulating quantum local asymptotic normality. Note that when $\operatorname{supp}\sigma \subset \operatorname{supp}\rho$, $\mathcal{R}(\sigma|\rho)$ coincides with the conventional noncommutative Radon–Nikodym derivative,[2] showing that our formulation strictly generalises the traditional one.

3.5 Local asymptotic normality in the quantum domain

The notion of quantum local asymptotic normality (qLAN) is naturally defined in close parallel to its classical counterpart, by considering an expansion in the local parameter h for the logarithm of the square of the square-root likelihood ratio introduced in the previous subsection. Let $\mathcal{S}^{(n)} = \{\rho_\theta^{(n)} : \theta \in \Theta \subset \mathbb{R}^d\}$ be a sequence of quantum statistical models. We say that $\mathcal{S}^{(n)}$ is qLAN at $\theta_0 \in \Theta$ if

$$R_h^{(n)} := \mathcal{R}(\rho^{(n)}_{\theta_0+h/\sqrt{n}}|\rho^{(n)}_{\theta_0})$$

admits the expansion

$$\log\left[\left(R_h^{(n)} + o_{L^2}(\rho_{\theta_0}^{(n)})\right)^2\right] = h^i\Delta_i^{(n)} - \frac{1}{2}\left(J_{ij}h^ih^j\right)I^{(n)} + o_D(h^i\Delta_i^{(n)}, \rho_{\theta_0}^{(n)}), \tag{4}$$

[1] While the decomposition (3) into the absolutely continuous part $R\rho R$ and the singular part τ is unique, the choice of the operator R realising the absolutely continuous component is generally not unique.

[2] Introduced by Sakai in the context of normal positive linear functionals on a von Neumann algebra [16].

with

$$\Delta^{(n)} \overset{\rho_{\theta_0}}{\rightsquigarrow} N(0, J).$$

Here, the term $o_{L^2}(\rho_{\theta_0}^{(n)})$, which vanishes in L^2 in the asymptotic limit, is included for technical reasons (in particular, to ensure that the argument of the logarithm is strictly positive), while the final term on the right-hand side vanishes weakly as $n \to \infty$. The notation $N(0, J)$ denotes a hybrid classical/quantum Gaussian state[3] with complex covariance matrix J.

To illustrate that the definition is nontrivial and applies beyond regular full-rank settings, we include the following singular example. Let $\mathcal{H} = \mathbb{C}^2$ and consider the two-dimensional spin-$\frac{1}{2}$ pure-state family

$$\begin{aligned}\tilde{\rho}_\theta &:= e^{\frac{1}{2}(\theta^1\sigma_1+\theta^2\sigma_2-\psi(\theta))} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} e^{\frac{1}{2}(\theta^1\sigma_1+\theta^2\sigma_2-\psi(\theta))} \\ &= \frac{1}{2}\left\{ I + \frac{\tanh\|\theta\|}{\|\theta\|}(\theta^1\sigma_1 + \theta^2\sigma_2) + \frac{1}{\cosh\|\theta\|}\sigma_3 \right\},\end{aligned}$$

where σ_i are the standard Pauli matrices, $\theta = (\theta^1, \theta^2) \in \mathbb{R}^2$, and $\psi(\theta) := \log\cosh\|\theta\|$. Now define a perturbed model ρ_θ by

$$\rho_\theta := e^{-\|\theta\|^\alpha}\tilde{\rho}_\theta + \left(1 - e^{-\|\theta\|^\alpha}\right)\begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix},$$

where α is a fixed positive constant. Intuitively, larger α means that the additional mixing term vanishes faster as $\theta \to 0$. Geometrically, this model is tangential to the Bloch sphere at the north pole $\rho_0 = \tilde{\rho}_0$, and it has a singularity at $\theta = 0$ in the sense that the rank drops there. Nevertheless, it is shown in Section 7.3 of [7] that the i.i.d. extension $\{\rho_\theta^{\otimes n}\}_\theta$ is qLAN at $\theta = 0$ if and only if $\alpha > 2$. This example illustrates the flexibility of our definition of qLAN.

The notion of contiguity also extends naturally to the quantum setting, enabling a quantum version of Le Cam's third lemma. Specifically, suppose that $\{\rho_\theta^{(n)}\}$ is qLAN at θ_0, and that the observable $\Delta^{(n)}$ in (4) and $X^{(n)}$ jointly satisfy

$$\begin{pmatrix} X^{(n)} \\ \Delta^{(n)} \end{pmatrix} \overset{\rho_{\theta_0}^{(n)}}{\rightsquigarrow} N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \Sigma & \tau \\ \tau^* & J \end{pmatrix}\right).$$

Then, under the shifted states $\rho_{\theta_0+h/\sqrt{n}}^{(n)}$,

$$X^{(n)} \overset{\rho_{\theta_0+h/\sqrt{n}}^{(n)}}{\rightsquigarrow} N((\operatorname{Re}\tau)h, \Sigma).$$

In other words, the limiting states form a family of quantum Gaussian states whose mean shifts in the direction of $\operatorname{Re}\tau$ as h varies. This mirrors the classical case: a

[3] For the reader's convenience, Appendix B provides a brief overview of quantum Gaussian states and quantum weak convergence.

qLAN model is locally asymptotically equivalent to a quantum Gaussian shift model. A full proof is technically involved and relies on the machinery of quantum contiguity and quantum weak convergence; we therefore refer the reader to [7] for the complete argument.

Finally, under an additional technical condition, D-extendibility, an asymptotic representation theorem[4] in the quantum domain is established in [8]. This theorem guarantees the existence of a POVM on the limiting von Neumann CCR algebra, specified by the imaginary part of Σ, that faithfully captures the asymptotic behaviour of the sequence of POVMs. This result serves as a cornerstone in proving that no sequence of POVMs can attain superefficiency. In particular, it rules out quantum analogues of estimators such as Hodges' estimator or the James–Stein estimator. We refer to [8] for the precise statement, assumptions, and proof.

4 Conclusions

In this paper, we have argued that quantum local asymptotic normality can be formulated most naturally within the framework of quantum information geometry. Alternative approaches, most notably those based on the Connes cocycle derivative, yield valuable insights into sufficiency, but they lack the flexibility required for a comprehensive theory of quantum asymptotic statistics. In particular, the conventional noncommutative Radon–Nikodym derivative applies only under the restrictive support-inclusion condition, which, for example, forces equality in the pure-state case (see Subsection 3.4). Such operator-algebraic formulations offer no natural route to analogues of the Gaussian shift models, which lie at the core of LAN theory.

By contrast, the SLD-based dualistic structure of quantum information geometry directly motivates a broader, operationally meaningful notion of absolute continuity between quantum states, leading to the quantum Lebesgue decomposition and the square-root likelihood ratio. These constructions enable the extension of classical asymptotic tools, such as contiguity and Le Cam's third lemma, to the quantum domain, and culminate in a formulation of qLAN [7] that closely parallels its classical counterpart. The same machinery also underpins the asymptotic representation theorem under D-extendibility, which is established in [8]. In particular, that theorem yields an exact quantum Gaussian shift limit and excludes superefficient estimation schemes.

In summary, quantum information geometry is not merely a convenient language for reexpressing quantum statistical concepts in geometric terms. It provides an essential structural foundation for linking finite-dimensional quantum models to their quantum Gaussian limits, thereby ensuring both conceptual clarity and statistical coherence.

Acknowledgements The present study was supported by JSPS KAKENHI Grant Numbers JP22340019, JP17H02861, JP23H01090, and JP23K25787.

Author Contributions A.F. is the sole author of this manuscript and contributed to all aspects of the work.

Funding Open Access funding provided by The University of Osaka

[4] For the classical counterpart, see Theorem 7.10 and Section 9.2 of [1].

Data availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The author is a member of the Editorial Board of Information Geometry. This did not affect the editorial handling of the manuscript. The author declares no other conflict of interest.

Appendix

A Proof of sufficiency theorem when $\mathcal{A}$ is commutative

In the general (noncommutative) case, providing the sufficiency theorem requires more elaborate machinery from operator algebras [13, Section 9.5]. Here, for clarity and accessibility, we present a self-contained proof for the case when $\mathcal{A}$ is commutative, using only the information-geometric structure of quantum and classical statistical models.

Proof Assume first that there exists a selfadjoint $X \in \mathcal{A}$ such that $\sigma = \rho X$. Then

$$D(\sigma, \rho) = \mathrm{Tr}\, \sigma \log X = \mathrm{Tr}\, (\sigma|_{\mathcal{A}}) \log X = D(\sigma|_{\mathcal{A}}, \rho|_{\mathcal{A}}),$$

which proves the "if" part.

Suppose now that $D(\sigma, \rho) = D(\sigma|_{\mathcal{A}}, \rho|_{\mathcal{A}})$. Let $\rho_t := (1-t)\sigma + t\rho$ be the mixture geodesic connecting σ and ρ, and let

$$p_t(\omega) := (1-t)\, \{\mathrm{Tr}\, \sigma E(\omega)\} + t\, \{\mathrm{Tr}\, \rho E(\omega)\}$$

be the corresponding classical mixture geodesic of outcome distributions for the projective measurement E generating $\mathcal{A}$. It is well known that the Bogoliubov quantum metric on the quantum state space $\mathcal{S}(\mathcal{H})$ induces a dually flat structure, whose canonical divergence coincides with the Umegaki relative entropy $D(\sigma, \rho)$, and consequently,[5]

$$D(\sigma, \rho) = \int_0^1 t J^B(t) dt,$$

[5] This integral representation follows directly from Eq. (3.54) in [9].

where J^B is the Bogoliubov Fisher information along the geodesic ρ_t. Similarly, the classical Kullback–Leibler divergence $D(\sigma|_{\mathcal{A}}, \rho|_{\mathcal{A}})$ can be expressed as

$$D(\sigma|_{\mathcal{A}}, \rho|_{\mathcal{A}}) = \int_0^1 t J^E(t) dt,$$

where $J^E(t)$ denotes the classical Fisher information along p_t.

A fundamental fact is that $J^B(t) \geq J^E(t)$ for all t, with equality for all t if and only if the Umegaki relative entropy is preserved under the restriction to $\mathcal{A}$. To make the equality condition explicit, we insert the symmetric logarithmic derivative (SLD) quantum Fisher information $J^S(t)$ into the chain of equalities:

$$J^B(t) \geq J^S(t) \geq J^E(t).$$

Thus, equality of relative entropies occurs if and only if

$$J^B(t) = J^S(t) = J^E(t), \qquad (t \in (0, 1]).$$

The first equality $J^B(t) = J^S(t)$ holds if and only if σ and ρ commute, while the second equality $J^S(t) = J^E(t)$ holds if and only if the SLD operator L_t^S lies in $\mathcal{A}$ for all t. In particular, at $t = 1$, setting $X := I - L_1^S$ yields

$$\rho X = \rho - \rho L_1^S = \rho - \left.\frac{d\rho_t}{dt}\right|_{t=1} = \rho - (\rho - \sigma) = \sigma,$$

which completes the proof of the "only if" part. □

Note that the condition $L_t^S \in \mathcal{A}$ carries a clear operational interpretation: it ensures that the SLD, which specifies the optimal measurement for estimating t along the geodesic ρ_t, can be realised entirely within the algebra $\mathcal{A}$. Thus, no information about the parameter t is lost when restricting the measurement to $\mathcal{A}$.

B Gaussian states on degenerate CCR algebras

This appendix provides a brief account of degenerate canonical commutation relations (CCR) and the associated hybrid classical/quantum Gaussian states.

Let V be a real symplectic space with nonsingular symplectic form Δ. The unital $*$-algebra generated by V satisfying

$$fg - gf = \sqrt{-1}\,\Delta(f, g), \quad f^* = f, \qquad (\forall f, g \in V)$$

is called the *canonical commutation relations (CCR) algebra*.

There is another, closely related notion of CCR: let $\mathcal{H}$ be a separable Hilbert space and let $W : V \to B(\mathcal{H})$ satisfy

$$W(f)W(g) = e^{-\frac{\sqrt{-1}}{2}\Delta(f,g)} W(f + g), \quad W(f)^* = W(-f), \qquad (f, g \in V).$$

These are called the *Weyl form* of the CCR, from which it follows that each $W(f)$ is unitary and $W(0) = 1$.

One would like to represent the CCR using selfadjoint operators on $\mathcal{H}$. We first treat the case when V is two-dimensional with symplectic basis $\{e_1, f_1\}$ satisfying $\Delta(e_1, f_1) = 1$. Then the Weyl relations yield

$$W(te_1)W(sf_1) = e^{-\frac{\sqrt{-1}}{2}st}\, W(te_1 + sf_1) = e^{-\sqrt{-1}st}\, W(sf_1)W(te_1).$$

Let $U(t) := W(te_1)$ and $V(s) := W(sf_1)$ be one-parameter unitary groups on $\mathcal{H}$. By Stone's theorem, there is a one-to-one correspondence between strongly continuous one-parameter unitary groups and selfadjoint operators. We define Q and P such that

$$U(t) := e^{\sqrt{-1}tQ}, \quad V(s) = e^{\sqrt{-1}sP},$$

which fulfill the Weyl form of the CCR:

$$e^{\sqrt{-1}tQ}e^{\sqrt{-1}sP} = e^{-\sqrt{-1}st}e^{\sqrt{-1}sP}e^{\sqrt{-1}tQ}.$$

Formally differentiating in t and s at $s = t = 0$ yields the Heisenberg form:

$$QP - PQ = \sqrt{-1}\, I.$$

The operators Q and P are called the *canonical observables* of the CCR.

There are many choices of $\mathcal{H}$ and of irreducible representations of canonical observables. By the Stone–von Neumann theorem, all such representations are unitarily equivalent [17]. In this paper, we keep the Schrödinger representation on $\mathcal{H} = L^2(\mathbb{R})$ in mind. The von Neumann algebra generated by $\{e^{\sqrt{-1}tQ}\}$ is $L^\infty(\mathbb{R})$, and that generated by $\{e^{\sqrt{-1}(tQ+sP)}\}$ is $B(\mathcal{H})$.

The above formulation extends directly to any even-dimensional symplectic space V. Given a regular $(2k) \times (2k)$ real skew-symmetric matrix $S = (S_{ij})$, let $\mathrm{CCR}(S)$ be the von Neumann algebra generated by

$$\left\{e^{\sqrt{-1}\,t_1 X_1}, \dots, e^{\sqrt{-1}\,t_{2k} X_{2k}}\right\}$$

satisfying

$$e^{\sqrt{-1}\,t_i X_i}e^{\sqrt{-1}\,t_j X_j} = e^{\sqrt{-1}\,t_i t_j S_{ij}}e^{\sqrt{-1}(t_i X_i + t_j X_j)},$$

where $X = (X_1, \dots, X_{2k})$ are the canonical observables of CCR(S). Choosing a regular matrix T with

$$T^\top S T = \frac{1}{2}\begin{bmatrix} 0 & -1 & & & & & \\ 1 & 0 & & & & & \\ & & 0 & -1 & & & \\ & & 1 & 0 & & & \\ & & & & \ddots & & \\ & & & & & 0 & -1 \\ & & & & & 1 & 0 \end{bmatrix}$$

provides a symplectic basis $\{e_i, f_i\}_{1\le i\le k}$ generating $\{Q_i, P_i\}_{1\le i\le k}$ so that each X_i is an $\mathbb{R}$-linear combination of them.

For a possibly degenerate $d \times d$ real skew-symmetric matrix $S = (S_{ij})$, choose T so that

$$T^\top S T = \frac{1}{2}\left[\begin{array}{ccc|ccccccc} 0 & & & & & & & & & \\ & \ddots & & & & & & & & \\ & & 0 & & & & & & & \\ \hline & & & 0 & -1 & & & & & \\ & & & 1 & 0 & & & & & \\ & & & & & 0 & -1 & & & \\ & & & & & 1 & 0 & & & \\ & & & & & & & \ddots & & \\ & & & & & & & & 0 & -1 \\ & & & & & & & & 1 & 0 \end{array}\right].$$

The basis then splits into

$$\{\tilde{e}_1, \dots, \tilde{e}_r\} \sqcup \{e_i, f_i\}_{1\le i\le k},$$

where $r + 2k = d$. The $\tilde{e}$-sector corresponds to commuting observables, while the (e, f)-sector corresponds to a non-degenerate CCR. Thus, CCR(S) decomposes naturally into a *commutative part* ($\cong L^\infty(\mathbb{R}^r)$), representing classical variables, and a *noncommutative part*, behaving like an ordinary CCR algebra of dimension $2k$. This decomposition provides the algebraic foundation for hybrid classical/quantum Gaussian states.

The above discussion can also be phrased more directly. Extending $\{\tilde{e}_1, \dots, \tilde{e}_r\}$ to $\{\tilde{e}_i, \tilde{f}_i\}_{1\le i\le r}$ yields a symplectic basis $\{\tilde{e}_i, \tilde{f}_i\}_{1\le i\le r} \sqcup \{e_i, f_i\}_{1\le i\le k}$ for a $2(r+k)$-dimensional symplectic space V, with canonical observables $\{\tilde{Q}_i, \tilde{P}_i\}_{1\le i\le r} \sqcup \{Q_i, P_i\}_{1\le i\le k}$. We then define CCR($S$) as the von Neumann subalgebra generated by

$$\{e^{\sqrt{-1}\,\tilde{t}_i \tilde{Q}_i}\}_{1\le i\le r} \sqcup \{e^{\sqrt{-1}\,t_i Q_i}, e^{\sqrt{-1}\,s_i P_i}\}_{1\le i\le k}.$$

In summary, for any (possibly degenerate) $d \times d$ real skew-symmetric matrix S, $\mathrm{CCR}(S)$ is the algebra generated by $X = (X_1, \ldots, X_d)$ satisfying the Weyl form

$$e^{\sqrt{-1}\,\xi^i X_i} e^{\sqrt{-1}\,\eta^j X_j} = e^{\sqrt{-1}\,\xi^\top S \eta} e^{\sqrt{-1}\,(\xi+\eta)^i X_i} \qquad (\xi, \eta \in \mathbb{R}^d),$$

which is formally rewritten in the Heisenberg form

$$\frac{\sqrt{-1}}{2}[X_i, X_j] = S_{ij} \qquad (1 \le i, j \le d).$$

The above formulation is useful in handling hybrid classical/quantum Gaussian states. Let S be as above, and $X = (X_1, \ldots, X_d)$ the canonical observables of $\mathrm{CCR}(S)$. A state ϕ on $\mathrm{CCR}(S)$ is called a *quantum Gaussian state*, denoted $\phi \sim N(\mu, \Sigma)$, if its characteristic function

$$\mathcal{F}_\xi\{\phi\} := \phi(e^{\sqrt{-1}\xi^i X_i})$$

takes the form

$$\mathcal{F}_\xi\{\phi\} = e^{\sqrt{-1}\xi^i \mu_i - \frac{1}{2}\xi^i \xi^j V_{ij}},$$

where $\xi = (\xi^i)_{i=1}^d \in \mathbb{R}^d$, $\mu = (\mu_i)_{i=1}^d \in \mathbb{R}^d$, and $V = (V_{ij})$ is a real symmetric $d \times d$ matrix such that $\Sigma := V + \sqrt{-1}S$ is positive semidefinite as a Hermitian matrix. When needed, we write $(X, \phi) \sim N(\mu, \Sigma)$ to specify the observables X.

To transfer a Gaussian state between CCR algebras, we use the *quasi-characteristic function* [18]: For $(X, \phi) \sim N(\mu, \Sigma)$ and $\{\xi_t\}_{t=1}^r \subset \mathbb{R}^d$,

$$\phi\left(\prod_{t=1}^r e^{\sqrt{-1}\xi_t^i X_i}\right) = \exp\left(\sum_{t=1}^r \left(\sqrt{-1}\xi_t^i \mu_i - \frac{1}{2}\xi_t^i \xi_t^j \Sigma_{ji}\right) - \sum_{t=1}^r \sum_{u=t+1}^r \xi_t^i \xi_u^j \Sigma_{ji}\right). \tag{5}$$

Expression (5) admits analytic continuation to $\{\xi_t\}_{t=1}^r \subset \mathbb{C}^d$.

The quasi-characteristic function is central in discussing the quantum analogue of weak convergence [18, 19]. Let $\rho^{(n)}$ be a quantum state and $X^{(n)} = (X_1^{(n)}, \ldots, X_d^{(n)})$ observables on a finite dimensional Hilbert space $\mathcal{H}^{(n)}$. We say $(X^{(n)}, \rho^{(n)})$ *weakly converges* to $(X, \phi) \sim N(\mu, \Sigma)$, written

$$(X^{(n)}, \rho^{(n)}) \rightsquigarrow (X, \phi) \qquad \text{or} \qquad X^{(n)} \overset{\rho^{(n)}}{\rightsquigarrow} N(\mu, \Sigma)$$

if for any $r \in \mathbb{N}$ and subset $\{\xi_t\}_{t=1}^r$ of $\mathbb{R}^d$,

$$\lim_{n\to\infty} \mathrm{Tr}\, \rho^{(n)} \left(\prod_{t=1}^r e^{\sqrt{-1}\xi_t^i X_i^{(n)}}\right) = \phi\left(\prod_{t=1}^r e^{\sqrt{-1}\xi_t^i X_i}\right).$$

References

1. van der Vaart, A. W.: Asymptotic Statistics: Cambridge Series in Statistical and Probabilistic Mathematics 3. (Cambridge University Press 1998)

2. Guţă, M., Kahn, J.: Local asymptotic normality for qubit states. Phys. Rev. A **73**, 052108 (2006)
3. Kahn, J., Guţă, M.: Local asymptotic normality for finite dimensional quantum systems. Commun. Math. Phys. **289**, 597–652 (2009)
4. Gill, R. D., Guţă, M.: On asymptotic quantum statistical inference. in From Probability to Statistics and Back: High-Dimensional Models and Processes, Inst. Math. Stat. (IMS) Collect., **9**, 105–127 (2013)
5. Guţă, M., Jenčová, A.: Local asymptotic normality in quantum statistics. Commun. Math. Phys. **276**, 341–379 (2007)
6. Yamagata, K., Fujiwara, A., Gill, R.D.: Quantum local asymptotic normality based on a new quantum likelihood ratio. Ann. Stat. **41**, 2197–2217 (2013)
7. Fujiwara, A., Yamagata, K.: Noncommutative Lebesgue decomposition and contiguity with application to quantum local asymptotic normality. Bernoulli **26**, 2105–2142 (2020)
8. Fujiwara, A., Yamagata, K.: Efficiency of estimators for locally asymptotically normal quantum statistical models. Ann. Stat. **51**, 1159–1182 (2023)
9. Amari, S.-I., Nagaoka, H.: Methods of Information Geometry, Translations of Mathematical Monographs 191. (American Mathematical Society, Providence, RI, 2000); Originally published in Japanese (Iwanami Shoten, Tokyo, 1993)
10. Fujiwara, A.: Foundations of Information Geometry. Kyoritsu Shuppan, Tokyo (2021). (**in Japanese**)
11. Williams, D.: Probability with Martingales: Cambridge Mathematical Textbooks. (Cambridge University Press, 1991)
12. Ohya, M., Petz, D.: Quantum Entropy and Its Use (Springer, 1993)
13. Petz, D.: Quantum Information Theory and Quantum Statistics (Springer, 2008)
14. Torgersen, E.: Comparison of Statistical Experiments (Cambridge University Press, 1991)
15. Petz, D.: Monotone metrics on matrix spaces. Linear Algebra Appl. **224**, 81–96 (1996)
16. Sakai, S.: A Radon-Nikodym theorem in W^*-algebras. Bull. Am. Math. Soc. **71**, 149–151 (1965)
17. Holevo, A. S.: Probabilistic and Statistical Aspects of Quantum Theory (Edizioni della Normale, 2011); originally published by North-Holland (1982)
18. Jakšić, V., Pautrat, Y., Pillet, C.A.: A quantum central limit theorem for sums of independent identically distributed random variables. J. Math. Phys. **51**, 015208 (2010)
19. Jakšić, V., Pautrat, Y., Pillet, C.A.: A non-commutative Lévy-Cramér continuity theorem. Markov Process. Related Fields **16**, 59–78 (2010)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.