POLITECNICO MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE E DELL'INFORMAZIONE

Executive Summary of the Thesis

Multi-Parameter Quantum Estimation

Laurea Magistrale in Mathematical Engineering - Ingegneria Matematica

Author: Riccardo Bonetti

Advisor: Prof. Michele Correggi

Co-advisor: Angelo Lucia

Academic year: 2025-2026

1. Introduction

Quantum metrology studies how the laws of quantum mechanics constrain, and can enhance, the precision with which unknown physical parameters can be estimated. A generic metrological protocol consists of preparing a probe state, imprinting the parameter(s) through a parameter-dependent dynamics, and extracting information from repeated measurement outcomes. The ultimate attainable precision is captured by Cramér–Rao type inequalities: in the classical setting, through the Fisher information of the induced statistical model; in the quantum setting, through the quantum Fisher information matrix (QFIM), obtained by optimizing over all possible measurements [2, 3]. For single-parameter problems, the quantum bound is typically saturable under mild regularity assumptions, and the main question becomes how to engineer probe states and dynamics that yield a favorable scaling with available resources [1].

Multi-parameter estimation is qualitatively different. When several parameters are encoded in the same probe, the measurements that are optimal for different parameters can be incompatible, so that no single measurement may simultaneously attain all single-parameter quantum limits. This leads to a genuinely quantum trade-off between the achievable precisions, which is naturally described at the matrix level via the QFIM and multi-parameter quantum Cramér–Rao bounds [3–5]. Understanding when the quantum bound is attainable, and constructing explicit measurements that achieve it (at least locally around a working point), are therefore central questions in quantum multi-parameter metrology.

This thesis develops a local and constructive framework for multi-parameter quantum metrology, with an emphasis on collective spin probes relevant for magnetometry. On the theoretical side, we review and exploit saturation conditions for the matrix quantum Cramér–Rao bound, highlighting the role of compatibility (in particular, weak commutativity conditions at the working point) in determining whether the QFIM is jointly attainable. On the constructive side, we provide an explicit recipe to build locally optimal positive operator-valued measures (POVMs) from the geometry of the tangent space of the quantum statistical model at a chosen working point, and we instantiate this construction for paradigmatic probe families in the symmetric subspace.

The practical performance of the resulting measurements is assessed numerically. We compare single-parameter and two-parameter estimation settings using representative probes (including symmetric Dicke states and GHZ-type probes), and we validate local optimality by Monte Carlo simulations, showing agreement between the empirical covariance matrix and the corresponding Cramér–Rao benchmarks in the asymptotic regime. For the two-parameter setting, we present compact visual diagnostics (confidence ellipses and eigenvalue-based summaries) that make the attainability and anisotropy of the precision trade-off transparent.

The executive summary is organized as follows. Section 2 recalls only the essential notions of classical and quantum parameter estimation (POVMs, Fisher information, and Cramér–Rao bounds). Section 3 outlines the local constructive methodology and the models studied. Sections 4 and 5 summarize the main results for the single-parameter and two-parameter scenarios, respectively, using the most representative figures and tables from the thesis. Finally, Section 6 reports the main conclusions and discusses limitations and directions for future work.

2. Background essentials

This section collects only the basic notions required to interpret the results reported later: (i) classical and quantum statistical models induced by measurements, (ii) Fisher information and Cramér–Rao bounds, and (iii) the key difference between single-parameter and multi-parameter estimation.

2.1. Classical parameter estimation and Fisher information

Let m denote the outcome of an experiment with probability distribution $p(m|\boldsymbol{\theta})$, where $\boldsymbol{\theta} \in \mathbb{R}^d$ is the unknown parameter vector. An estimator $\hat{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}}(m)$ has covariance matrix

$$\mathrm{Cov}(\hat{\boldsymbol{\theta}}) = \mathbb{E}\big[(\hat{\boldsymbol{\theta}} - \mathbb{E}[\hat{\boldsymbol{\theta}}])(\hat{\boldsymbol{\theta}} - \mathbb{E}[\hat{\boldsymbol{\theta}}])^T\big]. \tag{1}$$

For a (locally) unbiased estimator at $\boldsymbol{\theta}$, the classical Cramér–Rao bound states

$$\mathrm{Cov}(\hat{\boldsymbol{\theta}}) \geq \mathbf{F}_C(\boldsymbol{\theta})^{-1}, \tag{2}$$

$$\big[\mathbf{F}_C(\boldsymbol{\theta})\big]_{ij} = \sum_m \frac{\partial_i \ln p(m|\boldsymbol{\theta}) \cdot \partial_j \ln p(m|\boldsymbol{\theta})}{p(m|\boldsymbol{\theta})}, \tag{3}$$

where $\partial_i \equiv \partial/\partial\theta_i$ and $\geq$ denotes the Loewner order on symmetric matrices. For independent repetitions (n shots) of the same experiment, Fisher information adds: $\mathbf{F}_C^{(n)} = n\,\mathbf{F}_C$, hence the bound scales as $1/n$.

2.2. Quantum statistical models and POVMs

In quantum metrology the unknown parameter(s) are encoded in a quantum state $\rho_{\boldsymbol{\theta}}$ (density operator on a Hilbert space $\mathcal{H}$). A general measurement is described by a POVM $\{\Pi_m\}_m$ with

$$\Pi_m \geq 0, \qquad \sum_m \Pi_m = I. \tag{4}$$

The measurement induces a classical statistical model via the Born rule

$$p(m|\boldsymbol{\theta}) = \mathrm{Tr}\big(\rho_{\boldsymbol{\theta}} \Pi_m\big). \tag{5}$$

Given $\{\Pi_m\}$, one can compute the corresponding classical Fisher information matrix $\mathbf{F}_C(\boldsymbol{\theta}|\{\Pi_m\})$ from Eq. (3) using the probabilities in Eq. (5).

2.3. Quantum Fisher information and the quantum Cramér–Rao bound

The quantum Fisher information matrix (QFIM) $\mathbf{F}_Q(\boldsymbol{\theta})$ is defined as the maximum of the classical Fisher information over all POVMs,

$$\mathbf{F}_Q(\boldsymbol{\theta}) = \sup_{\{\Pi_m\}} \mathbf{F}_C(\boldsymbol{\theta}|\{\Pi_m\}). \tag{6}$$

The supremum is meant in the Loewner sense (least upper bound): $\mathbf{F}_Q(\boldsymbol{\theta}) \geq \mathbf{F}_C(\boldsymbol{\theta}|\{\Pi_m\})$ for all POVMs. It then yields the measurement-independent quantum Cramér–Rao bound (QCRB) for locally unbiased estimators at $\boldsymbol{\theta}$

$$\mathrm{Cov}(\hat{\boldsymbol{\theta}}) \geq \mathbf{F}_Q(\boldsymbol{\theta})^{-1} \tag{7}$$

In the symmetric logarithmic derivative (SLD) formulation, the SLDs L_i are defined implicitly by

$$\partial_i \rho_{\boldsymbol{\theta}} = \frac{1}{2}\big(\rho_{\boldsymbol{\theta}} L_i + L_i \rho_{\boldsymbol{\theta}}\big), \tag{8}$$

and the QFIM can be written as

$$\big[\mathbf{F}_Q(\boldsymbol{\theta})\big]_{ij} = \mathrm{Re}\,\mathrm{Tr}\big(\rho_{\boldsymbol{\theta}} L_i L_j\big). \tag{9}$$

In this thesis we mostly consider unitary encodings around a working point,

$$\rho_{\boldsymbol{\theta}} = U(\boldsymbol{\theta})\,\rho_0\,U(\boldsymbol{\theta})^\dagger, \tag{10}$$

$$U(\boldsymbol{\theta}) = \exp\Big(-i\sum_{i=1}^d \theta_i G_i\Big), \tag{11}$$

where G_i are Hermitian generators. For pure probes $\rho_0 = |\psi_0\rangle\langle\psi_0|$, the QFIM has the convenient expression

$$\big[\mathbf{F}_Q(\boldsymbol{\theta})\big]_{ij} = 4\,\mathrm{Cov}_{\psi_{\boldsymbol{\theta}}}(G_i, G_j) \tag{12}$$
$$= 4\Big(\langle G_i G_j\rangle - \langle G_i\rangle\langle G_j\rangle\Big)_{\psi_{\boldsymbol{\theta}}}, \tag{13}$$

with expectations taken on $|\psi_{\boldsymbol{\theta}}\rangle = U(\boldsymbol{\theta})|\psi_0\rangle$.

2.4. Single-parameter vs multi-parameter: attainability and incompatibility

For $d = 1$, it is often possible (under standard regularity assumptions) to find measurements and estimators that attain the QCRB asymptotically, so that the QFIM provides the relevant benchmark for precision. For $d > 1$, however, two distinct issues arise:

(i) Joint attainability. Even if each parameter has an individually optimal measurement, those measurements may be incompatible. As a result, there may exist no single POVM such that $\mathbf{F}_C(\boldsymbol{\theta}|\{\Pi_m\}) = \mathbf{F}_Q(\boldsymbol{\theta})$ as matrices.

(ii) Local vs global design. A common and practically relevant strategy is to target *local* optimality at a working point $\boldsymbol{\theta}^\star$, where the experiment is operated and where the measurement is designed. In this local approach, one seeks a POVM $\{\Pi_m\}$ such that the induced Fisher information saturates the quantum bound at $\boldsymbol{\theta}^\star$ (or, equivalently, the covariance approaches $\mathbf{F}_Q(\boldsymbol{\theta}^\star)^{-1}$ in the large-shot regime).
A widely used compatibility condition ensuring local attainability of the SLD-based bound is *weak commutativity* at the working point,

$$\mathrm{Tr}\big(\rho_{\boldsymbol{\theta}^\star}[L_i, L_j]\big) = 0 \qquad \forall\, i, j, \tag{14}$$

which expresses that the SLDs commute "on average" with respect to $\rho_{\boldsymbol{\theta}^\star}$. When Eq. (14) holds, one can construct measurements that are jointly optimal at $\boldsymbol{\theta}^\star$ within a local (tangent-space) description of the model. This is the regime exploited in the two-parameter constructions and numerical validations reported in Sections 4–5.

3. Methodology and models

We adopt a local approach around a working point $\boldsymbol{\theta}^\star$. For a fixed probe ρ_0 and encoding map $\boldsymbol{\theta} \mapsto \rho_{\boldsymbol{\theta}}$, we compute the quantum benchmark $\mathbf{F}_Q(\boldsymbol{\theta}^\star)$ and design (or select) a measurement $\{\Pi_m\}$ such that the induced classical Fisher information $\mathbf{F}_C(\boldsymbol{\theta}^\star|\{\Pi_m\})$ matches $\mathbf{F}_Q(\boldsymbol{\theta}^\star)$ whenever compatibility allows it.

3.1. Local design principle

At $\boldsymbol{\theta}^\star$, the model is characterized by its first-order variations. In the pure-state cases considered in this thesis, the relevant information is contained in the two-dimensional tangent space spanned by the orthogonal components of $\partial_i|\psi_{\boldsymbol{\theta}}\rangle|_{\boldsymbol{\theta}^\star}$. When the tangent geometry at $\boldsymbol{\theta}^\star$ is real (as ensured by the compatibility conditions discussed in the thesis), we construct a locally optimal POVM by: (i) orthonormalizing the tangent directions (whitening), (ii) implementing a minimal symmetric POVM on the resulting plane, and (iii) completing it with a residual element on the orthogonal complement. This yields a measurement that is locally optimal at $\boldsymbol{\theta}^\star$ by construction in the compatible regime.

3.2. Numerical validation

Numerical evidence is obtained by simulating repeated measurements and forming an estimator from outcome frequencies. Performance is reported through matrix-level comparisons (e.g. eigenvalues and traces) between the empirical covariance $\mathrm{Cov}(\hat{\boldsymbol{\theta}})$ and the relevant Cramér–Rao benchmark. In the two-parameter case, confidence ellipses extracted from these matrices provide a compact visualization of attainability and anisotropy.

3.3. Models considered

We consider an ensemble of N identical spin-$\frac{1}{2}$ particles (two-level systems) employed as quantum probes to estimate unknown parameters of a magnetic field. The parameter dependence is encoded through a collective unitary evolution acting on an initial probe state ρ_0,

$$\rho_{\boldsymbol{\theta}} = U_{\boldsymbol{\theta}}\,\rho_0\,U_{\boldsymbol{\theta}}^\dagger, \tag{15}$$
$$U_{\boldsymbol{\theta}} = \exp(-i\boldsymbol{\gamma}(\boldsymbol{\theta})\cdot\mathbf{J})\,, \tag{16}$$

where $J_\alpha = \frac{1}{2}\sum_{n=1}^{N}\sigma_\alpha^{(n)}$ denote the components of the collective (total) spin operator, and T is the interrogation time.
Assuming that individual particles cannot be addressed or detected separately remains confined

to the totally symmetric sector.
We can thus formulate the model in the collective spin-j representation associated to $\mathbf{J}$

Single-parameter setting. We consider phase (field) estimation with a single unknown parameter, focusing on:
• *symmetric Dicke probes*, i.e. states with m qubits in $|1\rangle$ and $N-m$ in $|0\rangle$:

$$|D_N^{(m)}\rangle = \binom{N}{m}^{-\frac{1}{2}} \sum_{k=1}^{N} P_k \left(|1\rangle^{\otimes m} \otimes |0\rangle^{\otimes (N-m)}\right);$$

• *GHZ probes*, i.e., superpositions of states with all qubits in $|0\rangle$ and all qubits in $|1\rangle$:

$$|\mathrm{GHZ}\rangle = \frac{1}{\sqrt{2}} \left(|0\rangle^{\otimes N} + |1\rangle^{\otimes N}\right).$$

Results are summarized together with the corresponding readout strategies adopted in the thesis, and compared against the single-parameter quantum benchmark.

Two-parameter setting. We consider a two-parameter model where two generators jointly encode $\boldsymbol{\theta} = (\theta_1, \theta_2)$. The focus is on (i) verifying compatibility at the working point and (ii) exhibiting an explicit locally optimal POVM via the tangent-space construction. Results are presented through the associated confidence ellipses and spectral diagnostics of the relevant information and covariance matrices.

4. Single-parameter estimation

We first benchmark the single-parameter setting at fixed system size $N = 50$, reporting both bias and variance as a function of the number of shots per experiment ν for the four strategies considered (Models A, B, Copt, Ccoh). The four strategies share the same parameter-encoding dynamics, Monte Carlo sampling procedure, and performance metrics; they differ only in the choice of probe state and in the corresponding readout, i.e., the measurement and the estimator used to infer the parameter from the measurement outcomes. This controlled setup enables a direct, like-for-like comparison.
Model A employs a maximally entangled GHZ state in the symmetric spin-j manifold and a parity-type (or equivalent optimal) collective measurement.
Model B (symmetric Dicke) uses a Dicke state (with fixed total magnetization, typically chosen near $m \simeq 0$ for metrological advantage) together with a collective population-imbalance readout.
Model Copt (optimized entangled probe) is defined by a numerically optimized symmetric probe state obtained by maximizing a chosen information figure of merit for the given encoding and constraints.
Model Ccoh (coherent-state baseline) uses a spin-coherent state polarized in the z direction (product state) as a classical reference, with a standard collective readout such as J_z.
For each model we compare Monte Carlo (MC) estimates with the corresponding theoretical predictions, and we use the quantum Cramér–Rao bound (QCRB) as the ultimate variance benchmark.

4.1. Variance and attainability of the QCRB

Figure 1 shows $\mathrm{Var}(\hat{\theta})$ versus ν (markers) together with the corresponding QCRB (dashed curves). The data display the expected $1/\nu$ scaling across the explored range, and, for each model, the MC points track the dashed QCRB curve closely, indicating that the adopted measurement–estimator pair is asymptotically efficient in the regime considered. In particular, the relative ordering of the variances is stable across ν: Model A yields the smallest variance, Models B and Copt are intermediate and close to each other, while Model Ccoh is consistently larger. This confirms, at the working configuration used in the simulation, a clear separation in achievable precision between the coherent baseline and the other strategies, and it provides a compact quantitative summary of the single-parameter performance at $N = 50$.

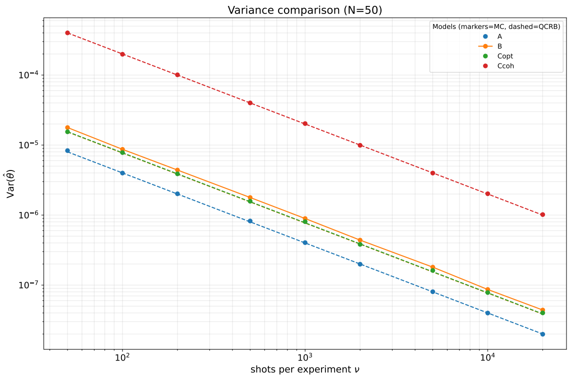

Figure 1: Variance comparison at $N = 50$ as a function of the shots per experiment ν. Markers: Monte Carlo estimates of $\mathrm{Var}(\hat{\theta})$. Dashed curves: corresponding QCRB. Models: A, B, Copt, Ccoh.

4.2. Bias and finite-sample effects

Figure 2 reports the bias $\mathbb{E}[\hat{\theta} - \theta]$ versus ν, comparing MC (markers) with theoretical predictions (dashed). Across all models the bias approaches zero as ν increases, consistent with the expected asymptotic behavior of the estimators used in the large-shot regime. At low and intermediate ν, finite-sample effects are visible and differ among models: the coherent strategy (Ccoh) shows the largest deviations from zero at small ν, while the other strategies exhibit smaller and more rapidly suppressed biases. As ν reaches the largest values explored, the MC points cluster near zero and align with the theoretical curves, supporting the internal consistency between the analytical approximations and the simulated estimation pipeline.

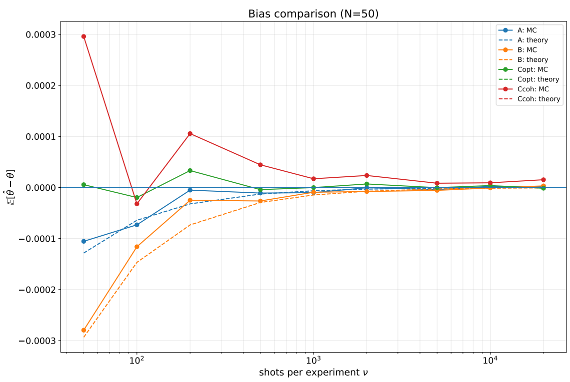

Figure 2: Bias comparison at $N = 50$ as a function of the shots per experiment ν. Solid markers: Monte Carlo estimates of $\mathbb{E}[\hat{\theta} - \theta]$. Dashed curves: corresponding theoretical predictions. Models: A, B, Copt, Ccoh.

4.3. Takeaway at fixed N

At $N = 50$, the single-parameter results provide a clean separation between (i) asymptotically efficient strategies whose variance closely follows the QCRB with a stable advantage (Model A and, at intermediate level, Models B and Copt) and (ii) a coherent baseline (Ccoh) with markedly larger variance and stronger finite-ν bias at small shot numbers. These plots serve as the reference point for the multi-parameter analysis, where attainability becomes conditional on compatibility and must be assessed at the matrix level.

5. Results: two-parameter estimation

We now summarize the two-parameter setting at the working point $\boldsymbol{\theta}^\star = (0, 0)$. In this regime the model is compatible in the sense discussed in the thesis, so that the SLD-based quantum bound is jointly attainable locally and an explicit locally optimal POVM can be constructed via the tangent-space recipe outlined in Section 3. The numerical goal is therefore matrix-level saturation:

$$\mathrm{Cov}(\hat{\boldsymbol{\theta}}) \approx \mathbf{F}_Q(\boldsymbol{\theta}^\star)^{-1}/\nu \qquad \text{for large } \nu, \tag{17}$$

which we assess through confidence ellipses and spectral diagnostics.

5.1. Attainability visualized by confidence ellipses

Figures 3 and 4 compare, at fixed $N = 14$ and $\nu = 10^5$ shots, the error ellipse obtained from the empirical covariance matrix (Monte Carlo) with the ellipse predicted by the QCRB. In both cases the two ellipses are nearly indistinguishable, providing a direct visual confirmation that the constructed measurement achieves local saturation at $\boldsymbol{\theta}^\star$ in the asymptotic regime. The different shapes of the ellipses across probes reflect the anisotropic sensitivity of the two-parameter model, i.e. different principal directions corresponding to different achievable variances.

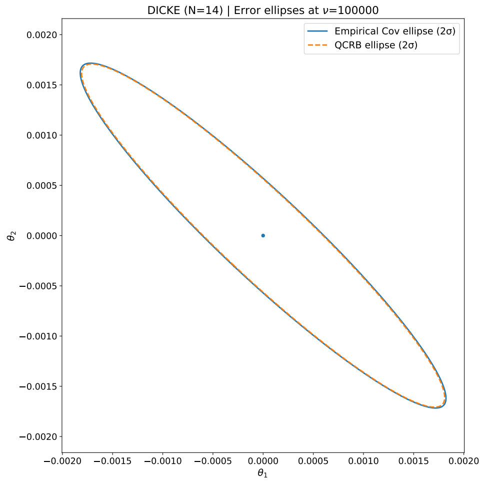

Figure 3: Dicke probe ($N = 14$): empirical 2σ error ellipse (from Cov) versus QCRB ellipse at $\nu = 10^5$.

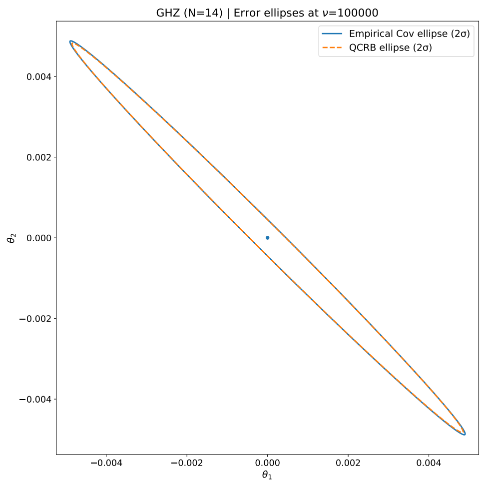

Figure 4: GHZ-type probe ($N = 14$): empirical 2σ error ellipse (from Cov) versus QCRB ellipse at $\nu = 10^5$.

5.2. Spectral convergence to the QCRB

A complementary, quantitative diagnostic is provided by the eigenvalues of the matrix difference $\mathrm{Cov}(\hat{\boldsymbol{\theta}}) - \mathbf{F}_Q(\boldsymbol{\theta}^\star)^{-1}/\nu$. Figures 5 and 6 report the minimum and maximum eigenvalues of this difference as functions of ν (log scale on the horizontal axis). For both probes, the eigenvalues converge to zero as ν increases, indicating convergence of the full covariance matrix to the QCRB prediction rather than only of a scalar summary. At small shot numbers, deviations are larger and probe-dependent, reflecting finite-sample effects and estimator non-idealities; these become negligible in the large-ν regime where local asymptotic theory applies.

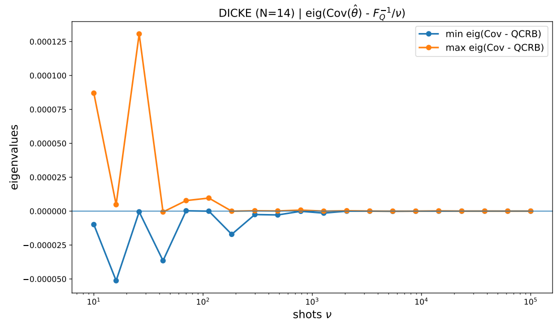

Figure 5: Dicke probe ($N = 14$): eigenvalues of $\mathrm{Cov}(\hat{\boldsymbol{\theta}}) - \mathbf{F}_Q^{-1}/\nu$ versus ν.

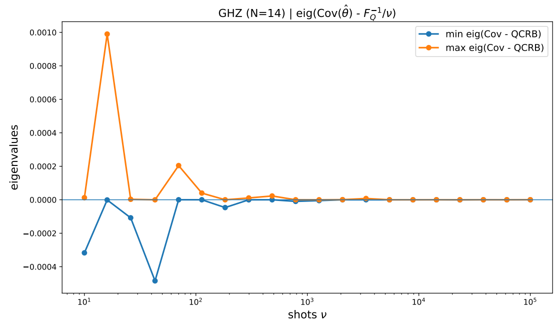

Figure 6: GHZ-type probe ($N = 14$): eigenvalues of $\mathrm{Cov}(\hat{\boldsymbol{\theta}}) - \mathbf{F}_Q^{-1}/\nu$ versus ν.

6. Future Perspectives

Among the extensions we plan to consider, we mention the investigation of robustness beyond a single working point, inclusion of noise and decoherence, development of tailored multiparameter maximum-likelihood estimators, and translation of the locally optimal POVMs into experimentally feasible measurement schemes under realistic constraints.

7. Acknowledgements

I would like to thank my advisor, Prof. Michele Correggi, for his guidance and support throughout this thesis. I am also grateful to my co-advisor, Angelo Lucia, for his helpful feedback and valuable discussions.

References

[1] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Quantum-enhanced mea-

surements: Beating the standard quantum limit. *Science*, 306(5700):1330–1336, 2004.

[2] Carl W. Helstrom. *Quantum Detection and Estimation Theory*. Academic Press, New York, 1976.

[3] Matteo G. A. Paris. Quantum estimation for quantum technology. *International Journal of Quantum Information*, 7(supp01):125–137, 2009.

[4] Magdalena Szczykulska, Tillmann Baumgratz, and Animesh Datta. Multi-parameter quantum metrology. *Advances in Physics: X*, 1(4):621–639, 2016.

[5] Jing Yang, Shengshi Pang, Yiyu Zhou, and Andrew N. Jordan. Optimal measurements for quantum multiparameter estimation with general states. 2019. arXiv:1806.07337v3 (dated: September 10, 2019).