

Quantum Tomography

Asymptotic Theory and Statistical Methodology



University of
Nottingham
UK | CHINA | MALAYSIA

Anirudh Acharya

School of Mathematical Sciences
University of Nottingham

This dissertation is submitted for the degree of
Doctor of Philosophy

February 2018

I dedicate this thesis to my parents, who have been a constant and unfailing source of support and love, and to whom I owe everything.

Abstract

Recent experimental progress in the preparation and control of quantum systems has brought to light the importance of *Quantum State Tomography* (QST) in validating the results. In this thesis we investigate several aspects of QST, whose central problem is to devise estimation schemes for the recovery of an unknown state, given an ensemble of n independent identically prepared systems.

The key issues in tackling QST for large dimensional systems is the construction of physically relevant low dimensional state models, and the design of appropriate measurements. Inspired by compressed sensing tomography, in chapters 4, 5 we consider the statistical problem of estimating low rank states ($r \ll d$) in the set-up of *Multiple Ions Tomography* (MIT), where r and d are the rank and the dimension of the state respectively. We investigate how the estimation error behaves with a reduction in the number of measurement settings, compared to ‘full’ QST in two setups - Pauli and random bases measurement designs. We study the estimation errors in this ‘incomplete’ measurement setup in terms of a concentration of the Fisher information matrix. For the random bases design we demonstrate that $O(r \log d)$ settings suffice for the *mean square error* w.r.t the Frobenius norm to achieve the optimal $O(1/n)$ rate of estimation.

When the error functions are locally quadratic, like the Frobenius norm, then the expected error (or risk) of standard procedures achieves this optimal rate. However, for fidelity based errors such as the Bures distance we show that no ‘compressive’ recovery exists for states close to the boundary, and it is known that even with conventional ‘full’ tomography schemes the risk scales as $O(1/\sqrt{n})$ for such states and error functions. For qubit states this boundary is the surface of the Bloch sphere. Several estimators have been proposed to improve this scaling with ‘adaptive’ tomography. In chapter 6 we analyse this problem from the perspective of the maximum Bures risk over all qubit states. We propose two adaptive estimation strategies, one based on

local measurements and another based on *collective* measurements utilising the results of quantum local asymptotic normality. We demonstrate a scaling of $O(1/n)$ for the maximum Bures risk with both estimation strategies, and also discuss the construction of a minimax optimal estimator.

In chapter 7 we return to the MIT setup and systematically compare several tomographic estimators in an extensive simulation study. We present and analyse results from this study, investigating the performance of the estimators across various states, measurement designs and error functions. Along with commonly used estimators like maximum likelihood, we propose and evaluate a few new ones. We finally introduce two web-based applications designed as tools for performing QST simulations online.

Publications

This thesis consists of both published and unpublished works. The results in chapters 4, 5, 6 have appeared in the following three papers respectively:

- **Anirudh Acharya**, Theodore Kypraios and Mădălin Guță (2016). Statistically efficient tomography of low rank states with incomplete measurements. *New Journal of Physics*, 18(4):043018. <http://stacks.iop.org/1367-2630/18/i=4/a=043018>
- **Anirudh Acharya** and Mădălin Guță (2017). Statistical analysis of compressive low rank tomography with random measurements. *Journal of Physics A: Mathematical and Theoretical*, 50(19):195301. <http://stacks.iop.org/1751-8121/50/i=19/a=195301>
- **Anirudh Acharya** and Mădălin Guță (2017). Minimax estimation of qubit states with Bures risk. arXiv:1708.04941. <https://arxiv.org/abs/1708.04941>

A paper based on chapter 7 is in the final stages of its preparation.

आचार्यात् पादमादत्ते पादं शिष्यः स्वमेधया ।

पादं सब्रह्मचारिभ्यः पादं कालक्रमेण च ॥

A student receives a quarter of his learning from his teacher, a quarter by virtue of his own intelligence, a quarter from his fellow learners, and a quarter from the passage of time.

Acknowledgements

This thesis is the product of research carried out under the supervision of Dr.Mădălin Guță and Dr.Theodore Kypraios, to both of whom I am indebted for all that they have taught me. It would not have been possible to write this thesis without their encouragement, support and guidance. I would specially like to thank Mădălin, who first introduced me to the field of quantum information science in his wonderful MSc module several years ago. He has been instrumental in making it possible for me to pursue a PhD here at Nottingham.

I would also like to thank my colleagues and staff from the department who have helped make this experience enjoyable. In particular I wish to thank Vladimir, who has been a friend, tennis partner and walking companion with whom I have explored much of the city. I would also like to express my gratitude to the friends, old and new, who have helped me through some of the more challenging times. RS, NG, VC, MA and BS, I thank you for everything.

Finally, I cannot begin to express the gratitude I feel for the endless love and support of my parents. This work is dedicated to them.

Contents

List of Figures	xi
List of Tables	xiv
Nomenclature	xv
I	1
1 Introduction	2
2 Quantum State Tomography	7
2.1 Notations and mathematical background	7
2.1.1 Quantum states and measurements	7
2.1.2 Composite systems, entanglement and collective measurements	9
2.2 Example: Continuous variable systems	10
2.3 Estimation of finite-dimensional quantum states	13
2.3.1 Multiple Ions Tomography	15
2.3.2 The Least Squares Estimator	16
2.3.3 Maximum Likelihood Estimator	18
2.4 Fisher Information and Local Asymptotic Normality	19
2.4.1 Classical Fisher Information	19
2.4.2 Classical Local Asymptotic Normality	22
2.4.3 Quantum Fisher Information	23
2.4.4 Quantum Local Asymptotic Normality	24
2.5 The Minimax estimator	26
2.5.1 Loss functions and quadratic approximations	29
3 Thesis rationale	32
3.1 Quantum tomography via compressed sensing	32

3.1.1	Coarse grained statistics	35
3.2	The projection estimator	36
3.2.1	Alternating projections	39
3.3	Compressive tomography with fine grained statistics	40
3.3.1	A problem at the boundary	41
3.4	Adaptive protocols and the minimax rate	42
3.5	A simulation study	46
II		48
4	Statistically efficient tomography of low rank states with in-	
	complete measurements	49
4.1	Introduction	49
4.2	Multiple Ions Tomography with Incomplete Measurements	52
4.3	Numerical Simulations	55
4.4	A Concentration Bound for the MSE	57
4.4.1	Pauli settings	59
4.5	Coarse vs Fine Grained Models	63
4.6	Conclusions	65
4.7	Appendix	66
4.8	Proof of Theorem 2	66
5	Statistical analysis of compressive low rank tomography with	
	random measurements	72
5.1	Introduction	72
5.2	Quantum tomography with random basis measurements	75
5.3	Bounds for the MSE	77
5.4	The Single Qubit Model	82
5.4.1	Estimation Error in Terms of Quantum Infidelity	87
5.5	Conclusions	91
5.6	Appendix	93
5.6.1	Proof of Lemma 1	93
5.6.2	Proof of Theorem 4	94
6	Minimax estimation of qubit states with Bures risk	98
6.1	Introduction	98
6.2	Estimator based on local adaptive measurements	102

6.2.1	An n^{-1} scaling upper bound scaling	104
6.3	Irreducible representations, collective measurements and local asymptotic normality	106
6.3.1	Preliminary localisation and parametrisation	106
6.3.2	The ‘which block’ measurement stage	108
6.3.3	Local asymptotic normality	110
6.4	Minimax upper and lower bounds	112
6.4.1	The upper bound	113
6.4.2	The lower bound	115
6.5	The minimax optimal estimator	117
6.6	Quantum Relative Entropy	120
6.7	Conclusion	122
6.8	Appendix	123
6.8.1	Proof of Lemma 5	123
6.8.2	Expansion of the square Bures distance	123
6.8.3	Proof of Theorem 8	124
6.8.4	Expansion of Quantum Relative Entropy	129
7	Comparison of estimation methods in quantum state tomography: a simulation study	131
7.1	Introduction	131
7.2	A review of multiple ions tomography	133
7.3	Estimators	134
7.3.1	Maximum Likelihood (ML)	134
7.3.2	Least Squares (LS)	135
7.3.3	Generalised Least Squares (GLS)	137
7.3.4	Thresholded Least Squares (TLS)	138
7.3.5	Thresholded Generalised Least Squares (TGLS)	139
7.3.6	Positive Least Squares (PLS)	139
7.3.7	Positive Generalised Least Squares (PGLS)	140
7.4	The Fisher information matrix	140
7.5	Numerical simulations	142
7.5.1	Error/Loss functions	143
7.5.2	Implementation of the estimators	144
7.6	Simulation results and analysis	145
7.6.1	Square Frobenius norm	146
7.6.2	Square Bures distance	148

7.6.3	Square Hellinger distance and Trace norm	149
7.6.4	Special states	149
7.7	The Apps	172
7.7.1	Dashboard: Comparing Estimators	172
7.7.2	State Estimation	175
7.8	Conclusion	175
8	Conclusions and directions for future work	177
	Bibliography	179

List of Figures

2.1	Fisher information for a Bernoulli trial	20
3.1	The error of the ‘projection’ estimator plotted for various number of settings measured	38
3.2	The standard deviation ellipses for different rebit states	43
4.1	Asymptotic MSE for 4 ion states of various ranks, plotted against the number of measured settings	51
4.2	The MSE of the MLE relative to its asymptotic value for the incomplete measurement setup	56
4.3	Concentration of the Fisher information matrix for the reduced measurement setup	60
4.4	The relative error for the random settings and pure states of 3-6 qubits	61
4.5	The relative error for the Pauli settings and a random 4 qubit pure state	62
4.6	Box-plots of the minimum eigenvalues of the Fisher matrix for the Pauli settings	63
4.7	Coarse grained vs fine grained models	64
5.1	Dependance of the Fisher information concentration on the eigenvalue of the state	81
5.2	Concentration of the MSE for qubit states of various eigenvalues	83
5.3	The difference in scaling between the MSE and the MINF for a pure qubit state	89
6.1	Hellinger risk for the estimation of the Poisson parameter	119
7.1	Mean square Frobenius error (MSE) for the Pauli measurement design	151

7.2	Mean square Frobenius error (MSE) for the random bases measurement design and 3 qubits	152
7.3	Mean square Frobenius error (MSE) for the random bases measurement design and 4 qubits	153
7.4	Mean square Bures error (MSB) for the Pauli measurement design	154
7.5	Mean square Bures error (MSB) for the random bases measurement design and 3 qubits	155
7.6	Mean square Bures error (MSB) for the random bases measurement design and 4 qubits	156
7.7	Mean square Hellinger error (MSH) for the Pauli measurement design	157
7.8	Mean square Hellinger error (MSH) for the random bases measurement design and 3 qubits	158
7.9	Mean square Hellinger error (MSH) for the random bases measurement design and 4 qubits	159
7.10	Mean trace norm error for the Pauli measurement design	160
7.11	Mean trace norm error for the random bases measurement design and 3 qubits	161
7.12	Mean trace norm error for the random bases measurement design and 4 qubits	162
7.13	Estimation errors for the Pauli measurement design and the state with exponentially decaying eigenvalues	163
7.14	Estimation errors for the random bases measurement design and the state with exponentially decaying eigenvalues	164
7.15	Estimation errors for the Pauli measurement design and the state with two dominant eigenvalues	165
7.16	Estimation errors for the random measurement design and the state with two dominant eigenvalues	166
7.17	Estimation errors for the Pauli measurement design and the state with four dominant eigenvalues	167
7.18	Estimation errors for the random measurement design and the state with four dominant eigenvalues	168
7.19	Estimation errors for the Pauli measurement design and the GHZ state	169
7.20	Estimation errors for the random bases measurement design and the GHZ state	170

7.21	Estimation errors for the Pauli measurement design and a random pure state	171
7.22	An application for performing tomographic simulations online .	173
7.23	An application for performing tomographic simulations online - output	174
7.24	An application to estimate an unknown state from the outcome dataset of the tomographic experiment	175
7.25	An application to estimate an unknown state from the outcome dataset of the tomographic experiment - output	176

List of Tables

- 5.1 The mean and range of the elements of the Fisher matrix $I(\rho|\mathbf{s})$ as functions of λ_2 . Note that the expressions for the means in the table above are valid for all $\lambda_2 < 0.5$. When $\lambda_2 = 0.5$, then all diagonal elements $\bar{I}^{rr/dd/ii}$ have the same value of $4/3$ 84
- 7.1 The different error functions used. The Bures distance is defined only for states $\hat{\rho}, \rho \in \mathbb{S}_d$, and its classical analogue the Hellinger distance is defined between two probability distributions. . . . 143

Nomenclature

Acronyms / Abbreviations

CLT Central Limit Theorem

CRLB Cramér-Rao Lower Bound

CS Compressed Sensing

FIM Fisher Information Matrix

GHZ Greenberger-Horne-Zeilinger

GLS Generalised Least Squares

KL Kullback-Leibler

LAN Local Asymptotic Normality

LSE Least Squares Estimator

LS Least Squares

MINF Mean Infidelity

MIT Multiple Ions Tomography

MLE Maximum Likelihood Estimator

ML Maximum Likelihood

MSB Mean Square Bures

MSE Mean Square Error

MSH Mean Square Hellinger

ONB Orthonormal Basis

PGLS Positive Generalised Least Squares

PLS Positive Least Squares

POVM Positive Operator-Valued Measure

PVM Projection-Valued Measure

QFI Quantum Fisher Information

QRE Quantum Relative Entropy

QST Quantum State Tomography

RIP Restricted Isometry Property

SLD Symmetric Logarithmic Derivative

TGLS Thresholded Generalised Least Squares

TLS Thresholded Least Squares

Part I

Chapter 1

Introduction

Quantum information tasks such as the ones involved in quantum computation, cryptography and communication, require the preparation, control, evolution and transformation of individual quantum systems. An important element in such tasks is the validation or determination of the resulting quantum systems. Complete information about a system is encoded in the mathematical description of the ‘state’ of the system. The process of determining this state by performing multiple measurements on the quantum system is known as *Quantum State Tomography* (QST). A key feature of QST arises from the fact that a measurement performed on a quantum system ‘disturbs’ the original state, thereby restricting the information that can be gained from subsequent measurements. This is unlike the classical case, where it is possible to perform multiple measurements without disturbing the state of the system. Thus, complete information about the state of a generic quantum system can only be determined from measurements performed on several identical copies of the original state.

The importance of QST has become more evident in recent times, as remarkable experimental progress in the preparation and control of large dimensional systems has required experimentalists to deal with the challenges of validating such prepared states. The example most relevant to the work in this thesis is that of the ion-trap experiments, where large scale entangled states of multiple ions are created and manipulated. The validation of such states using full tomography is challenging due to the exponential increase in the dimensionality of the estimation problem. As an example, the landmark experiment in 2005 that successfully generated an 8-ion entangled state, used full tomography to characterise the prepared state. This involved more than 656,100 measurements and a total measurement time of 10 hours [62]. The final estimation of the state from the raw experimental data is reported to

have required several weeks of post processing [54]. A more recent experiment has demonstrated the successful creation of entangled states of 14 ions [90], and the complete characterisation of such large dimensional systems is not computationally feasible.

Therefore the main challenge of QST is to devise efficient state estimation strategies that best utilise the available resources. This involves addressing the difficulties of both choosing an optimum measurement design, and implementing a robust estimator that reconstructs the state from the measurement outcomes. In practice, choice in measurement designs is limited by the difficulties in experimental implementation, and is therefore prescribed by the current state of technology. On the other hand, for a given measurement design there are often several estimators that can be implemented to recover the state from the data. The recovery of states from measurement outcomes is essentially a classical parameter estimation problem, and therefore the various estimation strategies employed in QST can be studied using the statistical tools and methodologies typically used in such investigations. Of course, some additional caveats might need to be considered to account for certain features of quantum states such as positivity.

This thesis is composed of two parts. Part I serves as a mathematical introduction and motivates the results presented in part II. We begin in chapter 2 by introducing QST with a brief review of the mathematical representation of quantum states and measurements. We consider examples of QST in both finite and infinite dimensional quantum systems. This chapter also introduces some commonly used tomographic estimators such as the *Maximum Likelihood Estimator* (MLE). Chapter 2 also serves as a toolbox that introduces and defines the various statistical techniques that find repeated use through the rest of the thesis. It is worth highlighting some of the key concepts. We introduce the classical *Fisher Information Matrix* (FIM) which characterises the variance of certain ‘efficient’ estimators when the number of copies of the state is large, and the related Cramér-Rao Lower Bound (CRLB) which describes their errors. The *Quantum Fisher Information* (QFI) which is an analogue of the classical Fisher information is introduced. Another key tool of asymptotic theory is also introduced - *Local Asymptotic Normality* (LAN). Its quantum equivalent (qLAN) finds extensive use in chapter 6. Chapter 2 ends with an introduction to ‘Bayesian’ and ‘frequentist’ approaches to quantifying the performance of state estimators. We list the various loss functions which serve to define the relevant figures of merit, such as the Frobenius norm, trace

norm and the various fidelity based loss functions (square Bures distance and infidelity). Chapter 3 involves a more detailed discussion of the motivation behind the work presented in part II, and places the results of this thesis in its broader context.

We present the main results of the thesis in part II. In chapters 4 and 5, inspired by ‘compressed sensing’ tomography [54, 45, 98] we consider the problem of QST with incomplete measurements. The primary motivation for considering a reduced measurement design is that full QST becomes challenging for large dimensional systems. The key idea here is that it is possible to exploit some underlying ‘sparsity’ of the state for more efficient estimation. The nature of sparsity we consider in these chapters is the low rank ($r \ll d$) structure of states, where r and d are the rank and dimension of the state respectively. Several states that are of interest in quantum information processing tasks are pure (examples include the GHZ state [90] and the W states [62]). Therefore these states when prepared experimentally are likely to be low-rank or close to it. In chapter 4 we consider the statistical problem of estimating such low rank states in the setup of *Multiple Ions Tomography*. We investigate how the estimation error behaves with a reduction in the number of measurement settings, compared with the standard full ion tomography setup. We present extensive simulation results showing that the error is robust with respect to the choice of states of a given rank, the random selection of settings, and that the number of settings can be significantly reduced with only a negligible increase in error. We present an argument to explain these findings based on a concentration inequality for the Fisher information matrix (FIM). In the more general setup of random basis measurements we use this argument to show that for *certain* rank- r states it suffices to measure in $O(r \log d)$ bases to achieve the average Fisher information over all bases.

In chapter 5 we extend these results to hold for all rank- r states. We generalise and extend previous results, and show that the *Mean Square Error* (MSE) associated with the Frobenius norm attains the optimal estimation rate of $O(1/n)$ with only $O(r \log d)$ random basis measurements for all states, where n is the number of copies of the state available as a resource for the tomographic experiment. As in chapter 4, an important tool in the analysis is the concentration of the FIM. We demonstrate that although a concentration of the MSE follows from a concentration of the FIM for most states, the FIM fails to concentrate for states with eigenvalues close to zero. We analyse this phenomenon in the case of a single qubit and demonstrate a concentration

of the MSE about its optimum despite a lack of concentration of the FIM for states close to the boundary of the Bloch sphere. We also consider the estimation error in terms of a different metric - the quantum infidelity. We show that a concentration in the mean infidelity (MINF) does not exist uniformly over all states, highlighting the importance of loss function choice. Specifically, we show that for states that are nearly pure, the MINF scales as $1/\sqrt{n}$ but the constant converges to zero as the number of settings is increased. This demonstrates a lack of ‘compressive’ recovery for nearly pure states with this loss function.

This poor $O(1/\sqrt{n})$ scaling of the MINF is in general observed for states that are close to the boundary of the positive semi-definite cone of density matrices, i.e, for states that have very small eigenvalues. However, the poor scaling is not observed for states away from this boundary. As long as the loss function is *locally* quadratic in the neighbourhood of the true state, the expected loss (or *risk*) of standard tomographic procedure achieves the $O(1/n)$ scaling. The failure of standard tomography in achieving this scaling for states near the boundary is due to the fact that loss functions like the infidelity and the square Bures distance - which are defined only over the space of states - are *not* locally quadratic near the boundary, while other loss functions like the Frobenius norm continue to be quadratic in this region. Several estimators have been proposed in the literature to improve this poor scaling of the fidelity based loss functions [105, 9, 89, 70].

In chapter 6 we analyse this problem from the perspective of the maximum Bures risk of estimators over all qubit states. We propose two qubit estimation strategies; one based on *local* measurements, and a second based on *collective* measurements. In the case of local measurements, we consider a two-step adaptive strategy that is similar to already proposed estimators [9, 89]. The estimator based on collective measurements extensively utilises the results of quantum LAN results [57, 59]. We demonstrate that the maximum Bures risk of both these estimation strategies achieve the $O(1/n)$ scaling. In this chapter we also discuss the construction of a minimax optimal estimator for the Bures risk. Finally, we consider *Quantum Relative Entropy* (QRE) and show that the risk of the estimator based on collective measurements achieves a rate $O(n^{-1} \log n)$ under this loss function. Furthermore, we show that no estimator can achieve faster rates, in particular the ‘standard’ rate $O(1/n)$.

In chapter 7 we return to QST in the MIT setup, and present results from an extensive simulation study comparing the performance of several estimators

across a range of states, ranks, number of copies n and number of qubits N , as well for several loss functions - the Frobenius norm, trace norm, the square Bures distance and the square Hellinger distance. We consider the commonly used estimators like the MLE, and also introduce and define a few new ones. The simulation study offers further insight into questions considered in previous chapters, such as the behaviour of the risk of estimators for multi-qubit states that are near the boundary, and also the suitability of asymptotic theory for values of n that are typical in tomographic experiments. In addition to these results, this chapter also introduces two web-based applications that are designed for performing tomographic simulations online. These applications make available all of the estimators studied in this chapter, and allow the user to both reproduce the results of the study and more importantly, to perform simulations for arbitrary multi-qubit states.

Chapter 2

Quantum State Tomography

2.1 Notations and mathematical background

In this chapter we introduce *Quantum State Tomography* (QST), along with the necessary mathematical and statistical tools that will be used through the rest of this thesis. We begin first with the mathematical representation of quantum states and measurements.

2.1.1 Quantum states and measurements

Associated with every closed quantum system is a complex Hilbert space \mathcal{H} , known as the state space of the system. The Hilbert space is a complex inner product space that is complete with respect to the norm induced by the inner product. The state of a quantum system is described by a unit vector in Hilbert space, represented in ket notation as $|\psi\rangle \in \mathcal{H}$. A more general description of the state is in terms of the *density matrix* ρ which is a positive semi-definite matrix of unit trace and acts as an operator on \mathcal{H} . A quantum state that is described by a single ket vector $|\psi\rangle$ is known as a *pure state*, and its associated density matrix is the one dimensional projector $\rho = |\psi\rangle\langle\psi|$. General density matrices are called *mixed states*, and can be thought of as statistical mixtures of pure states

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|, \quad (2.1)$$

where the weights p_i can be thought of as probabilities associated with the pure states $|\psi_i\rangle$. As the density matrix is positive semi-definite and of trace one, the weights p_1, \dots, p_r sum to one and do constitute a probability density. The above decomposition of a density matrix is however not unique. In infinite-dimensional Hilbert space, the density matrix belong to the space of trace-class

operators on \mathcal{H} defined as

$$\mathcal{T}(\mathcal{H}) := \left\{ \sigma \in \mathcal{B}(\mathcal{H}) : \|\sigma\|_1 := \text{Tr} \left[\sqrt{\sigma^* \sigma} \right] < \infty \right\} \quad (2.2)$$

where $\mathcal{B}(\mathcal{H})$ is the space of bounded operators on \mathcal{H} .

A measurement on a quantum system is described by a set of *measurement operators* $\{E_m\}$ that satisfy the completeness relation

$$\sum_m E_m^* E_m = \mathbf{1}, \quad (2.3)$$

where E_m^* is the Hermitian adjoint, and $\mathbf{1}$ is the identity matrix. The index m represents the various possible measurement outcomes that could occur. Given a state ρ , the probability of obtaining an outcome indexed by m is

$$\mathbb{P}_\rho(m) := \text{Tr} [\rho E_m^* E_m], \quad (2.4)$$

and the post-measurement state of the system is

$$\frac{E_m \rho E_m^*}{\text{Tr} [\rho E_m^* E_m]}. \quad (2.5)$$

The completeness relation imposed on the measurement operators E_m follows from the fact that the probabilities need to sum to one $\sum_m \mathbb{P}_\rho(m) = 1$. A common formalism used to describe these general measurements applies when the post-measurement state is of little interest and does not need to be described. Defining the positive operators

$$M_m \equiv E_m^* E_m, \quad (2.6)$$

the probability of obtaining an outcome m is simply $\mathbb{P}_\rho(m) = \text{Tr} [\rho M_m]$, and the completeness relation is $\sum_m M_m = \mathbf{1}$. As the operators M_m are positive, the measurement defined by the set $\{M_m\}$ is known as a *Positive Operator-Valued Measure* (POVM), and each operator M_m is called a POVM element.

More generally POVMs may also be defined in cases where the outcomes of the measurements are not discrete as above but take values in some continuous space, and the Hilbert space is infinite dimensional. In this case we define POVMs as follows [74, 69].

Definition 1. A POVM over a measure space (Ω, \mathcal{A}) is a set $\{M(A)\}_{A \in \mathcal{A}}$ of bounded operators on \mathcal{H} such that:

- $M(\Omega) = \mathbf{1}$
- $M(A)$ is positive
- for any countable family $\{A_i\}_{i=1}^{\infty}$ of mutually disjoint sets, we have $M(\bigcup_{i=1}^{\infty} A_i) = \sum_{i=1}^{\infty} M(A_i)$

Here, each $M(A)$ is a POVM element, and the probability distribution of the outcomes for each POVM element is $\mathbb{P}_{\rho}(A) = \text{Tr}[\rho M(A)]$.

If the measurement operators $M(A)$ are taken to be orthogonal projectors for each $A \in \mathcal{A}$, then this special class of measurements is known as ‘Projection Valued Measure’ (PVMs). A measurement of this kind, where all the POVM elements are projectors, is associated with an *observable*. In quantum mechanics observables are described by self-adjoint operators on \mathcal{H} , and represent measurable physical quantities such as energy, position etc. Although by this token, POVMs might not seem to be valid physical measurements, they can be thought of as observables on a larger multipartite system, and are therefore also called *generalised observables* [111].

2.1.2 Composite systems, entanglement and collective measurements

Composite quantum systems are made up of several distinct physical systems and have an associated state space that is the tensor product of the individual Hilbert spaces. Consider n individual systems, with an isolated state in each system denoted $|\psi_i\rangle$ from $i = 1, \dots, n$. The composite state the total system is given by $|\psi_1\rangle \otimes \dots \otimes |\psi_n\rangle$, and the associated state space is $\mathcal{H} = \mathcal{H}_1 \otimes \dots \otimes \mathcal{H}_n$. Such composite systems also allows for states that cannot be written as a tensor product of isolated states in each system. Consider as an example the composite Hilbert space $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2$ of a bipartite system. If we fix $\{|i\rangle_1\}$ as the basis for \mathcal{H}_1 and $\{|j\rangle_2\}$ as the basis for \mathcal{H}_2 , then any pure state in \mathcal{H} can be expressed in the form

$$|\psi\rangle = \sum_{i,j} c_{ij} |i\rangle_1 \otimes |j\rangle_2. \quad (2.7)$$

If there exist vectors \mathbf{c}^1 and \mathbf{c}^2 such that $c_{ij} = c_i^1 c_j^2 \forall i, j$, then the vector $|\psi\rangle$ is *seperable* and can be written as $|\psi\rangle = \sum_i c_i^1 |i\rangle_1 \otimes \sum_j c_j^2 |j\rangle_2$. However, if no such vectors \mathbf{c}^1 and \mathbf{c}^2 exist, then the composite state cannot be decomposed into a tensor product of individual pure states, and is said to be *entangled*. In

terms of the density matrix ρ , a mixed state is separable if and only if it can be expressed in the form

$$\rho = \sum_i c_i \rho_i^1 \otimes \rho_i^2, \quad (2.8)$$

where c_i are positive valued and sum to one, and ρ_i^1 and ρ_i^2 are individual density matrices on the subsystems.

This concept also applies to measurements. In the example of the bipartite system, a measurement $\{E_m\}$ is said to be separable if each measurement operator E_m can be expressed as

$$E_m = \sum_i E_{m,i}^1 \otimes E_{m,i}^2, \quad (2.9)$$

where $E_{m,i}^1$ and $E_{m,i}^2$ are positive operators on the individual subsystems. A simple instance of separable measurements are *independent measurements*

$$E_m = E_{m_1}^1 \otimes E_{m_2}^2, \quad (2.10)$$

where $E_{m_1}^1$ and $E_{m_2}^2$ are measurement operators acting on each of the subsystems independently, and m_1, m_2 index the pair of outcomes from each subsystem constituting m , while all independent measurements are necessarily separable, the converse is not true. A measurement on the composite system that cannot be expressed in either of the above forms is called a *collective measurement* [66].

2.2 Example: Continuous variable systems

In order to illustrate an infinite dimensional quantum system we consider the example of a *continuous variable* (cv) system [1, 83, 115]. A single mode cv system has a Hilbert space $\mathcal{H} = L^2(\mathbb{R})$ that is infinite dimensional and is spanned by a special orthonormal basis $\{|n\rangle\}_{n=0}^{\infty}$ called the Fock (or number state) basis. Associated with the system are a pair of operators $\{a, a^*\}$ called the annihilation and creation operator respectively. The action of these operators on the Fock basis states is given by

$$a^*|n\rangle = \sqrt{n+1}|n+1\rangle, \quad a|n\rangle = \sqrt{n}|n-1\rangle \quad (2.11)$$

with $a|0\rangle = 0$. Another pair of operators P, Q that describe the one mode system are called the *quadratures*, defined as

$$Q = \frac{1}{\sqrt{2}}(a + a^*), \quad P = \frac{-i}{\sqrt{2}}(a - a^*). \quad (2.12)$$

The quadrature operators represent the canonical observables of the system and satisfy the commutation relation $[Q, P] = i\mathbf{1}$. The operators Q and P can represent the position and momentum of a particle in the case of the quantum harmonic oscillator, or the electric and magnetic fields of a light pulse in Bosonic systems [83]. Let $|p\rangle$ and $|q\rangle$ denote the eigenstates of the observables such that

$$Q|q\rangle = q|q\rangle, \quad P|p\rangle = p|p\rangle, \quad (2.13)$$

with $q, p \in \mathbb{R}$ being continuous eigenvalues. These *quadrature states* are orthogonal $\langle q|q'\rangle = \delta(q - q')$, $\langle p|p'\rangle = \delta(p - p')$, and complete

$$\int |q\rangle\langle q|dq = \int |p\rangle\langle p|dp = 1. \quad (2.14)$$

We introduce another important operator called the *displacement operator*, defined as

$$D(\alpha) := \exp(\alpha a^* - \bar{\alpha} a), \quad (2.15)$$

where $\alpha = (q + ip)/2$ is a complex number. The displacement operator acts on the vacuum state $|0\rangle$ to generate a *coherent state* $|\alpha\rangle = D(\alpha)|0\rangle$. These coherent states are the eigenstates of the annihilation operator $a|\alpha\rangle = \alpha|\alpha\rangle$, and can be expanded in the Fock basis as

$$|\alpha\rangle = \exp\left(-\frac{1}{2}|\alpha|^2\right) \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle. \quad (2.16)$$

An arbitrary density operator ρ in such a cv system is an infinite dimensional matrix. However, density matrices also admit an equivalent representation in terms of quasi-probability distribution called the *Wigner Function*. We define the Weyl operator

$$\tilde{W}(u, v) := \exp(-iuQ - ivP), \quad (2.17)$$

where $(u, v) \in \mathbb{R}^2$. Then, a density operator ρ is equivalent to its *quantum characteristic function*

$$\tilde{W}_\rho(u, v) = \text{Tr}[\rho \exp(-iuQ - ivP)], \quad (2.18)$$

and by a Fourier transform to the Wigner function

$$W_\rho(q, p) = \frac{1}{(2\pi)^2} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \tilde{W}_\rho(u, v) \exp(iuq + ivp) du dv, \quad (2.19)$$

which is normalised to one like a true probability distribution is. However the Wigner function is not strictly positive and can take negative values, which is why it is called a quasi-probability distribution. The Wigner function contains information about the statistical moments of the quantum states. In particular, the first moment or mean of the quadratures

$$\langle P \rangle_\rho = \text{Tr}[\rho P], \quad \langle Q \rangle_\rho = \text{Tr}[\rho Q] \quad (2.20)$$

and the second moment or the covariance matrix V ,

$$V := \begin{pmatrix} \langle Q^2 \rangle_\rho & \langle Q \circ P \rangle_\rho \\ \langle P \circ Q \rangle_\rho & \langle P^2 \rangle_\rho \end{pmatrix} - \begin{pmatrix} \langle Q \rangle_\rho^2 & \langle Q \rangle_\rho \langle P \rangle_\rho \\ \langle P \rangle_\rho \langle Q \rangle_\rho & \langle P \rangle_\rho^2 \end{pmatrix} \quad (2.21)$$

where $Q \circ P = (QP + PQ)/2$.

With these definitions in place, we now define an important class of states called *Gaussian states*. These states are completely characterised by the first two moments, and their Wigner representation is

$$W_\rho(p, q) = \frac{1}{(2\pi)^2 \sqrt{\det V}} \exp \left[-(\mathbf{w} - \langle \mathbf{w} \rangle)^T V^{-1} (\mathbf{w} - \langle \mathbf{w} \rangle) \right] \quad (2.22)$$

where $\mathbf{w} = (q, p)^T \in \mathbb{R}^2$ and $\langle \mathbf{w} \rangle = (\langle Q \rangle_\rho, \langle P \rangle_\rho)^T$ is the mean vector. Clearly this is of the form of a multivariate Gaussian distribution. The simplest example of a Gaussian state is the vacuum state $|0\rangle\langle 0|$, which can be shown to have a mean zero vector and a covariance matrix proportional to identity. The coherent state $|\alpha\rangle\langle \alpha|$ being a displaced vacuum state $D(\alpha)|0\rangle\langle 0|D(-\alpha)$, is also Gaussian with mean proportional to the displacement and an unchanged variance. We define a zero mean thermal equilibrium state as

$$\phi^{\mathbf{0}} := (1-p) \sum_{k=0}^{\infty} p^k |k\rangle\langle k| \quad (2.23)$$

with $p \leq 1$. The zero photon state is seen to be a special case with $p = 1$. These states are displaced by the action of the displacement operator

$$\phi^\alpha := D(\alpha) \phi^{\mathbf{0}} D(-\alpha), \quad (2.24)$$

and are called *displaced thermal equilibrium states*.

Homodyne and Heterodyne detection

The most common measurements on cv systems are of the Homodyne and Heterodyne detection schemes. A homodyne detection consists of the measurements of the Q , P (or other) quadratures of the single mode. Its measurement operators are the projectors $|q\rangle\langle q|$ and $|p\rangle\langle p|$ of the corresponding quadrature basis. The probability distributions of the outcomes are given by the marginal probability densities [115]

$$\mathbb{P}_\rho(q) = \int W_\rho(q,p) dp, \quad \mathbb{P}_\rho(p) = \int W_\rho(q,p) dq \quad (2.25)$$

The heterodyne detection corresponds to a projection onto coherent states (2.16), and has POVM elements given by [74]

$$M(A) := \frac{1}{\pi} \int_A |\alpha\rangle\langle\alpha| d\alpha, \quad (2.26)$$

where the outcomes $\alpha \in \mathbb{C}$ contain information about both $q, p \in \mathbb{R}$. The probability distribution of the outcomes is a convolution of the Wigner function of ρ with a Gaussian.

2.3 Estimation of finite-dimensional quantum states

Through most of this thesis we shall be interested in finite dimensional quantum systems. The states of such systems exist in a finite dimensional Hilbert space $\mathcal{H}^d \equiv \mathbb{C}^d$, where d is the dimension of the space. The density matrix ρ is a $d \times d$ positive-semidefinite matrix of trace one, and let $\mathbb{S}_d \subset M(\mathbb{C}^d)$ denote the space of density matrices. One of the simplest and most familiar examples of a finite dimensional quantum system is that of a two level system of a ‘quantum bit’ or ‘qubit’. There are many different physical systems that are used to realise and prepare qubits for tasks such as quantum computing [5, 65, 11, 86]. The Hilbert space associated with a qubit system is $\mathcal{H}^2 \equiv \mathbb{C}^2$, where the dimension $d = 2$. The density matrix of an arbitrary qubit state is given by

$$\rho = \frac{1}{2}(\mathbf{1} + \mathbf{r} \cdot \boldsymbol{\sigma}), \quad (2.27)$$

where $\boldsymbol{\sigma} = (\sigma_x, \sigma_y, \sigma_z)^T$ is a vector of the standard Pauli observables, and $\mathbf{r} \in \mathbb{R}^3$ is the Bloch vector associated with the state. The components of the Bloch vector are given by $r_i = \text{Tr}(\rho\sigma_i)$ for $i = (x, y, z)$. A pure state has a Bloch vector of length $|\mathbf{r}| = 1$, while mixed states have $0 \leq |\mathbf{r}| < 1$. A more general example is a quantum system of *multiple qubits*. The joint Hilbert space of N ions is given by the tensor product $(\mathbb{C}^2)^{\otimes N} = \mathbb{C}^d$, where $d = 2^N$.

Given a ‘true’ and unknown state associated with a quantum system, the aim of quantum tomography is to statistically reconstruct its density matrix $\rho \in \mathbb{S}_d$ from the outcomes of measurements performed on the system. We can describe a measurement in general terms as a POVM with elements $\{M_m\}_{m=1}^k$, such that $M_m \in M(\mathbb{C}^d)$, and m indexes the various possible outcomes. The probability distribution of the outcomes is given by the map

$$M : \rho \mapsto \{\mathbb{P}_\rho(m) := \text{Tr}[\rho M_m]; m = 1, \dots, k\}. \quad (2.28)$$

However, as the measurement necessarily ‘disturbs’ the state of the system with some back-action, it is generally impossible to gain further information about the true state with subsequent measurements. It is due to this limitation that multiple identical copies of the quantum state are necessary. Suppose that we have available n identical copies of the state available as a resource. The POVM measurement is repeated on each copy of the state and therefore we have the model of n *independent and identically distributed* (iid) random variables X_1, \dots, X_n with probability distributions given by (2.28). Let x_1, \dots, x_n be the observed measurement data, then the aim of quantum tomography is to build an estimator $\hat{\rho}_n = \hat{\rho}(x_1, \dots, x_n)$ of the true state ρ .

It is often convenient to model the tomographic problem as a classical parameter estimation problem. Let the state ρ_θ be parameterised by a finite dimensional vector of parameters $\theta \in \Theta$. The parameters can, for example, simply be the elements of the density matrix, or for qubit states the Bloch vector components in (2.27). Additionally, when some prior information about the state is available θ can parameterise a lower dimensional model. For the given POVM $\{M_m\}_{m=1}^k$, the distribution of the outcomes is also dependent on the parameter, and is given by the model

$$\mathcal{P}_M := \{\mathbb{P}_\theta := \mathbb{P}_{\rho_\theta}; \theta \in \Theta\}. \quad (2.29)$$

Therefore the aim is to construct an estimator $\hat{\boldsymbol{\theta}}_n$ of the unknown parameter from the outcomes of the iid random variables X_1, \dots, X_n , and to thereby arrive at an estimator $\hat{\rho}_{\hat{\boldsymbol{\theta}}_n}$ of the true state.

In this section we consider the state estimation problem in greater detail, and consider the specific instance of *Multiple Ions Tomography* (MIT). We will introduce the measurement procedure and statistical model of the MIT setup, and review some estimators such as the Least Squares Estimator (LSE) and the Maximum Likelihood Estimator (MLE).

2.3.1 Multiple Ions Tomography

In the MIT setup as in the ion-trap experiments [90, 62, 82], the aim is to estimate an unknown joint state of N ions (modelled as qubits) from the outcomes of measurements performed on identically prepared systems. Let $\rho \in \mathbb{S}_d$ be the $d \times d$ density matrix of the unknown state, where $d = 2^N$ is the dimension of the Hilbert space. We consider measurements of two types - the standard tomographic measurements in the Pauli basis, and measurements that are drawn randomly from the uniform measure over orthonormal bases (ONB). In the case of the Pauli basis, one measures an arbitrary Pauli observable $\sigma_x, \sigma_y, \sigma_z$ on each of the N ions simultaneously. Therefore, each measurement is labelled by a sequence $\mathbf{s} = (s_1, \dots, s_N) \in \{x, y, z\}^N$, and there are 3^N possible measurement bases. In the uniformly random measurement set up, a measurement can be implemented by first rotating the state ρ by a random unitary $U \in M(\mathbb{C}^d)$, after which each ion is measured in the σ_z eigenbasis.

Let $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_k\}$ be the measurement design consisting of k measurement settings. In the case of the Pauli set up the total number $k = 3^N$, while in the random measurement setup the number of settings measured can be chosen freely. A measurement in a particular setting produces a ± 1 outcome from each ion, and we let $\mathbf{o} \in \{+1, -1\}^N$ be a vector record of outcomes from each of the N ions. The probability of obtaining a particular outcome \mathbf{o} is given by $p_\rho(\mathbf{o}|\mathbf{s}) := \text{Tr}(\rho P_{\mathbf{o}}^{\mathbf{s}})$, where the one-dimensional projection matrix is given by

$$P_{\mathbf{o}}^{\mathbf{s}} = |e_{o_1}^{s_1}\rangle\langle e_{o_1}^{s_1}| \otimes \dots \otimes |e_{o_N}^{s_N}\rangle\langle e_{o_N}^{s_N}|$$

For each setting \mathbf{s} , measurements are repeated on m identical copies of the state, and the counts of the outcomes $N(\mathbf{o}|\mathbf{s})$ are recorded, where $N(\mathbf{o}|\mathbf{s})$ represents the number of times a given outcome record \mathbf{o} is observed given that measurements were performed in a chosen setting \mathbf{s} . The total number

of quantum samples used is therefore $n = m \times k$. The resulting dataset \mathcal{D} of counts is a $2^N \times k$ table whose columns are independent and contain all the counts in a given setting. The probability of observing a given dataset of counts is given by the product of multinomials

$$p_\rho(\mathcal{D}|\mathcal{S}) = \prod_{\mathbf{s}} \frac{m!}{\prod_{\mathbf{o}} N(\mathbf{o}|\mathbf{s})!} \prod_{\mathbf{o}} p_\rho(\mathbf{o}|\mathbf{s})^{N(\mathbf{o}|\mathbf{s})} \quad (2.30)$$

The goal is the statistical reconstruction of the density matrix from this dataset of counts. There are several estimators known in the literature for this purpose, but here we introduce one of the simplest estimators - the *Least Squares Estimator* (LSE).

2.3.2 The Least Squares Estimator

In order to specify the LSE, we need to first establish a choice of parametrisation, thereby converting the tomographic problem into a problem of parameter estimation. To better describe the linear estimation problem, we consider the true probability vector

$$\mathbf{y} = (p_\rho(\mathbf{o}_1|\mathbf{s}_1), \dots, p_\rho(\mathbf{o}_d|\mathbf{s}_1), \dots, p_\rho(\mathbf{o}_d|\mathbf{s}_k))^T \in \mathbb{R}^{kd} \quad (2.31)$$

In the standard basis, each element of this vector can be expressed in terms of the density matrix elements and the corresponding one-dimensional projections as

$$p_\rho(\mathbf{o}|\mathbf{s}) = 2 \sum_{j>i} \text{Re}(\rho_{ij}) \text{Re}(P_{\mathbf{o}}^{\mathbf{s}})_{i,j} + 2 \sum_{j>i} \text{Im}(\rho_{ij}) \text{Im}(P_{\mathbf{o}}^{\mathbf{s}})_{i,j} + \sum_i^d \rho_{ii} (P_{\mathbf{o}}^{\mathbf{s}})_{i,i} \quad (2.32)$$

We now choose to parameterise the state by the elements of the density matrix ρ that appear in the above equation. Let $\boldsymbol{\theta} \in \mathbb{R}^{d^2}$ be a vector of parameters defined as

$$\boldsymbol{\theta} := \left(\text{Re}\rho_{1,2}, \dots, \text{Re}\rho_{d-1,d}, \text{Im}\rho_{1,2}, \dots, \text{Im}\rho_{d-1,d}, \rho_{1,1}, \dots, \rho_{d,d} \right)^T. \quad (2.33)$$

With this choice of parameterisation, equation (2.32) can be expressed as an inner product between vectors $p_\rho(\mathbf{o}|\mathbf{s}) = X(\mathbf{o}|\mathbf{s})^T \boldsymbol{\theta}$, where

$$X(\mathbf{o}|\mathbf{s})^T := \left(2\operatorname{Re}(P_{\mathbf{o}}^{\mathbf{s}})_{1,2}, \dots, 2\operatorname{Re}(P_{\mathbf{o}}^{\mathbf{s}})_{d-1,d}, 2\operatorname{Im}(P_{\mathbf{o}}^{\mathbf{s}})_{1,2}, \dots, \right. \\ \left. 2\operatorname{Im}(P_{\mathbf{o}}^{\mathbf{s}})_{d-1,d}, (P_{\mathbf{o}}^{\mathbf{s}})_{1,1}, \dots, (P_{\mathbf{o}}^{\mathbf{s}})_{d,d} \right).$$

This notion allows us to express the tomographic system of equations in matrix form as

$$\mathbf{y} = X\boldsymbol{\theta}, \quad (2.34)$$

where X is a $kd \times d^2$ matrix whose rows are given by $X(\mathbf{o}|\mathbf{s})^T$ for each pair of outcome \mathbf{o} and setting \mathbf{s} . Of course, in reality, we do not have access to the true probability vector. Instead from the $d \times k$ dataset of counts, we have access to the empirical probabilities $f(\mathbf{o}|\mathbf{s}) := N(\mathbf{o}|\mathbf{s})/m$, whose expectations are $\mathbb{E}f(\mathbf{o}|\mathbf{s}) = p_\rho(\mathbf{o}|\mathbf{s})$. Replacing the probability vector \mathbf{y} by the vector of empirical frequencies we have

$$\mathbf{f} = X\boldsymbol{\theta} + \boldsymbol{\epsilon} \quad (2.35)$$

where $\boldsymbol{\epsilon} \in \mathbb{R}^{dk}$ is a mean zero vector of statistical noise. The least-squares solution to the above system of equations is defined as the minimiser of the following optimisation problem

$$\hat{\boldsymbol{\theta}} := \arg \min_{\boldsymbol{\tau} \in \mathbb{R}^{d^2}} \|\mathbf{f} - X\boldsymbol{\tau}\|^2 \quad (2.36)$$

and has the well known explicit form $\hat{\boldsymbol{\theta}} = (X^T X)^{-1} \cdot X^T \cdot \mathbf{f}$. The final estimate of the density matrix $\hat{\rho}_{\text{LS}}$ is then constructed from the estimated parameter vector $\hat{\boldsymbol{\theta}}$.

Although the LS estimator is computationally simple, and can prove to be good starting point for several other refinements and estimation methods, it suffers from a number of drawbacks. The most serious of which is that LS estimate is not necessary a density matrix, i.e, it often produces estimates that are neither positive semi-definite nor of trace one. In other words, the LS estimator is blind to the physical properties of the state. On the other hand estimators such as the extensively used *Maximum Likelihood Estimator* (MLE) produce an estimate which is a density matrix.

2.3.3 Maximum Likelihood Estimator

The *Maximum Likelihood Estimator* (MLE) is one of the most commonly used estimation methods in statistics. It estimates the unknown parameters from the measurement data by finding values for the parameter that maximise the *likelihood* of having obtained the measurement outcomes. Let X_1, \dots, X_n be n iid random variables with the joint distribution that depends on an unknown parameter $\boldsymbol{\theta} \in \Theta$

$$\mathbb{P}_{\boldsymbol{\theta}}^n(X_1 = x_1, \dots, X_n = x_n) = \prod_{i=1}^n \mathbb{P}_{\boldsymbol{\theta}}(X_i = x_i) \quad (2.37)$$

Given the observations (x_1, \dots, x_n) , we construct the *likelihood function* which has the same form as the joint distribution

$$\mathcal{L}(\boldsymbol{\theta}; x_1, \dots, x_n) = \prod_{i=1}^n \mathbb{P}_{\boldsymbol{\theta}}(X_i = x_i) \quad (2.38)$$

except that this function varies with the parameter $\boldsymbol{\theta}$, and the outcomes are taken to be fixed. The MLE $\hat{\boldsymbol{\theta}}_n$ is obtained by maximising this likelihood function over the entire parameter space

$$\hat{\boldsymbol{\theta}}_n = \arg \max_{\boldsymbol{\tau} \in \Theta} \mathcal{L}(\boldsymbol{\tau}; x_1, \dots, x_n). \quad (2.39)$$

It is more common and convenient to maximise the natural logarithm of the likelihood function. In the MIT setup, the probability distribution over outcomes is given by (2.30); discarding the constant factorial term, we arrive at the following form of the MLE

$$\hat{\rho}_{ML} := \arg \max_{\boldsymbol{\tau} \in \mathbb{S}_d} \sum_{\mathbf{o}, \mathbf{s}} N(\mathbf{o}|\mathbf{s}) \log p_{\boldsymbol{\tau}}(\mathbf{o}|\mathbf{s}), \quad (2.40)$$

where the maximum is taken over the space of states ($\boldsymbol{\tau} \geq 0, \text{Tr}[\boldsymbol{\tau}] = 1$). The maximum likelihood can be seen to be independent of the choice of parameterisation. The MLE is commonly used in quantum tomography [13, 19, 71, 101, 50], with several methods proposed in the literature such as the iterative algorithm proposed by Hradil [101]. Unlike estimators like the LS estimator, The MLE is guaranteed to produce a quantum state. It additionally satisfies certain desirable statistical properties in the asymptotic regime such as attaining the Cramér-Rao lower bound. However, as we shall point out in the following

section, the asymptotic theory does not apply when the parameter $\boldsymbol{\theta}$ lies at the boundary of the parameter space Θ . Along with this, the MLE also suffers from the drawback that it has a tendency to produce rank deficient estimates when the true state has some small eigenvalues.

2.4 Fisher Information and Local Asymptotic Normality

2.4.1 Classical Fisher Information

Consider the problem of estimating an unknown single parameter $\theta \in \Theta$ from the outcomes of n *independent and identically distributed* (iid) random variables $\mathbf{X} = (X_1, \dots, X_n)$ with the joint distribution (dependent on θ)

$$\mathbb{P}_\theta^n(X_1 = x_1, \dots, X_n = x_n) = \prod_{i=1}^n \mathbb{P}_\theta(X_i = x_i). \quad (2.41)$$

As the random variables are assumed to be iid, they all have the same dependence on the parameter θ . A simple example of such an estimation problem is the task of determining the *bias* of a coin from n coin flips. Let X_i be the i th trial of such an experiment. There are two possible outcomes $x_i = 0$ for tails and $x_i = 1$ for heads. The probabilities of the outcomes are given by the Bernoulli distribution with parameter θ

$$p_\theta(x) = \theta^x(1 - \theta)^{1-x} \quad x \in \{0, 1\} \quad (2.42)$$

In this case the parameter $\theta \in [0, 1]$. Let $\hat{\theta}_n$ be an estimator constructed as some function of the outcomes of the n trials, such that

$$\mathbb{E}[\hat{\theta}_n] = \sum_{\mathbf{x}} p_\theta(\mathbf{x}) \hat{\theta}_n(\mathbf{x}) = \theta, \quad (2.43)$$

where \mathbf{x} is a vector of n outcomes. An estimator that recovers the true value of the parameter in expectation is said to be *unbiased*. In order to quantify how ‘good’ such an estimator is, we may consider the variance of $\hat{\theta}_n$. Certainly an unbiased estimator with small variance is preferable to one with a large variance. We may also wish to know how small this value can be irrespective of the particular choice of estimator. This question is answered by the *Cramér-Rao Lower Bound* (CRLB), which bounds the variance of all possible unbiased

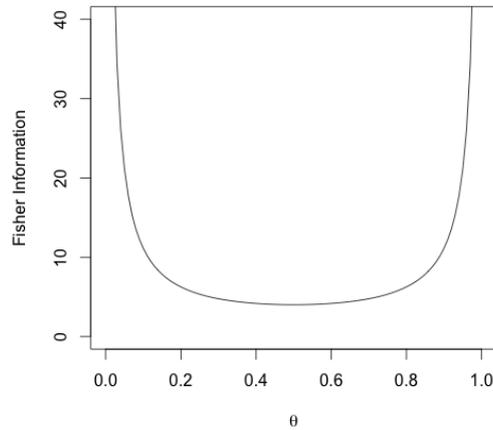


Figure 2.1 The Fisher information $I(\theta) = \frac{1}{\theta(1-\theta)}$ for a single trial of the coin flip experiment, plotted for the range of $\theta \in [0, 1]$.

estimators $\hat{\theta}_n$ from below

$$\text{Var}[\hat{\theta}_n] \geq \frac{1}{nI(\theta)}, \quad (2.44)$$

where the quantity $I(\theta)$ is called the *Fisher information*, and is defined as

$$I(\theta) = \mathbb{E} \left[\left(\frac{\partial}{\partial \theta} \log \mathbb{P}_\theta(X_i) \right)^2 \right] = \sum_{x_i} \frac{1}{p_\theta(x_i)} \left(\frac{\partial p_\theta(x_i)}{\partial \theta} \right)^2. \quad (2.45)$$

The quantity $\log \mathbb{P}_\theta(X_i)$ is the *log-likelihood* function and its derivative is known as the *score*. The Fisher information represents the average ‘sensitivity’ of the score to the variation of the parameter θ , and can be thought of as being related to the mean curvature of the log-likelihood function. The greater this curvature, the larger the fisher information will be. Notice that while the Fisher information has no dependence on a particular estimator, it is a function of θ and also depends on the choice of ‘measurements’, i.e, the random variable X_i .

In the example of the coin flip experiment, X_i are independent coin tosses with $p_\theta(x_i = 1) = \theta$, and the associated Fisher information is easily evaluated to be $I(\theta) = \frac{1}{\theta(1-\theta)}$. The CRLB for n iid trails is therefore $\theta(1-\theta)/n$. The standard unbiased estimator

$$\hat{\theta}_n = \frac{1}{n} \sum_{i=1}^n X_i, \quad (2.46)$$

has variance $\text{Var}(\hat{\theta}_n) = \theta(1-\theta)/n$, and therefore achieves the CRLB. Looking at the plot in figure 2.1, we see that Fisher information has its lowest value at $\theta = 1/2$. This implies that random variables $\mathbf{X} = (X_1, \dots, X_n)$ carry less information about the parameter when $\theta = 1/2$ than when it is away from this midpoint. Notice however that the Fisher information diverges at the boundary of the parameter space $\theta \in [0, 1]$. This is explained by the fact that the CRLB applies only for values of θ in the interior of the parameter space. Additionally, the parameter and the log-likelihood function need to satisfy certain regularity and ‘smoothness’ conditions for the CRLB to hold.

In the multi-parameter case, the aim is to estimate an unknown vector of parameters $\boldsymbol{\theta} = (\theta_1, \dots, \theta_p) \in \Theta$ from the outcomes of n iid random variables $\mathbf{X} = (X_1, \dots, X_n)$ with distribution $\mathbb{P}_{\boldsymbol{\theta}}^n$. In this case the Fisher information is not a single number but is given by a $p \times p$ positive matrix whose elements are defined as

$$I(\boldsymbol{\theta})_{k,l} = \mathbb{E} \left[\left(\frac{\partial}{\partial \theta_k} \log \mathbb{P}_{\boldsymbol{\theta}}(X_i) \right) \left(\frac{\partial}{\partial \theta_l} \log \mathbb{P}_{\boldsymbol{\theta}}(X_i) \right) \right]. \quad (2.47)$$

The multi-parameter CRLB is expressed as a matrix inequality

$$n\mathbb{E} \left[(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta})^T (\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}) \right] \geq I(\boldsymbol{\theta})^{-1}. \quad (2.48)$$

Often it is more convenient to express the bound in terms of a single number such as the *Mean Square Error* (MSE)

$$n\mathbb{E} \left[(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta})G(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta})^T \right] \geq \text{Tr} \left[I(\boldsymbol{\theta})^{-1}G \right] \quad (2.49)$$

where G is a $p \times p$ positive *weight matrix* that depends on the parameters, so that the overall MSE is independent of the choice of parametrisation.

In general, there may not be any unbiased estimators $\hat{\boldsymbol{\theta}}_n(\mathbf{X})$ that achieve the CRLB for a finite number of iid trials n . However, it is known that certain ‘efficient’ estimators (such as the MLE under some regularity conditions [114]) are normally distributed

$$\sqrt{n}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}) \rightarrow \mathcal{N}(\mathbf{0}, I(\boldsymbol{\theta})^{-1}), \quad (2.50)$$

in the limit $n \rightarrow \infty$, where $\mathcal{N}(\mathbf{0}, I(\boldsymbol{\theta})^{-1})$ is the multivariate normal distribution with zero mean vector and a covariance matrix given by the inverse of the Fisher information matrix. Such estimators therefore achieve the CRLB *asymptotically*.

It is important to note however, that the CRLB does not hold for values of $\boldsymbol{\theta}$ that are at the boundary of the parameter space. Hence the above definition of efficiency of estimators holds only when $\boldsymbol{\theta}$ is within the parameter space Θ .

2.4.2 Classical Local Asymptotic Normality

Returning once again to the example of the coin toss experiment, we know that the standard unbiased estimator (2.46) is a ‘good’ estimator of the probability θ of getting heads as it achieves the CRLB for values of θ not at the boundary of the parameter space. From both this fact and the *central limit theorem* (CLT) we have that the error $\hat{\theta}_n - \theta$ has the following Gaussian distribution in the asymptotic limit

$$\sqrt{n}(\hat{\theta}_n - \theta) \rightarrow \mathcal{N}(0, \theta(1 - \theta)). \quad (2.51)$$

as $n \rightarrow \infty$ [59]. Also the mean square error is easily seen to be $\mathbb{E}[(\theta - \hat{\theta}_n)^2] = \theta(1 - \theta)/n$. Imagine that instead of estimating θ over the entire parameter space $[0, 1]$, we had some prior knowledge of the unknown parameter. Namely that it lies in a local neighbourhood of size $\frac{1}{\sqrt{n}}$ around a known value θ_0 . Let $u \in \mathbb{R}$ be a local parameter, such that we may redefine $\theta = \theta_0 + \frac{u}{\sqrt{n}}$, and denote the distributions of the random variables X_i as $\mathbb{P}_{\theta_0 + u/\sqrt{n}}(X_i = x_i)$. The estimator \hat{u}_n of the local parameter is such that

$$\hat{u}_n = \sqrt{n}(\hat{\theta}_n - \theta_0) \rightarrow \mathcal{N}(u, \theta_0(1 - \theta_0)) \quad (2.52)$$

as $n \rightarrow \infty$. This means that the estimator \hat{u}_n is asymptotically distributed as a random Gaussian variable with the above mean and variance. Therefore in the asymptotic limit, the problem of estimating the local parameter $u \in \mathbb{R}$ from n iid random variables $\mathbf{X} = (X_1, \dots, X_n)$ is statistically equivalent to the problem of estimating the mean of a Gaussian distribution from a single random variable $Y \sim \mathcal{N}(u, 1/I(\theta_0))$. This asymptotic property is termed *Local Asymptotic Normality* (LAN).

LAN essentially means that in the limit of large n , the iid statistical experiment is approximated by Gaussian experiments after a re-parameterisation. Let $(\mathbb{P}_{\boldsymbol{\theta}_0 + \mathbf{u}/\sqrt{n}}^n : \mathbf{u} \in \mathbb{R}^p)$ be a sequence of experiments consisting of observing n iid random variables $\mathbf{X} = (X_1, \dots, X_n)$ with distribution $\mathbb{P}_{\boldsymbol{\theta}_0 + \mathbf{u}/\sqrt{n}}^n$, where $\mathbf{u} \in \mathbb{R}^p$ is a vector of unknown local parameters. Provided that the map $\boldsymbol{\theta} \mapsto \mathbb{P}_{\boldsymbol{\theta}}$

is sufficiently ‘smooth’, we have that the experiments

$$\{\mathbb{P}_{\boldsymbol{\theta}_0 + \mathbf{u}/\sqrt{n}}^n : \mathbf{u} \in \mathbb{R}^p\} \quad \text{and} \quad \{\mathcal{N}(\mathbf{u}, I(\boldsymbol{\theta}_0)^{-1}) : \mathbf{u} \in \mathbb{R}^p\} \quad (2.53)$$

have the same statistical properties in the limit $n \rightarrow \infty$ [75]. That is, the iid experiment asymptotically converges to the limit experiment where a single sample is observed from a multivariate Gaussian distribution with unknown mean $\mathbf{u} \in \mathbb{R}^p$ and a fixed covariance matrix $I(\boldsymbol{\theta}_0)^{-1}$.

We shall not formulate LAN in precise terms here, but note that there are two notions of convergence, and two corresponding formulations of the LAN principle. The *weak* version is based on the CLT as in the binomial example above, and is a *convergence in distribution of the finite dimensional marginals of the likelihood ratio process* [114, 112], while the *strong* version is a convergence in norm rather than in distribution and is with respect to the Le Cam distance [75, 58].

2.4.3 Quantum Fisher Information

We now introduce the quantum analogue of the classical Fisher information and the CRLB for unbiased estimators. In the quantum case the probability distribution $\mathbb{P}_{\boldsymbol{\theta}}$ with its dependence on the underlying unknown parameter vector $\boldsymbol{\theta} \in \mathbb{R}^p$ is replaced by a density matrix $\rho_{\boldsymbol{\theta}}$ that depends ‘smoothly’ on the unknown parameter.

The classical Fisher information was defined by the derivatives of the log-likelihood known as the score $\frac{\partial}{\partial \theta_k} \log p_{\boldsymbol{\theta}}(x_i)$, for $k = 1, \dots, p$. In the quantum case the density matrix $\rho_{\boldsymbol{\theta}}$ in general does not commute with its derivatives, and therefore the definition of the *quantum score* is not a straightforward extension of the classical one. In fact there are several definitions of the quantum score [69, 15, 21], and here we introduce the most common version called the *Symmetric Logarithmic Derivative* (SLD) $\mathcal{L}_{\boldsymbol{\theta}, k}$ which is defined as the solution to the equation

$$\frac{\partial \rho_{\boldsymbol{\theta}}}{\partial \theta_k} = \mathcal{L}_{\boldsymbol{\theta}, k} \circ \rho_{\boldsymbol{\theta}} \quad k = 1, \dots, p, \quad (2.54)$$

where we define $A \circ B = (AB + BA)/2$. Analogous to the definition of classical Fisher information in (2.47), the *Quantum Fisher information* (QFI) is defined

as the $p \times p$ positive matrix with elements

$$F(\boldsymbol{\theta})_{k,l} = \text{Tr} \left[\rho_{\boldsymbol{\theta}} \mathcal{L}_{\boldsymbol{\theta},k} \circ \mathcal{L}_{\boldsymbol{\theta},l} \right]. \quad (2.55)$$

The quantum Fisher matrix $F(\boldsymbol{\theta})$ does not depend on the measurement choice, and only depends on the parameterisation chosen. This is unlike the classical fisher information $I(\boldsymbol{\theta})$ which depends on both the parameters and the measurements. Consider a POVM with elements $\{M_m\}_{m=1}^k$, and outcomes indexed by m . The probability distribution of the m outcomes is $\mathbb{P}_{\boldsymbol{\theta}}(m) = \text{Tr}[\rho_{\boldsymbol{\theta}} M_m]$. Let $I_M(\boldsymbol{\theta})$ be the classical Fisher information associated with this POVM measurement. Then we have the following matrix inequality [69, 67]

$$I_M(\boldsymbol{\theta}) \leq F(\boldsymbol{\theta}). \quad (2.56)$$

If $\hat{\boldsymbol{\theta}}$ is an unbiased estimator of the unknown parameter vector, then we have the Quantum CRLB

$$\text{Var} \left[\hat{\boldsymbol{\theta}} \right] \geq I_M(\boldsymbol{\theta})^{-1} \geq F(\boldsymbol{\theta})^{-1}. \quad (2.57)$$

The inequality on the right establishes an absolute lower bound on the variance for *any* unbiased estimator and any measurement. One may wonder if the quantum CRLB is achievable by any estimator-measurement pair. In the case of multi-parameter estimation problems this lower bound is asymptotically attainable if and only if the SLDs satisfy certain commutation relations [48]. However it can be shown that the quantum CRLB is asymptotically achievable for one dimensional parameters $\theta \in \mathbb{R}$ [21].

2.4.4 Quantum Local Asymptotic Normality

The classical principle of LAN establishes the following: given n iid random variables X_1, \dots, X_n with distribution $\mathbb{P}_{\boldsymbol{\theta}_0 + \mathbf{u}/\sqrt{n}}^n$ depending ‘smoothly’ on the local parameter $\mathbf{u} \in \mathbb{R}^p$, then in the asymptotic limit, this model is equivalent to a single draw $Y \in \mathbb{R}^p$ from a multivariate Gaussian distribution with unknown mean \mathbf{u} and with fixed variance $I(\boldsymbol{\theta}_0)^{-1}$.

In the quantum case, we have n identical copies of the true state $\rho_{\boldsymbol{\theta}}$. Consider that the unknown parameter $\boldsymbol{\theta}$ lies in a local neighbourhood of size $n^{-1/2+\eta}$ around a fixed $\boldsymbol{\theta}_0$. That is, let $\mathbf{u} \in \mathbb{R}^p$ be a vector of local parameters such that $\|\mathbf{u}\| \leq n^\eta$, and let $\rho_{\boldsymbol{\theta}}$ be parameterised locally as $\rho_{\boldsymbol{\theta}_0 + \mathbf{u}/\sqrt{n}}$. The quantum

version of LAN demonstrates that the joint state $\rho_{\boldsymbol{\theta}_0 + \mathbf{u}/\sqrt{n}}^{\otimes n}$ converges in the asymptotic limit to a displaced Gaussian state of a continuous variable system.

We now make this convergence more concrete and consider the quantum LAN results for qubit states. Let ρ_0 be a fixed qubit state that is away from the boundary of the Bloch sphere

$$\rho_0 := \begin{pmatrix} 1 - \lambda_0 & 0 \\ 0 & \lambda_0 \end{pmatrix}, \quad (2.58)$$

where $0 < \lambda_0 \leq 1/2$ is the minimum eigenvalue of the state. We parameterise states in a local neighbourhood of ρ_0 as $\rho_{\mathbf{u}/\sqrt{n}}$, where $\mathbf{u} = (u_x, u_y, u_z) \in \mathbb{R}^3$, $\|\mathbf{u}\| \leq n^\eta$ is a vector of local parameters. Of the three parameters the first two $\mathbf{w} = (u_x, u_y) \in \mathbb{R}^2$ account for unitary rotations around ρ_0 and $u_z \in \mathbb{R}$ accounts for changes in the eigenvalues

$$\rho_{\mathbf{u}/\sqrt{n}} := U \left(\frac{\mathbf{w}}{\sqrt{n}} \right) \begin{pmatrix} 1 - \lambda_0 - \frac{u_z}{\sqrt{n}} & 0 \\ 0 & \lambda_0 + \frac{u_z}{\sqrt{n}} \end{pmatrix} U \left(\frac{\mathbf{w}}{\sqrt{n}} \right)^* \quad (2.59)$$

where the unitary $U(\mathbf{w}) := \exp[i(u_x \sigma_x + u_y \sigma_y)]$. Now consider n identical copies of the true state, each of which is known to be in the local neighbourhood of size $n^{-1/2+\eta}$ around the fixed state ρ_0 . Then in the limit of large n , the joint state $\rho_{\mathbf{u}/\sqrt{n}}^{\otimes n}$ approaches a Gaussian state $\phi^{\mathbf{w}} \otimes N^{\mathbf{u}}$, where $N^{\mathbf{u}}$ is the classical one-dimensional Gaussian distribution

$$N^{\mathbf{u}} := \mathcal{N}(u_z, \lambda_0(1 - \lambda_0)), \quad (2.60)$$

centred around u_z , and $\phi^{\mathbf{w}}$ is a displaced thermal equilibrium state with displacement proportional to $\mathbf{w} = (u_x, u_y)$. Let $\mathcal{H} = L^2(\mathbb{R})$ be the Hilbert space of a one mode cv system, and let

$$\phi^{\mathbf{0}} := (1 - p) \sum_{k=0}^{\infty} p^k |k\rangle \langle k|, \quad p = \frac{\lambda_0}{1 - \lambda_0} \quad (2.61)$$

be the centred thermal equilibrium state of a quantum harmonic oscillator, where $\{|k\rangle\}_{k=0}^{\infty}$ is the Fock basis. The state $\phi^{\mathbf{w}}$ is then defined as

$$\phi^{\mathbf{w}} := D(\sqrt{1 - 2\lambda_0} \alpha_{\mathbf{w}}) \phi^{\mathbf{0}} D(-\sqrt{1 - 2\lambda_0} \alpha_{\mathbf{w}}) \quad (2.62)$$

where $\alpha_{\mathbf{w}} = -u_y + iu_x \in \mathbb{C}$, and the operator $D(\alpha)$ is the displacement operator (2.15). We now state the convergence precisely in the following theorem.

Theorem 1. (Theorem 3.1 of [57]). *Let $\rho_{\mathbf{u}/\sqrt{n}}^{\otimes n}$ be a family of states defined by (2.59) on the Hilbert space $(\mathbb{C}^2)^{\otimes n}$, let $N^{\mathbf{u}}$ be the family of Gaussian distributions (2.60), and let $\phi^{\mathbf{w}}$ be the family of displaced thermal equilibrium states of a quantum harmonic oscillator (2.62) on the Hilbert space $\mathcal{H} = L^2(\mathbb{R})$. Then for each n and $\|\mathbf{u}\| \leq n^\eta$ there exists quantum channels (trace preserving CP maps)*

$$T_n : \mathcal{T}((\mathbb{C}^2)^{\otimes n}) \rightarrow L^1(\mathbb{R}) \otimes \mathcal{T}(\mathcal{H}) \quad (2.63)$$

$$S_n : L^1(\mathbb{R}) \otimes \mathcal{T}(\mathcal{H}) \rightarrow \mathcal{T}((\mathbb{C}^2)^{\otimes n}) \quad (2.64)$$

with $\mathcal{T}(\mathcal{H})$ being the trace-class operators on \mathcal{H} such that for any $0 \leq \eta < 1/6$

$$\lim_{n \rightarrow \infty} \sup_{\|\mathbf{u}\| \leq n^\eta} \|N^{\mathbf{u}} \otimes \phi^{\mathbf{w}} - T_n(\rho_{\mathbf{u}/\sqrt{n}}^{\otimes n})\|_1 = 0 \quad (2.65)$$

$$\lim_{n \rightarrow \infty} \sup_{\|\mathbf{u}\| \leq n^\eta} \|\rho_{\mathbf{u}/\sqrt{n}}^{\otimes n} - S_n(N^{\mathbf{u}} \otimes \phi^{\mathbf{w}})\|_1 = 0. \quad (2.66)$$

The above theorem shows that statistically the joint qubit state is asymptotically equivalent to the limit Gaussian system. The information about the eigenvalue is encoded in the classical Gaussian distribution, while information about the eigenvectors is contained in the state $\phi^{\mathbf{w}}$.

Although the LAN principle might seem to involve some prior assumptions about the state, namely that it exists in a local neighbourhood around a fixed and known state, it is possible to localise a completely unknown state in this manner. The idea is to perform measurements on a small and vanishing fraction $\tilde{n} \ll n$ of the total resource in order to construct a rough estimate ρ_0 of the true state. It can then be established – using standard concentration inequalities – that true state is within the required local neighbourhood of ρ_0 with high probability [57]. A detailed study of quantum LAN and its applications can be found in these references [57, 59, 60, 58, 56, 35, 75, 74].

2.5 The Minimax estimator

Let us return once again the problem of estimating an unknown parameter $\theta \in \Theta$ for the outcomes of n iid random variables $\mathbf{X} = (X_1, \dots, X_n)$. Typically there might be several possible estimators $\hat{\theta}_n$ to choose from. We may wish to

judge and compare the quality of these estimators in a way that is independent of the particular realisation of outcomes of the random variables (x_1, \dots, x_n) . Therefore, instead of considering something like the mean square error (MSE)- which is evaluated for a particular estimate - we look to the *risk* of the estimators. Given an estimator $\hat{\boldsymbol{\theta}}_n$ of the true parameter $\boldsymbol{\theta}$, the risk is defined as

$$R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}_n) = \mathbb{E} \left[D(\hat{\boldsymbol{\theta}}_n, \boldsymbol{\theta}) \right] \quad (2.67)$$

where the expectation is over all possible outcomes of the random variable \mathbf{X} . The function $D(\hat{\boldsymbol{\theta}}_n, \boldsymbol{\theta})$ is called the *loss function*, and examples for such loss functions are the L^2 norm, L^1 norm and the Hellinger distance. We shall return to a discussion of loss functions later in this section.

Estimators with a smaller risk might be considered ‘better’, and preferred over others. However, the risk $R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}_n)$ is a function of the true parameter, and in general there will be no single estimator whose risk will be uniformly lower than all other possible estimators over all values of $\boldsymbol{\theta} \in \Theta$. Additionally, as the true parameter is taken to be unknown, it is impossible to compare the risk of the various estimators at this value. How then can estimators be compared in a meaningful way? We could restrict our attention to estimators that satisfy certain desirable properties such as unbiasedness and from among all such estimators prefer ones that have the least variance over all values of $\boldsymbol{\theta}$, or achieve the CRLB. However, it might not always be preferable to consider only unbiased estimators, and estimators achieving the CRLB may not exist. Also, as we have seen the CRLB is not defined at the boundary of the parameter space. Another way to choose the ‘best’ estimator is to convert the risk into a single quantity instead of studying it as a function over the entire parameter space Θ . There are two methods that do precisely this, and the resulting estimators are called the *Bayes estimator* and the *Minimax estimator*.

The idea behind the Bayes principle is to assume that the parameter $\boldsymbol{\theta}$ is a random variable with a prior distribution π over the parameter space Θ . Then the *Bayes risk* with respect to the prior distribution π is defined as the expected risk

$$R_B(\pi, \hat{\boldsymbol{\theta}}_n) := \int_{\Theta} R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}_n) \pi(d\boldsymbol{\theta}). \quad (2.68)$$

This provides a ‘summary’ of the risk function, and a meaningful quantity with which to compare the various estimators. The estimator $\hat{\boldsymbol{\theta}}_n$ that minimises the Bayes risk for a given prior π is called the *Bayes estimator*. In the frequen-

tist model the true parameter is not considered to come from an underlying distribution, and in this case the risk is ‘summarised’ by the *maximum risk*

$$R_{\max}(\hat{\boldsymbol{\theta}}_n) := \sup_{\boldsymbol{\theta} \in \Theta} R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}_n). \quad (2.69)$$

An estimator that achieves the smallest possible maximum risk over all estimators is called the *Minimax estimator*. That is, such an estimator $\hat{\boldsymbol{\theta}}_n$ satisfies

$$\sup_{\boldsymbol{\theta} \in \Theta} R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}_n) = \inf_{\tilde{\boldsymbol{\theta}}_n} \sup_{\boldsymbol{\theta} \in \Theta} R(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}_n). \quad (2.70)$$

The term on the right in the above equation is called the *minimax risk*. The two estimators, Minimax and Bayes are linked through the notion of the *least favourable prior*. If $\hat{\boldsymbol{\theta}}_n$ is the Bayes estimator for some prior π such that

$$R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}_n) \leq R_B(\pi, \hat{\boldsymbol{\theta}}_n) \quad \forall \boldsymbol{\theta} \in \Theta, \quad (2.71)$$

then π is called a least favourable prior, and for some well-behaved models the estimator $\hat{\boldsymbol{\theta}}_n$ is minimax [17, 99].

The challenge in statistical decision theory is to both determine the minimax risk (2.70) and to determine an estimator that achieves this risk. For ‘good’ estimators and certain choices of the loss function, the minimax risk decreases at a rate of $1/n$, uniformly over all $\boldsymbol{\theta} \in \Theta$. However, it may not be possible to determine an exact minimax estimator for a fixed and finite n , instead we consider the rescaled asymptotic minimax risk

$$\limsup_{n \rightarrow \infty} \inf_{\tilde{\boldsymbol{\theta}}_n} n R_{\max}(\tilde{\boldsymbol{\theta}}_n), \quad (2.72)$$

and determine the estimator that achieves this rescaled risk. Such an estimator is said to be *asymptotically minimax*.

Let us now consider the quantum statistical model, where the aim is to determine an unknown density matrix $\rho \in \mathbb{S}_d$. Given n identically prepared copies of the state, we estimate the state from the outcomes of measurements performed on the copies. Given a measurement design M and the observed outcomes, the risk of the measurement-estimator pair is

$$R(M, \rho) := \mathbb{E}[D(\hat{\rho}_n, \rho)]. \quad (2.73)$$

The minimax risk for a given measurement design M , in keeping with (2.70), is defined as

$$R_{\text{minimax}}(M, \rho) = \inf_{\hat{\rho}_n} \sup_{\rho \in \mathbb{S}_d} \mathbb{E}[D(\hat{\rho}_n, \rho)], \quad (2.74)$$

where the supremum is taken over the space of all density matrices and the loss function $D(\hat{\rho}_n, \rho)$ is defined over complex matrices. Since the risk is defined for a given choice of loss function, the corresponding minimax and Bayes estimators may differ according to this choice.

2.5.1 Loss functions and quadratic approximations

The simplest choice for the loss function is perhaps the *square Frobenius* or L^2 matrix norm. The Frobenius distance between the density matrix ρ and its estimate $\hat{\rho}_n$ (not necessarily a state) is defined as

$$\|\rho - \hat{\rho}_n\|_F^2 := \text{Tr}[(\rho - \hat{\rho}_n)(\rho - \hat{\rho}_n)^*] = \sum_{i=1}^d \sum_{j=1}^d |(\rho - \hat{\rho}_n)_{i,j}|^2. \quad (2.75)$$

The right most equality in the above equation gives us a quadratic expansion of the Frobenius norm. More concretely, if we choose to parameterise the matrix $\rho \in \mathbb{S}_d$ by its matrix elements as in (2.33)

$$\begin{aligned} \boldsymbol{\theta} &= (\boldsymbol{\theta}^{(r)}, \boldsymbol{\theta}^{(i)}, \boldsymbol{\theta}^{(d)}) \\ &:= (\text{Re}\rho_{1,2}, \dots, \text{Re}\rho_{d-1,d}, \text{Im}\rho_{1,2}, \dots, \text{Im}\rho_{d-1,d}, \rho_{1,1}, \dots, \rho_{d,d})^T \in \mathbb{R}^{d^2} \end{aligned} \quad (2.76)$$

and denote its estimate as $\hat{\boldsymbol{\theta}}_n$, then the distance between the true state and the estimated state $\rho_{\hat{\boldsymbol{\theta}}_n}$ can be expressed as

$$\begin{aligned} \|\rho_{\boldsymbol{\theta}} - \rho_{\hat{\boldsymbol{\theta}}_n}\|_F^2 &= 2\|\boldsymbol{\theta}^{(r)} - \hat{\boldsymbol{\theta}}_n^{(r)}\|^2 + 2\|\boldsymbol{\theta}^{(i)} - \hat{\boldsymbol{\theta}}_n^{(i)}\|^2 + \|\boldsymbol{\theta}^{(d)} - \hat{\boldsymbol{\theta}}_n^{(d)}\|^2 \\ &= (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_n)^T G_F (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_n) \end{aligned} \quad (2.77)$$

where the matrix weight matrix G_F is the constant and diagonal $G_F = \text{Diag}(2 \cdot \mathbf{1}_{d(d-1)/2}, 2 \cdot \mathbf{1}_{d(d-1)/2}, \mathbf{1}_d)$, with $\mathbf{1}_{(\cdot)}$ specifying an identity matrix of dimension $(\cdot) \times (\cdot)$. A different choice of parameterisation will lead to a different form of the weight matrix. With this quadratic expansion, we can express the Frobenius risk as

$$\mathbb{E}[\|\rho - \hat{\rho}_n\|_F^2] = \mathbb{E}[(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_n)^T G_F (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_n)], \quad (2.78)$$

which we recognise from (2.49) to be the mean square error MSE. In general, for any given choice of parameterisation $\boldsymbol{\theta} \in \Theta$, if a loss function admits a quadratic expansion then it is seen to be equivalent to the MSE of the estimate, with a corresponding weight matrix that is defined by the choice of parametrisation. However, in general a loss function might not be quadratic over the entire parameter space Θ , but could be expanded quadratically for a choice of *local parameterisation*.

Let us consider the case where the state ρ is in a local neighbourhood around a fixed and known state ρ_0 , and let us choose a parameterisation of states $\rho_0 = \rho_{\boldsymbol{\theta}_0}$. The unknown parameter $\boldsymbol{\theta}$ of the true state ρ is taken to lie in a local neighbourhood of size $1/\sqrt{n}$ around the fixed parameter $\boldsymbol{\theta}_0$, and we define the local parameters $\mathbf{u} := \sqrt{n}(\boldsymbol{\theta} - \boldsymbol{\theta}_0)$. A loss function is *locally quadratic* if its risk can be approximated as

$$nR(M, \rho) = \mathbb{E} \left[(\mathbf{u} - \hat{\mathbf{u}}_n)^T G (\mathbf{u} - \hat{\mathbf{u}}_n) \right] + O \left(\|\mathbf{u} - \hat{\mathbf{u}}_n\|^3 \right). \quad (2.79)$$

As long as the fixed parameter $\boldsymbol{\theta}_0$ does not lie at the boundary of the parameter space, we know from classical LAN that the experiment is asymptotically equivalent to the Gaussian limit model of estimating the local parameter \mathbf{u} from a single observation of the random variable $Y \sim \mathcal{N}(\mathbf{u}, I_M(\boldsymbol{\theta}_0)^{-1})$. This implies that the rescaled local asymptotic risk is

$$\limsup_{n \rightarrow \infty} nR(M, \rho) = \text{Tr} \left[I_M(\boldsymbol{\theta}_0)^{-1} G \right] + O \left(\|\mathbf{u} - \hat{\mathbf{u}}_n\|^3 \right). \quad (2.80)$$

In the limit Gaussian model the minimum variance estimator is the observed value of Y itself. Therefore, the right side of the above equation gives us the rescaled local asymptotic minimax risk for the given measurement design M . As efficient estimators like the MLE achieve this rate asymptotically for parameters that lie away from the boundary of the parameter space, such estimators are asymptotically and locally minimax for locally quadratic loss functions.

We now introduce the common distance measures that will be used throughout this thesis. We have already introduced the Frobenius norm $\|A\|_F^2 := \text{Tr}[A^*A]$ for all matrices $A \in M(\mathbb{C}^d)$. We also consider the *Trace norm* $\|A\|_1 := \text{Tr}[\sqrt{AA^*}]$, and the *operator norm* $\|A\| := \sqrt{\lambda_{\max}(A^*A)} = \sigma_{\max}(A)$ which is the maximum singular value of the matrix. We also consider the

several commonly used measures of distance based on the *fidelity* between quantum states.

Fidelity based loss functions

The fidelity in itself is not a proper metric on the space of density matrices, but gives rise to a useful metric [120, 92]. The fidelity between two states $\rho, \sigma \in \mathbb{S}_d$ is defined as

$$F(\rho, \sigma) := \text{Tr} \left[\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}} \right]^2. \quad (2.81)$$

The *infidelity* is then defined as $1 - F(\rho, \sigma)$. It is easy to see that $0 \leq F(\rho, \sigma) \leq 1$, $\forall \sigma, \rho \in \mathbb{S}_d$. The fidelity for pure states $\rho = |\psi\rangle\langle\psi|$ and $\sigma = |\phi\rangle\langle\phi|$ is simply $F(\rho, \sigma) = |\langle\psi|\phi\rangle|^2$. The fidelity is used to define a proper metric on the space of density matrices, the *square Bures distance*

$$D_B(\rho, \sigma)^2 := 2 \left(1 - \sqrt{F(\rho, \sigma)} \right). \quad (2.82)$$

For states ρ and σ that are diagonal in the same basis, i.e

$$\rho = \sum_i^d r_i |i\rangle\langle i|, \quad \sigma = \sum_i^d s_i |i\rangle\langle i| \quad (2.83)$$

for some orthonormal basis $\{|i\rangle\}$, the square Bures distance is simply the classical *square Hellinger distance* between two probability distributions

$$\begin{aligned} D_B(\rho, \sigma)^2 &= 2 \left(1 - \sum_i^d \sqrt{r_i s_i} \right) \\ &= \sum_i^d (\sqrt{r_i} - \sqrt{s_i})^2 = D_H(\mathbf{r}, \mathbf{s})^2 \end{aligned} \quad (2.84)$$

where $\mathbf{r}, \mathbf{s} \in \mathbb{R}^d$ are the eigenvalue (or probability) vectors.

Chapter 3

Thesis rationale

This chapter is intended to briefly explain the motivation behind the work presented in part II of this thesis. We begin with a brief review of the compressed sensing paradigm employed in the estimation of low-rank states. This will serve as a good starting point from which to introduce the research questions investigated in chapters 4 and 5. A discussion related to the work in chapters 6 and 7 follows in sections 3.4 and 3.5 respectively.

3.1 Quantum tomography via compressed sensing

Imagine that we have an experimentally prepared quantum state of N qubits that we wish to estimate. The density matrix corresponding to this unknown state is a $d \times d$ density matrix $\rho \in \mathbb{S}_d$, where $d = 2^N$ is the dimension of the Hilbert space \mathcal{H}^d . Simple parameter counting informs us that there are $d^2 - 1$ unknown parameters that completely specify the state. In the MIT setup (section 2.3.1) the standard measurements involve measuring one of the three Pauli observables $\{\sigma_x, \sigma_y, \sigma_z\}$ on each individual qubit simultaneously. We index the set of all such measurement by the settings

$$\{\mathbf{s} = (s_1, \dots, s_N) \in \{x, y, z\}^N\}. \quad (3.1)$$

As there are 3^N settings in total, this gives a total set of $3^N \times 2^N$ projectors, which is far larger than the number of unknown parameters $d^2 - 1$. Thus this measurement design is highly over-complete in $M(\mathbb{C}^2)$. The state can therefore be estimated from the outcomes of repeated measurements in the 3^N settings, and as we have briefly seen in the previous chapter, several estimators such as the MLE or the LSE could be used to recover the state.

We see that this is a general method. There are no assumptions made about the unknown state, nor has any prior information been incorporated. If however, we have prior knowledge that the prepared state is likely to be pure, or close to pure, then we know that the unknown state ρ has low rank. A simple counting of parameters tells us that a state of rank r is completely specified by $O(rd)$ parameters, and for low rank states this is a much smaller number than $d^2 - 1$. This implies that in principle a smaller number of measurements would be sufficient to recover the unknown low-rank state. It is easy to see that if *we knew the structure of the density matrix beforehand* (for example the subspace on which ρ is supported), this is certainly true. However, this begs the question of how a general tomographic technique is to be implemented when we do not know the structure of the density matrix, but only that it is of low rank. Inspired by the techniques and results of classical compressed sensing (CS) [29, 88, 30, 40, 14, 31] and matrix recovery problems [27, 97, 28], the results in [45, 54, 85] first demonstrated such a general tomographic technique based on the idea that the low rank (or *sparsity*) of states can be exploited to estimate them with a reduced number of ‘measurements’.

The CS tomographic setup as in [45, 54, 85], involves the measurement of the Pauli observables

$$\sigma_{\mathbf{b}} := \sigma_{b_1} \otimes \dots \otimes \sigma_{b_N} \quad b_i \in \{0, x, y, z\}, \quad (3.2)$$

where $\sigma_0 := \mathbf{1}$ is the identity matrix. There are 4^N such observables in total, and in the CS setup only a fraction of these observables are measured. Anticipating the discussion that follows, only $O(rd \text{ poly log } d)$ randomly chosen observables are sufficient to robustly reconstruct an unknown and arbitrary rank- r density matrix. Let k denote the number of observables that are chosen randomly from the full set of 4^N . Measurements of each observable is repeated on m identical copies of the state, producing a binary outcome $\{+1, -1\}$ on each measurement. The outcomes are then averaged to approximate the expected value $\langle \sigma_{\mathbf{b}} \rangle := \text{Tr}[\sigma_{\mathbf{b}}\rho]$. The essential problem in the CS setup is to solve the following under-determined system of linear equations

$$\mathbf{f} = \mathcal{A}(\rho) + \boldsymbol{\epsilon} \quad (3.3)$$

where $\mathbf{f} \in \mathbb{R}^k$ is a vector of the estimated expected values, and $\boldsymbol{\epsilon} \in \mathbb{R}^k$ is a statistical noise vector. The linear map $\mathcal{A} : \mathcal{H}^d \mapsto \mathbb{R}^k$ is defined for a particular

choice of Pauli observable as

$$(\mathcal{A}(\rho))_{\mathbf{b}} := \langle \sigma_{\mathbf{b}} \rangle_{\rho}. \quad (3.4)$$

The key idea here is that it is possible to recover a low-rank state from $k \ll 4^N$ Pauli observables as the operator \mathcal{A} embeds the manifold of rank- r matrices in $M(\mathbb{C}^d)$ into a lower dimensional vector space, while preserving distances between matrices in the lower dimensional space of vectors [45]. This property of the map is called the *restricted isometry property* (RIP).

Definition 2. (see [45, 85]) *The operator \mathcal{A} is said to satisfy the rank- r restricted isometry property if there exists some constant $0 \leq \delta_r < 1$ such that, for all $X \in M(\mathbb{C}^d)$ with rank r*

$$(1 - \delta_r) \|X\|_F \leq \|\mathcal{A}(X)\| \leq (1 + \delta_r) \|X\|_F \quad (3.5)$$

In [85], it has been show that with high probability the RIP property holds for the randomly chosen Pauli observables described above, provided that the number of observables measured is of the order $O(rd \text{ poly log } d)$. The randomness in the choice of observables is a necessary component of the proof. In particular it has been shown that if the map \mathcal{A} satisfies the RIP, then the intersection between the space of all solutions to linear system $y = \mathcal{A}(\rho)$ (noiseless scenario), and the space of rank- r density matrices is unique [54].

Given the vector $f \in \mathbb{R}^k$ of the estimated expected values, and the map \mathcal{A} satisfying the RIP with $k \approx O(rd \text{ poly log } d)$, how is the unknown low-rank state recovered? The CS literature proposes the following estimator, also called the matrix Dantzig selector [45]

$$\hat{\rho}_{\text{DS}} := \arg \min_X \|X\|_1 \quad \text{s.t.} \quad \|\mathcal{A}^*(\mathcal{A}(X) - \mathbf{f})\| \leq c \quad (3.6)$$

where c is a set upper-bound on the noise, and \mathcal{A}^* is the adjoint of the operator \mathcal{A} . The choice of minimising the *trace norm* is a very particular one. Since the unknown state ρ is known to be sparse, in that it is of low-rank, one might imagine estimating ρ by implementing a search over all low-rank matrices that best fit the data. However, this search is known to be NP-hard, and the closest convex relaxation is the programme above (3.6) [104]. An intuition for why this might be is gained by remembering that the trace norm is the sum of the singular values of a matrix. As the rank of a matrix equals the number of

non-zero singular values, the problem of rank minimisation is well approximated by the trace norm minimisation.

The performance of the Dantzig selector has been shown to be robust in the presence of noise, and bounds for the error in the trace norm are obtained in [45]. Closely related work in [34], proposes an iterative thresholding estimator that uses a computationally efficient, explicit algorithm instead of the convex optimisation problem (3.6). This method imposes a property, closely related to the RIP, on the linear map \mathcal{A} . This estimator is also shown to perform well, and bounds for both the entry-wise error of the resulting estimate and its Frobenius error are derived.

3.1.1 Coarse grained statistics

As it turns out, the CS model as described above does not make the most efficient use of the outcome statistics produced from the tomographic experiment. To consider this in more detail, let us compare the measurement of the Pauli observables

$$\sigma_{\mathbf{b}} = \sigma_{b_1} \otimes \dots \otimes \sigma_{b_N} \quad b_i \in \{0, x, y, z\}, \quad (3.7)$$

to the standard Pauli settings defined earlier

$$\mathbf{s} = (s_1, \dots, s_N) \in \{x, y, z\}^N. \quad (3.8)$$

There are 3^N such settings in total, and clearly the set of all observables $\sigma_{\mathbf{s}}$, indexed by the settings \mathbf{s} , is a subset of the set of all 4^N Pauli observables $\sigma_{\mathbf{b}}$. However, it is possible to estimate the expectation values $\langle \sigma_{\mathbf{b}} \rangle_{\rho}$ of all 4^N Pauli observables by measuring the full 3^N settings. To see this, recall that the outcome of a measurement in setting \mathbf{s} is a vector record $\mathbf{o} = (o_1, \dots, o_N) \in \{+1, -1\}^N$ of the ± 1 outcomes from each of the N ions, while the outcomes in the CS setup are simply binary $\{\pm 1\}$, and do not keep a record of the individual outcomes from each of the ions. In fact from single setting \mathbf{s} , it is possible to compute the expectations

$$\langle \sigma_{\mathbf{b}} \rangle_{\rho} = \text{Tr}(\rho \sigma_{\mathbf{b}}) = \sum_{\mathbf{o}} \left(\prod_{i: b_i \neq 0} o_i \right) p_{\rho}(\mathbf{o} | \mathbf{s}). \quad (3.9)$$

for all Pauli observables in the set $\{\sigma_{\mathbf{b}} | \forall i: b_i \neq 0, b_i = s_i\}$, by simply estimating the probabilities $p_{\rho}(\mathbf{o} | \mathbf{s})$ from the data. It is easy to see that a single setting \mathbf{s} therefore generates information about d such expectation values. The Pauli

expectations are a ‘coarse graining’ of the ‘raw’ statistical data that is obtained from measuring the Pauli settings. This raises the question if ‘compressive behaviour’ can be demonstrated in the standard MIT setup with measurements in the standard Pauli settings. In a sense this is certainly true as one could randomly pick $O(rd \text{ polylog } d)$ Pauli observables $\sigma_{\mathbf{b}}$, and estimate their expectations from a (possibly) smaller set of settings \mathbf{s} . However this procedure would still involve the coarse-graining of data and relies on the map \mathcal{A} satisfying the RIP. The question is to determine if compressive behaviour exists with the raw data *as is*.

3.2 The projection estimator

In this section we propose a simple estimator that demonstrates ‘compressive recovery’ of low-rank states with incomplete measurements in the MIT setup. This estimator uses the LSE of the true state as a starting point, and we therefore begin by defining the linear regression model for incomplete measurements. As before (section 2.3.1), let us consider an unknown state $\rho \in \mathbb{S}_d$ of N qubits, and we define $\mathcal{S} := (\mathbf{s}_1, \dots, \mathbf{s}_k)$ as the set of k randomly chosen settings $\mathbf{s} \in \{x, y, z\}^N$ from the full 3^N settings. Let n be the total number of copies of the state, and we let $m = n/k$ be the number of times a measurement is repeated in a given setting \mathbf{s} . The measurement outcomes are collected in a dataset \mathcal{D} of counts $N(\mathbf{o}|\mathbf{s})$. We would like to describe the linear estimation problem in terms of the basis formed by the Pauli observables $\sigma_{\mathbf{b}}$ (3.2), and we follow the convention in [25]. Consider the expansion of ρ in the Pauli tensor product basis in $M(\mathbb{C}^d)$

$$\rho = \sum_{\mathbf{b}} \rho_{\mathbf{b}} \sigma_{\mathbf{b}}, \quad \rho_{\mathbf{b}} = \frac{1}{d} \text{Tr}[\rho \sigma_{\mathbf{b}}]. \quad (3.10)$$

In this basis, the probabilities of obtaining a given outcome \mathbf{o} in a setting \mathbf{s} are expressed as

$$p_{\rho}(\mathbf{o}|\mathbf{s}) = \sum_{\mathbf{b}} \rho_{\mathbf{b}} \text{Tr}[\sigma_{\mathbf{b}} P_{\mathbf{o}}^{\mathbf{s}}] = \sum_{\mathbf{b}} \rho_{\mathbf{b}} A_{\mathbf{b}}(\mathbf{o}|\mathbf{s}), \quad (3.11)$$

where,

$$A_{\mathbf{b}}(\mathbf{o}|\mathbf{s}) = \prod_{j \notin E_{\mathbf{b}}} o_j \mathbb{1}(b_j = s_j), \quad \text{with } E_{\mathbf{b}} := \{j : b_j = 0\}, \quad (3.12)$$

with $\mathbb{1}(b_j = s_j)$ being the indicator function. We let $\boldsymbol{\theta} \in \mathbb{R}^{4^n}$ be a parameter vector of the coefficients $\rho_{\mathbf{b}}$, and $\mathbf{y} \in \mathbb{R}^{2^n \times k}$ be the vector of probabilities. This

allows us to express the equations (3.11) in matrix notation as

$$\mathbf{y} = A\boldsymbol{\theta}, \quad (3.13)$$

where A is a $(2^n \times k) \times 4^n$ design matrix. In an experimental set up, instead of the probabilities $p_\rho(\mathbf{o}|\mathbf{s})$, we have access to the empirical frequencies $f(\mathbf{o}|\mathbf{s}) := \frac{N(\mathbf{o}|\mathbf{s})}{m}$. Therefore, replacing \mathbf{y} by the vector of frequencies \mathbf{f} the problem of determining an unknown quantum state can be cast as

$$\mathbf{f} = A\boldsymbol{\theta} + \boldsymbol{\epsilon} \quad (3.14)$$

where, $\boldsymbol{\epsilon}$ is a mean zero vector of statistical noise. When all 3^N measurements are made, i.e. ($k = 3^n$), the above equation admits the usual least squares solution (section (2.3.2))

$$\hat{\boldsymbol{\theta}} = (A^T A)^{-1} A^T \mathbf{f} \quad (3.15)$$

The estimate $\hat{\rho}_{\text{LS}}$ of the density matrix is then reconstructed using the Pauli basis expansion (3.10). The matrix $A^T A$ is invertible and diagonal when all 3^N settings are measured. However, in the incomplete measurement setup, this matrix is not invertible and has a number of its diagonal entries equal to zero. To see this, we note that for a given $\sigma_{\mathbf{b}}$, the term $A_{\mathbf{b}}(\mathbf{o}|\mathbf{s})$ is non-zero for $3^{|\mathcal{E}_{\mathbf{b}}|}$ settings \mathbf{s} . If all of these $3^{|\mathcal{E}_{\mathbf{b}}|}$ settings are not measured then the diagonal entries $(A^T A)_{\mathbf{b},\mathbf{b}}$ are seen to be zero. Let us define the set $\mathcal{B} := \{\mathbf{b} | (A^T A)_{\mathbf{b},\mathbf{b}} = 0\}$.

Since the matrix $A^T A$ is non-invertible for $k < 3^N$, we may instead implement a pseudo-inverse in (3.15), and invert only the non-zero diagonal elements of the matrix. With this, it can be shown that the estimate $\hat{\rho}_{\text{LS}}$ constructed from the estimated parameter vector $\hat{\boldsymbol{\theta}}$ is such that

$$\text{Tr}[\hat{\rho}_{\text{LS}}\sigma_{\mathbf{b}}] = 0 \quad \forall \mathbf{b} \in \mathcal{B}. \quad (3.16)$$

Therefore the LSE in the reduced measurement setup is not an estimate of the whole state, and lies in a subspace orthogonal to the one spanned by the observables $\{\sigma_{\mathbf{b}} | \mathbf{b} \in \mathcal{B}\}$. Following this, the Frobenius error of the LSE can be expressed as the sum

$$\|\hat{\rho}_{\text{LS}} - \rho\|_F^2 = \left\| \sum_{\mathbf{b} \in \mathcal{B}} \rho_{\mathbf{b}} \sigma_{\mathbf{b}} \right\|_F^2 + \left\| \sum_{\mathbf{b} \notin \mathcal{B}} \rho_{\mathbf{b}} \sigma_{\mathbf{b}} - \hat{\rho}_{\text{LS}} \right\|_F^2. \quad (3.17)$$

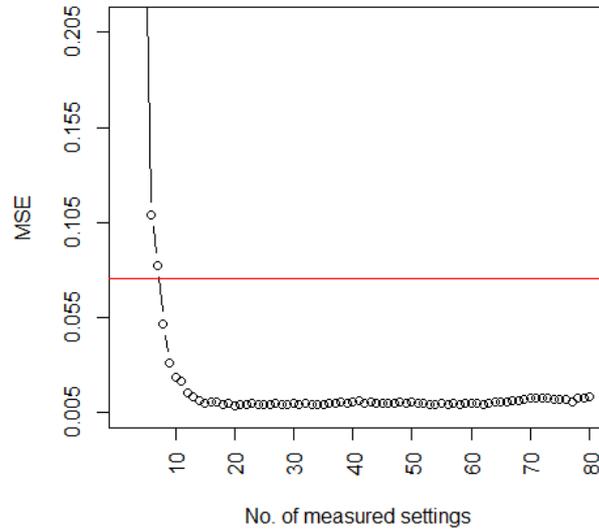


Figure 3.1 The average Frobenius error of the projection estimator for a random 4-qubit pure state, and total sample size $n = m \cdot 3^4 = 8100$. For each number of settings k (horizontal axis) we chose 30 random sets of settings \mathcal{S} , and each plotted point is the average of the 30 different resulting estimates. The red line is the Frobenius error of the full LSE with 3^N settings measured.

As fewer settings \mathbf{s} are measured, the contribution to the error from the first term on the right increases as a greater number of coefficients will not be estimated. The problem now becomes one of estimating the remaining coefficients $\rho_{\mathbf{b}} = \langle \sigma_{\mathbf{b}} \rangle / d$ of the observables $\{\sigma_{\mathbf{b}} | \mathbf{b} \in \mathcal{B}\}$. Although this is similar to the CS problem of estimating an unknown state from the incomplete set of expectations $\langle \sigma_{\mathbf{b}} \rangle$ of Pauli observables, there are a few key differences that are worth keeping in mind.

1. Even though we expand the LSE in terms of the coefficients $\rho_{\mathbf{b}}$ and therefore in terms of the expectations $\langle \sigma_{\mathbf{b}} \rangle$, there is no ‘coarse graining’ of the outcome statistics in the actual estimation. This is seen from the system of linear equations (3.15), where the vector \mathbf{f} is a vector of frequencies $N(\mathbf{o}|\mathbf{s})/m$ and not a vector of estimated expected values as in (3.3). In fact, in chapter 4 we shall demonstrate that any estimator using the Pauli expectations will have a significantly larger estimation error than if the ‘fine grained’ outcome statistics are used.

2. In the CS setup we have seen that low-rank states can be compressively recovered provided that $O(rd \text{ polylog } d)$ randomly chosen expectation values are estimated accurately. These expectation values are picked randomly and *uniformly* from a set of all 4^N such expectations, whereas in the MIT setup with incomplete measurements, the coefficients $\rho_{\mathbf{b}}$ that are estimated are not picked randomly from the full set. Instead, we pick uniformly from the set of 3^N measurable settings \mathbf{s} , and each setting generates information about d coefficients. Therefore the randomness over the coefficients $\rho_{\mathbf{b}}$ here is different than the randomness employed in the CS setup.
3. Because we use ‘fine grained’ outcome statistics and each setting \mathbf{s} generates information about several coefficients, we may reasonably expect the minimum number of ‘measurements’ needed to be smaller than in the CS setup.

3.2.1 Alternating projections

We now introduce the following projection estimator that uses $\hat{\rho}_{\text{LS}}$ to build an estimate of the whole state ρ . The procedure involves repeated, alternating projections - starting from the estimate $\hat{\rho}_{\text{LS}}$ - onto the subspace of rank- r density matrices. The following algorithm outlines the basic procedure, where L is taken to be the number of iterations of the procedure.

Algorithm 1: The Projection Estimator

Input : The least squares estimator $\hat{\rho}_{\text{LS}}$, the rank r of the state ρ , and the number of iterations L .

Output : The projection estimate $\hat{\rho}_{\text{proj}}$

```

1  $l \leftarrow 1$ 
2  $\hat{\rho}_{l=1} \leftarrow \hat{\rho}_{\text{LS}}$ 
3 while  $l \leq L$  do
4   | Project the estimate  $\hat{\rho}_l$  onto a rank- $r$  dimensional subspace by
   | retaining only its  $r$ -largest eigenvalues, and setting the rest to zero
   |  $\hat{\rho}_l^r \leftarrow \sum_{i=1}^r \hat{\lambda}_i |\hat{\lambda}_i\rangle \langle \hat{\lambda}_i|$ 
5   |  $\hat{\rho}_{l+1} \leftarrow \hat{\rho}_{\text{LS}} + \sum_{\mathbf{b} \in \mathcal{B}} \frac{1}{d} \text{Tr}[\hat{\rho}_l^r \sigma_{\mathbf{b}}] \sigma_{\mathbf{b}}$ 
6   |  $l \leftarrow l + 1$ 
7 end
8  $\hat{\rho}_{\text{proj}} \leftarrow \hat{\rho}_{l=L}$ 

```

Figure 3.1 plots the average Frobenius error $\|\hat{\rho}_{\text{proj}} - \rho\|_F^2$ for a random 4-qubit pure state. The number of repetitions in each setting is $m = n/k$, where $n = 8100$ is the total number of copies of the state. The total number of iterations $L = 10$, and we chose 30 random sets of settings \mathcal{S} for each k . The plotted points correspond to the Frobenius error averaged over the 30 different estimates. We note that the estimator defined in Algorithm 1 assumes knowledge of the true rank of the state, and is therefore an ‘oracle’ estimator. Its performance, even with incomplete measurements, remains better than the LSE obtained with full measurement settings (red line). Additionally the estimation errors remain fairly constant, and this shows that a 4-qubit pure state can be estimated by using only 10-15 settings, out of a total of 81. This demonstrates ‘compressive’ recovery with the Pauli settings, and although the true rank is typically unknown, this can be estimated from the outcome data itself. This is done with techniques like Cross-Validation [25].

3.3 Compressive tomography with fine grained statistics

Numerical simulations with the projection estimator demonstrates the possibility of ‘compressively’ recovering a low-rank state in the MIT setup. However, as this estimator was proposed only as an illustrative example, we do not provide any bounds for the estimation error, nor derive any rate for the minimum number of settings \mathbf{s} that need to be measured. In chapters 4 and 5 we study this ‘compressive’ behaviour in more general terms. Instead of considering a particular estimator (like the Dantzig selector or the projection estimator), we would like to consider the estimation errors of the general class of ‘efficient’ estimators, i.e, ones which achieve the CRLB asymptotically. The two central questions that need to be answered are as follows. Can we demonstrate the recovery of low-rank states with an incomplete measurement design and such ‘efficient’ estimators? Secondly, can we say something meaningful about the scaling of the minimum number of settings that need to be measured?

As we have seen in chapter 2, in the asymptotic regime efficient estimators $\hat{\rho}_n$ achieve the CRLB, and their MSE is given by

$$\mathbb{E}\|\hat{\rho}_n - \rho\|_F^2 \approx \frac{1}{n} \text{Tr} [I(\rho|\mathcal{S})^{-1} G_F], \quad (3.18)$$

where the classical Fisher information matrix $I(\rho|\mathcal{S})$ depends on both the true state and the measurement design \mathcal{S} , and the constant weight matrix G_F is

particular to the choice of the Frobenius norm. Since the asymptotic MSE depends only on the local properties of the model, the Fisher information above is defined by a choice of local parameterisation of rank- r states.

From (3.18) we see that in the asymptotic regime, the problem of studying the MSE of efficient estimators can be converted into one of studying the behaviour of the Fisher information matrix. In chapters 4 and 5, we do precisely this and find that even as the number of settings k are reduced, the asymptotic MSE for low-rank states remains robust. In order to explain this ‘compressive’ behaviour of the MSE, we look to derive a concentration bound for the Fisher information matrix, and to show that $I(\rho|\mathcal{S})$ remains close to its optimal value even when the number of settings k making up the measurement design \mathcal{S} is small $k \ll 3^N$.

However, deriving such a concentration bound requires knowledge of certain spectral properties of the Fisher information matrix. This proves to be challenging for the Pauli settings case, and remains an open problem. Instead we relax the Pauli settings setup and consider *random basis measurements* as introduced in section 2.3.1. In this setup the necessary spectral properties of the Fisher information are derived and we show that certain “least sparse” rank- r states can be estimated with only $O(r \log d)$ measurements, with the resulting MSE demonstrating only a small increase relative to the setup where a large number of random bases are measured.

3.3.1 A problem at the boundary

This concentration bound of the Fisher information matrix is a powerful tool, in that similar ‘compressive’ results can immediately be derived for any distance measure that admits a locally quadratic expansion. We simply replace the G_F in (3.18) with the weight matrix that corresponds to the chosen distance measure. However, in chapter 5 we show that for states that have very small eigenvalues, the Fisher information $I(\rho|\mathcal{S})$ does not concentrate. In fact, we extend the results in chapter 4 to show that for states with arbitrary spectrums, the minimum number of settings scales as $O\left(\frac{1}{\lambda_{\min}(\rho)} \frac{r+1}{r} \log d\right)$. This would suggest that as $\lambda_{\min}(\rho) \rightarrow 0$ the minimum number of measurements required for the Fisher information to concentrate tends to infinity $k \rightarrow \infty$! It is natural to wonder if this lack of concentration in the Fisher information also implies a lack of concentration in the MSE for states that have very small eigenvalues.

We show that even at the boundary of the parameter space, the asymptotic MSE (3.18) for an arbitrary rank- r state scales as $O(1/n)$, with a constant that is bounded by roughly the number of unknown parameters given that $O(r \log d)$ random bases are measured. In the case of a single qubit, we show that the MSE concentrates as $\lambda_{\min} \rightarrow 0$, even if the Fisher information does not.

However, not all distances show a similar concentration in the mean error as $\lambda \rightarrow 0$. We consider a simple illustration of the problem. Let $\rho_0 = \text{Diag}(1 - \lambda_0, \lambda_0)$ be a qubit state diagonal in its own eigenbasis, and let $\boldsymbol{\theta} := (\lambda, u, v) \in \mathbb{R}^3$ be a local parameterisation of states such that any state $\rho_{\boldsymbol{\theta}}$ in this local neighbourhood is given by

$$\rho_{\boldsymbol{\theta}} = U(u, v) \begin{pmatrix} 1 - \lambda & 0 \\ 0 & \lambda \end{pmatrix} U(u, v)^*, \quad (3.19)$$

and is obtained by small changes in the eigenvalues and a small rotation of the eigenbasis of ρ_0 . In this parameterisation let us consider the infidelity between states $1 - F(\rho_{\boldsymbol{\theta}}, \rho_{\boldsymbol{\theta} + \delta\boldsymbol{\theta}})$, and perform a one sided Taylor expansion (with $\delta\lambda > 0$) around $\lambda = 0$. We get

$$1 - F(\rho_{\boldsymbol{\theta}}, \rho_{\boldsymbol{\theta} + \delta\boldsymbol{\theta}}) = \delta\lambda + O(\|\delta\boldsymbol{\theta}\|^2). \quad (3.20)$$

This shows that the distance is no longer quadratic near the boundary of the Bloch sphere. This is true not just for the infidelity but also for the Bures distance. This illustrates the fact that the infidelity (and the Bures distance) are sensitive to the misestimation of small eigenvalues. In chapter 5 we demonstrate that due to this linearity in the parameters near the boundary, a concentration in the mean infidelity does not hold. Additionally, since in general the parameters are estimated with uncertainty of the order $O(1/\sqrt{n})$, (3.20) implies that the mean infidelity for states near the boundary scales only as $O(1/\sqrt{n})$. We will also demonstrate that in general one can always significantly decrease the estimation error by measuring more random bases, and this means a lack of ‘compressive’ behaviour for low-rank states with eigenvalues close to zero.

3.4 Adaptive protocols and the minimax rate

The work in chapters 4 and 5 focuses on questions of ‘compressive’ recovery. In chapter 6, we do not consider a reduced measurement setup, but instead study

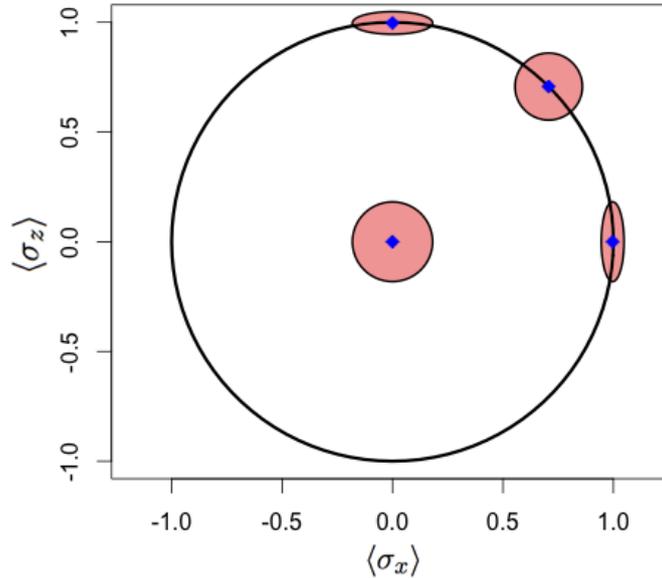


Figure 3.2 The standard deviation ellipses for several rebit states (blue diamonds). The red ellipses represent the uncertainly ‘balls’ as determined by (3.22) for large values of n . The four rebit states we consider are the fully mixed state $(r_x, r_z) = (0, 0)$, the state $(r_x, r_z) = (\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$ and the states $(r_x, r_z) = (0, 1 - n^{-1/2})$, $(r_x, r_z) = (1 - n^{-1/2}, 0)$.

the problem of estimating nearly pure qubit states with full tomography. As briefly hinted at in the previous section, the infidelity and Bures distance are no longer quadratic close to the surface of the Bloch sphere, and this has important implications for the rate at which nearly pure states can be estimated.

To better illustrate this behaviour at the boundary, let us continue with the qubit example and consider an arbitrary mixed state $\rho \in \mathbb{S}_2$. For the purposes of this discussion let us fix the parametrisation $\boldsymbol{\theta} = (r_x, r_y, r_z)$ with the Bloch vector components. The standard tomographic measurement procedure involves the repeated measurement of a Pauli observable in each setting $s \in \{x, y, z\}$ on $m = n/3$ copies of the state. The classical Fisher information for this measurement design and parametrisation can be easily evaluated to be

$$I(\boldsymbol{\theta}) := \begin{pmatrix} \frac{1}{1-\langle \sigma_x \rangle^2} & 0 & 0 \\ 0 & \frac{1}{1-\langle \sigma_y \rangle^2} & 0 \\ 0 & 0 & \frac{1}{1-\langle \sigma_z \rangle^2} \end{pmatrix}. \quad (3.21)$$

We have seen that the distribution of efficient estimators is asymptotically Gaussian, i.e., $\sqrt{m}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \rightarrow \mathcal{N}(0, I(\boldsymbol{\theta})^{-1})$. This implies that the standard deviation matrix for such estimators is given by

$$\sqrt{\Sigma} = \sqrt{\frac{3}{n}} \begin{pmatrix} \sqrt{1 - \langle \sigma_x \rangle^2} & 0 & 0 \\ 0 & \sqrt{1 - \langle \sigma_y \rangle^2} & 0 \\ 0 & 0 & \sqrt{1 - \langle \sigma_z \rangle^2} \end{pmatrix}. \quad (3.22)$$

This standard deviation matrix gives us the distribution of the estimates around the true state, and in figure 3.2 we plot the standard deviation ellipses around rebit states (setting $\langle \sigma_y \rangle = 0$) for large n and different values of $\langle \sigma_x \rangle$, $\langle \sigma_z \rangle$. We mention the following points about the three regions in figure 3.2.

1. When the state is well within the Bloch sphere, the standard deviation ‘ball’ is of the order $O(1/\sqrt{n})$. This implies that the estimates $\hat{\rho}_n = \rho_{\hat{\boldsymbol{\theta}}_n}$ lie in a region of uncertainty of $O(1/\sqrt{n})$ around the true state. This together with the fact that distance measures such as the infidelity and the Bures distance (along with several other norms) are locally quadratic in this region, gives us the standard rate of estimation $O(1/n)$ from the CRLB.
2. When the state is nearly pure but is not aligned along one of the measurement axis, then the standard deviation ‘ball’ is still of the order $O(1/\sqrt{n})$ as both $r_x, r_z \not\approx 1$, however some of the estimates are seen to lie outside the Bloch sphere and do not represent physical states. This fact has important consequences.

We first note that the Fisher information matrix is still defined in this region. This is because we have not constrained the parameters to represent states ($\|\mathbf{r}\| \leq 1$), and secondly the parameters r_x, r_z are not close to the boundary of the parameter space. Along with the fact that distances such as the squared Frobenius norm –which are valid even for unphysical estimates– are locally quadratic in this region, this implies that the corresponding MSE of certain ‘efficient’ estimators asymptotically achieves the CRLB with the standard $O(1/n)$ scaling.

However there are implications for estimators such as the MLE that produce only certifiable states [102]. We shall return to this point in a later section. For now, we note that the CRLB is not meaningful for

distance measures such as the infidelity or the Bures distance which are defined only over the space of density matrices. In fact, these distances are no longer quadratic close to the boundary (3.20), and this results in the poor $O(1/\sqrt{n})$ scaling for nearly pure states.

3. When the state is nearly pure and its Bloch vector is aligned along one of the measurement axes, then the standard deviation is no longer $O(1/\sqrt{n})$ in all directions. This can easily be seen by considering a state along the σ_z direction with $|\mathbf{r}| = r_z = 1 - O(1/\sqrt{n})$. Evaluating the corresponding terms in the matrix (3.22), we see that $\mathbb{E}[(\hat{r}_z - r_z)^2] = O(n^{-3/2})$. As we shall shortly explain, this fact makes it possible to improve the poor scaling of the fidelity based metrics, and recover the standard $O(1/n)$ scaling (at least in the asymptotic limit).

We also note that the Fisher information is only defined for values of the parameter within the parameter space. Therefore the Fisher information matrix and the CRLB are not valid in the limit $r_x, r_z \rightarrow 1$, and certainly not for values of $r_x, r_z > 1$.

The above discussion shows that for states near the boundary of the Bloch sphere and fidelity based distances, the CRLB does not hold. In fact, the estimation error scales as $O(1/\sqrt{n})$ for nearly pure states. However for regions within the Bloch sphere where the distances are locally quadratic, we have the standard scaling of $O(1/n)$ from the CRLB. This poor scaling of the rate for states near the boundary is true not only of the standard tomographic procedure, but holds for all fixed bases measurement designs. We recommend the papers [89, 102] for a further discussion of this problem, and the work in [42] which considers a classical analogue in the estimation of the bias of a noisy coin.

We also see that if the state is along one of the measurement axes then the corresponding Bloch vector component is estimated more accurately and with smaller uncertainty. This suggests a method to recover the $O(1/n)$ scaling for all nearly pure states; instead of the standard tomographic measurements we measure along the eigenbasis of the state. It can be shown that for such measurements the infidelity scales as $O(1/n)$ even for states at the boundary [89]. Clearly this measurement protocol is impossible to implement in practice as the state and its eigenbasis is unknown. However, based on this idea several papers propose two-step adaptive protocols that work by first producing a preliminary estimate with a fraction of the copies of the state n_1 , and then performing

measurements along the estimated eigenbasis on the remaining $n - n_1$ copies [49, 100, 9, 89, 70]. Adaptive protocols involving collective measurements have also been proposed [8, 7, 57, 59]. For a good review of the theoretical and experimental results with such adaptive schemes see [105].

In chapter 6, we consider this problem from the perspective of the maximum risk over all qubit states. Specifically, we consider the following asymptotic rescaled maximum Bures risk

$$r_{\max}(\hat{\rho}_n) := \limsup_{n \rightarrow \infty} \sup_{\rho \in \mathbb{S}_2} n \mathbb{E} \left[D_B(\rho, \hat{\rho}_n)^2 \right]. \quad (3.23)$$

We propose and analyse two adaptive estimation strategies, one based on separable adaptive measurements, and the other based on collective measurements which uses results of quantum LAN. For both estimation methods, we demonstrate a scaling of $O(1/n)$ of the maximum Bures risk. We also discuss how to construct a minimax optimal estimator in the setup with the collective measurements. Finally, we consider another distance measure, the quantum relative entropy, and show that no estimator can have maximum risk converging faster than $O(n^{-1} \log n)$ under this loss function.

3.5 A simulation study

As discussed in the previous section, the Fisher information matrix $I(\boldsymbol{\theta})$ is defined even for states near the boundary as long as the parameters $\boldsymbol{\theta}$ are well within the parameter space Θ . In the qubit case this boundary is at the ‘surface’ of the Bloch sphere, and we saw the Fisher information is defined as long as the Bloch vector parameters r_x, r_y, r_z are away from 1.

Since $I(\boldsymbol{\theta})$ is defined even for states near the boundary (under certain conditions), we may expect the variance of certain unbiased and ‘efficient’ estimators to be given by the inverse $I(\boldsymbol{\theta})^{-1}$ for suitably large n . As we have seen in the previous section, if the state is close enough to the boundary this implies that some of the estimates of such estimators will lie outside the space of states. However, this cannot reflect the distribution of estimators such as the MLE, since its estimates are constrained by positivity to produce states. Therefore the MLE is *not* an unbiased estimator of states near the boundary. In fact, for such states the constraint of positivity provides more information than encoded by the Fisher matrix, and it is reasonable to expect the risk of the MLE to break the CRLB, while for states away from the boundary, the risk

of the MLE is approximated by the asymptotic value $\text{Tr} [I(\boldsymbol{\theta})^{-1}G]$ for large n . This discussion is of course meaningful only when the loss function continues to be quadratic at the boundary (e.g the square Frobenius norm), as the CRLB is not meaningful at the boundary for distances that are defined only over the space of states.

This non-applicability of the asymptotic theory at the boundary is an important feature of the MLE. The above discussion however suggests that the asymptotic theory *should* apply to an ‘unconstrained’ version of the MLE, where the estimates are not required to be positive. In chapter 7 we shall investigate this behaviour of the MLE and its ‘unconstrained’ version for multi-qubit states near the boundary, in an extensive numerical simulation study.

The aim of the study is however *far more general-* to systematically compare and analyse the performance of several tomographic estimation methods across a range of different multi-qubit states, ranks, measurement designs, number of copies n , and the number of qubits N . We also consider the risk of the estimators in terms of several different loss functions, such as the square Frobenius norm, trace norm, square Bures and Hellinger distances. Apart from the several commonly used and studied estimators such as the MLE and the linear regression estimators, we also introduce and define estimators such as the *Generalised Least Squares* (GLS) and the *Generalised Positive Least Squares* (PGLS). We show that for sufficiently large n the ‘unconstrained’ MLE is equivalent to the GLS estimator, and that both the PGLS and the ML estimates can be obtained by projecting the GLS estimate onto the space of density matrices.

In addition to the results of the simulation study, we also introduce two web-based applications that are designed as tools for performing tomographic simulations online. These applications make available all of the estimators studied in chapter 7, and thus enable a user to perform their own simulations with arbitrary multi-qubit states. These applications also serve to complement the results presented in chapter 7.

Part II

Chapter 4

Statistically efficient tomography of low rank states with incomplete measurements

4.1 Introduction

Recent years have witnessed great experimental progress in the study and control of individual quantum systems [64, 116]. A common feature of many experiments is the use of *Quantum State Tomography* (QST) methods as a key tool for validating the results [62, 47]. The aim of QST is to statistically reconstruct an unknown state from the outcomes of repeated measurements performed on identical copies of the state. Among the proposed estimation methods we mention, e.g. variations of maximum likelihood [13, 19, 71, 101, 50], linear inversion [113], Bayesian inference [8, 6], estimation with incomplete measurements [51, 108, 109], and continuous variables tomography [87].

However, for composite systems such as trapped ions, full state tomography becomes challenging due to the exponential increase in dimension [90]. Therefore, there has been a significant interest in developing tomography methods that are efficient for certain *lower dimensional families of physically relevant states*. For instance, the estimation of low rank states has been considered in the context of compressed sensing (CS) [45, 54, 34, 33, 76] (chapter 3), model selection [61], and spectral thresholding [3, 25]. The estimation of the permutationally invariant part of the density matrix as an approximation to the true state is also relevant in certain physical models [110, 91, 103]. Similarly, the estimation of matrix product states [39] is particularly relevant for many-body systems, but also for estimating dynamical parameters of open systems [35, 60].

In this chapter we build on the fruitful CS idea that the sparsity of low rank states can be exploited in order to identify and estimate the state with a *reduced number of ‘measurements’*, in contrast to standard, informationally complete QST. Recall that a rank- r joint state of N qubits can be characterised

by $O(rd)$ parameters, where $d = 2^N$ is the dimension of the associated Hilbert space. In the original CS proposal it is shown that such a state can be recovered from the expectation values of $O(rd \log d)$ randomly chosen Pauli observables. More recent work concentrates on error bounds [34, 45] and confidence intervals [33] of CS estimators.

However as briefly explained in section 3.1.1, from a statistical and experimental viewpoint the estimation based on Pauli expectations does not make the most efficient use of the measurement data available in ion trap experiments. Indeed, the Pauli expectations can be seen as ‘coarse grained’ statistics of the ‘raw data’ which consists of counts for individual outcomes of a measurement in an orthonormal basis. This coarse graining leads to loss of information and a significant increase in estimation error, as shown in section 4.5.

In contrast, here we consider the statistical problem of estimating low rank states in the set-up of multiple ions tomography (MIT), where the input is the counts dataset provided by the experiment. The goal is to investigate the possibility of using a *reduced number of measurement settings* (Pauli bases), without a significant loss of statistical accuracy, in comparison to standard, full settings MIT. For this, we consider the behaviour of the *Mean Square Error* (MSE) with respect to the Frobenius distance between the true state and the estimator $\mathbb{E}\|\hat{\rho} - \rho\|_F^2$, in the limit of large number of measurement samples. According to asymptotic theory [118], in this regime the MSE of efficient estimators (e.g. maximum likelihood) $\hat{\rho}$ takes the following expression

$$\mathbb{E}\|\hat{\rho} - \rho\|_F^2 = \frac{1}{n} \text{Tr}(I(\rho|\mathcal{S})^{-1}G) + o(n^{-1}). \quad (4.1)$$

Above, $I(\rho|\mathcal{S})$ is the classical Fisher information associated with the chosen measurement design \mathcal{S} and a local parametrisation of rank- r states, and n is the total number of quantum samples available as a resource. G is the positive weight matrix associated with the quadratic approximation of the Frobenius distance in the local parameters.

In the following section we review the MIT set-up, and formulate the ‘reduced settings hypothesis’ in statistical terms. After this, we present the results of extensive numerical simulations testing this hypothesis, which are summarised in Figure 4.1. We find that the asymptotic MSE given by (4.1) is very robust with respect to a reduction the number of settings, with a random choice of settings making up the measurement design \mathcal{S} . For instance, 4 ion states of rank 3 can be estimated by using 20 settings (out of a total of 81)

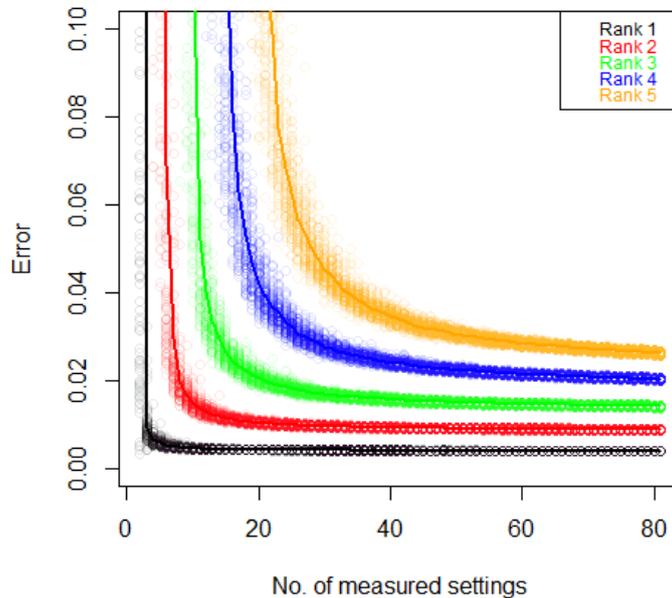


Figure 4.1 Asymptotic MSE for 4 ion states with ranks $r = 1, \dots, 5$, and total sample size $n = m \cdot 3^4 = 8100$. For each rank we chose 10 random states, and for each number of settings k (horizontal axis) we chose 10 random sets of settings \mathcal{S} . The MSE for each such combination is represented with a circle, and the lines are the average values.

with a negligible increase in estimation error. Also, to test the validity of the asymptotic theory for low rank states, we compared the theoretical prediction (4.1) with the actual MSE of the maximum likelihood estimator and found a very good agreement for $m = 100$ repetitions per setting, a typical value used in experiments [62].

To explain the observed robustness, we outline an argument based on a concentration inequality [2] for the Fisher information matrix of an experiment with randomly chosen Pauli settings. Transforming the argument into a mathematical proof requires control over certain spectral properties of the Fisher information matrix, and remains an open problem. However, by relaxing the Pauli measurement setup, and allowing for measurements with respect to *random bases*, we can prove that states of rank r can be estimated by using $O(r \log d)$ settings, with only a small increase in the MSE, relative to the setup in which a large number of settings is probed, cf. Theorem 2. For Pauli measurements we present numerical evidence on the lowest eigenvalue of the

Fisher information matrix, which strongly suggests that the MSE of random low rank states concentrates for a small number of measurement settings.

From a CS viewpoint, our question is closely related to the work [81, 73] inspired by the PhaseLift problem [32, 53] which considers the case where the incomplete ‘measurements’ are expectations of rank-one projections sampled randomly from a Gaussian distribution, or a projective t-design. In [80] the analysis is extended to the physically relevant case of random orthonormal basis measurements, and it is shown that rank- r states become identifiable with a large probability for only $O(r \log^3 d)$ ‘sufficiently random’ measurements. These results are in broad agreement with our findings, calling for a better understanding of the connections between the CS estimators and statistical approaches considered here.

4.2 Multiple Ions Tomography with Incomplete Measurements

In this chapter we shall consider the multiple ions tomographic (MIT) setup as in the ion-trap experiments [62]. In MIT the goal is to statistically reconstruct the joint state of N ions (modelled as two-level systems), from counts data generated by performing a large number of measurements on identically prepared systems. The unknown state ρ is a $d \times d$ density matrix (complex, positive trace-one matrix) where $d = 2^N$ is the dimension of the Hilbert space of N ions. The experimenter can measure an arbitrary Pauli observable σ_x, σ_y or σ_z of each ion, simultaneously on all N ions. Thus, each measurement setting is labelled by a sequence $\mathbf{s} = (s_1, \dots, s_N) \in \{x, y, z\}^N$ out of 3^N possible choices. The measurement produces an outcome $\mathbf{o} = (o_1, \dots, o_N) \in \mathcal{O}_N := \{+1, -1\}^N$, whose probability is

$$p_\rho(\mathbf{o}|\mathbf{s}) := \text{Tr}(\rho P_{\mathbf{o}}^{\mathbf{s}}) = \langle \lambda_{\mathbf{o}}^{\mathbf{s}} | \rho | \lambda_{\mathbf{o}}^{\mathbf{s}} \rangle, \quad (4.2)$$

where $P_{\mathbf{o}}^{\mathbf{s}}$ is the one dimensional projection $P_{\mathbf{o}}^{\mathbf{s}} = |\lambda_{o_1}^{s_1}\rangle\langle\lambda_{o_1}^{s_1}| \otimes \dots \otimes |\lambda_{o_N}^{s_N}\rangle\langle\lambda_{o_N}^{s_N}|$, and, $|\lambda_{\pm}^s\rangle$ is an eigenvector of the Pauli matrix σ_s , with a corresponding ± 1 eigenvalue.

The measurement procedure and statistical model can be summarised as follows. For each setting \mathbf{s} the experimenter performs m repeated measurements and collects the counts of different outcomes $N(\mathbf{o}|\mathbf{s})$, so that the total number of quantum samples used is $n := m \times 3^N$. The resulting dataset is a $2^N \times 3^N$

table whose columns are independent and contain all the counts in a given setting. The overall measurement is informationally complete, and the state can be estimated by using a number of methods proposed in the literature [101, 25, 3].

Now, there are several reasons to consider a set-up in which a *smaller* number of measurement settings are used for estimating the state; switching measurement settings may be more costly than repeating a measurement in the same setting, and smaller datasets may be easier to handle computationally. However, by removing even a single setting, the state becomes unidentifiable. This is because the corresponding tensor of Pauli operators is linearly independent from all the one dimensional projections of the remaining settings, and therefore its expectation value cannot be estimated. This can be remedied if some prior information about the state is available. The relevant example here is from compressed sensing [45, 54, 34, 33, 76] which shows that low rank states are uniquely determined by the Pauli expectations associated with a reduced number of settings. However, the existing compressed sensing literature does not address the statistical problem of estimating the state directly from the *raw* measurement data (i.e. the counts $N(\mathbf{o}|\mathbf{s})$), as it typically employs *coarse grained* statistics such as Pauli expectations. Our goal is to investigate the statistical efficiency of estimating low rank states with reduced measurement settings. We will consider an asymptotic scenario in which the number m of measurement repetitions per setting is large and the mean square error can be characterised in terms of the classical Fisher information, as discussed above. As we show below this regime is already attained for $m = 100$, which is of the order of repetitions cycles used in experiments [62].

As stated above, we assume that the prepared state ρ belongs to the space of rank r states $\mathbb{S}_r \subset M(\mathbb{C}^d)$, for a fixed rank $r < d$. Since the asymptotic mean square error depends only on the local properties of the statistical model, it suffices to consider a parametrisation $\boldsymbol{\theta} \mapsto \rho_{\boldsymbol{\theta}}$ of a neighbourhood of ρ in \mathbb{S}_r , which can be chosen as follows. In its own eigenbasis ρ is the diagonal matrix of eigenvalues $\text{Diag}(\lambda_1, \dots, \lambda_r, 0, \dots, 0)$, and any sufficiently close state is uniquely determined by its matrix elements in the first r rows (or columns). Intuitively this can be understood by noting that any rank- r state ρ' in the neighbourhood of ρ can be obtained by perturbing the eigenvalues and performing a small rotation of the eigenbasis; in the first order of approximation, these transformations

leave the $(d-r) \times (d-r)$ lower-right corner unchanged so

$$\rho' = \begin{pmatrix} \text{Diag}(\lambda_1, \dots, \lambda_r) & 0 \\ & 0 \end{pmatrix} + \begin{pmatrix} \Delta_{diag} & \Delta_{off} \\ \Delta_{off}^\dagger & O(\|\Delta\|^2) \end{pmatrix}. \quad (4.3)$$

We therefore choose the (local) parametrisation $\rho' = \rho_{\boldsymbol{\theta}}$ with

$$\begin{aligned} \boldsymbol{\theta} &:= (\boldsymbol{\theta}^{(d)}; \boldsymbol{\theta}^{(r)}; \boldsymbol{\theta}^{(i)}) \\ &= (\rho'_{2,2}, \dots, \rho'_{r,r}; \text{Re}\rho'_{1,2}, \dots, \text{Re}\rho'_{r,d}; \text{Im}\rho'_{1,2}, \dots, \text{Im}\rho'_{r,d}) \in \mathbb{R}^{2rd-r^2-1} \end{aligned} \quad (4.4)$$

where, in order to enforce a trace-one normalisation, we constrain the first diagonal matrix element to be $\rho'_{1,1} = 1 - \sum_{i=2}^d \rho'_{i,i}$. In this parametrisation we denote $\rho = \rho_{\boldsymbol{\theta}_0}$, with $\boldsymbol{\theta}_0 := (\lambda_2, \dots, \lambda_r; 0 \dots 0; 0 \dots 0)$. The Frobenius distance is locally quadratic in $\boldsymbol{\theta}$ so that

$$\|\rho_{\boldsymbol{\theta}_1} - \rho_{\boldsymbol{\theta}_2}\|_F^2 = (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2)^T G (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2) + o(\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|^2) \quad (4.5)$$

where

$$G_{a,b} = \text{Tr} \left[\frac{\partial \rho_{\boldsymbol{\theta}}}{\partial \theta_a} \cdot \frac{\partial \rho_{\boldsymbol{\theta}}}{\partial \theta_b} \right] \quad (4.6)$$

is a constant weight matrix whose expression can be found in the appendix to this chapter, below equation (4.19). After fixing the parametrisation, we now define the statistical model of multiple ions tomography with incomplete settings. Let $\mathcal{S} \subset \{x, y, z\}^N$ be a set of k randomly chosen settings, and consider the modified scenario in which ions prepared in the unknown state ρ are repeatedly measured $m = n/k$ times for all settings in \mathcal{S} , so that the overall number of quantum samples is always n . The classical Fisher information associated with a single chosen setting \mathbf{s} is defined as

$$I(\rho|\mathbf{s})_{a,b} := \sum_{\mathbf{o}} \frac{1}{p_{\rho}(\mathbf{o}|\mathbf{s})} \frac{\partial p_{\rho}(\mathbf{o}|\mathbf{s})}{\partial \theta_a} \cdot \frac{\partial p_{\rho}(\mathbf{o}|\mathbf{s})}{\partial \theta_b}. \quad (4.7)$$

For a set \mathcal{S} of k settings the Fisher information matrix associated with a single measurement sample from each setting $\mathbf{s} \in \mathcal{S}$ is given by the sum of the individual Fisher matrices $I(\rho|\mathbf{s})$, and for later purposes we will denote the average $I(\rho|\mathcal{S}) = \frac{1}{k} \sum_{\mathbf{s} \in \mathcal{S}} I(\rho|\mathbf{s})$. The individual matrices can be computed by using definition (4.7) together with equation (4.2) and the parametrisation (4.4).

Since the outcomes from m repeated measurements in a setting \mathbf{s} are i.i.d, when the number of repetitions m is sufficiently large, efficient estimators of $\boldsymbol{\theta}$ (and hence of ρ) from these outcomes have an asymptotically Gaussian distribution [118]

$$\sqrt{m}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \approx \mathcal{N}(0, I(\rho|\mathcal{S})^{-1}) \quad (4.8)$$

where the covariance matrix $I(\rho|\mathcal{S})^{-1}$ is the Fisher information associated with a single measurement sample of the set \mathcal{S} . From this behaviour and the local expansion of the Frobenius distance, we see that for (reasonably) large m , the mean square error of an efficient estimator (e.g. maximum likelihood) scales as

$$\text{MSE} := \mathbb{E}(\|\hat{\rho} - \rho\|_F^2) \approx \frac{1}{n} \text{Tr}(I(\rho|\mathcal{S})^{-1}G). \quad (4.9)$$

Compare this equation with (4.1) in the introduction. The trace expression is a measure of the sensitivity of the chosen set of settings \mathcal{S} at ρ . Since the settings are chosen randomly we need to study the fluctuations of $\text{Tr}(I(\rho|\mathcal{S})^{-1}G)$. In the next section we present extensive simulation results which essentially show that one can significantly reduce the number of settings without affecting the MSE.

4.3 Numerical Simulations

In Figure 4.1 we plot the values of the asymptotic MSE $\text{Tr}(I(\rho|\mathcal{S})^{-1}G)/n$ for various ranks, choices of 4 ion states, and choices of measurement designs \mathcal{S} (sets of settings). For each rank $r = 1, \dots, 5$ we generated 10 states by using the Cholesky decomposition $\rho = T^*T$, cf. [18]. For each state, the MSE values are calculated over a range of measurements with reduced settings, starting from the ‘full’ measurement with 3^N settings, as follows. For a given number of reduced measurements k , we generated 10 independent sets \mathcal{S} of randomly chosen settings. For each pair (ρ, \mathcal{S}) we evaluated the Fisher information $I(\rho|\mathcal{S})$ and the weight matrix G in the parametrisation described above. In these simulations, the total number of copies of the state is kept constant as a resource. Therefore, a smaller number of measurement settings k leads to a larger number of repetitions $m = n/k$ per setting. The simulations show that asymptotic risk for low rank states demonstrates only a gradual increase even over a significant reduction in the number of settings measured. For example, for states of rank 3, one can reduce the number of settings from 81 to 20 with a negligible increase in the MSE. Moreover, for a given k , the fluctuations

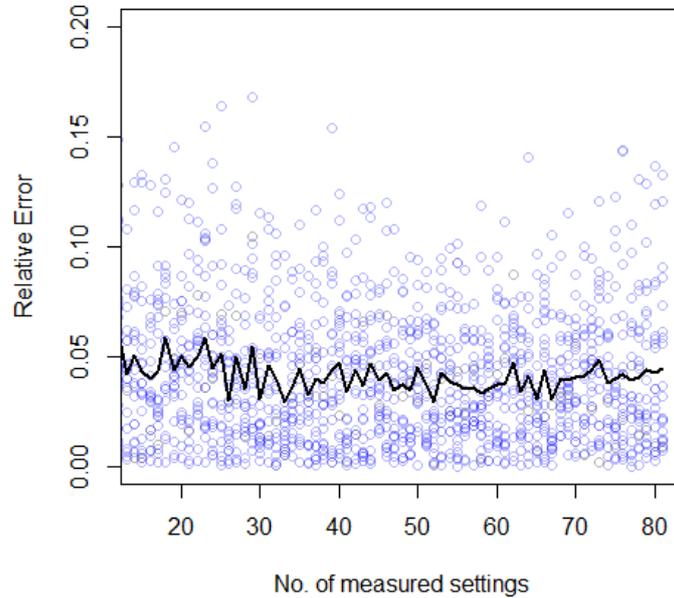


Figure 4.2 The error of the (estimated) MSE of the Maximum Likelihood estimate relative to the asymptotic Fisher MSE (see 4.10), for a 4 ion pure state. The circles plot the RE for 20 random sets \mathcal{S} chosen for each k number of settings, and the line plots the average relative error. The total sample size is $n = 8100$.

of the MSE over choices of states and settings are rather small, showing the robustness of the procedure.

In the previous section we argued that the theoretical value (4.9) is close to the actual error of an efficient estimator, when the number of samples is reasonably large. To verify this we have computed the *Maximum Likelihood* (ML) estimator and studied its performance in this reduced measurement settings MIT setup. The ML estimator implemented is a modified form of the iterative $R\rho R$ method in [101] - for the estimates generated at each iteration of the algorithm, only the r largest eigenvalues are retained. This modification ensures that the ML estimator has knowledge about the rank of the density matrix. The results of the comparison with the Fisher prediction (4.9) are shown in Figure 4.2, and show a very good agreement between the two. For a given random set \mathcal{S} of k settings, the MSE of the MLE $\mathbb{E}\|\hat{\rho}_{ML} - \rho\|_F^2$ is estimated by averaging the square error over 30 ML estimates. The relative

error

$$\left| 1 - \frac{n \cdot \mathbb{E} \|\hat{\rho}_{ML} - \rho\|_2^2}{\text{Tr}(I(\rho|\mathcal{S})^{-1}G)} \right| \quad (4.10)$$

is then plotted for each choice of \mathcal{S} as a single circle. On average, the relative error is of order of 5%. In conclusion, the simulations indicate that low rank states can be estimated with a significantly smaller number of measurement settings than the total of 3^N currently used in experiments, with a negligible loss of statistical accuracy.

4.4 A Concentration Bound for the MSE

Why is the MSE robust with respect to the reduction of the number of measured settings? In this section we provide an intuitive explanation based on a concentration bound for the asymptotic MSE, i.e. the random function $\mathcal{S} \mapsto \text{Tr}(I(\rho|\mathcal{S})^{-1}G)$. Analysing the observed MSE concentration for MIT with Pauli measurements is difficult due to the special, discrete set of bases which contribute to the average. Much like the problem of proving the RIP property in compressed sensing [85, 45], it is easier to analyse a more random set-up, namely one where the measurement bases making up the design \mathcal{S} are drawn randomly from the *uniform measure* over orthonormal bases (ONB). We therefore begin by considering this general setup of random measurements and return to the Pauli measurements later in the section.

Physically, this random setup could be implemented by first rotating the state ρ by a random unitary U , after which each atom is measured in the σ_z eigenbasis. We therefore let $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_k\}$ be the altered design with randomly, uniformly distributed measurement bases. Since the settings in \mathcal{S} are independent, the Fisher information matrices $I(\rho|\mathbf{s})$ are independent, and for k large enough the average information per setting approaches the mean information over all random settings

$$I(\rho|\mathcal{S}) = \frac{1}{k} \sum_{\mathbf{s} \in \mathcal{S}} I(\rho|\mathbf{s}) \approx \bar{I} := \mathbb{E}_{\mathbf{s}} [I(\rho|\mathbf{s})]. \quad (4.11)$$

Since we are interested in the behaviour of the MSE for the randomly chosen designs \mathcal{S} , we look at the relative error

$$RE(\rho|\mathcal{S}) := \text{Tr}(I(\rho|\mathcal{S})^{-1}G) / \text{Tr}(\bar{I}^{-1}G). \quad (4.12)$$

and would like to determine the number of settings k required for the MSE to be concentrated close the optimal value of $\text{Tr}(\bar{I}^{-1}G)$.

To investigate this MSE concentration for states of rank r in this setup, we focus our attention on states with equal eigenvalues, i.e. $\rho_0 := \text{Diag}\left(\frac{1}{r}, \dots, \frac{1}{r}, 0, \dots, 0\right)$, with respect to its eigenbasis; due to the unitary symmetry of the random settings design, the eigenbasis can be chosen to be the standard basis. The reason for choosing this particular spectrum is that such states represent the ‘least sparse’ state of rank r . Indeed, rank- r states which have some eigenvalues close to zero can be approximated by states of lower ranks, and we expect that they require even smaller number of measurement settings. The following Theorem shows that in order to keep the relative error (4.12) close to 1 it suffices to take a number of random settings k which scales as $O(r \log(2rd))$ with respect to the rank and Hilbert space dimension. Taking into account that one setting provides d probabilities, the total number of expectations is of the order $O(rd \log(2rd))$ which roughly agrees with the number of Pauli expectations required in compressed sensing. We will come back to this comparison in the following section.

Theorem 2. *Let $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_k\}$ be a design with randomly, uniformly distributed measurement bases. Let $I_{\mathcal{S}} := I(\rho_0|\mathcal{S})$ be the associated Fisher information, and let \bar{I} be the mean Fisher information over all possible bases, both calculated at ρ_0 (as defined above). For a sufficiently small $\epsilon \geq 0$, the following inequality holds*

$$(1 - \epsilon)\text{Tr}[\bar{I}^{-1}G] \leq \text{Tr}[I_{\mathcal{S}}^{-1}G] \leq (1 + \epsilon)\text{Tr}[\bar{I}^{-1}G] \quad (4.13)$$

with probability $1 - \delta$, provided that the number of measurements performed is $k = C(r + 1) \log(2D/\delta)$, with $D = 2rd - r^2 - 1$ the dimension of the space of rank- r states.

The proof of this theorem is detailed in the appendix, and uses a matrix Chernoff bound [2], to bound the deviation of $G^{-1/2}I_{\mathcal{S}}G^{-1/2}$ from the mean $G^{-1/2}\bar{I}G^{-1/2}$. This is then recast in terms of a bound on the MSE as in the theorem above. The two bounds show that with probability $1 - \delta$, the relative error $RE(\rho_0|\mathcal{S})$ is in the interval $[1 - \epsilon, 1 + \epsilon]$, so using design \mathcal{S} induces at most an ϵ increase of MSE. Similar results can be derived along the lines of the proof for states with arbitrary spectrum, and we shall consider this in detail in chapter 5.

Figure 4.3 illustrates this concentration in two ways; by plotting the relative error $RE(\rho_0|\mathcal{S})$ and by plotting the eigenvalues of $G^{-1/2}I_{\mathcal{S}}G^{-1/2}$, for various values of k . The concentration in the spectrum of the eigenvalues demonstrates the rate at which $I_{\mathcal{S}}$ approximates the mean Fisher information \bar{I} . We see that for pure states all eigenvalues concentrate around 1, this is because $G^{-1/2}\bar{I}G^{-1/2}$ is an identity matrix for pure states, while for ranks 2 and 3 this matrix is no longer identity and has eigenvalues that are either 1 or $r/(r+1)$. We see in the plots for these ranks that the lower band in the eigenvalues spectrum approaches the minimum eigenvalue of $r/(r+1)$, while the remaining eigenvalues concentrate around 1. The explicit form of the $G^{-1/2}\bar{I}G^{-1/2}$ matrix is detailed in the appendix.

The above theorem guarantees that for a 4 ion pure state, the MSE is within 5% of the optimal, with a probability of failure $\delta = 0.1$, provided that we measure $k \approx 7100$ settings. However, the bottom-right plot in Figure 4.3 shows that the MSE concentrates much earlier, well within $k = 100$ settings. This indicates that studying the concentration of $G^{-1/2}I_{\mathcal{S}}G^{-1/2}$ to bound the MSE provides a highly pessimistic estimate for k . Note however, that although the value $k \approx 7100$ is much larger than the full set of measurements for a 4 ion state in the MIT setup, the theorem demonstrates a significant reduction in the number of settings needed when we consider larger states of $N \geq 9$ ions.

In Figure 4.4, we plot the relative error $RE(\rho_0|\mathcal{S})$ for random pure states of differing dimensions. Interestingly, we see that the error does not seem to depend on the dimension of the state. Whereas from the above theorem, we would expect a dependence because of the $\log(2D/\delta)$ term that appears in the concentration. An explanation for this has been suggested in [10], where the authors conjecture that measuring only a few random bases correspond to *strictly complete POVMs* for low rank quantum states. Meaning that states of a given low rank can be recovered by measuring a small number of random bases, independent of dimension. This special feature of the random measurements suggests that a better scaling in the number of measurement settings is possible in the concentration bound for this setup.

4.4.1 Pauli settings

We now return to the more physical set-up in which the settings are chosen from the set $\{x, y, z\}^N$ of Pauli measurements. Figure 4.5 plots the error $RE(\rho|\mathcal{S})$ of the MSEs for the reduced settings, relative to the MSE of the average

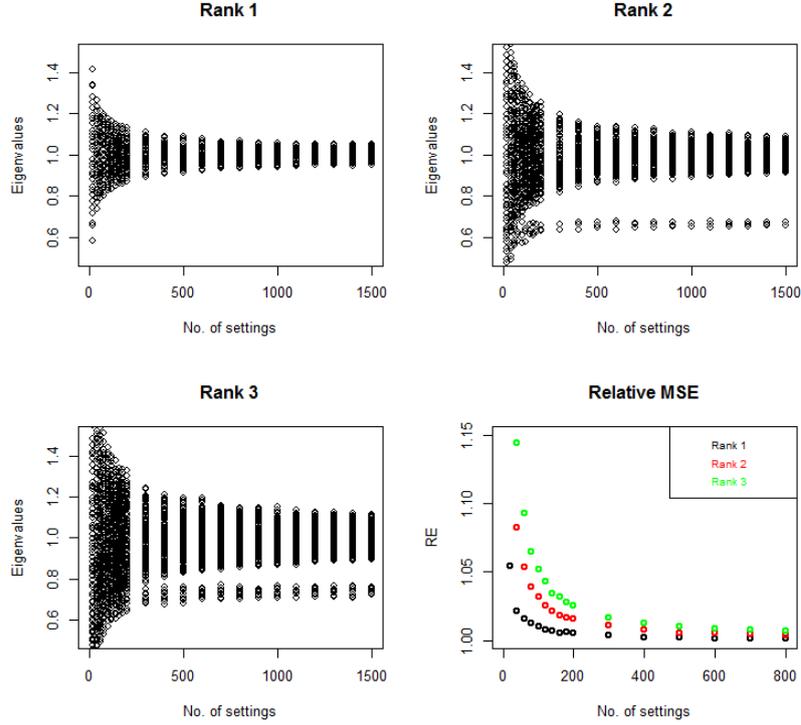


Figure 4.3 Plots of the eigenvalues of $G^{-1/2}I_{\mathcal{S}}G^{-1/2}$ for k random settings. For a given rank r , we chose a random 4 ion state ρ_0 with r equal eigenvalues. We observe a concentration of the eigenvalues as $I_{\mathcal{S}}$ approximates the mean information \bar{I} (See Thm.2). The bottom right plot graphs the relative error $RE(\rho_0|\mathcal{S})$ for the different ranks.

information for all 3^N settings $\bar{I} = (3^N)^{-1} \sum_{\mathbf{s} \in \{x,y,z\}^N} I(\rho_0|\mathbf{s})$. The numerical simulations show that even for $k = 20$ settings, the average MSE is only 5% higher than the MSE of the full settings experiment, while when the variance is taken into account, most MSEs are less than 10% higher. We note that in the simulations, the different settings making up the measurement design \mathcal{S} are chosen without replacement, while an application of the concentration bound in the theorem would use a slightly altered setup in which the different settings are chosen independently and with equal probabilities (drawing with replacement). For a discussion on the relation between the two set-ups we refer to [55].

The key step in establishing a concentration bound as in Theorem 2 is to control the ratio

$$\frac{\mu_{\max}}{\mu_{\min}} := \frac{\max_{\mathbf{s}} \lambda_{\max}(G^{-1/2}I(\rho_0|\mathbf{s})G^{-1/2})}{\lambda_{\min}(G^{-1/2}\bar{I}G^{-1/2})} \quad (4.14)$$

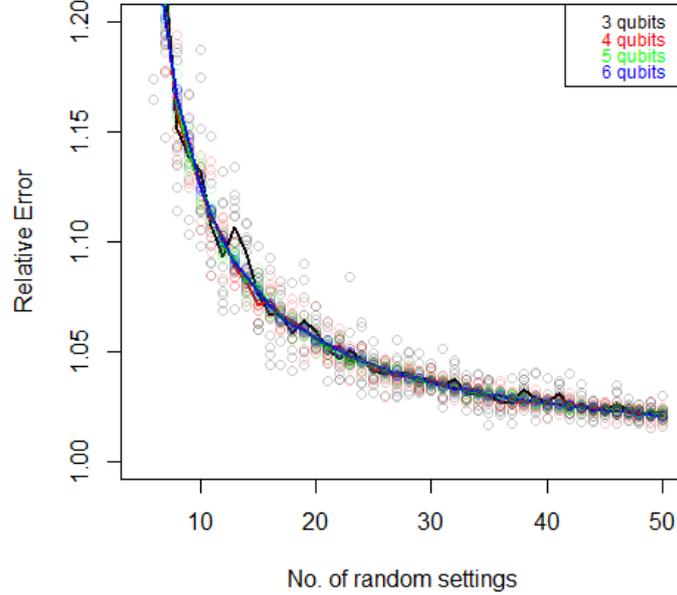


Figure 4.4 Relative error $RE(\rho_0|\mathcal{S})$ for the random settings, and pure states of 3-6 qubits. The coloured thick lines plot the mean relative error over 10 different random states of a given dimension. The light coloured circles plot the errors for a particular state and a given k number of random settings.

between the largest maximum eigenvalue of $G^{-1/2}I(\rho_0|\mathbf{s})G^{-1/2}$ over all measurements and the minimum eigenvalue of $G^{-1/2}\bar{I}G^{-1/2}$. In the case of the uniformly distributed settings, \bar{I} can be computed explicitly by using analytic expressions for moments of random unitaries [37], which gives $\mu_{\min} = \frac{r}{r+1}$ for $r > 1$, and $\mu_{\min} = 1$ for pure states, while μ_{\max} can be upper bounded by using the inequality between the quantum and classical Fisher informations [24], as $\mu_{\max} \leq 2r$ for $r > 1$ and $\mu_{\max} \leq 4$ for $r = 1$. Together these give a $\frac{\mu_{\max}}{\mu_{\min}} = 2(r+1)$ which determines the number of measurement settings k in Theorem 2.

For the Pauli measurements set-up, the same upper bound holds for the maximum eigenvalue, but at the moment we do not have a similar lower bound for $\lambda_{\min}(G^{-1/2}\bar{I}G^{-1/2})$, where \bar{I} is the average Fisher information over *Pauli settings*. However, there is strong numerical evidence that the smallest eigenvalue of $\lambda_{\min}(G^{-1/2}\bar{I}G^{-1/2})$ for random states remains well bounded away from zero. Figure 4.6 plots the minimum eigenvalues for 100 such states of 4 to 8 ions, over three different ranks. The boxes in the plot mark the interquartile range, with whiskers extending to extreme points which are no more than

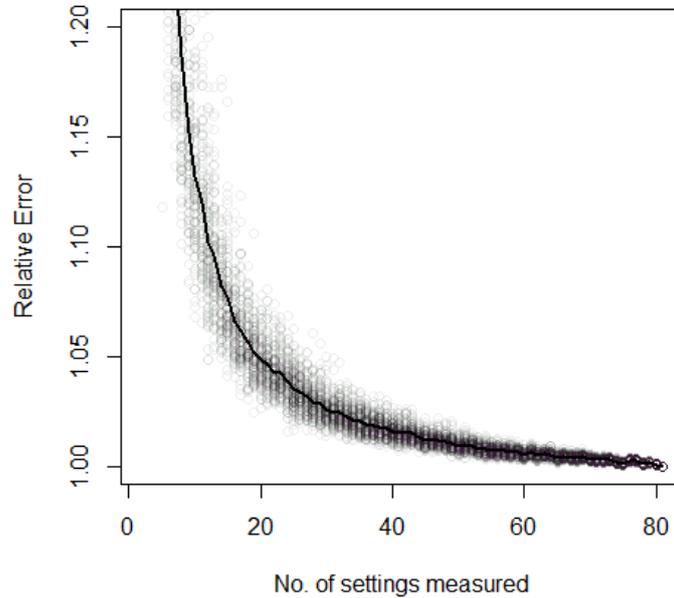


Figure 4.5 Relative Error $RE(\rho|\mathcal{S})$ for the Pauli settings, for a randomly chosen 4 ion pure state. The thick line plots the mean relative error over 100 different choices of k settings, with the light grey circles plotting the relative errors for different choices of the settings.

1.5 times the interquartile range. We notice that the minimum eigenvalue for each rank is well concentrated away from zero and for ranks $r > 1$ clearly demonstrates an increase with the dimension of the space. While the full dependence of μ_{\min} on r and d is unclear, the simulations strongly suggest that for any fixed rank, μ_{\min} is larger than a fixed constant with high probability for random states of rank r , of arbitrarily many ions. If this was true, it would imply that random states of *fixed* rank r can be estimated efficiently with $O(\log d)$ settings.

For now, as a step in the direction of proving and demonstrating the concentration as in Theorem 2 for reduced settings, we will prove a weaker result based on a rough lower bound for μ_{\min} . From Theorem 2 in [25] we have that for full 3^N settings, the MSE of an optimal estimator $\hat{\rho}$ is upper bounded as

$$\mathbb{E}\|\hat{\rho} - \rho\|_F^2 \leq C' \frac{rd}{n} \log(2d), \quad (4.15)$$

with $C' > 0$ being an absolute constant. Asymptotically, the MSE is lower bounded by $1/n \cdot \text{Tr}(\bar{I}^{-1}G)$ which implies $1/\mu_{\min} \leq C' rd \log(2d)$. This gives us

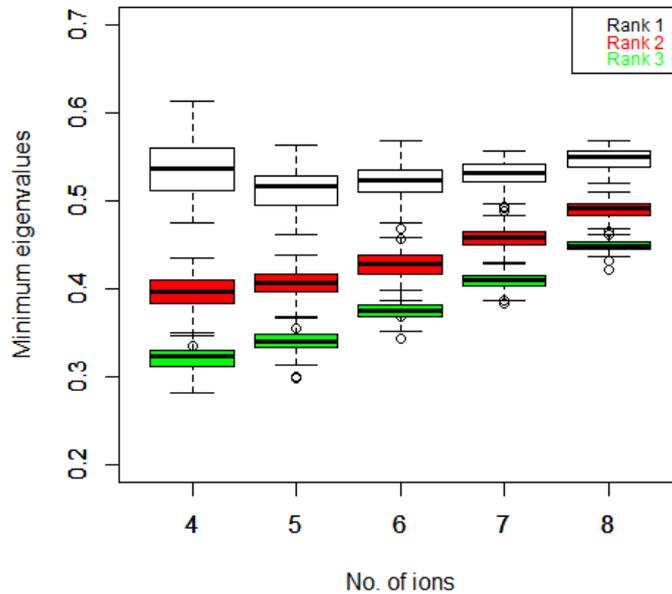


Figure 4.6 Box-plots of the minimum eigenvalues of $G^{-1/2} \bar{I} G^{-1/2}$ for the Pauli settings. For a given rank and ion number, we chose randomly 100 different states ρ_0 with r equal eigenvalues.

a rough lower bound on the minimum eigenvalue. Plugging this value into the concentration bound of Theorem 2 gives us that the minimum number of settings k scales as $O(r^2 d \log^2(2D))$, which despite being far from optimal, demonstrates a better scaling than the 3^N of the ‘full settings’ setup.

4.5 Coarse vs Fine Grained Models

As mentioned in the introduction, a similar reduction in the number of ‘measurements’ has been found in compressed sensing (CS) estimators [45, 54, 34, 33, 76], which use $O(rd \log d)$ expectations of Pauli operators to recover the unknown state. CS techniques provide computationally efficient estimators whose estimation errors scale optimally with the number of parameters and with the errors in the estimation of the Pauli expectations. However, from the statistical viewpoint the Pauli expectations are not the most efficient starting point in estimation, as they are ‘coarse grained’ statistics of the ‘raw’, or ‘fine grained’ measurement data given by the counts $N(\mathbf{o}|\mathbf{s})$.

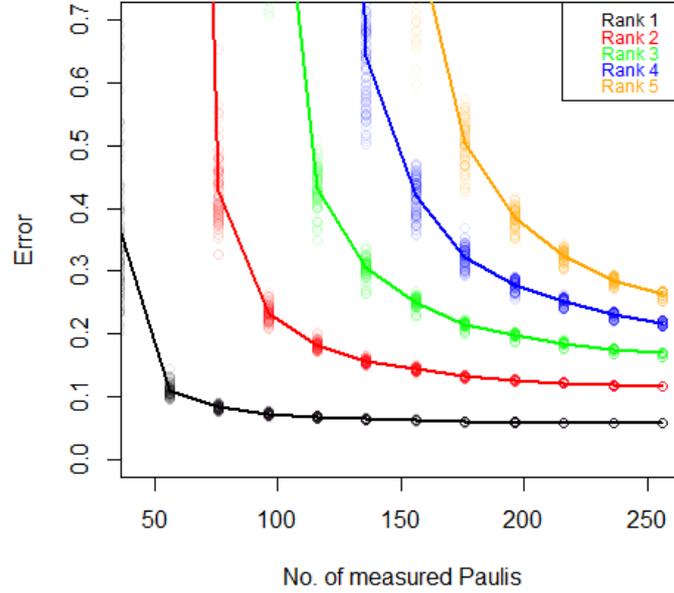


Figure 4.7 The plot of $\text{Tr}(I^{-1}G)$ for the ‘coarse grained’ model for multiple ranks, with 4 ions and 256 total measured Paulis \mathbf{b} . The total number of copies of the states is 8100. The experiment is repeated over 10 different states, and 10 random choices of Pauli measurements for each state.

A single measurement in the ‘coarse grained’ model is defined by a Pauli observable $\sigma_{\mathbf{b}} := (\sigma_{b_1} \otimes \dots \otimes \sigma_{b_n})$ with $b_i \in \{0, x, y, z\}$, where σ_0 is the identity matrix. To compute its expectation one needs to measure $\sigma_{\mathbf{b}}$ to obtain a binary outcome $\{\pm 1\}$ and average over the results. The outcomes probabilities are $p_{\rho}(\pm 1|\mathbf{b}) = \text{Tr}(\rho P_{\pm}^{\mathbf{b}})$ where $P_{\pm}^{\mathbf{b}}$ are the two spectral projections of $\sigma_{\mathbf{b}}$, and the Fisher information of this model can be computed in much the same way as that of the Pauli bases measurements. In Figure 4.7 the asymptotic MSE $\text{Tr}(I(\rho|\mathcal{B})^{-1}G)/n$ is plotted for different sets of randomly chosen Pauli observables $\mathcal{B} := \{\mathbf{b}_1, \dots, \mathbf{b}_k\}$. On comparison with Figure 4.1 we see that the minimum number of measurements that need to be measured in the Pauli bases model is much smaller. Additionally, the risk for a full set of measurements is an order of magnitude larger in the ‘coarse grained’ model. This increase in the asymptotic risk has also been pointed out in [61].

The discrepancy can be explained by noting that the measurement of $\sigma_{\mathbf{b}}$ is a coarse graining of a finer, ONB measurement of a setting such that $s_i = b_i$ whenever $b_i \neq 0$. Indeed, using the spectral decomposition of $\sigma_{\mathbf{b}}$, we can

compute its expectation as

$$\langle \sigma_{\mathbf{b}} \rangle_{\rho} = \text{Tr}(\rho \sigma_{\mathbf{b}}) = \sum_{\mathbf{o}} \left(\prod_{i: b_i \neq 0} o_i \right) p_{\rho}(\mathbf{o}|\mathbf{s}). \quad (4.16)$$

By replacing the probabilities $p_{\rho}(\mathbf{o}|\mathbf{s})$ in the above formula by the empirical frequencies $N(\mathbf{o}|\mathbf{s})/m$ we obtain the estimate of the Pauli expectations. However, by constructing this statistic we lose a large amount of information contained in the frequencies, which explains the increase in the MSE.

4.6 Conclusions

In this chapter we investigated the statistical performance of reduced settings measurements in ion tomography. We did not focus on a particular estimation method but rather on how the accuracy of efficient estimators (which achieve the asymptotic scaling (4.1) of the MSE) depends on the state and the measurement design. We found that for low rank states, the experimenter can measure a small proportion of randomly selected settings without a significant increase in the MSE. Furthermore we presented a possible line of argument for a mathematical proof based on concentration inequality for the Fisher information. In the case of measurements with respect to random bases we showed that certain states of rank r can be estimated with $O(r \log d)$ settings with an ϵ increase in MSE compared with designs with a large number of settings. It remains an open question whether the same scaling of the size of the measurement design holds for the Pauli measurements, but we presented strong numerical evidence that for random states the Fisher information may satisfy the required spectral properties.

4.7 Appendix

4.8 Proof of Theorem 2

As briefly mentioned in the main text of the chapter, the proof of the theorem utilises the following matrix Chernoff bound [2], where the random matrices X_i are given by $G^{-1/2}I(\rho_0|\mathbf{s}_i)G^{-1/2}$, with \mathbf{s}_i random bases.

Theorem 3. (*Matrix Chernoff*) Consider a finite sequence X_1, \dots, X_k of independent, random, positive matrices with dimension D , such that $\lambda_{\max}(X) \leq R$. For $\mathbb{E}X = M \geq \mu \mathbf{1}$ and $0 \leq \epsilon \leq \frac{1}{2}$,

$$\mathbb{P} \left\{ \frac{1}{k} \sum_{i=1}^k X_i \notin \left[(1-\epsilon)M, (1+\epsilon)M \right] \right\} \leq 2D \cdot \exp \left(-k \cdot \frac{\epsilon^2 \mu}{2R \cdot \log 2} \right) \quad (4.17)$$

We note that $G^{-1/2}I_{\mathcal{S}}G^{-1/2}$ is a sum of k independent, random, positive matrices. In order to apply the above bound, we need to upper bound the largest eigenvalue of $G^{-1/2}I(\rho_0|\mathbf{s})G^{-1/2}$ over all measurements. We also need to lower bound the smallest eigenvalue of the expected Fisher information $G^{-1/2}\bar{I}(\rho_0)G^{-1/2}$. We will first derive these bounds and then derive the result by applying the Chernoff bound.

As in the text, we work with the local parametrisation

$$\boldsymbol{\theta} = (\boldsymbol{\theta}^{(d)}, \boldsymbol{\theta}^{(r)}, \boldsymbol{\theta}^{(i)}) = (\rho_{2,2}, \dots, \rho_{r,r}; \text{Re}\rho_{1,2}, \dots, \text{Re}\rho_{r,d}; \text{Im}\rho_{1,2}, \dots, \text{Im}\rho_{r,d})$$

where $\rho_{1,1}$ is constrained to enforce the trace-one normalisation. The Fisher information therefore, has the following block structure

$$I(\rho) = \begin{pmatrix} I^{dd}(\rho) & I^{dr}(\rho) & I^{di}(\rho) \\ I^{rd}(\rho) & I^{rr}(\rho) & I^{ri}(\rho) \\ I^{id}(\rho) & I^{ir}(\rho) & I^{ii}(\rho) \end{pmatrix} \quad (4.18)$$

with the superscripts identifying the parameters considered; diagonal, real and imaginary. The weight matrix G also has the same block structure with elements

$$G_{a,b} = \text{Tr} \left[\frac{\partial \rho_{\boldsymbol{\theta}}}{\partial \theta_a} \cdot \frac{\partial \rho_{\boldsymbol{\theta}}}{\partial \theta_b} \right] \quad (4.19)$$

In the parametrisation described above, the weight matrix G has the following block diagonal form:

1. The *diagonal-diagonal* block:

$$(a) \quad G_{a,b}^{dd} = 1 + \delta_{a,b}$$

2. The *real-real* and *imaginary-imaginary* block:

$$(a) \quad G_{a,b}^{rr/ii} = 2 \cdot \delta_{a,b}$$

with the other blocks being zero. We note that both the Fisher, and the weight matrix are of dimension $D := 2rd - r^2 - 1$.

Bound on the largest eigenvalue—We use the inequality $I(\rho_0|\mathbf{s}) \leq F$ between the classical and quantum Fisher informations to bound the largest eigenvalue of $G^{-1/2}I(\rho_0|\mathbf{s})G^{-1/2}$ over all measurements by the largest eigenvalue of $G^{-1/2}F(\rho_0)G^{-1/2}$. The derivation of the quantum Fisher matrix presented here follows [72]. We calculate the quantum Fisher information in the local parametrisation described above, and evaluate it at the diagonal state ρ_0 .

We begin by considering a state ρ_θ locally around some arbitrary rank- r state ρ' , and write the spectral decomposition as:

$$\rho_\theta = \sum_{i=1}^r p_i |\psi_i\rangle \langle \psi_i| \quad (4.20)$$

The quantum Fisher information matrix is defined as:

$$F_{a,b} = \text{Tr} \left[\rho_\theta (L_\theta^a \circ L_\theta^b) \right] = \frac{1}{2} \text{Tr} \left[\rho_\theta \left(L_\theta^a L_\theta^b + L_\theta^b L_\theta^a \right) \right] \quad (4.21)$$

where the symmetric logarithmic derivatives are defined through the equation:

$$\partial_{\theta_a} \rho_\theta = L_\theta^a \circ \rho_\theta = \frac{1}{2} (L_\theta^a \rho_\theta + \rho_\theta L_\theta^a) \quad (4.22)$$

We determine the elements of this matrix in the ONB formed by the eigenbasis set $\{|\psi_i\rangle\}$

$$\begin{aligned} \langle \psi_i | \partial_{\theta_a} \rho_\theta | \psi_j \rangle &= \frac{1}{2} \langle \psi_i | L_\theta^a \rho_\theta | \psi_j \rangle + \frac{1}{2} \langle \psi_i | \rho_\theta L_\theta^a | \psi_j \rangle \\ &= \frac{1}{2} (p_j + p_i) \langle \psi_i | L_\theta^a | \psi_j \rangle \end{aligned} \quad (4.23)$$

As pointed out in [72], L_{θ}^a (and L_{θ}^b) is in principle supported on the full space, but its entries for $i, j > r$ are arbitrary. However, the Fisher information does not use values for which $i, j > r$. This can be seen by expanding (4.21) in the following way

$$\begin{aligned} F_{a,b} &= \frac{1}{2} \text{Tr} \left[\sum_i p_i |\psi_i\rangle \langle \psi_i| (L_{\theta}^a L_{\theta}^b) + \sum_i p_i |\psi_i\rangle \langle \psi_i| (L_{\theta}^b L_{\theta}^a) \right] \\ &= \frac{1}{2} \sum_i^r \sum_j^d p_i (L_{\theta;i,j}^a L_{\theta;j,i}^b + L_{\theta;i,j}^b L_{\theta;j,i}^a) \end{aligned} \quad (4.24)$$

Since the index $i \leq r$, (4.23) can be inverted inside the expansion of the Fisher information as

$$L_{\theta;i,j}^a = \frac{2(\partial_{\theta_a} \rho_{\theta})_{i,j}}{p_i + p_j} \quad (4.25)$$

The quantum Fisher matrix therefore becomes

$$F_{a,b} = \sum_i^r \sum_j^d \frac{4p_i}{(p_i + p_j)^2} \text{Re} \left[(\partial_{\theta_a} \rho_{\theta})_{i,j} (\partial_{\theta_b} \rho_{\theta})_{j,i} \right] \quad (4.26)$$

where we used the fact that $\partial_{\theta_a/b} \rho_{\theta}$ is self-adjoint. Since ρ_{θ} is parameterised by its matrix elements in the eigenbasis $\{|\lambda_i\rangle\}$ of the state ρ' , we can use this to write the partial derivatives out explicitly. Using the notation that r_a, c_a represents the row and column indices for the parameter θ_a , the quantum Fisher matrix now becomes:

$$\begin{aligned} F_{a,b}^{d,d} &= \sum_i^r \sum_j^d \frac{4p_i}{(p_i + p_j)^2} \text{Re} \left[\langle \psi_i | \left(|\lambda_{r_a}\rangle \langle \lambda_{r_a}| - |\lambda_1\rangle \langle \lambda_1| \right) | \psi_j \rangle \right. \\ &\quad \left. \langle \psi_j | \left(|\lambda_{r_b}\rangle \langle \lambda_{r_b}| - |\lambda_1\rangle \langle \lambda_1| \right) | \psi_i \rangle \right] \end{aligned} \quad (4.27)$$

for the **diagonal-diagonal** block, and for the rest

$$\begin{aligned} F_{a,b} &= \sum_i^r \sum_j^d \frac{4p_i}{(p_i + p_j)^2} \text{Re} \left[\langle \psi_i | \left(|\lambda_{r_a}\rangle \langle \lambda_{c_a}| + |\lambda_{c_a}\rangle \langle \lambda_{r_a}| \right) | \psi_j \rangle \right. \\ &\quad \left. \langle \psi_j | \left(|\lambda_{r_b}\rangle \langle \lambda_{c_b}| + |\lambda_{c_b}\rangle \langle \lambda_{r_b}| \right) | \psi_i \rangle \right] \end{aligned} \quad (4.28)$$

We now evaluate these last two equations at $\boldsymbol{\theta} = \boldsymbol{\theta}_0$, for our special state that is diagonal with entries given by $1/r$. At this state $|\psi_i\rangle = |\lambda_i\rangle$. The *diagonal-diagonal* block of the Fisher matrix has elements:

$$F_{a,b}^{d,d}\Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} = r(1 + \delta_{r_a,r_b}), \quad (4.29)$$

while the *real-real* and *imaginary-imaginary* blocks are diagonal with elements

$$F_{a,b}\Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} = \frac{4}{p_{r_a} + p_{c_a}} (\delta_{r_a,r_b} \cdot \delta_{c_a,c_b}) \quad (4.30)$$

It is easy to see that the *real-diagonal*, *identity-diagonal* blocks are all zero. The *real-imaginary* blocks are zero since we consider only $Re\left[(\partial_{\theta_a}\rho\boldsymbol{\theta})_{i,m}(\partial_{\theta_b}\rho\boldsymbol{\theta})_{m,i}\right]$. Therefore, the elements of the quantum Fisher matrix are:

1. For the **Diagonal-Diagonal** block with $r > 1$,

- (a) $F_{a,a}^{dd}\Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} = 2r$ when $r_a \leq r$

- (b) $F_{a,b}^{dd}\Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} = r$ when $r_a, r_b \leq r$, and $a \neq b$

2. For the **Real-Real** and **Imaginary-Imaginary** blocks:

- (a) $F_{a,a}^{rr/ii}\Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} = 2r$ when $r_a < c_a \leq r$

- (b) $F_{a,a}^{rr/ii}\Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} = 4r$ when $r_a \leq r, c_a > r$

On comparing this with the weight matrix G , we notice that both G and F have the same block diagonal structure, with the off-diagonal blocks being zero. So we can write

$$G^{-1/2}FG^{-1/2} = G^{dd-1/2}F^{dd}G^{dd-1/2} \oplus G^{rr-1/2}F^{rr}G^{rr-1/2} \oplus G^{ii-1/2}F^{ii}G^{ii-1/2}$$

We notice that $F^{dd} = r \cdot G^{dd}$, which gives us

$$G^{-1/2}FG^{-1/2} = r\mathbf{1}_{(r-1)} \oplus \frac{1}{2}F^{rr} \oplus \frac{1}{2}F^{ii} \quad (4.31)$$

The maximum eigenvalue of this matrix comes from the diagonal block matrices $F^{rr/ii}/2$, and is $2r$ for $r > 1$, and 4 for $r = 1$.

Bound on the smallest eigenvalue— We are now interested in evaluating the smallest eigenvalue of the average Fisher information $G^{-1/2}\bar{I}(\rho_0)G^{-1/2}$. As in [25] we let

$$\mathbf{B}_U := \left\{ |\mathbf{o}; U\rangle := U|\mathbf{o}\rangle : \mathbf{o} = 1, \dots, 2^N \right\} \quad (4.32)$$

denote the ONB basis obtained by rotating the standard basis by a random unitary U . With this notation, we get that for randomly chosen basis

$$G^{-1/2}\bar{I}(\rho_0)G^{-1/2} := G^{-1/2} \cdot \int \mu(dU) I(\rho_0 | \mathbf{B}_U) \cdot G^{-1/2} \quad (4.33)$$

where $\mu(dU)$ is the Haar measure over unitaries used for generating the random basis. The integral in the above equation has been evaluated in [25], and we do not reproduce the calculation here. However, we point out that the integral in [25] has been evaluated for a slightly different parametrisation of the state. Since we constrain the $\rho_{1,1}$ element, the partial derivatives in our parametrisation become

$$\frac{\partial p_\rho(\mathbf{o} | \mathbf{B}_U)}{\partial \rho_{i,i}} = |\langle \mathbf{o}, U|i \rangle|^2 - |\langle \mathbf{o}, U|1 \rangle|^2 \quad (4.34)$$

$$\frac{\partial p_\rho(\mathbf{o} | \mathbf{B}_U)}{\partial \text{Re} \rho_{i,j}} = 2 \text{Re}(\langle i | \mathbf{o}, U \rangle \langle \mathbf{o}, U | j \rangle) \quad (4.35)$$

$$\frac{\partial p_\rho(\mathbf{o} | \mathbf{B}_U)}{\partial \text{Im} \rho_{i,j}} = 2 \text{Im}(\langle i | \mathbf{o}, U \rangle \langle \mathbf{o}, U | j \rangle) \quad (4.36)$$

Going through the calculation with this change gives us

1. The **Diagonal-Diagonal** block with $r > 1$:

$$(a) \quad \bar{I}_{a,a}^{dd} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} = \frac{2r}{r+1} \text{ when } r_a \leq r$$

$$(b) \quad \bar{I}_{a,b}^{dd} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} = \frac{r}{r+1} \text{ when } r_a, r_b \leq r, \text{ and } a \neq b$$

2. The **Real-Real** and **Imaginary-Imaginary** blocks are diagonal with:

$$(a) \quad \bar{I}_{a,a}^{rr/ii} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} = \frac{2r}{r+1} \text{ when } r_a < c_a \leq r$$

$$(b) \quad \bar{I}_{a,a}^{rr/ii} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} = 2 \text{ when } r_a \leq r, c_a > r$$

On comparing this with the weight matrix G , we once again notice that both G and $\bar{I}(\rho_0)$ have the same block diagonal structure, with the off-diagonal blocks being zero. So we can write

$$G^{-1/2}\bar{I}G^{-1/2} = G^{dd-1/2}\bar{I}^{dd}G^{dd-1/2} \oplus G^{rr-1/2}\bar{I}^{rr}G^{rr-1/2} \oplus G^{ii-1/2}\bar{I}^{ii}G^{ii-1/2}$$

We notice that $\bar{I}^{dd} = \frac{r}{r+1} \cdot G^{dd}$, which gives us

$$G^{-1/2} \bar{I} G^{-1/2} = \frac{r}{r+1} \mathbf{1}_{(r-1)} \oplus \frac{1}{2} \bar{I}^{rr} \oplus \frac{1}{2} \bar{I}^{ii} \quad (4.37)$$

The minimum eigenvalue of this matrix is $r/r+1$ for $r > 1$ and 1 for pure states.

Putting it all together– We can now substitute these values into the matrix Chernoff bound. While the value of the minimum eigenvalue differs for $r > 1$ and $r = 1$, the final bound remains the same because the upper bounds are different in these cases. Therefore here we calculate the bound for the case when $r > 1$. Writing $P_S = G^{-1/2} I_S G^{-1/2}$ and $\bar{P} = G^{-1/2} \bar{I} G^{-1/2}$ for notational simplicity, we have for $r > 1$

$$\mathbb{P} \left\{ P_S \notin \left[(1-\epsilon) \bar{P}, (1+\epsilon) \bar{P} \right] \right\} \leq 2D \cdot \exp \left(-k \frac{\epsilon^2}{4(r+1) \cdot \log 2} \right) := \delta \quad (4.38)$$

Therefore, with probability $1 - \delta$ we have that

$$(1-\epsilon) \bar{P} \leq P_S \leq (1+\epsilon) \bar{P} \quad (4.39)$$

This can be re-written in the form of inequalities of Mean Square Errors with $\epsilon > 0$ sufficiently small

$$(1-\epsilon) \text{Tr}(\bar{P}^{-1}) \leq \text{Tr}[P_S^{-1}] \leq (1+\epsilon) \text{Tr}(\bar{P}^{-1}) \quad (4.40)$$

For a fixed value of ϵ and δ , we see that the minimum number of settings k required for the above bound to hold with probability greater than $1 - \delta$ is

$$k = C \cdot (r+1) \log \left(\frac{2D}{\delta} \right) \quad (4.41)$$

where $C := 4(\log 2/\epsilon^2)$ and $D := 2rd - r^2 - 1$.

□

Chapter 5

Statistical analysis of compressive low rank tomography with random measurements

5.1 Introduction

In chapters 3 and 4 we have seen that the compressed sensing (CS) paradigm is motivated by the fact that full tomography often becomes challenging for large dimensional states. There is significant interest in addressing this challenge posed by dimensionality and as a result, extensive work has been done in developing tomography methods for certain lower dimensional families of physically relevant states. Including CS tomography of low rank states with incomplete measurements, pertinent examples include model selection [61], spectral thresholding [3, 25] and the estimation of matrix product states [39] which is relevant for many-body systems and also for estimating dynamical parameters of open systems [35, 60].

Similar to the work presented in chapter 4, several papers [81, 73] consider the problem of estimating low rank states from random measurements. Inspired by the PhaseLift problem, the papers [81, 32, 53] consider the case of estimating low rank states from expectations of rank-one projections sampled randomly from a Gaussian distribution, or a projective t-design, and demonstrate stable compressive recovery with estimation errors of the order of the number of unknown parameters. Compressive quantum process tomography has been considered in this context for unitary 2-designs [77]. In [80] the analysis is extended to the physically relevant case of random orthonormal basis measurements, and it is shown that a rank- r state can be identified with a large probability for only $O(r \log^3 d)$ such random measurements. Related to this question of low-rank state estimation, work in [10] conjectures that only a few random bases correspond to *strictly complete POVMs* for low rank states,

implying that states of a given low rank can be compressively recovered by measuring a small number of random bases, independent of dimension.

In this chapter we build on the work in chapter 4, and consider the statistical problem of estimating low-rank states in the set up of random basis measurements. Instead of choosing a particular estimator, the idea is to investigate the statistical efficiency of an arbitrary optimal estimator, and find whether rank- r states can be estimated from only a few random bases measurements. For this, we consider the behaviour of the *Mean Square Error* (MSE) with respect to the Frobenius distance between the true state and the estimator $\|\hat{\rho} - \rho\|_F^2$ in the limit of large numbers of measurement samples. According to asymptotic theory [118], in the regime of large number of repetitions the MSE of efficient estimators (e.g. maximum likelihood) $\hat{\rho}$ takes the following expression

$$\mathbb{E}\|\hat{\rho} - \rho\|_F^2 = \frac{1}{n} \text{Tr}(I(\rho|\mathcal{S})^{-1}G_F) + o(n^{-1}). \quad (5.1)$$

Above, $I(\rho|\mathcal{S})$ is the classical Fisher information associated with the chosen measurement design \mathcal{S} and a local parametrisation of rank- r states, n is the total number of measured systems, and G_F is the positive weight matrix associated with the quadratic approximation of the Frobenius distance in the local parameters.

In chapter 4 we showed that the asymptotic MSE (5.1) remains robust even with only a few random basis measurements making up the design \mathcal{S} . This robustness was explained using an argument based on a concentration inequality [2] for the Fisher information matrix. We also demonstrated that certain ‘least sparse’ states of rank- r can be estimated by using only $O(r \log d)$ settings with only a small increase in the MSE, relative to the setup in which a large number of settings is probed. In this chapter the argument using the concentration of the Fisher information is extended to hold for all rank- r states (Theorem 4), incorporating the results in the previous chapter. However, we discuss drawbacks of using a concentration in the Fisher information to derive a corresponding concentration in the MSE. Specifically, for rank- r states that are close to pure with small eigenvalues, we show that such a concentration of the Fisher information does not hold. This difficulty is overcome by proving an *upper bound* on the MSE that holds for all states independently of their spectrum. We show that $\text{Tr}(I(\rho|\mathcal{S})^{-1}G_F)$ is bounded from above by roughly the number of unknown parameters given that $O(r \log d)$ random bases constitute the measurement design \mathcal{S} . As an illustrative example, we consider a single

qubit state and analyse the failure of the Fisher concentration for states that are close to pure. We argue that despite a lack of concentration in the Fisher information for such states, the MSE demonstrates the necessary concentration.

The lack of concentration of the Fisher information occurs as the elements of the matrix corresponding to the small eigenvalues of the state diverge. However for states that are known to be pure and have $d - 1$ eigenvalues exactly zero, the Fisher information matrix contains entries that correspond only to the ‘rotation’ parameters. Therefore it concentrates following Theorem 4, and Theorem 2 in chapter 4 given that $O(\log d)$ random basis are measured. Related work in a different context [84] shows that (local) informationally complete measurements for pure states require the number of outcomes to scale linearly with the dimension d , specifically $2d - 1$. In our measurement design, the total number of outcomes is $O(d \log d)$ for pure states.

For the single qubit case, we also investigate the problem of ‘compressive’ state estimation using the quantum infidelity $1 - F(\hat{\rho}, \rho) = 1 - \text{Tr}(\sqrt{\sqrt{\rho}\hat{\rho}\sqrt{\rho}})^2$ as the error metric. For this we consider the asymptotic mean infidelity (MINF),

$$\mathbb{E}[1 - F(\hat{\rho}, \rho)] = \frac{1}{n} \text{Tr}(I(\rho|\mathcal{S})^{-1}G_{INF}), \quad (5.2)$$

with the Fisher information as defined in (5.1), and G_{INF} being the weight matrix corresponding to the quadratic approximation of the infidelity. Unlike the Frobenius distance, the quantum infidelity is very sensitive to the misestimation of small eigenvalues. In particular, for states that are close to pure with small eigenvalues, the local expansion of the infidelity in the asymptotic regime is linear in the estimation error of these eigenvalues [89]. This means that the MINF is no longer given by quadratic expression (5.2) for such states. We show that for states with eigenvalues well away from zero, a concentration in the MINF given by (5.2) can be demonstrated using a concentration of the Fisher information matrix, while for nearly pure states and random measurements both the Fisher information and the MINF demonstrate a lack of concentration. For such states the MINF scales as $O(1/\sqrt{n})$, and additionally there is no finite number of settings such that the state can be estimated ‘compressively’.

5.2 Quantum tomography with random basis measurements

In this chapter we consider the problem of estimating an unknown quantum state represented by a $d \times d$ density matrix ρ (complex, positive trace-one matrix), where d is the dimension of the associated Hilbert space \mathcal{H}^d . The unknown state is reconstructed from the outcomes of projective measurements on identical copies of the state. The measurement settings are chosen by randomly drawing an orthonormal basis (ONB) from the *uniform measure*, or equivalently by rotating a fixed (standard) ONB with a random unitary U drawn from the Haar measure over the unitaries on \mathcal{H}^d . We denote measurement settings by \mathbf{s} and the corresponding ONBs by $\{|e_{\mathbf{s}}^{\mathbf{o}}\rangle\}$ where $\mathbf{o} \in \{1, \dots, d\}$ is the label of a measurement outcome. Its probability is $p_{\rho}(\mathbf{o}|\mathbf{s}) := \text{Tr}(\rho P_{\mathbf{o}}^{\mathbf{s}})$, where $P_{\mathbf{o}}^{\mathbf{s}} = |e_{\mathbf{s}}^{\mathbf{o}}\rangle\langle e_{\mathbf{s}}^{\mathbf{o}}|$ is the one-dimensional projection corresponding to the outcome \mathbf{o} , in the measurement setting \mathbf{s} . Because of the cyclicity of the trace, this measurement design is mathematically equivalent to fixing a particular measurement basis and rotating the state ρ with a known random unitary corresponding to the measurement design.

This design is motivated by the multiple ion tomography (MIT) set up of ion-trap experiments [62] considered in the previous few chapters. The aim of MIT is to determine the unknown density matrix $\rho \in \mathcal{H}^d$ of the joint state of a system of N ions, where $d = 2^N$ is the dimension of the associated Hilbert space. A random measurement setting \mathbf{s} in the MIT setup can be thought of as a rotation of the fixed $\sigma_z^{\otimes N}$ basis by a random unitary drawn from the Haar measure over the whole Hilbert space \mathcal{H}^{2^N} .

The measurement procedure and statistical model are summarised below, following the notation in the previous chapter. For each given setting, the measurement is repeated on m copies of the state. This is repeated for all k settings, and the total number of copies of the state utilised as a resource is $n = m \times k$. This procedure results in information about the number of times a particular outcome was observed for a given setting $N(\mathbf{o}|\mathbf{s})$. This information can be thought of as a $d \times k$ dataset of counts whose columns are independent. In this chapter we investigate the statistical efficiency of estimating low rank states from such measurement outcomes. We work in the asymptotic regime in which the number of repetitions m in each setting is large, and characterise the estimation errors in terms of the classical Fisher information matrix as explained below.

We first introduce a parametrisation of the state ρ , and assume that the state has rank $r \leq d$, and therefore belongs to the space of rank r states $\mathbb{S}_r \subset M(\mathbb{C}^d)$. In the asymptotic scenario the estimation error is characterised by a local statistical model. Therefore, we consider a local parametrisation $\boldsymbol{\theta} \rightarrow \rho_{\boldsymbol{\theta}}$ of the state ρ in the space \mathbb{S}_r . In this space, any rank- r state ρ' in the local neighbourhood of ρ can be obtained by a perturbation of the eigenvalues of ρ , along with a small rotation of the eigenbasis. In the first order this transformation leaves the $(d-r) \times (d-r)$ lower-right corner unchanged, so that in the eigenbasis of the state ρ we have

$$\rho' = \begin{pmatrix} \text{Diag}(\lambda_1, \dots, \lambda_r) & 0 \\ 0 & 0 \end{pmatrix} + \begin{pmatrix} \Delta_{diag} & \Delta_{off} \\ \Delta_{off}^\dagger & O(\|\Delta\|^2) \end{pmatrix}. \quad (5.3)$$

We therefore choose to parametrise such a state $\rho' = \rho'_{\boldsymbol{\theta}}$ with

$$\begin{aligned} \boldsymbol{\theta} &:= (\boldsymbol{\theta}^{(d)}; \boldsymbol{\theta}^{(r)}; \boldsymbol{\theta}^{(i)}) \\ &= (\rho'_{2,2}, \dots, \rho'_{r,r}; \text{Re}\rho'_{1,2}, \dots, \text{Re}\rho'_{r,d}; \text{Im}\rho'_{1,2}, \dots, \text{Im}\rho'_{r,d}) \in \mathbb{R}^{2rd-r^2-1} \end{aligned} \quad (5.4)$$

where, the first diagonal matrix element does not appear in the parametrisation as it is fixed by the trace normalisation of density matrices. We can now describe the statistical model in this parametrisation, and define the classical Fisher information matrix associated with a given setting \mathbf{s} as

$$I(\rho|\mathbf{s})_{a,b} := \sum_{\mathbf{o}: p_{\rho}(\mathbf{o}|\mathbf{s}) > 0} \frac{1}{p_{\rho}(\mathbf{o}|\mathbf{s})} \frac{\partial p_{\rho}(\mathbf{o}|\mathbf{s})}{\partial \theta_a} \cdot \frac{\partial p_{\rho}(\mathbf{o}|\mathbf{s})}{\partial \theta_b}, \quad (5.5)$$

where $\theta_{a,b}$ are labelled elements of the parameter vector $\boldsymbol{\theta}$. Following the measurement procedure described above, we define the set of k measurement settings as \mathcal{S} . The Fisher information matrix associated with a single measurement from each setting $\mathbf{s} \in \mathcal{S}$ is given by the sum of the individual Fisher matrices above. The average Fisher information for the measurement design \mathcal{S} is denoted as $I(\rho|\mathcal{S}) = \frac{1}{k} \sum_{\mathbf{s} \in \mathcal{S}} I(\rho|\mathbf{s})$. The individual matrices are computed using definition (5.5) together with parametrisation (5.4).

The measurement in each setting is repeated m times on identical copies of the state, and the outcomes are i.i.d. When this number m is sufficiently large, efficient estimators of $\boldsymbol{\theta}$ (and hence of ρ) from these outcomes have an

asymptotically Gaussian distribution [118]

$$\sqrt{m}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \approx \mathcal{N}(0, I(\rho|\mathcal{S})^{-1}) \quad (5.6)$$

where the covariance matrix $I(\rho|\mathcal{S})^{-1}$ is the Fisher information associated with a single measurement sample of the set \mathcal{S} . In the following section, this asymptotic behaviour of the estimate $\hat{\boldsymbol{\theta}}$ is combined with local expansions of the Frobenius distance in order to characterise the MSE in terms of the classical Fisher information matrix. In section 5.4.1 the asymptotic mean infidelity (MINF) is investigated using similar techniques.

5.3 Bounds for the MSE

We now consider the Frobenius distance, and characterise the efficiency of any efficient estimator (such as maximum likelihood) in its terms. This distance has a locally quadratic expansion around the state ρ , given by

$$\|\rho_{\boldsymbol{\theta}} - \rho_{\boldsymbol{\theta} + \delta\boldsymbol{\theta}}\|_F^2 = (\delta\boldsymbol{\theta})^T G_F(\delta\boldsymbol{\theta}) + o(\|\delta\boldsymbol{\theta}\|^2), \quad (5.7)$$

where G_F is a constant weight matrix that reproduces the Frobenius norm. The explicit form of this weight matrix can be found in the appendix. From this and the asymptotic behaviour of efficient estimators, we see that for (reasonably) large m , the mean square error scales as

$$\text{MSE} := \mathbb{E}(\|\hat{\rho} - \rho\|_F^2) \approx \frac{1}{n} \text{Tr}(I(\rho|\mathcal{S})^{-1} G_F). \quad (5.8)$$

The expression of the right side of the above equation is a measure of the sensitivity of the chosen set of settings \mathcal{S} at ρ . We therefore consider the behaviour of the MSE in terms of this trace expression, and study the effect of the measurement design \mathcal{S} and the number of settings k on the error. We first present a preliminary concentration bound for this quantity $\text{Tr}(I(\rho|\mathcal{S})^{-1} G_F)$, which extends the results in chapter 4.

The bound determines the number of settings k required for the MSE $\text{Tr}(I(\rho|\mathcal{S})^{-1} G_F)$ to be concentrated close its optimal value. This result is derived from a concentration of the Fisher information matrix around the mean Fisher information, where the main ingredient is a matrix Chernoff bound for sums of bounded random Hermitian matrices. Since the settings in \mathcal{S} are independent, the Fisher information matrices $I(\rho|\mathbf{s})$ are independent and this

bound is applicable. The Chernoff bound determines how quickly the average information per setting $\frac{1}{k} \sum_{k \in \mathcal{S}} I(\rho|\mathbf{s})$ approaches the mean information \bar{I} over all random settings. In terms of the MSE, this translates to determining the number of settings k required for the MSE $\text{Tr}(I(\rho|\mathcal{S})^{-1}G_F)$ to be concentrated close the optimal value of $\text{Tr}(\bar{I}(\rho)^{-1}G_F)$. We consider states with arbitrary spectrums $\rho := \text{Diag}(\lambda_1, \dots, \lambda_r, \dots, 0)$, diagonal with respect to its eigenbasis. Due to the unitary symmetry of the random settings design, the eigenbasis can be chosen to be the standard basis.

Theorem 4. *Let $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_k\}$ be a design with randomly, uniformly distributed measurement bases. Let $I_{\mathcal{S}} := I(\rho|\mathcal{S})$ be the associated Fisher information, and let \bar{I} be the mean Fisher information over all possible bases, both calculated at the true state ρ . For a sufficiently small $\epsilon \geq 0$, the following inequality holds*

$$(1 - \epsilon) \text{Tr}[\bar{I}^{-1}G_F] \leq \text{Tr}[I_{\mathcal{S}}^{-1}G_F] \leq (1 + \epsilon) \text{Tr}[\bar{I}^{-1}G_F]$$

with probability $1 - \delta$, provided that the number of measurements performed is $k = \frac{C_1}{\lambda_{\min}(\rho)} \frac{(r+1)}{r} \log(\frac{2D}{\delta})$, with $D = 2rd - r^2 - 1$ the dimension of the space of rank- r states, and $C_1 = 4(\log 2/\epsilon^2)$.

The proof of this theorem and further details can be found in the appendix. As mentioned earlier, the main ingredient is a matrix Chernoff bound [2], which is used to bound the deviation of $G_F^{-1/2}I(\rho|\mathcal{S})G_F^{-1/2}$ from the mean $G_F^{-1/2}\bar{I}(\rho)G_F^{-1/2}$. The number of uniformly random settings k required in the theorem above depends on the following ratio

$$\frac{\mu_{\max}}{\mu_{\min}} := \frac{\max_{\mathbf{s}} \lambda_{\max} \left(G_F^{-1/2} I(\rho|\mathbf{s}) G_F^{-1/2} \right)}{\lambda_{\min} \left(G_F^{-1/2} \bar{I}(\rho) G_F^{-1/2} \right)} \quad (5.9)$$

between the largest maximum eigenvalue of $G_F^{-1/2}I(\rho|\mathbf{s})G_F^{-1/2}$ over all possible measurements and the minimum eigenvalue of $G_F^{-1/2}\bar{I}(\rho)G_F^{-1/2}$. Details of the explicit values of this ratio is left to the appendix. The numerator μ_{\max} is upper bounded by using the inequality between the quantum and classical Fisher informations [24], as $\mu_{\max} \leq 2/\lambda_{\min}(\rho)$ for $r > 1$ and $\mu_{\max} \leq 2$ for $r = 1$, while the minimum eigenvalue of $G_F^{-1/2}\bar{I}(\rho)G_F^{-1/2}$ is lower bounded using the following lemma.

Lemma 1. *For any rank- r state ρ with an arbitrary spectrum, and the rank- r state ρ_0 which has equal non-zero eigenvalues $1/r$ and the same eigenvectors as ρ , the following inequality holds between their average Fisher information matrices, evaluated over all possible random measurement settings.*

$$\bar{I}(\rho_0) \leq \bar{I}(\rho) \quad (5.10)$$

The proof is left to the appendix. The matrix $\bar{I}(\rho_0)$ for the equal eigenvalue state has been computed explicitly by using analytic expressions for moments of random unitaries [37], which gives $\mu_{\min} \geq \frac{r}{r+1}$ for $r > 1$, and $\mu_{\min} \geq 1$ for pure states. Together these give $\frac{\mu_{\max}}{\mu_{\min}} \leq 2 \frac{(r+1)}{r} \frac{1}{\lambda_{\min}(\rho)}$ which determines the number of measurement settings in the theorem above. When the state ρ is the equal eigenvalue state ρ_0 , we get $\lambda_{\min}(\rho_0) = 1/r$ and we recover the rate presented in the previous chapter.

We noted in chapter 4 that deriving a concentration in the MSE via a concentration of Fisher average $I(\rho|\mathcal{S})$ provides a pessimistic estimate of the number of settings needed. Simulations in chapter 4 demonstrated that the MSE concentrates for a much smaller number of settings k than predicted. In the theorem presented above, we note that the dependence of the number of settings on the minimum eigenvalue of ρ suggests a lack of concentration as $\lambda_{\min}(\rho)$ is made arbitrarily small. The number of required settings $k \rightarrow \infty$ in the limit that $\lambda_{\min}(\rho) \rightarrow 0$. This is because the maximum eigenvalue of the Fisher information $I(\rho|\mathbf{s})$ over all settings \mathbf{s} becomes arbitrarily large when the rank- r state ρ is arbitrarily close to being pure. However, as we shall demonstrate, this does not reflect the behaviour of the MSE concentration. Instead of deriving a concentration about $\bar{I}(\rho)$ as in the above theorem, we derive a useful upper bound for the MSE that is independent of the spectrum of the state.

Theorem 5. *Let $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_k\}$ be a design with randomly, uniformly distributed measurement bases. Let $I_{\mathcal{S}} := I(\rho|\mathcal{S})$ be the associated Fisher information evaluated at ρ . For a sufficiently small $\epsilon \geq 0$, the following inequality holds*

$$\text{Tr}[I(\rho|\mathcal{S})^{-1}G_F] \leq 2(1+\epsilon)\frac{r+1}{r}D$$

with probability $1 - \delta$, provided that the number of measurements performed is $k = C_1(r+1)\log(2D/\delta)$, with $D = 2rd - r^2 - 1$ the dimension of the space of rank- r states, and $C_1 = 4(\log 2/\epsilon^2)$.

The upper bound is roughly twice the number of unknown parameters, and although not optimal, it demonstrates that the MSE concentrates below a meaningful threshold given a fixed $O(r \log D)$ scaling in the number of settings. The proof of this theorem follows now. A key element in the proof is to overcome the potential unboundedness of the maximum eigenvalue of $I(\rho|\mathbf{s})$. This is done by bounding $I(\rho|\mathbf{s})$ from below over all possible settings \mathbf{s} by matrices whose spectrums are well behaved. This in turn gives us an upper bound for the inverse of the sum $I(\rho|\mathcal{S})^{-1}$.

To this end, we define a new state $\tilde{\rho}$ such that over all possible settings \mathbf{s} , we have the following inequality in the Fisher matrices

$$I(\rho|\mathbf{s}) \geq \frac{1}{2}I(\tilde{\rho}|\mathbf{s}). \quad (5.11)$$

The state $\tilde{\rho}$ is defined to be $\tilde{\rho} := (\rho + \rho_0)/2$, where ρ_0 is the rank- r state with equal $1/r$ eigenvalues, and the same eigenvectors as ρ . It is easy to see that $\tilde{\rho}$ has eigenvalues bounded between $(1 + 1/r)/2$ and $1/2r$, and has the same eigenvectors as ρ by construction. The above inequality then follows from the fact that $\rho \leq 2\tilde{\rho}$, and from the definition of the Fisher information matrix (5.5). For any given measurement design $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_k\}$, this inequality in the Fisher matrices implies that $I(\rho|\mathcal{S}) \geq I(\tilde{\rho}|\mathcal{S})/2$. Since the matrix $I(\tilde{\rho}|\mathbf{s})$ has eigenvalues that are well behaved over all possible settings \mathbf{s} , we can use Theorem 4 to meaningfully bound the deviation $G_F^{-1/2}I(\tilde{\rho}|\mathcal{S})G_F^{-1/2}$ from its mean. In fact, we get that for a sufficiently small $\epsilon \geq 0$, the following inequality holds

$$(1 - \epsilon)\text{Tr}[\bar{I}(\tilde{\rho})^{-1}G_F] \leq \text{Tr}[I(\tilde{\rho}|\mathcal{S})^{-1}G_F] \leq (1 + \epsilon)\text{Tr}[\bar{I}(\tilde{\rho})^{-1}G_F] \quad (5.12)$$

with probability $1 - \delta$, provided that the number of settings $k = C_1(r + 1)\log(2D/\delta)$. The upper bound in the equation above, combined with the inequality $I(\rho|\mathcal{S}) \geq I(\tilde{\rho}|\mathcal{S})/2$ gives the stated upper bound

$$\text{Tr}[I(\rho|\mathcal{S})^{-1}G_F] \leq 2(1 + \epsilon)\text{Tr}[\bar{I}(\tilde{\rho})^{-1}G_F] \leq 2(1 + \epsilon)\frac{r+1}{r}D. \quad (5.13)$$

□

Theorem 5 derives a uniform upper bound for all rank- r states irrespective of the eigenvalue spectrum. This demonstrates that sensible bounds exist in the limit of $\lambda_{\min}(\rho) \rightarrow 0$ for a finite number of measurement settings k . It

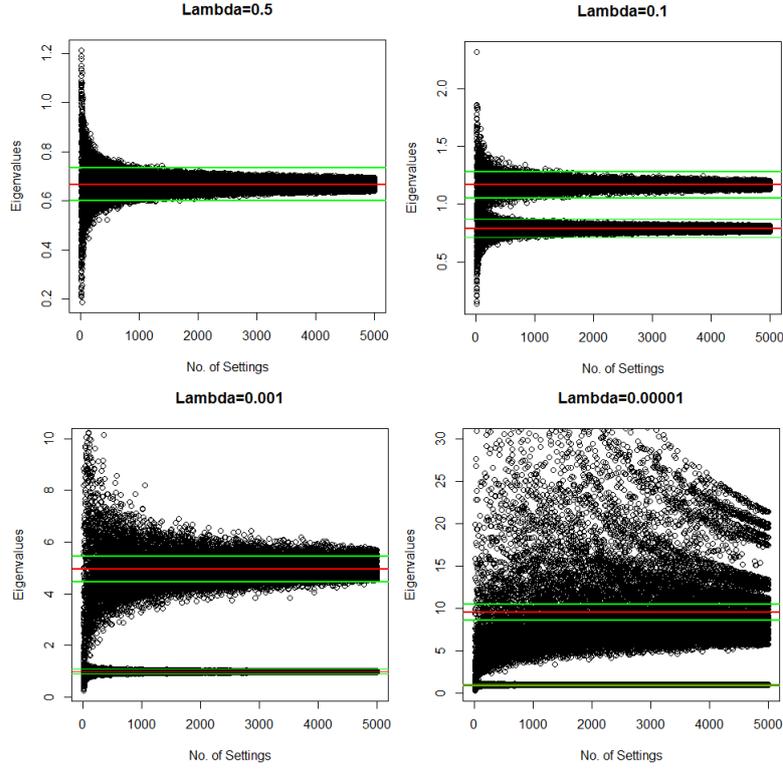


Figure 5.1 Plots of the eigenvalues of $G_F^{-1/2} I(\rho|\mathcal{S}) G_F^{-1/2}$ for various k random settings. We chose 40 random single qubit states for each of the four values of λ_2 . The red line indicates the eigenvalues of $\bar{I}(\rho)/2$, with the green marking the $(1 \pm \epsilon)\bar{I}/2$ deviations ($\epsilon = 0.1$). We observe that as the state becomes purer, the number of settings needed for concentration increases, and in the limit $\lambda_2 \rightarrow 0$ there is a lack of concentration of the largest eigenvalue.

is clear that the divergence of the maximum eigenvalue $\max_{\mathbf{s}} \lambda_{\max}(I(\rho|\mathbf{s}))$ as $\lambda_{\min}(\rho) \rightarrow 0$ does not cause a similar divergence in the MSE. Therefore Theorem 4 does not sensibly define a rate for the required number of measured settings k in the limit $\lambda_{\min}(\rho) \rightarrow 0$.

Although Theorem 5 derives a uniform upper bound for the MSE, it does not demonstrate a concentration in the MSE. However, for the simplified model for a rank-2 qubit state, we show that a concentration in the MSE does in fact hold in the limit $\lambda_{\min}(\rho) \rightarrow 0$.

5.4 The Single Qubit Model

In this section we work with the simple model of a rank-2 qubit state to show that a concentration in the MSE about its optimal holds in the limit $\lambda_{\min}(\rho) \rightarrow 0$ without requiring the sum $I(\rho|\mathcal{S})$ to concentrate about \bar{I} .

Lemma 2. *Let ρ be a single qubit rank-2 state, and let $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_k\}$ be a uniformly random measurement design. Let $I_{\mathcal{S}} := I(\rho|\mathcal{S})$ be the associated Fisher information, and let $\bar{I}(\rho)$ be the mean Fisher information over all possible measurement bases. For any $\epsilon > 0$, there exists a finite k such that the following inequality holds for all ρ with high probability*

$$\text{Tr}[I(\rho|\mathcal{S})^{-1}G_F] \leq (1 + \epsilon)\text{Tr}[\bar{I}(\rho)^{-1}G_F]. \quad (5.14)$$

In order to investigate the behaviour of the MSE concentration as the spectrum is varied, we consider the generic state $\rho := \lambda_1|0\rangle\langle 0| + \lambda_2|1\rangle\langle 1|$ diagonal in its eigenbasis. We consider the same local parametrisation as in the previous sections and denote $\boldsymbol{\theta} := (\lambda_2, \text{Re}\rho_{1,2}, \text{Im}\rho_{1,2})$. The measurement design consists of random, uniformly distributed measurement bases, and without loss of generality we set the projection vector corresponding to the +1 outcome for a given setting \mathbf{s} as:

$$|e_{\mathbf{s}}^{+1}\rangle := \cos\frac{\phi}{2}|0\rangle + e^{i\omega}\sin\frac{\phi}{2}|1\rangle \quad 0 \leq \phi \leq \pi, \quad 0 \leq \omega \leq 2\pi \quad (5.15)$$

The orthogonal vector corresponds to the -1 outcome. Therefore, the probabilities $p_{\rho}(\mathbf{o}|\mathbf{s})$ corresponding to the two outcomes are $p_{\rho}(+1|\mathbf{s}) = (1 - \lambda_2)\cos^2\frac{\phi}{2} + \lambda_2\sin^2\frac{\phi}{2}$ and $p_{\rho}(-1|\mathbf{s}) = (1 - \lambda_2)\sin^2\frac{\phi}{2} + \lambda_2\cos^2\frac{\phi}{2}$. From equation (5.5), we evaluate the elements of the Fisher information matrix for a given random measurement setting \mathbf{s} .

$$\begin{aligned} I(\rho|\mathbf{s}) &= \begin{pmatrix} I^{dd}(\rho|\mathbf{s}) & I^{rd}(\rho|\mathbf{s}) & I^{id}(\rho|\mathbf{s}) \\ I^{rd}(\rho|\mathbf{s}) & I^{rr}(\rho|\mathbf{s}) & I^{ri}(\rho|\mathbf{s}) \\ I^{id}(\rho|\mathbf{s}) & I^{ri}(\rho|\mathbf{s}) & I^{ii}(\rho|\mathbf{s}) \end{pmatrix} \\ &= \frac{2}{1 - \cos^2(\phi)(1 - 2\lambda_2)^2} \begin{pmatrix} 2\cos^2(\phi) & -\cos(\omega)\sin(2\phi) & \sin(\omega)\sin(2\phi) \\ -\cos(\omega)\sin(2\phi) & 2\cos^2(\omega)\sin^2(\phi) & -\sin(2\omega)\sin^2(\phi) \\ \sin(\omega)\sin(2\phi) & -\sin(2\omega)\sin^2(\phi) & 2\sin^2(\omega)\sin^2(\phi) \end{pmatrix} \end{aligned} \quad (5.16)$$

As before \mathcal{S} is the set of k randomly chosen settings \mathbf{s} , and as the settings in \mathcal{S} are independent, the Fisher information matrices $I(\rho|\mathbf{s})$ are independent.

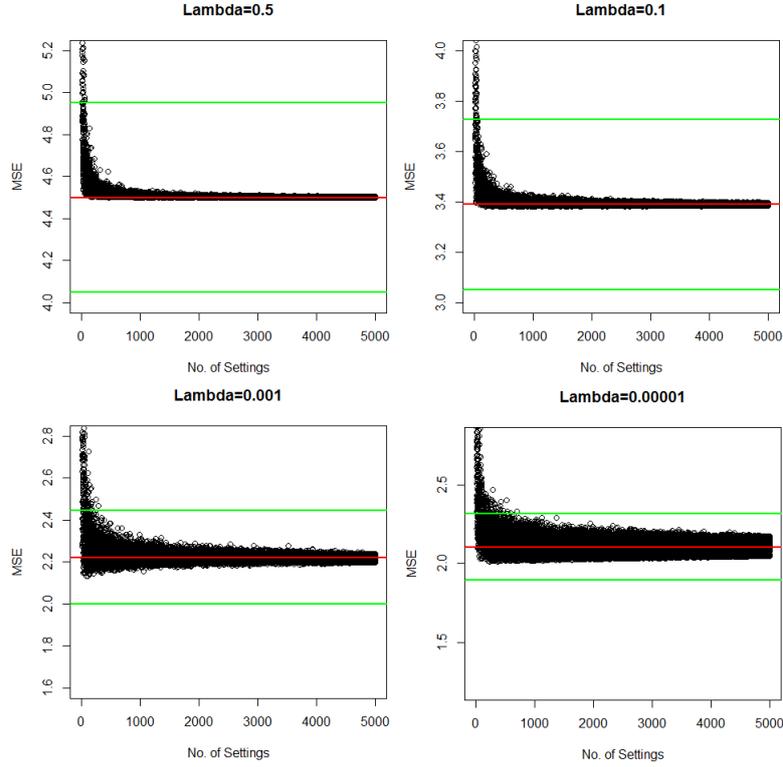


Figure 5.2 Plots of the MSE $\text{Tr}[I(\rho|\mathcal{S})^{-1}G_F]$ for various k random settings. We chose 40 random single qubit states for each of the four values of λ_2 . The red line indicates the theoretical optimal MSE $\text{Tr}[\bar{I}^{-1}G_F]$, with the green marking the $(1 \pm \epsilon)\text{Tr}[\bar{I}^{-1}G_F]$ deviations ($\epsilon = 0.1$). It is easier to observe concentration in the MSE, despite a lack of concentration of $I(\rho|\mathcal{S})$ (see Figure 5.1). Although the number of settings needed for concentration within a prescribed relative error increases with a decrease in λ_2 , there is a limiting value of k as $\lambda_2 \rightarrow 0$ (see text).

The concentration of the quantity $I(\rho|\mathcal{S}) := \frac{1}{k} \sum_{\mathbf{s} \in \mathcal{S}} I(\rho|\mathbf{s})$ around the mean Fisher matrix $\bar{I}(\rho)$ is given by Theorem 4. We recall that the number of settings k required to bound the deviation from its mean $\bar{I}(\rho)$ depends on the ratio of the eigenvalues

$$\frac{\mu_{\max}}{\mu_{\min}} := \frac{\max_{\mathbf{s}} \lambda_{\max} \left(G_F^{-1/2} I(\rho|\mathbf{s}) G_F^{-1/2} \right)}{\lambda_{\min} \left(G_F^{-1/2} \bar{I}(\rho) G_F^{-1/2} \right)}. \quad (5.17)$$

For the simple qubit model these can be explicitly evaluated. For settings with $\phi = 0$, the maximum eigenvalue is $\frac{1}{2\lambda_2(1-\lambda_2)}$. This implies that $\mu_{\max} \geq \frac{1}{4\lambda_2(1-\lambda_2)}$. This value is a contribution from the I^{dd} element of the Fisher matrix, and

Fisher Element	Mean	Range
I^{dd}	$\frac{2\ln[2(1-\lambda_2)] - 2\ln[2\lambda_2]}{(1-2\lambda_2)^3} - \frac{4}{(1-2\lambda_2)^2}$	$\left[0, \frac{1}{\lambda_2(1-\lambda_2)}\right]$
I^{rd}, I^{id}	0	$\left[\frac{-2}{\sqrt{\lambda_2(1-\lambda_2)}}, \frac{2}{\sqrt{\lambda_2(1-\lambda_2)}}\right]$
I^{ri}	0	$[-2, 2]$
I^{rr}, I^{ii}	$\frac{\ln[2(1-\lambda_2)] - \ln[2\lambda_2]}{(1-2\lambda_2)} - \frac{\bar{I}^{dd}}{2}$	$[0, 4]$

Table 5.1 The mean and range of the elements of the Fisher matrix $I(\rho|\mathbf{s})$ as functions of λ_2 . Note that the expressions for the means in the table above are valid for all $\lambda_2 < 0.5$. When $\lambda_2 = 0.5$, then all diagonal elements $\bar{I}^{rr/dd/ii}$ have the same value of $4/3$.

tends to infinity as $\lambda_2 \rightarrow 0$. The minimum eigenvalue μ_{\min} is a contribution from the \bar{I}^{rr} and the \bar{I}^{ii} term, and tends to a limiting value of 1 when $\lambda_2 \rightarrow 0$. The explicit expressions can be found in Table 5.1. Taken together this implies that the ratio becomes unbounded as $\lambda_2 \rightarrow 0$. This is precisely the difficulty characterised in the previous section, and is illustrated in Figure 5.1, where we plot the eigenvalues of the sum $G_F^{-1/2}I(\rho|\mathcal{S})G_F^{-1/2}$ for various values of λ_2 and choices of measurement designs \mathcal{S} .

However we are not interested in the concentration of the Fisher matrix itself, but rather the quantity $\text{Tr}[I(\rho|\mathcal{S})^{-1}G_F]$, and in Figure 5.2 it is seen that the MSE exhibits clear concentration about the optimal. Although the number of settings needed for the MSE to be within $(1 \pm \epsilon)$ of the optimal is seen to increase for smaller values of λ_2 , we shall show that there exists a limiting value of k as $\lambda_2 \rightarrow 0$. To demonstrate this, we consider the concentration of the individual Fisher elements, and directly bound the deviation of $\text{Tr}[I(\rho|\mathcal{S})^{-1}G_F]$ from its optimal.

It is clear from Table 5.1 that the Fisher matrix elements I^{rr}, I^{ii}, I^{ri} have bounded means and spread even in the limit $\lambda_2 \rightarrow 0$. Their sums can therefore be shown to concentrate around their means using one of several concentration inequalities. For example, we apply Hoeffding's inequality below.

Fact 1. *Let X_1, \dots, X_k be independent random variables such that each X_i is bounded as $a \leq X_i \leq b$, and let $\mu := \mathbb{E}[X]$. Let $S_k := \frac{1}{k} \sum_i^k X_i$, and $C := b - a$, then for any $t \geq 0$ and $\tau > 0$ the following inequalities hold,*

1. *Hoeffding's inequality : $\mathbb{P}(|S_k - \mu| \geq t) \leq 2e^{-2kt^2/C^2}$*

$$2. \text{ Markov's Inequality : } \mathbb{P}(|S_k| \geq \tau) \leq \frac{\mathbb{E}|X|}{\tau}$$

$$3. \text{ Chebyshev's Inequality : } \mathbb{P}(|S_k - \mu| \geq \tau) \leq \frac{\text{Var}(X)}{\tau^2}$$

From Table 5.1 we see that $C = 4$ for the I^{rr}, I^{ii}, I^{ri} matrix elements. Thus we derive that for any $t \geq 0$, their empirical means are within $\pm t$ of the true value with probability $(1 - \delta)$, provided that the number of settings $k \geq (8/t^2) \ln(2/\delta)$. Therefore the concentration for these elements is well behaved in the limit $\lambda_2 \rightarrow 0$. While the same inequality can be applied to I^{rd}, I^{id} matrix elements when λ_2 is away from zero, it fails in the limit $\lambda_2 \rightarrow 0$ because their ranges become infinite. However, we make a ‘*weak law of large numbers*’ argument to show that even in this limit, there exists a finite but ‘sufficiently large’ k , such that $I^{rd}(\rho|\mathcal{S})$ and $I^{id}(\rho|\mathcal{S})$ concentrate around their mean.

The key point is that the random variables I^{rd}, I^{id} remain absolutely integrable in limit $\lambda_2 \rightarrow 0$. This is combined with a truncation trick, to show that although the range of these variables is unbounded in the limit, for ‘sufficiently large’ k their empirical means converge in probability to their expected value. We follow the argument presented in [107] to demonstrate this. The idea of the truncation method is to split the random variable I^{rd} as

$$\begin{aligned} I^{rd} &:= I_{\leq T}^{rd} + I_{> T}^{rd} \\ &= I^{rd} \mathbb{1}(|I^{rd}| \leq T) + I^{rd} \mathbb{1}(|I^{rd}| > T), \end{aligned} \quad (5.18)$$

with T being a ‘truncation parameter’ that is chosen appropriately. We shall not be interested in the actual value of T , but endeavour only to show that such a method demonstrates the existence of a finite k for which I^{rd} converges in probability to zero. We similarly split the sum

$$I^{rd}(\rho|\mathcal{S}) = \frac{1}{k} \sum_{i=1}^k \left[I_{\leq T}^{rd}(\rho|\mathbf{s}_i) + I_{> T}^{rd}(\rho|\mathbf{s}_i) \right] =: I_{\leq T}^{rd}(\rho|\mathcal{S}) + I_{> T}^{rd}(\rho|\mathcal{S}) \quad (5.19)$$

We now bound these two sums using different inequalities. Since the random variable is absolutely integrable even in the limit $\lambda_2 \rightarrow 0$, we can always choose the truncation parameter T such that $\mathbb{E}|I_{> T}^{rd}|$ is made small, say some $\delta_2 > 0$, so that from Markov’s inequality (Fact 1) we get

$$\mathbb{P}\left(|I_{> T}^{rd}(\rho|\mathcal{S})| \geq \tau\right) \leq \frac{\delta_2}{\tau}. \quad (5.20)$$

The variable $I_{\leq T}^{rd}$ has bounded spread by construction, and therefore has bounded variance. This allows us to use Chebyshev's inequality (Fact 1), from which we see that

$$\mathbb{P}\left(|I_{\leq T}^{rd}(\rho|\mathcal{S})| \geq \tau\right) \leq \frac{\text{Var}(I_{\leq T}^{rd})}{k\tau^2}, \quad (5.21)$$

where we use the fact that $\mathbb{E}(I_{\leq T}^{rd}) = 0$, since the distribution is symmetric about zero. Clearly $\text{Var}(I_{\leq T}^{rd})$ is bounded, and there exists a finite k such that (5.20) and (5.21) together imply

$$\mathbb{P}\left(|I^{rd}(\rho|\mathcal{S})| \geq \tau\right) \leq \frac{\delta_2}{\tau} + \frac{1}{\tau^2}. \quad (5.22)$$

The term δ_2 can be made arbitrarily small by choosing T appropriately, which demonstrates that the sum converges in probability to zero for some finite, but 'sufficiently large' k . Although the above argument was demonstrated with $I^{rd}(\rho|\mathcal{S})$, the same holds for $I^{id}(\rho|\mathcal{S})$. This leaves the term $I^{dd}(\rho|\mathcal{S})$, which due to the non-integrability, infinite mean and range of I^{dd} in the limit $\lambda_2 \rightarrow 0$, does not concentrate around any finite value. The term I^{dd} contributes the maximum eigenvalue of the Fisher matrix over all settings \mathbf{s} , and as mentioned earlier its divergence in the limit $\lambda_2 \rightarrow 0$ is why a concentration inequality of the form of Theorem 4 does not hold. Collecting the individual bounds for the other matrix elements, we have that for any value of λ_2 there exists a finite k for which, with large probability, the matrix sum $I(\rho|\mathcal{S})$ has elements

$$\begin{pmatrix} \sum_{i=1}^k I_i^{dd}/k & [-\tau, +\tau] & [-\tau, +\tau] \\ [-\tau, +\tau] & [\mu - t, \mu + t] & [-t, +t] \\ [-\tau, +\tau] & [-t, +t] & [\mu - t, \mu + t] \end{pmatrix}, \quad (5.23)$$

where $\mu := \mathbb{E}(I^{rr/ii})$. We can now explicitly evaluate $\text{Tr}[I(\rho|\mathcal{S})^{-1}G_F]$, making the simplifying assumption that k is large enough to ignore terms quadratic in the off-diagonal elements, i.e, in τ and t . Going through the calculation, we get that provided $\sum_{i=1}^k I_i^{dd}/k > 1$,

$$\text{Tr}[I(\rho|\mathcal{S})^{-1}G_F] \leq \frac{2k}{\sum_{i=1}^k I_i^{dd}} + \frac{4}{\mu - t}. \quad (5.24)$$

In order to show that the MSE is close to optimal as in Lemma 2, we require that the term on the right in the above equation is smaller than

$(1 + \epsilon)\text{Tr}[\bar{I}(\rho)^{-1}G_F]$. That is, for some $\epsilon > 0$,

$$\frac{k}{\sum_{i=1}^k I_i^{dd}} + \frac{2}{\mu - t} \leq (1 + \epsilon) \left[\frac{1}{\bar{I}^{dd}} + \frac{2}{\mu} \right]. \quad (5.25)$$

When λ_2 is sufficiently large, the random variable I^{dd} is bounded and therefore the sum $\sum_{i=1}^k I_i^{dd}/k$ concentrates about its mean. In the limit $\lambda_2 \rightarrow 0$ however, $1/\bar{I}^{dd} \rightarrow 0$, which implies that the sum $\sum_{i=1}^k I_i^{dd}/k$ does not need to concentrate about its (infinite) mean, but only needs to be larger than a value dependent on ϵ and t . In the limit $\lambda_2 \rightarrow 0$, the Fisher element I^{dd} has a limiting distribution which can be explicitly evaluated. Setting $\lambda_2 = 0$ in (5.16) we have that $I^{dd} = 4 \cot^2 \phi$. Inverting this we get $\phi = \cot^{-1}(\sqrt{I^{dd}(\phi)}/2)$, and as the projection vectors are drawn uniformly over the unit sphere, ϕ is distributed as $f_\Phi(\phi) d\phi = \sin \phi d\phi$. Performing a change of variables then gives the limiting distribution $f_I(I^{dd}) dI^{dd} = \frac{2}{\sqrt{I^{dd}(I^{dd}+4)^{3/2}}} dI^{dd}$. From this distribution and the truncation method it is easy to show that for any value C , there exists a finite number of settings k such that $\sum I^{dd}/k > C$. This implies that for a given $\epsilon > 0$, and for all values of $\lambda_2 \in (0, 0.5]$ there always exists a finite number of settings k such that the required concentration holds.

5.4.1 Estimation Error in Terms of Quantum Infidelity

In this section we consider the problem of 'compressive' state estimation in terms of a different metric, the quantum infidelity

$$1 - F(\hat{\rho}, \rho) = 1 - \text{Tr} \left(\sqrt{\sqrt{\rho} \hat{\rho} \sqrt{\rho}} \right)^2. \quad (5.26)$$

As briefly hinted at in the introduction, a local expansion of this metric is not quadratic uniformly over all states. In particular for states that are well in the interior of the state space the expansion is locally quadratic, while for states with eigenvalues that are close to zero, the infidelity becomes linear [89]. This linear expansion highlights the sensitivity of the infidelity to misestimation of small eigenvalues, and we show that in our setup with uniformly random basis measurements, 'compressive' estimation for all states in the sense of Lemma 2 does not hold for this metric. To demonstrate this we continue considering the single qubit model from the previous section. We derive a theorem for the concentration of the mean infidelity (MINF) for states well within the Bloch sphere, and then demonstrate a lack of concentration for nearly pure states.

As before, we consider the state $\rho = \text{Diag}(1 - \lambda_2, \lambda_2)$ diagonal in its eigenbasis. For qubits, the infidelity can be expressed as [16]

$$1 - F(\hat{\rho}, \rho) = 1 - \text{Tr}(\hat{\rho}\rho) - 2\sqrt{\det\hat{\rho} \cdot \det\rho}. \quad (5.27)$$

A Taylor expansion of the infidelity about ρ demonstrates that for states within the Bloch sphere (i.e. λ_2 is well away from zero), the infidelity is locally quadratic in the (local) parameters

$$1 - F(\rho_{\boldsymbol{\theta}}, \rho_{\boldsymbol{\theta} + \delta\boldsymbol{\theta}}) = (\delta\boldsymbol{\theta})^T G_{INF}(\delta\boldsymbol{\theta}) + O(\|\delta\boldsymbol{\theta}\|^3), \quad (5.28)$$

where $G_{INF} = \text{Diag}(1/2\lambda_2(1 - \lambda_2), 2, 2)$ is the weight matrix reproducing the infidelity. In general for states of arbitrary dimension that have eigenvalues away from zero, the local expansion remains quadratic [89], and a concentration of the MINF is readily established using the techniques in the previous sections. Here we formulate this concentration for the single qubit state considered. Combining the above local expansion with the asymptotic normality of efficient estimators (5.6), the MINF is given by an expression similar to (5.8)

$$\text{MINF} := \mathbb{E}(1 - F(\hat{\rho}, \rho)) \approx \frac{1}{n} \text{Tr}(I(\rho|\mathcal{S})^{-1} G_{INF}). \quad (5.29)$$

A concentration of this error term can be demonstrated using the same tools used to establish Theorem 4. Concretely, we derive the following theorem.

Lemma 3. *Let $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_k\}$ be a design with randomly, uniformly distributed measurement bases. Let $I_{\mathcal{S}} := I(\rho|\mathcal{S})$ be the associated Fisher information, and let \bar{I} be the mean Fisher information over all possible bases, both calculated at the single qubit state ρ . For a sufficiently small $\epsilon \geq 0$, the following inequality holds*

$$(1 - \epsilon) \text{Tr}[\bar{I}^{-1} G_{INF}] \leq \text{Tr}[I_{\mathcal{S}}^{-1} G_{INF}] \leq (1 + \epsilon) \text{Tr}[\bar{I}^{-1} G_{INF}]$$

with probability $1 - \delta$, provided that the number of measurements performed is $k = \frac{C_2}{\lambda_2(1 - \lambda_2)} \log(\frac{2D}{\delta})$, with $D = 3$ the dimension of the space of rank-2 qubit states, and C_2 being a constant depending on ϵ that can be arbitrarily set.

Due to the dependence of the number of settings k on the minimum eigenvalue of the true state, the above lemma sensibly demonstrates concentration only

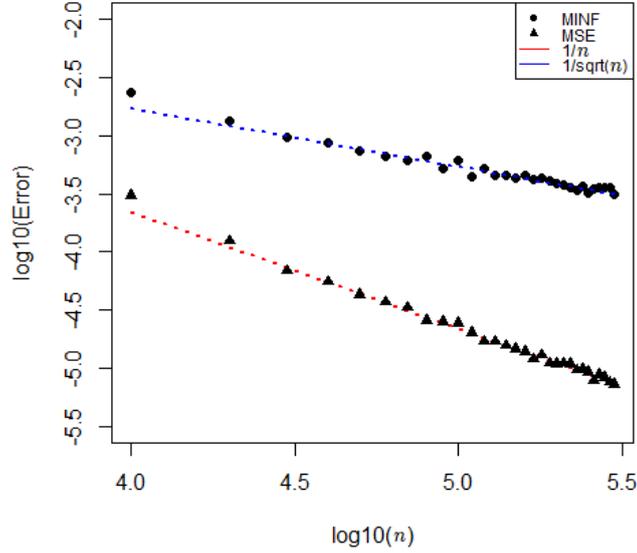


Figure 5.3 Plots of the MINF and MSE of the maximum likelihood estimate of a randomly chosen single qubit pure state and random basis measurements. The total number of samples of the state is $n = k \times m$, where k is the number of settings measured and $m = 1000$ is the number of repetitions per setting. The number of random basis measured k is varied between 10 and 300. The expected error is approximated over 300 different choices of k randomly chosen settings. The MSE demonstrates a $O(1/n)$ scaling, while for the same estimates the MINF scales as $O(1/\sqrt{n})$.

when λ_2 is away from zero. This is similar to the dependence of the number of settings on $\lambda_{\min}(\rho)$ in Theorem 4.

In the case of the MSE we demonstrated that for qubits, concentration does occur as $\lambda_2 \rightarrow 0$, even if Theorem 4 does not hold in this limit. However, a similar concentration of the MINF for qubits in this limit does not occur. To show this, we first notice that in the limit $\lambda_2 \rightarrow 0$ the local expansion of the infidelity becomes linear in the leading order

$$1 - F(\rho_{\theta}, \rho_{\theta + \delta\theta}) = |\delta\theta^d| + O(\|\delta\theta\|^2). \quad (5.30)$$

Clearly, for estimates $\hat{\rho}$ in the local neighbourhood of the pure state $\rho = |0\rangle\langle 0|$, the MINF is no longer given by the quadratic expression as in (5.29), but is $\mathbb{E}(1 - F(\hat{\rho}, \rho)) = \mathbb{E}(\hat{\theta}^d) = \mathbb{E}(\langle 1|\hat{\rho}|1\rangle)$. Since the dominant error term is linear in the diagonal element of the estimate (in the eigenbasis of the true state), we note that the infidelity is highly sensitive to the misestimation of small

eigenvalues [89]. The errors in the estimation of the ‘rotation parameters’ θ^r, θ^i however remain quadratic, and therefore exhibit a $O(1/n)$ scaling as in the previous sections. As the interesting contribution to the infidelity is from the estimation errors of the eigenvalue, we consider a simplified single parameter model and assume that only θ^d is unknown. When the number of repetitions m in a setting \mathbf{s} is sufficiently large, efficient estimators of θ^d from the outcomes of these measurements have an asymptotically Gaussian distribution $\sqrt{m}(\hat{\theta}^d - \theta^d) \approx \mathcal{N}(0, \text{Var}(\hat{\theta}^d))$. Therefore, in this asymptotic limit the MINF $\mathbb{E}(1 - F(\hat{\rho}, \rho))$ is given by

$$\mathbb{E}(\hat{\theta}^d) = \frac{1}{\sqrt{2\sigma^2\pi}} \int_0^\infty \hat{\theta}^d \cdot \exp\left(-\frac{(\hat{\theta}^d)^2}{2\sigma^2}\right) d\hat{\theta}^d = \sqrt{\frac{2}{\pi}}\sigma, \quad (5.31)$$

where negative estimates of the parameter are set to zero to ensure that $\hat{\rho}$ is a density matrix, and the standard deviation $\sigma = \text{Var}(\hat{\theta}^d)^{1/2}$. From this asymptotic behaviour of efficient estimators we see that for a large number of repetitions m , the MINF scales as

$$\mathbb{E}(1 - F(\hat{\rho}, \rho)) \approx \sqrt{\frac{2}{\pi n}} \sqrt{I^{dd}(\rho|\mathcal{S})^{-1}}, \quad (5.32)$$

where the Fisher information I^{dd} corresponding to the diagonal parameter is found in the previous section. From Table 5.1 and the discussion in the previous section, we know that in the limit $\lambda_2 \rightarrow 0$ the mean Fisher information \bar{I}^{dd} diverges. The Fisher information $I(\rho|\mathcal{S})$ for any finite sample of random measurements will therefore not concentrate within $1 \pm \epsilon$ of the optimal, implying a lack of concentration in the MINF. In the case of the Frobenius norm, in the limit $\lambda_2 \rightarrow 0$ the dominant error terms contributing to the MSE correspond to the rotation parameters, and this fact ensures a concentration of the MSE even in the pure state limit, while for the infidelity we see that the dominant error terms comes from the estimation of the small eigenvalues, and a concentration of the MINF does not exist in the sense of Lemma 2. In general, the local expansion of the infidelity around any rank- r state that is close to pure is linear in the diagonal terms of the estimate [89]. The MINF for such states therefore demonstrates a similar lack of concentration in the corresponding diagonal elements of the Fisher information matrix.

Furthermore, from (5.32) it is clear that with uniform random measurements the MINF scales as $O(1/\sqrt{n})$ for states that are close to pure, while for states

well within the Bloch sphere, Lemma 3 demonstrates a scaling of $O(1/n)$. This poor scaling is observed in Figure 5.3, which plots the expected error in terms of the MINF and the MSE for pure states. As discussed in the previous sections, it is seen that the MSE scales as $O(1/n)$ for all states, while it is clear from Figure 5.3 that the MINF demonstrates a $O(1/\sqrt{n})$ scaling for pure states. This scaling has also been demonstrated for the closely related Bures distance error metric. In [78, 117], the minimax Bures error for estimators based on Pauli expectations is shown to scale as $O(1/\sqrt{n})$. This poor scaling along with a lack of concentration is important as many quantum information tasks utilise states that are pure [89]. Several adaptive measurement protocols have been suggested and implemented [89, 96, 79, 52] to improve this scaling. The aim of such adaptive strategies is to make measurements that are close to the eigenbasis of the true state. In our qubit model, for measurements with angle ϕ smaller than $O(1/\sqrt{n})$ the Fisher information $I^{dd}(\rho)$ scales as $O(n)$. From (5.32), this gives a $O(1/n)$ scaling of the infidelity even in the limit $\lambda_2 \rightarrow 0$.

5.5 Conclusions

In this chapter we investigated the asymptotic behaviour of the error for an arbitrary optimal estimator in the random measurement setup. Specifically we looked at how the accuracy of efficient estimators depends on the measurement design and the state. We considered two distance measures, the Frobenius norm and the quantum infidelity. In the case of the Frobenius norm, we extended the concentration results in chapter 4, and demonstrated that the MSE attains the optimal rate (up to a constant) with only $O(r \log D)$ random basis measurements for all states of rank r . Furthermore, to investigate the behaviour of the MSE concentration for states that are close to pure, we considered the model of a single qubit. We presented an argument to show that concentration in the MSE occurs for all qubit states, despite a lack of concentration in the Fisher information matrix for states close to the surface of the Bloch sphere.

It remains an open problem if a similar scaling of the MSE exists in the Pauli measurement setup used in standard multiple ions tomography. The application of the tools in the chapter to the Pauli setup requires control of the eigenvalues in equation (5.9), specifically a lower bound on the minimum eigenvalue $\lambda_{\min}(\bar{I})$. Strong numerical evidence in chapter 4 suggests that for random measurements the Fisher information may satisfy the required spectral

properties. Concentration results for distances other than the Frobenius norm can be in principle derived using similar arguments as long as their local expansions are quadratic in the parameters (see (5.7)). However, for the infidelity (an important measure of error for quantum tomography), it is known that while the scaling is quadratic for states deep in the Bloch sphere, for states close to pure this scaling is linear in the parameters [89]. We demonstrated with a single qubit model that for such nearly pure states and random measurements, the mean infidelity (MINF) does not concentrate around the optimal for any finite number of settings. This implies a lack of ‘compressive’ recovery of such low rank states, and therefore by increasing the number of measurement settings one can always significantly decrease the corresponding estimation error.

The FIM has been an important tool in our investigation of both the MINF and the MSE. We noticed that the FIM fails to concentrate when one of the eigenvalues of the state approaches zero. Related work in establishing and using continuity relations of the Quantum Fisher information (QFI) [94, 4] also shows a dependence on the smallest eigenvalues of the state, and therefore interesting behaviour occurs when eigenvalues approach zero. It is a possible direction for future research to see if our results about the concentration failure of the FIM are more deeply connected to the work in [94, 4].

5.6 Appendix

5.6.1 Proof of Lemma 1

Lemma. *For any rank- r state ρ with an arbitrary spectrum, and the rank- r state ρ_0 which has equal non-zero eigenvalues $1/r$ and the same eigenvectors as ρ , the following inequality holds between their average Fisher information matrices, evaluated over all possible random measurement settings.*

$$\bar{I}(\rho_0) \leq \bar{I}(\rho) \quad (5.33)$$

Proof. For a given random measurement setting \mathbf{s} , the probabilities of occurrence of an outcome \mathbf{o} for the two states ρ_0 and ρ are given by

$$p_0(\mathbf{o}|\mathbf{s}) = \sum_{i=1}^r \frac{1}{r} |\langle e_{\mathbf{s}}^{\mathbf{o}} | \lambda_i \rangle|^2 \quad ; \quad p_{\rho}(\mathbf{o}|\mathbf{s}) = \sum_{i=1}^r \lambda_i |\langle e_{\mathbf{s}}^{\mathbf{o}} | \lambda_i \rangle|^2 \quad (5.34)$$

where λ_i and $|\lambda_i\rangle$ are the eigenvalues and the eigenvectors of the state ρ respectively. We now consider states ρ' , that are constructed by permuting the r non-zero eigenvalues λ_i of the state ρ , while keeping the eigenvectors fixed. Let \mathcal{P} denote the set of $r!$ such permuted states. Averaging the probabilities $p_{\rho'}(\mathbf{o}|\mathbf{s})$ over all the permuted states $\rho' \in \mathcal{P}$ recovers the probability $p_0(\mathbf{o}|\mathbf{s})$ corresponding to the state with the uniform spectrum. That is,

$$\frac{1}{|\mathcal{P}|} \sum_{\rho' \in \mathcal{P}} p_{\rho'}(\mathbf{o}|\mathbf{s}) = p_0(\mathbf{o}|\mathbf{s}). \quad (5.35)$$

From the convexity of the function $f(x) = 1/x$ in the interval $(0, +\infty)$, the above equation together with Jensen's inequality implies,

$$\frac{1}{p_0(\mathbf{o}|\mathbf{s})} \leq \frac{1}{|\mathcal{P}|} \sum_{\rho' \in \mathcal{P}} \frac{1}{p_{\rho'}(\mathbf{o}|\mathbf{s})} \quad (5.36)$$

where we assumed that $p_{\rho'}(\mathbf{o}|\mathbf{s}) > 0$ for all \mathbf{o} . From 5.5, we see that for a setting \mathbf{s} , the Fisher matrix in our parametrisation can be written as a sum of d matrices

$$I(\rho'|\mathbf{s}) = \sum_{\mathbf{o}: p_{\rho'}(\mathbf{o}|\mathbf{s}) > 0} \frac{1}{p_{\rho'}(\mathbf{o}|\mathbf{s})} |V_{\mathbf{s}}^{\mathbf{o}}\rangle \langle V_{\mathbf{s}}^{\mathbf{o}}| \quad (5.37)$$

where $|V_{\mathbf{s}}^{\mathbf{o}}\rangle \in \mathbb{R}^D$, with $D = 2rd - r^2 - 1$, are vectors that depend only on the measurement vectors $|e_{\mathbf{s}}^{\mathbf{o}}\rangle$, and the eigenvectors $|\lambda_i\rangle$ of the state. Since by

construction the eigenvectors for all the states considered above are the same, together with 5.36, we get for all settings \mathbf{s}

$$\begin{aligned} I(\rho_0|\mathbf{s}) &= \sum_{\mathbf{o}: p_0(\mathbf{o}|\mathbf{s}) > 0} \frac{1}{p_0(\mathbf{o}|\mathbf{s})} |V_{\mathbf{s}}^{\mathbf{o}}\rangle \langle V_{\mathbf{s}}^{\mathbf{o}}| \leq \frac{1}{|\mathcal{P}|} \sum_{\mathbf{o}} \sum_{\rho' \in \mathcal{P}} \frac{1}{p_{\rho'}(\mathbf{o}|\mathbf{s})} |V_{\mathbf{s}}^{\mathbf{o}}\rangle \langle V_{\mathbf{s}}^{\mathbf{o}}| \\ &= \frac{1}{|\mathcal{P}|} \sum_{\rho' \in \mathcal{P}} I(\rho'|\mathbf{s}). \end{aligned} \quad (5.38)$$

The inequality holds for settings \mathbf{s} such that $p_{\rho'}(\mathbf{o}|\mathbf{s}) > 0$ for all “permuted” states ρ' and all outcomes \mathbf{o} , which holds with probability one under the Haar measure over settings. Since each ρ' is an unitary rotation of the state ρ , we arrive at the required inequality of the average Fisher matrices by integrating both sides of the above equation over all possible random measurement settings \mathbf{s} . \square

5.6.2 Proof of Theorem 4

The proof of this theorem is similar to the one presented in chapter 4. Here we present the important elements of the proof, and refer to appendix in chapter 4 for details. As briefly mentioned in the main text of the chapter, the proof of the theorem utilises the following matrix Chernoff bound [2], where the random matrices X_i are given by $G_F^{-1/2} I(\rho|\mathbf{s}_i) G_F^{-1/2}$, with \mathbf{s}_i random bases.

Theorem 6. (Matrix Chernoff Bound) Consider a finite sequence X_1, \dots, X_k of independent, random, positive matrices with dimension D , such that $\lambda_{\max}(X) \leq R$. For $\mathbb{E}X = M \geq \mu \mathbf{1}$ and $0 \leq \epsilon \leq \frac{1}{2}$,

$$\mathbb{P} \left\{ \frac{1}{k} \sum_{i=1}^k X_i \notin \left[(1-\epsilon)M, (1+\epsilon)M \right] \right\} \leq 2D \cdot \exp \left(-k \cdot \frac{\epsilon^2 \mu}{2R \cdot \log 2} \right) \quad (5.39)$$

We note that $G_F^{-1/2} I_S G_F^{-1/2}$ is a sum of k independent, random, positive matrices. In order to apply the above bound, we need to upper bound the largest eigenvalue of $G_F^{-1/2} I(\rho|\mathbf{s}) G_F^{-1/2}$ over all measurements, denoted μ_{\max} . We also need to lower bound the smallest eigenvalue of the expected Fisher information $G_F^{-1/2} \bar{I}(\rho) G_F^{-1/2}$, denoted μ_{\min} . We will first derive these bounds and then obtain the result by applying the Chernoff bound. As in the text, we

work with the local parametrisation

$$\boldsymbol{\theta} = (\boldsymbol{\theta}^{(d)}, \boldsymbol{\theta}^{(r)}, \boldsymbol{\theta}^{(i)}) = (\rho_{2,2}, \dots, \rho_{r,r}; \operatorname{Re}\rho_{1,2}, \dots, \operatorname{Re}\rho_{r,d}; \operatorname{Im}\rho_{1,2}, \dots, \operatorname{Im}\rho_{r,d}) \quad (5.40)$$

where $\rho_{1,1}$ is constrained to enforce the trace-one normalisation. The Fisher information therefore, has the following block structure

$$I(\rho) = \begin{pmatrix} I^{dd}(\rho) & I^{dr}(\rho) & I^{di}(\rho) \\ I^{rd}(\rho) & I^{rr}(\rho) & I^{ri}(\rho) \\ I^{id}(\rho) & I^{ir}(\rho) & I^{ii}(\rho) \end{pmatrix} \quad (5.41)$$

with the superscripts identifying the parameters considered; diagonal, real and imaginary. The weight matrix G_F also has the same block structure with elements

$$G_{F_{a,b}} = \operatorname{Tr} \left[\frac{\partial \rho_{\boldsymbol{\theta}}}{\partial \theta_a} \cdot \frac{\partial \rho_{\boldsymbol{\theta}}}{\partial \theta_b} \right] \quad (5.42)$$

In the parametrisation described above, the weight matrix G_F has the following block diagonal form:

1. The *diagonal-diagonal* block:

$$(a) \quad G_{F_{a,b}}^{dd} = 1 + \delta_{a,b}$$

2. The *real-real* and *imaginary-imaginary* block:

$$(a) \quad G_{F_{a,b}}^{rr/ii} = 2 \cdot \delta_{a,b}$$

with the other blocks being zero. We note that both the Fisher, and the weight matrix are of dimension $D := 2rd - r^2 - 1$.

Lower bound on the smallest eigenvalue—As mentioned in the main text, we use Lemma 1 to bound the the smallest eigenvalue from below as

$$G_F^{-1/2} \bar{I}(\rho_0) G_F^{-1/2} \leq G_F^{-1/2} \bar{I}(\rho) G_F^{-1/2}, \quad (5.43)$$

where ρ_0 is the state with r equal eigenvalues and the same eigenvectors as the state ρ . The explicit form of $\bar{I}(\rho_0)$ is known, and has been evaluated in chapter 4, and from it, we see that the minimum eigenvalue is lower bounded as $\mu_{\min} \geq r/r + 1$ for $r > 1$ and $\mu_{\min} \geq 1$ for pure states.

Upper bound on the largest eigenvalue—We use the inequality $I(\rho|\mathbf{s}) \leq F(\rho)$ between the classical and quantum Fisher informations to bound the largest eigenvalue of $G_F^{-1/2}I(\rho|\mathbf{s})G_F^{-1/2}$ over all measurements by the largest eigenvalue of $G_F^{-1/2}F(\rho)G_F^{-1/2}$. The quantum Fisher information is calculated in the local parameterisation described above and evaluated at the state $\rho = \text{Diag}(\lambda_1, \dots, \lambda_r, \dots, 0)$, diagonal in its eigenbasis. The details of this calculation can be found in the appendix of chapter 4, and we therefore avoid the repetition and merely state the elements of the matrix. Denoting r_a, c_a to be the row and column positions of the element a of the parameter vector $\boldsymbol{\theta}$, we have

1. For the **Diagonal-Diagonal** block with $r > 1$,

$$\begin{aligned} \text{(a)} \quad F_{a,a}^{dd} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} &= \frac{1}{\lambda_{r_a}} + \frac{1}{\lambda_1} \text{ when } r_a \leq r \\ \text{(b)} \quad F_{a,b}^{dd} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} &= \frac{1}{\lambda_1} \text{ when } r_a, r_b \leq r, \text{ and } a \neq b \end{aligned}$$

2. For the **Real-Real** and **Imaginary-Imaginary** blocks:

$$\begin{aligned} \text{(a)} \quad F_{a,a}^{rr/ii} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} &= \frac{4}{\lambda_{r_a} + \lambda_{c_a}} \text{ when } r_a < c_a \leq r \\ \text{(b)} \quad F_{a,a}^{rr/ii} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} &= \frac{4}{\lambda_{r_a}} \text{ when } r_a \leq r, c_a > r \end{aligned}$$

The off-diagonal blocks are zero. It is easy to see that the quantum Fisher matrix is upper bounded by the matrix $\frac{1}{\lambda_{\min}(\rho)}G_F^{dd} \oplus \frac{2}{\lambda_{\min}(\rho)}G_F^{rr} \oplus \frac{2}{\lambda_{\min}(\rho)}G_F^{ii}$. So we can write

$$G_F^{-1/2}F G_F^{-1/2} \leq \frac{1}{\lambda_{\min}(\rho)}\mathbf{1}_{(r-1)} \oplus \frac{2}{\lambda_{\min}(\rho)}\mathbf{1}_{(2rd-r^2+r)} \oplus \frac{2}{\lambda_{\min}(\rho)}\mathbf{1}_{(2rd-r^2+r)} \quad (5.44)$$

The maximum eigenvalue μ_{\max} is therefore upper bounded by $2/\lambda_{\min}(\rho)$ for $r > 1$, and 2 for $r = 1$.

Combining the bounds to prove concentration— We can now substitute these values into the matrix Chernoff bound. While the value of the minimum/maximum eigenvalues differ for $r > 1$ and $r = 1$, we calculate the bound for the case when $r > 1$, as this will provide a general bound for the number of settings required that holds even in the case of pure states. Writing $P_S = G_F^{-1/2}I_S G_F^{-1/2}$ and $\bar{P} = G_F^{-1/2}\bar{I}G_F^{-1/2}$ for notational simplicity, we have for $r > 1$

$$\mathbb{P}\left\{P_S \notin \left[(1-\epsilon)\bar{P}, (1+\epsilon)\bar{P}\right]\right\} \leq 2D \cdot \exp\left(-k \frac{r\epsilon^2 \lambda_{\min}(\rho)}{4(r+1) \cdot \log 2}\right) := \delta \quad (5.45)$$

Therefore, with probability $1 - \delta$ we have that

$$(1 - \epsilon)\bar{P} \leq P_S \leq (1 + \epsilon)\bar{P} \quad (5.46)$$

This can be re-written in the form of inequalities of the MSE with $\epsilon > 0$ sufficiently small

$$(1 - \epsilon)\text{Tr}(\bar{P}^{-1}) \leq \text{Tr}[P_S^{-1}] \leq (1 + \epsilon)\text{Tr}(\bar{P}^{-1}) \quad (5.47)$$

For a fixed value of ϵ and δ , we see that the minimum number of settings k required for the above bound to hold with probability greater than $1 - \delta$ is

$$k = \frac{C_1}{\lambda_{\min}(\rho)} \cdot \frac{(r+1)}{r} \log\left(\frac{2D}{\delta}\right) \quad (5.48)$$

where $C_1 := 4(\log 2/\epsilon^2)$ and $D := 2rd - r^2 - 1$.

□

Chapter 6

Minimax estimation of qubit states with Bures risk

6.1 Introduction

The aim of state tomography is the estimation of an unknown density matrix ρ from the outcomes of measurements performed on n identical copies of the state available as a resource. The quality of the resulting estimate $\hat{\rho}_n$ is quantified in terms of its average error, or risk. Given a measurement design M , and the corresponding set of outcomes X , the risk of the measurement-estimator pair is

$$R(\rho, \hat{\rho}_n) := \mathbb{E}[D(\hat{\rho}_n(X), \rho)], \quad (6.1)$$

where the expectation is taken with respect to the measurement outcomes X , given the unknown state ρ . The risk depends on the choice of the error, or loss function $D(\hat{\rho}_n, \rho)$, which is a measure of the deviation of the estimated state from the true state ρ . Examples of commonly used loss functions are square Frobenius (norm-two) distance, infidelity, the trace-norm distance, and the Bures distance. The risk is a function of the resource size n , and one is interested in its behaviour in the limit of large n . Typically, for a ‘good’ estimator and particular choices of loss functions, (e.g. locally quadratic functions) the optimal risk exhibits a rate of $O(1/n)$ *uniformly* over all states ρ .

However, as pointed out in chapter 5 for certain loss functions (e.g. square Bures distance, or infidelity), the risk is known to behave differently for states of unknown purity [89]. This is readily illustrated in the qubit case, where the fidelity between the state ρ and its estimate $\hat{\rho}_n$ is defined as $F(\rho, \hat{\rho}_n) := [\text{Tr}(\sqrt{\sqrt{\rho}\hat{\rho}_n\sqrt{\rho}})]^2$ and can be expressed as

$$F(\rho, \hat{\rho}_n) := \frac{1}{2}(1 + \sqrt{1 - |r|}\sqrt{1 - |\hat{r}|} + \mathbf{r} \cdot \hat{\mathbf{r}}), \quad (6.2)$$

where $\mathbf{r}, \hat{\mathbf{r}} \in \mathbb{R}^3$ are the Bloch vectors of the two states. For ρ within the Bloch sphere, the fidelity is locally quadratic in the components $\delta_i := r_i - \hat{r}_i$, with $i = x, y, z$. However, for states close to the boundary of the Bloch sphere, the fidelity becomes linear in $|\delta_i|$. Standard tomographic estimation of the Bloch vector components by measuring the spin operators $\sigma_x, \sigma_y, \sigma_z$ gives an accuracy of the order of $n^{-1/2}$ in estimating δ_i . This implies that for a loss function such as the infidelity $1 - F(\hat{\rho}_n, \rho)$, or the square Bures distance $D_B(\rho, \hat{\rho}_n)^2 := 2(1 - \sqrt{F(\rho, \hat{\rho}_n)})$, the risk scales as $O(1/n)$ for states within the Bloch sphere, but only as $O(1/\sqrt{n})$ for nearly pure states.

This poor scaling for nearly pure states is significant as the preparation and estimation of pure states is ubiquitous in quantum information processing tasks. Although many papers discuss the issues surrounding quantum tomography for Bures risk, we consider the problem in the context of minimax estimation (see also [42, 43]), i.e. where the figure of merit is the *maximum risk* over all states

$$R_{max}(\hat{\rho}_n) := \sup_{\rho} R(\rho, \hat{\rho}_n) = \sup_{\rho} \mathbb{E} \left[D_B(\rho, \hat{\rho}_n)^2 \right]. \quad (6.3)$$

Our aim is to show that *adaptive, separable* measurement strategies can achieve the n^{-1} scaling of the maximum risk. We also consider collective measurements, and derive an upper bound to the asymptotic constant of the maximum risk. Our analysis shows that the problem of finding *minimax* estimators, i.e. estimators with smallest possible asymptotic constant, reduces to that of finding minimax estimator for the ‘classical’ problem of estimating the coin toss probability with respect to the square Hellinger distance risk.

Several estimation methods have been proposed in the literature, involving both global and local measurement strategies, with the aim of improving the poor scaling of fidelity based risks for nearly pure states. However, to the best of our knowledge a scaling of n^{-1} of the maximum risk has not been demonstrated for any of these estimators. Two-steps adaptive quantum tomography [49, 100, 9, 89, 70] involves using a fraction n_1 of the available resource n to obtain a preliminary estimate of the eigenbasis of ρ , and then performing measurements along the estimated eigenbasis on the remaining $n - n_1$ copies of the state. In [9] it is shown that using a vanishing fraction $n_1 = n^\alpha$, with $1/2 < \alpha < 1$ for the preliminary estimate gives a rate of $O(1/n)$ for the average infidelity with respect to certain distributions over states. However, it has been pointed out in [89], that for certain states a vanishing fraction is insufficient, and that for almost pure states the worst case infidelity scales

as $O(n^{-5/6})$. Numerical results in [89] suggest that using a fixed fraction $n_0 = \beta n$ instead gives the $O(1/n)$ scaling for nearly pure states. The two step adaptive protocols have been experimentally implemented [89, 70], showing a quadratic improvement in scaling for nearly pure states. The extension of the two-step adaptive protocol to a fully adaptive one has been considered in [100, 93, 46, 41], where the measurement basis is aligned according to a current estimate after every measurement step. In the Bayesian framework, ‘self-learning’ measurement protocols have been considered in [44, 52, 63, 79]. A detailed review of various adaptive protocols and experimental results is found in [105].

Protocols considering collective (or joint rather than separable) measurements have also been considered [8, 7, 57, 59]. It is known that joint measurements perform better than separable measurements in the case of mixed states [7]. In a Bayesian framework, [8] showed that with certain optimal joint measurements, the asymptotic infidelity averaged over a prior distribution achieved a value of $\frac{3+2\langle r \rangle}{4n}$ for mixed qubit states, where $\langle r \rangle$ is the mean purity over the prior distribution. Work in [57] proposes a two-step adaptive estimation strategy that is shown to be *locally* optimal, achieving an infidelity risk of $\frac{1+4\lambda_{\max}(\rho)}{4n}$ for mixed qubit states, where $\lambda_{\max}(\rho)$ is the maximum eigenvalue of ρ . However these theoretical results cannot be directly used to derive the n^{-1} scaling of the minimax risk.

We propose two different estimators, one based on adaptive local measurements similar to [9, 89], and a second based on global collective measurements and *Local Asymptotic Normality* (LAN) as in [57, 59]. In terms of local measurements, we consider a two step adaptive strategy much in line with already proposed estimators [49, 100, 9, 89, 70]. A *fixed fraction* of the total sample size n is used to obtain a preliminary estimate $\tilde{\mathbf{r}}$ of the Bloch vector \mathbf{r} of the state, by performing standard tomographic measurements of the spin observables. The remaining copies of the state are used to estimate the eigenvalues of the state by performing measurements along the estimated direction $\tilde{\mathbf{r}}/|\tilde{\mathbf{r}}|$. The final estimate $\hat{\rho}_n$ of the state is then constructed from the estimated eigenvalues in the adaptive step and the preliminary estimate. For this estimator, we upper-bound the maximum Bures risk 6.3 and demonstrate a scaling of n^{-1} .

The estimator based on global collective measurements uses established LAN results for qubit states [57, 59]. The measurement strategy involves two stages. The first stage involves the standard tomographic measurements of the spin components on a *vanishing number* of copies of the state $\tilde{n} \ll n$. A preliminary

estimate $\tilde{\rho}$ is constructed from the outcomes. The second measurement stage depends on this preliminary estimate, and for technical purposes related to the asymptotic analysis we consider the following two cases: $|\tilde{\mathbf{r}}| < \delta$ and $|\tilde{\mathbf{r}}| \geq \delta$ for some small constant $\delta > 0$.

When $\tilde{\rho}$ is close to the fully mixed state, the standard tomographic measurements are performed on the remaining copies of the state. When the preliminary estimate $\tilde{\rho}$ is away from the fully mixed state, a joint measurement is performed on the remaining copies of the state. The joint state $\rho_n^\theta := \rho_\theta^{\otimes n}$ has a block-diagonal form following the Weyl decomposition of the underlying space $(\mathbb{C}^2)^{\otimes n}$. Information about the eigenvalue parameter λ is encoded in a probability distribution over the different blocks of the decomposition, while information about the local parameters (u, v) is encoded in the block states. We consider a parameterisation of states ρ_θ , with $\theta = (\lambda, u, v) \in \mathbb{R}^3$, where λ parametrises the smallest eigenvalue of the states and $\mathbf{w} = (u, v)$ are certain local rotation parameters around a fixed state ρ_0 . The parameter λ is then estimated from the outcomes of a ‘‘which-block’’ measurement, while the local parameters u, v are optimally estimated by exploiting the local asymptotic normality of the block states. The LAN results in [57] establish that in the limit $n \rightarrow \infty$, the block states converge to Gaussian states $\phi^\mathbf{w}$, with displacement proportional to parameters (u, v) (Theorem 7). The optimal estimator of u, v is then the optimal estimator of the displacement of a Gaussian state $\phi^\mathbf{w}$, which is known to be the heterodyne measurement. We derive minimax upper and lower bounds for the risk (6.3), and demonstrate that the maximum Bures risk for the estimator $\hat{\rho}_n$ scales as C/n , with C being a constant. We obtain lower bound of $5/4$ for this constant, and an upper bound of $3/2$.

An important element in the derivation of the upper-bounds for the maximum risk of both estimators is the fact that square Bures distance can locally approximated as a sum of contributions from the eigenvalue parameter and the ‘rotation’ parameters. More explicitly, we have that

$$D_B(\rho, \hat{\rho}_n)^2 \approx D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}})^2 + \frac{1}{4} \frac{(1-2\lambda)(1-2\hat{\lambda})}{\sqrt{(1-\lambda)(1-\hat{\lambda})} + \sqrt{\lambda\hat{\lambda}}} \Phi^2, \quad (6.4)$$

where $D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}})^2 := \|\sqrt{\boldsymbol{\lambda}} - \sqrt{\hat{\boldsymbol{\lambda}}}\|^2$ with $\boldsymbol{\lambda} = (\lambda, 1-\lambda)$ is the classical square Hellinger distance, and Φ is the angle between the Bloch vectors of the two states. The optimal rate of estimation of the ‘rotation’ parameters is easily shown to be $1/n$ for all states. The problem of establishing minimax results for

the square Bures distance therefore converts a problem of establishing minimax results for the Hellinger risk of estimating the eigenvalue parameter. To the best of our knowledge such minimax results for the Hellinger risk are not known. Instead, we upper-bound this risk by the Kullback-Leibler risk, and use known results about the minimax estimator in this case. However, in section 6.5 we propose that a minimax optimal estimator of the classical parameter λ under the Hellinger risk gives a minimax optimal estimator $\hat{\rho}_n$ for qubit states.

The chapter is organised as follows. In section 6.2 we consider an estimator based on local measurements and detail a two step adaptive measurement strategy. We demonstrate that the proposed estimator achieves a minimax rate of $1/n$. In section 6.3 we propose a second estimator based on global collective measurements. We begin in section 6.3.1 by describing the preliminary measurement stage and introduce our parametrisation of states. In section 6.3.2, we describe the second measurement stage, and overview the block decomposition of the joint state and results of LAN. The minimax bounds for this estimator are derived in section 6.4, and in section 6.5 we discuss and state the proposition that a minimax estimator for the Hellinger loss function implies a minimax optimal estimator for qubit states. Finally in section 6.6 we consider the quantum relative entropy and establish that the minimax rate under this loss function scales as $O(n^{-1} \log n)$.

6.2 Estimator based on local adaptive measurements

We let ρ be an arbitrary density matrix associated with a single qubit state. Given n identical copies of the state as a resource, we wish to construct an estimator of the state. As briefly discussed in the introduction, in this section we propose a two-step adaptive measurement strategy based on local measurements. While the idea of an adaptive local measurement strategy is not new and has been treated in various instances in the literature [49, 100, 9, 89, 70], we are interested in analysing the performance of the proposed estimator $\hat{\rho}_n$ in terms of the *maximum risk* with respect to the *square Bures distance* defined in equation (6.3) and its asymptotic rescaled version

$$r_{\max}(\hat{\rho}) = \lim_{n \rightarrow \infty} \sup_{\rho} \sup_{\rho} n \mathbb{E} \left[D_B(\rho, \hat{\rho})^2 \right] = \lim_{n \rightarrow \infty} \sup n R_{\max}(\hat{\rho}_n). \quad (6.5)$$

We will derive an upper bound for the latter risk, thereby demonstrating a n^{-1} scaling for maximum risk over all states ρ . Since the maximum risk of

any estimation procedure cannot scale faster than n^{-1} , this implies that the existence of a non-trivial scaling constant for the *minimax risk* given by

$$R_{\min\max} := \limsup_{n \rightarrow \infty} \inf_{\hat{\rho}_n} n R_{\max}(\hat{\rho}_n). \quad (6.6)$$

Finding the value of the minimax constant remains an open problem. We will come back to this problem in section 6.5 where it is shown that the minimax qubit estimation problem reduces to that of minimax estimation of a coin probability with respect to the square Hellinger distance risk.

The estimator we propose is constructed as follows. The first stage is a preliminary localisation step involving standard projective measurements of Pauli observables $\sigma_x, \sigma_y, \sigma_z$ on a fixed fraction n_1 of the total number of qubits n . An estimate of the direction vector $\tilde{\mathbf{r}}/|\tilde{\mathbf{r}}|$ is constructed from the outcomes of these measurements. The following lemma shows that with high probability the estimated directional vector is within an angle of $O(n_1^{-1/2+\epsilon_1})$ of the true vector, where ϵ_1 is a fixed (small) positive constant.

Lemma 4. *Let X_i, Y_i, Z_i be the outcomes of measurements of $\sigma_x, \sigma_y, \sigma_z$ performed on independent qubits in state ρ with Bloch vector \mathbf{r} , where $i = 1, \dots, n_1/3$. Let $\tilde{\mathbf{r}}$ be the estimate of the Bloch vector, where each Bloch vector component is obtained by averaging the outcome results, e.g. $\tilde{r}_x := \frac{3}{n_1} \sum_i X_i$. Then we have that for $\epsilon_1 > 0$,*

$$\mathbb{P}\left(\|\mathbf{r} - \tilde{\mathbf{r}}\|_2^2 > 6n_1^{-1+2\epsilon_1}\right) \leq 6 \exp\left(-\frac{2n_1^{2\epsilon_1}}{3}\right). \quad (6.7)$$

The proof of this lemma follows directly from Hoeffding's inequality applied to the binomial distribution corresponding to each of the Bloch vector components. The concentration inequality implies that when $|\mathbf{r}|$ is bounded away from zero, the magnitude of the angle Φ between the directional vectors $\mathbf{r}/|\mathbf{r}|$ and $\tilde{\mathbf{r}}/|\tilde{\mathbf{r}}|$ is of the order $O(n_1^{-1/2+\epsilon_1})$ with high probability.

The second adaptive stage involves performing measurements along this estimated direction. That is, projective measurements of the observable $\Xi := \vec{\sigma} \cdot \tilde{\mathbf{r}}/|\tilde{\mathbf{r}}|$ are performed on the remaining $n_2 := n - n_1$ copies of the state. Let k be the total number of +1 outcomes from these measurements. It is easy to see that k is distributed binomially $B_{n_2, p}(k)$ with binomial parameter $p := (1 + |\mathbf{r}| \cos \Phi)/2$. We estimate this parameter as \hat{p} from the measurement

outcomes using the ‘add-beta’ estimator [22, 23] defined as follows,

$$\hat{p}_{n_2} = \begin{cases} \frac{1/2}{n_2+5/4}, & k = 0, \\ \frac{2}{n_2+7/4}, & k = 1, \\ \frac{k+3/4}{n_2+3/2}, & k = 2, \dots, n_2 - 2, \\ \frac{n_2-1/4}{n_2+7/4}, & k = n_2 - 1, \\ \frac{n_2+3/4}{n_2+5/4}, & k = n_2. \end{cases} \quad (6.8)$$

While this estimator is not in any sense optimal, it is known to be the minimax estimator for the Kulback-Leibler risk, which will be used below in deriving the upper bound for qubit tomography.

The final estimate of the state puts together the estimate \hat{p} and the estimated Bloch vector $\tilde{\mathbf{r}}$ as follows

$$\hat{\rho}_n = \frac{1}{2} \left(I + \frac{2\hat{p}_{n_2} - 1}{|\tilde{\mathbf{r}}|} \tilde{\mathbf{r}} \cdot \vec{\sigma} \right). \quad (6.9)$$

It is easy to see that $\hat{p}_{n_2} = \hat{\lambda}$ by construction, where $\hat{\lambda}$ is the eigenvalue of the estimate $\hat{\rho}_n$.

6.2.1 An n^{-1} scaling upper bound scaling

We now look at deriving an upper bound for the Bures risk of this estimator, and demonstrate that the maximum over all states scales as n^{-1} . Recall that the square Bures distance between the final estimate $\hat{\rho}_n$ and the true state ρ is defined as $D_B(\rho, \hat{\rho}_n)^2 := 2 \left[1 - \sqrt{F(\rho, \hat{\rho}_n)} \right]$, where $F(\rho, \hat{\rho}_n)$ is the fidelity, expressed in terms of the Bloch vectors $\hat{\mathbf{r}}$ and \mathbf{r} as

$$\begin{aligned} F(\rho, \hat{\rho}_n) &= \frac{1}{2} \left(1 + \sqrt{1 - |\mathbf{r}|^2} \sqrt{1 - |\hat{\mathbf{r}}|^2} + \mathbf{r} \cdot \hat{\mathbf{r}} \right) \\ &= \frac{1}{2} \left(1 + \sqrt{1 - |\mathbf{r}|^2} \sqrt{1 - |\hat{\mathbf{r}}|^2} + |\mathbf{r}| |\hat{\mathbf{r}}| \cos \Phi \right), \end{aligned} \quad (6.10)$$

where Φ is the angle between the Bloch vectors, or equivalently the angle between the vectors $\tilde{\mathbf{r}}$ and \mathbf{r} by construction. From Lemma 4, this angle is known to be small and of the order $O(n^{-1/2+\epsilon_1})$ with high probability. This implies that the cosine term in (6.10) can be expanded to leading order in Φ .

In this case, the square Bures distance is expressed as

$$D_B(\rho, \hat{\rho}_n)^2 = D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}})^2 + \frac{1}{4} \frac{(1-2\lambda)(1-2\hat{\lambda})}{\sqrt{(1-\lambda)(1-\hat{\lambda})} + \sqrt{\lambda\hat{\lambda}}} \Phi^2 + O(\Phi^4), \quad (6.11)$$

where $\lambda = (1 - |\mathbf{r}|)/2$, $\hat{\lambda} = (1 - |\hat{\mathbf{r}}|)/2$ are the smallest eigenvalues of ρ and $\hat{\rho}$, respectively, and

$$D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}})^2 := \left(\sqrt{\lambda} - \sqrt{\hat{\lambda}}\right)^2 + \left(\sqrt{1-\lambda} - \sqrt{1-\hat{\lambda}}\right)^2 \quad (6.12)$$

is the square Hellinger distance between the probability distributions $\boldsymbol{\lambda} = (\lambda, 1-\lambda)$ and $\hat{\boldsymbol{\lambda}} = (\hat{\lambda}, 1-\hat{\lambda})$. The proof of this approximation can be found in Appendix 6.8.2. Identifying $\hat{\boldsymbol{\lambda}} = (\hat{p}_{n_2}, 1-\hat{p}_{n_2})$, we upper bound the square Bures distance as

$$\begin{aligned} D_B(\rho, \hat{\rho}_n)^2 &\leq D_H(\boldsymbol{\lambda}, \hat{\mathbf{p}}_{n_2})^2 + \frac{1}{4} \Phi^2 + O(\Phi^4) \\ &\leq 2D_H(\boldsymbol{\lambda}, \mathbf{p})^2 + 2D_H(\mathbf{p}, \hat{\mathbf{p}}_{n_2})^2 + \frac{1}{4} \Phi^2 + O(\Phi^4), \end{aligned} \quad (6.13)$$

where the second inequality is established using the fact that the square Hellinger distance satisfies the triangle inequality. Using the inequality $D_H^2(\boldsymbol{\lambda}, \mathbf{p}) \leq 2|\lambda - p|$ and $p = (1 - |\mathbf{r}| \cos \Phi)/2$ we further upper bound the risk as

$$D_B(\rho, \hat{\rho}_n)^2 \leq 2D_H(\mathbf{p}, \hat{\mathbf{p}}_{n_2})^2 + \left(\frac{1}{4} + |\mathbf{r}|\right) \Phi^2 + O(\Phi^4). \quad (6.14)$$

Taking expectation with respect to the measurement outcomes given the true state ρ , we have

$$\mathbb{E} \left[D_B(\rho, \hat{\rho}_n)^2 \right] \leq 2\mathbb{E} \left[D_H(\mathbf{p}, \hat{\mathbf{p}}_{n_2})^2 \right] + \frac{5}{4} \mathbb{E} \left[\Phi^2 \right] + O(n_1^{-2+4\epsilon_1}). \quad (6.15)$$

The maximum risk of the estimator is therefore bounded from above as

$$\begin{aligned} \sup_{\rho} \mathbb{E} \left[D_B(\rho, \hat{\rho}_n)^2 \right] &\leq \sup_{\rho} 2\mathbb{E} \left[D_H(\mathbf{p}, \hat{\mathbf{p}}_{n_2})^2 \right] + O(n_1^{-1}) + O(n_1^{-2+4\epsilon_1}) \\ &\leq \sup_{\rho} 2\mathbb{E} \left[D_{KL}(\mathbf{p}, \hat{\mathbf{p}}_{n_2}) \right] + O(n_1^{-1}) + O(n_1^{-2+4\epsilon_1}) \end{aligned} \quad (6.16)$$

In the first inequality we upper bounded $\mathbb{E}[\Phi^2]$ as $O(n_1^{-1})$; this follows from the concentration inequality of Lemma 4. The second step employs the inequality between the square Hellinger distance and the Kullback-Leibler (KL) distance,

defined as $D_{KL}(\hat{\mathbf{p}}, \mathbf{p}) := \hat{p} \log \hat{p}/p + (1 - \hat{p}) \log (1 - \hat{p})/(1 - p)$. The reason for employing this inequality is that to the best of the authors' knowledge the asymptotic minimax optimal Hellinger risk of estimating the binomial parameter p is not known in the literature. A discussion related to the difficulty in obtaining the asymptotic minimax Hellinger risk is left to section 6.5. However, the minimax optimal rate for the KL loss function under the binomial distribution is known. The minimax optimal estimator is precisely the 'add beta' estimator defined in (6.8), and is known to achieve the asymptotic rate $\frac{1}{2n_2}(1 + o(1))$ [22, 23]. Choosing n_1 to be constant fraction of the total number of samples n establishes an overall rate of $O(1/n)$, for example by choosing $n_1 = n_2 = n/2$.

6.3 Irreducible representations, collective measurements and local asymptotic normality

In this section we propose a two stage estimator of the qubit state, which uses collective rather than separable measurements. The first stage, much like in the previous section, is a preliminary localisation stage. However, this stage does not use a fixed fraction of the total number of copies of the state, but a vanishing fraction \tilde{n} of the overall ensemble of identical qubits n . The second measurement stage involves performing joint measurements of the remaining $n - \tilde{n}$ copies of the state, and is based on the techniques of Local Asymptotic Normality (LAN) established in [57, 59]. This section is structured as follows. In subsection 6.3.1 we describe the preliminary measurement stage, and define a choice of parameterisation of states. In subsection 6.3.2 we describe the joint measurement strategy. In subsection 6.3.3 we provide a brief review of established LAN results and techniques which can be used to derive asymptotic estimation bounds. Minimax results for the proposed estimator are then established in section 6.4.

6.3.1 Preliminary localisation and parametrisation

The first stage is a preliminary localisation step that involves performing standard projective measurements of the Pauli observables $\sigma_x, \sigma_y, \sigma_z$ on a vanishing fraction \tilde{n} of the overall ensemble of n identically prepared qubits. An estimate $\tilde{\rho}$ of the state is constructed from the outcomes of these measurements. The following lemma shows that *with high probability* the true state ρ lies within a ball of radius $O(n^{-1/2+\epsilon_2})$ of this estimate $\tilde{\rho}$. This allows us to restrict our

attention to a local neighbourhood of the preliminary estimator in the second stage of the estimation.

Lemma 5. *Let X_i, Y_i, Z_i be independent outcomes of measurements of $\sigma_x, \sigma_y, \sigma_z$ performed on independent qubits in state ρ with Bloch vector \mathbf{r} , where $i = 1, \dots, \tilde{n}/3$. Let $\tilde{\rho}$ be the estimator with Bloch vector $\tilde{\mathbf{r}}$ obtained by averaging the outcome results, e.g. $\tilde{r}_x := \frac{3}{\tilde{n}} \sum_i X_i$. In order to obtain a physical state, the final estimate of the state is constructed as*

$$\tilde{\rho} := \arg \min_{\tau \in \mathbb{S}_2} \|\tau - (\mathbf{1} + \tilde{\mathbf{r}} \cdot \boldsymbol{\sigma})/2\|_1^2 \quad (6.17)$$

where the minimisation is over all the space of all 2×2 density matrices \mathbb{S}_2 . For this estimator $\tilde{\rho}$, we have that for all $\epsilon_2 > 0$,

$$\mathbb{P}\left(\|\tilde{\rho} - \rho\|_1^2 > 3n^{2\epsilon_2 - 1}\right) \leq 6 \exp\left(-\frac{2\tilde{n}n^{2\epsilon_2 - 1}}{3}\right), \quad \text{for all } \rho \in \mathbb{S}_2. \quad (6.18)$$

The proof of this lemma follows from an application of Hoeffding's inequality, and can be found in Appendix 6.8.1. Setting $\tilde{n} = n^{1-\kappa}$, with $0 < \kappa < 2\epsilon_2$, the probability of failure is exponentially small. The preliminary measurement stage therefore places the estimate $\tilde{\rho}$ in a local neighbourhood around the true state. Since $\|\tilde{\rho} - \rho\|_1^2 = \|\tilde{\mathbf{r}} - \mathbf{r}\|^2$, the angle between the two normalised Bloch vectors $\mathbf{r}/|\mathbf{r}|$ and $\tilde{\mathbf{r}}/|\tilde{\mathbf{r}}|$ is of the order $O(n^{-1/2+\epsilon_2})$ with high probability.

For technical reasons related to local asymptotic normality theory and the derivation of certain error bounds, the subsequent measurement depends on $\tilde{\rho}$, and we distinguish the following two cases.

i) If $\tilde{\rho}$ is within a fixed but small ball of radius $\delta > 0$ around the fully mixed state (i.e., $|\tilde{\mathbf{r}}| \leq \delta$), the secondary measurement stage consists of the standard tomographic measurements in the σ_i , $i = x, y, z$ bases. For each i , measurements of σ_i are performed on $(n - \tilde{n})/3$ identical copies of the state. The final estimate of the state $\hat{\rho}_n$ is constructed from the outcomes of these measurements, and is detailed in section 6.4.

ii) If $\tilde{\rho}$ is away from the fully mixed state, we can apply the tools of LAN. The remaining $n - \tilde{n}$ copies of the state available for the second stage are rotated such that the estimated Bloch vector $\tilde{\mathbf{r}}$ is pointing along the z -axis. From Lemma 5, the angle between the directional vectors $\mathbf{r}/|\mathbf{r}|$ and $\tilde{\mathbf{r}}/|\tilde{\mathbf{r}}|$ is known to be of the order $O(n^{-1/2+\epsilon_2})$ with high probability. This allows us to consider a restricted parametrisation of states which we describe now for an

arbitrary but fixed state ρ_0 (which plays the role of $\tilde{\rho}$) with its Bloch vector along the z -axis

$$\rho_0 = \begin{pmatrix} 1 - \lambda_0 & 0 \\ 0 & \lambda_0 \end{pmatrix}, \quad (6.19)$$

with $0 < \lambda_0 < 1/2$. We consider a parametrisation $\boldsymbol{\theta} \rightarrow \rho_{\boldsymbol{\theta}}$ of states obtained by small unitary rotations of ρ_0 , and different choices of the eigenvalue. We choose the parameter vector $\boldsymbol{\theta} := (\lambda, \mathbf{w})$, where $\mathbf{w} = (u, v) \in \mathbb{R}^2$ corresponds to the small unitary rotations of the eigenvectors, and λ is the smallest eigenvalue. That is, any state ρ described by $\boldsymbol{\theta} = (\lambda, u, v)$ is of the form

$$\rho_{\boldsymbol{\theta}} := U\left(\frac{\mathbf{w}}{\sqrt{n}}\right) \begin{pmatrix} 1 - \lambda & 0 \\ 0 & \lambda \end{pmatrix} U\left(\frac{\mathbf{w}}{\sqrt{n}}\right)^*, \quad (6.20)$$

where the unitary $U\left(\frac{\mathbf{w}}{\sqrt{n}}\right)$ is given by

$$\begin{aligned} U\left(\frac{\mathbf{w}}{\sqrt{n}}\right) &:= \exp\left(\frac{i}{\sqrt{n}}(u\sigma_x + v\sigma_y)\right) \\ &= \begin{pmatrix} \cos|\mathbf{w}|/\sqrt{n} & -\exp(-i\varphi)\sin|\mathbf{w}|/\sqrt{n} \\ \exp(i\varphi)\sin|\mathbf{w}|/\sqrt{n} & \cos|\mathbf{w}|/\sqrt{n} \end{pmatrix}, \end{aligned} \quad (6.21)$$

with $\varphi = \text{Arg}(-v + iu)$. Note that in this parametrisation we have $\rho_0 = \rho_{\boldsymbol{\theta}_0}$ with $\boldsymbol{\theta}_0 = (\lambda_0, 0, 0)$. The aim of the second measurement stage is then to estimate the unknown parameter vector $\boldsymbol{\theta} = (\lambda, u, v) = (\lambda, \mathbf{w})$ corresponding to the true state ρ .

6.3.2 The ‘which block’ measurement stage

The second measurement stage involves a joint measurement on the $n - \tilde{n}$ remaining copies of the state. We therefore consider the joint states $\rho_n^{\boldsymbol{\theta}} := \rho_{\boldsymbol{\theta}}^{\otimes n}$ on n identical qubits, with the parametrisation around the preliminary estimator $\rho_0 = \tilde{\rho}$ described above. It is known that the states $\rho_n^{\boldsymbol{\theta}}$ have a block-diagonal form with respect to the decomposition of the underlying space $(\mathbb{C}^2)^{\otimes n}$ in irreducible representations of the groups $SU(2)$ and $S(n)$ [57, 8, 36]. The representation π_n of $SU(2)$ is given by $\pi^{(n)}(u) = u^{\otimes n}$ for any $u \in SU(2)$, and the representation $\tilde{\pi}_n$ of the symmetric group $S(n)$ is given by the permutation of factors

$$\tilde{\pi}^{(n)}(\tau) : v_1 \otimes \dots \otimes v_n \rightarrow v_{\tau^{-1}(1)} \otimes \dots \otimes v_{\tau^{-1}(n)}, \quad \tau \in S(n). \quad (6.22)$$

The Hilbert space can be decomposed in the form

$$(\mathbb{C}^2)^{\otimes n} = \bigoplus_{j=0,1/2}^{n/2} \mathcal{H}_j \otimes \mathcal{H}_n^j, \quad (6.23)$$

where the lower limit in the direct sum is 0 for even n and $1/2$ for odd n . The two group representations decompose into direct sums of irreducible representations as $\pi^{(n)}(u) = \bigoplus_j \pi_j(u) \otimes \mathbf{1}$ and $\tilde{\pi}^{(n)}(\tau) = \bigoplus_j \mathbf{1} \otimes \tilde{\pi}_j(\tau)$ where π_j is the irreducible representation of $SU(2)$ with total angular momentum $J^2 = j(j+1)$ which acts on $\mathcal{H}_j \cong \mathbb{C}^{2j+1}$, and $\tilde{\pi}_j$ is the irreducible representation of the symmetric group $S(n)$ acting on $\mathcal{H}^j \cong \mathbb{C}^{n_j}$ with

$$n_j = \binom{n}{n/2-j} - \binom{n}{n/2-j-1}. \quad (6.24)$$

The density matrix ρ_n^θ is invariant under permutations and can be decomposed as

$$\rho_n^\theta = \bigoplus_{j=0,1/2}^{n/2} p_{n,\lambda}(j) \rho_{j,n}^w \otimes \frac{\mathbf{1}}{n_j}, \quad (6.25)$$

where the probability distribution $p_{n,\lambda}(j)$ is given by [57, 59, 8]

$$p_{n,\lambda}(j) := \frac{n_j}{1-2\lambda} \lambda^{n/2-j} (1-\lambda)^{n/2+j+1} (1-p^{2j+1}), \quad (6.26)$$

with $p = \frac{\lambda}{1-\lambda}$. The above distribution can be written in the form

$$p_{n,\lambda}(j) := B_{n,\lambda}(n/2-j) \times K(j, n, \lambda), \quad (6.27)$$

where $B_{n,\lambda}(k) = \binom{n}{k} \lambda^k (1-\lambda)^{n-k}$ is the binomial distribution and the term $K(j, n, \lambda)$ is given by

$$K(j, n, \lambda) := (1-p^{2j+1}) \frac{n + (2(j-j_n) + 1)/(1-2\lambda)}{n + (j-j_n + 1)/(1-\lambda)}, \quad j_n := n(1/2 - \lambda). \quad (6.28)$$

The binomially distributed variable $n/2 - j$ concentrates around its mean value of $n\lambda$ with high probability

$$\mathbb{P} \left[n\lambda - n^{1/2+\epsilon_3} \leq n/2 - j \leq n\lambda + n^{1/2+\epsilon_3} \right] \geq 1 - 2 \exp(-2n^{2\epsilon_3}), \quad (6.29)$$

where $\epsilon_3 > 0$ is an arbitrary constant. This follows from a straightforward application of Hoeffding's inequality (6.68) to the binomial distribution. The mass of the distribution $B_{n,\lambda}(n/2 - j)$ therefore concentrates over values of j in the interval

$$\mathcal{J}_n := \{j \mid j_n - n^{1/2+\epsilon_3} \leq j \leq j_n + n^{1/2+\epsilon_3}\}. \quad (6.30)$$

For all $j \in \mathcal{J}_n$, the factor $K(j, n, \lambda) = 1 + O(n^{-1/2+\epsilon_3})$ provided that λ is bounded away from $1/2$, which is one of the reasons we chose to treat the two cases above separately. Additionally we note that the factor $K(j, n, \lambda)$ remains bounded over all values of j as long as $\lambda < 1/2$. From the concentration of the binomial distribution over values of $j \in \mathcal{J}_n$, and the value of $K(j, n, \lambda)$ in this interval it follows that

$$p_{n,\lambda}(\mathcal{J}_n) = 1 - O(n^{-1/2+\epsilon_3}). \quad (6.31)$$

A ‘‘which block’’ measurement corresponds to an output of a particular value of j from the distribution (6.27), and an associated posterior state $\rho_{j,n}^w$. This value of j lies in the set \mathcal{J}_n with high probability. The eigenvalue parameter λ is estimated from this value of j . As in the case of the local adaptive estimator in section 6.2, in order to derive a minimax upper bound, we shall define $\hat{\lambda}$ as the ‘add-beta’ estimator, identifying $n/2 - j$ with k in (6.8). However a discussion regarding this choice for the estimator $\hat{\lambda}$ is discussed later in section 6.5. We note that a possible physical implementation of such a measurement is detailed in [57], and involves coupling the joint states to different bosonic field and performing a homodyne measurement.

Information about the ‘rotation’ parameters are contained in the block state $\rho_{j,n}^w$. These parameters are estimated using established LAN results [57] which we recall in section 6.3.3 below.

6.3.3 Local asymptotic normality

The block state $\rho_{j,n}^w$ encodes information about the rotation parameters $\mathbf{w} = (u, v)$. The optimal estimation strategy for these parameters has been established using results about the LAN of qubit states [57]. This shows that for large n , the block states $\rho_{j,n}^w$ approach a Gaussian state ϕ^w of a one-mode continuous variables system *uniformly* over all $j \in \mathcal{J}_n$ and $\|\mathbf{w}\| \leq n^\eta$. The rotation parameters (u, v) are encoded linearly into the mean of the Gaussian

state $\phi^{\mathbf{w}}$. So the problem of the optimal estimation of these parameters for the block state can be translated into one of estimating the displacement of $\phi^{\mathbf{w}}$. These ideas have been treated in detail in [57], and we only include a brief overview here. The block states $\rho_{j,n}^{\mathbf{w}}$ depend on the parameters (u, v) in the following way

$$\rho_{j,n}^{\mathbf{w}} = U_j \left(\frac{\mathbf{w}}{\sqrt{n}} \right) \rho_{j,n}^{\mathbf{0}} U_j \left(\frac{\mathbf{w}}{\sqrt{n}} \right)^*, \quad (6.32)$$

where the unitaries are defined as $U_j(\mathbf{w}) := \exp(i(uJ_{j,x} + vJ_{j,y}))$, with $J_{j,l}$ being the generators of rotations in the irreducible representation π_j of $SU(2)$. The state $\rho_{j,n}^{\mathbf{0}}$ is expressed as

$$\rho_{j,n}^{\mathbf{0}} = \frac{1-p}{1-p^{2j+1}} \sum_{m=-j}^j p^{j-m} |j, m\rangle \langle j, m| \quad (6.33)$$

with $p = \lambda/(1-\lambda)$ as before. The set $\{|j, m\rangle : m = -j, \dots, j\}$ is an orthonormal basis on \mathcal{H}_j such that $J_{j,z}|j, m\rangle = m|j, m\rangle$. It has been demonstrated [57] that the family of states $\mathcal{F}_n := \{\rho_{j,n}^{\mathbf{w}}, \|\mathbf{w}\| \leq n^\eta, j \in \mathcal{J}_n\}$ is asymptotically Gaussian. This means that as $n \rightarrow \infty$ the family of states $\rho_{j,n}^{\mathbf{w}}$ ‘‘converges’’ to a family of Gaussian states $\phi^{\mathbf{w}}$ of a one-mode continuous variables system, for all $j \in \mathcal{J}_n$ and $\|\mathbf{w}\| \leq n^\eta$. In order to make this convergence more precise, we let

$$\phi^{\mathbf{0}} := (1-p) \sum_{k=0}^{\infty} p^k |k\rangle \langle k| \quad (6.34)$$

be a centred Gaussian state of a one mode continuous variables system, with $\{|k\rangle : k \geq 0\}$ denoting the Fock basis. The states $\phi^{\mathbf{w}}$ are defined as

$$\phi^{\mathbf{w}} := D(\sqrt{1-2\lambda}\alpha_{\mathbf{w}}) \phi^{\mathbf{0}} D(-\sqrt{1-2\lambda}\alpha_{\mathbf{w}}), \quad (6.35)$$

where $\alpha_{\mathbf{w}} = -v + iu \in \mathbb{C}$. The operator $D(\alpha) := \exp(\alpha a^* - \bar{\alpha} a)$ is the displacement operator that for every $\alpha \in \mathbb{C}$ maps the vacuum vector $|0\rangle$ to the coherent state $|\alpha\rangle$, with a^*, a being the creation and annihilation operators satisfying $[a, a^*] = 1$. The convergence of the block state $\rho_{n,j}^{\mathbf{w}}$ to the Gaussian state $\phi^{\mathbf{w}}$ is formalised in the following theorem.

Theorem 7. *Let $V_j : \mathcal{H}_j \rightarrow L^2(\mathbb{R})$ be the isometry*

$$V_j : |j, m\rangle \rightarrow |j - m\rangle \quad (6.36)$$

that maps the orthonormal basis of \mathcal{H}_j into the Fock basis of $L^2(\mathbb{R})$. Then for the family of block states $\rho_{n,j}^{\mathbf{w}}$ defined by (6.32), and the family of Gaussian states $\phi^{\mathbf{w}}$ defined by (6.35), the following convergence holds for any $0 \leq \eta \leq 1/6$ and $0 < \epsilon_3 < 1/2$

$$\sup_{\|\mathbf{w}\| \leq n^\eta} \max_{j \in \mathcal{J}_n} \|V_j \rho_{j,n}^{\mathbf{w}} V_j^* - \phi^{\mathbf{w}}\|_1 = O(n^{-1/4+\eta+\epsilon_3}) \quad (6.37)$$

over the set $\mathcal{J}_n = \{j \mid j_n - n^{1/2+\epsilon_3} \leq j \leq j_n + n^{1/2+\epsilon_3}\}$. The convergence is uniform over $\lambda \geq 1/2(1+\delta)$ for an arbitrary fixed $\delta > 0$.

The interpretation is that the block state $\rho_{j,n}^{\mathbf{w}}$ can be mapped by means of physical transformations (in this case an isometric embedding) into the Gaussian state $\phi^{\mathbf{w}}$ with vanishing norm-one error, uniformly over the unknown parameter \mathbf{w} and over the block index j . A possible physical implementation is detailed in [57]; the ensemble of qubits is coupled with a Bosonic field such that the state is transferred to the field after some time.

In order to estimate the rotation parameters $\hat{\mathbf{w}} = (\hat{u}, \hat{v})$, one first maps the qubit state via the isometry V_j , and then performs a heterodyne measurement, which is optimal for estimating displacement. In the next section we discuss the Bures risk of the estimation procedure described above.

6.4 Minimax upper and lower bounds

The overall measurement procedure we propose can be briefly summarised as consisting of a preliminary localisation stage, where a vanishing number \tilde{n} of copies of the state is used to localise the state ρ . This estimate $\tilde{\rho}$ informs the choice of measurements in the second stage. When $\tilde{\rho}$ is within a fixed ball of radius $\delta > 0$ around the fully mixed state, standard tomographic measurements are performed on the remaining copies of the state. However, if $\tilde{\rho}$ lies outside this ball, the measurements performed in second stage uses techniques based on the principle of LAN to estimate the parameter vector $\boldsymbol{\theta} = (\lambda, u, v)$. In this section we look at the Bures risk of the measurement estimator pair $R(\rho, \hat{\rho}_n) := \mathbb{E} [D_B(\rho, \hat{\rho}_n)^2]$. As we will show below, the risk of a good estimator scales as $1/n$, and we would like to find asymptotic upper and lower bounds for the rescaled maximum risk

$$R_{max}(\hat{\rho}) = \limsup_{n \rightarrow \infty} \sup_{\rho \in \mathbb{S}_2} nR(\rho, \hat{\rho}_n). \quad (6.38)$$

6.4.1 The upper bound

We now make concrete our final estimator for the state ρ . The first stage involves using a vanishing number of copies $\tilde{n} := n^{1-\kappa}$ (with $\kappa > 0$) to get a rough estimate $\tilde{\rho}$. This estimate informs the second measurement stage. The subsequent measurement stage differs depending on whether the state is estimated to be close to the fully mixed state.

If the estimate $\tilde{\rho}$ lies in a small ball of radius $\delta > 0$ around the fully mixed state, then measurements in the standard σ_i , $i = x, y, z$ basis are performed on $(n - \tilde{n})/3$ copies of the state. The outcomes of each measurement ± 1 , and the associated probabilities are $p_\rho(\pm 1 | \sigma_i) = p_i(\pm 1) := \text{Tr}(\rho P_i^{\pm 1})$, where the projectors $P_i^{\pm 1}$ are defined via $\sigma_i = P_i^{+1} - P_i^{-1}$. The total number n_i of ± 1 outcomes obtained by $n/3$ measurements of σ_i is binomially distributed $B_{n/3, p_i(\pm 1)}(n_i)$. The final estimate of the state is constructed as the maximum likelihood (ML) estimate from these measurement outcomes

$$\hat{\rho}_n = \arg \max_{\tau \in \mathbb{S}_2} \sum_{i=x,y,z} n_i \log \text{Tr}(\tau P_i^{+1}) + (n/3 - n_i) \log \text{Tr}(\tau P_i^{-1}), \quad (6.39)$$

where the maximisation is over the space of all 2×2 density matrices \mathbb{S}_2 .

On the other hand, if the preliminary estimate $\tilde{\rho}$ lies away from the fully mixed state, we perform the following measurements to estimate the parameter vector $\theta = (\lambda, u, v)$. A ‘which block’ measurement outputs a value of j from which the eigenvalue λ is estimated, cf. section 6.3.2. Similarl to the separable measurements strategy, we consider the following ‘add-beta’ estimator for the eigenvalue λ [22, 23]

$$\hat{\lambda}_n = \begin{cases} \frac{1/2}{n+5/4}, & \frac{n}{2} - j = 0, \\ \frac{2}{n+7/4}, & \frac{n}{2} - j = 1, \\ \frac{n/2-j+3/4}{n+3/2}, & \frac{n}{2} - j = 2, \dots, n-2, \\ \frac{n-1/4}{n+7/4}, & \frac{n}{2} - j = n-1, \\ \frac{n+3/4}{n+5/4}, & \frac{n}{2} - j = n \end{cases} \quad (6.40)$$

The range of possible values of j is $[0, n/2]$, and therefore only some of the rules of the estimator described above are used. However, we describe the estimator over the range $[0, n]$ as this will be used shortly to upper bound the minimax risk.

Conditional on j , we are left with the block state $\rho_{j,n}^w$. Using Theorem 7 we can isometrically map this state onto the Fock space, close to the Gaussian state ϕ^w . In order to estimate the displacement parameter w we perform a heterodyne measurement with outcome \hat{w}_n . The final estimate of our the state $\hat{\rho}_n$ is constructed from the estimated parameter vector $\hat{\theta}_n = (\hat{\lambda}_n, \hat{w}_n)$ as

$$\hat{\rho}_n = U \left(\frac{\hat{w}_n}{\sqrt{n}} \right) \begin{pmatrix} 1 - \hat{\lambda}_n & 0 \\ 0 & \hat{\lambda}_n \end{pmatrix} U \left(\frac{\hat{w}_n}{\sqrt{n}} \right)^*. \quad (6.41)$$

We now state precisely the an upper bound for the minimax risk of the square Bures distance for the measurement strategy described above.

Theorem 8. *Let $\hat{\rho}_n$ be the estimator described above. The asymptotic rescaled maximum risk is bounded from above as*

$$\limsup_{n \rightarrow \infty} \sup_{\rho} nR(\rho, \hat{\rho}_n) \leq \frac{3}{2}. \quad (6.42)$$

The proof of this theorem is detailed in the appendix (6.8.3), and here we only provide an outline for the arguments employed. The choice of measurements in the second stage depend on whether the preliminary estimate $\tilde{\rho}$ lies inside or outside a small ball of radius $\delta > 0$ around the fully mixed state. In keeping with this, let us therefore denote $\hat{\rho}_n^1$ as the estimator (6.39) chosen when $|\tilde{r}| \leq \delta$, and let $\hat{\rho}_n^2$ be the LAN based estimator (6.41), when $|\tilde{r}| > \delta$. The minimax risk can then be bounded from above as follows

$$\limsup_{n \rightarrow \infty} \sup_{\rho} nR(\rho, \hat{\rho}_n) \leq \max \left\{ \limsup_{n \rightarrow \infty} \sup_{\rho \in B_1} n\mathbb{E} \left[D_B(\rho, \hat{\rho}_n^1)^2 \mid |\tilde{r}| \leq \delta \right], \limsup_{n \rightarrow \infty} \sup_{\rho \notin B_2} n\mathbb{E} \left[D_B(\rho, \hat{\rho}_n^2)^2 \mid |\tilde{r}| > \delta \right] \right\}, \quad (6.43)$$

where B_1 and B_2 are balls of radius $\delta + n^{-1/2+\epsilon_2}$ and $\delta - n^{-1/2+\epsilon_2}$ respectively. The two terms are evaluated explicitly in section (6.8.3) of the appendix. The term corresponding to the estimator $\hat{\rho}_n^1$ is straightforward to evaluate as the square Bures distance is locally quadratic for states in B_1 . From this quadratic expansion and the efficiency of the maximum likelihood estimator in the asymptotic regime, the risk can be expressed as

$$\mathbb{E} \left[D_B(\rho, \hat{\rho}_n^1)^2 \right] \approx \frac{1}{n} \text{Tr} \left(I(\rho)^{-1} G \right), \quad (6.44)$$

where G is the weight matrix reconstructing the quadratic approximation of the square Bures distance, and I is the Fisher information matrix. From the explicit form of G and I , the asymptotic risk can be bounded as

$$\limsup_{n \rightarrow \infty} \sup_{\rho \in B_1} n\mathbb{E} \left[D_B(\rho, \hat{\rho}_n^1)^2 \right] \leq \frac{3}{4} \left(1 + \frac{\delta}{1 - \delta} \right). \quad (6.45)$$

The other term in (6.43) corresponding to the estimator $\hat{\rho}_n^2$ uses the local parameterisation of states $\boldsymbol{\theta} = (\boldsymbol{\lambda}, \boldsymbol{w})$, and the approximation of the square Bures distance used in section 6.2, and detailed in the appendix (6.8.2). We therefore get that the risk can be bounded as

$$\mathbb{E} \left[D_B(\rho, \hat{\rho}_n^2)^2 \right] \leq \mathbb{E} \left[D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n)^2 + \frac{1}{4} \Phi^2 \right] + O(\Phi^4) \quad (6.46)$$

$$\leq \mathbb{E} \left[D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n)^2 \right] + \frac{1}{n} \mathbb{E} \left[(u - \hat{u}_n)^2 + (v - \hat{v}_n)^2 \right] + O(n^{-2}) \quad (6.47)$$

The term corresponding to the rotation parameters has been evaluated in [57] using LAN based techniques. Since LAN holds in the limit of large n , the problem of estimating the rotation parameters is translated to one of determining the displacement of a Gaussian state $\phi^{\boldsymbol{w}}$. The heterodyne measurement is known to be the optimal measurement in this case [57, 59]. The first term corresponding to the Hellinger risk is bounded from above by the Kullback-Leibler (KL) risk of estimating a binomial parameter λ from outcomes k distributed as $B_{n,\lambda}(k)$. The ‘add-beta’ estimator is known to be minimax optimal in this case, and its rate is known in the literature. Together with the LAN results for the rotation parameters, we bound risk as

$$\limsup_{n \rightarrow \infty} \sup_{\rho \notin B_2} nR(\rho, \hat{\rho}_n^2) \leq \limsup_{n \rightarrow \infty} \sup_{\lambda} n\mathbb{E}_{\text{Binom}} \left[D_{KL}(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n) \right] + 1 \leq \frac{3}{2}. \quad (6.48)$$

Comparing this bound with (6.45), we arrive at the stated upper bound of 3/2 in Theorem 8 provided $\delta < 1/2$.

6.4.2 The lower bound

In this section we derive a lower bound on the asymptotic rescaled risk with respect to the square Bures distance. The key idea is to restrict the attention to a smaller state space region where the state is ‘hardest’ to estimate, and

evaluate the minimax risk over this region, thus obtaining a lower bound for the overall minimax risk.

Let us consider that the true state ρ lies in a local neighbourhood of size $n^{-1/2+\epsilon}$ around an arbitrary but fixed state ρ_0 as defined in equation (6.19), whose smallest eigenvalue satisfies $0 < \lambda_0 < 1/2$. For any estimation procedure $\hat{\rho}_n$ we have the lower bound for the maximum risk

$$\begin{aligned} \limsup_{n \rightarrow \infty} \sup_{\rho \in \mathbb{S}_2} nR(\rho, \hat{\rho}_n) &\geq \limsup_{n \rightarrow \infty} \sup_{\|\rho - \rho_0\|_1 \leq n^{-1/2+\epsilon}} nR(\rho, \hat{\rho}_n) \\ &\geq \limsup_{n \rightarrow \infty} \inf_{\hat{\rho}_n} \sup_{\|\rho - \rho_0\|_1 \leq n^{-1/2+\epsilon}} nR(\rho, \hat{\rho}_n) \\ &:= R_{\minimax}(\rho_0), \end{aligned} \quad (6.49)$$

where the right side is the *local minimax* risk at ρ_0 .

Since the state ρ_0 is taken to be away from the boundary of the Bloch sphere, we can parametrise its local neighbourhood using the local parameter $\boldsymbol{\theta} = (h, \mathbf{u})$

$$\rho = \rho_{\boldsymbol{\theta}} := U \left(\frac{\mathbf{w}}{\sqrt{n}} \right) \begin{pmatrix} 1 - \lambda_0 - h/\sqrt{n} & 0 \\ 0 & \lambda_0 + h/\sqrt{n} \end{pmatrix} U \left(\frac{\mathbf{w}}{\sqrt{n}} \right)^*. \quad (6.50)$$

The square Bures distance is locally quadratic

$$D_B(\rho_{\boldsymbol{\theta}}, \rho_{\boldsymbol{\theta}'})^2 = \frac{1}{n} (\boldsymbol{\theta} - \boldsymbol{\theta}')^T \Gamma_0 (\boldsymbol{\theta} - \boldsymbol{\theta}') + O(n^{-3/2}), \quad (6.51)$$

where Γ_0 is the weight matrix

$$\Gamma_0 = \begin{pmatrix} \frac{1}{4\lambda_0(1-\lambda_0)} & 0 & 0 \\ 0 & (1-2\lambda_0)^2 & 0 \\ 0 & 0 & (1-2\lambda_0)^2 \end{pmatrix}. \quad (6.52)$$

In this case we can apply the LAN theory [57] to obtain the local minimax risk for the square Bures distance. The upshot of the theory is that the classical statistical model given by the distribution over blocks (cf. equation (6.26)) can be approximated by a one-dimensional Gaussian model $N(h, v_0)$ with fixed variance $v_0 = \lambda_0(1 - \lambda_0)$ and mean equal to the unknown local parameter h . Additionally, the quantum statistical model described by the quantum state of the irreducible block can be approximated by a quantum Gaussian shift model (independent of the classical one), as described in Theorem 7. The optimal

measurement here is the heterodyne, and after rescaling by a constant factor we obtain the unbiased estimator $\hat{\mathbf{u}}$ which has a two-dimensional Gaussian distribution $\hat{\mathbf{u}} \sim N(\mathbf{u}, w_0 \cdot I_2)$ with $w_0 = (1 - \lambda_0)/(2(1 - 2\lambda_0)^2)$. The local minimax risk is the sum of the contribution from the classical and respectively the quantum part of the Gaussian model, weighted with the matrix Γ_0

$$\begin{aligned} R_{\min\max}(\rho_0) &= \Gamma_{00}\mathbb{E}[(\hat{h} - h)^2] + \Gamma_{11}\mathbb{E}[(\hat{u} - u)^2] + \Gamma_{22}\mathbb{E}[(\hat{v} - v)^2] \\ &= \frac{1}{4\lambda_0(1 - \lambda_0)}v_0 + 2(1 - 2\lambda_0)^2w_0 \\ &= \frac{1}{4} + (1 - \lambda_0) = \frac{5}{4} - \lambda_0. \end{aligned} \tag{6.53}$$

As the state ρ_0 defining the local neighbourhood is chosen arbitrarily, we see that the right side of the above equation achieves its maximum as $\lambda_0 \rightarrow 0$, and we therefore get the asymptotic lower bound for the rescaled maximum risk of any estimator.

$$\limsup_{n \rightarrow \infty} \sup_{\rho \in \mathbb{S}_2} nR(\rho, \hat{\rho}_n) \geq \frac{5}{4}. \tag{6.54}$$

As expected, the above lower bound is smaller than the $3/2$ upper bound derived in section 6.4.1.

6.5 The minimax optimal estimator

In deriving the minimax bounds for both the proposed estimators, the key observation was that the Bures risk decomposes locally into contributions from the Hellinger risk of estimating the eigenvalue parameter λ and a quadratic risk corresponding to the estimation of the rotation parameters (see (6.11) and appendix 6.8.2). The Hellinger risk was then bounded from above by the Kullback-Leibler (KL) risk of estimating the binomial parameter. The estimator of the binomial parameter achieving the minimax rate for the KL risk is known to be ‘add-beta’ estimator, and both the local and global estimators for the state ρ proposed using this estimator for the eigenvalue parameter (6.40,6.8).

The reason why we were not able to prescribe an asymptotically minimax estimator is that we could not devise a minimax estimator for the binomial parameter λ , with respect to the square Hellinger distance. The following proposition follows immediately from the asymptotic analysis of section 6.4

and shows that the original optimal state estimation problem reduces to the ‘classical’ one of estimating the binomial parameter λ .

Proposition 1. *Let $\hat{\lambda}_{opt}$ be an asymptotically minimax estimator of the binomial parameter λ under the square Hellinger loss function. The estimators defined by replacing the ‘add-beta’ estimators for λ in equation (6.40) with $\hat{\lambda}_{opt}$ will then be asymptotically minimax optimal for qubit states.*

Although we were not able to devise a minimax estimator under the square Hellinger distance, we would like to make some comments on this problem, emphasising that it is crucial to study what happens at the boundary when $\lambda \approx 0$. Indeed, for values of λ away from this boundary, the local asymptotic minimax rate is easily derived as the square Hellinger distance is locally quadratic and the classical asymptotic efficiency theory [118] applies. The standard estimator $\hat{\lambda} = k/n$ is a natural first choice as it is unbiased and achieves the Cramer Rao lower bound with variance $\text{Var}(\hat{\lambda}) = (nI)^{-1}$, where $I = \frac{1}{\lambda(1-\lambda)}$ is the Fisher information. In the region where $\lambda > 0$, using a locally quadratic approximation for the Hellinger risk, we have

$$\mathbb{E}_{\text{Binom}} [D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}})^2] = \frac{1}{4\lambda(1-\lambda)} \text{Var}(\hat{\lambda}) + o(n^{-1}) = \frac{1}{4n} + o(n^{-1}). \quad (6.55)$$

This holds for every *fixed* $\lambda \in (0, 1/2]$, and gives the same rate as the one in the lower bound (6.53). However the convergence is not uniform over λ close to the zero, which affects the constant in the asymptotic maximum risk. To see this, consider the case when λ is n dependent such that $n\lambda \rightarrow \mu$, with $\mu > 0$ being a fixed constant. The Hellinger risk is given by

$$\mathbb{E}_{\text{Binom}} [D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}})^2] = \mathbb{E}_{\text{Binom}} \left[\left(\hat{\lambda}^{1/2} - \lambda^{1/2} \right)^2 \right] + \mathbb{E}_{\text{Binom}} \left[\left((1-\hat{\lambda})^{1/2} - (1-\lambda)^{1/2} \right)^2 \right]. \quad (6.56)$$

The second term in the above equation is bounded as

$$\mathbb{E}_{\text{Binom}} \left[\left((1-\hat{\lambda})^{1/2} - (1-\lambda)^{1/2} \right)^2 \right] \leq \mathbb{E}_{\text{Binom}} \left[(\lambda - \hat{\lambda})^2 \right] / (1-\lambda) = \mu/n^2. \quad (6.57)$$

Substituting (6.57) in (6.56), we have

$$\mathbb{E}_{\text{Binom}} [D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}})^2] = \mathbb{E}_{\text{Binom}} \left[(\hat{\lambda}^{1/2} - \lambda^{1/2})^2 \right] + O(n^{-2}) \quad (6.58)$$

$$= \frac{1}{n} \mathbb{E}_{\text{Po}(\mu)} \left[\left(K^{1/2} - \mu^{1/2} \right)^2 \right] + O(n^{-2}) \quad (6.59)$$

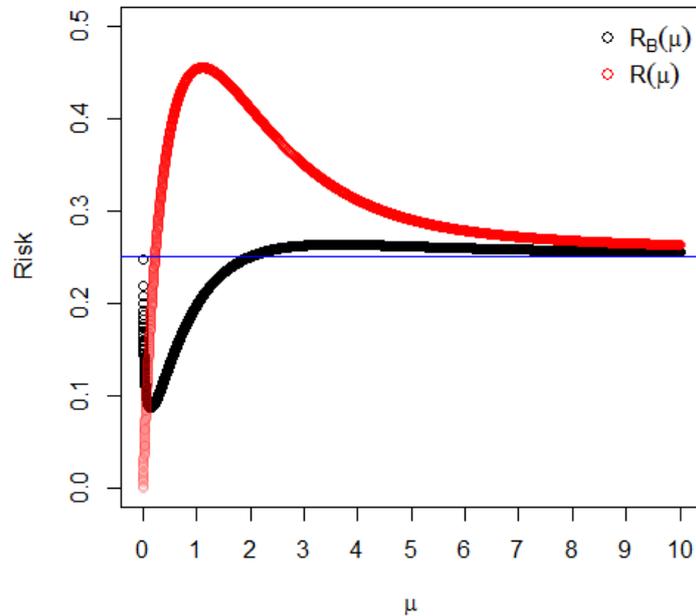


Figure 6.1 Plots of the Hellinger risk functions $R(\mu)$ and $R_B(\mu)$ for various values of the ‘true’ Poisson parameter μ . The horizontal blue line marks a value of $1/4$. See main text for details.

In the last equality we used the fact that under the scaling $n\lambda \rightarrow \mu$, the Binomial random variable converges to a Poisson random variable $K \sim \text{Po}(\mu)$. Therefore the risk in this case is given by the function $R(\mu) = \mathbb{E}_{\text{Po}(\mu)} \left[\left(K^{1/2} - \mu^{1/2} \right)^2 \right]$, where the expectation is taken with respect to the Poisson distribution with parameter μ . If this function was bounded such that $R(\mu) \leq 1/4$, then it would suggest that the standard estimator $\hat{\lambda}$ might be globally asymptotically minimax. However, plotting the function $R(\mu)$ numerically, we see from Figure 6.1 that it attains a maximum value of $\max R(\mu) \approx 0.455$ around $\mu = 1.11$, and converges to $1/4$ for large values of μ which corresponds to λ away from zero. This shows that the standard estimator doesn’t achieve a minimax constant of $1/4$ for all λ , and illustrates that the difficulty in deriving a minimax rate for the square Hellinger distance lies in the Poisson range, i.e for values of λ such that $n\lambda \rightarrow \mu$.

As an alternative, we consider the Bayes estimator $\hat{\mu}_B^{1/2}$ for $\mu^{1/2}$, and plot numerically the function $R_B(\mu) := \mathbb{E}_{\text{Po}(\mu)} \left[\left(\hat{\mu}_B^{1/2}(K) - \mu^{1/2} \right)^2 \right]$. We now describe the Bayes estimator $\hat{\mu}_B^{1/2}$ of the Poisson parameter. It is known that a conjugate

family for the Poisson model is the Gamma family of priors, i.e

$$f_{\alpha,\beta}(t) = \frac{t^{\alpha-1} \exp(-t/\beta)}{\beta^\alpha \Gamma(\alpha)}, \quad t > 0, \quad (6.60)$$

where α, β are the shape and scale parameters respectively. Given an outcome $K = k$ from the Poisson distribution $\text{Po}(\mu)$, the posterior distribution is easily calculated to be $\text{Gamma}(\alpha + k, \frac{\beta}{\beta+1})$. Then the Bayes estimator is of the following form

$$\hat{\mu}_B^{1/2} := \frac{\Gamma(k + \alpha + 1/2)}{\Gamma(k + \alpha)} \left(\frac{\beta}{\beta + 1} \right)^{1/2}. \quad (6.61)$$

In Figure 6.1, we plot the ‘frequentist’ risk $R_B(\mu)$ of the Bayes estimator for a particular prior with $\alpha = 0.41$ and $\beta = 200$, and a range of ‘true’ values for μ . We see that the risk remains upper-bounded by a value only slightly greater than $1/4$, and for large values of μ the risk tends to a limiting value of $1/4$. This supports the conjecture that the minimax constant for the square Hellinger distance is $1/4$.

6.6 Quantum Relative Entropy

In our derivation of the minimax upper bounds, we bounded the Hellinger risk of estimating the eigenvalues by the Kullback-Leibler (KL) risk for which the asymptotic minimax rate is known. As the KL distance is the classical analogue of the quantum relative entropy $S(\rho||\rho') = \text{Tr}[\rho(\log \rho - \log \rho')]$, a question naturally arises - can the techniques used in this chapter be applied to derive the minimax rate for the *Quantum Relative Entropy* (QRE)? A key element would be to decompose the QRE locally. Similar to decomposition of the Bures distance in (6.11), the QRE risk can be shown to locally decompose into a sum of contributions from the KL risk and a term involving the ‘rotation parameters’. For qubit states, the QRE between the state ρ and the estimate $\hat{\rho}_n$ is represented in terms of the Bloch vectors as [38]

$$S(\rho||\hat{\rho}_n) = \frac{1}{2} \left[\log(1 - |\mathbf{r}|^2) - \log(1 - |\hat{\mathbf{r}}_n|^2) + |\mathbf{r}| \log \left(\frac{1 + |\mathbf{r}|}{1 - |\mathbf{r}|} \right) - |\mathbf{r}| \cos \Phi \log \left(\frac{1 + |\hat{\mathbf{r}}_n|}{1 - |\hat{\mathbf{r}}_n|} \right) \right], \quad (6.62)$$

where Φ is the angle between the Bloch vectors of the two states. This can be rewritten as

$$\begin{aligned} S(\rho \|\hat{\rho}_n) &= D_{KL}(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n) + \frac{1-2\lambda}{2}(1 - \cos \Phi) \log \left(\frac{1 - \hat{\lambda}_n}{\hat{\lambda}_n} \right) \\ &= D_{KL}(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n) + \frac{1-2\lambda}{4}(\Phi^2 + O(\Phi^4)) \log \left(\frac{1 - \hat{\lambda}_n}{\hat{\lambda}_n} \right), \end{aligned} \quad (6.63)$$

where $D_{KL}(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}})$ is the Kullback-Leibler distance between the two distributions $\boldsymbol{\lambda} = (\lambda, 1 - \lambda)$ and $\hat{\boldsymbol{\lambda}} = (\hat{\lambda}, 1 - \hat{\lambda})$, and in the second equality we expanded the cosine term to leading order in Φ . The proof of this decomposition can be found in appendix 6.8.4.

We will show that the global estimator discussed in section 6.4 achieves the rate $O(n^{-1} \log n)$ and no estimator can achieve faster rates, in particular the ‘standard’ rate n^{-1} . Consider the estimator defined by equations (6.40) and (6.41), and note that the classical component estimator $\hat{\lambda}_n$ is the minimax optimal estimator for the binomial model and is always larger than c/n for some fixed constant $c > 0$. Using the same arguments as in Theorem 8 we find

$$\sup_{\rho} \mathbb{E}[D_{KL}(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n)] = O(n^{-1}). \quad (6.64)$$

On the other hand, since $\hat{\lambda}_n \geq c/n$, the second term in equation (6.63) is bounded by $c' \log n (\Phi^2 + O(\Phi^4))$; since Φ is estimated at standard rate, the second term is therefore upper bounded as

$$\sup_{\rho} \mathbb{E} \left[\frac{1-2\lambda}{4} (\Phi^2 + O(\Phi^4)) \log \left(\frac{1 - \hat{\lambda}_n}{\hat{\lambda}_n} \right) \right] = O(n^{-1} \log n), \quad (6.65)$$

which determines the rate.

We will now show that no estimator can have maximum risk converging faster than $n^{-1} \log n$. Since we are interested in the maximum risk, we will set $\lambda = 0$ (pure states), and show that the risk cannot decrease faster than $n^{-1} \log n$ even if we know that the state is pure! As a consequence of local asymptotic normality, any estimator will have the property that $\mathbb{P}[1 - \cos \Phi \geq c/n] \geq \epsilon$ for some constants c, ϵ . Therefore we will consider the contribution to the risk

conditional on $1 - \cos \Phi \geq c/n$. By expanding the $D_{KL}(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}})$ term we have

$$S(\rho \parallel \hat{\rho}_n) \geq \log \frac{1}{1 - \hat{\lambda}_n} + \frac{c}{4n} \log \left(\frac{1 - \hat{\lambda}_n}{\hat{\lambda}_n} \right). \quad (6.66)$$

However, the righthand side achieves its minimum at $\hat{\lambda}_n = c/4n$, so the risk is larger than $c' \log n/n$. This shows that the minimax risk for the quantum relative entropy scales as $\log n/n$.

6.7 Conclusion

In this chapter we proposed two adaptive estimators for the qubit mixed state, one based on local measurements and the other on collective global measurements. In section 6.2 we upper-bounded the minimax Bures distance risk for the estimator based on local measurements and showed that it scales as $1/n$. In section 6.3, we proposed an estimator based on collective measurements and used LAN theory to obtain upper and lower bounds for the risk of $3/2n$ and $5/4n$ respectively. A key element in obtaining the upper bounds was the local decomposition of the Bures risk into contributions from the Hellinger risk of estimating the eigenvalue, and a quadratic contribution from the risk of estimating the ‘rotation parameters’. While the contribution to the Bures risk from the ‘rotation terms’ is easily shown to scale as $O(1/n)$, we noticed that the difficulty in establishing minimax results for the Bures distance is encapsulated in the challenges of establishing minimax results for the Hellinger risk. Finally in section 6.5, we considered these challenges and proposed that a minimax optimal estimator for the mixed qubit state ρ is immediately obtained given a minimax optimal estimator for the binomial parameter under the Hellinger loss function.

We also briefly considered the derivation of minimax bounds for the quantum relative entropy (QRE) risk. We derived a local decomposition of the QRE similar to the one obtained for the Bures distance, and demonstrated that the global estimator proposed achieves a rate of $O(n^{-1} \log n)$. We also showed that no estimator can achieve faster rates and established that the minimax QRE risk scales as $O(n^{-1} \log n)$. A possible direction for future work is to extend the results presented in this chapter to the multi-qubit case.

6.8 Appendix

6.8.1 Proof of Lemma 5

The proof is a straightforward application of Hoeffding's inequality.

Theorem 9 (Hoeffding's inequality). *Let R_1, \dots, R_m be independent random variables with $a_i \leq R_i \leq b_i$. Let $S = \sum_{i=1}^m R_i$, and $\mu = \mathbb{E}[S]$. Then for all $t > 0$*

$$\mathbb{P}(|S - \mu| \geq nt) \leq 2 \exp^{-2n^2 t^2 / \sum_i (b_i - a_i)}. \quad (6.67)$$

We have $X_i \in [-1, 1]$, and $\mathbb{E} \sum_i X_i = \frac{\tilde{n}}{3} r_x$. Applying Hoeffding's inequality we get

$$\mathbb{P}\left(\left|\frac{3}{\tilde{n}} \sum_i X_i - r_x\right|^2 \geq t^2\right) \leq 2 \exp^{-2t^2 \tilde{n}/3} \quad (6.68)$$

and similarly for the other spin components. Applying the three inequalities together, with $t^2 = n^{2\epsilon_2} - 1$ and $\epsilon_2 > 0$, we have

$$\mathbb{P}\left(\sum_{j=x,y,z} |r'_j - r_j|^2 \geq 3n^{2\epsilon_2 - 1}\right) \leq 6 \exp^{-2\tilde{n}n^{2\epsilon_2 - 1}/3}. \quad (6.69)$$

The estimate $\tilde{\rho}$ is then the closest state in trace distance to the matrix $\frac{1}{2}(\mathbf{1} + \mathbf{r}' \cdot \boldsymbol{\sigma})$. As $\|\rho - \tilde{\rho}\|_1^2 = \sum_{i=x,y,z} |\tilde{r}_i - r_i|^2$, (6.69) implies the stated bound. \square

6.8.2 Expansion of the square Bures distance

Here we derive the expansion (6.11) of the square Bures distance $D_B(\rho, \rho')^2 := 2[1 - \sqrt{F(\rho, \rho')}]$. We know that for qubits the fidelity between two states can be expressed in terms of the Bloch vectors as

$$\begin{aligned} F(\rho, \rho') &:= \frac{1}{2} \left(1 + \sqrt{1 - |\mathbf{r}|^2} \sqrt{1 - |\mathbf{r}'|^2} + \mathbf{r} \cdot \mathbf{r}'\right) \\ &= \frac{1}{2} \left(1 + \sqrt{1 - |\mathbf{r}|^2} \sqrt{1 - |\mathbf{r}'|^2} + |\mathbf{r}| |\mathbf{r}'| \cos \Phi\right) \end{aligned} \quad (6.70)$$

where Φ is the angle between the Bloch vectors. For two sufficiently close states, the angle Φ is small and the cosine term can be expanded as

$$\begin{aligned} F(\rho, \rho') &= \frac{1}{2} \left(1 + \sqrt{1 - |\mathbf{r}|^2} \sqrt{1 - |\mathbf{r}'|^2} + |\mathbf{r}||\mathbf{r}'| - \frac{|\mathbf{r}||\mathbf{r}'|}{2} \Phi^2 + \frac{|\mathbf{r}||\mathbf{r}'|}{24} \Phi^4 \right) \\ &= \left(\sqrt{(1 - \lambda)(1 - \lambda')} + \sqrt{\lambda\lambda'} \right)^2 - \frac{|\mathbf{r}||\mathbf{r}'|}{4} \Phi^2 + \frac{|\mathbf{r}||\mathbf{r}'|}{24} \Phi^4, \end{aligned} \quad (6.71)$$

where we have use the fact that $|\mathbf{r}| = 1 - 2\lambda$. Therefore the square Bures distance is given by

$$\begin{aligned} D_B(\rho, \rho')^2 &= 2 \left[1 - \sqrt{\left(\sqrt{(1 - \lambda)(1 - \lambda')} + \sqrt{\lambda\lambda'} \right)^2 - \frac{|\mathbf{r}||\mathbf{r}'|}{4} \Phi^2 + \frac{|\mathbf{r}||\mathbf{r}'|}{24} \Phi^4} \right] \\ &= 2 \left[1 - \left(\sqrt{(1 - \lambda)(1 - \lambda')} + \sqrt{\lambda\lambda'} \right) + \frac{1}{8} \frac{|\mathbf{r}||\mathbf{r}'|}{\sqrt{(1 - \lambda)(1 - \lambda')} + \sqrt{\lambda\lambda'}} \Phi^2 + O(\Phi^4) \right] \\ &= D_H(\boldsymbol{\lambda}, \boldsymbol{\lambda}')^2 + \frac{1}{4} \frac{(1 - 2\lambda)(1 - 2\lambda')}{\sqrt{(1 - \lambda)(1 - \lambda')} + \sqrt{\lambda\lambda'}} \Phi^2 + O(\Phi^4), \end{aligned} \quad (6.72)$$

where $D_H(\boldsymbol{\lambda}, \boldsymbol{\lambda}')^2$ is the square Hellinger distance between the binary distributions $\boldsymbol{\lambda} = (\lambda, 1 - \lambda)$ and $\boldsymbol{\lambda}' = (\lambda', 1 - \lambda')$. □

6.8.3 Proof of Theorem 8

Since the first measurement stage is the localisation of the true state by the estimate $\tilde{\rho}$, we write the risk as a sum of two terms

$$\begin{aligned} R(\rho, \hat{\rho}_n) &= \mathbb{E} \left[D_B(\rho, \hat{\rho}_n)^2 \right] \\ &= \mathbb{P}(|\tilde{r}| \leq \delta) \cdot \mathbb{E} \left[D_B(\rho, \hat{\rho}_n)^2 \mid |\tilde{r}| \leq \delta \right] + \mathbb{P}(|\tilde{r}| > \delta) \cdot \mathbb{E} \left[D_B(\rho, \hat{\rho}_n)^2 \mid |\tilde{r}| > \delta \right] \end{aligned} \quad (6.73)$$

The expectation is taken over the measurement outcomes given the true state ρ . The final estimate $\hat{\rho}_n$ is defined by either (6.39) or (6.41) depending on the preliminary estimate $\tilde{\rho}$. Specifically, if the estimate $\tilde{\rho}$ is within a ball of radius $\delta > 0$ of the fully mixed state, we perform standard tomographic measurements on the remaining copies of the state. The final estimate is then the maximum likelihood (ML) estimate given by (6.39), while in the other instance the technology of LAN is utilised and the final estimate is (6.41). In order to make the difference between the two estimators explicit, we let $\hat{\rho}_n^1$ denote the final estimator in the case when $|\tilde{r}| \leq \delta$ and $\hat{\rho}_n^2$ be the final estimate

when $|\tilde{r}| > \delta$. Therefore, we have

$$\begin{aligned} R(\rho, \hat{\rho}_n) &= \mathbb{P}(|\tilde{r}| \leq \delta) \cdot \mathbb{E} \left[D_B(\rho, \hat{\rho}_n^1)^2 \mid |\tilde{r}| \leq \delta \right] + \mathbb{P}(|\tilde{r}| > \delta) \cdot \mathbb{E} \left[D_B(\rho, \hat{\rho}_n^2)^2 \mid |\tilde{r}| > \delta \right] \\ &= R_1 + R_2. \end{aligned} \quad (6.74)$$

We consider the contribution to the risk from the term R_1 first. Let B_1 be a ball of radius $\delta + n^{-1/2+\epsilon_2}$ around the centre of the Bloch sphere. When the true state $\rho \notin B_1$, we note that the probability $\mathbb{P}(|\tilde{r}| \leq \delta)$ goes to zero exponentially fast in n . This is because the estimate $\tilde{\rho}$ lies within a ball of radius $O(n^{-1/2+\epsilon_2})$ around the true state with high probability (Lemma 5). This along with the fact that the square Bures distance is bounded as $D_B(\sigma, \pi) \leq 2$ for any pair of density matrices σ, π implies that when $\rho \notin B_1$, the term R_1 can be neglected. However, when $\rho \in B_1$, the term R_1 has a non zero contribution and is written as

$$R_1 = \begin{cases} \mathbb{E} \left[D_B(\rho, \hat{\rho}_n^1)^2 \mid |\tilde{r}| \leq \delta \right] \cdot \mathbb{P}(|\tilde{r}| \leq \delta), & \rho \in B_1 \\ o(1) & \rho \notin B_1 \end{cases} \quad (6.75)$$

The term R_2 can be treated similarly. Let B_2 be a ball of radius $\delta - n^{-1/2+\epsilon_2}$ around the centre of the Bloch sphere. As the probability $\mathbb{P}(|\tilde{r}| > \delta)$ decays exponentially if $\rho \in B_2$, the term R_2 is relevant only when $\rho \notin B_2$

$$R_2 = \begin{cases} \mathbb{E} \left[D_B(\rho, \hat{\rho}_n^2)^2 \mid |\tilde{r}| > \delta \right] \cdot \mathbb{P}(|\tilde{r}| > \delta), & \rho \notin B_2 \\ o(1) & \rho \in B_2 \end{cases} \quad (6.76)$$

Substituting (6.75), (6.76) in (6.74) we see that the minimax risk is bounded from above as

$$\begin{aligned} \limsup_{n \rightarrow \infty} \sup_{\rho} nR(\rho, \hat{\rho}_n) &\leq \limsup_{n \rightarrow \infty} \max \left\{ \sup_{\rho \in B_1} n \mathbb{E} \left[D_B(\rho, \hat{\rho}_n^1)^2 \mid |\tilde{r}| \leq \delta \right], \right. \\ &\quad \left. \sup_{\rho \notin B_2} n \mathbb{E} \left[D_B(\rho, \hat{\rho}_n^2)^2 \mid |\tilde{r}| > \delta \right] \right\} \end{aligned} \quad (6.77)$$

Case 1 : $\rho \in B_1$ and $|\tilde{r}| \leq \delta$

We now evaluate the risk when the state ρ is in B_1 while the estimate $\tilde{\rho}$ is within a ball of radius $\delta > 0$ around the fully mixed state. The final estimate $\hat{\rho}_n^1$ is the ML estimate and given by (6.39). The outcomes from the $n/3$ repeated measurements in a setting σ_i are i.i.d, this implies that the ML estimate of the Bloch vector parameters r_x, r_y, r_z from the outcomes of the standard

tomographic measurements are asymptotically Gaussian in distribution

$$\lim_{n \rightarrow \infty} \sqrt{n}(\hat{r}_i - r_i) = \mathcal{N}(0, I(\rho)^{-1}), \quad i = x, y, z \quad (6.78)$$

where the covariance matrix is the inverse of the Fisher information matrix $I(\rho)$. The elements of the matrix $I(\rho)$ are defined for each $i, j \in \{x, y, z\}$ as

$$I(\rho)_{i,j} = \frac{1}{(1+r_i)(1-r_i)} \delta_{i,j}. \quad (6.79)$$

The local expansion of the square Bures distance for states away from the boundary of the Bloch sphere is quadratic in the Bloch vector components

$$D_B(\rho, \hat{\rho}_n^1)^2 = (\mathbf{r} - \hat{\mathbf{r}})^T G(\mathbf{r} - \hat{\mathbf{r}}) + O(\|\mathbf{r} - \hat{\mathbf{r}}\|^3), \quad (6.80)$$

where G is the weight matrix of the square Bures distance

$$G_{j,k} := \frac{1}{4} \left(1 + \frac{r_i^2}{(1-|r|^2)} \right) \delta_{j,k}. \quad (6.81)$$

The asymptotic behaviour of the ML estimator (6.39) together with this local expansion of the square Bures distance, implies that the risk of the ML estimate scales as follows for large n

$$\mathbb{E} [D_B(\rho, \hat{\rho}_n^1)^2] = \frac{1}{n} \text{Tr} (I(\rho)^{-1} G) + o(n^{-1}) \quad (6.82)$$

It is easy to see that $I(\rho) \geq \mathbf{1}$, and therefore we have that asymptotically the minimax Bures risk is upper bounded by

$$\limsup_{n \rightarrow \infty} \sup_{\rho \in B_1} n \mathbb{E} [D_B(\rho, \hat{\rho}_n^1)^2] \leq \frac{3}{4} \left(1 + \frac{\delta}{1-\delta} \right) \quad (6.83)$$

where we used the fact that $|r| \leq \delta + O(n^{-1/2+\epsilon_2})$.

Case 2: $|\tilde{r}| > \delta$, $\rho \notin B_2$

We now consider the case when $\tilde{\rho}$ is away from the fully mixed state. Since the state $\tilde{\rho}$ is within the ball of radius $O(n^{-1/2+\epsilon_2})$ of the true state $\rho \notin B_2$, we consider the local parametrisation of the states $\boldsymbol{\theta} = (\lambda, \mathbf{w})$, and perform the secondary measurements on the joint state $\rho_n^\theta = \rho_\theta^{\otimes n}$ of n qubits. Using the

approximation of the square Bures distance in (6.11), the risk is expressed as

$$\begin{aligned} \mathbb{E} \left[D_B(\rho, \hat{\rho}_n^2)^2 \right] &= \mathbb{E} \left[D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n)^2 + \frac{1}{4} \frac{(1-2\lambda)(1-2\hat{\lambda})}{\left(\sqrt{(1-\lambda)(1-\hat{\lambda}_n)} + \sqrt{\lambda\hat{\lambda}_n} \right)} \Phi^2 \right] + O(n^{-2}) \\ &\leq \mathbb{E} \left[D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n)^2 \right] + \frac{1}{n} \mathbb{E} \left[(u - \hat{u}_n)^2 + (v - \hat{v}_n)^2 \right] + O(n^{-2}). \end{aligned} \quad (6.84)$$

The second term on the right corresponding to the rotation parameters has been evaluated in [57]. Since LAN holds in the limit of large n , the problem of estimating the rotation parameters is translated to one of determining the displacement of a Gaussian state $\phi^{\mathbf{w}}$. The heterodyne measurement is known to be the optimal measurement in this case [57, 59]. The estimation of these parameters is described in detail in [57], and here we only note that both $\mathbb{E} \left[(u - \hat{u}_n)^2 \right]$ and $\mathbb{E} \left[(v - \hat{v}_n)^2 \right]$ are bounded from above by $(1-\lambda)/(2(1-2\lambda)^2) \leq 1/2$. Substituting these values, the minimax risk becomes

$$\begin{aligned} \limsup_{n \rightarrow \infty} \sup_{\rho \notin B_2} nR(\rho, \hat{\rho}_n^2) &= \limsup_{n \rightarrow \infty} \sup_{\boldsymbol{\theta}} n\mathbb{E} \left[D_B(\rho, \hat{\rho}_n^2) \right] \\ &\leq \limsup_{n \rightarrow \infty} \sup_{\lambda} n\mathbb{E} \left[D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n)^2 \right] + 1. \end{aligned} \quad (6.85)$$

The minimax risk is upper bounded by 1 plus the minimax risk of the square Hellinger distance. We now deal with this term. The expectation is taken over the probability distribution $p_{n,\lambda}(j)$ defined in equation (6.27),

$$\begin{aligned} \mathbb{E} \left[D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n)^2 \right] &= \sum_{j=0,1/2}^{n/2} D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n(j))^2 B_{n,\lambda}(n/2-j) \times K(j, n, \lambda) \\ &\leq \sum_{j=0,1/2}^{n/2} D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n(j))^2 B_{n,\lambda}(n/2-j) \times |1 - K(j, n, \lambda)| \\ &\quad + \sum_{j=0,1/2}^{n/2} D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n(j))^2 B_{n,\lambda}(n/2-j) = E_1 + E_2. \end{aligned} \quad (6.86)$$

We now consider the term E_1 separately. We split the sum in E_1 over the values of $j \in \mathcal{J}_n$, and $j \notin \mathcal{J}_n$, where \mathcal{J}_n is interval defined in (6.30). Thus, we have

$$\begin{aligned}
E_1 &= \sum_{j \in \mathcal{J}_n} D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n(j))^2 B_{n,\lambda}(n/2-j) \times |1 - K(j, n, \lambda)| \\
&\quad + \sum_{j \notin \mathcal{J}_n} D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n(j))^2 B_{n,\lambda}(n/2-j) \times |1 - K(j, n, \lambda)| \\
&\leq \max_{j \in \mathcal{J}_n} |1 - K(j, n, \lambda)| \sum_{j \in \mathcal{J}_n} D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n(j))^2 B_{n,\lambda}(n/2-j) \\
&\quad + \max_{j \notin \mathcal{J}_n} |1 - K(j, n, \lambda)| \sum_{j \notin \mathcal{J}_n} D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n(j))^2 B_{n,\lambda}(n/2-j) \\
&\leq O(n^{-1/2+\epsilon_3}) \sum_{j \in \mathcal{J}_n} D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n(j))^2 B_{n,\lambda}(n/2-j) \\
&\quad + \max_{j \notin \mathcal{J}_n} |1 - K(j, n, \lambda)| \sum_{j \notin \mathcal{J}_n} 2B_{n,\lambda}(n/2-j). \quad (6.87)
\end{aligned}$$

In the last inequality, we used the fact that $K(j, n, \lambda) = 1 + O(n^{-1/2+\epsilon_3})$ on the values of $j \in \mathcal{J}_n$, and that $D_H(\mathbf{p}, \mathbf{q})^2 \leq 2$ for any pair of probability distributions \mathbf{p}, \mathbf{q} . The value $\max_{j \notin \mathcal{J}_n} |1 - K(j, n, \lambda)|$ is uniformly bounded for λ away from $1/2$. This along with the fact that the mass of the binomial distribution $B_{n,\lambda}(n/2-j)$ is concentrated on values of $j \in \mathcal{J}_n$ implies that the second term in (6.87) goes to zero exponentially fast in n . Therefore,

$$E_1 \leq O(n^{-1/2+\epsilon_3}) \sum_{j \in \mathcal{J}_n} D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n(k))^2 B_{n,\lambda}(k). \quad (6.88)$$

Substituting this back in (6.86), we have that the Hellinger risk is upper bounded by

$$\begin{aligned}
\mathbb{E} \left[D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n)^2 \right] &\leq O(n^{-1/2+\epsilon_3}) \sum_{j \in \mathcal{J}_n} D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n(j))^2 B_{n,\lambda}(n/2-j) \\
&\quad + \sum_{j=0,1/2}^{n/2} D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n(j))^2 B_{n,\lambda}(n/2-j) \\
&\leq O(n^{-1/2+\epsilon_3}) \sum_{k=0}^n D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n(k))^2 B_{n,\lambda}(k) + \sum_{k=0}^n D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n(k))^2 B_{n,\lambda}(k) \\
&= \left(1 + O(n^{-1/2+\epsilon_3})\right) \mathbb{E}_{\text{Binom}} \left[D_H(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n(k))^2 \right] \quad (6.89) \\
&\leq \left(1 + O(n^{-1/2+\epsilon_3})\right) \mathbb{E}_{\text{Binom}} \left[D_{KL}(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n(k)) \right]. \quad (6.90)
\end{aligned}$$

In the second line we expand the sums over the entire support of the binomial distribution $k \in \{0, \dots, n\}$, and let $\mathbb{E}_{\text{Binom}}$ mark expectation with respect to the binomial distribution. The last inequality employs the inequality between the square Hellinger distance and the Kullback–Leibler (KL) distance, defined as

$D_{KL}(\mathbf{p}, \mathbf{q}) := p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q}$. Substituting (6.90) in (6.85), and noting that the ‘add-beta’ estimator $\hat{\lambda}_n$ was defined in (6.40) over the full support, we have

$$\begin{aligned} \limsup_{n \rightarrow \infty} \sup_{\rho \notin B_2} nR(\rho, \hat{\rho}_n^2) &\leq \limsup_{n \rightarrow \infty} \sup_{\lambda} n\mathbb{E}_{\text{Binom}} \left[D_{KL}(\boldsymbol{\lambda}, \hat{\boldsymbol{\lambda}}_n) \right] \left(1 + O(n^{-1/2+\epsilon_3}) \right) + 1 \\ &\leq \frac{1}{2} + 1. \end{aligned} \quad (6.91)$$

The rate follows from the fact that the ‘add beta’ estimator of the binomial parameter, defined in equation (6.40), is known to be asymptotically minimax for the KL risk, achieving the rate $\frac{1}{2n}(1+o(1))$ [22, 23]. Comparing this rate with the one in (6.83), we see that $3/2$ is the larger value provided $\delta < 1/2$. This gives the upper bound stated in Theorem 8. \square

6.8.4 Expansion of Quantum Relative Entropy

We derive the expansion of the quantum relative entropy $S(\rho||\rho') = \text{Tr}[\rho(\log \rho - \log \rho')]$. For qubits the relative entropy between two states can be expressed in terms of the Bloch vector components as

$$\begin{aligned} S(\rho||\rho') &= \frac{1}{2} \left[\log(1 - |\mathbf{r}|^2) - \log(1 - |\mathbf{r}'|^2) + |\mathbf{r}| \log \left(\frac{1 + |\mathbf{r}|}{1 - |\mathbf{r}|} \right) \right. \\ &\quad \left. - |\mathbf{r}| \cos \Phi \log \left(\frac{1 + |\mathbf{r}'|}{1 - |\mathbf{r}'|} \right) \right], \end{aligned} \quad (6.92)$$

where Φ is the angle between the Bloch vectors of the two states. For sufficiently close states, the angle Φ is small and the cosine term in the above equation can be expanded as

$$\begin{aligned} S(\rho||\rho') &= \frac{1}{2} \left[\log(1 - |\mathbf{r}|^2) - \log(1 - |\mathbf{r}'|^2) + |\mathbf{r}| \log \left(\frac{1 + |\mathbf{r}|}{1 - |\mathbf{r}|} \right) \right. \\ &\quad \left. - |\mathbf{r}| \log \left(\frac{1 + |\mathbf{r}'|}{1 - |\mathbf{r}'|} \right) \right] + \frac{|\mathbf{r}|}{4} (\Phi^2 + O(\Phi^4)) \log \left(\frac{1 + |\mathbf{r}'|}{1 - |\mathbf{r}'|} \right). \end{aligned} \quad (6.93)$$

Using the fact that $|\mathbf{r}| = 1 - 2\lambda$ and simplifying, we get

$$\begin{aligned}
S(\rho||\rho') &= \frac{1}{2} \left[\log \left(\frac{\lambda(1-\lambda)}{\lambda'(1-\lambda')} \right) - (1-2\lambda) \log \left(\frac{\lambda(1-\lambda')}{\lambda'(1-\lambda)} \right) \right] \\
&\quad + \frac{|\mathbf{r}|}{4} (\Phi^2 + O(\Phi^4)) \log \left(\frac{1+|\mathbf{r}'|}{1-|\mathbf{r}'|} \right) \\
&= \lambda \log \left(\frac{\lambda}{\lambda'} \right) + (1-\lambda) \log \left(\frac{1-\lambda}{1-\lambda'} \right) + \frac{1-2\lambda}{4} (\Phi^2 + O(\Phi^4)) \log \left(\frac{1-\lambda'}{\lambda'} \right) \\
&\hspace{15em} (6.94)
\end{aligned}$$

$$= D_{KL}(\boldsymbol{\lambda}, \boldsymbol{\lambda}') + \frac{1-2\lambda}{4} (\Phi^2 + O(\Phi^4)) \log \left(\frac{1-\lambda'}{\lambda'} \right), \quad (6.95)$$

where $D_{KL}(\boldsymbol{\lambda}, \boldsymbol{\lambda}')$ is the Kullback-Leibler distance between the two distributions $\boldsymbol{\lambda} = (\lambda, 1-\lambda)$ and $\boldsymbol{\lambda}' = (\lambda', 1-\lambda')$.

□

Chapter 7

Comparison of estimation methods in quantum state tomography: a simulation study

7.1 Introduction

In previous chapters we considered estimation schemes such as simple linear inversion and maximum likelihood with standard full tomography in chapter 2, ‘compressive’ tomography with incomplete measurements in chapter 4, 5, adaptive tomography with local and collective measurement designs in chapter 6. In this chapter we consider finite dimensional full state tomography in the case of composite systems such as trapped ions. This is the by now familiar *multiple ions tomography* (MIT) setup from earlier chapters. As before, we consider two measurement designs in the MIT setup - the Pauli bases measurements and the more general random bases measurements. The tomographic procedure involves repeating each measurement on identical copies of the state, resulting in a dataset of outcome counts from which the density matrix is reconstructed. Since the measurements are performed on n copies of the state, and no evolution of the state is assumed, the tomographic problem can be recast as a classical parameter estimation problem familiar in statistics, where the parameters of the density matrix are to be determined. Several familiar and well-studied estimators can be employed for this purpose. In chapter 2 we introduced examples of two such estimators - the *Least Squares Estimator* (LSE) and the commonly used *Maximum Likelihood Estimator* (MLE).

The various estimators in general produce differing estimates of the state, and we may wish to determine which estimation strategy is ‘best’. An obvious method of comparing the performance of estimators is in terms of their estimation errors for a given state and measurement design. These errors may be quantified in terms of various loss functions such as the Frobenius norm, trace norm and the fidelity based Bures distance, which measure the ‘distance’

between the true state and the estimates. Since these estimates necessarily depend on the counts data obtained from the experiment, it is common to consider the expected loss (or risk) of an estimator, where the expectation is taken with respect to all possible outcome datasets.

In the case of certain efficient estimators, asymptotic theory informs us that their risk achieves the *Cramér-Rao lower bound* (CRLB) in the limit $n \rightarrow \infty$, and they are therefore ‘optimal’ for sufficiently large n . The MLE is known to be asymptotically optimal in this sense, and its risk decreases at a rate of $O(1/n)$ for most states. However, as we have seen in chapter 3 and 6, this asymptotic theory does not meaningfully apply when the underlying parameters are close to the boundary of the parameter space. In terms of the density matrix, this means that for states of high purity and very small eigenvalues, such estimators cannot be said to be ‘optimal’. Additionally, since in experiments only a finite number of copies of the state are available as a resource, it is not immediately clear what the implications of asymptotic theory are in practice. For instance, estimators that are asymptotically unbiased may not be so for finite n [106].

An additional difficulty arises from the fact that the performance of estimators can be studied in terms of several different loss functions, and therefore the answer to the question of how well an estimator performs depends on this choice. These loss functions usually satisfy certain smoothness conditions and are locally quadratic for most states. However, as we have seen in chapters 3, 5 and 6, the fidelity based loss functions are not quadratic near the boundary, while other loss functions such as the Frobenius norm continue to be so. This implies that for states near the boundary ‘efficient’ estimators can be shown to achieve the CRLB in terms of Frobenius norm but the asymptotic theory is not valid for fidelity based loss functions [106]. This feature at the boundary is known to give rise to a poor scaling of $O(1/\sqrt{n})$ of the risk for fidelity based loss functions [89].

The aim of this chapter is to better understand and compare tomographic estimators in a systematic way. To this end we present results from an extensive and systematic simulation study comparing the performance of several estimators across a range of different variables - types of states, ranks, measurement design, number of copies of the state as a resource and the number of qubits. Along with the commonly used and well studied estimators like the MLE, a few new estimators are introduced and their performance investigated. We quantify the errors of the estimators in terms of several loss functions such as

the Bures distance, trace and Frobenius norms. We analyse and interpret the simulation results, highlighting the performance of the estimators for states near the boundary of the parameter space and their differing behaviour across the several loss functions. The performance of all of the estimators are studied in both the asymptotic regime with large n , and realistic values of n typically used in tomographic experiments [62]. Along with the numerical results presented in this chapter we introduce two web-based applications that makes available all of the estimators studied here for use online. These applications will allow the user to reproduce the results presented here, and more importantly they enable the user to perform their own simulations with arbitrary states, and thus complement the results presented in this chapter.

The chapter is structured in the following way. In section 7.2 we begin with a brief review of the MIT setup and introduce the tomographic problem. In section 7.3 we define all of the estimators considered in the simulation study. These include commonly used estimators such as the LSE, MLE, positive LSE along with some new estimators such as the *Thresholded Generalised Least Squares* (TGLS) and the *Positive Generalised Least Squares* (PGLS) estimators. We then define the Fisher information matrix and review the essential results of asymptotic theory in section 7.4. In sections 7.5 and 7.6 we describe, present and analyse the results of the simulation study. Finally in section 7.7 we introduce and demonstrate the web-based applications.

7.2 A review of multiple ions tomography

In the MIT setup, the aim of quantum tomography is to estimate an unknown state of N ions from the outcomes of measurements performed on identically prepared systems. Let $\rho \in \mathbb{S}_d$ be the density matrix associated with the unknown state, where $d = 2^N$ is the dimension of the associated Hilbert space \mathcal{H}^d . We consider measurements of two types - the standard tomographic measurements in the Pauli basis, and measurements that are drawn randomly from the uniform measure over orthonormal bases (ONB). In the case of the Pauli basis, one measures an arbitrary Pauli observable $\sigma_x, \sigma_y, \sigma_z$ on each of the N ions simultaneously. Therefore, each measurement is labelled by a sequence $\mathbf{s} = (s_1, \dots, s_N) \in \{x, y, z\}^N$, and there are 3^N possible measurement bases. In the uniformly random measurement set up, a measurement can be implemented by first rotating the state ρ by a random unitary $U \in M(\mathbb{C}^d)$, after which each ion is measured in the σ_z eigenbasis.

Let $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_k\}$ be the measurement design consisting of k measurement settings. In the case of the Pauli set up the total number $k = 3^N$, while in the random measurement setup the number of settings measured can be chosen freely. A measurement in a particular setting produces a ± 1 outcome from each ion, and we let $\mathbf{o} \in \{+1, -1\}^N$ be a vector record of outcomes from each of the N ions. The probability of obtaining a particular outcome \mathbf{o} is given by $p_\rho(\mathbf{o}|\mathbf{s}) := \text{Tr}(\rho P_{\mathbf{o}}^{\mathbf{s}})$, where the one-dimensional projection matrix is given by

$$P_{\mathbf{o}}^{\mathbf{s}} = |e_{o_1}^{s_1}\rangle\langle e_{o_1}^{s_1}| \otimes \dots \otimes |e_{o_N}^{s_N}\rangle\langle e_{o_N}^{s_N}|. \quad (7.1)$$

For each setting \mathbf{s} , measurements are repeated on m identical copies of the state, and the counts of the outcomes $N(\mathbf{o}|\mathbf{s})$ are recorded, where $N(\mathbf{o}|\mathbf{s})$ represents the number of times a given outcome record \mathbf{o} is observed given that measurements were performed in a chosen setting \mathbf{s} . The total number of quantum samples used is therefore $n = m \times k$. The resulting dataset \mathcal{D} of counts is a $2^N \times k$ table whose columns are independent and contain all the counts in a given setting. The probability of observing a given dataset of counts is given by the product of multinomials

$$p_\rho(\mathcal{D}|\mathcal{S}) = \prod_{\mathbf{s}} \frac{m!}{\prod_{\mathbf{o}} N(\mathbf{o}|\mathbf{s})!} \prod_{\mathbf{o}} p_\rho(\mathbf{o}|\mathbf{s})^{N(\mathbf{o}|\mathbf{s})} \quad (7.2)$$

Our goal is the statistical reconstruction of the density matrix from this dataset of counts. As described in the introduction, there are many estimators known in the literature for this purpose, and in this chapter we consider several of the more commonly used estimators. We also introduce and define new estimators such as the *Threshold Generalised Least Squares* (TGLS) and the *Positive Generalised Least Squares* (PGLS).

7.3 Estimators

7.3.1 Maximum Likelihood (ML)

We describe the estimators considered over the next few subsections, and begin with the commonly used *Maximum Likelihood Estimator* (MLE)

$$\hat{\rho}_{\text{ML}} := \arg \max_{\tau \in \mathbb{S}_d} p_\tau(\mathcal{D}|\mathcal{S}), \quad (7.3)$$

where the minimisation is over the space of $d \times d$ density matrices. It is more common and convenient to maximise the natural logarithm of the likelihood function instead. Discarding the constant factorial terms in (7.2), we have the following form of the estimator

$$\hat{\rho}_{\text{ML}} := \arg \max_{\tau \in \mathbb{S}_d} \sum_{\mathbf{o}, \mathbf{s}} N(\mathbf{o}|\mathbf{s}) p_{\tau}(\mathbf{o}|\mathbf{s}). \quad (7.4)$$

The MLE is commonly used in quantum tomography [13, 19, 71, 101, 50], with several methods proposed in the literature such as the iterative algorithm proposed by Hradil [101]. We leave a discussion of our specific implementation until a later section. The MLE is known to possess certain desirable statistical properties. Specifically, under certain regularity conditions it is known to be asymptotically normal, meaning that in the limit of large m it has a Gaussian distribution with covariance matrix given by the inverse of the classical Fisher information matrix [118, 61]. While this is true for most states $\rho \in \mathbb{S}_d$, the asymptotic theory does not hold for states at the boundary of \mathbb{S}_d , and as we shall see in a later section this is because the MLE is constrained to produce only state estimates, while the maximum of the likelihood may lie outside the convex space \mathbb{S}_d . Anticipating a later discussion we note that as long as $p_{\rho}(\mathbf{o}|\mathbf{s}) > 0$ for all \mathbf{o}, \mathbf{s} the asymptotic theory does apply for the ‘unconstrained’ version of the MLE, where the estimates are not constrained to lie in \mathbb{S}_d [61].

7.3.2 Least Squares (LS)

There are a number of techniques that can be used to estimate the state from the outcome statistics in the dataset. An example of such a method is the linear *Least Squares* (LS) estimator [95, 25]. To better describe the linear estimation problem, we consider the true probability vector

$$\mathbf{y} = (p(\mathbf{o}_1|\mathbf{s}_1), \dots, p(\mathbf{o}_d|\mathbf{s}_1), \dots, p(\mathbf{o}_d|\mathbf{s}_d))^T \in \mathbb{R}^{kd}. \quad (7.5)$$

In the standard basis, each element of this vector can be expressed in terms of the density matrix elements and the corresponding one-dimensional projections as

$$p(\mathbf{o}|\mathbf{s}) = 2 \sum_{j>i} \text{Re}(\rho_{ij}) \text{Re}(P_{\mathbf{o}}^{\mathbf{s}})_{i,j} + 2 \sum_{j>i} \text{Im}(\rho_{ij}) \text{Im}(P_{\mathbf{o}}^{\mathbf{s}})_{i,j} + \sum_i^d \rho_{ii} (P_{\mathbf{o}}^{\mathbf{s}})_{i,i} \quad (7.6)$$

We now choose to parameterise the state by the elements of the density matrix ρ that appear in the above equation. Let $\boldsymbol{\beta} \in \mathbb{R}^{d^2}$ be a vector of parameters defined as

$$\boldsymbol{\beta} := \left(\operatorname{Re}\rho_{1,2}, \dots, \operatorname{Re}\rho_{d-1,d}, \operatorname{Im}\rho_{1,2}, \dots, \operatorname{Im}\rho_{d-1,d}, \rho_{1,1}, \dots, \rho_{d,d} \right)^T. \quad (7.7)$$

With this choice of parameterisation (7.6) can be expressed as an inner product between vectors $p(\mathbf{o}|\mathbf{s}) = X(\mathbf{o}|\mathbf{s})^T \boldsymbol{\beta}$, where

$$X(\mathbf{o}|\mathbf{s})^T := \left(2\operatorname{Re}(P_{\mathbf{o}}^{\mathbf{s}})_{1,2}, \dots, 2\operatorname{Re}(P_{\mathbf{o}}^{\mathbf{s}})_{d-1,d}, 2\operatorname{Im}(P_{\mathbf{o}}^{\mathbf{s}})_{1,2}, \dots, \right. \\ \left. 2\operatorname{Im}(P_{\mathbf{o}}^{\mathbf{s}})_{d-1,d}, (P_{\mathbf{o}}^{\mathbf{s}})_{1,1}, \dots, (P_{\mathbf{o}}^{\mathbf{s}})_{d,d} \right).$$

This notion allows us to express the tomographic system of equations in matrix form as

$$\mathbf{y} = X\boldsymbol{\beta}, \quad (7.8)$$

where X is a $kd \times d^2$ matrix whose rows are given by $X(\mathbf{o}|\mathbf{s})^T$ for each pair of outcome \mathbf{o} and setting \mathbf{s} . In an experimental set up, we do not have access to the true probability vector. Instead from the $d \times k$ dataset of counts, we have access to the empirical probabilities $f(\mathbf{o}|\mathbf{s}) := N(\mathbf{o}|\mathbf{s})/m$, whose expectations are $\mathbb{E}f(\mathbf{o}|\mathbf{s}) = p(\mathbf{o}|\mathbf{s})$. Replacing the probability vector \mathbf{y} by the vector of empirical frequencies we have

$$\mathbf{f} = X\boldsymbol{\beta} + \boldsymbol{\epsilon}, \quad (7.9)$$

where $\boldsymbol{\epsilon} \in \mathbb{R}^{dk}$ is a mean zero vector of statistical noise. The least-squares solution to the above system of equations is defined as the minimizer of the following optimisation problem

$$\hat{\boldsymbol{\beta}} := \arg \min_{\boldsymbol{\tau} \in \mathbb{R}^{d^2}} \|\mathbf{f} - X\boldsymbol{\tau}\|^2, \quad (7.10)$$

and has the well known explicit form $\hat{\boldsymbol{\beta}} = (X^T X)^{-1} \cdot X^T \cdot \mathbf{f}$. The final estimate of the density matrix $\hat{\rho}_{\text{LS}}$ is then constructed from the estimated parameter vector $\hat{\boldsymbol{\beta}}$. We note that the least squares estimator produces an estimate that is not necessarily a density matrix.

7.3.3 Generalised Least Squares (GLS)

The error term in equation (7.9) is such that given the measurements design \mathcal{S} it has mean zero $\mathbb{E}[\boldsymbol{\epsilon}|\mathcal{S}] = 0$, and a covariance matrix we denote as $\text{Cov}[\boldsymbol{\epsilon}|\mathcal{S}] = \Omega$. The covariance is not a diagonal matrix because for each given measured setting \mathbf{s} , the errors corresponding to the d possible outcomes are correlated. However, as the outcomes from different settings are independent, Ω is a $dk \times dk$ block diagonal matrix where each block is given by the $d \times d$ covariance matrix for a given setting \mathbf{s}

$$\text{Cov}[\boldsymbol{\epsilon}_{\mathbf{s}}|\mathbf{s}] = \begin{cases} p(\mathbf{o}_i|\mathbf{s})(1 - p(\mathbf{o}_i|\mathbf{s})), & i = j \\ -p(\mathbf{o}_j|\mathbf{s})p(\mathbf{o}_i|\mathbf{s}), & i \neq j. \end{cases} \quad (7.11)$$

In a situation where there is correlation in the error terms and the covariance matrix Ω is known beforehand, it is more natural to consider the *Generalised Least Squares* (GLS) [12] over the ordinary Least Squares (LS) estimator

$$\hat{\boldsymbol{\beta}}_{\text{GLS}} = \arg \min_{\boldsymbol{\tau} \in \mathbb{R}^{d^2}} (\mathbf{f} - X\boldsymbol{\tau})^T \Omega^{-1} (\mathbf{f} - X\boldsymbol{\tau}) \quad (7.12)$$

$$= \arg \min_{\boldsymbol{\tau} \in \mathbb{R}^{d^2}} \|\Omega^{-1/2} (\mathbf{f} - X\boldsymbol{\tau})\|^2, \quad (7.13)$$

which has the following explicit form

$$\hat{\boldsymbol{\beta}}_{\text{GLS}} = (X^T X)^{-1} X^T \Omega^{-1} \mathbf{f}. \quad (7.14)$$

The GLS estimator is known to be unbiased, efficient and asymptotically normal [12]

$$\sqrt{m}(\hat{\boldsymbol{\beta}}_{\text{GLS}} - \boldsymbol{\beta}) \rightarrow \mathcal{N}(\mathbf{0}, (X^T \Omega^{-1} X)^{-1}). \quad (7.15)$$

However in practice the covariance matrix is unknown since it depends on the true probabilities determined by the unknown state ρ . We therefore propose to use an estimate $\hat{\Omega}$ of the covariance matrix instead. We shall describe the computation of this estimate in a later section. In addition to this difficulty, the matrix Ω as defined is singular due the fact that for each setting the frequencies are constrained to sum to one $\sum_{\mathbf{o}} f(\mathbf{o}|\mathbf{s}) = 1$, and the system of equations (7.9) is overdetermined. This is easily remedied by removing one equation for each setting from the linear system. Let $\tilde{\mathbf{f}}$ be a truncated vector of frequencies, where for each setting \mathbf{s} a randomly selected row corresponding to one outcome

is omitted. For example we may have

$$\tilde{\mathbf{f}} = (f(\mathbf{o}_2|\mathbf{s}_1), \dots, f(\mathbf{o}_d|\mathbf{s}_1), f(\mathbf{o}_1|\mathbf{s}_2), f(\mathbf{o}_3|\mathbf{s}_2), \dots, f(\mathbf{o}_{d-1}|\mathbf{s}_d))^T \in \mathbb{R}^{kd-k}. \quad (7.16)$$

Similarly let \tilde{X} and $\tilde{\boldsymbol{\epsilon}}$ be truncated versions of the design matrix and error vector such that $\tilde{\mathbf{f}} = \tilde{X}\boldsymbol{\beta} + \tilde{\boldsymbol{\epsilon}}$. Also let $\tilde{\Omega} = \text{Cov}[\tilde{\boldsymbol{\epsilon}}|\mathcal{S}]$. The GLS estimate is now given by

$$\hat{\boldsymbol{\beta}}_{\text{GLS}} = (\tilde{X}^T \tilde{X})^{-1} \tilde{X}^T (\hat{\tilde{\Omega}})^{-1} \tilde{\mathbf{f}}, \quad (7.17)$$

where $\hat{\tilde{\Omega}}$ is the estimated truncated covariance matrix. We denote $\hat{\rho}_{\text{GLS}}$ as estimate of the density matrix constructed from its vectorised form $\hat{\boldsymbol{\beta}}_{\text{GLS}}$.

7.3.4 Thresholded Least Squares (TLS)

The LSE suffers from the disadvantage that it does not necessarily produce a density matrix, i.e, a positive semi-definite estimate of trace one. Additionally as pointed out in [25], the LS estimator does a poor job of estimating zero eigenvalues of the true state, and may estimate them as negative, while it estimates the large non-zero eigenvalues fairly well. The *Thresholded Least Squares* (TLS) estimator proposed in [25], improves the LS estimate $\hat{\rho}_{\text{LS}}$ by setting the eigenvalues of $\hat{\rho}_{\text{LS}}$ that are below a certain statistical noise threshold to zero and rescaling the remaining eigenvalues so that they sum to one. The resulting estimate is therefore a density matrix. The choice of the statistical noise threshold is informed by the accuracy of the LS estimate, and a theoretical choice for this threshold is detailed in [25, 26]. The following algorithm 2 describes the thresholding and rescaling of the eigenvalues of the LS estimate.

The algorithm takes as its input the eigenvalues of the normalised LS estimate $\hat{\rho}_{\text{LS}}/\text{Tr}(\hat{\rho}_{\text{LS}})$ sorted in descending order. Then beginning with the smallest eigenvalue, the algorithm checks if it is above the noise threshold. If so, then the TLS estimate is simply $\hat{\rho}_{\text{TLS}} = \hat{\rho}_{\text{LS}}/\text{Tr}(\hat{\rho}_{\text{LS}})$. On the other hand if the eigenvalue is below the threshold, it is set to zero and the remaining eigenvalues are suitably rescaled. The algorithm then iteratively checks each of the remaining eigenvalues in ascending order. The final estimate $\hat{\rho}_{\text{TLS}}$ is constructed by replacing the eigenvalues of $\hat{\rho}_{\text{LS}}$ with these thresholded eigenvalues. Constructed in this way $\hat{\rho}_{\text{TLS}}$ is the closest state to the normalised LS estimate with positive (non-zero) eigenvalues larger than ν .

The noise threshold ν in the above procedure depends on the LS estimate $\hat{\rho}_{\text{LS}}$. Specifically, this threshold is proportional to an upper bound on the operator

Algorithm 2: Algorithm to threshold the eigenvalues of the LS estimate

Input : Eigenvalues of $\hat{\rho}_{\text{LS}}/\text{Tr}(\hat{\rho}_{\text{LS}})$ sorted in descending order
 $\hat{\lambda}_1 \geq \dots \geq \hat{\lambda}_d$, and noise threshold ν

Output : Eigenvalues of thresholded estimate $\hat{\rho}_{\text{TLS}}$

```

1 for  $p = 1, \dots, d$  do
2   | if  $\hat{\lambda}_{d-p+1} \geq \nu$  then
3   |   | STOP;
4   | else
5   |   |  $\hat{\lambda}_{d-p+1} \leftarrow 0$ ;
6   |   | for  $j = 1, \dots, d-p$  do
7   |   |   |  $\hat{\lambda}_j \leftarrow \hat{\lambda}_j + \frac{1}{d-p} (1 - \sum_{m=1}^{d-p} \hat{\lambda}_m)$ 
8   |   | end
9   | end
10 end
```

norm error of the LS estimate. Let $\nu := C\mu$, where $\|\hat{\rho}_{\text{LS}} - \rho\| \leq \mu$, and C is an absolute constant. In [25, 26], for the case of the standard tomographic setup and Pauli bases measurements, this upper bound μ is shown to be proportional to $\sqrt{\frac{4}{m} \log 2^{N+1}}$. In practice, we use cross validation to choose the value of this constant C , and therefore of the noise threshold ν from the data. We describe this cross validation procedure in section 7.5.2.

7.3.5 Thresholded Generalised Least Squares (TGLS)

This estimator is obtained by using the GLS estimate $\hat{\rho}_{\text{GLS}}$ instead of the LS estimate as a starting point for the thresholding procedure. The constant for thresholding is chosen in the same way by cross validation.

7.3.6 Positive Least Squares (PLS)

The *Positive Least Squares* (PLS) estimator is the restriction of the minimisation in (7.10) to hold only over $\tau \in \mathbb{R}^{d^2}$ that correspond to density matrices. Let us define the map $\mathcal{X} : M(\mathbb{C}^d) \mapsto \mathbb{R}^{kd}$ that is a one-to-one map from the space of Hermitian matrices to \mathbb{R}^{dk} . For a given outcome \mathbf{o} and setting \mathbf{s} the map is defined by

$$(\mathcal{X}(A))_{\mathbf{o}, \mathbf{s}} = \text{Tr}[AP_{\mathbf{o}}^{\mathbf{s}}] \quad A \in M(\mathbb{C}^d). \quad (7.18)$$

Therefore \mathcal{X} maps the density matrix $\rho \in \mathbb{S}_d$ to its vector of probabilities \mathbf{y} . If we parameterise the state ρ by the vector $\boldsymbol{\beta}$ as in (7.7), then the map \mathcal{X} can be

expressed in matrix notation and is precisely the design matrix X . The PLS estimator is defined in terms of this map as

$$\hat{\rho}_{\text{PLS}} = \arg \min_{\tau \in \mathbb{S}_d} \|\mathcal{X}(\tau) - \mathbf{f}\|^2. \quad (7.19)$$

7.3.7 Positive Generalised Least Squares (PGLS)

This is defined in much the same way as the PLS estimator and is a restriction of the minimisation in (7.13) to hold only over parameters that produce density matrices. In keeping with the discussion in the section 7.3.3, we consider the truncated frequency vector $\tilde{\mathbf{f}}$ as in (7.16), and a correspondingly truncated map $\tilde{\mathcal{X}}: \mathbb{S}_d \mapsto \mathbb{R}^{kd-k}$. The *Positive Generalised Least Squares* (PGLS) estimator is defined as

$$\hat{\rho}_{\text{PGLS}} = \arg \min_{\tau \in \mathbb{S}_d} \|\hat{\Omega}^{-1/2} (\tilde{\mathcal{X}}(\tau) - \tilde{\mathbf{f}})\|^2. \quad (7.20)$$

7.4 The Fisher information matrix

As we will be interested in the performance of the estimators for reasonably large values of m , it is meaningful to consider the Fisher information matrix and the implications of asymptotic theory. For reasonably large values of m , the estimates of ‘good’ estimators lie in a local neighbourhood around the true state, and therefore for practical purposes we may study the estimation errors only in terms of a local parameterisation of states. Let us denote $\boldsymbol{\theta} \in \Theta$ to be a vector of local parameters such that $\rho \equiv \rho_{\boldsymbol{\theta}}$, with Θ being an open set and the parameter $\boldsymbol{\theta}$ varying smoothly. We shall make this choice of parameterisation explicit shortly. The square Frobenius norm can be locally expanded in such parameters as

$$\|\rho_{\boldsymbol{\theta}} - \rho_{\boldsymbol{\theta} + \delta\boldsymbol{\theta}}\|_F^2 = (\delta\boldsymbol{\theta})^T G_F(\boldsymbol{\theta})(\delta\boldsymbol{\theta}) + O(\|\delta\boldsymbol{\theta}\|^2), \quad (7.21)$$

where $G_F(\boldsymbol{\theta})$ is a positive and constant weight matrix. It is known that estimators such as the MLE (under certain conditions) are efficient, meaning that for sufficiently large m and values of $\boldsymbol{\theta}$ within the parameter space, such estimators have a Gaussian distribution [118]

$$\sqrt{m}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \approx \mathcal{N}(\mathbf{0}, I(\boldsymbol{\theta})^{-1}), \quad (7.22)$$

where the covariance matrix $I(\boldsymbol{\theta})^{-1}$ is the inverse of the classical Fisher information matrix. The Fisher matrix depends on both the parameterisation and the measurement design. In particular (7.22) together with the locally quadratic expansion (7.21) gives

$$\sqrt{m}\mathbb{E}\left[\|\rho_{\boldsymbol{\theta}} - \rho_{\hat{\boldsymbol{\theta}}}\|_F^2\right] \approx \text{Tr}\left[G_F(\boldsymbol{\theta})I(\boldsymbol{\theta})^{-1}\right] \quad (7.23)$$

for such efficient estimators and reasonably large m . The right side of the above equation shows that the asymptotic mean square error can be characterised in terms of Fisher information, and in general this quantity serves as a meaningful lower bound on the mean errors in the limit of large m . Therefore we shall compare the performance of the estimators introduced in the previous sections against this theoretical ‘benchmark’ quantity.

In order to define the Fisher information matrix in the MIT setup, we first need to choose a local parametrisation of states. Let us assume that the unknown state ρ belongs to the space of rank r states $\mathbb{S}_r \subset M(\mathbb{C}^d)$, for a fixed rank $r \leq d$. In its own eigenbasis ρ is the diagonal matrix of eigenvalues $\text{Diag}(\lambda_1, \dots, \lambda_r, 0, \dots, 0)$, and any sufficiently close state is uniquely determined by its matrix elements in the first r rows (or columns). Intuitively this can be understood by noting that any rank- r state ρ' in the neighbourhood of ρ can be obtained by perturbing the eigenvalues and performing a small rotation of the eigenbasis; in the first order of approximations these transformation leave the $(d-r) \times (d-r)$ lower-right corner unchanged so

$$\rho' = \begin{pmatrix} \text{Diag}(\lambda_1, \dots, \lambda_r) & 0 \\ 0 & 0 \end{pmatrix} + \begin{pmatrix} \Delta_{diag} & \Delta_{off} \\ \Delta_{off}^\dagger & O(\|\Delta\|^2) \end{pmatrix}. \quad (7.24)$$

We therefore choose the (local) parametrisation $\rho' = \rho_{\boldsymbol{\theta}}$ with

$$\begin{aligned} \boldsymbol{\theta} &:= (\theta^{(d)}; \theta^{(r)}; \theta^{(i)}) \\ &= (\rho'_{2,2}, \dots, \rho'_{r,r}; \text{Re}\rho'_{1,2}, \dots, \text{Re}\rho'_{r,d}; \text{Im}\rho'_{1,2}, \dots, \text{Im}\rho'_{r,d}) \in \mathbb{R}^{2rd-r^2-1}, \end{aligned} \quad (7.25)$$

where, in order to enforce a trace-one normalisation, we constrain the first diagonal matrix element to be $\rho'_{1,1} = 1 - \sum_{i=2}^d \rho'_{i,i}$. With this parameterisation fixed, the classical Fisher information associated with a single chosen setting \mathbf{s}

is defined as

$$I(\rho|\mathbf{s})_{a,b} := \sum_{\mathbf{o}:p(\mathbf{o}|\mathbf{s})>0} \frac{1}{p_\rho(\mathbf{o}|\mathbf{s})} \frac{\partial p_\rho(\mathbf{o}|\mathbf{s})}{\partial \theta_a} \cdot \frac{\partial p_\rho(\mathbf{o}|\mathbf{s})}{\partial \theta_b}. \quad (7.26)$$

For a set \mathcal{S} of k settings the Fisher information matrix associated with a single measurement sample from each setting $\mathbf{s} \in \mathcal{S}$ is given by the sum of the individual Fisher matrices $I(\rho|\mathbf{s})$, and we denote the average $I(\rho|\mathcal{S}) = \frac{1}{k} \sum_{\mathbf{s} \in \mathcal{S}} I(\rho|\mathbf{s})$. The individual matrices can be computed by using definition (7.26) together with the parametrisation (7.25). With the Fisher information thus defined, we can express the mean square error (7.23) more concretely as

$$\text{MSE} := \mathbb{E} \left[\|\rho - \hat{\rho}\|_F^2 \right] \approx \frac{1}{n} \text{Tr}(I(\rho|\mathcal{S})^{-1}G). \quad (7.27)$$

It is worth noting some important aspects of the Fisher information matrix as we define it here. Firstly, since we choose to parameterise a local rank- r neighbourhood of the true state, the asymptotic MSE defined by (7.27) should be thought of as characterising the performance of estimators that have prior knowledge of the rank of the state. Whereas both in experiments as well as in our simulations the true rank of the state is unknown. Secondly we note that the chosen parametrisation is not constrained by the positivity of the states. This means that while the state ρ may be close to the boundary of the positive semi-definite cone of matrices, the parameters $\boldsymbol{\theta}$ need not necessarily be close to the boundary of the parameter space Θ .

7.5 Numerical simulations

In this section we detail the simulation study, and begin by describing the various states and variables of the study. The performance of the estimators presented in the previous sections are compared for two categories of states. The first category of states we consider are equal eigenvalue rank- r states. These states have a fixed eigenvalue spectrum of r equal eigenvalues of magnitude $1/r$ each, and have randomly generated eigenvectors. The performance of the estimators is compared for such states for all ranks $r = 1, \dots, d$. The reason for choosing such states is that we expect them to be representative of the behaviour of rank- r states and to have relatively large errors amongst states of rank r . Additionally, having such a fixed spectrum allows for a more consistent comparison of the estimators across several ranks. The second category of

Error function	Definition
Square Frobenius norm	$D_F(\hat{\rho}, \rho) := \ \hat{\rho} - \rho\ _F^2 = \text{Tr} [(\hat{\rho} - \rho)^*(\hat{\rho} - \rho)]$
Trace norm	$D_{\text{Tr}}(\hat{\rho}, \rho) := \ \hat{\rho} - \rho\ _1 = \text{Tr} \left[\sqrt{(\hat{\rho} - \rho)^*(\hat{\rho} - \rho)} \right]$
Bures distance	$D_B(\hat{\rho}, \rho) := \sqrt{2} \left(1 - \text{Tr} \left[\sqrt{\sqrt{\rho} \hat{\rho} \sqrt{\rho}} \right] \right)^{1/2}$
Hellinger distance	$D_H(\hat{\boldsymbol{\lambda}}, \boldsymbol{\lambda}) := \sqrt{2} \left(1 - \sum_i^d \sqrt{\hat{\lambda}_i \lambda_i} \right)^{1/2}$

Table 7.1 The different error functions used. The Bures distance is defined only for states $\hat{\rho}, \rho \in \mathbb{S}_d$, and its classical analogue the Hellinger distance is defined between two probability distributions.

states we consider are ones of fixed rank and spectrum, such as the GHZ state, and states with decaying eigenvalue spectrums.

We generate the above mentioned states for 3 or 4 qubits, and for a particular ‘true state’ we simulate a dataset \mathcal{D} of counts from which the state is to be reconstructed. The outcome statistics depend on a few variables that we may vary, namely the type of measurement design (random basis vs Pauli), the number of repetitions per settings m , and in the case of the random basis measurements the total number of settings measured k . This allows us to study the performance of the estimators across several different combinations of variables - types of states, ranks, measurement design, number of repetitions per setting m , the total number of settings k and the number of ions N .

7.5.1 Error/Loss functions

Let us denote the mean error (or risk) of an estimator $\hat{\rho}$ as $\mathbb{E}[D(\hat{\rho}, \rho)]$, where $D(\hat{\rho}, \rho)$ represents the choice of error function. In the simulation study we estimate this mean error for several choices of the error function $D(\hat{\rho}, \rho)$, which are tabulated in Table 7.1. We would like to note that the Bures distance is defined only over the space of density matrices, and therefore applies only in the case where the estimates $\hat{\rho}$ are certifiably states. The classical analogue of the Bures distance is the Hellinger distance, which is defined between two probability distributions. We shall have reason to explain the behaviour of the square Bures distance of the estimators $D_B(\hat{\rho}, \rho)^2$ in terms of the square Hellinger distance between their eigenvalues $D_H(\hat{\boldsymbol{\lambda}}, \boldsymbol{\lambda})^2$.

7.5.2 Implementation of the estimators

We now list certain practical details about the implementation of the estimators described in the previous sections. In particular we discuss how we estimate the covariance matrix $\hat{\hat{\Omega}}$ used in the GLS, TGLS, PGLS estimators, and we also describe the cross-validation procedure we use to select a constant C for the thresholded in the TLS and TGLS estimators.

1. The estimate of covariance matrix $\hat{\hat{\Omega}}$ for the generalised estimators (GLS, TGLS and PGLS) is computed from the data in the following way. For a given dataset of counts \mathcal{D} , we first obtain the LS estimate $\hat{\rho}_{\text{LS}}$ and then construct the TLS estimate (see Algorithm 2) with threshold $\nu = 0$. From this we obtain an estimate of the probabilities $\hat{p}(\mathbf{o}|\mathbf{s}) = \text{Tr}[\hat{\rho}_{\text{TLS}} P_{\mathbf{o}}^{\mathbf{s}}]$. The matrix $\hat{\hat{\Omega}}$ is then constructed from these estimated probabilities via (7.11). The generalised estimates (GLS/TGLS/PGLS) are then evaluated using $\hat{\hat{\Omega}}$ and the same dataset \mathcal{D} .
2. The MLE and the positive estimators (PLS and PGLS) are all implemented using the CVX package for disciplined convex programming on MATLAB¹.
3. As mentioned briefly in section 7.3.4, the constant C in the threshold for the TLS and TGLS estimators is selected using cross-validation. We describe this cross-validation method below [25].
 - For a particular number of repetitions per setting m , we simulate data in 5 independent batches and we denote the corresponding datasets as $\mathcal{D}_1, \dots, \mathcal{D}_5$. With $m/5$ repetitions per setting in each batch, so that the number of repetitions overall is m . The total dataset of counts is therefore simply the sum $\mathcal{D} = \sum_{i=1}^5 \mathcal{D}_i$.
 - We choose a vector of constants C . For each value of C , and for each $j \in \{1, \dots, 5\}$ we compute the following estimators. We hold out the dataset \mathcal{D}_j , and compute the TLS/TGLS estimate $\hat{\rho}_{\text{T(G)LS}}^{-j}(C)$ for the dataset $\mathcal{D}_{-j} = \sum_{i \neq j} \mathcal{D}_i$, with constant $\nu = C\mu = C\sqrt{\frac{4}{m} \log 2^{N+1}}$. For each \mathcal{D}_j the LS/GLS estimate $\hat{\rho}_{\text{(G)LS}}^j$ is also evaluated.

¹<http://cvxr.com/cvx/>

- For all values of C , the empirical discrepancy is evaluated

$$CV_D(C) = \frac{1}{5} \sum_{i=1}^5 D(\hat{\rho}_{\text{T(G)LS}}^{-j}(C), \hat{\rho}_{\text{(G)LS}}^j)$$

for a choice of error function $D(\hat{\rho}, \rho)$.

- This function $CV_D(C)$ is then minimised over all values C

$$\hat{C}_D = \arg \min_C CV_D(C)$$

this gives an estimate for the thresholding constant, which is then used to evaluate the TLS or the TGLS estimators with threshold $\nu = \hat{C}_D \mu$.

Notice that the cross-validation procedure picks different constants for different choices of the error function. An important caveat here is that the Bures distance is not defined for the LS/GLS estimates $\hat{\rho}_{\text{(G)LS}}^j$, and therefore the above procedure cannot apply. Instead in the simulations we estimate the thresholding constant \hat{C}_{D_B} using the ML estimate as

$$\hat{C}_{D_B} = \arg \min_C D_B(\hat{\rho}_{\text{T(G)LS}}, \hat{\rho}_{\text{ML}})^2.$$

7.6 Simulation results and analysis

We first present plots of the estimated mean errors $\mathbb{E}[D(\hat{\rho}, \rho)]$ of the estimators for the equal eigenvalues states. The results for the special states follows in section 7.6.4. For each given rank r and number of qubits N , we generate a state with equal eigenvalues $(\frac{1}{r}, \dots, \frac{1}{r}, 0, \dots, 0)$ and random eigenbasis. Then for each choice of measurement design and values of k and m , the several estimates of the true state are evaluated. The error of each resulting estimate is computed using all the error functions listed in Table 7.1, and the corresponding mean errors are estimated from 100 different runs of the experiment.

In order to make the results of the simulation study more accessible, we have made the plots presented here available online via an interactive R Shiny application at this address: <https://rudhacharya.shinyapps.io/plots/>

7.6.1 Square Frobenius norm

We notice in figures 7.1, 7.2 and 7.3 that the performance of the least squares estimator is clearly the poorest amongst all the estimators for ranks $r < d$, additionally its mean square error (MSE) demonstrates no significant dependence on the rank of the true state and remains fairly constant, while the remaining estimators all show a scaling of the MSE with the rank of the true state. On comparison with the Fisher MSE, we see that the performance of several of the estimators matches well with the asymptotic value. This is remarkable as the Fisher MSE (7.27) was defined for a rank- r parameterisation of states, while none of the estimators have any prior knowledge of the rank.

We also note that for relatively small values of $n = m \times k$ the TLS, TGLS, PLS, PGLS, ML estimators all appear to perform better than the Fisher error (for example the values $k = m = 100$ in figure 7.3 and $k = 81, m = 100$ in figure 7.1). Although this might seem surprising at first, it can be understood from the following observation. The Fisher information matrix (7.26) is defined in terms of a local parametrisation that is not constrained by the positivity of states, while all of the listed estimators produce estimates that are density matrices. The eigenvalues $\lambda = 1/r$ of states of high ranks are all small and close to zero. For such eigenvalues and relatively small values of n , the estimates $\hat{\lambda}$ of unbiased estimators can be thought of as being roughly normally distributed with mean $1/r$, and *variance that is comparable or larger than the magnitude of the eigenvalues themselves*. Therefore without the constraint of positivity, unbiased estimators would produce estimates with $\hat{\lambda} < 0$. Whereas estimators like the MLE are not unbiased for small n as the estimates are constrained to be positive. This requirement of positivity provides additional information when the eigenvalues are small, and explains the observed difference in the performance of such estimators when compared with the Fisher error. For large values of n however, the uncertainty in the eigenvalues is very small and the Fisher MSE acts as a lower bound for all of the estimators. For a more detailed discussion on the applicability of the asymptotic bound and the unbiasedness of the MLE for finite n in the qubit case, we recommend [106].

Across both the Pauli and the random measurement design we note that the performance of the PGLS and the ML estimators is very similar, and for large m almost identical. We can demonstrate that this similarity in behaviour is to be expected and in fact show that the PGLS and the ML estimators are equivalent in the limit of large m . In this limit and for probabilities not close

to the boundary, the multinomial distribution (7.2) is well approximated by the multivariate Gaussian, i.e $\text{Multi}(m, \mathbf{y}) \rightarrow \mathcal{N}(m\mathbf{y}, m\Omega)$, where $\mathbf{y} \in \mathbb{R}^{dk}$ is the vector of probabilities (7.5), and Ω is the covariance matrix of the multinomial distribution (7.11). Therefore the distribution of the frequency vector \mathbf{f} in this limit is given by the Gaussian density function

$$\frac{1}{\sqrt{(2\pi)^{dk}} \sqrt{|\Omega|}} \exp \left[-\frac{m}{2} (\mathbf{f} - \mathbf{y})^T \Omega^{-1} (\mathbf{f} - \mathbf{y}) \right]. \quad (7.28)$$

This is also by definition the form of the likelihood function that needs to be maximised in order to obtain the MLE. Taking the natural logarithm of (7.28) and discarding the constant terms we may therefore express the MLE in the large m limit as the following minimisation problem

$$\hat{\rho}_{\text{ML}} = \arg \min_{\tau \in \mathbb{S}_d} \|\Omega^{-1/2} (\mathbf{f} - \mathcal{X}(\tau))\|^2. \quad (7.29)$$

This we recognise from (7.20) to be the definition of the PGLS estimator. Therefore in the large m limit, the two estimators are equivalent. There is also a similar equivalence between the non-positive GLS estimator and the ‘unconstrained’ ML estimate. Let us define the unconstrained MLE as the following relaxation of the above minimisation

$$\hat{\rho}_{\text{uML}} := \arg \min_{\tau \in M(\mathbb{C}^d)} \|\Omega^{-1/2} (\mathbf{f} - \mathcal{X}(\tau))\|^2. \quad (7.30)$$

On comparison with (7.13) we recognise this to simply be the GLS estimator. Since the GLS is known to be asymptotically normal $\sqrt{m}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \rightarrow \mathcal{N}(\mathbf{0}, I(\boldsymbol{\theta})^{-1})$, with $\boldsymbol{\theta} \in \Theta$ being some choice of parameterisation of $M(\mathbb{C}^d)$, it follows that the unconstrained ML has the same asymptotic distribution. It is important to note that this distribution is not only over states, but over self-adjoint matrices. This means that for states ρ that are close to boundary of the positive semi-definite cone, the GLS and the unconstrained ML may produce estimates that are not valid density matrices. Additionally this Gaussian distribution holds as long as the probabilities $p_\rho(\mathbf{o}|\mathbf{s}) > 0$ for all \mathbf{o} and \mathbf{s} . This was also noted with the Fisher information matrix in section 7.4, where the state being at the boundary of \mathbb{S}_d does not necessarily imply that parameter $\boldsymbol{\theta}$ is at the boundary of the parameter space Θ .

However this is not true of the PGLS and ML estimators (7.29). For most states the MLE (and therefore the PGLS) is known to be asymptotically nor-

mally distributed with covariance given by the inverse of the Fisher information matrix. But since the MLE produces only density matrices as estimates, its distribution for states ρ close to the boundary is not Gaussian. In fact it can be shown that for such states the ML/PGLS estimate is obtained by projecting the unconstrained ML estimate onto the space of density matrices. To make the form of this projection concrete, we first note that since Ω is a positive matrix the function $\|\Omega^{-1/2}(\mathbf{a} - \mathbf{b})\|^2$ is a valid distance $\forall \mathbf{a}, \mathbf{b} \in \mathbb{R}^{dk}$, and (7.30) describes an orthogonal projection of \mathbf{f} onto the linear subspace defined by the map \mathcal{X} . Using this fact we may expand the right side of (7.29) by the Pythagorean theorem for any $\tau \in \mathbb{S}_d$ as

$$\|\Omega^{-1/2}(\mathbf{f} - \mathcal{X}(\tau))\|^2 = \|\Omega^{-1/2}(\mathbf{f} - \mathcal{X}(\hat{\rho}_{\text{uML}}))\|^2 + \|\Omega^{-1/2}(\mathcal{X}(\hat{\rho}_{\text{uML}}) - \mathcal{X}(\tau))\|^2. \quad (7.31)$$

The term $\|\Omega^{-1/2}(\mathbf{f} - \mathcal{X}(\hat{\rho}_{\text{uML}}))\|^2$ is fixed for a particular estimate $\hat{\rho}_{\text{uML}}$, and therefore the minimisation (7.29) is equivalently expressed as

$$\hat{\rho}_{\text{ML}} = \arg \min_{\tau \in \mathbb{S}_d} \|\Omega^{-1/2}(\mathcal{X}(\hat{\rho}_{\text{uML}}) - \mathcal{X}(\tau))\|^2. \quad (7.32)$$

This is seen to be a projection of $\hat{\rho}_{\text{uML}}$ onto the closest density matrix. Therefore the ML estimate can be obtained by projecting the unconstrained ML estimate $\hat{\rho}_{\text{uML}}$ onto the space of density matrices.

7.6.2 Square Bures distance

As the Bures distance $D_B(\hat{\rho}, \rho)$ is defined only over density matrices, we plot the *Mean Square Bures* (MSB) error only for the ML, TLS, TGLS, PLS and the PGLS estimators. The most noticeable feature across both the Pauli and the random bases measurement design is that for large values of n (for example the plots with $m = 1000$ in figures 7.4, 7.5, 7.6) the MSB is seen to be larger for states of middling ranks than for the full rank states. Contrast this with the behaviour for smaller values of n (plots with $m = 100$ in figure 7.4).

This behaviour is explained by the fact that the Bures distance is known to be sensitive to the misestimation of very small eigenvalues [89]. In chapter 6 we demonstrated that for qubit states in a local neighbourhood close to the boundary of the Bloch sphere, the leading order contribution to the square Bures error is from the classical Hellinger component. In fact, comparing the plots in this section with the corresponding plots for the square Hellinger distance (section 7.6.3), we see that for large n the *Mean Square Hellinger*

(MSH) error of the estimators $\mathbb{E} [D_H(\hat{\lambda}, \lambda)^2]$ is approximately equal to the MSB. Furthermore the MSH also demonstrates the same behaviour with the rank. For large n the eigenvalues are all seen to be estimated well and the resulting square Hellinger errors are substantially smaller than $1/d$. This implies that for full rank states the error $(\lambda_i^{1/2} - \hat{\lambda}_i^{1/2})^2 = (\frac{1}{\sqrt{d}} - \hat{\lambda}_i^{1/2})^2$ for each $i \in \{1, \dots, d\}$ is very small, while for states of lower ranks, the estimators may misestimate the rank and assign small positive values for some of the zero eigenvalues, i.e. $\hat{\lambda}_i > 0$ for $i > r$. Since the error for these components is simply $(\lambda_i^{1/2} - \hat{\lambda}_i^{1/2})^2 = \hat{\lambda}_i$, the estimated value of $\hat{\lambda}_i$ can be fairly small and still result in a total error that is larger than in the case of the full rank states.

As discussed in chapters 6 and 3, this sensitivity of the Bures distance to the misestimation of small eigenvalues is a feature of the quantum fidelity, and this behaviour affects the rate at which states with small eigenvalues can be estimated [105, 89]. For states near the boundary of the Bloch sphere and fidelity based distances, the estimation errors scale only as $O(1/\sqrt{n})$ for fixed bases measurements, instead of the standard rate of $O(1/n)$ as in (7.27).

7.6.3 Square Hellinger distance and Trace norm

Since like the Bures distance, the Hellinger distance is valid only between probability vectors, we plot the MSH only for estimators that produce certifiable states. As explained above, it is insightful to compare the plots (figures 7.7, 7.8, 7.9) to those for the square Bures distance. For large values of n and states with eigenvalues close to zero, the magnitude and behaviour of the MSH is similar to that of the MSB.

The relative performance of the estimators in terms of the trace norm (figures 7.10, 7.11, 7.12) is largely similar to their performance in terms of the square Frobenius norm. We specially note the behaviour of the mean errors as the rank of the true state is varied, and the comparable performance of the PGLS and the ML estimators.

7.6.4 Special states

Apart from the states with equal eigenvalues, we also compare the performance of the estimators in the estimation of certain states of fixed ranks and eigenvalue spectrums. We consider the GHZ state $(|0\rangle^{\otimes N} + |1\rangle^{\otimes N})/\sqrt{2}$, which is commonly used in several quantum information processing tasks [119, 20, 68], along with states that have exponentially decaying eigenvalue spectrums, and full rank

states with a few dominant eigenvalues and the rest sampled as small uniform noise. Simulations with such states allows us to better understand how the spectrum of the eigenvalues affect the performance of the estimators. For each given state of $N = 4$ qubits, the several estimates are computed for the two different measurement designs, with $m = 100$ repetitions per setting (which is typical in tomographic experiments [62]) and 100 different runs of the experiment.

In figures 7.13, 7.14 we plot the estimation errors in terms of the various loss functions for a randomly generated full-rank state with an exponentially decaying eigenvalue spectrum. We also plot the mean eigenvalues of the estimates and compare them against the eigenvalues of the true state. For both measurement designs the errors are comparable to estimation errors for the equal eigenvalue state of rank 3/4 for $m, k = 100$. Apart from the LS and GLS estimators, the others are seen to estimate the true eigenvalues well. However, both the LS and GLS prove to be useful starting points for the thresholded estimators. Figures 7.15, 7.16, 7.17, 7.18 plot the estimation errors for a random full-rank state with 2 and 4 dominant eigenvalues. The remaining eigenvalues are small and are generated as uniform noise. Unlike their performance for the state with decaying eigenvalues, all of the estimators are seen to mis-estimate several of the smaller eigenvalues.

Plots for the GHZ state are presented in figures 7.19 and 7.20. It is worth noting that the magnitude of the square Bures errors is comparable to that of the square Hellinger errors for all the estimators and both measurement designs. This is in keeping with the discussion in section 7.6.2, as the estimators have no knowledge that the true state is pure and therefore any misestimation of zero eigenvalues carries a large contribution to the Hellinger component of the Bures error. In the plots for the Pauli measurement design (figure 7.19), the PLS estimator is seen to perform the worst in terms of the square Hellinger and Bures error. In the eigenvalue plots we see that the PLS estimate is not close to pure and has a few large eigenvalues. However, on comparison with the estimation errors for a random pure state and Pauli measurement design in figure 7.21, one notices that rather than the PLS performing poorly, the TGLS, PGLS and ML estimators all perform better for the GHZ state and Pauli settings. This behaviour is however not seen in the square Frobenius and trace norm errors, nor in the plots for the random bases measurement design 7.20.

Mean square Frobenius error (MSE) for the Pauli measurement design

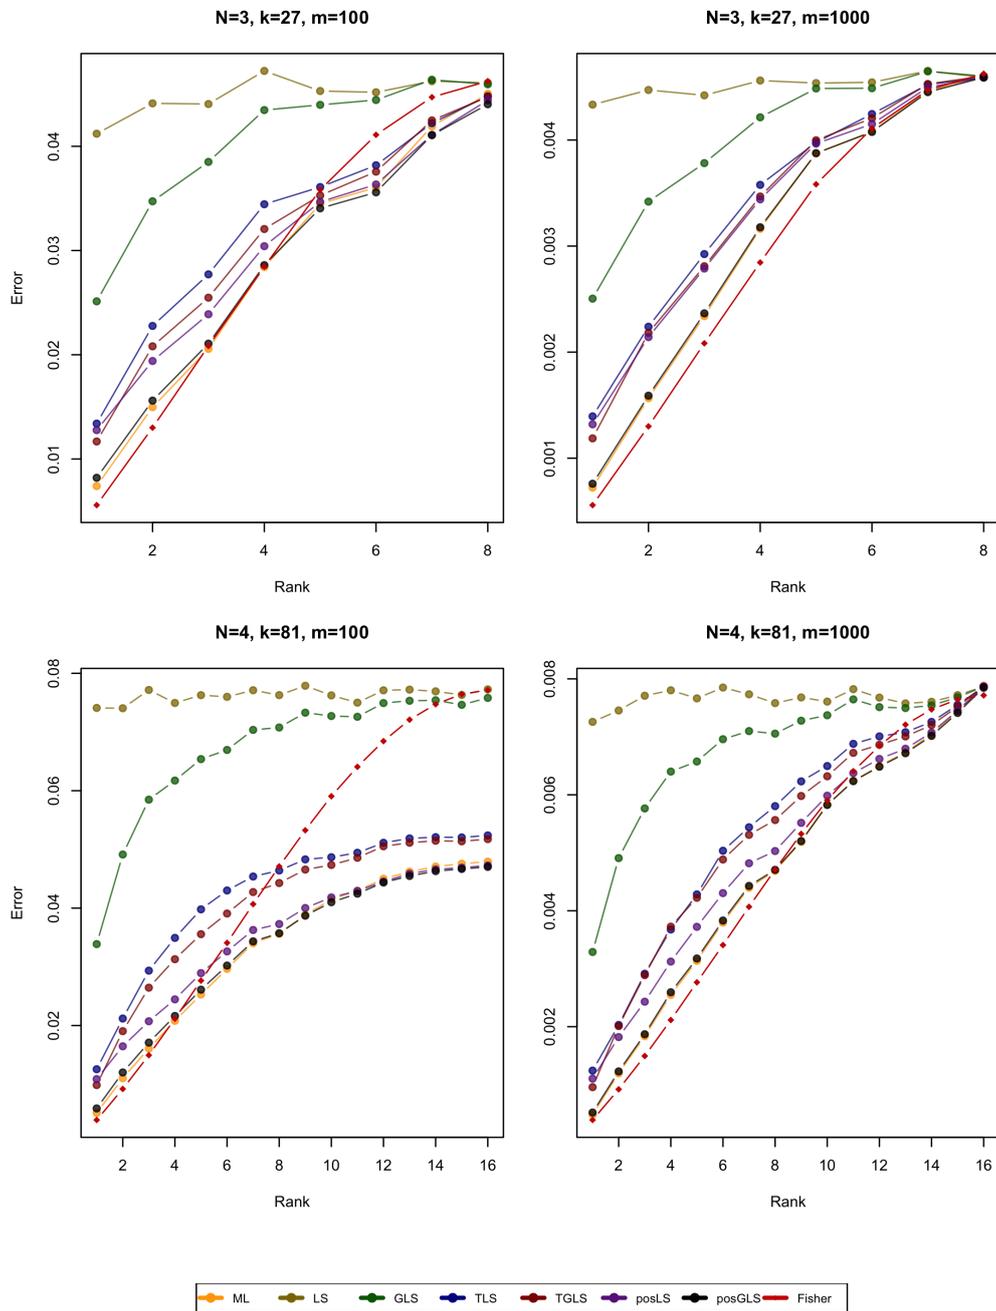


Figure 7.1 The estimated mean square Frobenius error (MSE) $\mathbb{E}[D_F(\hat{\rho}, \rho)]$ of the estimators for randomly generated 3 and 4 qubit rank- r states of equal eigenvalues, and the Pauli measurement settings.

Mean square Frobenius error (MSE) for the random bases measurement design and 3 qubits

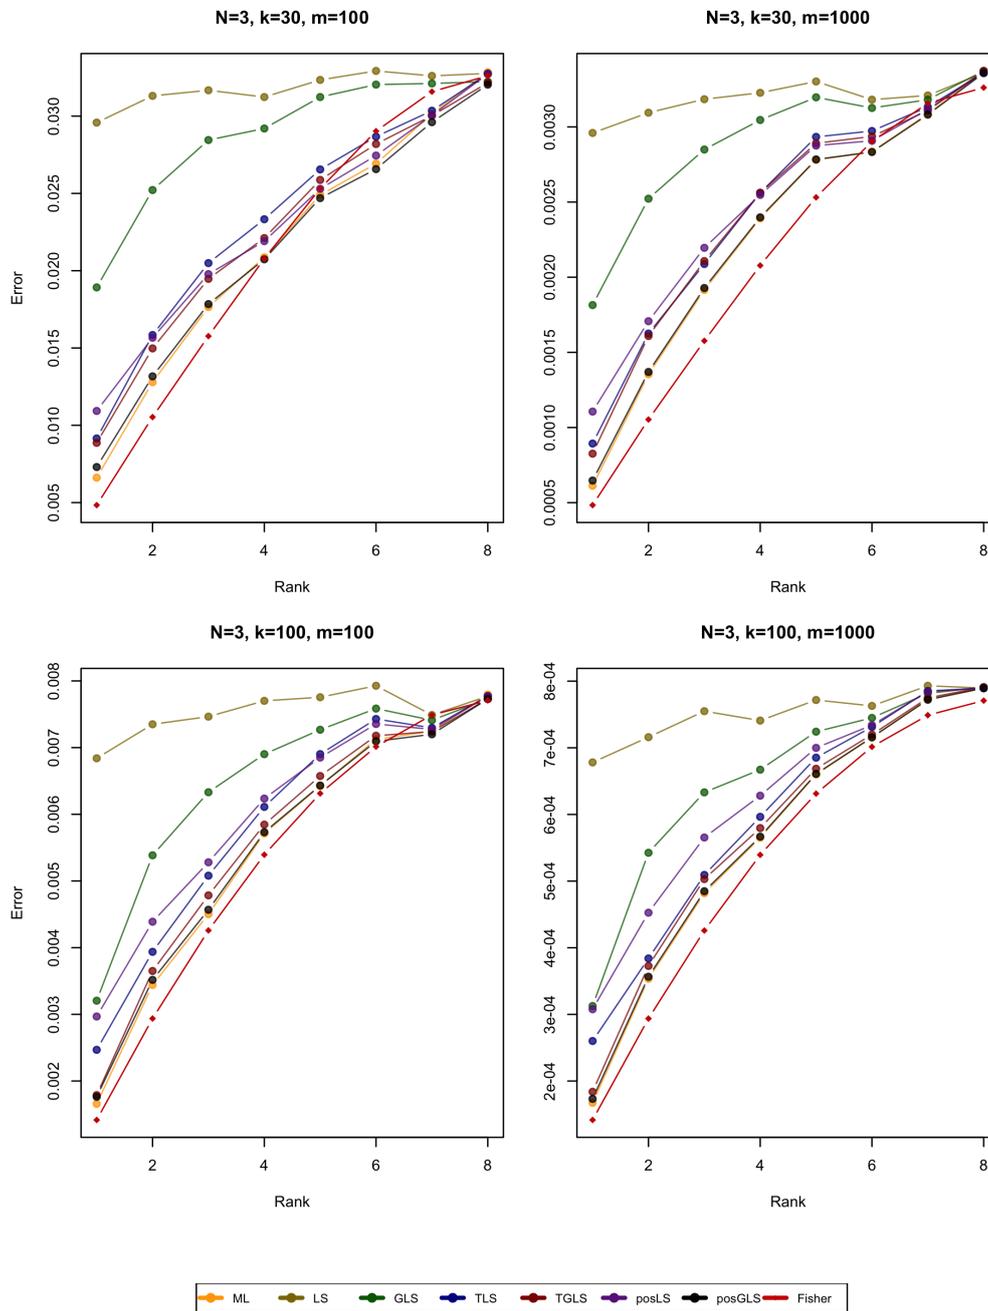


Figure 7.2 The estimated mean square Frobenius error (MSE) $\mathbb{E}[D_F(\hat{\rho}, \rho)]$ of the estimators for randomly generated 3 qubit rank- r states of equal eigenvalues, and the random bases measurement design.

Mean square Frobenius error (MSE) for the random bases measurement design and 4 qubits

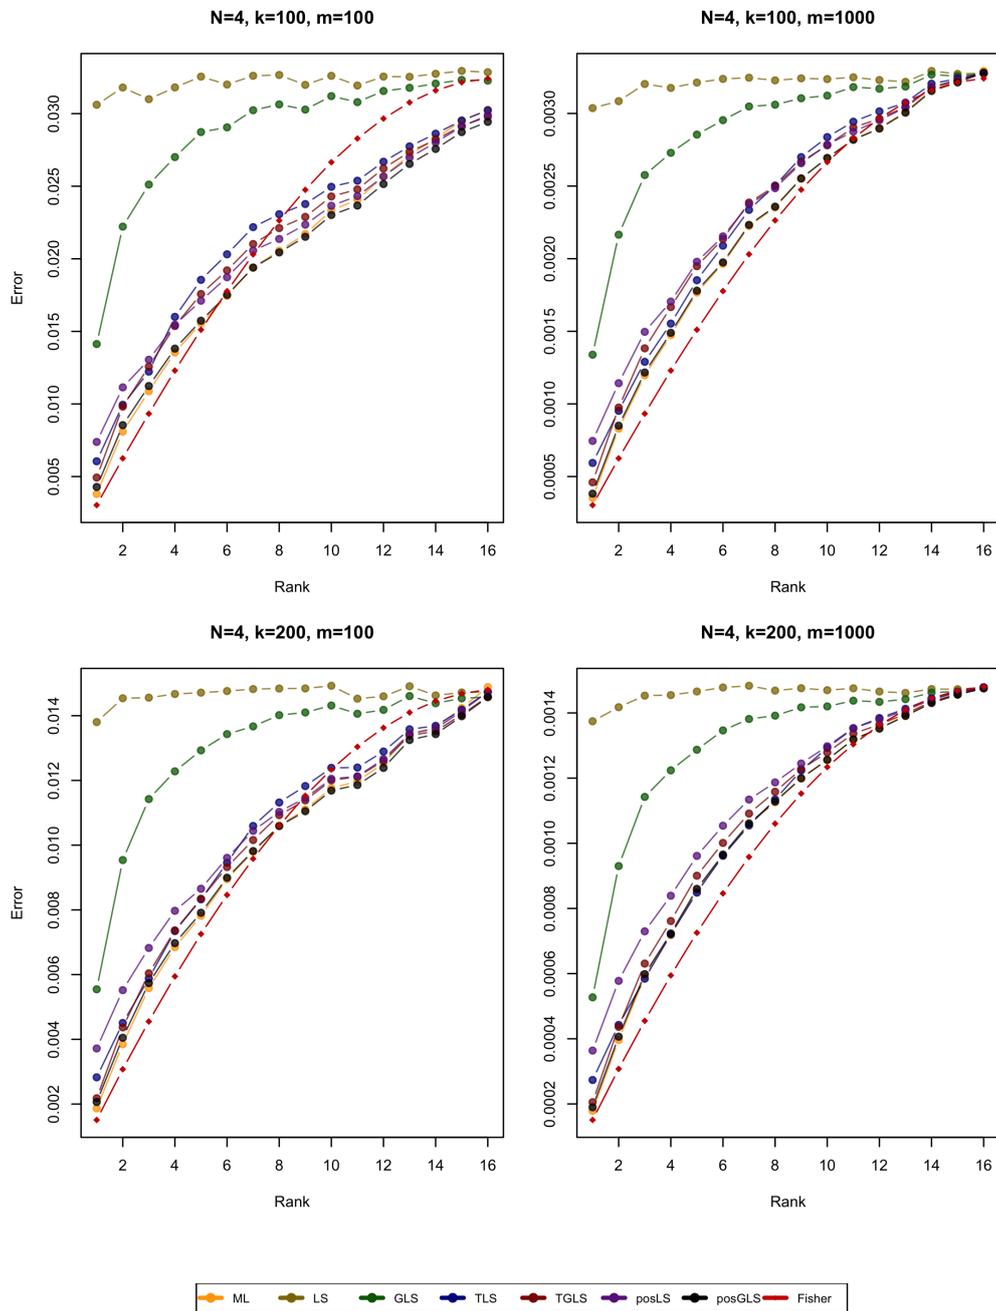


Figure 7.3 The estimated mean square Frobenius error (MSE) $\mathbb{E}[D_F(\hat{\rho}, \rho)]$ of the estimators for randomly generated 4 qubit rank- r states of equal eigenvalues, and the random bases measurement design.

Mean square Bures error (MSB) for the Pauli measurement design

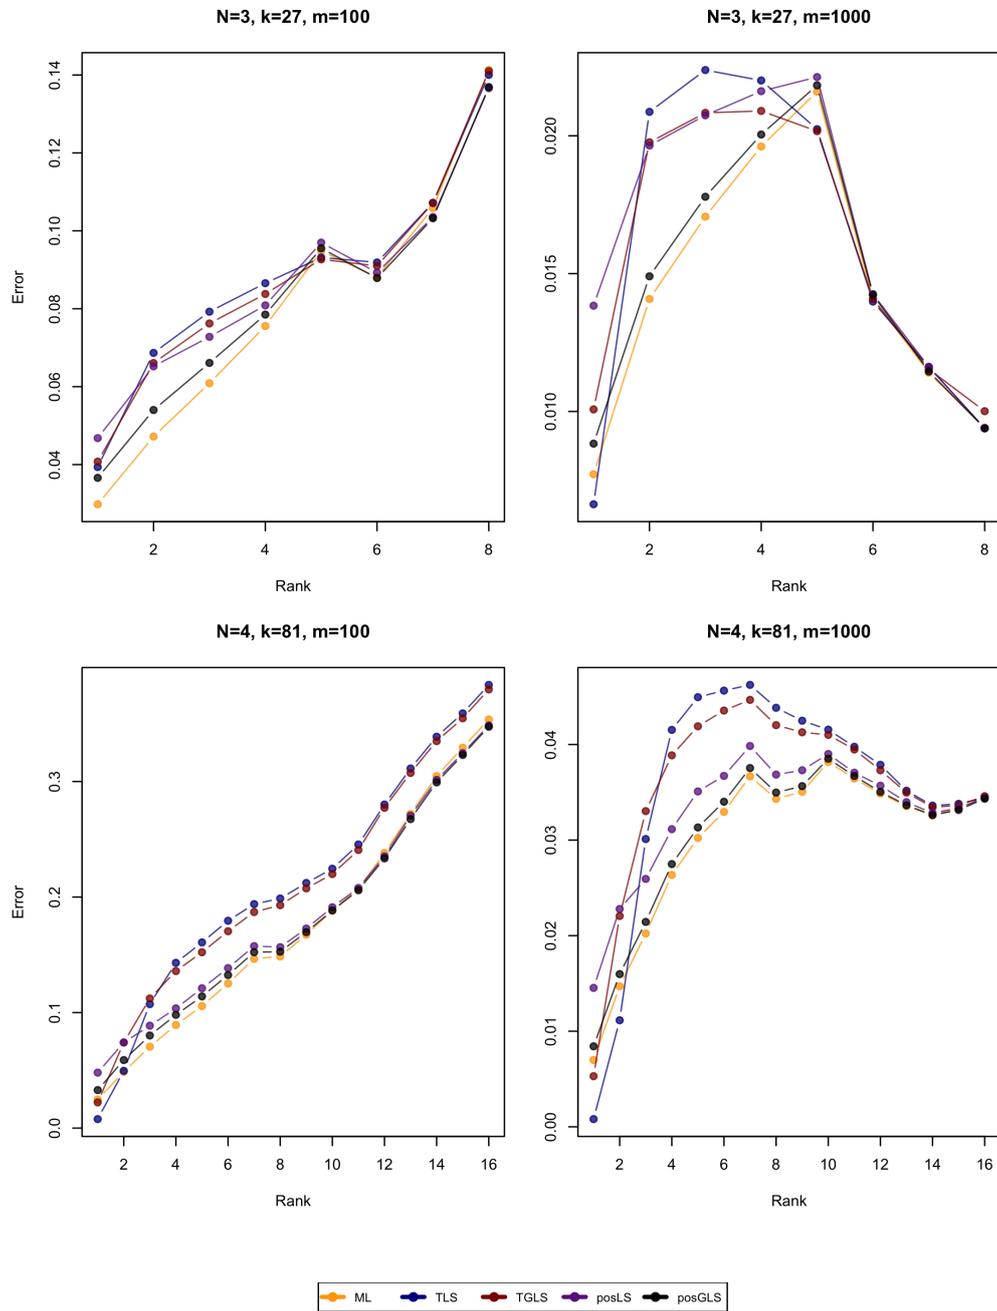


Figure 7.4 The estimated mean square Bures error (MSB) $\mathbb{E} \left[D_B(\hat{\rho}, \rho)^2 \right]$ of the estimators for randomly generated 3 and 4 qubit rank- r states of equal eigenvalues, and the Pauli measurement settings.

Mean square Bures error (MSB) for the random bases measurement design and 3 qubits

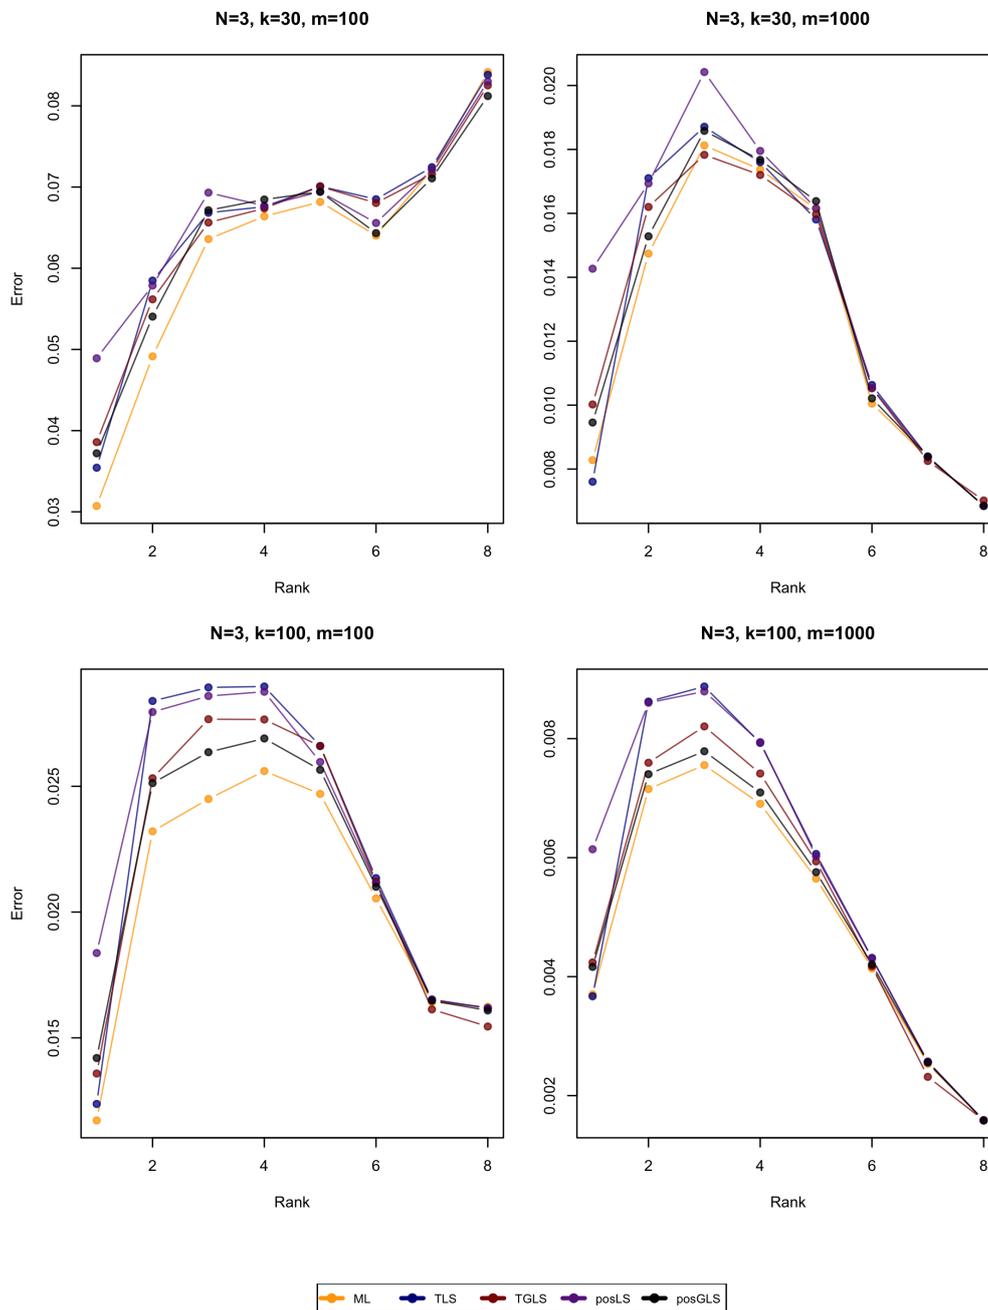


Figure 7.5 The estimated mean square Bures error (MSB) $\mathbb{E} [D_B(\hat{\rho}, \rho)^2]$ of the estimators for randomly generated 3 qubit rank- r states of equal eigenvalues, and the random bases measurement design.

Mean square Bures error (MSB) for the random bases measurement design and 4 qubits

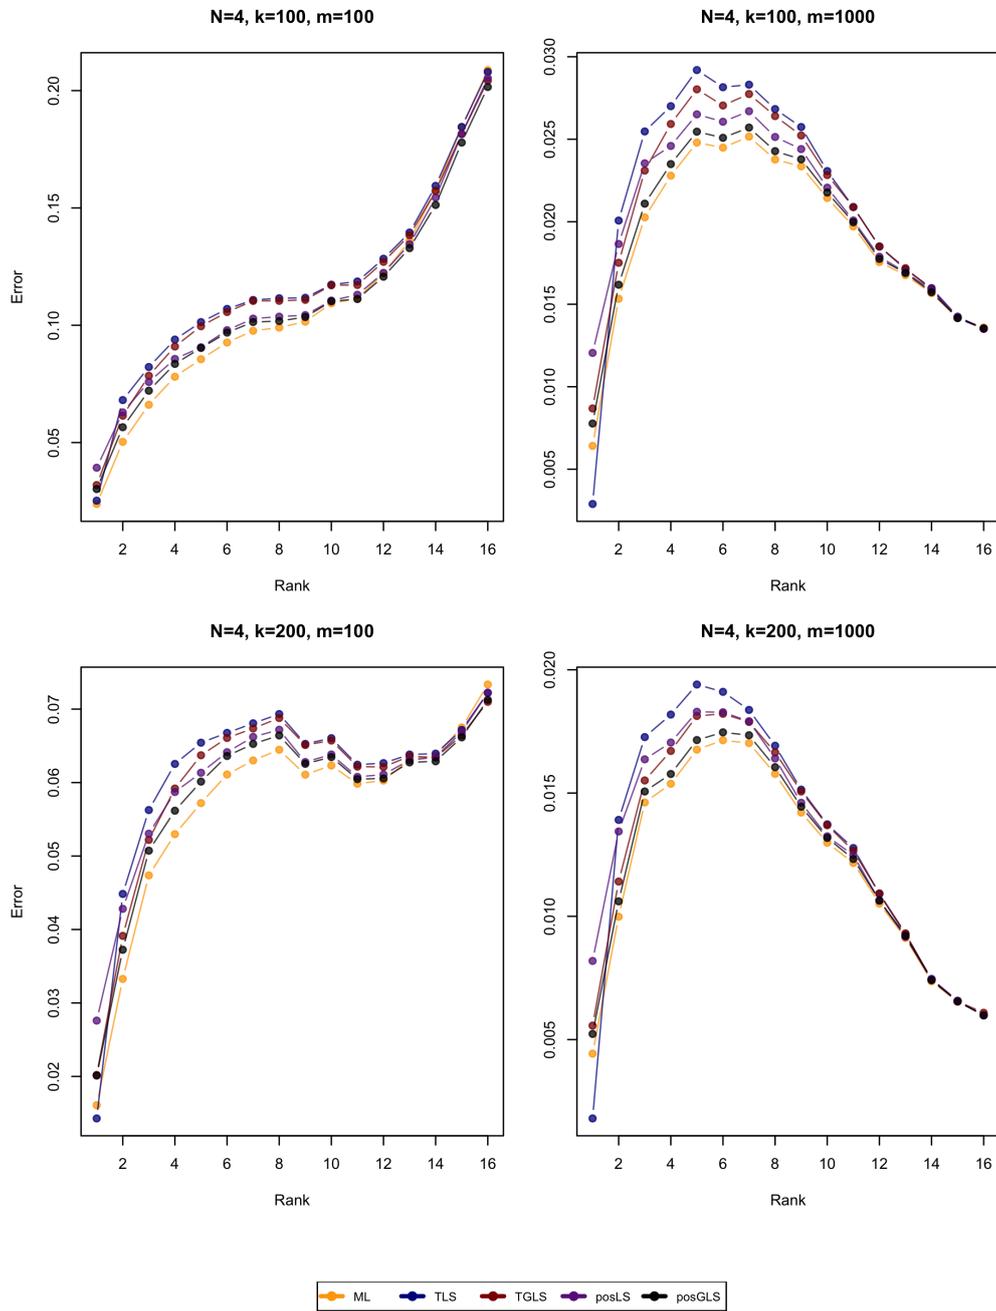


Figure 7.6 The estimated mean square Bures error (MSB) $\mathbb{E} [D_B(\hat{\rho}, \rho)^2]$ of the estimators for randomly generated 4 qubit rank- r states of equal eigenvalues, and the random bases measurement design.

Mean square Hellinger error (MSH) for the Pauli measurement design

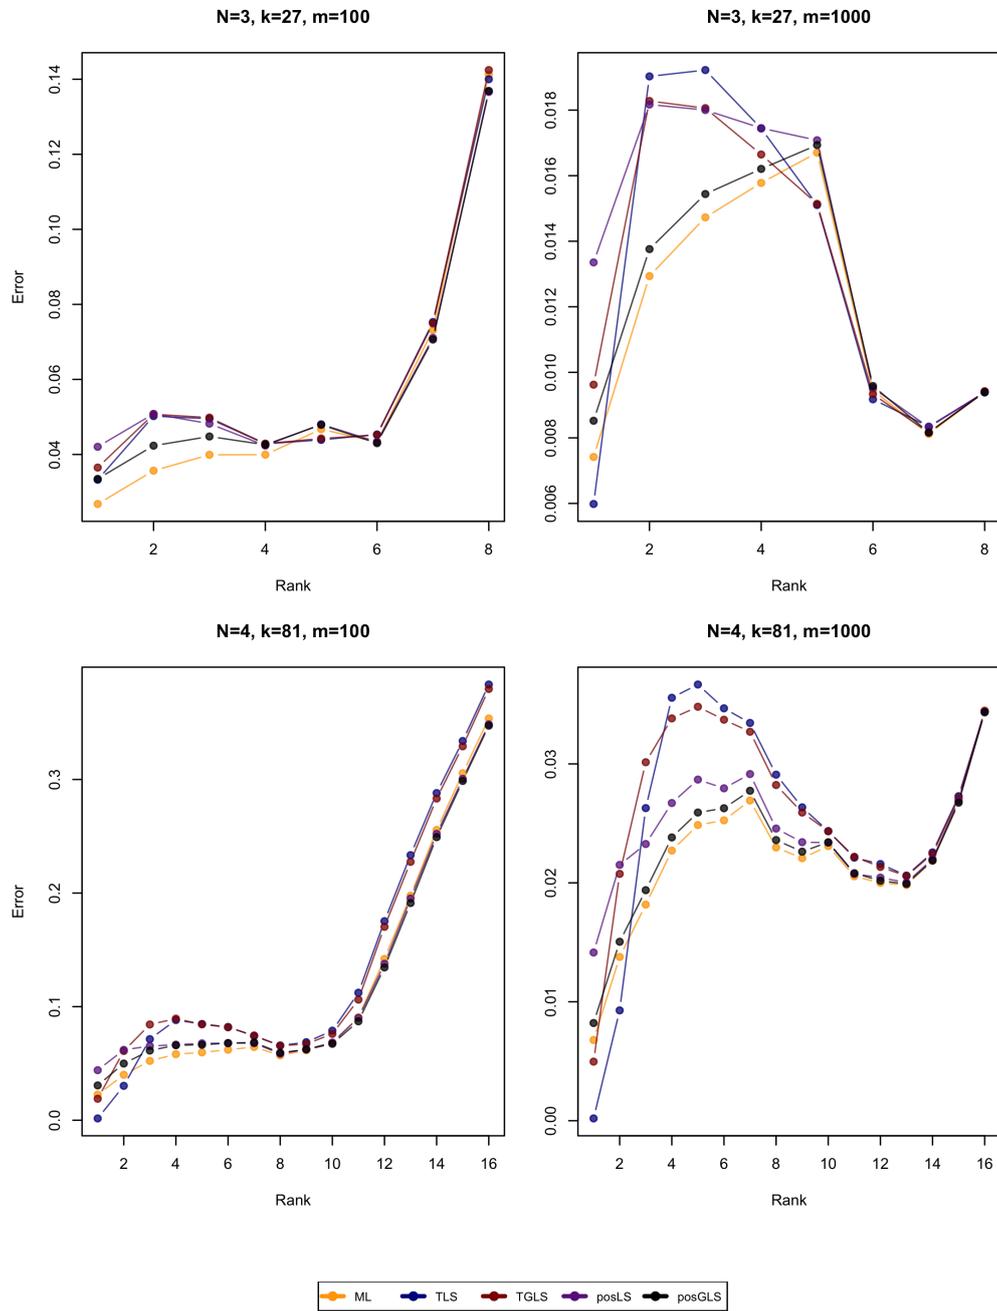


Figure 7.7 The estimated mean square Hellinger error (MSH) $\mathbb{E} [D_H(\hat{\lambda}, \lambda)^2]$ between the eigenvalues of the estimates and those of a randomly generated 3 and 4 qubit rank- r state with equal eigenvalues, for the Pauli measurement design.

Mean square Hellinger error (MSH) for the random bases measurement design and 3 qubits

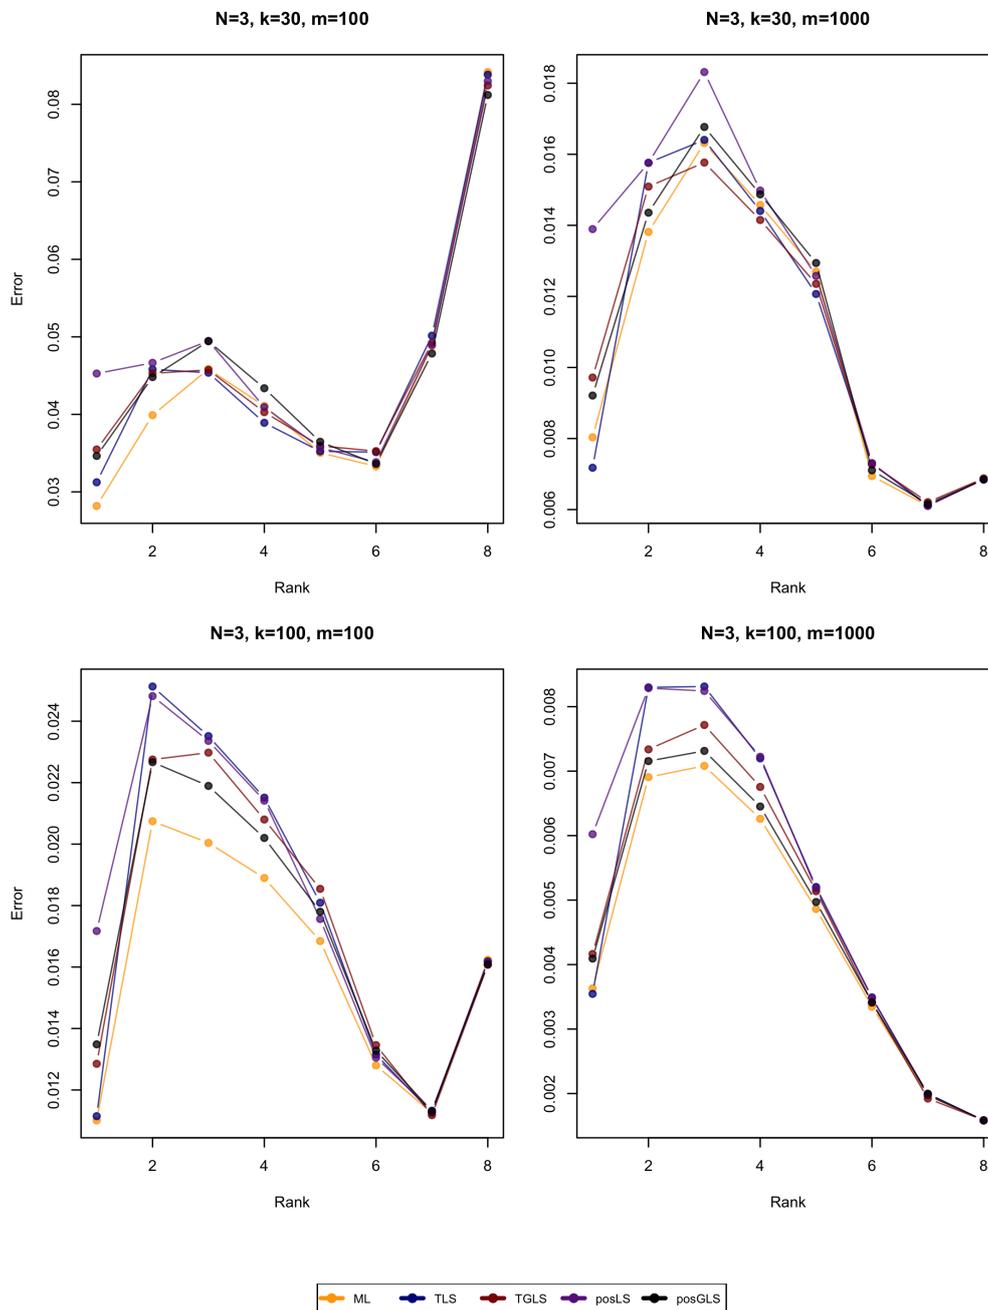


Figure 7.8 The estimated mean square Hellinger error (MSH) $\mathbb{E} \left[D_H(\hat{\lambda}, \lambda)^2 \right]$ between the eigenvalues of the estimates and those of a randomly generated 3 qubit rank- r state with equal eigenvalues, for the random bases measurement design.

Mean square Hellinger error (MSH) for the random bases measurement design and 4 qubits

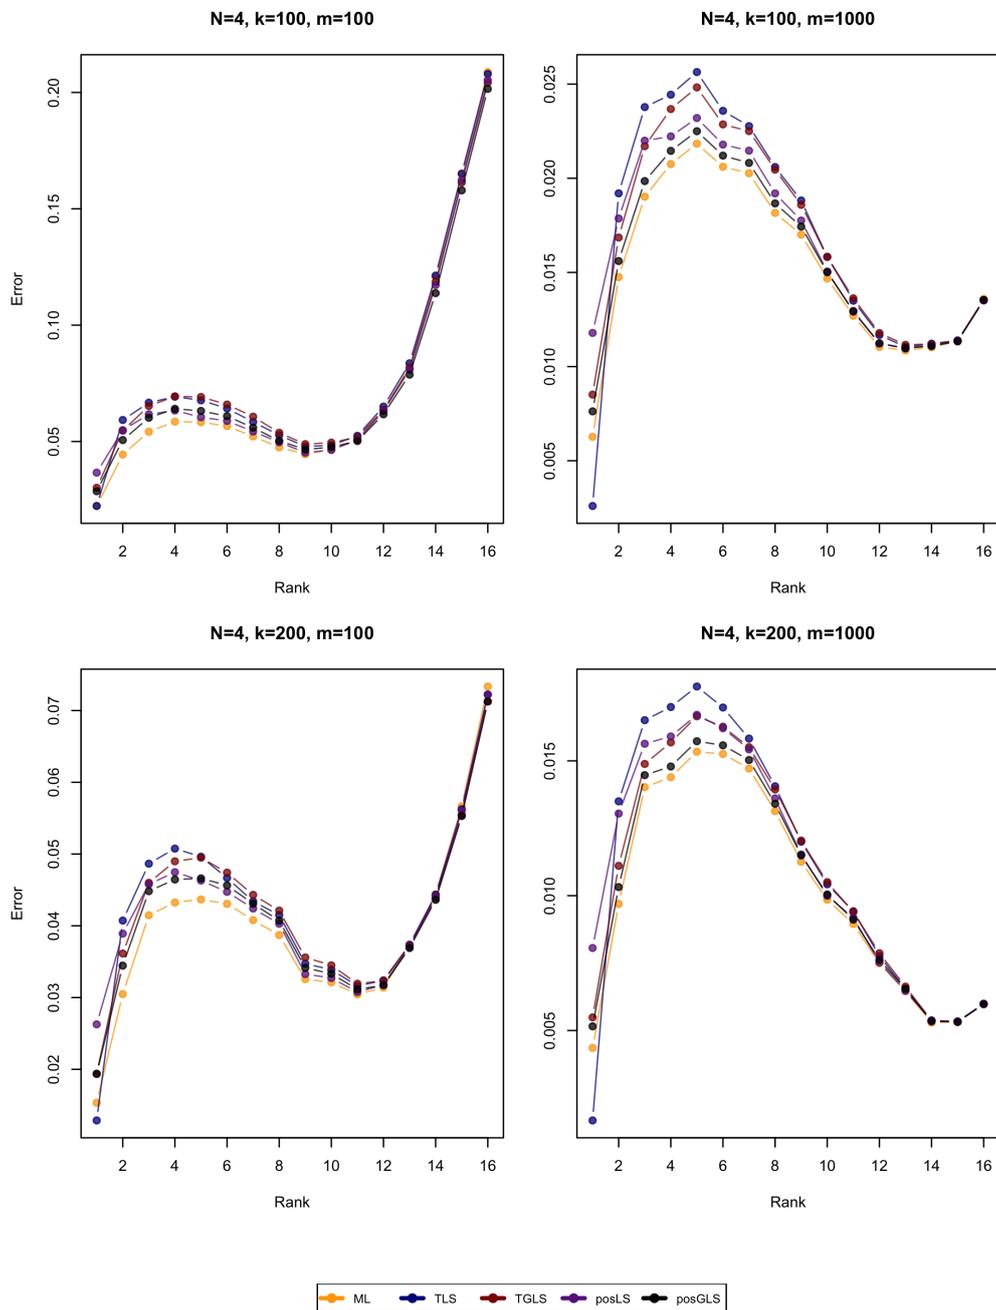


Figure 7.9 The estimated mean square Hellinger error (MSH) $\mathbb{E} [D_H(\hat{\lambda}, \lambda)^2]$ between the eigenvalues of the estimates and those of a randomly generated 4 qubit rank- r state with equal eigenvalues, for the random bases measurement design.

Mean trace norm error for the Pauli measurement design

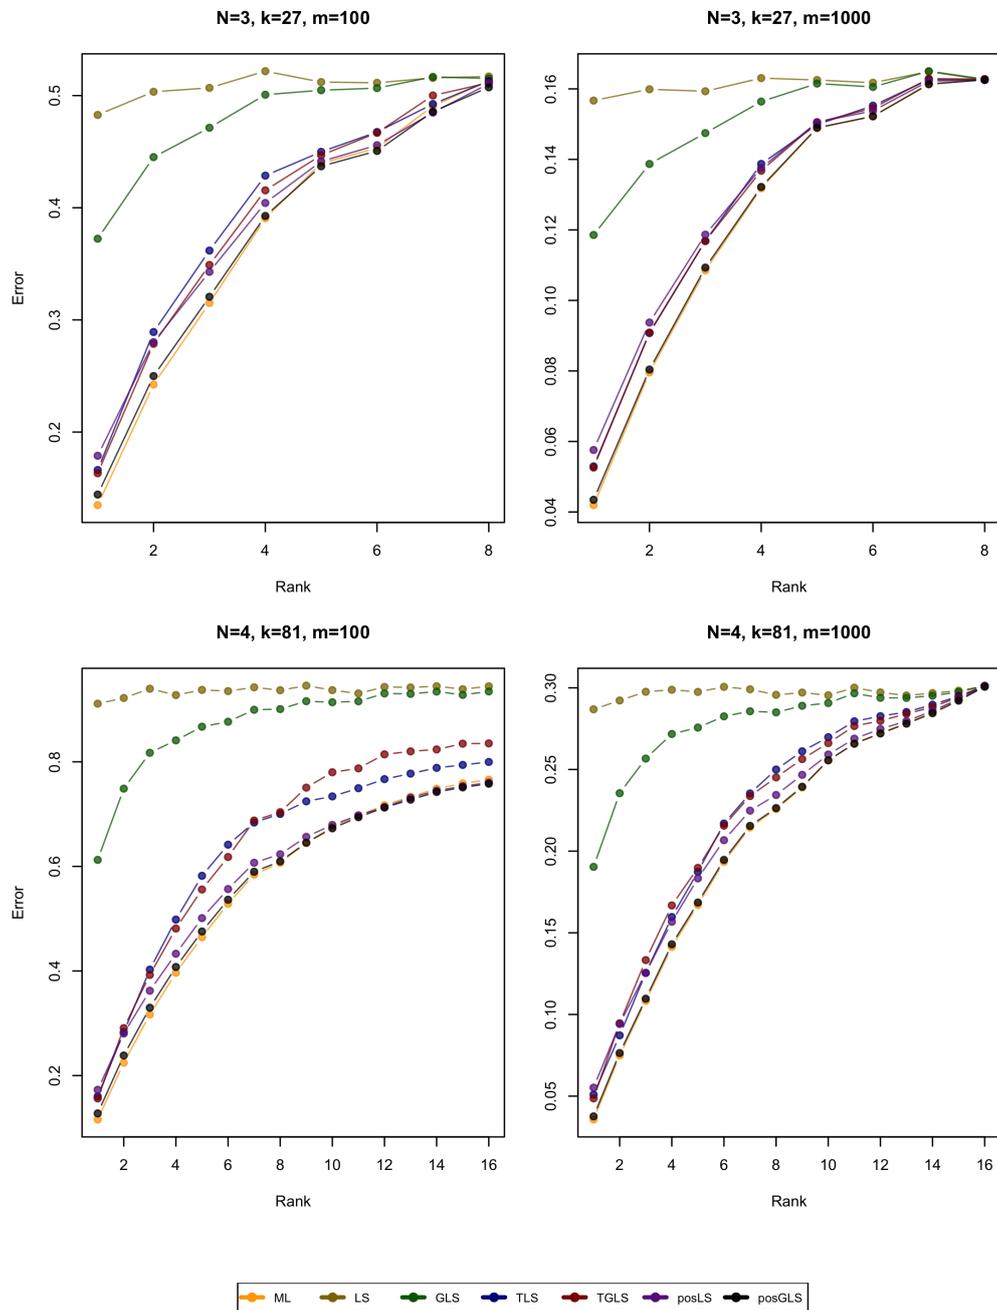


Figure 7.10 The estimated mean trace norm error $\mathbb{E}[\|\hat{\rho} - \rho\|_1]$ of the estimators for randomly generated 3 and 4 qubit rank- r states of equal eigenvalues, and the Pauli measurement design.

Mean trace norm error for the random bases measurement design and 3 qubits

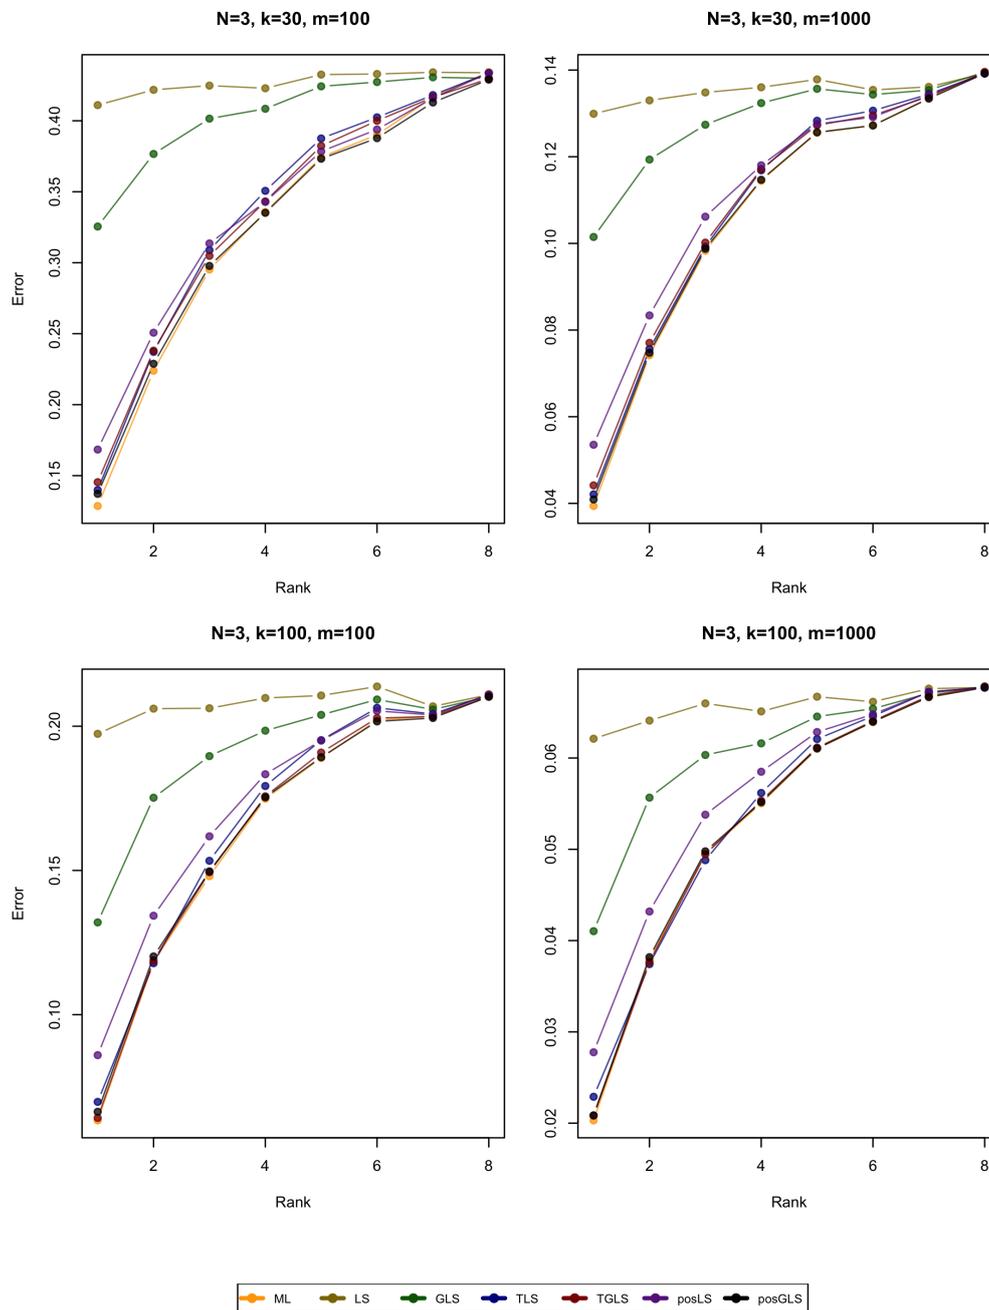


Figure 7.11 The estimated mean trace norm error $\mathbb{E}[\|\hat{\rho} - \rho\|_1]$ of the estimators for randomly generated 3 qubit rank- r states of equal eigenvalues, and the random bases measurement design.

Mean trace norm error for the random bases measurement design and 4 qubits

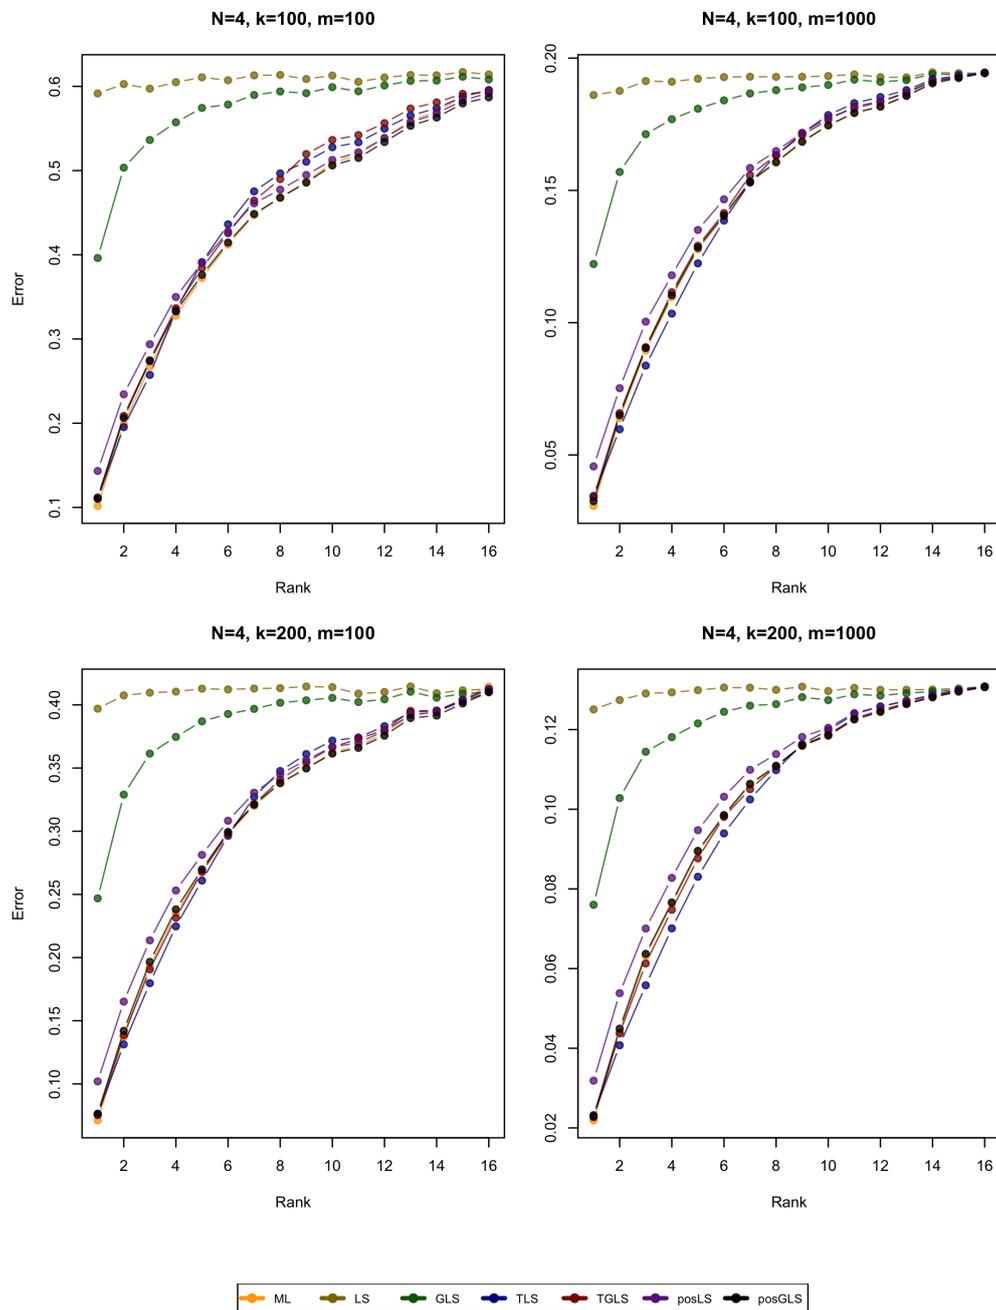


Figure 7.12 The estimated mean trace norm error $\mathbb{E}[\|\hat{\rho} - \rho\|_1]$ of the estimators for randomly generated 4 qubit rank- r states of equal eigenvalues, and the random bases measurement design.

Estimation errors for the Pauli measurement design and the state with exponentially decaying eigenvalues

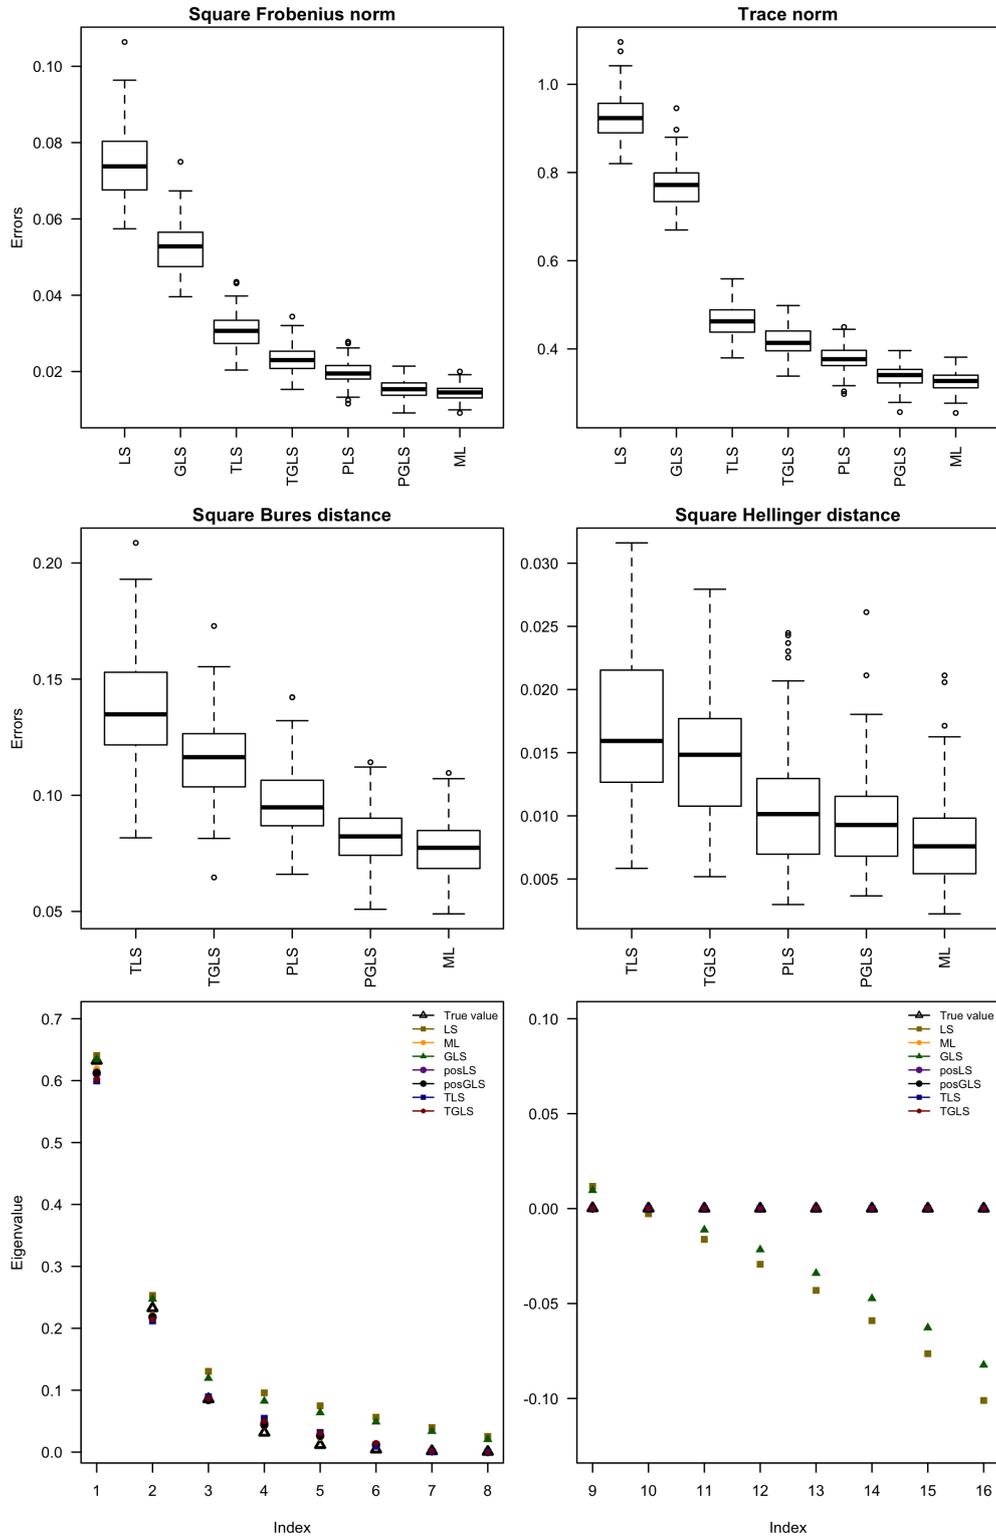


Figure 7.13 Box-plots of the various mean errors for all the estimators, and a 4-qubit random state with exponentially decaying eigenvalues. The number of repetitions per setting is $m = 100$, and we perform 100 iterations of the simulation. The two plots in the final row plot the mean eigenvalues (in descending order) of the various estimates in comparison to those of the true state.

Estimation errors for the random bases measurement design and the state with exponentially decaying eigenvalues

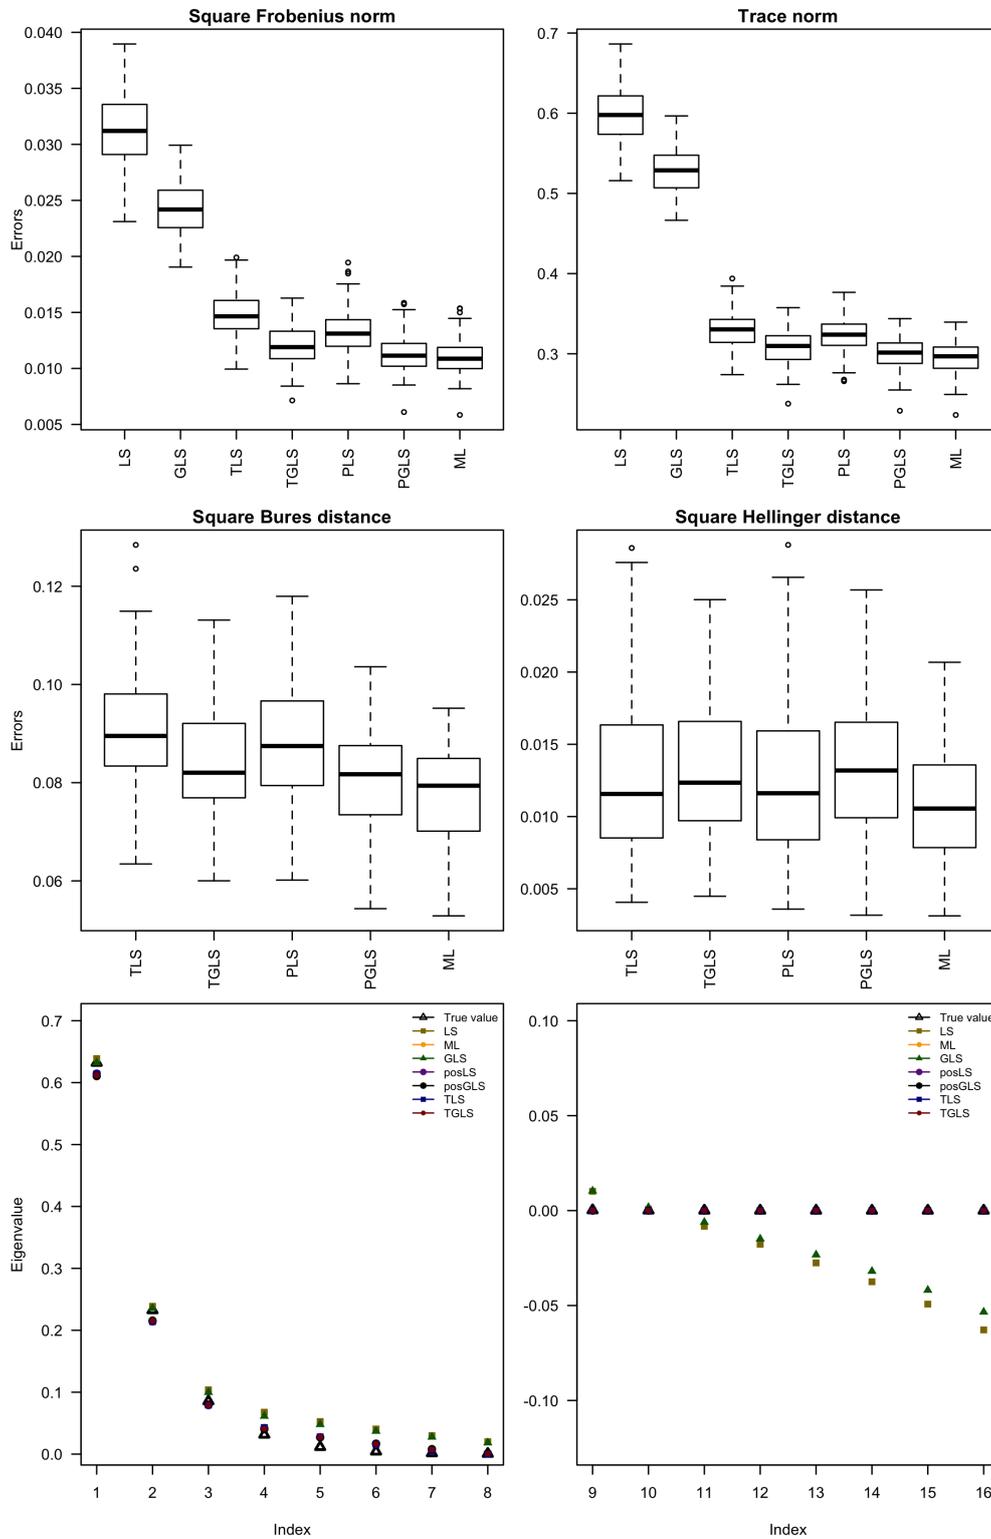


Figure 7.14 Box-plots of the various mean errors for all the estimators, and a 4-qubit random state with exponentially decaying eigenvalues. The number of settings is $k = 100$, repetitions per setting is $m = 100$, and we perform 100 iterations of the simulation. The two plots in the final row plot the mean eigenvalues (in descending order) of the various estimates in comparison to those of the true state.

Estimation errors for the Pauli measurement design and the state with two dominant eigenvalues

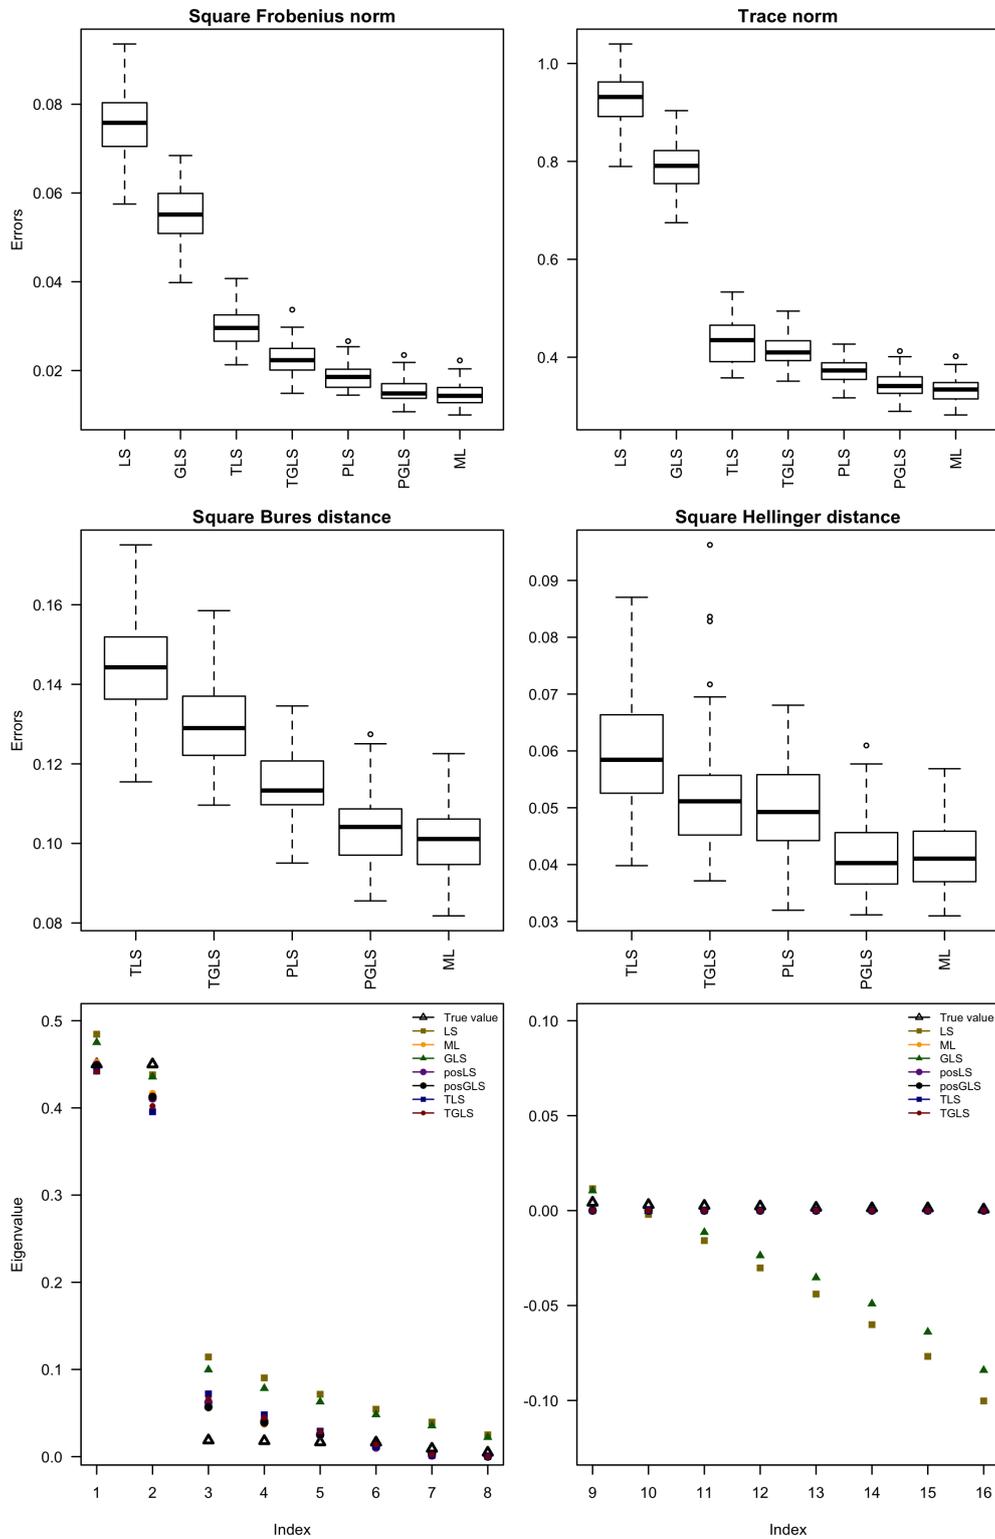


Figure 7.15 Box-plots of the various mean errors for all the estimators, and a 4-qubit random state with two dominant eigenvalues. The number of repetitions per setting is $m = 100$, and we perform 100 iterations of the simulation. The two plots in the final row plot the mean eigenvalues (in descending order) of the various estimates in comparison to those of the true state.

Estimation errors for the random measurement design and the state with two dominant eigenvalues

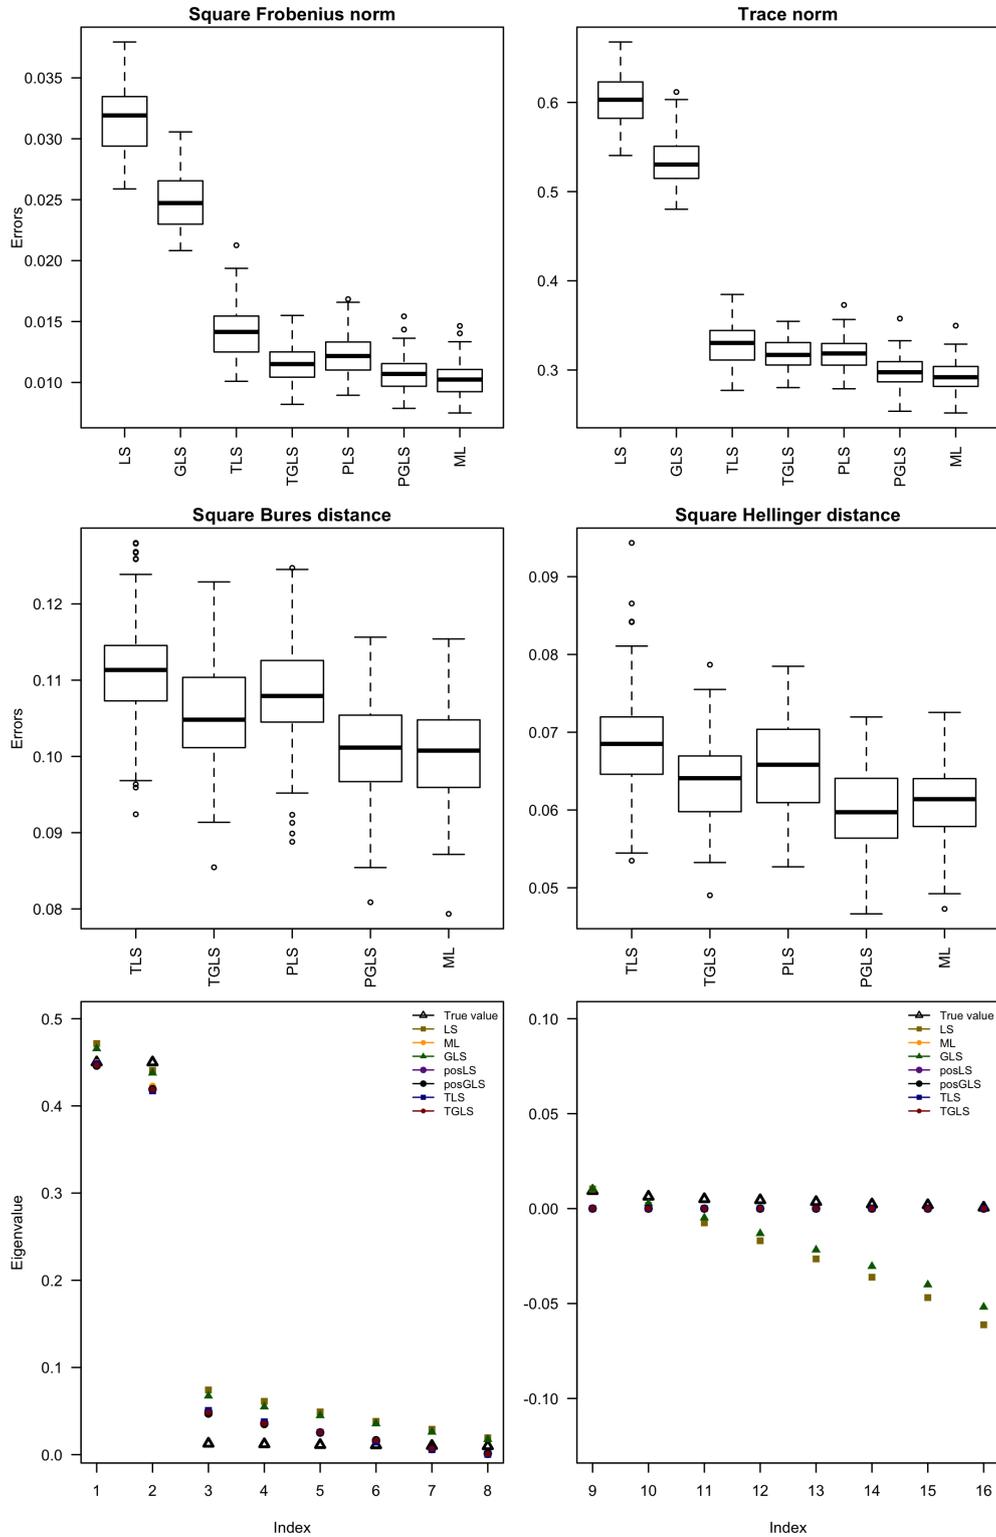


Figure 7.16 Box-plots of the various mean errors for all the estimators, and a 4-qubit random state with two dominant eigenvalues. The number of settings is $k = 100$, repetitions per setting is $m = 100$, and we perform 100 iterations of the simulation. The two plots in the final row plot the mean eigenvalues (in descending order) of the various estimates in comparison to those of the true state.

Estimation errors for the Pauli measurement design and the state with four dominant eigenvalues

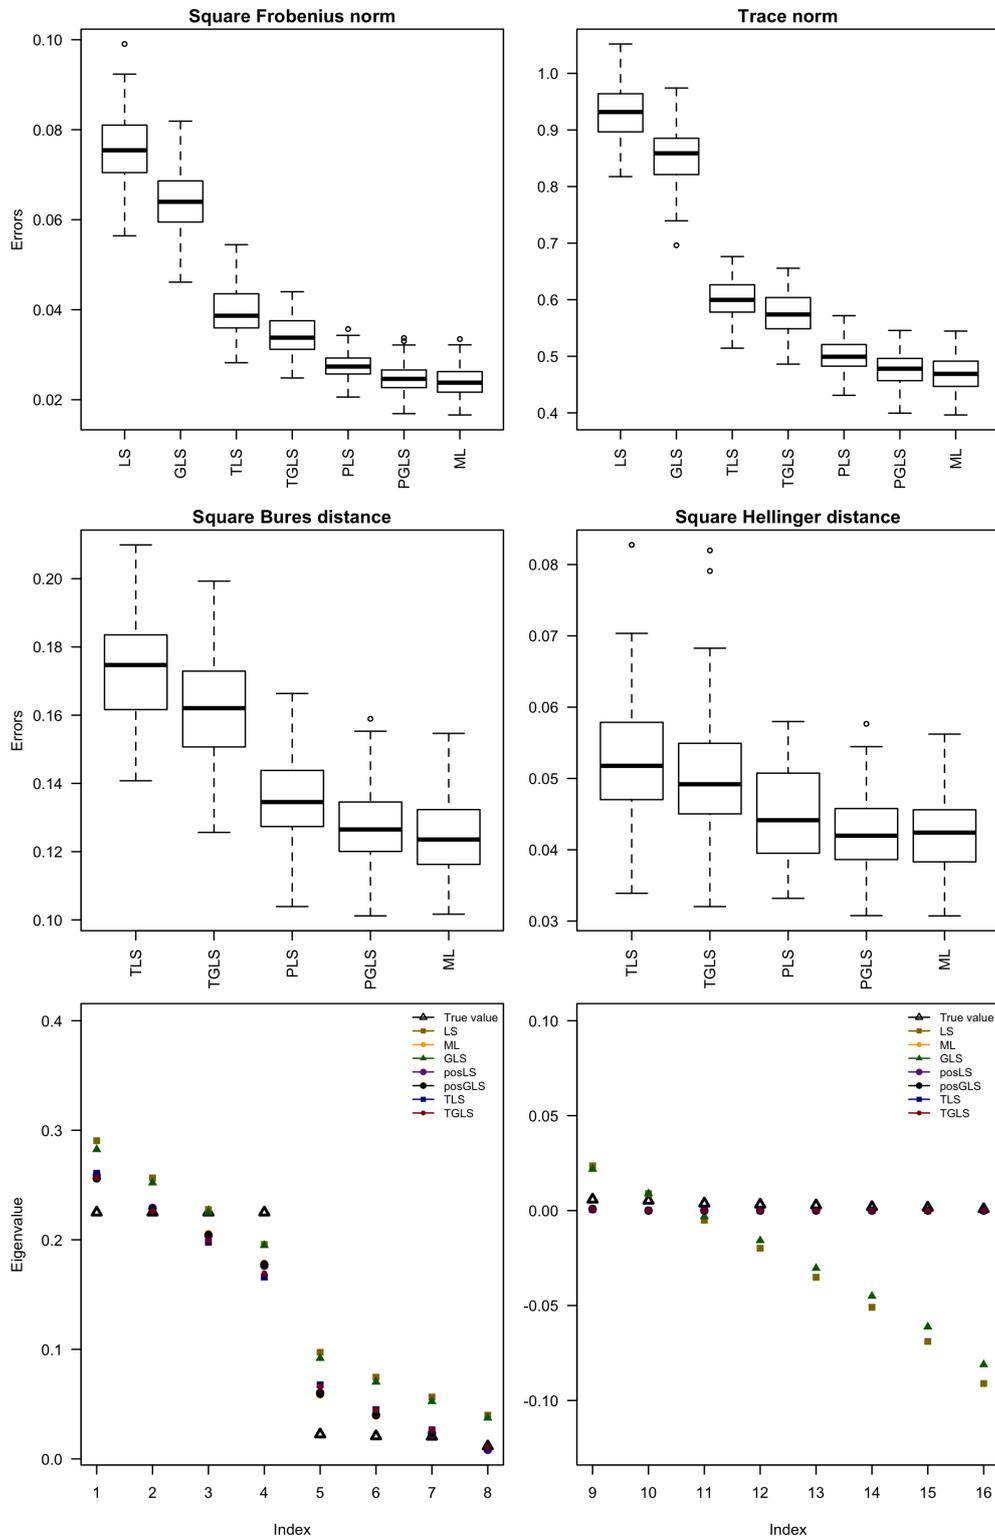


Figure 7.17 Box-plots of the various mean errors for all the estimators, and a 4-qubit random state with four dominant eigenvalues. The number of repetitions per setting is $m = 100$, and we perform 100 iterations of the simulation. The two plots in the final row plot the mean eigenvalues (in descending order) of the various estimates in comparison to those of the true state.

Estimation errors for the random measurement design and the state with four dominant eigenvalues

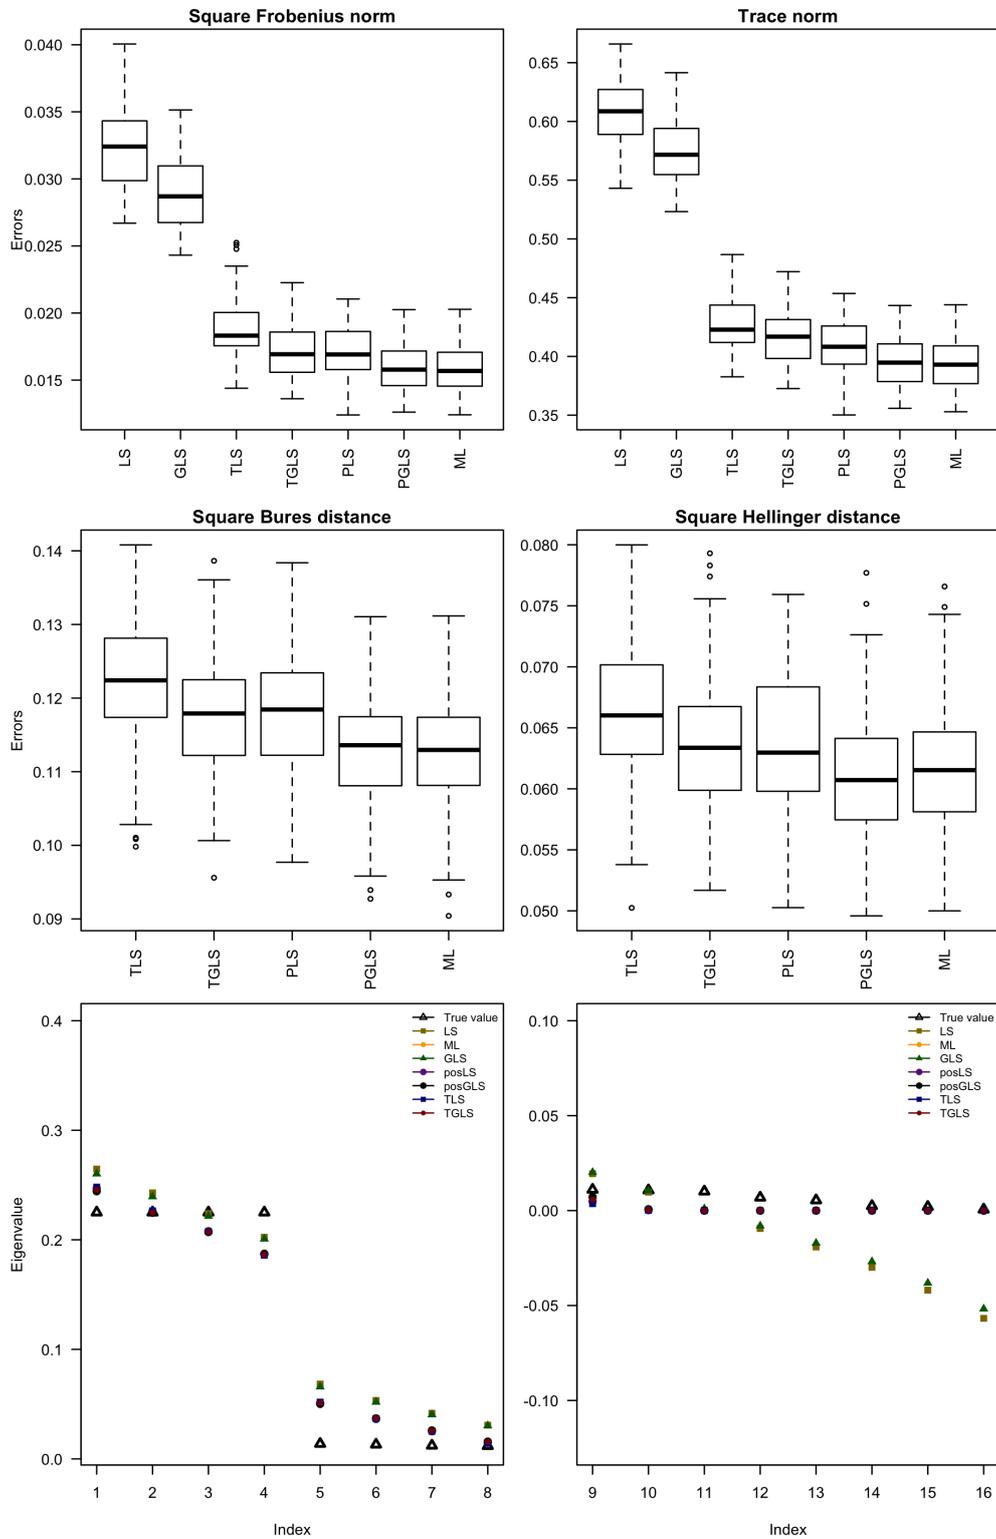


Figure 7.18 Box-plots of the various mean errors for all the estimators, and a 4-qubit random state with four dominant eigenvalues. The number of settings is $k = 100$, repetitions per setting is $m = 100$, and we perform 100 iterations of the simulation. The two plots in the final row plot the mean eigenvalues (in descending order) of the various estimates in comparison to those of the true state.

Estimation errors for the Pauli measurement design and the GHZ state

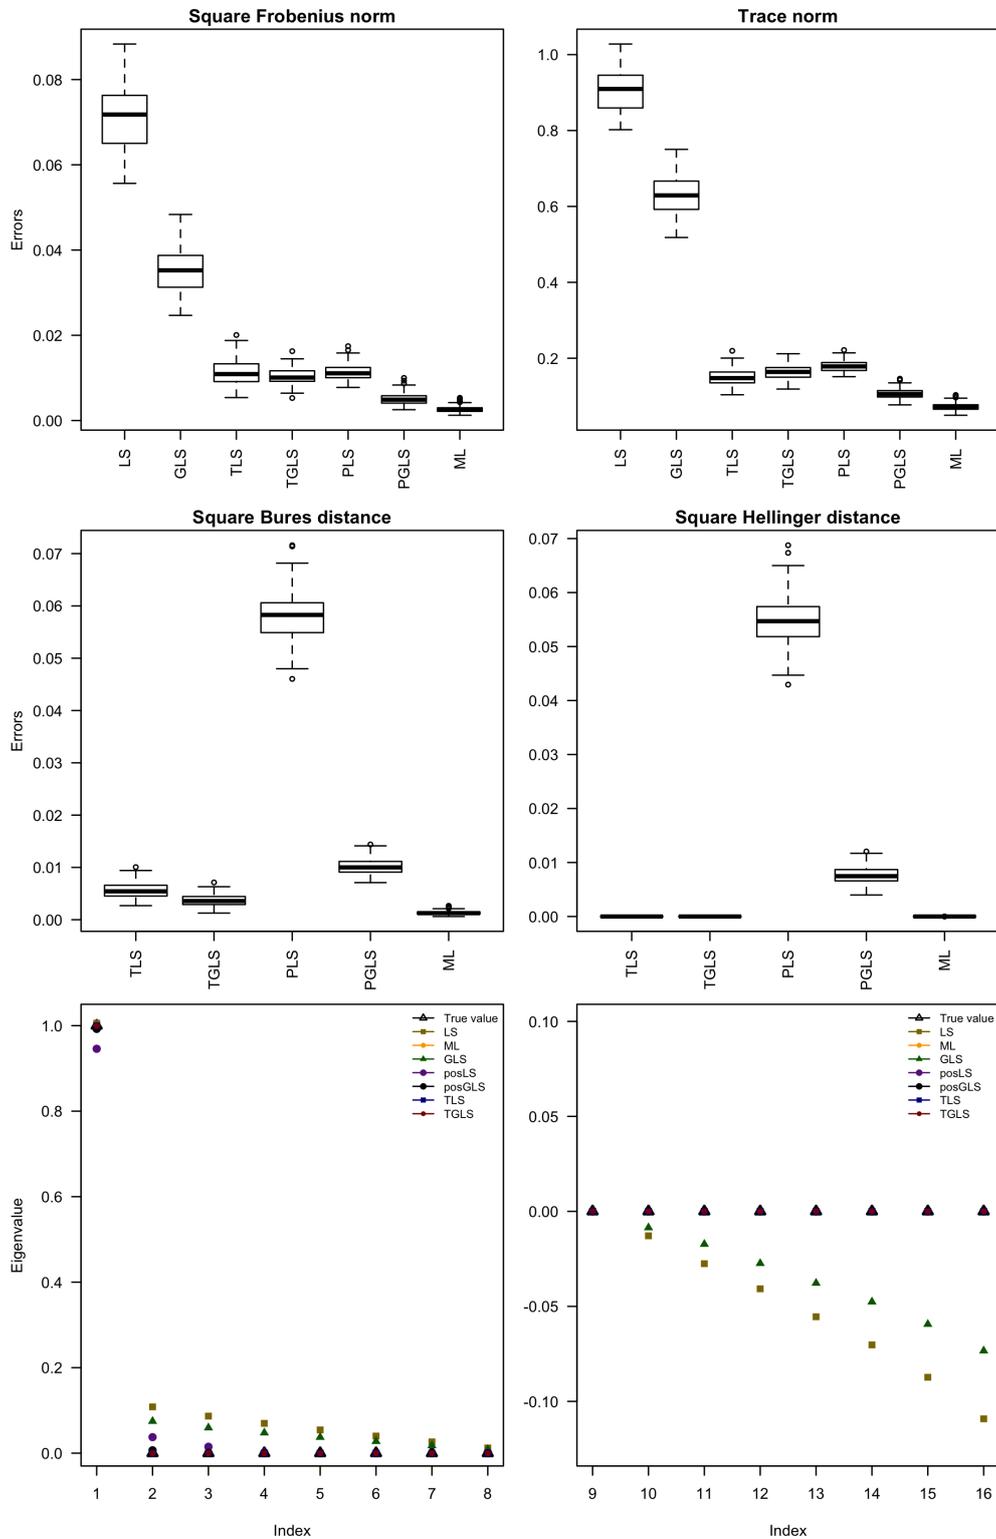


Figure 7.19 Box-plots of the various mean errors for all the estimators, and a 4-qubit GHZ state. The number of repetitions per setting is taken to be $m = 100$, and we perform 100 iterations of the simulation. The two plots in the final row plot the mean eigenvalues (in descending order) of the various estimates in comparison to those of the true state.

Estimation errors for the random bases measurement design and the GHZ state

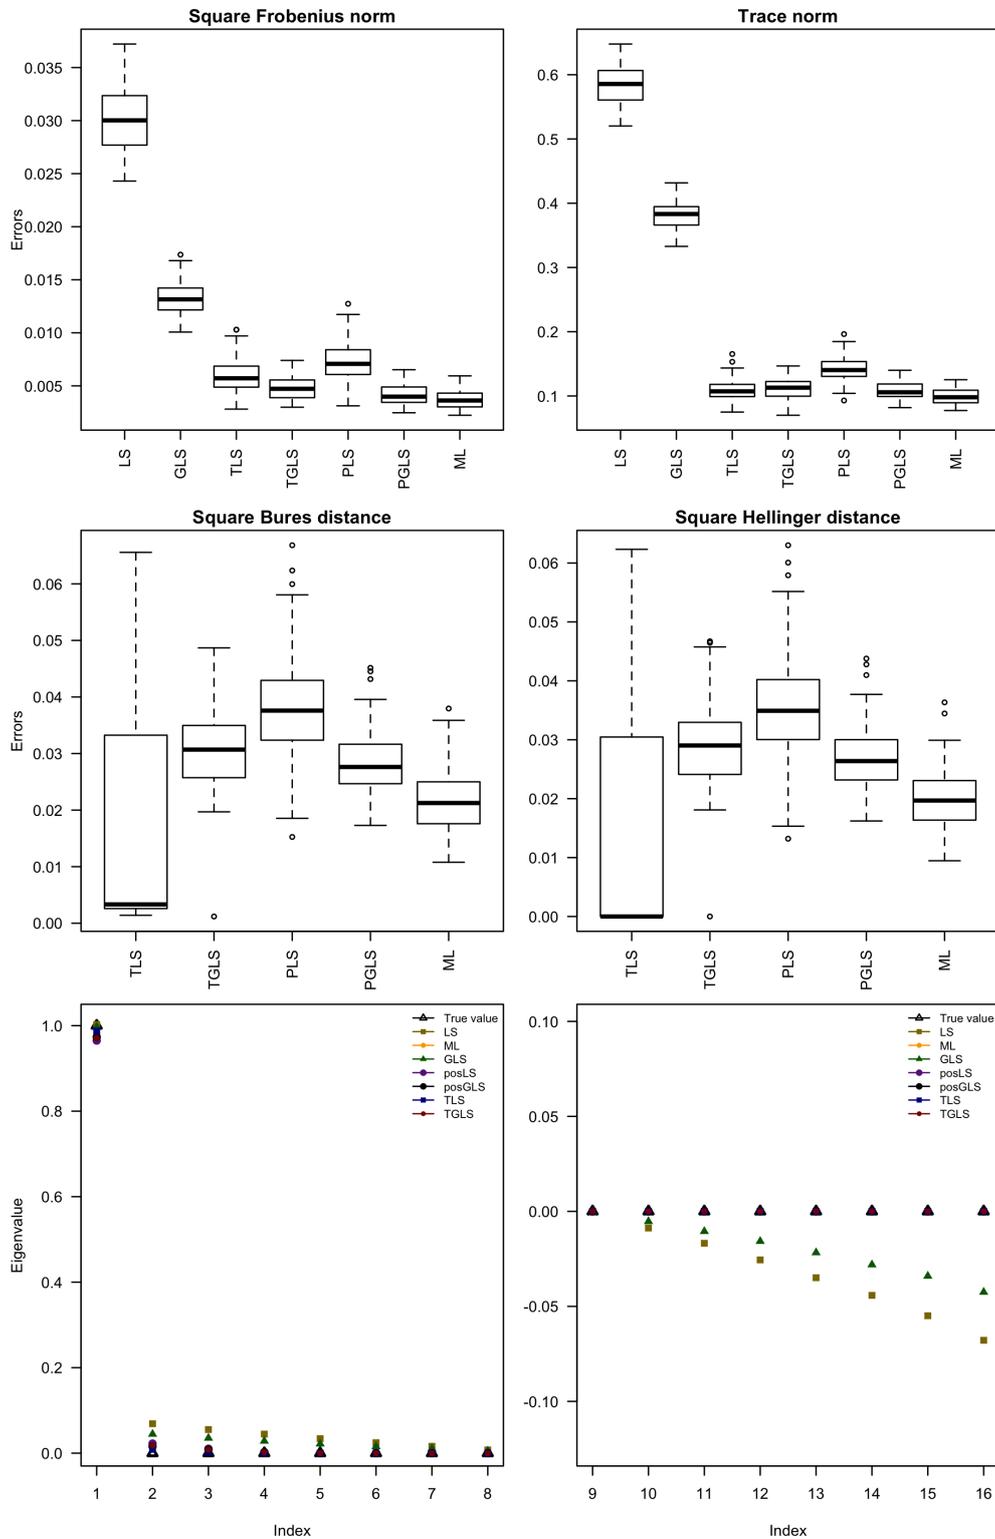


Figure 7.20 Box-plots of the various mean errors for all the estimators, and a 4-qubit GHZ state. The number of settings is $k = 100$, repetitions per setting $m = 100$, and we perform 100 iterations of the simulation. The two plots in the final row plot the mean eigenvalues (in descending order) of the various estimates in comparison to those of the true state.

Estimation errors for the Pauli measurement design and a random pure state

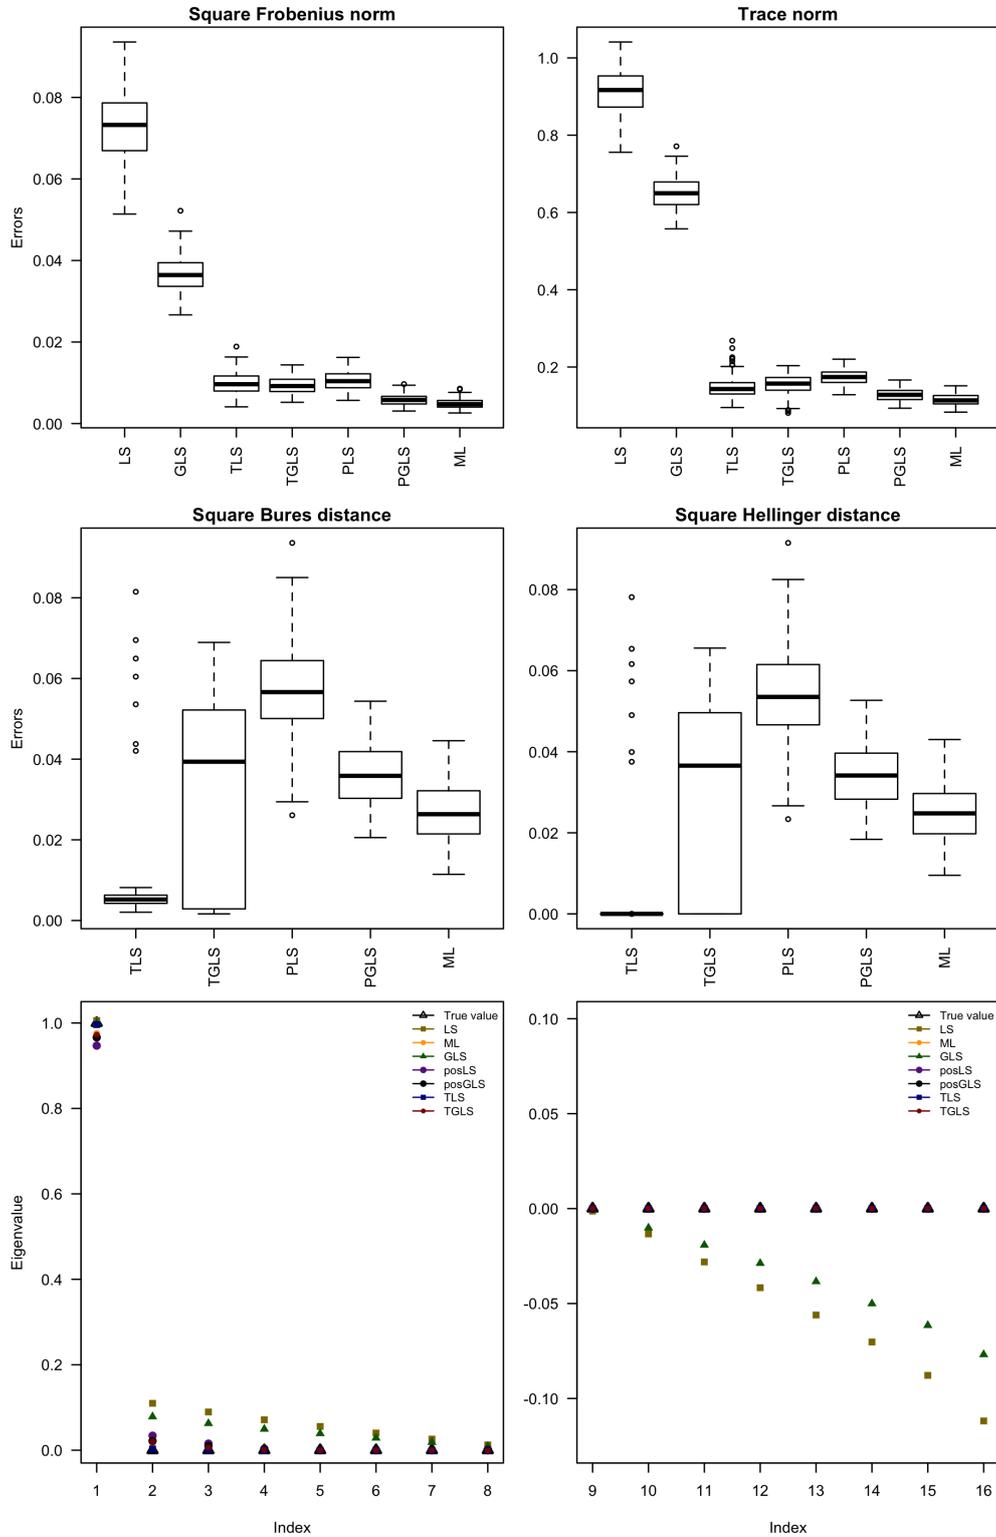


Figure 7.21 Box-plots of the various mean errors for all the estimators, and a random 4-qubit pure state. The number of repetitions per setting is $m = 100$, and we perform 100 iterations of the simulation. The two plots in the final row plot the mean eigenvalues (in descending order) of the various estimates in comparison to those of the true state.

7.7 The Apps

We now introduce two web-based applications we have developed for performing tomographic simulations online. The first application ‘Dashboard: Comparing Estimators’ makes it possible for the user to reproduce the simulations of standard tomography presented in this chapter. More importantly however, it allows the user to upload and carry out simulations for *any* multi-qubit state, provided that its dimension is no larger than 2^4 . Our aim in developing this application is to provide a tool with a user-friendly interface that makes it convenient to compare the performance of several estimators for any states of interest. It also serves to complement the results of the simulation study presented in this chapter.

The second application ‘State Estimation’, generates estimates of an unknown true state from the counts dataset of the standard tomographic experiment. This dataset of counts is uploaded by the user, and may either be simulated or from an actual tomographic experiment. Then for a particular choice of estimator, the final estimated density matrix is made available for download.

Documentation for the applications is made available online, along with sample files for the simulations. Both applications are designed with R Shiny², and are powered by R and MATLAB. They are hosted on the Shiny server maintained by the UoN School of Mathematical Sciences.

7.7.1 Dashboard: Comparing Estimators

Web address : https://shiny.maths.nottingham.ac.uk/shiny/qt_dashboard/

This application is designed to allow the user to perform simulations of the standard tomographic procedure with the Pauli measurement design for any state, and any combination of estimators. The application supports arbitrary states of upto 4 qubits. States can either be uploaded in a CSV format file, or chosen from the options in the dropdown menu. The user can choose from options of randomly generated states, GHZ states, fully mixed states and the equal eigenvalue states as considered in section 7.6. For a particular choice, the application then allows the user to specify the number of qubits and the rank of the state, see figure 7.22.

²<https://shiny.rstudio.com/>

Dashboard: Compare Estimators

[READ ME](#)

Input Panel

Choose/Upload State

Random State

No. of qubits

Choose Rank

Estimators

No. of repetitions per setting

Global Iterations

Compute

Hello,

Welcome to our 'Shiny' interface! This page allows the user to run quantum state tomography simulations for a number of different estimators. The key elements of the interface are as follows:

- Choose/Upload State:** This allows the user to choose from a list of states or upload a density matrix. The uploaded state must be in a CSV file with each entry corresponding to a density matrix element. Please ensure no row/column names or headers are present. A sample file can be downloaded [here](#).
- Estimators:** Details about the estimators can be found in the brief [notes](#).
- Repetitions per setting:** This value indicates the number of times measurements in a particular setting are performed. When choosing the Thresholded LS estimator, please ensure that this number is in multiples of 5. This is for purposes of cross validation.
- Global Repetitions:** This value indicates the number of global iterations of the simulation. A larger value here implies a better estimate of the risk of the estimator.
- Output Panel:** Once the simulations have been completed, a pop-up box will inform the user that the plots can be generated using this panel. The panel generates box-plots for a chosen error function. The definitions of the error functions can be found in the [notes](#).

Output Panel

Choose Error/Loss Function

Y-axis max

Y-axis min

Plot

Figure 7.22 An application for performing tomographic simulations online. The application allows the user to choose or upload a multi-qubit states of various ranks, and compare the performance of several estimators.

In addition to all of the estimators studied in this chapter, the application also makes available the Lasso estimator, defined via the minimisation

$$\hat{\rho} := \arg \min_{\tau \geq 0} \|\mathcal{X} - \mathbf{f}\|^2 + 2\mu \|\tau\|_1. \quad (7.33)$$

The final estimate is the obtained as $\hat{\rho}_{\text{Lasso}} := \hat{\rho} / \text{Tr} \hat{\rho}$. The penalty μ is defined by the user, see figure 7.23. Any combination of estimators can be chosen from the checkbox menu. For a given true state of N qubits and specified value of the number of repetitions m , the counts data is simulated resulting in a $2^N \times 3^N$ dataset \mathcal{D} from which the state is estimated. The number of 'Global Iterations' corresponds to the number of times a dataset \mathcal{D} is simulated and the state estimates evaluated. A larger value here corresponds to a better approximation of the mean error or risk of an estimator.

Dashboard: Compare Estimators

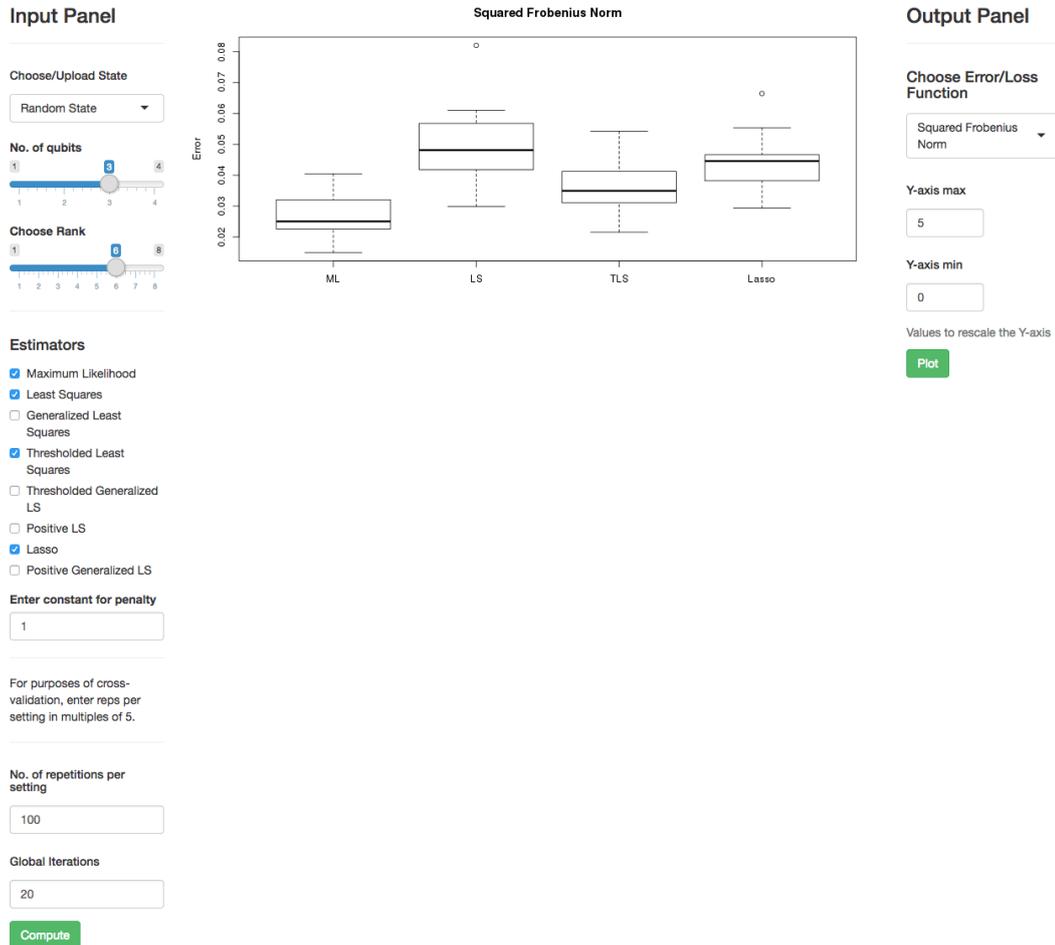


Figure 7.23 Once the simulations have completed, box-plots of the estimation errors can be generated for several different error/loss functions from the output panel. The min and max values of the y-axis can be specified to suitably scale the plots.

Once the simulations have completed, a popup informs the user that the ‘Output Panel’ on the right is ready to use. The panel allows the user to generate box-plots in terms of several error functions- the square Frobenius norm, the trace norm, the square Bures distance and the fidelity/infidelity. The panel also provides the option to specify the minimum and maximum values for the y-axis of the plot, allowing the user to focus on a particular range of the errors.

7.7.2 State Estimation

Web address : https://shiny.maths.nottingham.ac.uk/shiny/state_estimation/

This application is designed to generate state estimates directly from a $2^N \times 3^N$ dataset of counts uploaded by the user. Therefore the outcome data from the tomographic experiment is not simulated from a chosen state. The application is designed to comfortably handle datasets of size $2^4 \times 3^4$. For a particular dataset of counts, the file is first processed by clicking the ‘Process file’ button. This prepares the dataset for the estimators, and once the file has been successfully processed a popup appears informing the user that an estimator may be chosen. As before, all of the estimators presented in this chapter along with the Lasso estimator (7.33) are made available. For the Lasso, TLS and TGLS estimators, the penalty and the constants for the threshold are to be specified by the user. Once the estimation has completed the density matrix of the estimate is made available for download in CSV format.

State Estimation

[Upload the file](#)

No file selected

Default max. file size is 5MB

Separator

Comma

Semicolon

Tab

Space

Process file only for each new upload.

Choose Estimator

[READ ME](#)

This interface allows the user to upload a table of counts of outcomes observed from the standard tomography experiment on an n-qubit unknown state. The user can then choose a method to estimate the unknown state.

1. The input counts table needs to be structured as follows. It needs to have 3^n columns corresponding to the different possible measurement settings, and have 2^n rows corresponding to the outcomes per setting. The number of times a given outcome in a setting was observed is therefore index by its row and column position, and all columns sum up to the number of repetitions per settings. A sample data table of counts is available for download [here](#) in CSV format.
2. For each uploaded data set the 'process file' button needs to be clicked only once. This generates all the files necessary for the estimators to run.
3. A description of the estimators can be found in the [brief notes](#).
4. The estimated density matrix can be downloaded in CSV format once the processing has completed.

Figure 7.24 An application to estimate an unknown state from the outcome dataset of the tomographic experiment.

7.8 Conclusion

In this chapter we presented results from an extensive and systematic simulation study comparing the performance of several tomographic estimators for several different combinations of states, ranks, measurement designs, number of copies

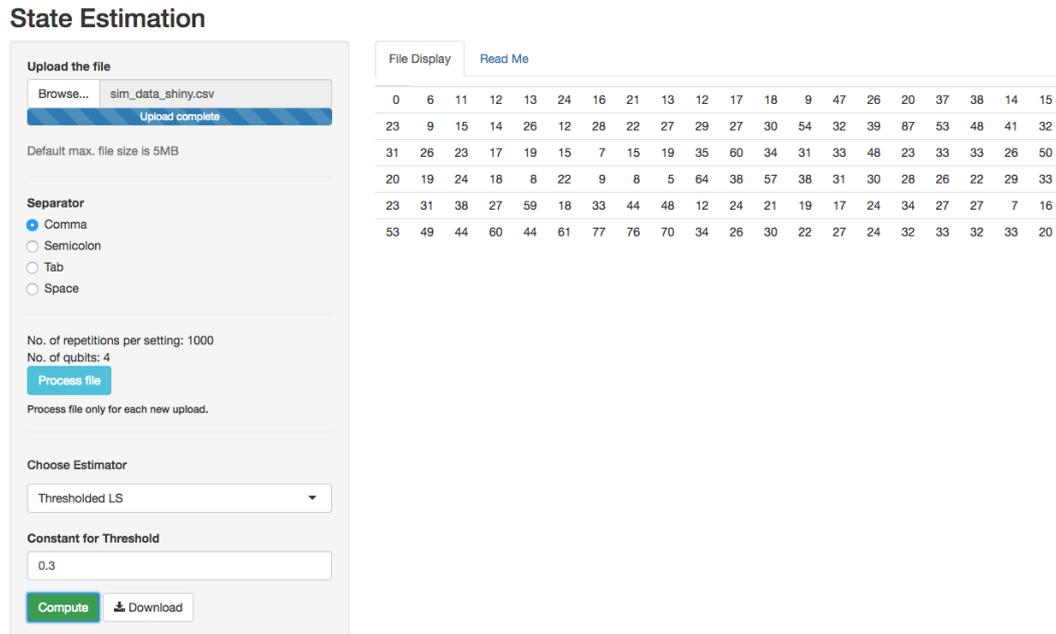


Figure 7.25 Once the state has been estimated for a particular choice of estimator, the resulting density matrix can be downloaded in CSV format.

of the state and loss functions. Apart from the fairly standard and well known estimators like the maximum likelihood (ML) and the least squares (LS), we also introduced some new estimators such as the positive generalised least squares (PGLS) and the thresholded generalised least squares (TGLS). We outlined the equivalence of the PGLS and the ML estimators in the limit of large number of samples, and considered the implications of asymptotic theory for states near the boundary of the positive semi-definite cone of matrices. We analysed the errors for several loss functions and highlighted the difference in performance of the estimators across the various loss functions, specifically illustrating the sensitivity of the Bures distance to the misestimation of small eigenvalues.

We also introduced two web-based applications with the aim of providing tools that will enable their users to perform tomographic simulations online and compare the performance of estimators for any states of interest. These applications also serve to complement and verify the results of the simulation study presented here.

Chapter 8

Conclusions and directions for future work

In this thesis we considered various aspects of quantum state estimation, whose central problem is the devising of estimation schemes for the recovery of an unknown quantum state from an ensemble of n independent and identically prepared systems. We have investigated the performance and efficiency of several estimation methods and measurement designs. In chapters 4 and 5 we explored the possibility of ‘compressive’ recovery of low-rank states with ‘fine-grained’ or ‘raw data’ from the standard tomographic procedure. We demonstrated such compressive recovery both theoretically and numerically, and obtained rates for the minimum number of measurement settings required, by establishing a concentration inequality for the Fisher information matrices associated with the reduced measurement design. There are several open questions and possible avenues for future research related to the research presented in these chapters. As already mentioned in chapter 4, deriving a concentration bound for the Pauli setting case remains an open problem as it requires control over the spectral properties of the Fisher information matrix. In the case of the random basis measurements, it would be fruitful to further study the links between the results in chapter 5 and related work demonstrating compressive recovery with rank-one projectors sampled from unitary t -designs [81, 32, 53], and recovery using *rank- r strictly complete* POVMS [10]. Furthermore, as pointed out in chapter 4, establishing a concentration of the Fisher information provides the correct rate for the minimum number of measurements required, but provides a pessimistic estimate for the constant in front. Also, we have seen in chapter 5 that despite a lack of concentration of the Fisher information for states with small eigenvalues, the *Mean Square Error* (MSE) still concentrates about its optimal. For these reasons, it might be possible to derive stronger results by studying the concentration of the MSE directly.

In terms of the *Mean Infidelity* (MINF) however, we have seen in chapter 5 that compressive recovery of states with small eigenvalues is not possible. More importantly, it is known that for *any* fixed basis measurement design and nearly pure states the MINF demonstrates only a $O(1/\sqrt{n})$ rate of estimation, as opposed to the optimal $O(1/n)$ rate. This motivated the work in chapter 6, where we pose the problem in terms of the maximum Bures risk. We proposed two *adaptive* qubit estimation strategies, one based on local measurements and the other on collective global measurements. For both these estimation strategies we demonstrated a $O(1/n)$ scaling for the maximum Bures risk. We also investigated the possibility of deriving a minimax optimal estimator for the Bures risk. We demonstrated that a minimax optimal estimator for the mixed qubit state can be immediately obtained given a minimax optimal estimator for the binomial parameter under the Hellinger loss function. To the best of our knowledge, such an estimator for the Hellinger loss function is not currently known. This suggests an area of future research that is of great importance not just for classical estimation theory, but for quantum state estimation as well. Another avenue for future work is to extend and generalise the results presented in this chapter to the multi-qubit case.

Finally in chapter 7 we systematically compared the performance of several estimation schemes in an extensive simulation study. This work serves to highlight the suitability of various estimators- both standard and newly defined- for states of different ranks, purity and dimensions. To complement this study, we have developed two web-based applications that will enable the performing of tomographic simulations online for arbitrary user-defined states and for all of the estimators defined in chapter 7. We aim to update these apps with more estimators of interest and make it possible to perform simulations with different measurement designs.

Bibliography

- [1] Adesso, G., Ragy, S., and Lee, A. R. (2014). Continuous variable quantum information: Gaussian states and beyond. *Open Systems and amp; Information Dynamics*, 21(01n02):1440001.
- [2] Ahlswede, R. and Winter, A. (2002). Strong converse for identification via quantum channels. *IEEE Transactions on Information Theory*, 48:569–579.
- [3] Alquier, P., Butucea, C., Hebiri, M., Meziani, K., and Morimae, T. (2013). Rank-penalized estimation of a quantum system. *Phys. Rev. A*, 88:032113.
- [4] A.T.Rezakhani and S.Alipour (2015). On continuity of quantum fisher information. *arXiv:1507.01736*.
- [5] Aude Craik, D. P. L., Linke, N. M., Sepiol, M. A., Harty, T. P., Goodwin, J. F., Ballance, C. J., Stacey, D. N., Steane, A. M., Lucas, D. M., and Allcock, D. T. C. (2017). High-fidelity spatial and polarization addressing of $^{43}\text{Ca}^+$ qubits using near-field microwave control. *Phys. Rev. A*, 95:022337.
- [6] Audenaert, K. M. R. and Scheel, S. (2009). Quantum tomographic reconstruction with error bars: A kalman filter approach. *New Journal of Physics*, 11:023028.
- [7] Bagan, E., Baig, M., Muñoz Tapia, R., and Rodriguez, A. (2004). Collective versus local measurements in a qubit mixed-state estimation. *Phys. Rev. A*, 69:010304.
- [8] Bagan, E., Ballester, M. A., Gill, R. D., Monras, A., and Muñoz Tapia, R. (2006a). Optimal full estimation of qubit mixed states. *Phys. Rev. A*, 73:032301.
- [9] Bagan, E., Ballester, M. A., Gill, R. D., Muñoz Tapia, R., and Romero-Isart, O. (2006b). Separable measurement estimation of density matrices and its fidelity gap with collective protocols. *Phys. Rev. Lett.*, 97:130501.
- [10] Baldwin, C. H., Deutsch, I. H., and Kalev, A. (2016). Strictly-complete measurements for bounded-rank quantum-state tomography. *Phys. Rev. A*, 93:052105.
- [11] Ballance, C. J., Harty, T. P., Linke, N. M., Sepiol, M. A., and Lucas, D. M. (2016). High-fidelity quantum logic gates using trapped-ion hyperfine qubits. *Phys. Rev. Lett.*, 117:060504.
- [12] Baltagi, B. (2011). *Econometrics*. Number 978-3-642-20059-5. Springer-Verlag Berlin Heidelberg, 5 edition.

- [13] Banaszek, K., D'Ariano, G. M., Paris, M. G. A., and Sacchi, M. F. (1999). Maximum-likelihood estimation of the density matrix. *Physical Review A*, 61.
- [14] Baraniuk, R., Davenport, M., DeVore, R., and Wakin, M. (2008). A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263.
- [15] Belavkin, V. P. (1976). Generalized uncertainty relations and efficient measurements in quantum systems. *Theoretical and Mathematical Physics*, 26(3):213–222.
- [16] Bengtsson, I. and Życzkowski, K. (2008). *Geometry of Quantum States: An Introduction to Quantum Entanglement*. Number 9780521891400. Cambridge University Press.
- [17] Bernardo, J. M. and Smith, A. F. M. (2000). *Bayesian Theory*. Number 978-0-471-49464-5. Wiley.
- [18] Bhatia, R. (1997). *Matrix Analysis*, volume 169 of *Graduate Texts in Mathematics (ISSN 0072-5285)*. Springer New York.
- [19] Blume-Kohout, R. (2010). Hedged maximum likelihood quantum state estimation. *Phys. Rev. Lett.*, 105:200504.
- [20] Bouwmeester, D., Pan, J.-W., Daniell, M., Weinfurter, H., and Zeilinger, A. (1999). Observation of three-photon greenberger-horne-zeilinger entanglement. *Phys. Rev. Lett.*, 82:1345–1349.
- [21] Bowles, P. (2012). *Applications of Local Asymptotic Normality in Quantum Information Theory*. PhD thesis, University of Nottingham.
- [22] Braess, D. and Dette, H. (2004). The asymptotic minimax risk for the estimation of constrained binomial and multinomial probabilities. *Sankhyā: The Indian Journal of Statistics (2003-2007)*, 66(4):707–732.
- [23] Braess, D. and Sauer, T. (2004). Bernstein polynomials and learning theory. *Journal of Approximation Theory*, 128(2):187 – 206.
- [24] Braunstein, S. L. and Caves, C. M. (1994). Statistical distance and the geometry of quantum states. *Phys. Rev. Lett.*, 72:3439–3443.
- [25] Butucea, C., Guță, M., and Kypraios, T. (2015). Spectral thresholding quantum tomography for low rank states. *New Journal of Physics*, 17(11):113050.
- [26] Butucea, C., Guță, M., and Kypraios, T. (2016). Corrigendum: Spectral thresholding quantum tomography for low rank states (2015 new j. phys. 17 113050). *New Journal of Physics*, 18(6):069501.
- [27] Candès, E. and Plan, Y. (2011). Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Transactions on Information Theory*, 57(4):57.

- [28] Candès, E. and Recht, B. (2009). Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(717).
- [29] Candès, E., Romberg, J., and Tao, T. (2006). Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59(8):1207–1223.
- [30] Candès, E. and Wakin, M. (2008). An introduction to compressive sampling. *IEEE Signal Processing Magazine*, 25(2):21 – 30.
- [31] Candès, E. J. (2008). The restricted isometry property and its implications for compressed sensing. *Comptes rendus de l'Academie des Sciences, Serie I*:346(9–10):589–592.
- [32] Candès, E. J., Strohmer, T., and Voroninski, V. (2013). Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Comm. Pure Appl. Math.*, 66(8):1241–1274.
- [33] Carpentier, A., Eisert, J., Gross, D., and Nickl, R. (2015). Uncertainty quantification for matrix compressed sensing and quantum tomography problems. *arXiv:1504.03234v1 [math.ST]*.
- [34] Carpentier, A. and Kim, A. (2015). An iterative hard thresholding estimator for low rank matrix recovery with explicit limiting distribution. *arXiv:1502.04654v2 [math.ST]*.
- [35] Catana, C., Bouten, L., and Guță, M. (2015). Fisher informations and local asymptotic normality for continuous-time quantum markov processes. *Journal of Physics A: Mathematical and Theoretical*, 48(36):365301.
- [36] Cirac, J. I., Ekert, A. K., and Macchiavello, C. (1999). Optimal purification of single qubits. *Phys. Rev. Lett.*, 82:4344–4347.
- [37] Collins, B. and Śniady, P. (2006). Integration with respect to the haar measure on unitary, orthogonal and symplectic group. *Communications in Mathematical Physics*, 264(3):773–795.
- [38] Cortese, J. A. (2003). *Quantum Information Theory: Classical Communication over Quantum Channels*. PhD thesis, California Institute of Technology, Pasadena, California.
- [39] Cramer, M., Plenio, M. B., Flammaria, S. T., Somma, R., Gross, D., Bartlett, S. D., Landon-Cardinal, O., Poulin, D., and Liu, Y.-K. (2010). Efficient quantum state tomography. *Nat Commun*, 1:149.
- [40] Davenport, M., Duarte, M., Eldar, Y., and Kutyniok, G. (2012). *Compressed Sensing: Theory and Applications*. Cambridge University Press.
- [41] Ferrie, C. (2014). Self-guided quantum tomography. *Phys. Rev. Lett.*, 113:190404.
- [42] Ferrie, C. and Blume-Kohout, R. (2012). Estimating the bias of a noisy coin. *AIP Conf. Proc.*, 1443(1):14–21.

- [43] Ferrie, C. and Blume-Kohout, R. (2016). Minimax quantum tomography: Estimators and relative entropy bounds. *Phys. Rev. Lett.*, 116:090407.
- [44] Fischer, D. G., Kienle, S. H., and Freyberger, M. (2000). Quantum-state estimation by self-learning measurements. *Phys. Rev. A*, 61:032306.
- [45] Flammia, S. T., Gross, D., Liu, Y.-K., and Eisert, J. (2012). Quantum tomography via compressed sensing: Error bounds, sample complexity and efficient estimators. *New Journal of Physics*, 14:095022.
- [46] Fujiwara, A. (2006). Strong consistency and asymptotic efficiency for adaptive quantum estimation problems. *Journal of Physics A: Mathematical and General*, 39(40):12489.
- [47] Gao, W.-B., Lu, C.-Y., Yao, X.-C., Xu, P., Guhne, O., Goebel, A., Chen, Y.-A., Peng, C.-Z., Chen, Z.-B., and Pan, J.-W. (2010). Experimental demonstration of a hyper-entangled ten-qubit schrodinger cat state. *Nat Phys*, pages 331–335.
- [48] Gill, R. D. and Guță, M. (2013). On asymptotic quantum statistical inference. *From Probability to Statistics and Back: High-Dimensional Models and Processes*, 9(10.1214/12-IMSCOLL909):105–127.
- [49] Gill, R. D. and Massar, S. (2000). State estimation for large ensembles. *Phys. Rev. A*, 61:042312.
- [50] Goncalves, D. S., Gomes-Ruggiero, M. A., Lavor, C., Farias, O. J., and Ribeiro, P. H. S. (2012). Local solutions of maximum likelihood estimation in quantum state tomography. *Quantum Information and Computation*, 12(9):775–790.
- [51] Gonçalves, D. S., Lavor, C., Gomes-Ruggiero, M. A., Cesário, A. T., Vianna, R. O., and Maciel, T. O. (2013). Quantum state tomography with incomplete data: Maximum entropy and variational quantum tomography. *Phys. Rev. A*, 87:052140.
- [52] Granade, C., Ferrie, C., and Flammia, S. T. (2017). Practical adaptive quantum tomography. *New Journal of Physics*, 19(11):113017.
- [53] Gross, D., Krahmer, F., and Kueng, R. (2015). A partial derandomization of phaselift using spherical designs. *Journal of Fourier Analysis and Applications*, 21(2):229–266.
- [54] Gross, D., Liu, Y.-K., Flammia, S. T., Becker, S., and Eisert, J. (2010). Quantum state tomography via compressed sensing. *Phys. Rev. Lett.*, 105:150401.
- [55] Gross, D. and Nemes, V. (2010). Note on sampling without replacing from a finite collection of matrices. *CoRR*, abs/1001.2738.
- [56] Guță, M. (2011). Fisher information and asymptotic normality in system identification for quantum markov chains. *Phys. Rev. A*, 83:062324.

- [57] Guță, M., Janssens, B., and Kahn, J. (2008). Optimal estimation of qubit states with continuous time measurements. *Communications in Mathematical Physics*, 277(1):127–160.
- [58] Guță, M. and Jenčová, A. (2007). Local asymptotic normality in quantum statistics. *Communications in Mathematical Physics*, 276(2):341–379.
- [59] Guță, M. and Kahn, J. (2006). Local asymptotic normality for qubit states. *Phys. Rev. A*, 73:052108.
- [60] Guță, M. and Kiukas, J. (2015). Equivalence classes and local asymptotic normality in system identification for quantum markov chains. *Communications in Mathematical Physics*, 335(3):1397–1428.
- [61] Guță, M., Kypraios, T., and Dryden, I. (2012). Rank-based model selection for multiple ions quantum tomography. *New Journal of Physics*, 14:105002.
- [62] Häffner, H., Hänsel, W., Roos, C. F., Benhelm, J., Chek-al kar, D., Chwalla, M., Körber, T., Rapol, U. D., Riebe, M., Schmidt, P. O., Becher, C., Gühne, O., Dür, W., and Blatt, R. (2005). Scalable multiparticle entanglement of trapped ions. *Nature*, 438:643–646.
- [63] Hannemann, T., Reiss, D., Balzer, C., Neuhauser, W., Toschek, P. E., and Wunderlich, C. (2002). Self-learning estimation of quantum states. *Phys. Rev. A*, 65:050303.
- [64] Haroche, S. and Raimond, J.-M. (2013). *Exploring the Quantum: Atoms, Cavities, and Photons*. Oxford University Press.
- [65] Harty, T. P., Allcock, D. T. C., Ballance, C. J., Guidoni, L., Janacek, H. A., Linke, N. M., Stacey, D. N., and Lucas, D. M. (2014). High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit. *Phys. Rev. Lett.*, 113:220501.
- [66] Hayashi, M. (2017). *Quantum Information Theory: Mathematical Foundation*. Number 978-3-662-49723-4 in Graduate Texts in Physics. Springer, Berlin, Heidelberg.
- [67] Helstrom, C. W. (1976). *Quantum Detection and Estimation Theory*, volume 123. Academic Press.
- [68] Hillery, M., Bužek, V., and Berthiaume, A. (1999). Quantum secret sharing. *Phys. Rev. A*, 59:1829–1834.
- [69] Holevo, A. S. (2011). *Probabilistic and Statistical Aspects of Quantum Theory*, volume 1. Edizioni della Normale, 1 edition.
- [70] Hou, Z., Zhu, H., Xiang, G.-Y., Li, C.-F., and Guo, G.-C. (2016). Achieving quantum precision limit in adaptive qubit state tomography. *Npj Quantum Information*, 2:16001 EP.
- [71] James, D. F. V., Kwiat, P. G., Munro, W. J., and White, A. G. (2001). Measurement of qubits. *Phys. Rev. A*, 64:052312.

- [72] Jing, L., Xiao-Xing, J., Wei, Z., and Xiao-Guang, W. (2014). Quantum fisher information for density matrices with arbitrary ranks. *Communications in Theoretical Physics*, 61(1):45.
- [73] Kabanava, M., Kueng, R., Rauhut, H., and Terstiege, U. (2016). Stable low-rank matrix recovery via null space properties. *Information and Inference: A Journal of the IMA*, 5(4):405–441.
- [74] Kahn, J. (2008). *Quantum Local Asymptotic Normality and Other Questions of Quantum Statistics*. PhD thesis, Leiden University.
- [75] Kahn, J. and Guță, M. (2009). Local asymptotic normality for finite dimensional quantum systems. *Communications in Mathematical Physics*, 289(2):597–652.
- [76] Kalev, A., Kosut, R. L., and Deutsch, I. H. (2015). Quantum tomography protocols with positivity are compressed sensing protocols. *Npj Quantum Information*, 1.
- [77] Kimmel, S. and Liu, Y.-K. (2017). Phase retrieval using unitary 2-designs. In *2017 International Conference on Sampling Theory and Applications (SampTA)*, pages 345–349.
- [78] Koltchinskii, V. and Xia, D. (2015). Optimal estimation of low rank density matrices. *Journal of Machine Learning Research*, 16:1757–1792.
- [79] Kravtsov, K. S., Straupe, S. S., Radchenko, I. V., Houlsby, N. M. T., Huszár, F., and Kulik, S. P. (2013). Experimental adaptive bayesian tomography. *Phys. Rev. A*, 87:062122.
- [80] Kueng, R. (2015). Low rank matrix recovery from few orthonormal basis measurements. In *2015 International Conference on Sampling Theory and Applications*, pages 402–406. IEEE.
- [81] Kueng, R., Rauhut, H., and Terstiege, U. (2017). Low rank matrix recovery from rank one measurements. *Applied and Computational Harmonic Analysis*, 42(1):88 – 116.
- [82] Leibbrandt, D. R., Labaziewicz, J., Clark, R. J., Chuang, I. L., Epstein, R. J., Ospelkaus, C., Wesenberg, J. H., Bollinger, J. J., Leibfried, D., Wineland, D. J., Stick, D., Sterk, J., Monroe, C., Pai, C.-S., Low, Y., Frahm, R., and Slusher, R. E. (2009). Demonstration of a scalable, multiplexed ion trap for quantum information processing. *Quantum Info. Comput.*, 9(11):901–919.
- [83] Leonhardt, U. (2010). *Essential Quantum Optics From Quantum Measurements to Black Holes*. Number 9780521869782. Cambridge University Press, 1 edition.
- [84] Li, N., Ferrie, C., Gross, J. A., Kalev, A., and Caves, C. M. (2016). Fisher-symmetric informationally complete measurements for pure states. *Phys. Rev. Lett.*, 116:180402.

- [85] Liu, Y. (2011). Universal low-rank matrix recovery from pauli measurements. In *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011. Proceedings of a meeting held 12-14 December 2011, Granada, Spain.*, pages 1638–1646.
- [86] Lucas, D. M., Donald, C. J. S., Home, J. P., McDonnell, M. J., Ramos, A., Stacey, D. N., Stacey, J. P., Steane, A. M., and Webster, S. C. (2003). Oxford ion-trap quantum computing project. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 361(1808):1401–1408.
- [87] Lvovsky, A. I. and Raymer, M. G. (2009). Continuous-variable optical quantum-state tomography. *Rev. Mod. Phys.*, 81:299–332.
- [88] M.A.Davenport (2010). *Random Observations on Random Observations: Sparse Signal Acquisition and Processing*. Thesis, Rice University.
- [89] Mahler, D. H., Rozema, L. A., Darabi, A., Ferrie, C., Blume-Kohout, R., and Steinberg, A. M. (2013). Adaptive quantum state tomography improves accuracy quadratically. *Phys. Rev. Lett.*, 111:183601.
- [90] Monz, T., Schindler, P., Barreiro, J. T., Chwalla, M., Nigg, D., Coish, W. A., Harlander, M., Hänsel, W., Hennrich, M., and Blatt, R. (2011). 14-qubit entanglement: Creation and coherence. *Phys. Rev. Lett.*, 106:130506.
- [91] Moroder, T., Hyllus, P., Tóth, G., Schwemmer, C., Niggelbaum, A., Gaile, S., Gühne, O., and Weinfurter, H. (2012). Permutationally invariant state reconstruction. *New Journal of Physics*, 14(10):105001.
- [92] Nielson, M. A. and Chuang, I. L. (2011). *Quantum Computation and Quantum Information*. Number 9781107002173. Cambridge University Press.
- [93] Okamoto, R., Iefuji, M., Oyama, S., Yamagata, K., Imai, H., Fujiwara, A., and Takeuchi, S. (2012). Experimental demonstration of adaptive quantum state estimation. *Phys. Rev. Lett.*, 109:130404.
- [94] Oszmaniec, M., Augusiak, R., Gogolin, C., Kołodyński, J., Acín, A., and Lewenstein, M. (2016). Random bosonic states for robust quantum metrology. *Phys. Rev. X*, 6:041044.
- [95] Qi, B., Hou, Z., Li, L., Dong, D., Xiang, G., and Guo, G. (2013). Quantum state tomography via linear regression estimation. *Scientific Reports*, 3.
- [96] Qi, B., Hou, Z., Wang, Y., Dong, D., Zhong, H.-S., Li, L., Xiang, G.-Y., Wiseman, H. M., Li, C.-F., and Guo, G.-C. (2017). Adaptive quantum state tomography via linear regression estimation: Theory and two-qubit experiment. *npj Quantum Information*, 3(1):19.
- [97] Recht, B., Fazel, M., and Parrilo, P. A. (2010). Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52(3):471–501.

- [98] Riofrío, C. A., Gross, D., Flammia, S. T., Monz, T., Nigg, D., Blatt, R., and Eisert, J. (2017). Experimental quantum compressed sensing for a seven-qubit system. *Nature Communications*, 8.
- [99] Robert, C. (2007). *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*. Number 978-0-387-95231-4. Springer-Verlag New York.
- [100] Řeháček, Jaroslav and Englert, Berthold-Georg and Kaszlikowski, Dagomir (2004). Minimal qubit tomography. *Phys. Rev. A*, 70:052321.
- [101] Řeháček, Jaroslav and Hradil, Zdeněk and Knill, E. and Lvovsky, A. I. (2007). Diluted maximum-likelihood algorithm for quantum tomography. *Phys. Rev. A*, 75:042108.
- [102] Scholten, T. L. and Blume-Kohout, R. (2017). Behavior of the maximum likelihood in quantum state tomography. *arXiv:1609.04385*.
- [103] Schwemmer, C., Tóth, G., Niggelbaum, A., Moroder, T., Gross, D., Gühne, O., and Weinfurter, H. (2014). Experimental comparison of efficient tomography schemes for a six-qubit state. *Phys. Rev. Lett.*, 113:040503.
- [104] Steffens, A., Riofrío, C. A., McCutcheon, W., Roth, I., Bell, B. A., McMillan, A., Tame, M. S., Rarity, J. G., and Eisert, J. (2017). Experimentally exploring compressed sensing quantum tomography. *Quantum Science and Technology*, 2(2):025005.
- [105] Straupe, S. S. (2016). Adaptive quantum tomography. *JETP Letters*, 104(7):510–522.
- [106] Sugiyama, T., Turner, P. S., and Murao, M. (2012). Effect of non-negativity on estimation errors in one-qubit state tomography with finite data. *New Journal of Physics*, 14(8):085005.
- [107] Tao, T. (2010). 254 notes 1: Concentration of measure. <https://terrytao.wordpress.com/>.
- [108] Teo, Y. S., Englert, B.-G., Řeháček, J., Hradil, Z., and Mogilevtsev, D. (2012). Verification of state and entanglement with incomplete tomography. *New Journal of Physics*, 14:105020.
- [109] Teo, Y. S., Zhu, H., Englert, B.-G., Řeháček, J., and Hradil, Z. c. v. (2011). Quantum-state reconstruction by maximizing likelihood and entropy. *Phys. Rev. Lett.*, 107:020404.
- [110] Tóth, G., Wieczorek, W., Gross, D., Krischek, R., Schwemmer, C., and Weinfurter, H. (2010). Permutationally invariant quantum tomography. *Phys. Rev. Lett.*, 105:250403.
- [111] Tumulka, R. (2009). *POVM (Positive Operator Value Measure)*, pages 480–484. Springer Berlin Heidelberg, Berlin, Heidelberg.
- [112] Vaart, A. W. v. d. (1998). *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.

-
- [113] Vogel, K. and Risken, H. (1989). Determination of quasiprobability distributions in terms of probability distributions for the rotated quadrature phase. *Phys. Rev. A*, 40:2847–2849.
- [114] Wang, Q. and Jing, B. (2000). Local asymptotic normality and asymptotical minimax efficiency of the mle under random censorship. *Science in China Series A: Mathematics*, 43(6):591–600.
- [115] Weedbrook, C., Pirandola, S., García-Patrón, R., Cerf, N. J., Ralph, T. C., Shapiro, J. H., and Lloyd, S. (2012). Gaussian quantum information. *Rev. Mod. Phys.*, 84:621–669.
- [116] Wiseman, H. M. and Milburn, G. J. (2014). *Quantum Measurement and Control*. Cambridge University Press.
- [117] Xia, D. and Koltchinskii, V. (2016). Estimation of low rank density matrices: Bounds in Schatten norms and other distances. *Electronic Journal of Statistics*, 10(2):2717–2745.
- [118] Young, G. A. and Smith, R. L. (2005). *Essentials of Statistical Inference*. Cambridge University Press.
- [119] Zhang, X., Weng, J., Lu, W., Li, X., Luo, W., and Tan, X. (2017). Greenberger-horne-zeilinger states-based blind quantum computation with entanglement concentration. *Scientific Reports*, 7(1):11104.
- [120] Życzkowski, Karol and Sommers, Hans-Jürgen (2005). Average fidelity between random quantum states. *Phys. Rev. A*, 71:032313.