JONAS KAHN

Normalité asymptotique locale quantique et autres questions de statistiques quantiques, 2009

N° D'ORDRE : xxxx

UNIVERSITÉ PARIS–SUD — FACULTÉ DES SCIENCES D'ORSAY

THÈSE

présentée pour obtenir le grade de

DOCTEUR EN SCIENCES DE L'UNIVERSITÉ PARIS XI

Spécialité : **Mathématiques**

par

Jonas KAHN

NORMALITÉ ASYMPTOTIQUE LOCALE QUANTIQUE et AUTRES QUESTIONS DE STATISTIQUES QUANTIQUES

Soutenue publiquement le **30 septembre 2009** devant la commission d'examen :

M.	Stéphane	**ATTAL**	Université Lyon–I	Président
Mme.	Cristina	**BUTUCEA**	Université Lille–I	Rapporteur
M.	Richard	**GILL**	Université Leiden	Directeur
M.	Mădălin	**GUŢĂ**	Université Nottingham	Examinateur
M.	Alexander	**HOLEVO**	Institut Steklov, Moscou	Rapporteur
M.	Pascal	**MASSART**	Université Paris–XI, Orsay	Directeur

À qui s'en va

Tous les chapitres de cette thèse correspondent à des articles publiés.

Chapitre 2 (Model selection for quantum homodyne tomography)

J. Kahn. Model selection for quantum homodyne tomography. *Accepté par ESAIM : Probability and Statistics.* arXiv :0712.2912.

Chapitre 3 (Discrimination)

G. M. D'Ariano, M. F. Sacchi, and J. Kahn. Minimax quantum state discrimination. *Phys. Rev. A*, 72 :032310, 2005. arXiv :quant-ph/0504048. DOI :10.1103/PHYSREVA.72.032310.

G. M. D'Ariano, M. F. Sacchi, and J. Kahn. *Phys. Rev. A*, 72 :052302, 2005. arXiv :quant-ph/0507081. DOI :10.1103/PHYSREVA.72.052302.

Chapitre 4 (Fast estimation of unitary operations)

Fast rate estimation of unitary operations in $SU(d)$. *Phys. Rev. A*, 75 :022326, 2007. arXiv :quant-ph/0603115. DOI :10.1103/PHYSREVA.75.022326.

Chapitre 5 (Clean positive operator valued measures)

Clean positive operator valued measures for qubits and similar cases. *J. Phys. A, Math. Theor.*, 40 :4817–4832, 2007. arXiv :quant-ph/0603117. DOI :10.1088/1751-8113/40/18/009.

Chapitre 6 (Complementary subalgebras)

J. Kahn and D Petz. Complementary reductions for two qubits. *J. Math. Phy.*, 48 : 012107, 2007. arXiv :quant-ph/0608227. DOI :10.1063/1.2424883.

Chapitre 7 (QLAN for qubits)

M. Guţă and J. Kahn. Local asymptotic normality for qubit states. *Phys. Rev. A*, 73 : 052108, 2006. arXiv :quant-ph/0512075. DOI :10.1103/PHYSREVA.73.052108.

Chapitre 8 (Optimal estimation of qubit states with continuous time measurements)

M. Guţă, B. Janssens, and J. Kahn. Optimal estimation of qubit states with continuous time measurements. *Comm. Math. Phy.*, 277(1) :127 – 160, 2008. arXiv : quant-ph/0608074. DOI :10.1007/S00220-007-0357-5.

Chapitre 9 (QLAN finite dimension)

M. Guţă and J. Kahn. Local asymptotic normality for finite-dimensional systems. *Comm. Math. Phys.*, pages 79–+, March 2009. arXiv :0804.3876 10.1007/S00220-009-0787-3.

Remerciements Acknowledgments Ringraziamenti Köszönetnyilvánítás

This thesis has been written under the scientific direction of Richard Gill. I thank him for introducing me to the world of quantum statistics, and to local asymptotic normality. The periods I spent with him were always enlightening, be it on quantum statistics, or on classical statistics, or the way a statistician should behave. I also admire his will of introducing more poeple to this field, and his always suggesting problems of interest.

Pascal Massart a été mon deuxième directeur de thèse. Son aide dans la bataille administrative a été des plus précieuses. Il a également su me fournir la bibliographie idoine quand l'occasion s'y prêtait.

Je ne peux hélas pas remercier Alexander Holevo dans sa langue, lui qui a pratiquement créé le domaine des statistiques quantiques, et qui me fait aujourd'hui l'honneur d'être mon rapporteur.

Cristina Butucea a été mon interlocuteur en France, la seule avec laquelle je puisse dialoguer sur le sujet des statistiques quantiques. Je la remercie également de bien avoir voulu écrire un rapport sur mon travail durant une période où elle avait fort peu de temps.

Je remercie aussi Stéphane Attal de bien vouloir être dans mon jury. J'avais apprécié les brèves discussions que nous avions pu avoir à Lyon.

Je considère Mădă lin Guţă comme un troisième directeur de thèse et mon principal collaborateur. Nous avons mis au point la normalité asymptotique quantique locale forte ensemble. J'ai beaucoup apprécié mes séjours à Nottingham, et sa manière de trouver de nouveaux problèmes de statistiques quantiques à partir de problèmes de statistiques classiques.

Je ne peux évoquer la normalité asymptotique locale quantique (QLAN) sans remercier aussi Bas Janssens pour notre travail sur les équations differentielles stochas-

tiques quantiques. Les discussions avec Anna Jenčová ont approfondi ma compréhension de QLAN, en créant l'autre côté, à savoir l'équivalence faible.

Nelle due settimane che ho passato in Pavia, all'inizio della tesi, Professore d'Ariano mi ha dato parecchio problemi, e abbiamo lavorato insieme, con Massimiliano Sacchi. Ringrazio tutta la squadra locale, che ho ancora visto in congressi durante gli anni seguenti.

Köszönöm Petz Dénes felhívásot Budapestre. Nem csak készültünk el a cíkkeimmel, de sokat tanított is a „CCR algebrák"-ról.

Je regrette de ne pas pouvoir remercier Masahito Hayashi et Kenji Matsumoto en japonais, pour mes séjours dans leur pays en Mars 2007 et Novembre 2008, et pour les discussions que nous avons pu avoir sur les statistiques quantiques, au premier chef sur les bornes de Cramér-Rao.

Je remercie aussi tous ceux qui ont été autour de moi, élèves comme professeurs, quand j'étais à l'ENS. Les discussions scientifiques permanentes, les problèmes que nous nous posions, restent la plus belle formation que j'aie pu avoir.

Merci Yan Pautrat, pour les paroles échangées sur la physique statistique quantique du point de vue mathématique, et pour avoir invité Vladimir Jakšić à donner un cours inspiré.

Merci à ceux qui ont relu ma thèse, Cristina, Patricia Reynaud, Borg, et surtout mon voisin de bureau Sylvain Arlot, dont je me suis outrageusement inspiré dans l'espoir chimérique d'approcher sa clarté.

Merci enfin à Valérie Lavigne, qui m'a été d'une grande aide pour survivre dans l'enfer administratif lié à la thèse.

J'ai essayé d'être aussi neutre que possible dans les lignes précédentes, bien que plusieurs des personnes évoquées méritent le titre d'ami. Mais j'espère que tous ceux qui comptent pour moi, autrement plus que les magnifiques statues de glace que j'étudie, n'ont pas besoin que je les cite pour le savoir.

Table des matières

Chapitre 1

Introduction

Les statistiques, étymologiquement sciences de l'État, peuvent être vues comme l'art de tirer des informations de données. Quoiqu'ils puissent prendre des formes très variées, tout problème de statistiques peut se décomposer en trois morceaux : l'objet étudié, les opérations que nous pouvons effectuer, et la question mathématique précise. En d'autres termes, ce que nous avons, ce que nous pouvons faire, et ce que nous voulons savoir.

Les statistiques quantiques diffèrent des statistiques classiques sur le premier point, ce que nous avons. Par ricochet, elles en diffèrent aussi sur le second, ce que nous pouvons faire.

En statistiques classiques, nous partons en général du résultat des mesures physiques, qui sont modélisées par des variables aléatoires et leurs lois de probabilité correspondantes. En effet, si nous pouvons mesurer les quantités A et B, nous pouvons en théorie mesurer les deux simultanément. Les expériences mesurent souvent toutes les quantités utiles et accessibles. En théorie, «ce que nous pouvons faire» est appliquer n'importe quelle transformation mathématique aux données, éventuellement avec une composante aléatoire supplémentaire. En pratique, la puissance de calcul peut être un facteur limitant.

Dans certains cas, cependant, nous devons considérer d'ores et déjà l'objet étudié, et choisir quelle mesure effectuer. Par exemple, si nous voulons comprendre le fonctionnement d'une boîte noire, nous devons la sonder avec différentes entrées, une nouvelle entrée à chaque fois. Cette thématique relève des «plans d'expérience». «Ce que nous pouvons faire» dépend largement du problème spécifique. Dans le cas de la boîte noire, nous pouvons choisir notre entrée. La description mathématique de ce choix peut varier d'une boîte noire à une autre, cependant. Toutefois, une fois la mesure effectuée, nous avons de nouveau des probabilités, et sommes de retour au paragraphe précédent.

En statistiques quantiques, le plan d'expérience est inévitable. En effet, si nous pouvons mesurer A ou B, les lois même de la physique nous interdisent de mesurer simultanément A et B, en général. Nous devons donc choisir quelle mesure nous apporte les informations les plus utiles. Néanmoins, la mécanique quantique fournit un cadre parallèle à celui des statistiques classiques, qui nous dit exactement «ce que nous pouvons faire». Initialement, «ce que nous avons» est un objet quantique, modélisé par un état quantique. «Ce que nous pouvons faire» est mesurer l'état, et obtenir une variable aléatoire classique, ou bien plus généralement transformer l'état quantique. Les ensembles de mesures et transformations possibles sont précisément définis mathématiquement, ce qui permet un traitement unifié de nombreuses questions.

«Ce que nous voulons savoir» ne diffère guère en statistiques quantiques et classiques. Le plus souvent, nous souhaitons soit résumer les informations contenues dans les données (inférence statistique), soit infirmer une hypothèse ou choisir la meilleure hypothèse dans un ensemble fini (test), soit deviner avec précision le phénomène qui a généré les données (estimation). Les réponses à ces questions sont toutes décrites par un paramètre classique. L'exception est quand nous cherchons à obtenir un objet intrinsèquement quantique, comme par exemple quand nous essayons de cloner le plus précisément possible un état.

La Partie I de cette thèse est consacrée à l'étude d'un certain nombre de systèmes particuliers. Spécifiquement, nous commençons au Chapitre 2.5.3 par un cas où la mesure est déjà effectuée, si bien que le problème devient classique : nous évaluons un état de la lumière par tomographie homodyne. Au Chapitre 3, nous nous demandons comment décider au mieux dans lequel d'un ensemble fini d'états se trouve notre système; au Chapitre 4, nous donnons une procédure d'estimation rapide $(1/n)$ d'une transformation unitaire boîte noire. Les Chapitres 5 et 6 s'attachent davantage à la structure générale des expériences quantiques : le premier est consacré à une relation d'ordre sur les mesures quantiques, et le second à la recherche de sous-systèmes "aussi différents que possible d'un même système quantique, dans le cas le plus simple.

D'un autre côté, nous pouvons avoir des questions très différentes sur un système donné. Pour un tel système, «ce que nous avons» et «ce que nous pouvons faire» restent fixes. Nous pouvons donc nous interroger sur ce que l'on peut dire sur le système lui-même, sans référence à une question particulière. La théorie de la convergence d'expériences en statistiques classiques nous dit avec quelle précision nous pouvons approcher une expérience par une autre. Ainsi nous pouvons traduire toutes les procédures pour une expérience en une procédure pour l'autre expérience. Si bien que nous obtenons une réponse à «ce que nous voulons savoir» dans les deux expériences dès que l'on sait répondre pour l'une d'entre elles.

La Partie 1.7, principale contribution de cette thèse, généralise au monde quantique

le cas le plus basique de convergence d'expériences, à savoir la normalité asymptotique locale. Nous prouvons qu'une expérience assez lisse d'états quantiques indépendants identiquement distribués (i.i.d.) converge vers une *expérience de décalage gaussienne quantique.* L'important est que cette expérience est très bien connue, et tout ce que nous savons à son sujet peut être traduit pour la classe très large des expériences i.i.d lisses.

Le reste de cette introduction commence par préciser les règles des statistiques classiques et quantiques, puis introduit chacun des chapitres de la thèse, et les problématiques correspondantes dans l'ordre donné ci-dessus.

1.1 Statistiques

Nous présentons une autre introduction aux statistiques quantiques à l'usage du statisticien, plus condensée, en Appendice 2.A du Chapitre 2.5.3.

1.1.1 Statistiques Classiques

On pourra consulter Le Cam (1986) et van der Vaart (1998) comme références supplémentaires, entre autres nombreux livres de statistiques. Nous résumons dans le Tableau 1.1, page 26, les ingrédients de base des statistiques classiques. Le Tableau 1.2 adjacent donne les notions quantiques correspondantes.

Ce que nous avons

En statistiques classiques, on nous donne les données, qui peuvent être modélisées par une variable aléatoire X de loi p. On sait par avance que p est dans un ensemble

$$\mathcal{E} = \{p_\theta, \theta \in \Theta\}, \tag{1.1}$$

sans contrainte en général sur l'espace de paramètres Θ. Les lois p_θ sont toutes définies sur le même espace de probabilités $(\Omega, \mathcal{A})$. Cet $\mathcal{E}$ est appelé *expérience* ou *modèle statistique.*

Remarques :

- Les données proviennent souvent de plusieurs mesures, générant autant de variables aléatoires $X_1, \ldots, X_n$, de lois $p_1, \ldots, p_n$ sur des espaces de probabilités potentiellement différents. Toutefois, nous pouvons toujours considérer toutes ces données comme une seule variable aléatoire $X = (X_1, \ldots, X_n)$, de loi $p = p_1 \otimes \cdots \otimes p_n$, et nous restons dans le cadre ci-dessus.
- Quoiqu'il n'y ait pas de contrainte sur Θ à ce point de la théorie, cet ensemble est souvent soit fini soit un sous-ensemble raisonnable de $\mathbb{R}^d$. Le premier cas mène aux statistiques discrètes, et à certaines familles de tests en particulier, et le second aux statistiques paramétriques. Quand Θ est de dimension infinie, nous sommes dans le complexe royaume des statistiques non paramétriques, thème privilégié de la recherche ces dernières années

Exemples : expérience de Bernoulli, expérience de décalage gaussienne

L'espace de probabilité non trivial le plus simple est l'espace à deux éléments $\{0, 1\}$. Une expérience de pile ou face s'écrit

$$\mathcal{E}_{Ber} = \{p_\theta = (\theta, 1-\theta), \theta \in [0,1]\}. \tag{1.2}$$

Une alternative consiste à lancer la pièce n fois. Si on note $X = (X_1, \ldots, X_n)$ le résultat, nous obtenons cette expérience sur $\{0,1\}^{\otimes n}$:

$$\mathcal{E}_{Bin} = \left\{p_\theta : \{X\} \mapsto \theta^{\sum X_i}(1-\theta)^{n-\sum X_i}, \theta \in [0,1]\right\}. \tag{1.3}$$

Quant aux fonctions continues, l'exemple type est le gaussienne. Nous nous intéresserons en particulier aux *expériences de décalage gaussiennes*, où la variance de la gaussienne est fixée et où le paramètre est la moyenne :

$$\mathcal{E}_{gs} = \left\{\mathcal{N}(\theta, \mathcal{I}^{-1}), \theta \in \mathbb{R}^d\right\}, \tag{1.4}$$

où $\mathcal{N}$ est la loi normale, et $\mathcal{I}$ toute matrice définie positive fixée[1].

Ce que nous pouvons faire

Une fois acquises nos données X, comment les traitons-nous ?

La procédure la plus générale consiste à tirer une nouvelle variable aléatoire Y de loi p_X dépendant seulement de X, mesurable en tant que fonction de X.

[1]Nous utilisons cette étrange notation car cette matrice est l'inverse de la matrice d'information de Fisher (1.13).

Nous pouvons voir ce protocole de deux manières. La première est de considérer Y comme la solution à «ce que nous voulons savoir». Alors Y est un *estimateur* (randomisé), typiquement un estimateur de θ, auquel cas nous le dénoterons également $\hat{\theta}$.

Mais nous pouvons également considérer Y comme une nouvelle variable aléatoire, et que nous avons transformé notre expérience. Notre nouvelle expérience est donc constituée de Y de loi q dans un ensemble $\{q_\theta, \theta \in \Theta\}$ sur un espace $(\Omega_1, \mathcal{B})$, de densité[2]

$$q_\theta(y) = T(p_\theta)(y) \hat{=} \int_\Omega p_X(y) \mathrm{d}p_\theta(X). \tag{1.5}$$

La transformation T est un *noyau de Markov*.

Dans le cas classique, ces deux notions sont les mêmes. Toutefois, j'insiste pour les séparer dès maintenant car elles seront différentes dans le cas quantique.

Exemples

Revenons à notre n-échantillon (1.3) de Bernoulli $\mathcal{E}_{Bin}$. Notre espace de probabilité est $\{0,1\}^{\otimes n}$. Nous pouvons utiliser un noyau de Markov de cet espace dans $[0,n] \cap \mathbb{N}$ qui envoie $X = (X_1, \ldots, X_n)$ sur $Y = \sum X_i$. Ici, les p_X sont simplement des pics de Dirac. Nous obtenons alors une loi binomiale pour Y, c'est-à-dire $q_\theta = \mathcal{B}(n, \theta)$. L'expérience correspondante est $\mathcal{E} = \{q_\theta, \theta \in \Theta\}$.

De même, nous pourrions souhaiter construire un estimateur $\hat{\theta}$. Le plus évident est de prendre $X \mapsto \sum X_i/n = Y$. La loi de notre estimateur est alors la binomiale ci-dessus, divisée par n.

Pour ce qui est de trouver un estimateur dans l'expérience (1.4) de décalage gaussienne $\mathcal{E}_{gs}$, la première idée est encore plus simple : on garde X. Le noyau de Markov correspondant est l'identité.

Ce que nous voulons savoir

Nous souhaitons en général obtenir de l'information sur le processus sous-jacent inconnu qui a généré nos données. En d'autres termes, nous voulons deviner le

[2]Nous pourrions aussi bien travailler avec des ensembles non dominés de lois, mais cela ne ferait qu'alourdir les notations. Nous faisons donc l'hypothèse que toutes les lois ont une densité, et utilisons la même lettre pour la loi et la densité.

paramètre[3] θ.

Nous pouvons donner notre solution soit sous la forme d'un intervalle de confiance, soit par une estimation de notre quantité, éventuellement assortie d'estimations de la variance de cette estimation. Cette estimation correspond à la donnée d'un estimateur $\hat{\theta}$ de θ.

Nous voulons construire un bon estimateur. Nous avons donc besoin de pouvoir jauger les estimateurs. En théorie de la décision, nous considérons une *fonction de coût* $c(\theta, \hat{\theta})$. C'est le coût que l'on doit payer si notre estimateur renvoie $\hat{\theta}$ quand le vrai paramètre est θ. Ainsi, les fonctions de coût sont en général nulles sur la diagonale, et augmentent quand θ et $\hat{\theta}$ s'éloignent dans un certain sens.

Une fonction de coût typique quand Θ est discret dénombrable serait $c(\theta, \hat{\theta}) = \delta_{\theta,\hat{\theta}}$. Quand Θ est un sous-ensemble ouvert de $\mathbb{R}^d$, la fonction de coût la plus facile à traiter mathématiquement est le carré de la distance euclidienne $c(\theta, \hat{\theta}) = \|\theta - \hat{\theta}\|_2^2$, ou plus généralement toute fonction de coût quadratique $(\theta - \hat{\theta})^\top G(\theta - \hat{\theta})$ pour une matrice définie positive G, éventuellement dépendant de θ.

Comme $\hat{\theta}$ est une variable aléatoire, nous voulons minimiser l'espérance du coût, appelée le *risque au point* θ :

$$r_\theta(\hat{\theta}) = \int_{\Omega_1} c(\theta, \hat{\theta}) \mathrm{d}q_\theta(\hat{\theta}). \tag{1.6}$$

Cependant, nous ne pouvons minimiser directement cette expression, comme la meilleure stratégie dépend de θ, qui est inconnu. Nous devons donc trouver le moyen de choisir un estimateur efficace pour θ que nous risquons de rencontrer. Il y a essentiellement deux approches. Les physiciens favorisent le paradigme bayésien, où nous admettons l'existence d'une loi *a priori* sur le paramètre θ. Les mathématiciens vont en général préférer les critères minimax, où une stratégie est évaluée par son cas le pire.

Critère bayésien

Nous avons considéré des données X de loi p. Jusqu'ici, nous étions parti du principe que notre seule information était l'expérience, l'ensemble dont nous savons qu'il contient p.

[3] Plus généralement, on peut être intéressé seulement par une fonction f de θ. Cependant, on peut toujours utiliser $(\theta, f(\theta))$ comme paramètre. On choisira dès lors les fonctions de coût introduites ci-dessous pour qu'elles ne dépendent que de $f(\theta)$.

Supposons maintenant que nous avons davantage d'informations. Plus précisément, on nous a dit avant l'expérience que θ est choisi au hasard suivant une loi π. Alors, en moyenne, le meilleur estimateur sera celui qui minimise la moyenne du risque (1.6), c'est-à-dire :

$$\begin{aligned} R_\pi(\hat{\theta}) &= \int_\Theta \pi(\mathrm{d}\theta) r_\theta(\hat{\theta}) \\ &= \int_\Theta \int_{\Omega_1} c(\theta, \hat{\theta}) \mathrm{d}q_\theta(\hat{\theta}) \pi(\mathrm{d}\theta). \end{aligned} \tag{1.7}$$

À partir du risque de Bayes d'un estimateur spécifique $\hat{\theta}$, nous pouvons écrire le risque de Bayes associé à la loi *a priori* π comme l'infimum des risques sur tous les estimateurs $\hat{\theta}$:

$$R_\pi = \inf_{\hat{\theta}} R_\pi(\theta). \tag{1.8}$$

La faiblesse de cette approche vient de ce qu'il n'y aucune raison pour avoir une loi de probabilité *a priori* sur Θ, mis à part la fonction de Dirac sur le vrai θ... qui est exactement ce que nous souhaitons trouver. Nous avons donc à choisir une loi *a priori* et à considérer que c'est la vraie. Le risque de l'estimateur final sera sous-estimé, cependant.

La plus grande force des estimateurs bayésiens est qu'ils utilisent de manière optimale l'information des mesures, à loi *a priori* donnée. La loi *a priori* correspond à de l'information *a priori* en général fausse. De ce fait, les meilleures lois *a priori* sont choisies pour minimiser l'information qu'elles contiennent[4]. Pour un Θ fini, on choisira d'habitude l'équiprobabilité *a priori* sur chaque θ possible. Sur un sous-ensemble précompact ouvert de $\mathbb{R}^d$, on choisira souvent la loi *a priori* de Jeffrey Jeffreys (1946), proportionnelle à la racine carrée de l'information de Fisher (1.13) donnée ci-dessous. Une analyse à θ fixé montrent que ces estimateurs sont très bons en général.

Les estimateurs bayésiens peuvent être calculés en déterminant les lois *a posteriori*. Dans certains cas simples, ces calculs peuvent être réalisés explicitement, et l'estimateur sera le barycentre des θ pondérés par leurs vraisemblances. Dans les situations plus complexes, on utilisera les chaînes de Markov Monte-Carlo.

Critères minimax

Soit qu'il est pessimiste ou mégalomane, le mathématicien part du principe qu'il joue contre le Diable. Aussi, il veut mettre au point une stratégie efficace quel que

[4]Les bayésiens subjectivistes considèrent les lois de probabilité comme des degrés de croyance. Ils peuvent donc utiliser toute loi *a priori* basée sur les informations d'experts.

soit le vrai θ. Un estimateur $\hat{\theta}$ est donc évalué par sa valeur dans le pire des cas :

$$R_M(\hat{\theta}) = \sup_{\theta} r_\theta(\hat{\theta}). \tag{1.9}$$

Le risque minimax est le risque du meilleur estimateur, dit estimateur minimax :

$$R_M = \inf_{\hat{\theta}} R_M(\hat{\theta}) = \inf_{\hat{\theta}} \sup_{\theta} r_\theta(\hat{\theta}). \tag{1.10}$$

Le défaut de cette méthode est qu'elle peut conduire à affaiblir l'estimation sur intuitivement beaucoup de valeurs possibles de θ afin d'être efficace dans quelques cas particuliers. Ce problème est contourné en réclamant d'être adaptatif, c'est-à-dire d'être minimax sur toute une classe de sous-ensembles de $\{p_\theta\}$. Cette dernière technique s'utilise surtout en statistiques non paramétriques.

L'intérêt de ces méthodes est qu'elle ne font aucune hypothèse. Elles donnent une efficacité dont nous savons qu'elle est atteinte, à partir du moment où le modèle (ou l'expérience) lui-même est juste.

Liens entre critères bayésiens et minimax

Le lien principal entre ces deux critères vient de la remarque suivante. Si une stratégie $\hat{\theta}$ est optimale au sens bayésien pour une loi *a priori* quelconque, et si le risque de $\hat{\theta}$ ne dépend pas de θ, alors $\hat{\theta}$ est optimale au sens minimax.

En effet, pour tout π, le risque de Bayes est plus faible que le risque minimax :

$$R_\pi(\hat{\theta}) \leq \sup_{\theta} r_\theta(\hat{\theta}) = R_M(\theta), \tag{1.11}$$

avec égalité si et seulement si le risque au point θ est le même π-presque partout.

Sous certaines conditions, l'énoncé inverse est vrai : un estimateur minimax est optimal pour une loi *a priori* précise, celle pour laquelle le risque bayésien est maximal. Nous discuterons de questions similaires au Chapitre 3.

Exemple

Nous calculons le risque de l'estimateur susmentionné pour la famille de décalage gaussienne (1.4). La loi de $\hat{\theta}$ est la loi des données originales, c'est-à-dire la loi

normale $\mathcal{N}(\theta, \mathcal{I}^{-1})$. Donc

$$\begin{aligned} r_\theta(\hat{\theta}) &= \mathbb{E}_\theta \left[(\theta - \hat{\theta})^\top G (\theta - \hat{\theta}) \right] \\ &= \mathrm{Tr}(G\mathcal{I}^{-1}). \end{aligned} \tag{1.12}$$

Ce risque au point θ ne dépend pas de θ, si bien que cette même valeur est aussi les risques minimax et bayésiens pour toute loi *a priori* de cet estimateur. Nous verrons plus bas que cet estimateur est aussi minimax pour le modèle.

Le reste de cette section résume brièvement les risques que l'on peut attendre dans les cas suffisamment réguliers, pour des fonctions de coût quadratiques.

Information de Fisher

Les risques que nous donnons ci-dessus dépendent de la question (la fonction de coût) et de l'expérience $\{p_\theta, \theta \in \Theta\}$, mais pas d'un estimateur particulier. Nous pouvons donc les lire directement sur l'expérience.

La notion la plus importante à cette fin est celle de matrice d'*information de Fisher*. C'est une notion locale, qui peut être interprétée comme une mesure de la vitesse à laquelle nous pouvons distinguer p_θ des $p_{\theta+\mathrm{d}\theta}$ environnants. La borne de Cramér-Rao décrite dans la prochaine section explicite cette interprétation. Notons que pour ce qui suit, il faut que le modèle soit assez régulier. Deux fois différentiable en θ est plus que suffisant.

L'information de Fisher au point $\theta = (\theta_\alpha)_{\alpha=1\ldots d}$ est donnée partons

$$\mathcal{I}_{\alpha,\beta}(\theta) = \int_\Omega \frac{\partial \ln(p_\theta(X))}{\partial \theta_\alpha} \frac{\partial \ln(p_\theta(X))}{\partial \theta_\beta} \mathrm{d}p_\theta(X). \tag{1.13}$$

La matrice d'information de Fisher est définie positive, et définit une métrique sur Θ, qui est invariante par changement de variables lisse. Ce fait peut être vu comme le lien le plus basique entre statistiques et géométrie différentielle. La géométrie différentielle peut être utilisée pour étudier les asymptotiques d'ordre supérieur, comme par exemple dans le livre d'Amari (1985).

En développant le logarithme des produits, nous constatons facilement qu'avoir n échantillons de données multiplie l'information de Fisher par n, c'est-à-dire $\mathcal{I}^{(n)}(\theta) = n\mathcal{I}^{(1)}(\theta)$ où $\mathcal{I}^{(n)}$ est la matrice d'information de Fisher de l'expérience $\mathcal{E}^{(n)} = \{p_\theta^{\otimes n}, \theta \in \Theta\}$.

Borne de Cramér-Rao

Nous pouvons utiliser la matrice d'information de Fisher pour trouver une borne inférieure sur la matrice de variance des estimateurs localement non biaisés :

$$\int_{\Omega_1} (\theta - \hat{\theta})(\theta - \hat{\theta})^\top \mathrm{d}q_\theta(\hat{\theta}) \geq \mathcal{I}^{-1}(\theta). \tag{1.14}$$

Cette borne tient[5] pour tous les estimateurs localement non biaisés $\hat{\theta}$, c'est-à-dire aussi longtemps que $\int \hat{\theta} \mathrm{d}q_\theta(\hat{\theta}) = \theta$ et $\partial/\partial\theta_i \int \hat{\theta}_j \mathrm{d}q_\theta(\hat{\theta}) = \delta_{i,j}$.

Comme conséquence immédiate, pour une fonction de coût quadratique $(\theta - \hat{\theta})^\top G(\theta - \hat{\theta})$ et tous les estimateurs localement non biaisés, nous obtenons cette borne inférieure sur le risque au point θ :

$$r_\theta(\hat{\theta}) \geq \mathrm{Tr}(G\mathcal{I}^{-1}). \tag{1.15}$$

Cette borne est asymptotiquement saturée. En effet, une expérience de n-échantillon ressemble de plus en plus à une expérience de décalage gaussienne, pour laquelle la borne est saturée. L'explication précise vient de la théorie de la convergence d'expériences de Le Cam, que nous esquissons plus avant à la Section 1.7.1.

Exemples

Calculons l'information de Fisher pour l'expérience de Bernoulli, en un point θ différent de 0 et de 1. L'expression se simplifie légèrement comme nous n'avons qu'un paramètre.

$$\begin{aligned}
\mathcal{I}(\theta) &= \theta \left(\frac{\mathrm{d}\ln(\theta)}{\mathrm{d}\theta} \right)^2 + (1-\theta) \left(\frac{\mathrm{d}\ln(1-\theta)}{\mathrm{d}\theta} \right)^2 \\
&= \frac{1}{\theta} + \frac{1}{1-\theta} \\
&= \frac{1}{\theta(1-\theta)}.
\end{aligned}$$

De ceci et notre remarque précédente sur les n-échantillons, nous déduisons que $\mathcal{I}(\theta) = n/(\theta(1-\theta))$ dans l'expérience binomiale $\mathcal{E}_{bin}$.

Un calcul un peu plus pénible montrerait que la matrice d'information de Fisher d'une expérience de décalage gaussienne est l'inverse de la variance des gaussiennes.

[5] Les estimateurs supereefficaces tel l'estimateur de Stein montrent qu'on ne peut pas simplement éliminer la condition d'être localement non biaisé. Cependant, cette condition peut être supprimée au prix de modifications techniques, consistant essentiellement à considérer l'efficacité sur tout un voisinage de θ, soit dans une approche minimax, soit bayésienne.

D'où notre choix de notation dans l'équation (1.4). De plus, après comparaison entre la borne (1.15) et le risque (1.12) de l'estimateur consistant à prendre X lui-même, nous obtenons l'optimalité de ce dernier estimateur dan la classe des estimateurs localement non biaisés.

Nous allons maintenant nous attacher à donner des équivalents de ces notions dans le monde quantique.

1.1.2 Objets et Opérations Quantiques

Les livres de Helstrom (1976) et Holevo (1982) sont les références habituelles en statistiques quantiques. Nous pouvons également ajouter l'article de revue plus récent de Barndorff-Nielsen et al. (2003). Comme nous l'avons déjà mentionné, nous avons résumé dans le Tableau 1.2, page 27, les ingrédients de base des statistiques quantiques, avec le Tableau classique correspondant 1.1 en regard.

États, opérateurs densité

L'objet de base des statistiques quantiques est l'état. L'état est l'équivalent d'une loi de probabilité.

Nous le définissons sur un espace de Hilbert $\mathcal{H}$. Son expression mathématique est donnée par l'opérateur densité.

Definition 1.1.1. *Un* opérateur densité ρ *sur un espace de Hilbert* $\mathcal{H}$ *est un opérateur à classe de trace doté des propriétés suivantes :*
- *Auto-adjonction :* ρ *est auto-adjoint.*
- *Positivité :* ρ *est positif.*
- *Normalisation :* $\mathrm{Tr}(\rho) = 1$.

Ces conditions sont les équivalentes de celles qui régissent les mesures de probabilité : ces dernières sont réelles (= auto-adjointes), positives et normalisées à un.

Pour les espaces de Hilbert de dimension finie, les opérateurs sont des matrices, et les matrices densité satisfont également aux conditions ci-dessus. La variété des états est de dimension réelle $d^2 - 1$ si $\mathcal{H}$ est de dimension complexe d.

Exemple : Qubits

La situation la plus élémentaire correspond à $\dim(\mathcal{H}) = 2$. Physiquement, ce système pourrait être le spin d'un électron. Ces états sont appelés états qubit, et sont largement utilisés en information quantique.

Nous définissons les *matrices de Pauli* comme

$$\sigma_x = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \qquad \sigma_y = \begin{bmatrix} 0 & i \\ -i & 0 \end{bmatrix}, \qquad \sigma_z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \tag{1.16}$$

Comme une matrice densité est auto-adjointe, elle sera une combinaison linéaire réelle de ces trois matrices et de l'identité $\mathbf{1}$. La positivité et la normalisation imposent de plus :

$$\rho = \frac{1}{2}\left(\mathbf{1} + \vec{\theta} \cdot \vec{\sigma}\right), \qquad \|\vec{\theta}\| \leq 1, \tag{1.17}$$

avec $\vec{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$ un vecteur de matrices.

Nous voyons que nous avons déjà besoin de trois paramètres réels pour décrire les états qubit, *confer* le paramètre unique dont nous avons besoin pour décrire une loi sur un espace classique à deux éléments.

États purs

L'ensemble des mesures de probabilité peut être vue comme l'enveloppe convexe des fonctions delta. De même, les états sont l'enveloppe convexe des états purs.

Les états purs sont caractérisés par le fait d'être des opérateurs de rang un, de valeur propre un. Nous pouvons les écrire $|\psi\rangle\langle\psi|$, où $|\psi\rangle$ est un vecteur de norme un dans $\mathcal{H}$. Les états purs peuvent donc être vus comme des points de l'espace projectif associé à $\mathcal{H}$.

Ils sont extrêmement importants : de nombreuses descriptions de la mécanique quantique traitent uniquement les états purs. Les états généraux sont des mélanges classiques d'états purs. Un état qui n'est pas pur est dit *mélangé.*

Contrairement aux fonctions delta, où il suffit de tirer une fois la variable aléatoire pour identifier la loi inconnue, il n'existe pas de mesure permettant d'identifier sans ambiguïté n'importe quel état pur, quand bien même nous saurions auparavant que l'état est pur. Cette différence fondamentale avec le cas classique est une marque de la non-commutatiivté entre les différents états. L'étude des états purs est déjà un problème difficile.

Pour les qubits paramétrés comme ci-dessus, les états purs correspondent à $\|\vec{\theta}\| = 1$ Cette paramétrisation par une sphère, appelée *sphère de Bloch*, nous donne une intuition graphique pour les problèmes sur les qubits.

La dimension réelle des états purs est de $2(d-1)$ si $\dim(\mathcal{H}) = d$.

États cohérents

Les qubits sont l'exemple-type des états de dimension finie. Les états cohérents[6] forment l'autre famille fondamentale d'états.

Ces états vivent dans l'espace de Fock[7] $\mathcal{F}(\mathcal{C})$, c'est-à-dire l'espace de Hilbert de dimension infinie $\ell^2(\mathbb{N})$. Nous notons par $\{|k\rangle\}_{k\in\mathbb{N}}$ la base canonique de $\ell^2(\mathbb{N})$. Les physiciens appellent $|k\rangle$ le k-ième état de Fock.

Les états sur l'espace de Fock sont ceux de l'oscillateur harmonique, comme par exemple l'état de la lumière monochromatique, *i.e.* l'état d'un laser. Nos sommes donc sur le terrain de l'optique quantique. Parmi ces états, les états cohérents sont en un sens les plus classiques : ils saturent les relations d'incertitude de Heisenberg.

Ils sont donnés par un coefficient θ complexe, soit deux paramètres réels. Comme ce sont des états purs, nous pouvons les décrire par un vecteur de $\mathcal{F}(\mathbb{C})$, plutôt que par un opérateur[8] :

$$|\theta) = \exp(-|\theta|^2/2) \sum_{k=0}^{n} \frac{\theta^k}{\sqrt{k!}} |k\rangle \,. \tag{1.18}$$

États multipartites, états intriqués

Considérons deux objets quantiques ρ_1 et ρ_2 sur $\mathcal{H}_1$ et $\mathcal{H}_2$. Ils peuvent être vus comme un seul objet quantique sur l'espace $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2$, d'état $\rho = \rho_1 \otimes \rho_2$.

Tout état sur pareil espace de Hilbert produit est appelé *état multipartite*. Maintenant certains états multipartites ne peuvent pas être écrits comme une combinaison convexe $\sum c_i \rho_1^i \otimes \rho_2^i$, avec des c_i positifs. Nous pourrions avoir besoin de c_i

[6] Plus généralement, tous les états gaussiens éventuellement compressés jouent un rôle majeur en optique quantique et, comme nous allons le voir, en statistiques quantiques. Dans l'exemple, nous nous restreignons aux états cohérents par souci de simplicité.

[7] Les états cohérents de dimension supérieure à deux sont des produits tensoriels d'états cohérents sur l'espace de Fock produit $\mathcal{F}(\mathbb{C}^d) = \mathcal{F}(\mathbb{C})^{\otimes d}$.

[8] Nous utiliserons la notation $|\theta)$ au lieu du ket habituel $|\theta\rangle$ afin d'éviter la confusion avec les états de Fock, en particulier quand θ est un entier positif.

strictement négatifs. En d'autres termes, ces états ne sont pas un mélange statistique classique de paires d'états. Ils contiennent un couplage intrinsèquement quantique. De tels états sont appelés *états intriqués.*

Commençons par prouver leur existence. Nous écrivons $\dim \mathcal{H}_1 = d_1$ et $\dim \mathcal{H}_2 = d_2$. Donc $\dim \mathcal{H} = d_1 d_2$. Les états multipartites pur sont des états purs sur $\mathcal{H}$, donc constituent une variété de dimension $2(d_1 d_2 - 1)$. D'un autre côté, un état pur de la forme $\sum c_i \rho_1^i \otimes \rho_2^i$ avec les c_i positifs impose que la somme ne contienne qu'un seul terme, avec ρ_1 et ρ_2 tous deux purs. La dimension de la variété de ces états produit est $2(d_1 + d_2 - 2) < 2(d_1 d_2 - 1)$. Il y a donc de nombreux états purs intriqués.

Un exemple typique sont les *états d'intrication maximale*, c'est-à-dire les états de la forme $|\Psi\rangle\langle\Psi|$, avec $|\Psi\rangle = \frac{1}{\sqrt{d}} \sum |\psi^i\rangle \otimes |\psi^i\rangle$, où $\mathcal{H}_1 = \mathcal{H}_2$ et $\{|\psi^i\rangle\}$ est une base orthonormale de $\mathcal{H}_1$. Comme leur nom l'indique, ces états sont aussi intriqués que possible.

L'intrication est peut-être la ressource la plus basique et la plus essentielle de toute l'information quantique. Elle joue un rôle au cœur de la téléportation quantique, de la plupart des protocoles de cryptographie quantique et dans les algorithmes accélérés des ordinateurs quantiques. La littérature qui y est consacrée est trop immense pour être seulement esquissée. En statistiques quantiques, les états intriqués peuvent être utilisés pour accélérer l'estimation de transformations quantiques

Actions sur les états

Dans le cas classique, nous avons remarqué que donner un estimateur ed θ, ou plus généralement de n'importe quelle fonction de θ, était équivalent à la transformation de nos données initiales pour obtenir une nouvelle variable aléatoire Y de loi $T(p_\theta)$.

Dans le cas quantique, les deux notions sont bien distinctes. En effet, transformer les données signifie obtenir un nouvel état quantique, c'est-à-dire un nouvel opérateur sur un espace de Hilbert. Les états sont transformés quand ils sont envoyés à travers un *canal.* Un estimateur d'un paramètre classique, en revanche, est une quantité classique. Nous obtenons donc une variable aléatoire classique. Ces données classiques sont obtenues en effectuant une *mesure* de l'état.

Si nous souhaitons simplement considérer les estimateurs, pourquoi s'intéresser aux canaux ? En effet, l'application de canaux successifs suivie d'une mesure peut être résumée à une mesure plus complexe.

La première raison est que nous pourrions vouloir transformer nos états en une nouvelle famille pour laquelle nous savons quelle mesure effectuer. En fait, tout le but

de la normalité asymptotique locale quantique forte, dont l'étude forme l'essentiel de cette thèse, est de transformer des expériences en d'autres expériences quasi-équivalentes, et plus simples et mieux connues.

Deuxièmement, les canaux décrivent des transformations quantiques. Nous pourrions souhaiter étudier la transformation elle-même, plutôt que l'état. Typiquement, cette transformation pourrait être générée par une force que nous voulons mesurer. Nous nous étendrons davantage sur le sujet au Chapitre 4.

Nous appelons *instrument* une fonction qui retourne *à la fois* des données classiques et quantiques en prenant un état quantique en entrée. Les véritables instruments de mesure sont en fait des instruments, quand bien même l'état de sortie peut être oublié. En particulier, les mesures en temps continu sont communes en pratique. Typiquement, nous mesurons le champ électromagnétique par son interaction avec la matière, comme au Chapitre 8. Ces mesures peuvent être vues comme une suite d'instruments infinitésimaux. Écrire les équations correspondantes est le but du filtrage quantique, créé par Davies et Belavkin (Bouten et al., 2006, for an introduction).

Mesures, POVMs

Si nous voulons effectuer de l'inférence statistique classique sur les paramètres inconnus, il nous faut traduire notre information quantique en information classique. À cette fin, nous effectuons une mesure. Comme les états mélangés sont des mélanges classiques d'états purs, nous exigeons que cette transformation soit linéaire. De plus, la sortie doit toujours suivre une loi de probabilité classique. De ceci, nous déduisons la forme suivante pour les mesures physiquement permises :

Definition 1.1.2. *Une* mesure à valeur dans les opérateurs positifs, *ou POVM, de l'acronyme anglais, sur un espace mesuré* $(\Omega, \mathcal{A})$ *est une ensemble* $\{M(A)\}_{A\in\mathcal{A}}$ *d'opérateurs bornés sur* $\mathcal{H}$ *tels que :*
- $M(\Omega) = \mathbf{1}_{\mathcal{H}}$.
- $M(A)$ *est positif.*
- *Pour toute collection dénombrable* $(A_i)_{i\in\mathbb{N}}$ *d'*A_i *disjoints, nous avons* $M(\bigcup A_i) = \sum M(A_i)$.

On remarquera que ce sont exactement les axiomes habituels d'une mesure de probabilité, à ceci près que nous utilisons des opérateurs au lieu de nombres réels. Nous appelons chaque $M(A)$ un *élément de POVM.*

Appliquer une mesure M à un état ρ génère une loi P_ρ sur $(\Omega, \mathcal{A})$, donnée par la *règle de Born* :

$$P_\rho(A) = \mathrm{Tr}(\rho M(A)). \tag{1.19}$$

Au Chapitre 5, nous examinerons une relation d'ordre spécifique sur les POVMs.

Quelques remarques s'imposent. Tout d'abord, nous pouvons inclure tout traitement classique des données dans la POVM. En effet, effectuer la mesure M, puis appliquer le noyau de Markov T (défini par (1.5)) à la variable aléatoire de sortie est équivalent à effectuer la mesure N sur $(\Omega_1, \mathcal{B})$ donnée par $N(B) = \int_\Omega p_\omega(B) M(\mathrm{d}\omega)$. Si bien que travailler avec des POVMs est équivalent à travailler avec des estimateurs.

Deuxièmement, en général, nous ne pouvons pas mesurer simultanément M_1 et M_2 sur $(\Omega_1, \mathcal{A}_1)$ et $\Omega_2, \mathcal{A}_2)$. Contrairement au cas classique, où l'on peut obtenir simultanément les résultats de l'application des noyaux T_1 et T_2. En effet, mesurer à la fois M_1 et M_2 signifie mesurer N sur $(\Omega_1 \otimes \Omega_2)$ avec $N(A_1 \times \Omega_2) = M_1(A_1)$ et $N(\Omega_1 \times A_2) = M_2(A_2)$. Un contre-exemple simple illustrant le rôle de la non-commutativité est donné par M_1 et M_2 toutes deux définies sur $\{0,1\}$, avec

$$M_1(0) = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \qquad M_1(1) = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix},$$
$$M_2(0) = \frac{1}{2}\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}, \qquad M_2(1) = \frac{1}{2}\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}.$$

Toutes ces matrices sont de rang un. Il faut maintenant $N(0,0) + N(0,1) = M_1(0)$. Comme tous les éléments de POVM sont positifs, nous obtenons $M_1(0) \geq N(0,0)$. Comme de plus $M_1(0)$ est de rang un, nous avons $N(0,0) = c_1 M_1(0)$ pour un certain $0 \leq c_1 \leq 1$. De même $N(0,0) + N(1,0) = M_2(0)$. Si bien que $N(0,0) = c_2 M_2(0)$. La seule solution est $c_1 = c_2 = 0$ et $N(0,0) = 0$. Le même raisonnement tient pour $N(0,1)$, $N(1,0)$ et $N(1,1)$. Par ailleurs, il faut que $N(\{0,1\}^2) = \mathbf{1}_{\mathbb{C}^2}$. Contradiction.

Finalement, on croit que toutes ces mesures sont permises par les lois de la physique. Mais elles peuvent être très dures à implémenter en pratique. En particulier si l'état est multipartite, il peut être raisonnable de se restreindre à des classes de mesure plus petites. Notamment, si différentes personnes possèdent différentes particules en des lieux différents, elles ne pourront pas implémenter une mesure générale, même s'ils coopèrent. Le mieux qu'elles puissent faire est : l'une d'elles mesure sa particule (éventuellement avec un état quantique non trivial en sortie), donne le résultat aux autres, qui choisissent quel mesure effectuer sur leurs particules, garde l'état de sortie et donnent le résultat aux autres, et on itère. De telles mesures, qui utilisent uniquement des opérations quantiques locales et les communications classiques, sont appelées LOCC : *Local Operations, Classical Communication.*

En information quantique, quand le système (souvent intriqué) est divisé entre plusieurs personnes, nous nous restreignons naturellement aux opérations LOCC. En estimation quantique avec n copies de l'état initial, nous sommes intéressés par ce que nous pouvons réaliser avec des mesures LOCC, beaucoup plus simples à implémenter en pratique, que les mesures générales, dites collectives. Nous pouvons

généralement améliorer la précision de la mesure par des mesures collectives, ce qui peut paraître surprenant pour des physiciens, puisque les n copies sont totalement indépendantes. Dans certains cas, en particulier quand nous savons que l'état est pur (Matsumoto, 2002), les mesures collectives n'améliorent guère les mesures LOCC. Cela peut surprendre les mathématiciens, comme l'espace des mesures collectives est beaucoup plus grand que celui des mesures LOCC.

Exemple : Spin z

Considérons la mesure binaire sur les qubits donnée par

$$M(\uparrow) = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} = \frac{1}{2}(\mathbf{1} + \sigma_z), \qquad M(\downarrow) = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} = \frac{1}{2}(\mathbf{1} - \sigma_z).$$

Cette mesure appliquée à l'état $\rho = \frac{\mathbf{1}+\vec{\theta}\cdot\vec{\sigma}}{2}$ renvoie $\uparrow$ avec probabilité

$$\operatorname{Tr}(\rho M(\uparrow)) = \frac{1}{2}\Big(\operatorname{Tr}(\mathbf{1}M(\uparrow)) + \sum_{\alpha=x,y,z} \theta_\alpha \operatorname{Tr}(\sigma_\alpha M(\uparrow))\Big) = \frac{1}{2}(1 + \theta_z).$$

En particulier, si $\theta_z = 1$, la sortie est toujours $\uparrow$. À l'inverse, si $\theta_z = -1$, la sortie est toujours $\downarrow$. D'autre part, si $\theta_x = 1$, si bien que $\theta_z = 0$, la sortie sera ou bien $\uparrow$ ou bien $\downarrow$ avec probabilité un demi, alors que l'état ρ est pur.

Les mesures de ce genre, où tous les éléments de POVM sont des projecteurs, sont aussi appelées *observables*. Elles génèrent de l'information uniquement sur la base où tous les éléments de POVM sont diagonaux. Accessoirement, les axiomes usuels de la mécaniques quantiques restreignent les mesures aux observables. Nos récupérons l'ensemble des POVMs en appliquant une observable à un état multipartite dont notre état n'est qu'une partie ; c'est le théorème de Naimark.

Mesure hétérodyne

La mesure hétérodyne tire son nom de la technique utilisée pour l'implémenter en laboratoire, avec des lasers déphasés. Cette POVM de sortie dans $\mathbb{C}$ se décrit mathématiquement par :

$$M(A) = \frac{1}{\pi}\int_A |z)(z|dz, \tag{1.20}$$

où $|z)$ est un état cohérent (1.18).

La loi du résultat de la mesure de ρ a donc pour densité $(z|\rho|z)$ par rapport à la mesure de Lebesgue, au point z. En particulier, la loi du résultat de la mesure d'un état cohérent est une gaussienne :

$$q_\theta(\mathrm{d}z) = \frac{1}{\pi}(z|\theta)(\theta|z) = \frac{1}{\pi}\exp(-|\theta - z|^2). \tag{1.21}$$

Si nous considérons tous les θ complexes, nous reconnaissons une expérience de décalage gaussienne (1.4) sur $\mathbb{R}^2$.

Plus généralement, la densité de la loi du résultat de la mesure d'un état ρ est appelée fonction de *Husimi* de l'état :

$$H_\rho(\mathrm{d}z) = \frac{1}{\pi}(z|\rho|z). \tag{1.22}$$

Les états dont la fonction de Husimi est une gaussienne sont appelés *états gaussiens.*

Canaux

Nous décrivons maintenant comment obtenir un nouvel état quantique à partir d'un premier. Notez que l'état original est détruit au cours du processus.

Une transformation physique d'un objet prend un état et renvoie un aure état, éventuellement sur un espace de Hilbert diffèrent. Elle est décrite par un canal, l'équivalent d'un noyau de Markov.

Pour rappel, un superopérateur positif $\mathcal{E}$ est une application qui associe à chaque opérateur A positif un résultat $\mathcal{E}(A)$ également positif.

Definition 1.1.3. *Un* canal $\mathcal{E}$ *est une application de l'ensemble* $\mathcal{T}(\mathcal{H}_1)$ *des opérateurs à classe de trace, dans* $\mathcal{T}(\mathcal{H}_2)$*, doté des propriétés suivantes :*
- *Linéarité :* $\mathcal{E}$ *est linéaire.*
- *Positivité complète : pour tout espace auxiliaire* $\mathcal{H}_3$*, le superopérateur* $\mathcal{E} \otimes Id$ *:* $\mathcal{T}(\mathcal{H}_1 \otimes \mathcal{H}_3) \to \mathcal{T}(\mathcal{H}_2 \otimes \mathcal{H}_3)$ *donné par* $(\mathcal{E} \otimes Id)(\rho \otimes \sigma) = \mathcal{E}(\rho) \otimes \sigma$ *est positif.*
- *Préservation de la trace :* $\mathrm{Tr}(\mathcal{E}(A)) = \mathrm{Tr}(A)$.

Notez que les noyaux de Markov satisfont à ces critères, quand on remplace les opérateurs par des mesures[9].

[9]Dans le cadre plus général des C^*-algèbres, les espaces de fonctions sont des C^*-algèbres commutatives et tous les superopérateurs positifs sur ces espaces sont complètement positifs

La nécessité de la linéarité peut être prouvée par l'axiome d'évolution unitaire[10] et en incluant l'observateur dans le système.

Nous voulons que l'image d'un état soit un état, donc un opérateur positif doit être envoyé sur un opérateur positif. Pour comprendre d'où vient l'exigence de positivité complète, considérons un état éventuellement intriqué sur $\mathcal{H}_1 \otimes \mathcal{H}_3$. Si nous transformons l'état sur $\mathcal{H}_1$, nous transformons aussi l'état sur $\mathcal{H}_1 \otimes \mathcal{H}_3$, par le canal $\mathcal{E} \otimes Id$. Donc cette dernière transformation doit aussi être positive. D'où la requête de complète positivité.

Finalement, la sortie est un état si l'entrée est un état, et tous deux ont trace un, donc la trace doit être conservée.

Nous considérerons souvent des canaux du point de vue (pré)dual, c'est-à-dire comme agissant sur les éléments de $\mathcal{B}(\mathcal{H})$. Nous définissons $\mathrm{Tr}(\mathcal{E}(\rho)A) = \mathrm{Tr}(\rho\mathcal{E}_*(A))$ pour tout état ρ et tout opérateur borné A. Dans ce cas, $\mathcal{E}_*$ est aussi une application linéaire complètement positive, mais nous devons remplacer la préservation de la trace par la *préservation de l'identité*, c'est-à-dire $\mathcal{E}_*(\mathbf{1}) = \mathbf{1}$.

Notations : Nous utilisons d'habitude les lettres $\mathcal{E}$ ou $\mathcal{F}$ pour les canaux. Par abus de notation, nous ne noterons en général pas l'étoile pour le prédual et écrirons également $\mathcal{E}$ dans ce cas. Cependant, ces notations standard sont également les notations standard pour les expériences ; Aussi , dans les chapitres où nous utilisons cette dernière notion, nous désignerons les canaux de la même façon que les noyaux de Markov, à savoir par T, T_n, S, S_n.

Représentation de Kraus, théorème de Stinespring

La définition donnée ci-dessus ne permet pas une manipulation simple des canaux. Heureusement, nous disposons de deux théorèmes de représentation qui décrivent les applications complètement positives de manière plus pratique. Le livre de Paulsen (1987) est une bonne référence sur ces sujets.

La représentation de Kraus (1983) est l'outil principal quand l'espace de Hilbert est de dimension finie.

Theorem 1.1.4. *Une application complètement positive $\mathcal{E}$ de $M(\mathbb{C}^{d_1})$ dans $M(\mathbb{C}^{d_2})$ peut s'écrire*

$$\mathcal{E}(A) = \sum_\alpha R_\alpha A R_\alpha^*, \tag{1.23}$$

[10]La mécanique quantique affirme que l'évolution d'un système est donnée par $\rho(t) = U(t)\rho(0)U^*(t)$, où $U(t)$ est un opérateur unitaire qui peut être calculé à partir de l'opérateur auto-adjoint $\mathcal{H}$ appelé le hamiltonien. Si le hamiltonien ne dépend pas du temps, alors $U(t) = e^{itH}$.

où α va de 1 à $d_1 d_2$ au plus, et $R_\alpha \in M_{d_2,d_1}(\mathbb{C})$. L'étoile représente l'adjonction.

De plus, le canal préserve la trace si et seulement si $\sum R_\alpha^* R_\alpha = \mathbf{1}_{\mathbb{C}^{d_1}}$.

Cette décomposition n'est pas unique. Le canal dual est donné par $A \mapsto \sum R_\alpha^* A R_\alpha$.

En dimension infinie, nous utiliserons de préférence le plus puissant théorème de dilatation[11] de Stinespring (1955).

Theorem 1.1.5. *Soit $\mathcal{E} : \mathcal{B}(\mathcal{H}_1) \to \mathcal{B}(\mathcal{H}_2)$ une application complètement positive. Alors il existe un espace de Hilbert $\mathcal{K}$ et un *-homomorphisme (ou représentation) $\pi : \mathcal{B}(\mathcal{H}_1) \to \mathcal{B}(\mathcal{H}_2)$ tels que*

$$\mathcal{E}(A) = V\pi(A)V^*, \tag{1.24}$$

où $V : \mathcal{K} \to \mathcal{H}$ est un opérateur borné.

De plus, si $\mathcal{E}$ préserve l'identité, alors V est une isométrie, c'est-à-dire $VV^* = \mathbf{1}_{\mathcal{H}}$.

Si de plus nous imposons que $\mathcal{K}$ soit la fermeture de l'espace vectoriel engendré par $\pi(A)V^*\mathcal{H}$, alors la dilatation est unique à des transformations unitaires près.

Instruments

Nous donnons les représentations d'instruments en dimension finie[12]. Pour simplifier davantage les notations, nous nous placerons dans le cas où la mesure a un nombre fini d'issues.

Definition 1.1.6. *Un instrument est donné par un ensemble $\{N_{\omega,k}\}$ de matrices de $\mathcal{H}_1$ dans $\mathcal{H}_2$, tel que*

$$\sum_\omega \sum_k N_{\omega,k}^* N_{\omega,k} = \mathbf{1}_{\mathcal{H}_1}.$$

La mesure correspondante est donnée par

$$M(\omega) = \sum_k N_{\omega,k}^* N_{\omega,k},$$

[11] En fait le théorème de Stinespring a été prouvé pour toute C^*-algèbre unitaire. On peut prouver qu'il implique le théorème de représentation de Kraus, ainsi que la représentation GNS, une base de la théorie des algèbres d'opérateurs.

[12] En dimension infinie, il faut se placer dans le cadre des C^*-algèbres, et un instrument est alors simplement un canal entre C^*-algèbres

et l'état de sortie quand le résultat de la mesure est ω est donné par

$$\mathcal{N}(\rho, \omega) = \frac{\sum_k N_{\omega,k} \rho N_{\omega,k}^*}{\mathrm{Tr}(\rho M(\omega))}.$$

L'état de sortie vit dans $\mathcal{H}_2$.

Nous avons désormais une nouvelle manière de comprendre pourquoi nous ne pouvons pas mesurer deux POVMs simultanément : après avoir mesuré M, l'objet quantique, donc notre donnée, a en général été perturbé. En fait, si la mesure est suffisamment riche, l'état de sortie ne dépend que du résultat ω de la mesure, et plus du tout de l'état d'entrée.

Nous avons désormais tous les outils pour transposer les statistiques classiques au monde quantique.

1.1.3 Statistiques quantiques

Nous travaillons d'habitude sur les états quantiques ; à l'occasion, nous voudrons obtenir des informations sur un canal. Nous traitons les deux cas séparément.

États : ce que nous avons, ce que nous pouvons faire, ce que nous voulons savoir

De manière analogue au cas classique, nous disposons d'habitude d'un état quantique ρ, que nous savons être dans l'ensemble

$$\mathcal{E} = \{\rho_\theta, \theta \in \Theta\}. \tag{1.25}$$

Nous appellerons également cet ensemble *expérience* ou *modèle.*

Dans l'exemple des qubits, les modèles usuels seront le modèle complet de mélange, à trois dimensions, $\mathcal{E}_m = \{\rho_\theta, \|\theta\| < 1\}$ et le modèle à deux dimensions des états purs $\mathcal{E}_p = \{\rho_\theta, \|\theta\| = 1\}$, où nous avons utilisé la paramétrisation précédente (1.17) de l'état ρ_θ. Si nous avons n copies de l'état, nous remplaçons ρ_θ par $\rho_\theta^{\otimes n}$.

Un autre exemple typique est le modèle $\mathcal{E}_t = \{\rho_\theta, \theta \in \{\theta_1, \theta_2\}\}$, où la question habituelle est de discriminer entre les deux θ possibles. Nous étudions ce genre de problèmes dans la Section 1.3 et le Chapitre 3.

Nous pouvons *a priori* utiliser n'importe quelle suite d'instruments sur l'état. Si nous voulons simplement obtenir des renseignements sur θ, nous pouvons nous restreindre aux mesures M, les POVMs. Nous associons alors à M un estimateur, disons $\hat{\theta}$, dont la loi dépend du vrai paramètre θ de la manière suivante :

$$q_\theta(B) \hat{=} \mathbb{P}_\theta\left[\hat{\theta} \in B\right] = \mathrm{Tr}(\rho_\theta M(B)).$$

Selon les circonstances, nous pourrons permettre toutes les mesures physiques, ou nous restreindre à des ensembles plus petits, comme les mesures séparées ou les mesures LOCC.

Enfin, ce que nous voulons savoir est la même chose que dans le cas classique. Nous voulons connaître une fonction du paramètre θ. Nous voulons donc estimer θ, et évaluer notre estimateur $\hat{\theta}$ au travers d'une fonction de coût $c(\theta, \hat{\theta})$. Comme auparavant, les fonctions de coût les plus communes sont $(1 - \delta_{\theta,\hat{\theta}})$„ si l'ensemble de paramètres est fini, et les fonctions de coût quadratique $(\hat{\theta} - \theta)^\top G(\hat{\theta} - \theta)$ pour une matrice G définie positive, si le paramètre vit dans un sous-ensemble ouvert de $\mathbb{R}^d$. La matrice de poids G peut dépendre de θ.

Nous pouvons à nouveau écrire le risque (1.6) d'un estimateur au point θ. Comme nous ne connaissons pas θ, nous pouvons utiliser soit le risque bayésien (1.7) pour une loi *a priori* adaptée, soit le risque minimax (1.9), et optimiser (1.8, 1.10) sur les estimateurs disponibles. Notons que la dernière étape dépend de l'ensemble d'estimateurs que nous nous autorisons.

Information de Fisher quantique et bornes de Cramér-Rao

Nous pouvons essayer d'imiter la définition de l'information de Fisher classique et obtenir des ornes similaires sur la variance des estimateurs. En fait, nous pouvons construire pareil équivalent pour tout choix de dérivée logarithmique. Nous choisissons la dérivée logarithmique à droite (RLD), définie pour chaque θ et chaque coordonnée θ_α comme la matrice $\lambda_{\alpha,\theta}$ telle que :

$$\frac{\partial \rho_\theta}{\partial \theta_\alpha} = \rho_\theta \lambda_{\alpha,\theta} \tag{1.26}$$

sur le support de ρ_θ.

Alors l'examen de la définition (1.13) et le rappel du fait que la règle de Born (1.19) est l'équivalent de l'espérance classique rendent naturelle la définition suivante de la matrice d'information de Fisher quantique :

$$\mathcal{J}_{\alpha,\beta}(\theta) \hat{=} \mathrm{Tr}(\rho_\theta \lambda_{\beta,\theta} \lambda^*_{\alpha,\theta}). \tag{1.27}$$

Helstrom (1976) a prouvé que la matrice de covariance de tout estimateur localement non biaisé $\hat{\theta}$ était plus grande que l'inverse de la matrice d'information de Fisher quantique. De ce fait, pour toute fonction de coût quadratique $(\theta - \hat{\theta})^\top G(\theta - \hat{\theta})$, nous avons cette borne sur le risque (1.6) :

$$r_\theta(\hat{\theta}) \geq \mathrm{Tr}\left(\mathrm{Re}(G^{1/2}\mathcal{J}^{-1}(\theta)G^{1/2}) + \left|\mathrm{Im}(G^{1/2}\mathcal{J}^{-1}(\theta)G^{1/2})\right|\right). \tag{1.28}$$

Remarquons que nous n'écrivons pas le membre de droite $\mathrm{Tr}(G\mathcal{J}^{-1}(\theta))$, car notre matrice d'information de Fisher est auto-adjointe, mais pas réelle.

Holevo (1982) a amélioré[13] cette borne pour un paramètre de dimension p et un système vivant dans un espace de Hilbert de dimension d :

$$r_\theta(\hat{\theta}) \geq \inf_{\vec{X}} \mathrm{Tr}\left(\mathrm{Re}(G^{1/2}Z(\vec{X})G^{1/2}) + \left|\mathrm{Im}(G^{1/2}Z(\vec{X})G^{1/2})\right|\right), \tag{1.29}$$

où $Z_{i,j} = \mathrm{Tr}(\rho_\theta X_i X_j)$, et $\vec{X} = (X_1, \ldots, X_p)$ est un vecteur de $d \times d$ matrices auto-adjointes satisfaisant la contrainte $\partial/\partial\theta_i(\mathrm{Tr}(\rho X_j)) = \delta_{i,j}$.

Cette borne s'applique pour tous les estimateurs localement non biaisés. Hayashi et Matsumoto (2004) ont prouvé que cette borne était asymptotiquement saturée pour les modèles de qubits. Comme dans le cas classique, la raison sous-jacente est la convergence vers une expérience de décalage gaussienne quantique. Dans la Partie II, nous construisons une théorie qui montre que toute fonction raisonnable d'un modèle de qubits converge vers sa valeur dans une expérience de décalage gaussienne quantique.

Cette borne a l'air horrible, mais elle est souvent calculable. Par exemple, si le paramètre θ est de dimension $d(d-1)$, il n'y a qu'un seul $\vec{X}$ admissible. C'est le cas quand notre expérience est le modèle complet de mélange. De plus, on peut prouver que cette borne est multipliée par n quand nous avons n échantillons. Nous retrouvons donc la vitesse de convergence en racine carrée des modèles classiques réguliers.

Ces bornes sont valides pour toutes les mesures permises par la physique. Si nous nous restreignons à des classes plus petites, nous pouvons obtenir de meilleures bornes (Nagaoka, 1991; Hayashi, 2005a; Gill et Massar, 2000).

[13] La matrice d'information de Fisher (1.27) est un $Z(\vec{X})$ convenable, ce qui implique à la fois l'existence du membre de droite de l'équation (1.29), et que cette borne est meilleure que la borne de Helstrom (1.28).

Exemple : Expérience de décalage cohérente

Considérons l'expérience quantique suivante sur l'espace de Fock :

$$\mathcal{E}_{qgs} = \{|\theta)(\theta|, \theta \in \mathbb{C}\}.$$

Alors Yuen et Lax, M. (1973) et Holevo (1982)[14] ont calculé la borne de Cramér-Rao (1.28) et obtenu $\mathrm{Tr}(G)/2 + \sqrt{\det(G)}$. Si $G = \mathbf{1}$, cela vaut 2.

Par une mesure hétérodyne (1.20), nous transformons notre expérience quantique en une expérience de décalage gaussienne classique $\mathcal{E}_{gs} = \{\mathcal{N}(\theta, 2 \cdot \mathbf{1}), \theta \in \mathbb{C}\}$. Donc, avec $G = \mathbf{1}$, nous lisons sur notre calcul pour le cas classique (1.12) que le risque au point θ vaut 2.

De ce fait, la mesure hétérodyne sature la borne de Cramér-Rao pour la matrice de poids identité. De légères modifications de cette mesure, faisant usage d'états dit cohérents compressés au lieu des états cohérents (1.18), atteignent l'optimalité pour toute matrice de poids. Il faut remarquer, cependant, que contrairement au cas quantique, la mesure optimale dépend de la matrice de poids.

Exemple 2 : Modèle complet de mélange pour les qubits

Dans le modèle complet de mélange pour les qubits $\mathcal{E}_m$, la borne de Cramér-Rao[15] pour la fonction de coût $(\theta - \hat{\theta})^T(\theta - \hat{\theta})$ est connue pour valoir $3 - 2\,\|\theta\|$.

D'autre part, nous savons aussi (Hayashi et Matsumoto, 2004, pour cette forme précise) que, quand seules les mesures locales sont permises, cette borne devient $(2\sqrt{1 - \|\theta\|})^2$. Nous avons ici un exemple où l'utilisation des mesures collectives améliore la vitesse d'approche, pour tout $\|\theta\| \leq 1$, c'est-à-dire pour tous les états mélangés.

Canaux : Ce que nous avons, ce que nous pouvons faire

Nous avons mis au point notre cadre pour quand des états quantiques nous sont donnés. Dans d'autres applications, nous voulons étudier des machines qui transforment les états quantiques. En statistiques classiques, ce problème correspond à essayer de comprendre ce que fait une boîte noire. Mathématiquement ces machines sont des

[14] Pour une matrice de poids G arbitraire.

[15] Hayashi et Matsumoto (2004) l'ont calculée pour une matrice de poids arbitraire, et montré qu'elle pouvait être atteinte dans tous les cas.

canaux quantiques. Ballester (2005a) a notamment consacré sa thèse à l'estimation de canaux unitaires, correspondant à l'évolution naturelle d'un système quantique. Ji et al. (2006) nous fournit une autre ressource récente.

Dans ce cas, on ne nous donne pas une «loi de probabilité quantique» ρ, mais plutôt un canal $T : \mathcal{B}(\mathcal{H}_1) \to \mathcal{B}(\mathcal{H}_2)$ dans un ensemble

$$\mathcal{E} = \{T_\theta, \theta \in \Theta\}.$$

Pour obtenir des informations sur T, nous devons envoyer un état au travers, et nous obtenons une expérience quantique plus habituelle. Néanmoins, nous avons plusieurs méthodes à disposition. La plus évidente est d'envoyer un état bien choisi ρ. Nous obtenons en sortie l'état $T(\rho)$ et nous retrouvons avec le modèle

$$\mathcal{E}_\rho^1 = \{T_\theta(\rho), \theta \in \Theta\}.$$

Cependant, nous pouvons aussi utiliser un système auxiliaire : au lieu de sonder T, nous sondons de manière équivalente $T \otimes Id : \mathcal{B}(\mathcal{H}_1 \otimes \mathcal{H}_3) \to \mathcal{B}(\mathcal{H}_2 \otimes \mathcal{H}_3)$. Nous envoyons un état ρ multipartite, intriqué et obtenons :

$$\mathcal{E}_\rho^2 = \{(T_\theta \otimes Id)(\rho), \theta \in \Theta\}.$$

Si nous avons le droit de sonder plusieurs fois le canal, le premier réflexe sera d'envoyer n copies du même état. Nous obtenons :

$$\mathcal{E}_\rho^3 = \left\{(T_\theta(\rho))^{\otimes n}, \theta \in \Theta\right\}.$$

Cependant il pourrait être plus efficace d'envoyer un grand état intriqué $\rho \in \mathcal{B}(\mathcal{H}_1)^{\otimes n}$. Nous obtenons alorss l'expérience très générale :

$$\mathcal{E}_\rho^4 = \left\{(T_\theta)^{\otimes n}(\rho), \theta \in \Theta\right\}.$$

Pour finir, on peut décider d'ajouter un système auxiliaire à l'entrée précédente :

$$\mathcal{E}_\rho^5 = \left\{((T_\theta)^{\otimes n} \otimes Id)(\rho), \theta \in \Theta\right\}.$$

Toutes ces distinctions ne sont pas superflues[16]. La première stratégie est plus simple que la seconde, mais Fujiwara (2001) a prouvé qu'envoyer la moitié d'un état d'intrication maximale au travers d'un canal qubit inconnu et garder l'autre moitié

[16]Des stratégies encore plus complexes existent, où l'on renvoie en entrée l'état de sortie...

Classique	Exemple classique simple
Espace de probabilité $(\Omega, \mathcal{A})$	$\{0,1\}$
Mesure de probabilité p_θ	$\left(\frac{1}{2}(1+\theta), \frac{1}{2}(1-\theta)\right)$ avec $-1 \leq \theta \leq 1$.
Mesure de Dirac	$(1,0)$ or $(0,1)$ donnée par $\theta = -1$ ou 1.
Estimateur à valeurs dans l'espace mesuré $(\mathcal{X}, \mathcal{A})$ $X : \Omega \otimes \Omega_2 \to \mathcal{X}$ où $(\Omega_2, \mathcal{B}, q)$ est un espace de probabilité avec q connu.	$X : i \mapsto X_i(\omega_2)$ with $X_i : \Omega_2 \to \mathcal{X}$ for $i = 0, 1$, où $(\Omega_2, \mathcal{B}, q)$ est un espace de probabilité avec q connu.
Loi de probabilité de l'estimateur $\mathbb{P}_\theta[X \in A] = (p_\theta \otimes q)(X^{-1}(A))$.	$\mathbb{P}_\theta[X \in A] = \frac{1}{2}(1-\theta)q(X_0^{-1}(A)) + \frac{1}{2}(1+\theta)q(X_1^{-1}(A))$.
Noyau de Markov (donné par (1.5)) τ	$p_\theta \mapsto p_\theta(0)\tau_0 + p_\theta(1)\tau_1$ avec τ_0 et τ_1 des lois de probabilité sur le même espace.

FIG. 1.1 – Correspondances entre notions classiques et quantiques de base

Quantique	Exemple quantique simple
Espace de Hilbert $\mathcal{H}$	$\mathbb{C}^2$
État (donné par la Définition 1.1.1) ρ_θ	$\frac{1}{2}\left(\mathbf{1}_{\mathbb{C}^2} + \sum_{i=1}^{3} \theta_i \sigma_i\right)$ avec σ_i donné par (1.16) et $\|\theta\| = 1$.
État pur $\|\psi\rangle\langle\psi\|$ avec $\langle\psi\|\psi\rangle = 1$.	ρ_θ de rang un, équivalant à $\|\theta\| = 1$ dans la formule précédente.
POVM (donnée par la Définition 1.1.2), à valeur dans l'espace mesuré $(\mathcal{X}, \mathcal{A})$ $M = \{M(A)\}_{A\in\mathcal{A}}$	Pas de simplification
Loi de probabilité de la mesure $\mathbb{P}_\theta[X \in A] = \mathrm{Tr}(\rho_\theta M(A))$.	Pas de simplification
Canal (donné par la Définition 1.1.3) $\mathcal{E} : \mathcal{T}(\mathcal{H}) \to \mathcal{T}(\mathcal{K})$.	Si $\dim(\mathcal{K}) = d < \infty$, alors $\mathcal{E}(\rho_\theta) = \sum_{\alpha=1}^{2d} R_\alpha \rho_\theta R_\alpha^*$ avec $R_\alpha \in M_{d,2}(\mathbb{C})$ et $\sum_\alpha R_\alpha^* R_\alpha = \mathbf{1}_{\mathbb{C}^2}$.

FIG. 1.2 – Correspondances entre notions classiques et quantiques de base

comme système auxiliaire permettait une estimation asymptotiquement trois fois plus rapide que toute stratégie du premier et du troisième types.

De manière encore beaucoup plus impressionnante, l'utilisation de l'intrication (quatrième et cinquième stratégies) permet d'estimer des opérations unitaires avec une erreur quadratique au carré en $1/n^2$. Par contraste, toutes les premières stratégies fourniraient n copies d'un état, et la borne de Cramér-Rao (1.29) nous assure que la vitesse ne peut être meilleure que $1/n$.

En tout cas, choisir ce que nous permettons n'est qu'une partie du problème. La question la plus difficile reste de trouver quel état envoyer. L'expérience quantique obtenue en sortie dépend beaucoup de ce choix. Quand on utilise seulement un système auxiliaire, les états d'intrication maximale sont le choix naturel. Quand nous utilisons les immenses états intriqués en entrée de la quatrième expérience, nous sommes guidés par la théorie des groupes.

Nous étudions la discrimination entre deux canaux de Pauli au Chapitre 3.

Au Chapitre 4, nous traitons l'estimation de canaux unitaires sur des espaces de dimension finie, et la Section 1.4 correspondante dans l'introduction s'étend plus longuement sur l'histoire et les références.

1.2 Tomographie homodyne

1.2.1 Motivation

Pour pouvoir communiquer à un niveau quantique, il faut à la fois pouvoir coder l'information dans l'état d'un système, transmettre ce système, et récupérer l'information.

La lumière étant très facile à transmettre, elle est un candidat naturel à servir de vecteur de communication. Toutefois, pour transmettre davantage d'information quantique que la polarisation, il faut pouvoir identifier l'état d'un mode de la lumière.

La première méthode permettant de reconstituer complètement l'état d'un oscillateur harmonique est la *tomographie quantique homodyne*, proposée par Vogel et Risken (1989) et implémentée par Smithey et al. (1993).

Sa relative facilité de mise en œuvre et sa relative rapidité, dans la mesure où de nombreuses copies de l'état peuvent être examinées à chaque seconde, en font un outil de base de l'information et de l'optique quantique actuelles. Cette technique

peut tout d'abord être utilisée au terme d'une expérience pour vérifier que l'état voulu a bien été créé (par exemple par Ourjoumtsev, 2007). Elle pourra aussi servir à certaines formes d'expériences de Bell (Daffer et Knight, 2005). Et comme évoqué plus haut, en transmission d'information quantique. Le livre de Leonhardt (1997) et l'article de revue de Lvovsky et Raymer (2005) détaillent davantage l'intérêt de cette méthode de mesure.

Du point de vue du mathématicien, comme l'instrument de mesure est donné, nous récupérons un problème de statistiques classiques. L'aspect quantique du problème n'est plus présent que dans la forme étrange de l'espace des paramètres. Ainsi la tomographie homodyne peut servir d'introduction en douceur au monde des statistiques quantiques. Quant aux problèmes soulevés, il nous faut, après récupération des données, déterminer l'état. Nous restons avec une inversion de transformée de Radon à effectuer, d'où le nom de tomographie. C'est un problème inverse mal posé, en dimension infinie, qui n'est donc pas sans poser certains défis. Mettre au point des méthodes d'estimation plus efficaces permet aux physiciens d'acquérir moins de données, et donc de gagner du temps ou augmenter les débits.

1.2.2 Résultats antérieurs

Mathématiquement, nous voulons estimer un état de l'oscillateur harmonique, sur l'espace de Fock, tel que décrit dans la sous-section 1.1.2. Nos données sont des échantillons de loi p_ρ sur $\mathbb{R} \times [0, \pi]$, qui sont la transformée de Radon d'une autre représentation usuelle en optique quantique, à savoir la *fonction de Wigner*. On peut notamment voir la fonction de Husimi (1.22) comme la fonction de Wigner convoluée avec une gaussienne de variance $\hbar/2$.

Les premières méthodes d'estimation utilisaient la fonction de Wigner comme représentation de l'état. Les données étaient collectées en histogrammes et lissées, puis on inversait la transformée de Radon. Toutefois, ce lissage introduit un biais difficile à contrôler.

La première méthode d'estimation sans biais remonte à D'Ariano et al. (1994), qui ont introduit les fonctions motif, des bases biduales pour les entrées matricielles de l'opérateur densité ρ. (D'Ariano et al., 1995) les ont ensuite généralisé pour gérer le bruit au niveau des détecteurs. Notons que ce bruit est très handicapant, car extrêmement lisse. Banaszek et al. (1999) ont également appliqué l'estimateur du maximum de vraisemblance, avec une bien plus grande efficacité. Le revers de la médaille est le temps de calcul.

Dans chacun de ces cas, néanmoins, il faut fixer à l'avance quelles coordonnées, typiquement celles correspondant à un petit nombre de photons, pourront ne pas

être nulles, sinon la variance de l'estimateur est infinie. Artiles et al. (2005) ont toutefois pu établir la consistance de ces estimateurs utilisés avec un tamis.

Enfin, Butucea et al. (2007), ont mis au point une méthode d'estimation par noyaux de la fonction de Wigner asymptotiquement optimale au sens L^2, sur de larges classes d'états, y compris en tenant compte du bruit de détection.

Par ailleurs, D'Ariano et al. (2004) ont utilisé la tomographie homodyne dans une procédure de calibration d'un compteur de photons.

Mentionnons pour finir la thèse de Meziani (2008), qui porte notamment sur des procédures de test basées sur la fonction de Wigner.

1.2.3 Contributions de la thèse

Je me suis attaqué au problème évoqué, ci-dessus, à savoir que l'on doit choisir à quel niveau on coupe l'estimation des paramètres dans les procédures d'estimation par fonctions motif ou maximum de vraisemblance. Il ne s'agit de rien d'autre que d'un choix de modèle. J'ai donc appliqué les techniques de sélection de modèles et de pénalisation à la Birgé-Massart pour créer un estimateur complètement automatique de l'état d'un oscillateur harmonique.

Au passage, j'ai prouvé que nous récupérions une vitesse d'estimation polynomiale bien que nous soyions en dimension infinie, du moment que nous avons une borne sur l'énergie.

J'ai de même utilisé ces méthodes pour la calibration du compteur de photons telle qu'évoquée par D'Ariano et al. (2004), que j'ai aussi interprétée comme un problème classique de données manquantes.

1.3 Discrimination

1.3.1 Motivation

Alice et Bertrand veulent établir et partager une clé cryptographique sûre. Alice envoie alors une suite de particules à Bertrand, où chaque particule est soit dans l'état $|\psi_1\rangle$ soit dans l'état $|\psi_2\rangle$. Ces états ne sont pas orthogonaux. Pourtant Bertrand peut mesurer chacun d'entre eux et obtenir l'un des trois résultats suivants : l'état est $|\psi_1\rangle$, $|\psi_2\rangle$, ou «je ne connais pas l'état». Quand il a une réponse explicite, l'état

est toujours correctement identifié. Quand il obtient le résultat douteux, Bertrand téléphone simplement à Alice pour lui dire de jeter ce bit là. Pour une efficacité maximale, Bertrand doit réaliser une mesure qui donne une solution explicite aussi souvent que possible.

Or, Ève les espionne. Si elle ne veut pas être remarquée, elle doit envoyer un état à Bertrand, quelle que soit la conclusion de sa mesure. Contrairement à Bertrand, elle n'a pas le droit de dire «je ne sais pas». Donc sa meilleure stratégie consiste à réaliser la mesure qui lui donnera le plus souvent raison, même si elle n'est pas certaine d'avoir correctement identifié l'état. Comme les états ne sont de toute façon pas orthogonaux, elle finira par faire une erreur et sera repérée.

Ce protocole de distribution de clé quantique a été proposé par Bennett et al. (1992). Il contient les deux exemples les plus élémentaires de discrimination quantique. Le cadre général est le suivant. Nous avons un objet quantique, en général un état. Nous savons qu'il appartient à un ensemble fini. Nous devons deviner duquel il s'agit. Pour choisir une stratégie optimale, il nous faut une fonction de coût. Les plus naturelles sont les deux qui apparaissent dans l'exemple ci-dessus. Le critère de Bertrand est appelé *discrimination optimale sans ambiguïté*, celui d'Ève *discrimination avec erreur minimale.*

Historiquement, la discrimination avec erreur minimale fut étudiée en premier, dès Helstrom (1976). En effet, elle correspond aux tests d'hypothèses, un sujet majeur en statistiques classiques. Ivanovic (1987) a introduit la discrimination sans ambiguïté. Contrairement à la discrimination avec erreur minimale, le problème classique correspondant est trivial. Cependant, il y a de nombreux liens avec d'autres sujets d'information quantique, tels le clonage exact (Chefles et Barnett, 1998b) ou la distillation d'intrication (Chefles et Barnett, 1997).

1.3.2 Résultats antérieurs

Chefles (2000) et Bergou et al. (2004) ont récemment écrit deux articles de revue, qui sont mes sources principales pour cette partie historique.

En premier point, remarquons que tous les travaux précédents ont été effectués dans un cadre bayésien. On peut alors expliciter le problème d'Ève ainsi : elle essaie de trouver une POVM $P = (P_1, P_2)$ qui minimise la probabilité moyenne d'erreur, ou de manière équivalente qui maximise la probabilité de succès :

$$p_S = \pi_1 \operatorname{Tr}(\rho_1 P_1) + \pi_2 \operatorname{Tr}(\rho_2 P_2), \tag{1.30}$$

où π est la loi *a priori* et $\rho_i = |\psi_i\rangle\langle\psi_i|$.

Bertrand doit quant à lui maximiser la même expression (1.30), mais avec une POVM $P = (P_1, P_2, P_?)$, et la condition supplémentaire $\mathrm{Tr}(\rho_2 P_1) = \mathrm{Tr}(\rho_1 P_2) = 0$. Ici $P_?$ correspond à la réponse ambigüe. Avec notre définition d'un problème statistique en trois points (ce que nous avons, ce que nous pouvons faire, ce que nous voulons savoir), la différence se situe sur le second point : ce que nous pouvons faire.

Commençons par suivre Helstrom (1976) sur la discrimination avec erreur minimale. Comme $P_2 = 1 - P_1$, en écrivant $\rho_1 = |\psi_1\rangle \langle\psi_1|$ et $|\psi_2\rangle \langle\psi_2|$, nous obtenons

$$p_S = \pi_2 \mathrm{Tr}(\rho_2) + \mathrm{Tr}(P_1(\pi_1\rho_1 - \pi_2\rho_2)).$$

Ainsi une POVM optimale est donnée par P_1 la projection sur le support de la partie positive de $\pi_1\rho_1 - \pi_2\rho_2$. En particulier, la POVM est observable. Ceci résout la discrimination avec erreur minimale pour deux états, même s'ils sont mélangés La même stratégie fonctionnerait si nous ajoutions des poids aux différentes erreurs.

Les difficultés pour l'erreur minimale commencent quand nous avons plus de deux états possibles, disons N. Nous pouvons écrire la probabilité de succès sur le modèle de l'équation (1.30), c'est-à-dire $\sum_i \pi_i \mathrm{Tr}(P_i\rho_i)$. Cependant, nous ne pouvons plus utiliser l'astuce du remplacement de P_1 par $1-P_2$, et il n'y a pas de solution générale au problème de maximisation. Résumons ce que nous savons, néanmoins.

Pour commencer, Eldar (2003) a montré que l'une des POVMs optimales était toujours une observable, du moment que les ρ_i sont linéairement indépendants. Via les multiplicateurs de Lagrange, Holevo (1973) et Yuen et al. (1975b) ont donné une solution implicite ; les conditions suivantes sont nécessaires et suffisantes pour qu'une POVM soit optimale :

$$P_i(\pi_i\rho_i - \pi_j\rho_j)P_j = 0,$$
$$\sum_{k=1}^{N}(\pi_k\rho_k)P_k - \pi_i\rho_i \geq 0,$$

pour tout $1 \leq i, j \leq N$.

Nous avons des solutions analytiques dans quelques cas particuliers (Barnett, 2001; Yuen et al., 1975b; Andersson et al., 2002). Le plus intéressant est celui où le problème est covariant, c'est-à-dire quand $\pi_i = 1/N$ pour tout i, et il y a un opérateur unitaire V tel que $V^N = I$ et $\rho_i = V^{i-1}\rho_1 V^{1-i}$. Nous pouvons alors appliquer Holevo (1982) et chercher une solution de la forme $P_i = V^i \Xi V^{-i}$, où Ξ est appelée le germe de la POVM. Ce point de départ a permis d'abord à Ban et al. (1997) pour les états purs, puis à Eldar et al. (2004) et Chou et Hsu (2003) pour les états mélangés généraux, de dériver une solution. Ils ont obtenu la «mesure racine carrée», qui pour

des états purs $|\psi_1\rangle$ devient :

$$P_i = B^{-1/2}|\psi_i\rangle\langle\psi_i|B^{-1/2}$$
$$\text{with} \quad B = \sum_i |\psi_i\rangle\langle\psi_i|.$$

Quoique nous avons une solution explicite pour deux états, il est difficile de dire à quelle vitesse nos prédictions s'améliorent qui nous avons n copies du même état, si bien que nous devons discriminer entre $\rho_1^{\otimes n}$ et $\rho_2^{\otimes n}$. Les travaux récents se sont concentrés sur cette vitesse, et sur quelle classe de mesures peut l'atteindre (Hayashi, 2002b; Nagaoka et Hayashi, 2007; Nussbaum et Szkola, 2006; Audenaert et al., 2007; Kargin, 2005). Ils utilisent essentiellement des bornes de Chernoff quantiques ou le théorème de Sanov, c'est-à-dire la théorie des grandes déviations quantiques. Ces résultats sont également exprimés dans le cadre minimax.

Enfin, puisque nous essayons de minimiser une fonctionnelle linéaire sous des contraintes linéaires, à savoir que P doit être une POVM, la programmation linéaire semi-définie permet un traitement numérique efficace (Jezek et al., 2002).

Riis et Barnett (2001) ont implémenté expérimentalement la situation d'Ève, c'est-à-dire la discrimination entre deux qubits, tandis que Clarke et al. (2001b) a réalisé la discrimination des états trines et tétrades, *i.e.* trois et quatre états purs qui forment les sommets d'un triangle et d'un tétraèdre réguliers.

Revenons-en au problème de Bertrand, la discrimination sans ambiguïté de deux états $|\psi_1\rangle$ et $|\psi_2\rangle$. Pour la loi *a priori* équiprobable $\pi_1 = \pi_2 = 1/2$, Ivanovic (1987); Dieks (1988) et Peres (1988) ont trouvé la mesure optimale. La probabilité correspondante d'obtenir un résultat explicite est alors appelée la limite IDP :

$$p_S = 1 - |\langle\psi_1|\psi_2\rangle|. \tag{1.31}$$

Comment arrive-t-on à ce résultat ? Tout d'abord, la seule partie pertinente de l'espace est celle générée par les deux vecteurs $|\psi_1\rangle$ et $|\psi_2\rangle$, et est donc de dimension deux. Nous pouvons ainsi considérer la base biorthogonale à (ψ_1, ψ_2), c'est-à-dire la base non orthogonale (ω_1, ω_2) caractérisée par $\langle\omega_i|\psi_j\rangle = \delta_{ij}$ pour $1 \leq i, j \leq 2$. De plus, l'élément de POVM P_1 doit satisfaire à $\mathrm{Tr}(P_1\rho_2) = 0$, ou de manière équivalente, son support doit être orthogonal à $|\psi_2\rangle$. Donc $P_1 = c_1|\omega_1\rangle\langle\omega_1|$. De même, $P_2 = c_2|\omega_2\rangle\langle\omega_2|$. Nous devons donc simplement trouver les c_1 et c_2 qui maximisent (1.30) tout en gardant $P_1 + P_2 \leq I$. Alors $P_? = I - P_1 - P_2$. Par symétrie, si $\pi_1 = \pi_2$, nous devons avoir $c_1 = c_2$. Nous prenons donc le c_1 le plus grand tel que $P_1 + P_2 \leq I$. Les calculs mènent à (1.31).

La discrimination sans ambiguïté, contrairement à la discrimination avec erreur minimale, se généralise assez bien à plusieurs états purs. En revanche, discriminer même entre deux états mélangés est difficile.

Jaeger et Shimony (1995) ont généralisé au cas $\pi_1 \neq \pi_2$. Pour plus de deux états purs, nous pouvons commencer notre raisonnement de la même manière : nous écrivons $P_i = c_i|\omega_i\rangle\langle\omega_i|$, avec $\{\omega_i\}_{1\leq i\leq N}$ la base biorthogonale de $\{\psi_i\}_{1\leq i\leq N}$. Nous avons donc à gérer uniquement N coefficients. Mais nous n'avons pas en général de solution explicite. Les cas particuliers résolus incluent le cas covariant, où $|\psi_i\rangle = V^{i-1}|\psi_N\rangle$, et $V^N = I = VV^*$ (Chefles et Barnett, 1998a). Les principaux résultats pour plusieurs états purs sont des bornes supérieures et inférieures sur la probabilité de succès. Zhang et al. (2001) ont prouvé que :

$$p_S \leq 1 - \frac{1}{N-1} \sum_{\substack{1\leq j,k\leq N \\ j\neq k}} \sqrt{\pi_j \pi_k}|\langle\psi_i|\psi_j\rangle|.$$

Remarquons que la limite IDP sature cette borne. De l'autre côté, Sun et al. (2002) ont montré que p_S était plus grande que la plus petite valeur propre de la matrice $N \times N$ dont les éléments sont les produits scalaires $\langle\psi_i|\psi_j\rangle$. Ils ont utilisé des travaux antérieurs de Duan et Guo (1998), portant sur le clonage.

Cependant, l'essentiel de la littérature tourne autour de la discrimination de deux états mélangés, ou davantage. Je serai bref comme je n'ai pas abordé ce cas là. Rudolph et al. (2003) ont donné des bornes inférieures et supérieures sur la probabilité de succès p_S, et montré qu'elles correspondent souvent. En résultat annexe, ils donnent une solution quand le rang des matrices densité est la dimension de l'espace de Hilbert moins un. De plus, Raynal et al. (2003) ont montré qu'ils pouvaient réduire l'étude de la discrimination à celles de deux matrices densité de même rang sur un espace de Hilbert de dimension deux fois ce rang. De plus, Feng et al. (2005) a donné des bornes supérieures pour la discrimination entre N états mélangés, et Qiu (2007) une borne inférieure. Herzog et Bergou (2005); Raynal et Lütkenhaus (2005); Herzog (2007) ont fourni des solutions explicites dans plusieurs cas particuliers.

Comme pour la discrimination avec erreur minimale, Eldar (2003) a montré que nous pouvions utiliser les techniques de programmation linéaire semi-définie. Qui plus est, Huttner et al. (1996); Clarke et al. (2001a) ont implémenté expérimentalement la situation de Bertrand, c'est-à-dire la discrimination entre deux états purs. Mohseni et al. (2004) a également mis en pratique le cas plus compliqué de la discrimination entre un état pur et un état mélangé.

Jusqu'ici, nous avons uniquement évoqué la discrimination entre états. On peut souhaiter discriminer entre d'autres objets quantiques, par exemple entre des canaux. Nous avons un canal $\mathcal{E}$ dont nous savons qu'il fait partie d'un ensemble fini $\{\mathcal{E}_i\}_{1\leq i\leq k}$.

Nous devons sonder la boîte noire $\mathcal{E}$ avec un état ρ. Nous obtenons $\mathcal{E}(\rho)$ en sortie, et devons ensuite discriminer entre les états $\mathcal{E}_i(\rho)$. Nous sommes de retour à la situation précédente, à ceci près que nous devons choisir l'état d'entrée pour obtenir les sorties les plus faciles à distinguer. Le choix de l'entrée est le point le plus difficile, et soulève ses propres questions, comme de déterminer si un système auxiliaire peut s'avérer utile.

Childs et al. (2000b) ont été les premiers à étudier la discrimination avec erreur minimale pour des canaux unitaires, en insistant sur les applications en information quantique, telles l'algorithme de Grover (1996) pour les recherches en bases de données. Sacchi (2005b) ont considéré des canaux de Pauli, comme exemple élémentaire de canaux non unitaires. La discrimination sans ambiguïté a été abordée plus récemment. Wang et Ying (2006) a trouvé sous quelles conditions deux canaux peuvent être distingués sans ambiguïté, soit avec une entrée, soit avec plusieurs. Dans ce dernier cas, intriquer les états en entrée améliore en général les résultats. Enfin, Chefles et al. (2007) ont rassemblé les résultats connus en discrimination sans ambiguïté, et en ont ajouté d'autres, dans un article clairement motivé par l'informatique quantique. Davantage de travail est nécessaire sur ces questions.

Quoiqu'ils n'apparaissent pas dans cette thèse, la discrimination recouvre d'autres aspects. Une première classe de problèmes vient de l'utilisation d'autres critères d'optimalité (par exemple Fiurasek et Jezek, 2003; Touzel et al., 2007; Sasaki et al., 2002). Herzog et Bergou (2002) ont aussi exploré la discrimination entre classes d'états, ou filtrage. Une extension à la mode est la suivante : nous avons jusqu'ici supposé que nous pouvions utilisé n'importe quelle mesure permise par la physique. Mais si nous partons d'un état produit, nous ne pouvons pas forcément effectuer des mesures collectives, et devons nous restreindre aux mesures LOCC. Une application possible est le partage du secret : trouver un état qu'Alice et Bertrand pourront identifier s'ils coopèrent, mais pas individuellement. Un tel état devrait être symétrique. Un point de départ pour la bibliographie est l'article de revue de Bergou et al. (2004), et ses références, ou le travail plus contemporain d'Owari et Hayashi (2008).

1.3.3 Contributions de la thèse

Comme je l'ai déjà mentionné, tous les travaux antérieurs utilisaient le cadre bayésien, exigeant une probabilité *a priori*. Mon travail, réalisé en collaboration avec G.M. d'Ariano et M.F. Sacchi, a consisté en l'étude du cadre minimax, particulièrement utile s'il n'y a pas de raison physique de choisir une loi *a priori*.

À partir du lien entre risques minimax et bayésiens, donné dans la Section 1.1.1, nous avons exprimé la solution quand les états sont covariants. Cette solution est

la même que pour la loi *a priori* uniforme. Relevons une différence importante avec le scénario bayésien : même pour la discrimination avec erreur minimale entre deux états, la mesure optimale n'est en général pas une observable.

Nous avons aussi prouvé qu'il y a toujours une solution minimax à la discrimination avec erreur minimale pour tout ensemble fini d'états éventuellement mélangés ρ_i, où tous les états ont la même probabilité d'être identifiée, autrement dit où $\mathrm{Tr}(\rho_i P_i)$ ne dépend pas de i.

La discrimination sans ambiguïté minimax se révèle plus simple que la discrimination bayésienne pour plusieurs états purs : nous avons toujours une solution explicite. De manière similaire à nos explications sous l'équation (1.31), nous pouvons prouver que les éléments de POVM doivent être de la forme $P_i = c_i|\omega_i\rangle\langle\omega_i|$, où $\{\omega_i\}$ est une base biorthogonale à $\{\psi_i\}$. Alors les c_i sont tous donnés par la valeur propre la plus basse d'une matrice dépendant des ω_i. Quand il y a plusieurs solutions, nous pouvons raffiner le critère minimax pour en choisir une unique.

Nous avons également étudié la discrimination avec erreur minimale entre deux canaux de Pauli. Quand nous pouvons utiliser un système auxiliaire, nous avons montré que l'efficacité maximale était obtenue avec un état d'intrication maximale, tout comme dans le cas bayésien. Nous avons aussi caractérisé les canaux de Pauli pour lesquels l'usage d'un système auxiliaire améliore les chances de succès. Point intéressant, si dans le cas bayésien nous pouvons toujours choisir un état propre de l'une des matrices de Pauli en entrée, ces choix peuvent ne pas être optimaux au sens minimax.

1.4 Estimation Rapide d'Opérations Unitaires

1.4.1 Motivation

L'évolution d'un système quantique en l'absence de mesure est unitaire. De ce fait, considérer cette évolution comme une boîte noire à estimer signifie estimer un opérateur unitaire. Cela peut nous fournir des informations sur la physique du système.

Il y a aussi de nombreux cas en information quantique où nos devons estimer une opération unitaire, le plus souvent car cela correspond à l'orientation des vecteurs de base, c'est-à-dire la partie purement quantique d'un état.

Ces deux catégories de problèmes en tête, nous pouvons donner davantage de détails sur les diverses applications. Certaines requièrent juste l'estimation d'un paramètre :

Horloges quantiques L'évolution d'un système est donnée par $U_t = e^{itH}$. Une horloge quantique consiste à estimer le paramètre libre t, c'est-à-dire le temps. Nous devons donc réaliser de l'estimation dans une famille d'unitaires à un paramètre (Buzek et al., 1999).

Mesures de précision Plus généralement, les petites forces de forme connus et d'intensité inconnue vont apparaître dans l'opérateur d'évolution comme $U = e^{i\phi H}$. Déterminer ϕ revient à déterminer la force. Un usage notable concerne les accéléromètres (Yurke, 1986).

D'autres applications exigent de déterminer tout l'opérateur :

Transmission de repères de référence Quand Alice et Bertrand veulent communiquer en échangeant des qubits, ou plus généralement des états à d dimensions, ils doivent se mettre d'accord sur les axes de la mesure, c'est-à-dire le repère de référence (Holevo, 1982). Ces axes vont tourner durant la transmission des particules d'Alice à Bertrand. Ce dernier doit donc estimer la rotation des axes, soit l'évolution unitaire des objets. Remarquons néanmoins l'existence de protocoles basés sur les représentations de groupe, où les repères de référence sont superflus (Bartlett et al., 2003).

Estimation d'états d'intrication maximale Les états d'intrication maximale sont des ressources fondamentales en téléportation quantique (Bennett et al., 1993) ou en information quantique (Ekert, 1991). Pour atteindre à une efficacité maximale, néanmoins, Alice et Bob doivent savoir *quel* état d'intrication maximale ils partagent, soit quel est l'unitaire U tel que $|\psi\rangle = \frac{1}{d}\sum |i\rangle \otimes U\,|i\rangle$.

1.4.2 Résultats antérieurs

À ma connaissance, Yurke (1986) a été le premier a remarquer qu'un paramètre d'une évolution quantique pouvait être estimé à un taux $1/N^2$ (pour l'erreur au carré), où N est le nombre d'états qui subissent l'évolution. C'est extrêmement remarquable, puisque les paramètres ne peuvent en général être estimés qu'à vitesse $1/N$ dans les situations classiques.

Ce genre d'estimation rapide, qui utilise l'intrication entre les états en entrée, sature ce que les physiciens appellent la limite de Heisenberg, la limite fondamentale à la précision des mesures quantiques. Giovannetti et al. (2004) ont récemment écrit un article de revue de ce genre d'estimation accélérée, et font mention d'expériences. La plupart des méthodes pratiques reposent soit sur des photons obtenus par conversion paramétrique (*e.g.* Eisenberg et al., 2005), soit sur des pièges à ions (*e.g* Dalvit et al., 2006) soit sur des atomes en cavité QED (*e.g.* Vitali et al., 2006).

Acin et al. (2001) ont les premiers donné la forme générale de l'état d'entrée optimal, avec des coefficients non spécifiés dépendant de la fonction de coût, pour

tout problème d'optimisation bayésienne uniforme avec une fonction de coût $SU(d)$-covariante. Quand nous avons le droit d'envoyer N particules à travers l'opérateur unitaire, elle s'écrit :

$$|\Phi\rangle = \bigoplus_{\vec{\lambda}:|\vec{\lambda}|=N} \frac{c(\vec{\lambda})}{\sqrt{\mathcal{D}(\vec{\lambda})}} \sum_{i=1}^{\mathcal{D}(\vec{\lambda})} |\psi_i^{\vec{\lambda}}\rangle \otimes |\psi_i^{\vec{\lambda}}\rangle, \tag{1.32}$$

où nous utilisons les notations du Chapitre 4 pour les représentations de groupe. Les coefficients $c(\vec{\lambda})$ dépendent de la fonction d'optimisation, et les $|\psi_i^{\vec{\lambda}}\rangle$ forment une base orthonormale de l'espace $\mathcal{H}^{\lambda}$. Seules les N premières particules, correspondant à la droite du produit tensoriel, sont envoyées à travers l'opérateur unitaire ? Comme notre problème initial est entièrement invariant sous l'action de $SU(d)$, il n'est pas surprenant que la solution le soit également. Plus tard, Chiribella et al. (2005) ont généralisé cette équation à d'autres symétries, et spécifié les coefficients comme coordonnées d'un vecteur propre d'une matrice dépendant des coefficients de Clebsch-Gordan.

Les travaux ultérieurs se sont concentrés sur $SU(2)$. Peres et Scudo (2001) ont été les premiers à donner une stratégie convergeant à vitesse $1/N^2$ avec la fidélité comme fonction d'évaluation, bien que l'état d'entrée n'ait pas été covariant. Bagan et al. (2004a) a alors trouvé les bons coefficients dans l'équation (1.32) et atteint la même vitesse, avec la constante optimale π^2/N^2. Puis Bagan et al. (2004b) ont Chiribella et al. (2004) noté que l'on pouvait se passer de système auxiliaire. On a alors moitié moins de particules à préparer. Ils remplacent l'intrication avec des particules extérieures par une «auto-intrication», basée sur le fait que la multiplicité $\mathcal{M}(\vec{\lambda})$ de la plupart des représentations irréductibles est suffisamment élevée dans la représentation N-tensorielle.

Hayashi (2004) ont établi des résultats similaires avec des critères minimax. En ce qui concerne $SU(d)$, Ballester (2005b) a juste donné une indication que cette vitesse de $1/N^2$ pouvait être atteinte. Il a trouvé un état d'entrée telle que la matrice d'information de Fisher quantique (1.27) se comporte en $1/N^2$. Il n'a pas trouvé de procédure d'estimation complète, cependant.

Notons que ces hautes vitesses ne peuvent pas être généralisées à des canaux arbitraires. En effet, de nombreuses familles continues de canaux peuvent être programmées par une famille continue d'états ρ_θ, c'est-à-dire que nous pouvons choisir une opération unitaire agissant sur $\sigma \otimes \rho_\theta$, et observer uniquement l'effet sur σ. Alors, estimer θ pour les canaux revient à l'estimer pour les états ρ_θ. À cause de la borne de Cramér-Rao classique (1.15), cette dernière estimation est toujours plus lente que $1/N$ (Ji et al., 2006). Fujiwara et Imai (2003) ont explicitement dérivé de taux maximum de $1/N$ pour les canaux de Pauli généralisés, et mentionnent une remarque équivalente de Hayashi (2006).

1.4.3 Contributions de la thèse

Acin et al. (2001) et Chiribella et al. (2005) ont fourni une forme générale pour estimer de manière optimale une opération unitaire. Cependant, on ne peut y lire la vitesse. Mon travail a consisté à trouver des coefficients $c(\vec{\lambda})$ dans l'état (1.32) pour lesquels les calculs sont possibles, et prouver que nous atteignions encore la vitesse $1/N^2$, à la fois dans les cadres bayésiens et minimax. Imai et Fujiwara (2007) ont depuis indépendamment donné une interprétation de géométrie différentielle à ce taux.

L'idée est la suivante : les calculs montrent que $c(\vec{\lambda})$ doit être presque égal à $c(\vec{\lambda})'$ pour $\vec{\lambda}$ et $\vec{\lambda}'$ correspondent à des tableaux de Young qui ne diffèrent que d'une boîte. Quand $\lambda_i = \lambda_{i+1}$ pour un i donné, nous devons aussi prendre un petit $c(\vec{\lambda})$. Nous choisissons alors nos coefficients proportionnels à

$$c(\vec{\lambda}) = \prod_{i=1}^{d} (\lambda_i - \lambda_{i+1}),$$

et vérifions que nous avons la bonne vitesse.

1.5 Mesures à Valeurs dans les Opérateurs Positifs Propres

Nous avons un appareil de mesure $\mathbf{P}$. Nous pourrions vouloir réutiliser ce coûteux appareil pour des effectuer des mesures différentes. À cette fin, nous pouvons transformer ρ avant d'utiliser notre appareil. Cette combinaison d'une transformation et d'une mesure correspond à un nouvel instrument de mesure $\mathbf{Q}$.

Ce scénario, illustré par la Figure 1.5, soulève quelques questions naturelles. Mathématiquement, nous partons d'une POVM $\mathbf{P}$ et obtenons une nouvelle POVM $\mathbf{Q} = \mathcal{E}(\mathbf{P})$ par l'application préalable d'un canal $\mathcal{E}$ à l'état ρ. Nous disons alors que $\mathbf{P}$ est *plus propre* que $\mathbf{Q}$. C'est un pré-ordre, noté $\mathbf{P} \succcurlyeq \mathbf{Q}$. On peut se demander, cependant, à $\mathbf{P}$ et $\mathbf{Q}$ fixés, s'il y a un canal $\mathcal{E}$ tel que $\mathbf{Q} = \mathcal{E}(\mathbf{P})$. À $\mathbf{P}$ fixé, quelles sont les POVMs $\mathbf{Q}$ équivalente en propreté à $\mathbf{P}$, *i.e.* telles qu'on ait à la fois $\mathbf{P} \succcurlyeq \mathbf{Q}$ et $\mathbf{Q} \succcurlyeq \mathbf{P}$? Néanmoins, la première étape pour la compréhension de cette relation d'ordre est la connaissance de ses points maximaux : quelles sont les *POVMs propres*, *i.e.* les POVMs $\mathbf{P}$ telles que $\mathbf{Q} \succcurlyeq \mathbf{P}$ implique $\mathbf{P} \succcurlyeq \mathbf{Q}$?

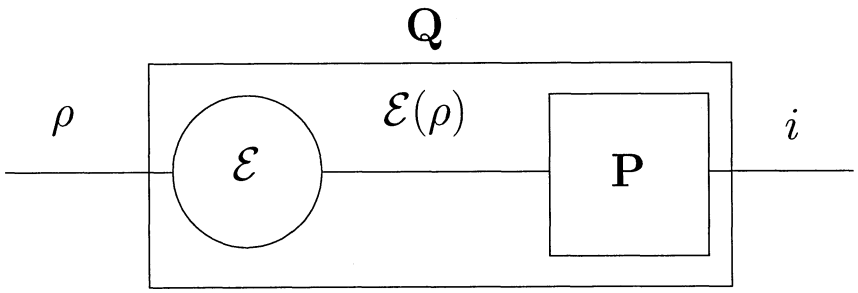

FIG. 1.3 – Nous appliquons le canal $\mathcal{E}$ à ρ avant de le mesurer avec la POVM $\mathbf{P}$. L'opération globale, qui renvoie des données classiques i à partir de l'état ρ, peut être vue comme la mesure de ρ avec une POVM $\mathbf{Q}$. Nous disons que $\mathbf{P}$ est plus propre que $\mathbf{Q}$.

1.5.1 Résultats antérieurs

Le pré-ordre «plus propre que» a été introduit par Buscemi et al. (2005), pour formaliser le traitement *a priori* des POVMs, par opposition à leur traitement *a posteriori*, c'est-à-dire le traitement classique des données classiques de sortie.

Pour nous donner une certaine perspective, , mentionnons quelques autres ordres habituels sur les POVMs (Heinonen, 2005) :
- Une POVM $\mathbf{P}$ *donne plus d'information* qu'une POVM $\mathbf{Q}$ si elle permet de distinguer toutes les paires d'états que $\mathbf{Q}$ permet de distinguer. Une POVM permet de distinguer deux états si les distributions de probabilité des sorties sont différentes. Les POVMs maximales pour dette relation sont dites *infocomplètes* (Prugorevčki, 1977).
- La relation d'ordre plus faible «*avoir une plus grande force de détermination d'état que*» a également les POVMs infocomplètes comme éléments maximaux. Une POVM *détermine* un état si la loi de la sortie ne peut être obtenue qu'avec cet état en entrée (Busch et Lahti, 1989; Davies, 1970).
- Une POVM $\mathbf{Q}$ est une *version floue* (Martens et de Muynck, 1990) de $\mathbf{P}$ si elle peut être obtenue par un traitement *a posteriori* de la sortie de $\mathbf{P}$. Les POVMs maximales sont alors les POVMs de rang un Buscemi et al. (2005), *i.e.* dont tous les éléments sur les singletons sont de rang un au plus.

Remarquons que si $\mathbf{Q}$ est une version floue de $\mathbf{P}$, alors $\mathbf{P}$ donne plus d'information que $\mathbf{Q}$. Cependant, il n'y a pas de relations entre les éléments maximaux. Remarquons aussi que les POVMs de rang un sont les points extrémaux de l'ensemble convexe des POVMs. Si la fonction d'optimisation est convexe, ce qui est souvent le cas, les solutions correspondantes sont de rang un (Helstrom, 1976).

La relation «plus propre que» se révèle assez dénuée de liens avec les relations précédentes. La caractérisation de ses points maximaux est aussi un problème difficile.

Nous connaissons déjà certains résultats partiels, néanmoins. Buscemi et al. (2005) ont prouvé que les POVMs de rang un sont propres, de même que les POVMs où la valeur propre maximale de chaque élément est un. Ce dernier cas part du principe que $\mathbf{P}$ doit avoir le même nombre de valeurs possibles en sortie que $\mathbf{Q}$. Si nous admettons que $\mathbf{P}$ puisse en avoir plus, alors ces dernières POVMs ne sont pas propres, à moins d'être de rang un. En effet, aucun traitement *a priori* ne pourra augmenter le nombre de valeurs possibles en sortie, tandis qu'une observable de rang un traitée *a priori* peut générer toute POVM à d sorties : il suffit de mesurer $\mathbf{Q}$ et préparer l'état propre i comme entrée de l'observable.

Buscemi et al. (2005) ont également prouvé que si $\mathbf{Q}$ est infocomplète et $\mathbf{P} \succcurlyeq \mathbf{Q}$, alors $\mathbf{P}$ est aussi infocomplète, et qu'un effet, c'est-à-dire une POVM à deux sorties, $\mathbf{P} = \{P_1, 1 - P_1\}$ est plus propre qu'un effet $\mathbf{Q} = \{Q_1, 1 - Q_1\}$ si et seulement si $[\lambda_m(P_1), \lambda_M(P_1)] \supset [\lambda_m(Q_1), \lambda_M(Q_1)]$, où λ_m et λ_M sont les plus petites et plus grandes valeurs propres.

Le reste de leur travail s'appuie sur les notions de relations d'ordre et d'équivalence liées.

La plus élémentaire est l'équivalence unitaire. Les POVMs $\mathbf{P}$ et $\mathbf{Q}$ sont unitairement équivalentes si on peut obtenir $\mathbf{Q}$ à partir de $\mathbf{P}$ en utilisant un canal unitaire, *i.e.* $UP_iU^* = Q_i$ pour tous les éléments de POVM. Nous pouvons alors revenir à $\mathbf{P}$ par application du canal unitaire inverse. Ainsi, l'équivalence unitaire implique l'équivalence en propreté. L'inverse n'est pas vrai : prenons par exemple deux effets en dimension trois, avec $P_1 = |\phi\rangle\langle\phi| = 1 - Q_1$. Alors nous n'avons pas équivalence unitaire, mais $\lambda_m(P_1) = 0 = \lambda_m(Q_1)$ et $\lambda_M(P_1) = 1 = \lambda_M(Q_1)$, donc $\mathbf{P}$ et $\mathbf{Q}$ sont équivalentes en propreté. Néanmoins les équivalences unitaires et en propreté sont les mêmes dans un certain nombre de cas particuliers : pour les POVMs infocomplètes, pour les qubits, *i.e.* quand l'espace de Hilbert est de dimension 2, et pour POVMs de rang un.

Pour donner une idée des méthodes, prouvons la dernière assertion sur les POVMs de rang un. Nous pouvons écrire $Q_i = \lambda_M(Q_i)|\psi_i\rangle\langle\psi_i|$ avec $|\psi_i\rangle$ normalisé. Nous pouvons écrire $\lambda_M(Q_i) = \mathrm{Tr}(Q_i|\psi_i\rangle\langle\psi_i|) = \mathrm{Tr}(P_i\mathcal{E}(|\psi_i\rangle\langle\psi_i|))$. Comme $\mathcal{E}(|\psi_i\rangle\langle\psi_i|)$ est un état, cette dernière expression vaut moins que $\lambda_M(P_i) \leq \mathrm{Tr}(P_i)$. Comme les POVMs sont normalisées, nous savons que $\sum_i \lambda_M(Q_i) = d = \sum_i \mathrm{Tr}(P_i)$, où d est la dimension de l'espace de Hilbert. Donc $\mathrm{Tr}(P_i) = \lambda_M(Q_i) = \lambda_M(P_i)$, si bien que $P_i = \lambda_M(Q_i)|\phi_i\rangle\langle\phi_i|$ pour un certain $|\phi_i\rangle$ de norme un. Donc $\mathcal{E}(|\psi_i\rangle\langle\psi_i|) = |\phi_i\rangle\langle\phi_i|$. Donc $\mathcal{E}(Id) = \sum_i \lambda_M(Q_i)\mathcal{E}(|\psi_i\rangle\langle\psi_i|) = \sum_i P_i = Id$, autrement dit $\mathcal{E}$ préserve à la fois la trace et l'identité. Donc son dual aussi, qui envoie $|\phi_i\rangle$ sur $|\psi_i\rangle$. Nous terminons la preuve en rappelant qu'il existe deux canaux envoyant un ensemble d'états purs sur un autre, et réciproquement, si et seulement si ces deux ensembles sont unitairement équivalents (Chefles et al., 2003).

L'autre relation d'ordre qu'ils utilisent est «avoir une plus grande portée», notée $\mathbf{P} \supset_r \mathbf{Q}$, où la portée est l'ensemble des distributions de probabilité possibles à obtenir en sortie, *i.e.* $\{(\mathrm{Tr}(\rho P_i))_i : \rho \text{ état}\}$. Comme nous pouvons fournir $\mathcal{E}(\rho)$ en entrée de $\mathbf{P}$ et obtenir le même résultat que si nous avions fourni ρ en entrée de $\mathbf{Q}$, la relation «plus propre que» est plus forte que «avoir une plus grande portée». L'inverse n'est pas vrai. Toutefois, s'il existe une POVM infocomplète $\mathbf{M}$ sur le même espace de Hilbert, telle que $\mathbf{P} \otimes \mathbf{M} \supset_r \mathbf{Q} \otimes \mathbf{M}$, alors $\mathbf{P} \succcurlyeq \mathbf{Q}$. La présence de $\mathbf{M}$ assure que l'application définie sur l'espace engendré par les éléments de POVM $\{P_i\}$ par $\mathcal{E}(P_i) = Q_i$ est complètement positive, et donc peut être étendue à tout l'espace, par le théorème d'extension d'Arveson (1969).

Enfin, Buscemi et al. (2005) ont aussi prouvé que l'ensemble $\mathcal{C}_{\mathbf{P},\mathbf{Q}}$ des canaux $\mathcal{E}$ qui vérifient $\mathcal{E}(\mathbf{P}) = \mathbf{Q}$ est un ensemble convexe. Nous avons peu de résultats généraux qui concernent également les POVMs non propres.

1.5.2 Contributions de la thèse

Nous avons vu que nous n'avions pas, à l'heure actuelle, de caractérisation des POVMs propres. Cette thèse donne une condition suffisante, et prouve que cette condition est aussi nécessaire pour une certaine catégorie de POVMs, qui inclut notamment toutes les POVMs sur les qubits. Nous avons donc caractérisé les POVMs pour les qubits.

Nous utilisons deux idées principales. Commençons par nous donner une POVM $\mathbf{P}$. Nous voulons prouver qu'elle est propre. En d'autres termes, étant donnée $\mathbf{Q}$ telle que $\mathbf{Q} \succcurlyeq \mathbf{P}$, nous voulons prouver que l'inverse $\mathbf{P} \succcurlyeq \mathbf{Q}$ est aussi vrai. Le cas le plus simple est celui où $\mathbf{P} = \mathcal{E}(\mathbf{Q})$ avec $\mathcal{E}$ unitaire. Nous essayons donc de trouver une condition sur $\mathbf{P}$ sous laquelle $\mathcal{E}$ est unitaire pour tout $\mathbf{Q}$.

Maintenant, par la représentation de Kraus (1.23), nous savons que $P_i = \sum_\alpha R_\alpha^* Q_i R_\alpha$. Tous les éléments de la somme sont positifs, donc $P_i \geq R_\alpha^* Q_i R_\alpha$ pour tout i et α. En particulier, le support de $R_\alpha^* Q_i R_\alpha$ doit être inclus dans le support de P_i, en tant qu'opérateurs sur l'espace de Hilbert $\mathcal{H}$. Ceci nous fournit $d - \dim(\mathrm{Supp}(P_i))$ équations linéaires homogènes sur les éléments matriciels de R_α, pour chaque vecteur dans le support de Q_i. Si nous obtenons par ce biais $d^2 - 1$ équations indépendantes, les matrices R_α sont déterminées à constante près, et la contrainte $\sum R_\alpha^* R_\alpha = Id$ suffira à prouver que $\mathcal{E}$ est unitaire.

La difficulté dans ce scénario réside dans la dépendance de ces équations en $\mathbf{Q}$. J'introduis donc la définition suivante : un ensemble de sous-espaces de $\mathcal{H}$ *détermine totalement* $\mathcal{H}$ s'ils fournissent suffisamment d'équations pour tout $\{Q_i\}$ quand ils sont les supports des P_i. Il se révèle qu'un ensemble de vecteurs $\{|\psi_i\rangle\}$, soit un

ensemble de supports de dimension un, détermine totalement $\mathcal{H}$ si et seulement si, pour toute paire de sous-espaces propres supplémentaire $\mathcal{V}$ et $\mathcal{W}$, il existe un i tel que $|\psi_i\rangle \notin \mathcal{V}$ et $|\psi_i\rangle \notin \mathcal{W}$.

Ceci fournit une condition suffisante pour qu'une POVM soit propre, et elle peut être vérifiée algorithmiquement. J'ai également prouvé qu'être de rang un ou vérifier cette condition était nécessaire si tous les éléments de POVM sont soit de rang un, soit de rang plein. J'appelle pareilles POVMs des *POVMs quasi-qubit*, comme toutes les POVMs qubit sont quasi-qubit.

La nécessité est prouvée en considérant des canaux $\mathcal{E}$ qui sont proches de l'identité, et dont l'inverse est une application positive. On peut alors considérer $\mathbf{Q} = \mathcal{E}^{-1}(\mathbf{P})$ et il nous faut prouver que $\mathbf{Q}$ est une POVM. Un choix précis de $\mathcal{V}$ et $\mathcal{W}$ tels que donnés dans le paragraphe précédent permet de s'en assurer.

Pour les qubits, les POVMs propres sont donc les POVMs de rang un d'une part, et les POVMs ayant au moins trois éléments de rang un non colinéaires deux à deux. Cette dernière condition est une traduction plus intuitive de «détermine totalement » dans le cas des qubits.

1.6 Sous-algèbres complémentaires

1.6.1 Motivation

Nous avons deux qubits intriqués. Nous les laissons évoluer comme nous le souhaitons, puis mesurons l'un d'entre eux. Comment les laisser évoluer si nous voulons reconstruire l'état de ces deux qubits avec aussi peu d'évolutions, et aussi efficacement que possible ?

Formellement, ce problème se traduit comme celui d'estimer un état sur $\mathbb{C}^2 \otimes \mathbb{C}^2$. Nous avons quinze paramètres réels à estimer. Nous avons le droit de mesurer les états réduits sur un sous-espace de dimension 2, c'est-à-dire sur les deux premières coordonnées de $W\mathbb{C}^4$, où W est unitaire, correspondant à l'évolution. Chaque W génère un état réduit, correspondant à trois paramètres. Nous voulons utiliser aussi peu de transformations W que nous le pouvons.

Nous avons à l'évidence besoin d'au moins cinq W différents. Nous pouvons dans un premier temps nous demander si ce nombre est suffisant. Nous pouvons aussi essayer de trouver un ensemble optimal de W. Nous répondons à ces deux questions en remarquant que connaître un état, c'est connaître ses valeurs moyennes sur toute

l'algèbre des observables $M_2(\mathbb{C}) \otimes M_2(\mathbb{C})$. Connaître l'état réduit sur différents sous-espaces, c'est connaître l'état original sur les sous-algèbres $\mathcal{A}_i = W_i(M_2(\mathbb{C}) \otimes Id)W_i^*$, pour différents W_i. Donc les états réduits déterminent l'état d'origine si et seulement si les sous-algèbres $\mathcal{A}_i$ génèrent linéairement, en tant qu'espace vectoriel, l'algèbre initiale $M_2(\mathbb{C}) \otimes M_2(\mathbb{C})$.

Intuitivement, nous obtenons autant d'informations que possible si les sous-algèbres $\mathcal{A}_i$ diffèrent autant que possible l'une de l'autre. Mathématiquement, nous traduisons cela en demandant que les sous-algèbres soient *complémentaires*, c'est-à-dire que $(\mathcal{A}_i - \mathbb{C}\mathbf{1})$ soit orthogonale à $(\mathcal{A}_j - \mathbb{C}\mathbf{1})$ pour $i \neq j$ et le produit scalaire $\langle A|B\rangle = \mathrm{Tr}(A^*B)$ sur $M_4(\mathbb{C})$.

En résumé, nous cherchons cinq sous-algèbres de $M_4(\mathbb{C})$ chacune isomorphe à $M_2(\mathbb{C})$, et complémentaires deux à deux.

1.6.2 Résultats antérieurs

Petz, Hangos, Szántó, et Szöllősi (2006) ont introduit les notions et problème précédents. Ils étaient également motivés par l'analogie avec les observables complémentaires, telles la position et l'impulsion. Schwinger (1960) ont peut-être été les premiers à en donner une approche rigoureuse dans les espaces de dimension finie. Deux observables d'un espace de Hilbert de dimension d sont complémentaires si leurs bases propres vérifient $\langle\phi|\psi\rangle = 1/d$ pour tout ϕ dans la première base et ψ dans la seconde. Ces bases sont fréquemment utilisées en information quantique, que ce soit pour la discrimination d'états (Ivanovic, 1981), pour le «problème du Méchant Roi» (Kimura et al., 2006) ou en cryptographie quantique (Bruss, 1998). Maintenant, nous pouvons associer à une observable l'algèbre commutative des éléments diagonaux dans leur base propre. Deux observables sont complémentaires si et seulement si les algèbres commutatives correspondantes sont complémentaires. L'importance des observables complémentaires peut donner quelque espoir d'utilité pour les sous-algèbres $M_2(\mathbb{C})$ complémentaires.

Revenons au problème initial. Petz et al. (2006) ont prouvé que cinq sous-algèbres suffisaient effectivement pour générer $M_2(\mathbb{C}) \otimes M_2(\mathbb{C})$. Ils ont exhibé quatre sous-algèbres $M_2(\mathbb{C})$ complémentaires. Mais ils n'ont pas pu en trouver cinq. Ils ont également considéré n qubits, avec $M_2(\mathbb{C})^{\otimes n}$ comme algèbre correspondante. On a alors besoin d'au moins $(2^{2n} - 1)/3$ sous-algèbres isomorphes à $M_2(\mathbb{C})$ pour générer l'algèbre de départ. Ils ont prouvé que, si l'on se restreignait aux sous-algèbres générées par des éléments de la forme $\sigma_1 \otimes \sigma_2 \otimes \cdots \otimes \sigma_n$, où chaque σ est une matrice de Pauli (1.16), alors cette borne n'est pas saturée, et il nous faut au moins une sous-algèbre supplémentaire.

Comme choisir des sous-algèbres avec pareils générateurs est la manière la plus simple d'obtenir des sous-algèbres complémentaires, ceci suggère l'impossibilité de générer tout $M_2(\mathbb{C})^{\otimes n}$ avec des sous-algèbres complémentaires isomorphes à $M_2(\mathbb{C})$.

1.6.3 Contributions de la thèse

Ce travail est commun avec Dénes Petz. Nous avons prouvé que le nombre maximal de sous-algèbres complémentaires isomorphes à $M_2(\mathbb{C})$ dans $M_2(\mathbb{C}) \otimes M_2(\mathbb{C})$ était de quatre.

L'idée est la suivante : nous considérons une base orthonormale d'une sous-algèbre $\mathcal{A}$ isomorphe à $M_2(\mathbb{C})$ de la forme $\mathbf{1}, A_1, A_2, A_3$. Comme cette base est orthonormale, les A_i sont de trace nulle. Prenons également $\mathbf{1}, B_1, B_2, B_3$ comme base orthonormale de $\mathbf{1} \otimes M_2(\mathbb{C})$. Si $\mathcal{A}$ est complémentaire à $M_2(\mathbb{C}) \otimes \mathbf{1}$, alors $\sum_{i,j} |\operatorname{Tr}(A_i^* B_j)| \geq 1$. D'autre part, pour $\{C_i\}_{i \leq 16}$ base orthonormale de $M_2(\mathbb{C}) \otimes M_2(\mathbb{C})$, nous avons $\sum_{i,j} |Tr(C_i^* B_j)| = 3$. Donc, il y a au plus trois sous-algèbres complémentaires isomorphes à $M_2(\mathbb{C})$, qui sont aussi complémentaires à $M_2(\mathbb{C}) \otimes \mathbf{1}$.

Pour être complet, je précise que depuis la publication de ce travail, Petz (2006) a prouvé que l'espace orthogonal aux quatre sous-algèbres, plus l'identité, était encore une sous-algèbre, mais commutative.

1.7 Normalité asymptotique locale quantique

1.7.1 Normalité asymptotique locale classique

Comme point de référence et motivation, nous passons rapidement en revue le théorie de la distance entre expériences et de la convergence d'expériences de Le Cam (1986), en insistant sur la normalité asymptotique locale.

Wald (1943) a été le premier à avoir l'idée d'approcher une suite d'expériences par des expériences gaussiennes. Le Cam (1960, 1964) ont alors donné un ensemble précis de conditions sous lesquelles ces approximations peuvent être faites, défini une notion de distance entre expériences, et exploré les conséquences de cette approximation.

Commençons avec deux expériences $\mathcal{E} = \{p_\theta : \theta \in \Theta\}$ et $\mathcal{F} = \{q_\theta : \theta \in \Theta\}$ ayant le même ensemble de paramètres Θ. Nous pouvons définir le *défaut* de $\mathcal{E}$ par rapport à $\mathcal{F}$ à partir d'idées de théorie de la décision. Considérons des fonctions de coût $c(\theta, \theta')$ bornées entre 0 et 1. Le défaut est défini comme l'infimum des ϵ tels que pour toute

semblable fonction de coût, pour tout estimateur $\hat{\theta}_{\mathcal{F}}$ dans la seconde expérience $\mathcal{F}$, il y a un estimateur $\hat{\theta}_{\mathcal{E}}$ dans la première expérience qui vérifie :

$$r_\theta(\hat{\theta}_{\mathcal{E}}) \leq r_\theta(\hat{\theta}_{\mathcal{F}}) + \epsilon \qquad \forall \theta \in \Theta,$$

où nous avons utilisé les notations précédentes (1.6) pour le risque d'un estimateur au point θ.

En d'autres termes, à ϵ près, nous pouvons faire aussi bien dans l'expérience $\mathcal{E}$ que dans l'expérience $\mathcal{F}$ pour toute question que l'on pourrait poser, quelle que soit la vraie valeur du paramètre. Le défaut est noté $\delta(\mathcal{E}, \mathcal{F})$.

Considérons maintenant un noyau de Markov T (donné par l'équation (1.5)) tel que $\|T(p_\theta) - q_\theta\|_1 = 2\epsilon$ pour tout $\theta \in \Theta$. Cela signifie approcher les lois de $\mathcal{F}$ par celles de $\mathcal{E}$. Alors pour toute fonction de coût c comme ci-dessus, et tout estimateur $\hat{\theta}_{\mathcal{F}}$, nous pouvons considérer l'estimateur $\hat{\theta}_{\mathcal{E}}$ défini par l'application de $\hat{\theta}_{\mathcal{F}}$ à la variable aléatoire de loi $T(p_\theta)$. Nous obtenons :

$$\begin{aligned} r_\theta(\hat{\theta}_{\mathcal{E}}) - r_\theta(\hat{\theta}_{\mathcal{F}}) &= \int c(\theta, \hat{\theta}(x)) T(p_\theta)(\mathrm{d}x) - \int c(\theta, \hat{\theta}(x)) q_\theta(\mathrm{d}x) \\ &\leq (\sup \; c(\theta, \theta')) \int (T(p_\theta) - q_\theta)^+(\mathrm{d}x) \\ &\leq 1 \times \|T(p_\theta) - q_\theta\|_1 / 2 \\ &\leq \epsilon. \end{aligned}$$

Si bien que le défaut est inférieur à ϵ. En fait, l'inverse est aussi vrai[17] Nous pouvons trouver un noyau de Markov qui transforme tout p_θ en q_θ à deux fois le défaut près. Nous pouvons donc écrire :

$$\delta(\mathcal{E}, \mathcal{F}) = \frac{1}{2} \inf_T \sup_\theta \|T(p_\theta) - q_\theta\|_1 .$$

Quand nous symétrisons le défaut, nous obtenons une distance, appelée *distance de Le Cam* $\Delta(\mathcal{E}, \mathcal{F})$. Nous pouvons alors considérer une suite d'expériences $\mathcal{E}_n = \{p_{n,\theta}\}$ qui converge vers l'expérience limite $\mathcal{F}$ pour cette distance. En d'autres termes, il y a deux familles de noyaux de Markov T_n et S_n tels que $\|T_n(p_{n,\theta}) - q_\theta\|_1 \to 0$ et $\|p_{n,\theta} - S_n(q_\theta)\|_1 \to 0$ uniformément en θ.

Cette convergence avec noyaux est appelée *convergence forte*. Il existe une autre convergence, dite *convergence faible*, basée sur les *rapports de vraisemblances*.

[17] À strictement parler, sans hypothèse de domination, il faut utiliser des objets légèrement plus généraux que les noyaux de Markov, appelés *transitions*. Les idées restent les mêmes.

Considérons l'expérience $\mathcal{E} = \{p_\theta\}$, sur un ensemble de paramètres fini Θ. Alors les rapports de vraisemblance sont les processus stochastiques $\Lambda_\Theta(\mathcal{E}) = \left\{\frac{p_\theta}{\sum_\theta p_\theta}\right\}_{\theta\in\Theta}$. Avec des ensembles de paramètres Θ infinis, nous disons que $\mathcal{E}_n$ converge faiblement vers $\mathcal{F}$ si la loi des processus $\Lambda_\mathcal{I}(\mathcal{E}_n)$ converge faiblement en loi vers $\Lambda_\mathcal{I}(\mathcal{F})$ pour tout sous-ensemble fini $\mathcal{I}$ de Θ.

La convergence faible se révèle équivalente à la convergence forte sur les ensembles de paramètres finis. Donc pour les ensembles dénombrables. Avec quelques conditions de régularité, cette équivalence peut être étendue aux ensembles Θ non dénombrables.

Pourquoi tant de définitions ? La définition basée sur les fonctions de coût exprime la vraie motivation : si $\mathcal{E}$ converge vers $\mathcal{F}$, nous pouvons asymptotiquement répondre aux questions concernant $\mathcal{E}_n$ de la même manière qu'à celles concernant $\mathcal{F}$. La convergence forte, avec les noyaux de Markov, donne un moyen direct de traduire les estimateurs d'une expérience à l'autre : nous transformons la première expérience, et appliquons l'estimateur de la seconde expérience. Cela nous assure d'obtenir les mêmes risques. D'un autre côté, exhiber ces noyaux de Markov pour des expériences réelles n'est pas forcément évident. La convergence des rapports de vraisemblance, par contre, est assez simple à établir. Et elle suffit à prouver l'existence des noyaux de Markov. Même si nous ne connaissons pas ces noyaux, et ne pouvons donc pas traduire directement les méthodes d'une expérience à l'autre, nous savons que les risques optimaux sont les mêmes pour tous les problèmes, que nous soyons dans un cadre bayésien ou minimax.

Les bénéfices pratiques de cette théorie sont au plus haut quand l'expérience limite est simple et bien comprise. Par exemple, les données indépendantes identiquement distribuées (i.i.d.) sont extrêmement fréquentes en statistiques, et peuvent être vues comme des variables aléatoires de loi $p_\theta^{\otimes n}$. Sous certaines conditions de régularité, elles convergent vers une *expérience de décalage gaussienne*, qui est en effet très bien connue.

Theorem 1.7.1. Normalité asymptotique locale(Le Cam, 1960)

Soit Θ un sous-ensemble ouvert de $\mathbb{R}^k$. Soit

$$\mathcal{E}_n = \left\{p_{\theta_0+h/\sqrt{n}}^{\otimes n} : h \in \mathbb{R}^k\right\}.$$

Alors si une famille $\{p_\theta\}$ est suffisamment régulière[18] *autour de 0, la suite d'expériences $\mathcal{E}_n$ converge faiblement vers une expérience de décalage gaussienne*

$$\mathcal{F} = \left\{\mathcal{N}(h, \mathcal{I}_{\theta_0}^{-1}) : h \in \mathbb{R}^k\right\},$$

[18]La bonne condition est la *différentiabilité en moyenne quadratique.* Deux fois différentiable en θ est plus que suffisant.

où $\mathcal{N}(h, \mathcal{I}_{\theta_0}^{-1})$ *la loi normale sur* $\mathbb{R}^k$, *de moyenne* h *et matrice de covariance* $\mathcal{I}_{\theta_0}^{-1}$ *l'information de Fisher* (1.13) *au point* θ_0.

Il y a deux différences avec un théorème de la limite centrale. Tout d'abord, la convergence vers la limite est uniforme[19] sur les ensembles qui ne croissent pas trop vite. Deuxièmement, la matrice de covariance est la même pour toutes les gaussiennes de l'expérience limite. Le nom «expérience de décalage» vient de cette observation : le paramètre est simplement la moyenne de la gaussienne.

Pourquoi est-ce appréciable ? Parce que nous connaissons la réponse à la plupart des questions statistiques habituelles pour les expériences de décalage gaussiennes. En particulier, nous connaissons un estimateur minimax optimal pour les fonctions de coût quadratiques, et nous pouvons traduire cet estimateur pour les expériences *i.i.d.* Cette observation constitue la manière habituelle de prouver l'optimalité de l'estimateur du maximum de vraisemblance dans ces conditions, par exemple. C'est là le théorème que nous aimerions généraliselser au cas quantique.

Le lecteur attentif aura noté que la fonction de coût quadratique n'est pas bornée en général, et que nous changeons l'échelle du paramètre h dans nos définitions de $\mathcal{E}_n$. Le théorème précédent est essentiellement local par nature. C'est déjà suffisant pour démontrer que les bornes de Cramér-Rao (1.15) ne peuvent pas être meilleures que dans l'expérience limite. Cependant, nous ne pouvons pas directement traduire la stratégie utilisée dans l'expérience de décalage gaussienne pour l'expérience initiale.

En pratique, nous contournons ces difficultés en utilisant une stratégie en deux temps : nous utilisons une proportion négligeable de notre n-échantillon pour effectuer une estimation préliminaire grossière, puis utilisons l'estimateur optimal fourni par la normalité asymptotique locale au point estimé. Nous devons pour finir vérifier que la partie non bornée de la fonction de coût contribue de manière négligeable à l'erreur globale.

Le Cam a bien davantage développé sa théorie de la convergence d'expériences, pour d'autres conditions de régularité, ce qui fournit des approximations différentes, et dans des cadres très généraux, basés sur les treillis de Banach. La largeur et la profondeur de cette théorie est suggérée par le volume de son livre de 1986.

1.7.2 Motivation

Dans une expérience physique, nous obtenons souvent en sortie n copies d'un état préparées de la même manière, et voulons apprendre certaines caractéristiques de cet état, typiquement l'identifier.

[19] Pour que ce point ait du sens, il faut employer une version exprimant la convergence forte.

Une normalité asymptotique locale quantique nous permettrait de répondre à toutes ces questions sur les expériences répétées en étudiant juste une expérience, que nous espérons plus simple. Par analogie avec le cas classique, nous nous attendons à obtenir une *expérience de décalage gaussienne quantique*, qui est en effet bien comprise.

Comme pour la convergence forte avec les noyaux de Markov, nous voudrions trouver des canaux qui transforment approximativement les états qui nous sont donnés en états gaussiens, et réciproquement.

Un inconvénient de cette stratégie est que l'équivalence tient quand nous avons le droit d'utiliser tout ce que permet la physique, c'est-à-dire les mesures et procédures collectives. Celles-ci peuvent être difficile à implémenter en pratique. De plus, nous ne pourrons pas étudier les mesures LOCC ou séparées directement par la normalité asymptotique locale quantique.

Un avantage d'exhiber les canaux est de permettre, pour peu que ceux-ci puissent être implémentés en laboratoire, d'appliquer en pratique les méthodes des expériences gaussiennes à l'expérience initiale.

1.7.3 Résultats antérieurs et liés

Le premier pas vers des résultats similaires au cas classique en quantique remonte à Dyson (1956), qui a remarqué que les fluctuations des composantes du spin total orthogonales à l'axe z de n spins $\uparrow$ purs se comportaient comme l'état fondamental d'un oscillateur harmonique quantique. De manière générale, les physiciens traitent les *états spin cohérents* (Holtz et Hanus, 1974) comme des gaussiennes. Kitagawa et Ueda (1993); Geremia et al. (2004) étendent cette situation aux types d'intrication qui ressemblent à des états compressés.

Ce genre de résultats peut être vu comme autant de théorèmes de la limite centrale quantique, dont la première preuve rigoureuse nous vient de Cushen et Hudson (1971). Hayashi (2003); Hayashi et Matsumoto (2004) ont prouvé une certaine régularité locale de ces limites et les ont utilisé pour donner la première méthode d'estimation optimale pour des pour un qubit totalement inconnu, ou des sous-modèles paramétriques, quand les mesures collectives sont autorisées.

Trouver et expliquer pareilles procédures optimales pour des problèmes variés est une grande motivation de la normalité asymptotique locale quantique. Le problème de l'estimation d'un qubit à partir de plusieurs copies a généré une immense bibliographie, comme il est élémentaire. Les études vont des mesures séparées aux mesures adaptatives et aux mesures collectives. Les références bayésiennes pour les états purs

comptent (Jones, 1994; Massar et Popescu, 1995; Latorre et al., 1998; Fisher et al., 2000; Hannemann et al., 2002b; Bagan et al., 2002; Embacher et Narnhofer, 2004; Bagan et al., 2005), et pour les états mélangés (Cirac et al., 1999; Vidal et al., 1999; Mack et al., 2000; Keyl et Werner, 2001; Bagan et al., 2004c; Zyczkowski et Sommers, 2005; Bagan et al., 2006). L'approche à point fixé est illustrée dans (Hayashi, 2002a; Gill et Massar, 2000; Barndorff-Nielsen et Gill, R., 2000; Matsumoto, 2002; Barndorff-Nielsen et al., 2003; Hayashi et Matsumoto, 2004). Les principaux points à retenir sont les suivants : pour les états purs et pas spécifiquement pour les qubits, les mesures séparées, qui sont faciles à implémenter, sont asymptotiquement aussi efficaces que les mesures collectives (Matsumoto, 2002) ; en revanche, pour des états mélanges généraux, nous pouvons obtenir une réelle accélération par des mesures collectives (Gill et Massar, 2000) ; les méthodes bayésiennes font en général usage de la théorie des groupes, et ne sont donc valides que pour les lois *a priori* uniformes ; Bagan et al. (2006) donnent une mesure optimale avec la fidélité comme fonction de coût, et prouvent que cette méthode est aussi asymptotiquement optimale au sens minimax.

Cependant cette dernière mesure covariante pourrait bien être très difficile à implémenter en pratique.

À un niveau plus fondamental, Petz et Jenčová (2006) ont caractérisé *l'exhaustivité quantique*. Dans le monde classique, une expérience $\mathcal{E}$ est exhaustive par rapport à une autre $\mathcal{F}$ si son défaut $\delta(\mathcal{E}, \mathcal{F})$ est nul. Petz et Jenčová ont caractérisé l'exhaustivité notamment par des canaux, équivalents des noyaux de Markov, et par l'usage de *cocycles de Connes*, qui peuvent être vus comme équivalents aux rapports de vraisemblance.

Sur la base de ce travail, Guţă et Jenčová (2007) ont prouvé la normalité asymptotique locale au sens de la convergence des cocycles de Connes, correspondant à la normalité asymptotique locale classique faible. Pour être précis, une expérience d'états définis sur un espace de dimension finie, dépendant de manière lisse du paramètre θ dans un ouvert $\Theta \in \mathbb{C}^d$ converge vers une expérience de décalage gaussienne quantique[20] de dimension d. Cette dernière expérience est une expérience où l'état est gaussien[21] sur l'espace de Fock $\mathcal{F}(\mathbb{C}^d)$, dont la fonction de Husimi (1.22) a θ pour moyenne et une matrice de covariance fixée.

Nous avons vu à la Section 1.1.3 que la mesure hétérodyne saturait la borne de Holevo (1.29) pour les expériences de décalage gaussiennes quantiques. Cependant, il n'y a pour l'heure aucun lien établi entre la normalité asymptotique locale quantique

[20] Pour être parfaitement exact, une partie de l'expérience quantique peut dégénérer en expérience de décalage gaussienne classique, ce qui correspond à déterminer des valeurs propres à vecteurs propres fixés

[21] Les états gaussiens peuvent être vus comme des mélanges gaussiens d'états cohérents.

faible et la théorie de la décision, aussi nous ne pouvons pas immédiatement utiliser ces bornes pour les expériences en dimension finie.

1.7.4 Contributions de la thèse

Avec Mădălin Guţă, j'ai établi la normalité asymptotique locale forte pour les qubits (2006). Spécifiquement, nous avons exhibé des familles de canaux T_n et S_n de $M_2(\mathbb{C})^{\otimes n}$ dans $\mathcal{T}(\mathcal{F}(\mathbb{C})) \otimes L^1(\mathbb{R})$, et réciproquement, qui envoient les matrices densité *i.i.d.* $\rho_{\theta_0+h/\sqrt{n}}^{\otimes n}$ près du produit d'une gaussienne classique à une dimension, correspondant aux valeurs propres, avec une gaussienne quantique à une dimension, correspondant aux vecteurs propres. La dérivation de ces canaux, qui repose sur la théorie des groupes, est lourdement inspirée par les travaux de Hayashi et Matsumoto (2004).

Nous avons prouvé que la convergence en norme d'opérateurs L^1 était uniforme pour $\|h\| \leq n^{1/4-\epsilon}$. Ce grand domaine de validité nous assure de pouvoir utiliser les stratégies en deux temps pour traduire les procédures de l'expérience limite pour l'expérience initiale.

Nous avons explicité cette stratégie en deux temps, avec Guţă et Bas Janssens (2008), en considérant une interaction en temps continu des qubits avec le champ électromagnétique. En utilisant les équations différentielles stochastiques quantiques (Hudson et Parthasarathy, 1984), nous avons prouvé que l'état du champ, ou lumière monochromatique, était la partie quantique de $T_n(\rho^{\otimes n})$ pour des temps plus longs que $\ln n$.

Nous pouvons dès lors utiliser la mesure hétérodyne sur cette lumière et obtenir une estimation optimale de la partie quantique. La partie classique reste dans les qubits, et peut être récupérée via une mesure du spin total. En pratique, cela sera réalisé par un autre couplage au champ et une mesure hétérodyne.

Cette stratégie d'estimation est asymptotiquement globalement optimale, à la fois dans un sens minimax et bayésien pour des lois *a priori* covariantes, aussi longtemps que nous sommes éloignés de l'état totalement mélangé. Nous pensons qu'elle doit être possible à implémenter en pratique.

Enfin Mădălin Guţă et moi avons généralisé cette construction des canaux aux qudits, pour toute dimension (2009). Là encore, le paramètre local h a le droit de grandir comme une faible fonction puissance, permettant la traduction des résultats de l'expérience limite pour l'expérience initiale, notamment en matière d'estimation de paramètres. Pour montrer que l'optimalité asymptotique de la méthode et calculer les risques explicites pour des pertes localement quadratiques, nous avons

également établi des théorèmes de représentation asymptotique et asymptotique minimax quantiques très généraux.

1.7.5 Perspectives

Des recherches supplémentaires pourraient suivre plusieurs voies :

Équivalence entre convergences d'expériences faibles et fortes

Les expériences limite sont les mêmes pour les convergences faibles et fortes. Le fragment de la normalité asymptotique locale classique le plus important à manquer encore à son équivalent quantique est la quasi équivalence de ces deux notions. Comme la convergence faible est relativement plus simple à prouver, nous en tirerions les mêmes bénéfices que dans le cas classique.

Éliminer les singularités de la normalité asymptotique locale quantique forte

Nos preuves reposent sur les représentations de groupe. Elles introduisent une singularité pour les valeurs propres égales, qui n'est pas importante au niveau des algèbres, utilisées pour la convergence faible. C'est pourquoi nous exigeons pour la convergence forte que les valeurs propres soient deux à deux différentes, bien que ce soit très probablement un artefact de la preuve.

Obtenir une preuve de la convergence forte en utilisant uniquement les C^*-algèbres semble difficile. Néanmoins, cela nous donnerait automatiquement un équivalent de la condition classique de «différentiabilité en moyenne quadratique».

Par contre, la singularité générée par les valeurs propres égales a une signification physique dans notre «implémentation pratique». Elle correspond à l'égalité des niveaux d'énergie pour les qubits. Comme la lumière monochromatique est donnée par les transitions entre deux niveaux d'énergie, ce couplage dégénère.

Traiter d'autres cas

D'autres directions de recherche recouvrent l'explicitation de la convergence d'expériences pour d'autres cas, non *i.i.d.*, tels les états spin cohérents, ou les chaînes de Markov quantiques.

Convergence d'expériences quantiques avec opérations locales

Un but plus ambitieux serait de définir une distance LOCC entre expériences, et la convergence correspondante. En d'autres termes définir une équivalence entre modèles quand nous n'avons le droit qu'aux méthodes LOCC, e pas à toutes les mesures collectives. La prolifération des scénarios utilisant les opérations LOCC, en information quantique en particulier, et le fait que ces méthodes sont plus faciles à implémenter en pratique, feraient tout le prix de cette théorie.

Implémentation pratique

Pour finir sur une idée plus réalisable, il devrait être assez simple de convertir «l'implémentation pratique» de la normalité asymptotique locale quantique du cas des qubits à celui des qudits.

Première partie

Divers Problèmes de Statistiques Quantiques

Chapitre 2

Model selection for quantum homodyne tomography

Ce chapitre dérive de l'article (Kahn).

Résumé : Nous nous intéressons à un problème de statistique non-paramétrique issu de la physique, et plus précisément à la tomographie quantique, c'est-à-dire la détermination de l'état quantique d'un mode de la lumière via une mesure homodyne. Nous appliquons plusieurs procédures de sélection de modèles : des estimateurs par projection pénalisés, où on peut utiliser soit des fonctions motif, soit des ondelettes, et l'estimateur du maximum de vraisemblance pénalisé. Dans chaque cas, nous obtenons une inégalité oracle. Nous prouvons également une vitesse de convergence polynomiale pour ce problème non-paramétrique, pour les estimateurs par projection. Nous appliquons ensuite des idées à la calibration d'un photocompteur, l'appareil dénombrant le nombre de photons dans un rayon lumineux. Le problème mathématique se réduit dans ce cas à un problème non-paramétrique à données manquantes. Nous obtenons à nouveau des inégalités oracle, qui nous assurent des vitesses de convergence d'autant meilleures que le photocompteur est bon.

2.1 Introduction

Quantum mechanics introduces intrinsic randomness in physics: the result of a measurement, or any macroscopic interaction, on a physical system is not deterministic.

Therefore, a host of statistical problems can stem from it. Some are (almost) specifically quantum, notably any question about which measurement yields the maximum information, or whether simultaneously measuring n samples is more efficient than measuring them sequentially (Gill, 2001). However, once we have chosen the measurement we carry out on our physical system, we are left with an entirely classical statistical problem. This chapter aims at applying model selection methods *à la* Birgé-Massart to one such instance, which is of interest both practical, as physicists use this measurement quite often (the underlying physical system is elementary; it is the particle with one degree of freedom), and mathematical, as it yields a non-parametric inverse problem with uncommon features.

Moreover, as this classical problem stemming from quantum mechanics could be seen as an easy introduction to the subject to classical statisticians, we have added more general notions on quantum statistics at the beginning of the appendix. The interested reader can get further acquaintance with these concepts through the textbooks by Helstrom (1976) and Holevo (1982) or the review article by Barndorff-Nielsen et al. (2003).

More precisely, the problem we are interested in is quantum homodyne tomography. As an aside, we apply the results we get to the calibration of a photocounter, using a quantum tomographer as a tool. The word "Homodyne" refers to the experimental technique used for this measurement, first implemented by Smithey et al. (1993), where the state of one mode of electromagnetic radiation, that is a pulse of laser light at a given frequency, is probed using a reference laser beam at the same ("homo") frequency. Respectively, "Tomography" is used because one of the physicists' favourite representations of the state, the Wigner function, can be recovered from the data by inverting a Radon transform.

Mathematically, our data are samples from a probability distribution p_ρ on $\mathbb{R}\times[0,\pi]$. From this data, we want to recover the "density operator" ρ of the system. This is the most common representation of the state, that is a mathematical object which encodes all the information about the system. Perfect knowledge of the state means knowing how the system will evolve and the probability distribution of the result of any measurement we might carry out on the system. These laws of evolution and measurement can be expressed naturally enough within the density operator framework (see Appendix). The density operator is a non-negative trace-one self-adjoint operator ρ on $L^2(\mathbb{R})$ (in our particular case). We know the linear transform $\mathcal{T}$ which takes ρ to p_ρ and can make it explicit in particular bases such as the Fock basis. We may also settle for the Wigner function W, another representation of the state. That is a two-dimensional real function with integral one, and p_ρ is the Radon transform of W.

The first reconstruction methods used the Wigner function as an intermediate representation: after collecting the data in histograms and smoothing, one inverted the

Radon transform to get an estimate of W. This smoothing, however, introduces hard-to-control bias. Pattern functions (bidual bases) for the entries of the density operator ρ were introduced by D'Ariano et al. (1994), yielding an unbiased estimator of those individual entries. They were later extended to allow for low noise in the measurement. Maximum likelihood procedures are used since the work of Banaszek et al. (1999). For both these estimators, we need an arbitrary cut-off of the density operator, so that the model is finite-dimensional. Artiles et al. (2005) have established consistency of these two estimators used with a sieve. Then, a sharp adaptive kernel estimator for the Wigner function was devised by Butucea et al. (2007), and this even if there is noise in the measurement (see subsection 2.3.6).

In this chapter, we devise penalized estimators that fulfill oracle-type inequalities among the L^2-projections on submodels, analyze the penalized maximum likelihood estimator and apply these estimators to the calibration of a photocounter. Hence, we provide automatic cut-offs for the estimators formerly mentioned. We can also cast in the L^2 projection framework wavelets estimators used for inverting the Radon transform on classical probability densities, to whom the Wigner function does not belong. We also have finer granularity for pattern functions, since we threshold them one by one, instead of keeping a whole submatrix. We get an explicit polynomial rate of convergence for this estimator. Notice that all our results are derived for finite samples (all the previous works considered only the asymptotic regime). We have mainly worked under the idealized hypothesis where there is no noise, however.

The appendix is not logically necessary for the chapter. We have inserted it for background and as an invitation to this field. It first features a general introduction to quantum statistics with a public of classical statisticians in mind. We then describe what quantum homodyne tomography precisely is. This latter subsection is largely based on the article by Butucea et al. (2007).

Section 2.2 formalizes the statistical problem at hand, with no need of the appendix, except the equations explicitly referred to therein.

Section 2.3 aims at devising a model selection procedure to choose between L^2-projection estimators. We first give general theorems (2.3.2 and 2.3.4) leading to oracle-type inequalities for hard-thresholding estimators. We then apply them to two bases. One is the Fock basis and the corresponding pattern functions physicists have used for a while. For it we also prove a polynomial convergence rate for any state with finite energy. The other is a wavelet basis for the Wigner function. We finish with a short subsection describing what changes are entailed by the presence of noise. Especially, we do not need to adapt our theorems if the noise is low enough, as long as we change the dual basis.

Section 2.4 similarly applies a classical theorem (2.4.2) to solve the question of which (size of) model is best to use a maximum likelihood estimator on.

Section 2.5 switches to the determination of a kind of measurement apparatus (and not any more on the state that is sent in) using a known state and this same tomographer that was studied in the previous sections. The law of our samples are then very similar and we apply the same type of techniques (penalized projection and maximum likelihood estimators). The fact that the POVM (mathematical modelling of a measurement) is a projective measurement (see Appendix) enables us to work with L^1-operator norm, however.

2.2 The mathematical problem

We now describe the mathematical problem at hand.

We are given n independent identically distributed random variables $Y_i = (X_i, \Phi_i)$ with density p_ρ on $[0, \pi) \times \mathbb{R}$.

This data is the result of a measurement on a physical system. Now the "state" of a system is described by a mathematical object, and there are two favourites for physical reasons: one is the *density operator* ρ, the other is the *Wigner function* W_ρ. We describe them below.

Therefore we are not actually interested in p_ρ, but rather in W_ρ or (maybe preferably) ρ. The probability distribution p_ρ of our samples can be retrieved if we know either ρ or W_ρ.

In other words we aim at estimating as precisely as possible ρ or W_ρ from the data $\{Y_i\}$. By " as precisely as possible", we mean that with a suitable notion of distance, we shall minimize $\mathbb{E}[d(\rho, \hat{\rho})]$. Our choice of distance will be partly dictated by mathematical tractability.

We now briefly explain what W_ρ and ρ stand for.

The Wigner function $W_\rho : \mathbb{R}^2 \to \mathbb{R}$ is the inverse Radon transform of p_ρ. In fact we would rather say that p_ρ is the Radon transform of W_ρ. Explicitly:

$$p_\rho(x, \phi) = \int_{-\infty}^{\infty} W(x \cos\phi + y \sin\phi, x \sin\phi - y \cos\phi) dy.$$

Figure 2.1 might be of some help. An important remark is that the Wigner function is not a probability density, but only a quasi-probability density: a function with integral 1, but that may be negative at places. However its Radon transform is a true probability density, as it is p_ρ.

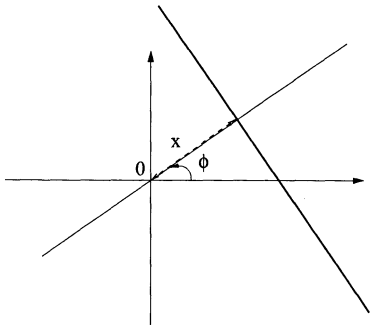

Figure 2.1: The value of p_ρ at (x, ϕ) is the integral of the Wigner function over the bold line

Retrieving W_ρ from P_ρ then amounts to inverting the Radon transform, hence the name of tomography: that is the same mathematical problem as with the brain imagery technique called Positron Emission Tomography.

As for ρ, this is a density operator on the Hilbert space $L^2(\mathbb{R})$, that is a *self-adjoint positive* operator with *trace* 1. We denote the set of such operators by $\mathcal{S}(L^2(\mathbb{R}))$. There is a linear transform $\mathbf{T}$ that takes ρ to p_ρ. We give it explicitly using a basis of $L^2(\mathbb{R})$ known as the *Fock basis*. This orthonormal basis, which has many nice physical properties, is defined by:

$$\psi_k(x) \;=\; H_k(x)e^{-x^2/2} \tag{2.1}$$

where H_k is the kth Hermite polynomial normalized such that $\|\psi_k\|_2 = 1$. The matrix entries of ρ in this basis are $\rho_{j,k} = \langle \psi_j, \rho\psi_k \rangle$. Then $\mathbf{T}$ can be written:

$$\begin{aligned} \mathbf{T} : \mathcal{S}(L^2(\mathbb{R})) &\longrightarrow L^1(\mathbb{R} \times [0, \pi]) \\ \rho &\mapsto \left(p_\rho : (x, \phi) \mapsto \sum_{j,k=0}^{\infty} \rho_{j,k}\psi_j(x)\psi_k(x)e^{-i(j-k)\phi} \right). \end{aligned}$$

Notice that as we have defined precisely the set of possible ρ, this mapping yields the set of possible p_ρ and W_ρ.

The relations between ρ, W_ρ and p_ρ are further detailed in subsection 2.A.2.

Anyhow we may now state our problem as consisting in inverting either the Radon transform or $\mathbf{T}$ from empirical data.

This is a classical problem of non-parametric statistics, that we want to treat non-asymptotically. We then take estimators based on a *model*, that is a subset of the operators on $L^2(\mathbb{R})$, or equivalently of the two-dimensional real functions. These

models are usually vector spaces, which may not be the domain of the object to be estimated. To choose a candidate within a given model, there are different methods, two of which we study, projection estimators and maximum likelihood estimators. Once we have a candidate within each model, we then use model selection methods to choose (almost) the best.

We first study projection estimators, for which the most convenient distance comes from the L^2 norm

$$\|\tau\|_2 = \sqrt{\sum |\lambda_i(\tau)|^2} = \sqrt{\sum_{j,k} |\tau_{j,k}|^2},$$

where the λ_i are the eigenvalues of τ, and the second equality holds for τ written in any orthonormal basis. Notice that there is an isometry (up to a constant) between the space of density operators with L^2-operator norm and the space of Wigner functions with L^2-Lebesgue norm, that is:

$$\|W_\rho - W_\tau\|_2^2 = \iint |W_\rho(q,p) - W_\tau(q,p)|^2 \, dp \, dq = \frac{1}{2\pi}\|\rho - \tau\|_2^2.$$

For maximum likelihood estimators, we have to make do with the weaker Hellinger distance (see later (2.24)) on $L^1\big(\mathbb{R} \times [0,\pi]\big)$, to which p_ρ belongs.

2.3 Projection estimators

In this section, which owes much to Massart's book (2006), we apply penalization procedures to projection estimators. The first subsection explains that we want to obtain oracle-type inequalities. In the second we obtain a general inequality where the left-hand side corresponds to an oracle inequality, and where the remainder term in the right-hand side depends on the penalty and on the large deviations of empirical coefficients. The two following subsections give two ways to choose the penalty term large enough for this remainder term to be small enough. In section 2.3.3 this penalty is deterministic. We design it and prove that it is a "good choice" by keeping Hoeffding's inequality in mind. In section 2.3.4, the penalty is random, and designed by taking Bernstein's inequality into account.

We next express these theorems in terms of two specific bases. For the Fock basis, we obtain polynomial worst-case convergence rates, using the structure of states. For a wavelet basis, we notice we obtain a usual estimator in classical tomography. We finish by saying what can be done if there is noise, that is (mainly) convolution of the law of the sample by a Gaussian. We multiply the Fourier transform of the dual basis with the inverse of the Fourier transform of the Gaussian, and as long as we still have well-defined functions, and we can re-use our theorems without changes.

2.3.1 Aim of model selection

Let's assume we are given a (countable) L^2-basis $(e_i)_{i\in\mathcal{I}}$ of a space in which $\mathcal{S}(L^2(\mathbb{R}))$ is included (typically $\mathcal{T}(L^2(\mathbb{R}))$, the trace-class operators on $L^2(\mathbb{R})$). We may then try and find the coefficients of ρ in this basis. The natural way to do so is to find a dual basis $(f_i)_{i\in\mathcal{I}}$ such that $\langle \mathbf{T}(e_i), f_j\rangle = \delta_{i,j}$ for all i and j. Then, if $\rho = \sum_i \rho_i e_i$ we get $\langle p_\rho, f_i\rangle = \rho_i$ for all i. And if the f_i are well enough behaved, then $\frac{1}{n}\sum_{k=1}^n f_i(X_k, \Phi_k) = \hat{\rho}_i$ tends to ρ_i by the law of large numbers.

Now if we took $\sum_i \hat{\rho}_i e_i$ as an estimator of ρ, we would have an infinite risk as the variance would be infinite. We must therefore restrict ourselves to models $m \in \mathcal{M}$, that is Vect $(e_i, i \in m)$, where m is a finite set, and $\mathcal{M}$ is a set of models (we might take $\mathcal{M}$ smaller than the set of all finite sets of $\mathbb{N}$).

We may then write the loss as

$$\|\hat{\rho}_m - \rho\|^2 = \sum_{i\notin m} |\rho_i|^2 + \sum_{i\in m} |\rho_i - \hat{\rho}_i|^2$$

where the first term is a bias (modelling error) and the second term is an estimation error. The risk would have this expression:

$$\mathbb{E}\left[\|\hat{\rho}_m - \rho\|^2\right] = \sum_{i\notin m} |\rho_i|^2 + \sum_{i\in m} \mathbb{E}\left[|\rho_i - \hat{\rho}_i|^2\right]$$

where the expectation is taken with respect to p_ρ, since $\hat{\rho}_i$ depends on the (X_k, Φ_k).

If we use an arbitrary model m, we probably do not strike a good balance between the bias term and the variance term. The whole point of penalisation is to have a data-driven procedure to choose the "best" model. We are aiming at choosing a model with (almost) the lowest error. We would dream of obtaining:

$$\hat{m} = \arg\inf_{m\in\mathcal{M}} \|\hat{\rho}_m - \rho\|^2 .$$

That is of course too ambitious. Instead, we shall obtain the following kind of bound, called an oracle inequality:

$$\mathbb{E}\left[\left\{\|\hat{\rho}_{\hat{m}} - \rho\|^2 - \left(C \inf_{m\in\mathcal{M}} \left(d^2(\rho, m) + \mathrm{pen}(m)\right)\right)\right\} \vee 0\right] \leq \epsilon_n \qquad (2.2)$$

where $d^2(\rho, m)$ is the bias of the model m, C is some constant, independent of ρ, pen(m) is a *penalty* associated to the model m (the bigger the model, the bigger the penalty) and ϵ_n depends only on n the number of observations, and goes to 0 when

n is going to infinity. We shall try to take the penalty of the order of the variance of the model.

Notice that we have given in (2.2) an unusual form of oracle inequality. These inequalities are more often written as

$$\mathbb{E}\left[\|\hat{\rho}_{\hat{m}} - \rho\|^2\right] \leq \left(C \inf_{m \in \mathcal{M}} \left(d^2(\rho, m) + \mathbb{E}\left[\mathrm{pen}(m)\right]\right)\right) + \epsilon_n.$$

Our form implies the latter.

The strategy is the following:

First, rewrite the projection estimators as *minimum contrast estimators*, that is minimizers of a function (called the *empirical contrast* function, and written γ_n), which is the same for all models. We also demand that, for any m, this empirical contrast function converges to a *contrast* function γ, the minimizer in m of which is the projection of ρ on m.

Second, find a penalty function that overestimates with high enough probability $(\gamma - \gamma_n)(\hat{\rho}_m)$ for all m simultaneously. Use of concentration inequalities is pivotal at this point.

Next section makes all this more explicit.

2.3.2 Risk bounds and choice of the penalty function

First we notice that the minimum of

$$\begin{aligned}\gamma(\tau) &= \|\tau\|^2 - 2\langle \tau, \rho \rangle \\ &= \|\rho - \tau\|^2 - \|\rho\|^2\end{aligned}$$

over a model m is attained at the projection of ρ on m. Moreover

$$\gamma_n(\tau) = \|\tau\|^2 - 2\sum_i \frac{1}{n}\sum_{k=1}^{n} \tau_i f_i(X_k, \Phi_k)$$

converges in probability to γ for any m (and all τ such that $\|\tau\| = 1$ simultaneously), as there is only a finite set of i such that $\tau_i \neq 0$ for $\tau \in m$.

Now the minimum of γ_n over m is attained by

$$\tau = \sum_{i \in m} \frac{1}{n}\sum_{k=1}^{n} f_i(X_k, \Phi_k) e_i.$$

So we have succeeded in writing projection estimators as minimum contrast estimators. We then define our final estimator by:

$$\hat{\rho}^{(n)} = \hat{\rho}_{\hat{m}}$$

with

$$\hat{m} = \arg\min_{m \in \mathcal{M}} \gamma_n(\hat{\rho}_m) + \mathrm{pen}_n(m)$$

where pen_n is a suitably chosen function depending on n, m and possibly the data.

We then get, for any m, for any $\tau_m \in m$,

$$\gamma_n(\hat{\rho}^{(n)}) + \mathrm{pen}_n(\hat{m}) \leq \gamma_n(\hat{\rho}_m) + \mathrm{pen}_n(m) \leq \gamma_n(\tau_m) + \mathrm{pen}_n(m). \tag{2.3}$$

What's more, for any m, for any $\tau_m \in m$,

$$\gamma_n(\tau_m) = \|\rho - \tau_m\|^2 - \|\rho\|^2 - 2\nu_n(\tau_m) \tag{2.4}$$

with

$$\begin{aligned} \nu_n(\tau) &= \langle \tau, \rho \rangle - \sum_i \sum_{k=1}^{n} \tau_i f_i(X_k, \Phi_k) \\ &= \sum_{i \in m} \tau_i(\rho_i - \hat{\rho}_i) + \sum_{i \notin m} \tau_i \rho_i. \end{aligned}$$

Putting together (2.3) and (2.4), we get, for all m and $\tau_m \in m$:

$$\left\|\hat{\rho}^{(n)} - \rho\right\|^2 \leq \|\tau_m - \rho\|^2 + 2\nu_n(\hat{\rho}^{(n)} - \tau_m) + \mathrm{pen}_n(m) - \mathrm{pen}_n(\hat{m}).$$

We then want to take penalties big enough to dominate the fluctuations ν_n. Some manipulations will make this expression more tractable. First we bound $\nu_n(\hat{\rho}^{(n)} - \tau_m)$ by $\left\|\hat{\rho}^{(n)} - \tau_m\right\| \chi_n(m \cup \hat{m})$, with

$$\chi_n(m) = \sup_{\substack{\tau \in m \\ \|\tau\|=1}} \nu_n(\tau).$$

Now the triangle inequality gives $\left\|\hat{\rho}^{(n)} - \tau_m\right\| \leq \left\|\hat{\rho}^{(n)} - \rho\right\| + \|\rho - \tau_m\|$, so that:

$$\begin{aligned} \left\|\hat{\rho}^{(n)} - \rho\right\|^2 \leq\ & \|\rho - \tau_m\|^2 + 2\chi_n(m \cup \hat{m}) \left\|\rho - \hat{\rho}^{(n)}\right\| \\ & + 2\chi_n(m \cup \hat{m}) \|\rho - \tau_m\| - \mathrm{pen}_n(\hat{m}) + \mathrm{pen}_n(m). \end{aligned}$$

For all $\alpha > 0$, the following holds:

$$2ab \leq \alpha a^2 + \alpha^{-1} b^2 \tag{2.5}$$

Using this twice, we get, for all $\epsilon > 0$:

$$\frac{\epsilon}{2+\epsilon} \left\| \rho - \hat{\rho}^{(n)} \right\|^2 \leq \left(1 + \frac{2}{\epsilon}\right) \|\rho - \tau_m\|^2 + (1+\epsilon)\chi_n^2(m \cup \hat{m}) - \mathrm{pen}_n(\hat{m}) + \mathrm{pen}_n(m).$$

Noticing that $\chi_n(m \cup \hat{m}) \leq \chi_n(m) + \chi_n(\hat{m})$ and putting our estimate of the error in the left-hand side:

$$\begin{aligned} &\frac{\epsilon}{2+\epsilon} \left\| \rho - \hat{\rho}^{(n)} \right\|^2 - \left\{ \left(1 + \frac{2}{\epsilon}\right) \|\rho - \tau_m\|^2 + 2\,\mathrm{pen}(m) \right\} \\ &\leq (1+\epsilon)(\chi_n^2(\hat{m}) + \chi_n^2(m)) - \mathrm{pen}_n(\hat{m}) - \mathrm{pen}_n(m). \end{aligned}$$

Now what we want to avoid is that our penalty is less than the fluctuations, so we separate this event and take its expectation:

$$\begin{aligned} \mathbb{E}\left[\left\{ \frac{\epsilon}{2+\epsilon} \left\| \rho - \hat{\rho}^{(n)} \right\|^2 - \left(\left(1 + \frac{2}{\epsilon}\right) \|\rho - \tau_m\|^2 + 2\,\mathrm{pen}_n(m) \right) \right\} \vee 0 \right] \\ \leq \mathbb{E}\left[\left\{ (1+\epsilon)(\chi_n^2(\hat{m}) + \chi_n^2(m)) - \mathrm{pen}(\hat{m}) - \mathrm{pen}(m) \right\} \vee 0 \right] \\ \leq 2\mathbb{E}\left[\sup_m \left\{ (1+\epsilon)\chi_n^2(m) - \mathrm{pen}(m) \right\} \vee 0 \right]. \end{aligned} \tag{2.6}$$

Thus stated, our problem is to take a penalty large enough to make the right-hand side negligible, that is vanishing like $1/n$.

We shall use this form of $\chi_n(m)$:

$$\chi_n(m) = \sup_{\substack{(\tau_i)_{i \in m} \\ \sum \tau_i^2 = 1}} \sum_{i \in m} \tau_i (\rho_i - \hat{\rho}_i) = \sqrt{\sum_{i \in m} |\rho_i - \hat{\rho}_i|^2}$$

so that

$$\chi_n(m)^2 = \sum_{i \in m} |\rho_i - \hat{\rho}_i|^2 = \sum_{i \in m} \left| \rho_i - \frac{1}{n} \sum_{k=1}^{n} f_i(X_k, \Phi_k) \right|^2. \tag{2.7}$$

2.3.3 Deterministic penalty

First we may try to craft a deterministic penalty.

We plan to use Hoeffding's inequality, recalling that $\hat{\rho}_i$ is a sum of independent variables:

Lemma 2.3.1. : Hoeffding's inequality (1964) *Let $X_1, \ldots, X_n$ be independent random variables, such that X_i takes his values in $[a_i, b_i]$ almost surely for all $i \leq n$. Then for any positive x,*

$$\mathbb{P}\left[\sum_{i=1}^{n}\left(X_i - \mathbb{E}\left[X_i\right]\right) \geq x\right] \leq \exp\left(-\frac{2x^2}{\sum_{i=1}^{n}(b_i - a_i)^2}\right).$$

We may also apply this inequality to $-X_i$ so as to get a very probable lower bound on the sum of X_i.

This is enough to prove:

Theorem 2.3.2. *Let ρ be a density operator. Assume that each f_i is bounded, where $(f_i)_{i\in\mathcal{I}}$ is the dual basis of $(e_i)_{i\in\mathcal{I}}$, as defined at the beginning of this section. Let $M_i = \sup_{(x,\phi)\in\mathbb{R}\times[0,\pi]} f_i(x,\phi) - \inf_{(x,\phi)\in\mathbb{R}\times[0,\pi]} f_i(x,\phi)$. Let $(x_i)_{i\in\mathcal{I}}$ be a family of positive real numbers such that $\sum_{i\in\mathcal{I}} \exp(-x_i) = \Sigma < \infty$. Let*

$$\mathrm{pen}_n(m) = \sum_{i\in\mathcal{I}_m}(1+\epsilon)\left(\ln(M_i) + \frac{x_i}{2}\right)\frac{M_i^2}{n}. \tag{2.8}$$

Then the penalized projection estimator satisfies:

$$\mathbb{E}\left[\frac{\epsilon}{2+\epsilon}\left\|\hat{\rho}^{(n)} - \rho\right\|^2\right] \leq \inf_{m\in\mathcal{M}}\left(1 + \frac{2}{\epsilon}\right)d^2(\rho, m) + 2\,\mathrm{pen}_n(m) + \frac{(1+\epsilon)\Sigma}{n}. \tag{2.9}$$

Remark: Here the penalty depends only on the subspace spanned by the model m. So it is the same whether $\mathcal{M}$ is small or large. The best we can do is then to take $\mathcal{M} = \mathcal{P}(\mathcal{I})$, that is to choose for every vector e_i whether to keep the estimated coordinate $\hat{\rho}_i$ or to put it to zero. In other words we get a hard-thresholding estimator:

$$\hat{\rho}^{(n)} = \sum_{i\in\mathcal{I}} \hat{\rho}_i \mathbf{1}_{|\hat{\rho}_i| > \alpha_i} \mathbf{e_i},$$

with

$$\alpha_i = \sqrt{(1+\epsilon)\left(\ln(M_i) + \frac{x_i}{2}\right)}\frac{M_i}{\sqrt{n}}. \tag{2.10}$$

Proof. Considering (2.6), we have only to bound appropriately

$$\mathbb{E}\left[\sup_m \left((1+\epsilon)\chi_n^2(m) - \mathrm{pen}(m)\right) \vee 0\right].$$

Now, by (2.7) and (2.8), both $\chi_n^2(m)$ and pen_m are a sum of terms over m. As the positive part of a sum is smaller than the sum of the positive parts, we obtain:

$$\begin{aligned}
&\mathbb{E}\left[\sup_m \left\{(1+\epsilon)\chi_n^2(m) - \mathrm{pen}(m)\right\} \vee 0\right] \\
&\leq \mathbb{E}\left[\sup_m \left\{\sum_{i\in m} \left((1+\epsilon)\left(\hat{\rho}_i - \rho_i\right)^2 - \alpha_i^2\right\} \vee 0\right)\right] \\
&= \sum_{i\in\mathcal{I}} \mathbb{E}\left[\left\{(1+\epsilon)\left(\frac{1}{n}\sum_{k=1}^{n} f_i(X_k,\Phi_k) - \rho_i\right)^2 - (1+\epsilon)\left(\ln(M_i) + \frac{x_i}{2}\right)\frac{M_i^2}{n}\right\} \vee 0\right].
\end{aligned}$$

Each of the expectations is evaluated using the following formula, valid for any positive function f:

$$\mathbb{E}[f] = \int_0^\infty \mathbb{P}[f(x) \geq y]\,\mathrm{d}y. \tag{2.11}$$

Remembering (2.10) we notice that the inequality

$$\left\{(1+\epsilon)\left(\frac{1}{n}\sum_{k=1}^{n} f_i(X_k,\Phi_k) - \rho_i\right)^2 - (1+\epsilon)\left(\ln(M_i) + \frac{x_i}{2}\right)\frac{M_i^2}{n}\right\} \vee 0 \;\geq\; y$$

is equivalent to

$$\left|\frac{1}{n}\sum_{k=1}^{n} f_i(X_k,\Phi_k) - \rho_i\right| \;\geq\; \sqrt{\frac{\alpha_i^2 + y}{1+\epsilon}}.$$

We may then conclude, using Hoeffding's inequality on the second line and the value

(2.10) of α_i on the fourth line:

$$\begin{aligned}
& \mathbb{E}\left[\sup_m \left\{(1+\epsilon)\chi_n^2(m) - \mathrm{pen}(m)\right\} \vee 0\right] \\
\leq & \sum_{i\in\mathcal{I}} \int_0^\infty \mathbb{P}\left[\left|\frac{1}{n}\sum_{k=1}^n f_i(X_k,\Phi_k) - \rho_i\right| \geq \sqrt{\frac{\alpha_i^2+y}{1+\epsilon}}\right] dy \\
= & \sum_{i\in\mathcal{I}} \int_0^\infty 2\exp\left(-\frac{2n(\alpha_i^2+y)}{(1+\epsilon)M_i^2}\right) \mathrm{d}y \\
= & \sum_{i\in\mathcal{I}} 2\exp\left(-\frac{2n\alpha_i^2}{(1+\epsilon)M_i^2}\right) \frac{(1+\epsilon)M_i^2}{2n} \\
= & \frac{1+\epsilon}{n}\sum_{i\in\mathcal{I}} \exp(-x_i) \\
= & \frac{(1+\epsilon)\Sigma}{n}.
\end{aligned}$$

□

2.3.4 Random penalty

The most obvious way to improve on Theorem 2.3.2 is to use sharper inequalities than Hoeffding's. Indeed the range of f_i might be much larger than its standard deviation, so that we gain much by using Bernstein's inequality:

Lemma 2.3.3. : Bernstein's inequality (1964) *Let $X_1, \ldots, X_n$ be independent, bounded, random variables. Then with*

$$M = \sup_i \|X_i\|_\infty, \qquad v = \sum_{i=1}^n \mathbb{E}\left[X_i^2\right],$$

for any positive x

$$\mathbb{P}\left[\sum_{i=1}^n (X_i - \mathbb{E}\left[X_i\right]) \geq \sqrt{2vx} + \frac{M}{3}x\right] \leq \exp(-x).$$

With this tool, we may devise a hard-thresholding estimator where the thresholds are data-dependent:

Theorem 2.3.4. *Let $(y_i)_{i\in\mathcal{I}}$ be positive numbers such that $\sum_{i\in\mathcal{I}} e^{-y_i} = \Sigma < \infty$. Let then*

$$x_i \;=\; 2\ln(\|f_i\|_\infty) + y_i.$$

Let the penalty be a sum of penalties over the vectors we admit in the model. That is, for any $\delta \in (0,1)$, for any $i \in \mathcal{I}$, define

$$\mathrm{pen}_n^i = \frac{1+\epsilon}{n}\left(\sqrt{\frac{2}{1-\delta}x_i\left(\mathbb{P}_n\left[f_i^2\right] + \frac{1}{n}\|f_i\|_\infty^2\left(\frac{1}{3}+\frac{1}{\delta}\right)x_i\right)} + \frac{\|f_i\|_\infty}{3\sqrt{n}}x_i\right)^2 \tag{2.12}$$

and the penalty of the model m:

$$\mathrm{pen}_n(m) \;=\; \sum_{i\in m}\mathrm{pen}_n^i\,.$$

Then there is a constant C such that:

$$\mathbb{E}\left[\left(\frac{\epsilon}{2+\epsilon}\left\|\hat{\rho}^{(n)}-\rho\right\|^2 - \left(\inf_{m\in\mathcal{M}_n}\left(1+\frac{2}{\epsilon}\right)d^2(\rho,m) + 2\,\mathrm{pen}_n(m)\right)\right)\vee 0\right] \;\le\; \frac{C\Sigma}{n}$$

where $\mathcal{M}_n$ is the set of models m for which $i \in m \rightarrow x_i \le n$.

Remark: As with the deterministic penalty, we end up with a hard-thresholding estimator. Morally, that is, forgetting all the small δ whose origin is technical, the threshold is

$$\sqrt{\frac{2\mathbb{P}_n\left[f_i^2\right]\ln\|f_i\|_\infty^2}{n}}.$$

Proof. Once again we have to dominate the right-hand side of (2.6). We first subtract and add, inside that expression, what could be seen as a target for the penalty. Writing

$$M_i = \|f_i\|_\infty\,, \qquad v_i = n\mathbb{E}\left[f_i^2\right], \qquad \alpha_i = \sqrt{2v_ix_i} + \frac{M_i}{3}x_i \tag{2.13}$$

we have

$$\begin{aligned}&\mathbb{E}\left[\sup_m\left((1+\epsilon)\chi_n^2(m) - \mathrm{pen}(m)\right)\vee 0\right]\\ &\le \mathbb{E}\left[\sup_m(1+\epsilon)\left(\chi_n^2(m) - \sum_{i\in m}\frac{1}{n^2}\alpha_i^2\right)\vee 0\right] + \mathbb{E}\left[\left(\sum_{i\in m}\frac{1+\epsilon}{n^2}\alpha_i^2 - \mathrm{pen}(m)\right)\vee 0\right].\end{aligned} \tag{2.14}$$

Using (2.7), we bound the first term as a sum of expectations.

$$\begin{aligned}&\mathbb{E}\left[\sup_m(1+\epsilon)\left(\chi_n^2(m)-\sum_{i\in m}\frac{1}{n^2}\alpha_i^2\right)\vee 0\right]\\ \leq&(1+\epsilon)\sum_{i\in m}\mathbb{E}\left[\left(\left|\rho_i-\frac{1}{n}\sum_{k=1}^n f_i(X_k,\Phi_k)\right|^2-\frac{1}{n^2}\alpha_i^2\right)\vee 0\right].\end{aligned}$$

We now bound each of these expectations using (2.11).

$$\mathbb{E}\left[\left(\left|\rho_i-\frac{1}{n}\sum_{k=1}^n f_i(X_k,\Phi_k)\right|^2-\frac{1}{n^2}\alpha_i^2\right)\vee 0\right] \tag{2.15}$$

$$=\int_0^\infty \mathbb{P}\left[\left|\rho_i-\frac{1}{n}\sum_{k=1}^n f_i(X_k,\Phi_k)\right|\geq\sqrt{y+\frac{\alpha_i^2}{n^2}}\right]\mathrm{d}y. \tag{2.16}$$

We change variables in the integral, choosing ξ defined by:

$$\sqrt{y+\frac{\alpha_i^2}{n^2}} = \frac{\sqrt{2v_i\xi}+\frac{M_i}{3}\xi}{n^2}. \tag{2.17}$$

Using Bernstein's inequality, the integrand in (2.16) is upper bounded by $2\exp(-\xi)$. Given the value of α_i (2.13), the range of the integral is now from x_i to ∞. Finally, taking the square on both sides of (2.17), then using (2.5), we get:

$$\begin{aligned}\mathrm{d}y &= 2\frac{\sqrt{2v_i\xi}+\frac{M_i}{3}\xi}{n^2}\left(\frac{M_i}{3}+\frac{\sqrt{2v_i}}{2\sqrt{\xi}}\right)\mathrm{d}\xi\\ &= \frac{2}{n^2}\left(v_i+\frac{M_i^2}{9}\xi+\frac{M_i}{2}\sqrt{2v_i}\sqrt{x}\right)\mathrm{d}\xi\\ &\leq \frac{2}{n^2}\left(2v_i+\frac{11M_i^2}{18}\xi\right)\mathrm{d}\xi.\end{aligned}$$

Hence

$$\begin{aligned}\mathbb{E}\left[\left(\left|\rho_i-\frac{1}{n}\sum_{k=1}^n f_i(X_k,\Phi_k)\right|^2-\frac{1}{n^2}\alpha_i^2\right)\vee 0\right] &\leq \frac{4}{n^2}\int_{x_i}^\infty \exp(-\xi)\left(2v_i+\frac{11M_i^2}{18}\xi\right)\mathrm{d}\xi\\ &= \frac{4}{n^2}\left(2v_i+\frac{11M_i^2}{18}(1+x_i)\right)\exp(-x_i).\end{aligned} \tag{2.18}$$

Let us now look over the second term of (2.14). We notice, through (2.12) and (2.13), that this term is of the form:

$$\frac{1+\epsilon}{n^2}\sum_{i\in m}\mathbb{E}\left[\left(\left(a_i+\frac{M_i x_i}{3}\right)^2-\left(b_i+\frac{M_i x_i}{3}\right)^2\right)\vee 0\right]\leq\frac{1+\epsilon}{n^2}\sum_{i\in m}\mathbb{E}\left[2\left(a_i^2-b_i^2\right)\vee 0\right],$$

with

$$a_i^2-b_i^2 \;=\; 2v_i x_i-\frac{2}{1-\delta}\left(n\mathbb{P}_n\left[f_i^2\right]x_i+M_i^2\left(\frac{1}{3}+\frac{1}{\delta}\right)x_i^2\right).$$

Using again (2.11), we end up with:

$$\begin{aligned}&\mathbb{E}\left[\left(\sum_{i\in m}\frac{1+\epsilon}{n^2}\alpha_i^2-\mathrm{pen}(m)\right)\vee 0\right]\\&\leq\frac{1+\epsilon}{n^2}\sum_{i\in m}\frac{2}{1-\delta}x_i\int_0^\infty\mathbb{P}\left[(1-\delta)v_i-\left(n\mathbb{P}_n\left[f_i^2\right]+M_i^2\left(\frac{1}{3}+\frac{1}{\delta}\right)x_i\right)\geq y\right]\mathrm{d}y.\end{aligned}\tag{2.19}$$

We can again make use of Bernstein's inequality:

$$\mathbb{P}\left[v_i-\sum_{k=1}^n f_i^2(X_k,\Phi_k)\geq\sqrt{2n\mathbb{E}\left[f_i^4\right]\xi}+\frac{\|f_i^2\|_\infty\xi}{3}\right]\;\leq\;\exp(-\xi).$$

Noticing that f_i^2 is non-negative everywhere, so that $\mathbb{E}\left[f_i^4\right]\leq\mathbb{E}\left[f_i^2\right]\|f_i^2\|_\infty$, and using (2.5), we get:

$$\mathbb{P}\left[(1-\delta)v_i\geq n\mathbb{P}_n\left[f_i^2\right]+M_i^2\left(\frac{1}{3}+\frac{1}{\delta}\right)\xi\right]\;\leq\;\exp(-\xi).$$

Recalling (2.19), we get

$$\begin{aligned}&\int_0^\infty\mathbb{P}\left[(1-\delta)v_i-\left(n\mathbb{P}_n\left[f_i^2\right]+M_i^2\left(\frac{1}{3}+\frac{1}{\delta}\right)x_i\right)\geq y\right]\mathrm{d}y\\&=\int_0^\infty\exp\left(-x_i-\frac{y}{M_i^2\left(\frac{1}{3}+\frac{1}{\delta}\right)}\right)\mathrm{d}y\\&=\exp(-x_i)M_i^2\left(\frac{1}{3}+\frac{1}{\delta}\right)\exp\left(-\frac{x_i}{M_i^2\left(\frac{1}{3}+\frac{1}{\delta}\right)}\right)\\&\leq\exp(-y_i)\left(\frac{1}{3}+\frac{1}{\delta}\right).\end{aligned}$$

With that and (2.18), we are left with:

$$\mathbb{E}\left[\sup_m \left\{(1+\epsilon)\chi_n^2(m) - \mathrm{pen}(m)\right\} \vee 0\right] \leq \frac{C}{n^2}\sum_{i\in\mathcal{I}} e^{-x_i}(v_i + M_i^2(1+x_i)) + x_i e^{-y_i}.$$

Replacing x_i with its value, and overestimating v_i by nM_i^2 we obtain (under the condition that $2\ln M_i + y_i \leq n$):

$$\mathbb{E}\left[\sup_m \left\{(1+\epsilon)\chi_n^2(m) - \mathrm{pen}(m)\right\} \vee 0\right] \leq C\left(\frac{\Sigma}{n} + \frac{\Sigma}{n^2}\right).$$

□

Remark: The logarithmic factor in the penalty (that would not be here if we took only the variance) comes from the multitude of models allowed by a hard-thresholding estimator. By selecting fewer models (for example the square matrices obtained by truncating the density operator) and using a random penalty, we can get rid of this term. However, crafting the penalty requires much more work and more powerful inequalities (Talagrand's). An interested reader may have a look at the section 3.4 of the author's master thesis (2004).

2.3.5 Applications with two bases

We now give two bases that are reasonable when applying these theorems. As can be seen from (2.2), a good basis should approximate well any density operator (so that the bias term gets low fast when m is big), with dual vectors having a low variance. With the first of the two bases, we have this interesting phenomenon that we obtain a polynomial convergence rate under the mere physical hypothesis that the state has finite energy.

Photon basis

Here we shall take as our $(e_i)_{i\in\mathcal{I}}$ a slight variation of the matrix entries of our density operator with respect to the Fock basis (2.1).

More precisely, we worked in the previous subsections with real coefficients. To apply Theorems 2.3.4 and 2.3.2, we then need to parametrize ρ with real coefficients. The matrix entries are *a priori* complex. However, using the fact that ρ is self-adjoint, we may separate the real and imaginary parts.

We use a double index for i and define the orthonormal basis, denoting by $E_{j,k}$ the null matrix except for a 1 in case (j,k):

$$e_{j,k} = \begin{cases} \frac{1}{\sqrt{2}}(E_{j,k} + E_{k,j}) & \text{if } j < k \\ \frac{i}{\sqrt{2}}(E_{k,j} - E_{j,k}) & \text{if } k < j \\ E_{j,j} & \text{if} j = k \end{cases}.$$

Then, using a tilde to distinguish it from the matrix entries, with $\tilde{\rho}_{j,k} = \langle \rho, e_{j,k} \rangle$,we have

$$\langle \psi_j, \rho \psi_k \rangle = \begin{cases} \frac{1}{\sqrt{2}}(\tilde{\rho}_{j,k} + i\tilde{\rho}_{k,j}) & \text{if } j < k \\ \frac{1}{\sqrt{2}}(\tilde{\rho}_{k,j} - i\tilde{\rho}_{j,k}) & \text{if } j > k \\ \tilde{\rho}_{j,j} & \text{if } j = k. \end{cases}$$

The associated $\tilde{f}_{j,k}$ are well-known. They are a slight variation of the usual "pattern functions" (see Appendix 2.A.2, and (2.38) therein), the behaviour of which may be found in (Artiles et al., 2005). Notably, we know that:

$$\sum_{j,k=0}^{N} \|f_{j,k}\|_\infty^2 \leq CN^{7/3}. \tag{2.20}$$

As the upper bounds on the supremum of $\tilde{f}_{j,k}$ may not be sharp, the best way to apply the above theorems (especially Theorem (2.3.2)) would probably be to tabulate these maxima for the (j,k) we plan to use.

The interest of this basis is that it is a priori adapted to the structure of our problem: if we have a bound on the energy (let's say it is lower than $H + \frac{1}{2}$), we get worst-case estimates on the convergence speed with the deterministic penalty: indeed, the energy of a state ρ may be written $\frac{1}{2} + \sum_j j\rho_{j,j}$, so that

$$\sum_{j \geq N} \tilde{\rho}_{j,j} \leq \frac{H}{N}.$$

Moreover, by positivity of the operator,

$$\tilde{\rho}_{j,k}^2 + \tilde{\rho}_{k,j}^2 \leq \tilde{\rho}_{j,j}\tilde{\rho}_{k,k}.$$

If we look at the models N such that $\mathcal{I}_N = \{(j,k) : j < N, k < N\}$, we can get:

$$\begin{aligned} d^2(\rho, N) &\leq \sum_{j,k=0}^{\infty} \tilde{\rho}_{j,k}^2 - \sum_{j,k=0}^{N} \tilde{\rho}_{j,k}^2 \\ &\leq (\sum_{j\geq 0} \tilde{\rho}_{j,j})^2 - (\sum_{j=0}^{N} \tilde{\rho}_{j,j})^2 \\ &\leq 1 - (1 - \frac{H}{N})^2 \\ &\leq \frac{2H}{N} \end{aligned}$$

where we have used that the density operator has trace one.

We substitute in (2.9) and get:

$$\mathbb{E}\left[\left\|\hat{\rho}^{(n)} - \rho\right\|^2\right] \leq C\left(\frac{H}{N} + \mathrm{pen}_n(N) + \frac{1}{n}\right).$$

Now, using the bounds on infinite norms (2.20), the penalty is less than:

$$\mathrm{pen}_n(N) = C\frac{N^{7/3}\ln(N)}{n}.$$

Optimizing in N ($N = C(Hn)^{3/10}$), we get

$$\mathbb{E}\left[\left\|\hat{\rho}^{(n)} - \rho\right\|^2\right] \leq CH^{7/10}\ln(H)n^{-3/10}\ln(n). \tag{2.21}$$

This estimate holds true for any state and is non-asymptotic. It is generally rather pessimistic, though. For many classical states, such as squeezed states or thermal states, $\rho_{j,j} \equiv A\exp(-B/n)$, the same calculation yields a rate for the square of the L^2-distance as $n^{-1}\ln(n)^\beta$ for some β. In such a case, the penalized estimator automatically converges at this latter rate.

Wavelets

Another try could be to use functions known for their good approximations properties. To this end we look at the Wigner function and write it in a wavelet basis.

Recall that wavelets on $\mathbb{R}$ are an orthonormal basis such that all functions are scaled translations of a same function, the mother wavelet. In multiscale analysis,

we use a countable basis $\psi_{j,k} : x \mapsto 2^{j/2}\psi_{0,0}(2^j x + k)$, for j and k integers. Let $\mathcal{V}_i = \{\psi_{j,k} : j \leq i\}$. There is a ϕ, called father wavelet, such that the $\phi_k(x) = \phi(x+k)$ for $k \in \mathbb{Z}$ are a basis of the vector space generated by all the wavelets of larger or equal scale, that is $\mathcal{V}_0$. We may choose them with compact support, or localized both in frequency and position. So they harvest local information and can fetch this whatever the regularity of the function to be approximated, as they exist at several scales.

From a one-dimensional wavelet basis $\psi_{j,k} : x \mapsto 2^{j/2}\psi_{0,0}(2^j x+k)$, C^3 and zero mean, with a father wavelet $\phi_{j,k}$, also C^3, we shall make a tensor product basis on $L^2(\mathbb{R}^2)$: let $I = (j, k, \epsilon)$ be indices, with j integer (scale), $k = (k_x, k_y) \in \mathbb{Z}^2$ (position), and $\epsilon \in 0, 1, 2, 3$. Let

$$\Psi_I(x, y) = \begin{cases} \phi_{j,k}(x)\phi_{j,k}(y) & \text{if } \epsilon = 0 \\ \phi_{j,k}(x)\psi_{j,k}(y) & \text{if } \epsilon = 1 \\ \psi_{j,k}(x)\phi_{j,k}(y) & \text{if } \epsilon = 2 \\ \psi_{j,k}(x)\psi_{j,k}(y) & \text{if } \epsilon = 3 \end{cases}$$

We may then define a multiscale analysis from the one-dimensional one (written $\mathcal{V}, \mathcal{W}$): $V_0 = \overline{\mathcal{V}_0 \otimes \mathcal{V}_0}$ and for all $j \in \mathbb{Z}$, $V_{j+1} = V_j \oplus W_j$, so that $W_{j+1} = \overline{\mathcal{V}_j \otimes \mathcal{W}_j} \oplus \overline{\mathcal{W}_j \otimes \mathcal{V}_j} \oplus \overline{\mathcal{V}_j \otimes \mathcal{W}_j}$.

For any j, $V_j \cup \bigcup_{k \geq j} W_k$ is then an orthonormal basis of $L^2(\mathbb{R}^2)$. We hereafter choose our models as subspaces spanned by finite subsets of the basis vectors for well-chosen j.

It can be shown that:

$$\gamma_I(x, \phi) = \frac{1}{4\pi} \int_{-\infty}^{\infty} |u| \, \hat{\Psi}_I(u \cos\phi, u \sin\phi) e^{ixu} du$$

fulfills this property:

$$[\gamma_I, Kf] = \langle \Psi_I, f \rangle.$$

Noticing that

$$\gamma_I(x, \phi) = 2^j \gamma_{0,0,\epsilon}(2^j x - k_x \cos\phi - k_y \sin\phi, \phi),$$

we see that these functions have the same dilation properties as wavelets, and they are "translated" in a way that depends on ϕ, through sinusoids. Their normalizations, though, explode with j; this derives from inverting the Radon transform being an ill-posed problem.

We can now apply Theorem 2.3.4. Before doing so, though, we restrict ourselves to a finite subdomain of $\mathbb{R}^2$, which we denote $\mathcal{D}$, and put the Wigner function to zero

outside this domain, that we should choose big enough to ensure this does not cost too much.

Then, $\mathcal{M}$ is the set of all models characterized by

$$\begin{aligned} m &= \{(j_1,k,0) : 2^{j_1}k \in \mathcal{D}\} \\ &\quad \cup \{(j,k,\epsilon) : (j,k,\epsilon) \in \mathcal{I}'_m \subset \{(j,k,\epsilon) : \epsilon = 1;2;3, j_1 < j < j_0, 2^j k \in \mathcal{D}\}\}. \end{aligned}$$

To have good approximating properties, we choose $2^{j_1} \equiv n^{1/7}$ and $2^{j_0} \equiv \frac{n}{(\ln n)^2}$. The projection estimator within a model is then:

$$\hat{f} = \sum_{I \in m} \alpha_I \Psi_I$$

with

$$\alpha_I = \frac{1}{n}\sum_{i=1}^{n} \gamma_I(X_i, \Phi_{i/}.$$

Denoting $B_\epsilon = \|\gamma_{0,0,\epsilon}\|_\infty$, the translation of Theorem 2.3.4 gives (notice that applying (2.3.2) would be awkward, as the variance of γ_I is like 2^j whereas its maximum is like 2^{2j}):

Theorem 2.3.5. *Let y_I be such that $\sum_I \exp(-y_I) = \Sigma \leq \infty$. For example $y_I = j$. Let then:*

$$x_I = 2(j + \ln(B_\epsilon)) + y_I.$$

We choose an $\alpha \in (0,1)$ and the penalty (and restraining ourselves to the m such that $I \in m \rightarrow x_I \leq n$):

$$\text{pen}(m) = \frac{1+\epsilon'}{n} \sum_{I \in \mathcal{M}} 2\left(\sqrt{\frac{2}{1-\alpha}x_I\left(\mathbb{P}_n[\gamma_I^2] + \frac{1}{n}2^{2j}B_\epsilon^2\left(\frac{1}{3}+\frac{1}{\alpha}\right)x_I\right)} + \frac{2^j B_\epsilon}{3\sqrt{n}}x_I\right)^2.$$

Then there is a constant C such that:

$$\mathbb{E}\left[\left\{\frac{\epsilon}{2+\epsilon}\left\|\rho - \hat{\rho}^{(n)}\right\|^2 - \left(\inf_{m \in \mathcal{M}}\left(1+\frac{2}{\epsilon}\right)d^2(\rho,m) + 2\,\text{pen}_n(m)\right)\right\} \vee 0\right] \leq \frac{C\Sigma}{n} + C\frac{1}{n}2^{2j_1}. \tag{2.22}$$

Proof. First it's easily checked that $x_I = 2\ln(\|\gamma_I\|_\infty) + y_I$. Second $\sum_I \exp(-j) = C\sum_j 2^j \exp(-j) < \infty$ implies that $y_I = j$ does indeed the work, as there are at most $C2^j$ wavelets at scale j whose support meet $\mathcal{D}$.

The last term is the variance of $\hat{a}_{j_1,k,0}$, corresponding to the vectors that are in every model.:

$$\begin{aligned}\frac{1}{n}\mathbb{V}\left[\sum_{2^{j_1}k\in\mathcal{D}}\gamma_{j_1,k,0}\right] &\leq \frac{1}{n}\mathbb{E}\left[\sum_{2^{j_1}k\in\mathcal{D}}\gamma^2_{j_1,k,0}\right]\\ &\leq \frac{1}{n}\sum_{2^{j_1}k\in\mathcal{D}}\int_{\mathbb{R}\times[0,\pi]}\gamma^2_{j_1,k,0}(x,\phi)dx\frac{d\phi}{\pi}p_\rho(x,\phi)\\ &= \frac{1}{n}\sum_{2^{j_1}k\in\mathcal{D}}\int_{\mathbb{R}}\gamma^2_{j_1,k,0}(x,0)\int_0^\pi p_\rho(x-k_x\cos\phi-k_y\sin\phi,\phi)dx\frac{d\phi}{\pi}\\ &= C\frac{1}{n}2^{2j_1}\end{aligned}$$

where we have used that for all x and k, $\int_0^\pi p_\rho(x-k_x\cos\phi-k_y\sin\phi,\phi)\frac{d\phi}{\pi}$ is less than a constant about 1.086. Indeed, the translation of a Wigner function is still the Wigner function of a state, so that we may take $k=0$. Then

$$\int_0^\pi p_\rho(x-k_x\cos\phi-k_y\sin\phi,\phi)\frac{d\phi}{\pi} \leq \sup_{i,x}|\psi_i(x)|^2$$

and the upper bound on this supremum is due to Cramér (Erdélyi, 1953, 10.18.19). □

Remarks: As the variance of γ_I goes like 2^j the threshold might be seen as $C2^{j/2}\sqrt{\frac{j}{n}}$. This yields the wavelets estimator studied by Cavalier et Koo (2002), for a general Radon transform on usual (non-negative) probability densities (*i.e.* not on Wigner functions).

The role of the approximation speed is apparent in (2.22). Articles like that by Cavalier et Koo show that this strategy is asymptotically (quasi)-optimal for approximating a function in a Besov ball. However, this is no proof of the efficiency in our case, as the set of Wigner functions is not a Besov ball. There is still some work in approximation theory needed there. In particular, we do not know if a statement similar to (2.21) can be proven.

Finally, notice that we may combine projection estimators: as the contrast function is the same for any basis we are working with, keeping the same penalizations, we could find an estimator that is almost the best among those built with the photon basis and those with the wavelet basis simultaneously (just add a $\ln(2)$ to Σ). In other words, we do not have to choose beforehand which basis we use. Moreover an estimator allowing for the two bases would satisfy (2.21).

2.3.6 Noisy observations

The situation we have studied was very idealized: we did not take any noise into account. In practice, a number of photons fail to be detected. These losses may be quantified by one single coefficient η between 0 (no detection) and 1 (ideal case). We suppose it to be known.

There are several methods to recover the state from noisy observations. One consists in calculating the density matrix as if there was no noise, and then apply the Bernoulli transformation with factor η^{-1}. We can also use modified pattern functions (D'Ariano et al., 1995). Or we can approximate the Wigner function with a kernel estimator that performs both the inverse Radon transform and the deconvolution (Butucea et al., 2007). The former two methods fail if the observations are too noisy ($\eta \leq \frac{1}{2}$), but the latter is asymptotically optimal for all η over wide classes of Wigner functions.

This noise can be seen as a convolution of the result (X, Φ) with a Gaussian of variance depending on η:

$$p_\rho^\eta(y,\phi) = \frac{1}{\sqrt{\pi(1-\eta)}} \int_{-\infty}^{\infty} p_\rho(x,\phi) \exp\left(-\frac{\eta}{1-\eta}\left(x-\eta^{-1/2}y\right)^2\right) dx$$

or equivalently in terms of generating functions

$$\int p_\rho^\eta(x,\phi)e^{irx}dx = e^{-\frac{1-\eta}{4\eta}r^2} \int p_\rho(x,\phi)e^{irx}dx.$$

We can use the methods described above and then use the Bernoulli transform. For free, we may also use the modified pattern functions $f_{j,k}^\eta$ knowing $f_{j,k}$. Explicitly we see that the dual basis of the matrix entry $\rho_{j,k}$ becomes:

$$f_{j,k}^\eta(x,\phi) = \frac{1}{2\pi} \int dr e^{\frac{1-\eta}{4\eta}r^2} \int dy f_{j,k}(y,\phi)e^{iry}.$$

The reason why one needs $\eta > \frac{1}{2}$ is for this Fourier transform to be well defined.

And we can again apply Theorems 2.3.2 and 2.3.4 with the dual basis $\tilde{f}_{j,k}^\eta$.

Obtaining results with high noise $\eta \leq \frac{1}{2}$ is harder. We would need to introduce a cut-off h within the inverse Fourier transform (and therefore a bias). Using the same h as in (Butucea et al., 2007) would ensure this bias $b(\rho, h)$ is asymptotically reasonable. We could then reuse Theorems 2.3.2 and 2.3.4 to have an "almost best" approximation of $\rho + b(\rho, h)$ within a set of models, for finite samples. Careful examination would then be required to check the variance (or the penalties) go to

0 as n and $h(n)$ go to infinity. Moreover, we would need to translate conditions on the Wigner function into conditions on the density operator to see whether we can reproduce the asymptotic optimality results of Butucea *et al.* with model selection in the Fock basis (or any other basis chosen and studied *a priori*).

2.4 Maximum likelihood estimator

Projection estimators are not devoid of defects. Notably, the variance of empirical coefficients might be high, the convergence therefore rather slow, and the estimator is not a true density matrix. Especially, the trace is probably not one, though this could be fixed easily enough. We can diagonalize the estimated density matrix, replace the negative eigenvalues with 0, and divide by the trace.

Anyhow, there are other types of estimator that automatically yield density matrices. One such estimator is the maximum likelihood estimator, which selects the nearest point of the empirical probability measure in a given model for the Kullback-Leibler distance (which is not a true distance as it is not symmetric). Recall that the Kullback-Leibler distance between two probability measures is:

$$K(p,q) = \int \ln\left(\frac{p(x)}{q(x)}\right) p(x)dx.$$

In other words, the maximum likelihood estimator is

$$\arg\min_{\tau\in\mathcal{Q}} \sum_{l=1}^{n} -\ln p_\tau(X_l,\Phi_l)$$

where $\mathcal{Q}$ is any set of density operators (such that the minimum exists). This way, it is automatically a true density operator. A practical drawback is that calculating it is very power-consuming.

As $\gamma_n(\cdot) \to -\int \ln(p_\cdot)d_{p_\rho}$, we have defined a minimum contrast estimator in the sense of section 2.3.1. Much like for projection estimators, the Kullback distance thus estimated might be overly optimistic, and all the more when $\mathcal{Q}$ is big. Indeed, if $\mathcal{Q}$ is the set of all density operators, then there is no minimizer of the distance with the empirical distribution; however when we take only finite-dimensional models, such as

$$\mathcal{Q}(N) = \left\{\tau \in \mathcal{S}(L^2(\mathbb{R})) : \tau_{j,k} = 0 \text{ for all } j > N \text{ or } k > N\right\}, \qquad (2.23)$$

then the minimum is attained by compactness. Here the matrix entries $\tau_{j,k}$ are taken in the Fock basis (2.1).

We then have to define a penalty for choosing (almost) the best model. To do so, we make use of a (slightly simplified but sufficient for our needs) version of a theorem by Massart (2006), but we need a few definitions before stating it.

First we need a distance with which to express our results, and it is not the Kullback-Leibler, but the Hellinger distance. The Hellinger distance between two probability densities is defined in relation with the L^2-distance of the square roots of these densities:

$$h^2(p,q) = \frac{1}{2}\int \left(\sqrt{p}-\sqrt{q}\right)^2. \tag{2.24}$$

This distance does not depend on the underlying measure. The following relations are well-known:

$$\begin{aligned} \frac{1}{8}\|p-q\|_1^2 &\leq h^2(p,q) \leq \frac{1}{2}\|p-q\|_1 \\ h^2(p,q) &\leq \frac{1}{2}K(p,q). \end{aligned} \tag{2.25}$$

The penalty to be defined shall depend on the size of the model, that we have to estimate. The right tool is the metric entropy, and more precisely the metric entropy with bracketing of the model.

Definition 2.4.1. *Let $\mathcal{G}$ a function class. Let $N_{B,2}(\delta,\mathcal{G})$ be the smallest p such that there are couples of functions $[f_i^L, f_i^U]$ for i from 1 to p that fulfill $\left\|f_i^L - f_i^U\right\|_2 \leq \delta$ for every j, and for any $f \in \mathcal{G}$, there is an $i \in [1,p]$ such that:*

$$f_i^L \leq f \leq f_i^U.$$

Then $H_{B,2}(\delta,\mathcal{G}) = \ln N_{B,2}(\delta,\mathcal{G})$ is called the δ-bracketing entropy of $\mathcal{G}$

Remarks:

- Notice that the f_i^U and f_i^L need not be in $\mathcal{G}$.
- The 2 in $H_{B,2}$ stands for L^2 distance.

Looking closely at definition 2.4.1, we see that the concept of entropy depends only on those of positivity and norms. We may then define a similar bracketing entropy for any space with a norm and a partial order, such as the L^1 δ-bracketing entropy of $\mathcal{Q}(N)$: we must find couples of Hermitian operators $[\tau_i^L, \tau_i^U]$ such that $\left\|\tau_i^U - \tau_i^L\right\|_1 \leq \delta$ and such that for any $\tau \in \mathcal{Q}(N)$, there is an i such that $\tau_i^L \leq \tau \leq \tau_i^U$.

We are chiefly interested in the L^2 entropy of square roots of density (denoted by $H_{B,2}(\delta, \mathcal{P}^{\frac{1}{2}})$):

$$\mathcal{P}^{1/2}(N) \;=\; \left\{\sqrt{p_\rho} : p_\rho \in \mathcal{P}(N)\right\}.$$

Now the Theorem by Massart (2006):

Theorem 2.4.2. *Let $X_1, \ldots, X_n$ be independent, identically distributed variables with unknown density s with respect to some measure μ. Let $(S_m)_{m\in\mathcal{M}}$ be an at most countable collection of models, where for each $m \in \mathcal{M}$, the elements of S_m are assumed to be densities with respect to μ. We consider the corresponding collection of maximum likelihood estimators $\hat{s}_m$. Let* $\mathrm{pen} : \mathcal{M} \longrightarrow \mathbb{R}$ *and consider the random variable $\hat{m}$ such that:*

$$\mathbb{P}_n\left[-\ln(\hat{s}_{\hat{m}})\right] + \mathrm{pen}(\hat{m}) \;=\; \inf_{m\in\mathcal{M}} \mathbb{P}_n\left[-\ln(\hat{s}_m)\right] + \mathrm{pen}(m).$$

Let $(x_m)_{m\in\mathcal{M}}$ a collection of numbers such that

$$\sum_{m\in\mathcal{M}} e^{-x_m} \;=\; \Sigma \leq \infty.$$

For each m, we consider a function ϕ_m of $\mathbb{R}^{+*}$, nondecreasing, and such that $x \mapsto \frac{\phi_m(x)}{x}$ is nonincreasing, and:

$$\phi_m(\sigma) \;\geq\; \int_0^\sigma \sqrt{H_{B,2}(\epsilon, S_m^{\frac{1}{2}})}d\epsilon.$$

We then define each σ_m as the one positive solution of

$$\phi_m(\sigma) \;=\; \sqrt{n}\sigma^2.$$

Then there are absolute constants κ and C such that if for all $m \in \mathcal{M}$,

$$\mathrm{pen}(m) \;\geq\; \kappa\left(\sigma_m^2 + \frac{x_m}{n}\right),$$

then

$$\mathbb{E}\left[h^2(s, \hat{s}_{\hat{m}})\right] \;\leq\; C\left(K(s, S_m) + \mathrm{pen}(m) + \frac{\Sigma}{n}\right)$$

where, for every $m \in \mathcal{M}$, $K(s, S_m) = \inf_{t\in S_m} K(s, t)$.

We notice that what is bounded *in fine* is the Hellinger distance and not the Kullback. Indeed our evaluation of the estimation error, which depends upon the size of the model (its bracketing entropy), dominates the Hellinger distance but maybe not the Kullback-Leibler distance.

In our case, we have parametrized the models m by N, through definition (2.23).

To apply Theorem 2.4.2, we need to find suitable ϕ_m, and this calls for dominating the entropy integral. We reproduce here the article by Artiles et al. (2005).

By (2.25), it is sufficient to control $H_{B,1}(\delta, \mathcal{P}(N))$. Moreover, the linear extension of the morphism $\mathbf{T}$ sends a positive matrix to a positive function, and is contractive. So any covering of $\mathcal{Q}(N)$ by δ-brackets is sent upon a covering of $\mathcal{P}(N)$ by L^1 δ-brackets, that is $[p_j^L, p_j^U] = [p_{\tau_j^L}, p_{\tau_j^U}]$. Thus

$$H_{B,1}(\delta, \mathcal{P}(N)) \leq H_{B_1}(\delta, \mathcal{Q}(N)),$$

so that

$$H_{B,2}(\delta, \mathcal{P}^{\frac{1}{2}}(N)) \leq CH_{B,1}(\delta^2, \mathcal{Q}(N)).$$

Moreover:

Lemma 2.4.3.

$$H_{B,1}(\delta, \mathcal{Q}(N)) \leq CN^2 \ln \frac{N}{\delta}$$

where C is a constant not depending on δ or N, and can be put to $1 + \ln(5)$.

Proof. Let $\{\rho_j : j = 1, \ldots, c(\delta, N)\}$ a maximal set of density matrices in $\mathcal{Q}(N)$ such that for all $j \neq k$, $\|\rho_j - \rho_k\|_1 \geq \frac{\delta}{2N}$. Define the brackets $[\rho_j^L, \rho_j^U]$ as

$$\rho_j^L = \rho_j - \frac{\delta}{2N}\mathbf{1} \qquad \rho_j^U = \rho_j + \frac{\delta}{2N}\mathbf{1}.$$

Then $\|\rho_j^L - \rho_j^U\|_1 = \delta$. Moreover for any ρ in the ball $B_1(\rho_j, \frac{\delta}{2N})$, as $\|\rho - \rho_j\|_1 \leq \frac{\delta}{2N}\mathbf{1}$, we have

$$\rho_j^L \leq \rho \leq \rho_j^U$$

and as $\{\rho_j\}$ was a maximal set, this set of brackets cover $\mathcal{Q}(N)$.

So $H_{B,1}(\delta, \mathcal{Q}(N)) \leq c(\delta, N)$.

Notice that $B_1(\rho_j, \frac{\delta}{4N})$ are disjoint and included in the shell $B_1(0, 1+\frac{\delta}{4N}) - B_1(0, 1-\frac{\delta}{4N})$, so that

$$\begin{aligned} c(\delta, N) &\leq \left(\frac{4N}{\delta}\right)^{N^2} \left(\left(1+\frac{\delta}{4N}\right)^{N^2} - \left(1-\frac{\delta}{4N}\right)^{N^2}\right) \\ &\leq \left(1+\frac{4N}{\delta}\right)^{N^2} \\ &\leq \left(\frac{5N}{\delta}\right)^{N^2}, \end{aligned} \tag{2.26}$$

concluding the demonstration.

□

From this, we can obtain:

Corollary 2.4.4. *There is a constant C such that:*

$$H_{B,2}(\delta, \mathcal{P}^{\frac{1}{2}}(N)) \leq CN^2 \ln \frac{N}{\delta^2}.$$

Writing

$$\phi_N(\sigma) = \int_0^\sigma \sqrt{H_{B,2}(\epsilon, \mathcal{P}^{\frac{1}{2}}(N))} d\epsilon$$

and $\sigma_N(n)$ the only σ such that

$$\phi_N(\sigma) = \sqrt{n}\sigma^2$$

we get

$$\sigma_N(n) \leq \sqrt{\frac{C}{n}} N \left(1 + \sqrt{0 \vee \ln \frac{n}{N}}\right). \tag{2.27}$$

Indeed

$$\begin{aligned}
\phi_N(\sigma) &\leq CN\int_0^{\sigma}\sqrt{\ln\left(\frac{N}{\epsilon^2}\right)}d\epsilon \\
&= CN^{\frac{3}{2}}\int_{\sqrt{\ln\frac{N}{\sigma^2}}}^{\infty} xe^{-\frac{x^2}{2}}dx \\
&= CN^{\frac{3}{2}}\left(\int_{\sqrt{\ln\frac{N}{\sigma^2}}}^{\infty} e^{-\frac{x^2}{2}}dx - \left[xe^{-\frac{x^2}{2}}\right]_{\sqrt{\ln\frac{N}{\sigma^2}}}^{\infty}\right) \\
&\leq CN\sigma\left(1+\sqrt{\ln\frac{N}{\sigma^2}}\right)
\end{aligned}$$

where we have made use of, in each line in turn,

- Corollary 2.4.4
- the change of variables $x = \sqrt{\ln(N\epsilon^{-2})^2}$, with $\frac{d\epsilon}{dx} = -\sqrt{N}xe^{-\frac{x^2}{2}}$
- integration by parts, with x seen as a primitive and $xe^{-\frac{x^2}{2}}$ as a derivative
- the upper bound $Ce^{-\frac{x^2}{2}}$ of $\int_x^{\infty} e^{-x^2/2}dx$ for x positive when evaluating the first term.

We are looking for an upper bound on σ_N, solution of the equation

$$\sqrt{n}\sigma_N^2 = CN\sigma\left(1+\sqrt{\ln\frac{N}{\sigma_N^2}}\right).$$

We lower bound the second term by 0, and get

$$\sigma_N \geq C\frac{N}{\sqrt{n}} \equiv \sigma_m.$$

Hence the upper bound

$$\begin{aligned}
\sigma_N &= CNn^{-\frac{1}{2}}\left(1+\sqrt{\ln\frac{N}{\sigma_N^2}}\right) \\
&\leq CNn^{-\frac{1}{2}}\left(1+\sqrt{\ln\frac{N}{\sigma_m^2}}\right) \\
&= C\frac{N}{\sqrt{n}}\left(1+\sqrt{\ln\frac{n}{C^2N}}\right).
\end{aligned}$$

We may absorb the C^2 in the first multiplicative constant to find (2.27). Of course we take only the positive part of the logarithm. This will always be the case hereafter.

Applying Theorem 2.4.2 we get:

Theorem 2.4.5. *Consider the collection of maximum likelihood estimators $(\hat{\rho}_N)_{N\in\mathbb{N}}$, that is for any integer N,*

$$\mathbb{P}_n\left[-\ln(p_{\hat{\rho}_N})\right] = \inf_{\rho\in\mathcal{Q}(N)} \mathbb{P}_n\left[-\ln(p_{\hat{\rho}})\right]$$

Let pen : $\mathbb{N}\mapsto\mathbb{R}_+$ *and consider a random variable $\hat{N}$ such that*

$$\mathbb{P}_n\left[-\ln(p_{\hat{\rho}_{\hat{N}}})\right] + \mathrm{pen}(\hat{N}) = \inf_{N\in\mathbb{N}} (\mathbb{P}_n\left[-\ln(p_{\hat{\rho}_N})\right] + \mathrm{pen}(N))$$

Let $(x_N)_{N\in\mathbb{N}}$ a family of positive numbers such that

$$\sum_{N\in\mathbb{N}} e^{-x_N} = \Sigma < \infty$$

Then there are absolute constants κ and C such that if

$$\mathrm{pen}(N) \geq \kappa(\frac{N^2}{n}(1+(0\vee\ln\frac{n}{N}))+\frac{x_N}{n})$$

then

$$\mathbb{E}[h^2(p_\rho, p_{\hat{\rho}_{\hat{N}}})] \leq C\left(\inf_{N\in\mathbb{N}}(\mathbb{E}[K(\rho,\mathcal{Q}(N))]+\mathrm{pen}(N))+\frac{\Sigma}{n}\right)$$

with $K(\rho,\mathcal{Q}(N)) = \inf_{\tau\in\mathcal{Q}(N)} K(p_\rho, p_\tau)$.

Remarks:

- When designing the penalty, what stands out in this theorem is the general form of the penalty. Now the constant κ that can be explicitly computed would be very pessimistic. The best thing to do is therefore to keep the general formula for the penalty and calibrate κ using cross-validation, the slope heuristic (Massart, 2006) or any other appropriate method.

- If we wanted an explicit convergence rate for a given state, as for the photon basis in section 2.3.5, we would first need to know how the Kullback-Leibler distance $K(\rho,\mathcal{Q}(N))$ is decreasing with N. One thing that is obvious, however, is that if we add noise we convolve with the same function p_ρ and p_σ for all σ in $\mathcal{Q}(N)$, so the Kullback-Leibler distance is decreasing with the noise, so

convergence is faster when there is noise... The reason for this is that we are looking at convergence in Hellinger distance, that is a distance between the law of the result of the measurement p_ρ and p_σ. This does not tell us directly anything about what we are really interested in, that is the distance between ρ and σ (as operators). Indeed we may bound the L^2 or L^1 norm between elements of $\mathcal{Q}(N)$ by the Hellinger distance, times something depending on the sum of the L^2 or L^∞ norms of the $f^\eta_{j,k}$. And these norms are going (very fast) to infinity when there is noise, so that low Hellinger distance gives no indication on the operator norms.

2.5 Quantum calibration of a photocounter

This section features a scheme to calibrate an apparatus M measuring the number of photons in a beam with the help of a photocounter.

The physical motivation is given in Appendix 2.A.3.

The first subsection states the mathematical problem. In the two others, projection estimators and maximum likelihood estimators are respectively studied.

2.5.1 Statistical problem

The practical problem of calibration of a photocounter turns out to be mathematically speaking an entirely classical missing data problem. However, to the best of our knowledge, it has never been studied. We now describe this missing data problem.

We are given samples (I, X) in $\mathbb{N} \times \mathbb{R}$ from a probability density of the form

$$p(i,x) = \sum_{k=0}^{\infty} b_k^2 P_i^k \psi_k(x)^2. \tag{2.28}$$

In this expression, the real numbers b_k^2 satisfy $\sum_m b_k^2 = 1$. The ψ_k are the Fock basis functions given in Equation (2.1). For any k, the P_i^k are a probability measure, that is they are non-negative and $\sum_{i=0}^{\infty} P_i^k = 1$.

We know the b_k^2, and we want to retrieve the P_i^k, which we do not know. We write $P = (P_i^k)_{i,k}$.

To make clearer that this is a missing data problem, we give the following way to obtain this experiment. First we choose $K \in \mathbb{N}$ with probability given by b_k^2. We forget K, which is the missing data. Our data consists in (I, X), with i having law given by P_i^k and x with law $\psi_k(x)$.

Notice that the experimentalist has some control on the b_k^2, but usual techniques will yield b_k^2 proportional to ξ^k. This means that the low k are probed faster.

We propose below two types of estimators $\hat{P}$ for P. To get results on their efficiency, we must first find meaningful distance $d(P, \hat{P})$. Since $\sum_i P_i^k = 1$ for all $k \in \mathbb{N}$, distances like $d_2^2(P, Q) = \sum_{i,k}(P_i^k - Q_i^k)^2$ are bound to yield infinite errors on our estimators. We then must weight them, using $(a_k)_{k\in\mathbb{N}}$ of our choice. We shall use, depending on the estimator, either $d_2^2(P, Q) = \sum_{i,k} a_k^2(P_i^k - Q_i^k)^2$ with $\sum a_k^2 = 1$, or $d_1(P, Q) = \sum_{i,k} a_k|P_i^k - Q_i^k|$, with $\sum_k a_k = 1$. Then these distances are bounded by 2 on the set of all P such that $\{P_i^k\}_{i\in\mathbb{N}}$ is a probability measure for every k.

Varying the choice of a_k corresponds to putting the emphasis on different k, that is deciding which P_i^k we demand to know with the more precision. If we take the a_k decreasing, it means physically that we are more interested in the behaviour of our photocounter for a low number of photons. This is usually the case for a physicist. A possible choice is to take a_k or a_k^2 equal to b_k^2.

In the next subsection, we use projection estimators, and in the following, maximum likelihood estimators.

2.5.2 Using projection estimators

As in the tomography problem, the parameter space is contained in an infinite-dimensional vector space, and a natural type of estimators are projections of the empirical law on finite-dimensional subspaces. The problem we are left with is then again finding the best subspace.

Concretely, we consider the distance $d_2^2(P, Q) = \sum_{i,k} a_k^2(P_i^k - Q_i^k)^2$ and write $E_i^k = a_k P_i^k$. Similarly we shall write $\hat{E}_i^k = a_k \hat{P}_i^k$ for our estimator. Then

$$d_2^2(P, \hat{P}) = \sum_{i,k}(E_i^k - \hat{E}_i^k)^2,$$

and the law of our samples can be rewritten as

$$p(i, x) = \sum_k E_i^k \frac{b_k^2}{a_k} \psi_k(x)^2. \tag{2.29}$$

We may then consider $\{(b_k^2/a_k)\psi_k \mathbf{1}_{i=l}\}_{k,i}$ as a basis of our functions on $\mathbb{N} \times \mathbb{R}$. We want to use the general constructions of section 2.3. We first need a dual basis $\{g_{i,k}\}$. Now, the dual basis of $\{\psi_k^2\}$ as functions on $\mathbb{R}$ is well-known. Those are the "pattern functions" $f_{k,k}$ introduced by D'Ariano et al. (1994) (see (2.38)). From this, we deduce:

$$g_{i,k}(l,x) = \frac{a_k}{b_k^2} f_{k,k}(x) \mathbf{1}_{i=l}.$$

With these dual functions, we can define the minimum contrast function:

$$\gamma_n(Q) = d_2^2(Q,0) - 2\left(\sum_{\alpha=1}^{n} \frac{g_{i,k}(L_\alpha, X_\alpha)}{a_k}\right)\left(\sum_{i,k} a_k^2 Q_i^k\right),$$

where the (L_α, X_α) are our data, that is n independent samples with law p.

Our models $m \in \mathcal{M}$ consist in the subsets of $\mathbb{N}^2$. If $(i,k) \notin m$, then $\hat{P}_i^k = 0$. In a model m, the estimator $\hat{P}^{(m)}$ given by minimizing the contrast function is then

$$\hat{P}_i^k = \frac{1}{n}\sum_{\alpha=1}^{n} \frac{g_{i,k}(L_\alpha, X_\alpha)}{a_k} \text{ for } (i,k) \in m.$$

The penalized estimator is as always the projection estimator of the model $\hat{m}$ such that:

$$\hat{m} = \arg\min_{m \in \mathcal{M}} \gamma_n(\hat{P}^{(m)}) + \mathrm{pen}_n(m).$$

We also use the usual notation for the distance to a model:

$$d_2(P,m) = \inf_{Q \in m} d_2(P,Q).$$

We then obtain from the general theorems of section 2.3:

Theorem 2.5.1. *Let P be a photocounter and (a_k) and (b_k) with $\sum_k a_k^2 = \sum_k b_k^2 = 1$. Let $(x_{i,k})_{(i,k)\in\mathbb{N}^2}$ such that $\sum_{i,k} e^{-x_{i,k}} = \Sigma < \infty$. We define a penalty as*

$$\mathrm{pen}_n(m) = \sum_{(i,k)\in m} (1+\epsilon)\left(\ln(M_{i,k}) + \frac{x_{i,k}}{2}\right)\frac{M_{i,k}^2}{n}$$

with

$$M_{i,k} = \frac{a_k}{b_k^2}(\sup_x f_{k,k}(x) - \inf_x f_{k,k}(x)).$$

Then the penalized estimator fulfills

$$\mathbb{E}\left[\frac{\epsilon}{2+\epsilon} d_2^2(P,\hat{P})\right] \leq \inf_{m\in\mathcal{M}}\left(1+\frac{2}{\epsilon}\right) d_2^2(P,m) + 2\,\mathrm{pen}_n(m) + \frac{(1+\epsilon)\Sigma}{n}.$$

Theorem 2.5.2. *Let P be a photocounter and (a_k) and (b_k) with $\sum_k a_k^2 = \sum_k b_k^2 = 1$. Let $(y_{i,k})_{(i,k)\in\mathbb{N}^2}$ such that $\sum_{i,m} e^{-y_{i,m}} = \Sigma < \infty$. Let then*

$$x_{i,k} = 2\ln\left(\frac{a_k}{b_k^2}\|f_{k,k}\|_\infty\right) + y_{i,k}.$$

For any $\delta \in (0,1)$, with

$$\mathrm{pen}_n(m) = \sum_{(i,k)\in m} \mathrm{pen}_n^{(i,k)},$$

$$\frac{n\,\mathrm{pen}_n^{(i,k)}}{2(1+\epsilon)} = \left(\sqrt{\frac{2}{1-\delta}x_{i,k}\left(\mathbb{P}_n[g_{i,k}^2] + \frac{1}{n}\frac{a_k^2}{b_k^4}\|f_{k,k}\|_\infty^2\left(\frac{1}{3}+\frac{1}{\delta}\right)x_{i,k}\right)} + \frac{a_k\|f_{k,k}\|_\infty}{3b_k^2\sqrt{n}}x_{i,k}\right)^2,$$

there is a constant C such that:

$$\mathbb{E}\left[\left(\frac{\epsilon}{2+\epsilon}d_2^2(P,\hat{P}) - \left(\left(1+\frac{2}{\epsilon}\right)\inf_{m\in\mathcal{M}_n} d_2^2(P,m) + 2\,\mathrm{pen}_n(m)\right)\right)\vee 0\right] \leq \frac{C\Sigma}{n}$$

where $\mathcal{M}_n$ is the set of models m for which $(i,k)\in m$ implies $x_{i,k} < n$.

Remarks:

- As with the estimation of states with tomography in section 2.3, we choose with high efficiency the best subspace. It should be noticed that convergence is fast if the photocounter is good, and could be slower if it is bad. In the latter case, we know it is bad, though. Indeed, the dependence of the convergence rate on the photocounter P lies in the approximation properties of the models – subspaces – m, that is on how fast $d_2^2(P,m)$ decrease when m gets bigger. Now for an ideal photocounter, we need only the (i,i) to be in m. The penalty would be as low as possible when neglecting what happens to beams with more than a given number k of photons. For a worse photocounter, to have a good approximation of how a k-photons beam is read, we might need many i, and the penalty would include all the $\mathrm{pen}^{i,k}$.

- The estimator depends only weakly on (a_k) (unlike the distance), which is good news as it is somewhat arbitrary. Indeed, the empirical $\hat{P}_i^k$ does not depend of this sequence at all, nor do the main terms in the threshold on $\hat{P}_i^k$ of both theorems. For Theorem 2.5.1, this main term is $a_k^{-1}\sqrt{(1+\epsilon)\ln(B_{i,k})}B_{i,k}/\sqrt{n}$. Now $B_{i,k}$ depends linearly on a_k, so the only a_k left in this expression is in the logarithm which can be developed as $\ln(B_{i,k}/a_k) + \ln(a_k)$. In this way, we see that we only get another term in the penalty. For Theorem 2.5.2, the threshold is essentially $a_k^{-1}\sqrt{8(1+\epsilon)\mathbb{P}_n\left[g_{i,k}^2\right]\ln(\|g_{i,k}\|_\infty)/((1-\delta)n)}$; and as $g_{i,k}$ is proportional to a_k, the situation is the same.

- The process by which we get our data includes a tomographer and the laws $p(i,x)$ were given in the ideal case when there is no noise. If there is noise, as briefly sketched in section 2.3.6, these laws are different. However we may characterize the noise with a single $0 < \eta < 1$. We then have for free the same theorems for $\eta > \frac{1}{2}$: we only need to replace $f_{k,k}$ with $f^{\eta}_{k,k}$.

2.5.3 Maximum likelihood procedure

In this case, our results are easier expressed with the distance

$$\begin{aligned} d_1(P,\hat{P}) &= \sum_{i,k} a_k \left| P_i^m - \hat{P}_i^k \right| \\ &= \sum_{i,k} \left| E_i^k - \hat{E}_i^k \right| \end{aligned}$$

with $E_i^k = a_k P_i^k$ and $\sum_k a_k = 1$. We denote $w_i = \sum_k E_i^k$. Notice that $\sum_i w_i = 1$.

Recall that our data consists in n independent samples (L_α, X_α) with law p given by Eq. (2.28).

The main difficulty with applying here Theorem 2.4.2 lies in that the Kullback distance to the models is usually infinite (if we have $\hat{E}_i^k = 0$ for all k for some i, then $\hat{p}(i,\mathbb{R}) = 0$ and this is generally not the case for $p(i,\mathbb{R})$). The easiest way around is to keep independence and restrict attention to some set of i.

Explicitly, we take an ordering on the possible results i of the photocounter (typically, if we expect that one result corresponds roughly to a given number of photons, we can order them in increasing order. The idea is that the results that interest us most should come first). We then choose, still beforehand, $I_+ \in \mathbb{N}$, and we restrict our attention to the first $i \in [0, I_+]$. We just throw away the part of the data where the photocounter gave a result more than I_+. We are left with data size n_{I_+}, with law p_{I_+} on $[0, I_+] \times \mathbb{R}$:

$$p_{I_+} = \frac{p_{|[0,I_+]\times\mathbb{R}}}{\int_{[0,I_+]\times\mathbb{R}} p}.$$

This law is the probability measure associated to the apparatus $\tilde{P}$ for which $\tilde{P}_i^k = \frac{1}{\sum_{l\le I_+} w_l} P_i^k \mathbf{1}_{i\le I_+}$.

The models $m_{I,K}$ we work with are indexed by $K \in \mathbb{N}$ and $I \leq I_+$. They are given by the constraints:

$$\begin{aligned} \hat{E}_i^k &= 0 && \text{if } i > I_+ \\ \hat{E}_i^k &= 0 && \text{if } i > I_+ \text{ and } k \leq K \\ \sum_{i \leq I} \hat{E}_i^k &= a_k && \text{for } k \leq K \\ \hat{E}_i^k &= \frac{a_k}{I_+ + 1} && \text{for } k > K \text{ and } i \leq I_+. \end{aligned} \tag{2.30}$$

Any such element gives a probability measure on $([0, I_+] \times \mathbb{R})$. Similarly to equation (2.29), the corresponding probability law reads $\hat{p}(l, x) = \sum_{i,k} b_k^2 a_k^{-1} \hat{E}_i^k \psi_k(x)^2 \mathbf{1}_{i=l}$. The fourth condition (2.30) does not increase the complexity of the model and ensures that the Kullback distance remains finite.

We can now use an empirical maximum likelihood procedure to select within each model an estimator. It minimizes on each $m_{I,K}$ the contrast function

$$\gamma_n(Q) \;=\; \sum_{\alpha=1}^{n} - \ln q(L_\alpha, X_\alpha).$$

where Q is an element of the model $m_{I,K}$ and q the associated probability law.

We then use Theorem 2.4.2 to select the model of which we keep the estimator, through a penalization procedure. We obtain the following theorem.

Theorem 2.5.3. *Consider the collection $(\hat{P}_{I,K})_{I \leq I_+, K \in \mathbb{N}}$ of maximum likelihood estimators , defined as minimizers of*

$$\gamma_n(\hat{P}_{I,K}) \;=\; \inf_{P \in m_{I,K}} \gamma_n(P)$$

Let pen $: [0, I_+] \times \mathbb{N} \to \mathbb{R}$ *be a penalty function and define $(\hat{I}, \hat{K})$ by*

$$\gamma_n(\hat{P}_{(\hat{I},\hat{K})}) + \mathrm{pen}(\hat{I}, \hat{K}) \;=\; \inf_{I \leq I_+, K \in \mathbb{N}} \gamma_n(\hat{P}_{I,K}) + \mathrm{pen}(I, K).$$

Let $(x_{I,K})$ be a family of numbers such that

$$\sum_{I \leq I_+, K \in \mathbb{N}} e^{-x_{I,K}} \;=\; \Sigma \;<\; \infty.$$

Then there are absolute constants κ and C such that if

$$\mathrm{pen}(I, K) \;\geq\; \kappa \left((I+1)(K+1) \frac{\ln(n_{I_+})}{n_{I_+}} + \frac{x_{I,K}}{n_{I_+}} \right),$$

then

$$\mathbb{E}\left[d_1(P,\hat{P}_{(\hat{I},\hat{K})})\right]$$
$$\leq \sum_{i>I_+} w_i + \sum_{k\in\mathbb{N}} \left(2a_k \wedge \left(C\frac{a_k}{b_k^2}\|f_{k,k}\|_\infty \sqrt{\inf_{\substack{I\leq I_+\\ K\in\mathbb{N}}} K(p_{I_+}, m_{I,K}) + \mathrm{pen}(I,K) + \frac{\Sigma}{n_{I_+}}}\right)\right),$$

where $K(p_{I_+}, m_{I,K}) = \inf_{Q\in m_{I,K}} K(p_{I_+}, q)$, *intended as the Kullback distance on* $[0, I_+] \times \mathbb{R}$.

Remarks:

- As with projection estimators, we can expect fairly quick approximation if the photocounter is good. Indeed, for $K = I_+$ and the ideal photocounter, the distance $K(p_{I_+}, m_{I_+,K}) = 0$.
- Like projection estimators, the maximum likelihood strategy can also be used with noise. If $\eta > \frac{1}{2}$, we get the same theorem changing $f_{k,k}$ in $f^\eta_{k,k}$. Just notice that the infinite norm $\|f_{k,k}\|_\infty$ is exploding.
- As in section 2.4, an explicit computation of κ would be over-pessimistic and it is best to estimate it with a data-driven procedure.

Proof. First we rewrite and bound the distance d_1 in a way that suits our purpose. We separate the entries corresponding to measurement results bigger than I_+, and we recall at the third line that $\sum_{i\in\mathbb{N}} E_i^k = a_k$. Then

$$\begin{aligned}
&d_1(P,\hat{P})\\
=&\sum_{i,k}\left|E_i^k - \hat{E}_i^k\right|\\
=&\sum_{i>I_+}\sum_k E_i^k + \sum_k\sum_{i\leq I_+}\left|\hat{E}_i^k - E_i^k\right|\\
\leq&\sum_{i>I_+}\sum_k E_i^k + \sum_k\left(2a_k \wedge \left(\sum_{i\leq I_+}\left|\hat{E}_i^k - \frac{1}{\sum_{i\leq I_+} w_i}E_i^k\right| + \left(\frac{1}{\sum_{i\leq I_+} w_i} - 1\right)E_i^k\right)\right)\\
=&\sum_{i>I_+} w_i + \sum_{i\leq I_+}\frac{\sum_{i>I_+} w_i}{\sum_{i\leq I_+} w_i}\sum_k E_i^k + \sum_k\left(2a_k \wedge \sum_{i\leq I_+}\left|\hat{E}_i^k - \frac{1}{\sum_{i\leq I_+} w_i}E_i^k\right|\right)\\
=&2\sum_{i>I_+} w_i + \sum_k\left(2a_k \wedge \sum_{i\leq I_+}\left|\hat{E}_i^k - \frac{1}{\sum_{i\leq I_+} w_i}E_i^k\right|\right).
\end{aligned}$$

Let us now work a little on the last term:

$$\begin{aligned}\frac{1}{\sum_{i\leq I_+} w_i} E_i^k &= \int \frac{a_k}{b_k^2} f_{k,k}(x)\mathbf{1}_{i=l} dp_{I_+}(l,x),\\ \hat{E}_i^k &= \int \frac{a_k}{b_k^2} f_{k,k}(x)\mathbf{1}_{i=l} d\hat{p}(l,x).\end{aligned}$$

So that

$$\begin{aligned}\left|\frac{1}{\sum_{i\leq I_+} w_i} E_i^k - \hat{E}_i^k\right| &= \left|\int f_{k,k}(x)\mathbf{1}_{i=l} d(p_{I_+} - \hat{p})(l,x)\right|\\ &\leq \frac{a_k}{b_k^2}\|f_{k,k}\|_\infty \int \mathbf{1}_{i=l} d|p_{I_+} - \hat{p}|(l,x).\end{aligned}$$

Summing over i, we get:

$$\sum_{i\leq I_+}\left|\frac{1}{\sum_{i\in I_+} w_i} E_i^k - \hat{E}_i^k\right| \leq \frac{a_k}{b_k^2}\|f_{k,k}\|_\infty \int d|p_{I_+} - \hat{p}|(l,x).$$

We may then bound the distance between the POVM we calibrate and our estimator by

$$d_1(P,\hat{P}) = 2\sum_{i>I_+} w_i + \sum_{k\in\mathbb{N}}\left(2a_k \wedge \left(\frac{a_k}{b_k^2}\|f_{k,k}\|_\infty \int d|p_{I_+} - \hat{p}|(l,x)\right)\right).$$

Finishing the proof of our theorem amounts to controlling $\int d|p_{I_+} - \hat{p}|(l,x)$. We first apply Theorem 2.4.2 (assuming that our penalty is big enough, which we check below). We get:

$$\mathbb{E}\left[h^2(p_{I_+}, \hat{p}_{(\hat{I},\hat{K})})\right] \leq C\left(\inf_{I\leq I_+, K\in\mathbb{N}} K(p_{I_+}, m_{I,K}) + \mathrm{pen}(I,K) + \frac{\Sigma}{n_{I_+}}\right).$$

We then use the bound (2.25) of the square of the L^1-distance in the Hellinger distance, and finish with Jensen, using the concavity of both the function $x \mapsto (C\wedge x)$

and the square root.

$$
\begin{aligned}
& \mathbb{E}\left[d_1(P, \hat{P}_{(\hat{I},\hat{K})})\right] \\
\leq\ & \mathbb{E}\left[\sum_{i>I_+} w_i + \sum_{k\in\mathbb{N}} \left(2a_k \wedge \left(C\frac{a_k}{b_k^2}\|f_{k,k}\|_\infty \int d|p_{I_+} - \hat{p}_{(\hat{I},\hat{K})}|(l,x)\right)\right)\right] \\
\leq\ & \sum_{i>I_+} w_i + \sum_{k\in\mathbb{N}} \mathbb{E}\left[\left(2a_k \wedge \left(C\frac{a_k}{b_k^2}\|f_{k,k}\|_\infty \sqrt{h^2\left(p_{I_+} - \hat{p}_{\hat{I},\hat{K}}\right)}\right)\right)\right] \\
\leq\ & \sum_{i>I_+} w_i + \sum_{k\in\mathbb{N}} \left(2a_k \wedge \left(C\frac{a_k}{b_k^2}\|f_{k,k}\|_\infty \sqrt{\mathbb{E}\left[h^2\left(p_{I_+} - \hat{p}_{\hat{I},\hat{K}}\right)\right]}\right)\right) \\
\leq\ & \sum_{i>I_+} w_i + \sum_{k\in\mathbb{N}} \left(2a_k \wedge \left(C\frac{a_k}{b_k^2}\|f_{k,k}\|_\infty \sqrt{\inf_{\substack{I\leq I_+ \\ K\in\mathbb{N}}} K(p_{I_+}, m_{I,K}) + \mathrm{pen}(I,K) + \frac{\Sigma}{n_{I_+}}}\right)\right).
\end{aligned}
$$

The only thing we still have to check is our penalty. We must dominate $H_{B,2}(\delta, \mathcal{P}^{1/2}(I,\mathcal{M}))$ where

$$\mathcal{P}^{1/2}(I,K) \ = \ \{\sqrt{q}, Q \in m_{I,K}\}.$$

With the same reasoning as in section 2.4, it is sufficient to dominate $H_{B,1}(\delta^2, m_{I,K})$. We then mimic lemma 2.4.3. All the elements of $m_{I,K}$ are on the L^1-sphere of radius $\sum_{k\leq K} a_k$ of a vector space of dimension $(K+1)(I+1)$. We can then associate a maximal collection of brackets to a maximal collection (P_j) of $P \in m_{I,K}$ separated by $\delta^2/(2(K+1)(I+1))$. The balls $B_1(P_j, \frac{\delta^2}{(K+1)(I+1)})$ are disjoint and in the shell $B_1(0, \sum_{k\leq K} a_k + \frac{\delta^2}{(K+1)(I+1)}) - B_1(0, \sum_{k\leq K} a_k - \frac{\delta^2}{(K+1)(I+1)})$. And as with equation (2.26), we obtain

$$H_{B,1}(\delta^2, m_{I,K}) \leq C(K+1)(I+1)\ln\left(\frac{(K+1)(I+1)}{\delta^2}\right)$$

Imitating the calculation in the proof of corollary 2.4.4, we find that the solution $\sigma_{I,K}$ of the equation

$$\sqrt{n_{I_+}}\sigma_{I,K}^2 = \int_0^{\sigma_{I,K}} \sqrt{H_{B,2}(\delta, \mathcal{P}^{1/2}(I,K))}$$

admits this upper bound:

$$\sigma_{I,K} \leq C\sqrt{\frac{(K+1)(I+1)}{n_{I_+}}}(1+\sqrt{\ln n_{I_+}})$$

We may absorb the latter 1 in the constant, as long as $n_{I_+} \geq 2$...

This ends the proof.

□

2.A Background in quantum mechanics

Subsection 2.A.1 gives parallel developments of classical statistics and quantum statistics, so that any quantum notion is linked with a classical equivalent.

Subsection 2.A.2 describes both the experimental setup of quantum homodyne tomography and some basic mathematics playing a role in it. More precisely, it highlights several different representations of the state to be recovered (our unknown) and the links between them.

Subsection 2.A.3 is background for section 2.5. Notably, it explains where the formulas such as (2.29) come from.

2.A.1 Statistics: classical and quantum

We have here three different parts. The aim is to highlight the equivalences in classical and quantum formalism. The first part lies then upon the classical world, the second part recast this construction as a special case of what will be our quantum formalism, and the third part describes these quantum statistics. Bold numbers refer to the same number in the other sections. They might be repeated inside a section if the same object is introduced under different forms.

In this short introduction to the subject, we shall restrict ourselves more or less to describing what physical measurements can be done and how they can be encoded mathematically. In other words, we characterize what information can be retrieved from a system.

Classical

In the classical setting of statistics, we are working with probability measures p { **1** } on a probability space $(\mathcal{X}, \mathcal{A})$ { **2** }. For comparison, we recall that probability measures are normalized { **3** } real { **4** } non-negative { **5** } measures. Similarly measures are elements of $\mathcal{M}(\mathcal{X}, \mathcal{A})$ { **6** }, the dual of $L^\infty(\mathcal{X}, \mathcal{A})$ { **7** }.

Notice that the probability measures form a convex set, the extremal points of which are the Dirac measures { **8** } on x for $x \in (\mathcal{X}, \mathcal{A})$. They may then be described by x { **9** }. If we want to draw on the analogy with physics $(\mathcal{X}, \mathcal{A})$ may be viewed as a phase space, and the x would be the pure states. A general probability measure would describe a mixed state. These are systems that have a probability to be in this or that pure state. Any mixed state (probability measure) can be decomposed in a unique way over pure states (Dirac).

A statistical *model* { **10** } consists in a set of probability measures p_θ on a probability space $(\mathcal{X}, \mathcal{A})$ indexed by a parameter θ, for $\theta \in \Theta$ { **11** } the *parameter space.* A statistical problem consists in determining as precisely as possible, with a meaning depending on the instance, a function of θ.

Now we must gain access at information on these θ in some way. What we have access at are random variables.

The aforementioned space $L^\infty(\mathcal{X}, \mathcal{A})$ is the space of real bounded random variables f { **12** }. By analogy with the quantum case, we call these f *observables.* They correspond to the set of physical measurements that can be carried out on the system, to what can be "observed".

"Measuring" an observable f yields a result $f(x)$ { **13** }, with law:

$$\mathbb{P}_p\left[f \in B\right] = \int_{\mathcal{X}} \mathbf{1}_{f(x)\in B} dp(x) \qquad \text{for } B \in \mathcal{B} \text{ \{ 14 \}} \tag{2.31}$$

where $\mathcal{B}$ is the Borelian σ-algebra of $\mathbb{R}$. Notice that this result is not random for a pure state.

Notice also that the way we could see the probability measures p as elements of the dual of $L^\infty(\mathcal{X}, \mathcal{A})$ was by writing $p(f) = \int_{\mathcal{X}} f(x)\mathrm{d}p(x)$ { **15** }.

The most general type of statistic or estimator we can extract from data, including random strategies, is obtained by associating to each x a probability measure on an auxiliary space $(\mathcal{X}_a, \mathcal{A}, a)$ { **16** } and draw a final result according to this probability measure. This is equivalent (at the price of changing the auxiliary space) to measuring a function f { **17** } on a space $(\mathcal{X} \otimes \mathcal{X}_a, \mathcal{A} \otimes \mathcal{A}_a)$ { **18** } according to a probability measure $p_\theta \otimes s$ { **19** } with s independent of θ.

If we write (2.31) in this case, we get

$$\mathbb{P}_\theta [f \in B] = \int_{\mathcal{X}} \int_{\mathcal{X}_a} \mathbf{1}_{f(x,x_a)\in B} dp_\theta(x) ds(x_a) \qquad \text{for } B \in \mathcal{B}.$$

If we integrate out $\mathcal{X}_a$, this yields

$$\mathbb{P}_\theta [f \in B] = \int_{\mathcal{X}} f_B(x) dp_\theta(x) \qquad \text{for } B \in \mathcal{B} \text{ } \{ \textbf{ 20 } \}$$

where

- $f_{\mathbb{R}} = 1$ { **21** }
- $0 \leq f_B \leq 1$ { **22** }
- For countable disjoint B_i, $\sum_i f_{B_i} = f_{\bigcup_i B_i}$ { **23** }.

As a remark, the result $f(x)$ is essentially a label. We could write the same formula for functions with values in other measure spaces $(\mathcal{Y}, \mathcal{B})$ than $\mathbb{R}$. Just let $\mathcal{B}$ be the σ-algebra on this space. In this way, we retrieve in particular estimators in $\mathbb{R}^d$.

Another very important remark is that if we have access to two statistics f and g, we have access to both { **24** }. Indeed suppose that f was taking its values in $(\mathcal{Y}, \mathcal{B})$ and g in $(\mathcal{Z}, \mathcal{C})$. Then take a new statistic with values in the product space $(\mathcal{Y} \otimes \mathcal{Z}, \mathcal{B} \otimes \mathcal{C})$, characterized by $h_{B\otimes C} = f_B * g_C$ as real functions on $(\mathcal{X}, \mathcal{A})$. We see that the three conditions are satisfied, and that the marginals of h are f and g.

From classical to quantum

The above description was already somewhat non-conventional, with the parallel with quantum formalism in mind. In this subsection, we take one further step, by setting classical probability as a special case of what will be our quantum probability theory.

To have something easy to understand, we start from a finite probability space $(\mathcal{X}, \mathcal{A}) = \{1, \ldots, d\}$ { **2** }. We associate to it the Hilbert space of complex valued functions on this space, that is $\mathcal{H} = \mathbb{C}^d$ { **2** }. We are here endowed with a distinguished orthonormal basis $\{|e_i\rangle\}_{1\leq i\leq d}$ with $|e_i\rangle$ the function whose value is one on i and zero elsewhere.

Notice by the way the notation $|\psi\rangle$: this is a physicist's notation for vectors, elements of $\mathcal{H}$. They call this a "ket". The associated linear form, that is, the adjoint of the

vector, is called a "bra" and denoted $\langle\psi|$. Thus $\langle\phi|\psi\rangle$ is the scalar product of $|\phi\rangle$ and $|\psi\rangle$ (a "bracket").

Now to the probability measure $p = (p_1, \ldots, p_d)$ **{ 1 }** on $\{1, \ldots, d\}$, we associate the matrix ρ **{ 1 }** diagonal in our special orthonormal basis **{ 6 }**, with diagonal entries $(p_1, \ldots, p_d)$. As this is a diagonal matrix in an orthonormal basis, with non-negative elements, this is a self-adjoint **{ 4 }** non-negative **{ 5 }** matrix. Moreover, as $\sum_i p_i = 1$ **{ 3 }**, it has trace 1 **{ 3 }**.

We see that the extremal points of our set are of matrices are the orthogonal projectors on the lines spanned by our special eigenvectors, that is $|e_i\rangle\langle e_i|$ **{ 8 }**. They correspond to the Dirac measures on i. We may represent any of these *pure states* by the eigenvector $|e_i\rangle$ **{ 9 }**. We may also rewrite $\rho = \sum_i p_i |e_i\rangle\langle e_i|$.

A statistical *model* **{ 10 }** consists in a set of non-negative matrices ρ_θ with trace 1, on a Hilbert space $\mathcal{H}$, diagonal in the $\{|e_i\rangle\}_i$ basis, indexed by a parameter θ, for $\theta \in \Theta$ **{ 11 }** the *parameter space.* A statistical problem consists in determining as precisely as possible, with a meaning depending on the instance, a function of θ.

As we have done for probability measures, we identify $f \in L^\infty(\{1, \ldots, d\})$ **{ 12,7 }** with the diagonal matrix $O \in M(\mathbb{C}^d)$ **{ 12,7 }** whose diagonal elements are the $O_{i,i} = f(i)$. This is still the dual of the set of matrices diagonal on our special basis. We view the action of ρ by taking the trace of the product with ρ. That is $p(f) = \mathrm{Tr}(\rho O)$ **{ 15 }**. One can see that we have only rewritten the classical formula for the expectation.

Equivalently, measuring an observable O yields as a result an eigenvalue of O **{ 13 }**. The law of the result is given by:

$$\mathbb{P}_\rho [O \in B] = \mathrm{Tr}(\rho P_{O,B}) \qquad \text{for } B \in \mathcal{B} \text{ { 14 }}$$

where $P_{O,B}$ is the projection upon the space spanned by the eigenspaces of O corresponding to those eigenvalues λ of O such that $\lambda \in B$. In other words, in our case, $O = \sum_i f(i)|e_i\rangle\langle e_i|$. Then $P_{O,B} = \sum_{i|f(i)\in B} |e_i\rangle\langle e_i|$. This $P_{O,B}$ is playing the role of $\mathbf{1}_{f(x)\in B}$ in the classical setting. And we take note that $\mathrm{Tr}(\rho P_{O,B}) = \sum_{i|f(i)\in B} p_i$, as we should obtain from the classical formula.

We can encode in the same framework the general strategies for estimators, provided that $\mathcal{X}_a$ is also finite **{ 16 }**. The auxiliary space is then identified to $\mathcal{H}_a = \mathbb{C}^{d_a}$. We have matrices $\rho_\theta \otimes \sigma$ **{ 19 }**, with σ independent of θ. We are allowed to use as observable O **{ 17 }** any matrix diagonal in the same basis as these $\rho_\theta \otimes \sigma$. The procedure equivalent to the partial integration on $\mathcal{X}_a$ is then taking partial trace on $\mathcal{H}_a$ in $\mathbb{P}_\theta[O \in B] = \mathrm{Tr}((\rho_\theta \otimes \sigma) P_{O,B})$. And this yields $\mathrm{Tr}(\rho_\theta M(B))$ **{ 20 }** with

- $M(\mathbb{R}) = 1_{\mathcal{H}}$ **{ 21 }**
- $M(B)$ is non-negative and diagonal in the $\{|e_i\rangle\}$ basis **{ 22 }**
- For countable disjoint B_i, $\sum_i M(B_i) = M(\bigcup_i B_i)$ **{ 23 }**.

Here again, we see that if we have access to O_1 and O_2 characterized by the families $M_1(B)$ and $M_2(C)$, we have access to both **{ 24 }**. Our new measurement would be characterized by $N(B \otimes C) = M_1(B)M_2(C)$ as multiplication of matrices. Notice that this set of matrices still satisfies the three above conditions. Especially, the fact that they are still non-negative stems from that they are diagonal in the same eigenbasis.

Going from classical to quantum now means throwing away our special eigenbasis $\{|e_i\rangle\}$. The immediate consequence will be that we shall deal with objects that do not commute. And of course, we did not restrain to finite probability spaces in the classical case. Likewise, we do not restrain to finite-dimensional Hilbert spaces in the quantum case. We shall therefore deal with operators rather than matrices. Keeping the finite-dimensional example firmly in mind should be a guide to the intuition of those less proficient in operator theory.

Quantum

A quantum system is described by a *density operator* ρ **{ 1 }** over a Hilbert space $\mathcal{H}$ **{ 2 }**, that is:

Definition 2.A.1. : Density operator

A density operator, usually denoted by ρ, is a trace-class linear operator on a (complex, separable) Hilbert space $\mathcal{H}$ that satisfies:

- *ρ is self-adjoint* **{ 4 }**.
- *ρ is non-negative (notice that this implies self-adjointness)* **{ 5 }**.
- Tr $\rho = 1$ **{ 3 }**.

If $\mathcal{H}$ is finite-dimensional, those are just the (self-adjoint) non-negative matrices with trace 1.

We denote by $\mathcal{S}(\mathcal{H})$ the set of density operators on $\mathcal{H}$.

Density operators are a convex set, too. The extremal points are called "pure states". They are the orthogonal projectors on 1-dimensional spaces { 8 }. Thus we can represent them by a norm 1 element of $\mathcal{H}$, denoted by $|\psi\rangle$ { 9 }. The corresponding density matrix is then $\rho = |\psi\rangle\langle\psi|$. Notice that it would be more precise to speak of $|\psi\rangle$ as an element of the projective space $\mathcal{PH}$, but we conform here to the usage of physicists. Notice also that there are infinitely many pure states even in the finite-dimensional case, unlike in the classical framework. Let us finally signal that the decomposition of a mixed state on pure states is *not* unique. It is essentially unique if we further impose that the pure states of the decomposition are all orthogonal, though.

A quantum statistical *model* { 10 } consists in a set of density operators ρ_θ on a Hilbert space $\mathcal{H}$ indexed by a parameter θ, for $\theta \in \Theta$ { 11 } the *parameter space*. A statistical problem consists in determining as precisely as possible, with a meaning depending on the instance, a function of θ.

Now the role of random variables is played by observables. Those are the elements O { 12 } of $\mathcal{B}_{sa}(\mathcal{H})$ { 7 }, the bounded self-adjoint operators upon $\mathcal{H}$. If we are dealing with finite-dimensional $\mathcal{H}$, those are the self-adjoint matrices.

As a remark, the dual of $\mathcal{B}_{sa}(\mathcal{H})$ is the set of self-adjoint trace-class operators, which ρ is in. This duality is given by the formula of the expectation of measuring O on ρ, also called *Born's rule*:

$$\mathbb{E}_\rho[O] = \mathrm{Tr}(\rho O) \quad \{ 15 \} \tag{2.32}$$

When measuring O, the result is an element of the spectrum of O { 13 }, that is in the finite-dimensional picture, an eigenvalue of O. The law of the result when measuring O on ρ is:

$$\mathbb{P}_\rho[O \in B] = \mathrm{Tr}(\rho P_{O,B}) \qquad \text{for } B \in \mathcal{B} \{ 14 \} \tag{2.33}$$

where $P_{O,B}$ is coming from the spectral measure of O. This is an object associated to self-adjoint operators through the spectral theorem, whose main property is that the expectation of the law above is given by the Born's rule for any density operator ρ. We only give the derivation for finite-dimensional $\mathcal{H}$. Then, as O is self-adjoint, we can diagonalize it in an orthonormal basis, and write $O = \sum_i \lambda_i |\psi_i\rangle\langle\psi_i|$. Then $P_{O,B} = \sum_{i|\lambda_i \in B} |\psi_i\rangle\langle\psi_i|$. We see that in this case the law of the measurement is coherent with the expectation given by Born's rule (2.32).

Generally $\{P_{O,B}\}_B$ is a projector valued measure, the definition of which we give below. To each projector valued measure corresponds an observable, and to each observable corresponds a projector valued measure. We may then consider that this concept is also a definition of an observable.

Definition 2.A.2. : Projector valued measure { 12 }

A projector operator valued measure $\{P(B)\}_{B\in\mathcal{B}}$ *is a set of operators on* $\mathcal{H}$ *such that:*

- $P(B)$ *is an orthogonal projector.*
- $P(\mathbb{R}) = \mathbf{1}_{\mathcal{H}}$.
- *For disjoint countable* B_i, $\sum_i P(B_i) = P(\bigcup_i B_i)$.

Notice that these are the axioms of a probability measure, except that we do not deal with real numbers but with projection operators.

Combining this definition with the definition of a density operator, we can check that formula (2.33) yields a true probability measure. Indeed, as both ρ and $P_{O,B}$ are non-negative, the probability of any event is non-negative. With the countable additivity property of projector valued measure and linearity of product and trace, we get the countable additivity of a probability measure. Finally, the probability of the universe is $\mathrm{Tr}(\rho P_{O,\mathbb{R}}) = \mathrm{Tr}(\rho \mathbf{1}_{\mathcal{H}}) = 1$.

Remark: - even for a pure state, the result of the measurement is random, unless the pure state is an eigenvector of O.

Now what is the most general estimation strategy, or measurement? The right analogy is that of the auxiliary space. We measure observables O **{ 17 }** on a Hilbert space $\mathcal{H} \otimes \mathcal{H}_a$ **{ 18 }** under the density operator $\rho_\theta \otimes \sigma$ **{ 19 }**, with σ independent of θ. Now we may take partial trace in (2.33) along $\mathcal{H}_a$, and we obtain equivalence of this scheme with measuring a *positive operator valued measure* (POVM).

Definition 2.A.3. : Measurement (POVM) { 17 }

A measurement M *on a quantum system, taking values* x *in a measurable space* $(\mathcal{X}, \mathcal{A})$ *is specified by a* positive operator valued probability measure *or* POVM *for short, that is a collection of self-adjoint matrices* $M(A) : A \in \mathcal{A}$ *such that:*

- $M(\mathcal{X}) = \mathbf{1}$, *the identity matrix* **{ 21 }**
- *Each* $M(A)$ *is non-negative* **{ 22 }**
- *For disjoint countable* A_i, $\sum_i M(A_i) = M(\bigcup A_i)$ **{ 23 }**.

The $M(A)$ *are called the* POVM elements.

The law of measuring M on ρ is given by

$$\mathbb{P}_\rho[O \in A] = \mathrm{Tr}(\rho M(A)) \qquad \text{for } A \in \mathcal{A} \text{ { 20 }}. \tag{2.34}$$

With the same reasoning as for projector valued measure (which are a special case of these POVMs), this is a genuine probability measure.

A special case of POVM is that of a POVM dominated by σ-finite measure ν on $(\mathcal{X}, \mathcal{A})$, that is

$$M(A) = \int_A m(x)\mathrm{d}\nu(x) \text{ for all } A \in \mathcal{A} \tag{2.35}$$

where $m(x)$ is positive for all x and $\int_\mathcal{X} m(x)\mathrm{d}\nu(x) = \mathbf{1}_\mathcal{H}$. The POVM associated to homodyne tomography is dominated by the Lebesgue measure.

The very important difference with the classical world is that if we can have access to M_1 or M_2, in general, we cannot have access to both simultaneously { 24 }. We cannot copy what we have done in the former paragraph, since $M_1(A)M_2(B) + M_2(B)M_1(A)$ might not be non-negative if $M_1(A)$ and $M_2(B)$ do not commute. More generally, there is usually no way to create a new POVM N with values in $(\mathcal{X}\otimes\mathcal{Y}, \mathcal{A}\otimes\mathcal{B})$ such that the marginals are M_1 and M_2. Notably, two observables that do not commute can never be measured simultaneously. As an example, consider that M_1 and M_2 are two projector valued measures on $\mathbb{C}^2$, each with values in $\{0, 1\}$, corresponding to observables diagonal in different bases $\{e_0, e_1\}$ and $\{f_0, f_1\}$. Then $N(0, 0)$ should be proportional both to $|e_0\rangle\langle e_0|$ and $|f_0\rangle\langle f_0|$. So that it is null. Same remark for the other $N(i, j)$. Thus $N(\{O, 1\}^{\otimes 2}) = 0 \neq \mathbf{1}$. So that it is null.

The truly quantum feature of quantum statistics lies in that we should decide which measurement is to be carried out. Once we have chosen our measurement, we are left through (2.34) with a classical statistical experiment. This is the case in this chapter.

As a last remark on the subject, we could have developed a slightly more general formalism, based on C^*-algebras, that would have been parallel to Le Cam formulation of statistics. In practical applications, the formalism above is usually sufficient.

2.A.2 Quantum homodyne tomography

The system we work with is the harmonic oscillator. Both in classical or quantum mechanics, the harmonic oscillator is a basic and pervading system. It describes,

notably, a particle on a line, or a mode of the electromagnetic field (that is monochromatic light), as in our case.

The state of a quantum harmonic oscillator is described by an operator on $L^2(\mathbb{R})$ (this is the Hilbert space $\{\ \mathbf{1}\ \}$). There are two important observables corresponding to the canonical coordinates of the particle. If we know the expectation of measuring on a state ρ any operator in the algebra they generate, then we know ρ. Those observables are $\mathbf{P}$,the magnetic field, and $\mathbf{Q}$, the electric field. They satisfy the (canonical) commutation relations:

$$\begin{aligned}[\mathbf{Q},\mathbf{P}] &= \mathbf{QP} - \mathbf{PQ} \\ &= i\mathbf{1}.\end{aligned}$$

They are realized as:

$$\begin{aligned}(\mathbf{Q}\psi_1)(x) &= x\psi_1(x) \\ (\mathbf{P}\psi_2)(x) &= -i\frac{d\psi_2(x)}{dx}.\end{aligned} \tag{2.36}$$

As they do not commute, they cannot be measured simultaneously. However, any linear combination can theoretically be measured. These $\mathbf{X}_\phi = \sin(\phi)\mathbf{Q} + \cos(\phi)\mathbf{P}$ are called *quadratures.*

Using an experimental setup proposed by Vogel et Risken (1989), each of these quadratures could be experimentally measured on a laser beam (Smithey et al., 1993). The technique is called *quantum homodyne tomography.*

The optical set-up sketched in figure 2.2 consists of an additional laser of high intensity $|z| \gg 1$ called the local oscillator, a beam splitter through which the cavity pulse prepared in state ρ is mixed with the laser, and two photodetectors each measuring one of the two beams and producing currents $I_{1,2}$ proportional to the number of photons. An electronic device produces the result of the measurement by taking the difference of the two currents and rescaling it by the intensity $|z|$. A simple quantum optics computation by Leonhardt (1997) shows that if the relative phase between the laser and the cavity pulse is chosen to be ϕ then $(I_1 - I_2)/|z|$ has density $p_\rho(x|\phi)$ corresponding to measuring $\mathbf{X}_\phi$.

Knowledge of $P_\rho(x|\phi)$, the law of the result of the measurement $\mathbf{X}_\phi$ on ρ, for all ϕ, is enough to reconstruct the state ρ. As we have seen, the experimentalist may choose ϕ when measuring. We assume that the measurement carried out on each of the n systems in state ρ is the following: first choose ϕ uniformly at random, then

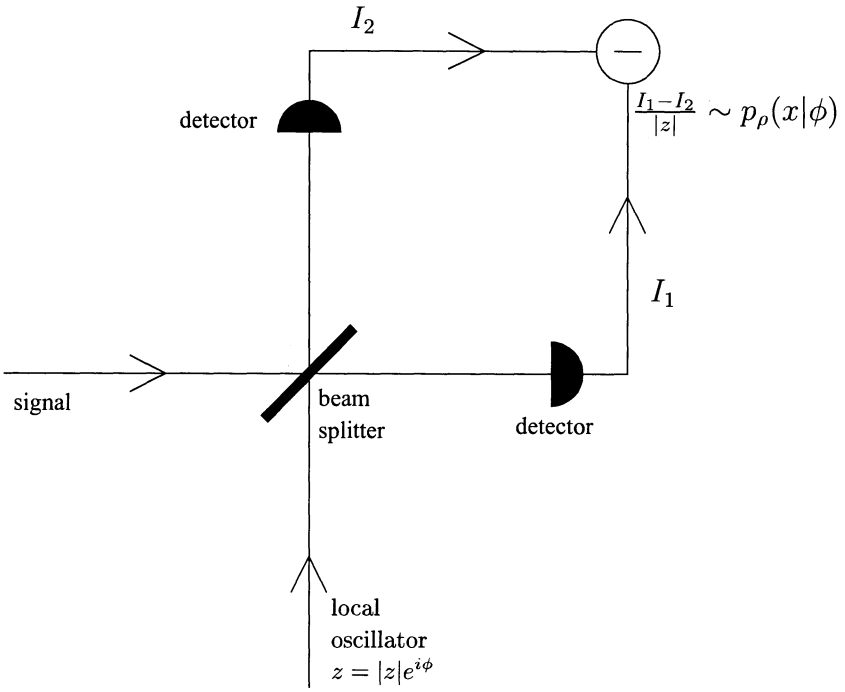

Figure 2.2: Quantum Homodyne Tomography measurement set-up

measure $\mathbf{X}_\phi$. We get a random variable $\mathbf{Y} = (\mathbf{X}, \Phi)$ with values in $\mathbb{R} \times [0, \pi)$ whose density with respect to the Lebesgue measure is $p_\rho(x, \phi) = \frac{1}{\pi} p_\rho(x|\phi)$.

Now we make explicit the links between ρ, $p_\rho(x, \phi)$ and the Wigner function W_ρ. First we write ρ in a particular basis, physically very meaningful, the *Fock basis*, already given in Sec. 2.2:

$$\psi_k(x) = H_k(x)e^{-x^2/2},$$

where H_k is the k-th Hermite polynomial, normalized so that the L^2-norm of ψ_k is 1. The projector on ψ_k is the pure state with precisely k photons. We also denote this state by the ket $|k\rangle$.

The matrix entries of p_ρ in this basis are $\rho_{j,k} = \langle \psi_j, \rho\psi_k \rangle$. We can then derive from (2.32) and (2.36) the formula we gave in Sec. 2.2:

$$\begin{aligned} \mathbf{T} : \mathcal{S}(L^2(\mathbb{R})) &\longrightarrow L^1(\mathbb{R} \times [0, \pi]) \\ \rho &\mapsto \left(p_\rho : (x, \phi) \mapsto \sum_{j,k=0}^{\infty} \rho_{j,k} \psi_j(x) \psi_k(x) e^{-i(j-k)\phi} \right). \end{aligned} \tag{2.37}$$

The mapping $\mathbf{T}$ associating P_ρ to ρ is invertible, so we may hope to find ρ from the independent identically distributed results $Y_1, Y_2, \ldots, Y_n$ of the measurements of the

n systems in state ρ. This implies notably that p_ρ is another representation of the state.

More explicitly, there are *pattern functions* $f_{j,k}$ (D'Ariano et al., 1994) against which to integrate p_ρ to find any matrix entry of ρ in the Fock basis, that is:

$$\rho_{j,k} = \int_{-\infty}^{\infty} dx \int_0^\pi \frac{d\phi}{\pi} p_\rho(x,\phi) f_{j,k}(x) e^{i(j-k)\phi}.$$

These $f_{j,k}$ are bounded real functions. That inverting the Radon transform is an ill-posed problem can be seen in the behaviour of $f_{j,k}$ when j and k go to infinity. Several formulas were found for these functions (Leonhardt et al., 1995), among which:

$$f_{j,k}(x) = \frac{d}{dx}(\chi_j(x)\phi_k(x)) \tag{2.38}$$

for $k \geq j$, where χ_j and ϕ_k are respectively the square-integrable and the unbounded solutions of the Schrödinger equation:

$$\left[-\frac{1}{2}\frac{d^2}{dx^2} + \frac{1}{2}x^2\right]\psi = \omega\psi, \quad \omega \in \mathbb{R}.$$

Another one, maybe more practical when it comes to theoretical calculations, or when we add noise (see section 2.3.6) is:

$$f_{j,k}(x,\phi) = \sqrt{\frac{j!}{k!}} \int_{-\infty}^{\infty} |r| e^{-\frac{r^2}{2}+2irx} r^{k-j} L_j^{k-j}(r^2) dr$$

where the L_j^d are the Laguerre polynomials, that is the orthogonal polynomials with respect to the measure $e^{-x}x^d$ on $\mathbb{R}^+$.

Let's now have a look at the Wigner function. This is a real function of two variables, with integral 1, but that may be negative in places. It can be interpreted as a generalized joint probability density of the electric and magnetic fields q and p. As both cannot be measured simultaneously, the negative patches are not nonsense. On the other hand, any projection on a line of the Wigner function must be a true probability density, as it is the law of $\mathbf{X}_\phi$, which is an observable. In fact, the Wigner function may be seen as the probability density on $\mathbb{R}^2$ resulting from (2.34) when measuring on ρ a "POVM" whose elements are not non-negative, but whose marginals on each line $\mathbb{R}$ are the X_ϕ.

As we have already said in the introduction, p_ρ is the Radon transform of the Wigner function. The Wigner function can be defined by its Fourier transform.

This definition tells how to find the Wigner function W of the state from its density matrix ρ:

$$\mathcal{F}_2 W(u,v) = \mathrm{Tr}(\rho e^{-iu\mathbf{Q}-iv\mathbf{P}}). \tag{2.39}$$

On the other hand, the generating function of $p_\rho(\cdot|\phi)$ is

$$\mathbb{E}\left[e^{itX_\phi}\right] = \mathrm{Tr}(\rho e^{it\mathbf{X}_\phi}).$$

In other words, $\mathcal{F}_2 W(t\cos\phi, t\sin\phi) = \mathcal{F}[p_\rho(\cdot,\phi)](t)$. These relations are known to imply that $p_\rho = \mathbf{R}(W)$ (Deans, 1983) where $\mathbf{R}$ is the Radon transform. Explicitly:

$$p_\rho(x,\phi) = \int_{-\infty}^{\infty} W(x\cos\phi + y\sin\phi, x\sin\phi - y\cos\phi)dy.$$

The Radon transform is illustrated by Fig. 2.1, given in Sec. 2.2.

Finding the Wigner function from the data means then inverting the Radon transform, hence the name of tomography: that is the same mathematical problem as with the brain imagery technique called Positron Emission Tomography.

2.A.3 Physical origin of the photocounter calibration problem

An experiment usually ends with a measurement. We need, however, an apparatus to measure. And we first have to know what is the meaning of the result the apparatus is giving us: it is not at all obvious a priori that if our new thermometer says "31° C", the temperature cannot be "32° C". That is why we must *calibrate* our measurement apparatus. In quantum mechanics, this means associating with each result i of our measurement the positive operator $P(i)$, such that P is the POVM (see definition 2.A.3) corresponding to our measurement.

D'Ariano et al. (2004) have introduced a general calibration procedure. The procedure relies on comparing with an already calibrated apparatus, using entangled states. Let us describe this more precisely in the special case of the photocounter.

A photocounter is an apparatus that aims at counting the photons in a beam. The ideal detector D has therefore POVM elements given by $D(i) = |i\rangle\langle i|$ in the Fock basis. Recall we use the physicists' notation, where $|\cdot\rangle$ is a vector and $\langle\cdot|$ is the associated linear form. Moreover $|i\rangle$ is the vector corresponding to the pure state with i photons, that is the function ψ_i on $L^2(\mathbb{R})$, that we had defined in (2.1).

Models of the noise (non-unit efficiency and dark current) leave the POVM diagonal in this basis. Thus, we are only interested in the diagonal elements of P_i in the

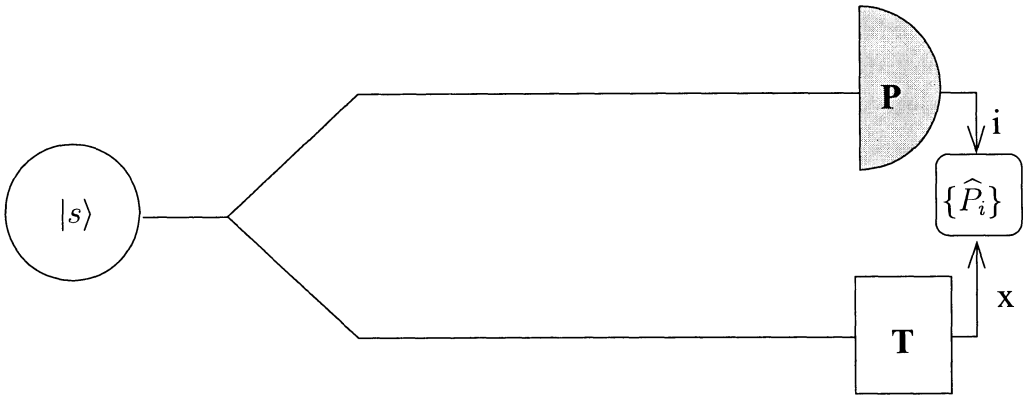

Figure 2.3: Experimental set-up to determine the POVM associated to an unknown photocounter **P**. We use it to measure a known bipartite state $|s\rangle$, jointly with a tomographer **T**. The photocounter gives a result i and the tomographer a result x. From these samples, we construct an estimator $\{\hat{P}_i\}$ of the self-adjoint operators associated to the results $\{i\}$ by the photocounter **P**.

Fock basis. To obtain those we send a twin beam state, one of the beams in the photocounter, the other in a homodyne tomographer. We get a result I from the photo-counter, and X from the tomographer (figure 2.3; as we are only interested in the diagonal elements, we shall see that we do not need the phase ϕ, as long as the experimentalist chooses it randomly). We then have to process these outcomes (I, X) to find P.

Mathematically, the twin beam is a system in a state $|s\rangle = \sum_{k=0}^{\infty} b_k |k\rangle \otimes |k\rangle$. This notation (where we may choose the b_k non-negative) means that the underlying Hilbert space is $L^2(\mathbb{R}) \otimes L^2(\mathbb{R})$, and that ρ is the pure state that projects on the line spanned by this vector. Here again, $|k\rangle$ is the vector corresponding to the pure state with k photons. Finally $\sum_k b_k^2 = 1$, so that the vector state $|s\rangle$ is normalized and the density operator is $\rho = |s\rangle\langle s|$.

Now, what is the law $p(i, x)$ of the samples we get? By (2.37) we see that the POVM associated to the tomographer is dominated by the Lebesgue measure on $\mathbb{R} \times [0, \pi)$, as in (2.35). That is $\langle j|t_{x,\phi}|k\rangle = \psi_j(x)\psi_k(x)e^{-i(j-k)\phi}$, where we have denoted $t_{x,\phi}$ the self-adjoint operator associated to the result (x, ϕ) for the POVM of the tomographer. If we forget about ϕ after having chosen it randomly, we then get $\langle j|t_x|k\rangle = \psi_k(x)^2 \mathbf{1}_{j=k}$. We have now all the ingredients for calculating our law, given the notation $\langle k|M_i|k\rangle = M_i^k$.

$$\begin{aligned} p(i,x) &= \mathrm{Tr}(\rho(P_i \otimes t_x)) \\ &= \langle s|(P_i \otimes t_x)|s\rangle \\ &= \sum_{k_1,k_2} b_{k_1} b_{k_2} (\langle k_1| \otimes \langle k_1|)(P_i \otimes t_x)(|k_2\rangle \otimes |k_2\rangle) \\ &= \sum_{k_1,k_2} b_{k_1} b_{k_2} \langle k_1|P_i|k_2\rangle \langle k_1|t_x|k_2\rangle \\ &= \sum_{k=0}^{\infty} b_k^2 P_i^k \psi_k(x)^2. \end{aligned}$$

(As a remark, the fourth line shows that the use of the phase would be to retrieve the non-diagonal elements, in which we are not interested.)

We have thus recovered (2.28), and explained how we got the data with which we want to estimate the M_i^m.

Chapitre 3

Discrimination

Ce chapitre est la fusion de (D'Ariano et al., 2005a) et (D'Ariano et al., 2005b).

Résumé : Nous dérivons la mesure optimale pour la discrimination des états quantiques, ainsi que pour la discrimination entre des canaux de Pauli, dans un cadre minimax. Pour les états, nous considérons à la fois les problèmes de discrimination avec erreur minimale, et de discrimination sans ambiguïté. Nous présentons les relations entre les mesures optimales résultant de ces deux critères. Nous montrons qu'il y a des cas où le risque minimal ne peut être atteint par une observable, et que ce trait est fréquent dans l'estimation minimax.

Pour les canaux de Pauli, nous considérons uniquement le problème de discrimination avec erreur minimale, c'est-à-dire que nous maximisons la plus faible des probabilités d'identifier correctement le canal. Nous trouvons l'état d'entrée optimal et montrons sous quelles conditions l'usage de l'intrication améliore strictement les résultats. Enfin, nous comparons les stratégies minimax et bayésiennes.

3.1 Introduction

The concept of distinguishability applies to quantum states (Wootters, 1981; Braunstein et Caves C. M., 1994) and quantum processes (Gilchrist et al., 2004; Belavkin et al., 2005), and is strictly related to quantum nonorthogonality, a basic feature of quantum mechanics. The problem of discriminating nonorthogonal quantum states has been extensively addressed (Bergou et al., 2004, and references therein), also

with experimental demonstrations. Typically, two discrimination schemes are considered: the minimal-error probability discrimination (Helstrom, 1976), where each measurement outcome selects one of the possible states and the error probability is minimized, and the optimal unambiguous discrimination (Ivanovic, 1987), where unambiguity is paid by the possibility of getting inconclusive results from the measurement. The problem has been analyzed also in the presence of multiple copies (Acin et al., 2005), and for bipartite quantum states, and global joint measurements have been compared to LOCC measurements, i.e. local measurements with classical communication (Walgate et al., 2000; Virmani et al., 2001; Ji et al., 2005).

The problem of discrimination can be addressed also for quantum operations (Sacchi, 2005a). This may be of interest in quantum error correction (Knill et al., 2002, and references therein), since knowing which error model is the proper one influences the choice of the coding strategy as well as the error estimation employed. Clearly, when a repeated use of the quantum operation is allowed, a full tomography can identify it. On the other hand, a discrimination approach can be useful when a restricted number of uses of the quantum operation is available. Differently from the case of discrimination of unitary transformations (Childs et al., 2000b), for quantum operations there is the possibility of improving the discrimination by means of ancillary-assisted schemes such that quantum entanglement can be exploited (Sacchi, 2005a). Notably, entanglement can enhance the distinguishability even for entanglement-breaking channels (Sacchi, 2005c). The use of an arbitrary maximally entangled state turns out to be always an optimal input when we are asked to discriminate two quantum operations that generalize the Pauli channel in any dimension. Moreover, in the case of Pauli channels for qubits, a simple condition reveals if entanglement is needed to achieve the ultimate minimal error probability (Sacchi, 2005a,b). All the previous statements refer to a Bayesian approach.

We address here the problem of optimal discrimination of quantum states, and of two Pauli channels, in the minimax game-theoretical scenario. In this strategy no prior probabilities are given. The relevance of this approach is both conceptual, since for a frequentist statistician the *a priori* probabilities have no meaning, and practical, because the prior probabilities may be actually unknown, as in a non cooperative cryptographic scenario. We shall derive the optimal measurement for minimax state discrimination both for minimal-error and unambiguous discrimination problems. We shall also provide the relation between the optimal measurements according to the minimax and the Bayesian strategies. We shall show that, quite unexpectedly, there are instances in which the minimum risk can be achieved only by non orthogonal POVM measurement, and this is a common feature of the minimax estimation strategy. Similarly, for channels discrimination, we shall give the optimal input states and measurements whether or not we allow using an ancilla, and show that in the latter case, the optimal input state might differ from the usual Bayesian ones.

In more detail, in Section 3.2, we pose the problem of discrimination of two quantum states in the minimax scenario. Such an approach is equivalent to a minimax problem, where one should maximise the smallest of the two probabilities of correct detection over all measurement schemes. For simplicity we will consider equal weights (i.e. equal prices of misidentifying the states), and we will provide the optimal measurement for the minimax discrimination, along with the connection with the optimal Bayesian solution. As mentioned, a striking result of this section is the existence of couples of mixed states for which the optimal minimax measurement is unique and *non orthogonal*.

In Section 3.3 we generalize the results for two-state discrimination to the case of $N \geq 2$ states and arbitrary weights. First, we consider the simplest situation of covariant state discrimination problem. Then, we address the problem in generality, resorting to the related convex programming method.

In Section 3.4 we provide the solution of the minimax discrimination problem in the scenario of unambiguous discrimination. We refine, if need be, the minimax criterion, so that the solution becomes unique.

From Section 3.5, we turn our attention from states to Pauli channels. We first briefly review the problem of discrimination of two Pauli channels in the Bayesian framework, where the channels are supposed to be given with assigned *a priori* probabilities. We report the result for the optimal discrimination, along with the condition for which entanglement with an ancillary system at the input of the channel strictly enhances the distinguishability.

In Section 3.6 we study the problem of discrimination of two Pauli channels in the minimax approach. We show that when an entangled-input strategy is adopted, the optimal discrimination can always be achieved by sending a maximally entangled state into the channel, as it happens in the Bayesian approach. On the contrary, the optimal input state for a strategy where no ancillary system is used can be different in the minimax approach with respect to the Bayesian one. In the latter the optimal input can always be chosen as an eigenstate of one of the Pauli matrices, whereas in the former this may not be the case.

3.2 Optimal minimax discrimination of two quantum states

We are given two states ρ_1 and ρ_2, generally mixed, and we want to find the optimal measurement to discriminate between them in a minimax strategy. The measurement is described by a positive operator-valued measurement (POVM) with two

outcomes, namely $\vec{P} \equiv (P_1, P_2)$, where P_i for $i = 1, 2$ are non-negative operators satisfying $P_1 + P_2 = I$.

In the usually considered Bayesian approach to the discrimination problem, the states are given with *a priori* probability distribution $\vec{\pi} \equiv (\pi_1, \pi_2)$, respectively, and one looks for the POVM that minimizes the average error probability

$$p_E = \pi_1 \mathrm{Tr}[\rho_1 P_2] + \pi_2 \mathrm{Tr}[\rho_2 P_1]. \tag{3.1}$$

The solution can then be achieved by taking the orthogonal POVM made by the projectors on the support of the positive and negative part of the Hermitian operator $\pi_1\rho_1 - \pi_2\rho_2$, and hence one has (Helstrom, 1976)

$$p_E^{(Bayes)} = \frac{1}{2}\left(1 - \|\pi_1\rho_1 - \pi_2\rho_2\|_1\right), \tag{3.2}$$

where $\|A\|_1$ denotes the trace norm of A.

In the minimax problem, one does not have *a priori* probabilities. However, one defines the error probability $\varepsilon_i(\vec{P}) = \mathrm{Tr}[\rho_i(I - P_i)]$ of failing to identify ρ_i. The optimal minimax solution consists in finding the POVM that achieves the minimax

$$\varepsilon = \min_{\vec{P}} \max_{i=1,2} \varepsilon_i(\vec{P}), \tag{3.3}$$

or equivalently, that maximizes the worst probability of correct detection

$$1 - \varepsilon = \max_{\vec{P}} \min_{i=1,2} [1 - \varepsilon_i(\vec{P})] = \max_{\vec{P}} \min_{i=1,2} \mathrm{Tr}[\rho_i P_i]. \tag{3.4}$$

The minimax and Bayesian strategies of discrimination are connected by the following theorem.

Theorem 3.2.1. *If there is an* a priori *probability* $\vec{\pi} = (\pi_1, \pi_2)$ *for the states* ρ_1 *and* ρ_2*, and a measurement* $\vec{P}$ *that achieves the optimal Bayesian average error for* $\vec{\pi}$*, with equal probabilities of correct detection, i.e.*

$$\mathrm{Tr}[\rho_1 P_1] = \mathrm{Tr}[\rho_2 P_2], \tag{3.5}$$

then $\vec{P}$ *is also the solution of the minimax discrimination problem.*

Proof. In fact, suppose on the contrary that there exists a POVM $\vec{P}$ such that $\min_{i=1,2} \mathrm{Tr}[\rho_i P_i] > \min_{i=1,2} \mathrm{Tr}[\rho_i B_i]$. Due to assumption (3.5) one has $\mathrm{Tr}[\rho_i P_i] > \mathrm{Tr}[\rho_i B_i]$ for both $i = 1, 2$, whence

$$\sum_i \pi_i \,\mathrm{Tr}(\rho_i P_i) > \sum_i \pi_i \,\mathrm{Tr}(\rho_i B_i) \tag{3.6}$$

which contradicts the fact that $\vec{P}$ is optimal for $\vec{a}$. □

The existence of an optimal $\vec{P}$ as in Theorem 3.2.1 will be shown in the following.

First, by labeling with $\vec{P}^{(\pi)}$ an optimal POVM for the Bayesian problem with prior probability distribution $\vec{\pi} = (\pi, 1-\pi)$, and defining

$$\chi(\pi, \vec{P}) \doteq \pi \operatorname{Tr}(\rho_1 P_1) + (1-\pi)\operatorname{Tr}(\rho_2 P_2), \tag{3.7}$$

we have the lemma:

Lemma 3.2.2. *The function* $f(\pi) \doteq \operatorname{Tr}(\rho_1 P_1^{(\pi)}) - \operatorname{Tr}(\rho_2 P_2^{(\pi)})$ *is monotonically nondecreasing, with minimum value* $f(0) \leq 0$, *and maximum value* $f(1) \geq 0$.

In fact, consider $\vec{P}^{(\pi)}$ and $\vec{P}^{(\varpi)}$ for two values π and ϖ with $\pi < \varpi$ and define $\vec{D} = \vec{P}^{(\varpi)} - \vec{P}^{(\pi)}$. Then

$$\begin{aligned} \chi(\pi, \vec{P}^{(\varpi)}) &= \chi(\pi, \vec{P}^{(\pi)}) + \chi(\pi, \vec{D}) \\ \chi(\varpi, \vec{P}^{(\pi)}) &= \chi(\varpi, \vec{P}^{(\varpi)}) - \chi(\varpi, \vec{D}). \end{aligned} \tag{3.8}$$

Now, since $\chi(\pi, \vec{P}^{(\pi)})$ is the optimal probability of correct detection for prior π, and analogously $\chi(\varpi, \vec{P}^{(\varpi)})$ for prior ϖ, then $\chi(\pi, \vec{D}) \leq 0$ and $\chi(\varpi, \vec{D}) \geq 0$, and hence

$$0 \leq \chi(\varpi, \vec{D}) - \chi(\pi, \vec{D}) = (\varpi - \pi)[\operatorname{Tr}(\rho_1 D_1) - \operatorname{Tr}(\rho_2 D_2)].$$

It follows that $\operatorname{Tr}(\rho_1 D_1) \geq \operatorname{Tr}(\rho_2 D_2)$, namely

$$\operatorname{Tr}(\rho_1 P_1^{(\varpi)}) - \operatorname{Tr}(\rho_1 P_1^{(\pi)}) \geq \operatorname{Tr}(\rho_2 P_2^{(\varpi)}) - \operatorname{Tr}(\rho_2 P_2^{(\pi)}) \tag{3.9}$$

or, equivalently

$$\operatorname{Tr}(\rho_1 P_1^{(\varpi)}) - \operatorname{Tr}(\rho_2 P_2^{(\varpi)}) \geq \operatorname{Tr}(\rho_1 P_1^{(\pi)}) - \operatorname{Tr}(\rho_2 P_2^{(\pi)}). \tag{3.10}$$

Equation (3.10) states that the function $f(\pi)$ is monotonically nondecreasing. Moreover, for $\pi = 0$ the POVM detects only the state ρ_2, whence $\operatorname{Tr}(\rho_2 P_2^{(0)}) = 1$, and one has $f(0) = -1 + \operatorname{Tr}[\rho_1 P_1^{(0)}] \leq 0$. Similarly one can see that $f(1) \geq 0$. □

We can now prove the theorem:

Theorem 3.2.3. *An optimal* $\vec{P}$ *as in Theorem 3.2.1 always exists.*

Proof. Consider the value π_0 of π where $f(\pi)$ changes its sign from negative to positive, and there take the left and right limits

$$\vec{P}^{(\mp)} = \lim_{\pi \to \pi_0^{\mp}} \vec{P}^{(\pi)}. \tag{3.11}$$

For $f(\pi_0^+) = f(\pi_0^-) = 0$ just define $\vec{P} = \vec{P}^{(\pi_0)}$.

For $f(\pi_0^+) > f(\pi_0^-)$ define the POVM $\vec{P}$

$$\vec{P} = \frac{f(\pi_0^+)\vec{P}^{(-)} - f(\pi_0^-)\vec{P}^{(+)}}{f(\pi_0^+) - f(\pi_0^-)}. \tag{3.12}$$

In fact, one has

$$\begin{aligned} &\mathrm{Tr}[\rho_1 P_1] - \mathrm{Tr}[\rho_2 P_2] = [f(\pi_0^+) - f(\pi_0^-)]^{-1} \times \\ &\{\mathrm{Tr}[\rho_1 P_1^{(-)} - \rho_2 P_2^{(-)}] f(\pi_0^+) - \\ &\mathrm{Tr}[\rho_1 P_1^{(+)} - \rho_2 P_2^{(+)}] f(\pi_0^-)\} = 0\,, \end{aligned} \tag{3.13}$$

namely Eq. (3.5) holds. □

Notice that the value π_0 is generally not unique, since the function $f(\pi)$ can be locally constant. However, on the Hilbert space $\mathrm{Supp}(\rho_1) \cup \mathrm{Supp}(\rho_2)$, the optimal POVM for the minimax problem is unique, apart from the very degenerate case in which $D = \pi_0\rho_1 - (1-\pi_0)\rho_2$ has at least two-dimensional kernel. In fact, upon denoting by Π_+ and K the projector on the strictly positive part and the kernel of D, respectively, any Bayes optimal POVM writes ($P_1 = \Pi_+ + K'$, $P_2 = I - P_1$), with $K' \leq K$. Since for the optimal minimax POVM we need $\mathrm{Tr}[\rho_1 P_1] = \mathrm{Tr}[\rho_2 P_2]$, one obtains $\mathrm{Tr}[(\rho_1 + \rho_2)K'] = 1 - \mathrm{Tr}[(\rho_1 + \rho_2)\Pi_+]$, which has a unique solution $K' = \alpha K$ if K is a one-dimensional projector.

Corollary 3.2.4. *There are couples of mixed states for which the optimal minimax POVM is unique and non orthogonal.*

For example, consider the following states in dimension two

$$\rho_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad \rho_2 = \begin{bmatrix} \frac{1}{2} & 0 \\ 0 & \frac{1}{2} \end{bmatrix}. \tag{3.14}$$

Then an optimal minimax POVM is given by

$$P_1 = \begin{bmatrix} \frac{2}{3} & 0 \\ 0 & 0 \end{bmatrix}, \quad P_2 = \begin{bmatrix} \frac{1}{3} & 0 \\ 0 & 1 \end{bmatrix}. \tag{3.15}$$

In fact, clearly there is an optimal POVM of the diagonal form. We need to maximize $\min_{i=1,2} \mathrm{Tr}[\rho_i P_i]$, whence, according to Theorem 3.2.3, we need to maximize $\mathrm{Tr}[\rho_1 P_1]$ with the constraints $\mathrm{Tr}[\rho_1 P_1] = \mathrm{Tr}[\rho_2 P_2]$ and $P_2 = I - P_1$. Such an optimal POVM is unique, otherwise there would exists a convex combination $\pi_0\rho_1 - (1-\pi_0)\rho_2$ with kernel at least two-dimensional, which is impossible in the present example (see comments after the proof of Theorem 3.2.3). □

Notice that when the optimal POVM for the minimax strategy is unique and non-orthogonal, then there is a prior probability distribution $\vec{\pi}$ for which the optimal POVM for the Bayes problem is not unique, and the non-orthogonal POVM which optimizes the minimax problem is also optimal for the Bayes' one. In the example of remark 3.2.4 the optimal POVM (3.15) is also optimal for the Bayes problem with $\vec{\pi} = (\frac{1}{3}, \frac{2}{3})$ as one can easily check. However, in the Bayes case one can always choose an optimal orthogonal POVM, whereas in the minimax case you may have to choose a non-orthogonal POVM.

Finally, notice that, unlike in the Bayesian case, the optimal POVM for the minimax strategy may be also not extremal.

3.3 Optimal minimax discrimination of $N \geq 2$ quantum states

We now consider the easiest case of discrimination with more than two states, namely the discrimination among a covariant set. In a fully covariant state discrimination, one has a set of states $\{\rho_i\}$ with $\rho_i = U_i \rho_0 U_i^\dagger \ \forall i$, for fixed ρ_0 and $\{U_i\}$ a (projective) unitary representation of a group. In the Bayesian case full covariance requires that the prior probability distribution $\{\pi_i\}$ is uniform. Then, one can easily prove (see, for example, Ref. (Holevo, 1982)) that also the optimal POVM is covariant, namely it is of the form $P_i = U_i K U_i^\dagger$, for suitable fixed operator $K \geq 0$.

Theorem 3.3.1. *For a fully covariant state discrimination problem, there is an optimal measurement for the minimax strategy that is covariant, and coincides with an optimal Bayesian measurement.*

Proof. A covariant POVM $\{P_i\}$ gives a probability $p = \mathrm{Tr}[\rho_i P_i]$ independent of i. Moreover, there always exists an optimal Bayesian POVM that is covariant and maximizes p, which then is also the maximum over all POVM's of the average probability of correct estimation $\overline{\mathrm{Tr}[\rho_i P_i]}$ for uniform prior distribution (Holevo, 1982). Now, suppose by contradiction that there exists an optimal minimax POVM $\{P_i'\}$ maximizing $p' = \min_i \mathrm{Tr}[\rho_i P_i']$, for which $p' > p$. Then, one has $p < p' \leq \overline{\mathrm{Tr}[\rho_i P_i']}$, contradicting the assertion that an optimal Bayesian POVM maximizes $\overline{\mathrm{Tr}[\rho_i P_i]}$ over all POVM's. Therefore, $p = p'$, and the covariant Bayesian POVM also solves the minimax problem. □Notice that in the covariant case also for any optimal minimax POVM $\{P_i\}$ one has $\mathrm{Tr}[\rho_i P_i]$ independent of i, since the average probability of correct estimation is equal to the minimum one.

As an immediate consequence of Theorem 3.3.1 we derive the case of optimal discrimination of two pure states:

Corollary 3.3.2. *For two pure states the optimal POVM for the minimax discrimination is orthogonal and unique (up to trivial completion of* $\mathrm{Span}\{|\psi_i\rangle\}_{i=1,2}$ *to the full Hilbert space of the quantum system).*

Proof. Any set of two pure states $\{|\psi_i\rangle\}_{i=1,2}$ is trivially covariant under the group $\{I, U\}$ with $|\psi_2\rangle = U|\psi_1\rangle$. Then, there exists an optimal POVM for the minimax discrimination which coincides with the optimal Bayesian POVM, which is orthogonal. Uniqueness of the minimax optimal POVM follows from the assertion after Theorem 3.2.3 when restricting to the subspace spanned by the two states.

In the following we generalize Theorem 3.2.1 for two states to the case of $N \geq 2$ states and arbitrary weights. We have

Theorem 3.3.3. *For any set of states* $\{\rho_i\}_{2\leq i\leq N}$ *and any set of weights* w_{ij} *(price of misidentifying* i *with* j*) the solution of the minimax problem*

$$R_M = \inf_{\vec{P}} \sup_i \sum_j w_{ij} \operatorname{Tr}[\rho_i P_j] \tag{3.16}$$

is equivalent to the solution of the problem

$$R_M = \max_{\vec{\pi}} R_B(\pi), \tag{3.17}$$

where $R_B(\vec{\pi})$ *is the Bayesian risk*

$$R_B(\vec{\pi}) \doteq \max_{\vec{P}} \sum_i \pi_i \sum_j w_{ij} \operatorname{Tr}[\rho_i P_j]. \tag{3.18}$$

Proof. The minimax problem in Eq. (3.16) is equivalent to look for the minimum of the real function $\delta = f(\vec{P})$ over $\vec{P}$, with the constraints

$$\begin{aligned} \textstyle\sum_j w_{ij} \operatorname{Tr}[\rho_i P_j] \leq \delta, \quad & \forall i \\ P_j \geq 0, \quad & \forall j \\ \textstyle\sum_j P_j = I. & \end{aligned} \tag{3.19}$$

Upon introducing the Lagrange multipliers:

$$\begin{aligned} \mu_i \in \mathbb{R}^+, \quad & \forall i \\ 0 \leq Z_i \in M_d(\mathbb{C}), \quad & \forall i \\ Y^\dagger = Y \in M_d(\mathbb{C}), & \end{aligned} \tag{3.20}$$

$M_d(\mathbb{C})$ denoting the $d \times d$ matrices on the complex field, the problem is equivalent to

$$R_M = \inf_{\vec{P},\delta} \sup_{\vec{\mu},\vec{Z},Y}{}' \, l(\vec{P},\delta,\vec{\mu},\vec{Z},Y),$$
$$l(\vec{P},\delta,\vec{\mu},\vec{Z},Y) \doteq \delta + \sum_i [\mu_i (\sum_j w_{ij} \operatorname{Tr}[\rho_i P_j] - \delta)]$$
$$- \sum_i \operatorname{Tr}[Z_i P_i] + \operatorname{Tr}[Y(I - \sum_i P_i)], \tag{3.21}$$

where sup′ denotes the supremum over the set defined in Eqs. (3.20). The problem is convex (namely both the function δ and the constraints (3.19) are convex) and meets Slater's conditions (Boyd et Vandenberghe, 2004) (namely one can find values of $\vec{P}$ and δ such that the constraints are satisfied with strict inequalities), and hence in Eq. (3.21) one has

$$\inf_{\vec{P},\delta} \sup_{\vec{\mu},\vec{Z},Y}{}' \, l(\vec{P},\delta,\vec{\mu},\vec{Z},Y) = \max_{\vec{\mu},\vec{Z},Y}{}' \inf_{\vec{P},\delta} \, l(\vec{P},\delta,\vec{\mu},\vec{Z},Y). \tag{3.22}$$

It follows that

$$R_M = \max_{\vec{\mu},\vec{Z},Y}{}' \operatorname{Tr} Y \tag{3.23}$$

under the additional constraints

$$\sum_i \mu_i = 1\,,$$
$$\sum_i w_{ij} \mu_i \rho_i - Z_j - Y = 0\,, \qquad \forall j. \tag{3.24}$$

The constraints can be rewritten as

$$\mu_i \geq 0\,, \qquad \sum_i \mu_i = 1\,,$$
$$Y \leq \sum_i w_{ij} \mu_i \rho_i\,, \qquad \forall j. \tag{3.25}$$

Now, notice that for the Bayesian problem with prior $\vec{\pi}$, along the same reasoning, one writes the equivalent problem

$$R_B(\vec{\pi}) = \max_{Y}{}' \operatorname{Tr} Y, \tag{3.26}$$

with the constraint

$$\sum_i w_{ij} \pi_i \rho_i - Z_j - Y = 0\,, \qquad \forall j \tag{3.27}$$
$$\pi_i \geq 0\,, \qquad \sum_i \pi_i = 1\,,$$
$$Y \leq \sum_i w_{ij} \pi_i \rho_i\,, \qquad \forall j, \tag{3.28}$$

which is the same as the minimax problem, with the role of the Lagrange multipliers $\{\mu_i\}$ now played by the prior probability distribution $\{\pi_i\}$. □Clearly, a POVM that attains R_M in the minimax problem (3.16) actually exists, being the infimum over a (weakly) compact set—the POVMs' convex set—of the (weakly) continuous function $\sup_i \sum_j w_{ij} \operatorname{Tr}[\rho_i P_j]$.

3.4 Optimal minimax unambiguous discrimination

In this section we consider the so-called unambiguous discrimination of states (Ivanovic, 1987), namely with no error, but possibly with an inconclusive outcome of the measurement. We focus attention on a set of N pure states $\{\psi_i\}_{i\in\mathsf{S}}$. In such a case, it is possible to have unambiguous discrimination only if the states of the set S are linearly independent, whence there exists a biorthogonal set of vectors $\{|\omega_i\rangle\}_{i\in\mathsf{S}}$, with $\langle\omega_i|\psi_j\rangle = \delta_{ij}$, $\forall i,j \in \mathsf{S}$. We shall conveniently restrict our attention to $\mathrm{Span}\{|\psi_i\rangle\}_{i\in\mathsf{S}} \equiv \mathsf{H}$ (otherwise one can trivially complete the optimal POVM for the subspace to a POVM for the full Hilbert space of the quantum system). While in the Bayes problem the probability of inconclusive outcome is minimized, in the minimax unambiguous discrimination we need to maximize $\min_i \langle\psi_i|P_i|\psi_i\rangle$ over the set of POVM's with $\langle\psi_i|P_j|\psi_i\rangle = 0$ for $i \neq j \in \mathsf{S}$, and the POVM element that pertains to the inconclusive outcome will be given by $P_{N+1} = I - \sum_{i\in\mathsf{S}} P_i$. We have the following theorem.

Theorem 3.4.1. *The optimal minimax unambiguous discrimination of N pure states $\{\psi_i\}_{i\in\mathsf{S}}$ is achieved by the POVM*

$$\begin{aligned} P_i &= \kappa|\omega_i\rangle\langle\omega_i|, \qquad i \in \mathsf{S}, \\ P_{N+1} &= I - \sum_{i\in\mathsf{S}} P_i, \end{aligned} \tag{3.29}$$

where κ is given by

$$\kappa^{-1} = \text{max eigenvalue of } \sum_{i\in\mathsf{S}} |\omega_i\rangle\langle\omega_i|. \tag{3.30}$$

Proof. We need to maximize $\min_i \langle\psi_i|P_i|\psi_i\rangle$ over the set of POVM's with $\langle\psi_i|P_j|\psi_i\rangle = 0$ for $i \neq j \in \mathsf{S}$, whence clearly $P_j = \kappa_j|\omega_j\rangle\langle\omega_j|$. Then the problem is to maximize $\min_{i\in\mathsf{S}} \kappa_i$. This can be obtained by taking $\kappa_i = \kappa$ independent of i and then maximizing κ. In fact, if there is a $\kappa_i > \kappa_j$ for some i,j, then we can replace κ_i with κ_j, and iteratively we get $\kappa_i = \kappa$ independently of i. Finally, the maximum κ giving $P_{N+1} \geq 0$ is the one given in the statement of the theorem. □

As regards the uniqueness of the optimal POVM, we can show the following.

Theorem 3.4.2. *The optimal POVM of Theorem 3.4.1 is non-unique if and only if* $|\omega_i\rangle \in \mathrm{Supp}(P_{N+1})$ *for some* $i \in \mathsf{S}$.

Proof. In fact, if there exists an $i \in \mathsf{S}$ such that $|\omega_i\rangle \in \mathrm{Supp}(P_{N+1})$, this means that there exists $\varepsilon > 0$ such that $\varepsilon|\omega_i\rangle\langle\omega_i| \leq P_{N+1}$. Then the following is a POVM

$$\begin{aligned} Q_j &= P_j, \quad \text{for } j \neq i \\ Q_i &= P_i + \varepsilon|\omega_i\rangle\langle\omega_i|, \\ Q_{N+1} &= P_{N+1} - \varepsilon|\omega_i\rangle\langle\omega_i|, \end{aligned} \tag{3.31}$$

and is optimal as well. Conversely, if there exists another equivalently optimal POVM $\{Q_j\}$, then there exists an $i \in \mathsf{S}$ such that $Q_i > P_i$ (since both are proportional to $|\omega_i\rangle\langle\omega_i|$, and $\min_i\langle\psi_i|Q_i|\psi_i\rangle$ has to be maximized). Then $|\omega_i\rangle \in \mathrm{Supp}(P_{N+1})$. □

When the optimal POVM according to Theorem 3.4.2 is not unique, one can refine the optimality criterion in the following way. Define the set $\mathsf{S}_1 \subset \mathsf{S}$ for which one has $|\omega_i\rangle \in \mathrm{Supp}(P_{N+1})$. Denote by $\mathfrak{P}_1$ the set of POVM's which are equivalently optimal to those of Theorem 3.4.1. Then define the set of POVM's $\mathfrak{P}_2 \subset \mathfrak{P}_1$ which maximizes $\min_{i\in\mathsf{S}_1}\langle\omega_i|P_i|\omega_i\rangle$. In this way one iteratively reach a unique optimal POVM, which is just the one given in Eqs. (3.29) and (3.30).

3.5 Bayesian discrimination of two Pauli channels

The problem of optimally discriminating two quantum operations $\mathcal{E}_1$ and $\mathcal{E}_2$ can be reformulated into the problem of finding the state ρ in the input Hilbert space $\mathcal{H}$, such that the error probability in the discrimination of the output states $\mathcal{E}_1(\rho)$ and $\mathcal{E}_2(\rho)$ is minimal. The possibility of exploiting entanglement with an ancillary system can increase the distinguishability of the output states (Sacchi, 2005a). In this case the output states to be discriminated will be of the form $(\mathcal{E}_1 \otimes \mathcal{I}_\mathcal{K})\rho$ and $(\mathcal{E}_2 \otimes \mathcal{I}_\mathcal{K})\rho$, where the input ρ is generally a bipartite state of $\mathcal{H} \otimes \mathcal{K}$, and the quantum operations act just on the first party whereas the identity map $\mathcal{I}_\mathcal{K}$ acts on the second.

We now make use of the expression for the Bayesian risk of discrimination between states (3.2). Upon denoting with $\mathcal{R}'_B(\pi)$ the minimal error probability when a strategy without ancilla is adopted, one has

$$\mathcal{R}'_B(\pi) = \frac{1}{2}\left(1 - \max_{\rho\in\mathcal{H}} \|\pi_1\mathcal{E}_1(\rho) - \pi_2\mathcal{E}_2(\rho)\|_1\right) . \tag{3.32}$$

On the other hand, by allowing the use an ancillary system, we have

$$\mathcal{R}_B(\pi) = \frac{1}{2}\left(1 - \max_{\xi\in\mathcal{H}\otimes\mathcal{K}} \|\pi_1(\mathcal{E}_1\otimes\mathcal{I})\xi - \pi_2(\mathcal{E}_2\otimes\mathcal{I})\xi\|_1\right) . \tag{3.33}$$

The maximum of the trace norm in Eq. (3.33) with the supremum over the dimension of $\mathcal{K}$ is equivalent to the norm of complete boundedness (Paulsen, 1987) of the map $\pi_1\mathcal{E}_1 - \pi_2\mathcal{E}_2$, and in fact for finite-dimensional Hilbert space the supremum is achieved for $\dim(\mathcal{K}) = \dim(\mathcal{H})$ (Paulsen, 1987), and in the following we shall drop the subindex $\mathcal{K}$ from the identity map. Moreover, due to linearity of quantum operations and convexity of the trace norm, the maximum in both Eqs. (3.32) and (3.33) is achieved on pure states.

Clearly, $\mathcal{R}_B(\pi) \leq \mathcal{R}'_B(\pi)$. In the case of discrimination between two unitary transformations U and V (Childs et al., 2000b), one has $\mathcal{R}_B(\pi) = \mathcal{R}'_B(\pi)$, namely there is no need of entanglement with an ancillary system to achieve the ultimate minimum error probability, which is given by

$$\begin{aligned}\mathcal{R}_B(\pi) &= \min_{|\psi\rangle\in\mathcal{H}} \frac{1}{2}\left(1 - \sqrt{1 - 4\pi_1\pi_2|\langle\psi|U^\dagger V|\psi\rangle|^2}\right) \\ &= \frac{1}{2}\left(1 - \sqrt{1 - 4\pi_1\pi_2 D^2}\right) ,\end{aligned} \tag{3.34}$$

where D is the distance between 0 and the polygon in the complex plane whose vertices are the eigenvalues of $U^\dagger V$.

In the case of discrimination of two Pauli channels for qubits, namely

$$\mathcal{E}_i(\rho) = \sum_{\alpha=0}^{3} q_\alpha^{(i)} \sigma_\alpha\rho\sigma_\alpha \qquad i = 1, 2 , \tag{3.35}$$

where $\sum_{\alpha=0}^{3} q_\alpha^{(i)} = 1$, $\sigma_0 = I$, and $\{\sigma_1 , \sigma_2 , \sigma_3\} = \{\sigma_x , \sigma_y , \sigma_z\}$ denote the customary spin Pauli matrices, the minimal error probability can be achieved by using a maximally entangled input state, and one obtains (Sacchi, 2005a)

$$\mathcal{R}_B(\pi) = \frac{1}{2}\left(1 - \sum_{\alpha=0}^{3} |r_\alpha|\right) , \tag{3.36}$$

with

$$r_\alpha = \pi_1 q_\alpha^{(1)} - p_2 q_\alpha^{(2)} = \pi(q_\alpha^{(1)} + q_\alpha^{(2)}) - q_\alpha^{(2)} , \tag{3.37}$$

where we fixed the *prior* $\pi = \pi_1$ and $\pi_2 = 1 - \pi_1$. For a strategy with no ancillary assistance one has (Sacchi, 2005a)

$$\mathcal{R}'_B(\pi) = \frac{1}{2}(1 - C) , \tag{3.38}$$

where

$$C = \max\left\{|r_0 + r_3| + |r_1 + r_2|\,, |r_0 + r_1| + |r_2 + r_3|\,, |r_0 + r_2| + |r_1 + r_3|\right\}\ , \quad (3.39)$$

and the three cases inside the brackets corresponds to using an eigenstate of σ_z, σ_x, and σ_y, respectively, as the input state of the channel. More generally, for pure input state $\rho = \frac{1}{2}(I + \vec{\sigma}\cdot\vec{n})$, with $\vec{n} = (\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta)$, the Bayes risk for discriminating the outputs will be (Sacchi, 2005a,b)

$$\mathcal{R}'_B(\pi, \vec{\sigma}\cdot\vec{n}) = \frac{1}{2}\left(1 - \max\left\{|a+b|, \sqrt{\cos^2\theta(a-b)^2 + \sin^2\theta(c^2 + d^2 + 2cd\cos(2\phi))}\right\}\right) \quad (3.40)$$

with $a = r_0 + r_3$, $b = r_1 + r_2$, $c = r_0 - r_3$, and $d = r_1 - r_2$. Notice that the term $|a+b| = |2\pi - 1|$ corresponds to the trivial guessing $\{\mathcal{E}_1$ if $\pi_1 = \pi > 1/2\,, \mathcal{E}_2$ if $\pi < 1/2\}$.

We can also rewrite Eq. (3.38) as

$$\mathcal{R}'_B(\pi) = \min_{i=1,2,3} \mathcal{R}'_B(\pi, \sigma_i)\ . \quad (3.41)$$

From Eqs. (3.36–3.39) one can see that entanglement is not needed to achieve the minimal error probability as long as $C = \sum_{i=0}^{3} |r_i|$, which is equivalent to the condition $\Pi_{i=0}^{3} r_i \geq 0$. On the other hand, we can find instances where the channels can be perfectly discriminated only by means of entanglement, for example in the case of two channels of the form

$$\mathcal{E}_1(\rho) = \sum_{\alpha\neq\beta} q_\alpha \sigma_\alpha \rho \sigma_\alpha\ , \qquad \mathcal{E}_2(\rho) = \sigma_\beta \rho \sigma_\beta\ , \quad (3.42)$$

with $q_\alpha \neq 0$, and arbitrary *a priori* probability.

3.6 Minimax discrimination of Pauli channels

As in the Bayesian approach, the minimax discrimination of two channels consists in finding the optimal input state such that the two possible output states are discriminated with minimum risk. Again, we will consider the two cases with and without ancilla, upon defining

$$\mathcal{R}_M = \min_{\xi\in\mathcal{H}\otimes\mathcal{K}} R_M((\mathcal{E}_1\otimes\mathcal{I})(\xi), (\mathcal{E}_2\otimes\mathcal{I})(\xi))\ ,$$
$$\mathcal{R}'_M = \min_{\rho\in\mathcal{H}} R_M(\mathcal{E}_1(\rho), \mathcal{E}_2(\rho))\ , \quad (3.43)$$

where $R_M(\rho_1, \rho_2)$ is given in Eq. (3.17). Since for all $\vec{M}$, ρ, and π, one has

$$\begin{aligned} &\max\{\mathrm{Tr}[(\mathcal{E}_1 \otimes \mathcal{I})(\rho)M_2], \mathrm{Tr}[(\mathcal{E}_2 \otimes \mathcal{I})(\rho)M_1]\} \\ &\geq \pi \,\mathrm{Tr}[(\mathcal{E}_1 \otimes \mathcal{I})(\rho)M_2] + (1-\pi)\,\mathrm{Tr}[(\mathcal{E}_2 \otimes \mathcal{I})(\rho)M_1]\,, \end{aligned} \tag{3.44}$$

then $\mathcal{R}_M \geq \mathcal{R}_B(\pi)$ for all π. Analogously, $\mathcal{R}'_M \geq \mathcal{R}'_B(\pi)$ for all π.

Theorems 3.2.3 and 3.3.3 can be immediately applied to state that the minimax discrimination of two unitaries is equivalent to the Bayesian one. In fact, the optimal input state in the Bayesian problem which achieves the minimum error probability of Eq. (3.34) does not depend on the *a priori* probabilities. Therefore it is also optimal for the minimax problem and there is no need of entanglement [and the minimax risk $\mathcal{R}_M$ will be equivalent to the Bayes risk $\mathcal{R}_B(1/2)$].

Let us now consider the problem of discriminating the Pauli channels of Eq. (3.35) in the minimax framework. In the following theorem, we show that an (arbitrary) maximally entangled state always allows to achieve the optimal minimax discrimination as in the Bayesian problem.

Theorem 3.6.1. *The minimax risk $\mathcal{R}_M$ for the discrimination of two Pauli channels can be achieved by using an arbitrary maximally entangled input state. Moreover, the minimax risk is then the Bayes risk for the worst* a priori *probability:*

$$\mathcal{R}_M = \max_{\pi} \mathcal{R}_B(\pi)\,. \tag{3.45}$$

Proof. Let us discriminate between the states $\rho_i = (\mathcal{E}_i \otimes \mathcal{I})(\xi^e)$, where ξ^e is a maximally entangled state. By Theorem 3.2.1 there are *a priori* probabilities $(\pi_*, 1-\pi_*)$ whose optimal Bayes measurement fulfills

$$\mathrm{Tr}[\rho_1 P_1] = \mathrm{Tr}[\rho_2 P_2]\,. \tag{3.46}$$

Since the input state ξ^e is always optimal in the Bayes problem we infer $\mathcal{R}_B(\pi_*) = \mathrm{Tr}[\rho_1 P_2]$, and moreover $R_M(\rho_1, \rho_2) = \mathcal{R}_B(\pi_*)$. Now, one has also $\mathcal{R}_M = R_M(\rho_1, \rho_2)$, since if it would not be true, then there would be an input state ρ and a measurement $\vec{M}$ for which $\max\{\mathrm{Tr}[(\mathcal{E}_1 \otimes \mathcal{I})(\rho)M_2], \mathrm{Tr}[(\mathcal{E}_2 \otimes \mathcal{I})(\rho)M_1]\} < R_B(\pi_*)$, and hence $\pi_* \,\mathrm{Tr}[(\mathcal{E}_1 \otimes \mathcal{I})(\rho)M_2] + (1-\pi_*)\,\mathrm{Tr}[(\mathcal{E}_2 \otimes \mathcal{I})(\rho)M_1] < R_B(\pi_*)$, which is a contradiction. Equation (3.45) simply comes from the relation $\mathcal{R}_M \geq \mathcal{R}_B(\pi)$ for all π, along with $\mathcal{R}_M = \mathcal{R}_B(\pi_*)$. □

Notice the nice correspondence between Eqs. (3.17) and (3.45). Theorem 3.6.1 holds true also in the case of generalized Pauli channels in higher dimension, since entangled states again achieve the optimal Bayesian discrimination, whatever the *a priori* probability (Sacchi, 2005a). More generally, Eq. (3.45) will hold in the

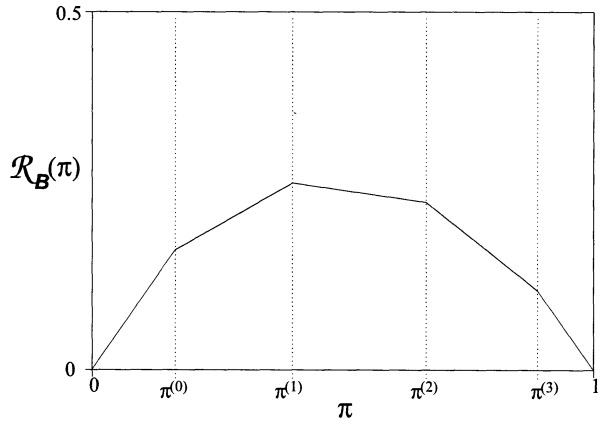

Figure 3.1: The optimal Bayes risk $\mathcal{R}_B(\pi)$ in the discrimination of two Pauli channels versus the *a priori* probability π will usually look like this. Notice that the rightmost and leftmost segments have slope 1 and (-1), respectively. The minimal risk for the minimax discrimination corresponds to $\mathcal{R}_M = \max_\pi \mathcal{R}_B(\pi)$, and is achieved at one of the breakpoints $\pi^{(\alpha)}$.

discrimination of any couple of quantum operations for which the minimal Bayes risk $\mathcal{R}_B(\pi)$ can be achieved by the same input state for any π.

Now we establish some visual images on which to read the minimax risks. We must look at the function $\mathcal{R}_B(\pi)$ given in Eq. (3.36) drawn on $[0,1]$. By Eq. (3.45), we know that its maximum is $\mathcal{R}_M$. As the r_α defined in (3.37) are increasing affine functions of π, their absolute value is a convex piecewise affine function, and hence $\mathcal{R}_B(\pi)$ is a concave piecewise affine function (see Fig. 3.1). The four breakpoints correspond to the four values of π for which each r_α vanishes. We define $t_\alpha = q_\alpha^{(1)} + q_\alpha^{(2)}$ as the slopes of the functions r_α and $\pi^{(\alpha)} = q_\alpha^{(2)}/\, t_\alpha$ as the value of π for which $r_\alpha = 0$. We denote by π_* the point at which $\mathcal{R}_B(\pi)$ reaches its maximum (the maximum will be attained at one of the breakpoints $\pi^{(\alpha)}$). We also reorder the index α such that $\pi^{(0)} \leq \pi^{(1)} \leq \pi^{(2)} \leq \pi^{(3)}$. In this way, $\mathcal{R}_B(\pi)$ rewrites

$$\mathcal{R}_B(\pi) = \frac{1}{2}\left(1 - \sum_{\alpha=0}^{3} t_\alpha |\pi - \pi^{(\alpha)}|\right) . \tag{3.47}$$

Let us now look at the discrimination strategy without any ancillary system. Another picture, that should be superimposed on Fig. 3.6, is the Bayes risk $\mathcal{R}'_B(\pi)$ of Eq. (3.38) versus π for the strategy with no ancillary system. One can see that $\mathcal{R}'_B(\pi)$ is the minimum of the three piecewise affine functions $\mathcal{R}'_B(\pi, \sigma_x)$, $\mathcal{R}'_B(\pi, \sigma_y)$, $\mathcal{R}'_B(\pi, \sigma_z)$, corresponding to the Bayes risks when sending an eigenstate of the Pauli matrices. Here again $\mathcal{R}'_B(\pi)$ is the minimum of concave functions, so it is concave as well, and the maximum will be attained at a breakpoint $\pi = \pi'_*$ (see Fig. 3.6).

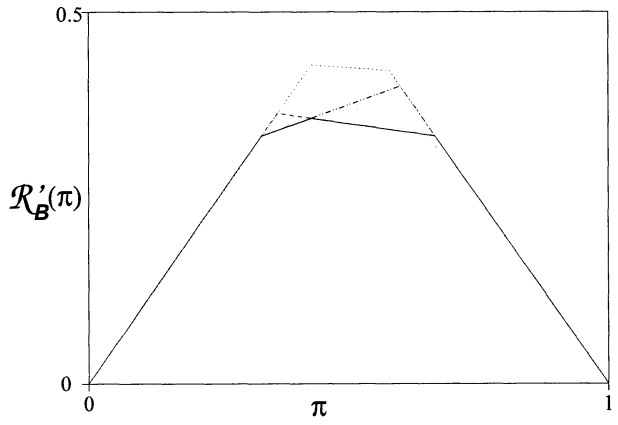

Figure 3.2: An example for the Bayes risks $\mathcal{R}'_B(\pi, \sigma_i)$ with $i = x, y, z$ versus the *a priori* probability π, for discrimination without ancilla. Each of the three different dotted lines correspond to the Bayes risk $\mathcal{R}'_B(\pi, \sigma_i)$ when sending an eigenstate of the Pauli matrix σ_i through the channel. The solid line is the optimal Bayes risk $\mathcal{R}'_B(\pi)$ without ancillary assistance, and corresponds at any π to the minimum of the three $\mathcal{R}'_B(\pi, \sigma_i)$. The minimal risk for the minimax discrimination with no ancilla corresponds to $\mathcal{R}'_M = \max_\pi \mathcal{R}'_B(\pi)$, and is achieved at one of the breakpoints of $\mathcal{R}'_B(\pi)$.

To "read" more on these pictures, once again we prove that the optimal minimax risk $\mathcal{R}'_M$ for discrimination without ancilla corresponds to the optimal Bayes risk without ancilla for the worst *a priori* probability π'_*:

Theorem 3.6.2. *The optimal minimax discrimination with no ancilla is equivalent to the solution of the problem*

$$\mathcal{R}'_M = \max_\pi \mathcal{R}'_B(\pi) \equiv \mathcal{R}'_B(\pi'_*) \; . \tag{3.48}$$

Proof. Notice again the similarity between equations (3.17), (3.45) and (3.48). For any ρ one has

$$R_M(\mathcal{E}_1(\rho), \mathcal{E}_2(\rho)) \geq \mathcal{R}'_M \geq \max_\pi \mathcal{R}'_B(\pi) \; . \tag{3.49}$$

If we find an input state $\rho_{\vec{n}} = \frac{1}{2}(I + \vec{\sigma} \cdot \vec{n})$ such that

$$\max_\pi \mathcal{R}'_B(\pi) = \max_\pi \mathcal{R}'_B(\pi, \vec{\sigma} \cdot \vec{n}) \tag{3.50}$$

from Eq. (3.17) of Theorem 3.3.3 it follows that

$$R_M(\mathcal{E}_1(\rho_{\vec{n}}), \mathcal{E}_2(\rho_{\vec{n}})) = \max_\pi \mathcal{R}'_B(\pi, \vec{\sigma} \cdot \vec{n}) \; , \tag{3.51}$$

which, along with Eqs. (3.49) and (3.50), provides the proof. Moreover, $\rho_{\vec{n}}$ will be the optimal input state for the minimax discrimination without ancilla.

Now we have just to find a state such that condition (3.50) holds. We already noticed that π'_* is a breaking point of $\mathcal{R}'_B(\pi)$. Either this breakpoint is also a breakpoint (and the maximum) of $\mathcal{R}'_B(\pi, \sigma_i)$ for some $i \in x, y, z$, or else at least two of the $\mathcal{R}'_B(\pi, \sigma_i)$ are crossing in π'_*, one increasing and the other decreasing (Fig. 3.6). In the first case Eq. (3.50) is immediately satisfied, and an eigenstate of σ_i will be the optimal input state. In the second case, we show that when two $\mathcal{R}'_B(\pi, \sigma_i)$ are crossing at π'_* we can find a state $\rho_{\vec{n}}$ such that

$$\begin{aligned} &\mathcal{R}'_B(\pi'_*, \vec{\sigma}\cdot\vec{n}) = \mathcal{R}'_B(\pi'_*, \sigma_i) \,, \\ &\partial_\pi \mathcal{R}'_B(\pi, \vec{\sigma}\cdot\vec{n})|_{\pi=\pi'_*} = 0 \,, \end{aligned} \tag{3.52}$$

and therefore has the maximum at π'_* by concavity. In fact, the crossing, and therefore non-equality of the $\mathcal{R}'_B(\pi, \sigma_i)$ in a neighborhood of π'_*, implies that for each of the two $\mathcal{R}'_B(\pi, \sigma_i)$, the maximum in (3.40) for π'_* is attained by the square root term (since the term $|a+b|$ is just a function of π). Let us assume that the σ_i that give such a crossing are σ_x and σ_y. Then looking at (3.40), we have at point π'_*

$$\begin{aligned} &|c+d| = |c-d| \,, \\ &\partial_\pi|c+d|\,\partial_\pi|c-d| < 0 \end{aligned} \tag{3.53}$$

(notice that all functions are linear, i.e. differentiable in π'_*). Indeed, the first of Eqs. (3.53) implies that any linear combination of eigenstate of σ_x and σ_y satisfies the first of Eqs. (3.52). By taking an input state with $\theta = \pi/2$ and ϕ such that

$$\tan^2\phi = -\left.\frac{\partial_\pi|c+d|}{\partial_\pi|c-d|}\right|_{\pi=\pi'_*} , \tag{3.54}$$

the second equation in (3.52) is satisfied as well. Similarly, if the σ_i are σ_z, σ_x one can take the input state with $\phi = 0$ or π and θ such that

$$\tan^2\theta = -\left.\frac{\partial_\pi|a-b|}{\partial_\pi|c+d|}\right|_{\pi=\pi'_*} . \tag{3.55}$$

Finally, for σ_z, σ_y one has $\phi = \pm\pi/2$ and

$$\tan^2\theta = -\left.\frac{\partial_\pi|a-b|}{\partial_\pi|c-d|}\right|_{\pi=\pi'_*} . \tag{3.56}$$

As a remark, no eigenstate of σ_i for $i = x, y, z$ can be an optimal input in the minimax sense in this situation. This is a typical result of the minimax discrimination. As in

the case of discrimination of states, when the correspondent Bayes problem presents a kind of degeneracy and have multiple solutions, in the minimax problem the degeneracy is partially or totally removed. In the present situation, if we have the maximum of $\mathcal{R}'_B(\pi)$ at the crossing point of exactly two $\mathcal{R}'_B(\pi, \sigma_i)$, one increasing and the other decreasing, we find just four optimal input states: two non-orthogonal states and their respective orthogonal states. We shall give an explicit example at the end of the section. □

If we want to find in which case entanglement is not necessary for optimal minimax discrimination, then we have just to characterize when $\mathcal{R}'_B(\pi'_*) = \mathcal{R}_B(\pi_*)$. We already noticed that we can choose π_* to be one of the $\pi^{(\alpha)}$. The corresponding r_α is zero, and hence $C = \sum_\alpha |r_\alpha|$, namely $\mathcal{R}'_B(\pi_*) = \mathcal{R}_B(\pi_*)$. Since one has

$$\mathcal{R}'_B(\pi'_*) = \mathcal{R}'_M \geq \mathcal{R}_M = \mathcal{R}_B(\pi_*) = \mathcal{R}'_B(\pi_*) \,, \tag{3.57}$$

we only have to check that π_* is a maximum of $\mathcal{R}'_B(\pi)$, recalling that the function is concave (see Fig. 3.6).

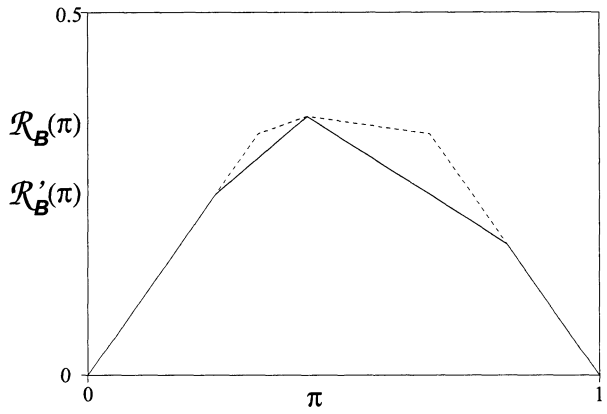

Figure 3.3: Optimal Bayes risks versus the *a priori* probability π for the discrimination of the Pauli channels with parameters given in Eq. (3.64). The solid line gives $\mathcal{R}_B(\pi)$ for an entanglement-assisted strategy; the dotted lines gives $\mathcal{R}'_B(\pi)$ for strategy without ancilla. The minimal risk in the optimal minimax discrimination corresponds in both strategies to $\mathcal{R}'_M = \max_\pi \mathcal{R}'_B(\pi) = \max_\pi \mathcal{R}_B(\pi) = \mathcal{R}_M$, namely there is no need of an ancillary system.

Ultimately, we shall have to list down cases. Reading them might be clearer with the quantities appearing in Eqs. (3.36–3.39) explicitly written as a function of π. The most useful segmentation of $[0, 1]$ is based on the $\pi^{(\alpha)}$, that is the points where the r_α vanish, and $\mathcal{R}_B(\pi)$ breaks. Recall that $r_\alpha = t_\alpha(\pi - \pi^{(\alpha)})$, and $r_\alpha > 0$ for

$\pi > \pi^{(\alpha)}$. As we have four α, we have five segments (they may get degenerated). Remember that knowing C in Eq. (3.39) and $\sum_\alpha |r_\alpha|$ is tantamount to knowing $\mathcal{R}'_B(\pi)$ or $\mathcal{R}_B(\pi)$. Here is a list of the signs of the r_α and the value of C on each open segment (so that all $r_\alpha \neq 0$):

- $(0, \pi^{(0)})$: $\sum_\alpha |r_\alpha| = -\sum_\alpha r_\alpha = C$. Notice that $\mathcal{R}'_B(\pi) = \mathcal{R}_B(\pi)$ and that their common slope is 1.
- $(\pi^{(0)}, \pi^{(1)})$: $\sum_\alpha |r_\alpha| = r_0 - r_1 - r_2 - r_3$, so that $C = r_0 - r_1 - r_2 - r_3 - 2\inf_{\alpha=1,2,3} |r_\alpha|$. On this segment, $\mathcal{R}'_B(\pi) > \mathcal{R}_B(\pi)$.
- $(\pi^{(1)}, \pi^{(2)})$: $\sum_\alpha |r_\alpha| = r_0 + r_1 - r_2 - r_3 = C$, so that $\mathcal{R}'_B(\pi) = \mathcal{R}_B(\pi)$.
- $(\pi^{(2)}, \pi^{(3)})$: $\sum_\alpha |r_\alpha| = r_0 + r_1 + r_2 - r_3$, so that $C = r_0 + r_1 + r_2 - r_3 - 2\inf_{\alpha=0,1,2} r_\alpha$ and $\mathcal{R}'_B(\pi) > \mathcal{R}_B(\pi)$.
- $(\pi^{(3)}, 1)$: $\sum_\alpha |r_\alpha| = \sum_\alpha r_\alpha = C$ and $\mathcal{R}'_B(\pi) = \mathcal{R}_B(\pi)$. Their common slope is (-1).

A close look at these expressions, as we shall show in the following, proves that $\mathcal{R}'_B(\pi)$ is derivable at $\pi^{(\alpha)}$ unless there is $\beta \neq \alpha$ such that $\pi^{(\alpha)} = \pi^{(\beta)}$. With this in mind, we see that π_* cannot be a maximum of $\pi^{(\alpha)}$ unless several r_α are null at the same point (with supplementary conditions) or $\pi_* = \pi^{(1)}$ and the segment $(\pi^{(1)}, \pi^{(2)})$ is flat. Here is the full-fledged study, using repeatedly the list above. It is complete as any other case can be handled by symmetry (switching channels, that is mapping π to $1 - \pi$).

- $\pi_* = \pi^{(0)} < \pi^{(1)}$: At $\pi^{(0)}$, we have $r_0 = 0$ and $r_\alpha < 0$ for $\alpha \neq 0$. So that $\inf_\alpha |r_\alpha| = |r_0|$ on a neighborhood of $\pi^{(0)}$. On that neighborhood, we deduce $C = -\sum_\alpha r_\alpha$, and hence $\partial_\pi \mathcal{R}'_B(\pi)|_{\pi=\pi^{(0)}} = 1$, so that $\pi^{(0)}$ is not a maximum of $\mathcal{R}'_B(\pi)$. Entanglement is then necessary for optimal discrimination.
- $\pi_* = \pi^{(0)} = \pi^{(1)} < \pi^{(2)}$: On $(0, \pi^{(0)}) \cup (\pi^{(1)}, \pi^{(2)})$, equality $\mathcal{R}'_B(\pi) = \mathcal{R}_B(\pi)$ holds. Thus, the two functions are equal on a neighborhood of π_*, and since π_* is a (local) maximum of $\mathcal{R}_B(\pi)$, it is also a local maximum of $\mathcal{R}'_B(\pi)$. In this case an unentangled strategy is then as efficient as any entangled one.
- $\pi_* = \pi^{(0)} = \pi^{(1)} = \pi^{(2)} < \pi^{(3)}$: The risk $\mathcal{R}'_B(\pi)$ is nondecreasing on the left of π_* (slope 1), we then want it to be non-increasing on a right neighborhood of π_*. Now this is part of the segment $(\pi^{(2)}, \pi^{(3)})$, where $C = r_0 + r_1 + r_2 - r_3 - 2\inf_{\alpha=0,1,2} r_\alpha$. Recall that $r_\alpha = t_\alpha(\pi - \pi^{(\alpha)})$. Since $r_\alpha = 0$ for $\alpha \neq 3$ at π_*, and they are all nondecreasing, $\inf_{\alpha=0,1,2} r_\alpha$ is the one with the smallest slope t_α. It

follows that the slope of $\mathcal{R}'_B(\pi)$ on the right of π_* is $t_3 - t_0 - t_1 - t_2 + 2\inf_{\alpha=0,1,2} t_\alpha$, and so entanglement is not needed if and only if

$$t_3 + 2 \inf_{\alpha=0,1,2} t_\alpha \leq \sum_{\alpha=0,1,2} t_\alpha \tag{3.58}$$

- $\pi_* = \pi^{(0)} = \pi^{(1)} = \pi^{(2)} = \pi^{(3)}$: This is the trivial case where both channels are the same. Of course, entanglement is useless.

- $\pi^{(0)} < \pi_* = \pi^{(1)} < \pi^{(2)}$: In this case $\mathcal{R}'_B(\pi)$ is derivable at π_*. Indeed, on $(\pi^{(1)}, \pi^{(2)})$, we have $C = r_0 + r_1 - r_2 - r_3$ whereas on $(\pi^{(0)}, \pi^{(1)})$, $C = r_0 - r_1 - r_2 - r_3 - 2\inf_{\alpha=1,2,3} |r_\alpha|$. In a neighborhood of π_*, one has $\inf_{\alpha=1,2,3} |r_\alpha| = r_1$, as it is the only one which is 0 at π_*; hence $C = r_0 + r_1 - r_2 - r_3$ also on a left neighborhood of π_* and the slope of $\mathcal{R}'_B(\pi)$ at π_* is $t_3 + t_2 - t_1 - t_0$. Since $\pi*$ is a maximum if and only if this slope is null, we get the condition

$$t_0 + t_1 = t_2 + t_3 \ . \tag{3.59}$$

- $\pi^{(0)} < \pi_* = \pi^{(1)} = \pi^{(2)} < \pi^{(3)}$: On the left of π_*, we are on the segment $(\pi^{(0)}, \pi^{(1)})$, so that $C = r_0 - r_1 - r_2 - r_3 - 2\inf_{\alpha=1,2,3} |r_\alpha|$. On the right, we are on the segment $(\pi^{(2)}, \pi^{(3)})$ and $C = r_0 + r_1 + r_2 - r_3 - 2\inf_{\alpha=0,1,2} r_\alpha$. In a neighborhood of π_*, the r_α with the smallest absolute value will be either r_1 or r_2 (more precisely, the one with the smallest slope t_α), so that we can write in a neighborhood of π_* for both sides $C = r_0 - r_3 + |r_2 - r_1|$. The slope of $\mathcal{R}'_B(\pi)$ is then $t_3 - t_0 + |t_2 - t_1|$ and $t_3 - t_0 - |t_2 - t_1|$ on the left and on the right of π_*, respectively. Entanglement is not necessary when π_* is a maximum of $\mathcal{R}'_B(\pi)$, and hence we get the necessary and sufficient condition

$$|t_0 - t_3| \leq |t_1 - t_2| \ . \tag{3.60}$$

We can summarize the above discussion as follows

Theorem 3.6.3. *The minimax risk without using ancilla is strictly greater than the minimax risk using entanglement, except in the following cases:*

- *the trivial situation where both channels are the same, so that* $\pi_* = \pi^{(\alpha)} = \frac{1}{2}$ *for all* α.

- *if* $\pi_* = \pi^{(0)} \leq \pi^{(1)} < \pi^{(2)}$

- *if* $\pi_* = \pi^{(0)} = \pi^{(1)} = \pi^{(2)} < \pi^{(3)}$ *and*

$$t_3 + 2 \inf_{\alpha=0,1,2} t_\alpha \leq \sum_{\alpha=0,1,2} t_\alpha \tag{3.61}$$

- *if* $\pi^{(0)} < \pi_* = \pi^{(1)} < \pi^{(2)}$ *and*

$$t_0 + t_1 = t_2 + t_3 \tag{3.62}$$

- *if* $\pi^{(0)} < \pi_* = \pi^{(1)} = \pi^{(2)} < \pi^{(3)}$ *and*

$$|t_0 - t_3| \leq |t_1 - t_2| \tag{3.63}$$

- *The symmetric cases (obtained by exchanging channels 1 and 2, i.e. exchanging indexes 0 and 1 with 3 and 2, respectively, both in* $\pi^{(\alpha)}$ *and* t_α*.*

Differently from the Bayesian result, we notice that when entanglement is not necessary to achieve the optimal minimax discrimination, the optimal input state may not be an eigenstate of the Pauli matrices. Consider, for example, the two Pauli channels featured in Fig. 3.6 that correspond to the parameters

$$\begin{array}{llll} q_0^{(1)} = 0.3 & q_1^{(1)} = 0.4 & q_2^{(1)} = 0.2 & q_3^{(1)} = 0.1 \\ q_0^{(2)} = 0.1 & q_1^{(2)} = 0.3 & q_2^{(2)} = 0.15 & q_3^{(2)} = 0.45 \end{array} \tag{3.64}$$

We can compute $\pi^{(\alpha)} = q_\alpha^{(2)}/(q_\alpha^{(1)} + q_\alpha^{(2)})$ and get $\pi^{(\alpha)} = (1/4, 3/7, 3/7, 9/11)$. Here $\pi_* = 3/7$, and we are in the situation of Eq. (3.63), since $t_\alpha = (q_\alpha^{(1)} + q_\alpha^{(2)}) = (0.4, 0.7, 0.35, 0.55)$. Hence, entanglement is not necessary to achieve the optimal minimax risk, but the state to be used is not an eigenstate of the Pauli matrices. In fact, we are in the case of the proof of Theorem 3.6.2, where $\mathcal{R}'_B(\pi, \sigma_x)$ and $\mathcal{R}'_B(\pi, \sigma_y)$ are crossing in π_*. The optimal input state for the minimax discrimination will be given by $\theta = \pi/2$ and ϕ as in Eq. (3.54), which gives $\tan^2\phi = 2/5$. Then, we have four optimal input states that lie on the equator of the Bloch sphere, with $\vec{n} = (\pm\sqrt{5/7}, \pm\sqrt{2/7}, 0)$.

Chapitre 4

Fast estimation of unitary operations

Ce chapitre dérive de l'article (Kahn, 2007b).

Résumé : NOus donnons une procédure explicite, basée sur un état intriqué en entrée, pour estimer une opération U dans $SU(d)$, dont le taux de convergence est de $1/N^2$ quand on envoie N particules dans l'appareil. Nous prouvons l'optimalité de ce taux. Nous évaluons également la constante C telle que le taux asymptotique soit C/N^2. Toutefois, d'autres stratégies pourraient permettre d'obtenir une meilleure constante C.

4.1 Introduction

The question that we are investigating in this chapter is: "What is the best way of estimating a unitary operation U?"

By "unitary operation", we mean a device (or a *channel*) that sends a density operator ρ_0 on $\mathbb{C}^d$ to another density operator $\rho = U\rho_0 U^*$, where $U \in SU(d)$, a special unitary matrix.

We immediately stress that the solution to this estimation problem can be divided into two parts: what is the input state, and which measurement (POVM) to apply on the output state? Indeed, in order to estimate the channel U, we have to let it act on a state (the input state). And once we have the output state, the problem consists in discriminating states in the family of possible output states.

This estimation of unitary operation has been extensively studied over the last few years.

The first invitation was (Childs et al., 2000a), featuring numerous special cases. In most of those, the unitary U is known to belong to some subset of $SU(2)$.

Then Acin et al. (2001) provided the form of an optimal state to be sent in with non-specified coefficients depending on the cost function (we give the formula of this state in equation (4.2)). In that paper the authors consider the situation where the unitary operation is performed independently on N systems. That study applied to any $SU(d)$, and any covariant loss function, in particular fidelity, in a Bayesian framework. The proposed input state uses an ancilla, that is an auxiliary system that is not sent through the unitary channel with Hilbert space $(\mathbb{C}^d)^{\otimes N}$. The state is prepared as a superposition of maximally entangled states, one for each irreducible representation of $SU(d)$ appearing in $(\mathbb{C}^d)^{\otimes n}$. We emphasize that the state is an entangled state of $(\mathbb{C}^d)^{\otimes N} \otimes (\mathbb{C}^d)^{\otimes N}$: we do not send N copies of an entangled state through the device, but all the N systems that are sent through the channel together with the N particles of the ancilla are part of the same entangled state, yielding the most general possible strategy. There was no evaluation of the rate of convergence, though.

Subsequent works mainly focused on $SU(2)$, as the case is simpler and yields many applications, e.g. transmission of reference frames in quantum communication. Indeed, the latter is equivalent to the estimation of a $SU(2)$ operation. The first strategy to be proved to converge (in fidelity) at $1/N^2$ rate was not covariant (Peres, 1993). It made no use of an ancilla. Later, Bagan et al. (2004a) achieved the same rate for a covariant measurement with an ancilla through a judicious choice of the coefficients left free in the state proposed by Acin et al. (2001). The optimal constant (π^2/N^2 for the fidelity) was also computed. It was almost simultaneously noticed (Bagan et al., 2004b; Chiribella et al., 2004) that asymptotically the ancilla is unnecessary. Indeed what we need is entangling different copies of the same irreducible representation. Now each irreducible representation appears with multiplicity in $(\mathbb{C}^d)^{\otimes N}$, most of them with higher multiplicity than dimension, which is the condition we need. This method was dubbed "self-entanglement". The advantage is that we need to prepare half the number of particles, as we do not need an ancilla. In all these articles, the Bayesian paradigm with uniform prior was used. The same $1/N^2$ rate was shown to hold true in a minimax sense, in pointwise estimation (Hayashi, 2004). We stress the importance of this $1/N^2$ rate, proving how useful entanglement can be. Indeed, in classical data analysis, we cannot expect a better rate than $1/N$. Similarly the $1/N$ bound holds for any strategy where the N particles we send through the device are not entangled "among themselves" (that is, even if there is an ancilla for each of these N particles).

Another popular theme has been the determination of the phase ϕ for unitaries of the form $U_\phi = e^{i\phi H}$. This very special case already has many applications, especially in interferometry or measurement of small forces, as featured in the review article

by Giovannetti et al. (2004) and references therein. A common feature of the most efficient techniques is the need for entangled states of many particles, and much experimental work has aimed at generating such states. These methods essentially involve either manipulation of photons obtained through parametric down-conversion (for example (Eisenberg et al., 2005)), ions in ion traps (for example (Dalvit et al., 2006)) or atoms in cavity QED (for example (Vitali et al., 2006)).

In recent years, there has been renewed interest in the $SU(d)$ case. Notably, Chiribella et al. (2005) takes off from (Acin et al., 2001), allowing for more general symmetries and making explicit for natural cost functions both the free coefficients – as the coordinates of the eigenvector of a matrix – and the POVM (see Theorem 4.2.1 below). With a completely different strategy, aiming rather at pointwise estimation (and therefore minimax theorems), an input state for $U^{\otimes n}$ was found (Ballester, 2005b,a) such that the Quantum Fisher Information matrix is scaling like $1/N^2$, yielding hopes of getting as fast an estimator for $SU(d)$. No associated measurement was found in that paper.

Given the state of the art, a natural question is whether we can obtain, as for $SU(2)$, this dramatic increase in performance when using entanglement for general $SU(d)$. That is, do we have an estimation procedure whose rate is $1/N^2$, instead of $1/N$? Neither Chiribella et al. (2005), who do not study the asymptotics for $SU(d)$, nor Ballester (2005b), who does not give any measurement, answer this question.

In this chapter, we first prove that we cannot expect a better rate than $1/N^2$. This kind of bound based on the laws of quantum physics, without any *a priori* on the experimental device, is traditionally called the *Heisenberg limit* of the problem. Then we choose a completely explicit input state of the form (4.2) (as in (Acin et al., 2001)), by specifying the coefficients. By using the associated POVM, the estimator of a unitary quantum operation $U \in SU(d)$ converges at rate $1/N^2$. The constant is not optimal, but is briefly studied at the end of the chapter. We obtain these results with fidelity as a cost function, both in a Bayesian setting, with a uniform prior, and in a minimax setting. Notice that we shall not need an ancilla.

The next section consists in formulating the problem and restating Theorem 2 of (Chiribella et al., 2005) within our framework. Section 4.3 then shows that it is impossible to converge at rate faster than $O(N^{-2})$. In section 4.4, we write a general formula for the risk of a strategy as described in Theorem 4.2.1, and in section 4.5 we specify our estimators by choosing our coefficients in (4.2). We then prove that the risk of this estimator is $O(N^{-2})$. The last section (4.6) consists in finding the precise asymptotic speed of our procedure, that is the constant C in CN^{-2}. We finish by stating in Theorem 4.6.1 the results of the chapter.

4.2 Description of the problem

We are given an unknown unitary operation $U \in SU(d)$ and must estimate it "as precisely as possible". We are allowed to let it act on N particles, so that we are discriminating between the possible $U^{\otimes N}$. We shall work both with pointwise estimation (as preferred by mathematicians) and with a Bayes uniform prior (a favorite of physicists).

Any estimation procedure can be described as follows (see Figure 4.1): the unitary channel $U^{\otimes N}$ acts as

$$U^{\otimes N} \otimes \mathbf{1} : (\mathbb{C}^d)^{\otimes N} \otimes \mathcal{K} \to (\mathbb{C}^d)^{\otimes N} \otimes \mathcal{K},$$

on the space of the N systems together with a possible ancilla. The input state $\rho_n \in M((\mathbb{C}^d)^{\otimes n} \otimes \mathcal{K}_n)$ is mapped into an output state on which we perform a measurement M whose result is the estimator $\hat{U} \in SU(d)$.

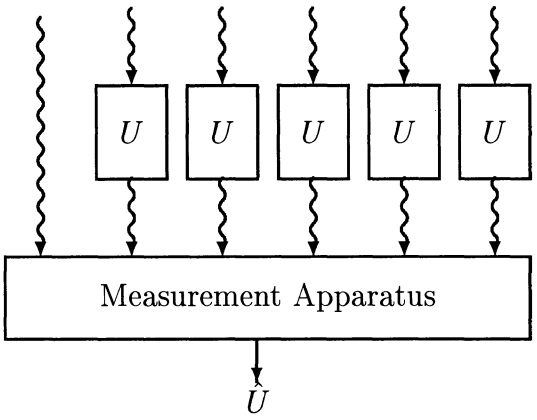

Figure 4.1: Most general estimation scheme of U when n copies are available at the same time, and using entanglement.

In order to evaluate the quality of an estimator $\hat{U}$, we fix a cost function $\Delta(U, V)$. The global pointwise risk of the estimator is

$$R_P(\hat{U}) = \sup_{U \in SU(d)} \mathbb{E}_U[\Delta(U, \hat{U})].$$

The probability distribution of $\hat{U}$ depends on U, and we take expectation with respect to this probability distribution.

On the other hand, the Bayes risk with uniform prior is:

$$R_B(\hat{U}) = \int_{SU(d)} \mathbb{E}_U[\Delta(U,\hat{U})]\mathrm{d}\mu(U).$$

where μ is the Haar measure on $SU(d)$.

As cost function, we choose the fidelity F (or rather $1-F$), which for an element of $SU(d)$ is defined as:

$$\begin{aligned}\Delta(U,\hat{U}) &= 1 - \frac{|\operatorname{Tr}(U^{-1}\hat{U})|^2}{d^2} \\ &= 1 - \frac{|\chi_{\square}(U^{-1}\hat{U})|^2}{d^2}\end{aligned}$$

where $\chi_{\square}$ is the character of the defining representation of $SU(d)$, whose Young tableau consists in only one box. In other words, $\chi_{\square}(U) = \operatorname{Tr}(U)$.

Before really addressing the problem, we make a few remarks on why this choice of distance is suitable for mathematical analysis.

Firstly, this cost function is covariant, i.e. $\Delta(U,\hat{U}) = \Delta(\mathbf{1}_{\mathbb{C}^d}, U^{-1}\hat{U})$.

Secondly, a useful feature within the Bayesian framework is that Δ is of the form (4.1), as required in Theorem 4.2.1. Indeed we can rewrite:

$$\Delta(U,\hat{U}) = 1 - \chi_{\square}(U^{-1}\hat{U})\chi^*_{\square}(U^{-1}\hat{U})/d^2.$$

Now the conjugate of a character is the character of the adjoint representation, the product of two characters is again the character of a possibly reducible representation π. This character is equal to the sum of the characters of the irreducible representations appearing in the Clebsch-Gordan development of π, in which all coefficients are non-negative. Therefore $\Delta = 1 - (\sum_{\vec{\lambda}} a_{\vec{\lambda}}\chi^*_{\vec{\lambda}})$ where $a_{\vec{\lambda}} \geq 0$ and $\vec{\lambda}$ runs over all irreducible representations of $SU(d)$. That is the condition (4.1) that we shall need for applying Theorem 4.2.1, given at the end of the section.

On the other hand, the theory of pointwise estimation deals usually with the variance of the estimated parameters when we use a smooth parameterization of $SU(d)$. As we want to use the Quantum Cramér-Rao Bound (4.9), we need Δ to be quadratic in the parameters to the first order, and positive lower bounded for $\hat{U}$ outside a neighborhood of U. As Δ is covariant, it is sufficient to check this with $U = \mathbf{1}_{\mathbb{C}^d}$. Now an example of a smooth parameterization in a neighborhood of the identity is $U(\theta) = \exp(\sum_\alpha \theta_\alpha T_\alpha)$ where $\theta \in \mathbb{R}^{d^2-1}$ and the T_α are generators of the Lie algebra, so that $\operatorname{Tr}(T_\alpha) = 0$. Now $\operatorname{Tr}[\exp(\sum_\alpha \theta_\alpha T_\alpha)] = d + \sum_\alpha \theta_\alpha \operatorname{Tr}(T_\alpha) + O(\|\theta\|^2)$, so that the trace minus d, and consequently Δ, is quadratic in θ to the first order.

As stated at the beginning of this section, we are working with $U^{\otimes N}$. The Clebsch-Gordan decomposition of the n-th tensor product representation is

$$U^{\otimes N} = \bigoplus_{\vec{\lambda}:|\vec{\lambda}|=N} U^{\vec{\lambda}} \otimes \mathbf{1}_{\mathbb{C}^{\mathcal{M}(\vec{\lambda})}}$$

acting on $\bigoplus_{\vec{\lambda}:|\vec{\lambda}|=N} \mathcal{H}^{\vec{\lambda}} \otimes \mathbb{C}^{\mathcal{M}(\vec{\lambda})}$, where $\mathcal{H}^{\vec{\lambda}} = \mathbb{C}^{\mathcal{D}(\vec{\lambda})}$ is the representation space of $\vec{\lambda}$, $\mathcal{M}(\vec{\lambda})$ is the multiplicity of $\vec{\lambda}$ in the n-th tensor product representation, and $\mathcal{D}(\vec{\lambda})$ the dimension of $\vec{\lambda}$. We refer to $\mathbb{C}^{\mathcal{M}(\vec{\lambda})}$ as the multiplicity space of $\vec{\lambda}$. We have indexed the irreducible representations of $SU(d)$ by $\vec{\lambda} = (\lambda_1, \ldots, \lambda_d)$, and written $|\vec{\lambda}| = \sum_{i=1}^{d} \lambda_i$. Notice that this labelling of irreducible representations is redundant, but that if $|\vec{\lambda}^1| = |\vec{\lambda}^2|$, then $\vec{\lambda}^1$ and $\vec{\lambda}^2$ are equivalent (denoted $\vec{\lambda}^1 \equiv \vec{\lambda}^2$) if and only if $\vec{\lambda}^1 = \vec{\lambda}^2$.

The starting point of our argument will be the following reformulation of the results of (Chiribella et al., 2005), with less generality, and without the formula for the risk whose form is not adapted to our subsequent analysis:

Theorem 4.2.1. (Chiribella et al., 2005) *Let $U \in SU(d)$ be a unitary operation to be estimated, through its action on N particles. We may use entanglement and/or an ancilla.*

Then, for a uniform prior and any cost function of the form

$$c(U, \hat{U}) = a_0 - \sum_{\vec{\lambda}} a_{\vec{\lambda}} \chi^*_{\vec{\lambda}}(U^{-1}\hat{U}), \tag{4.1}$$

we can find as optimal input state a pure state of the form

$$|\Psi\rangle = \bigoplus_{\vec{\lambda}:|\vec{\lambda}|=N} \frac{c(\vec{\lambda})}{\sqrt{\mathcal{D}(\vec{\lambda})}} \sum_{i=1}^{\mathcal{D}(\vec{\lambda})} |\psi_i^{\vec{\lambda}}\rangle \otimes |\phi_i^{\vec{\lambda}}\rangle \tag{4.2}$$

with $c(\vec{\lambda}) \geq 0$, and the normalization condition,

$$\sum_{\vec{\lambda}} c(\vec{\lambda})^2 = 1. \tag{4.3}$$

Moreover $|\psi_i^{\vec{\lambda}}\rangle$ is an orthonormal basis of $\mathcal{H}^{\lambda}$ and $|\phi_i^{\vec{\lambda}}\rangle$ are orthonormal vectors of the multiplicity space, which may be augmented by an ancilla if necessary (see remark below on the dimensions).

The corresponding measurement is the covariant POVM with seed $\Xi = |\eta\rangle\langle\eta|$ given by:

$$|\eta\rangle = \bigoplus_{\vec{\lambda}|c(\vec{\lambda})\neq 0} \sqrt{\mathcal{D}(\vec{\lambda})} \sum_{i=1}^{\mathcal{D}(\vec{\lambda})} |\psi_i^{\vec{\lambda}}\rangle \otimes |\phi_i^{\vec{\lambda}}\rangle, \tag{4.4}$$

that is a POVM whose density with respect to the Haar measure is given by $m(U) = U|\eta\rangle\langle\eta|U^*$ *with*

$$U|\eta\rangle = \bigoplus_{\vec{\lambda}|c(\vec{\lambda})\neq 0} \sqrt{\mathcal{D}(\vec{\lambda})} \sum_{i=1}^{\mathcal{D}(\vec{\lambda})} U^{\vec{\lambda}} |\psi_i^{\vec{\lambda}}\rangle \otimes |\phi_i^{\vec{\lambda}}\rangle.$$

Remark: We use $\mathcal{D}(\vec{\lambda})$ orthonormal vectors in the multiplicity space of $\vec{\lambda}$. This requires $\mathcal{M}(\vec{\lambda}) \geq \mathcal{D}(\vec{\lambda})$. If this is not the case, we must increase the dimension of the multiplicity space by using an ancilla in $\mathbb{C}^\delta$. Then the action of U is $U^{\otimes N} \otimes \mathbf{1}_{\mathbb{C}^\delta}$ whose Clebsch-Gordan decomposition is $\bigoplus_{\vec{\lambda}||\vec{\lambda}|=N} U^{\vec{\lambda}} \otimes \mathbf{1}_{\mathbb{C}^{\delta\mathcal{M}(\vec{\lambda})}}$. With big enough δ, we have $\delta\mathcal{M}(\vec{\lambda}) \geq \mathcal{D}(\vec{\lambda})$. Notice that an ancilla is not necessary if $c(\vec{\lambda}) = 0$ for all $\vec{\lambda}$ such that $\mathcal{D}(\vec{\lambda}) > \mathcal{M}(\vec{\lambda})$.

Another remark is that, as defined, our POVM is not properly normalized: $M(SU(d)) \neq \mathbf{1}$, but is equal to the projection on the space spanned by the $U|\Psi\rangle$. As this is the only subspace of importance, we can complete the POVM (through the seed, for example) *ad libitum.*

Our estimator $\hat{U}$ is the result of the measurement with POVM defined by (4.4) and input state of the form (4.2), with specific $c(\vec{\lambda})$. Such an estimator is covariant, that is $p_U(\hat{U}) = p_{\mathbf{1}_{\mathbb{C}^d}}(U^{-1}\hat{U})$, where p_U is the probability distribution of $\hat{U}$ when we are estimating U. The cost function is also covariant, so that $\mathbb{E}_U[\Delta(U,\hat{U})]$ does not depend on U. This implies that the Bayesian risk and the pointwise risk coincide. With the second equality true for all $U \in SU(d)$, we have:

$$R_B(\hat{U}) = R_P(\hat{U}) = \mathbb{E}_U[\Delta(U,\hat{U})]. \tag{4.5}$$

Theorem 4.2.1 states that there exists an optimal (Bayes uniform) estimator $\hat{U}_o$ of this form (corresponding to the optimal choice of $c(\vec{\lambda})$), so that it obeys (4.5). From this we first prove that no estimator whatsoever can have a better rate than $1/N^2$.

4.3 Why we cannot expect better rate than $1/N^2$

For proving this result, we need the Bayesian risk for priors π other than the uniform prior:

$$R_\pi(\hat{U}) = \mathbb{E}_\pi[\mathbb{E}_U[\Delta(U,\hat{U})]].$$

As $\hat{U}_o$ is Bayesian optimal for the uniform prior, we only have to prove that $R_B(\hat{U}_o) = O(N^{-2})$. This is also sufficient for pointwise risk as, for any estimator $\hat{U}$, we have

$R_B(\hat{U}) \leq R_P(\hat{U})$. Moreover, as $\mathbb{E}_U[\Delta(U, \hat{U}_o)]$ does not depend on U, $R_\pi(\hat{U}_o) = R_B(\hat{U}_o)$. It is then sufficient to prove, for a π of our choice, that:

$$R_\pi(\hat{U}_o) = O(N^{-2}). \tag{4.6}$$

The idea is to find a Cramér-Rao bound that we can apply to some π. We shall combine the Braunstein and Caves information inequality (4.8) and the Van Trees inequality (4.7) to obtain the desired Quantum Cramér-Rao Bound, much in the spirit of Gill (2005b). This bound will yield an explicit rate through a result of Ballester (2005b).

Van Trees' inequality states that given a classical statistical model smoothly parameterized by $\theta \in \Theta \subset \mathbb{R}^p$, and a smooth prior with compact support $\Theta_0 \subset \Theta$, then for any estimator $\hat{\theta}$, we have:

$$\mathbb{E}_\pi[\mathrm{Tr}(V_\theta(\hat{\theta}))] \geq \frac{p^2}{\mathbb{E}_\pi[\mathrm{Tr}(I(\theta))] - \mathcal{I}_\pi}, \tag{4.7}$$

where $I(\theta)$ is the Fisher information matrix of the model at point θ, $\mathcal{I}_\pi$ is a finite (for reasonable π) constant depending on π (quantifying in some way the prior information), and $V_\theta(\hat{\theta}) \in M_p(\mathbb{R})$ is the mean square error (MSE) of the estimator $\hat{\theta}$ at point θ given by:

$$V_\theta(\hat{\theta})_{\alpha,\beta} = \mathbb{E}[(\theta_\alpha - \hat{\theta}_\alpha)(\theta_\beta - \hat{\theta}_\beta)].$$

This form of Van Trees inequality is obtained by setting $N = 1$, $G = C = Id$ and $\psi = \theta$ in (12) of (Gill, 2005b).

Now the Braunstein et Caves C. M. (1994) information inequality yields an upper bound on the information matrix $I_M(\theta)$ of any classical statistical model obtained by applying the measurement M to a quantum statistical model. For any family of quantum states parameterized by a p-dimensional parameter $\theta \in \Theta \in \mathbb{R}^p$, for any measurement M on these states, the following holds:

$$I_M(\theta) \leq H(\theta), \tag{4.8}$$

where $H(\theta)$ is the quantum Fisher information information matrix at point θ.

Now it was proved by Ballester (2005b) that for a smooth parameterization of an open set of $SU(d)$, and for any input state, the quantum Fisher information of the output states fulfils:

$$H(\theta) = O(N^2).$$

Inserting in (4.7) together with (4.8) we get as quantum Cramér-Rao bound

$$\mathbb{E}_\pi[\mathrm{Tr}(V_\theta(\hat{\theta}))] = O\left(\frac{1}{N^2}\right). \tag{4.9}$$

We now want to apply this bound to obtain (4.6). There are a few small technical difficulties. First of all, we cannot use the uniform prior for π as $SU(d)$ is not homeomorphic to an open set of $\mathbb{R}^p$. We then have to define two neighborhoods of the identity $\Theta_0 \subset \Theta$, allowing to use the Van Trees inequality. Now our estimator $\hat{U}_o$ need not be in Θ, so that we shall in fact apply Van Trees inequality to a modified estimator $\tilde{U}$. Finally, this bound is on the variance, and we must relate it to Δ.

Our first task consists in restricting our attention to a neighborhood Θ of $\mathbf{1}_{\mathbb{C}^d}$. It corresponds to a neighborhood Θ (we use the same notation) of $0 \in \mathbb{R}^p$ through $U = \exp(\sum_\alpha \theta_\alpha T_\alpha)$. This holds if the neighborhood is small enough, so we define it by $U \in \Theta$ if and only if $\Delta(\mathbf{1}_{\mathbb{C}^d}, U) < \epsilon$ for a fixed small enough ϵ. We define Θ_0 through $U \in \Theta_0$ for $\Delta(\mathbf{1}_{\mathbb{C}^d}, U) \leq \epsilon/3$, and take a smooth fixed prior π with support in Θ_0, such that $\mathcal{I}_\pi < \infty$.

Now we modify our estimator $\hat{U}_o$ into an estimator $\tilde{U}$ given by $\tilde{U} = \hat{U}_o$ for $\hat{U}_o \in \Theta$ and $\tilde{U} = \mathbf{1}_{\mathbb{C}^d}$ for $\hat{U}_o \notin \Theta$. Then, by the triangle inequality, for any $U \in \Theta_0$, we have $\Delta(U, \hat{U}_o) \geq \Delta(U, \tilde{U})$.

The fundamental point of the reasoning (used at (4.10)) is that, as Δ is quadratic at the first-order, there is a positive constant c such that, for any $U^1, U^2 \in \Theta$, corresponding to θ^1, θ^2, we have $\Delta(U_1, U_2) \geq c \sum_\alpha (\theta^1_\alpha - \theta^2_\alpha)^2$.

Finally we get

$$\begin{aligned} R_\pi(\hat{U}_o) &= \mathbb{E}_\pi[\mathbb{E}_U[\Delta(U, \hat{U}_o)]] \\ &\geq \mathbb{E}_\pi[\mathbb{E}_U[\Delta(U, \tilde{U})]] \\ &\geq c\mathbb{E}_\pi[V_{\tilde{\theta}}] \\ &= O(N^{-2}). \end{aligned} \tag{4.10}$$

We have thus proved (4.6), and hence our bound on the efficiency of any estimator.

We now write formulas for the risk of any estimator of the form given in Theorem 4.2.1.

4.4 Formulas for the risk

By (4.5), our risk $R_P(\hat{U})$ is equal to the pointwise risk at $\mathbf{1}_{\mathbb{C}^d}$, with which we shall work:

$$\int_{SU(d)} p_{\mathbf{1}_{\mathbb{C}^d}}(\hat{U}) \left\{ 1 - \frac{|\chi_\square(\hat{U})|^2}{d^2} \right\} d\mu(\hat{U}). \tag{4.11}$$

Now we compute the probability distribution of $\hat{U}$ for a given $|\Psi\rangle$ of the form (4.2), that is

$$\begin{aligned} p_{\mathbf{1}_{\mathbb{C}^d}}(\hat{U}) &= \langle\Psi|\hat{U}\Xi\hat{U}^*|\Psi\rangle \\ &= \left|\sum_{\vec{\lambda}:|\vec{\lambda}|=N} \frac{c(\vec{\lambda})}{\mathcal{D}(\vec{\lambda})}\mathcal{D}(\vec{\lambda}) \sum_{i=1}^{\mathcal{D}(\vec{\lambda})} \langle\psi_i^{\vec{\lambda}}|U|\psi_i^{\vec{\lambda}}\rangle\right|^2 \\ &= \left|\sum_{\vec{\lambda}:|\vec{\lambda}|=N} c(\vec{\lambda})\chi_{\vec{\lambda}}(\hat{U})\right|^2, \end{aligned}$$

where we have used that the character $\chi_{\vec{\lambda}}$ of $\vec{\lambda}$ is the trace of U in the representation.

Then, using (4.11), recalling that $p_{\mathbf{1}_{\mathbb{C}^d}}$ is a probability density for Haar measure μ on $SU(d)$, and that $\chi_{\vec{\lambda}^1}\chi_{\vec{\lambda}^2} = \chi_{\vec{\lambda}^1\otimes\vec{\lambda}^2}$ (for the second term), we get:

$$R_P(\hat{U}) = 1 - \frac{1}{d^2}\int_{SU(d)} \left|\sum_{\vec{\lambda}:|\vec{\lambda}|=N} c(\vec{\lambda})\chi_{\vec{\lambda}\otimes\square}(\hat{U})\right|^2 d\mu(\hat{U}). \tag{4.12}$$

In order to evaluate the second term, we use the following orthogonality relations for characters:

$$\int_{SU(d)} d\mu(U)\chi_{\vec{\lambda}_1}(U)\chi_{\vec{\lambda}_2}(U)^* = \delta_{\vec{\lambda}_1\equiv\vec{\lambda}_2}. \tag{4.13}$$

To do so we need the Clebsch-Gordan series of $\vec{\lambda}\otimes\square$:

$$\vec{\lambda}\otimes\square = \oplus_{\{1\le i\le d|\lambda_i>\lambda_{i+1}\}}\vec{\lambda} + e_i, \tag{4.14}$$

where conventionally $\lambda_{d+1} = 0$. Here we see $\vec{\lambda}$ as a d-dimensional vector and e_i as the i-th basis vector.

We then reorganize the sum of characters as:

$$\sum_{\vec{\lambda}:|\vec{\lambda}|=N} c(\vec{\lambda})\chi_{\vec{\lambda}\otimes\square}(\hat{U}) = \sum_{\vec{\lambda}':|\vec{\lambda}'|=N+1}\sum_{i\in\mathcal{S}(\vec{\lambda}')} c(\vec{\lambda}'-e_i)\chi_{\vec{\lambda}'}(\hat{U}),$$

where $\mathcal{S}(\vec{\lambda}')$ is the set of i between 1 and d such that $\vec{\lambda}'-e_i$ is still a representation, that is $\lambda'_i > \lambda'_{i+1}$. We shall write $\#\mathcal{S}(\vec{\lambda}')$ for its cardinality.

Inserting in (4.12) and remembering (4.13), we are left with

$$R_P(\hat{U}) = 1 - \frac{\sum_{\vec{\lambda}':|\vec{\lambda}'|=N+1}|\sum_{i\in\mathcal{S}(\vec{\lambda}')} c(\vec{\lambda}'-e_i)|^2}{d^2}. \tag{4.15}$$

To go any further, we must work with specific $c(\vec{\lambda})$.

4.5 Choice of the coefficients $c(\vec{\lambda})$ and proof of their efficiency

We now have to choose the coefficients $c(\vec{\lambda})$ so that the right-hand side of (4.15) is small.

It appears useful to introduce subsets of the set of all irreducible representations. Let $\mathcal{P}_N = \{\vec{\lambda}|\ |\vec{\lambda}| = N; \lambda_1 > \cdots > \lambda_d > 0\}$. Obviously, if $\vec{\lambda}' \in \mathcal{P}_{N+1}$, then $\#\mathcal{S}(\vec{\lambda}') = d$, and the converse is true. We can see them intuitively as points on a $(d-1)$-dimensional surface, and with this picture in mind, we shall speak of the border of $\mathcal{P}_N$ (when $\lambda_i = \lambda_{i+1} + 1$ for some i), or of being far from the border (without precise mathematical meaning).

We are ready to give heuristic arguments on how good coefficients should behave.

We must try to get the fraction in (4.15) close to one. Now

$$\begin{aligned}
&\frac{\sum_{\vec{\lambda}':|\vec{\lambda}'|=N+1} |\sum_{i\in\mathcal{S}(\vec{\lambda}')} c(\vec{\lambda}' - e_i)|^2}{d^2} \\
&\leq \sum_{\vec{\lambda}':|\vec{\lambda}'|=N+1} \frac{\#\mathcal{S}(\vec{\lambda}')}{d} \frac{\sum_{i\in\mathcal{S}(\vec{\lambda}')} |c(\vec{\lambda}' - e_i)|^2}{d} \\
&\leq \sum_{\vec{\lambda}':|\vec{\lambda}'|=N+1} \frac{\sum_{i\in\mathcal{S}(\vec{\lambda}')} |c(\vec{\lambda}' - e_i)|^2}{d} \\
&\leq \sum_{\vec{\lambda}:|\vec{\lambda}|=N} |c(\vec{\lambda})|^2 = 1.
\end{aligned}$$

The first inequality was obtained using Cauchy-Schwarz inequality for each inner sum. There is equality if $c(\vec{\lambda}' - e_i)$ does not depend on i. From this, we deduce that for most $\vec{\lambda}'$, the $c(\vec{\lambda}' - e_i)$ must be approximately equal, especially if they are large. The second inequality follows from $\#\mathcal{S}(\vec{\lambda}') \leq d$. From this we deduce that for $\vec{\lambda} \notin \mathcal{P}_{N+1}$, the coefficients $c(\vec{\lambda} - e_i)$ must be small. Remark that about $1/N$ of the $\vec{\lambda}'$ such that $|\vec{\lambda}'| = N+1$ are not in $\mathcal{P}_{N+1}$, so that if all $c(\vec{\lambda})$ were equal, these border terms would cause our rate to be $1/N$. The key of the third inequality is to notice that each $c(\vec{\lambda})$ is appearing in the sum once for each term in its Clebsch-Gordan series (4.14), and that there are at most d terms. Please note that there are d terms if $\vec{\lambda} \in \mathcal{P}_N$, and if $\vec{\lambda}'$ is in $\mathcal{P}_{N+1}$, far from the border, then $\vec{\lambda}' - e_i$ is in $\mathcal{P}_N$, far from the border.

The conclusion of these heuristics is that we must choose coefficients "locally" approximately equal (at most $1/N$ variation in ratio), and that the coefficients must go to 0 when we are approaching the border of $\mathcal{P}_N$.

One weight satisfying these heuristics is the following.

$$c(\vec{\lambda}) = \mathcal{N} \prod_{i=1}^{d} p_i, \tag{4.16}$$

where $\mathcal{N}$ is a normalization constant to ensure that (4.3) is satisfied and $p_i = \lambda_i - \lambda_{i+1}$. We shall use it below, and prove that it delivers the $1/N^2$ rate.

A first remark about these weights is that $c(\vec{\lambda}) = 0$ if $\vec{\lambda} \notin \mathcal{P}_N$. Now, for any $\vec{\lambda} \in \mathcal{P}_N$, we have $\mathcal{D}(\vec{\lambda}) \geq \mathcal{M}(\vec{\lambda})$, so that we do not need an ancilla.

Indeed, using hook formulas (see (Schensted, 1976)), we get

$$\mathcal{M}(\vec{\lambda})/\mathcal{D}(\vec{\lambda}) = N! \prod_{i=1}^{d} \frac{(\lambda_i + d - i)!}{(d-i)!}.$$

Now for $\vec{\lambda} \in \mathcal{P}_N$, we know that $\lambda_i \neq 0$. Under this constraint and $\sum \lambda_i = N$, the maximum is attained by $\lambda_1 = N - d + 1$ and $\lambda_i = 1$ for $i \neq 1$. We end up with exactly 1.

We shall now use (4.16) and express the numerator of (4.15) with our choice of p_i. Notice first that if p_j characterize $\vec{\lambda}'$ then those which characterize $\vec{\lambda}' - e_i$ are given by $p_j^{(i)} = p_j + \delta_{j,i-1} - \delta_{j,i}$. So

$$\mathcal{N}^{-1} c(\vec{\lambda}' - e_i) = \prod_{j=1}^{d} p_j + r_{\vec{\lambda}'}(i),$$

with

$$r_{\vec{\lambda}'}(i) = -\prod_{j \neq i} p_j + \delta_{j>1} \left(\prod_{j \neq i-1} p_j - \prod_{j \neq i,i-1} p_j \right).$$

Introducing another notation will make this slightly more compact. For a vector $\vec{x}$ with d components and $\mathcal{E}$ a subset of $\{1, \ldots, d\}$, define:

$$x_{\mathcal{E}} = \prod_{j \neq \mathcal{E}} x_j. \tag{4.17}$$

Then

$$r_{\vec{\lambda}'}(i) = -p_{\{i\}} + \delta_{j>1} \left(p_{\{i-1\}} - p_{\{i,i-1\}} \right).$$

Notice now that for $\vec{\lambda} \in \mathcal{P}_N$, there are exactly d irreducible representations appearing in the Clebsch-Gordan decomposition of $\vec{\lambda} \otimes \square$ (4.14). So that $c(\vec{\lambda})^2$ appears exactly d times in $\sum_{\vec{\lambda}':|\vec{\lambda}'|=N+1} \sum_{i \in \mathcal{S}(\vec{\lambda}')} c(\vec{\lambda}' - e_i)^2$. We may then rewrite the renormalization constant $\mathcal{N}$ as

$$d^{-1} \sum_{\vec{\lambda}':|\vec{\lambda}'|=N+1} \sum_{i \in \mathcal{S}(\vec{\lambda}')} \prod_{j=1}^{d} p_j^{(i)2}.$$

Therefore, rewriting the second term in (4.15) with our values of $c(\vec{\lambda})$, we aim at proving:

$$\frac{\sum_{\vec{\lambda}':|\vec{\lambda}'|=N+1} \left(\sum_{i \in \mathcal{S}(\vec{\lambda}')} \prod_{j=1}^{d} p_j + r_{\vec{\lambda}'}(i) \right)^2}{d \sum_{\vec{\lambda}':|\vec{\lambda}'|=N+1} \sum_{i \in \mathcal{S}(\vec{\lambda}')} \left(\prod_{j=1}^{d} p_j + r_{\vec{\lambda}'}(i) \right)^2} = 1 + O(N^{-2}). \tag{4.18}$$

Let us expand the numerator:

$$\sum_{\vec{\lambda}':|\vec{\lambda}'|=N+1} \left(\sum_{i \in \mathcal{S}(\vec{\lambda}')} \prod_{j=1}^{d} p_j + r_{\vec{\lambda}'}(i) \right)^2 = C_t \left(1 + t_1 + t_2\right),$$

with

$$\begin{aligned} C_t &= \sum_{\vec{\lambda}'} (\#\mathcal{S}(\vec{\lambda}'))^2 \prod_{j=1}^{d} p_j^2, \\ t_1 &= \frac{2 \sum_{\vec{\lambda}'} \sum_{i \in \mathcal{S}(\vec{\lambda}')} \#\mathcal{S}(\vec{\lambda}') r_{\vec{\lambda}'}(i) \prod_{j=1}^{d} p_j}{C_t}, \\ t_2 &= \frac{\sum_{\vec{\lambda}'} \left(\sum_{i \in \mathcal{S}(\vec{\lambda}')} r_{\vec{\lambda}'}(i) \right)^2}{C_t}. \end{aligned}$$

Similarly the denominator can be read as:

$$d \sum_{\vec{\lambda}':|\vec{\lambda}'|=N+1} \sum_{i \in \mathcal{S}(\vec{\lambda}')} \left(\prod_{j=1}^{d} p_j + r_{\vec{\lambda}'}(i) \right)^2 = C_u \left(1 + u_1 + u_2\right),$$

with

$$C_u = \sum_{\vec{\lambda}'} d\#\mathcal{S}(\vec{\lambda}') \prod_{j=1}^{d} p_j^2,$$
$$u_1 = \frac{2d \sum_{\vec{\lambda}'} \sum_{i\in\mathcal{S}(\vec{\lambda}')} r_{\vec{\lambda}'}(i) \prod_{j=1}^{d} p_j}{C_u},$$
$$u_2 = \frac{\sum_{\vec{\lambda}'} d \sum_{i\in\mathcal{S}(\vec{\lambda}')} r_{\vec{\lambda}'}(i)^2}{C_u}.$$

With these notations, we aim at proving the set of estimates given in Lemma 4.5.1. Indeed they imply:

$$\frac{\sum_{\vec{\lambda}':|\vec{\lambda}'|=N+1} \left(\sum_{i\in\mathcal{S}(\vec{\lambda}')} \prod_{j=1}^{d} p_j + r_{\vec{\lambda}'}(i)\right)^2}{d \sum_{\vec{\lambda}':|\vec{\lambda}'|=N+1} \sum_{i\in\mathcal{S}(\vec{\lambda}')} \left(\prod_{j=1}^{d} p_j + r_{\vec{\lambda}'}(i)\right)^2} = 1 + t_2 - u_2 + O(N^{-3}) \tag{4.19}$$

with $(t_2 - u_2)$ of order N^{-2}. By (4.18), the risk of the estimator is then $u_2 - t_2 + O(N^{-3})$. Thus proving Lemma 4.5.1 amounts at proving $1/N^2$ rate.

We shall make use of the notation $\Theta(f)$, meaning that there are universal positive constants m and M such that:

$$mf \leq \Theta(f) \leq Mf.$$

Lemma 4.5.1. *With the above notations,*

$$C_u = C_t = d^2 \sum_{\vec{\lambda}':|\vec{\lambda}'|=N+1} \left(\prod_{j=1}^{d} p_j\right)^2$$
$$= \Theta(N^{3d-1})$$
$$t_1 = u_1 = O(N^{-1})$$
$$t_2 = O(N^{-2})$$
$$u_2 = O(N^{-2}).$$

Proof. We first prove the first line.

Indeed for $\vec{\lambda}' \in \mathcal{P}_{N+1}$, all i are in $\mathcal{S}(\vec{\lambda}')$, and

$$\left(\sum_{i\in\mathcal{S}(\vec{\lambda}')} \prod_{j=1}^{d} p_j\right)^2 = d \sum_{i\in\mathcal{S}(\vec{\lambda}')} \prod_{j=1}^{d} p_j^2 = d^2 \prod_{j=1}^{d} p_j^2.$$

But if $\vec{\lambda}' \notin \mathcal{P}_{N+1}$, there is at least one p_j equal to zero, so they do not contribute to the sum. So that $C_u = C_t = d^2 \sum_{\vec{\lambda}':|\vec{\lambda}'|=N+1} \left(\prod_{j=1}^{d} p_j\right)^2$.

We have then equality of the denominators of t_1 and u_1. The same argument gives equality of the numerators. On $\mathcal{P}_{N+1}$, $\#\mathcal{S}(\vec{\lambda}') = d$ so that

$$\sum_{i\in\mathcal{S}(\vec{\lambda}')} \#\mathcal{S}(\vec{\lambda}') r_{\vec{\lambda}'}(i) \prod_{j=1}^{d} p_j = d \sum_{i\in\mathcal{S}(\vec{\lambda}')} r_{\vec{\lambda}'}(i) \prod_{j=1}^{d} p_j,$$

and outside $\mathcal{P}_{N+1}$, $\prod_{j=1}^{d} p_j = 0$ so that the equality still holds. Therefore $t_1 = u_1$.

Now $p_j \leq N+1$ so that $\prod_{j=1}^{d} p_j \leq (N+1)^d$ and $|r_{\vec{\lambda}'}(i)| \leq 2(N+1)^{d-1}$. Moreover, as $1 \leq \lambda_i \leq N+1$ and λ_d is known if the other λ_i are known, the number of elements $\vec{\lambda}'$ in $\mathcal{P}_{N+1}$ satisfies $\#\mathcal{P}_{N+1} \leq (N+1)^{d-1}$. Thus the numerator of t_1 and u_1 is $O(N^{3d-2})$ and that of t_2 and u_2 is $O(N^{3d-3})$. To end the proof of the lemma, it is then sufficient to show that $C_u = \Theta(N^{3d-1})$.

Let us write $N+1 = a(1+d(d+1))/2 + b$ with a and b natural integers and $b < (1+d(d+1))$. We then select h_i for $i = 1$ to d such that $\sum h_i = a/2$. The number of ways of partitioning $a/2$ in d parts is $\binom{a/2+d-1}{d-1}$, and this is $\Theta(a^{d-1}) = \Theta(N^{d-1})$. To each of these partitions, we associate a different $\vec{\lambda}'$ in $\mathcal{P}_{N+1}$ through $\lambda_i = (d-i+1)a + \delta_{i=1}b + h_i$. For each of these $\vec{\lambda}'$, we have $p_j = \lambda_j - \lambda_{j+1} \geq a/2$, so that $\prod_{j=1}^{d} p_j^2 = \Theta(N^{2d})$. We may lower bound C_u by the sum over these $\vec{\lambda}'$ of $\prod_{j=1}^{d} p_j^2$, so that we have proved $C_u = \Theta(N^{3d-1})$. □

4.6 Evaluation of the constant in the speed of convergence and final result

The strategy we study is asymptotically optimal up to a constant, but a better constant can probably be obtained. Anything like $c(\vec{\lambda}) = (\prod p_j)^\alpha$ with $\alpha \geq 1/2$ should yield the same rate, though it would be more cumbersome to prove. Polynomials in the p_j could also bring some improvement. All the same we give in this section a quick evaluation of the constant, that may serve as a benchmark for more precise strategies.

Write $p_j = (N+1)x_j$. Then, recalling our notation 4.17,

$$\prod_{j=1}^{d} p_j^2 = (N+1)^{2d} \prod_{j=1}^{d} x_j^2$$
$$r_{\vec{\lambda}'}(i) = (N+1)^{d-1} \left(-x_{\{i\}} + \delta_{i>1} x_{\{i-1\}} + O(N^{-1})\right).$$

Similarly, the set of allowed $\vec{x} = (x_1, \ldots, x_n)$ may be described as

$$\mathcal{S}_{N+1} = \left\{ \vec{x} \,|\, x_j(N+1) \in \mathbb{N}; \sum_{j=1}^{d} (d-j+1)x_j = 1 \right\}.$$

We may then rewrite:

$$u_2 = \frac{\sum_{\vec{x}\in\mathcal{S}_{N+1}} d \sum_{i=1}^{d} \left(x_{\{i\}} - \delta_{i>1} x_{\{i-1\}}\right)^2}{d^2(N+1)^2 \sum_{\vec{x}\in\mathcal{S}_{N+1}} \prod_{j=1}^{d} x_j^2} + O(N^{-3})$$
$$t_2 = \frac{\sum_{\vec{x}\in\mathcal{S}_{N+1}} \left(x_{\{i\}} - \delta_{i>1} x_{\{i-1\}}\right)^2}{d^2(N+1)^2 \sum_{\vec{x}\in\mathcal{S}_{N+1}} \prod_{j=1}^{d} x_j^2} + O(N^{-3}).$$

Subtracting, we obtain (the first sums being on $\mathcal{S}_{N+1}$)

$$u_2 - t_2 + O(N^{-3}) = \tag{4.20}$$
$$\frac{\sum_{\vec{x}} 2d \left(\sum_{i=1}^{d} (x_{\{i\}})^2 - \sum_{i=2}^{d} x_{\{i\}} x_{\{i-1\}}\right) - (d+1)(x_{\{d\}})^2}{n^2\, d^2 \sum_{\vec{x}} \prod_{j=1}^{d} x_j^2}. \tag{4.21}$$

Now $\mathcal{S}_{N+1}$ is the intersection $\mathcal{S}$ of the lattice in $[0,1]^d$ with mesh size $1/(N+1)$ with the hyperplane given by the equation $\sum(d-j+1)x_j = 1$. Therefore the points of $\mathcal{S}_{N+1}$ are a regular paving of a flat $(d-1)$-dimensional volume, with more and more points (we know that $\#\mathcal{S}_{N+1} = O(N^{d-1})$). Therefore both denominator and numerator of (4.20) are Riemannian sums with respect to the Lebesgue measure, with a multiplicative constant that is the same for both. Therefore we have proved:

Theorem 4.6.1. *The estimator $\hat{U}$ corresponding to (4.16) has the following risk:*

$$R_B(\hat{U}) = R_P(\hat{U}) = \mathbb{E}_{\mathbf{1}_{\mathbb{C}^d}} \left[\Delta(\mathbf{1}_{\mathbb{C}^d}, \hat{U})\right] = CN^{-2} + O(N^{-3})$$

where C is the fraction

$$\frac{\int_{\mathcal{S}} 2d \left(\sum_{i=1}^{d} (x_{\{i\}})^2 - \sum_{i=2}^{d} x_{\{i\}} x_{\{i-1\}}\right) - (d+1)(x_{\{d\}})^2 \mathrm{d}\vec{x}}{d^2 \int_{\mathcal{S}} \prod_{j=1}^{d} x_j^2 \mathrm{d}\vec{x}}.$$

Up to a multiplicative constant, this risk is asymptotically optimal, both for a Bayes uniform prior and for global pointwise estimation.

Numerical estimation, up to two digits, for the low dimensions yields:

$$\begin{aligned} 10 \text{ for } d &= 2 \\ 75 \text{ for } d &= 3 \\ 2.7 \times 10^2 \text{ for } d &= 4. \end{aligned}$$

4.7 Conclusion

We have given a strategy for estimating an unknown unitary channel $U \in SU(d)$, and proved that the convergence rate of this strategy is $1/N^2$. We have further proved that this rate is optimal, even if the constant may be improved.

The interest of this result lies in that such rates are much faster than the $1/N$ achieved in classical estimation and, though they had already been obtained for $SU(2)$, they were never before shown to hold for general $SU(d)$.

Chapitre 5

Clean positive operator valued measures

This chapter is derived from the article (Kahn, 2007a).

Résumé : Dans un article récent Buscemi et al. (2005) ont défini une notion de propreté des mesures à valeurs dans les opérateurs positifs (POVMs). Nous caractérisons les POVMs propres dans une classe que nous appelons quasi-qubit, c'est-à-dire les POVMs dont les éléments sont tous de rang un ou de rang plein. Nous donnons un algorithme qui vérifie si une POVM quasi-qubit satisfait à la condition de propreté. Nous décrivons explicitement toutes les POVMs propres pour les qubits. Au passage, nous donnons une condition suffisante pour qu'une POVM générale soit propre.

5.1 Introduction

The laws of quantum mechanics impose restrictions on what measurements can be carried out on a quantum system. All the possible measurements can be described mathematically by "positive operator-valued measures", POVMs for short. Apart from measuring a state, we can also transform it via a quantum channel. Now suppose we have at our disposal a POVM $\mathbf{P}$ and a channel $\mathcal{E}$. We may first send our state through $\mathcal{E}$ and then feed the transformed state in our measurement apparatus $\mathbf{P}$. This procedure is a new measurement procedure, and can therefore be encoded by a POVM $\mathbf{Q}$. Now transforming the state with $\mathcal{E}$ can be seen as a kind of noise

on the POVM $\mathbf{P}$. We may then view $\mathbf{Q}$ as a disturbed version of $\mathbf{P}$, and we say that $\mathbf{P}$ is *cleaner* than $\mathbf{Q}$. Now, what are the maximal elements for this order relation?

The order relation "cleaner than" has been introduced in a recent article of Buscemi et al. (2005). Herein they look at which POVMs can be obtained from another, either by pre-processing (the situation we just described, where we first send our state through a channel) or by classical post-processing of the data. Especially, they try to find which POVMs are biggest for these order relations (in the former case, the POVM is said to be *clean*; there is no "extrinsic" noise). For pre-processing they get a number of partial answers. One of those is that a POVM on a d-dimensional space with n outcomes, with $n \leq d$, is clean if and only if it is an observable. They do not get a complete classification, though.

The object of the present chapter is to characterize which POVMs are clean in a special class of measurements. Namely, we are interested in POVMs such that all their elements (see definition below) are either full-rank or rank-one. We call these POVMs *quasi-qubit POVMs*. Notice that all the POVMs for qubits satisfy to this condition.

On the way we prove a sufficient condition for a POVM to be clean, that is usable also for POVMs that are not quasi-qubit.

It turns out that cleanness for quasi-qubit POVMs can be read on the span of the rank-one elements. Moreover,if a (non necessarily quasi-qubit) POVM is cleaner than a clean quasi-qubit POVM, the latter was in fact obtained by a channel that is a unitary transform. In other words, for quasi-qubit POVMs, cleanness-equivalence is unitary equivalence.

We give an algorithm to check whether a quasi-qubit POVM is clean or not. This algorithm may be the main contribution of the chapter, as almost all the following theorems can be summed up by saying the algorithm is valid.

In the end we apply these results to the qubit, for which all POVMs are quasi-qubit. We are then left with a very explicit characterization of clean POVMs for qubits.

Section 5.2 gives precise definitions of all the objects we cited in this introduction.

We define the algorithm, give heuristically the main ideas and define the important notion "totally determined" (Definition 5.3.2) in Section 5.3.

Section 5.4 gives a sufficient condition for a POVM to be clean, namely that the supports of the elements of the POVM "totally determine" the space (see Definition 5.3.2). We use this condition to show that when the algorithm exits with a positive result, the quasi-qubit POVM is really clean.

Section 5.5 proves that the above sufficient condition is in fact necessary for quasi-qubit POVMs. It checks that when the algorithm exits with a negative result, the POVM is truly not clean.

Section 5.6 gathers the results relative to quasi-qubit POVMs in Theorem 5.6.1 and deals with the qubit case in Corollary 5.6.2.

Ultimately section 5.7 gives a very rough idea for making explicit more explicit the sufficient condition for a POVM to be clean we have given in section 5.4.

If one wishes to look for the results of this chapter without bothering with the technical proofs, the best would be to read the algorithm of section 5.3 and then to read Theorem 5.6.1 and Corollary 5.6.2. You would also need Lemma 5.5.3 that you could use as a definition of "totally determined" if you are only interested in quasi-qubit POVMs.

If you also want the supplementary results that apply to other POVMs, further read Definitions 5.3.1 and 5.3.2, and Theorem 5.4.1.

5.2 Definitions and notations

We consider POVMs on a Hilbert space $\mathcal{H}$ of dimension $d \geq 2$. Dimension 2 is the qubit case. The set $\{|e_i\rangle\}_{1\leq i\leq d}$ will be an orthonormal basis of $\mathcal{H}$. If $\mathcal{V}$ is a subspace of $\mathcal{H}$ then $\mathcal{V}^\perp$ is the subspace orthogonal to $\mathcal{V}$ in $\mathcal{H}$. If we are given vectors $\{v_i\}_{i\in I}$, we denote by $\mathrm{Span}(v_i, i \in I)$ the space they generate. The set of operators on $\mathcal{H}$ is denoted by $\mathcal{B}(\mathcal{H})$.

A POVM $\mathbf{P}$ (with finite outcomes, case to which we restrict) is a set $\{P_i\}_{i\in I}$ of *non-negative* operators on $\mathcal{H}$, with I finite, such that $\sum_{i\in I} P_i = \mathbf{1}$. The P_i are called *POVM elements*. We write $\mathrm{Supp}(P_i)$ for the support of this element. This support is defined by its orthogonal. The set of $|\phi\rangle \in \mathrm{Supp}(P_i)^\perp$ is exactly the set of $|\phi\rangle$ such that $\langle\phi|P_i|\phi\rangle = 0$. The rank of a POVM element is its rank as an operator. In particular, rank-one elements are of the form $\lambda_i|\psi_i\rangle\langle\psi_i|$ and full-rank POVMs are invertible. Special cases of POVMs are *rank-one POVMs*, that is POVMs whose elements are all rank-one, and *full-rank POVMs*, that is POVMs whose elements are all full-rank. We are especially interested in a class of POVMs that includes both:

Definition 5.2.1. Quasi-qubits POVMs

A POVM $\mathbf{P}$ *is a* quasi-qubit *POVM if all its elements* P_i *are either full-rank or rank-one.*

Similarly, we shall speak of strict quasi-qubit *POVMs for quasi-qubit POVMs which are neither rank-one nor full-rank.*

A channel $\mathcal{E}$ is a completely positive identity-preserving map on $\mathcal{B}(\mathcal{H})$ the set of bounded operators on $\mathcal{H}$ (in this chapter, channels are always intended as going from $\mathcal{B}(\mathcal{H})$ to the *same* $\mathcal{B}(\mathcal{H})$). As a remark, this implies that the subspace of self-adjoint operators $\mathcal{B}_{sa}(\mathcal{H})$ is stable by $\mathcal{E}$. We know we can write it using Kraus (1983) decomposition, that is we can find a finite number of operators $R_\alpha \in \mathcal{B}(\mathcal{H})$ such that

$$\mathcal{E}(A) = \sum_\alpha R_\alpha^* A R_\alpha, \quad \text{with} \quad \sum_\alpha R_\alpha^* R_\alpha = \mathbf{1}. \tag{5.1}$$

Here the star is the adjoint.

We shall write $\mathcal{E} = \{R_\alpha\}_\alpha$. This decomposition is not unique.

Using the channel $\mathcal{E}$ before the measurement $\mathbf{P}$ is the same as using the POVM $\mathbf{Q} = \mathcal{E}(\mathbf{P})$ defined by its POVM elements $Q_i = \mathcal{E}(P_i)$.

Definition 5.2.2. *A POVM* $\mathbf{P}$ *is* cleaner *than a POVM* $\mathbf{Q}$ *if and only if there exists a channel* $\mathcal{E}$ *such that* $\mathcal{E}(\mathbf{P}) = \mathbf{Q}$. *We shall also write* $\mathbf{P} \succ \mathbf{Q}$.

Definition 5.2.3. Clean POVM

A POVM $\mathbf{P}$ *is clean if and only if, for any* $\mathbf{Q}$ *such that* $\mathbf{Q} \succ \mathbf{P}$, *then* $\mathbf{P} \succ \mathbf{Q}$ *also holds.*

We shall further say that two POVMs are cleanness-equivalent if both $\mathbf{Q} \succ \mathbf{P}$ and $\mathbf{P} \succ \mathbf{Q}$ hold. A special case of this (but not the general case, as proved in (Buscemi et al., 2005)) is *unitary equivalence*, when there is a unitary operator U such that for any $i \in I$, we have $UP_iU^* = Q_i$.

5.3 Algorithm and Ideas

5.3.1 Algorithm

We propose the following algorithm to check whether a quasi-qubit POVM $\mathbf{P}$ is clean or not.

(i) We check whether $\mathbf{P}$ is rank-one. If it is, exit with result "$\mathbf{P}$ *is clean*". Otherwise:

(ii) Write the rank-one elements $P_i = \lambda_i |\psi_i\rangle\langle\psi_i|$ for $1 \leq i \leq n$. Check whether these $|\psi_i\rangle$ generate $\mathcal{H}$. If not, exit with result "**P** *is not clean*". Else:

(iii) We can find a basis of $\mathcal{H}$ as a subset of those $|\psi_i\rangle$. We assume that this basis consists of $|\psi_i\rangle$ for $1 \leq i \leq d$. We define a variable $C = \{V_j\}_{j \in J}$, consisting in a collection of subspaces whose direct sum is the Hilbert space $\mathcal{H} = \bigoplus_j V_j$. We initialize C with $V_i = \mathrm{Span}(|\psi_i\rangle)$ for $1 \leq i \leq d$.

(iv) For i from $d+1$ to n, do:

(v) Write $|\psi_i\rangle = \sum_j v_j$ with $v_j \in V_j$. Call $J(i) = \{j | v_j \neq 0\}$.

(vi) Update $\{V_j\}$: Suppress all V_j for $j \in J(i)$. Add $V_i = \bigoplus_{j \in J(i)} V_j$.

(vii) Check whether $C = \{\mathcal{H}\}$. If so, exit with result "**P** *is clean*". Otherwise:

(viii) End of the "For" loop.

(ix) Exit with result "**P** *is not clean*".

Notice that the algorithm terminates: every stage is finite and we enter the loop a finite number of times.

5.3.2 Heuristics: what the algorithm really tests

In the Kraus decomposition (5.1), each of the terms $R_\alpha^* A R_\alpha$ is non-negative if A is non-negative, so that $\mathcal{E}(A) \geq R_\alpha^* A R_\alpha$ for any α. Hence if $\mathcal{E}(\mathbf{Q}) = \mathbf{P}$, then $R_\alpha^* Q_i R_\alpha$ must have support included in $\mathrm{Supp}(P_i)$ for all α and $e \in E$.

The central idea of the chapter is the following: the condition $\mathrm{Supp}(R_\alpha^* Q_i R_\alpha) \subset \mathrm{Supp}(P_i)$ yields $d - \dim(\mathrm{Supp}(P_i))$ homogeneous linear equations on the matrix entries of R_α, where you should remember that $d = \dim(\mathcal{H})$. Now R_α is determined up to a constant by $d^2 - 1$ homogeneous independent linear equations. In such a case, the additional condition $\sum R_\alpha^* R_\alpha = \mathbf{1}$ yields all R_α are proportional to the same unitary U, so that the channel $\mathcal{E}$ is unitary, and $\mathbf{P} \succ \mathbf{Q}$.

There is still one difficulty: the equations mentioned above depend not only on **P**, but also on **Q**. We would then like conditions on the supports of P_i such that the system of equations mentioned above is at least of rank $d^2 - 1$ for all **Q**. We formalize this requirement with the following definitions.

Definition 5.3.1. Corresponding

Let $\mathcal{V}$ be a Hilbert space and $\{F_i\}_{i\in I}$ a collection of subspaces of $\mathcal{V}$. Let $\{v_i\}_{i\in I}$ be a collection of vectors of $\mathcal{V}$. This set of vectors corresponds to *$\{F_i\}_{i\in I}$ if for any $i \in I$, there is a linear transform R_i such that $R_i(v_i) \neq 0$ and, for all $j \in I$, the transform is taking v_j within F_j, that is $R_i(v_j) \in F_j$.*

In the text, we usually drop the reference to $\{F_i\}_{i\in I}$ and write that the $\{v_i\}_{i\in I}$ are a corresponding *collection of vectors.*

Definition 5.3.2. Totally determined

Let $\mathcal{V}$ be a Hilbert space and $\{F_i\}_{i\in I}$ a collection of subspaces of $\mathcal{V}$.

If for all corresponding collections of vectors $\{v_i\}_{i\in I}$ there is only one (up to a complex multiplicative constant) linear transform R such that $R(v_i) \in F_i$ for all $i \in I$, we say that $\mathcal{V}$ is totally determined *by $\{F_i\}_{i\in I}$, or alternatively that $\{F_i\}_{i\in I}$* totally determines *$\mathcal{V}$.*

If F_i is one-dimensional with support vector w_i, this means there is only one R such that $R(v_i)$ is colinear to w_i for all $i \in I$.

What the algorithm does is checking that a quasi-qubit POVM $\mathbf{P}$ is rank-one (stage (i)), or that $\mathbf{P}$ totally determines $\mathcal{H}$.

More precisely, Proposition 5.4.9 states that each of the V_j belonging to C (appearing at stage (iii) and updated at stage (vi)) is totally determined by the $|\psi_i\rangle$ such that $|\psi_i\rangle \in V_j$. When the algorithm exits at stage (vii), then $C = \{\mathcal{H}\}$, so $\mathcal{H}$ is totally determined. If the algorithm does not exit at stage (vii), on the other hand, then C has at least two elements at the last stage, and each $|\psi_i\rangle$ is included in one of those two elements, which entails, from Lemma 5.5.3, that $\{\mathrm{Supp}(P_i)\}$ does not totally determine $\mathcal{H}$.

The equivalence with cleanness for quasi-qubit POVMs is still needed to get validity of the algorithm. This equivalence stems from Theorem 5.4.1 and Theorem 5.5.1. The former is the sufficient condition, for any POVM, not necessarily quasi-qubit. We have given the intuition for this theorem at the beginning of the section. Complementarily, Theorem 5.5.1 states that a strict quasi-qubit POVM is not clean if its supports do not totally determine $\mathcal{H}$.

The proof of Theorem 5.5.1 features the last important idea of the chapter. A channel $\mathcal{E}$ which is near enough the identity may be inverted as a positive map on $\mathcal{B}(\mathcal{H})$, even though $\mathcal{E}^{-1}$ is not a channel. Now if we denote $\mathbf{Q} = \mathcal{E}^{-1}(\mathbf{P})$, we have $\mathcal{E}(\mathbf{Q}) = \mathbf{P}$. We are then left with two questions: is $\mathbf{Q}$ a POVM, and can we find a channel $\mathcal{F}$ such that $\mathcal{F}(\mathbf{P}) = \mathbf{Q}$?

The main possible obstacle to $\mathbf{Q}$ being a POVM is the need for each of the Q_i to be non-negative. Now, if $\mathcal{E}$ is near enough the identity, if P_i was full-rank, then Q_i is still full-rank non-negative. The remaining case is $Q_i = \mathcal{E}^{-1}(P_i) = \lambda_i \mathcal{E}^{-1}(|\psi_i\rangle\langle\psi_i|)$. Now, we shall see that we may use the set of subspaces $C = \{V_j\}$ given by the algorithm to build channels ensuring that these Q_i are still rank-one non-negative matrices. Furthermore, these Q_i will have a bigger first eigenvalue than P_i, so that we are sure $\mathbf{Q}$ is strictly cleaner than $\mathbf{P}$, as channels are spectrum-width decreasing (see Lemma 5.5.2).

We now turn to the fully rigorous treatment.

5.4 Sufficient condition

We start by proving the following theorem, announced in the previous section.

Theorem 5.4.1. *If the supports* $\{\mathrm{Supp}(P_i)\}_{i\in I}$ *of the elements* P_i *of a POVM* $\mathbf{P}$ *totally determine* $\mathcal{H}$*, then* $\mathbf{P}$ *is clean and any cleanness-equivalent POVM* $\mathbf{Q}$ *is in fact unitarily equivalent to* $\mathbf{P}$.

Proof. It is enough to prove that if $\mathbf{Q} \succ \mathbf{P}$, then $\mathbf{Q}$ is unitarily equivalent to $\mathbf{P}$.

Let $\mathbf{Q}$ be a POVM and $\mathcal{E} = \{R_\alpha\}_\alpha$ a channel such that $\mathcal{E}(\mathbf{Q}) = \mathbf{P}$.

For all $i \in I$, we may write $Q_i = \sum_k \mu_{i,k} |\phi_i^k\rangle\langle\phi_i^k|$. Then we have

$$P_i = \sum_\alpha \sum_k \mu_{i,k} R_\alpha^* |\phi_i^k\rangle\langle\phi_i^k| R_\alpha.$$

Now $\mu_{i,k} R_\alpha^* |\phi_i^k\rangle\langle\phi_i^k| R_\alpha \geq \mathbf{0}$ for all k and α, and consequently $\mu_{i,k} R_\alpha^* |\phi_i^k\rangle\langle\phi_i^k| R_\alpha \leq P_i$. Hence $R_\alpha^* |\phi_i^k\rangle \in \mathrm{Supp}(P_i)$.

Moreover P_i is nonzero. So that there is at least one $k(i)$ and one $\alpha(i)$ for each i such that $R_\alpha^* |\phi_i^{k(i)}\rangle$ is nonzero. Thus $\{\phi_i^{k(i)}\}_{i\in I}$ corresponds to $\{\mathrm{Supp}(P_i)\}_{i\in I}$. As $\{\mathrm{Supp}(P_i)\}_{i\in I}$ totally determines $\mathcal{H}$, there is only one R, up to a constant, such that $R|\phi_i^{k(i)}\rangle \in \mathrm{Supp}(P_i)$ for all i. So that $R_\alpha = c(\alpha)R$ for all α. Since $\sum_\alpha R_\alpha^* R_\alpha = \mathbf{1}$, there is a constant such that λR_1 is unitary, and $\mathcal{E} = \{\lambda R_1\}$. So that $\mathbf{P}$ and $\mathbf{Q}$ are unitarily equivalent.

□

Before proving in Theorem 5.4.9 that "when the algorithm exits at stage (vii), then the supports of the POVM $\mathbf{P}$ totally determine $\mathcal{H}$", we need a few more tools.

We first need the notion of *projective frame*. Indeed, in the algorithm, we are dealing with supports of rank-one POVMs, that is essentially projective lines. And we want them to totally determine the space, that is essentially fix it. Projective frames are the most basic mathematical object meeting these requirements. We redefine them here, and reprove what basic properties we need; further information on projective frames may be found in most geometry or algebra textbooks, e.g. (Audin, 2002).

Definition 5.4.2. *A* projective frame $\{v_i\}_{1\leq i\leq d+1}$ *of a vector space* $\mathcal{V}$ *is a set of* $(\dim(\mathcal{V})+1)$ *vectors in general position, that is, such that any subset of* $\dim(\mathcal{V})$ *vectors is a basis of* $\mathcal{V}$.

Remark 5.4.3. *Equivalently we may say that* $\{v_i\}_{1\leq i\leq n}$ *is a basis of* $\mathcal{V}$ *and* $v_{d+1} = \sum_{i=1}^{n} c_i v_i$ *with all* $c_i \neq 0$.

Proposition 5.4.4. *A projective frame* $\Psi = \{e_i\}_{1\leq i\leq (n+1)}$ *of* $\mathcal{V}$ *totally determines* $\mathcal{V}$.

Proof. First we prove that if $\Phi = \{v_i\}_{1\leq i\leq (n+1)}$ is not a projective frame, the set of vectors $\{v_i\}_{1\leq i\leq (n+1)}$ does not correspond to Ψ. Indeed, as Φ is not a projective frame, we may find n vectors, say the n first, such that $\sum_{i=1}^{n} a_i v_i = 0$ with at least one a_i non-zero, say a_1. Then for any R such that $R(v_i)$ is colinear to e_i for all i, we still have $\sum_{i=1}^{n} a_i R(v_i) = 0$. As $\{e_i\}_{1\leq i\leq n}$ is a basis, $a_i R(v_i) = 0$ for all i, so that $R(v_1) = 0$. Hence $\{v_i\}_{1\leq i\leq n+1}$ does not correspond to $\{e_i\}_{1\leq i\leq n+1}$.

Let now $\Phi = \{v_i\}_{1\leq i\leq (n+1)}$ be corresponding to Ψ. Notably, this implies that Φ is a projective frame. Furthermore, there is a nonzero linear transform R such that $R(v_i)$ is colinear to e_i for all i. We must show that R is unique up to a constant.

We know that $\{e_i\}_{1\leq i\leq n}$ and $\{v_i\}_{1\leq i\leq n}$ are both bases of $\mathcal{V}$. Hence there is a unique transfer matrix X from the latter basis to the former. Since $R(v_i) = D_i e_i$ for some D_i, we know that R is of the form DX where D is a diagonal matrix with diagonal values D_i.

We still have not used our $(n+1)$th condition. We are dealing with projective frames, so that $e_{n+1} = \sum_{i=1}^{n} b_i e_i$ and $v_{n+1} = \sum_{i=1}^{n} c_i v_i$ with all b_i and c_i non-zero. Now $R(v_{n+1}) = \sum_{i=1}^{n} c_i R(v_i) = \sum_{i=1}^{n} c_i D_i e_i$, so that $c_i D_i / b_i$ must be independent on i and D and hence R is fixed up to a complex multiplicative constant.

□

We now turn to a few observations about totally determined spaces.

Remark 5.4.5. *If $\{F_i\}_{i\in I}$ totally determines $\mathcal{H}$, and if $\{v_i\}_{i\in I}$ corresponds to $\{F_i\}$, then the up to a constant unique nonzero R such that $Rv_i \in F_i$ for all $i \in I$ is invertible.*

Proof. Let us define $\Pi_{(\ker R)^\perp}$ the projector on the orthogonal of the kernel of R along its kernel, and $\Pi_{\ker R}$ the projector on the kernel of R along $(\ker R)^\perp$. We have $R = R\Pi_{(\ker R)^\perp}$, so that $R\Pi_{(\ker R)^\perp}v_i = Rv_i$. Thus $\{\Pi_{(\ker R)^\perp}v_i\}_{i\in I}$ is corresponding to $\{F_i\}_{i\in I}$. On the other hand, $\Pi_{\ker R}\Pi_{(\ker R)^\perp} = 0$, so that $(R+\Pi_{\ker R})(\Pi_{(\ker R)^\perp}v_i) = R(\Pi_{(\ker R)^\perp}v_i) \in F_i$. As $\{\Pi_{(\ker R)^\perp}\}$ is corresponding to $\{F_i\}$, the latter equality implies that R is proportional to $(R + \Pi_{\ker R})$. This is only possible if $\Pi_{\ker R} = 0$. Hence R is invertible. □

Remark 5.4.6. *If $\{v_l\}_{l\in I\cup J}$ is corresponding to $\{F_l\}_{l\in I\cup J}$, then $\{v_i\}_{i\in I}$ (resp. $\{v_j\}_{j\in J}$) is corresponding to $\{F_i\}_{i\in I}$ (resp. $\{F_j\}_{j\in J}$.*

Proof. The set I is a subset of $I \cup J$, thus, for all $i \in I$, there is an R_i such that $R_iv_i \neq 0$ and $R_iv_l \in F_l$ for all $l \in I \cup J$. *A fortiori* $R_iv_k \in F_k$ for all $k \in I$. Hence $\{v_i\}_{i\in I}$ is corresponding to $\{F_i\}_{i\in I}$. The same proof yields the result for J. □

Remark 5.4.7. *If $\{v_i\}_{i\in I}$ is corresponding to $\{F_i\}_{i\in I}$, then there exists R such that $Rv_i \in F_i$ and $Rv_i \neq 0$ for all i simultaneously.*

Proof. By the definition of "corresponding to", we have a set $\{R_i\}_{i\in I}$ of transforms such that $R_iv_i \neq 0$ and $R_iv_j \in F_j$ for all $j \in I$. Now, for any set of coefficients $\{a_i\}_{i\in I}$ the matrix $R = \sum_i a_iR_i$ fulfils $Rv_i \in F_i$ for all i. If we choose appropriately $\{a_i\}$ we also have $Rv_i \neq 0$. For example, we may write all the R_iv_i in the same basis, take note of all coordinates, and choose the a_i as any real numbers algebraically independent of those coordinates. □

Lemma 5.4.8. *If $\mathcal{V}$ and $\mathcal{W}$ are both totally determined by sets of subspaces $\{F_i\}_{i\in I}$ and $\{F_j\}_{j\in J}$ and if $\mathcal{V}$ and $\mathcal{W}$ intersect (apart from the null vector), then their sum $\mathcal{U} = \mathcal{V} + \mathcal{W}$ is totally determined by $\{F_l\}_{l\in I\cup J}$.*

Proof. Let $\{u_l\}_{l\in I\cup J}$ vectors of $\mathcal{U}$ correspond to $\{F_l\}_{l\in I\cup J}$. In other words, there is an R^* such that $R^*u_l \in F_l$ for all $l \in I \cup J$. By Remark 5.4.7, we may assume that $R^*u_l \neq 0$ for all l. We must show that R^* is unique up to a constant. Notice that the restriction $R^*u_l \neq 0$ does not play a role: if we find another R non proportional to R^*, such that $Ru_l \in F_l$ for all l, then $R^* + aR$ for appropriate a also fulfils $0 \neq (R^* + aR)u_l \in F_l$ for all l, and is not proportional to R^*.

We need a few notations. First, we consider the space $\mathcal{X} = \mathcal{V} \cap \mathcal{W}$. We also define $\mathcal{Y}$ by $\mathcal{V} = \mathcal{Y} \oplus \mathcal{X}$ and $\mathcal{Z}$ by $\mathcal{W} = \mathcal{Z} \oplus \mathcal{X}$. We write $I_\mathcal{V}$ and $I_\mathcal{W}$ for the natural inclusions

of $\mathcal{V}$ and $\mathcal{W}$ in $\mathcal{U}$. We also denote by $\Pi_{\mathcal{V}}$ for the projector on $\mathcal{V}$ along $\mathcal{Z}$, by $\Pi_{\mathcal{W}}$ the projector on $\mathcal{W}$ along $\mathcal{Y}$, and by $\Pi_{\mathcal{X}}$ the projector on $\mathcal{X}$ along $\mathcal{Y}+\mathcal{Z}$.

Please be aware that we do not define $\Pi_{\mathcal{V}}$ and $\Pi_{\mathcal{W}}$ as endomorphisms of $\mathcal{U}$, but as applications from $\mathcal{U}$ to $\mathcal{V}$ and $\mathcal{W}$, respectively. The corresponding endomorphisms are $I_{\mathcal{V}}\Pi_{\mathcal{V}}$ and $I_{\mathcal{W}}\Pi_{\mathcal{W}}$.

As a first step, we show that $I_{\mathcal{V}}\Pi_{\mathcal{V}}R^*$ is unique up to a constant.

The rank of $I_{\mathcal{V}}\Pi_{\mathcal{V}}R^*$ is at most $\dim(\mathcal{V})$, so we can factorize it by $\mathcal{V}$: there exists two linear applications $L^{\mathcal{U}}_{\mathcal{V}}$ from $\mathcal{U}$ to $\mathcal{V}$ and $L^{\mathcal{V}}_{\mathcal{U}}$ from $\mathcal{V}$ to $\mathcal{U}$, such that $I_{\mathcal{V}}\Pi_{\mathcal{V}}R^*L^{\mathcal{V}}_{\mathcal{U}}L^{\mathcal{U}}_{\mathcal{V}} = I_{\mathcal{V}}\Pi_{\mathcal{V}}R^*$.

Now for all $i \in I$, we have $R^*u_i \in F_i \subset \mathcal{V}$, so that $R^*u_i = I_{\mathcal{V}}\Pi_{\mathcal{V}}R^*u_i = I_{\mathcal{V}}\Pi_{\mathcal{V}}R^*L^{\mathcal{V}}_{\mathcal{U}}L^{\mathcal{U}}_{\mathcal{V}}u_i$, so that for all $i \in I$ we have the inclusion $0 \neq (\Pi_{\mathcal{V}}R^*L^{\mathcal{V}}_{\mathcal{U}})(L^{\mathcal{U}}_{\mathcal{V}}u_i) \in F_i$, where we have used $R^*u_l \neq 0$.. Thus $\{L^{\mathcal{U}}_{\mathcal{V}}u_i\}_{i\in I}$ is corresponding to $\{F_i\}_{i\in I}$. On the other hand, we know that $\{F_i\}_{i\in I}$ totally determine $\mathcal{V}$. Hence there is a nonzero constant $\lambda_{\mathcal{V}}$, and a $R_{\mathcal{V}}$ depending only on $\{F_i\}_{i\in I}$, such that $\Pi_{\mathcal{V}}R^*L^{\mathcal{V}}_{\mathcal{U}} = \lambda_{\mathcal{V}}R_{\mathcal{V}}$. Moreover, by Remark 5.4.5, $R_{\mathcal{V}}$ is invertible. So that finally $I_{\mathcal{V}}\Pi_{\mathcal{V}}R^* = \lambda_{\mathcal{V}}I_{\mathcal{V}}R_{\mathcal{V}}L^{\mathcal{U}}_{\mathcal{V}}$, with image $\mathrm{im}(\lambda_{\mathcal{V}}I_{\mathcal{V}}R_{\mathcal{V}}L^{\mathcal{U}}_{\mathcal{V}}) = \mathcal{V}$. Replacing $\mathcal{V}$ with $\mathcal{W}$, we get similarly $I_{\mathcal{W}}\Pi_{\mathcal{W}}R^* = \lambda_{\mathcal{W}}I_{\mathcal{W}}R_{\mathcal{W}}L^{\mathcal{U}}_{\mathcal{W}}$.

The last step consists in proving that the two constants $\lambda_{\mathcal{V}}$ and $\lambda_{\mathcal{W}}$ are proportional, independently of R^*.

We notice that $\Pi_{\mathcal{X}}I_{\mathcal{V}}\Pi_{\mathcal{V}} = \Pi_{\mathcal{X}} = \Pi_{\mathcal{X}}I_{\mathcal{W}}\Pi_{\mathcal{W}}$. Hence $\lambda_{\mathcal{V}}\Pi_{\mathcal{X}}I_{\mathcal{V}}R_{\mathcal{V}}L^{\mathcal{U}}_{\mathcal{V}} = \lambda_{\mathcal{W}}\Pi_{\mathcal{X}}I_{\mathcal{W}}R_{\mathcal{W}}L^{\mathcal{U}}_{\mathcal{W}}$. As $\mathcal{X} \subset \mathcal{V}$ and $\mathrm{im}(\lambda_{\mathcal{V}}I_{\mathcal{V}}R_{\mathcal{V}}L^{\mathcal{U}}_{\mathcal{V}}) = \mathcal{V}$, we know that $\lambda_{\mathcal{V}}\Pi_{\mathcal{X}}I_{\mathcal{V}}R_{\mathcal{V}}L^{\mathcal{U}}_{\mathcal{V}} \neq 0$. The equality $\lambda_{\mathcal{V}}\Pi_{\mathcal{X}}I_{\mathcal{V}}R_{\mathcal{V}}L^{\mathcal{U}}_{\mathcal{V}} = \lambda_{\mathcal{W}}\Pi_{\mathcal{X}}I_{\mathcal{W}}R_{\mathcal{W}}L^{\mathcal{U}}_{\mathcal{W}}$ then yields the proportionality of $\lambda_{\mathcal{W}}$ and $\lambda_{\mathcal{V}}$.

We conclude by recalling that $\mathcal{V}+\mathcal{W} = \mathcal{U}$, so that knowing both $I_{\mathcal{V}}\Pi_{\mathcal{V}}R^*$ and $I_{\mathcal{W}}\Pi_{\mathcal{W}}R^*$ is equivalent to knowing R^*. As our only free parameter is the multiplicative constant $\lambda_{\mathcal{V}}$, we have proved uniqueness of R^*, up to a constant.

□

Lemma 5.4.8 and Proposition 5.4.4 are the two ingredients for proving the following proposition, central for the validity of the algorithm.

Proposition 5.4.9. *In the algorithm, the spaces in the set $C = \{V_j\}_{j\in J}$ are always totally determined by the supports $K(j) = \{\mathrm{Span}(|\psi_i\rangle) : |\psi_i\rangle \in V_j\}$ of the one-dimensional POVM elements they contain.*

Proof. We prove the proposition by induction on the stronger property *Prop* = " all V_j are totally determined by $K(j)$, and they are spanned by vectors of the initial basis, that is, they are of the form $\mathrm{Span}(|\psi_i\rangle : i \in I(j))$, where $I(j)$ is a subset of $\{1, \ldots, d\}$".

Initialization: We initialize C at step (iii). At this stage V_j is defined for $j \in \{1, \ldots, d\}$ by $V_j = \mathrm{Span}(|\psi_j\rangle)$. So that on the one hand V_j is of the form $\mathrm{Span}(|\psi_i\rangle : i \in I(j))$, where $I(j)$ is a subset of $\{1, \ldots, d\}$, and on the other hand V_j is totally determined by $K(j)$, as it is one-dimensional and $|\psi_j\rangle$ is nonzero.

Update: We update C at stage (vi). We must prove that $V_i = \bigoplus_{j \in J(i)} V_j$ still fulfils *Prop*.

For one thing, the space V_i is a sum of spaces of the form $\mathrm{Span}(|\psi_i\rangle : i \in I(j))$, where $I(j)$ is a subset of $\{1, \ldots, d\}$, hence V_i is also of this form with $I(i) = \bigcup_{j \in J(i)} I(j)$.

Now let us consider the set $I_{int} = \{j : j \in \{1 \ldots d\}, \langle\psi_i|\psi_j\rangle \neq 0\}$, and the space $V_{int} = \mathrm{Span}(|\psi_j\rangle : j \in I_{int})$. Since the $|\psi_j\rangle$ for $j \in I_{int}$ are part of the initial basis $\{|\psi_j\rangle\}_{1 \leq j \leq d}\}$, they are independent. The definition of I_{int} also ensures $|\psi_i\rangle = \sum_{j \in I_{int}} c_j |\psi_j\rangle$ with j nonzero, hence, by Remark (5.4.3), the set $\{|\psi_k\rangle : k = k \in I_{int} \cup \{i\}\}$ is a projective frame of V_{int}. So that, by Proposition 5.4.4, the space V_{int} is totally determined by $\{|\psi_j\rangle\}_{j \in I_{int} \cup \{i\}}$. We initialize $K_{int} = I_{int} \cup \{i\}$.

Finally, by definition of $J(i)$, we know that $V_{int} \cap V_j \neq 0$ for all $j \in J(i)$. Both are totally determined, by $K(j)$ and K_{int}. Hence by Lemma 5.4.8, $V_{int} \cup V_j$ is totally determined by $K(j) \cup K_{int}$. We update $V_{int} = V_{int} \cup V_j$ and $K_{int} = K_{int} \cup K(j)$. We iterate the latter step for all $j \in J(i)$ and we end up with $V_{int} = V_i$ totally determined by $\bigcup_{j \in j(i)} K(j) \cup I_{int} \cup \{i\} \subset I(i)$.

□

Corollary 5.4.10. *When the algorithm ends at stage (vii), the POVM* $\mathbf{P}$ *is clean.*

Proof. The algorithm ends at stage (vii) only if $C = \{\mathcal{H}\}$. By the above proposition, this condition implies that $\mathcal{H}$ is totally determined by $\{\mathrm{Span}(|\psi_j\rangle) : |\psi_j\rangle \in \mathcal{H}\}$. This amounts at saying that $\mathcal{H}$ is totally determined by the supports of the POVM elements P_i, and we conclude by Theorem 5.4.1. □

This section aims at giving sufficient conditions for a POVM to be clean, and at proving that one of these conditions is fulfilled if the algorithm exits with result "$\mathbf{P}$ is clean". We thus conclude the section with the case when the algorithm exits at stage (i). In other words, we must show that a rank-one POVM is clean. Now, this has already been proved as Theorem 11.2 of (Buscemi et al., 2005):

Theorem 5.4.11. *(Buscemi et al., 2005) If* $\mathbf{P}$ *is rank-one, then* $\mathbf{Q} \succ \mathbf{P}$ *if and only if* $\mathbf{P}$ *and* $\mathbf{Q}$ *are unitarily equivalent. Thus, rank-one POVMs are clean.*

For a quasi-qubit POVM $\mathbf{P}$, we prove in the following section that $\mathbf{P}$ is clean only if it fulfils the conditions either of Theorem 5.4.11 or of Theorem 5.4.1.

5.5 Necessary condition for quasi-qubit POVMs

This section proves that a clean quasi-qubit POVM either is rank-one, or the supports of its elements totally determine the space:

Theorem 5.5.1. *A non-rank-one quasi-qubit POVM where* $\{\mathrm{Supp}(P_i)_{i\in I}\}$ *does not determine* $\mathcal{H}$ *is not clean.*

We need a few more tools to prove the theorem.

To begin with, we need a way to prove in specific situations that a POVM is not cleaner than another. Using the fact that channels are *spectrum-width decreasing* is the easiest method. This is Lemma 3.1 of (Buscemi et al., 2005):

Lemma 5.5.2. *If the minimal (resp. maximal) eigenvalue of* X *is denoted* $\lambda_m(X)$ *(resp.* $\lambda_M(X)$*), then* $\lambda_m(X) \leq \lambda_m(\mathcal{E}(X)) \leq \lambda_M(\mathcal{E}(X)) \leq \lambda_M(X)$ *for any channel* $\mathcal{E}$*.*

This lemma implies that existence of $\mathbf{Q} \succ \mathbf{P}$ such that for some $i \in I$, either $\lambda_m(Q_i) < \lambda_m(P_i)$ or $\lambda_M(Q_i) > \lambda_M(P_i)$ entails that $\mathbf{Q}$ is strictly cleaner than $\mathbf{P}$, so that $\mathbf{P}$ is not clean.

We now give a characterization of the fact that $\mathcal{H}$ is totally determined by $\{F_j\}_{j\in J}$ when all the F_j are one-dimensional, that is of when the F_j can be seen as vectors. This characterization applies to $\{\mathrm{Supp}(P_i)\}_{i\in I}$ for quasi-qubit POVMs, and may be more intuitive than Definition 5.3.2. Moreover it is more adapted to our strategy of proof.

Lemma 5.5.3. *A set of vectors* $\{|\psi_j\rangle\}_{j\in J}$ *totally determine the space* $\mathcal{H}$*, if and only if, for any two supplementary proper subspaces* $\mathcal{V}$ *and* $\mathcal{W}$*, there is a* $j \in J$ *such that* $|\psi_j\rangle \notin \mathcal{V}$ *and* $|\psi_j\rangle \notin \mathcal{W}$*.*

Moreover, when the algorithm exits with result "$\mathbf{P}$ *is not clean", the supports of* $\mathbf{P}$ *do not totally determine* $\mathcal{H}$*.*

Proof. The proof is made of four steps:

(a) For any finite set of vectors $\{|\psi_j\rangle\}_{j\in J}$, there is a POVM whose supports of the rank-one elements are these vectors.

(b) if we feed into the algorithm a non-rank-one quasi-qubit POVM whose supports of rank-one elements are the $|\psi_j\rangle$ and if $\{|\psi_j\rangle\}$ does not totally determine $\mathcal{H}$, then the algorithm exits with result "**P** is not clean".

(c) if the algorithm exits with result "**P** is not clean", then we can find two supplementary proper subspaces such that $|\psi_j\rangle \in \mathcal{V}$ or $|\psi_j\rangle \in \mathcal{W}$ for all supports of rank-one elements.

(d) finding two supplementary proper subspaces such that $|\psi_j\rangle \in \mathcal{V}$ or $|\psi_j\rangle \in \mathcal{W}$ for all $j \in J$ implies that $\{|\psi_j\rangle\}_{j\in J}$ does not totally determine $\mathcal{H}$.

The equivalence in the lemma is then proved by contraposition, and the last statement by combining (c) and (d).

Step (a): A valid example is given by $P_j = \frac{1}{2\#J}|\psi_j\rangle\langle\psi_j|$ for $j \in J$ and $P_{\#J+1} = \mathbf{1} - \sum_j P_j$. Indeed the latter element is positive since $\sum_j P_j \leq \frac{1}{2\#J}\#J\mathbf{1} = \frac{1}{2}\mathbf{1}$.

Step (b): Since the quasi-qubit POVM is assumed not to be rank-one, we do not exit at stage (i). The only other possible exit with result "**P** is clean" is at stage (vii). Now the proof of Corollary 5.4.10 states that the algorithm exits at stage (vii) only if the supports of the rank-one elements totally determine $\mathcal{H}$. Hence, the algorithm exits with result "**P** is not clean".

Step (c): Exiting at stage (ii) means that the $|\psi_j\rangle$ do not generate $\mathcal{H}$. Then, if $J = \varnothing$, we may choose any two supplementary proper subspaces $\mathcal{V}$ and $\mathcal{W}$. Anyhow $|\psi_j\rangle \in \mathcal{V}$ for all $j \in J$. If $J \neq \varnothing$, then $\mathcal{V} = \mathrm{Span}(|\psi_i\rangle, i \in I)$ is a proper subspace of $\mathcal{H}$. Since $|\psi_j\rangle \in \mathcal{V}$ for all $j \in J$, any supplementary subspace $\mathcal{W}$ of $\mathcal{V}$ will turn the trick.

If the algorithm does not exit at stage (ii), then there is a basis included in $\{|\psi_j\rangle\}_{j\in J}$. We assume that it corresponds to $1 \leq j \leq d$.

Since the algorithm exits with result, "**P** is not clean", it exits at stage (ix). We end the algorithm with a collection $C = \{V_k\}$ of subspaces such that $\bigoplus_k V_k = \mathcal{H}$. Since we have not exited at stage (vii), we know that $C \neq \{\mathcal{H}\}$. Hence C counts at least two non-trivial elements. We take $\mathcal{V} = V_1$ and $\mathcal{W} = \bigoplus_{k\neq 1} V_k$.

The V_k are direct sums of the original $V_j = \mathrm{Span}(|\psi_j\rangle)$ for $1 \leq j \leq d$. Hence, for $1 \leq j \leq d$, either $|\psi_j\rangle \in \mathcal{V}$ or $|\psi_j\rangle \in \mathcal{W}$. On the other hand if $|\psi_j\rangle$ is not one of the

original basis vectors, it was used in the "For" loop. At the end of this loop, C was then containing a space $V = \bigoplus_{k \in J(j)} V_k$. And $|\psi_j\rangle$ was included in this space. This V is then included in one of the final V_j and a fortiori either in $\mathcal{V}$ or in $\mathcal{W}$. We have thus proved that when the algorithm exits with a negative value we may find two supplementary proper subspaces $\mathcal{V}$ and $\mathcal{W}$ such that for all $i \in I$, either $|\psi_i\rangle \in \mathcal{V}$ or $|\psi_i\rangle \in \mathcal{W}$.

Step (d): Since $\mathbf{1}|\psi_j\rangle = |\psi_j\rangle$ for all j, by Definition 5.3.1 the set of vectors $\{|\psi_j\rangle\}_{j\in J}$ is corresponding to the subspaces $\{|\psi_j\rangle\}_{j\in J}$. On the other hand, denoting by $\Pi_{\mathcal{V}}$ the projection on $\mathcal{V}$ parallel to $\mathcal{W}$, we get that $\Pi_{\mathcal{V}}|\psi_j\rangle$ is colinear to $|\psi_j\rangle$ for all $j \in J$. Moreover $\Pi_{\mathcal{V}}$ is not proportional to $\mathbf{1}$, so that, by definition 5.3.2, the set of vectors $\{|\psi_j\rangle\}$ does not totally determine $\mathcal{H}$.

□

Finally, as explained in Section 5.3, we want to build our cleaner POVMs as $\mathcal{E}^{-1}(\mathbf{P})$ where the channel is inverted as a positive map. We need to know some conditions under which a channel can be inverted. This is the purpose of Lemma 5.5.4, for which we need the following norms.

The Hilbert-Schmidt norm on $\mathcal{B}(\mathcal{H})$ is defined as $\|M\|_{HS}^2 = \mathrm{Tr}(MM^*)$. Notably, in any orthogonal basis,

$$\|M\|_{HS}^2 = \sum_{1\leq i,j\leq d} |M_{i,j}|^2.$$

Moreover $\|M\|_{HS} = \|M^*\|_{HS}$.

We also define a norm on $\mathcal{B}(\mathcal{B}(\mathcal{H}))$, space to which the channels belong:

$$\|\mathcal{O}\|_1 = \sup_{\{M | \|M\|_{HS}=1\}} \|\mathcal{O}(M)\|_{HS}.$$

Lemma 5.5.4. *If in the Kraus representation of a channel $\mathcal{E} = \{R_\alpha\}$ one of the R_α fulfils*

$$\|\mathbf{1} - R_\alpha\|_{HS} \leq \epsilon,$$

then

$$\|\mathbf{1} - \mathcal{E}\|_1 \leq 2(1 + \sqrt{d})\epsilon + 2\epsilon^2 = f(\epsilon) \underset{\epsilon\to 0}{\longrightarrow} 0. \tag{5.2}$$

As a consequence, if $f(\epsilon) < 1$, then $\mathcal{E}$ is invertible (as a map on $\mathcal{B}(\mathcal{H})$) and $\|\mathcal{E}^{-1} - \mathbf{1}\|_1 \leq f(\epsilon)/(1 - f(\epsilon))$. This inverse lets $\mathcal{B}_{sa}(\mathcal{H})$ stable.

This in turn shows that for any $X \in \mathcal{B}_{sa}(\mathcal{H})$ such that $\lambda_m(X) \geq 0$, the spectrum of the image by the inverse is bounded through

$$\lambda_m(X) - \lambda_M(X) f(\epsilon)\sqrt{d}/(1 - f(\epsilon)) \leq \lambda_m(\mathcal{E}^{-1}(X)). \tag{5.3}$$

So that for all $X > 0$, when ϵ small enough, $\mathcal{E}^{-1}(X) \geq 0$.

Remark: The bound 5.2 is probably far from sharp, but sufficient for our needs.

Proof. Without loss of generality, we assume that

$$\|\mathbf{1} - R_1\|_{HS} \leq \epsilon.$$

We write $S = R_1 - \mathbf{1}_{\mathcal{H}}$ and $\mathcal{O} = \mathcal{E} - \mathbf{1}_{\mathcal{B}(\mathcal{H})}$.

Then

$$\mathcal{O} : M \mapsto S^*MS + S^*M + MS + \sum_{\alpha \neq 1} R_\alpha^* M R_\alpha.$$

And

$$\begin{aligned}
\|\mathcal{O}\|_1 &= \sup_{\{M | \|M\|_{HS}=1\}} \left\| S^*MS + S^*M + MS + \sum_{\alpha \neq 1} R_\alpha^* M R_\alpha \right\|_{HS} \\
&\leq \sup_{\{M | \|M\|_{HS}=1\}} \|S^*\|\|M\|\|S\| + \|S^*\|\|M\| \\
&\qquad + \|M\|\|S\| + \sum_{\alpha \neq 1} \|R_\alpha^*\|\|M\|\|R_\alpha\| \\
&= \|S\|_{HS}^2 + 2\|S\|_{HS} + \sum_{\alpha \neq 1} \|R_\alpha\|_{HS}^2.
\end{aligned}$$

Now, for one thing, by hypothesis, $\|S\|_{HS} \leq \epsilon$. Furthermore

$$\sum_{\alpha \neq 1} \|R_\alpha\|_{HS}^2 = \sum_{\alpha \neq 1} \mathrm{Tr}(R_\alpha^* R_\alpha) = \mathrm{Tr}(\mathbf{1} - R_1^* R_1) = -\mathrm{Tr}(S^*S + S + S^*).$$

We finish our proof of 5.2 with the observation that $-\mathrm{Tr}(S + S^*) \leq 2\sqrt{d}\|S\|_{HS} = 2\sqrt{d}\epsilon$.

If $\|\mathcal{O}\|_1 < 1$, we know that $\mathcal{E} = \mathbf{1} + \mathcal{O}$ is invertible and $\mathcal{E}^{-1} = \sum_{n \geq 0} (-\mathcal{O})^n$. By taking the norm, $\|\mathcal{E}^{-1} - \mathbf{1}\|_1 \leq \sum_{n \geq 1} \|\mathcal{O}\|_1^n = f(\epsilon)/(1 - f(\epsilon))$.

Channels stabilize $\mathcal{B}_{sa}(\mathcal{H})$; as $\mathcal{E}$ is furthermore invertible, equality of dimension shows that $\mathcal{E}(\mathcal{B}_{sa}(\mathcal{H})) = \mathcal{B}_{sa}(\mathcal{H})$ and $\mathcal{E}^{-1}(\mathcal{B}_{sa}(\mathcal{H})) = \mathcal{B}_{sa}(\mathcal{H})$.

Now, X is positive, so that $\|X\|_{HS} \leq \sqrt{d}\lambda_M(X)$. This implies $\|(\mathcal{E}^{-1} - \mathbf{1})(X)\|_{HS} \leq \sqrt{d}\lambda_M(X) f(\epsilon)/(1 - f(\epsilon))$, and in turn $\mathcal{E}^{-1}(X) \geq X - \sqrt{d}\lambda_M(X) f(\epsilon)/(1 - f(\epsilon))\mathbf{1}$. Taking the bottom of the spectrum ends the proof.

□

We are now ready to prove Theorem 5.5.1.

Proof of Theorem 5.5.1. We aim at exhibiting a channel $\mathcal{E}$ and a POVM $\mathbf{Q}$ such that $\mathcal{E}(\mathbf{Q}) = \mathbf{P}$ and Q_i has a wider spectrum than P_i for some $e \in E$. Then Lemma 5.5.2 proves that $\mathbf{Q}$ is strictly cleaner than $\mathbf{P}$, and in turn that $\mathbf{P}$ is not clean.

The building blocks are the subspaces supplied by Lemma 5.5.3. Since $\mathcal{H}$ is not determined by $\{\mathrm{Supp}(P_i)\}_{i\in I}$, there are two supplementary proper subspaces $\mathcal{V}$ and $\mathcal{W}$ such that each rank-one element has support included either in $\mathcal{V}$ or in $\mathcal{W}$.

We shall write explicitly several matrices in the forthcoming proof. All of them shall be written on an orthonormal basis $\{e_j\}_{1\leq j\leq d}$ of $\mathcal{H}$, chosen so that $\{e_j\}_{1\leq j\leq \dim(\mathcal{V})}$ is a basis of $\mathcal{V}$. We shall express the matrices as two-by-two block matrices, the blocks corresponding to the subspaces $\mathcal{V}$ and $\mathcal{V}^\perp$.

We study separately the following cases:

(a) All POVM elements P_i are proportional to the identity, that is $P_i = \mu_i \mathbf{1}$.

(b) The POVM is not full-rank, each rank-one element has support either in $\mathcal{V}$ or in $\mathcal{V}^\perp$, and all POVM elements are block-diagonal in $\mathcal{V}$ and $\mathcal{V}^\perp$.

(c) Each rank-one element has support either in $\mathcal{V}$ or $\mathcal{V}^\perp$, and at least one POVM element is not block-diagonal.

(d) At least one rank-one element has support neither in $\mathcal{V}$ nor in $\mathcal{V}^\perp$.

As a sanity check, let us prove we did not forget any case. Either our POVM is full-rank, or it is not. In the latter situation, either there is a rank-one element whose support is not included in $\mathcal{V}$ nor in $\mathcal{V}^\perp$ – and we are in case (d) –, or all rank-one elements are included in $\mathcal{V}$ or $\mathcal{V}^\perp$. Then either there is a POVM element that is not block-diagonal – and we are in case (c) – or all POVM elements are block-diagonal – and we are in case (b). On the other hand, if $\mathbf{P}$ is full-rank, we may choose the subspaces $\mathcal{V}$ and $\mathcal{W}$ any way we like. Notably, if one POVM element P_i is not proportional to the identity, so that it has non-trivial eigenspaces, we may choose $\mathcal{V}$ such that P_i is not block-diagonal in $\mathcal{V}$ and $\mathcal{V}^\perp$ – and we are in case (c). Finally, if on the contrary, all POVM elements are proportional to the identity, we are in case (a).

Case (a): If all POVM elements are of the form $P_i = \mu_i \mathbf{1}$, then, for any $\mathcal{E} = \{R_\alpha\}$, we have $\mathcal{E}(P_i) = \sum_\alpha R_\alpha^*(\mu_i \mathbf{1})R_\alpha = \mu_i \sum_\alpha R_\alpha^* R_\alpha = \mu_i \mathbf{1} = P_i$. No channel can change the wholly uninformative measurement $\mathbf{P}$.

On the other hand, many POVMs can be degraded to $\mathbf{P}$. Consider for example the POVM given by $Q_1 = \mu_1 |e_1\rangle\langle e_1| + \sum_{j=2}^{d} |e_j\rangle\langle e_j|$ and $Q_i = \mu_i |e_1\rangle\langle e_1|$ for $i > 1$. Then $\mathbf{Q} \neq \mathbf{P}$, so that $\mathbf{P} \not\succ \mathbf{Q}$. Yet, with $R_\alpha = |e_1\rangle\langle e_\alpha|$ for $1 \leq \alpha \leq d$, we have $\mathcal{E}(\mathbf{Q}) = \mathbf{P}$, and $\mathbf{Q} \succ \mathbf{P}$. Hence $\mathbf{P}$ is not clean.

Case (b): Since all rank-one elements are included either in $\mathcal{V}$ or in $\mathcal{V}^\perp$, we take $\mathcal{W} = \mathcal{V}^\perp$. We further choose $\mathcal{V}$ to be the smaller of the two subspaces, that is $\dim(\mathcal{V}) \leq d/2 \leq \dim(\mathcal{W})$. Then there is a matrix $A : \mathcal{V} \to \mathcal{W}$ such that $AA^* = \mathbf{1}_\mathcal{V}$. If all rank-one elements have support in $\mathcal{W}$, we further impose that at least one of these supports is not included in the kernel of A.

We then define $R^*_\mathcal{V}$ and $R^*_\mathcal{W}$ as:

$$R^*_\mathcal{V}(\epsilon) = \left[\begin{array}{c|c} \mathbf{1}_\mathcal{V} & \epsilon A \\ \hline 0 & 0 \end{array}\right],$$

$$R^*_\mathcal{W}(\epsilon) = \left[\begin{array}{c|c} 0 & 0 \\ \hline 0 & \mathbf{1}_\mathcal{W} \end{array}\right].$$

Their images are respectively $\mathcal{V}$ and $\mathcal{W}$.

From $R_\mathcal{V}(\epsilon)$ and $R_\mathcal{W}(\epsilon)$, we define the channel $\mathcal{E}_\epsilon = \{R_1(\epsilon), R_2(\epsilon), R_3(\epsilon)\}$:

$$R^*_1(\epsilon) = \sqrt{\tfrac{\epsilon^2}{1+\epsilon^2}} R^*_\mathcal{V}(\epsilon) + \sqrt{\tfrac{1-\epsilon^2}{1+\epsilon^2}} R^*_\mathcal{W}(\epsilon) = \left[\begin{array}{c|c} \sqrt{\frac{\epsilon^2}{1+\epsilon^2}}\mathbf{1}_\mathcal{V} & \sqrt{\frac{\epsilon^4}{1+\epsilon^2}} A \\ \hline 0 & \sqrt{\frac{1-\epsilon^2}{1+\epsilon^2}}\mathbf{1}_\mathcal{W} \end{array}\right],$$

$$R^*_2(\epsilon) = \sqrt{\tfrac{\epsilon^2}{1+\epsilon^2}} R^*_\mathcal{W}(\epsilon) = \left[\begin{array}{c|c} 0 & 0 \\ \hline 0 & \sqrt{\frac{\epsilon^2}{1+\epsilon^2}}\mathbf{1}_\mathcal{W} \end{array}\right],$$

$$R^*_3(\epsilon) = \sqrt{\tfrac{1-\epsilon^2}{1+\epsilon^2}} R^*_\mathcal{V}(\epsilon) - \sqrt{\tfrac{\epsilon^2}{1+\epsilon^2}} R^*_\mathcal{W}(\epsilon) = \left[\begin{array}{c|c} \sqrt{\frac{1-\epsilon^2}{1+\epsilon^2}}\mathbf{1}_\mathcal{V} & \sqrt{\frac{\epsilon^2-\epsilon^4}{1+\epsilon^2}} A \\ \hline 0 & -\sqrt{\frac{\epsilon^2}{1+\epsilon^2}}\mathbf{1}_\mathcal{W} \end{array}\right].$$

Since $AA^* = \mathbf{1}_\mathcal{V}$, we have $\sum_\alpha R^*_\alpha R_\alpha = \mathbf{1}$, hence these matrices $\{R_\alpha\}$ define a genuine channel. A few calculations show that the effect of this channel is:

$$\mathcal{E}_\epsilon : \left[\begin{array}{c|c} B & C \\ \hline C^* & D \end{array}\right] \to \left[\begin{array}{c|c} \frac{1}{1+\epsilon^2}\left(B + \epsilon(AC^* + CA^*) + \epsilon^2 ADA^*\right) & 0 \\ \hline 0 & D \end{array}\right]. \tag{5.4}$$

Now, for any $w \in \mathcal{W}$, we have

$$\left[\begin{array}{c} -\epsilon Aw \\ \hline w \end{array}\right] \left[\begin{array}{c} -\epsilon Aw \\ \hline w \end{array}\right]^* = \left[\begin{array}{c|c} \epsilon^2 Aww^*A^* & -\epsilon Aww^* \\ \hline -\epsilon ww^*A^* & ww^* \end{array}\right],$$

so that for any sequence of $w_j \in \mathcal{W}$, the matrix $\sum_{j,k} \left[\begin{array}{c|c} \epsilon^2 A w_j w_k^* A^* & -\epsilon A w_j w_k^* \\ \hline -\epsilon w_j w_k^* A^* & w_j w_k^* \end{array}\right]$ is non-negative. As any non-negative endomorphism D of $\mathcal{W}$ can be written $\sum_{j,k} w_j w_k^*$ for appropriate w_j, we get that for any non-negative D, the matrix $\left[\begin{array}{c|c} \epsilon^2 ADA^* & -\epsilon AD \\ \hline -\epsilon DA^* & D \end{array}\right]$ is non-negative. Moreover applying equation (5.4) yields that its image by $\mathcal{E}_\epsilon$ is $\left[\begin{array}{c|c} 0 & 0 \\ \hline 0 & D \end{array}\right]$.

Similarly, if $B \in \mathcal{B}(\mathcal{V})$ is non-negative, then $\left[\begin{array}{c|c} (1+\epsilon^2)B & 0 \\ \hline 0 & 0 \end{array}\right]$ is non-negative and its image by $\mathcal{E}_\epsilon$ is $\left[\begin{array}{c|c} B & 0 \\ \hline 0 & 0 \end{array}\right]$.

We use these observations to define a map (not a channel) $\mathcal{F}_\epsilon$ on the block-diagonal matrices:

$$\mathcal{F}_\epsilon : \left[\begin{array}{c|c} B & 0 \\ \hline 0 & D \end{array}\right] \to \left[\begin{array}{c|c} (1+\epsilon^2)B + \epsilon^2 ADA^* & -\epsilon AD \\ \hline -\epsilon DA^* & D \end{array}\right]. \tag{5.5}$$

We get that $\mathcal{E}_\epsilon(\mathcal{F}_\epsilon(M)) = M$ for all block-diagonal M and that if furthermore $M \geq 0$ then $\mathcal{F}_\epsilon(M) \geq 0$.

We now isolate one full-rank element of $\mathbf{P}$, say P_1. For all $i \neq 1$, we define $Q_i(\epsilon) = \mathcal{F}_\epsilon(P_i)$. They are non-negative and fulfil $\mathcal{E}_\epsilon(Q_i(\epsilon)) = P_i$. Define now $Q_1(\epsilon) = \mathbf{1} - \sum_{i\neq 1} Q_i(\epsilon)$. The closure relation ensures that $\mathcal{E}_\epsilon(Q_1(\epsilon)) = P_1$. What's more, recalling that $\sum_i B_i = \mathbf{1}_\mathcal{V}$ and $\sum_i D_i = \mathbf{1}_\mathcal{W}$, we obtain:

$$\begin{aligned} Q_1(\epsilon) &= \left[\begin{array}{c|c} \mathbf{1}_\mathcal{V} - (1+\epsilon^2)\sum_{i\neq 1} B_i - \epsilon^2 A(\sum_{i\neq 1} D_i)A^* & \epsilon A \sum_{i\neq 1} D_i \\ \hline -\epsilon \sum_{i\neq 1} D_i A^* & \mathbf{1}_\mathcal{W} - \sum_{i\neq 1} D_i \end{array}\right] \\ &= \left[\begin{array}{c|c} (1+\epsilon^2)B_1 + \epsilon^2 AD_1A^* - 2\epsilon^2 \mathbf{1}_\mathcal{V} & \epsilon A(\mathbf{1}_\mathcal{W} - D_1) \\ \hline \epsilon(\mathbf{1}_\mathcal{W} - D_1)A^*) & D_1 \end{array}\right] \\ &\xrightarrow[\epsilon\to 0]{} \left[\begin{array}{c|c} B_1 & 0 \\ \hline 0 & D_1 \end{array}\right] \\ &= P_1. \end{aligned}$$

Since P_1 is positive, this convergence entails the non-negativity of $Q_1(\epsilon)$ for ϵ small enough. As $Q_1(\epsilon)$ has been chosen so that $\sum_i Q_i(\epsilon) = \mathbf{1}$, we have defined a genuine POVM $\mathbf{Q}(\epsilon) = \{Q_i(\epsilon)\}_{i\in I}$ such that $\mathcal{E}_\epsilon(\mathbf{Q}(\epsilon)) = \mathbf{P}$, hence $\mathbf{Q} \succ \mathbf{P}$.

We end the study of this case by considering a rank-one element $P_i = \mu_i |\psi_i\rangle\langle\psi_i|$ whose support is not in the kernel of A. Using formula (5.5), if $|\psi_i\rangle \in \mathcal{V}$, we get $\mathrm{Tr}(Q_i(\epsilon)) = (1+\epsilon^2)\,\mathrm{Tr}(P_i) > \mathrm{Tr}(P_i)$, else $|\psi_i\rangle \in \mathcal{W}$ and we get $\mathrm{Tr}(Q_i(\epsilon)) = \mathrm{Tr}(P_i) + \epsilon^2\,\mathrm{Tr}(A|\psi_i\rangle\langle\psi_i|A^*) > \mathrm{Tr}(P_i)$. In both cases, bigger trace implies that the

spectrum of $Q_i(\epsilon)$ is wider than that of P_i and Lemma 5.5.2 yields $\mathbf{P} \not\succ \mathbf{Q}$. So that $\mathbf{P}$ is not clean.

Case (c): Since all rank-one elements are included either in $\mathcal{V}$ or in $\mathcal{V}^\perp$, we take $\mathcal{W} = \mathcal{V}^\perp$.

We now define the channel $\mathcal{E}_\epsilon$ through:

$$R_1(\epsilon) = \epsilon \Pi_{\mathcal{V}}, \quad R_2(\epsilon) = \epsilon \Pi_{\mathcal{W}} = \epsilon \Pi_{\mathcal{V}^\perp}, \quad R_3(\epsilon) = \sqrt{1-\epsilon^2}\mathbf{1},$$

where Π denotes here orthogonal projection.

For ϵ small enough, by Lemma 5.2, the channel is invertible as a positive map. We then define $Q_i = \mathcal{E}_\epsilon^{-1}(P_i)$.

Through the formula $\mathcal{E}_\epsilon(Q_i) = P_i$, we check:

$$\text{If } P_i = \left[\begin{array}{c|c} B & C \\ \hline C^* & D \end{array}\right], \quad \text{then } Q_i(\epsilon) = \left[\begin{array}{c|c} B & (1-\epsilon^2)^{-1}C \\ \hline (1-\epsilon^2)^{-1}C^* & D \end{array}\right]. \tag{5.6}$$

The first remark is that the closure relation ensures $\sum Q_i(\epsilon) = \mathbf{1}$.

We also notice that, since rank-one elements have support either in $\mathcal{V}$ or in $\mathcal{W} = \mathcal{V}^\perp$, the rank-one elements are block-diagonal and $Q_i(\epsilon) = P_i$.

We know that at least one POVM element is not block-diagonal. So that there is an $i \in I$ such that P_i is full-rank and C is non-zero (say $[C]_{j,k} \neq 0$). Then, writing $n = \dim(\mathcal{V})$, there is an $\epsilon_+ \in (0,1)$ such that

$$\begin{aligned} [Q_i(\epsilon_+)]_{j,j}[Q_i(\epsilon_+)]_{n+k,n+k} &= [B]_{j,j}[D]_{k,k} \\ &< \tfrac{1}{1-\epsilon_+^2}|[C]_{j,k}|^2 = [Q_i(\epsilon_+)]_{j,n+k}[Q_i(\epsilon_+)]_{n+k,j} \end{aligned}$$

so that we cannot have positivity of $Q_i(\epsilon_+)$.

We define the bottom of the spectrum of the images Q_i of the full-rank elements of $\mathbf{P}$:

$$\lambda_m(\epsilon) = \inf_{i|P_i \text{ full-rank}} \lambda_m(Q_i(\epsilon)).$$

Equation (5.6) implies that the matrix $Q_i(\epsilon)$ is a continuous function of ϵ for $\epsilon \in [0,1)$. Hence its spectrum is also a continuous function of ϵ. Accordingly, the function $\lambda_m(\epsilon)$ is the minimum of a finite number of continuous function of ϵ, therefore $\lambda_m(\epsilon)$ is continuous. Its value in 0 is the bottom of the spectrum of the full-rank elements of $\mathbf{P}$, that is $\lambda_m(0) = \inf_{i|P_i \text{ full-rank}} \lambda_m(P_i(\epsilon)) > 0$. Moreover we

have just proved that $\lambda_m(\epsilon_+) < 0$. Thus, by the intermediate value Theorem, there is an $\epsilon_+ > \epsilon > 0$ such that $0 < \lambda_m(\epsilon) < \lambda_m(0)$.

As $\lambda_m(\epsilon) > 0$, the $Q_i(\epsilon) = \mathcal{E}_\epsilon(P_i)$ for P_i full-rank are non-negative, and valid POVM elements. Likewise, we already know that $Q_i(\epsilon) = P_i$ is a valid POVM element if P_i is rank-one. Since we have also shown that $\sum Q_i(\epsilon) = \mathbf{1}$, we have proved that $\mathbf{Q}(\epsilon)$ is a POVM. Furthermore $\mathcal{E}_\epsilon(\mathbf{Q}(\epsilon)) = \mathbf{P}$, thus $\mathbf{Q}(\epsilon) \succ \mathbf{P}$.

As $\lambda_m(\epsilon) < \lambda_m(0)$, there is a full-rank element P_i such that $\lambda_m(Q_i(\epsilon)) < \lambda_m(P_i)$. Hence, using Lemma 5.5.2, we get $\mathbf{P} \not\succ \mathbf{Q}(\epsilon)$ and $\mathbf{P}$ is not clean.

Hence $\lambda_m(\epsilon_+) \leq 0 < \lambda_m$. By the intermediate value Theorem, we can find an $\epsilon_0 \in (0, \epsilon_+)$ such that $\lambda_m(\epsilon_0) = 0$. As $0 \leq \lambda_m(\epsilon_0) < \lambda_m$ we have proved that $\mathbf{Q}(\epsilon_0) \succ \mathbf{P}$ and that $\mathbf{P}$ is not clean.

Case (d): As $\mathcal{V}$ and $\mathcal{W}$ are supplementary we may choose a matrix $A \in M_{\dim(\mathcal{V}),d-\dim(\mathcal{V})}(\mathbb{C})$ such that the non-zero columns of the following block matrix form an orthogonal (though not orthonormal) basis of $\mathcal{W}$:

$$R^*_{\mathcal{W}} = \left[\begin{array}{c|c} 0 & A \\ \hline 0 & \mathbf{1} \end{array}\right].$$

We know that the image of a matrix is spanned by its columns, so the image of $R^*_{\mathcal{W}}$ is $\mathcal{W}$.

We then define

$$B(\epsilon) = \sqrt{\mathbf{1} - \left(\frac{\epsilon^4}{1-\epsilon^2} + \frac{\epsilon^2}{(1-\epsilon^2)^2}\right) AA^*}. \tag{5.7}$$

This definition is valid if the matrix under the square root is positive. Now $\left(\frac{\epsilon^4}{1-\epsilon^2} + \frac{\epsilon^2}{(1-\epsilon^2)^2}\right)$ is going to 0 with ϵ, so that

$$\lim_{\epsilon \to 0} \mathbf{1} - \left(\frac{\epsilon^4}{1-\epsilon^2} + \frac{\epsilon^2}{(1-\epsilon^2)^2}\right) AA^* = \mathbf{1}.$$

From this we conclude that $\mathbf{1} - \left(\frac{\epsilon^4}{1-\epsilon^2} + \frac{\epsilon^2}{(1-\epsilon^2)^2}\right) AA^*$ is positive for ϵ small enough.

Accordingly, we can define

$$R^*_{\mathcal{V}}(\epsilon) = \left[\begin{array}{c|c} B(\epsilon) & -\frac{A}{1-\epsilon^2} \\ \hline 0 & 0 \end{array}\right].$$

Notice that the image of $R^*_{\mathcal{V}}$ is included in $\mathcal{V}$.

We may now define our channel $\mathcal{E}_\epsilon$ by

$$R_1^*(\epsilon) = \epsilon R_{\mathcal{V}}^*(\epsilon) = \left[\begin{array}{c|c} \epsilon B(\epsilon) & -\frac{\epsilon}{1-\epsilon^2}A \\ \hline 0 & 0 \end{array}\right] \tag{5.8}$$

$$R_2^*(\epsilon) = \epsilon R_{\mathcal{W}}^* = \left[\begin{array}{c|c} 0 & \epsilon A \\ \hline 0 & \epsilon \mathbf{1} \end{array}\right] \tag{5.9}$$

$$R_3^*(\epsilon) = \sqrt{1-\epsilon^2}\,(R_{\mathcal{V}}^*(\epsilon) + R_{\mathcal{W}}^*) = \left[\begin{array}{c|c} \sqrt{1-\epsilon^2}B(\epsilon) & -\frac{\epsilon^2}{\sqrt{1-\epsilon^2}}A \\ \hline 0 & \sqrt{1-\epsilon^2}\mathbf{1} \end{array}\right]. \tag{5.10}$$

Notice that $\sum_{\alpha=1}^{3} R_\alpha^*(\epsilon)R_\alpha(\epsilon) = \mathbf{1}$ so that $\mathcal{E}(\epsilon)$ is indeed a channel.

Moreover $\lim_{\epsilon\to 0} R_3(\epsilon) = \mathbf{1}_{\mathcal{H}}$. Hence, for ϵ small enough, $\|R_3 - \mathbf{1}\|_{HS}$ is as small as we want. So Lemma 5.5.4 allows us to invert the channel $\mathcal{E}_\epsilon$ as a map on $\mathcal{B}_{sa}(\mathcal{H})$. We define $\mathbf{Q}(\epsilon)$ by its elements $Q_i(\epsilon) = \mathcal{E}_\epsilon^{-1}(P_i)$. Let us check that for ϵ small enough, $\mathbf{Q}(\epsilon)$ is still a *bona fide* POVM.

First the closure relation still holds, as $\sum_{i\in I} Q_i = \sum_{i\in I} \mathcal{E}^{-1}(P_i) = \mathcal{E}^{-1}(\mathbf{1})$. Now $\mathcal{E}(\mathbf{1}) = \sum_\alpha R_\alpha^* R_\alpha = \mathbf{1}$ and taking the inverse $\mathcal{E}^{-1}(\mathbf{1}) = \mathbf{1}$.

Remains then to be shown that all $Q_i(\epsilon)$ are non-negative.

If P_i is full-rank, then its spectrum is included in $[\lambda_m, 1]$, with $\lambda_m > 0$. If R_3 is near enough of the identity, that is, if ϵ is small enough, the inequality (5.3) then ensures that $Q_i(\epsilon)$ is still positive.

If P_i is rank-one $P_i = \lambda_i|\psi_i\rangle\langle\psi_i|$, then by hypothesis $|\psi_i\rangle \in \mathcal{V}$ or $|\psi_i\rangle \in \mathcal{W}$. As R_3 is invertible for ϵ small enough, we may consider $|\phi_i\rangle$ non-zero colinear to $(R_3^*(\epsilon))^{-1}|\psi_i\rangle$. Then $R_3^*(\epsilon)|\phi_i\rangle$ is colinear to $|\psi_i\rangle$, and non-zero. Notice that $|\phi_i\rangle$ depends on ϵ, even if we drop it in the notation. Now

$$R_3(\epsilon)^*|\varphi\rangle = \sqrt{1-\epsilon^2}\quad (R_{\mathcal{V}}^*(\epsilon)|\varphi\rangle + R_{\mathcal{W}}^*|\varphi\rangle)$$
$$\text{with } R_{\mathcal{V}}^*(\epsilon)|\phi\rangle \in \mathcal{V} \text{ and } R_{\mathcal{W}}^*|\varphi\rangle \in \mathcal{W}.$$

Since $\mathcal{V}$ and $\mathcal{W}$ are supplementary, the latter equality implies that $R_{\mathcal{V}}^*(\epsilon)|\varphi\rangle = 0$ when $R_3^*(\epsilon)|\varphi\rangle \in \mathcal{W}$ and $R_{\mathcal{W}}^*(\epsilon)|\varphi\rangle = 0$ when $R_3^*(\epsilon)|\varphi\rangle \in \mathcal{V}$. Definitions (5.8, 5.9, 5.10) then yield $\mathcal{E}_\epsilon(|\phi_i\rangle\langle\phi_i|) = R_{\mathcal{W}}^*(|\phi_i\rangle\langle\phi_i|)R_{\mathcal{W}}$ if $|\psi_i\rangle \in \mathcal{W}$ and $\mathcal{E}_\epsilon(|\phi_i\rangle\langle\phi_i|) = R_{\mathcal{V}}^*(\epsilon)(|\phi_i\rangle\langle\phi_i|)R_{\mathcal{V}}(\epsilon)$ if $|\psi_i\rangle \in \mathcal{V}$. In both cases, the output matrix is of the form $\mathcal{E}_\epsilon(|\phi_i\rangle\langle\phi_i|) = C_i|\psi_i\rangle\langle\psi_i|$. So that $Q_i(\epsilon) = (\lambda_i/C_i)|\phi_i\rangle\langle\phi_i|$ and is non-negative.

Thus, for ϵ small enough, all $Q_i(\epsilon)$ are non-negative. We have proved that $\mathbf{Q}(\epsilon)$ is a POVM. Furthermore, since $\mathcal{E}_\epsilon(\mathbf{Q}(\epsilon)) = \mathbf{P}$, we know $\mathbf{Q}(\epsilon) \succ \mathbf{P}$.

We must still show that $\mathbf{Q}(\epsilon)$ is strictly cleaner $\mathbf{P}$.

By hypothesis, there is a rank-one element $P_i = \lambda_i |\psi_i\rangle\langle\psi_i|$ such that $|\psi_i\rangle \in \mathcal{W}$ and $|\psi_i\rangle \notin \mathcal{V}^\perp$. As above, we write $|\phi_i\rangle$ such that $Q_i(\epsilon) = (\lambda_i/C_i)|\phi_i\rangle\langle\phi_i|$. We start by proving that C_i is less than one.

We write $|\phi_i\rangle = v_i + v_i^\perp$ with $v_i \in \mathcal{V}$ and $v_i^\perp \in \mathcal{V}^\perp$. Since $|\psi_i\rangle \in \mathcal{W}$, we get:

$$\mathcal{E}_\epsilon(|\phi_i\rangle\langle\phi_i|) = R_{\mathcal{W}}^*(|\phi_i\rangle\langle\phi_i|)R_{\mathcal{W}} = \begin{bmatrix} Av_i^\perp \\ \hline v_i^\perp \end{bmatrix} \begin{bmatrix} Av_i^\perp \\ \hline v_i^\perp \end{bmatrix}^* .$$

As the latter expression is also equal to $C_i|\psi_i\rangle\langle\psi_i|$, we obtain that C_i is the square of the norm of $\begin{bmatrix} Av_i^\perp \\ \hline v_i^\perp \end{bmatrix}$. Therefore $C_i = \|Av_i^\perp\|^2 + \|v_i^\perp\|^2$. Notice that the squared norm of $|\phi_i\rangle$ is $1 = \|v_i\|^2 + \|v_i^\perp\|^2$. On the other hand, the image of $|\phi_i\rangle$ by $R_{\mathcal{V}}^*(\epsilon)$ is 0, so that $B(\epsilon)v_i - 1/(1-\epsilon^2)Av_i^\perp = 0$. From this we get:

$$Av_i^\perp = (1-\epsilon^2)B(\epsilon)v_i.$$

Since $|\psi_i\rangle \notin \mathcal{V}^\perp$, this equality shows that $v_i \neq 0$. Now, as AA^* is non-negative we see by (5.7) that $B(\epsilon) \leq \mathbf{1}$. A fortiori, for any $\epsilon > 0$, we have $(1-\epsilon^2)B(\epsilon) < \mathbf{1}$. So that:

$$\|v_i\| > \|(1-\epsilon^2)B(\epsilon)v_i\| = \|Av_i^\perp\|.$$

Thus, we finally obtain

$$C_i = \|Av_i^\perp\|^2 + \|v_i^\perp\|^2 < \|v_i\|^2 + \|v_i^\perp\|^2 = 1.$$

Hence the biggest eigenvalue of $Q_i(\epsilon) = (\lambda_i/C_i)|\phi_i\rangle\langle\phi_i|$, that is λ_i/C_i, is strictly bigger than the biggest eigenvalue of P_i, that is λ_i. Lemma 5.5.2 then gives $\mathbf{P} \not\succ \mathbf{Q}(\epsilon)$, and consequently $\mathbf{P}$ is not clean.

□

5.6 Summary for quasi-qubit POVMs and a special case

We now gather all our results specific to quasi-qubit POVMs.

Theorem 5.6.1. *A quasi-qubit POVM* $\mathbf{P}$ *is clean if and only if it is rank-one or the supports of its rank-one elements totally determine* $\mathcal{H}$*. The algorithm of section 5.3 figures out if this is the case. Moreover if* $\mathbf{Q}$ *is cleanness-equivalent to* $\mathbf{P}$*, the two POVMs are even unitarily equivalent.*

Proof. Rank-one POVMs are known to be clean (Theorem 5.4.11). If the support of the rank-one elements of $\mathbf{P}$ totally determine $\mathcal{H}$, we also know that $\mathbf{P}$ is clean by Theorem 5.4.1. In both cases the theorems state that for these clean POVMs, cleanness-equivalence is the same as unitary equivalence.

Conversely, if $\mathbf{P}$ is neither rank-one nor have rank-one elements that totally determine $\mathcal{H}$, then Theorem 5.5.1 applies and $\mathbf{P}$ is not clean.

Stage (i) of the algorithm checks whether $\mathbf{P}$ is rank-one, in which case it does say that $\mathbf{P}$ is clean. If $\mathbf{P}$ is not rank-one, the fact that it is clean or not depends on the support of its rank-one elements. The only remaining positive exit of the algorithm is at stage (vii) and Lemma 5.4.9 proves that in this case the rank-one elements of $\mathbf{P}$ totally determine $\mathcal{H}$.

Conversely, if the algorithm exits with a negative value, Lemma 5.5.3 ensures that $\mathcal{H}$ is not totally determined.

□

To get further feeling of these conditions we finish by making more explicit the qubit case, where the nice thing is that all POVMs are quasi-qubit.

Corollary 5.6.2. *A POVM $\mathbf{P}$ for a qubit is clean if and only if it is rank-one or if one can find three rank-one elements whose supports are two-by-two non-colinear (that is if they make a projective frame). For these POVMs cleanness-equivalence is the same as unitary equivalence.*

Proof. A POVM $\mathbf{P}$ for a qubit has non-zero elements which can be either of rank one, or of rank two, as $d = 2$. In the latter case, they are full-rank, so we may apply Theorem 5.6.1 to $\mathbf{P}$.

The only question is when do the supports of the rank-one elements totally determine $\mathcal{H}$? They do by Proposition 5.4.4 if they include a projective frame, that is a basis and a vector with all coefficients non-zero in this basis. As the space is of dimension 2, this amounts to saying a basis and a vector non-colinear to any basis vector, that is three vectors two-by-two non-colinear.

Conversely, if we cannot find a projective frame, then we can find two vectors v and w such that the support of any rank-one element is v or w, and we can apply Lemma 5.5.3 to obtain that $\mathcal{H}$ is not totally determined by the supports of the rank-one elements of $\mathbf{P}$. Thus $\mathbf{P}$ is not clean.

□

5.7 Outlook

We have solved the problem of cleanness for quasi-qubit POVMs. The obvious continuation would be to solve it in the general case. However we do not think that the condition of Theorem 5.4.1 is then necessary. Moreover it must be made explicit.

The heuristics in Section 5.3.2 suggest that, if the support of P_i are in "general position" then it is sufficient for $\mathbf{P}$ to be clean that $\sum_{i\in I} d - \dim[\mathrm{Supp}(P_i)] \geq d^2 - 1$. Yet, we still need to appropriately define the "general position" for general subspaces.

Chapitre 6

Complementary subalgebras

Ce chapitre dérive de l'article (Kahn et Petz, 2007).

Résumé : La réduction d'un système quantique à un sous-système donne une information partielle sur l'état du système total. En lien avec la détermination optimale de l'état de deux qubits, la question a été posée de savoir quel était le nombre maximum de réductions complémentaires deux à deux. Le principal résultat de ce chapitre est de montrer que ce nombre est de 4, c'est-à-dire que si $\mathcal{A}^1, \mathcal{A}^2, \ldots, \mathcal{A}^k$ sont des sous-algèbres deux à deux complémentaires de $M_4(\mathbb{C})$, et qu'elles sont isomorphes à $M_2(\mathbb{C})$, alors $k \leq 4$. La preuve est basée sur la décomposition de Cartan de $SU(4)$. Au passage, nous apportons quelques contributions sur la structure des réductions complémentaires.

6.1 Introduction

There is an obvious correspondence between bases of an m-dimensional Hilbert space $\mathcal{H}$ and maximal Abelian subalgebras of the algebra $\mathcal{A} \equiv B(\mathcal{H}) \simeq M_m(\mathbb{C})$. Given a basis, the linear operators diagonal in this basis form a maximal Abelian (or commutative) subalgebra. Conversely if $|e_i\rangle\langle e_i|$ are minimal projections in a maximal Abelian subalgebra, then $(|e_i\rangle)_i$ is a basis. From the points of view of quantum mechanics, a basis can be regarded as a measurement. Wootters et Fields (1989) argued that two measurements corresponding to the bases $\xi_1, \xi_2, \ldots, \xi_m$ and $\eta_1, \eta_2, \ldots, \eta_m$ yield the largest amount of information about the true state of the system in the average if

$$|\langle \xi_i, \eta_j \rangle|^2 = \frac{1}{m} \qquad (1 \leq i, j \leq m).$$

Two bases satisfying this condition are called **mutually unbiased**. Mutually unbiased bases are interesting from many point of view, for example in quantum information theory, tomography and cryptography (Kraus, 1987; Bandyopadhyay et al., 2002; Kimura et al., 2006). The maximal number of such bases is not known for arbitrary m. Nevertheless, $(m^2-1)/(m-1) = m+1$ is a bound being checked easily (Parthasarathy, 2004; Pittenger et Rubin, 2004).

The concept of mutually unbiased (or complementary) maximal Abelian subalgebras can be extended to more general subalgebras. In particular, a 4-level quantum system can be regarded as the composite system of two qubits, $M_4(\mathbb{C}) \simeq M_2(\mathbb{C}) \otimes M_2(\mathbb{C})$. A density matrix $\rho \in M_4(\mathbb{C})$ describes a state of the composite system and ρ determines the "marginal" or reduced states on both tensor factors. Since the decomposition $M_2(\mathbb{C}) \otimes M_2(\mathbb{C})$ is not unique, there are many reductions to different subalgebras, they provide partial quantum information about the composite system. It seems that the reductions provide the largest amount of information if the corresponding subalgebras are quasi-orthogonal or complementary in a different terminology. In (Petz et al., 2006) the state ρ was to be determined by its reductions. 4 pairwise complementary subalgebras were given explicitly, but the question remained open to know if 5 such subalgebras exist. The main result of this paper is to prove that at most 4 pairwise complementary subalgebras exist.

6.2 Preliminaries

In this paper an algebraic approach and language is used. A k-level quantum system is described by operators of the algebra $M_k(\mathbb{C})$ of $k \times k$ matrices. Although the essential part of the paper focuses on a 4-level quantum system, certain concepts can be presented slightly more generally. Let $\mathcal{A}$ be an algebra corresponding to a quantum system. The normalized trace τ gives the Hilbert-Schmidt inner product $\langle A, B\rangle := \tau(B^*A)$ on $\mathcal{A}$ and we can speak about orthogonality with respect to this inner product.

The projections in $\mathcal{A}$ may be defined by the algebraic properties $P = P^2 = P^*$ and the partial ordering $P \leq Q$ means $PQ = QP = P$. We consider subalgebras of $\mathcal{A}$ such that their minimal projections have the same trace. (A maximal Abelian subalgebra and a subalgebra isomorphic to a full matrix algebra have this property.) Let $\mathcal{A}^1$ and $\mathcal{A}^2$ be two such subalgebras of $\mathcal{A}$. Then the following conditions are equivalent:

(i) If $P \in \mathcal{A}^1$ and $Q \in \mathcal{A}^2$ are minimal projections, then $\operatorname{Tr} PQ = \operatorname{Tr} P \operatorname{Tr} Q$.

(ii) The traceless subspaces of $\mathcal{A}^1$ and $\mathcal{A}^2$ are orthogonal with respect to the Hilbert-Schmidt inner product on $\mathcal{A}$.

The subalgebras $\mathcal{A}^1$ and $\mathcal{A}^2$ are called **complementary** (or quasi-orthogonal) if these conditions hold. This terminology was used in the maximal Abelian case (Accardi, 1984; Kraus, 1987; Ohya et Petz, D., 2004; Parthasarathy, 2004) and the case of noncommutative subalgebras appeared in (Petz et al., 2006). More details about complementarity are presented in (Petz, 2006).

Given a density matrix $\rho \in \mathcal{A}$, its reduction $\rho_1 \in \mathcal{A}_1$ to the subalgebra $\mathcal{A}_1 \subset \mathcal{A}$ is determined by the formula

$$\operatorname{Tr} \rho A = \operatorname{Tr} \rho_1 A \qquad (A \in \mathcal{A}_1).$$

In most cases ρ_1 is given by the partial trace but an equivalent way is based on the conditional expectation (P. Busch et Mittelstaedt, 1991). The orthogonal projection $E : \mathcal{A} \to \mathcal{A}_1$ is called conditional expectation. $\rho_1 = E(\rho)$ and

$$E(AB) = AE(B) \qquad (A \in \mathcal{A}_1, B \in \mathcal{A})$$

is an important property.

The situation we are interested in is the algebra $M_4(\mathbb{C})$. In the paper $M_4(\mathbb{C})$ is regarded as a Hilbert space with respect to the inner product

$$\langle A, B \rangle = \frac{1}{4} \operatorname{Tr} A^* B = \tau(A^* B). \tag{6.1}$$

$M_4(\mathbb{C})$ has a natural orthonormal basis:

$$\sigma_i \otimes \sigma_j \qquad (0 \leq i, j \leq 3),$$

where $\sigma_1, \sigma_2, \sigma_3$ are the Pauli matrices and σ_0 is the identity I:

$$\sigma_0 := \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \sigma_1 := \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad \sigma_2 := \begin{bmatrix} 0 & -\mathrm{i} \\ \mathrm{i} & 0 \end{bmatrix}, \quad \sigma_3 := \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}.$$

6.3 Complementary subalgebras

Any subalgebra $\mathcal{A}^1$ of $M_4(\mathbb{C})$ isomorphic to $M_2(\mathbb{C})$ can be written $\mathbb{C}I \otimes M_2(\mathbb{C})$ in some basis, hence there is a unitary operator W such that $\mathcal{A}^1 = W(\mathbb{C}I \otimes M_2(\mathbb{C}))W^*$.

This section is organized as follows: we first give a characterization of the W such that $\mathcal{A}^1$ is complementary to $\mathcal{A}^0 = W(\mathbb{C}I \otimes M_2(\mathbb{C}))W^*$ (Theorem 6.3.1 for a general form and Theorem 6.3.2 for a form specific to our problem). The second stage consists in proving, using the form of W, that any such $\mathcal{A}^1$ has "a large component" along $\mathcal{B} = M_2(\mathbb{C}) \otimes \mathbb{C}I$. Theorem 6.3.4 gives the precise formulation. It entails that no more than four complementary subalgebras con be found (Theorem 6.3.5), which was our initial aim, and hence is our conclusion.

Although our main interest is $M_4(\mathbb{C})$, our first theorem is more general. E_{ij} stand for the matrix units.

Theorem 6.3.1. *Let $W = \sum_{i,j=1}^n E_{ij} \otimes W_{ij} \in M_n(\mathbb{C}) \otimes M_n(\mathbb{C})$ be a unitary. The subalgebra $W(\mathbb{C}I \otimes M_n(\mathbb{C}))W^*$ is complementary to $\mathbb{C}I \otimes M_n(\mathbb{C})$ if and only if $\{W_{ij} : 1 \leq i, j \leq n\}$ is an orthonormal basis in $M_n(\mathbb{C})$ (with respect to the inner product $\langle A, B\rangle = \operatorname{Tr} A^*B$).*

Proof. Assume that $\operatorname{Tr} B = 0$. Then the condition

$$W(I \otimes A^*)W^* \perp (I \otimes B)$$

is equivalently written as

$$\operatorname{Tr} W(I \otimes A)W^*(I \otimes B) = \sum_{i,j=1}^n \operatorname{Tr} W_{ij}AW_{ij}^*B = 0.$$

This implies

$$\sum_{i,j=1}^n \operatorname{Tr} W_{ij}AW_{ij}^*B = (\operatorname{Tr} A)(\operatorname{Tr} B)\,. \tag{6.2}$$

We can transform this into another equivalent condition in terms of the left multiplication and right multiplication operators. For $A, B \in M_n(\mathbb{C})$, the operator R_A is the right multiplication by A and L_B is the left multiplication by B: $R_A, L_B : M_n(\mathbb{C}) \to M_n(\mathbb{C})$, $R_BX = XB$, $L_AX = AX$. Equivalently, $L_A|e\rangle\langle f| = |Ae\rangle\langle f|$ and $R_B|e\rangle\langle f| = |e\rangle\langle B^*f|$. From the latter definition one can deduce that $\operatorname{Tr} R_AL_B = \operatorname{Tr} A \operatorname{Tr} B$. Let $|e_i\rangle$ be a basis. Then $|e_i\rangle\langle e_j|$ form a basis in $M_n(\mathbb{C})$ and

$$\begin{aligned}\operatorname{Tr} R_AL_B &= \sum_{ij} \langle |e_i\rangle\langle e_j|, R_AL_B|e_i\rangle\langle e_j|\rangle = \sum_{ij} \langle |e_i\rangle\langle e_j|, |Be_i\rangle\langle A^*e_j|\rangle \\ &= \sum_{ij} \langle e_i, Be_i\rangle\langle e_j, Ae_j\rangle.\end{aligned}$$

The equivalent form of (6.2) is the equation

$$\sum_{i,j=1}^n \langle W_{ij}, R_AL_BW_{ij}\rangle = \operatorname{Tr} A \operatorname{Tr} B = \operatorname{Tr} R_AL_B$$

for every $A, B \in M_n(\mathbb{C})$. Since the operators $R_A L_B$ linearly span the space of all linear operators on $M_n(\mathbb{C})$, we can conclude that W_{ij} form an orthonormal basis. □

We shall call any unitary satisfying the condition in the previous theorem a useful unitary and we shall denote the set of all $n^2 \times n^2$ useful unitaries by $\mathrm{i}(n^2)$.

We try to find a useful 4×4 unitary W, that is we require that the subalgebra

$$W \begin{bmatrix} A & 0 \\ 0 & A \end{bmatrix} W^* \qquad (A \in M_2(\mathbb{C}))$$

is complementary to $\mathcal{A}^0 \equiv \mathbb{C}I \otimes M_2(\mathbb{C})$. We shall use the **Cartan decomposition** of W given by

$$W = (L_1 \otimes L_2) N (L_3 \otimes L_4)\,,$$

where L_1, L_2, L_3 and L_4 are 2×2 unitaries and

$$N = \exp(\alpha \mathrm{i}\, \sigma_1 \otimes \sigma_1) \exp(\beta \mathrm{i}\, \sigma_2 \otimes \sigma_2) \exp(\gamma \mathrm{i}\, \sigma_3 \otimes \sigma_3) \tag{6.3}$$

is a 4×4 unitary in a special form, see equation (11) in (Zhang et al., 2003) or (D'Alessandro et Albertini, 2005). The subalgebra

$$W(\mathbb{C}I \otimes M_2(\mathbb{C}))W^* = (L_1 \otimes L_2) N (\mathbb{C}I \otimes M_2(\mathbb{C})) N^* (L_1^* \otimes L_2^*)$$

does not depend on L_3 and L_4, therefore we may assume that $L_3 = L_4 = I$.

The orthogonality of $\mathbb{C}I \otimes M_2(\mathbb{C})$ and $W(\mathbb{C}I \otimes M_2(\mathbb{C}))W^*$ does not depend on L_1 and L_2. Therefore, the equations

$$\operatorname{Tr} N(I \otimes \sigma_i) N^* (I \otimes \sigma_j) = 0$$

should be satisfied, $1 \leq i, j \leq 3$. We know from Theorem 6.3.1 that these conditions are equivalent to the property that the matrix elements of N form a basis.

A simple computation gives that

$$N = \sum_{i=0}^{3} c_i\, \sigma_i \otimes \sigma_i\,,$$

where

$$\begin{aligned}
c_0 &= \cos\alpha \cos\beta \cos\gamma + \mathrm{i} \sin\alpha \sin\beta \sin\gamma\,, \\
c_1 &= \cos\alpha \sin\beta \sin\gamma + \mathrm{i} \sin\alpha \cos\beta \cos\gamma\,, \\
c_2 &= \sin\alpha \cos\beta \sin\gamma + \mathrm{i} \cos\alpha \sin\beta \cos\gamma\,, \\
c_3 &= \sin\alpha \sin\beta \cos\gamma + \mathrm{i} \cos\alpha \cos\beta \sin\gamma\,.
\end{aligned}$$

Therefore, we have

$$N = \begin{bmatrix} c_0 + c_3 & 0 & 0 & c_1 - c_2 \\ 0 & c_0 - c_3 & c_1 + c_2 & 0 \\ 0 & c_1 + c_2 & c_0 - c_3 & 0 \\ c_1 - c_2 & 0 & 0 & c_0 + c_3 \end{bmatrix}$$

$$= \begin{bmatrix} e^{i\gamma}\cos(\alpha-\beta) & 0 & 0 & ie^{i\gamma}\sin(\alpha-\beta) \\ 0 & e^{-i\gamma}\cos(\alpha+\beta) & ie^{-i\gamma}\sin(\alpha+\beta) & 0 \\ 0 & ie^{-i\gamma}\sin(\alpha+\beta) & e^{-i\gamma}\cos(\alpha+\beta) & 0 \\ ie^{i\gamma}\sin(\alpha-\beta) & 0 & 0 & e^{i\gamma}\cos(\alpha-\beta) \end{bmatrix}. \quad (6.4)$$

Since the 2×2 blocks form a basis (see Theorem 6.3.1), we have

$$\overline{(c_0 + c_3)}(c_0 - c_3) + \overline{(c_0 - c_3)}(c_0 + c_3) = 0\,,$$

$$\overline{(c_1 - c_2)}(c_1 + c_2) + \overline{(c_1 + c_2)}(c_1 - c_2) = 0\,,$$

$$|c_0 + c_3|^2 + |c_0 - c_3|^2 = 1\,,$$

$$|c_1 + c_2|^2 + |c_1 - c_2|^2 = 1\,.$$

These equations give

$$|c_0|^2 = |c_1|^2 = |c_2|^2 = |c_3|^2 = \frac{1}{4}$$

and we arrive at the following solution. Two of the values of $\cos^2\alpha, \cos^2\beta$ and $\cos^2\gamma$ equal $1/2$ and the third one may be arbitrary. Let $\mathcal{N}$ be the set of all matrices such that the parameters α, β and γ satisfy the above condition, in other words two of the three values are of the form $\pi/4 + k\pi/2$. (k is an integer.)

The conclusion of the above argument can be formulated as follows.

Theorem 6.3.2. *$W \in \mathcal{M}(4)$ if and only if $W = (L_1 \otimes L_2)N(L_3 \otimes L_4)$, where L_i are 2×2 unitaries ($1 \leq i \leq 4$) and $N \in \mathcal{N}$.*

We now turn to the "second stage", that is proving that any such $W(\mathbb{C}I \otimes M_2(\mathbb{C})$ is far from being complementary to $M_2(\mathbb{C}) \otimes \mathbb{C}I$. To get a quantitative result (Theorem 6.3.4), recall that we consider $M_4(\mathbb{C})$ as a Hilbert space with Hilbert-Schmidt inner product (see (6.1)). For the proof of Theorem 6.3.4, we shall need the following obvious lemma:

Lemma 6.3.3. *Let $\mathcal{K}_1$ and $\mathcal{K}_2$ be subspaces of a Hilbert space $\mathcal{K}$ and denote by $\mathbf{P}_i : \mathcal{K} \to \mathcal{K}_i$ the orthogonal projection onto $\mathcal{K}_i$ ($i = 1, 2$). If $\xi_1, \xi_2, \ldots, \xi_r$ is an orthonormal basis in $\mathcal{K}_1$ and $\eta_1, \eta_2, \ldots, \eta_s$ is such a basis in $\mathcal{K}_2$, then*

$$\mathrm{Tr}\,\mathbf{P}_1\mathbf{P}_2 = \sum_{i,j} |\langle \xi_i, \eta_j \rangle|^2.$$

□

Theorem 6.3.4. *Let $\mathcal{A}^0 \equiv \mathbb{C}I \otimes M_2(\mathbb{C})$ and $\mathcal{B} \equiv M_2(\mathbb{C}) \otimes \mathbb{C}I$. Assume that the subalgebra $\mathcal{A}^1 \subset M_2(\mathbb{C}) \otimes M_2(\mathbb{C})$ is isomorphic to $M_2(\mathbb{C})$ and complementary to $\mathcal{A}^0$. If $\mathbf{P}$ is the orthogonal projection onto the traceless subspace of $\mathcal{A}^1$ and $\mathbf{Q}$ is the orthogonal projection onto the traceless subspace of $\mathcal{B}$, then*

$$\mathrm{Tr}\,\mathbf{PQ} \geq 1.$$

Proof. There is a unitary $W = (L_1 \otimes L_2)N$ such that $\mathcal{A}^1 = W\mathcal{A}^0W^*$, L_1, L_2 are 2×2 unitaries and $N \in \mathcal{M}(4)$. In the traceless subspace of $\mathcal{B}$,

$$(L_1\sigma_iL_1^*) \otimes I \qquad (1 \leq i \leq 3)$$

form a basis, while

$$(L_1 \otimes L_2)N(I \otimes \sigma_i)N^*(L_1^* \otimes L_2^*) \qquad (1 \leq i \leq 3)$$

is a basis in the traceless part of $\mathcal{A}^1$. Therefore, we have to show

$$\sum_{ij} \left|\langle (L_1 \otimes L_2)N(I \otimes \sigma_i)N^*(L_1^* \otimes L_2^*), L_1^*\sigma_jL_1 \otimes I\rangle\right|^2 = \Big(\tau(N(I \otimes \sigma_i)N^*(\sigma_j \otimes I))\Big)^2 \geq 1.$$

In the computation we can use the conditional expectation $E : M_4(\mathbb{C}) \to \mathcal{B}$. Recall that it is defined as the linear operator which sends $\sigma_i \otimes \sigma_j$ to $\sigma_i \otimes I$, for all $0 \leq i, j \leq 3$.

Two of its main properties are that it preserves τ, and that $E(AB) = E(A)B$ when $B \in \mathcal{B}$. Hence

$$\tau\Big(N(I \otimes \sigma_i)N^*(\sigma_j \otimes I)\Big) = \tau\left(E\Big(N(I \otimes \sigma_i)N^*\Big)(\sigma_j \otimes I)\right).$$

Elementary computation in the basis $\sigma_i \otimes \sigma_j$ gives the following formulas:

$$\begin{aligned} E(N(I \otimes \sigma_1)N^*) &= \sin 2\beta \sin 2\gamma\,(\sigma_1 \otimes I), \\ E(N(I \otimes \sigma_2)N^*) &= \sin 2\alpha \sin 2\gamma\,(\sigma_2 \otimes I), \\ E(N(I \otimes \sigma_3)N^*) &= \sin 2\alpha \sin 2\beta\,(\sigma_2 \otimes I), \end{aligned}$$

where α, β and γ are from (6.3) and (6.4). Therefore,

$$\mathrm{Tr}\,\mathbf{PQ} = \sin^2 2\beta \sin^2 2\gamma + \sin^2 2\alpha \sin^2 2\gamma + \sin^2 2\alpha \sin^2 2\beta.$$

Recall that two of the parameters α, β and γ have rather concrete values, hence one of the three terms equals 1, and the proof is complete. □

Our main results says that there are at most four pairwise complementary subalgebras of $M_4(\mathbb{C})$ if they are assumed to be isomorphic to $M_2(\mathbb{C})$. Given such a family of subalgebras, we may assume that the above defined $\mathcal{A}^0$ belongs to the family.

Theorem 6.3.5. *Assume that* $\mathcal{A}^0 \equiv \mathbb{C}I \otimes M_2(\mathbb{C})$, $\mathcal{A}^1, \ldots, \mathcal{A}^r$ *are pairwise complementary subalgebras of* $M_4(\mathbb{C})$ *and they are isomorphic to* $M_2(\mathbb{C})$. *Then* $r \leq 3$.

Proof. Let $\mathbf{P}_i$ be the orthogonal projection onto the traceless subspace of $\mathcal{A}^i$ from $M_4(\mathbb{C})$, $1 \leq i \leq r$. Under these conditions $\sum_i \mathbf{P}_i \leq I$. As in Theorem 6.3.4, let $\mathbf{Q}$ the orthogonal projection on the traceless subspace of $\mathcal{B} \equiv M_2(\mathbb{C}) \otimes \mathbb{C}I$. The estimate

$$3 = \operatorname{Tr}\mathbf{Q} \geq \operatorname{Tr}(\mathbf{P}_1 + \mathbf{P}_2 + \cdots + \mathbf{P}_r)\mathbf{Q} = \sum_{i=1}^{r} \operatorname{Tr}\mathbf{P}_i\mathbf{Q} \geq r$$

yields the proof. □

Deuxième partie

Normalité Asymptotique Locale Quantique

Chapitre 7

Quantum local asymptotic normality for qubits

Ce chapitre dérive de l'article (Guţă et Kahn, 2006).

Résumé : Nous considérons n qubits identiquement préparés et étudions les propriétés asymptotiques de l'état joint $\rho^{\otimes n}$. Nous montrons que pour chaque état ρ situé dans un voisinage de rayon $1/\sqrt{n}$ autour d'un état fixé ρ^0, l'état joint converge vers un état thermique déplacé d'un oscillateur harmonique. La signification précise de cette convergence est l'existence de transformations physiques T_n (canaux présrvant la trace) qui envoient les états des qubits près des états correspondants de l'oscillateur, uniformément sur tous les états d'un voisinage.

Nous dérivons quelques conséquences de ce résultat. Nous montrons que la mesure optimale dans le cadre bayésien est également optimale dans une approche minimax. De plus, cette mesure converge la mesure hétérodyne qui est la mesure jointe optimale de la position et de l'impulsion d'un oscillateur harmonique quantique. Le problème de la discrimination locale est aussi résolu à travers la normalité asymptotique locale.

7.1 Introduction

Quantum measurement theory brings together the quantum world of wave functions and incompatible observables with the classical world of random phenomena studied in probability and statistics. These fields have come ever closer due to the

technological advances making it possible to perform measurements on individual quantum systems. Indeed, the engineering of a novel quantum state is typically accompanied by a verification procedure through which the state, or some aspect of it, is reconstructed from measurement data (Schiller et al., 1996).

An important example of such a technique is that of quantum homodyne tomography in quantum optics (Vogel et Risken, H., 1989). This allows the estimation with arbitrary precision of the whole density matrix (D'Ariano et al., 1995; Leonhardt et al., 1995, 1996; Artiles et al., 2005) of a monochromatic beam of light by repeatedly measuring a sufficiently large number of identically prepared beams (Smithey et al., 1993; Schiller et al., 1996; Zavatta et al., 2004).

In contrast to this "semi-classical" situation in which one fixed measurement is performed repeatedly on independent systems, the state estimation problem becomes more "quantum" if one is allowed to consider *joint measurements* on n identically prepared systems with joint state $\rho^{\otimes n}$. It is known (Gill et Massar, 2000) that in the case of unknown *mixed* states ρ, joint measurements perform strictly better than separate measurements in the sense that the asymptotic convergence rate of the optimal estimator $\hat{\rho}_n$ to ρ goes in both case as $C/\sqrt{n}$ with a strictly smaller constant C in the case of joint measurements.

Let us look at this problem in more detail: we dispose of a number of n copies of an unknown state ρ and the task is to estimate ρ as well as possible. The first step is to specify a cost function $d(\hat{\rho}_n, \rho)$ which quantifies the deviation of the estimator $\hat{\rho}_n$ from the true state. Then one tries to devise a measurement and an estimator which minimizes the mean cost or risk in statistics jargon:

$$R(\rho, \hat{\rho}_n) := \langle d(\hat{\rho}_n(X), \rho) \rangle ,$$

with the average taken over the measurement results X. Since this quantity still depends on the unknown state one may choose a Bayesian approach and try to optimize the average risk with respect to some prior distribution π over the states

$$R_{n,\pi} = \int R(\rho, \hat{\rho}_n) \pi(d\rho).$$

Results of this type have been obtained in both the pure state case (Jones, 1994; Massar et Popescu, 1995; Latorre et al., 1998; Fisher et al., 2000; Hannemann et al., 2002b; Bagan et al., 2002; Embacher et Narnhofer, 2004; Bagan et al., 2005) and the mixed state case (Cirac et al., 1999; Vidal et al., 1999; Mack et al., 2000; Keyl et Werner, 2001; Bagan et al., 2004c; Zyczkowski et Sommers, 2005; Bagan et al., 2006). However most of these papers use methods of group theory that depend on the symmetry of the prior distribution and the form of the cost function, and thus cannot be extended to arbitrary priors.

In the pointwise approach (Hayashi, 2002a; Gill et Massar, 2000; Barndorff-Nielsen et Gill, R., 2000; Matsumoto, 2002; Barndorff-Nielsen et al., 2003; Hayashi et Matsumoto, 2004) one tries to minimize $R(\rho, \hat{\rho}_n)$ for each fixed ρ. We can argue that even for a completely unknown state, as n becomes large the problem ceases to be global and becomes a local one as the error in estimating the state parameters is of the order $\frac{1}{\sqrt{n}}$. For this reason it makes sense to parametrize the state as $\rho := \rho(\theta)$ with θ belonging to some set in $\mathbb{R}^k$ and to replace the original cost with its quadratic approximation at θ:

$$d(\theta, \hat{\theta}_n) = (\theta - \hat{\theta}_n)^T G(\theta)(\theta - \hat{\theta}_n),$$

where G is a $k \times k$ positive, real symmetric weight matrix.

Although seemingly different, the two approaches can be compared (Gill, 2005a), and in fact for large n the prior distribution π of the Bayesian approach should become increasingly irrelevant and the optimal Bayesian estimator should be close to the maximum likelihood estimator. An instance of this asymptotic equivalence is proven in Subsection 7.7.2.

In this chapter we change the perspective and instead of trying to devise optimal measurements and estimators for a particular statistical problem, we concentrate our attention on the *family* of joint states $\rho(\theta)^{\otimes n}$ which is the primary "carrier" of statistical information about θ. As suggested by the locality argument sketched above, we consider a neighborhood of size $\frac{1}{\sqrt{n}}$ around a fixed but arbitrary parameter θ_0, whose points can be written as $\theta = \theta_0 + \mathbf{u}/\sqrt{n}$ with $\mathbf{u} \in \mathbb{R}^k$ the "local parameter" obtained by zooming into the smaller and smaller balls by a factor of $\sqrt{n}$. Very shortly, the principle of *local asymptotic normality* says that for large n the local family

$$\rho_n^{\mathbf{u}} := \rho\left(\theta_0 + \mathbf{u}/\sqrt{n}\right)^{\otimes n}, \qquad \|\mathbf{u}\| < C,$$

converges to a family of displaced Gaussian states $\phi^{\mathbf{u}}$ of a of a quantum system consisting of a number of coupled quantum and classical harmonic oscillators.

The term local asymptotic normality comes from mathematical statistics (van der Vaart, 1998) where the following result holds. We are given independent variables $X_1, \ldots, X_n \in \mathcal{X}$ drawn from the same probability distribution $P^{\theta_0 + \mathbf{u}/\sqrt{n}}$ over $\mathcal{X}$ depending smoothly on the unknown parameter $\mathbf{u} \in \mathbb{R}^k$. Then the statistical information contained in our data is asymptotically identical with the information contained in a *single* normally distributed $Y \in \mathbb{R}^k$ with mean $\mathbf{u}$ and variance $I(\theta_0)^{-1}$, the inverse Fisher information matrix. This means that for any statistical problem we can replace the original data $X_1, \ldots, X_n \in \mathcal{X}$ by the simpler Gaussian one Y with the same asymptotic results!

For the sake of clarity let us consider the case of qubits with states parametrized by their Bloch vectors $\rho(\vec{r}) = \frac{1}{2}(\mathbf{1} + \vec{r}\vec{\sigma})$ where $\vec{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$ are the Pauli

matrices. Define now the two-dimensional family of identical spin states obtained by rotating the Bloch vector $\overrightarrow{r_0} = (0, 0, 2\mu - 1)$ around an axis in the x-y plane

$$\rho_n^{\mathbf{u}} = \left[U\left(\frac{\mathbf{u}}{\sqrt{n}}\right) \begin{pmatrix} \mu & 0 \\ 0 & 1-\mu \end{pmatrix} U\left(\frac{\mathbf{u}}{\sqrt{n}}\right)^* \right]^{\otimes n}, \quad \mathbf{u} \in \mathbb{R}^2, \tag{7.1}$$

with unitary $U(\mathbf{v}) := \exp(i(v_x\sigma_x + v_y\sigma_y))$ and $\frac{1}{2} < \mu \leq 1$.

Consider now a quantum harmonic oscillator with position and momentum operators Q and P on $L^2(\mathbb{R})$ satisfying the commutation relations $[Q, P] = i\mathbf{1}$. We denote by $\{|n\rangle, n \geq 0\}$ the eigenbasis of the number operator and define the thermal equilibrium state

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

where $p = \frac{1-\mu}{\mu}$. We translate the state $\phi^{\mathbf{0}}$ by using the displacement operators $D(\mathbf{z}) = \exp(\mathbf{z}a^* - \bar{\mathbf{z}}a)$ with $\mathbf{z} \in \mathbb{C}$ which map the ground state $|0\rangle$ into the coherent state $|\mathbf{z}\rangle$:

$$\phi^{\mathbf{u}} := D(\sqrt{2\mu - 1}\alpha_{\mathbf{u}})\phi^{\mathbf{0}}D(\sqrt{2\mu - 1}\alpha_{\mathbf{u}})^*, \tag{7.2}$$

where $\alpha_{\mathbf{u}} := -u_y + iu_x$.

Theorem 7.1.1. *Let $\rho_n^{\mathbf{u}}$ be the family of states (7.1) on the Hilbert space $(\mathbb{C}^2)^{\otimes n}$ and $\phi^{\mathbf{u}}$ the family (7.2) of displaced thermal equilibrium states of a quantum oscillator. Then for each n there exist quantum channels (trace preserving CP maps)*

$$\begin{aligned} T_n &: M\left((\mathbb{C}^2)^{\otimes n}\right) \to \mathcal{T}(L^2(\mathbb{R})), \\ S_n &: \mathcal{T}(L^2(\mathbb{R})) \to M\left((\mathbb{C}^2)^{\otimes n}\right), \end{aligned} \tag{7.3}$$

with $\mathcal{T}(L^2(\mathbb{R}))$ the trace-class operators, such that

$$\begin{aligned} \lim_{n\to\infty} \sup_{\mathbf{u}\in I^2} \|\phi^{\mathbf{u}} - T_n(\rho_n^{\mathbf{u}})\|_1 &= 0, \\ \lim_{n\to\infty} \sup_{\mathbf{u}\in I^2} \|\rho_n^{\mathbf{u}} - S_n(\phi^{\mathbf{u}})\|_1 &= 0. \end{aligned} \tag{7.4}$$

for an arbitrary bounded interval $I \subset \mathbb{R}$.

Let us make a few comments on the significance of the above result.

i) The "convergence" (7.4) of the qubit states holds in a strong way (uniformly in $\mathbf{u}$) with direct statistical and physical interpretation. Indeed the channels T_n and S_n represent physical transformations which are analogues of randomizations of classical data (van der Vaart, 1998). The meaning of (7.4) is that the two quantum models are asymptotically equivalent from a statistical point of view.

ii) Indeed for any measurement M on $L^2(\mathbb{R})$ we can construct the measurement $M \circ T_n$ on the spin states by first mapping them to the oscillator space and then performing M. Then the optimal solution of any statistical problem concerning the states $\rho_n^{\mathbf{u}}$ can be obtained by solving the same problem for $\phi^{\mathbf{u}}$ and pulling back the optimal measurement M as above. We illustrate this in Section 7.7 for the estimation problem and for hypothesis testing.

iii) The proposed technique may be useful for applications in the domain of coherent spin states (Holtz et Hanus, 1974) and squeezed spin states (Kitagawa et Ueda, 1993). Indeed, it has been known since Dyson (1956) that n spin-$\frac{1}{2}$ particles prepared in the spin up state $|\uparrow\rangle^{\otimes n}$ behave asymptotically as the ground state of a quantum oscillator when considering the fluctuations of properly normalized total spin components in the directions orthogonal to z. Our Theorem extends this to spin directions making an "angle" $\mathbf{u}/\sqrt{n}$ with the z axis, as well as to mixed states, and gives a quantitative expression to heuristic pictures common in the physics literature (see Section 7.3). We believe that a similar approach can be followed in the case of spin squeezed states and continuous time measurements with feedback control (Geremia et al., 2004).

Next Section gives an introduction to the statistical ideas motivating our work. In Section 7.3 we give a heuristic picture of our main result based on the total spin vector representation of spin coherent states familiar in the physics literature.

The proof of Theorem 7.1.1 extends over the Sections 7.4,7.5,7.6 and uses methods of group theory and some ideas from (Hayashi et Matsumoto, 2004; Ohya et Petz, D., 2004; Accardi et Bach, A., 1987, 1985).

Section 7.7 describes a few applications of our main result. In Subsection 7.7.2 we compute the local asymptotic minimax risk for the statistical problem of qubit state estimation. An estimation scheme which achieves this risk asymptotically is optimal in the pointwise approach. We show that this figure of merit coincides with the risk of the heterodyne measurement and that it is achieved by the optimal Bayesian measurement for the $SU(2)$-invariant prior (Bagan et al., 2006; Hayashi et Matsumoto, 2004). This proves the asymptotic equivalence of the Bayesian and pointwise approaches.

In Subsection 7.7.1 we continue the investigation of the optimal Bayesian measurement and show that it converges locally to the heterodyne measurement on the oscillator, which is an optimal joint measurement of position and momentum (Holevo, 1982).

Another application is the problem discriminating between two states $\rho_n^{\pm\mathbf{u}}$ which asymptotically converge to each other at rate $1/\sqrt{n}$. In this case the optimal mea-

surement for the parameter $\mathbf{u}$ is not optimal for the testing problem, showing in particular that the quantum Fisher information in general does not encode all statistical information.

7.2 Local asymptotic normality in statistics and its extension to quantum mechanics

In this Section we introduce some statistical ideas which provide the motivation for deriving the main result.

Quantum statistical problems can be seen as a game between a statistician or physicist in our case, and Nature. The latter tries to codify some information by preparing a quantum system in a state which depends on some parameter $\mathbf{u}$ unknown to the former. The physicist tries to guess the value of the parameter by devising measurements and estimators which work well for *all* choices of parameters that Nature may make. In a Bayesian set-up Nature may build her strategy by randomly choosing a state with some prior distribution. In order to solve the problem the physicist is allowed to use the laws of quantum physics as well as those of classical stochastics and statistical inference. In particular he may transform the quantum state by applying an arbitrary quantum channel T and obtain a new family $T(\rho^{\mathbf{u}})$. In general such transformation goes with a loss of information so one should have a good reason to do it but there are non trivial situations when no such loss occurs (Petz et Jenčová, 2006), that is when there exists a channel S which reverses the effect of T *restricted* to the states of interest $S(T(\rho^{\mathbf{u}})) = \rho^{\mathbf{u}}$. If this is the case the we consider the two families of states $\rho^{\mathbf{u}}$ and $T(\rho^{\mathbf{u}})$ as statistically equivalent.

In statistics such transformations are called *randomizations* and a useful particular example is a *statistic*, which is just a function of the data which we want to analyze. When this statistic contains all information about the unknown parameter we say that it is sufficient, because knowing the value of this statistic alone suffices and given this information, the rest of the data is useless. For example if $X_1, \dots X_n \in \{0, 1\}$ are results of independent coin tosses with a biased coin, then $\bar{X} = \frac{1}{n}\sum_i X_i$ is sufficient statistic and may be used for any statistical decision without loss of efficiency.

Quantum randomizations through quantum channels allows us to compare seemingly different families of states and thus opens the possibility of solving a particular problem by casting it in a more familiar setting. The example of this chapter is that of state estimation for n identical copies of a state which can be cast *asymptotically* into the problem of estimating the center of a quantum Gaussian which has a rather simple solution (Holevo, 1982). The term "asymptotically" means that for large n

we can find quantum channels T_n, S_n which almost map the families of states into each other as in equation (7.4).

The second main idea that we want to introduce is that of local asymptotic normality. Back in the coin toss example we have that $\bar{X}$ is a good estimator of the probability μ of obtaining a 1 and by the Central Limit Theorem the error $\bar{X}-\mu$ has asymptotically a Gaussian distribution

$$\sqrt{n}(\bar{X}-\mu) \rightsquigarrow N(0, 1/\mu(1-\mu)),$$

in particular the mean error is $\langle(\bar{X}-\mu)^2\rangle = 1/(n\mu(1-\mu))$. Now, if for each n the unknown parameter μ is restricted to a local neighborhood of a fixed μ_0 of size $1/\sqrt{n}$, one might expect an improvement in the error because we know more about the parameter and we can use that information to built better estimators. However this is not entirely true. Indeed if we write $\mu = \mu_0 + u/\sqrt{n}$ then the estimator of the local parameter u is

$$\hat{u}_n = \sqrt{n}(\bar{X}-\mu_0) \rightsquigarrow N(u, 1/\mu_0(1-\mu_0))$$

which says that the problem of estimating μ in the local parameter model is as difficult as the original problem, i.e. the variance of the estimator is the same. The reason for this is that the additional information about the location of the parameter is nothing new as we could guess that directly form the data with very high probability. Thus without changing the difficulty of the original problem we can look at it locally and then we see that it transforms into that of estimating the center of a Gaussian with fixed variance $N(u, 1/\mu_0(1-\mu_0))$, which is a classical statistical problem.

In general we can formulate the following principle: given $X_1, \ldots, X_n \in \mathcal{X}$ independent with distribution $P^{\theta_0+\mathbf{u}/\sqrt{n}}$ depending smoothly on the unknown parameter $\mathbf{u} \in \mathbb{R}^k$, then asymptotically this model is statistically equivalent (there exist explicit randomizations in both directions) with that of a single draw $Y \in \mathbb{R}^k$ from the Gaussian distribution $N(\mathbf{u}, I(\theta_0)^{-1})$ with fixed variance equal to the inverse of the Fisher information matrix (van der Vaart, 1998).

In the quantum case we replace the randomizations by quantum channels and the Gaussian limit model by its quantum equivalent which in the simplest case is a family of displaced thermal states of a quantum oscillator (see Theorem 7.1.1), but in general is a Gaussian state on a number of coupled quantum and classical oscillators, with canonical variables satisfying general commutation relations (Petz, 1990).

A simple extension of Theorem 7.1.1 is obtained by adding an additional local parameter $t \in \mathbb{R}$ for the density matrix eigenvalues such that $\mu = \mu_0 + t/\sqrt{n}$. This leads

to a Gaussian limit model in which we are given a quantum oscillator is in state $\phi^{\mathbf{u}}$ and additionally, a classical Gaussian variable with distribution $N(t, 1/\mu_0(1-\mu_0))$. The meaning of this quantum-classical coupling is the following: asymptotically the problem of estimating the eigenvalues decouples from that of estimating the direction of the Bloch vector and becomes a *classical* statistical problem (identical with the coin toss discussed above), while that of estimating the direction remains quantum and converges to the estimation of a Gaussian state of a quantum oscillator. Bagan et al. (2006); Hayashi et Matsumoto (2004) have also observed this decoupling.

7.3 The big ball picture of coherent spin states

In this section we give a heuristic argument for why Theorem 7.1.1 holds which will guide our intuition in later computations.

It is customary to represent the state of two dimensional quantum system by a vector $\overrightarrow{r}$ in the Bloch sphere such that the corresponding density matrix is

$$\rho = \frac{1}{2}(\mathbf{1} + \overrightarrow{r}\,\overrightarrow{\sigma}) = \frac{1}{2}(\mathbf{1} + r_x\sigma_x + r_y\sigma_y + r_z\sigma_z),$$

where σ_i represent the Pauli matrices and satisfy the commutation relations $[\sigma_i, \sigma_j] = 2i\epsilon_{ijk}\sigma_k$. In particular if $\overrightarrow{r} = (0, 0, \pm 1)$ then the state is given by the spin up $|\uparrow\rangle$ and respectively spin down $|\downarrow\rangle$ basis vectors of $\mathbb{C}^2$, and the z-component of the spin σ_z takes value ± 1. As for the x and y spin components, each one may take the values ± 1 with equal probabilities such that on average $\langle\sigma_x\rangle = \langle\sigma_y\rangle = 0$ but the variances are $\langle\sigma_x^2\rangle = \langle\sigma_y^2\rangle = 1$. Moreover σ_x and σ_y do not commute and thus cannot be measured simultaneously.

What happens with the Bloch sphere picture when we have more spins? Consider for the beginning n identical spins prepared in a coherent spin up state $|\uparrow\rangle^{\otimes n}$, then we can think of the whole as a single spin system and define the global observables $L_i^{(n)} = \sum_{k=1}^n \sigma_i^{(k)}$ for $i \in x, y, z$, where $\sigma_i^{(k)}$ is the spin component in the direction i of the k's spin. Intuitively, we can represent the joint state by a vector of length n pointing to the north pole of a large sphere as in Figure 7.1. However due to the quantum character of the spin observables, the x and y components cannot be equal to zero and it is more instructive to think in terms of a vector whose tip lies on a small blob of the size of the uncertainties in x and y, sitting on the top of the sphere. Exactly how large is this blob? By using the Central Limit Theorem we conclude that in the limit $n \to \infty$ the distribution of the "fluctuation operator"

$$S_x^{(n)} := \frac{1}{\sqrt{2n}} L_x^{(n)} = \frac{1}{\sqrt{2n}} \sum_{k=1}^{n} \sigma_x^{(k)},$$

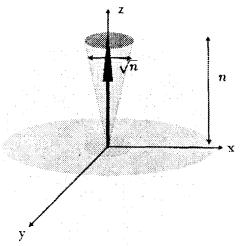

Figure 7.1: (Color online) Quasiclassical representation of n spin up qubits

converges to a $N(0,1/2)$ Gaussian, that is $\langle S_x \rangle = 0$ and $\langle S_x^2 \rangle \approx 1/2$, and similarly for the component $S_y^{(n)}$. The width of the blob is thus of the order $\sqrt{n}$ in both x and y directions.

Now, the two fluctuations do not commute with each other

$$[S_x^{(n)}, S_y^{(n)}] = \frac{i}{n} L_z^{(n)} \approx i\mathbf{1}, \tag{7.5}$$

which is the well know commutation relation for canonical variables of the quantum oscillator. In fact the quantum extension of the Central Limit Theorem (Ohya et Petz, D., 2004) makes this more precise

$$\lim_{n\to\infty} {}^{\otimes n}\langle \uparrow | \prod_{k=1}^{p} S_{i_k}^{(n)} | \uparrow \rangle^{\otimes n} = \langle \Omega, \prod_{k=1}^{p} X_{i_k} \Omega \rangle, \quad \forall i_k \in \{x, y\},$$

where $X_x := Q$ and $X_y := P$ satisfy $[Q, P] = i\mathbf{1}$ and Ω is the ground state of the oscillator.

The above description is not new in physics and goes back to Dyson's (1956) theory of spin-wave interaction. More recently squeezed spin states (Kitagawa et Ueda, 1993) for which the variances $\langle S_x^2 \rangle$ and $\langle S_y^2 \rangle$ of spin variables are different have been found to have important applications various fields such as magnetometry (Geremia et al., 2004), entanglement between many particles (Stockton et al., 2003) The connection with such applications will be discussed in more detail in Section 7.7.

We now rotate all spins by the same small angle for each particle as in Figure 7.2. As we will see, it makes sense to scale the angle by the factor $\frac{1}{\sqrt{n}}$ i.e. to consider

$$\psi_n^{\mathbf{u}} = \left[\exp\left(\frac{i}{\sqrt{n}} (u_x \sigma_x + u_y \sigma_y) \right) | \uparrow \rangle \right]^{\otimes n}, \qquad \mathbf{u} \in \mathbb{R}^2.$$

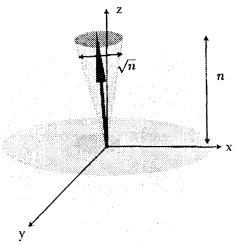

Figure 7.2: (Color online) Rotated coherent state of n qubits

Indeed for such angles the z component of the vector will change by a small quantity of the order $\sqrt{n} \ll n$ so the commutation relations (7.5) remain the same, while the uncertainty blob will just shift its center such that the new averages of the renormalized spin components are $\langle S_x^{(n)} \rangle \approx -\sqrt{2}u_y$ and $\langle S_y^{(n)} \rangle \approx \sqrt{2}u_x$. All in all, the spins state converges to the coherent state $|\alpha_{\mathbf{u}}\rangle$ of the oscillator where $\alpha_{\mathbf{u}} = (-u_y + iu_x) \in \mathbb{C}$ and in general

$$|\alpha\rangle := \exp\left(-|\alpha|^2/2\right) \sum_{j=0}^{\infty} \frac{\alpha^{\mathbf{j}}}{\sqrt{j!}} |j\rangle,$$

with $|j\rangle$ representing the j's energy level.

We consider now the case of qubits in individual mixed state $\mu|\uparrow\rangle\langle\uparrow| + (1-\mu)|\downarrow\rangle\langle\downarrow|$ with $< 1/2\mu < 1$. Then the "length" of L_z is $n(2\mu - 1)$ but the size of the blob is the same (see Figure 7.3). However the commutation relations of S_x and S_y do not

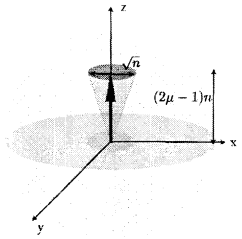

Figure 7.3: (Color online) Quasiclassical representation of n qubit mixed states

reproduce those of the harmonic oscillator and we need to renormalize the spin as

$$S_x^{(n)} := \frac{1}{\sqrt{2(2\mu-1)n}} L_x, \quad S_y^{(n)} := \frac{1}{\sqrt{2(2\mu-1)n}} L_y.$$

The limit state will be a Gaussian state of the quantum oscillator with variance $\langle Q^2\rangle = \langle P^2\rangle = \frac{1}{2(2\mu-1)} < \frac{1}{2}$, that is a thermal equilibrium state

$$\phi^{\mathbf{0}} = (1-p)\sum_{k=0}^{\infty} p^k |k\rangle\langle k|, \qquad p = \frac{1-\mu}{\mu}.$$

Finally the rotation by $\exp\left(\frac{i}{\sqrt{n}}(u_x\sigma_x + u_y\sigma_y)\right)$ produces a displacement of the thermal state such that $\langle Q\rangle = -\sqrt{2}(2\mu-1)u_y$ and $\langle P\rangle = \sqrt{2}(2\mu-1)u_x$.

7.4 Local asymptotic normality for mixed qubit states

We give now a rigorous formulation of the heuristics presented in the previous Section. Let

$$\rho^{\mathbf{0}} = \begin{pmatrix} \mu & 0 \\ 0 & 1-\mu \end{pmatrix} \tag{7.6}$$

be a density matrix on $\mathbb{C}^2$ with $\mu > 1/2$, representing a mixture of spin up and spin down states, and for every $\mathbf{u} = (u_x, u_y) \in \mathbb{R}^2$ consider the state

$$\rho^{\mathbf{u}} = U(\mathbf{u})\,\rho^{\mathbf{0}}\,U(\mathbf{u})^*$$

where

$$U(\mathbf{u}) := \exp(i(u_x\sigma_x + u_y\sigma_y)) = \begin{pmatrix} \cos|\mathbf{u}| & -e^{-i\varphi}\sin|\mathbf{u}| \\ e^{i\varphi}\sin|\mathbf{u}| & \cos|\mathbf{u}| \end{pmatrix},$$

with $\varphi = \mathrm{Arg}(-u_y + iu_x)$. We are interested in the asymptotic behavior as $n \to \infty$ of the family

$$\mathcal{F}_n := \left\{\rho_n^{\mathbf{u}} = \left(\rho^{\mathbf{u}/\sqrt{n}}\right)^{\otimes n}, \mathbf{u} \in I^2\right\}, \tag{7.7}$$

where $I = [-a, a]$ is a fixed finite interval.

The main result is that $\mathcal{F}_n$ is asymptotically normal, meaning that it converges as $n \to \infty$ to a limit family $\mathcal{G}_n := \{\phi^{\mathbf{u}}, \mathbf{u} \in I^2\}$ of Gaussian states of a quantum oscillator with creation and annihilation operators satisfying $[a, a^*] = \mathbf{1}$. Let

$$\phi^{\mathbf{0}} := (1-p)\sum_{k=0} p^k |k\rangle\langle k|, \tag{7.8}$$

be a thermal equilibrium state with $|k\rangle$ denoting the k's energy level of the oscillator and $p = \frac{1-\mu}{\mu} < 1$. For every $\mathbf{u} \in I^2$ define

$$\phi^{\mathbf{u}} := D(\sqrt{2\mu-1}\alpha_{\mathbf{u}})[\phi^{\mathbf{0}}]D(-\sqrt{2\mu-1}\alpha_{\mathbf{u}}), \tag{7.9}$$

where $D(\mathbf{z}) := \exp(\mathbf{z}a^* - \mathbf{z}^*a)$ is the displacement operator, mapping the vacuum vector $|\mathbf{0}\rangle$ to the coherent vector $|\mathbf{z}\rangle$ and $\alpha_{\mathbf{u}} = (-u_y + iu_x)$.

The exact formulation of the convergence is given in Theorem 7.1.1. Thus the state $\rho_n^{\mathbf{u}}$ of the n qubits which depends on the unknown parameter $\mathbf{u}$ can be manipulated by applying a quantum channel T_n such that its image converges to the Gaussian state $\phi^{\mathbf{u}}$, uniformly in $\mathbf{u} \in I^2$. Conversely by using the channel S_n, the state $\phi^{\mathbf{u}}$ can be mapped to a joint state of n qubits which is converges to $\rho_n^{\mathbf{u}}$ uniformly in $\mathbf{u} \in I^2$. By Stinespring's theorem we know that the channels are of the form

$$T(\rho) = \mathrm{Tr}_{\mathcal{K}} (V\rho V^*),$$
$$S(\phi) = \mathrm{Tr}_{\mathcal{K}'} (W\phi W^*),$$

where the partial traces are taken over some ancillary Hilbert spaces $\mathcal{K}, \mathcal{K}'$ and

$$V : \left(\mathbb{C}^2\right)^{\otimes n} \to L^2(\mathbb{R}) \otimes \mathcal{K},$$
$$W : L^2(\mathbb{R}) \to \left(\mathbb{C}^2\right)^{\otimes n} \otimes \mathcal{K}',$$

are isometries ($V^*V = \mathbf{1}$ and $W^*W = \mathbf{1}$).

Our task is now to identify the isometries V_n and W_n implementing the channels T_n and respectively S_n satisfying (7.4). The first step towards identifying these V_n is to use group representations methods so as to partially (block) diagonalize all the $\rho_n^{\mathbf{u}}$ simultaneously.

7.4.1 Block decomposition

In this Subsection we show that the states $\rho_n^{\mathbf{u}}$ have a block-diagonal form given by the decomposition of the space $(\mathbb{C}^2)^{\otimes n}$ into irreducible representations of the relevant symmetry groups. The main point is that for large n the weights of the different blocks concentrate around the representation with total spin $j_n = n(\mu - 1/2)$.

The space $(\mathbb{C}^2)^{\otimes n}$ carries a unitary representation π_n of the one spin symmetry group $SU(2)$ with $\pi_n(u) = u^{\otimes n}$ for any $u \in SU(2)$, and a unitary representation of the symmetric group $S(n)$ given by the permutation of factors

$$\pi_n(\tau) : v_1 \otimes \cdots \otimes v_n \mapsto v_{\tau^{-1}(1)} \otimes \cdots \otimes v_{\tau^{-1}(n)}, \qquad \tau \in S(n).$$

As $[\pi_n(u), \pi_n(\tau)] = 0$ for all $u \in SU(2), \tau \in S(n)$ we have the decomposition

$$\left(\mathbb{C}^2\right)^{\otimes n} = \bigoplus_{j=0,1/2}^{n/2} \mathcal{H}_j \otimes \mathcal{H}_n^j, \tag{7.10}$$

where the direct sum runs over all positive (half)-integers j up to $n/2$, and for each fixed j, $\mathcal{H}_j \cong \mathbb{C}^{2j+1}$ is a irreducible representation of $SU(2)$ with total angular momentum $J^2 = j(j+1)$, and $\mathcal{H}_n^j \cong \mathbb{C}^{n_j}$ is the irreducible representation of the symmetric group $S(n)$ with $n_j = \binom{n}{n/2-j} - \binom{n}{n/2-j-1}$. In particular the density matrix $\rho_n^{\mathbf{u}}$ is invariant under permutations and can be decomposed as a mixture of "block" density matrices

$$\rho_n^{\mathbf{u}} = \bigoplus_{j=0,1/2}^{n/2} p_n(j) \rho_{j,n}^{\mathbf{u}} \otimes \frac{\mathbf{1}}{n_j}, \tag{7.11}$$

with probability distribution $p_n(j)$ given by (Bagan et al., 2006):

$$p_n(j) := \frac{n_j}{2\mu - 1} (1-\mu)^{\frac{n}{2}-j} \mu^{\frac{n}{2}+j+1} \left(1 - p^{2j+1}\right), \tag{7.12}$$

where $p := \frac{1-\mu}{\mu}$. A key observation is that for large n and in the relevant range of j's, $p_n(j)$ is essentially a binomial distribution

$$B_{n,\mu}(k) := \binom{n}{k} \mu^k (1-\mu)^{n-k}, \qquad k = 0, \ldots, n.$$

Indeed we can rewrite $p_n(j)$ as

$$p_n(j) := B_{n,\mu}(n/2 + j) \times K(j, n, \mu) \tag{7.13}$$

where the factor $K(j, n, \mu)$ is given by

$$K(j, n, \mu) := \left(1 - p^{2j+1}\right) \frac{n + (2(j - j_n) + 1)/(2\mu - 1)}{n + (j - j_n + 1)/\mu}$$

and $j_n := n(\mu - 1/2)$. As $B_{n,\mu}$ is the distribution of the sum of n independent Bernoulli variables with individual distribution $(1-\mu, \mu)$ over $\{0, 1\}$, we can use the central limit Theorem to conclude that its mass concentrates around the average μn with a width of order $\sqrt{n}$, in other words of any $0 < \epsilon < 1/2$ we have

$$\lim_{n\to\infty} \sum_{p=-n^{1/2+\epsilon}}^{n^{1/2+\epsilon}} B_{n,\mu}(\mu n + p) = 1. \tag{7.14}$$

Let us denote by $\mathcal{J}_{n,\epsilon}$ the set of values j of the total angular momentum of n qubits which lie in the interval $[j_n - n^{1/2+\epsilon}, j_n + n^{1/2+\epsilon}]$. Then for large n, the factor $K(j, n, \mu)$ is close to 1 uniformly over $j \in \mathcal{J}_{n,\epsilon}$ and from formulas (7.13), (7.14) we conclude that $p_n(j)$ concentrates asymptotically in an interval of order $n^{1/2+\epsilon}$ around j_n:

$$\lim_{n\to\infty} p_n(\mathcal{J}_{n,\epsilon}) = 1. \tag{7.15}$$

This justifies the big ball picture used in the previous section.

7.4.2 Irreducible representations of $SU(2)$

Here we remind the reader some details about the representation π_j of $SU(2)$ on $\mathcal{H}_j$. Let $\sigma_x, \sigma_y, \sigma_z$ be the Pauli matrices and denote $\pi_j(\sigma_l) = J_{j,l}$ for $l = x, y, z$ then there exists an orthonormal basis $\{|j,m\rangle, m = -j, \ldots, j\}$ of $\mathcal{H}_j$ such that

$$J_{j,z}|j,m\rangle = m|j,m\rangle.$$

Moreover, with $J_{j,\pm} := J_{j,x} \pm iJ_{j,y}$ we have

$$\begin{aligned} J_{j,+}|j,m\rangle &= \sqrt{j-m}\sqrt{j+m+1}\,|j,m+1\rangle, \\ J_{j,-}|j,m\rangle &= \sqrt{j-m+1}\sqrt{j+m}\,|j,m-1\rangle. \end{aligned}$$

With these notations and $p = \frac{1-\mu}{\mu}$ as before, the state $\rho^{\mathbf{0}}_{j,n}$ can be written as (Hayashi et Matsumoto, 2004)

$$\rho^{\mathbf{0}}_{j,n} = c_j(p) \sum_{m=-j}^{j} p^{j-m} |j,m\rangle\langle j,m|,$$

where the normalizing factor is $c_j(p) = (1-p)/(1-p^{2j+1})$. The rotated block states can be obtained by applying the unitary transformation

$$\rho^{\mathbf{u}}_{j,n} = U_j(\mathbf{u}/\sqrt{n})\, \rho^{\mathbf{0}}_{j,n}\, U_j(\mathbf{u}/\sqrt{n})^*,$$

with $U_j(\mathbf{u}) = \exp\left(i(u_x J_{j,x} + u_y J_{j,y})\right)$. Finally, we define the vectors

$$|j,\mathbf{w}\rangle := U_j(\mathbf{w})|j,j\rangle \tag{7.16}$$

which will be used in later computations, and notice that their coordinates with respect to the $|j,m\rangle$ basis are given by (Hayashi et Matsumoto, 2004):

$$\langle j,m|j,\mathbf{w}\rangle = \sqrt{\binom{2j}{j+m}}\,\zeta^{j-m}(1-|\zeta|^2)^{\frac{j+m}{2}}. \tag{7.17}$$

where $\zeta = e^{i\varphi_w} \sin|\mathbf{w}|$ with $\varphi_w = \mathrm{Arg}(-w_y + iw_x)$.

7.5 Construction of the channels T_n

For each irreducible representation space $\mathcal{H}_j$ we define the isometry $V_j : \mathcal{H}_j \to L^2(\mathbb{R})$ by

$$V_j : |j,m\rangle \mapsto |j-m\rangle \tag{7.18}$$

where $\{|n\rangle, n \geq 0\}$ represents the energy eigenbasis of the quantum oscillator with eigenfunctions $\psi_n(x) = H_n(x)e^{-x^2/2}/\sqrt{\sqrt{\pi}2^n n!} \in L^2(\mathbb{R})$. Using the decomposition (7.10) we put together the different blocks we construct for each $n \in \mathbb{N}$ the "global" isometry

$$V_n := \bigoplus_{j=0,1/2}^{n/2} V_j \otimes \mathbf{1} : \bigoplus_{j=0,1/2}^{n/2} \mathcal{H}_j \otimes \mathbb{C}^{n_j} \to L^2(\mathbb{R}) \otimes \mathcal{K}_n,$$

where $\mathcal{K}_n := \bigoplus_{j=0,1/2}^{n/2} \mathbb{C}^{n_j}$. By tracing over $\mathcal{K}_n$ we obtain the channel $T_n(\rho) := \mathrm{Tr}_{\mathcal{K}_n}(V_n \rho V_n^*)$ mapping a joint state of n spins into a state of the quantum oscillator. This channel satisfies the convergence condition (7.4) as shown by the estimate

$$\begin{aligned}
\|T_n(\rho_n^{\mathbf{u}}) - \phi^{\mathbf{u}}\|_1 &= \left\| \sum_{j=0,1/2}^{n/2} p_n(j) V_j \rho_{n,j}^{\mathbf{u}} V_j^* - \phi^{\mathbf{u}} \right\|_1 \\
&\leq \sum_{j=0,1/2}^{n/2} p_n(j) \left\| V_j \rho_{n,j}^{\mathbf{u}} V_j^* - \phi^{\mathbf{u}} \right\|_1 \\
&\leq 2 \sum_{j \notin \mathcal{J}_{n,\epsilon}} p_n(j) + \sup_{\mathbf{u} \in I^2} \max_{j \in \mathcal{J}_{n,\epsilon}} \|V_j \rho_{j,n}^{\mathbf{u}} V_j^* - \phi^{\mathbf{u}}\|_1,
\end{aligned}$$

where the first term on the right side converges to 0 by (7.15), and for the second one we apply the following Proposition 7.5.1 which is the major technical contribution of this chapter.

Proposition 7.5.1. *The following uniform convergence holds*

$$\lim_{n \to \infty} \sup_{\mathbf{u} \in I^2} \max_{j \in \mathcal{J}_{n,\epsilon}} \|V_j \rho_{j,n}^{\mathbf{u}} V_j^* - \phi^{\mathbf{u}}\|_1 = 0.$$

where $\mathcal{J}_{n,\epsilon}$ is the set defined above equation (7.15).

The proof of the Proposition requires a few ingredients which in our opinion are important on their own for which reason we formulate them apart and refer to relevant papers for the proofs.

Theorem 7.5.2. *(Ohya et Petz, D., 2004) Let $a, b \in M(\mathbb{C}^d)$, satisfying $\mathrm{Tr}(a) = \mathrm{Tr}(b) = 0$ and define*

$$L(a,b) = \exp(ia)\exp(ib) - \exp(ia+ib)\exp\left(\frac{1}{2}[a,b]\right).$$

On $(\mathbb{C}^2)^{\otimes n}$ we define the fluctuation operator

$$F_n(a) = \frac{1}{\sqrt{n}} \sum a_i,$$

where $a_i = 1 \otimes \cdots \otimes a \otimes \cdots \otimes 1$ *with* a *acting on the* i*'s position of the tensor product. Notice that* $exp(iF_n(a)) = exp(ia/\sqrt{n})^{\otimes n}$ *and* $\sqrt{n}[F_n(a), F_n(b)] = F_n([a,b])$*. Then*

$$\lim_{n\to\infty} \| L\left(F_n(a), F_n(b)\right) \| = 0.$$

The convergence is uniform over $\|a\|, \|b\| < C$ *for some constant* C*.*

This Theorem is a key ingredient of the quantum central limit Theorem (Ohya et Petz, D., 2004) and it is not surprising that it plays an important role in our quantum local asymptotic normality result which is an extension of the latter. We apply the Theorem to two unitaries of the form $U(\mathbf{u}) = \exp(i(u_x\sigma_x + u_y\sigma_y))$. We thus get information on the effect of the $U_j(\mathbf{u})$ on the highest weight vectors $|j,j\rangle$ of an irreducible representation.

Corollary 7.5.3. *For any unitary* U *and state* τ *let* $\mathrm{Ad}[U](\tau) := U\tau U^*$ *and consider the rotated states*

$$\begin{aligned} \tau(\mathbf{u},\mathbf{v},j,n) &:= \mathrm{Ad}\left[U_j\left(\frac{\mathbf{u}}{\sqrt{n}}\right) U_j\left(\frac{\mathbf{v}}{\sqrt{n}}\right)\right](|jj\rangle\langle jj|) \\ \tau(\mathbf{u}+\mathbf{v},j,n) &:= \mathrm{Ad}\left[U_j\left(\frac{\mathbf{u}+\mathbf{v}}{\sqrt{n}}\right)\right](|jj\rangle\langle jj|)\,. \end{aligned}$$

Then the following uniform convergence holds

$$\lim_{n\to\infty} \sup_{\mathbf{u},\mathbf{v}\in I^2} \sup_{j\in\mathcal{J}_{n,\epsilon}} \|\tau(\mathbf{u},\mathbf{v},j,n) - \tau(\mathbf{u}+\mathbf{v},j,n)\|_1 = 0.$$

Proof. First notice that

$$[u_x\sigma_x + u_y\sigma_y, v_x\sigma_x + v_y\sigma_y] = 2(u_xv_y - u_yv_x)\sigma_z.$$

Applying Theorem 7.5.2 to $U(\mathbf{u})$, we get

$$\left\| U\left(\frac{\mathbf{u}}{\sqrt{n}}\right)^{\otimes n} U\left(\frac{\mathbf{v}}{\sqrt{n}}\right)^{\otimes n} - U\left(\frac{\mathbf{u}+\mathbf{v}}{\sqrt{n}}\right)^{\otimes n} \exp\left(\frac{u_xv_y - u_yv_x}{\sqrt{n}} F_n(\sigma_z)\right) \right\| \xrightarrow[n\to\infty]{} 0.$$

Now

□

The following Lemma is a slight strengthening of a theorem by Hayashi et Matsumoto (2004).

Lemma 7.5.4. *The uniform convergence holds*

$$\lim_{n\to\infty} \sup_{\mathbf{u}\in I^2} \sup_{j\in\mathcal{J}_{n,\epsilon}} \left\| V_j U_j \left(\frac{\mathbf{u}}{\sqrt{n}}\right) |jj\rangle - |\sqrt{2\mu-1}\alpha_{\mathbf{u}}\rangle \right\| = 0,$$

where $|\mathbf{z}\rangle$ *denotes a coherent state of the oscillator, and* $\alpha_{\mathbf{u}} := (-u_y + iu_x)$. *Moreover for any sequence* $j_n \to \infty$ *we have*

$$\lim_{n\to\infty} \left\| V_{j_n} \rho^{\mathbf{0}}_{j_n} V^*_{j_n} - \phi^{\mathbf{0}} \right\|_1 = 0. \tag{7.19}$$

The convergence holds uniformly over all sequences j_n *such that* $j_n/n > c$ *for some fixed constant* $c > 0$, *so in particular for* $j_n \in \mathcal{J}_{n,\epsilon}$.

Proof. We first prove the easier relation (7.19). As both density matrices are diagonal we get

$$\left\| V_{j_n} \rho^{\mathbf{0}}_{j_n} V^*_{j_n} - \phi^{\mathbf{0}} \right\|_1 = \frac{(1-p)p^{2j_n+1}}{1-p^{2j_n+1}} \sum_{k=0}^{2j_n} p^k -$$

$$(1-p) \sum_{k=2j_n+1}^{\infty} p^k \le \frac{p^{2j_n+1}}{1-p^{2j_n+1}} + p^{2j_n+1} \to 0,$$

as $n \to \infty$.

As for the first relation, let us denote $|\mathbf{u}, j, n\rangle := V_j U_j(\frac{\mathbf{u}}{\sqrt{n}})|j, j\rangle$, then by (7.17) and (7.18) we have

$$\langle k|\mathbf{u}, j, n\rangle = \sqrt{\binom{2j}{k}} (\sin(|\mathbf{u}|/\sqrt{n})e^{i\phi})^k (\cos(|\mathbf{u}|/\sqrt{n}))^{2j-k}.$$

Now, the following asymptotic relations hold uniformly over $j \in \mathcal{J}_{n,\epsilon}$:

$$\sin\left(\frac{|\mathbf{u}|}{\sqrt{n}}\right)^k = \left(\frac{|\mathbf{u}|}{\sqrt{n}}\right)^k \left(1 + O(|\mathbf{u}|^2 n^{-1})\right),$$

$$\cos\left(\frac{|\mathbf{u}|}{\sqrt{n}}\right)^{2j-k} = \exp(-\frac{(2\mu-1)|\mathbf{u}|^2}{2}) \left(1 + O(|\mathbf{u}|^2 n^{-\epsilon})\right),$$

$$\binom{2j}{k} = \frac{((2\mu-1)n)^k}{k!}(1 + O(n^{-\epsilon})),$$

and thus the coefficients converge uniformly to those of the corresponding coherent states as $n \to \infty$

$$\langle k|\mathbf{u}, j, n\rangle \to \exp\left(-\frac{(2\mu-1)|\mathbf{u}|^2}{2}\right) \frac{\left(e^{i\phi}|\mathbf{u}|\sqrt{2\mu-1}\right)^k}{\sqrt{k!}}.$$

□

Proof of Proposition 7.5.1. The main idea is to notice that ϕ^0 is a thermal equilibrium state of the oscillator and can be generated as a mixture of coherent states with centered Gaussian distribution over the displacements:

$$\phi^{\mathbf{0}} = \frac{1}{\sqrt{2\pi s^2}} \int e^{-|\mathbf{z}|^2/2s^2} |\mathbf{z}\rangle\langle\mathbf{z}| \, d^2\mathbf{z}. \tag{7.20}$$

The easiest way to see this is to think of the oscillator states in terms of their Wigner functions. Indeed, the Wigner function of a coherent state is

$$W_{\mathbf{z}}(q,p) = \exp\left(-(q - \sqrt{2}\mathrm{Re}\,\mathbf{z})^2 - (p - \sqrt{2}\mathrm{Im}\,\mathbf{z})^2\right),$$

and thus the state given by (7.20) has Wigner function which is the convolution of two centered Gaussians which is again a centered Gaussian with variance equal to the sum of their variances $2s^2 + 1/2$ which is equal to the variance of $\phi^{\mathbf{0}}$ for $s^2 := p/(2(1-p))$. Similarly,

$$\phi^{\mathbf{u}} = \frac{1}{2\pi s^2} \int e^{-|\mathbf{z}-\sqrt{2\mu-1}\alpha_{\mathbf{u}}|^2/2s^2} \left(|\mathbf{z}\rangle\langle\mathbf{z}|\right) d^2\mathbf{z}. \tag{7.21}$$

Let us first remark that

$$\begin{aligned}\left\|V_{j_n}\rho^{\mathbf{u}}_{j_n}V^*_{j_n} - \phi^{\mathbf{u}}\right\|_1 \leq \quad & \left\|\rho^{\mathbf{u}}_{j_n} - V^*_{j_n}\phi^{\mathbf{u}}V_{j_n}\right\|_1 + \\ & \left\|\phi^{\mathbf{u}} - P_{j_n}\phi^{\mathbf{u}}P_{j_n}\right\|_1,\end{aligned}$$

where $P_{j_n} = V_{j_n}V^*_{j_n}$ is the projection onto the image of V_{j_n}, and

$$\lim_{n\to\infty} \sup_{j_n\in\mathcal{J}_{n,\epsilon}} \sup_{\mathbf{u}\in I^2} \left\|\phi^{\mathbf{u}} - P_{j_n}\phi^{\mathbf{u}}P_{j_n}\right\|_1 = 0,$$

because $j_n \to \infty$ uniformly and P_{j_n} converges to the identity in strong operator topology (a tightness property). Thus it is enough to show that

$$\lim_{n\to\infty} \sup_{j_n\in\mathcal{J}_{n,\epsilon}} \sup_{\mathbf{u}\in I^2} \left\|\rho^{\mathbf{u}}_{j_n} - V^*_{j_n}\phi^{\mathbf{u}}V_{j_n}\right\|_1 = 0.$$

Now

$$\begin{aligned}&\left\|\rho^{\mathbf{u}}_{j_n} - V^*_{j_n}\phi^{\mathbf{u}}V_{j_n}\right\|_1 = \\ &\left\|\mathrm{Ad}\left[U_{j_n}\left(\frac{\mathbf{u}}{\sqrt{n}}\right)\right]\left(\rho^{\mathbf{0}}_{j_n}\right) - V^*_{j_n}\phi^{\mathbf{u}}V_{j_n}\right\|_1 \leq \\ &\left\|\rho^{\mathbf{0}}_{j_n} - V^*_{j_n}\phi^{\mathbf{0}}V_{j_n}\right\|_1 + \\ &\left\|\mathrm{Ad}\left[U_{j_n}\left(\frac{\mathbf{u}}{\sqrt{n}}\right)\right]\left(V^*_{j_n}\phi^{\mathbf{0}}V_{j_n}\right) - V^*_{j_n}\phi^{\mathbf{u}}V_{j_n}\right\|_1.\end{aligned}$$

The first term on the right side of the inequality converges to zero by Lemma 7.5.4, uniformly for any sequence (j_n) such that $j_n \in \mathcal{J}_{n,\epsilon}$ and does not depend on $\mathbf{u}$. Using (7.20) and (7.21) we bound the second term by

$$\frac{1}{s\sqrt{2\pi}} \int e^{-|\mathbf{z}|^2/2s^2} \|\Delta(\mathbf{u}, \mathbf{z}, j_n)\|_1 d^2\mathbf{z}$$

where the operator $\Delta(\mathbf{u}, \mathbf{z}, j_n)$ is given by

$$\begin{aligned} \Delta(\mathbf{u}, \mathbf{z}, j_n) \;:=\; & \mathrm{Ad}\left[U_{j_n}\left(\frac{\mathbf{u}}{\sqrt{n}}\right)\right](V_{j_n}^*|\mathbf{z}\rangle\langle\mathbf{z}|V_{j_n}) - \\ & V_{j_n}^*\left|\mathbf{z}+\sqrt{2\mu-1}\alpha_{\mathbf{u}}\right\rangle\left\langle\mathbf{z}+\sqrt{2\mu-1}\alpha_{\mathbf{u}}\right|V_{j_n} \end{aligned}$$

We analyze the expression under the integral. Let $\tilde{\mathbf{z}} \in \mathbb{R}^2$ be such that $\alpha_{\tilde{z}} = \mathbf{z}/\sqrt{2\mu-1}$, then

$$\begin{aligned} & \left\|\mathrm{Ad}\left[U_{j_n}\left(\frac{\mathbf{u}}{\sqrt{n}}\right)\right](V_{j_n}^*|\mathbf{z}\rangle\langle\mathbf{z}|V_{j_n}) - V_{j_n}^*|\mathbf{z}+\sqrt{2\mu-1}\alpha_{\mathbf{u}}\rangle\langle\mathbf{z}+\sqrt{2\mu-1}\alpha_{\mathbf{u}}|V_{j_n}\right\|_1 \leq \\ & \left\|\mathrm{Ad}\left[U_{j_n}\left(\frac{\mathbf{u}}{\sqrt{n}}\right)U_{j_n}\left(\frac{\tilde{\mathbf{z}}}{\sqrt{n}}\right)\right](|j_nj_n\rangle\langle j_nj_n|) - \mathrm{Ad}\left[U_{j_n}\left(\frac{\mathbf{u}+\tilde{\mathbf{z}}}{\sqrt{n}}\right)\right](|j_nj_n\rangle\langle j_nj_n|)\right\|_1 + \\ & \left\|V_{j_n}\mathrm{Ad}\left[U_{j_n}\left(\frac{\tilde{\mathbf{z}}}{\sqrt{n}}\right)\right](|j_nj_n\rangle\langle j_nj_n|)\,V_{j_n}^* - |\mathbf{z}\rangle\langle\mathbf{z}|\right\|_1 + \\ & \left\|V_{j_n}\mathrm{Ad}\left[U_{j_n}\left(\frac{\mathbf{u}+\tilde{\mathbf{z}}}{\sqrt{n}}\right)\right](|j_nj_n\rangle\langle j_nj_n|)\,V_{j_n}^* - |\mathbf{z}+\sqrt{2\mu-1}\alpha_{\mathbf{u}}\rangle\langle\mathbf{z}+\sqrt{2\mu-1}\alpha_{\mathbf{u}}|\right\|_1. \end{aligned}$$

By Corollary 7.5.3, the first term on the right side converges to zero uniformly in $(\mathbf{u}, j_n) \in I^2 \times \mathcal{J}_{n,\epsilon}$. By Lemma 7.5.4 we have that the last two terms converge to zero uniformly in $(\mathbf{u}, j_n) \in I^2 \times \mathcal{J}_{n,\epsilon}$. Thus if we denote

$$F_n(\mathbf{z}) := \sup_{j_n \in \mathcal{J}_{n,\epsilon}} \sup_{\mathbf{u} \in I^2} \|\Delta(\mathbf{u}, \mathbf{z}, j_n)\|_1$$

then $0 \leq F_n(\mathbf{z}) \leq 2$, $\lim_{n\to\infty} F_n(\mathbf{z}) = 0$ for all $\mathbf{z} \in \mathbb{R}^2$, and by the Lebesgue dominated convergence theorem we get

$$\lim_{n\to\infty} \frac{1}{s\sqrt{2\pi}} \int e^{-|z|^2/2s^2} F_n(\mathbf{z}) d^2\mathbf{z} = 0.$$

This implies the statement of the Proposition 7.5.1.

□

7.6 Construction of the inverse channel S_n

To complete our proof of asymptotic equivalence as defined by (7.4), we must now exhibit the inverse channel S_n which maps the displaced thermal states $\phi^{\mathbf{u}}$ of the oscillator into approximations of the rotated spin states. As the latter are block diagonal with weights $p_n(j)$ as defined in equation (7.12) , it is natural to look for S_n of the form

$$S_n(\phi) = \bigoplus_{j=0,1/2}^{n/2} p_n(j) S_n^j(\phi) \otimes \frac{\mathbf{1}}{n_j},$$

where S_n^j are channels with outputs in $\mathcal{H}_j$. Moreover because $V_j : \mathcal{H}_j \to L^2(\mathbb{R})$ is an isometry we can choose S_n^j such that

$$S_n^j \left(V_j \rho V_j^* \right) = \rho, \tag{7.22}$$

for all density matrices ρ on $\mathcal{H}_j$. This property does not fix the channel completely but it is sufficient for our purposes.

Theorem 7.6.1. *The following holds*

$$\lim_{n\to\infty} \sup_{u\in I^2} \|S_n(\phi^{\mathbf{u}}) - \rho_n^{\mathbf{u}}\|_1 = 0.$$

Proof. As both $\rho_n^{\mathbf{u}}$ and $\phi^{\mathbf{u}}$ are block-diagonal we may decompose their distance as

$$\begin{aligned}
\|S_n(\phi^{\mathbf{u}}) - \rho_n^{\mathbf{u}}\|_1 &= \sum_{j=0,1/2}^{n/2} p_n(j) \|S_n^j(\phi^{\mathbf{u}}) - \rho_{j,n}^{\mathbf{u}}\|_1 \\
&\leq \sum_{j\notin\mathcal{J}_{n,\epsilon}} 2p_n(j) + \sum_{j\in\mathcal{J}_{n,\epsilon}} p_n(j) \|S_n^j(\phi^{\mathbf{u}}) - S_n^j\left(V_j \rho_{j,n}^{\mathbf{u}} V_j^*\right)\|_1 \\
&+ \sum_{j\in\mathcal{J}_{n,\epsilon}} p_n(j) \|S_n^j\left(V_j \rho_{j,n}^{\mathbf{u}} V_j^*\right) - \rho_{j,n}^{\mathbf{u}}\|_1 \\
&\leq 2 \sum_{j\notin\mathcal{J}_{n,\epsilon}} p_n(j) + \sum_{j\in\mathcal{J}_{n,\epsilon}} p_n(j) \|\phi^{\mathbf{u}} - V_j \rho_{j,n}^{\mathbf{u}} V_j^*\|_1,
\end{aligned}$$

where we have used at the last line that S_n^j is a contraction and property (7.22) of S_n^j. Now the first sum is going to 0 by (7.15) and the second sum is also uniformly going to 0 by use of Proposition 7.5.1.

□

7.7 Applications

7.7.1 Local asymptotic equivalence of the optimal Bayesian measurement and the heterodyne measurement

In this subsection we begin a comparison of the pointwise (local) point of view with the global one used in the Bayesian approach. The result is that the optimal $SU(2)$ covariant measurement (Bagan et al., 2006; Hayashi et Matsumoto, 2004) converges locally to the optimal measurement for the family of displaced Gaussian states which is a heterodyne measurement (Holevo, 1982). Together with the results on the asymptotic local minimax optimality of this measurement, this closes a circle of ideas relating the different optimality notions and the relations between the optimal measurements.

Let us recall what are the ingredients of the state estimation problem in the Bayesian framework (Bagan et al., 2006). We choose as cost function the fidelity squared $F(\rho,\sigma)^2 = \mathrm{Tr}(\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}})^2$ and fix a prior prior distribution π over all states in $\mathbb{C}^2$ which is invariant under the $SU(2)$ symmetry group. Given n identical systems $\rho^{\otimes n}$ we would like to find a measurement M_n - whose outcome is the estimator $\hat{\rho}_n$ - which *maximizes*

$$R_{\pi,n} := \int \langle F(\hat{\rho}_n,\rho)^2\rangle \pi(d\rho).$$

By the $SU(2)$ invariance of π, the optimal measurement can be chosen to be $SU(2)$ covariant i.e.

$$UM_n(d\sigma)U^* = M_n(U^* d\sigma U),$$

and can be described as follows. First we use the decomposition (7.10) to make a "which block" measurement and obtain a result j and the conditional state $\rho_{j,n}$ as in (7.11). This part will provide us the eigenvalues of the estimator. Next we perform block-wise the covariant measurement $M_{j,n}(d\overrightarrow{s}) = m_{j,n}(\overrightarrow{s})d\overrightarrow{s}$ with

$$m_{j,n}(\overrightarrow{s}) := (2j+1)U_j(\overrightarrow{s})^*|j\rangle\langle j|U_j(\overrightarrow{s}) \otimes \mathbf{1}_j$$

whose result is a unit vector $\overrightarrow{s}$ where $U(\overrightarrow{s})$ is a unitary rotating the vector state $|\overrightarrow{s}\rangle$ to $|\uparrow\rangle$. The complete estimator is then $\hat{\rho}_n = \frac{1}{2}(\mathbf{1} + \frac{2j}{n}\overrightarrow{s}\overrightarrow{\sigma})$.

We pass now to the description of the heterodyne measurement for the quantum harmonic oscillator. This measurement has outcomes $\mathbf{u} \in \mathbb{R}^2$ and is covariant with respect to the translations induced by the displacement operators $D(\mathbf{z})$ such that $H(d\mathbf{u}) = h(\mathbf{u})d\mathbf{u}$ with

$$h(\mathbf{u}) := (2\mu-1)D(-\sqrt{2\mu-1}\alpha_{\mathbf{u}})|0\rangle\langle 0|D(\sqrt{2\mu-1}\alpha_{\mathbf{u}}).$$

Using Theorem 7.1.1 we can map H into a measurement on the n-spin system as follows: first we perform the which block step as in the case of the $SU(2)$-covariant measurements. Then we map $\rho_{j,n}$ into an oscillator state using the isometry V_j (see (7.18)), and subsequently we perform H. The result $\mathbf{u}$ will define our estimator for the local state, i.e.

$$\hat{\rho}_n = U\left(\frac{\mathbf{u}}{\sqrt{n}}\right)\begin{pmatrix} \frac{1}{2}+\frac{j}{n} & 0 \\ 0 & \frac{1}{2}-\frac{j}{n} \end{pmatrix} U\left(\frac{\mathbf{u}}{\sqrt{n}}\right)^* . \tag{7.23}$$

We denote by H_n the resulting measurement with values in the states on $\mathbb{C}^2$.

The next Theorem shows that in a *local neighborhood* of a fixed state $\rho^{\mathbf{0}}$, the $SU(2)$-covariant measurement M_n and the heterodyne type measurement H_n are asymptotically equivalent in the sense that the probability distributions $P(M_n,\rho)$ and $P(H_n,\rho)$ are close to each other uniformly over all local states ρ such that $\|\rho-\rho^{\mathbf{0}}\|_1 \leq \frac{C}{\sqrt{n}}$ for a fixed but arbitrary constant $C<\infty$.

Theorem 7.7.1. *Let $\rho^{\mathbf{0}}$ be as in (7.6), and let*

$$B_n(I) = \left\{\rho^{\mathbf{v}/\sqrt{n}} : \mathbf{v}\in I^2\right\}, \quad , |I|<\infty$$

be a local family of states around $\rho^{\mathbf{0}}$. Then

$$\lim_{n\to\infty}\sup_{\rho\in B_n(I)} \|P(M_n,\rho)-P(H_n,\rho)\|_1 = 0$$

Proof. Note first that both $P(M_n,\rho)$ and $P(H_n,\rho)$ are distributions over the Bloch sphere and the marginals over the length of the Bloch vectors are identical because by construction the first step of both measurements is the same. Then

$$\|P(M_n,\rho)-P(H_n,\rho)\|_1 = \sum_j p_n(j)\int |\mathrm{Tr}(\rho_{j,n}(m_{j,n}(\vec{s})-h_{j,n}(\vec{s})))|\, d\vec{s}.$$

According to (7.15) we can restrict the summation to the interval $\mathcal{J}_{n,\epsilon}$ around $j=n(\mu-\frac{1}{2})$. By Theorem 7.1.1 we can replace (whenever needed) the local states $\rho_{j,n}^{\mathbf{v}/\sqrt{n}}$ by their limits in the oscillator space $\phi^{\mathbf{v}}$ with an asymptotically vanishing error, uniformly over $\mathbf{v}\in I^2$.

We make now the change of variable $\vec{s}\to\mathbf{u}$ where $\mathbf{u}\in\mathbb{R}^2$ belongs to the ball $|\mathbf{u}|<2\sqrt{n}\pi$, and is the smallest vector such that $U\left(\frac{\mathbf{u}}{\sqrt{n}}\right)=U(\vec{s})$.

The density of the $SU(2)$ estimator with respect to the measure $d\mathbf{u}$ is

$$m_{j,n}(\mathbf{u}) := \frac{2j+1}{n} U_j\left(\frac{\mathbf{u}}{\sqrt{n}}\right)^* |j\rangle\langle j| U_j\left(\frac{\mathbf{u}}{\sqrt{n}}\right) J\left(\frac{\mathbf{u}}{\sqrt{n}}\right),$$

where J is the determinant of a Jacobian related with the change of variables such that $J(0) = 1$.

Similarly the density of the homodyne-type estimator becomes

$$h_{j,n}(\mathbf{u}) := \sum_{k\in\mathbb{N}} V_j^* h\left(\mathbf{u} + 2k\sqrt{n}\pi\frac{\mathbf{u}}{|\mathbf{u}|}\right) V_j \, |J_{k,n}(\mathbf{u})|,$$

because displacements in the same direction which differ by multiples of $2\sqrt{n}\pi$ lead to the same unitary on the qubits. Here $J_{k,n}(\mathbf{u})$ is again the determinant of the Jacobian of the map from the k-th ring to the disk, in particular $J_{0,n}(\mathbf{u}) = 1$.

The integral becomes then

$$\int_{|\mathbf{u}|\leq 2\pi\sqrt{n}} \left|\mathrm{Tr}\left(\rho_{j,n}^{\mathbf{v}/\sqrt{n}}(m_{j,n}(\mathbf{u}) - h_{j,n}(\mathbf{u}))\right)\right| d\mathbf{u}.$$

We bound this integral by the sum of two terms, the first one being

$$\int_{|\mathbf{u}|\leq 2\pi\sqrt{n}} \left|\mathrm{Tr}\left(\rho_{j,n}^{\mathbf{v}/\sqrt{n}}(m_{j,n}(\mathbf{u}) - \tilde{h}_j(\mathbf{u}))\right)\right| d\mathbf{u},$$

where $\tilde{h}_j(\mathbf{u})$ is just the term with $k = 0$ in $h_{j,n}(\mathbf{u})$. By Lemma 7.5.4, for any fixed $\mathbf{u}$ we have $m_{j,n}(\mathbf{u}) \to h(\mathbf{u})$ uniformly over $j \in \mathcal{J}_{n,\epsilon}$. Using similar estimates as in Lemma 7.5.4 it can be shown that the function under the integral is bounded by a fixed integrable function $g(\mathbf{u})$ uniformly over $\mathbf{v} \in I^2$, and then we can use dominated convergence to conclude that the integral converges to 0 uniformly over $\mathbf{v} \in I^2$ and $j \in \mathcal{J}_{n,\epsilon}$.

The second integral is

$$\int_{|\mathbf{u}|\leq 2\pi\sqrt{n}} \left|\mathrm{Tr}\left(\rho_{j,n}^{\mathbf{v}/\sqrt{n}}(\tilde{h}_j(\mathbf{u}) - h_{j,n}(\mathbf{u}))\right)\right| d\mathbf{u},$$

which is smaller than

$$\int_{|\mathbf{u}|>2\pi\sqrt{n}} \left|\mathrm{Tr}\left(\rho_{j,n}^{\mathbf{v}/\sqrt{n}} h\,(\mathbf{u})\right)\right| d\mathbf{u},$$

which converges uniformly to 0. This can be seen by replacing the states with $\phi^{\mathbf{v}}$ which are "confined" to a fixed region of the size I^2 in the phase space, while the terms $h(\mathbf{u})$ are Gaussians located at distance at least $2\pi\sqrt{n}$ from the origin.

Putting these two estimates together we obtain the desired result.

□

Remark. The result in the above theorem holds more generally for all states in a local neighborhood of ρ^0 but for the proof we need a slightly more general version of Theorem 7.1.1 where the eigenvalues of the density matrices are not fixed but allowed to vary in a local neighborhood of $(\mu, 1-\mu)$. This result will be presented in a future work concerning the general case of d-dimensional states.

7.7.2 The optimal Bayes measurement is also locally asymptotic minimax

In this subsection we will introduce some ideas from the pointwise approach to state estimation. We show that the measurement which is known to be optimal for a uniform prior in the Bayesian set-up, is also asymptotically optimal in the pointwise sense.

Using the jargon of mathematical statistics, we will call *quantum statistical experiment (model)* (Petz et Jenčová, 2006) a family $\{\rho^\theta \in M(\mathbb{C}^d) : \theta \in \Theta\}$ of density matrices indexed by a parameter belonging to a set Θ. The main examples of quantum statistical experiments considered so far are that of n identical qubits

$$\mathcal{F} := \left\{\rho^{\otimes n} : \rho \in M(\mathbb{C}^2)\right\},$$

the local model

$$\mathcal{F}_n^I := \left\{\rho_n^{\mathbf{u}} = \left(\rho^{\mathbf{u}/\sqrt{n}}\right)^{\otimes n}, \mathbf{u} \in I^2\right\},$$

and its "limit" model

$$\mathcal{G}^I := \{\phi^{\mathbf{u}}, \mathbf{u} \in I^2\},$$

where $I = [-a, a]$, and $\rho_n^{\mathbf{u}}$ and $\phi^{\mathbf{u}}$ are defined by (7.1) and (7.2). More generally we can replace the square I^2 by an arbitrary region K in the parameter space and obtain:

$$\mathcal{G}^K := \{\phi^{\mathbf{u}}, \mathbf{u} \in K \subset \mathbb{R}^2\}.$$

We shall also make use of

$$\mathcal{G} := \{\phi^{\mathbf{u}}, \mathbf{u} \in \mathbb{R}^2\}.$$

A natural choice of distance between density matrices is the fidelity square

$$F(\rho, \sigma)^2 = \left[\mathrm{Tr}\left(\sqrt{\rho}\sigma\sqrt{\rho}\right)^{1/2}\right]^2,$$

which is locally quadratic in first order approximation, i.e.

$$F(\rho_n^{\mathbf{u}}, \rho_n^{\mathbf{v}})^2 \approx \frac{1}{n}\|\mathbf{u} - \mathbf{v}\|^2.$$

As we expect that reasonable estimators are in a local neighborhood of the true state we will replace the fidelity square by the local distance

$$d(\mathbf{u}, \hat{\mathbf{u}}) = \|\hat{\mathbf{u}} - \mathbf{u}\|^2.$$

and define the risk of a measurement-estimator pair as $R_M(\mathbf{u}, \hat{\mathbf{u}}) = \langle d(\mathbf{u}, \hat{\mathbf{u}}) \rangle$, keeping in mind the factor $1/n$ relating the risks expressed in local and global parameters.

Similarly to the Bayesian approach, we are interested in estimators which have small risk *everywhere* in the parameter space and we define a worst case figure of merit called minimax risk.

Definition 7.7.2. *The minimax risk of a quantum statistical experiment $\mathcal{E}$ over the parameter space Θ is defined as*

$$C(\mathcal{E}) = \inf_{\hat{\mathbf{u}}} \sup_{\mathbf{u} \in \Theta} R_M(\mathbf{u}, \hat{\mathbf{u}}). \tag{7.24}$$

where the infimum is taken over all measurements and estimators $(M, \hat{\mathbf{u}})$.

The minimax risk tells us how difficult is the model and thus we expect that if two models are "close" to each other then their minimax risks are almost equal. The "statistical distance" between quantum experiments is defined in a natural way with direct physical interpretation and such a problem has been already addressed by Chefles et al. (2003) for the case of a quantum statistical experiment consisting of a finite family of pure states.

Definition 7.7.3. *Let $\mathcal{E} = \{\rho^\theta \in M(\mathbb{C}^d) : \theta \in \Theta\}$ and $\mathcal{F} = \{\tau^\theta \in M(\mathbb{C}^p) : \theta \in \Theta\}$ be two quantum statistical experiments (models) with the same parameter space Θ. We define the discrepancies*

$$\delta(\mathcal{E}, \mathcal{F}) = \inf_{T} \sup_{\theta \in \Theta} \|T(\rho^\theta) - \tau^\theta\|_1,$$
$$\delta(\mathcal{F}, \mathcal{E}) = \inf_{S} \sup_{\theta \in \Theta} \|\rho^\theta - S(\tau^\theta)\|_1,$$

where the infimum is taken over all trace preserving channels $T : M(\mathbb{C}^d) \to M(\mathbb{C}^p)$ and $S : M(\mathbb{C}^p) \to M(\mathbb{C}^d)$.

With this terminology, our main result states that for any bounded interval I:

$$\lim_{n \to \infty} \max\left(\delta(\mathcal{F}_n^I, \mathcal{G}^I), \delta(\mathcal{G}^I, \mathcal{F}_n^I)\right) = 0. \tag{7.25}$$

As suggested above, the discrepancy has a direct statistical interpretation: if we want to estimate θ in both statistical experiments $\mathcal{E}$ and $\mathcal{F}$ and we choose a bounded loss function $d(\theta, \hat{\theta}) \leq K$ then for any measurement and estimator $\hat{\theta}$ for $\mathcal{F}$ with risk

$R_M(\theta,\hat{\theta}) = \langle d(\theta,\hat{\theta})\rangle$ we can find a measurement N on $\mathcal{E}$ whose risk is at most $R_M(\theta,\hat{\theta}) + K\delta(\mathcal{E},\mathcal{F})$. Indeed if we choose T such that the infimum in the definition of $\delta(\mathcal{E},\mathcal{F})$ is achieved, we can map the state ρ^θ through the channel T and then perform M to obtain an estimator $\tilde{\theta}$ for which

$$R_N(\theta,\tilde{\theta}) = \langle d(\theta,\tilde{\theta})\rangle = \int_\Theta d(\theta,\tilde{\theta})\mathrm{Tr}\left(T(\rho^\theta)M(d\tilde{\theta})\right) \leq$$
$$\int_\Theta d(\theta,\tilde{\theta})\mathrm{Tr}\left(\tau^\theta M(d\tilde{\theta})\right) + \|d\|_\infty \|T(\rho^\theta) - \tau^\theta\|_1 \leq$$
$$R_M(\theta,\hat{\theta}) + K\delta(\mathcal{E},\mathcal{F}).$$

This means that asymptotically the difficulty of estimating the parameter θ in the two models is the same. With the above definition of the minimax risk and using the convergence (7.25) we obtain the following lemma.

Lemma 7.7.4. *Let $I = [-a,a]$ with $0 < a < \infty$, then*

$$\lim_{n\to\infty} C(\mathcal{F}_n^I) = C(\mathcal{G}^I)$$

The minimax risk for the local family $\mathcal{F}_n^I$ is a figure of merit for the "local difficulty" of the global model $\mathcal{F}_n$. It asymptotically converges to the minimax risk of a family of thermal states. However this quantity depends on the arbitrary parameter $I = [-a,a]$ which we would like to remove as our last step in defining the *local asymptotic minimax risk*:

$$C_{\text{l.a.m.}}(\mathcal{F}_n : n \in \mathbb{N}) := \lim_{a\to\infty}\lim_{n\to\infty} C(\mathcal{F}_n^I) = \lim_{a\to\infty} C(\mathcal{G}^I).$$

As one might expect, the minimax risks for the restricted families of thermal states will converge to that of the experiment with no restrictions on the paramaters. The proof of this fact is however non-trivial.

Lemma 7.7.5. *Let $I = [-a,a]$, then we have*

$$\lim_{a\to\infty} C(\mathcal{G}^I) = C(\mathcal{G})$$

Moreover the heterodyne measurement saturates $C(\mathcal{G})$, and thus $C(\mathcal{G})$ is equal to the Holevo bound.

Proof. The inequality in one direction is easy. For any estimator, $\sup_{\mathbf{u}\in I^2} R_M(\mathbf{u},\hat{\mathbf{u}}) \leq \sup_{\mathbf{u}\in\mathbb{R}^2} R_M(\mathbf{u},\hat{\mathbf{u}})$, so that $C(\mathcal{G}^I) \leq C(\mathcal{G})$ and the same holds for the limit. By the same reasoning, for any $K_1 \subset K_2 \subset \mathbb{R}^2$ we have $C(\mathcal{G}^{K_1}) \leq C(\mathcal{G}^{K_2})$.

When calculating minimax bounds we are interested in the worst risk of estimators within some parameter region K, and this worst risk is obviously higher than the

Bayes risk with respect to the probability distribution with constant density on K. We shall work on $B(0, c+b)$ the ball of center 0 and radius $(c+b)$, with $b > c$, and denote our measurement M with density $m(\hat{\mathbf{u}})\mathrm{d}\hat{\mathbf{u}}$. In general M need not have a density, but this will ease notations. Then

$$\sup_{\mathbf{u}\in B(0,c+b)} R_M(\mathbf{u}, \hat{\mathbf{u}}) \geq \int_{B(0,c+b)\times\mathbb{R}^2} \frac{\mathrm{d}\mathbf{u}\,\mathrm{d}\hat{\mathbf{u}}}{\pi(c+b)^2} \|\mathbf{u}-\hat{\mathbf{u}}\|^2 \operatorname{Tr}\left(\phi^{\mathbf{u}} m(\hat{\mathbf{u}})\right). \quad (7.26)$$

We fix the following notations

$$f(\mathcal{D}) = \int_{\mathcal{D}} \mathbf{dudv} \|\mathbf{x}-\mathbf{y}\|^2 \operatorname{Tr}\left(\phi^{\mathbf{u}} m(\mathbf{v})\right),$$
$$g(\mathcal{D}) = \int_{\mathcal{D}} \mathbf{dudv} \operatorname{Tr}\left(\phi^{\mathbf{u}} m(\mathbf{v})\right),$$

and define the domains

$$\mathcal{D}_1 = \{(\mathbf{u}, \hat{\mathbf{u}}) | \mathbf{u} \in B(0, c+b), \hat{\mathbf{u}} \in \mathbb{R}^2\}$$
$$\mathcal{D}_2 = \{(\mathbf{u}+\mathbf{k}, \mathbf{k}) | \mathbf{u} \in B(0, c), \mathbf{k} \in B(0, b)\}$$
$$\mathcal{D}_3 = \{(\mathbf{u}, \mathbf{u}+\mathbf{h}) | \mathbf{u} \in B(0, b-c), \mathbf{h} \in B(0, c)\}$$
$$\mathcal{D}_4 = \{(\mathbf{u}, \mathbf{u}+\mathbf{h}) | \mathbf{u} \in B(0, b-c), \mathbf{h} \in \mathbb{R}^2 \backslash B(0, c)\}.$$

Notice the following relations:

$$\mathcal{D}_3 \subset \mathcal{D}_2 \subset \mathcal{D}_1, \quad \mathcal{D}_4 \subset \mathcal{D}_1 \backslash \mathcal{D}_2. \quad (7.27)$$

Then (7.26) can be rewritten as

$$\sup_{\mathbf{u}\in B(0,c+b)} R_M(\mathbf{u}, \hat{\mathbf{u}}) \geq \frac{1}{\pi(b+c)^2} f(\mathcal{D}_1).$$

The following inequalities follow directly from the definitions:

$$f(\mathcal{D}_2) \leq c^2 g(\mathcal{D}_2) \qquad f(\mathcal{D}_3) \leq c^2 g(\mathcal{D}_3)$$
$$f(\mathcal{D}_4) \geq c^2 g(\mathcal{D}_4) \qquad g(\mathcal{D}_4) + g(\mathcal{D}_3) = \pi(b-c)^2.$$

Using these and (7.27), we may write:

$$\begin{aligned}
&\frac{1}{\pi(c+b)^2} f(\mathcal{D}_1) \geq \frac{1}{\pi(c+b)^2} \left(f(\mathcal{D}_2) + f(\mathcal{D}_4)\right) \\
&\geq \frac{1}{\pi(c+b)^2} \left(f(\mathcal{D}_2) + c^2 g(\mathcal{D}_4)\right) \\
&= \frac{(b-c)^2}{(b+c)^2} \left(\frac{f(\mathcal{D}_2)}{g(\mathcal{D}_2)} \frac{g(\mathcal{D}_2)}{\pi(b-c)^2} + c^2 - c^2 \frac{g(\mathcal{D}_3)}{\pi(b-c)^2} \right) \\
&\geq \frac{(b-c)^2}{(b+c)^2} \left(c^2 + \frac{g(\mathcal{D}_3)}{\pi(b-c)^2} \left(\frac{f(\mathcal{D}_2)}{g(\mathcal{D}_2)} - c^2 \right) \right) \\
&\geq \frac{(b-c)^2}{(b+c)^2} \frac{f(\mathcal{D}_2)}{g(\mathcal{D}_2)}.
\end{aligned} \tag{7.28}$$

We analyze now the expression $f(\mathcal{D}_2)/g(\mathcal{D}_2)$. By using the definition (7.2) of the displaced thermal states $\phi^{\mathbf{u}}$ we get that $\mathrm{Tr}\left[\phi^{\mathbf{u}+\mathbf{k}} m(\mathbf{l})\right] = \mathrm{Tr}\left[\phi^{\mathbf{k}} m_{\mathbf{u}}(\mathbf{l})\right]$, where

$$m_{\mathbf{u}}(\mathbf{l}) := D(-\sqrt{2\mu-1}\alpha_{\mathbf{u}}) m(\mathbf{l}) D(\sqrt{2\mu-1}\alpha_{\mathbf{u}}).$$

Then

$$g(\mathcal{D}_2) = \int_{B(0,c)\times B(0,b)} \mathbf{du}\mathbf{dk}\, \mathrm{Tr}\left[\phi^{\mathbf{u}+\mathbf{k}} m(\mathbf{k})\right] = \mathrm{Tr}\left[\tilde{\phi}_c \tilde{m}_b\right],$$

where we have written

$$\tilde{\phi}_c = \int_{B(0,c)} \phi^{\mathbf{u}} \mathbf{du}, \qquad \tilde{m}_b = \int_{B(0,b)} m_{\mathbf{k}}(\mathbf{k}) \mathbf{dk}.$$

Upon writing $v_c := \int_{B(0,c)} \|\mathbf{u}\|^2 \phi^{\mathbf{u}} \mathbf{du}$, we get similarly $f(\mathcal{D}_2) = \mathrm{Tr}\left[v_c \tilde{m}_b\right]$. Note that by rotational symmetry v_c and $\tilde{\phi}_c$ are diagonal in the number operator eigenbasis, so without restricting the generality we may assume that $\tilde{m}_b$ is also diagonal in that basis: $\tilde{m}_b = \sum_k p_k |k\rangle\langle k|$. We have then

$$\frac{f(\mathcal{D}_2)}{g(\mathcal{D}_2)} = \frac{\sum_{k\in\mathbb{N}} p_k \langle k|v_c|k\rangle}{\sum_{k\in\mathbb{N}} p_k \langle k|\tilde{\phi}_c|k\rangle} \geq \inf_{k\in\mathbb{N}} \frac{\langle k|v_c|k\rangle}{\langle k|\tilde{\phi}_c|k\rangle}.$$

The infimum on the right side is achieved by the vacuum vector. By Lemma 7.7.6, this fact follows from the inequality

$$\frac{\langle k|\phi^{\mathbf{u}_1}|k\rangle}{\langle k|\phi^{\mathbf{u}_2}|k\rangle} \geq \frac{\langle 0|\phi^{\mathbf{u}_1}|0\rangle}{\langle 0|\phi^{\mathbf{u}_2}|0\rangle}, \qquad \|\mathbf{u}_1\| \geq \|\mathbf{u}_2\|,$$

which can be checked by explicit calculations.

Letting now c and b go to infinity with $c = o(b)$ and using (7.28), we obtain that

$$\lim_{a\to\infty} C(\mathcal{G}_a) \geq \frac{\int_{\mathbb{R}^2} \langle 0|\phi^{\mathbf{u}}|0\rangle \, \|\mathbf{u}\|^2 d\mathbf{u}}{\int_{\mathbb{R}^2} \langle 0|\phi^{\mathbf{u}}|0\rangle \, d\mathbf{u}},$$

which is exactly the pointwise risk of the heterodyne measurement $H(d\mathbf{u}) = h(\mathbf{u})d\mathbf{u}$ whose density is

$$h(\mathbf{u}) = (2\mu - 1)D(-\sqrt{2\mu - 1}\alpha_{\mathbf{u}})|0\rangle\langle 0|D(-\sqrt{2\mu - 1}\alpha_{\mathbf{u}}).$$

By symmetry this pointwise risk does not depend on the point, so that $C(\mathcal{G}) \leq R_H(\mathbf{u}, \hat{\mathbf{u}})$. And we have our second inequality: $\lim_{a\to\infty} C(\mathcal{G}_a) \geq C(\mathcal{G})$.

Moreover, the heterodyne measurement is known to saturate the Holevo bound for $G = Id$ and the Cramér-Rao bound for locally unbiased estimators (Holevo, 1982; Hayashi et Matsumoto, 2004). We conclude that the local minimax risk for qubits is equal to the minimax risk for the limit Gaussian quantum experiment which is achieved by the heterodyne measurement.

□

Lemma 7.7.6. *Let p and q be two probability densities on $[0, 1]$ and assume that*

$$\frac{p(x_1)}{p(x_2)} \geq \frac{q(x_1)}{q(x_2)}, \qquad x_1 \geq x_2.$$

Then $\int x^2 p(x)dx \geq \int x^2 q(x)dx$.

Proof. It is enough to show that there exists a point $x_0 \in [0, 1]$ such that $p(x) \leq q(x)$ for $x \leq x_0$ and $p(x) \geq q(x)$ for $x \geq x_0$. Now, if $p(x) \leq q(x)$ then by using the assumption we get that $p(y) \leq q(y)$ for all $y \leq x$. Similarly, if $p(x) \geq q(x)$ then $p(y) \leq q(y)$ for all $y \geq x$. This implies the existence of the crossing point x_0.

□

We end this section with the conclusion that the optimal measurement from the Bayesian point of view is also asymptotically optimal from the pointwise point of view. Let us denote by $(M_n, \hat{\mathbf{u}})$ the measurement-estimator pair from (Bagan et al., 2006; Hayashi et Matsumoto, 2004).

Proposition 7.7.7. *The optimal measurement-estimator pair $(M_n, \hat{\mathbf{u}})$ is a local asymptotic minimax estimation scheme. That is*

$$\lim_{n\to\infty} R_{M_{cov}}(\mathbf{u}, \hat{\mathbf{u}}) = C_{\text{l.a.m}}(\mathcal{F}_n : n \in \mathbb{N}).$$

Proof. The pointwise risk of M_{cov} is known to converge to that of the heterodyne measurement (Bagan et al., 2006). The rest follows from Lemma 7.7.4 and Lemma 7.7.5.

□

7.7.3 Discrimination of states

Another illustration of the local asymptotic normality Theorem is the problem of discriminating between two states ρ^+ and ρ^-. When the two states are fixed, this problem has been solved by Helstrom (1976), and if we are given n systems in state $\rho_\pm^{\otimes n}$ then the probability of error converge to 0 exponentially. Here we consider the problem of distinguishing between two states $\rho_n^\pm$ which approach each other as $n \longrightarrow \infty$ with rate $\|\rho_n^+ - \rho_n^-\|_1 \approx \frac{1}{\sqrt{n}}$. In this case the probability of error does not go to 0 because the problem becomes more difficult as we have more systems, and converges to the limit problem of distinguishing between two fixed Gaussian states of a quantum oscillator.

This problem is interesting for several reasons. Firstly it shows that the convergence in Theorem 7.1.1 can be used for finding asymptotically optimal procedures for various statistical problems such as that of parameter estimation and hypothesis testing. Secondly, for any fixed n the optimal discrimination is performed by a rather complicated *joint* measurement and the hope is that the asymptotic problem of discriminating between two Gaussian states may provide a more realistic measurement which can be implemented in the lab. Thirdly, this example shows that a non-commuting one-parameter families of states is not "classical" as it is sometimes argued, but should be considered as a quantum "resource" which cannot be transformed into a classical one without loss of information. More explicitly, the optimal measurement for estimating the parameter is not optimal for other statistical problems such as the one considered here, and thus different statistical decision problems are accompanied by mutually incompatible optimal measurements.

Let is recall the framework of quantum hypothesis testing for two states $\rho^\pm$: we consider two-outcomes POVM's $M = (M_-, M_+)$ with $0 \leq M_+ \leq \mathbf{1}$ and $M_- = \mathbf{1} - M_+$ such that the probability of error when the state is ρ^- is given by $\mathrm{Tr}(M_+\rho^-)$,and similarly for ρ^+. As we do not know the state, we want to minimize our worst-case probability error. Our figure of merit (the lower, the better) is therefore:

$$R(\rho^+, \rho^-) = \inf_M \max\{\mathrm{Tr}(\rho_+ M_-), \mathrm{Tr}(\rho_+ M_-)\}$$

Now we are interested in the case when $\rho^\pm = \rho_n^{\pm\mathbf{u}}$ as defined in (7.1), and in the limit $\rho_\pm = \phi^{\pm\mathbf{u}}$ (recall that both $\rho_n^\mathbf{u}$ and $\phi^\mathbf{u}$ depend on μ). We then have:

Theorem 7.7.8. *The following limit holds*

$$\lim_{n\to\infty} R(\rho_n^{\mathbf{u}}, \rho_n^{-\mathbf{u}}) = R(\phi^{\mathbf{u}}, \phi^{-\mathbf{u}}).$$

Moreover for pure states this limit is equal to $\left(1-(1-e^{-4|\mathbf{u}|^2})^{1/2}\right)/2$ *which is strictly smaller than* $1/2 - erf(|\mathbf{u}|)$ *which is the limit if we do not use collective measurements on the qubits. Here we have used this convention for the error function:* $erf(x) = \int_0^x e^{-t^2}/\sqrt{\pi}\,\mathrm{d}t$.

Proof. Let M be the optimal discrimination procedure $\phi^{\pm\mathbf{u}}$. Then we use the channel T_n to send $\rho_n^{\pm\mathbf{u}}$ to states of the oscillator and then perform the measurement M. By Theorem 7.1.1, $\|\phi^{\pm\mathbf{u}} - T_n(\rho_n^{\pm\mathbf{u}})\|_1 \to 0$ so that $\mathrm{Tr}\,(T_n(\rho_n^{\pm\mathbf{u}})M_{\mp}) \to \mathrm{Tr}\,(\phi^{\pm\mathbf{u}}M_{\mp})$. Thus $M \circ T_n$ is asymptotically optimal for $\rho_n^{\pm\mathbf{u}}$.

Now for pure states $|\psi_+\rangle$ and $|\psi_-\rangle$ the optimal measurement is well-known (Guţă et Kahn, 2009; Chefles, 2000). It is unique on the span of these pure states and arbitrary on the orthogonal. If we choose the phase such that $\langle\psi_-|\psi_+\rangle > 0$, then M_+ is the projector on the vector

$$\frac{|\psi_+\rangle + |\psi_-\rangle}{2\sqrt{1+\langle\psi_-|\psi_+\rangle}} + \frac{|\psi_+\rangle - |\psi_-\rangle}{2\sqrt{1-\langle\psi_-|\psi_+\rangle}},$$

and the associated risk is

$$\frac{1}{2}(1-\sqrt{1-|\langle\psi_+|\psi-\rangle|^2}).$$

Now in our case, in the limit experiment, $\phi^{\mathbf{u}}$ is the coherent state

$$|\psi_{\mathbf{u}}\rangle = e^{-|\mathbf{u}|^2/2}\sum_n |\mathbf{u}|^n/\sqrt{n!}\ |n\rangle.$$

So that

$$\langle\psi_{\mathbf{u}}|\psi_{-\mathbf{u}}\rangle = e^{-|\mathbf{u}|^2}\sum_n \frac{(-|\mathbf{u}|^2)^n}{n!} = e^{-2|\mathbf{u}|^2},$$

and $R(\phi^{\mathbf{u}}, \phi^{-\mathbf{u}}) = \frac{1}{2}\left(1-\sqrt{1-e^{-4|\mathbf{u}|^2}}\right)$.

We would like to insist here that the best measurement for discrimination is not measuring the positive part of the position observable $\mathbf{Q}$ (we assume by symmetry that $\pm\mathbf{u}$ is on the first coordinate), as one might expect from the analogy with the classical problem. Indeed if we measure Q then we obtain a classical Gaussian variable with density $p(x) = e^{-(x-|\mathbf{u}|)^2}/\sqrt{\pi}$ and the best guess at the sign $\pm$ has in this case the risk $1/2 - erf(|\mathbf{u}|)$.

This may be a bit surprising considering that measuring Q preserves the quantum Fisher information. The conclusion is simply that the quantum Fisher information is not an exhaustive indicator of the statistical information in a family of states, as it may remain unchanged even when there is a clear degradation in the inference power. This example fits in a more general framework of a theory of quantum statistical experiments and quantum decisions (Guţă).

□

7.7.4 Spin squeezed states and continuous time measurements

In an emblematic experiment for the field of quantum filtering and control, Geremia et al. (2004) have shown how spin squeezed states can be prepared deterministically by using continuous time measurements performed in the environment and real time feedback on the spins. Without going in the details, the basic idea is to describe the evolution of identically prepared spins by passing first to the coherent state picture. There one can easily solve the stochastic Schrödinger equation describing the evolution (quantum trajectory) of the quantum oscillator conditioned on the continuous signal of the measurement device. The solution is a Gaussian state whose center evolves stochastically while one of the quadratures gets more and more squeezed as one obtains more information through the measurement. Using feedback one can then stabilize the center of the state around a fixed point.

This description is of course approximative and holds as long as the errors in identifying the spins with Gaussian states are not significant. The framework developed in the proof of Theorem 7.1.1 can then be used to make more precise statements about the validity of the results, including the squeezing process.

Perhaps more interesting for quantum estimation, such measurements may be used to perform optimal estimation of spin states. The idea would be to first localize the state in a small region by performing a weak measurement and then in a second stage one performs a heterodyne type measurement after rotating the spins so that they point approximately in the z direction. We believe that this type of procedure has better chances of being implemented in practice compared with the abstract covariant measurement of Bagan et al. (2006); Hayashi et Matsumoto (2004).

Chapitre 8

Optimal estimation of qubit states with continuous time measurements

Ce chapitre dérive de l'article (Guţă et al., 2008).

Résumé : Nous proposons une stratégie adaptative, en deux temps, pour estimer l'état mélangé d'un qubit. Nous montrons l'optimalité de cette stratégie dans un sens minimax local, pour la distance de la trace ainsi que pour d'autres fonctions de coût quadratiques. L'optimalité minimax locale signifie qu'étant donnés n qubits identiques, il n'existe aucun estimateur qui fasse mieux que celui proposé sur un voisinage de taille $n^{-1/2}$ autour d'un état quelconque. En particulier, l'estimateur est asymptotiquement optimal au sens bayésien pour une grande classe de distributions *a priori*.

Nous proposons une implémentation physique de cette stratégie d'estimation optimale, basée sur les mesures en temps continu de champs couplés aux qubits.

L'ingrédient fondamental de ce résultat est le concept de normalité asymptotique locale (ou LAN) pour les qubits. Elle signifie que, pour de grands n, le modèle statistique qui décrit n qubits préparés de manière identique est localement équivalent à un modèle qui décrit une distribution gaussienne classique couplée à un état gaussien sur un oscillateur harmonique quantique.

Le terme «local» fait référence au voisinage qui rétrécit autour d'un état fixé ρ_0. Un résultat essentiel est que le rayon de ce voisinage peut être choisi aussi proche que l'on veut de $n^{-1/4}$. Ceci nous permet d'utiliser une procédure en deux temps, où nous commençons par localiser l'état

dans un petit voisinage de rayon $n^{1/2+\epsilon}$, puis utilisons LAN pour effectuer l'estimation optimale.

8.1 Introduction

State estimation is a central topic in quantum statistical inference (Holevo, 1982; Helstrom, 1976; Barndorff-Nielsen et al., 2003; Hayashi, 2005b). In broad terms the problem can be formulated as follows : given a quantum system prepared in an unknown state ρ, one would like to reconstruct the state by performing a measurement M whose random result X will be used to build an estimator $\hat{\rho}(X)$ of ρ. The quality of the measurement-estimator pair is given by the *risk*

$$R_\rho(M, \hat{\rho}) = \mathbb{E}\left(d(\hat{\rho}(X), \rho)^2\right), \tag{8.1}$$

where d is a distance on the space of states, for instance the fidelity distance or the trace norm, and the expectation is taken with respect to the probability distribution $\mathbb{P}_\rho^M$ of X, when the measured system is in state ρ. Since the risk depends on the unknown state ρ, one considers a global figure of merit by either averaging with respect to a prior distribution π (Bayesian setup)

$$R_\pi(M, \hat{\rho}) = \int \pi(d\rho) R_\rho(M, \hat{\rho}), \tag{8.2}$$

or by considering a maximum risk (pointwise or minimax setup)

$$R_{\max}(M, \hat{\rho}) = \sup_\rho R_\rho(M, \hat{\rho}). \tag{8.3}$$

An optimal procedure in either setup is one which achieves the minimum risk.

Typically, one measurement result does not provide enough information in order to significantly narrow down on the true state ρ. Moreover, if the measurement is "informative" then the state of the system after the measurement will contain little or no information about the initial state (Janssens, 2006) and one needs to repeat the preparation and measurement procedure in order to estimate the state with the desired accuracy.

It is then natural to consider a framework in which we are given a number n of identically prepared systems and look for estimators $\hat{\rho}_n$ which are optimal, or become optimal in the limit of large n. This problem is the quantum analogue of the classical statistical problem (van der Vaart, 1998) of estimating a parameter θ from independent identically distributed random variables $X_1, \ldots, X_n$ with distribution

$\mathbb{P}_\theta$, and some of the methods developed in this chapter are inspired by the classical theory.

Various state estimation problems have been investigated in the literature and the techniques may be quite different depending on a number of factors : the dimension of the density matrix, the number of unknown parameters, the purity of the states, and the complexity of measurements over which one optimizes. A short discussion on these issues can be found in section 8.2.

In this chapter we give an asymptotically optimal measurement strategy for qubit states that is based on the technique of *local asymptotic normality* introduced by Guţă et Kahn (2006); Guţă et Jenčová (2007). The technique is a quantum generalisation of Le Cam's (1986) classical statistical result, and builds on previous work of Hayashi et Matsumoto (2004). We use an adaptive two stage procedure involving continuous time measurements, which could in principle be implemented in practice. The idea of adaptive estimation methods, which has a long history in classical statistics, was introduced in the quantum set-up by Barndorff-Nielsen et Gill, R. (2000), and was subsequently used by Gill et Massar (2000); Hayashi (2002a); Hayashi et Matsumoto (2005). The aim there is similar : one wants to first localize the state and then to perform a suitably tailored measurement which performs optimally around a given state. A different adaptive technique was proposed independently by Nagaoka (2005) and further developed by Fujiwara (2006).

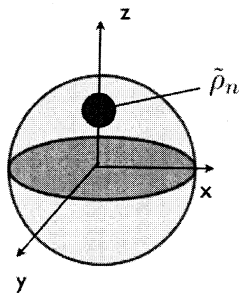

FIG. 8.1 – After the first measurement stage the state ρ lies in a small ball centered at $\tilde{\rho}_n$.

In the first stage, the spin components σ_x, σ_y and σ_z are measured separately on a small portion $\tilde{n} \ll n$ of the systems, and a rough estimator $\tilde{\rho}_n$ is constructed. By standard statistical arguments (see Lemma 8.2.1) we deduce that with high probability, the true state ρ lies within a ball of radius slightly larger than $n^{-1/2}$, say $n^{-1/2+\epsilon}$ with $\epsilon > 0$, centered at $\tilde{\rho}_n$. The purpose of the first stage is thus to localize the state within a small neighborhood as illustrated in Figure 8.1 (up to a unitary rotation) using the Bloch sphere representation of qubit states.

This information is then used in the second stage, which is a *joint* measurement on the remaining $n - \tilde{n}$ systems. This second measurement is implemented physically by two consecutive couplings, each to a bosonic field. The qubits are first coupled to the field via a spontaneous emission interaction and a continuous time heterodyne detection measurement is performed in the field. This yields information on the eigenvectors of ρ. Then the interaction is changed, and a continuous time homodyne detection is performed in the field. This yields information on the eigenvalues of ρ.

We prove that the second stage of the measurement is asymptotically optimal for all states in a ball of radius $n^{-1/2+\eta}$ around $\tilde{\rho}_n$. Here η can be chosen to be bigger that $\epsilon > 0$ implying that the two stage procedure as a whole is asymptotically optimal for any state as depicted in Figure 8.2.

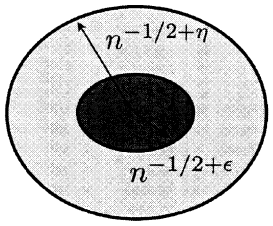

FIG. 8.2 – The smaller domain is the localization region of the first step. The second stage estimator is optimal for all states in the bigger domain.

The optimality of the second stage relies heavily on the principle of *local asymptotic normality* or LAN, see (van der Vaart, 1998), which we will briefly explain below, and in particular on the fact that it holds in a ball of radius $n^{-1/2+\eta}$ around $\tilde{\rho}_n$ rather than just $n^{-1/2}$ as it was the case in Guţă et Kahn's 2006 article.

Let ρ_0 be a fixed state. We parametrize the neighboring states as $\rho_{\mathbf{u}/\sqrt{n}}$, where $\mathbf{u} = (u_x, u_y, u_z) \in \mathbb{R}^3$ is a certain set of local parameters around ρ_0. Then LAN entails that the joint state $\rho_n^{\mathbf{u}} := \rho_{\mathbf{u}/\sqrt{n}}^{\otimes n}$ of n identical qubits converges for $n \to \infty$ to a Gaussian state of the form $N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}$, in a sense explained in Theorem 8.3.1. By $N^{\mathbf{u}}$ we denote a *classical* one-dimensional normal distribution centered at u_z. The second term $\phi^{\mathbf{u}}$ is a Gaussian state of a harmonic oscillator, i.e. a displaced thermal equilibrium state with displacement proportional to (u_x, u_y). We thus have the convergence

$$\rho_n^{\mathbf{u}} \rightsquigarrow N^{\mathbf{u}} \otimes \phi^{\mathbf{u}},$$

to a much simpler family of classical – quantum states for which we know how to optimally estimate the parameter $\mathbf{u}$ (Holevo, 1982; Yuen et Lax, M., 1973).

The idea of approximating a sequence of statistical experiments by a Gaussian one goes back to Wald (1943), and was subsequently developed by Le Cam (1986) who

coined the term local asymptotic normality. In quantum statistics the first ideas in the direction of local asymptotic normality for d-dimensional states appeared in a Japanese paper (Hayashi, 2003), as well as in Hayashi's conferences and were subsequently developed by Hayashi et Matsumoto (2004). In Theorem 8.3.1 we strengthen these results for the case of qubits, by proving a strong version of LAN in the spirit of Le Cam's pioneering work. We then exploit this result to prove optimality of the second stage. A different approach to local asymptotic normality has been developed by Guţă et Jenčová (2007) to which we refer for a more general exposition on the theory of quantum statistical models. A short discussion on the relation between the two approaches is given in the remark following Theorem 8.3.1.

From the physics perspective, our results put on a more rigorous basis the treatment of collective states of many identical spins, the keyword here being *coherent spin states* (Holtz et Hanus, 1974). Indeed, it has been known since Dyson (1956) that n spin-$\frac{1}{2}$ particles prepared in the spin up state $|\uparrow\rangle^{\otimes n}$ behave asymptotically as the ground state of a quantum oscillator, when considering the fluctuations of properly normalized total spin components in the directions orthogonal to z. We extend this to spin directions making an "angle" of order $n^{-1/2+\eta}$ with the z axis, as illustrated in Figure 8.3, as well as to mixed states. We believe that a similar approach can be followed in the case of spin squeezed states and continuous time measurements with feedback control (Geremia et al., 2004).

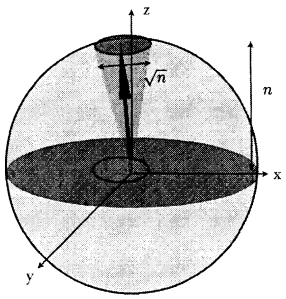

FIG. 8.3 – Total spin representation of the state of $n \gg 1$ spins : the quantum fluctuations of the x and y spin directions coincide with those of a coherent state of a harmonic oscillator.

In Theorem 8.4.1 we prove a dynamical version of LAN. The trajectory in time of the joint state of the qubits together with the field converges for large n to the corresponding trajectory of the joint state of the oscillator and field. In other words, time evolution preserves local asymptotic normality. This insures that for large n the state of the qubits "leaks" into a Gaussian state of the field, providing a concrete implementation of the convergence to the limit Gaussian experiment.

The punch line of the chapter is Theorem 8.6.1 which says that the estimator $\hat{\rho}_n$ is optimal in local minimax sense, which is the modern statistical formulation of optimality in the frequentist setup (van der Vaart, 1998). Also, its asymptotic risk is calculated explicitly.

The chapter is structured as follows : in section 8.2, we show that the first stage of the measurement sufficiently localizes the state. In section 8.3, we prove that LAN holds with radius of validity $n^{-1/2+\eta}$, and we bound its rate of convergence. sections 8.4 and 8.5 are concerned with the second stage of the measurement, i.e. with the coupling to the bosonic field and the continuous time field-measurements. Finally, in section 8.6, asymptotic optimality of the estimation scheme is proven.

The technical details of the proofs are relegated to the appendices in order to give the reader a more direct access to the ideas and results.

8.2 State estimation

In this section we introduce the reader to a few general aspects of quantum state estimation after which we concentrate on the qubit case.

State estimation is a generic name for a variety of results which may be classified according to the dimension of the parameter space, the kind or family of states to be estimated and the preferred estimation method. For an introduction to quantum statistical inference we refer to the books by Helstrom (1976) and Holevo (1982) and the more recent review paper by Barndorff-Nielsen et al. (2003). The collection (Hayashi, 2005b) is a good reference on quantum statistical problems, with many important contributions by the Japanese school.

For the purpose of this chapter, any quantum state representing a particular preparation of a quantum system, is described by a density matrix (positive selfadjoint operator of trace one) on the Hilbert space $\mathcal{H}$ associated to the system. The algebra of observables is $\mathcal{B}(\mathcal{H})$, and the expectation of an observable $a \in \mathcal{B}(\mathcal{H})$ with respect to the state ρ is $\mathrm{Tr}(\rho a)$. A measurement M with outcomes in a measure space $(\mathcal{X}, \Sigma)$ is completely determined by a σ-additive collection of positive selfadjoint operators $M(A)$ on $\mathcal{H}$, where A is an event in Σ. This collection is called a positive operator valued measure. The distribution of the results X when the system is in state ρ is given by $P_\rho(A) = \mathrm{Tr}(\rho M(A))$.

We are given n systems identically prepared in state ρ and we are allowed to perform a measurement M_n whose outcome is the estimator $\hat{\rho}_n$ as discussed in the Introduction.

The dimension of the density matrix may be finite, such as in the case of qubits or d-levels atoms, or infinite as in the case of the state of a monochromatic beam of light. In the finite or parametric case one expects that the risk converges to zero as n^{-1} and the optimal measurement-estimator sequence $(M_n, \hat{\rho}_n)$ achieves the *best constant* in front of the n^{-1} factor. In the non-parametric case the rates of convergence are in general slower that n^{-1} because one has to simultaneously estimate an infinite number of matrix elements, each with rate n^{-1}. An important example of such an estimation technique is that of quantum homodyne tomography in quantum optics (Vogel et Risken, H., 1989). This allows the estimation with arbitrary precision (D'Ariano et al., 1995; Leonhardt et al., 1995, 1996) of the whole density matrix of a monochromatic beam of light by repeatedly measuring a sufficiently large number of identically prepared beams (Smithey et al., 1993; Schiller et al., 1996; Zavatta et al., 2004). Artiles et al. (2005); Butucea et al. (2007) have shown how to formulate the problem of estimating infinite dimensional states without the need for choosing a cut-off in the dimension of the density matrix, and how to construct optimal minimax estimators of the Wigner function for a class of "smooth" states.

If we have some prior knowledge about the preparation procedure, we may encode this by parametrizing the possible states as $\rho = \rho_\theta$ with $\theta \in \Theta$ some unknown parameter. The problem is then to estimate θ optimally with respect to a distance function on Θ.

Indeed, one of the main problems in the finite dimensional case is to find optimal estimation procedures for a given family of states. It is known that if the state ρ is pure or belongs to a one parameter family, then separate measurements achieve the optimal rate of the class of joint measurements (Matsumoto, 2002). However for multi-dimensional families of mixed states this is no longer the case and joint measurements perform strictly better than separate ones (Gill et Massar, 2000).

In the Bayesian setup, one optimizes $R_\pi(M_n, \hat{\rho}_n)$ for some prior distribution π. We refer to (Jones, 1994; Massar et Popescu, 1995; Latorre et al., 1998; Fisher et al., 2000; Hannemann et al., 2002b; Bagan et al., 2002; Embacher et Narnhofer, 2004; Bagan et al., 2005) for the pure state case, and to (Cirac et al., 1999; Vidal et al., 1999; Mack et al., 2000; Keyl et Werner, 2001; Bagan et al., 2004c; Zyczkowski et Sommers, 2005; Bagan et al., 2006) for the mixed state case. The methods used here are based on group theory and can be applied only to invariant prior distributions and certain distance functions. In particular, the optimal covariant measurement in the case of completely unknown qubit states was found by Bagan et al. (2006) and Hayashi et Matsumoto (2004), but it has the drawback that it does not give any clue as to how it can be implemented in a real experiment.

In the pointwise approach (Hayashi, 2002a; Hayashi et Matsumoto, 2005; Gill et Massar, 2000; Barndorff-Nielsen et Gill, R., 2000; Fujiwara et Nagaoka, H., 1995;

Matsumoto, 2002; Barndorff-Nielsen et al., 2003; Hayashi et Matsumoto, 2004) one tries to minimize the risk for *each* unknown state ρ. As the optimal measurement-estimator pair cannot depend on the state itself, one optimizes the maximum risk $R_{\max}(M_n, \hat{\rho}_n)$, (see (8.3)), or a local version of this which will be defined shortly. The advantage of the pointwise approach is that it can be applied to arbitrary families of states and a large class of loss functions provided that they are locally quadratic in the chosen parameters. The underlying philosophy is that as the number n of states is sufficiently large, the problem ceases to be global and becomes a local one as the error in estimating the state parameters is of the order $n^{-1/2}$.

The Bayesian and pointwise approaches can be compared (Gill, 2005a), and in fact for large n the prior distribution π of the Bayesian approach becomes increasingly irrelevant and the optimal Bayesian estimator becomes asymptotically optimal in the minimax sense and vice versa.

8.2.1 Qubit state estimation : the localization principle

Let us now pass to the quantum statistical model which will be the object of our investigations. Let $\rho \in M_2(\mathbb{C})$ be an arbitrary density matrix describing the state of a qubit. Given n identically prepared qubits with joint state $\rho^{\otimes n}$, we would like to optimally estimate ρ based on the result of a properly chosen joint measurement M_n. For simplicity of the exposition we assume that the outcome of the measurement is an estimator $\hat{\rho}_n \in M_2(\mathbb{C})$. In practice however, the result X may belong to a complicated measure space (in our case the space of continuous time paths) and the estimator is a function of the "raw" data $\hat{\rho}_n := \hat{\rho}_n(X)$. The quality of the estimator at the state ρ is quantified by the risk

$$R_\rho(M_n, \hat{\rho}_n) := \mathbb{E}_\rho(d(\rho, \hat{\rho}_n)^2),$$

where d is a distance between states. The above expectation is taken with respect to the distribution $P_\rho(dx) := \mathrm{Tr}(\rho M(dx))$ of the measurement results, where $M(dx)$ represents the associated positive operator valued measure of the measurement M. In our exposition d will be the trace norm

$$\|\rho_1 - \rho_2\|_1 := \mathrm{Tr}(|\rho_1 - \rho_2|),$$

but similar results can be obtained using the fidelity distance. The aim is to find a sequence of measurements and estimators $(M_n, \hat{\rho}_n)$ which is asymptotically optimal in the *local minimax* sense : for any given ρ_0

$$\limsup_{n\to\infty} \sup_{\|\rho-\rho_0\|_1 \le n^{-1/2+\epsilon}} nR_\rho(M_n, \hat{\rho}_n) \le \limsup_{n\to\infty} \sup_{\|\rho-\rho_0\|_1 \le n^{-1/2+\epsilon}} nR_\rho(N_n, \check{\rho}_n),$$

for any other sequence of measurement-estimator pairs $(N_n, \check{\rho}_n)$. The factor n is inserted because typically $R_\rho(M_n, \hat{\rho}_n)$ is of the order $1/n$ and the optimization is about obtaining the smallest constant factor possible. The inequality says that one cannot find an estimator which performs better that $\hat{\rho}_n$ over a ball of size $n^{-1/2+\epsilon}$ centered at ρ_0, even if one has the knowledge that the state ρ belongs to that ball!

Here, and elsewhere in the chapter ϵ will appear in different contexts, as a generic strictly positive number and will be chosen to be sufficiently small for each specific use. At places where such notation may be confusing we will use additional symbols to denote small constants.

As set forth in the Introduction, our measurement procedure consists of two steps. The first one is to perform separate measurements of σ_x, σ_y and σ_z on a fraction $\tilde{n} = \tilde{n}(n)$ of the systems. In this way we obtain a rough estimate $\tilde{\rho}_n$ of the true state ρ which lies in a local neighborhood around ρ with high probability. The second step uses the information obtained in the first step to perform a measurement which is optimal precisely for the states in this local neighborhood. The second step ensures optimality and requires more sophisticated techniques inspired by the theory of local asymptotic normality for qubit states (Guţă et Kahn, 2006). We begin by showing that the first step amounts to the fact that, without loss of generality, we may assume that the unknown state is in a local neighborhood of a known state. This may serve also as an a posteriori justification of the definition of local minimax optimality.

Lemma 8.2.1. *Let M_i denote the measurement of the σ_i spin component of a qubit with $i = x, y, z$. We perform each of the measurements M_i separately on $\tilde{n}/3$ identically prepared qubits and define*

$$\tilde{\rho}_n = \frac{1}{2}(\mathbf{1} + \tilde{\mathbf{r}}\sigma), \qquad \text{if} \quad |\tilde{r}| \leq 1,$$

where $\tilde{\mathbf{r}} = (\tilde{r}_x, \tilde{r}_y, \tilde{r}_z)$ is the vector average of the measured components. If $|\tilde{r}| > 1$ then we define $\tilde{\rho}_n$ as the state which has the smallest trace distance to the right hand side expression. Then for all $\epsilon \in [0, 2]$, we have

$$\mathbb{P}\left(\|\tilde{\rho}_n - \rho\|_1^2 > 3n^{2\epsilon-1}\right) \leq 6\exp(-\tfrac{1}{2}\tilde{n}n^{2\epsilon-1}), \qquad \forall \rho.$$

Furthermore, for any $0 < \kappa < \epsilon/2$, if $\tilde{n} = n^{1-\kappa}$, the contribution to the risk $\mathbb{E}(\|\tilde{\rho}_n - \rho\|_1^2)$ brought by the event $E = [\,\|\tilde{\rho}_n - \rho\|_1 > \sqrt{3}n^{-1/2+\epsilon}\,]$ satisfies

$$\mathbb{E}\left(\|\tilde{\rho}_n - \rho\|_1^2 \chi_E\right) \leq 24\exp(-\tfrac{1}{2}n^{2\epsilon-\kappa}) = o(1).$$

Proof. For each spin component σ_i we obtain i.i.d coin tosses X_i with distribution $\mathbb{P}(X_i = \pm 1) = (1 \pm r_i)/2$ and average r_i.

Hoeffding's inequality (van der Vaart et Wellner, J.A., 1996) then states that for all $c > 0$, we have $\mathbb{P}(|X_i - \tilde{X}|^2 > c) \leq 2\exp(-\frac{1}{2}\tilde{n}c)$. By using this inequality three times with $c = n^{2\epsilon-1}$, once for each component, we get

$$\mathbb{P}\left(\sum_1^3 |\tilde{r}_i - r_i|^2 > 3n^{2\epsilon-1}\right) \leq 6\exp(-\tfrac{1}{2}\tilde{n}n^{2\epsilon-1}) \qquad \forall\rho,$$

which implies the statement for the norm distance since $\|\tilde{\rho}_n - \rho\|_1^2 = \sum_i |\tilde{r}_i - r_i|^2$. The bound on conditional risk follows from the previous bound and the fact that $\|\rho - \tilde{\rho}_n\|_1^2 \leq 4$.

□

In the second step of the measurement procedure we rotate the remaining $n - \tilde{n}$ qubits such that after rotation the vector $\tilde{r}$ is parallel to the z-axis. Afterwards, we couple the systems to the field and perform certain measurements in the field which will determine the final estimator $\hat{\rho}_n$. The details of this second step are given in sections 8.4 and 8.5, but at this moment we can already prove that the effect of errors in the the first stage of the measurement is asymptotically negligible compared to the risk of the second estimator. Indeed by Lemma 8.2.1 we get that if $\tilde{n} = n^{1-\kappa}$, then the probability that the first stage gives a "wrong" estimator (one which lies outside the local neighborhood of the true state) is of the order $\exp(-\frac{1}{2}n^{2\epsilon-\kappa})$ and so is the risk contribution. As the typical risk of estimation is of the order $1/n$, we see that the first step is practically "always" placing the estimator in a neighborhood of order $n^{-1/2+\epsilon}$ of the true state ρ, as shown in Figure 8.2. In the next section we will show that for such neighborhoods, the state of the remaining $n - \tilde{n}$ systems behaves asymptotically as a Gaussian state. This will allow us to devise an optimal measurement scheme for qubits based on the optimal measurement for Gaussian states.

8.3 Local asymptotic normality

The optimality of the second stage of the measurement relies on the concept of local asymptotic normality (van der Vaart, 1998; Guţă et Kahn, 2006). After a short introduction, we will prove that LAN holds for the qubit case, with radius of validity $n^{-1/2+\eta}$ for all $\eta \in [0, 1/4)$. We will also show that its rate of convergence is $O(n^{-1/4+\eta+\epsilon})$ for arbitrarily small ϵ.

8.3.1 Introduction to LAN and some definitions

Let ρ_0 be a fixed state, which by rotational symmetry can be chosen of the form

$$\rho_0 = \begin{pmatrix} \mu & 0 \\ 0 & 1-\mu \end{pmatrix}, \tag{8.4}$$

for a given $\frac{1}{2} < \mu < 1$. We parametrize the neighboring states as $\rho_{\mathbf{u}/\sqrt{n}}$ where $\mathbf{u} = (u_x, u_y, u_z) \in \mathbb{R}^3$ such that the first two components account for unitary rotations around ρ_0, while the third one describes the change in eigenvalues

$$\rho_{\mathbf{v}} := U(\mathbf{v}) \begin{pmatrix} \mu + v_z & 0 \\ 0 & 1-\mu-v_z \end{pmatrix} U(\mathbf{v})^*, \tag{8.5}$$

with unitary $U(\mathbf{v}) := \exp(i(v_x\sigma_x + v_y\sigma_y))$. The "local parameter" $\mathbf{u}$ should be thought of, as having a bounded range in $\mathbb{R}^3$ or may even "grow slowly" as $\|\mathbf{u}\| \leq n^\eta$.

Then, for large n, the joint state $\rho_n^{\mathbf{u}} := \rho_{\mathbf{u}/\sqrt{n}}^{\otimes n}$ of n identical qubits approaches a Gaussian state of the form $N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}$ with the parameter $\mathbf{u}$ appearing solely in the average of the two Gaussians. By $N^{\mathbf{u}}$ we denote a *classical* one-dimensional normal distribution centered at u_z which relays information about the eigenvalues of $\rho_{\mathbf{u}/\sqrt{n}}$. The second term $\phi^{\mathbf{u}}$ is a Gaussian state of a harmonic oscillator which is a displaced thermal equilibrium state with displacement proportional to (u_x, u_y). It contains information on the eigenvectors of $\rho_{\mathbf{u}/\sqrt{n}}$. We thus have the convergence

$$\rho_n^{\mathbf{u}} \rightsquigarrow N^{\mathbf{u}} \otimes \phi^{\mathbf{u}},$$

to a much simpler family of classical - quantum states for which we know how to optimally estimate the parameter $\mathbf{u}$. The asymptotic splitting into a classical estimation problem for eigenvalues and a quantum one for the eigenbasis has been also noticed by Bagan et al. (2006) and by Hayashi et Matsumoto (2004), the latter coming pretty close to our formulation of local asymptotic normality.

The precise meaning of the convergence is given in Theorem 8.3.1 below. In short, there exist quantum channels T_n which map the states $\rho_{\mathbf{u}/\sqrt{n}}^{\otimes n}$ into $N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}$ with vanishing error in trace norm distance, and uniformly over the local parameters $\mathbf{u}$. From the statistical point of view the convergence implies that a statistical decision problem concerning the model $\rho_n^{\mathbf{u}}$ can be mapped into a similar problem for the model $N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}$ such that the optimal solution for the latter can be translated into an asymptotically optimal solution for the former. In our case the problem of estimating the state ρ turns into that of estimating the local parameter $\mathbf{u}$ around the first stage estimator $\tilde{\rho}_n$ playing the role of ρ_0. For the family of displaced Gaussian states it is well known that the optimal estimation of the displacement is achieved by the heterodyne detection (Holevo, 1982; Yuen et Lax, M., 1973), while for the classical part it sufficient to take the observation as best estimator. Hence the second step will give an optimal estimator $\hat{\mathbf{u}}$ of $\mathbf{u}$ and an optimal estimator of the initial state $\hat{\rho}_n := \rho_{\hat{\mathbf{u}}/\sqrt{n}}$. The precise result is formulated in Theorem 8.6.1

8.3.2 Convergence to the Gaussian model

We describe the state $N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}$ in more detail. $N^{\mathbf{u}}$ is simply the classical Gaussian distribution

$$N^{\mathbf{u}} := N(u_z, \mu(1-\mu)), \tag{8.6}$$

with mean u_z and variance $\mu(1-\mu)$.

The state $\phi^{\mathbf{u}}$ is a density matrix on $\mathcal{H} = \mathcal{F}(\mathbb{C})$, the representation space of the harmonic oscillator. In general, for any Hilbert space $\mathfrak{h}$, the *Fock space* over $\mathfrak{h}$ is defined as

$$\mathcal{F}(\mathfrak{h}) := \bigoplus_{n=0}^{\infty} \mathfrak{h} \otimes_s \cdots \otimes_s \mathfrak{h}, \tag{8.7}$$

with $\otimes_s$ denoting the symmetric tensor product. Thus $\mathcal{F}(\mathbb{C})$ is the simplest example of a Fock space. Let

$$\phi := (1-p) \sum_{k=0} p^k |k\rangle\langle k|, \tag{8.8}$$

be a thermal equilibrium state with $|k\rangle$ denoting the k-th energy level of the oscillator and $p = \frac{1-\mu}{\mu} < 1$. For every $\alpha \in \mathbb{C}$ define the displaced thermal state

$$\phi(\alpha) := D(\alpha)\, \phi\, D(-\alpha),$$

where $D(\alpha) := \exp(\alpha a^* - \bar{\alpha} a)$ is the displacement operator, mapping the vacuum vector $|0\rangle$ to the coherent vector

$$|\alpha\rangle = \exp(-\alpha^2/2) \sum_{k=0}^{\infty} \frac{\alpha^k}{\sqrt{k!}} |k\rangle.$$

Here a^* and a are the creation and annihilation operators on $\mathcal{F}(\mathbb{C})$, satisfying $[a, a^*] = \mathbf{1}$. The family $\phi^{\mathbf{u}}$ of states in which we are interested is given by

$$\phi^{\mathbf{u}} := \phi(\sqrt{2\mu - 1}\alpha_{\mathbf{u}}), \qquad \mathbf{u} \in \mathbb{R}^3, \tag{8.9}$$

with $\alpha_{\mathbf{u}} := -u_y + i u_x$. Note that $\phi^{\mathbf{u}}$ does not depend on u_z.

We claim that the "statistical information" contained in the joint state of n qubits

$$\rho_n^{\mathbf{u}} := \rho_{\mathbf{u}/\sqrt{n}}^{\otimes n}, \tag{8.10}$$

is asymptotically identical to that contained in the couple $(N^{\mathbf{u}}, \phi^{\mathbf{u}})$. More precisely :

Theorem 8.3.1. *Let $\rho_n^{\mathbf{u}}$ be the family of states* (8.5) *on the Hilbert space $(\mathbb{C}^2)^{\otimes n}$, let $N^{\mathbf{u}}$ be the family* (8.6) *of Gaussian distributions, and let $\phi^{\mathbf{u}}$ be the family* (8.9) *of*

displaced thermal equilibrium states of a quantum oscillator. Then for each n there exist quantum channels (trace preserving CP maps)

$$T_n : \mathcal{T}((\mathbb{C}^2)^{\otimes n}) \to L^1(\mathbb{R}) \otimes \mathcal{T}(\mathcal{F}(\mathbb{C})),$$
$$S_n : L^1(\mathbb{R}) \otimes \mathcal{T}(\mathcal{F}(\mathbb{C})) \to \mathcal{T}((\mathbb{C}^2)^{\otimes n})$$

with $\mathcal{T}(\mathcal{H})$ the trace-class operators on $\mathcal{H}$, such that, for any $0 \leq \eta < 1/4$ and any $\epsilon > 0$,

$$\sup_{\|\mathbf{u}\| \leq n^\eta} \| N^{\mathbf{u}} \otimes \phi^{\mathbf{u}} - T_n(\rho_n^{\mathbf{u}}) \|_1 = O(n^{-1/4+\eta+\epsilon}), \tag{8.11}$$

$$\sup_{\|\mathbf{u}\| \leq n^\eta} \| \rho_n^{\mathbf{u}} - S_n(N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}) \|_1 = O(n^{-1/4+\eta+\epsilon}). \tag{8.12}$$

Moreover, for each $\epsilon_2 > 0$ there exists a function $f(n)$ of order $O(n^{-1/4+\eta+\epsilon})$ such that the above convergence rates are bounded by $f(n)$, with f independent of $\rho^{\mathbf{0}}$ as long as $|\frac{1}{2} - \mu| > \epsilon_2$.

Remark. Note that the equations (8.11) and (8.12) imply that the expressions on the left side converge to zero as $n \to \infty$. Following the classical terminology of Le Cam (1986), we will call this type of result *strong convergence* of quantum statistical models (experiments). Another local asymptotic normality result has been derived by Guţă et Jenčová (2007) based on a different concept of convergence, which is an extension of the *weak convergence* of classical (commutative) statistical experiments. In the classical set-up it is known that strong convergence implies weak convergence for arbitrary statistical models, and the two are equivalent for statistical models consisting of a finite number of distributions.

These two approaches to local asymptotic normality in quantum statistics are based on completely different methods and the results are complementary in the sense that the weak convergence of Guţă et Jenčová (2007) holds for the larger class of finite dimensional states while the strong convergence has more direct consequences as it is shown in this chapter for the case of qubits. Both results are part of a larger effort to develop a general theory of local asymptotic normality in quantum statistics. Several extensions are in order : from qubits to arbitrary finite dimensional systems (strong convergence), from finite dimensional to continuous variables systems, from identical system to correlated ones, and asymptotic normality in continuous time dynamical set-up.

Finally, let us note that the development of a general theory of convergence of quantum statistical models will set a framework for dealing with other important statistical decision problems such as quantum cloning (Werner, 1998) and quantum amplification (Caves, 1982), which do not necessarily involve measurements.

Remark. The construction of the channels T_n, S_n in the case of fixed eigenvalues ($u_z = 0$) is given in Theorem 1.1 of Guţă et Kahn (2006). It is also shown that a similar result holds uniformly over $\|\mathbf{u}\| < C$ for any fixed finite constant C. Guţă et Jenčová (2007) have shown that weak convergence also holds in the general case, with unknown eigenvalues. A classical component then appears in the limit statistical experiment. In the above result we extend the domain of validity of these Theorems from "local" parameters $\|\mathbf{u}\| < C$ to "slowly growing" local neighborhoods $\|\mathbf{u}\| \leq n^\eta$ with $\eta < 1/4$. Although this may be seen as merely a technical improvement, it is in fact essential in order to insure that the result of the first step of the estimation will, with high probability, fall inside a neighborhood $\|\mathbf{u}\| \leq n^\eta$ for which local asymptotic normality still holds (see Figure 8.2).

Proof. Following (Guţă et Kahn, 2006) we will first indicate how the channels T_n are constructed. The technical details of the proof can be found in Appendix 8.A.

The space $(\mathbb{C}^2)^{\otimes n}$ carries two unitary representations. 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

$$\tilde{\pi}_n(\tau) : v_1 \otimes \cdots \otimes v_n \mapsto v_{\tau^{-1}(1)} \otimes \cdots \otimes v_{\tau^{-1}(n)}, \qquad \tau \in S(n).$$

As $[\pi_n(u), \tilde{\pi}_n(\tau)] = 0$ for all $u \in SU(2), \tau \in S(n)$, we have the decomposition

$$\left(\mathbb{C}^2\right)^{\otimes n} = \bigoplus_{j=0,1/2}^{n/2} \mathcal{H}_j \otimes \mathcal{H}_n^j. \tag{8.13}$$

The direct sum runs over all positive (half)-integers j up to $n/2$. For each fixed j, $\mathcal{H}_j \cong \mathbb{C}^{2j+1}$ is an irreducible representation U_j of $SU(2)$ with total angular momentum $J^2 = j(j+1)$, and $\mathcal{H}_n^j \cong \mathbb{C}^{n_j}$ is the irreducible representation of the symmetric group $S(n)$ with $n_j = \binom{n}{n/2-j} - \binom{n}{n/2-j-1}$. The density matrix $\rho_n^{\mathbf{u}}$ is invariant under permutations and can be decomposed as a mixture of "block" density matrices

$$\rho_n^{\mathbf{u}} = \bigoplus_{j=0,1/2}^{n/2} p_{n,\mathbf{u}}(j)\, \rho_{j,n}^{\mathbf{u}} \otimes \frac{\mathbf{1}}{n_j}. \tag{8.14}$$

The probability distribution $p_{n,\mathbf{u}}(j)$ is given by (Bagan et al., 2006) :

$$p_{n,\mathbf{u}}(j) := \frac{n_j}{2\mu_{\mathbf{u}} - 1} \left(1 - \mu_{\mathbf{u}}\right)^{\frac{n}{2}-j} \mu_{\mathbf{u}}^{\frac{n}{2}+j+1} \left(1 - p_{\mathbf{u}}^{2j+1}\right), \tag{8.15}$$

with $\mu_{\mathbf{u}} := \mu + u_z/\sqrt{n}$, $p_{\mathbf{u}} := \frac{1-\mu_{\mathbf{u}}}{\mu_{\mathbf{u}}}$. We can rewrite $p_{n,\mathbf{u}}(j)$ as

$$p_{n,\mathbf{u}}(j) := B_{n,\mu_{\mathbf{u}}}(n/2 + j) \times K(j, n, \mu, \mathbf{u}), \tag{8.16}$$

where

$$B_{n,\nu}(k) := \binom{n}{k} \nu^k (1-\nu)^{n-k}, \qquad k = 0, \ldots, n$$

is a binomial distribution, and the factor $K(j, n, \mu, \mathbf{u})$ is given by

$$K(j, n, \mu, \mathbf{u}) := \left(1 - p_{\mathbf{u}}^{2j+1}\right) \frac{n + (2(j - j_n - \sqrt{n}u_z) + 1)/(2\mu_{\mathbf{u}} - 1)}{n + (j - j_n - \sqrt{n}u_z + 1)/\mu_{\mathbf{u}}},$$

for $j_n := n(\mu - 1/2)$.

Now $K(j, n, \mu, \mathbf{u}) = 1 + O(n^{-1/2+\epsilon})$ on the relevant values of j, i.e. the ones in an interval of order $n^{1/2+\epsilon}$ around j_n, as long as $\mu_{\mathbf{u}}$ is bounded away from $1/2$, which is automatically so for big n. As $B_{n,\mu_{\mathbf{u}}}(k)$ is the distribution of a sum of i.i.d. Bernoulli random variables, we can now use standard local asymptotic normality results (van der Vaart, 1998) to conclude that if j is distributed according to $p_{n,\mathbf{u}}$, then the centered and rescaled variable

$$g_n := \frac{j}{\sqrt{n}} - \sqrt{n}(\mu - 1/2),$$

converges in distribution to a normal $N^{\mathbf{u}}$, after an additional randomization has been performed. The latter is necessary in order to "smooth" the discrete distribution into a distribution which is continuous with respect to the Lebesgue measure, and will convergence to the Gaussian distribution in total variation norm.

The measurement "which block", corresponding to the decomposition (8.14), provides us with a result j and a posterior state $\rho_{j,n}^{\mathbf{u}}$. The function $g_n = g_n(j)$ (with an additional randomization) is the classical part of the channel T_n. The randomization consists of "smoothening" with a Gaussian kernel of mean $g_n(j)$ and variance $1/(2\sqrt{n})$, i.e. with $\tau_{n,j} := (n^{1/4}/\sqrt{\pi}) \exp\left(-\sqrt{n}(x - g_n(j))^2\right)$.

Note that this measurement is not disturbing the state $\rho_n^{\mathbf{u}}$ in the sense that the average state after the measurement is the same as before.

The quantum part of T_n is the same as in (Guţă et Kahn, 2006) and consists of embedding each block state $\rho_{j,n}^{\mathbf{u}}$ into the state space of the oscillator by means of an isometry $V_j : \mathcal{H}_j \to \mathcal{F}(\mathbb{C})$,

$$V_j : |j, m\rangle \mapsto |j - m\rangle,$$

where $\{|j, m\rangle : m = -j, \ldots, j\}$ is the eigenbasis of the total spin component $L_z := \sum_i \sigma_z^{(i)}$, cf. equation (5.1) of (Guţă et Kahn, 2006). Then the action of the channel T_n is

$$T_n : \bigoplus_j p_{n,\mathbf{u}}(j) \rho_{j,n}^{\mathbf{u}} \otimes \frac{\mathbf{1}}{n_j} \mapsto \sum_j p_{n,\mathbf{u}}(j)\, \tau_{n,j} \otimes V_j \rho_{j,n}^{\mathbf{u}} V_j^* .$$

The inverse channel S_n performs the inverse operation with respect to T_n. First the oscillator state is "cut-off" to the dimension of an irreducible representation and then a block obtained in this way is placed into the decomposition (8.13) (with an additional normalization from the remaining infinite dimensional block which is negligible for the states in which we are interested).

The rest of the proof is given in Appendix 8.A.

□

8.4 Time evolution of the interacting system

In the previous section, we have investigated the asymptotic equivalence between the states $\rho_n^{\mathbf{u}}$ and $N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}$ by means of the channel T_n. We now seek to implement this in a physical situation. The $N^{\mathbf{u}}$-part will follow in section 8.5.2, the $\phi^{\mathbf{u}}$-part will be treated in this section.

We couple the n qubits to a Bosonic field; this is the physical implementation of LAN. Subsequently, we perform a measurement in the field which will provide the information about the state of the qubits; this is the utilization of LAN in order to solve the asymptotic state estimation problem.

In this section we will limit ourselves to analyzing the joint evolution of the qubits and field. The measurement on the field is described in section 8.5.

8.4.1 Quantum stochastic differential equations

In the weak coupling limit (Gardiner et Zoller, 2004) the joint evolution of the qubits and field can be described mathematically by quantum stochastic differential equations (QSDE) (Hudson et Parthasarathy, 1984). The basic notions here are the Fock space, the creation and annihilation operators and the quantum stochastic differential equation of the unitary evolution. The Hilbert space of the field is the Fock space $\mathcal{F}(L^2(\mathbb{R}))$ as defined in (8.7). An important linearly complete set in $\mathcal{F}(L^2(\mathbb{R}))$ is that of the exponential vectors

$$e(f) := \bigoplus_{n=0}^{\infty} \frac{1}{\sqrt{n!}} f^{\otimes n} := \bigoplus_{n=0}^{\infty} \frac{1}{\sqrt{n!}} |f\rangle_n, \qquad f \in L^2(\mathbb{R}), \tag{8.17}$$

with inner product $\langle e(f), e(g)\rangle = \exp(\langle f, g\rangle)$. The normalized exponential states $|f\rangle := e^{-\langle f,f\rangle/2} e(f)$ are called coherent states. The vacuum vector is $|\Omega\rangle := e(0)$ and

we will denote the corresponding density matrix $|\Omega\rangle\langle\Omega|$ by Φ. The quantum noises are described by the creation and annihilation martingale operators $A_t^* := a^*(\chi_{[0,t]})$ and $A_t := a(\chi_{[0,t]})$ respectively, where $\chi_{[0,t]}$ is the indicator function for $[0,t]$ and

$$a(f) : e(g) \mapsto \langle f, g\rangle e(g).$$

The increments $dA_t := a(\chi_{[0,t+dt]}) - a(\chi_{[0,t]})$ and dA_t^* play the role of non-commuting integrators in quantum stochastic differential equations, in the same way as the one can integrate against the Brownian motion in classical stochastic calculus.

We now consider the joint unitary evolution for qubits and field defined by the quantum stochastic differential equation (Hudson et Parthasarathy, 1984; Bouten et al., 2004) :

$$dU_n(t) = (a_n dA_t^* - a_n^* dA_t - \frac{1}{2} a_n^* a_n dt) U_n(t),$$

where $U_n(t)$ is a unitary operator on $(\mathbb{C}^2)^{\otimes n} \otimes \mathcal{F}(L^2(\mathbb{R}))$, and

$$a_n := \frac{1}{\sqrt{2j_n}} \sum_{k=1}^{n} \sigma_+^{(k)}, \qquad \sigma_+^{(k)} := \mathbf{1} \otimes \cdots \otimes (\sigma_x + i\sigma_y)/2 \otimes \cdots \otimes \mathbf{1}, \quad j_n := (\mu - 1/2)n.$$

As we will see later, the "coupling factor" $1/\sqrt{j_n}$ of the order $n^{-1/2}$, is necessary in order to obtain convergence to the unitary evolution of the quantum harmonic oscillator and the field.

We remind the reader that the n-qubit space can be decomposed into irreducible representations as in (8.13), and the interaction between the qubits and field respects this decomposition

$$U_n(t) = \bigoplus_{j=0,1/2}^{n/2} U_{j,n}(t) \otimes \mathbf{1},$$

where $\mathbf{1}$ is the identity operator on the multiplicity space $\mathcal{H}_n^j$, and

$$U_{j,n}(t) : \mathcal{H}_j \otimes \mathcal{F}(L^2(\mathbb{R})) \to \mathcal{H}_j \otimes \mathcal{F}(L^2(\mathbb{R})),$$

is the restricted cocycle

$$dU_{j,n}(t) = (a_j dA_t^* - a_j^* dA_t - \frac{1}{2} a_j^* a_j dt) U_{j,n}(t), \tag{8.18}$$

with a_j acting on the basis $|j,m\rangle$ of $\mathcal{H}_j$ as

$$\begin{aligned} a_j|j,m\rangle &= \sqrt{j-m}\sqrt{(j+m+1)/2j_n}\,|j,m+1\rangle, \\ a_j^*|j,m\rangle &= \sqrt{j-m+1}\sqrt{j+m/2j_n}\,|j,m-1\rangle. \end{aligned}$$

Remark. We point out that the *lowering* operator for L_z acts as *creator* for our cut-off oscillator since the highest vector $|j,j\rangle$ corresponds by V_j to the vacuum of the oscillator. This choice does not have any physical meaning but is only related with our convention $\mu > 1/2$. Had we chosen $\mu < 1/2$, then the raising operator on the qubits would correspond to creation operator on the oscillator.

By (8.14) the initial state $\rho^{\otimes n}$ decomposes in the same way as the unitary cocycle, and thus the whole evolution decouples into separate "blocks" for each value of j. We do not have explicit solutions to these equations but based on the conclusions drawn from LAN we expect that as $n \to \infty$, the solutions will be well approximated by similar ones for a coupling between an oscillator and the field, at least for the states in which we are interested. As a warm up exercise we will start with this simpler limit case where the states can be calculated explicitly.

8.4.2 Solving the QSDE for the oscillator

Let a^* and a be the creation and annihilation operators of a quantum oscillator acting on $\mathcal{F}(\mathbb{C})$. We couple the oscillator with the Bosonic field and the joint unitary evolution is described by the family of unitary operators $U(t)$ satisfying the quantum stochastic differential equation

$$dU(t) = (adA_t^* - a^* dA_t - \frac{1}{2}a^* a dt)U(t).$$

We choose the initial (un-normalized) state $\psi(0) := e(\mathbf{z}) \otimes |\Omega\rangle$, where $\mathbf{z}$ is any complex number, and we shall find the explicit form of the vector state of the system and field at time t : $\psi(t) := U(t)\psi(0)$.

We make the following ansatz : $\psi(t) = e(\alpha_t) \otimes e(f_t)$, where $f_t(s) := f(s)\chi_{[0,t]}(s)$ for some $f \in L^2(\mathbb{R})$. For each $\beta \in \mathbb{C}$, $g \in L^2(\mathbb{R})$, define $I(t) := \langle e(\beta) \otimes e(g), \psi(t)\rangle$. We then have $I(t) = \exp(\bar{\beta}\alpha(t) + \langle g, f_t\rangle)$, so that it satisfies

$$dI(t) = \left(\bar{\beta}\tfrac{d}{dt}\alpha(t) + \bar{g}(t)f(t)\right) I(t)dt\,. \tag{8.19}$$

We now calculate $\frac{d}{dt}I(t)$ with the help of the QSDE. Since $A_t e(f) = \langle \chi_{[0,t]}, f\rangle e(f)$, we have, for continuous g, $dA_t e(g) = g(t)e(g)dt$. However, since $A_s e(f_t)$ is constant for $s \geq t$, we have $dA_t e(f_t) = 0$. Thus

$$dI(t) = \langle e(\beta) \otimes e(g), (adA_t^* - a^* dA_t - \tfrac{1}{2}a^* a dt)\psi(t)\rangle = (\bar{g}(t)\alpha(t) - \tfrac{1}{2}\bar{\beta}\alpha(t))I(t)dt\,. \tag{8.20}$$

Equating (8.19) with (8.20) for all t, β and continuous g, we find $f(s) = \alpha(s)$, $\frac{d}{dt}\alpha(t) = -\frac{1}{2}\alpha(t)$. Thus $\alpha(t) = \alpha(0)e^{-\frac{1}{2}t}$, $f_t(s) = \alpha(0)\chi_{[0,t]}(s)e^{-\frac{1}{2}s}$ with $\alpha(0) = \mathbf{z}$.

In conclusion $\psi(t) = e(\mathbf{z}e^{-\frac{1}{2}t}) \otimes e(\mathbf{z}e^{-\frac{1}{2}s}\chi_{[0,t]}(s))$. For later use we denote the *normalized* solution by $\psi_{\mathbf{z}}(t) := U(t)|\mathbf{z}\rangle \otimes |\Omega\rangle = e^{-|\mathbf{z}|^2/2}U(t)e(\mathbf{z}) \otimes |\Omega\rangle$.

8.4.3 QSDE for large spin

We consider now the unitary evolution for qubits and field :

$$dU_n(t) = (a_n dA_t^* - a_n^* dA_t - \frac{1}{2}a_n^* a_n dt)U_n(t).$$

It is no longer possible to obtain an explicit expression for the joint vector state $\psi_n(t)$ at time t. However we will show that for the states in which we are interested, a satisfactory explicit *approximate* solution exists.

The trick works for an arbitrary family of unitary solutions of a quantum stochastic differential equation $dU(t) = G_{dt}U(t)$, and the general idea is the following : if $\psi(t)$ is the true state $\psi(t) = U(t)\psi$ and $\xi(t)$ is a vector describing an approximate evolution ($\psi(0) = \xi(0)$) then with $U_{t+dt}^t := U(t+dt)U(t)^{-1}$ we get

$$\begin{aligned}
\psi(t+dt) - \xi(t+dt) &= \psi(t+dt) - U_{t+dt}^t\xi(t) + U_{t+dt}^t\xi(t) \\
&\quad -\xi(t) + \xi(t) - \xi(t+dt) \\
&= U_{t+dt}^t[\psi(t) - \xi(t)] + [U(t+dt) - U(t)]U(t)^{-1}\xi(t) \\
&\quad +[\xi(t) - \xi(t+dt)] \\
&= U_{t+dt}^t[\psi(t) - \xi(t)] + G_{dt}\xi(t) - d\xi(t).
\end{aligned}$$

By taking norms we get

$$d\|\psi(t) - \xi(t)\| \leq \|G_{dt}\xi(t) - d\xi(t)\|. \tag{8.21}$$

The idea is now to devise a family $\xi(t)$ such that the right side is as small as possible.

We apply this technique block-wise, that is to each unitary $U_{j,n}(t)$ acting on $\mathcal{H}_j \otimes \mathcal{F}(L^2(\mathbb{R}))$ (see equation (8.18)) for a "typical" $j \in \mathcal{J}_n$ (see equation (8.39)). By means of the isometry V_j we can embed the space $\mathcal{H}_j$ into the first $2j+1$ levels of the oscillator and for simplicity we will keep the same notions as before for the operators acting on $\mathcal{F}(\mathbb{C})$. As initial states for the qubits we choose the block states $\rho_{j,n}^{\mathbf{u}}$.

Theorem 8.4.1. *Let $\rho_{j,n}^{\mathbf{u}}(t) = U_{j,n}(t)\left[\rho_{j,n}^{\mathbf{u}} \otimes \Phi\right]U_{j,n}^*(t)$ be the j-th block of the state of qubits and field at time t. Let $\phi^{\mathbf{u}}(t) := U(t)\left[\phi^{\mathbf{u}} \otimes \Phi\right]U(t)^*$ be the joint state of the oscillator and field at time t. For any $\eta < 1/6$, for any $\epsilon > 0$,*

$$\sup_{j\in\mathcal{J}_n}\sup_{\|\mathbf{u}\|\leq n^\eta}\sup_t \|\rho_{j,n}^{\mathbf{u}}(t) - \phi^{\mathbf{u}}(t)\|_1 = O(n^{-1/4+\eta+\epsilon}, n^{-1/2+3\eta+\epsilon}). \tag{8.22}$$

Proof. From the proof of the local asymptotic normality Theorem 8.3.1 we know that the initial states of the two unitary evolutions are asymptotically close to each other

$$\sup_{j\in\mathcal{J}_n}\sup_{\|\mathbf{u}\|\leq n^\eta}\|\rho_{j,n}^{\mathbf{u}}-\phi^{\mathbf{u}}\|_1=O(n^{-1/4+\eta+\epsilon}). \tag{8.23}$$

The proof consists of two estimation steps. In the first one, we will devise another initial state $\tilde{\rho}_{j,n}^{\mathbf{u}}$ which is an approximation of $\phi^{\mathbf{u}}$ and thus also of $\rho_{j,n}^{\mathbf{u}}$:

$$\sup_{j\in\mathcal{J}_n}\sup_{\|\mathbf{u}\|\leq n^\eta}\|\tilde{\rho}_{j,n}^{\mathbf{u}}-\phi^{\mathbf{u}}\|_1=O(e^{-n^\epsilon}). \tag{8.24}$$

In the second estimate we show that the evolved states $\tilde{\rho}_{j,n}^{\mathbf{u}}(t)$ and $\phi^{\mathbf{u}}(t)$ are asymptotically close to each other

$$\sup_{j\in\mathcal{J}_n}\sup_{\|\mathbf{u}\|\leq n^\eta}\sup_t\|\tilde{\rho}_{j,n}^{\mathbf{u}}(t)-\phi^{\mathbf{u}}(t)\|_1=O(n^{-1/4+\eta+\epsilon},n^{-1/2+3\eta+\epsilon}). \tag{8.25}$$

This estimate is important because, the two trajectories are driven by different Hamiltonians, and in principle there is no reason why they should stay close to each other.

From (8.23), (8.24) and (8.25), and using triangle inequality we get

$$\sup_{j\in\mathcal{J}_n}\sup_{\|\mathbf{u}\|\leq n^\eta}\sup_t\|\rho_{j,n}^{\mathbf{u}}(t)-\phi^{\mathbf{u}}(t)\|_1=O(n^{-1/4+\eta+\epsilon},n^{-1/2+3\eta+\epsilon}).$$

The following diagram illustrates the above estimates. The upper line concerns the time evolution of the block state $\rho_{j,n}^{\mathbf{u}}$ and the field. The lower line describes the time evolution of the oscillator and the field. The estimates show that the diagram is "asymptotically commutative" for large n.

$$\begin{array}{ccccc}
\mathcal{S}(\mathcal{H}_j) & \xrightarrow{\mathrm{Id}_j\otimes\Phi} & \mathcal{S}(\mathcal{H}_j\otimes\mathcal{F}) & \xrightarrow{U_{j,n}(t)} & \mathcal{S}(\mathcal{H}_j\otimes\mathcal{F}) \\
{\scriptstyle V_j\cdot V_j^*}\downarrow & & \downarrow & & \downarrow \\
\mathcal{S}(\mathcal{F}(\mathbb{C})) & \xrightarrow{\mathrm{Id}\otimes\Phi} & \mathcal{S}(\mathcal{F}(\mathbb{C})\otimes\mathcal{F}) & \xrightarrow{U(t)} & \mathcal{S}(\mathcal{F}(\mathbb{C})\otimes\mathcal{F})
\end{array}$$

For the rest of the proof, we refer to Appendix 8.B.

□

We have shown how the mathematical statement of LAN (the joint state of qubits converges to a Gaussian state of a quantum oscillator plus a classical Gaussian random variable) can in fact be physically implemented by coupling the spins to the environment and letting them "leak" into the field. In the next section, we will use this for the specific purpose of estimating $\mathbf{u}$ by performing a measurement in the field.

8.5 The second stage measurement

We now describe the second stage of our measurement procedure. Recall that in the first stage a relatively small part $\tilde{n} = n^{1-\kappa}, 1 > \kappa > 0$, of the qubits is measured and a rough estimator $\tilde{\rho}_n$ is obtained. The purpose of this estimator is to localize the state within a small neighborhood such that the machinery of local asymptotic normality of Theorem 8.3.1 can be applied.

In Theorem 8.4.1 the local asymptotic normality was extended to the level of time evolution of the qubits interacting with a bosonic field. We have proven that at time t the joint state of the qubits and field is

$$\begin{aligned}\rho_n^{\mathbf{u}}(t) \;:=\; & \bigoplus_{j=0,1/2}^{n/2} p_{n,\mathbf{u}}(j) \frac{1}{2\pi s^2} \int_{\mathbb{C}} d\mathbf{z}\, e^{-|\mathbf{z}-\sqrt{2\mu-1}\alpha_{\mathbf{u}}|^2/2s^2} \exp(-|\mathbf{z}|^2) \times \\ & |e(\mathbf{z}e^{-t/2})_j\rangle\langle e(\mathbf{z}e^{-t/2})_j| \otimes |e(\mathbf{z}e^{-u/2}\chi_{[0,t]}(u))\rangle\langle e(\mathbf{z}e^{-u/2}\chi_{[0,t]}(u))| \\ & + O(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon}),\end{aligned}$$

for $\|\mathbf{u}\| \leq n^{\eta}$. The index j serves to remind the reader that the first exponential states live in different copies $\mathcal{F}(\mathbb{C})_j$ of the oscillator space, corresponding to $\mathcal{H}_j$ via the isometry V_j. We will continue to identify $\mathcal{H}_j$ with its image in $\mathcal{F}(\mathbb{C})_j$.

We can now approximate the above state by its limit for large t, since

$$\exp(-|\mathbf{z}|^2)\langle e(\mathbf{z}e^{-t/2})_j|\, j,j\rangle\langle e(\mathbf{z}e^{-u/2}\chi_{[0,t]}(u))\,|\, e(\mathbf{z}e^{-u/2})\rangle = \exp(-|\mathbf{z}|^2 e^{-t}). \quad (8.26)$$

As we are always working with $\|\mathbf{u}\| \leq n^{\eta}$, the only relevant $\mathbf{z}$ are bounded by $n^{\eta+\delta}$ for small δ. (The remainder of the Gaussian integral has an exponentially decreasing norm, as discussed before). Thus, for large enough time (i.e. for $t \geq \ln(n)$), we can write $\rho_n^{\mathbf{u}}(t) = \rho_n^{\mathbf{u}}(\infty) + O(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon})$ with

$$\begin{aligned}\rho_n^{\mathbf{u}}(\infty) := & \bigoplus_{j=0,1/2}^{n/2} p_{n,\mathbf{u}}(j)|j,j\rangle\langle j,j| \otimes \\ & \left[\frac{1}{2\pi s^2}\int_{\mathbb{C}} d\mathbf{z}\, e^{-|\mathbf{z}-\sqrt{2\mu-1}\alpha_{\mathbf{u}}|^2/2s^2} |e(\mathbf{z}e^{-u/2})\rangle\langle e(\mathbf{z}e^{-u/2})| \exp(-|\mathbf{z}|^2)\right]. \quad (8.27)\end{aligned}$$

Thus, the field is approximately in the state $\phi^{\mathbf{u}}$ depending on (u_x, u_y), which is carried by the mode $(u \mapsto e^{-u/2}\chi_{[0,\infty)}(u)) \in L^2(\mathbb{R})$ denoted for simplicity by $e^{-u/2}$. The atoms end up in a mixture of $|j,j\rangle$ states with coefficients $p_{n,\mathbf{u}}(j)$, which depend only on u_z, and are well approximated by the Gaussian random variable $N^{\mathbf{u}}$ as shown

in Theorem 8.3.1. Moreover since there is no correlation between atoms and field, the statistical problem decouples into one concerning the estimation of the displacement in a family of Gaussian states $\phi^{\mathbf{u}}$, and one for estimating the center of $N^{\mathbf{u}}$.

For the former problem, the optimal estimation procedure is known to be the heterodyne measurement (Holevo, 1982; Yuen et Lax, M., 1973) ; for the latter, we perform a "which block" measurement. These measurements are described in the next two subsections.

8.5.1 The heterodyne measurement

A heterodyne measurement is a "joint measurement" of the quadratures $\mathbf{Q} := (a + a^*)/\sqrt{2}$ and $\mathbf{P} := -i(a - a^*)/\sqrt{2}$ of a quantum harmonic oscillator which in our case represents a mode of light. Since the two operators do not commute, the price to pay is the addition of some "noise" which will allow for an approximate measurement of both operators. The light beam passes through a beamsplitter having a vacuum mode as the second input, and then one performs a homodyne (quadrature) measurement on each of the two emerging beams. If $\mathbf{Q}_v$ and $\mathbf{P}_v$ are the vacuum quadratures then we measure the following output quadratures $\mathbf{Q}_1 := (\mathbf{Q}+\mathbf{Q}_v)/\sqrt{2}$ and $\mathbf{P}_2 := (\mathbf{P} - \mathbf{P}_v)/\sqrt{2}$, with $[\mathbf{Q}_1, \mathbf{P}_2] = 0$. Since the two input beams are independent, the distribution of $\sqrt{2}\mathbf{Q}_1$ is the convolution between the distribution of $\mathbf{Q}$ and the distribution of $\mathbf{Q}_v$, and similarly for $\sqrt{2}\mathbf{P}_2$.

In our case we are interested in the mode $e^{-u/2}$ which is in the state $\phi^{\mathbf{u}}$, up to a factor of order $O(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon})$. From (8.9) we obtain that the distribution of $\mathbf{Q}$ is $N(\sqrt{2(2\mu-1)}u_x, 1/(2(2\mu - 1)))$, that of $\mathbf{P}$ is $N(\sqrt{2(2\mu-1)}u_y, 1/(2(2\mu-1)))$, and the joint distribution of the rescaled output

$$\left((\mathbf{Q} + \mathbf{Q}_v)/\sqrt{2(2\mu-1)}\,,\ (\mathbf{P} - \mathbf{P}_v)/\sqrt{2(2\mu-1)}\right),$$

is

$$N(u_x, \mu/(2(2\mu-1)^2)) \times N(u_y, \mu/(2(2\mu-1)^2)). \tag{8.28}$$

We will denote by $(\tilde{u}_x, \tilde{u}_y)$ the result of the heterodyne measurement rescaled by the factor $\sqrt{2\mu-1}$ such that with good approximation $(\tilde{u}_x, \tilde{u}_y)$ has the above distribution and is an unbiased estimators of the parameters (u_x, u_y).

Since we know in advance that the parameters (u_x, u_y) must be within the radius of validity of LAN we modify the estimators $(\tilde{u}_x, \tilde{u}_y)$ to account for this information and obtain the final estimator $(\hat{u}_x, \hat{u}_y)$:

$$\hat{u}_i = \begin{cases} \tilde{u}_i & \text{if } |\tilde{u}_i| \leq 3n^{\eta} \\ 0 & \text{if } |\tilde{u}_i| > 3n^{\eta} \end{cases} \tag{8.29}$$

Notice that if the true state ρ is in the radius of validity of LAN around $\tilde{\rho}$, then $\|\mathbf{u}\| \leq n^{\eta}$, so that $|\hat{u}_i - u_i| \leq |\tilde{u}_i - u_i|$. We shall use this when proving optimality of the estimator.

8.5.2 Energy measurement

Having seen the $\phi^{\mathbf{u}}$-part, we now move to the $N^{\mathbf{u}}$-part of the equivalence between $\rho_n^{\mathbf{u}}$ and $N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}$. This too is a coupling to a bosonic field, albeit a different coupling. We also describe the measurement in the field which will provide the information on the qubit states.

The final state of the previous measurement, restricted to the atoms alone (without the field), is obtained by a partial trace of equation (8.27) (for large time) over the field

$$\tau_n^{\mathbf{u}} = \sum_{j=0,1/2}^{n/2} p_{n,\mathbf{u}}(j)|j,j\rangle\langle j,j| + O(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon}) .$$

We will take this as the initial state of the second measurement, which will determine j.

A direct coupling to the J^2 does not appear to be physically available, but a coupling to the energy J_z is realizable. This suffices, because the above state satisfies $j = m$ (up to order $O(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon})$). We couple the atoms to a new field (in the vacuum state $|\Omega\rangle$) by means of the interaction

$$dU_t = \{J_z(dA_t^* - dA_t) - \tfrac{1}{2}J_z^2 dt\}U_t ,$$

with $J_z := \frac{1}{\sqrt{n}}\sum_{k=1}^{n}\sigma_z$. Since this QSDE is 'essentially commutative', i.e. driven by a single classical noise $B_t = (A_t^* - A_t)/i$, the solution is easily seen to be

$$U_t = \exp(J_z \otimes (A_t^* - A_t)) .$$

Indeed, we have $df(B_t) = f'(B_t)dB_t + \frac{1}{2}f''(B_t)dt$ by the classical Itô rule, so that

$$d\exp(iJ_z \otimes B_t) = \{iJ_z dB_t - \tfrac{1}{2}J_z^2 dt\}\exp(iJ_z \otimes B_t) .$$

For an initial state $|j,m\rangle \otimes |\Omega\rangle$, this evolution gives rise to the final state

$$\begin{aligned} U_t|j,m\rangle \otimes \Omega &= |j,m\rangle \otimes \exp((m/\sqrt{n})(A_t^* - A_t))\Omega \\ &= |j,m\rangle \otimes |(m/\sqrt{n})\chi_{[0,t]}\rangle , \end{aligned}$$

where $|f\rangle \in \mathcal{F}(L^2(\mathbb{R}))$ denotes the normalized vector $\exp(-\langle f,f\rangle/2)e(f)$. Applying this to the states $|j,j\rangle\langle j,j|$ in $\tau_n^{\mathbf{u}}$ yields

$$\begin{aligned} U_t\,\tau_n^{\mathbf{u}} \otimes \Phi\, U_t^* = &\sum_{j=0,1/2}^{n/2} p_{n,\mathbf{u}}(j)|j,j\rangle\langle j,j| \otimes |j/\sqrt{n}\chi_{[0,t]}\rangle\langle j/\sqrt{n}\chi_{[0,t]}| \\ &+ O(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon}) . \end{aligned}$$

The final state of the field results from a partial trace over the atoms; it is given by

$$\sum_{j=0,1/2}^{n/2} p_{n,\mathbf{u}}(j)\, |(j/\sqrt{n})\chi_{[0,t]}\rangle\langle(j/\sqrt{n})\chi_{[0,t]}| + O(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon})\,. \tag{8.30}$$

We now perform a homodyne measurement on the field, which amounts to a direct measurement of $(A_t + A_t^*)/2t$. In the state $|(j/\sqrt{n}\chi_{[0,t]}\rangle$, this yields the value of j with certainty for large time (i.e. $t \gg \sqrt{n}$). Indeed, for this state, $\mathbb{E}((A_t + A_t^*)/2t) = j/\sqrt{n}$, whereas $\mathbb{V}\mathrm{ar}(A_t + A_t^*)/2t) = 1/(4t)$. Thus the probability distribution $p_{n,\mathbf{u}}$ is reproduced up to order $O(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon})$ in L^1-distance.

The following is a reminder from the proof of Theorem 8.3.1. If we start with j distributed according to $p_n(j)$ and we smoothen $\frac{j}{\sqrt{n}} - \sqrt{n}(\mu - 1/2)$ with a Gaussian kernel, then we obtain a random variable g_n which is continuously distributed on $\mathbb{R}$ and converges in distribution to $N(u_z, \mu(1-\mu))$, the error term being of order $O(n^{\eta-1/2}) + O(n^{\epsilon-1/2})$. For j distributed according to the actual distribution, as measured by the homodyne detection experiment, we can therefore state that g_n is distributed according to

$$N(u_z, \mu(1-\mu)) + O(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon}) + O(n^{\eta-1/2}) + O(n^{\epsilon-1/2}). \tag{8.31}$$

As in the case of $(\hat{u}_x, \hat{u}_y)$, we take into account the range of validity of LAN by defining the final estimator

$$\hat{u}_z = \begin{cases} g_n & \text{if } |g_n| \leq 3n^\eta \\ 0 & \text{if } |g_n| > 3n^\eta\,. \end{cases} \tag{8.32}$$

Similarly, we note that if the true state ρ is in the radius of validity of LAN around $\tilde{\rho}$, then $\|\mathbf{u}\| \leq n^\eta$, so that $|\hat{u}_z - u_z| \leq |\tilde{u}_z - u_z|$.

8.6 Asymptotic optimality of the estimator

In order to estimate the qubit state, we have proposed a strategy consisting of the following steps. First, we use $\tilde{n} := n^{1-\kappa}$ copies of the state ρ to get a rough estimate $\tilde{\rho}_n$. Then we couple the remaining qubits with a field, and perform a heterodyne measurement. Finally, we couple to a different field, followed by homodyne measurement. From the measurement outcomes, we construct an estimator $\hat{\rho}_n := \rho_{\hat{\mathbf{u}}_n/\sqrt{n}}$.

This strategy is asymptotically optimal in a global sense : for *any* true state ρ even if we knew beforehand that the true state ρ is in a small ball around a known state ρ_0, it would be impossible to devise an estimator that could do better asymptotically, than our estimator $\hat{\rho}_n$ on a small ball around ρ. More precisely :

Theorem 8.6.1. *Let $\hat{\rho}_n$ be the estimator defined above. For any qubit state ρ_0 different from the totally mixed state, for any sequence of estimators $\hat{\varrho}_n$, the following local asymptotic minimax result holds for any $0 < \epsilon < 1/12$:*

$$\limsup_{n\to\infty} \sup_{\|\rho-\rho_0\|_1 \le n^{-1/2+\epsilon}} nR(\rho, \hat{\rho}_n) \le \limsup_{n\to\infty} \sup_{\|\rho-\rho_0\|_1 \le n^{-1/2+\epsilon}} nR(\rho, \hat{\varrho}_n). \tag{8.33}$$

Let $(\mu_0, 1-\mu_0)$ be the eigenvalues of ρ_0 with $\mu_0 > 1/2$. Then the local asymptotic minimax risk is

$$\limsup_{n\to\infty} \sup_{\|\rho-\rho_0\|_1 \le n^{-1/2+\epsilon}} nR(\rho, \hat{\rho}_n) = R_{\text{minimax}}(\mu_0) = 8\mu_0 - 4\mu_0^2. \tag{8.34}$$

Démonstration. We write the risk as the sum of two terms corresponding to the events E and E^c that $\tilde{\rho}_n$ is inside or outside the ball of radius $n^{-1/2+\epsilon}$ around ρ. Recall that LAN is valid inside the ball. Thus

$$R(\rho, \hat{\rho}_n) = \mathbb{E}(\|\rho - \hat{\rho}_n\|_1^2 \chi_{E^c}) + \mathbb{E}(\|\rho - \hat{\rho}_n\|_1^2 \chi_E),$$

where the expectation comes from $\hat{\rho}_n$ being random. The distribution of the result $\hat{\rho_n}$ of our measurement procedure applied to the true unknown state ρ depends on ρ. We bound the first part by R_1 and the second part by R_2 as shown below.

R_1 equals $\mathbb{P}(E^c)$ times the maximum error, which is 4 since for any pair of density matrices ρ and σ, we have $\|\rho - \sigma\|_1^2 \le 4$. Thus

$$R_1 = 4\mathbb{P}(\|\rho - \tilde{\rho}_n\|_1 \ge n^{-1/2+\epsilon}).$$

According to Lemma 8.2.1 this probability goes to zero exponentially fast, therefore the contribution brought by this term can be neglected.

We can now assume that $\tilde{\rho}_n$ is in the range of validity of local asymptotic normality and we can write $\rho^{\otimes n} = \rho_n^{\mathbf{u}}$ with $\mathbf{u}$ the local parameter around $\tilde{\rho}_n$. We get the following inequalities for the second term in the risk.

$$\begin{aligned}
\mathbb{E}(\|\rho - \hat{\rho}_n\|_1^2 \chi_E) &\le \mathbb{E}\left[\|\hat{\rho}_n - \rho\|_1^2 \;\middle|\; \|\tilde{\rho}_n - \rho\|_1 \le n^{-1/2+\epsilon}\right] \\
&\le \sup_{\|\rho-\rho_0\|<n^{-1/2+\epsilon}} \mathbb{E}\left[\|\hat{\rho}_n - \rho\|_1^2 \;\middle|\; \tilde{\rho}_n = \rho_0\right] \\
&\le \sup_{\|\rho-\rho_0\|<n^{-1/2+\epsilon}} \mathbb{E}_{\rho_n^{\mathbf{u}}(\infty)}\left[\|\hat{\rho}_n - \rho\|_1^2 \;\middle|\; \tilde{\rho}_n = \rho_0\right] \\
&\quad + \sup_{\|\rho-\rho_0\|<n^{-1/2+\epsilon}} \|\rho_n^{\mathbf{u}}(t) - \rho_n^{\mathbf{u}}(\infty)\|_1 \sup_{\hat{\mathbf{u}}_n} \|\hat{\rho}_n - \rho\|_1^2 \\
&\le \sup_{\|\rho-\rho_0\|<n^{-1/2+\epsilon}} \mathbb{E}_{\rho_n^{\mathbf{u}}(\infty)}\left[\|\hat{\rho}_n - \rho\|_1^2 \;\middle|\; \tilde{\rho}_n = \rho_0\right] \\
&\quad + cn^{-1+2\eta} \sup_{\|\rho-\rho_0\|<n^{-1/2+\epsilon}} \|\rho_n^{\mathbf{u}}(t) - \rho_n^{\mathbf{u}}(\infty)\|_1 = R_2.
\end{aligned} \tag{8.35}$$

The first two inequalities are trivial. In the third inequality we change the expectation from the one with respect to the probability distribution of our data $\mathbb{P}_{\rho_n^{\mathbf{u}}(t)}$ to the probability distribution $\mathbb{P}_{\rho_n^{\mathbf{u}}(\infty)}$. In doing so, an additional term $\|\mathbb{P}_{\rho_n^{\mathbf{u}}(t)} - \mathbb{P}_{\rho_n^{\mathbf{u}}(\infty)}\|_1$ appears which is bounded from above by $\|\rho_n^{\mathbf{u}}(t) - \rho_n^{\mathbf{u}}(\infty)\|_1$. In the last inequality we can bound $\|\hat{\rho}_n - \rho\|_1^2$ by $cn^{-1+2\eta}$ for some constant c. Indeed from definitions (8.29) and (8.32) we know that $\|\hat{\rho}_n - \rho_0\|_1 \leq c' n^{-1/2+\eta}$ and additionally we are under the assumption $\|\rho - \rho_0\|_1 \leq n^{-1/2+\epsilon}$ with $\epsilon < \eta$.

For the following, recall that all our LAN estimates are valid uniformly around any state $\rho^{\mathbf{0}} = \tilde{\rho}$ as long as $\mu - 1/2 \geq \epsilon_2 > 0$. As we are working with ρ different from the totally mixed state and $\|\rho - \tilde{\rho}\| \leq n^{-1/2+\epsilon}$, we know that for big enough n, $\tilde{\mu} - 1/2 \geq \epsilon_2$ for any possible $\tilde{\rho}$. We can then apply the uniform results of the previous sections.

The second term in R_2 is $O(n^{-5/4+3\eta+\delta}, n^{-3/2+5\eta+\delta})$ where $\delta > 0$ can be chosen arbitrarily small. Indeed in the end of section 8.4 we have proven that after time $t \geq \ln n$, the following holds : $\|\rho_n^{\mathbf{u}}(t) - \rho_n^{\mathbf{u}}(\infty)\|_1 = O(n^{-1/4+\eta+\delta}, n^{-1/2+3\eta+\delta})$. The contribution to $nR(\rho, \hat{\rho}_n)$ brought by this term will not count in the limit, as long as η and ϵ are chose such that $1/12 > \eta > \epsilon$.

We now deal with the first term in R_2. We write ρ in local parametrization around $\rho_0 = \tilde{\rho}$ as $\rho_{\mathbf{u}_n/\sqrt{n}}$. We have

$$\begin{aligned} \|\hat{\rho}_n - \rho\|_1^2 &= \|\rho_{\mathbf{u}/\sqrt{n}} - \rho_{\hat{\mathbf{u}}_n/\sqrt{n}}\|_1^2 \\ &= 4\frac{(u_z - \hat{u}_z)^2 + (2\mu - 1)^2((u_x - \hat{u}_x)^2 + (u_y - \hat{u}_y)^2)}{n} \\ &\quad + O(\|\mathbf{u} - \hat{\mathbf{u}}_n\|^3 n^{-3/2}). \end{aligned} \tag{8.36}$$

The remainder term $O(\|\mathbf{u} - \hat{\mathbf{u}}_n\|^3 n^{-3/2})$ is negligible. It is $O(n^{3\eta-3/2})$ which does not contribute to $nR(\rho, \hat{\rho}_n)$ for $\eta < 1/6$. This is because on the one hand we have asked for $\|\tilde{\rho}_n - \rho\| < n^{-1/2+\epsilon}$, and on the other hand, we have bounded our estimator $\hat{\mathbf{u}}_n$ by using (8.29) and (8.32).

We now evaluate $\mathbb{E}_{\rho_n^{\mathbf{u}}(\infty)}[d(\mathbf{u}, \hat{\mathbf{u}}_n)^2]$ with the notation

$$d(\mathbf{u}, \mathbf{v})^2 := 4\left[(u_z - v_z)^2 + (2\mu - 1)^2((u_x - v_x)^2 + (u_y - v_y)^2)\right]. \tag{8.37}$$

Note that the risk of $\hat{\mathbf{u}}_n$ is smaller than that of $\tilde{\mathbf{u}}_n$ (see discussion below (8.29) and (8.32)). Under the law $\mathbb{P}_{\rho_n^{\mathbf{u}}(\infty)}$ the estimator $\tilde{\mathbf{u}}_n$ has a Gaussian distribution as shown in (8.28) and (8.31) with fixed and known variance and unknown expectation. In statistics this type of model is known as a Gaussian shift experiment (van der Vaart, 1998). Using (8.28) and (8.31), we get $\mathbb{E}_{\rho_n^{\mathbf{u}}(\infty)}[(u_z - \hat{u}_z)^2] \leq \mu(1-\mu)$ and $\mathbb{E}_{\rho_n^{\mathbf{u}}(\infty)}[(u_i - \hat{u}_i)^2] \leq \mu/(2(2\mu-1)^2)$ for $i = x, y$. Substituting these bounds in (8.36), we obtain (8.34).

We will now show that the sequence $\hat{\rho}_n$ is optimal in the local minimax sense : for any ρ_0 and any other sequence of estimators $\hat{\varrho}_n$ we have

$$R_0 = \limsup_{n\to\infty} \sup_{\|\rho-\rho_0\|_1 \leq n^{-1/2+\epsilon}} nR(\rho, \hat{\varrho}_n) \geq 8\mu_0 - 4\mu_0^2.$$

We will first prove that the right hand side is the minimax risk $R_{\text{minimax}}(\mu_0)$ for the family of states $N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}$ which is the limit of the local families $\rho_n^{\mathbf{u}}$ of qubit states centered around ρ_0. We then extend the result to our sequence of quantum statistical models $\rho_n^{\mathbf{u}}$.

The minimax optimality for $N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}$ can be checked separately for the classical and the quantum part of the experiment. For the quantum part $\phi^{\mathbf{u}}$, the optimal measurement is known to be the heterodyne measurement. A proof of this fact can be found in Lemma 7.4 of (Guţă et Kahn, 2006). For the classical part, which corresponds to the measurement of L_z, the optimal estimator is simply the random variable $X \sim N^{\mathbf{u}}$ itself (van der Vaart, 1998).

We now end the proof by using the other direction of LAN. Suppose that there exists a better sequence of estimators $\hat{\varrho}_n$ such that

$$R_0 < R_{\text{minimax}}(\mu_0) = 8\mu_0 - 4\mu_0^2.$$

We will show that this leads to an estimator $\hat{u}$ of $\mathbf{u}$ for the family $N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}$ whose maximum risk is smaller than the minimax risk $R_{\text{minimax}}(\mu_0)$, which is impossible.

By means of a beamsplitter one can divide the state $\phi^{\mathbf{u}}$ into two independent Gaussian modes, using a thermal state $\phi := \phi^0$ as the second input. If r and t are the reflectivity and respective transmitivity of the beamsplitter ($r^2 + t^2 = 1$), then the transmitted beam has state $\phi_{tr}^{\mathbf{u}} = \phi^{t\mathbf{u}}$ and the reflected one $\phi_{ref}^{\mathbf{u}} = \phi^{r\mathbf{u}}$. By performing a heterodyne measurement on the latter, and observing the classical part $N^{\mathbf{u}}$, we can localize $\mathbf{u}$ within a big ball around the result $\tilde{\mathbf{u}}$ with high probability, in the spirit of Lemma 8.2.1. More precisely, for any small $\tilde{\epsilon} > 0$ we can find $a > 0$ big enough such that the risk contribution from unlikely $\tilde{\mathbf{u}}$'s is small

$$\mathbb{E}(\|\mathbf{u} - \tilde{\mathbf{u}}\|^2 \chi_{\|\mathbf{u}-\tilde{\mathbf{u}}\|>a}) < \tilde{\epsilon}.$$

Summarizing the localization step, we may assume that the parameter $\mathbf{u}$ satisfies $\|\mathbf{u}\| < a$ with an $\tilde{\epsilon}$ loss of risk, where $a = a(r, \tilde{\epsilon})$.

Now let n be large enough such that $n^{\epsilon} > a$, then the parameter $\mathbf{u}$ falls within the domain of convergence of the inverse map S_n of Theorem 8.3.1 and by (8.12) (with ϵ replacing η and δ replacing ϵ) we have

$$\|\rho_n^{t\mathbf{u}} - S(N^{t\mathbf{u}} \otimes \phi^{t\mathbf{u}})\|_1 \leq Cn^{-1/4+\epsilon+\delta},$$

for some constant C.

Next we perform the measurement leading to the estimator $\hat{\varrho}_n$ and equivalently to an estimator $\hat{\mathbf{u}}_n$ of $\mathbf{u}$. Without loss of risk we can implement the condition $\|\mathbf{u}\| < a$ into the estimator $\hat{\mathbf{u}}_n$ in a similar fashion as in (8.29) and (8.32). The risk of this estimation procedure for $\phi^{\mathbf{u}}$ is then bounded from above by the sum of three terms : the risk $nR_\rho(\hat{\varrho}_n)/t^2$ coming from the qubit estimation, the error contribution from the map S_n which is $a^2 n^{-1/4+\epsilon+\delta}$, and the localization risk contribution $\tilde{\epsilon}$. This risk bound uses the same technique as the third inequality of (8.35). The second contribution can be made arbitrarily small by choosing n large enough, for $\epsilon < 1/4$. From our assumption we have $R_0 < R_{minimax}(\mu_0)$ and we can choose t close to one such that $R_0/t^2 < R_{minimax}(\mu_0)$ and further choose $\tilde{\epsilon}$ such that $R_0/t^2 + \tilde{\epsilon} < R_{minimax}(\mu_0)$.

In conclusion, we get that the risk for estimating $\mathbf{u}$ is asymptotically smaller that the risk of the heterodyne measurement combined with observing the classical part which is known to be minimax (Guţă et Kahn, 2006). Hence no such sequence $\hat{\varrho}_n$ exists, and $\hat{\rho}_n$ is optimal.

□

Remark. In Theorem 8.33, we have used the risk function $R(\rho, \hat{\rho}) = \mathbb{E}(d^2(\rho, \hat{\rho}))$, with d the L_1-distance $d(\rho, \hat{\rho}) = \|\rho - \hat{\rho}\|_1$. However, the obtained results can easily be adapted to *any* distance measure $d^2(\rho_{\hat{\mathbf{u}}}, \rho_{\mathbf{u}})$ which is locally quadratic in $\hat{\mathbf{u}} - \mathbf{u}$, i.e.

$$d^2(\rho_{\hat{\mathbf{u}}}, \rho_{\mathbf{u}}) = \sum_{\alpha,\beta=x,y,z} \gamma_{\alpha\beta}(u_\alpha - \hat{u}_\alpha)(u_\beta - \hat{u}_\beta) + O(\|u - \hat{u}\|^3)\,.$$

For instance, one may choose $d^2(\hat{\rho}, \rho) = 1 - F^2(\hat{\rho}, \rho)$ with the fidelity $F(\hat{\rho}, \rho) := \mathrm{Tr}(\sqrt{\sqrt{\hat{\rho}}\rho\sqrt{\hat{\rho}}})$. For non-pure states, this is easily seen to be locally quadratic with

$$\gamma = \begin{pmatrix} (2\mu_0 - 1)^2 & 0 & 0 \\ 0 & (2\mu_0 - 1)^2 & 0 \\ 0 & 0 & \frac{1}{1-(2\mu_0-1)^2} \end{pmatrix}.$$

For the corresponding risk function $R_F(\rho, \hat{\rho}_n) := \mathbb{E}(1 - F^2(\rho, \hat{\rho}_n))$, this yields

$$\limsup_{n\to\infty} \sup_{\|\rho-\rho_0\|_1 \le n^{-1/2+\epsilon}} nR_F(\rho, \hat{\rho}_n) = \mu_0 + 1/4\,, \tag{8.38}$$

with the same asymptotically optimal $\hat{\rho}$. The asymptotic rate $R_F \sim \frac{4\mu_0+1}{4n}$ was found earlier by Bagan et al. (2006), using different methods.

8.7 Conclusions

In this chapter, we have shown two properties of quantum local asymptotic normality (LAN) for qubits. First of all, we have seen that its radius of validity is arbitrarily close to $n^{-1/4}$ rather than $n^{-1/2}$. And secondly, we have seen how LAN can be implemented physically, in a quantum optical setup.

We use these properties to construct an asymptotically optimal estimator $\hat{\rho}_n$ of the qubit state ρ, provided that we are given n identical copies of ρ. Compared with other optimal estimation methods (Bagan et al., 2006; Hayashi et Matsumoto, 2004), our measurement technique makes a significant step in the direction of an experimental implementation.

The construction and optimality of $\hat{\rho}_n$ are shown in three steps.

I In the preliminary stage, we perform measurements of σ_x, σ_y and σ_z on a fraction $\tilde{n} = n^{1-\kappa}$ of the n atoms. As shown in section 8.2, this yields a rough estimate $\tilde{\rho}_n$ which lies within a distance $n^{-1/2+\epsilon}$ of the true state ρ with high probability.

II In section 8.3, it is shown that local asymptotic normality holds within a ball of radius $n^{-1/2+\eta}$ around ρ ($\eta > \epsilon$). This means that locally, for $n \to \infty$, all statistical problems concerning the n identically prepared qubits are equivalent to statistical problems concerning a Gaussian distribution $N^{\mathbf{u}}$ and its quantum analogue, a displaced thermal state $\phi^{\mathbf{u}}$ of the harmonic oscillator.

Together, I and II imply that the principle of LAN has been extended to a global setting. It can now be used for a wide range of asymptotic statistical problems, including the global problem of state estimation. Note that this hinges on the rather subtle extension of the range of validity of LAN to neighborhoods of radius larger than $n^{-1/2}$.

III LAN provides an abstract equivalence between the n-qubit states $\rho_{\mathbf{u}/\sqrt{n}}^{\otimes n}$ on the one hand, and on the other hand the Gaussian states $N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}$. In sections 8.4 and 8.5 it is shown that this abstract equivalence can be implemented physically by two consecutive couplings to the electromagnetic field. For the particular problem of state estimation, homodyne and heterodyne detection on the electromagnetic field then yield the data from which the optimal estimator $\hat{\rho}_n$ is computed.

Finally, in section 8.6, it is shown that the estimator $\hat{\rho}_n$, constructed above, is optimal in a local minimax sense. Local here means that optimality holds in a ball of radius slightly bigger than $n^{-1/2}$ around *any* state ρ_0 except the tracial state. That is, even if we had known beforehand that the true state lies within this ball around ρ_0, we would not have been able to construct a better estimator than $\hat{\rho}_n$, which is of course independent of ρ_0.

For this asymptotically optimal estimator, we have shown that the risk R converges

to zero at rate $R(\rho, \hat{\rho}_n) \sim \frac{8\mu_0 - 4\mu_0^2}{n}$, with $\mu_0 > 1/2$ an eigenvalue of ρ. More precisely, we have

$$\limsup_{n\to\infty} \sup_{\|\rho-\rho_0\|_1 \leq n^{-1/2+\epsilon}} nR(\rho, \hat{\rho}_n) = 8\mu_0 - 4\mu_0^2.$$

The risk is defined as $R(\rho, \hat{\rho}) = \mathbb{E}(d^2(\rho, \hat{\rho}))$, where we have chosen $d(\hat{\rho}, \rho)$ to be the L_1-distance $\|\hat{\rho} - \rho\|_1 := \mathrm{Tr}(|\hat{\rho} - \rho|)$. This seems to be a rather natural choice because of its direct physical significance as the worst case difference between the probabilities induced by $\hat{\rho}$ and ρ on a single event.

Even still, we emphasize that the same procedure can be applied to a wide range of other risk functions. Due to the local nature of the estimator $\hat{\rho}_n$ for large n, its rate of convergence in a risk R is only sensitive to the lowest order Taylor expansion of R in local parameters $\hat{\mathbf{u}} - \mathbf{u}$. The procedure can therefore easily be adapted to other risk functions, provided that the distance measure $d^2(\rho_{\hat{\mathbf{u}}}, \rho_{\mathbf{u}})$ is locally quadratic in $\hat{\mathbf{u}} - \mathbf{u}$.

Remark. The totally mixed state ($\mu = 1/2$) is a singular point in the parameter space, and Theorem 8.3.1 does not apply in this case. The effect of the singularity is that the family of states (8.9) collapses to a single degenerate state of infinite temperature. However this phenomenon is only due to our particular parametrisation, which was chosen for its convenience in describing the local neighborhoods around arbitrary states, with the exception of the totally mixed state. Had we chosen a different parametrisation, e.g. in terms of the Bloch vector, we would have found that local asymptotic normality holds for the totally mixed state as well, but the limit experiment is different : it consists of a three dimensional *classical* Gaussian shift, each independent component corresponding to the local change in the Bloch vector along the three possible directions. Mathematically, the optimal measurement strategy in this case is just to observe the classical variables. However this strategy cannot be implemented by coupling with the field since this coupling becomes singular (see equation (8.18)).

These issues become more important for higher dimensional systems where the eigenvalues may exhibit more complicated multiplicities, and will be dealt with in that context.

8.A Appendix : Proof of Theorem 8.3.1

Here we give the technical details of the proof of local asymptotic normality with "slowly growing" local neighborhoods $\|\mathbf{u}\| \leq n^\eta$, with $\eta < 1/4$. We start with the map T_n.

8.A.1 Proof of Theorem 8.3.1 ; the map T_n

Let us define, for $0 < \epsilon < (1/4 - \eta)$ the interval

$$\mathcal{J}_n = \left\{ j \,:\, (\mu - 1/2)n - n^{1/2+\epsilon} \leq j \leq (\mu - 1/2)n + n^{1/2+\epsilon} \right\}. \tag{8.39}$$

Notice that $j \in \mathcal{J}_n$ satisfies $2j \geq \epsilon_2 n$ for all $\mu - 1/2 \geq \epsilon_2$ and n big enough, independently of μ.

Then $\mathcal{J}_n$ contains the relevant values of j, uniformly for $\mu - 1/2 \geq \epsilon_2$:

$$\lim_{n\to\infty} p_{n,\mathbf{u}}(\mathcal{J}_n) = 1 - O(n^{-1/2+\epsilon}). \tag{8.40}$$

This is a consequence of Hoeffding's inequality applied to the binomial distribution, and recalling that $p_{n,\mathbf{u}}(j) = B(n/2 + j)(1 + O(n^{-1/2+\epsilon}))$ for $j \in \mathcal{J}_n$.

We upper-bound $\|T_n(\rho_n^{\mathbf{u}}) - N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}\|$ by the sum

$$3 \sum_{j \notin \mathcal{J}_n} p_{n,j}^{\mathbf{u}} + \left\| N^{\mathbf{u}} - \sum_{j \in \mathcal{J}_n} p_{n,\mathbf{u}}(j) \tau_{n,j} \right\|_1 + \sup_{j \in \mathcal{J}_n} \|V_j \rho_{j,n}^{\mathbf{u}} V_j^* - \phi^{\mathbf{u}}\|_1. \tag{8.41}$$

The first two terms are "classical" and converge to zero uniformly over $\|\mathbf{u}\| \leq n^\eta$: for the first term, this is (8.40), while the second term converges uniformly on $\mu - 1/2 \geq \epsilon_2$ at rate $n^{\eta - 1/2}$ (Guţă et Kahn, 2009). The third term can be analyzed as in Proposition 5.1 of (Guţă et Kahn, 2006) :

$$\left\|V_j \rho_{n,j}^{\mathbf{u}} V_j^* - \phi^{\mathbf{u}}\right\|_1 \;\leq\; \left\|\rho_{n,j}^{\mathbf{u}} - V_j^* \phi^{\mathbf{u}} V_j\right\|_1 + \left\|\phi^{\mathbf{u}} - P_j \phi^{\mathbf{u}} P_j\right\|_1, \tag{8.42}$$

where $P_j := V_j V_j^*$ is the projection onto the image of V_j. We will show that both terms on the right side go to zero uniformly at rate $n^{-1/4+\eta+\epsilon}$ over $j \in \mathcal{J}_n$ and $\|\mathbf{u}\| \leq n^\eta$. The trick is to note that displaced thermal equilibrium states are Gaussian mixtures of coherent states

$$\phi^{\mathbf{u}} = \frac{1}{\sqrt{2\pi s^2}} \int e^{-|\mathbf{z} - \sqrt{2\mu - 1}\alpha_{\mathbf{u}}|^2/2s^2} \left(|\mathbf{z}\rangle\langle\mathbf{z}|\right) d^2\mathbf{z}, \tag{8.43}$$

where $s^2 := (1 - \mu)/(4\mu - 2)$.

The second term on the left side of (8.42) is bounded from above by

$$\frac{1}{\sqrt{2\pi s^2}} \int e^{-|\mathbf{z} - \sqrt{2\mu - 1}\alpha_{\mathbf{u}}|^2/2s^2} \| |\mathbf{z}\rangle\langle\mathbf{z}| - P_j |\mathbf{z}\rangle\langle\mathbf{z}| P_j \|_1 \, d^2\mathbf{z},$$

which after some simple computations can be reduced (up to a constant) to

$$\int e^{-|\mathbf{z}|^2/2s^2} \| P_j^{\perp} |\mathbf{z} + \sqrt{2\mu - 1}\alpha_{\mathbf{u}}\rangle \| \, d^2\mathbf{z}. \tag{8.44}$$

We now split the integral. The first part is integrating over $|\mathbf{z}| \geq n^{\eta+\delta}$ with $0 < \delta < 1/4 - \eta/2$. The integral is dominated by the Gaussian and its value is $O(e^{-n^{2(\eta+\delta)}/(2s^2)})$. The other part is bounded by the supremum over $|\mathbf{z}| \leq 2n^{\eta+\delta}$ (as $\|\mathbf{u}\| \leq n^{\eta}$) of $\|P_j^{\perp}|\mathbf{z}\rangle\|$. Now $\|P_j^{\perp}|\mathbf{z}\rangle\| \leq |\mathbf{z}|^j/\sqrt{j!} = O(e^{-n(1/2-\eta-2\delta)})$ uniformly on $j \in \mathcal{J}_n$, for any $\mu - 1/2 \geq \epsilon_2$ since then $2j \geq \epsilon_2 n$.

The same type of estimates apply to the first term

$$\begin{aligned} \left\| \rho_{n,j}^{\mathbf{u}} - V_j^* \phi^{\mathbf{u}} V_j \right\|_1 &= \left\| \mathrm{Ad}\left[U_j\left(\frac{\mathbf{u}}{\sqrt{n}}\right)\right](\rho_{n,j}^{\mathbf{0}}) - V_j^* \phi^{\mathbf{u}} V_j \right\|_1 \leq \\ \left\| \rho_{n,j}^{\mathbf{0}} - V_j^* \phi^{\mathbf{0}} V_j \right\|_1 &+ \left\| \mathrm{Ad}\left[U_j\left(\frac{\mathbf{u}}{\sqrt{n}}\right)\right](V_j^* \phi^{\mathbf{0}} V_j) - V_j^* \phi^{\mathbf{u}} V_j \right\|_1 . \end{aligned} \tag{8.45}$$

The first term on the right side does not depend on $\mathbf{u}$. From the proof of Lemma 5.4 of (Guţă et Kahn, 2006), we know that

$$\left\| \rho_{n,j}^{\mathbf{0}} - V_j^* \phi^{\mathbf{0}} V_j \right\|_1 \leq \left(\frac{p^{2j+1}}{1 - p^{2j+1}} + p^{2j+1} \right)$$

with $p = (1-\mu)/\mu$. Now the left side is of the order p^{2j+1} which converges exponentially fast to zero uniformly on $\mu - 1/2 \geq \epsilon_2$ and $j \in \mathcal{J}_n$.

The second term of (8.45) can be bounded again by a Gaussian integral

$$\frac{1}{\sqrt{2\pi s^2}} \int e^{-|\mathbf{z}|^2/2s^2} \| \Delta(\mathbf{u}, \mathbf{z}, j) \|_1 d^2\mathbf{z}, \tag{8.46}$$

where the operator $\Delta(\mathbf{u}, \mathbf{z}, j)$ is given by

$$\Delta(\mathbf{u}, \mathbf{z}, j) := \mathrm{Ad}\left[U_j\left(\mathbf{u}/\sqrt{n}\right)\right]\left(V_j^* |\mathbf{z}\rangle\langle\mathbf{z}| V_j\right) - V_j^* \, |\mathbf{z} + \sqrt{2\mu - 1}\alpha_{\mathbf{u}}\rangle\langle\mathbf{z} + \sqrt{2\mu - 1}\alpha_{\mathbf{u}}| \, V_j.$$

Again, we split the integral along $\|\mathbf{z}\| \geq n^{\eta+\delta}$. The outer part converges to zero faster than any power of n, as we have already seen. The inner integral, on the other hand, can be bounded uniformly over $\|\mathbf{u}\| \leq n^{\eta}$, $\mu - 1/2 \geq \epsilon_2$ and $j \in \mathcal{J}_n$ by the supremum of $\|\Delta(\mathbf{u}, \mathbf{z}, j)\|_1$ over $|\mathbf{z}| \leq 2n^{\eta+\delta}$, $\mu - 1/2 \geq \epsilon_2$, $j \in \mathcal{J}_n$ and $\|\mathbf{u}\| \leq n^{\eta}$.

Let $\tilde{\mathbf{z}} \in \mathbb{R}^2$ be such that $\alpha_{\tilde{\mathbf{z}}} = \mathbf{z}/\sqrt{2\mu - 1}$, and denote $\psi(n, j, \mathbf{v}) = V_j U_j(\mathbf{v}/\sqrt{n})|j, j\rangle$. Then, up to a $\sqrt{2}$ factor, $\|\Delta(\mathbf{u}, \mathbf{z}, j)\|_1$ is bounded from above by the

$$\begin{aligned} &\| \psi(n, j, \tilde{\mathbf{z}}) - |\mathbf{z}\rangle \| + \\ &\left\| \psi(n, j, \mathbf{u} + \tilde{\mathbf{z}}) - |\mathbf{z} + \sqrt{2\mu - 1}\alpha_{\mathbf{u}}\rangle \right\| + \\ &\left\| U_j\left(\frac{\mathbf{u}}{\sqrt{n}}\right) U_j\left(\frac{\tilde{\mathbf{z}}}{\sqrt{n}}\right) |jj\rangle - U_j\left(\frac{\mathbf{u} + \tilde{\mathbf{z}}}{\sqrt{n}}\right) |jj\rangle \right\|. \end{aligned} \tag{8.47}$$

This is obtained by adding and subtracting $|\psi(n,j,\tilde{\mathbf{z}})\rangle\langle\psi(n,j,\tilde{\mathbf{z}})|$ and $|\psi(n,j,\mathbf{u}+\tilde{\mathbf{z}})\rangle\langle\psi(n,j,\mathbf{u}+\tilde{\mathbf{z}})|$ and using the fact that $\||\psi\rangle\langle\psi| - |\phi\rangle\langle\phi|\|_1 = \sqrt{2}\|\psi-\phi\|$ for normalized vectors ψ, ϕ.

The two first terms are similar, we want to dominate them uniformly : we replace $\mathbf{u}+\tilde{\mathbf{z}}$ by $\tilde{\mathbf{z}}$ with $|\mathbf{z}| \leq 2n^{\eta+\delta}$. We then write :

$$\begin{aligned}\|\psi(n,j,\tilde{\mathbf{z}}) - |\mathbf{z}\rangle\|^2 &= \sum_{k=0}^{\infty} |\langle k|\psi(n,j,\tilde{\mathbf{z}})\rangle - \langle k|\mathbf{z}\rangle|^2 \\ &\leq \sum_{k=0}^{r-1} |\langle k|\psi(n,j,\tilde{\mathbf{z}})\rangle - \langle k|\mathbf{z}\rangle|^2 + 2\sum_{k=r}^{\infty} \left(|\langle k|\psi(n,j,\tilde{\mathbf{z}})\rangle|^2 + |\langle k|\mathbf{z}\rangle|^2\right). \quad (8.48)\end{aligned}$$

If $\mathbf{z} = |\mathbf{z}|e^{i\theta}$ then we have (Hayashi et Matsumoto, 2004)

$$\begin{aligned}\langle k|\psi(n,j,\tilde{\mathbf{z}})\rangle &= \sqrt{\binom{2j}{k}} \left(\sin(|\mathbf{z}|/\sqrt{n})e^{i\theta}\right)^k \left(\cos(|\mathbf{z}|\sqrt{n})\right)^{2j-k}, \\ \langle k|\mathbf{z}\rangle &= \exp\left(-\frac{(2\mu-1)|\mathbf{z}|^2}{2}\right) \frac{\left(e^{i\theta}|\mathbf{z}|\sqrt{2\mu-1}\right)^k}{\sqrt{k!}}.\end{aligned}$$

In (8.48) we choose $r = n^{2\eta+\epsilon_3}$ with ϵ_3 satisfying the conditions $2\delta + 2\eta + \epsilon < 2\eta + \epsilon_3 + \epsilon < 1/2$ and $\eta + \epsilon_3 < 1/4$. Then the tail sums are of the order

$$\begin{aligned}\sum_{k=r}^{\infty} |\langle k|\mathbf{z}\rangle|^2 &\leq \frac{|\mathbf{z}|^{2r}}{r!} \leq \frac{(2n^{(\eta+\delta)})^{2n^{2\eta+\epsilon_3}}}{(n^{2\eta+\epsilon_3})!} = o\left(\exp(-n^{2\eta+\epsilon_3})\right), \\ \sum_{k=r}^{\infty} |\langle k|\psi(n,j,\tilde{\mathbf{z}})\rangle|^2 &\leq \sum_{k=r}^{j} \left(\frac{|\mathbf{z}|^2}{n}\right)^k \frac{(2j)!}{(2j-k)!k!} \leq n\frac{|\mathbf{z}|^{2r}}{r!} = o\left(\exp(-n^{2\eta+\epsilon_3})\right).\end{aligned}$$

For the finite sums we use the following estimates which are uniform over all $|\mathbf{z}| \leq 2n^{\eta+\delta}$, $k \leq r$, $j \in \mathcal{J}_n$:

$$\begin{aligned}\sqrt{\binom{2j}{k}} &= \frac{((2\mu-1)n)^{k/2}}{\sqrt{k!}}(1 + O(n^{-1/2+\epsilon+2\eta+\epsilon_3})), \\ (\sin(|\mathbf{z}|/\sqrt{n}))^k &= (|\mathbf{z}|/\sqrt{n})^k(1 + O(n^{4\eta+\epsilon_3+2\delta-1})), \\ (\cos(|\mathbf{z}|/\sqrt{n}))^{2j-k} &= \exp\left(-\frac{(2\mu-1)|\mathbf{z}|^2}{2}\right)(1 + O(n^{2\eta-1/2+\epsilon+2\delta})),\end{aligned}$$

where we have used on the last line that $(1+x/n)^n = \exp(x)(1+O(n^{-1/2}x))$ for $x \leq n^{1/2-\epsilon_4}$ (cf. (Guţă et Kahn, 2009)). This is enough to show that the finite sum converges uniformly to zero at rate $O(n^{2\eta-1/2+\epsilon+\epsilon_3})$ (the worst if ϵ_3 is small enough) and thus the first second terms in (8.47) as the square root of this, that is $O(n^{\eta-1/4+\epsilon/2+\epsilon_3/2})$.

Notice that the errors terms depend on μ only through j, and that $2j \geq \epsilon n$ for $\mu - 1/2 \geq \epsilon_2$. Hence they are uniform in μ.

We pass now to the third term of (8.47). By direct computation it can be shown that if we consider two general elements $\exp(iX_1)$ and $\exp(iX_2)$ of $SU(2)$ with X_i selfadjoint elements of $M(\mathbb{C}^2)$ then

$$\exp(-i(X_1+X_2))\exp(iX_1)\exp(iX_2)\exp([X_1,X_2]/2) = \mathbf{1} + O(X_{i_1}X_{i_2}X_{i_3}), \quad (8.49)$$

where the $O(\cdot)$ contains only third order terms in X_1, X_2. If X_1, X_2 are in the linear span of σ_x and σ_y then all third order monomials are such linear combinations as well.

In particular we get that for $\mathbf{z}, \mathbf{u} \leq n^{\eta+\epsilon_3}$:

$$\begin{aligned} U(\beta) := \quad & U\left(-\frac{\mathbf{u}+\mathbf{v}}{\sqrt{n}}\right) U\left(\frac{\mathbf{u}}{\sqrt{n}}\right) U\left(\frac{\mathbf{v}}{\sqrt{n}}\right) \exp(i(u_x v_y - u_y v_x)\sigma_z/n) \\ = \quad & \begin{bmatrix} 1+O(n^{-2+4\eta+4\epsilon_3}) & O(n^{-3/2+3\eta+3\epsilon_3}) \\ O(n^{-3/2+3\eta+3\epsilon_3}) & 1+O(n^{-2+4\eta+4\epsilon_3}) \end{bmatrix}. \end{aligned} \quad (8.50)$$

Finally,using the fact that $|j,j\rangle$ is an eigenvector of L_z, the third term in (8.47) can be written as

$$\| |j,j\rangle\langle j,j| - U_j(\beta)|j,j\rangle\langle j,j|U_j(\beta)^*\|$$

and both states are pure, so it suffices to show that the scalar product converges to to one uniformly. Using (8.50) and the expression of $\langle j|U_j(\beta)|j\rangle$ (Hayashi et Matsumoto, 2004) we get, as $j \leq n$,

$$\langle j,j|U_j(\beta)|j,j\rangle = [U(\beta)_{1,1}]^j = 1 + O(n^{-1+4\eta+4\epsilon_3}),$$

which implies that the third term in (8.47) is of order $O(n^{-1+4\eta+4\epsilon_3})$. By choosing ϵ_3 and ϵ small enough, we obtain that all terms used in bounding (8.46) are uniformly $O(n^{-1/4+\eta+\epsilon})$ for any $\epsilon > 0$.

This ends the proof of convergence (8.11) from the n qubit state to the oscillator.

8.A.2 Proof of Theorem 8.3.1 ; the map S_n

The opposite direction (8.12) does not require much additional estimation, so will only give an outline of the argument.

Given the state $N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}$, we would like to map it into $\rho_n^{\mathbf{u}}$ or close to this state, by means of a completely positive map S_n.

Let X be the classical random variable with probability distribution $N^{\mathbf{u}}$. With X we generate a random $j \in \mathbb{Z}$ as follows

$$j(X) = [\sqrt{n}X + n(\mu - 1/2)].$$

This choice is evident from the scaling properties of the probability distribution $p_n^{\mathbf{u}}$ which we want to reconstruct. Let $q_n^{\mathbf{u}}$ be the probability distribution of $j(X)$. By classical local asymptotic normality results we have the convergence

$$\sup_{\|\mathbf{u}\| \leq n^{\eta}} \|q_n^{\mathbf{u}} - p_n^{\mathbf{u}}\|_1 = O(n^{\eta - 1/2}). \tag{8.51}$$

Now, if the integer j is in the interval $\mathcal{J}_n$ then we prepare the n qubits in block diagonal state with the only non-zero block corresponding to the j'th irreducible representation of $SU(2)$:

$$\tau_{n,j}^{\mathbf{u}} := \left(V_j^* \phi^{\mathbf{u}} V_j + \mathrm{Tr}(P_j^{\perp} \phi^{\mathbf{u}})\mathbf{1}\right) \otimes \frac{\mathbf{1}}{n_j}.$$

The transformation $\phi^{\mathbf{u}} \mapsto \tau_{n,j}^{\mathbf{u}}$ is trace preserving and completely positive (Guţă et Kahn, 2006).

If $j \notin \mathcal{J}_n$ then we may prepare the qubits in an arbitrary state which we also denote by $\tau_{n,j}^{\mathbf{u}}$. The total channel S_n then acts as follows

$$S_n : N^{\mathbf{u}} \otimes \phi^{\mathbf{u}} \mapsto \tau_n^{\mathbf{u}} := \bigoplus_{j=0,1/2}^{n/2} q_{n,j}^{\mathbf{u}} \tau_{n,j}^{\mathbf{u}}.$$

We estimate the error $\|\rho_n^{\mathbf{u}} - \tau_n^{\mathbf{u}}\|_1$ as

$$\|\rho_n^{\mathbf{u}} - \tau_n^{\mathbf{u}}\|_1 \leq \|q_n^{\mathbf{u}} - p_n^{\mathbf{u}}\|_1 + 2\mathbb{P}_{p_n^{\mathbf{u}}}(j \notin \mathcal{J}_n) + \sup_{j \in \mathcal{J}_n} \|\tau_{n,j}^{\mathbf{u}} - \rho_{n,j}^{\mathbf{u}}\|_1$$

The first term on the r.h.s. is $O(n^{\eta - 1/2})$ (see (8.51)), the second term is $O(n^{\epsilon - 1/2})$ (see (8.40)). As for the third term, we use the triangle inequality to write, for $j \in \mathcal{J}_n$,

$$\|\tau_{n,j}^{\mathbf{u}} - \rho_{n,j}^{\mathbf{u}}\|_1 \leq \|\tau_{n,j}^{\mathbf{u}} - V_j^* \phi^{\mathbf{u}} V_j^*\|_1 + \|V_j^* \phi^{\mathbf{u}} V_j^* - \rho_{n,j}^{\mathbf{u}}\|_1 .$$

The first term is $O(e^{-n(1/2 - \eta - 2\delta)})$, according to the discussion following equation (8.44). The second term on the right is $O(n^{-1/4 + \eta + \epsilon})$ according to equations (8.45) through (8.50).

Summarizing, we have $\|S_n(N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}) - \rho_n^{\mathbf{u}}\|_1 = O(n^{-1/4 + \eta + \epsilon})$, which establishes the proof in the inverse direction.

□

8.B Appendix : Proof of Theorem 8.4.1

First estimate. We build up the state $\tilde{\rho}^{\mathbf{u}}_{j,n}$ by taking linear combinations of number states $|m\rangle$ to obtain an approximate coherent state $|\mathbf{z}\rangle$, and finally mixing such states with a Gaussian distribution to get an approximate displaced thermal state. Consider the approximate coherent vector $P_{\tilde{m}}|\mathbf{z}\rangle$, for some fixed $\mathbf{z} \in \mathbb{C}$ and $\tilde{m} = n^{\gamma}$, with γ to be fixed later. Define the normalized vector

$$|\psi^n_{\mathbf{z},j}\rangle := \frac{1}{\|P_{\tilde{m}}|\mathbf{z}\rangle\|} \sum_{m=0}^{\tilde{m}} \frac{|\mathbf{z}|^m}{\sqrt{m!}} |m\rangle, \tag{8.52}$$

We mix the above states to obtain

$$\tilde{\rho}^{\mathbf{u}}_{j,n} := \frac{1}{\sqrt{2\pi s^2}} \int e^{-|\mathbf{z}-\sqrt{2\mu-1}\alpha_{\mathbf{u}}|^2/2s^2} \left(|\psi^n_{\mathbf{z},j}\rangle\langle\psi^n_{\mathbf{z},j}|\right) d^2\mathbf{z}.$$

Recall that $s^2 = (1-\mu)(4\mu-2)$, and

$$\phi^{\mathbf{u}} = \frac{1}{\sqrt{2\pi s^2}} \int e^{-|\mathbf{z}-\sqrt{2\mu-1}\alpha_{\mathbf{u}}|^2/2s^2} \left(|\mathbf{z}\rangle\langle\mathbf{z}|\right) d^2\mathbf{z}.$$

From the definition of $|\psi^n_{\mathbf{z},j}\rangle$ we have

$$\| |\psi^n_{\mathbf{z},j}\rangle - |\mathbf{z}\rangle \| \leq \sqrt{2} \frac{|\mathbf{z}|^{\tilde{m}}}{\sqrt{\tilde{m}!}} \wedge 2, \tag{8.53}$$

which implies

$$\|\tilde{\rho}^{\mathbf{u}}_{j,n} - \phi^{\mathbf{u}}\|_1 \leq \frac{\sqrt{2}}{\sqrt{\pi s^2}} \int e^{-|\mathbf{z}|^2/2s^2} \left(\frac{|\mathbf{z}+\sqrt{2\mu-1}\alpha_{\mathbf{u}}|^{\tilde{m}}}{\sqrt{\tilde{m}!}} \wedge \sqrt{2} \right) d^2\mathbf{z} = O(e^{-n^{2(\eta+\epsilon)}}),$$

for any $\epsilon > 0$, for any $\gamma \geq 2(\eta+\epsilon)$. Indeed we can split the integral into two parts. The integral over the domain $|\mathbf{z}| \geq n^{\eta+\epsilon}$ is dominated by the Gaussian factor and is $O(e^{-n^{2(\eta+\epsilon)}})$. The integral over the disk $|\mathbf{z}| \leq n^{\eta+\epsilon}$ is bounded by supremum of (8.53) since the Gaussian integrates to one, and is $O(e^{-(\gamma/2-\eta-\epsilon)n^{\gamma}})$. In the last step we use Stirling's formula to obtain $\log\left[(n^{\eta+\epsilon})^{n^{\gamma}}/\sqrt{n^{\gamma}!}\right] \approx (\eta+\epsilon-\gamma/2)n^{\gamma}\log n$. Note that the estimate is uniform with respect to $\mu - 1/2 > \epsilon_2$ for any fixed $\epsilon_2 > 0$.

Second estimate. We now compare the evolved qubits state $\tilde{\rho}^{\mathbf{u}}_{j,n}(t)$ and the evolved oscillator state $\phi^{\mathbf{u}}(t)$. Let $|\psi^n_{m,j}(t)\rangle = U_{j,n}(t)\,|m\rangle \otimes |\Omega\rangle$ be the joint state at time t when the initial state of the system is $|m\rangle$ corresponding to $|j, j-m\rangle$ in the L_z basis notation. We choose the following approximation of $|\psi^n_{m,j}(t)\rangle$

$$|\xi^n_{m,j}(t)\rangle := \sum_{i=0}^{m} c_n(m,i)\alpha_i(t)|m-i\rangle \otimes |e^{-1/2u}\chi_{[0,t]}(u)\rangle_i, \tag{8.54}$$

where $\alpha_i(t) = \exp((-m+i)t/2)$, $c_n(m,i) := c_n(m,i-1)\sqrt{\frac{2j-m+i}{2j_n}}\sqrt{\frac{m-i+1}{i}}$ with $c_n(m,0) := 1$, and $|f\rangle_n := f^{\otimes n}$ as defined in (8.17). In particular for $\mu - 1/2 > \epsilon_2$ and $j \in \mathcal{J}_n$ we have $c_n(m,i) \leq \sqrt{\binom{m}{i}(1+\frac{2}{\epsilon_2}n^{-1/2+\epsilon})^i}$.

We apply now the estimate (8.21). By direct computations we get

$$\begin{aligned} d|\xi^n_{m,j}(t)\rangle &= -\frac{1}{2}\sum_{i=0}^{m} c_n(m,i)\alpha_i(t)(m-i)|m-i\rangle \otimes |e^{-1/2u}\chi_{[0,t]}(u)\rangle_i dt \\ &+ \sum_{i=1}^{m} c_n(m,i)\alpha_{i-1}(t)|m-i\rangle \otimes |e^{-1/2u}\chi_{[0,t]}(u)\rangle_{i-1} \otimes_s |\chi_{[t,t+dt]}\rangle \quad (8.55) \end{aligned}$$

where

$$f^{\otimes i} \otimes_s g := \sum_{k=1}^{i+1} f \otimes f \otimes \cdots \otimes g \otimes \cdots \otimes f.$$

From the quantum stochastic differential equation we get

$$\begin{aligned} &G_{dt}\,|\xi^n_{m,j}(t)\rangle = \\ &-\frac{1}{2}\sum_{i=0}^{m} c_n(m,i)\alpha_i(t)(m-i)\frac{2j-m+i+1}{2j_n}|m-i\rangle \otimes |e^{-1/2u}\chi_{[0,t]}(u)\rangle_i dt \\ &+\sum_{i=0}^{m} c_n(m,i)\alpha_i(t)\sqrt{\frac{(m-i)(2j-m+i+1)}{2j_n(i+1)}}|m-i-1\rangle\otimes \\ &\qquad |e^{-1/2u}\chi_{[0,t]}(u)\rangle_i \otimes_s |\chi_{[t,t+dt]}\rangle. \quad (8.56) \end{aligned}$$

In the second term of the right side of (8.56) we can replace $c_n(m,i)\sqrt{\frac{(m-i)(2j-m+i+1)}{2j_n(i+1)}}$ by $c_n(m,i+1)$ and thus we obtain the same sum as in the second term of the left side of (8.55). Thus

$$\begin{aligned} &G_{dt}|\xi^n_{m,j}(t)\rangle - d|\xi^n_{m,j}(t)\rangle = \\ &\frac{1}{2}\sum_{i=0}^{m-1} c_n(m,i)\alpha_i(t)(m-i)\frac{2(j_n-j)+m-i-1}{2j_n}|m-i\rangle \otimes |e^{-1/2u}\chi_{[0,t]}(u)\rangle_i\, dt. \end{aligned}$$

Then using $c_n(m,i) \leq \sqrt{\binom{m}{i}(1+(2/\epsilon_2)n^{-1/2+\epsilon})^i}$ we get that $\|G_{dt}\xi^n_{m,j}(t) - d\xi^n_{m,j}(t)\|$ is bounded from above by

$$\frac{1}{2}\left[\sum_{i=0}^{m-1}\binom{m}{i}\frac{((1+n^{-1/2+\epsilon})(1-e^{-t}))^i}{e^{(m-i)t}}\left(\frac{(2(j_n-j)+m-i-1)(m-i)}{2j_n}\right)^2\right]^{1/2} dt.$$

We have

$$\frac{(2(j_n - j) + m - i - 1)(m - i)}{2j_n} = O(m(n^{-1/2+\epsilon} + n^{-1}m))$$

Inside the sum we recognize the binomial terms with the m'th term missing. Thus the sum is

$$\begin{aligned}&\left(1 + n^{-1/2+\epsilon} - e^{-t}n^{-1/2+\epsilon}\right)^m - \left((1 - e^{-t})(1 + n^{-1/2+\epsilon})\right)^m \\ &\leq (1 + n^{-1/2+\epsilon})^m(1 - (1 - e^{-t})^m) \leq (1 + n^{-1/2+\epsilon})^m me^{-t}.\end{aligned}$$

Then there exists a constant C (independent of μ if $\mu - 1/2 \geq \epsilon_2$) such that

$$\|G_{dt}\xi^n_{m,j}(t) - d\xi^n_{m,j}(t)\| \leq \frac{C}{2}e^{-t/2}m^{3/2}(n^{-1/2+\epsilon} + mn^{-1})\left(1 + \frac{2}{\epsilon_2}n^{-1/2+\epsilon}\right)^{m/2}$$

By integrating over t we finally obtain

$$\|\psi^n_{m,j}(t) - \xi^n_{m,j}(t)\| \leq Cm^{3/2}(n^{-1/2+\epsilon} + mn^{-1})\left(1 + \frac{2}{\epsilon_2}n^{-1/2+\epsilon}\right)^{m/2}. \tag{8.57}$$

Note that under the assumption $\gamma < 1/3 - 2\epsilon/3$, the right side converges to zero at rate $n^{3\gamma/2-1/2+\epsilon}$ for all $m \leq \tilde{m} = n^\gamma$. Summarizing, the assumptions which we have made so far over γ are

$$2\eta + 2\epsilon < \gamma < 1/3 - 2\epsilon/3.$$

Now consider the vector $|\psi^n_{\mathbf{z},j}\rangle$ as defined in (8.52) and let us denote $|\psi^n_{\mathbf{z},j}(t)\rangle = U_{j,n}(t)|\psi^n_{\mathbf{z},j}\rangle \otimes |\Omega\rangle$. Then based on (8.54) we choose the approximate solution

$$|\xi^n_{\mathbf{z},j}(t)\rangle = e^{-|\mathbf{z}|^2/2}\sum_{m=0}^{\tilde{m}}\frac{|\mathbf{z}|^m}{\sqrt{m!}}\sum_{i=0}^{m}c_n(m,i)\alpha_i(t)|m-i\rangle \otimes |e^{-1/2u}\chi_{[0,t]}(u)\rangle_i.$$

Note that the vectors $|\psi^n_{k,j}(t)\rangle$ and $|\xi^n_{k,j}(t)\rangle$ live in the "k-particle" subspace of $\mathcal{H}_j \otimes \mathcal{F}(L^2(\mathbb{R}))$ and thus are orthogonal to all vectors $|\psi^n_{p,j}(t)\rangle$ and $|\xi^n_{p,j}(t)\rangle$ with $p \neq k$. By (8.57), the error is

$$\begin{aligned}&\|\psi^n_{\mathbf{z},j}(t) - \xi^n_{\mathbf{z},j}(t)\| \\ &\leq Ce^{-|\mathbf{z}|^2/2}\left(\sum_{m=0}^{\tilde{m}}\frac{|\mathbf{z}|^{2m}}{m!}m^3(n^{-1/2+\epsilon} + mn^{-1})^2\left(1 + \frac{2}{\epsilon_2}n^{-1/2+\epsilon}\right)^m\right)^{1/2} \\ &\quad + \frac{|\mathbf{z}|^{2\tilde{m}}}{\tilde{m}!} \\ &\leq C\tilde{m}^{3/2}(n^{-1/2+\epsilon} + \tilde{m}n^{-1})\left(1 + \frac{2}{\epsilon_2}n^{-1/2+\epsilon}\right)^{\tilde{m}/2} + \frac{|\mathbf{z}|^{2\tilde{m}}}{\tilde{m}!}.\end{aligned} \tag{8.58}$$

We now compare the approximate solution $\xi^n_{\mathbf{z},j}(t)$ with the "limit" solution $\psi_{\mathbf{z}}(t)$ for the oscillator coupled with the field as described in section 8.4.2. We can write

$$\psi_{\mathbf{z}}(t) = e^{-|\mathbf{z}|^2/2} \sum_{m=0}^{\infty} \frac{|\mathbf{z}|^m}{\sqrt{m!}} \sum_{i=0}^{m} \sqrt{\binom{m}{i}} e^{-(m-i)t/2} |m-i\rangle \otimes |e^{-1/2u} \chi_{[0,t]}(u)\rangle_i.$$

Then

$$\|\xi^n_{\mathbf{z},j}(t) - \psi_{\mathbf{z}}(t)\|^2 =$$
$$e^{-|\mathbf{z}|^2} \sum_{m=0}^{\tilde{m}} \frac{|\mathbf{z}|^{2m}}{m!} \sum_{i=0}^{m} e^{-(m-i)t} \left| c_n(m,i) - \sqrt{\binom{m}{i}} \right|^2 (1-e^{-t})^i + e^{-|\mathbf{z}|^2} \sum_{m=\tilde{m}}^{\infty} \frac{|\mathbf{z}|^{2m}}{m!}.$$

Now

$$\begin{aligned} \left| c_n(m,i) - \sqrt{\binom{m}{i}} \right|^2 &\leq \left| c_n(m,i)^2 - \binom{m}{i} \right| \\ &\leq \binom{m}{i} \left| 1 - \prod_{p=1}^{i} \left(1 + \frac{2(j-j_n) - m + p}{2j_n} \right) \right| \\ &\leq C_2 \binom{m}{i} m n^{-1/2+\epsilon}, \end{aligned}$$

where C_2 does not depend on μ as long as $\mu - 1/2 \geq \epsilon_2$ (recall that the dependence in μ is hidden in $j_n = (2\mu-1)n$). Thus

$$\|\xi^n_{\mathbf{z},j}(t) - \psi_{\mathbf{z}}(t)\|^2 \leq C_2 n^{-1/2+\epsilon} e^{-|\mathbf{z}|^2} \sum_{m=0}^{\tilde{m}} \frac{m|\mathbf{z}|^{2m}}{m!} + \frac{|\mathbf{z}|^{2\tilde{m}}}{\tilde{m}!} \leq C_2 n^{-1/2+\epsilon} |\mathbf{z}|^2 + \frac{|\mathbf{z}|^{2\tilde{m}}}{\tilde{m}!}. \tag{8.59}$$

From (8.58) and (8.59) we get

$$\begin{aligned} \|\psi^n_{\mathbf{z},j}(t) - \psi_{\mathbf{z}}(t)\| \leq 2 \wedge \Bigg[& C\tilde{m}^{3/2}(n^{-1/2+\epsilon} + \tilde{m}n^{-1}) \left(1 + \frac{2}{\epsilon_2} n^{-1/2+\epsilon} \right)^{\tilde{m}/2} \\ & + \frac{|\mathbf{z}|^{2\tilde{m}}}{\tilde{m}!} + \left[C_2 n^{-1/2+\epsilon} |\mathbf{z}|^2 + \frac{|\mathbf{z}|^{2\tilde{m}}}{\tilde{m}!} \right]^{1/2} \Bigg] \\ := E(\tilde{m}, n, \mathbf{z}) \end{aligned}$$

We now integrate the coherent states over the displacements $\mathbf{z}$ as we did in the case of local asymptotic normality in order to obtain the thermal states in which we are interested

$$\tilde{\rho}^{\mathbf{u}}_{j,n} := \frac{1}{\sqrt{2\pi s^2}} \int e^{-|\mathbf{z} - \sqrt{2\mu-1}\alpha_{\mathbf{u}}|^2/2s^2} \left(|\psi^n_{\mathbf{z},j}\rangle\langle\psi^n_{\mathbf{z},j}| \right) d^2\mathbf{z}.$$

We define the evolved states

$$\tilde{\rho}^{\mathbf{u}}_{j,n}(t) := U_{j,n}(t)\tilde{\rho}^{\mathbf{u}}_{j,n}U_{j,n}(t)^*, \qquad \text{and} \qquad \phi^{\mathbf{u}}(t) := U(t)\phi^{\mathbf{u}}U(t)^*,$$

Then

$$\sup_{j\in\mathcal{J}_n}\sup_{\|\mathbf{u}\|\leq n^\eta}\|\tilde{\rho}^{\mathbf{u}}_{j,n}(t)-\phi^{\mathbf{u}}(t)\|_1 \leq \sup_{\|\mathbf{u}\|\leq n^\eta}\frac{1}{\sqrt{\pi s^2}}\int e^{-|\mathbf{z}-\sqrt{2\mu-1}\alpha_{\mathbf{u}}|^2/2s^2}E(\tilde{m},n,\mathbf{z})\,d^2\mathbf{z}.$$

Here again we cut the integral in two parts. On $|\mathbf{z}| \geq n^{\eta+\epsilon}$, the Gaussian dominates, and this outer part is less than $e^{-n^{\eta+\epsilon}}$. Now the inner part is dominated by $\sup_{|\mathbf{z}|\leq n^{\eta+\epsilon}} E(\tilde{m},n,\mathbf{z})$. Now we want $\tilde{m}$ to be not too big for (8.58) to be small, on the other hand, we want $\mathbf{z}^{2\tilde{m}}/\tilde{m}!$ to go to zero. A choice which satisfies the condition is $\gamma = 2\eta + 3\epsilon$. By renaming ϵ we then get

$$E(\tilde{m},n,\mathbf{z}) = O(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon}),$$

for any small enough $\epsilon > 0$. Hence we obtain (8.22).

□

Chapitre 9

Quantum local asymptotic normality for d-dimensional states

Ce chapitre dérive de l'article (Guţă et Kahn, 2009).

Résumé : Nous étendons la normalité asymptotique locale quantique forte à tous les systèmes de dimension finie. Comme au Chapitre 7, nous considérons les états de la forme $\rho_{\theta/\sqrt{n}}^{\otimes n}$, et exigeons que ρ_0 ait des valeurs propres différentes deux à deux. Nous construisons ensuite des canaux depuis et vers la famille limite. Cette famille limite est un produit d'une expérience de décalage gaussienne classique et d'une expérience de décalage gaussienne quantique, les états de cette dernière étant plus précisément le produit d'états thermiques déplacés dont la temperature ne dépend pas du paramètre θ. De plus, nous autorisons l'espace de paramètres à croître, et obtenons des vitesses de convergence polynomiales.

La preuve exige un travail très technique sur les tableaux de Young, et utilise un résultat intermédiaire intéressant en lui-même : la base générée par les tableaux de Young semi-standards d'une représentation de $SU(d)$ est «presque» orthonormale.

En application, nous mentionnons une méthode d'estimation asymptotiquement optimale. Nous établissons au cours de la preuve un théorème de représentation asymptotique quantique et un théorème minimax asymptotique quantique.

9.1 Introduction

Quantum statistics deals with problems of statistical inference arising in quantum mechanics. The first significant results in this area appeared in the seventies and tackled issues such as quantum Cramér-Rao bounds for unbiased estimators, optimal estimation for families of states possessing a group symmetry, estimation of Gaussian states, optimal discrimination between non-commuting states. It is impossible to list all contributions but the following references may give the flavour of these developments (Helstrom, 1969; Yuen et Lax, M., 1973; Yuen et al., 1975a; Belavkin, 1975, 1976; Holevo, 1982). The more recent theoretical advances (Hayashi, 2005b, 2006; Paris et Řeháček, 2004; Barndorff-Nielsen et al., 2003; Artiles, L et al., 2005; Audenaert et al.) are closely related to the rapid development of quantum information and quantum engineering, and are often accompanied by practical implementations (Armen et al., 2002; Hannemann et al., 2002a; Smith et al., 2006; Schiller et al., 1996).

An important topic in quantum statistics is that of optimal estimation of an unknown state using the results of measurements performed on n identically prepared quantum systems (Massar et Popescu, 1995; Cirac et al., 1999; Vidal et al., 1999; Gill et Massar, 2000; Keyl et Werner, 2001; Bagan et al., 2002; Hayashi et Matsumoto, 2004, 2005; Bagan et al., 2006; Gill, 2005a). In the case of two dimensional systems, or qubits, the problem has been solved explicitly in the *Bayesian set-up*, in the particular case of an invariant prior and figure of merit based on the fidelity distance between states (Bagan et al., 2006). However the method used there does not work for more general priors, loss functions, or higher dimensions. In the *pointwise approach*, Hayashi et Matsumoto (2004) have shown that the Holevo (1982) bound for the variance of locally unbiased estimators can be achieved asymptotically, and provided a sequence of measurements with this property. Their results, building on earlier work (Hayashi, 2003; Hayashi), indicate for the first time the emergence of a Gaussian limit in the problem of optimal state estimation for qubits. The extension to d-dimensional case is analysed by Matsumoto.

We (Guţă et Kahn, 2006; Guţă et al., 2008) performed a detailed analysis of this phenomenon (again for qubits), and showed that we deal with the quantum generalization of an important concept in mathematical statistics called *local asymptotic normality*. As a corollary, we devised a two steps adaptive measurement strategy for state estimation which is asymptotically optimal for a large class of loss functions and priors, and could be practically implemented using continuous-time measurements. In 'classical statistics', the idea of approximating a sequence of statistical models by a family of Gaussian distributions was first formulated by Wald (1950), and was fully developed by Le Cam (1986) who coined the term "local asymptotic normality". Among the many applications we mention its role in asymptotic opti-

mality theory and in proving the asymptotic normality of certain estimators such as the maximum likelihood estimator. The aim of this chapter is to extend our previous results (Guţă et Kahn, 2006; Guţă et al., 2008) to systems of *arbitrary dimension* $d < \infty$, and solve the *open problem* of optimal state estimation for d-dimensional quantum systems.

Before stating the main result of the chapter we shall explain briefly the meaning of local asymptotic normality for two dimensional systems (Guţă et Kahn, 2006; Guţă et al., 2008). We are given n qubits identically prepared in an unknown state ρ. Asymptotic normality means that for large n we can encode the statistical information contained in the state $\rho^{\otimes n}$ into a Gaussian model consisting of a classical random variable with distribution $N(u, I^{-1})$, and a quantum harmonic oscillator prepared in a (Gaussian) displaced thermal state ϕ_ζ. The term *local* refers to how ρ is related to the parameters $\theta = (u, \zeta)$, as explained below.

For a more precise formulation let us parametrise the qubit states by their Bloch vectors $\rho(\vec{r}) = \frac{1}{2}(\mathbf{1} + \vec{r}\,\vec{\sigma})$ where $\vec{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$ are the Pauli matrices. The neighbourhood of the state ρ_0 with $\vec{r_0} = (0, 0, 2\mu - 1)$ and $1/2 < \mu < 1$, is a three-dimensional ball parametrised by the deviation $u \in \mathbb{R}$ of diagonal elements and $\zeta \in \mathbb{C}$ of the off-diagonal ones

$$\rho_\theta = \begin{pmatrix} \mu + u & \zeta^* \\ \zeta & 1 - \mu - u \end{pmatrix}, \qquad \theta = (u, \zeta) \in \mathbb{R} \times \mathbb{C}.$$

Note that ρ_0 is to be considered fixed and known but otherwise arbitrary, and can be taken to be diagonal without any loss of generality. Consider now n identically prepared qubits whose individual states are in a neighbourhood of ρ_0 of size $1/\sqrt{n}$, so that their joint state is $\rho_\theta^n := \left[\rho_{\theta/\sqrt{n}}\right]^{\otimes n}$ for some unknown θ. We would like to understand the structure of the family (statistical experiment)

$$\mathcal{Q}_n := \{\rho_\theta^n : \|\theta\| \leq C\}, \tag{9.1}$$

as a whole, more precisely what is its asymptotic behavior as $n \to \infty$?

For this we consider a quantum harmonic oscillator with position and momentum operators satisfying the commutation relations $[\mathbf{Q}, \mathbf{P}] = i\mathbf{1}$. We denote by $\{|k\rangle, k \geq 0\}$ the eigenbasis of the number operator and define the thermal equilibrium state

$$\phi = (1 - e^{-\beta}) \sum_{k=0}^{\infty} e^{-k\beta} |k\rangle\langle k|, \qquad e^{-\beta} = \frac{1 - \mu}{\mu},$$

which has centered Gaussian distributions for both $\mathbf{Q}$ and $\mathbf{P}$ with variance $1/(4\mu - 2) > 1/2$. We define a family of displaced thermal equilibrium states

$$\phi^\zeta := D^\zeta(\phi) := W(\zeta/\sqrt{2\mu - 1})\, \phi\, W(\zeta/\sqrt{2\mu - 1})^*, \tag{9.2}$$

where $W(\zeta) := \exp(\zeta a^* - \zeta a)$ is the unitary displacement operator with $\zeta \in \mathbb{C}$. Additionally we consider a classical *Gaussian shift model* consisting of the family of normal distributions $N(u, \mu(1-\mu))$ with unknown center u and fixed known variance. The classical-quantum statistical experiment to which we alluded above is defined by the family of densities

$$\mathcal{R} := \{\phi^\theta := \mathcal{N}(u, \mu(1-\mu)) \otimes \phi^\zeta : \|\theta\| \leq C\} \tag{9.3}$$

where the unknown parameters $\theta = (u, \zeta) \in \mathbb{R} \times \mathbb{C}$ are the same as those of $\mathcal{Q}_n$.

Theorem 9.1.1. *(Guţă et Kahn, 2006; Guţă et al., 2008) Let $\mathcal{Q}_n$ be the quantum statistical experiment* (9.1) *and let $\mathcal{R}$ be the classical-quantum experiment* (9.3). *Then for each n there exist quantum channels (normalized completely positive maps)*

$$\begin{aligned} T_n &: M\left((\mathbb{C}^2)^{\otimes n}\right) \to L^1(\mathbb{R}) \otimes \mathcal{T}(L^2(\mathbb{R})), \\ S_n &: L^1(\mathbb{R}) \otimes \mathcal{T}(L^2(\mathbb{R})) \to M\left((\mathbb{C}^2)^{\otimes n}\right), \end{aligned}$$

with $\mathcal{T}(L^2(\mathbb{R}))$ the trace-class operators, such that

$$\begin{aligned} \lim_{n\to\infty} \sup_{\|\theta\|\leq C} \|\phi_\theta - T_n\left(\rho_\theta^n\right)\|_1 = 0, \\ \lim_{n\to\infty} \sup_{\|\theta\|\leq C} \|\rho_\theta^n - S_n\left(\phi_\theta\right)\|_1 = 0, \end{aligned}$$

for an arbitrary constant $C > 0$. The norm on trace class operators is $\|\tau\|_1 := \mathrm{Tr}(|\tau|)$.

The theorem shows that from a statistical point of view the joint qubits states are asymptotically indistinguishable from the limit Gaussian system. At the first sight one might object that the local nature of the result prevents us from drawing any conclusions for the original model of a completely unknown state ρ. However this is not a limitation, but reflects the correct normalisation of the parameters with $n \to \infty$. Indeed as n grows we have more information about the state which can be pinned down to a region of size slightly larger that $1/\sqrt{n}$ by performing rough measurements on a small proportion of the systems. After this 'localisation' step, we can use more sophisticated techniques to better estimate the state within the local neighbourhood of the first step estimator, and it is here where we use the local asymptotic normality result. Indeed, since locally the states are uniformly close to displaced Gaussian states we can pull back the optimal (heterodyne) measurement for estimating the latter to get an asymptotically optimal measurement for the former. Based on this insight we have proposed a realistic measurement set-up for this purpose using an atom-field interaction and continuous measurements in the field (Guţă et al., 2008).

This chapter deals with the extension of the previous result to d-dimensional systems. Like in the two-dimensional case we parametrise the neighbourhood of a fixed (diagonal) state ρ_0 by a vector $\vec{u} \in \mathbb{R}^{d-1}$ of diagonal parameters and $d(d-1)/2$ complex parameters $\vec{\zeta} = (\zeta_{j,k} : j < k)$, one for each off-diagonal matrix element (cf. (9.15) and (9.17)). We consider the same $1/\sqrt{n}$–scaling and look at the family

$$\mathcal{Q}_n = \left\{ \left[\rho_{\theta/\sqrt{n}}\right]^{\otimes n} : \theta = (\vec{u}, \vec{\zeta}) \in \Theta_n \subset \mathbb{R}^{d-1} \otimes \mathbb{C}^{d(d-1)/2} \right\},$$

where Θ_n is a ball of local parameters whose size is allowed to grow slowly with n.

As in the 2-dimensional case, the limit model is the product of a classical statistical model depending on the parameters $\vec{u}$ and a quantum model depending on $\vec{\zeta}$. Moreover the quantum part splits into a tensor product of displaced thermal states of quantum oscillators, one for each off-diagonal matrix element $\zeta_{j,k}$ with $j < k$. Thus

$$\phi^\theta = \mathcal{N}(\vec{u}, I_{\rho_0}^{-1}) \otimes \bigotimes_{j<k} \phi_{j,k}^{\zeta_{j,k}}, \qquad \theta = (\vec{u}, \vec{\zeta}).$$

Here, I_{ρ_0} is the Fisher information matrix of the multinomial model with parameters $(\mu_1, \ldots, \mu_d)$ described in Example 9.3.1, and $\phi_{j,k}^{\zeta_{j,k}}$ is the displaced thermal equilibrium state defined in (9.2) with inverse temperature $\beta = \ln(\mu_j/\mu_k)$.

Theorem 9.4.3 is the main result of the chapter and shows the convergence of $\mathcal{Q}_n$ to the Gaussian model

$$\mathcal{R}_n = \left\{ \phi^\theta : \theta \in \Theta_n \subset \mathbb{R}^{d-1} \otimes \mathbb{C}^{d(d-1)/2} \right\},$$

in the spirit of Theorem 9.1.1. On the technical side, the uniform convergence holds over local neighbourhoods Θ_n which are allowed to grow with n rather that being fixed balls. This is essential for constructing the two stage optimal measurement: first localise within a neighbourhood Θ_n, and then apply the optimal Gaussian measurement. The details of this construction are similar to the two dimensional case and are given in section 9.4.5.

Despite the similarity to the two dimensional case, the proof of the d-dimensional result has additional features which may be responsible for the fact that the optimal estimation problem has remained unsolved until now. The proof is based on the following observations:

- the n systems space $(\mathbb{C}^d)^{\otimes n}$ decomposes into a direct sum of irreducible representations of $SU(d)$, each representation being labelled by a Young diagram λ (cf. Theorem 9.4.1);

- the joint state $\rho_{\theta/\sqrt{n}}^{\otimes n}$ has the block diagonal form (9.21), the block weights $\lambda \to p_\lambda^{\theta,n}$ depend only on the diagonal parameters $\vec{u}$ and are closely related to the multinomial distribution of Example 9.3.1. This classical statistical model converges to the $(d-1)$-dimensional Gaussian shift model $N(\vec{u}, I_{\rho_0}^{-1})$;

- there exists an isometry V_λ mapping basis vectors $|\mathbf{m}, \lambda\rangle$ of the irreducible representation $\mathcal{H}_\lambda$ *almost* into number vectors $|\mathbf{m}\rangle$ of the multimode Fock space, where $\mathbf{m} = \{m_{j,k} : j < k\}$ is the collection of eigenvalues of the number operators for all oscillators.

- given a typical λ, the conditional block-state $\rho_\lambda^{\theta,n}$ can be mapped with V_λ into a multimode state which is close (in trace norm) to the Gaussian product state $\otimes_{j<k}\phi_{j,k}^{\zeta_{j,k}}$. This can be done *uniformly* over the typical diagrams whose normalised shapes have $1/\sqrt{n}$ fluctuations around $(\mu_1, \mu_2, \ldots, \mu_d)$, and over parameters $\theta \in \Theta_n$.

The first item is the well known Weyl duality which is extensively used in quantum statistics for i.i.d. states. The probability distribution of the second point has also been analysed the context of large deviations (Keyl et Werner, 2001) for the estimation of the state eigenvalues. The third point shows that the basis $|\mathbf{m}, \lambda\rangle$ is almost orthogonal for indices $\mathbf{m}$ which are not too big. This basis is obtained by projecting tensors of the form $f_\mathbf{a} := f_{a(1)} \otimes \cdots \otimes f_{a(n)}$ onto a subspace of $(\mathbb{C}^d)^{\otimes n}$ which is isomorphic to $\mathcal{H}_\lambda$ (cf. Theorem 9.5.2). Let us place the indices $\{a(i) : i = 1 \ldots n\}$ in the boxes of the diagram λ along rows, starting from the left end of the first row, to obtain a tableau $t_\mathbf{a}$. It turns out that we only need to consider $f_\mathbf{a}$ for which $t_\mathbf{a}$ is a semistandard tableau (nondecreasing along rows, increasing along columns). Then the label $\mathbf{m} := \{m_{i,j} : j > i\}$ is the collection of integers $m_{i,j}$ equal to the number of $j's$ on the row i, and is in one to one correspondence with $\mathbf{a}$. The following is an example of such semistandard tableau

$$t_\mathbf{m} = \begin{array}{|c|c|c|c|c|c|c|c|c|c|c|c|c|}\hline 1&1&1&1&1&1&1&1&2&2&2&3&3\\\hline 2&2&2&2&2&3\\\cline{1-6} 3&3&3\\\cline{1-3}\end{array}, \qquad \text{with } m_{1,2} = 3, m_{1,3} = 2, m_{2,3} = 1.$$

The relatively large number of i's in the row i is intentional, since it turns out that the 'relevant' vectors, i.e. those carrying the states $\rho_\lambda^{\theta,n}$, have indices $m_{i,j}$ small compared with the length of the rows ($\lambda_i \approx n\mu_i$ for typical representations λ). More precisely, in section 9.7.3 we prove the following quasi-orthogonality result which allows us to carry the block states over to the oscillator space: if $\mathbf{m} \neq \mathbf{l}$ and $|\mathbf{l}| \leq |\mathbf{m}| \leq n^\eta$ then

$$|\langle \mathbf{m}, \lambda | \mathbf{l}, \lambda \rangle| = O(n^{(9\eta-2)|\mathbf{m}-\mathbf{l}|/12}) \xrightarrow[n\to\infty]{} 0 \qquad \text{for } \eta < 2/9.$$

The proof of the fourth point involves a detailed analysis of the state $\rho_\lambda^{\theta,n}$ through its coefficients in the basis $|\mathbf{m}, \lambda\rangle$ of $\mathcal{H}_\lambda$. When $\theta = 0$ the state is diagonal and its coefficients approach uniformly those of the multidimensional thermal state $\phi^{\vec{0}} = \otimes_{j<k}\phi_{j,k}$ as shown in Lemma 9.6.3. The next step is to apply $SU(d)$ rotations and obtain the states $\rho_\lambda^{\theta,n}$. In Lemmas 9.6.4 and 9.6.5 it is shown that the unitary operations $\mathrm{Ad}[U_\lambda(\zeta/\sqrt{n})]$ act on $\rho_\lambda^{0,n}$ in the same way as the displacement operator $D^{\vec{\zeta}}$ acts on the thermal state $\phi^{\vec{0}}$. A remarkable fact is that in the limit the different off-diagonal parameters 'separate' into a product of shift experiments for quantum oscillators, one for each off-diagonal index $(j < k)$. This could be guessed from the Quantum Central Limit Theorem 9.4.6 which is related to the restriction of our result to $\theta = 0$.

Due to the apparent intricacy of the main result, the chapter is organised according to the 'onion peeling' principle. We start in section 9.2 with general classical statistical notions which motivate our investigation in quantum statistics. In particular we explain the relevance of the Le Cam distance between statistical models as a statistically meaningful way to describe convergence. Section 9.3 presents the classical version of local asymptotic normality with the multinomial model as example.

In section 9.4 we introduce the quantum statistical model consisting of n identical quantum systems with joint state $\rho^{\theta,n}$ described by diagonal and rotation parameters. We also introduce the multimode Gaussian states appearing in the limit. With this we can formulate the main result, Theorem 9.4.3. With the theorem, we immediately make explicit a two stage adaptive measurement strategy which is asymptotically optimal for both Bayesian and pointwise viewpoints, and for a large range of 'distances' on the state space, in 9.4.4 and below.

In section 9.5 we introduce the basis $|\mathbf{m}, \lambda\rangle$ and the isometry V_λ allowing us to define the channels T_n and S_n connecting the two statistical models.

In section 9.6 we break the proof of the main theorem into manageable lemmas, essentially by using triangle inequalities. Each lemma deals with a different aspect of the convergence and has an interest in its own.

Finally, the technical proofs are collected in section 9.7. Notably, subsection 9.7.2 and Lemma 9.7.11 contain the combinatorial substance of the chapter. Moreover, in the course of proving Theorem 9.4.4, subsection 9.7.1 contains important equivalents of the classical asymptotic representation and asymptotic minimax theorems.

Our investigation relies on the theory of representations of $SU(d)$. We refer to the books by Fulton (1997); Goodman R. et Wallach N.R. (1998); Fulton et Harris (1991) for proofs of standard results and more details.

Throughout, we will use the following symbols: φ, ψ for states, ϕ, ρ for density

matrices, T, S, M for channels (randomisations), $\mathcal{E}, \mathcal{P}, \mathcal{Q}$ for statistical models, θ, ζ, u for parameters, $\alpha, \beta, \gamma, \delta, \epsilon$ for positive constants, λ for Young diagrams.

9.2 Classical and quantum statistical experiments

In this section we introduce some basic notions from classical statistics with the aim of defining the Le Cam distance between statistical models and local asymptotic normality. In parallel, we shall define the quantum analogues and point out their relevance in quantum statistics. The reader may find the conceptual framework helpful in understanding the quantum version of the result, but otherwise the section can be skipped at the first reading.

Let X be a random variable with values in the measure space $(\mathcal{X}, \Sigma_{\mathcal{X}})$, and let us assume that its probability distribution P belongs to some family $\{P_\theta : \theta \in \Theta\}$ where the parameter θ is unknown. Statistical inference deals with the question of how to use the available data X in order to draw conclusions about some property of θ. We shall call the family

$$\mathcal{E} := \{P_\theta : \theta \in \Theta\}, \tag{9.4}$$

a *statistical experiment or statistical model* over $(\mathcal{X}, \Sigma_{\mathcal{X}})$ (Le Cam, 1986).

In quantum statistics the data is replaced by a quantum system prepared in a state φ which belongs to a family $\{\varphi_\theta : \theta \in \Theta\}$ of states over an algebra of observables. In order to make a statistical inference about θ one first has to measure the system, and then apply statistical techniques to draw conclusions from the data consisting of the measurement outcomes. An important difference with the classical case is that the experimenter has the possibility to choose the measurement set-up M, and each set-up will lead to a different classical model $\{P_\theta^{(M)} : \theta \in \Theta\}$, where $P_\theta^{(M)}$ is the distribution of outcomes when performing the measurement M on the system prepared in state φ_θ.

The guiding idea of this chapter is to investigate the structure of the family of quantum states

$$\mathcal{Q} := \{\varphi_\theta : \theta \in \Theta\},$$

which will be called a *quantum statistical experiment.* We shall show that in an important asymptotic set-up, namely that of a large number of identically prepared systems, the joint state can be approximated by a multidimensional quantum Gaussian state, for *all* possible preparations of the individual systems. This will bring a drastic simplification in the problem of optimal estimation for d-dimensional quantum systems, which can then be solved in the asymptotic framework, in section 9.4.5.

9.2.1 Classical and quantum randomisations

Any statistical decision (e.g. estimator, test) can be seen as data processing using a *Markov kernel*. Suppose we are given a random variable X taking values in $(\mathcal{X}, \Sigma_{\mathcal{X}})$ and we want to produce a 'decision' $y \in \mathcal{Y}$ based on the data X. The space $\mathcal{Y}$ may be for example the parameter space Θ in the case of estimation, or just the set $\{0,1\}$ in the case of testing between two hypotheses. For every value $x \in \mathcal{X}$ we choose y randomly with probability distribution given by $K_x(dy)$. Assuming that $K : \mathcal{X} \times \Sigma_{\mathcal{Y}} \to [0,1]$ is measurable with respect to x for all fixed $A \in \Sigma_{\mathcal{Y}}$, we can regard K as a map from probability distributions over $(\mathcal{X}, \Sigma_{\mathcal{X}})$ to probability distributions over $(\mathcal{Y}, \Sigma_{\mathcal{Y}})$ with

$$K(P)(A) = \int K_x(A)P(dx), \quad A \in \Sigma_{\mathcal{Y}}. \tag{9.5}$$

A *statistic* $S : \mathcal{X} \to \mathcal{Y}$ is a particular example of such a procedure, where K_x is simply the delta measure at $S(x)$.

Besides statistical decisions, there is another important reason why one would like to apply such treatment to the data, namely to summarize it in a more convenient and informative way for future purposes as illustrated in the following simple example. Consider n independent identically distributed random variables $X_1, \ldots, X_n$ with values in $\{0,1\}$ and distribution $P_\theta := (1-\theta, \theta)$ with $\theta \in \Theta := (0,1)$. The associated statistical experiment is

$$\mathcal{E}_n := \{P_\theta^n : \theta \in \Theta\}.$$

It is easy to see that $\bar{X}_n = \frac{1}{n}\sum_{i=1}^n X_i$ is an unbiased estimator of θ and moreover it is a *sufficient statistic* for $\mathcal{E}_n$, *i.e.* the conditional distribution $P_\theta^n(\cdot|\bar{X}_n = \bar{x})$ does not depend on θ! In other words the dependence on θ of the total sample $(X_1, X_2, \ldots, X_n)$ is completely captured by the statistic $\bar{X}_n$ which can be used as such for any statistical decision problem concerning $\mathcal{E}_n$. If we denote by $\bar{P}_\theta^n$ the distribution of $\bar{X}_n$ then the experiment

$$\bar{\mathcal{E}}_n = \{\bar{P}_\theta^n : \theta \in \Theta\},$$

is statistically equivalent to $\mathcal{E}_n$. To convince ourselves that $\bar{X}_n$ does contain the same statistical information as $(X_1, \ldots, X_n)$, we show that we can obtain the latter from the former by means of a randomised statistic. Indeed for every fixed value $\bar{x}$ of $\bar{X}_n$ there exists a measurable function

$$f_{\bar{x}} : [0,1] \to \{0,1\}^n,$$

such that the distribution of $f_{\bar{x}}(U)$ is $P_\theta^n(\cdot|\bar{X}_n = \bar{x})$. In other words

$$\lambda(f_{\bar{x}}^{-1}(x_1, \ldots, x_n)) = P_\theta^n(x_1, \ldots, x_n|\bar{X}_n = \bar{x}),$$

where λ is the Lebesgue measure on $[0,1]$. Then $F(\bar{X}_n, U) := f_{\bar{X}_n}(U)$, has distribution P_θ^n. To summarize, statistics, randomised statistics and Markov kernels, are ways to transform the available data for a specific purpose. The Markov kernel K defined in (9.5) maps the experiment $\mathcal{E}$ of equation (9.4) into the experiment

$$\mathcal{F} := \{Q_\theta : \theta \in \Theta\},$$

over $(\mathcal{Y}, \Sigma_\mathcal{Y})$ with $Q_\theta = K(P_\theta)$. For mathematical convenience it is useful to represent such transformations in terms of linear maps between linear spaces.

Definition 9.2.1. *A positive linear map*

$$T_* : L^1(\mathcal{X}, \Sigma_\mathcal{X}, P) \to L^1(\mathcal{Y}, \Sigma_\mathcal{Y}, Q)$$

is called a stochastic operator or transition if $\|T_*(g)\|_1 = \|g\|_1$ *for every* $g \in L^1_+(\mathcal{X})$.

Definition 9.2.2. *A positive linear map*

$$T : L^\infty(\mathcal{Y}, \Sigma_\mathcal{Y}, Q) \to L^\infty(\mathcal{X}, \Sigma_\mathcal{X}, P)$$

is called a Markov operator if $T\mathbf{1} = \mathbf{1}$, *and if for any* $f_n \downarrow 0$ *in* $L^\infty(\mathcal{Y})$ *we have* $Tf_n \downarrow 0$.

A pair (T_*, T) as above is called a dual pair if

$$\int fT(g)dP = \int T_*(f)gdQ,$$

for all $f \in L^1(\mathcal{X}, \Sigma_\mathcal{X}, P)$ and $g \in L^\infty(\mathcal{Y}, \Sigma_\mathcal{Y}, Q)$. It is a theorem that for any stochastic operator T_* there exists a unique dual Markov operator T and vice versa.

What is the relation between Markov operators and Markov kernels ? Roughly speaking, any Markov kernel defines a Markov operator when we restrict to families of dominated probability measures. Let us assume that all distributions P_θ of the experiment $\mathcal{E}$ defined in (9.4) are absolutely continuous with respect to a fixed probability distribution P, such that there exist densities $p_\theta := dP_\theta/dP : \mathcal{X} \to \mathbb{R}_+$. Such an experiment is called *dominated* and in concrete situations this condition is usually satisfied. Let $K_x(dy)$ be a Markov kernel (9.5) such that $Q_\theta = K(P_\theta)$, then we define associated Markov operator $(T(f))(x) := \int f(y)k_x(dy)$ and have

$$Q_\theta = P_\theta \circ T, \qquad \forall\theta. \tag{9.6}$$

When the probability distributions of two experiments are related to each other as in (9.6), we say that $\mathcal{F}$ is a *randomisation* of $\mathcal{E}$. From the duality between T

and T_* we obtain an equivalent characterization in terms of the stochastic operator $T_* : L^1(\mathcal{X}, \Sigma_\mathcal{X}, P) \to L^1(\mathcal{Y}, \Sigma_\mathcal{Y}, Q)$ such that

$$T_*(dP_\theta/dP) = dQ_\theta/dQ, \qquad \forall\theta\,.$$

The concept of randomisation is weaker than that of Markov kernel transformation, but under the additional condition that $(\mathcal{Y}, \Sigma_\mathcal{Y})$ is locally compact space with countable base and Borel σ-field, it can be shown that any randomisation can be implemented by a Markov kernel (Strasser, 1985).

What is the analogue of randomisations in the quantum case ? In the language of operator algebras $L^\infty(\mathcal{X}, \Sigma_\mathcal{X}, P)$ is a commutative von Neumann algebra and $L^1(\mathcal{X}, \Sigma_\mathcal{X}, P)$ is the space of (densities of) *normal* linear functionals on it. The stochastic operator T_* is the classical version of *quantum channel*, *i.e.* a completely positive normalized (trace-preserving) map

$$T_* : \mathcal{A}_* \to \mathcal{B}_*$$

where $\mathcal{A}_*, \mathcal{B}_*$ are the spaces of normal states on the von Neumann algebra $\mathcal{A}$ and respectively $\mathcal{B}$. Any normal state φ on $\mathcal{A}$ has a density ρ with respect to the trace such that $\varphi(A) = \mathrm{Tr}(\rho A)$ for all $A \in \mathcal{A}$. The dual of T_* is

$$T : \mathcal{B} \to \mathcal{A},$$

which is a unital completely positive map and has the property that $T_*(\varphi)(b) = \varphi(T(b))$ for all $b \in \mathcal{B}$ and $\varphi \in \mathcal{A}_*$. We interpret such quantum channels as possible physical transformations from input to output states.

A particular class of channels is that of measurements. In this case the input is the state of a quantum system described by an algebra $\mathcal{A}$, and the output is a probability distribution over the space of outcomes $(\mathcal{X}, \Sigma_\mathcal{X})$. Any measurement is described by a positive linear map

$$M : L^\infty(\mathcal{X}, \Sigma_\mathcal{X}, P) \to \mathcal{A},$$

which is completely specified by the image of characteristic functions of measurable sets, also called *positive operator valued measure* (POVM). This map $M : \Sigma_\mathcal{X} \to \mathcal{A}$ has following properties

1. Positive: $M(A) \geq 0, \qquad \forall\, A \in \Sigma_\mathcal{X}$;
2. Countably additive: $\sum_{i=1}^\infty M(A_i) = M(\cup_i A_i), \quad A_i \cap A_j = \emptyset, i \neq j$;
3. Normalized: $M(\mathcal{X}) = \mathbf{1}$.

The corresponding channel acting on states is a positive map $M_* : \mathcal{A}_* \to L^1(\mathcal{X}, \Sigma_\mathcal{X}, P)$ given by

$$M_*(\varphi)(A) = \varphi(M(A)) = \mathrm{Tr}(\rho M(A)),$$

where ρ is the density matrix of φ. By applying the channel M to the quantum statistical experiment consisting of the family of states $\mathcal{Q} = \{\varphi_\theta : \theta \in \Theta\}$ on $\mathcal{A}$ we obtain a classical statistical experiment

$$\mathcal{Q}_M := \{M_*(\varphi_\theta) : \theta \in \Theta\},$$

over the outcomes space $(\mathcal{X}, \Sigma_\mathcal{X})$.

As in the classical case, quantum channels can be seen as ways to compare quantum experiments. The first steps in this direction were made by Petz (1986); Petz et Jenčová (2006); Ohya et Petz, D. (2004) who developed the theory of *quantum sufficiency* dealing with the problem of characterizing when a sub-algebra of observables contains the same statistical information about a family of states, as the original algebra. More generally, two experiments $\mathcal{Q} := \{\mathcal{A}, \varphi_\theta : \theta \in \Theta\}$ and $\mathcal{R} := \{\mathcal{B}, \psi_\theta : \theta \in \Theta\}$ are called *statistically equivalent* if there exist channels $T : \mathcal{A} \to \mathcal{B}$ and $S : \mathcal{B} \to \mathcal{A}$ such that

$$\psi_\theta \circ T = \varphi_\theta \qquad \text{and} \qquad \varphi_\theta \circ S = \psi_\theta \qquad \forall\theta.$$

As consequence, for any measurement $M : L^\infty(\mathcal{X}, \Sigma_\mathcal{X}, P) \to \mathcal{A}$ there exists a measurement $T \circ M : L^\infty(\mathcal{X}, \Sigma_\mathcal{X}, P) \to \mathcal{B}$ such that the resulting classical experiments coincide $\mathcal{Q}_M = \mathcal{R}_{T\circ M}$. Thus for any statistical problem, and any procedure concerning the experiment $\mathcal{Q}$ there exists a procedure for $\mathcal{R}$ with the same risk (average error), and vice versa.

9.2.2 The Le Cam distance and its statistical meaning

We have seen that two experiments are statistically equivalent when they can be transformed into each other be means of quantum channels. When this cannot be done exactly, we would like to have a measure of how close the two experiments are when we allow any channel transformation. We define the *deficiency* of $\mathcal{R}$ with respect to $\mathcal{Q}$ as

$$\delta(\mathcal{R}, \mathcal{Q}) = \inf_T \sup_\theta \|\varphi_\theta - \psi_\theta \circ T\| \tag{9.7}$$

where the infimum is taken over all channels $T : \mathcal{A} \to \mathcal{B}$. The norm distance between two states on $\mathcal{A}$ is defined as

$$\|\varphi_1 - \varphi_2\| := \sup\{|\varphi_1(a) - \varphi_2(a)| : a \in \mathcal{A}, \|a\| \leq 1\},$$

and for $\mathcal{A} = \mathcal{B}(\mathcal{H})$ it is equal to $\|\rho_1 - \rho_2\|_1 := \mathrm{Tr}(|\rho_1 - \rho_2|)$, where ρ_i is the density matrix of the state φ_i. When $\delta(\mathcal{R}, \mathcal{Q}) = 0$ we say that $\mathcal{R}$ is more informative than $\mathcal{Q}$. Note that $\delta(\mathcal{R}, \mathcal{Q})$ is not symmetric but satisfies a triangle inequality of the form $\delta(\mathcal{R}, \mathcal{Q}) + \delta(\mathcal{Q}, \mathcal{T}) \geq \delta(\mathcal{R}, \mathcal{T})$. By symmetrizing we obtain a proper distance over the space of equivalence classes of experiments, called Le Cam's distance (Le Cam, 1986)

$$\Delta(\mathcal{Q}, \mathcal{R}) := \max\left(\delta(\mathcal{Q}, \mathcal{R})\,,\, \delta(\mathcal{R}, \mathcal{Q})\right). \tag{9.8}$$

What is the statistical meaning of the Le Cam distance ? We shall show that if $\delta(\mathcal{R}, \mathcal{Q}) \leq \epsilon$ then for any statistical decision problem with loss function between 0 and 1, any measurement procedure for $\mathcal{Q}$ can be matched by a measurement procedure for $\mathcal{R}$ whose risk will be at most ϵ larger than the previous one.

A decision problem is specified by a *decision space* $(\mathcal{X}, \Sigma_{\mathcal{X}})$ and a *loss function* $W_\theta : \mathcal{X} \to [0, 1]$ for each $\theta \in \Theta$. We are given a quantum system prepared in the state $\varphi_\theta \in \mathcal{A}_*$ with unknown parameter $\theta \in \Theta$ and would like to perform a measurement with outcomes in $\mathcal{X}$ such that the expected value of the loss function W_θ is small. Let

$$M : L^\infty(\mathcal{X}, \Sigma_{\mathcal{X}}, P) \to \mathcal{A},$$

be such a measurement, and $P_\theta^{(M)} = \varphi_\theta \circ M$, then the *risk* at θ is

$$R(M, \theta) := \int_{\mathcal{X}} W_\theta(x) P_\theta^{(M)}(dx).$$

Since the point θ is unknown one would like to obtain a small risk over all possible realizations

$$R_{max}(M) = \sup_{\theta \in \Theta} R(M, \theta).$$

The *minimax risk* is then

$$R_{minmax} := \inf_M R_{max}(M). \tag{9.9}$$

In the Bayesian framework one considers a prior distribution π over Θ and then averages the risk with respect to π

$$R_\pi(M) = \int_\Theta R(M, \theta)\pi(d\theta).$$

The optimal risk in this case is $R_\pi := \inf_M R_\pi(M)$.

Coming back to the experiments $\mathcal{Q}$ and $\mathcal{R}$ we shall compare their achievable risks for a given decision problem as above. Consider the measurement $N : L^\infty(\mathcal{X}, \Sigma_{\mathcal{X}}, P) \to$

$\mathcal{B}$ given by $N = T \circ M$ where $T : \mathcal{A} \to \mathcal{B}$ is the channel which achieves the infimum in (9.7). Then

$$\begin{aligned} R(N,\theta) &= \int_{\mathcal{X}} W_\theta(x) P_\theta^{(N)}(dx) = \psi_\theta(T \circ M(W_\theta)) \\ &\leq \|\psi_\theta \circ T - \varphi_\theta\| + \varphi_\theta(M(W_\theta)) \leq \delta(\mathcal{R},\mathcal{Q}) + R(M,\theta), \end{aligned}$$

where we have used the fact that $0 \leq W_\theta \leq 1$.

Lemma 9.2.3. *For every achievable risk $R(M,\theta)$ for $\mathcal{Q}$ there exists a measurement $N : L^\infty(\mathcal{X}, \Sigma_\mathcal{X}, P) \to \mathcal{B}$ for $\mathcal{R}$ such that*

$$R(N,\theta) \leq R(M,\theta) + \delta(\mathcal{R},\mathcal{Q}).$$

In consequence

$$R_{minmax}(\mathcal{R}) \leq R_{minmax}(\mathcal{Q}) + \delta(\mathcal{R},\mathcal{Q}).$$

9.3 Local asymptotic normality in statistics

In this section we describe the notion of local asymptotic normality and its significance in statistics (Le Cam, 1986; Torgersen, 1991; Strasser, 1985; van der Vaart, 1998). Suppose that we observe $X_1, \ldots, X_n$ where X_i take values in a measurable space $(\mathcal{X}, \Sigma_\mathcal{X})$ and are are independent, identically distributed with distribution P_θ indexed by a parameter θ belonging to an open subset $\Theta \subset \mathbb{R}^m$. The full sample is a single observation from the product P_θ^n of n copies of P_θ on the sample space (Ω^n, Σ^n). Local asymptotic normality means that for large n such statistical experiments can be approximated by Gaussian experiments after a suitable reparametrisation. Let θ_0 be a fixed point and define a local parameter $u = \sqrt{n}(\theta - \theta_0)$ characterizing points in a small neighbourhood of θ_0, and rewrite P_θ^n as $P_{\theta_0+u/\sqrt{n}}^n$ seen as a distribution depending on the parameter u. Local asymptotic normality means that for large n the experiments

$$\left\{ P_{\theta_0+u/\sqrt{n}}^n : u \in \mathbb{R}^m \right\} \qquad \text{and} \qquad \left\{ N(u, I_{\theta_0}^{-1}) : u \in \mathbb{R}^m \right\},$$

have the same statistical properties when the models $\theta \mapsto P_\theta$ are sufficiently 'smooth'. The point of this result is that while the original experiment may be difficult to analyse, the limit one is a tractable *Gaussian shift experiment* in which we observe a single sample from the normal distribution with unknown mean u and fixed variance matrix $I_{\theta_0}^{-1}$. Here

$$[I_{\theta_0}]_{ij} = \mathbb{E}_{\theta_0}\left[\ell_{\theta_0,i}\ell_{\theta_0,j}\right],$$

is the Fisher information matrix at θ_0, with $\ell_{\theta,i} := \partial \log p_\theta / \partial\theta_i$ the score function and p_θ is the density of P_θ with respect to a reference probability distribution P.

There exist two formulations of the result depending on the notion of convergence which one uses. In this chapter we only discuss the *strong* version based on convergence with respect to the Le Cam distance, and we refer to the book by (van der Vaart, 1998) for another formulation using the so called weak convergence (convergence in distribution of finite dimensional marginals of the likelihood ratio process), and to (Guţă et Jenčová, 2007) for its generalization to quantum statistical experiments.

Before formulating the theorem, we explain what sufficiently smooth means. The least restrictive condition is that p_θ is *differentiable in quadratic mean, i.e.* there exists a measurable function $\ell_\theta : \mathcal{X} \to \mathbb{R}$ such that as $u \to 0$

$$\int \left[p_{\theta+u}^{1/2} - p_\theta^{1/2} - u^t \ell_\theta p_\theta^{1/2}\right]^2 dP \to 0.$$

Note that ℓ_θ must still be interpreted as score function since under some regularity conditions we have $\partial p_\theta^{1/2}/\partial\theta_i = \frac{1}{2}(\partial \log p_\theta/\partial\theta_i)p_\theta^{1/2}$.

Theorem 9.3.1. *Let $\mathcal{E} := \{P_\theta : \theta \in \Theta\}$ be a statistical experiment with $\Theta \subset \mathbb{R}^d$ and $P_\theta \ll P$ such that the map $\theta \to p_\theta$ is differentiable in quadratic mean. Define*

$$\mathcal{E}_n = \{P^n_{\theta_0+u/\sqrt{n}} : \|u\| \leq C\}, \qquad \mathcal{F} = \{N(u, I_{\theta_0}^{-1}) : \|u\| \leq C\},$$

with I_{θ_0} the Fisher information matrix of $\mathcal{E}$ at point θ_0, and C a positive constant. Then $\Delta(\mathcal{E}_n, \mathcal{F}) \to 0$. In other words, there exist sequences of randomisations T_n and S_n such that:

$$\lim_{n\to\infty} \sup_{\|u\|\leq C} \left\| T_n(P^n_{\theta_0+u/\sqrt{n}}) - N(u, I_{\theta_0}^{-1}) \right\| = 0,$$

$$\lim_{n\to\infty} \sup_{\|u\|\leq C} \left\| P^n_{\theta_0+u/\sqrt{n}} - S_n(N(u, I_{\theta_0}^{-1})) \right\| = 0.$$

Remark 9.3.2. Note that the statement of the Theorem is not of Central Limit type which typically involves convergence *in distribution* to a Gaussian distribution at a *single* point θ_0. Local asymptotic normality states that the convergence is *uniform* around the point θ_0, and moreover the variance of the limit Gaussian is fixed whereas the variance obtained from the Central Limit Theorem depends on the point θ. Additionally, the randomisation transforming the data $(X_1, \ldots, X_n)$ into the Gaussian variable is the same for all $\theta = \theta_0 + u/\sqrt{n}$ and thus does not require *a priori* the knowledge of θ.

Remark 9.3.3. Local asymptotic normality is the basis of many important results in asymptotic optimality theory and explains the asymptotic normality of certain estimators such as the maximum likelihood estimator. The quantum version introduced in the next section plays a similar role for the case of quantum statistical model. An asymptotically optimal estimation strategy based on local asymptotic normality was derived by Guţă et al. (2008) for two-dimensional systems.

Remark 9.3.4. Let us define the real Hilbert space $L^2(\theta_0) = (\mathbb{R}^m, (\cdot,\cdot)_{\theta_0})$ with inner product

$$(u,v)_{\theta_0} = u^t I_{\theta_0} v.$$

By multiplying with I_{θ_0} we see that limit experiment can be equivalently chosen to be $N(I_{\theta_0}u, I_{\theta_0})$. The characteristic function of $X \sim N(I_{\theta_0}u, I_{\theta_0})$ is

$$F_u(w) := \mathbb{E}_{\theta_0}[\exp(iw^t X)] = \exp\left(-\frac{1}{2}\|w\|^2_{\theta_0} + i(w,u)_{\theta_0}\right). \tag{9.10}$$

A similar expression will be encountered in section 9.4 for the case of quantum Gaussian shift experiment.

Example 9.3.1. *Let $P_\mu = (\mu_1, \ldots, \mu_d)$ be a probability distribution with unknown parameters $(\mu_1, \ldots, \mu_{d-1}) \in \mathbb{R}^{d-1}_+$ satisfying $\mu_i > 0$ and $\sum_{i\leq d-1}\mu_i < 1$. The Fisher information at a point μ is*

$$I(\mu)_{ij} = \sum_{k=1}^{d-1} \mu_k(\delta_{ik}\mu_i^{-1}\cdot\delta_{jk}\mu_j^{-1}) + (1-\sum_{l=1}^{d-1}\mu_l)^{-1} = \delta_{ij}\mu_i^{-1} + (1-\sum_{l=1}^{d-1}\mu_l)^{-1}, \tag{9.11}$$

and its inverse is

$$V(\mu)_{ij} := [I(\mu)^{-1}]_{ij} = \delta_{ij}\mu_i - \mu_i\mu_j. \tag{9.12}$$

Thus the limit experiment in this case is $\mathcal{F} := (N(u, V(\mu)) : u \in \mathbb{R}^{d-1}, \|u\| \leq C)$.

This experiment will appear again in Theorem 9.4.3, as the classical part of the limit Gaussian shift experiment.

9.4 Local asymptotic normality in quantum statistics

In this section we present the main result of the chapter. Local asymptotic normality for d-dimensional quantum systems means roughly the following: the sequence $\mathcal{Q}_n$ of experiments consisting of joint states $\rho^{\otimes n}$ of n identical quantum systems prepared independently in the same state ρ, converges to a limit experiment $\mathcal{R}$ which is a quantum-classical Gaussian model involving displaced thermal equilibrium states of $d(d-1)/2$ oscillators and a $(d-1)$-dimensional classical Gaussian shift model. As in the classical case, the result has a local nature reflecting the $1/\sqrt{n}$ rate of convergence of state estimation. A neighbourhood of a fixed diagonal state $\rho_0 = \mathrm{Diag}(\mu_1, \ldots, \mu_d)$ is parametrised by (changes in the) diagonal parameters $\vec{u} \in \mathbb{R}^{d-1}$ and off-diagonal parameters $\vec{\zeta} \in \mathbb{C}^{d(d-1)/2}$. The latter can be implemented by small unitary rotations. The limit Gaussian model has a classical part $N(\vec{u}, V(\mu))$ with fixed known variance

$V(\mu)$, and a quantum part $\otimes_{j<k}\phi_{j,k}^{\zeta_{j,k}}$ with each $\phi_{j,k}^{\zeta_{j,k}}$ being a thermal equilibrium state with $\beta_{j,k} = \ln(\mu_j/\mu_k)$, displaced in phase space by an amount proportional to $\zeta_{j,k}$.

The reason for choosing the above parametrisation is twofold. Firstly, it unveils the important separation between 'classical' and 'quantum' parameters, and the further separation among the different off-diagonal parameters. Secondly, it is very convenient for the proof. However as we shall see in section 9.4.6, the limit experiment can be formulated in a 'coordinate-free' way in terms of quasifree states on a CCR-algebra. Although it is not needed in the main theorem, we include this formulation linking our result to the Quantum Central Limit Theorem. We stress again that local asymptotic normality is not a consequence of the Central Limit Theorem, indeed the latter is not even an ingredient in the proof but gives an indication as to what is the limit state when all parameters are zero.

9.4.1 The n-tuple of d-dimensional systems

As explained in section 9.3 for the classical case, our theory will be local in nature, so we shall be interested in a (shrinking) neighbourhood of an arbitrary but fixed faithful state

$$\rho_0 = \begin{bmatrix} \mu_1 & 0 & \dots & 0 \\ 0 & \mu_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \mu_d \end{bmatrix} \qquad \text{with } \mu_1 > \mu_2 > \cdots > \mu_d > 0, \tag{9.13}$$

which for technical reasons is chosen to have different eigenvalues. A sufficiently small neighbourhood of ρ_0 in the state space can be parametrised by $\theta := (\vec{u}, \vec{\zeta})$ as follows

$$\tilde{\rho}'_\theta := \begin{bmatrix} \mu_1 + u_1 & \zeta_{1,2}^* & \dots & \zeta_{1,d}^* \\ \zeta_{1,2} & \mu_2 + u_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \zeta_{d-1,d}^* \\ \zeta_{1,d} & \dots & \zeta_{d-1,d} & \mu_d - \sum_{i=1}^{d-1} u_i \end{bmatrix}, \qquad u_i \in \mathbb{R},\ \zeta_{j,k} \in \mathbb{C}. \tag{9.14}$$

Indeed, note that if θ is small enough then $\tilde{\rho}_\theta$ is a density matrix.

Later on, because of the limit experiments, the following normalisation will be easier,

and that's the one we shall use throughout the chapter:

$$\tilde{\rho}_\theta := \begin{bmatrix} \mu_1 + u_1 & \zeta_{1,2}^*\sqrt{\mu_1 - \mu_2} & \dots & \zeta_{1,d}^*\sqrt{\mu_1 - \mu_d} \\ \zeta_{1,2}\sqrt{\mu_1 - \mu_2} & \mu_2 + u_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \zeta_{d-1,d}\sqrt{\mu_{d-1} - \mu_d} \\ \zeta_{1,d}\sqrt{\mu_1 - \mu_d} & \dots & \zeta_{d-1,d}\sqrt{\mu_{d-1} - \mu_d} & \mu_d - \sum_{i=1}^{d-1} u_i \end{bmatrix}. \tag{9.15}$$

Let $\delta := \inf_{1\le i\le d} \mu_i - \mu_{i+1}$, with $\mu_{d+1} = 0$, be the separation between the eigenvalues. In the first order in $\theta/\sqrt{\delta}$, the family $\tilde{\rho}_\theta$ is obtained by first perturbing the diagonal elements of ρ_0 with $\vec{u}$ and then performing a small unitary transformation with

$$U(\vec{\zeta}) := \exp\left[i \left(\sum_{1\le j<k\le d} \frac{\mathrm{Re}(\zeta_{j,k}) T_{j,k} + \mathrm{Im}(\zeta_{j,k}) T_{k,j}}{\sqrt{\mu_j - \mu_k}} \right) \right] \tag{9.16}$$

where $T_{j,k}$ are generators of the Lie algebra of $SU(d)$ defined in (9.84). The advantage of the latter parametrisation is that we can fully exploit the machinery of irreducible group representations. For this reason, in all subsequent computations we shall work with the 'unitary' family

$$\rho_\theta := U(\vec{\zeta}) \begin{bmatrix} \mu_1 + u_1 & 0 & \dots & 0 \\ 0 & \mu_2 + u_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \mu_d - \sum_{i=1}^{d-1} u_i \end{bmatrix} U^*(\vec{\zeta}), \qquad u_i \in \mathbb{R},\ \zeta_{j,k} \in \mathbb{C}. \tag{9.17}$$

but we keep in mind the relationship with (9.15).

As in the classical case, the parameter θ will be scaled by the factor $1/\sqrt{n}$ meaning that we zoom in around ρ_0 with the rate equal to the typical estimation rate based on n samples. Let $\rho^{\theta,n} := \rho_{\theta/\sqrt{n}}^{\otimes n}$ and let $\mathcal{Q}_n$ be the sequence of statistical experiments

$$\mathcal{Q}_n := \left\{ \rho^{\theta,n} : \theta \in \Theta_n \right\}, \tag{9.18}$$

consisting of n systems, each one prepared in a state $\rho_{\theta/\sqrt{n}}$ situated in a local neighbourhood of ρ_0. The local parameter $\theta = (\vec{u}, \vec{\zeta})$ belongs to a neighbourhood Θ_n of the origin of $\mathbb{R}^{d-1} \times \mathbb{C}^{d(d-1)/2}$ which is allowed to grow slowly with n in a way that will be made precise later.

One of the principal tools in our result is the representation theory of the special unitary group $SU(d)$. Due to lack of space we shall not include any proofs and refer to the books by Fulton (1997); Goodman R. et Wallach N.R. (1998); Fulton et Harris (1991) for details. In particular we shall be working with the well known

tensor representation which will be analysed in increasing depth across the following sections.

The space $(\mathbb{C}^d)^{\otimes n}$ carries two commuting group representations: that of $SU(d)$ given by

$$\pi_n(U) : f_1 \otimes \cdots \otimes f_n \mapsto Uf_1 \otimes \cdots \otimes Uf_n, \qquad U \in SU(d), \tag{9.19}$$

and that of the permutation group $S(n)$ given by

$$\tilde{\pi}_d(\tau) : f_1 \otimes \cdots \otimes f_n \mapsto f_{\tau^{-1}(1)} \otimes \cdots \otimes f_{\tau^{-1}(n)}, \qquad \tau \in S(n). \tag{9.20}$$

Since the two group representations commute with each other, the representation space decomposes into a direct sum of tensor products of irreducible representations. It turns out that the irreducible representations of $SU(d)$ and $S(n)$ are indexed by *Young diagrams* with d rows for the former and n boxes for the latter. A Young diagram is defined by a tuple of ordered integers $\lambda = (\lambda_1 \geq \lambda_2 \cdots \geq \lambda_k)$ with λ_i the number of boxes on row i (see Figure 9.1). As we shall see later this pictorial

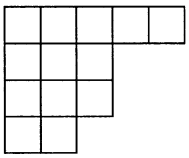

Figure 9.1: Young diagram with $\lambda = (5, 3, 3, 2)$.

representation will be very useful in understanding the structure of the irreducible representations $(\mathcal{H}_\lambda, \pi_\lambda)$ of $SU(d)$.

The following theorem called *Schur-Weyl duality* shows that the only tensor products appearing in the above mentioned direct sum are those of irreducible representations indexed by the same λ, and in particular the algebras generated by $\pi_n(u)$ and respectively $\tilde{\pi}_d(\tau)$ are each other's commutant!

Theorem 9.4.1. *Let π_n and $\tilde{\pi}_d$ be the representations of $SU(d)$ and respectively $S(n)$ on $(\mathbb{C}^d)^{\otimes n}$. Then the representation space decomposes into a direct sum of tensor products of irreducible representations of $SU(d)$ and $S(n)$ indexed by Young diagrams with d lines and n boxes:*

$$\begin{aligned}
(\mathbb{C}^d)^{\otimes n} &\cong \bigoplus_\lambda \mathcal{H}_\lambda \otimes \mathcal{K}_\lambda, \\
\pi_n &\equiv \bigoplus_\lambda \pi_\lambda \otimes \mathbf{1}_{\mathcal{K}_\lambda}, \\
\tilde{\pi}_d &\equiv \bigoplus_\lambda \mathbf{1}_{\mathcal{H}_\lambda} \otimes \tilde{\pi}_\lambda.
\end{aligned}$$

In particular $\rho^{\theta,n} = \rho_{\theta/\sqrt{n}}^{\otimes n}$ and $\tilde{\pi}_d(\tau)$ commute for all τ. Hence we have the block diagonal form for the joint states

$$\rho^{\theta,n} = \bigoplus_{\lambda} p_\lambda^{\theta,n} \rho_\lambda^{\theta,n} \otimes \frac{\mathbf{1}_{\mathcal{K}_\lambda}}{M_n(\lambda)}, \tag{9.21}$$

where $M_n(\lambda)$ is the dimension of $\mathcal{K}_\lambda$, $p_\lambda^{\theta,n}$ is a probability distribution over the Young diagrams, and $\rho_\lambda^{\theta,n}$ is a density matrix on $\mathcal{H}_\lambda$. >From (9.17) and the Schur-Weyl duality, we get the expression of the block states

$$\rho_\lambda^{\theta,n} = U_\lambda(\vec{\zeta}/\sqrt{n})\, \rho_\lambda^{u,0,n}\, U_\lambda(\vec{\zeta}/\sqrt{n})^*. \tag{9.22}$$

We interpret the decomposition (9.21) as follows: by doing a 'which block' measurement we obtain information about θ through the probability density $p_\lambda^{\theta,n}$. In fact it is easy to see that $p_\lambda^{\theta,n}$ does not depend on $\vec{\zeta}$, so it only gives information about the diagonal parameters $\vec{u}$. Later on we shall see that the model $p^{\theta,n}$ has the same limit as the classical multinomial model described in Example 9.3.1. Once this information has been obtained, one still possesses a conditional quantum state $\rho_\lambda^{\theta,n}$. It turns out that this state carries information about the rotation parameters $\vec{\zeta}$, and we shall show that the statistical model described by the conditional state converges to a 'purely quantum' Gaussian shift experiment.

9.4.2 Displaced thermal equilibrium states of a harmonic oscillator

The ground state of a quantum harmonic oscillator or the laser state of a monochromatic light pulse are well known examples of quantum Gaussian states. Both physical systems are described by the same algebra of observables generated by the canonical 'position' and 'momentum' observables $\mathbf{Q}$ and $\mathbf{P}$ satisfying the Heisenberg commutation relation

$$\mathbf{QP} - \mathbf{PQ} = i\mathbf{1}. \tag{9.23}$$

These observables can be represented on the Hilbert space $L^2(\mathbb{R})$ as

$$(\mathbf{Q}f)(x) = xf(x), \qquad (\mathbf{P}f)(x) = -i\frac{df}{dx}(x), \qquad f \in L^2(\mathbb{R}). \tag{9.24}$$

The space $L^2(\mathbb{R})$ has a special orthonormal basis $\{|0\rangle, |1\rangle, \ldots\}$ with the vector $|m\rangle$ given by

$$H_m(x)e^{-x^2/2}/(\sqrt{\pi}2^m m!)^{1/2},$$

where H_m are the Hermite polynomials (Erdélyi, 1953). These are the eigenvectors of the number operator $\mathbf{N} := \frac{1}{2}(\mathbf{Q}^2 + \mathbf{P}^2 - \mathbf{1})$ counting the number of 'excitations' of the oscillator or the number of photons in the case of the light beam, such that $\mathbf{N}\,|m\rangle = m\,|m\rangle$.

The creation and annihilation operators

$$\mathbf{a}^* = (\mathbf{Q} - i\mathbf{P})/\sqrt{2}, \qquad \mathbf{a} = (\mathbf{Q} + i\mathbf{P})/\sqrt{2},$$

satisfy $[\mathbf{a}, \mathbf{a}^*] = \mathbf{1}$ and act as 'ladder' operators on the number basis:

$$\mathbf{a}\,|m\rangle = \sqrt{m}\,|m-1\rangle\,, \qquad \mathbf{a}^*\,|m\rangle = \sqrt{m+1}\,|m+1\rangle\,.$$

In particular the following identity holds: $\mathbf{N} = \mathbf{a}^*\mathbf{a}$.

It can be easily checked that both $\mathbf{Q}$ and $\mathbf{P}$ have Gaussian distribution with respect to the vacuum state $|0\rangle$. In fact they are 'jointly Gaussian'

$$\langle 0|\exp(iu\mathbf{Q} + iv\mathbf{P})|\,|0\rangle = \exp\left(-\frac{1}{4}(u^2 + v^2)\right).$$

We shall often use the complex form of the unitary *Weyl operators*

$$W(z) := \exp(z\mathbf{a}^* - \bar{z}\mathbf{a}) = \exp(ip_0\mathbf{Q} - iq_0\mathbf{P}), \qquad z = (q_0 + ip_0)/\sqrt{2} \in \mathbb{C},$$

which satisfy the Weyl relations

$$W(z)^*W(z')W(z) = \exp\left(2i\mathrm{Im}(\bar{z}'z)\right)W(z').$$

The coherent (vector) states $|z\rangle$ are obtained by displacing the vacuum state with Weyl operators

$$|z\rangle := W(z)\,|0\rangle = \exp(-|z|^2/2)\sum_{m=0}^{\infty}\frac{z^m}{\sqrt{m!}}\,|m\rangle\,. \tag{9.25}$$

They are Gaussian states with the same variance as the vacuum, and means $\langle z|\,\mathbf{Q}\,|z\rangle = \sqrt{2}\mathrm{Re}(z)$ and $\langle z|\,\mathbf{P}\,|z\rangle = \sqrt{2}\mathrm{Im}(z)$:

$$\langle z|\,W(z')\,|z\rangle = \exp\left(-\frac{1}{2}|z - z'|^2 + 2i\mathrm{Im}(\bar{z}'z)\right).$$

Besides, coherent states, an important role in our discussion will be played by the thermal equilibrium states. For every $\beta > 0$ we define the Gaussian state

$$\varphi_\beta(W(z)) = \exp\left(-\frac{|z|^2}{2\tanh(\beta/2)}\right). \tag{9.26}$$

Its density matrix consisting of a mixture of k-photon states with geometrical weights

$$\phi_\beta = (1 - e^{-\beta}) \sum_{k=0}^{\infty} e^{-k\beta} |k\rangle \langle k| . \tag{9.27}$$

and can also be obtained by 'smearing' the coherent states with a Gaussian kernel:

$$\phi_\beta = \frac{e^\beta - 1}{\pi} \int_{\mathbb{C}} \exp\left(-(e^\beta - 1)|z|^2\right) |z\rangle \langle z| \, dz. \tag{9.28}$$

The thermal equilibrium states can be shifted in 'phase space' by means of displacement operations D^z which act by adjoining with unitaries $W(z)$, i.e.

$$D^z(\cdot) := \mathrm{Ad}[W(z)](\cdot) = W(z)^* \cdot W(z).$$

The result is a Gaussian state φ_β^z with the same variance as φ_β and the same means as $|z\rangle\langle z|$:

$$\varphi_\beta^z(W(z')) := \exp\left(-\frac{|z|^2}{2\tanh(\beta/2)} + 2i\mathrm{Im}(\bar{z}'z)\right), \qquad \phi_\beta^z := D^z(\phi_\beta) := W(z)^* \phi_\beta W(z). \tag{9.29}$$

9.4.3 The multimode Fock space and the limit Gaussian shift experiment

We now consider $d(d-1)/2$ commuting harmonic oscillators, with a joint state consisting of independent Gaussian states. Let us define the *multimode Fock space*

$$\mathcal{F} := \bigotimes_{1 \le j < k \le d} L^2(\mathbb{R}),$$

in which we identify the number basis

$$|\mathbf{m}\rangle = \bigotimes_{j<k} |m_{j,k}\rangle , \qquad \mathbf{m} = \{m_{j,k} \in \mathbb{N} : j < k\} . \tag{9.30}$$

For each of the oscillators we define the thermal equilibrium state

$$\phi_{j,k} := \phi_{\beta_{j,k}}, \qquad \beta_{j,k} = \ln(\mu_j/\mu_k), \tag{9.31}$$

where $\{\mu_1, \ldots, \mu_d\}$ are the eigenvalues of the density matrix ρ_0 (cf. (9.13)). We now use the Weyl operators to displace these states by an amount proportional to the off-diagonal elements $\zeta_{j,k}$ of ρ^θ (cf. (9.15) and (9.17))

$$\phi_{j,k}^{\zeta_{j,k}} := W(\zeta_{j,k})^* \, \phi_{j,k} \, W(\zeta_{j,k}) .$$

Next we define the joint state $\varphi^{\vec{\zeta}}$ of the oscillators with density matrix

$$\phi^{\vec{\zeta}} = \bigotimes_{j<k} \phi_{j,k}^{\zeta_{j,k}} \in \mathcal{T}_1(\mathcal{F}), \tag{9.32}$$

where $\mathcal{T}_1(\mathcal{F})$ is the space of trace-class operators on $\mathcal{F}$.

The states $\phi^{\vec{\zeta}}$ form the quantum part of the limit Gaussian experiment. The classical part is identical to the $(d-1)$-dimensional Gaussian shift model $N(\vec{u}, V(\mu))$ of Example 9.3.1, where $\mu = \{\mu_1, \ldots, \mu_d\}$.

Definition 9.4.2. *On the algebra $L^\infty(\mathbb{R}^{d-1}) \otimes \mathcal{B}(\mathcal{F})$ we define normal state φ^θ with density*

$$\phi^\theta := \mathcal{N}(\vec{u}, V(\mu)) \otimes \phi^{\vec{\zeta}} \in L^1(\mathbb{R}^{d-1}) \otimes \mathcal{T}_1(\mathcal{F}), \tag{9.33}$$

where $\mathcal{N}(\vec{u}, V(\mu))$ is the Gaussian density of Example 9.3.1. The quantum-classical Gaussian experiment $\mathcal{R}$ is defined by

$$\mathcal{R} = \{\phi^\theta : \theta = (\vec{u}, \vec{\zeta}) \in \mathbb{R}^{d-1} \times \mathbb{C}^{d(d-1)/2}\}.$$

9.4.4 The main theorem

We are now ready to formulate the main result of the chapter. In view of subsequent application to optimal state estimation, it is essential to consider (slowly) growing domains of the local parameters. For given $\beta, \gamma > 0$ we define

$$\Theta_{n,\beta,\gamma} = \left\{(\vec{\zeta}, \vec{u}) : \|\vec{\zeta}\|_\infty \leq n^\beta, \|\vec{u}\|_\infty \leq n^\gamma\right\}.$$

Recall that δ is the separation between the eigenvalues of ρ_0 given by equation (9.13). Though we use parametrisation (9.17) for density matrices ρ_θ, recall that in the first order this is approximated by $\tilde{\rho}_\theta$ defined in (9.15). In fact it can be shown that the same theorem holds for the latter parametrisation.

Theorem 9.4.3. *Let $\delta > 0$, let $\beta < 1/9$ and $\gamma < 1/4$. Let the quantum experiments*

$$\mathcal{Q}_n = \left\{\rho^{\theta,n} : \theta \in \Theta_{n,\beta,\gamma}\right\}, \qquad \mathcal{R}_n = \left\{\phi^\theta : \theta \in \Theta_{n,\beta,\gamma}\right\},$$

where $\rho^{\theta,n} = \rho_{\theta/\sqrt{n}}^{\otimes n}$ is the state on $M\left((\mathbb{C}^d)^{\otimes n}\right)$ given by equation (9.17), and ϕ^θ is given by (9.33).

Then, there exist channels (completely positive, normalised maps)

$$T_n : M(\mathbb{C}^d)^{\otimes n} \to L^1(\mathbb{R}^{d-1}) \otimes \mathcal{T}_1(\mathcal{F}) \tag{9.34}$$

$$S_n : L^1(\mathbb{R}^{d-1}) \otimes \mathcal{T}_1(\mathcal{F}) \to M(\mathbb{C}^d)^{\otimes n} \tag{9.35}$$

with $\mathcal{T}_1(\mathcal{F})$ is the space of trace-class operators on $\mathcal{F}$, such that

$$\sup_{\theta\in\Theta_{n,\beta,\gamma}} \left\|\phi^\theta - T_n(\rho^{\theta,n})\right\|_1 = O(n^{-\kappa}), \tag{9.36}$$

$$\sup_{\theta\in\Theta_{n,\beta,\gamma}} \left\|S_n(\phi^\theta) - \rho^{\theta,n}\right\|_1 = O(n^{-\kappa}), \tag{9.37}$$

where $\kappa > 0$ depends only on δ, β and γ. It is decreasing in each of them. In particular we have

$$\lim_{n\to\infty} \Delta(\mathcal{Q}_n, \mathcal{R}_n) = 0,$$

where $\Delta(\cdot,\cdot)$ is the Le Cam distance defined in (9.8).

In other words, we get polynomial speed of convergence of the approximation, which is enough to build two-step evaluation strategies, as shown in the next subsection. The main steps of the proof are given in a sequence of Lemmas in section 9.6 assembled into Theorem 9.6.7. The bound (9.37) follows easily from (9.36) as shown in section 9.6.2.

9.4.5 Application: Asymptotically optimal estimation procedure

This theorem allows us to pull back what we know from Gaussian shift experiments to *i.i.d.* experiments[1]. A first application is devising an optimal estimation procedure in the latter case.

We shall work with *well-behaved* loss functions $l(\rho,\hat{\rho})$, the typical example being the squared L^2 operator norm $l_\rho(\hat{\rho}) = \|\rho - \hat{\rho}\|_2^2$. Usual squared distances will also satisfy the conditions. The requirements are the following:

- Boundedness: $l(\rho,\hat{\rho})$ is bounded, by L.
- Lower semicontinuity: $l(\rho,\hat{\rho})$ is lower semicontinuous as a two-variable function.
- Estimation-fostering: $l \geq 0$ and is zero if and only if $\rho = \hat{\rho}$.
- Local quadraticity: If we use parametrisation (9.17) around any ρ_0 with positive distinct eigenvalues, and viewing θ as a d^2-dimensional real vector through

[1] The converse, too, but that's less often useful.

separating the real and imaginary parts of $\vec{\zeta}$, we have the approximation[2]

$$l_{\rho_\theta}(\rho_{\theta+h}) = h^* G_{\rho_0} h + O(\|h\|^3, \|\theta\|^3), \tag{9.38}$$

where G_{ρ_0} is a positive definite d^2-dimensional matrix possibly depending on ρ_0.

We shall work with *i.i.d.* experiments, that is with $\rho^{\otimes n}$. To take the scale into account, we shall rescale the risk, so that the risk of an estimator M_n in the n-sample experiment, with values in the states on $\mathbb{C}^d$, would have risk at point ρ given by:

$$R(M_n, \rho) = n \int_{\mathcal{T}_1^+(\mathbb{C}^d)} l(\rho, \hat{\rho}) P^{M_n}_{\rho^{\otimes n}}(\mathrm{d}\hat{\rho}),$$

where $P^{M_n}_{\rho^{\otimes n}}$ is the law of the result of the measurement M_n on $\rho^{\otimes n}$. An estimator of θ in some parametrisation can naturally be seen as an estimator of ρ.

We then define the minimax risk on a subset Θ of the states as

$$R^n_{minimax}(\Theta) = \inf_{M_n} \sup_{\rho \in \Theta} R(M_n, \rho).$$

We may now state:

Theorem 9.4.4. *Let $\epsilon > 0$. For any measurement procedures M_n on the experiments*

$$\mathcal{Q}_n = \left\{ \rho^{\theta, n} : \theta \in \Theta_{n,\epsilon,\epsilon} \right\},$$

the asymptotic maximum risk for a well-behaved loss function is at least the same as the minimax risk in the limit experiment

$$\mathcal{R} = \left\{ \phi^\theta : \theta \in \mathbb{R}^{d^2} \right\} \tag{9.39}$$

with loss function $r(\theta, \hat{\theta}) = (\theta - \hat{\theta})^* G_{\rho_0} (\theta - \hat{\theta})$. That is:

$$\liminf_n R^n_{minimax}(\Theta_{n,\epsilon,\epsilon}) \leq \inf_M \sup_\theta R(M, \phi^\theta) := R_{minimax} \tag{9.40}$$

There are measurement sequences that saturate this bound. Moreover, we do not need to know ρ_0 beforehand, that is, we do not need to know that ρ is in the small

[2] As a remark, we can allow bigger remainder terms $\|h\|^{2+\epsilon}$, by taking a smaller neighbourhood around ρ in the proof below. If the remainder term is too big, though, the first term in the risk (9.81) will get too big.

ball around ρ_0 permitted by $\Theta_{n,\epsilon,\epsilon}$. We may search among all states on $\mathbb{C}^d$ without any loss in performance, asymptotically[3].

This bound is the Holevo (1982) bound. Since its expression is very complicated, we do not reproduce it here, and instead give it in the special case when G is block-diagonal in the classical and each of the one-dimensional quantum Gaussian shift experiments parts:

$$\theta^* G \theta = \sum_{i=1}^{d-1} u^* G_c u + \sum_{1 \leq j < k \leq d} \zeta_{jk}^* G_{jk} \zeta_{jk} \tag{9.41}$$

for a nonnegative $(d-1)$-dimensional matrix G_c and nonnegative two-dimensional G_{jk}, the ζ_{jk} being seen as two-dimensional real vectors. In that case,

$$R_{minimax} = \mathrm{Tr}[V_\mu G_c] + \sum_{1 \leq j < k \leq d} \frac{1}{4} \frac{\mu_j + \mu_k}{\mu_j - \mu_k} \mathrm{Tr}(G_{jk}) + \frac{1}{2} \sqrt{\det G_{jk}}.$$

If $\zeta_{jk}^* G_{jk} \zeta_{jk} = \alpha_{jk} |\zeta_{jk}|^2$ for all j and k, this further simplifies as:

$$R_{minimax} = \mathrm{Tr}[V_\mu G_c] + \sum_{1 \leq j < k \leq d} \frac{\alpha_{jk} \mu_j}{\mu_j - \mu_k}.$$

We shall postpone almost all the proof to Section 9.7.1, and concentrate on making explicit the asymptotically optimal measurements.

The fact that is the minimax risk in the limit experiment is indeed a lower bound on the minimax risk of the finite-dimensional experiments is almost directly implied by the very general asymptotic minimax theorem 9.7.7[4]. We fill in some details in Section 9.7.1.

We thus only have to find measurement procedures that asymptotically attain the minimax risk of the limit experiment $R_{minimax}$.

We first use any rough procedure on $\tilde{n} = n^{1-\epsilon}$ copies of ρ, and get an estimate $\tilde{\rho}$.

We may then parametrise the states around $\tilde{\rho}$, that is use $\tilde{\rho} = \rho_0$ as in equation (9.13), and ρ_θ and $\rho^{\theta, n-\tilde{n}}$ defined accordingly. QLAN then yields a channel $T_{n-\tilde{n}}$ to approximately map the remaining states $\rho^{\otimes n-\tilde{n}}$ to a Gaussian state. We now apply on $T_{n-\tilde{n}}(\rho^{\otimes n-\tilde{n}})$ an optimal measurement for the limit experiment (9.39). This yields a result $\tilde{\theta}$.

[3] We give a global measurement strategy below.

[4] In the course of proving this theorem, I also establish other important results such as an asymptotic representation theorem.

If $\tilde{\theta}$ is more than $3\Gamma n^{\epsilon}$, where Γ is the ratio of the extreme eigenvalues of G_ρ, we estimate that we were outside the validity domain of QLAN. In other words, we had taken $\beta = \gamma = \epsilon$ for a small $\epsilon > 0$ in Theorem 9.4.3. We then take $\hat{\theta} = 0$. Else we take $\hat{\theta} = \tilde{\theta}$.

Our final estimator $\hat{\rho}$ will be the state corresponding to $\phi^{\hat{\theta}}$ in the limit experiment, that is $\hat{\rho} = \rho^{\hat{\theta}/\sqrt{n-\tilde{n}}}$ in the parametrisation around $\tilde{\rho}$.

The risk of this procedure may be divided in five parts. The first is what happens in the event that the rough estimation is really very bad, so that ρ is not in the validity domain of QLAN around $\tilde{\rho}$. The second is approximating the loss function r_n by the quadratic r. The third comes from QLAN, that is from the fact $T_n(\rho^{\theta,\otimes n})$ is not exactly ϕ^θ in general. The fourth comes from modifying the optimal limit estimator. All these will be shown to be negligible in the proof section. The final term is the risk of the minimax estimator in the limit experiment, that is $R_{minimax}$.

Let us now find an optimal measurement in the limit experiment. By the quantum Hunt-Stein theorem 9.7.9, we know that there is a minimax equivariant measurement for the action by translation of $\mathbb{R}^{d(d+1)/2}$ on the combined classical-quantum model. Such an equivariant measurement has constant bias, so we may subtract it if nonzero, and get better L^2 risk. Hence the measurement is unbiased.

Now, the optimal locally unbiased measurement is known (Holevo, 1982) to be a kind of heterodyne measurement. Its exact description for a general weight matrix G is hard to give, so we shall make it explicit only in the simpler case when the matrix is of the form 9.41.

An optimal equivariant estimator of a product experiment for a block-diagonal weight matrix is the product of the corresponding estimators. Indeed, the measurement for one of the subexperiments when the others are seen as an ancilla must still be equivariant, so its contribution cannot be better then that of the best equivariant estimator for that subexperiment.

So that the minimax risk is the sum of the minimax risk in the classical experiment and in each of the one-dimensional quantum Gaussian shift experiments.

For the classical experiment $\{\mathcal{N}(\vec{u}, V_\mu), \vec{u} \in \mathbb{R}^{d-1}\}$, the best unbiased estimator is known (van der Vaart, 1998, for example) to be the random variable itself, with law $\mathcal{N}(\vec{u}, V_\mu)$. So that the minimax risk on the classical experiment is $\mathrm{Tr}[V_\mu G_c]$.

For a one-dimensional quantum Gaussian shift experiment $\left\{\phi_\beta^\zeta\right\}$ with quadratic loss function, the best unbiased estimator is known to be a squeezed heterodyne measurement. Hayashi et Matsumoto (2004) give it a clear expression. The corresponding risk is $\frac{1}{4}\frac{\mu_j+\mu_k}{\mu_j-\mu_k}\mathrm{Tr}(G_{jk}) + \frac{1}{2}\sqrt{\det G_{jk}}$. In the frequent case when $G_{jk} = \alpha_{jk} Id$, this

expression reduces to $\frac{\alpha_{jk}\mu_j}{\mu_j-\mu_k}$, and the measurement to the usual heterodyne measurement.

Remark 9.4.5.

- *Most usual quadratic loss functions are products of the form 9.41. Notably, we deduce from that for rescaled L^2 operator loss, that is $\left\|\theta-\hat{\theta}\right\|_2^2$ with our parametrisation, the asymptotic minimax risk is $1-\sum_{j=1}^{d}\mu_i^2+\sum_{1\leq j\leq d-1}(d-j)\mu_j$.*
- *We may also use as loss function $n\left\|\rho_n-\hat{\rho}_n\right\|_1^2$. The corresponding weight matrix G is only easy in two dimensions, and we find again $8\mu_1-4\mu_1^2$, as in Chapter 8.*
- *We may use the square of the Bures distance, that is quadratic to the first order. Hübner (1992) has given a local development $d_B^2(\rho_0,\rho_0+\mathrm{d}\rho)=\frac{1}{2}\sum_{1\leq i,j\leq d}\frac{|\mathrm{d}\rho_{ij}|^2}{\mu_i+\mu_j}$. So that we get an asymptotic minimax risk of $d-1+\sum_{j<k}\frac{\mu_j}{\mu_j+\mu_k}$.*

Another remark is that since there are many dimensions, we can get second order improvement by using a Stein (1956) estimator to shrink the equivariant estimator.

9.4.6 The relation between LAN and CLT

One way to think of local asymptotic normality is the following: we would like to understand the asymptotic behaviour of the collective (fluctuation) observables (9.44) with respect to a *whole neighbourhood* of the state ρ, how the limit distribution changes as we change the reference state $\rho^{\otimes n}$.

The quantum Central Limit Theorem describes the asymptotic behaviour of the same observables with respect to a *fixed* state, and is one of the ingredients in the proof of a different version of LAN based on *weak convergence* (Guţă et Jenčová, 2007). However, in the case of strong convergence, which is the object of this chapter, CLT does not play any role since we are interested in convergence in norm rather than in distribution, and uniformly over a range of parameters.

The purpose of the section is to derive a 'coordinate free' version of the limit Gaussian experiment using the Central Limit Theorem and the notion of symmetric logarithmic derivative. The reader interested in the proof of main theorem may skip the following pages and continue with section 9.5.

Quantum Central Limit Theorem

Let ρ be the density matrix of a fixed faithful state on $M(\mathbb{C}^d)$. To ρ we associate an algebra of canonical commutation relations carrying a Gaussian state φ. The

Quantum Central Limit Theorem (Petz, 1990) says that φ is the limit distribution of certain multi-particle observables with respect to of product states $\rho^{\otimes n}$.

Let

$$(A,B)_\rho := \mathrm{Tr}(\rho\, A \circ B), \qquad \text{where } A \circ B := \frac{AB+BA}{2},$$

be a positive inner product on the real linear space of *selfadjoint operators* $M(\mathbb{C}^d)_{sa}$. We define the Hilbert space with inner product $(\cdot,\cdot)_\rho$.

$$L^2(\rho) = \{A \in M(\mathbb{C}^d)_{sa} : \mathrm{Tr}(A\rho) = 0\}.$$

Let σ be the *symplectic form* on $L^2(\rho)$

$$\sigma(A,B) = \frac{i}{2}\mathrm{Tr}(\rho\,[A,B]).$$

The C^*-algebra of canonical commutation relations $CCR(L^2(\rho),\sigma)$ is generated by the Weyl operators $W(A)$ satisfying the relations

$$W(A)^* = W(-A), \qquad W(A)W(B) = W(A+B)\exp(-i\sigma(A,B)), \quad A,B \in L^2(\rho).$$

On $CCR(L^2(\rho),\sigma)$ we define the Gaussian (quasifree) state

$$\varphi(W(A)) := \exp\left(-\frac{1}{2}\|A\|_\rho^2\right), \qquad \|A\|_\rho^2 = (A,A)_\rho. \tag{9.42}$$

The state φ is regular, i.e. there exists a representation $(\pi,\mathcal{H})$ of the algebra $CCR(L^2(\rho),\sigma)$ such that the one parameter family $t \mapsto \pi(W(tA))$ is weakly continuous and φ is extends to a normal state on the von Neumann algebra generated by $\pi(CCR(L^2(\rho),\sigma))$. This means that there exist selfadjoint 'field operators' $B(A)$ such that $\pi(W(tA)) = \exp(itB(A))$, and there exists a density matrix $\phi_\pi \in \mathcal{T}_1(\mathcal{H})$ such that

$$\varphi(W(A)) = \mathrm{Tr}\left(\exp(iB(A))\phi_\pi\right), \qquad A \in L^2(\rho).$$

The representation $(\pi,\mathcal{H})$ can be obtained through the GNS construction, or by 'diagonalising' the CCR algebra as we shall see in a moment. From (9.42) we deduce that the distribution of $B(A)$ with respect to φ is a centred normal distribution with variance $\|A\|_\rho^2$. From the Weyl relations it follows that the fields satisfy the following canonical commutation relations

$$[B(A),B(C)] = 2i\sigma(A,C)\mathbf{1}, \qquad A,C \in L^2(\rho).$$

Consider now the tensor product $\bigotimes_{k=1}^n M(\mathbb{C}^d)$ which is generated by elements of the form

$$A^{(k)} = \mathbf{1} \otimes \cdots \otimes A \otimes \cdots \otimes \mathbf{1}, \tag{9.43}$$

with A acting on the k-th position of the Hilbert space tensor product $\left(\mathbb{C}^d\right)^{\otimes n}$. We are interested in the asymptotics as $n \to \infty$ of the joint distribution under the state $\rho^{\otimes n}$, of 'fluctuation' elements of the form

$$F_n(A) := \frac{1}{\sqrt{n}} \sum_{k=1}^{n} A^{(k)}. \tag{9.44}$$

Theorem 9.4.6. [Quantum CLT] *Let $A_1, \ldots, A_s \in L^2(\rho)$. Then the following holds*

$$\lim_{n\to\infty} \mathrm{Tr}\left(\rho^{\otimes n}\left(\prod_{l=1}^{s} F_n(A_l)\right)\right) = \varphi\left(\prod_{l=1}^{s}(B(A_l))\right),$$
$$\lim_{n\to\infty} \mathrm{Tr}\left(\rho^{\otimes n}\left(\prod_{l=1}^{s} \exp(iF_n(A_l))\right)\right) = \varphi\left(\prod_{l=1}^{s} W(A_l)\right).$$

Although the algebra $CCR(L^2(\rho), \sigma)$ may look rather abstract, its structure can be easily understood by 'diagonalising' it. Let us assume that ρ is a diagonal matrix $\rho_0 = \mathrm{Diag}(\mu_1, \ldots, \mu_d)$. The Hilbert space $L^2(\rho_0)$ decomposes as direct sum of orthogonal subspaces $\mathcal{H}_{\rho_0} \oplus \mathcal{H}_{\rho_0}^{\perp}$ where

$$\mathcal{H}_{\rho_0} := \mathrm{Lin}\{A : [A, \rho_0] = 0, \mathrm{Tr}(A\rho_0) = 0\}, \quad \text{and} \quad \mathcal{H}_{\rho_0}^{\perp} = \mathrm{Lin}\{T_{j,k}, j \neq k\}, \tag{9.45}$$

with $T_{j,k}$ the generators of the $\mathfrak{su}(d)$ algebra defined in (9.84).

The elements $W(A)$ with $A \in \mathcal{H}_{\rho_0}$ generate the center of the algebra which is isomorphic to the algebra of bounded continuous functions $C_b(\mathbb{R}^{d-1})$. Explicitly, we identify the coordinates in $\mathbb{R}^{d-1}$ with the basis $\{d_i = -\mu\mathbf{1} + E_{i,i} : i = 1, \ldots d-1\}$ of $\mathcal{H}_{\rho_0}$, (see (9.84) for the definition of $E_{i,i}$). Then the covariance matrix for the basis vectors is

$$(d_i, d_j)_{\rho_0} = \mathrm{Tr}(\rho_0 d_i d_j) = \delta_{i,j}\mu_i - \mu_i\mu_j = [V(\mu)]_{i,j},$$

where V_μ is the covariance matrix (9.12).

Moreover

$$t_{j,k} := T_{j,k}/\sqrt{2(\mu_j - \mu_k)}, \qquad j \neq k, \tag{9.46}$$

form an *orthogonal and symplectic basis* of $\mathcal{H}_{\rho_0}^{\perp}$, i.e.

$$\sigma(t_{j,k}, t_{k,j}) = -1/2, \quad j < k, \quad \text{and } \sigma(t_{j,k}, t_{l,m}) = 0 \quad \text{for } \{j,k\} \neq \{l,m\}.$$

which means that $\{t_{j,k}, t_{k,j}\}$ generate isomorphic algebras of quantum harmonic oscillator which we denote by $CCR(\mathbb{C})$. From

$$\|t_{j,k}\|_{\rho_0}^2 = \mathrm{Tr}(\rho_0 t_{j,k}^2) = \frac{\mu_j + \mu_k}{2(\mu_j - \mu_k)}$$

and (9.26) we conclude that each of the oscillators is prepared independently in the thermal equilibrium state $\varphi_{j,k} = \varphi_{\beta_{j,k}}$ with $\beta_{j,k} = \ln(\mu_j/\mu_k)$.

Based on the discussion of sections 9.4.2 and 9.4.3 we can choose $\mathcal{H} := L^2(\mathbb{R}^{d-1}) \otimes \mathcal{F}$ and define the regular representation π of $CCR(L^2(\rho_0), \sigma)$ on this space in a straightforward way and its von Neumann completion is $L^\infty(\mathbb{R}^{d-1}) \otimes \mathcal{B}(\mathcal{F})$. The state φ decomposes as

$$\varphi \cong N(0, V_\mu) \otimes \bigotimes_{j<k} \varphi_{j,k}. \tag{9.47}$$

which is precisely the state φ^θ for $\theta = (\vec{u}, \vec{\zeta}) = (\vec{0}, \vec{0})$, defined in (9.33).

The quantum Gaussian shift experiment through Fisher information

We complete the family of states φ^θ of the experiment $\mathcal{R}$ by shifting φ^0 with the help of symmetric logarithmic derivatives. As in the classical case, this will be a family of Gaussian states with the same covariance, and mean proportional to the local parameter θ. The covariance is related to the Fisher information matrix as described in Remark 9.3.4. Thus we shall start by defining the quantum analogues of the score functions and the Fisher information matrix for the full quantum model ρ_θ.

Let us define the *symmetric logarithmic derivatives* (Helstrom, 1976; Holevo, 1982) as the solutions in $L^2(\rho_0)$ of

$$\mathcal{L}^{(re)}_{j,k} \circ \rho_0 = \left.\frac{\partial \rho_\theta}{\partial \operatorname{Re}\zeta_{j,k}}\right|_{\theta=0}, \quad \mathcal{L}^{(im)}_{j,k} \circ \rho_0 = \left.\frac{\partial \rho_\theta}{\partial \operatorname{Im}\zeta_{j,k}}\right|_{\theta=0}, \quad \ell_i \circ \rho_0 = \left.\frac{\partial \rho_\theta}{\partial u_i}\right|_{\theta=0},$$

Then with $T_{j,k}, E_{i,i}$ defined in (9.84)

$$\mathcal{L}^{(re)}_{j,k} = T_{k,j}/(\mu_j + \mu_k), \quad \mathcal{L}^{(im)}_{j,k} = T_{j,k}/(\mu_j + \mu_k), \quad \ell_i = E_{i,i}/\mu_i - E_{d,d}/\mu_d,$$

and the quantum Fisher information matrix consists of a 'classical block' that coincides with that of the classical multinomial model in (9.11)

$$[I_{\rho_0}]_{ij} := (\ell_i, \ell_j)_{\rho_0} = [I(\mu)]_{ij}, \qquad 1 \le i, j \le d-1,$$

and a 'purely quantum' block given by the diagonal matrix

$$H_{\rho_0} = \operatorname{Diag}\left(\|\mathcal{L}^{(re)}_{j,k}\|^2_{\rho_0}, \|\mathcal{L}^{(im)}_{k,j}\|^2_{\rho_0} : j < k\right) = \operatorname{Diag}\left((\mu_j + \mu_k)^{-1}, (\mu_j + \mu_k)^{-1} : j < k\right).$$

Lemma 9.4.7. *Let*

$$\mathcal{L}(\theta) := \sum_{j<k} \left(\operatorname{Re}(\zeta_{j,k})\mathcal{L}^{(re)}_{j,k} + \operatorname{Im}(\zeta_{j,k})\mathcal{L}^{(im)}_{j,k}\right) + \sum_i u_i \ell_i, \qquad \theta = (\vec{u}, \vec{\zeta}).$$

Consider the representation $(\pi, \mathcal{H})$ of $CCR(L^2(\rho_0), \sigma)$ and the normal state φ on $L^\infty(\mathbb{R}^{d-1}) \otimes \mathcal{B}(\mathcal{F})$ as defined in the previous section (cf. 9.47). Let $\tilde{\varphi}^\theta$ be the state defined by

$$\varphi^\theta(W(A)) := \exp\left(-\frac{1}{2}\|A, A\|_{\rho_0} + i(A, \mathcal{L}(\theta))_{\rho_0}\right), \qquad A \in L^2(\rho_0). \tag{9.48}$$

Then $\tilde{\varphi}^\theta$ is normal with respect to the representation $(\pi, \mathcal{H})$ and coincides with φ^θ (cf. (9.33)).

Remark 9.4.8. The expression (9.48) is clearly the quantum analogue of the characteristic function of the classical Gaussian shift experiment (9.10). Note in particular that the distribution of $B(A)$ with respect to φ_θ is the normal with variance $\|A\|^2_{\rho_0}$ centred at $(A, \mathcal{L}(\theta))_{\rho_0}$.

Proof. From (9.45) - (9.48), and by expressing A in the symplectic basis (9.46)

$$A = \sum_{j<k} (u_{j,k} t_{j,k} + v_{j,k} t_{k,j}) + \sum_i w_i \ell_i,$$

we get

$$\|A\|^2_{\rho_0} = w^T I_{\rho_0} w + \sum_{j<k} (u^2_{j,k} + v^2_{j,k}) \frac{\mu_j + \mu_k}{2(\mu_j - \mu_k)}, \tag{9.49}$$

$$(A, \mathcal{L}(\theta))_{\rho_0} = w^T I_{\rho_0} u + \sum_{j<k} \frac{u_{j,k}\mathrm{Re}(\zeta_{j,k}) + v_{j,k}\mathrm{Im}(\zeta_{j,k})}{\sqrt{2(\mu_j - \mu_k)}}, \tag{9.50}$$

which implies that the following decomposition holds

$$\varphi^\theta \cong N(I_{\rho_0} u, I_{\rho_0}) \otimes \bigotimes_{j<k} \varphi^{\zeta_{j,k}}_{j,k} := N(I_{\rho_0} u, I_{\rho_0}) \otimes \varphi^{\vec{\zeta}} \tag{9.51}$$

where we have used the following expression for the displaced thermal equilibrium states $\varphi^{\zeta_{j,k}}_{j,k} = \varphi^z_\beta$ defined in (9.29), with $\beta = \ln \mu_j/\mu_k$, $z = \zeta_{j,k}$

$$\varphi^{\zeta_{j,k}}_{j,k}\left(e^{i(u\mathbf{Q}+v\mathbf{P})}\right) = \exp\left(-(u^2+v^2)\frac{\mu_j + \mu_k}{4(\mu_j - \mu_k)} + i\frac{u\mathrm{Re}(\zeta_{j,k}) + v\mathrm{Im}(\zeta_{j,k})}{\sqrt{2(\mu_j - \mu_k)}}\right). \tag{9.52}$$

□

9.5 Explicit form of the channels and first steps of the proof

9.5.1 Second look at the irreducible representations of $SU(d)$

Before explaining the steps involved in the proof, let us take a closer look at the block states (9.22). Recall that we have the decomposition of Theorem 9.4.1 over Young diagrams with n boxes and

$$\rho^{\theta,n} = \bigoplus_{\lambda} \rho_{\lambda}^{\theta,n} \otimes \frac{\mathbf{1}_{\mathcal{K}_\lambda}}{M_n(\lambda)}.$$

Let $\{f_1, \ldots, f_d\}$ be the eigenvectors of ρ_0, i.e. the standard basis vectors of $\mathbb{C}^d$. Then the eigenvectors of $\rho_0^{\otimes n} = \rho^{0,n}$ are tensor products

$$f_{\mathbf{a}} := f_{a(1)} \otimes f_{a(2)} \otimes \cdots \otimes f_{a(n)},$$

and the eigenvalues $\prod_k \lambda_{a(k)}$ do not depend on the order of the vectors in the product.

Projecting onto a copy of $\mathcal{H}_\lambda$.

Our aim is to 'project' to an irreducible representation $\mathcal{H}_\lambda$ and obtain an explicit expression for the eigenvectors of the block components $\rho_\lambda^{\theta,n}$. Such a projection is not unique, in fact for any rank one operator $|v\rangle\langle u| \in \mathcal{B}(\mathcal{K}_\lambda)$ with $\langle u|v\rangle = 1$ we can define a (not necessarily orthogonal) projection $y = y^2$ on a copy of $\mathcal{H}_\lambda$

$$y_\lambda(u, v) := \mathbf{1}_{\mathcal{H}_\lambda} \otimes |v\rangle\langle u| : (\mathbb{C}^d)^{\otimes n} \to \mathcal{H}_\lambda \otimes |v\rangle.$$

However the action of $y_\lambda(u, v)$ on basis vectors $f_{\mathbf{a}}$ depends on a particular identification between $(\mathbb{C}^d)^{\otimes n}$ and the direct sum in Theorem 9.4.1. Therefore we need a direct way of defining such a projection and the key observation is that $y_\lambda(u, v)$ is a *minimal* projection in the algebra $\mathrm{Alg}(\tilde{\pi}_d(\tau) : \tau \in S(n))$, i.e. it cannot be decomposed into a sum of non-zero projections, and vice-versa any minimal projection is of this form. The following recipe (given without proof) shows how to construct minimal projections in the $S(n)$ group algebra. We recall that the group *-algebra $\mathcal{A}(S(n))$ is the linear space spanned by the group elements endowed with a product stemming from the group product, that is

$$a = \sum_{\tau \in S(n)} a(\tau)\tau, \quad b = \sum_{\varrho \in S(n)} b(\varrho)\varrho$$

implies

$$ab = \sum_{\tau,\varrho\in S(n)} a(\tau)b(\varrho)\tau\varrho = \sum_{\sigma\in S(n)} \left(\sum_{s\in S(n)} a(\sigma s^{-1})b(s) \right) \sigma,$$

and with adjoint $a^* = \sum_{\tau\in S(n)} a(\tau)\tau^{-1}$.

Let λ be a Young diagram with n boxes consider the (standard) Young tableau t in which the boxes are filled with the numbers $\{1,\ldots,n\}$ in increasing order from left to right along rows, starting with the top row and ending with the bottom row, as shown in the left-side tableau of Figure 9.2.

1	3	4	6
2	5		
7			

1	1	2	2	3
2	3	3		
3				

Figure 9.2: Left: a standard Young tableaux. Right: a semi-standard Young tableau for $d = 3$

Define the group algebra elements

$$P_\lambda = \sum_{\sigma\in\mathcal{R}_\lambda} \sigma, \qquad Q_\lambda = \sum_{\tau\in\mathcal{C}_\lambda} \mathrm{sgn}(\tau)\tau,$$

where $\mathcal{R}_\lambda$ is the $S(n)$-subgroup of permutation leaving the rows of t invariant, and $\mathcal{C}_\lambda$ is the subgroup of permutations leaving the columns of t invariant. Note that P_λ and Q_λ are self-adjoint elements of the $S(n)$ group algebra satisfying

$$P_\lambda P_\lambda = |\mathcal{R}_\lambda| P_\lambda = (\prod_{i=1}^{d} \lambda_i!) P_\lambda, \qquad Q_\lambda Q_\lambda = |\mathcal{C}(\lambda)| Q_\lambda = (\prod_{i=1}^{d} i^{\lambda_i - \lambda_{i+1}}) Q_\lambda. \tag{9.53}$$

The *Young symmetriser* is defined as

$$Y_\lambda := Q_\lambda P_\lambda.$$

Theorem 9.5.1. *Up to a scalar normalising factor, the Young symmetriser Y_λ is minimal projection in $\mathcal{A}(S(n))$ and $y_\lambda := q_\lambda p_\lambda = \tilde{\pi}_d(Q_\lambda)\tilde{\pi}_d(P_\lambda)$ projects onto a copy of $\mathcal{H}_\lambda \subset (\mathbb{C}^d)^{\otimes n}$.*

The action of the Young symmetriser y_λ on basis vectors $f_{\mathbf{a}} \in (\mathbb{C}^d)^{\otimes n}$ follows easily from the definition of Y_λ. For each $f_{\mathbf{a}}$ we fill the boxes of λ with the indices $a(k)$ going along rows from left to right, starting with the top row and finishing with the bottom one. For example, if $\lambda = \ydiagram{2,2}$ and $f_{\mathbf{a}} = f_2 \otimes f_2 \otimes f_1 \otimes f_2 \otimes f_1$ then $t_{\mathbf{a}} = \begin{array}{|c|c|c|}\hline 2&2&1\\\hline 2&1\\\cline{1-2}\end{array}$. $S(n)$ has an obvious action on the set of tableaux by permuting the *content* of the

boxes which are numbered from 1 to n in the standard way as in Figure 9.2. The action of the Young symmetriser $y_\lambda = q_\lambda p_\lambda$ on $f_{\mathbf{a}}$ is deduced from the action on the tableau $t_{\mathbf{a}}$: one first symmetrises with respect to components which are in the same row, and then antisymmetrises with respect to components in the same column. For example if $\lambda = \ydiagram{2,1}$ then

$$y_\lambda(f_2 \otimes f_1 \otimes f_3) = f_2 \otimes f_1 \otimes f_3 + f_1 \otimes f_2 \otimes f_3 - f_3 \otimes f_1 \otimes f_2 - f_3 \otimes f_2 \otimes f_1.$$

Finding a basis in $\mathcal{H}_\lambda$

By the previous Theorem the vectors $y_\lambda f_{\mathbf{a}}$ span $\mathcal{H}_\lambda$, but are not linearly independent. We show now how to select a basis (subset of linearly independent vectors spanning $\mathcal{H}_\lambda$). A *semistandard* Young tableau is a diagram filled with numbers in $\{1, \ldots, d\}$ such that the entries are non-decreasing along rows from left to right and increasing along columns from top to bottom, as in the right-side of Figure 9.2.

Theorem 9.5.2. *The vectors $y_\lambda f_{\mathbf{a}}$ for which $t_{\mathbf{a}}$ is a semistandard Young tableau form a (non-orthogonal) basis of the irreducible representation $(\pi_\lambda, \mathcal{H}_\lambda)$.*

Since the values in the rows are nondecreasing, there is a one-to-one correspondence between Young tableaux $t_{\mathbf{a}}$ and vectors $\mathbf{m} = (m_{i,j})_{1 \le i < j \le d}$ where $m_{i,j}$ is the number of j's appearing in line i of the Young tableau $t_{\mathbf{a}}$. Note that we need only consider $m_{i,j}$ for $j > i$, as there is no j in line i if $j < i$ (the columns are increasing), and the number of i in line i is $\lambda_i - \sum_{j=i+1}^{d} m_{i,j}$. For example, if $t_{\mathbf{a}} = \begin{array}{|c|c|c|c|c|}\hline 1&1&2&3&3\\\hline 2&3\\\cline{1-2} 3\\\cline{1-1}\end{array}$ then $\mathbf{m} = \{m_{1,2} = 1, m_{1,3} = 2, m_{2,3} = 1\}$.

By a slight abuse of notation we shall denote the corresponding vectors by $y_\lambda f_{\mathbf{m}}$ and the normalised vectors

$$|\mathbf{m}, \lambda\rangle := \mathcal{N}(\mathbf{m}, \lambda) y_\lambda f_{\mathbf{m}}, \tag{9.54}$$

where $\mathcal{N}(\mathbf{m}, \lambda) = 1/\|y_\lambda f_{\mathbf{m}}\|$. This constant is in general not easy to compute but we shall describe its asymptotic properties in section 9.7.4.

Using (9.53) we have

$$\langle y_\lambda f_{\mathbf{a}} | y_\lambda f_{\mathbf{b}} \rangle = \langle q_\lambda p_\lambda f_{\mathbf{a}} | q_\lambda p_\lambda f_{\mathbf{b}} \rangle = \langle p_\lambda f_{\mathbf{a}} | q_\lambda^2 p_\lambda f_{\mathbf{b}} \rangle = (\prod_{i=1}^{d} i^{\lambda_i - \lambda_{i+1}}) \langle p_\lambda f_{\mathbf{a}} | y_\lambda f_{\mathbf{b}} \rangle. \tag{9.55}$$

In order to get further simplifications, we examine some special vector states, that we shall call by analogy with the Fock spaces *finite-dimensional coherent states.*

The first is the special vector $|\mathbf{0}, \lambda\rangle$, the *highest weight vector* of the representation $(\pi_\lambda, \mathcal{H}_\lambda)$, which later on will play the role of the *finite-dimensional vacuum.* This

vector, as we have seen, corresponds to the semi-standard Young tableau where all the entries in row i are i. An immediate consequence is that

$$p_\lambda |f_\mathbf{0}\rangle = (\prod_{i=1}^{d} \lambda_i!)|f_\mathbf{0}\rangle. \tag{9.56}$$

Moreover $\langle f_\mathbf{0}|q_\lambda f_\mathbf{0}\rangle = 1$ since any column permutation produces a vector orthogonal to $f_\mathbf{0}$. Thus the normalised vector is:

$$|\mathbf{0}, \lambda\rangle = \frac{1}{\prod_{i=1}^{d} \lambda_i! \sqrt{i^{\lambda_i - \lambda_{i+1}}}} y_\lambda |f_\mathbf{0}\rangle. \tag{9.57}$$

The finite-dimensional coherent states are defined as $\pi_\lambda(U)|\mathbf{0}_\lambda\rangle$ for $U \in SU(d)$. From $[p_\lambda, \pi_\lambda(U)] = 0$ and (9.56), we get $p_\lambda \pi_\lambda(U)|\mathbf{0}_\lambda\rangle = (\prod_{i=1}^{d} \lambda_i!) U|\mathbf{0}_\lambda\rangle$, thus

$$\langle y_\lambda f_\mathbf{m}|\pi_\lambda(U)|\mathbf{0}, \lambda\rangle = \sqrt{\prod_{i=1}^{d} i^{\lambda_i - \lambda_{i+1}}} \langle p_\lambda f_\mathbf{m}|q_\lambda \pi_\lambda(U) f_\mathbf{0}\rangle \tag{9.58}$$

The latter expression holds for any linear combination of $f_\mathbf{m}$ on the left-hand side, in particular $\pi_\lambda(V)f_\mathbf{0}$ for another unitary operator V. In Lemma 9.7.11, we shall examine asymptotics of (9.58) for specific sequences of unitaries U when $n \longrightarrow \infty$. One of the main tools will be formula (9.88).

The following expressions of the dimensions of $\mathcal{K}_\lambda$ and $\mathcal{H}_\lambda$ are given without proof.

Let $g_{l,m}$ be the *hook length* of the box (l, m), defined as one plus the number of boxes under plus the number of boxes to the right. For example the diagram $(5, 3, 3)$ has the hook lengths : $\begin{array}{|c|c|c|c|c|}\hline 7&6&5&2&1\\\hline 4&3&2\\\cline{1-3} 3&2&1\\\cline{1-3}\end{array}$.

The dimension $M_n(\lambda)$ of $\mathcal{K}_\lambda$ is

$$\mathcal{M}(\vec{\lambda}) = \frac{n!}{\prod_{\substack{l=1\ldots d \\ m=1\ldots\lambda_l}} g_{l,m}},$$

and can be rewritten in the following form which is more adapted to our needs:

$$\mathcal{M}(\vec{\lambda}) = \binom{n}{\lambda_1, \ldots, \lambda_d} \prod_{\substack{l=1\ldots d \\ k=l+1\ldots d}} \frac{\lambda_l - \lambda_k + k - l}{\lambda_l + k - l}. \tag{9.59}$$

The dimension $\mathcal{D}(\lambda)$ of $\mathcal{H}_\lambda$ is:

$$\mathcal{D}(\lambda) = \prod_{\substack{i=1\ldots d \\ j=1\ldots\lambda_i}} \frac{j+d-i}{g_{i,j}}.$$

To summarise, we have defined a non-orthonormal basis $\{|\mathbf{m},\lambda\rangle\}$ of $\mathcal{H}_\lambda$ such that $|\mathbf{m},\lambda\rangle$ are eigenvectors of $\rho^{\vec{0},\vec{u},n}$ for all λ, with eigenvalues:

$$\langle \mathbf{m},\lambda|\rho^{\vec{0},\vec{u},n}|\mathbf{m},\lambda\rangle = \prod_{i=1}^{d}(\mu_i^{\vec{u},n})^{\lambda_i} \prod_{j=i+1}^{d}\left(\frac{\mu_j^{\vec{u},n}}{\mu_i^{\vec{u},n}}\right)^{m_{i,j}}, \tag{9.60}$$

where $\mu_i^{\vec{u},n} = \mu_i + u_i/\sqrt{n}$ for $1 \leq i \leq (d-1)$ and $\mu_d^{\vec{u},n} = \mu_d - (\sum_i u_i)/\sqrt{n}$.

The next step is to take into account the action of the unitary $U(\vec{\zeta})$. We define the automorphism

$$\Delta^{\vec{\zeta},n} : M((\mathbb{C}^d)^{\otimes n}) \to M((\mathbb{C}^d)^{\otimes n}),$$

by

$$\tau \mapsto \Delta^{\vec{\zeta},n}(\tau) = \mathrm{Ad}[U(\vec{\zeta},n)](\tau) := U(\vec{\zeta}/\sqrt{n})^{\otimes n}\,\tau\,U^*(\vec{\zeta}/\sqrt{n})^{\otimes n}. \tag{9.61}$$

Then we have $\rho^{\vec{\zeta},\vec{u},n} = \Delta^{\vec{\zeta},n}(\rho^{\vec{0},\vec{u},n})$. By Theorem 9.4.1 and using the decomposition (9.21), we get the blockwise action on irreducible components

$$\Delta^{\vec{\zeta},n}(\rho^{\otimes n}) = \bigoplus_\lambda \Delta_\lambda^{\vec{\zeta},n}(\rho_\lambda) \otimes \mathbf{1}_{\mathcal{K}_\lambda},$$

where $\Delta_\lambda^{\vec{\zeta},n} = \mathrm{Ad}[U_\lambda(\vec{\zeta},n)]$. In particular we have

$$\rho_\lambda^{\vec{\zeta},\vec{u},n} = \Delta_\lambda^{\vec{\zeta},n}(\rho_\lambda^{\vec{0},\vec{u},n}). \tag{9.62}$$

With these notations, we can set about building the channels T_n.

9.5.2 Description of T_n

We look for channels

$$T_n : M((\mathbb{C}^d)^{\otimes n}) \to L^1(\mathbb{R}^{d-1}) \otimes \mathcal{T}_1(\mathcal{F})$$

of the form:

$$T_n : \rho^{\theta,n} \longmapsto \sum_\lambda p_\lambda^{\theta,n}\,\tau_\lambda^n \otimes \left(V_\lambda \rho_\lambda^{\theta,n} V_\lambda^*\right). \tag{9.63}$$

Here, V_λ is an isometry from $\mathcal{H}_\lambda$ to $\mathcal{F}$, i.e. $V_\lambda^* V_\lambda = \mathbf{1}_{\mathcal{H}_\lambda}$. On the classical side, τ_λ^n is a probability law on $\mathbb{R}^{d-1}$. We may view τ^n as a Markov kernel (9.5) from the set of diagrams λ to $\mathbb{R}^{d-1}$.

The channel T_n can be described by the following sequence of operations. We first performs a 'which block' measurement over the irreducible representations and get a result λ. Then, on the one hand, we apply a classical randomisation to λ, and on the other hand we apply a channel depending on our result λ to the conditional state ρ_λ.

The underlying ideas are the following.

1). The probability distribution $p_\lambda^{\theta,n}$ is essentially a multinomial depending *only* on $\vec{u}$, as it can be deduced from (9.60) and (9.59). As we have seen in Example 9.3.1, this converges (in Le Cam sense) to a classical Gaussian shift experiment. Here, in order to obtain the strong norm convergence we need to smooth the discrete distribution into a continuous one with respect to the Lebesgue measure. We choose a particular smoothing distribution that will ensure the uniform L^1 convergence to the Gaussian model (Lemma 9.6.1).

Definition 9.5.3. *Let τ_λ^n be the probability density on $\mathbb{R}^{d-1}$ defined for all λ such that $\sum \lambda_i = n$, by:*

$$\tau_\lambda^n(\mathrm{d}x) = \tau_\lambda^n(x)\mathrm{d}x = \mathrm{d}x\, n^{(d-1)/2}\chi(A_{\lambda,n}), \tag{9.64}$$

where $A_{\lambda,n} = \{x \in \mathbb{R}^{d-1} : \ |n^{1/2}x_i + n\mu_i - \lambda_i| \leq 1/2, 1 \leq i \leq d-1\}$. We further denote

$$b_\lambda^{\theta,n} = p_\lambda^{\theta,n}\tau_\lambda^n,$$

depending on θ only through $\vec{u}$.

2). For the quantum part, we map the 'finite-dimensional vacuum' $|\mathbf{0}, \lambda\rangle$ to the Fock space vacuum $|\mathbf{0}\rangle$, and the basis vectors $|\mathbf{m}, \lambda\rangle$ of $\mathcal{H}_\lambda$ 'near' the basis vectors $|\mathbf{m}\rangle$ of the Fock space $\mathcal{F}$ (cf. definitions (9.54) and respectively (9.30)). Here we need to tackle the problem that $\{|\mathbf{m}, \lambda\rangle\}$ is not an orthonormal basis but only becomes so asymptotically. The following lemma provides the isometry V_λ appearing in (9.63).

Lemma 9.5.4. *Let $\eta < 2/9$. Suppose that $\lambda_i - \lambda_{i+1} \geq \delta n$ for all $1 \leq i \leq d$, with the convention $\lambda_{d+1} = 0$. Then for $n > n_0(\eta, \delta, d)$ there exists an isometry $V_\lambda : \mathcal{H}_\lambda \to \mathcal{F}$ such that, $V|\mathbf{0}, \lambda\rangle = |\mathbf{0}\rangle$ and for $0 < |\mathbf{m}| \leq n^\eta$,*

$$\langle \mathbf{m}|\, V_\lambda = \frac{1}{\sqrt{1 + (\tilde{C}n)^{(9\eta-2)/12}/\delta^{1/3}}}\, \langle \mathbf{m}, \lambda|$$

where $\tilde{C} = \tilde{C}(\eta, d)$ is a constant. More precisely, n_0 can be taken of the form $(C(d)/\delta^2)^{1/(1-3\eta)}$.

Proof. See section 9.7.3. The main tool is Lemma 9.7.13.

□

For Young diagrams which do not satisfy the assumption of the previous Lemma, the isometry V_λ can be defined arbitrarily. The reason is that those blocks have vanishing collective weight and can be neglected altogether (cf. Lemma 9.6.2).

From this operational description we conclude that T_n is a proper channel since τ^n is a Markov kernel and V_λ is an isometry. We then want to prove that $T_\lambda(\rho_\lambda^{\vec{0},\vec{u},n})$ is close to $\phi^{\mathbf{0}}$ and that the finite-dimensional operations $\Delta_\lambda^{\vec{\zeta},n}$ have almost the same action as the displacement operators D^ζ of the Fock space, cf. (9.29). Finite-dimensional coherent states and formula 9.28 will be the stepping stone to those results.

9.6 Main steps of the proof

9.6.1 Why T_n does the work

We shall break (9.36) in small manageable pieces. The result and brief explanatory remarks, repeating those in the derivation, are given from (9.67) on.

We introduce first a few shorthand notations: the restriction of T_n to the block λ is

$$T_\lambda : \rho_\lambda^{\theta,n} \mapsto V_\lambda \rho_\lambda^{\theta,n} V_\lambda^*,$$

so that

$$T_n : \rho^{\theta,n} \mapsto \sum_\lambda p_\lambda^{\theta,n} \tau_\lambda^n \otimes T_\lambda(\rho_\lambda^{\theta,n}) = \sum_\lambda b_\lambda^{\theta,n} \otimes \phi_\lambda^{\theta,n}.$$

We also define $T_\lambda^* : \phi \mapsto V_\lambda^* \phi V_\lambda$. and note that $T_\lambda^* T_\lambda = \mathrm{Id}_{\mathcal{H}_\lambda}$.

We expand (9.63) as

$$\begin{aligned} T_n(\rho^{\theta,n}) &= \sum_\lambda b_\lambda^{\theta,n} \otimes \phi_\lambda^{\theta,n} \\ &= \mathcal{N}(\vec{u}, V_\mu) \otimes \phi^{\vec{\zeta}} - \left(\mathcal{N}(\vec{u}, V_\mu) - \sum_\lambda b_\lambda^{\theta,n} \right) \otimes \phi^{\vec{\zeta}} - \sum_\lambda b_\lambda^{\theta,n} \otimes \left(\phi^{\vec{\zeta}} - \phi_\lambda^{\theta,n} \right). \end{aligned}$$

Proving (9.36) then amounts to proving

$$\sup_{\theta \in \Omega_{n,\epsilon}} \left\| \left(\mathcal{N}(\vec{u}, V_\mu) - \sum_\lambda b_\lambda^{\theta,n} \right) \otimes \phi^{\vec{\zeta}} + \sum_\lambda b_\lambda^{\theta,n} \otimes \left(\phi^{\vec{\zeta}} - \phi_\lambda^{\theta,n} \right) \right\|_1 \leq C n^{-\epsilon/\delta}.$$

We now use the triangle inequality to upper bound this norm by a sum of "elementary" terms to be treated separately in the following sections.

$$\begin{aligned}
&\left\| \left(\mathcal{N}(\vec{u}, V_\mu) - \sum_\lambda b_\lambda^{\theta,n} \right) \otimes \phi^{\vec{\zeta}} + \sum_\lambda b_\lambda^{\theta,n} \otimes \left(\phi^{\vec{\zeta}} - \phi_\lambda^{\theta,n} \right) \right\|_1 \leq \\
&\left\| \left(\mathcal{N}(\vec{u}, V_\mu) - \sum_\lambda b_\lambda^{\theta,n} \right) \otimes \phi^{\vec{\zeta}} \right\|_1 + \sum_\lambda \left\| b_\lambda^{\theta,n} \otimes \left(\phi^{\vec{\zeta}} - \phi_\lambda^{\theta,n} \right) \right\|_1 \leq \\
&\left\| \phi^{\vec{\zeta}} \right\|_1 \left\| \left(\mathcal{N}(\vec{u}, V_\mu) - \sum_\lambda b_\lambda^{\theta,n} \right) \right\|_1 + \sum_\lambda \left\| b_\lambda^{\theta,n} \right\|_1 \left\| \left(\phi^{\vec{\zeta}} - \phi_\lambda^{\theta,n} \right) \right\|_1 .
\end{aligned}$$

Since $\|\phi^{\vec{\zeta}}\|_1 = \|\mathcal{N}(\vec{u}, V_\mu)\|_1 = \|\phi_\lambda^{\theta,n}\| = 1$, we have $\left\| \left(\phi^{\vec{\zeta}} - \phi_\lambda^{\theta,n} \right) \right\|_1 \leq 2$. Similarly $\sum_\lambda \|b_\lambda^{\theta,n}\|_1 = 1$ because $\|b_\lambda^{\theta,n}\|_1 = p_\lambda^{\theta,n}$. We split the sum over λ in two parts, one for which it is expected that $\left\| \left(\phi^{\vec{\zeta}} - \phi_\lambda^{\theta,n} \right) \right\|_1$ is small, and the other on which the sum of all $\|b_\lambda^{\theta,n}\|_1$ is small. Specifically, define the set of *typical Young diagrams*

$$\Lambda_{n,\alpha} := \{\lambda : |\lambda_i - n\mu_i| \leq n^\alpha, 1 \leq i \leq d\}, \qquad \text{for } \alpha > 1/2, \tag{9.65}$$

then

$$\begin{aligned}
&\left\| T_n(\rho^{\theta,n}) - \mathcal{N}(\vec{u}, V_\mu) \otimes \phi^{\vec{\zeta}} \right\| \leq \\
&\left\| \mathcal{N}(\vec{u}, V_\mu) - \sum_\lambda b_\lambda^{\theta,n} \right\|_1 + \sup_{\lambda \in \Lambda_{n,\alpha}} \left\| \phi^{\vec{\zeta}} - \phi_\lambda^{\theta,n} \right\|_1 + 2 \sum_{\lambda \notin \Lambda_{n,\alpha}} \|b_\lambda^{\theta,n}\|_1 .
\end{aligned} \tag{9.66}$$

The first term corresponds to the convergence of the classical experiment in the Le Cam sense. If the second term is small, then on $\Lambda_{n,\alpha}$, the (purely quantum) family $\rho_\lambda^{\theta,n}$ is near the family $\phi^{\vec{\zeta}}$. The last term is small due to the concentration of $p_\lambda^{\theta,n}$ around the representations withshape $\lambda_i = n\mu_i$. In other words, the only representations that matter are those in $\Lambda_{n,\alpha}$.

The hardest term to dominate (notice that the two others are classical) is the second.

We transform it until we reach tractable fragments.

$$\begin{aligned}
\left\|\phi^{\vec{\zeta}} - \phi_\lambda^{\theta,n}\right\|_1 &= \left\|\phi^{\vec{\zeta}} - T_\lambda(\rho_\lambda^{\theta,n})\right\|_1 \\
&= \left\|D^{\vec{\zeta}}(\phi^{\vec{0}}) - [T_\lambda \Delta_\lambda^{\vec{\zeta},n} T_\lambda^*](T_\lambda(\rho_\lambda^{\vec{0},\vec{u},n}))\right\|_1 \\
&= \left\|D^{\vec{\zeta}}(\phi^{\vec{0}}) - D^{\vec{\zeta}}(T_\lambda(\rho_\lambda^{\vec{0},\vec{u},n})) + D^{\vec{\zeta}}(T_\lambda(\rho_\lambda^{\vec{0},\vec{u},n})) - [T_\lambda \Delta_\lambda^{\vec{\zeta},n} T_\lambda^*](T_\lambda(\rho_\lambda^{\vec{0},\vec{u},n}))\right\|_1 \\
&\leq \left\|D^{\vec{\zeta}}(\phi^{\vec{0}}) - D^{\vec{\zeta}}(T_\lambda(\rho_\lambda^{\vec{0},\vec{u},n}))\right\|_1 + \left\|[D^{\vec{\zeta}} - T_\lambda \Delta_\lambda^{\vec{\zeta},n} T_\lambda^*](T_\lambda(\rho_\lambda^{\vec{0},\vec{u},n}) - \phi^{\vec{0}})\right\|_1 \\
&\quad + \left\|[D^{\vec{\zeta}} - T_\lambda \Delta_\lambda^{\vec{\zeta},n} T_\lambda^*](\phi^{\vec{0}})\right\|_1 \\
&\leq 3\left\|T_\lambda(\rho_\lambda^{\vec{0},\vec{u},n}) - \phi^{\vec{0}}\right\|_1 + \left\|[D^{\vec{\zeta}} - T_\lambda \Delta_\lambda^{\vec{\zeta},n} T_\lambda^*](\phi^{\vec{0}})\right\|_1
\end{aligned}$$

where in the last inequality we have used the fact that the displacement operators are isometries.

Note that the first term does not depend on $\vec{\zeta}$ and the second term is small if the displacement operators $\Delta_\lambda^{\vec{\zeta},n}$ and $D^{\vec{\zeta}}$ have 'similar action' on an appropriate domain. Using the integral formula (9.28) for gaussian states ϕ_β and the fact that $\phi^{\vec{0}}$, is a tensor product of such states (cf. (9.51)) we bound the second term by

$$\left\|[D^{\vec{\zeta}} - T_\lambda \Delta_\lambda^{\vec{\zeta},n} T_\lambda^*](\phi^{\vec{0}})\right\|_1 \leq \int_{\mathbb{C}^{d(d-1)/2}} f(\vec{z}) \left\|[D^{\vec{\zeta}} - T_\lambda \Delta_\lambda^{\vec{\zeta},n} T_\lambda^*](|\vec{z}\rangle\langle\vec{z}|)\right\|_1 d\vec{z}$$

where

$$f(\vec{z}) = \prod_{i<j} \frac{\mu_i - \mu_j}{\pi \mu_j} \exp\left(-\frac{\mu_i - \mu_j}{\mu_j}|z_{i,j}|^2\right).$$

and $|\vec{z}\rangle\langle\vec{z}| = D^{\vec{z}}(|\mathbf{0}\rangle\langle\mathbf{0}|)$ is the multimode coherent state, so

$$[D^{\vec{\zeta}} - T_\lambda \Delta_\lambda^{\vec{\zeta},n} T_\lambda^*](|\vec{z}\rangle\langle\vec{z}|) = [D^{\vec{\zeta}} D^{\vec{z}} - T_\lambda \Delta_\lambda^{\vec{\zeta},n} T_\lambda^* D^{\vec{z}}](|\mathbf{0}\rangle\langle\mathbf{0}|).$$

Now, f is a probability density, and the norm in the integrand is dominated by two. By splitting the integral we obtain

$$\left\|[D^{\vec{\zeta}} - T_\lambda \Delta_\lambda^{\vec{\zeta},n} T_\lambda^*](\phi^{\vec{0}})\right\|_1 \leq 2\int_{\|\vec{z}\|>n^\beta} f(\vec{z})\mathrm{d}\vec{z} + \sup_{\|\vec{z}\|\leq n^\beta} \left\|[D^{\vec{\zeta}} D^{\vec{z}} - T_\lambda \Delta_\lambda^{\vec{\zeta},n} T_\lambda^* D^{\vec{z}}](|\mathbf{0}\rangle\langle\mathbf{0}|)\right\|_1.$$

By adding and subtracting additional terms

$$\begin{aligned}
D^{\vec{\zeta}} D^{\vec{z}} - T_\lambda \Delta_\lambda^{\vec{\zeta},n} T_\lambda^* D^{\vec{z}} =& D^{\vec{\zeta}+\vec{z}} - T_\lambda \Delta_\lambda^{\vec{\zeta}+\vec{z},n} T_\lambda^* \\
&+ T_\lambda \Delta_\lambda^{\vec{\zeta}+\vec{z},n} T_\lambda^* - T_\lambda \Delta_\lambda^{\vec{\zeta},n} \Delta_\lambda^{\vec{z},n} T_\lambda^* \\
&+ T_\lambda \Delta_\lambda^{\vec{\zeta},n} \Delta_\lambda^{\vec{z},n} T_\lambda^* - T_\lambda \Delta_\lambda^{\vec{\zeta},n} T_\lambda^* D^{\vec{z}}.
\end{aligned}$$

we deduce that

$$\begin{aligned}\left\|[D^{\vec{\zeta}} - T_\lambda \Delta_\lambda^{\vec{\zeta},n} T_\lambda^*](|\vec{z}\rangle\langle\vec{z}|)\right\|_1 \leq & \left\|[D^{\vec{\zeta}+\vec{z}} - T_\lambda \Delta_\lambda^{\vec{\zeta}+\vec{z},n} T_\lambda^*](|\mathbf{0}\rangle\langle\mathbf{0}|)\right\|_1 \\ & + \left\|[\Delta_\lambda^{\vec{\zeta}+\vec{z},n} - \Delta_\lambda^{\vec{\zeta},n}\Delta_\lambda^{\vec{z},n}](|\mathbf{0},\lambda\rangle\langle\mathbf{0},\lambda|)\right\|_1 \\ & + \left\|[\Delta_\lambda^{\vec{z},n} T_\lambda^* - T_\lambda^* D^{\vec{z}}](|\mathbf{0}\rangle\langle\mathbf{0}|)\right\|_1\end{aligned}$$

where the last two terms on the right side have been simplified using properties of $T_\lambda, T_\lambda^*, \Delta_\lambda^{\vec{\zeta},n}$. Notice that the first and third norms are essentially the same and the three terms are small if the action of $\Delta_\lambda^{\vec{\zeta}}$ is mapped into that of the displacement operators $D^{\vec{\zeta}}$.

Putting all this together, our 'expanded' form for (9.36) is

$$\sup_{\theta\in\Omega_{n,\beta,\gamma}} \left\|T_n(\rho^{\theta,n}) - \phi^{\vec{\zeta}} \otimes \mathcal{N}(\vec{u}, V_\mu)\right\| \tag{9.67}$$

$$\leq \sup_{\theta\in\Omega_{n,\beta,\gamma}} \left\|\left(\mathcal{N}(\vec{u}, V_\mu) - \sum_\lambda b_\lambda^{\theta,n}\right)\right\|_1 \tag{9.68}$$

$$+ 2 \sup_{\theta\in\Omega_{n,\beta,\gamma}} \sum_{\lambda\notin\Lambda_{n,\alpha}} \|b_\lambda^{\theta,n}\|_1 \tag{9.69}$$

$$+ 3 \sup_{\theta\in\Omega_{n,\beta,\gamma}} \sup_{\lambda\in\Lambda_{n,\alpha}} \left\|\phi^{\vec{0}} - T_\lambda(\rho_\lambda^{\vec{0},\vec{u},n})\right\|_1 \tag{9.70}$$

$$+ \sup_{\|\vec{z}\|\leq n^\beta} \sup_{\theta\in\Omega_{n,\beta,\gamma}} \sup_{\lambda\in\Lambda_{n,\alpha}} \left\|[D^{\vec{\zeta}+\vec{z}} - T_\lambda \Delta_\lambda^{\vec{\zeta}+\vec{z},n} T_\lambda^*](|\mathbf{0}\rangle\langle\mathbf{0}|)\right\|_1 \tag{9.71}$$

$$+ \sup_{\|\vec{z}\|\leq n^\beta} \sup_{\theta\in\Omega_{n,\beta,\gamma}} \sup_{\lambda\in\Lambda_{n,\alpha}} \left\|[D^{\vec{z}} - T_\lambda \Delta_\lambda^{\vec{z},n} T_\lambda^*](|\mathbf{0}\rangle\langle\mathbf{0}|)\right\|_1 \tag{9.72}$$

$$+ \sup_{\|\vec{z}\|\leq n^\beta} \sup_{\theta\in\Omega_{n,\beta,\gamma}} \sup_{\lambda\in\Lambda_{n,\alpha}} \left\|[\Delta_\lambda^{\vec{\zeta}+\vec{z},n} - \Delta_\lambda^{\vec{\zeta},n}\Delta_\lambda^{\vec{z},n}](|\mathbf{0},\lambda\rangle\langle\mathbf{0},\lambda|)\right\|_1 \tag{9.73}$$

$$+ 2 \int_{\|\vec{z}\|\geq n^\beta} f(\vec{z}) \mathrm{d}\vec{z}. \tag{9.74}$$

The last Gaussian tail term is less than $C\exp(-\delta n^{2\beta})$ where C depends only on the dimension d. Under the hypothesis $n^{2\beta} > 2/\delta$, this can be bounded again by $O(n^{-2\beta})$.

The following lemmas provide upper bounds for each of the terms. Before each lemma we remind the reader what is the significance of the bound. The proofs are gathered in section 9.7.

The classical part of the channel is a Markov kernel τ (see definition 9.5.3) mapping the 'which block' distribution $p_\lambda^{\theta,n}$ into the density $b_\lambda^{\theta,n}$ on $\mathbb{R}^{d-1}$ which is approaches

uniformly the gaussian shift experiment (9.68). Recall that $b_\lambda^{\theta,n}$ depends only on $\vec{u}$ and not on $\vec{\zeta}$, so that we have the same parameter set for the two classical experiments.

Lemma 9.6.1. *With the above definitions, for any ϵ, we have*

$$\sup_{\theta\in\Omega_{n,\beta,\gamma}} \left\| \mathcal{N}(\vec{u}, V_\mu) - \sum_\lambda b_\lambda^{\theta,n} \right\|_1 = O\left(n^{-1/4+\epsilon}/\delta, n^{-1/2+\gamma}/\delta\right).$$

The next lemma deals with (9.69) by showing concentration around Young diagrams λ in the 'typical subset' (9.65). This allows we to restrict to this set of diagrams in further estimates.

Lemma 9.6.2. *Let $\alpha - \gamma - 1/2 > 0$. Then, with the above definitions we have*

$$\sup_{\theta\in\Omega_{n,\beta,\gamma}} \sum_{\lambda\notin\Lambda_{n,\alpha}} \|b_\lambda^{\theta,n}\|_1 = O\left(n^{d^2}\exp(-n^{2\alpha-1}/2)\right),$$

with the $O(\cdot)$ term converging to zero.

The term (9.70) shows that when the rotation parameter is zero, the block states $\rho_\lambda^{\vec{0},\vec{u},n}$ are essentially thermal equilibrium states, as one would expect from the quantum Central Limit Theorem 9.4.6. However the convergence here is in norm rather than in distribution, and uniform over the various parameters.

Lemma 9.6.3. *Let $0 < \eta < 2/9$. With the above definitions, we have*

$$\sup_{\theta\in\Omega_{n,\beta,\gamma}} \sup_{\lambda\in\Lambda_{n,\alpha}} \left\| \phi^{\vec{0}} - T_\lambda(\rho_\lambda^{\vec{0},\vec{u},n}) \right\|_1 = O(n^{-1/2+\gamma+\eta}/\delta,\ n^{(9\eta-2)/24}/\delta^{1/6}, \exp(-\delta n^\eta)).$$

The terms (9.71) and (9.72) show that the 'finite dimensional coherent states' obtained by performing small rotations on the 'finite-dimensional vacuum' are uniformly close to their infinite dimensional counterparts, thus justifying the coherent state terminology.

Lemma 9.6.4. *Let $\epsilon > 0$ be such that $2\beta + \epsilon \leq \eta < 2/9$.*

Then,

$$\sup_{\|\vec{z}\|\leq n^\beta} \sup_{\|\vec{\xi}\|\leq n^{-1/2+2\beta}/\delta} \sup_{\theta\in\Omega_{n,\beta,\gamma}} \sup_{\lambda\in\Lambda_{n,\alpha}} \left\| [D^{\vec{\zeta}+\vec{z}} - T_\lambda \Delta_\lambda^{\vec{\zeta}+\vec{z},\vec{\xi},n} T_\lambda^*](|\mathbf{0}\rangle\langle\mathbf{0}|) \right\|_1 = R(n)$$

with

$$R(n)^2 = O\left(n^{(9\eta-2)/12}\delta^{-1/3}, n^{-1+2\beta+\eta}\delta^{-1}, n^{-1/2+3\beta+2\epsilon}\delta^{-3/2}, n^{-1+\alpha+2\beta}\delta^{-1}, n^{-1+\alpha+\eta}\delta^{-1}, n^{-1+3\eta}\delta^{-1}, n^{-\beta}\right) \quad (9.75)$$

For estimating the terms (9.71, 9.72), the case when $\vec{\xi} = \vec{0}$ is sufficient. This more general form is useful for the proof of Lemma 9.6.5. The unitary operation is defined as $\Delta_\lambda^{\vec{\zeta},\xi,n} := \mathrm{Ad}[U_\lambda(\vec{\zeta},\xi,n)]$ with $U(\vec{\zeta},\xi,n))$ the general $SU(d)$ element of (9.83).

Finally (9.73) shows that the 'finite-dimensional' displacement operators multiply as the corresponding displacement operators when acting on the vacuum.

Lemma 9.6.5. *With the above definitions, under the same hypotheses as in Lemma 9.6.4, we have*

$$\sup_{\|\vec{z}\|\leq n^\beta} \sup_{\theta\in\Omega_{n,\beta,\gamma}} \sup_{\lambda\in\Lambda_{n,\alpha}} \left\| [\Delta_\lambda^{\vec{\zeta}+\vec{z},n} - \Delta_\lambda^{\vec{\zeta},n}\Delta_\lambda^{\vec{z},n}](|\mathbf{0},\lambda\rangle\langle\mathbf{0},\lambda|) \right\|_1 = R(n)$$

with $R(n)$ given by equation (9.75).

>From the last three lemmas, together with the bound on the remainder integral (9.74) we obtain the following lemma which can be plugged into the bound (9.66):

Lemma 9.6.6. *With the above notations under the same hypotheses as in Lemma 9.6.4, we have*

$$\sup_{\theta\in\Omega_{n,\beta,\gamma}} \sup_{\lambda\in\Lambda_{n,\alpha}} \|\phi^{\vec{\zeta}} - \phi_\lambda^{\theta,n}\| = R(n) + O(n^{-1/2+\gamma+\eta}/\delta)$$

with $R(n)$ given by equation (9.75).

Gathering all these results and using the inequalities $\alpha - \gamma - 1/2 > 0$, $2\beta + \epsilon \leq \eta < 2/9$ we get the following relations between the error terms: $n^{-1/2+\beta+\eta/2}/\delta^{1/2} = o(n^{-1/2+3\eta/2}/\delta^{1/2})$ and $n^{-1/2+\alpha/2+\beta}/\delta^{1/2} = o(n^{-1/2+\alpha/2+\eta/2}/\delta^{1/2})$.

This yields the next theorem which provides the bound (9.36).

Theorem 9.6.7. *For any $\delta > 0$, $0 < \gamma < 1/4$, $\epsilon > 0$, $1/2 + \gamma < \alpha < 1$, $\eta < 2/9$, $0 < \beta < (\eta - \epsilon)/2$, the sequence of channels T_n satisfies*

$$\sup_{\theta\in\Omega_{n,\beta,\gamma}} \left\| T_n(\rho^{\theta,n}) - \phi \right\|_1 = O(n^{-1/4+3\beta/2+\epsilon}\delta^{-3/2} + n^{-1/2+\alpha/2+\eta/2}\delta^{-1/2} + n^{-1/2+3\eta/2}\delta^{-1/2} + n^{-\beta/2} + n^{-1/2+\gamma+\eta}/\delta + n^{(9\eta-2)/24}/\delta^{1/6} + \exp(-\delta n^\eta)) \quad (9.76)$$

For any given $0 < \delta < 1$, $\beta < 1/9$ and $\gamma < 1/4$, we can choose α, η, ϵ satisfying the above conditions, such that the right side is of order $O(n^{-\kappa})$, with $\kappa > 0$ depending on β, γ, δ.

9.6.2 Definition of S_n and proof of its efficiency

The channel S_n is essentially the inverse of T_n and as we shall see, (9.37) can be deduced from (9.36).

On the classical side we need a Markov kernel completing the equivalence between the family $p_\lambda^{\vec{u},n}$ and $\mathcal{N}(\vec{u}, V_\mu)$. Let σ^n be defined by

$$\sigma^n : x \in \mathbb{R}^{d-1} \mapsto \delta_{\lambda_x} \tag{9.77}$$

where λ_x is the Young diagram such that $\sum_1^d \lambda_i = n$, and $|n^{1/2}x_i + n\mu_i - \lambda_i| < 1/2$, for $2 \leq i \leq d$. No such diagram exists, we set λ_x to any admissible value, for example $(n, 0, \ldots, 0)$. Notice that with (9.64), $\sigma^n \circ \tau^n \circ \sigma^n = \sigma^n$. Moreover any probability on the λ such that $\sum_1^d \lambda_i = n$ is in the image of σ^n, so that $\sigma^n \circ \tau^n(p^{\theta,n}) = p^{\theta,n}$.

Lemma 9.6.8. *With the above definitions, for any ϵ, we have*

$$\sup_{\|\vec{u}\| \leq n^\gamma} \left\| \sigma^n \mathcal{N}(\vec{u}, V_\mu) - p^{\vec{u},n} \right\|_1 = O\left(n^{-1/2+\epsilon}/\delta, n^{-1/4+\gamma}/\delta\right).$$

Proof. See end of section 9.7.6. □

The channel S_n is given by the following sequence of operations acting on the two spaces of the product $L^1(\mathbb{R}^{d-1}) \otimes \mathcal{T}_1(\mathcal{F})$. Given a sample from the probability distribution $N(\vec{u}, V_\mu)$, we use the Markov kernel σ^n to produce a Young diagram λ. Conditional on λ we send the quantum part through the channel

$$S_\lambda : \phi \mapsto \tilde{S}_\lambda(\phi) \otimes \frac{\mathbf{1}_{\mathcal{K}_\lambda}}{M_n(\lambda)}$$

with

$$\tilde{S}_\lambda : \phi \mapsto T_\lambda^* \phi + (1 - \mathrm{Tr}(T_\lambda^*(\phi)))|\mathbf{0}, \lambda\rangle\langle\mathbf{0}, \lambda|.$$

The second term is rather arbitrary and ensures that $\tilde{S}_\lambda$ is trace preserving map. What is important is that for any density operator ρ_λ on the block λ, the operator $\tilde{S}_\lambda$ reverts the action of T_λ:

$$\begin{aligned}\tilde{S}_\lambda T_\lambda(\rho_\lambda) &= T_\lambda^* T_\lambda(\rho_\lambda) + (1 - \mathrm{Tr}(T_\lambda^* T_\lambda(\rho_\lambda)))|\mathbf{0}, \lambda\rangle\langle\mathbf{0}, \lambda| \\ &= \rho_\lambda + (1 - \mathrm{Tr}(\rho_\lambda))|\mathbf{0}, \lambda\rangle\langle\mathbf{0}, \lambda| \\ &= \rho_\lambda.\end{aligned}$$

Now

$$S_n(\mathcal{N}(\vec{u}, V_\mu) \otimes \phi^{\vec{\zeta}}) = \bigoplus [\sigma^n \mathcal{N}(\vec{u}, V_\mu)](\lambda) \tilde{S}_\lambda(\phi^{\vec{\zeta}}) \otimes \frac{\mathbf{1}_{\mathcal{K}_\lambda}}{M_n(\lambda)}.$$

and with the notation $\sigma^n \mathcal{N}_\lambda^{\vec{u}} := [\sigma^n \mathcal{N}(\vec{u}, V_\mu))](\lambda)$ and $q_\lambda^{\vec{u},n} := \min(\sigma^n \mathcal{N}_\lambda^{\vec{u}}, p_\lambda^{\vec{u},n})$ we have

$$S_n(\phi^{\vec{\zeta}} \otimes \mathcal{N}(\vec{u}, V_\mu)) - \rho^{\theta,n}$$
$$= \bigoplus_\lambda \left\{ q_\lambda^{\vec{u},n}(\tilde{S}_\lambda(\phi^{\vec{\zeta}}) - \rho_\lambda^{\theta,n}) + (\sigma^n \mathcal{N}_\lambda^{\vec{u}} - q_\lambda^{\vec{u},n})\tilde{S}_\lambda(\phi^{\vec{\zeta}}) - (p_\lambda^{\vec{u},n} - q_\lambda^{\vec{u},n})\rho_\lambda^{\theta,n} \right\} \otimes \frac{\mathbf{1}_{\mathcal{K}_\lambda}}{M_n(\lambda)}.$$

Taking L^1 norms, and using that all ϕ's and ρ's have trace 1 and that channels (such as $\tilde{S}_\lambda$) are trace preserving, we get the bound:

$$\begin{aligned}
&\left\| S_n(\phi^{\vec{\zeta}} \otimes \mathcal{N}(\vec{u}, V_\mu)) - \rho^{\theta,n} \right\|_1 \\
&\leq \sum_\lambda \left\| q_\lambda^{\vec{u},n}(\tilde{S}_\lambda(\phi^{\vec{\zeta}}) - \rho_\lambda^{\theta,n}) \right\|_1 + \sum_\lambda \left| \sigma \mathcal{N}_\lambda^{\vec{u}} - p_\lambda^{\vec{u},n} \right| \\
&\leq 2 \sum_{\lambda \notin \Lambda_{n,\alpha}} q_\lambda^{\vec{u},n} + \sup_{\lambda \in \Lambda_{n,\alpha}} \left\| \tilde{S}_\lambda(\phi^{\vec{\zeta}}) - \rho_\lambda^{\theta,n} \right\|_1 + \left\| \sigma^n \mathcal{N}(\vec{u}, V_\mu) - p^{\vec{u},n} \right\|_1 \\
&\leq 2 \sum_{\lambda \notin \Lambda_{n,\alpha}} q_\lambda^{\vec{u},n} + \sup_{\lambda \in \Lambda_{n,\alpha}} \left\| \phi^{\vec{\zeta}} - T_\lambda(\rho_\lambda^{\theta,n}) \right\|_1 + \left\| \sigma^n \mathcal{N}(\vec{u}, V_\mu) - p^{\vec{u},n} \right\|_1 .
\end{aligned}$$

Now the first term is smaller than the remainder term of the gaussian outside a ball whose radius is n^α. Hence this term is going to zero faster than any polynomial, independently on δ and $\vec{u}$ for $\|\vec{u}\| \leq n^\gamma$. The second term is treated in Lemma 9.6.6 (recalling that $\phi_\lambda^{\theta,n} = T_\lambda(\rho_\lambda^{\theta,n})$), and the third term is treated in Lemma 9.6.8.

This ends the proof of (9.37).

□

9.7 Technical proofs

9.7.1 Proof of Theorem 9.4.4

General needed theorems

We shall work in the algebraic setting of quantum mechanics. The main inspiration here is van der Vaart, for providing classical analogues, and Paulsen (1987) for technical results on C^*-algebras.

A general quantum-classical system can be characterized by a unital C^*-algebra $\mathcal{A}$, the algebra of observables. States are then positive norm-one functionals. They live

in the topological dual space $\mathcal{A}^*$, that is notably a Banach space. General channels are identity-preserving completely positive maps from $\mathcal{A}$ to another algebra $\mathcal{A}_2$.

This setting is slightly more general than the one in Section 9.2.1, since the bidual of a C^* algebra is a von Neumann algebra.

In this setting, we may define measurements with values in a Polish space $(\mathbb{D}, \mathcal{B})$ in the following way, that reduces to the usual definition if $\mathcal{A} = \mathcal{B}(H)$, and to Markov kernels in the classical case, where the algebra consists in continuous bounded functions $\mathcal{A} = C_b(\mathcal{X})$:

Definition 9.7.1. *A* measurement *from $\mathcal{A}$ on $(\mathbb{D}, \mathcal{B})$ is a set $\{M(B)\}_{B\in\mathcal{B}}$ of elements in $\mathcal{A}^{**}$, such that:*

- *$M(B)$ is non-negative for all B.*
- *$M(\mathbb{D})$ is the identity, that is $M(\mathbb{D})\phi = \phi(\mathbf{1}_{\mathcal{A}})$ for all $\phi \in \mathcal{A}^*$.*
- *For all pairwise disjoint $\{B_i\}_{i\in\mathbb{N}}$, we have $\sum M(B_i) = M(\bigcup B_i)$.*

The measurement M acts on $\phi \in \mathcal{A}^*$ by yielding a probability measure on $\mathbb{D}$ defined through:

$$M\phi(B) = M(B)\phi \qquad \text{for all } B \in \mathcal{B}. \tag{9.78}$$

As a technical tool, we shall also consider *generalized measurements*, that correspond to generalized procedures in classical Le Cam theory. A generalized measurement M on a Polish space is an identity-preserving completely positive map from $C_b(\mathbb{D})$ to $\mathcal{A}$. Alternatively, in the dual picture, it is a positive map from $\mathcal{A}^*$ to $C_b(\mathbb{D})^*$ such that $M\phi(\mathbf{1}_{\mathbb{D}}) = \phi(\mathbf{1}_A)$ for all ϕ. The latter condition will be called *trace-preserving.*

The reason why we consider generalized measurements is that we get compactness.

Lemma 9.7.2. *Every net $\{M_\alpha\}_{\alpha\in A_1}$ of generalized measurements admits a converging subnet $\{M_\alpha\}_{\alpha\in A_2}$ with limit M, such that for any $\phi \in \mathcal{A}^*$ and every $f \in C_b(\mathbb{D})$, we have:*

$$M_\alpha\phi(f) \underset{A_2}{\longrightarrow} M\phi(f).$$

Moreover, M is a generalized measurement.

Proof. This lemma essentially amounts to applying the following remark, a reformulation of Lemmas 7.1 and 7.2 in the book by Paulsen (1987).

Lemma 9.7.3. *For any two Banach spaces X and Y, with Y^* the dual of Y, the set of linear operators $L(X, Y^*)$ is compact for the* bounded weak *topology. A bounded net $\{L_\alpha\}$ converges to L in bounded weak topology if and only if*

$$L_\alpha(x)(y) \to L(x)(y), \qquad \forall x \in X,\, y \in Y.$$

We take as X the space $\mathcal{A}^*$ and as Y the bounded continuous functions $C_b(\mathbb{D})$. All generalized measurements have norm one, so the net is bounded. We therefore have a limit M in $L(\mathcal{A}^*, C_b(\mathbb{D}))^*$ satisfying the pointwise convergence.

We still need to establish positivity and preservation of trace. Since they are satisfied by all M_α, both follow immediately from the pointwise convergence. □

Any Polish space admits a Polish compactification $\overline{\mathbb{D}}$. Now, any continuous bounded function on $\overline{\mathbb{D}}$ can be restricted to a continuous bounded function on $\mathbb{D}$. The restriction is linear positive and unit-preserving, so that any generalized measurement M on $\overline{\mathbb{D}}$ yields a generalized measurement $M_{\overline{\mathbb{D}}}$ on $\overline{\mathbb{D}}$. By compactness, $C_b(\overline{\mathbb{D}}) = C_c(\overline{\mathbb{D}})$ the continuous functions with compact support. So that, by the Riesz representation theorem, $M_{\overline{\mathbb{D}}}\phi$ can be seen as a Radon signed measure on $\overline{\mathbb{D}}$, positive if ϕ is positive, and with the same norm as ϕ. In particular, if ϕ is a state, then $M_{\overline{\mathbb{D}}}\phi$ is a Radon probability measure.

Hence, for any Borelian B, and any ϕ, we may define $M_{\overline{\mathbb{D}}}(B)\phi = M_{\overline{\mathbb{D}}}\phi(\mathbf{1}_B)$. This $M_{\overline{\mathbb{D}}}(B)$ is clearly linear, bounded, and non-negative. Moreover, since $M_{\overline{\mathbb{D}}}\phi$ is a *bona fide* measure, we have the σ-additivity. Finally $M_{\overline{\mathbb{D}}}(\overline{\mathbb{D}})\phi = \phi(\mathbf{1}_{\mathcal{A}})$ by definition. So that $M_{\overline{\mathbb{D}}}$ is always a *bona fide* measurement on $\overline{\mathbb{D}}$, satisfying Definition 9.7.1.

We may now define $M_{\overline{\mathbb{D}}|\mathbb{D}}$ the restriction of $M_{\overline{\mathbb{D}}}$ to $\mathbb{D}$, by $M_{\overline{\mathbb{D}}|\mathbb{D}}\phi = M_{\overline{\mathbb{D}}}\phi_{|\mathbb{D}}$, that is restricting the output Radon measures to $\mathbb{D}$.

We can now characterise when M is a measurement satisfying Definition 9.7.1. The obstruction is having mass at infinity, or equivalently being additive but not σ-additive.

Lemma 9.7.4. *With the above definitions, the following statements are equivalent:*

(i) The generalized measurement M is a measurement.

(ii) For all $\phi \in \mathcal{A}^*$, the result $M\phi$ is a signed Radon measure.

(iii) $M_{\overline{\mathbb{D}}|\mathbb{D}} = M$.

(iv) There is no mass at infinity, that is $M_{\overline{\mathbb{D}}}(\overline{\mathbb{D}}\backslash\mathbb{D}) = 0$.

Proof.

$(i) \Rightarrow (ii)$: Immediate consequence of Definition 9.7.1.

$(ii) \Rightarrow (iii)$: Since $M\phi$ is a Radon measure, we may define $M(B)$ as we have defined $M_{\overline{\mathbb{D}}}(B)$. We immediately get $M = M_{\overline{\mathbb{D}}|\mathbb{D}}$.

$(iii) \Rightarrow (iv)$: For any ϕ, for any $f \in C_c(\overline{\mathbb{D}})$, we know that

$$M_{\overline{\mathbb{D}}}(\mathbb{D})\phi(f) = M_{\overline{\mathbb{D}}|\mathbb{D}}\phi(f_{|\mathbb{D}}) = M\phi(f_{|\mathbb{D}}) = M_{\overline{\mathbb{D}}}\phi(f) = M_{\overline{\mathbb{D}}}(\mathbb{D})\phi(f) + M_{\overline{\mathbb{D}}}(\overline{\mathbb{D}}\backslash\mathbb{D})\phi(f).$$

So that $M_{\overline{\mathbb{D}}}(\overline{\mathbb{D}}\backslash\mathbb{D}) = 0$.

$(iv) \Rightarrow (i)$: Since $M_{\overline{\mathbb{D}}}(\mathbb{D}) = M_{\overline{\mathbb{D}}}(\overline{\mathbb{D}}) - M_{\overline{\mathbb{D}}}(\overline{\mathbb{D}}\backslash\mathbb{D})$, taking $M'(B) = M_{\overline{\mathbb{D}}}(B)$ for all $B \in \mathcal{B}$ defines a measurement. We must prove it is the original measurement.

Equality on positive ϕ suffices, since any continuous linear functional may be decomposed in positive and negative parts. Now $M'\phi(f) = M\phi(f)$ for all $f = g_{|\mathbb{D}}, g \in C_c(\overline{\mathbb{D}})$. Let us consider $f \in C_b(\mathbb{D})$. Since it is a lower semicontinuous function on a dense subspace of $\overline{\mathbb{D}}$, it can be extended to a lower semicontinuous function on $\overline{\mathbb{D}}$. Ditto upper semicontinuous. So that there are sequences of functions g_n^- and g_n^+ in $C_c(\overline{\mathbb{D}})$ such that

- g_n^- is non-decreasing and g_n^+ non-increasing in n.
- g_n^- is bounded from below by $\inf_{\mathbb{D}} f$ and g_n^+ is bounded from above by $\sup_{\mathbb{D}} f$.
- g_n^- and g_n^+ converge pointwise to f on $\mathbb{D}$.

We write f_n^- and f_n^+ for the restrictions to $\mathbb{D}$. Since $M\phi$ is nonnegative, we have

$$\lim_n M\phi(f_n^-) \leq M\phi(f) \leq \lim_n M\phi(f_n^+).$$

By monotone convergence for Radon measures, we also have

$$\lim_n M'\phi(f_n^-) = M'\phi(f) = \lim_n M'\phi(f_n^+).$$

Since $M'\phi(f_n^-) = M\phi(f)$, we have $M\phi(f) = M'\phi(f)$ for all $f \in C_b(\mathbb{D})$.

For any experiment $\mathcal{E} = \left\{\phi^\theta, \theta \in \Theta\right\}$, for any finite subset I of Θ, we write $\mathcal{E}^I$ for the restricted experiment $\left\{\phi^\theta, \theta \in I\right\}$.

Theorem 9.7.5 (Asymptotic representation theorem). *Let us consider a net of experiments $\mathcal{E}_\alpha = \{\phi_\alpha^\theta, \theta \in \Theta\}$, and an experiment $\mathcal{E} = \{\phi^\theta \in \mathcal{A}^*, \theta \in \Theta\}$, such that for any finite subset I of the parameters set Θ, the deficiency of the restricted experiment $\mathcal{E}^I$ goes to zero, that is $\delta(\mathcal{E}^I, \mathcal{E}_\alpha^I) \underset{\alpha}{\to} 0$.*

Suppose that we have generalized measurements M_α to a Borel decision space $(\mathbb{D}, \mathcal{B})$, such that for any $\theta \in \Theta$, the results converge:

$$M_\alpha \phi_\alpha^\theta = Q^\theta \in C_b(\mathbb{D})^*.$$

Then there is a generalized measurement M on $\mathcal{A}^*$ such that $M\phi^\theta = Q^\theta$ for all θ.

Moreover, if all Q^θ are Radon measures, then we may take M as a measurement, in the sense of Definition 9.7.1.

Proof. By definition of the deficiency, since $\delta(\mathcal{E}^I, \mathcal{E}_\alpha^I) \underset{\alpha}{\to} 0$, there are channels S_α^I such that $\sup_{\theta \in I} \left\| S_\alpha^I(\phi^\theta) - \phi_\alpha^\theta \right\| \underset{\alpha}{\to} 0$ for all finite I.

We may then consider the generalized measurements on $\mathcal{A}^*$ defined by $P_\alpha^I = M_\alpha \circ S_\alpha^I$.

For each I, we thus obtain a have a net of generalized measurements $\left\{P_\alpha^I\right\}_\alpha$ indexed by α. By Lemma 9.7.2, it admits a converging subsequence to a generalized measurement P^I. We order the finite subsets by $I \leq J$ if $I \subset J$, and thus obtain a net of measurements $\left\{P^I\right\}_I$, so that it also has a converging subsequence to a generalized measurement M.

Now, for any θ, for any $f \in C_b(\mathrm{d})$, we have:

$$\begin{aligned} M\phi^\theta(f) &= \lim_I P_\alpha^I \phi^\theta(f) \\ &= \lim_I \lim_\alpha P_\alpha^I \phi^\theta(f) \\ &= \lim_I \lim_\alpha M_\alpha^I \phi_\alpha^\theta(f) + \epsilon(\alpha, I, \theta, f) \\ &= Q^\theta(f), \end{aligned}$$

where we have used

$$\begin{aligned} |\epsilon(\alpha, I, \theta, f)| &= \left| M_\alpha^I [S_\alpha^I(\phi^\theta) - \phi_\alpha^\theta](f) \right| \\ &\leq \left\| M_\alpha^I \right\| \|f\| \left\| S_\alpha^I(\phi^\theta) - \phi_\alpha^\theta \right\| \\ &\underset{\alpha}{\to} 0, \end{aligned}$$

since the first two norms are constant and the third vanishes.

Hence $M\phi^\theta = Q^\theta$ and we have established our main statement.

Suppose now that all Q^θ are Radon measures. Let us consider a Polish compactification $\overline{\mathbb{D}}$ of our Polish space $\mathbb{D}$, and define $M_{\overline{\mathbb{D}}}$ from M as we have above, that is restricting $M\phi$ to $C_c(\overline{\mathbb{D}}) \subset C_b(\mathbb{D})$. Then $M_{\overline{\mathbb{D}}}$ is a true measurement and $Q^\theta_{\overline{\mathbb{D}}} M_{\overline{\mathbb{D}}}\phi^\theta$ is a Radon measure for all θ. Now, for any $f \in C_c(\overline{\mathbb{D}})$, we have $Q^\theta_{\overline{\mathbb{D}}}(f) = Q^\theta(f_{|\mathbb{D}}) = \overline{Q^\theta}(f)$ where $\overline{Q^\theta}$ is defined as the Radon measure such that $\overline{Q^\theta}_{|\mathbb{D}} = Q^\theta$ and $\overline{Q^\theta}(\overline{\mathbb{D}}\backslash\mathbb{D}) = 0$. By uniqueness in the Riesz representation theorem, we obtain $Q^\theta_{\overline{\mathbb{D}}} = \overline{Q^\theta}$.

Hence, $M_{\overline{\mathbb{D}}}(\overline{\mathbb{D}}\backslash\mathbb{D})\phi^\theta = 0$ for all $\theta \in \Theta$, and we may modify $M_{\overline{\mathbb{D}}}$ by choosing any point $d \in \mathbb{D}$, and define

$$\begin{aligned} M'(\overline{\mathbb{D}}\backslash\mathbb{D}) &= 0 \\ M'(B) &= M(B) + M(\overline{\mathbb{D}}\backslash\mathbb{D}) && \text{for all } B \in \mathcal{B} \text{ with } d \in B, \\ M'(B) &= M(B) && \text{for all } B \in \mathcal{B} \text{ with } d \notin B. \end{aligned} \tag{9.79}$$

Then M' restricted to $\mathbb{D}$ is still a true measurement $M'_{|\mathbb{D}}$, and $M'_{|\mathbb{D}}\phi^\theta = M\phi^\theta = Q^\theta$ for all $\theta \in \Theta$. □

We shall use this representation theorem to prove a minimax theorem.

We consider loss functions as sets of functions $l_\theta : \mathbb{D} \to [0, \infty]$ for all $\theta \in \Theta$. We shall always require that all l_θ be lower semicontinuous.

A loss function is said to be subcompact if, for all θ, the sets $\{y : l_\theta(y) \leq c\}$ are either null, compact or the whole decision space $\mathbb{D}$, for all c.

Definition 9.7.6. *Let l be a loss function and an experiment $\mathcal{E} = \{\phi^\theta, \theta \in \Theta\}$. We define the* risk of a generalized measurement M at point θ *as*

$$R_\theta(M) = \sup_{\substack{f \in C_b(\mathbb{D}) \\ f \leq l_\theta}} M\phi^\theta(f),$$

and consequently the (maximum) risk *of M as*

$$R_{max}(M) = \sup_\theta R_\theta(M).$$

This definition agrees with the usual definition of the risk 9.9 for M a true measurement if l_θ is equal to the supremum of the smaller bounded functions, that is, if it is lower semicontinuous. In general, it is a lower bound to $\int_{\mathbb{D}} l_\theta M\phi^\theta(\mathrm{d}b)$ if it is defined.

Theorem 9.7.7 (Asymptotic minimax theorem). *Consider a sequence of experiments $\mathcal{E}_n$ and an experiment $\mathcal{E}$, such that $\mathcal{E}_n^I$ converges to $\mathcal{E}^I$ for all finite subset $I \in \Theta$.*

Then, for any sequence of generalized measurements M_n, the supremum of the limit risks of the experiments is more than the minimax risk of the limit experiment:

$$\sup_I \liminf_n \sup_{\theta \in I} R_\theta(M_n) \geq \inf_M \sup_{\theta \in \Theta} R_\theta(M), \tag{9.80}$$

where the second infimum is over all generalized measurements. This infimum is a achieved.

Moreover, if the loss function is subcompact, the infimum on the right-hand side is achieved for a true measurement.

Proof. Let us define by induction n_I as the smallest $n \in \mathbb{N}$ such that $\sup_{\theta \in I} R_\theta(M_{n_I}) < \liminf_n \sup_{\theta \in I} R_\theta(M_{n_I}) + 1/\#I$, and $n_I \geq n_J$ for all $J \subset I$. Then $\{n_I : I \subset \Theta\}$ is a subnet such that

$$\sup_I \liminf_n \sup_{\theta \in I} R_\theta(M_n) = \lim_I \sup_{\theta \in I} R_\theta(M_{n_I}).$$

Since the finite experiments $\mathcal{E}_n^I$ converge to $\mathcal{E}^I$, there are channels S_n^I such that $\left\| S_n^I(\phi^\theta) - \phi_n^\theta \right\| \underset{n}{\to} 0$ for all $\theta \in I$.

We may then build generalized measurements $P^I = M_{n_I}^I \circ S_{n_I}^I$. Using Lemma 9.7.2, this net admits a converging subsequence to a generalized measurement M.

We may then write:

$$\begin{aligned}
R_\theta(M) &= \sup_{\substack{f \in C_b(\mathbb{D}) \\ f \leq l_\theta}} M\phi^\theta(f) \\
&= \sup_{\substack{f \in C_b(\mathbb{D}) \\ f \leq l_\theta}} \lim_I P^I \phi^\theta(f) \\
&= \sup_{\substack{f \in C_b(\mathbb{D}) \\ f \leq l_\theta}} \lim_I M_{n_I} \phi_{n_I}^\theta(f) + M_{n_I}[S_n^I(\phi^\theta) - \phi_{n_I}^\theta](f) \\
&= \sup_{\substack{f \in C_b(\mathbb{D}) \\ f \leq l_\theta}} \lim_I M_{n_I} \phi_{n_I}^\theta(f) \\
&\leq \lim_I \sup_{\substack{f \in C_b(\mathbb{D}) \\ f \leq l_\theta}} M_{n_I} \phi_{n_I}^\theta(f) \\
&\leq \lim_I \sup_{\theta \in I} R_\theta(M_{n_I}).
\end{aligned}$$

We have thus proved our main statement.

The second statement is a mere application of the following complete class theorem. □

Theorem 9.7.8 (Complete class theorem). *If the loss function is subcompact, then the true measurements are a complete class, that is, for any generalized measurement N, there is a true measurement N such that $R_\theta(N) \leq R_\theta(M)$ for all $\theta \in \Theta$.*

Proof. Let us consider a Polish compactification $\overline{\mathbb{D}}$ of $\mathbb{D}$. Then, l_θ admits a maximal lower semicontinuous extension to $\overline{\mathbb{D}}$, given by

$$\bar{l}_\theta(y) = \sup_{\substack{f \in C_c(\overline{\mathbb{D}}) \\ f_{|\mathbb{D}} \leq l_\theta}} f(y).$$

Since l_θ is subcompact, we know that $\bar{l}_\theta(y) \geq \bar{l}_\theta(x)$ for all $y \in \overline{\mathbb{D}} \backslash \mathbb{D}$ and $x \in \mathbb{D}$.

Since $f \in C_c(\overline{\mathbb{D}}) \Rightarrow f_{|\mathbb{D}} \in C_b(\mathbb{D})$, we have, with the former notation $N_{\overline{\mathbb{D}}}$ for the corresponding restriction of N:

$$\begin{aligned} R_\theta(N) &= \sup_{\substack{f \leq l_\theta \\ f \in C_b(\mathbb{D})}} N\phi^\theta(f) \\ &\geq \sup_{\substack{f \leq \bar{l}_\theta \\ f_{|\mathbb{D}} \in C_c(\overline{\mathbb{D}})}} N_{\overline{\mathbb{D}}}\phi^\theta(f) \\ &= N_{\overline{\mathbb{D}}}\phi^\theta(\bar{l}_\theta), \end{aligned}$$

where we view $N_{\overline{\mathbb{D}}}\phi^\theta$ as a Radon measure on the last line, and use the monotone convergence theorem.

We know that $N_{\overline{\mathbb{D}}}$ is a true measurement on $\overline{\mathbb{D}}$. Let us now consider a modification $M_{\overline{\mathbb{D}}}$ of $N_{\overline{\mathbb{D}}}$, defined like equation (9.79). That is, we transfer the mass on $\overline{\mathbb{D}} \backslash \mathbb{D}$ to some point in $x \in \mathbb{D}$. Since $l_\theta(x) \leq l_\theta(y)$ for any $y \in \overline{\mathbb{D}} \backslash \mathbb{D}$, we get $N_{\overline{\mathbb{D}}}\phi^\theta(\bar{l}_\theta) \leq M_{\overline{\mathbb{D}}}\phi^\theta(\bar{l}_\theta)$.

Now $M_{\overline{\mathbb{D}}|\mathbb{D}} \hat{=} M$ is a true measurement on $\mathbb{D}$, and $M\phi^\theta$ is a Radon probability measure for all θ.

As a consequence, we may use the monotone convergence theorem to get to the second line, and the fact that $M(\overline{\mathbb{D}} \backslash \mathbb{D}) = 0$ to get to the third line, of the following

calculation:

$$\begin{aligned} R_\theta(M) &= \sup_{\substack{f \le l_\theta \\ f \in C_b(\mathbb{D})}} M\phi^\theta(f) \\ &= \sup_{\substack{f \le \bar{l}_\theta \\ f_{|\mathbb{D}} \in C_c(\overline{\mathbb{D}})}} M\phi^\theta(f_{|\mathbb{D}}) \\ &= \sup_{\substack{f \le \bar{l}_\theta \\ f_{|\mathbb{D}} \in C_c(\overline{\mathbb{D}})}} M_{\overline{\mathbb{D}}}\phi^\theta(f) \\ &= M_{\overline{\mathbb{D}}}\phi^\theta(\bar{l}_\theta) \\ &\le N_{\overline{\mathbb{D}}}\phi^\theta(\bar{l}_\theta) \\ &\le R_\theta(N). \end{aligned}$$

This ends the proof. □

We shall now prove a quantum Hunt-Stein theorem. The first version of this theorem, for compact groups in the density operator setting, is due to Holevo (1982). A version for non-compact groups was proved by Ozawa (1980). Since Ozawa usually works with convex structures, his theorem is probably even more general than what I need. However, I could not find the article, so I give my own proof below. The proof is an adaptation of the arguments by van der Vaart, with some inspiration from an article by Bondar et Milnes (1981).

We say that a group acts on a decision problem $(\mathcal{E}, \mathbb{D}, l)$ if:

- G acts on $\mathcal{A}^*$ and the restriction of the action to $\mathcal{E}$ is a bijection, so that Θ is G-homogeneous. We write $g\phi$ for the action of g on $\phi \in \mathcal{A}^*$, and $g\theta$, for that on $\theta \in \Theta$, hence $g\phi^\theta = \phi^{g\theta}$.
- G acts on $C_b(\mathbb{D})$. We write gf for the action of g on $f \in C_b(\mathbb{D})$.
- Both former actions are (weakly) continuous on bounded balls.
- The (nonnegative lower semicontinuous) loss function is G-superequivariant, that is, for all $\theta \in \Theta$, for all $g \in G$, for all nonnegative $f \in C_b(\mathbb{D})$ such that $f \le l_\theta$, we have $gf \le l_{g\theta}$.

In particular, if $\mathbb{D}$ is G-homogeneous, there is a natural action $gf(x) = f(g^{-1}x)$ for $f \in C_b(\mathbb{D})$ and $x \in \mathbb{D}$, and any equivariant loss functions given by $l_\theta(x) = l_{g\theta}(g^{-1}x)$ is notably superequivariant.

We may then define a contravariant action of G on generalized measurements by:

$$S_g M\phi(f) = M(g\phi)(gf).$$

That is, $S_{g_1 g_2} = S_{g_2} S_{g_1}$.

We shall call a generalized measurement equivariant if $S_g M = M$.

Since G acts continuously on states and bounded continuous functions, $S_g M$ is pointwise continuous, both in M and G.

A group G is said to have *the fixed-point property* if every representation $g \mapsto T_g$ of G, with $T_g(x)$ separately continuous in g and x, in a group of affine transformations of a compact convex subset of a locally convex topological vector space has a fixed point. Rickert (1967) has proved that this is equivalent to G satisfying Stein's condition, that is there is a finite chain of closed subgroups

$$G = G_m \supset G_{m-1} \supset \cdots \supset G_0 = e$$

such that each subgroup is a normal subgroup of the previous one and each G_i/G_{i-1} is either compact or commutative. Commutative groups obviously satisfy Stein's condition.

Theorem 9.7.9 (Quantum Hunt-Stein). *Suppose a group G with the fixed-point property acts on a decision problem $(\mathcal{E}, \mathbb{D}, l)$. Consider a G-invariant function $R : \Theta \to \mathbb{R}^+$. Suppose there is a generalized measurement P with lower risk function $R_\theta(P) \leq R(\theta)$ for all θ.*

Then there is an equivariant measurement M such that $R_\theta(M) \leq R(\theta)$ for all θ.

Moreover, if l is subcompact, if G is acting on $C_b(\mathbb{D})$ through an action on $\mathbb{D}$, if $\mathbb{D}$ is locally compact and there is a true equivariant measurement, then M may be chosen as a true measurement.

Corollary 9.7.10. *In the setting of the previous theorem, the minimax risk is attained by an equivariant generalized measurement, and by an equivariant true measurement with the same conditions as above.*

Proof of the corollary By the same arguments as in the asymptotic minimax theorem, the minimax risk is achieved by a generalized measurement. We then merely take $R(\theta) = R_{minimax}$. A constant is obviously G-invariant.

Let us now prove the quantum Hunt-Stein theorem.

Proof. Let us consider K the subset of generalized measurements M satisfying $R_\theta(M) \leq R(\theta)$ for all θ.

This is a closed set for pointwise convergence, hence compact. It is non-void, since $P \in K$. Moreover, for any M, we have $R_\theta(M) \leq R_\theta(S_g M)$ by G-superequivariance of l. Hence K is stable under S_g.

Since G has the fixed-point property, it admits a fixed point M in K. This M is the equivariant generalized measurement we were looking for.

If furthermore G was acting on $C_b(\mathbb{D})$ through an action on $\mathbb{D}$ and $(\mathbb{D}, \mathcal{B})$ is locally compact, it admits a one-point (Alexandrov) compactification $(\overline{\mathbb{D}}, \overline{B})$. We may extend continuously the group action on $\overline{\mathbb{D}}$ by $g(\infty) = \infty$. We consider $M\overline{\mathbb{D}}$ induced by the equivariant M above. It is a true measurement on $\overline{\mathbb{D}}$, characterized by $\{M_{\overline{\mathbb{D}}}(B)\}_{B \in \overline{\mathcal{B}}}$.

If there is also an equivariant true measurement N on $C_b(\mathbb{D})$, characterized by $\{N(B)\}_{B \in \mathcal{B}}$, we may define the equivariant true measurement M_t through:

$$M_t(B) = M_{\overline{\mathbb{D}}}(B) + M_{\overline{\mathbb{D}}}(\infty) N(B) \qquad \text{for all } B \in \mathcal{B}.$$

This is well-defined since the bidual of a C^*-algebra is still an algebra – even a Von Neumann algebra.

Moreover, since l is subcompact,

$$\begin{aligned} R_\theta(M_t) &= M_t \phi^\theta(l_\theta) \\ &= M\phi^\theta + M_{\overline{\mathbb{D}}}(\infty)\phi^\theta\left(N\phi^\theta(l_\theta) - l_\theta(\infty)\right) \\ &\leq R_\theta(M) \\ &\leq R(\theta). \end{aligned}$$

□

Missing parts in the proof of Theorem 9.4.4

Let us start with proving that the minimax risk in the limit experiment is a lower bound. We can almost use the asymptotic minimax theorem. Indeed, the risk on a finite subexperiment I is smaller than the risk on the whole parameter space, so equation (9.7.7) appears stronger than (9.40). Notice that since the parameter spaces $\Theta_{n,\epsilon,\epsilon}$ are increasing to infinity, the set I is a subset of the parameters for n big enough, so that bound (9.7.7) has a meaning.

However, we cannot apply immediately the theorem: the risk function is not the same in the different experiments. The theorem would hold if we used the loss function r in all experiments.

Let us prove that the limit of the minimax experiments is the same when using r in all experiments, or $r_n = nl(\rho, \hat{\rho})$ as is the case here.

By abuse of notation, we shall write $\rho \in \mathcal{Q}_n$ when $\rho^{\otimes n} \in \mathcal{Q}_n$.

For any $\delta > 0$ by local quadraticity, there is a two-variable open neighbourhood V_δ of (ρ_0, ρ_0) in which $(1-\delta)r \leq r_n \leq (1-\delta)r$. We choose $V_\delta = V_\delta^1 \otimes V_\delta^2$ in product form. Let us now consider a closed neighbourhood F of ρ_0 included in V_δ^1. Then $F \otimes \mathcal{T}_1^+(\mathbb{C}^d) \backslash V_\delta^2$ is compact. Since l is lower semicontinuous, it attains a minimum m on this compact. Since l is estimation-fostering and there is no diagonal element in K, the minimum m is positive. For n big enough, the states of the experiment $\mathcal{Q}_n$, that satisfy $\|\rho_0 - \rho\| \leq n^{-1/2+\epsilon}$, are included in F. So that for all element $\rho \in \mathcal{Q}_n$, for any estimate $\hat{\rho}$, either $r_n(\rho, \hat{\rho}) \geq m$, or $(\rho, \hat{\rho})$ is in V_δ. If n is small enough, there will be points $\hat{\rho}$ in V_δ^2 such that $r_n(\rho, \hat{\rho}) < m$ for all $\rho \in \mathcal{Q}_n$. So that an optimal measurement will always give an answer in V_δ^2. But in that case $r_n \geq (1-\delta)r$. So that the asymptotic minimax risk with loss function r cannot be worse than $(1-\delta)^{-1}$ the asymptotic minimax risk with loss function r_n. Since this is true for any δ, we have proved bound (9.40).

As a remark, local quadraticity and shrinking set $\mathcal{Q}_n$ imply that we may replace G_{ρ_0} with G_ρ in the cost functions up to negligible variations of the same order as r with respect to r_n, which we shall do when practical.

We now prove that the strategy suggested in Section 9.4.5 has asymptotic risk $R_{minimax}$.

Let us start with using shorter notations. Most depend silently on the true state ρ. We write E for the event that $\|\tilde{\rho} - \rho\| \leq n^{-1/2+\epsilon}$, and E^c for its complementary. On E, if n is big enough with respect to the eigenvalues separation δ, there is a smallest θ such that $\rho^{\otimes n-\tilde{n}} = \rho^{\theta, n-\tilde{n}}$ when $\rho_0 = \tilde{\rho}$. We write θ for it. We use this definition so as to have the same scale as $\tilde{n}$ as defined in Section 9.4.5. We write r_n for the loss function at ρ in the n-sample experiment, that is $r_n(\hat{\rho}) = nl(\rho, \hat{\rho})$. We write r for the loss function at θ in the limit experiment, that is $r(\hat{\theta}) = (\theta - \hat{\theta})^* G_\rho (\theta - \hat{\theta})$. We write $\mathbb{E}_n$ for the expectation with respect to the result of our measurement procedure, we write $\mathbb{E}_\infty$ for the expectation with respect to the measurement we apply on $T_{n-\tilde{n}}(\rho^{\otimes n-\tilde{n}})$, with output $\hat{\theta}$, but applied to ϕ^θ instead. We write $\mathbb{E}_\infty^m$ for the expectation of the same measurement applied on ϕ^θ, but without the modification $\hat{\theta} = 0$ if $\tilde{\theta}$ is too big, that is we always keep $\tilde{\theta}$. That measurement is nothing else than the optimal measurement in the limit experiment.

With those notations, he minimax risk $R_{minimax}$ is $\mathbb{E}_{\infty}^{m}[r]$ and the risk of our measurement at ρ is $\mathbb{E}_n[r_n]$. Let us develop the latter, making appear the five parts we had spoken about in Section 9.4.5.

$$\mathbb{E}_n[r_n] = \mathbb{E}_n[r_n\chi_{E^c}] + \mathbb{E}_n[(r_n - r)\chi_E] + (\mathbb{E}_n - \mathbb{E}_\infty)[r\chi_E] + (\mathbb{E}_\infty - \mathbb{E}_\infty^m)[r\chi_E] + \mathbb{E}_\infty^m[r\chi_E]. \tag{9.81}$$

For bounding the first term, we need to show that E^c is a rare event. We give an example of rough measurement that turns the trick.

Let us define, with $|\psi_i\rangle$ any orthonormal basis of $\mathbb{C}^d$, and $E_{ij} = |\psi_i\rangle\langle\psi_j|$:

$$\begin{aligned}
M_j &= \frac{1}{d}E_{jj} && \forall\ 1 \le k \le d,\\
M_{jk} &= \frac{1}{2d}(E_{jj} + E_{kk} + E_{jk} + E_{kj}) && \forall\ 1 \le j < k \le d,\\
M_{kj} &= \frac{1}{2d}(E_{jj} + E_{kk} + iE_{jk} - iE_{kj}) && \forall\ 1 \le j < k \le d,\\
M_r &= \sum_{j=1}^{d}\left[\frac{d-1}{2d}E_{jj} - \frac{1}{2d}\sum_{k=j+1}^{d}(1+i)E_{jk} - (1-i)E_{kj}\right].
\end{aligned}$$

We denote by $X_j^{\tilde{n}}, X_{jk}^{\tilde{n}}, X_{kj}^{\tilde{n}}, X_r^{\tilde{n}}$ the laws of the number of times we get result j, jk, kj or r when measuring when measuring $\tilde{n}$ copies of ρ. The laws of these random variables are binomials:

$$\begin{aligned}
X_j^{\tilde{n}} &= \mathcal{B}\left(\tilde{n}, \frac{\rho_{jj}}{\tilde{n}}\right),\\
X_{jk}^{\tilde{n}} &= \mathcal{B}\left(\tilde{n}, \frac{2\mathrm{Re}\rho_{jk} + \rho_{kk} + \rho_{jj}}{\tilde{n}}\right),\\
X_{kj}^{\tilde{n}} &= \mathcal{B}\left(\tilde{n}, \frac{2\mathrm{Re}\rho_{kj} + \rho_{kk} + \rho_{jj}}{\tilde{n}}\right).
\end{aligned}$$

We then use the following estimates for the matrix entries of $\tilde{\rho}$:

$$\begin{aligned}
\tilde{\rho}_{jj} &= \frac{d}{\tilde{n}}X_j^{\tilde{n}},\\
\mathrm{Re}\tilde{\rho}_{jk} &= \frac{d}{\tilde{n}}X_{jk}^{\tilde{n}} - \frac{d}{2\tilde{n}}(X_j^{\tilde{n}} + X_k^{\tilde{n}}),\\
\mathrm{Im}\tilde{\rho}_{jk} &= -\mathrm{Im}\tilde{\rho}_{kj} && j < k\\
&= \frac{d}{\tilde{n}}X_{kj}^{\tilde{n}} - \frac{d}{2\tilde{n}}(X_j^{\tilde{n}} + X_k^{\tilde{n}}),
\end{aligned}$$

Through repeated use of Hoeffding's inequality (9.117) on each of the matrix entries, for any $\epsilon \geq 0$, we get:

$$\mathbb{P}\left[\|\tilde{\rho} - \rho\|^2 \geq 4d^4 \tilde{n}^{2\epsilon-1}\right] \leq 2d^2 \exp(-2\tilde{n}^{2\epsilon}). \tag{9.82}$$

As a remark, if $\tilde{\rho}$ is not a genuine density matrix, we may project $\tilde{\rho}$ on the set of density matrices, to get a new $\tilde{\rho}$. Since the true ρ naturally belongs to the set of density matrices, we merely have to double $\|\tilde{\rho} - \rho\|$ for bound (9.82) to remain true.

So that

$$\begin{aligned}\mathbb{E}_n\left[r_n \chi_{E^c}\right] &\leq \sup_{\rho,\hat{\rho} \in \mathcal{T}_1^+(\mathbb{C}^d)} r_n(\hat{\rho}) \mathbb{P}\left[\|\rho - \tilde{\rho}\| n^{-1/2+\epsilon}\right] \\ &= O(n \exp(-n^{2\epsilon})).\end{aligned}$$

The second term of the risk (9.81) is bounded thanks to the local quadraticity of the loss function l, after noticing that $\|\rho - \tilde{\rho}\| \leq n^{1/2+\epsilon}$ on E, and that anyhow $\|\hat{\rho} - \tilde{\rho}\| \leq 3\Gamma n^{1/2+\epsilon}$, where Γ is the ratio of the extreme eigenvalues of G_ρ:

$$\begin{aligned}\mathbb{E}_n\left[(r_n - r)\chi_E\right] &\leq \sup_{\substack{\|\rho-\tilde{\rho}\| \leq n^{1/2+\epsilon} \\ \|\hat{\rho}-\tilde{\rho}\| \leq 3\Gamma n^{1/2+\epsilon}}} (r - r_n)(\hat{\rho}) \\ &= O(n(n^{-1/2} + \epsilon)^3) = O(n^{-1/2+3\epsilon}).\end{aligned}$$

The third term in (9.81) is bounded by noticing that we integrate the same function with respect to two probability laws that are very close in L^1 norm. In fact, that's there that we use Theorem 9.4.3:

$$\begin{aligned}(\mathbb{E}_n - \mathbb{E}_\infty)[r\chi_E] &\leq \sup_{\substack{\|\rho-\tilde{\rho}\| \leq n^{1/2+\epsilon} \\ \|\hat{\rho}-\tilde{\rho}\| \leq 3\Gamma n^{1/2+\epsilon}}} T_{n-\tilde{n}}(\rho^{\otimes n-\tilde{n}}) r(\hat{\rho}) \\ &= O(n^{-\kappa} n^{2\epsilon}).\end{aligned}$$

Since κ depends negatively on ϵ, this exponent is negative on for ϵ small enough.

The fourth term is negative. Indeed, the difference between $\mathbb{E}_\infty$ and $\mathbb{E}_\infty^m$ is that the mass outside a ball $B(0, 3\Gamma n^\epsilon)$ in $\mathbb{E}_\infty^m$ is displaced to 0 in $\mathbb{E}_\infty$. Now, since we are on E, we know that $\theta \in B(0, n^\epsilon)$. Hence $r(0)$ is smaller than $r(\tilde{\theta})$ for all $\tilde{\theta}$ outside $B(0, 3\Gamma n^\epsilon)$. So that:

$$(\mathbb{E}_\infty - \mathbb{E}_\infty^m)[r\chi_E] \leq 0$$

We end the proof with:

$$\begin{aligned}\mathbb{E}_\infty^m\left[r\chi_E\right] &\leq \mathbb{E}_\infty^m\left[r\right] \\ &= R_{minimax}\end{aligned}$$

9.7.2 Combinatorial and representation theoretical tools

Here we continue the analysis of the $SU(d)$ irreducible representations $(\pi_\lambda, \mathcal{H}_\lambda)$ started in section 9.5.1. The purpose of this section is to provide good estimates of quantities of the type $\langle \mathbf{m}, \lambda \mid \pi_\lambda(U) \mid \mathbf{l}, \lambda \rangle$ which will be needed in the proofs of Lemmas 9.7.13 and 9.6.4.

We shall use the following form of a general $SU(d)$ element and the shorthand notations

$$\begin{aligned} U(\vec{\zeta}, \vec{\xi}) &:= \exp\left[i \left(\sum_{i=1}^{d-1} \xi_i H_i + \sum_{1 \le j < k \le d} \frac{\mathrm{Re}(\zeta_{j,k}) T_{j,k} + \mathrm{Im}(\zeta_{j,k}) T_{k,j}}{\sqrt{\mu_j - \mu_k}} \right) \right], \\ U(\vec{\zeta}, \vec{\xi}, n) &:= U(\vec{\zeta}/\sqrt{n}, \vec{\xi}/\sqrt{n}), \quad U(\vec{\zeta}) := U(\vec{\zeta}, \vec{0}), \quad U(\vec{\zeta}, n) := U(\vec{\zeta}/\sqrt{n}) \end{aligned} \tag{9.83}$$

where H_i and $T_{i,j}$ are the generators of $SU(d)$ defined by

$$\begin{aligned} H_j &= E_{j,j} - E_{j+1,j+1} \qquad \text{for } j \le d-1; \\ T_{j,k} &= iE_{j,k} - iE_{k,j} \qquad \text{for } 1 \le j < k \le d; \\ T_{k,j} &= E_{j,k} + E_{k,j} \qquad \text{for1} \le j < k \le d. \end{aligned} \tag{9.84}$$

with $E_{i,j}$ the matrix with entry (i,j) equal to 1, and all others equal to 0.

We first introduce some new notations and remind the reader about the already existing ones.

1) We write $l(c)$ for the length of the column c in the Young diagram λ. There are then $\lambda_i - \lambda_{i+1}$ columns such that $l(c) = i$. An alternative definition is $l(c) = \inf\{i : \lambda_i \ge c\}$.

2) Recall that we denote by $f_{\mathbf{a}}$ the basis vectors $f_{a(1)} \otimes \cdots \otimes f_{a(n)}$, and to each vector we associate a Young tableau $t_{\mathbf{a}}$ where the indices $a(i)$ fill the boxes of a diagram λ in a particular way. We denote by $t_{\mathbf{a}}^c$ the column c of $t_{\mathbf{a}}$, i.e. the function $t_{\mathbf{a}}^c : \{1, \ldots, l(c)\} \to \{1, \ldots, d\}$ that associates to the row number r the value of the entry of that Young tableau in column c, row r. For example, if $t_{\mathbf{a}} = \begin{array}{|c|c|c|}\hline 2 & 2 & 1 \\ \hline 2 & 1 \\ \cline{1-2}\end{array}$ we get the values:

$$t_{\mathbf{a}}^1(1) = 2, \qquad t_{\mathbf{a}}^1(2) = 2, \qquad t_{\mathbf{a}}^2(1) = 2, \qquad t_{\mathbf{a}}^2(2) = 1, \qquad t_{\mathbf{a}}^3(1) = 1.$$

We shall often be interested in the image $t_{\mathbf{a}}^c(\{1, \ldots, l(c)\})$ as unordered set, or compare $t_{\mathbf{a}}^c$ to Id^c, the identity function on the integers $\{1, \ldots, l(c)\}$.

3) Recall also that $\mathcal{H}_\lambda$ is spanned by the vectors $y_\lambda f_{\mathbf{a}}$ for which $t_{\mathbf{a}}$ is a semistandard Young tableau, and $y_\lambda = q_\lambda p_\lambda$ is the Young symmetriser (cf. Theorem 9.5.2). If

$t_{\mathbf{a}}$ is semistandard then we can use the alternative notation $f_{\mathbf{m}}$ for $f_{\mathbf{a}}$ since $\mathbf{a}$ is in one-to-one correspondence with $\mathbf{m} = \{m_{i,j} : 1 \leq i < j \leq d\}$, where $m_{i,j}$ is the number of j's in the row i of $t_{\mathbf{a}}$. The normalised vectors are

$$|\mathbf{m}, \lambda\rangle := y_\lambda f_{\mathbf{m}} / \|y_\lambda f_{\mathbf{m}}\|.$$

4) Let $\mathcal{O}_\lambda(\mathbf{m})$ be the orbit of $f_{\mathbf{m}}$ under the subgroup $\mathcal{R}_\lambda$ of row permutations. This consists of vectors $f_{\mathbf{b}}$ which have exactly $m_{i,j}$ boxes with j in row i, and the rest are i. In particular, row i has no entries smaller than i. Since the action of permutations is transitive, we have

$$p_\lambda f_{\mathbf{m}} = \sum_{\sigma \in \mathcal{R}_\lambda} f_{\mathbf{a}\circ\sigma} = \sum_{f_{\mathbf{b}} \in \mathcal{O}_\lambda(\mathbf{m})} \frac{\#\mathcal{R}_\lambda}{\#\mathcal{O}_\lambda(\mathbf{m})} f_{\mathbf{b}}. \tag{9.85}$$

5) Since we antisymmetrize with q_λ, we are only interested in the $t_{\mathbf{a}}$ (not necessarily semistandard) which do not have two equal entries in the same column. Such tableaux $t_{\mathbf{a}}$ (or vectors $f_{\mathbf{a}}$) shall be called *admissible* and their set is denoted $\mathcal{V}$.

6) For any $f_{\mathbf{a}} \in \mathcal{O}_\lambda(\mathbf{m})$ we define

$$\Gamma(f_{\mathbf{a}}) := |\mathbf{m}| - \#\{1 \leq c \leq \lambda_1 : t^c_{\mathbf{a}} \neq \mathrm{Id}^c\},$$

and denote by $\mathcal{V}^\Gamma(\mathbf{m})$ the set of vectors $f_{\mathbf{a}} \in \mathcal{O}_\lambda(\mathbf{m}) \bigcap \mathcal{V}$ with $\Gamma(f_a) = \Gamma$. Then we have

$$\mathcal{O}_\lambda(\mathbf{m}) \bigcap \mathcal{V} = \bigcup_{\Gamma \in \mathbb{N}} \mathcal{V}^\Gamma(\mathbf{m}).$$

Note that $\Gamma(f_{\mathbf{a}}) \geq 0$ and is zero if and only if each column $t^c_{\mathbf{a}}$ is either Id^c or of the form $t^c_{\mathbf{a}}(r) = j\delta_{r=i} + r\delta_{r\neq i}$ for some $i \leq l(c) < j$. A $t^c_{\mathbf{a}}$ of this form will be called an (i,j)-*substitution.*

The following 'algorithm' shows how to build all the possible $f_{\mathbf{a}} \in \mathcal{V}^\Gamma(\mathbf{m})$, thus enabling us to estimate the size of $\mathcal{V}^\Gamma(\mathbf{m})$.

Algorithm

Let $(\mathbf{m}, \lambda)$ be fixed but otherwise arbitrary. In order to generate a particular admissible $f_{\mathbf{a}} \in \mathcal{O}_\lambda(\mathbf{m})$ we need to select the $m_{i,j}$ boxes on row i which are filled with j, for all $i < j$. The rest of the boxes are filled automatically with i's. The constraint is that no column should have two boxes filled with the same number.

Generating a diagram can be described intuitively as follows. We start with the 'vacuum' vector (tableau) $f_{\mathbf{0}} := f_{\mathbf{m}=\mathbf{0}}$ (row i is filled exclusively with i's), and with

a set of $|\mathbf{m}|$ bricks containing $m_{i,j}$ identical bricks labelled (i,j), for each pair $i<j$. To change the content of a box from i into j we place an (i,j)-brick in that box. This procedure is repeated until all bricks have been used, each box being modified at most once.

At this stage each column c may contain several bricks placed in the appropriate boxes, so that its configuration is uniquely defined by the set of bricks κ which shall be called a *column-modifier*. For example if $\kappa=\{(i,j),(f,l)\}$ then the column has entries

$$t^c_{\mathbf{a}}(k)=\begin{cases} j & \text{if } k=i;\\ l & \text{if } k=f;\\ k & \text{otherwise.}\end{cases}$$

Note that a column-modifier is not an arbitrary collection of bricks but one that can be used to produce a column with different entries. In the previous example, if $i<f$ this means either ($j\neq f$ and $j,l>l(c)$) or ($j=f$ and $l>l(c)$). The elementary one-brick column-modifier denoted $\kappa(i,j)$ can only be used in a column with $i\leq l(c)<j$, otherwise the entry j would appear twice.

Now, since the length of a column is at most d and all entries must be different, there are less than $d!$ different types of column-modifiers. Another important remark is that a column-modifier always increases the value of the modified cells, so that in this case $t^c_{\mathbf{a}}(\{1,\ldots,l(c)\})\neq\{1,\ldots,l(c)\}$.

Alternatively to the above scenario where the bricks are inserted sequentially, we can first cluster them into $|\mathbf{m}|-\Gamma$ column-modifiers, and then apply each column-modifier to a particular column. A given collection of column-modifiers is uniquely determined by $\{m_\kappa:\kappa\}$ where m_κ is the multiplicity of κ. This procedure is detailed in the following 3 stages:

I. Choose Γ bricks among our $|\mathbf{m}|$. As we have $d(d-1)/2$ different types of bricks (recall that $i>j$), and we do not distinguish between identical bricks, there are at most $[d(d-1)/2]^\Gamma$ possibilities. For $\Gamma=0$, we have only one choice.

II. Consider the remaining bricks as a set of elementary column-modifiers. Starting from these, we sequentially add each of the Γ bricks selected in the first stage, to one of these elementary column-modifiers to form non-elementary ones. At each step we have at most $d!$ different *types* of column modifiers to which we can attach the new brick. Note that we do not distinguish between column modifiers of the same type, but rather consider them as an unordered set. Hence, we have less that $(d!)^\Gamma$ possibilities. If $\Gamma=0$ there is only one possibility.

   Note that at the end of stage II at least $\max\{0,|\mathbf{m}|-2\Gamma\}$ of the column-modifiers are elementary, and that $m_{\kappa(i,j)}\leq m_{i,j}$.

III. Apply the column-modifiers to the columns of $f_{\mathbf{0}}$, so that no two modifiers are applied to the same column and the resulting $f_{\mathbf{a}} \in \mathcal{O}_\lambda(\mathbf{m})$ is admissible. By construction $\Gamma(f_{\mathbf{a}}) = \Gamma$ and all admissible tableaux can be generated in this way.

For counting the number of possibilities for the third stage we apply the column modifiers sequentially, but since some of them may be identical we need to divide by the combinatorial factor $\prod_\kappa m_\kappa!$, where m_κ is the number of column modifiers of type κ.

We distinguish between elementary column modifiers of type $\kappa(i,j)$ and composite ones. There are less than n possibilities of inserting a composite column-modifier κ. An elementary one of type $\kappa(i,j)$ can only be inserted in a column with at least i rows, and since the resulting vector has to be admissible, the column cannot contain another j, so its length is smaller than j. There are $\lambda_i - \lambda_j$ such columns. Hence the number of possibilities at stage three of the algorithm is upper bounded by

$$\prod_{\kappa \neq \kappa(i,j)} \frac{n^{m_\kappa}}{m_\kappa!} \cdot \prod_{i<j} \frac{(\lambda_i - \lambda_j)^{m_{\kappa(i,j)}}}{m_{\kappa(i,j)}!}. \tag{9.86}$$

When $\Gamma = 0$, for each elementary column modifier $\kappa(i,j)$ the number of available columns is at least $(\lambda_i - \lambda_j - |\mathbf{m}|)_+ := \max\{0, \lambda_i - \lambda_j - |\mathbf{m}|\}$. Thus we have the following lower bound

$$\prod_{i<j} \frac{(\lambda_i - \lambda_j - |\mathbf{m}|)_+^{m_{i,j}}}{m_{i,j}!}. \tag{9.87}$$

Notice that the upper bound (9.86) depends on the set of multiplicities $\{m_\kappa\}$.

We now return to our list of notations and definitions.

7) To each column of $t_{\mathbf{a}}$ we associated a column modifier which completely determines its content. If $m_\kappa^{\mathbf{a}}$ is the number of columns with column-modifer κ, we collect all multiplicities in $E := \{m_\kappa^{\mathbf{a}} : \kappa\}$. In particular Γ is a function of E

$$\Gamma(f_{\mathbf{a}}) = |\mathbf{m}| - \sum_\kappa m_\kappa^{\mathbf{a}}.$$

Vectors for which $\Gamma(f_{\mathbf{a}}) = 0$ have the same multiplicity set E^0 where $m_{\kappa(i,j)} = m_{i,j}$ for all $i < j$ and the other $m_\kappa = 0$. Similarly to $\mathcal{V}^\Gamma(\mathbf{m})$, we denote by $\mathcal{V}^E(\mathbf{m})$ the set of tableaux in $\mathcal{O}_\lambda(\mathbf{m}) \bigcap \mathcal{V}$ with $E(f_{\mathbf{a}}) = E$, in particular

$$\mathcal{V}^\Gamma(\mathbf{m}) = \bigcup_{E:\Gamma(E)=\Gamma} \mathcal{V}^E(\mathbf{m})$$

8) To each column c of $t_{\mathbf{a}}$ we associate two disjoint sets: the added entries $\{t^c_{\mathbf{a}}(1),\ldots,t^c_{\mathbf{a}}(l(c))\} \setminus \{1,\ldots,l(c)\}$ and the deleted entries $\{1,\ldots,l(c)\} \setminus \{t^c_{\mathbf{a}}(1),\ldots,t^c_{\mathbf{a}}(l(c))\}$. This data is placed into a single set by attaching a $\pm$ sign to each entry, indicating if it is added or deleted. It is easy to verify that if $t_{\mathbf{a}}$ is admissible, the set of added and deleted entries is uniquely determined by the column-modifer κ associated to c, and hence shall be denoted by $S(\kappa)$. For example $S(\kappa(i,j)) = \{(i,-),(j,+)\}$ and for $\kappa = \{(i,j),(j,k)\}$ we have $S(\kappa) = \{(i,-),(k,+)\}$. We define the multiplicities $m^{\mathbf{a}}_S = \sum_{\kappa:S(\kappa)=S} m^{\mathbf{a}}_\kappa$ and $F(f_{\mathbf{a}}) := \{m^{\mathbf{a}}_S : S\}$. To summarise, we have defined the maps

$$f_{\mathbf{a}} \longmapsto E(f_{\mathbf{a}}) \longmapsto F(f_{\mathbf{a}}).$$

We now state our estimates. The first point of the following lemma is an exact formula serving as the main tool to prove some of the bounds below.

Lemma 9.7.11.

1. *For any unitary operator $U \in M(\mathbb{C}^d)$, for any basis vectors $f_{\mathbf{a}}$ and $f_{\mathbf{b}}$, we have*

$$\langle f_{\mathbf{a}}|q_\lambda U^{\otimes n} f_{\mathbf{b}}\rangle = \prod_{1\le c\le\lambda_1} \det(U^{t^c_{\mathbf{a}},t^c_{\mathbf{b}}}), \tag{9.88}$$

where $U^{t^c_{\mathbf{a}},t^c_{\mathbf{b}}}$ is the $l(c)\times l(c)$ minor of U given by $[U^{t^c_{\mathbf{a}},t^c_{\mathbf{b}}}]_{i,j} = U_{t^c_{\mathbf{a}}(i),t^c_{\mathbf{b}}(j)}$.

Under the assumptions

$$\begin{aligned} |\mathbf{m}| &\le n^\eta, \\ \lambda &\in \Lambda_{n,\alpha}, \\ \inf_i |\mu_i - \mu_{i+1}| &\ge \delta, \\ \mu_d &\ge \delta, \\ \|\vec{\zeta}\|_1 &\le Cn^\beta, \qquad \beta \le 1/2, \\ \|\vec{\xi}\|_1 &\le n^{-1/2+2\beta}/\delta, \\ n &> \left(\frac{2}{\delta}\right)^{1/(1-\alpha)}. \end{aligned} \tag{9.89}$$

we have the following estimates with remainder terms uniform in the eigenvalues $\mu_\bullet$:

2. *The number of admissible $f_{\mathbf{a}} \in \mathcal{O}_\lambda(\mathbf{m})$ with $\Gamma(f_{\mathbf{a}}) = 0$ is*

$$\#\mathcal{V}^0(\mathbf{m}) = \prod_{j>i} \frac{(\lambda_i-\lambda_j)^{m_{i,j}}}{m_{i,j}!}(1+O(n^{-1+2\eta}/\delta)). \tag{9.90}$$

3. *Let* $E := \{m_\kappa : \kappa\}$ *with* $\Gamma(E) = \Gamma$. *The number of admissible* $f_{\mathbf{a}} \in \mathcal{O}_\lambda(\mathbf{m})$ *with* $E(f_{\mathbf{a}}) = E$ *is bounded by:*

$$\#\mathcal{V}^E(\mathbf{m}) \leq n^{-\Gamma+\sum_{i<j}(m_{i,j}-m_{\kappa(i,j)})} \prod_{j>i} \frac{(\lambda_i - \lambda_j)^{m_{\kappa(i,j)}}}{m_{\kappa(i,j)}!}. \tag{9.91}$$

4. *The number of admissible* $f_{\mathbf{a}} \in \mathcal{O}_\lambda(\mathbf{m})$ *with* $\Gamma(f_{\mathbf{a}}) = \Gamma$ *is bounded by:*

$$\#\mathcal{V}^\Gamma(\mathbf{m}) \leq C^\Gamma n^{-\Gamma} \delta^{-2\Gamma} |\mathbf{m}|^{2\Gamma} \prod_{j>i} \frac{(\lambda_i - \lambda_j)^{m_{i,j}}}{m_{i,j}!}, \tag{9.92}$$

for a constant $C = C(d)$.

5. *Let* $f_{\mathbf{a}} \in \mathcal{V}^{\Gamma^a}(\mathbf{l})$, *and consider* $\mathcal{V}^{\Gamma^b}(\mathbf{m}) \subset \mathcal{O}_\lambda(\mathbf{m})$ *for some fixed* Γ^b. *Then:*

$$\left| \left\langle f_{\mathbf{a}} \middle| q_\lambda \sum_{f_{\mathbf{b}} \in \mathcal{V}^{\Gamma^b}(\mathbf{m})} f_{\mathbf{b}} \right\rangle \right| \leq \begin{cases} 0 & \text{if } \Gamma^b \neq |\mathbf{m}| - |\mathbf{l}| + \Gamma^a \\ (C|\mathbf{m}|)^{\Gamma^b} & \text{otherwise} \end{cases}, \tag{9.93}$$

with $C = C(d)$.

6. *If* $f_{\mathbf{a}} \in \mathcal{V}^0(\mathbf{m})$, *then*

$$\left\langle f_{\mathbf{a}} \middle| q_\lambda \sum_{f_{\mathbf{b}} \in \mathcal{O}_\lambda(\mathbf{m})} f_{\mathbf{b}} \right\rangle = 1. \tag{9.94}$$

7. *If* $f_{\mathbf{a}} \in \mathcal{V}^0(\mathbf{m})$ *so that its set of elementary column-modifiers is* $E^0 = \{m_{\kappa(i,j)} = m_{i,j}\}$, *then*

$$\langle f_{\mathbf{a}} | q_\lambda U(\vec{\zeta}, \vec{\xi}, n)^{\otimes n} f_{\mathbf{0}} \rangle = \exp\left(i\phi - \frac{\|\vec{\zeta}\|_2^2}{2} \right) \prod_{i<j} \left(\frac{\zeta_{i,j}}{\sqrt{n}\sqrt{\mu_i - \mu_j}} \right)^{m_{i,j}} r(n), \tag{9.95}$$

with the phase and error factor

$$\phi = \sqrt{n} \sum_{i=1}^{d-1} (\mu_i - \mu_{i+1}) \xi_i,$$
$$r(n) = 1 + O\left(n^{-1+2\beta+\eta} \delta^{-1}, n^{-1/2+2\beta} \delta^{-1}, n^{-1+2\beta+\alpha} \delta^{-1} \right).$$

8. *If* $f_{\mathbf{a}} \in \mathcal{V}^E(\mathbf{m})$, *so that its set of column-modifiers is* $E = \{m_\kappa : \kappa\}$ *and* $\Gamma(E) = \Gamma$, *then*

$$\left| \langle f_{\mathbf{a}} | q_\lambda U(\vec{\zeta}, \vec{\xi}, n)^{\otimes n} f_{\mathbf{0}} \rangle \right|$$
$$\leq \exp\left(-\frac{\|\vec{\zeta}\|_2^2}{2} \right) \left(\frac{C\|\vec{\zeta}\|}{\sqrt{n\delta}} \right)^{-\Gamma+\sum_{i<j}(m_{i,j}-m_{\kappa(i,j)})} \prod_{i<j} \left(\frac{\zeta_{i,j}}{\sqrt{n}\sqrt{\mu_i - \mu_j}} \right)^{m_{\kappa(i,j)}} \tag{9.96}$$

with $C = C(d)$ *a constant and* $r(n)$ *as in point 7 above.*

9. *Under the further hypotheses that* $\|\vec{z}\| \leq n^{\beta}$, $m_{i,j} \leq 2|\zeta_{i,j} + z_{i,j}|n^{\beta+\epsilon}$ *for some* $\epsilon > 0$, *we have:*

$$\left\langle \sum_{f_{\mathbf{a}} \in \mathcal{O}_\lambda(\mathbf{m})} f_{\mathbf{a}} \middle| q_\lambda U(\vec{\zeta} + \vec{z}, \vec{\xi}, n) f_{\mathbf{0}} \right\rangle$$
$$= \exp\left(i\phi - \frac{\|\vec{\zeta} + \vec{z}\|_2^2}{2}\right) \prod_{i<j} \frac{\left((\zeta_{i,j} + z_{i,j})(\sqrt{n}\sqrt{\mu_i - \mu_j})\right)^{m_{i,j}}}{m_{i,j}!} r(n), \tag{9.97}$$

with

$$r(n) = 1 + O\left(n^{-1+2\beta+\eta}\delta^{-1}, n^{-1+2\beta+\alpha}\delta^{-1}, n^{-1+2\eta}\delta^{-1}, n^{-1+\alpha+\eta}\delta^{-1}, \delta^{-3/2}n^{-1/2+3\beta+2\epsilon}\right).$$

10. *Under the further hypotheses that* $|\mathbf{l}| \leq |\mathbf{m}|$ *and* $n^{1-3\eta} > 2C/\delta^2$, *where* $C = C(d)$,

$$\left| \left\langle \sum_{f_{\mathbf{a}} \in \mathcal{O}_\lambda(\mathbf{l})} f_{\mathbf{a}} \middle| q_\lambda \sum_{f_{\mathbf{b}} \in \mathcal{O}_\lambda(\mathbf{m})} f_{\mathbf{b}} \right\rangle \right| \leq (C|\mathbf{m}|)^{|\mathbf{m}|-|\mathbf{l}|} \prod_{i<j} \frac{(\lambda_i - \lambda_j)^{l_{i,j}}}{l_{i,j}!} \left(\frac{C|\mathbf{l}|^2|\mathbf{m}|}{n\delta^2}\right)^{\Gamma^a_{\min}(\mathbf{l},\mathbf{m})} \tag{9.98}$$

with

$$\Gamma^a_{\min}(\mathbf{l}, \mathbf{m}) \geq \frac{\left(|\mathbf{l} - \mathbf{m}| + 3|\mathbf{l}| - 3|\mathbf{m}|\right)_+}{6}. \tag{9.99}$$

11. *We have*

$$\left\langle \sum_{f_{\mathbf{a}} \in \mathcal{O}_\lambda(\mathbf{m})} f_{\mathbf{a}} \middle| q_\lambda \sum_{f_{\mathbf{b}} \in \mathcal{O}_\lambda(\mathbf{m})} f_{\mathbf{b}} \right\rangle = \prod_{i<j} \frac{(\lambda_i - \lambda_j)^{m_{i,j}}}{m_{i,j}!} \left(1 + O(n^{3\eta-1}/\delta)\right). \tag{9.100}$$

Proof.

Proof of (9.88). We first express $\langle f_{\mathbf{a}}|U^{\otimes n} f_{\mathbf{b}}\rangle$ as a product of matrix entries of U:

$$\langle f_{\mathbf{a}}|U^{\otimes n} f_{\mathbf{b}}\rangle = \prod_{1 \leq c \leq \lambda_1} \prod_{1 \leq r \leq l(c)} \langle f_{t^c_{\mathbf{a}}(r)}|U f_{t^c_{\mathbf{b}}(r)}\rangle$$
$$= \prod_{1 \leq c \leq \lambda_1} \prod_{1 \leq r \leq l(c)} U_{t^c_{\mathbf{a}}(r), t^c_{\mathbf{b}}(r)}.$$

Since the subgroup of column permutations $\mathcal{C}_\lambda$ is the product of the permutation groups of each column, each $\sigma \in \mathcal{C}_\lambda$ is $\sigma = s_1 \dots s_{\lambda_1}$ with s_c a permutation of column c which transforms $t^c_{\mathbf{b}}(r)$ into $t^c_{\mathbf{b}}(s_c(r))$. Then

$$\begin{aligned}\langle f_{\mathbf{a}} | q_\lambda U^{\otimes n} f_{\mathbf{b}} \rangle = \langle f_{\mathbf{a}} | U^{\otimes n} q_\lambda f_{\mathbf{b}} \rangle &= \sum_{\sigma \in \mathcal{C}_\lambda} \epsilon(\sigma) \prod_{1 \le c \le \lambda_1} \prod_{1 \le r \le l(c)} U_{t^c_{\mathbf{a}}(r), t^c_{\mathbf{b}}(s_c(r))} \\ &= \prod_{1 \le c \le \lambda_1} \sum_{s_c \in S_c} \epsilon(s_c) \prod_{1 \le r \le l(c)} U_{t^c_{\mathbf{a}}(r), t^c_{\mathbf{b}}(s_c(r))} \\ &= \prod_{1 \le c \le \lambda_1} \det(U^{t^c_{\mathbf{a}}, t^c_{\mathbf{b}}}).\end{aligned}$$

Proof of (9.90). The number of admissible $f_{\mathbf{a}}$ such that $\Gamma(f_{\mathbf{a}}) = 0$ is given by the products of the possibilities at each stage of the algorithm. For the first two stages, there is exactly one possibility when $\Gamma = 0$. Hence $\#\mathcal{V}^0$ is the number of possibilities at the third stage. Here the upper bound (9.86) reads as $\prod_{j>i} (\lambda_i - \lambda_j)^{m_{i,j}} / m_{i,j}!$. On the other hand, we may use (9.87) as a lower bound, recalling that $\lambda_i - \lambda_j \ge \delta n/2$ and $|\mathbf{m}| \le n^\eta$ (cf. (9.89)). This yields the result (9.90).

Proof of (9.91). The number of $f_{\mathbf{a}}$ in $\mathcal{V}^E$ is given by the third stage of the algorithm (the two first stages yield a particular E). We then obtain (9.91) by applying (9.86) and neglecting the $m_\kappa!$ factors, while noticing that $\sum_\kappa m_\kappa = |\mathbf{m}| - \Gamma$.

Proof of (9.92). The set $\mathcal{V}^\Gamma$ is the union of all $\mathcal{V}^E$ with $\Gamma(E) = \Gamma$. Now the first two stages of the algorithm imply that there are at most C^Γ different E with the latter property, with $C = C(d)$.

Now we use (9.91) to upper-bound $\mathcal{V}^E$ as follows. Since $\sum m_{\kappa(i,j)} \ge |\mathbf{m}| - 2\Gamma$, we may write $\prod_\kappa m_{\kappa(i,j)}! \ge \prod_{i<j} m_{i,j}! \sup_{i<j} m_{i,j}^{-2\Gamma}$. Moreover $\lambda_i - \lambda_j \ge \delta n/2$. By putting together we obtain

$$\#\mathcal{V}^E \le n^{-\Gamma} \delta^{-2\Gamma} |\mathbf{m}|^{2\Gamma} \prod_{j>i} \frac{(\lambda_i - \lambda_j)^{m_{i,j}}}{m_{i,j}!}, \qquad \forall E \text{ with } \Gamma(E) = \Gamma.$$

Multiplying by the number of possible E yields the result.

Proof of (9.93). We are applying (9.88) with $U = \mathbf{1}$. Since both $f_{\mathbf{a}}$ and $f_{\mathbf{b}}$ are product of basis vectors, the scalar product $\langle f_{\mathbf{a}} \mid q_\lambda f_{\mathbf{b}} \rangle$ is equal to -1 or 1 if $t^c_{\mathbf{a}}([1, l(c)]) = t^c_{\mathbf{b}}([1, l(c)])$ for all columns, and 0 otherwise. Here we denote by $t^c_{\mathbf{a}}([1, l(c)])$ the set of entries $\{t^c_{\mathbf{a}}(1), \dots, t^c_{\mathbf{a}}(l(c))\}$.

Now, since a modified column cannot satisfy $t^c_{\mathbf{a}}([1,l(c)]) = [1,l(c)]$ (and the same for $\mathbf{b}$), the vectors $f_{\mathbf{a}}$ and $f_{\mathbf{b}}$ are orthogonal unless they have the same number of modified columns. Finally, that number is $|\mathbf{l}| - \Gamma^a$ for $f_{\mathbf{a}}$ and $|\mathbf{m}| - \Gamma^b$ for $f_{\mathbf{b}}$. This yields the first line of (9.93).

We now concentrate on the case when $\Gamma^b = |\mathbf{m}| - |\mathbf{l}| + \Gamma^a$. Since $|\langle f_{\mathbf{a}} \mid q_\lambda f_{\mathbf{b}}\rangle| \leq 1$, we can bound the sum of scalar products by the number of non-zero inner products. The question is how many diagrams $f_{\mathbf{b}}$ have the same content (seen as an unordered set) in each column as $f_{\mathbf{a}}$: $t^c_{\mathbf{a}}([1,l(c)]) = t^c_{\mathbf{b}}([1,l(c)])$, or equivalently $S(\kappa^c_{\mathbf{a}}) = S(\kappa^c_{\mathbf{b}})$.

For building the relevant $f_{\mathbf{b}}$, we can follow the algorithm with the further condition that, at stage three, all the column-modifiers are applied in such a way that the unordered column content is identical to that of $f_{\mathbf{a}}$.

The first two stages of the algorithm are the same so they yield a C^{Γ^b} factor. We now have a collection $\{m_\kappa\}$ of column modifiers which have to be placed so that they match the column content of $f_{\mathbf{a}}$. For each S we identify the column modifiers $\kappa_1, \ldots, \kappa_{r(S)}$ such that $S(\kappa_i) = S$ for all $1 \leq i \leq r(S)$. The total number of such objects is $m_S := \sum_{i \leq r(S)} m_{\kappa_i}$ and the number of ways in which they can be inserted to produce *distinct* diagrams is

$$\begin{pmatrix} m_S \\ m_{\kappa_1} \ldots m_{\kappa_{r(S)}} \end{pmatrix}.$$

Recall that the number of elementary column-modifiers $\sum_{i<j} m_{\kappa(i,j)}$ is at least $|\mathbf{m}| - 2\Gamma^b$. Moreover, each elementary column-modifier $\kappa(i,j)$ corresponds to a different $S(\kappa(i,j)) = \{(i,-),(j,+)\}$. Thus

$$|\mathbf{m}| - 2\Gamma^b \leq \sum_{i<j} m_{\kappa(i,j)} \leq \sum_S \max_{\kappa : S(\kappa)=S} m_\kappa.$$

Since

$$\sum_S m_S = \sum_\kappa m_\kappa = |\mathbf{m}| - \Gamma^b,$$

we obtain

$$\sum_S \left(m_S - \max_{\kappa : S(\kappa)=S} m_\kappa \right) \leq \Gamma^b.$$

This implies

$$\prod_S \begin{pmatrix} m_S \\ m_{\kappa_1} \ldots m_{\kappa_{r(S)}} \end{pmatrix} \leq |\mathbf{m}|^{\Gamma^b}.$$

Multiplying by the C^{Γ^b} of the first stages, we get (9.93).

Proof of (9.94). As shown above the only non-zero contributions come from $f_{\mathbf{b}} \in \mathcal{V}^0 \subset \mathcal{O}_\lambda(\mathbf{m})$.

Since $\Gamma^b = 0$, the constant from the two first stages of the algorithm is 1, $m_S = m_{i,j} = m_{\kappa(i,j)}$ for all S corresponding to an elementary column-modifier, and 0 otherwise. So the combinatorial factor is again one: we do not have any choice in our placement of column-modifiers. In other words, the only $f_{\mathbf{b}}$ such that $\langle f_{\mathbf{a}} \mid q_\lambda f_{\mathbf{b}} \rangle \neq 0$ is $f_{\mathbf{a}}$. Finally, $\langle f_{\mathbf{a}} \mid q_\lambda f_{\mathbf{a}} \rangle = 1$.

Proof of (9.95). From (9.88) we deduce

$$\langle f_{\mathbf{a}} | q_\lambda U(\vec{\zeta}, \vec{\xi}, n)^{\otimes n} f_{\mathbf{0}} \rangle = \prod_{1 \leq c \leq \lambda_1} \det(U^{t^c_{\mathbf{a}}, \mathrm{Id}^c}), \qquad U = U(\vec{\zeta}, \vec{\xi}, n).$$

We shall use the Taylor expansion of the unitary $U(\vec{\zeta}, \vec{\xi}, n)$ to estimate the above determinants.

Entry-wise, for all $1 \leq i \leq d$ on the first line, and all $1 \leq i < j \leq d$ on the second and third lines:

$$\begin{aligned} U_{i,i}(\vec{\zeta}, \vec{\xi}, n) &= 1 + i\frac{\xi_i \delta_{i \neq d} - \xi_{i-1}\delta_{i \neq 1}}{\sqrt{n}} - \frac{1}{2n} \sum_{j \neq i} \frac{|\zeta_{i,j}|^2}{|\mu_i - \mu_j|} \\ &\quad + O(\|\vec{\zeta}\|^3 n^{-3/2} \delta^{-3/2}, \|\vec{\zeta}\| \|\vec{\xi}\| n^{-1} \delta^{-1/2}, \|\vec{\xi}\|^2 n^{-1}); \\ U_{i,j}(\vec{\zeta}, \vec{\xi}, n) &= -\frac{1}{\sqrt{n}} \frac{\zeta^*_{i,j}}{\sqrt{\mu_i - \mu_j}} + O(\|\vec{\zeta}\|^2 n^{-1} \delta^{-1}, \|\vec{\zeta}\| \|\vec{\xi}\| n^{-1} \delta^{-1/2}); \\ U_{j,i}(\vec{\zeta}, \vec{\xi}, n) &= \frac{1}{\sqrt{n}} \frac{\zeta_{i,j}}{\sqrt{\mu_i - \mu_j}} + O(\|\vec{\zeta}\|^2 n^{-1} \delta^{-1}, \|\vec{\zeta}\| \|\vec{\xi}\| n^{-1} \delta^{-1/2}). \end{aligned}$$

If $\vec{\zeta} = O(n^\beta)$, $\|\vec{\xi}\| \leq n^{-1/2+2\beta}/\delta$, and $\beta < 1/2$, the remainder terms are $O(n^{-3/2+3\beta}\delta^{-3/2})$ for the first line and $O(n^{-1+2\beta}\delta^{-1})$ for the last two lines.

Therefore, when our parameters are in this range, we can give precise enough evaluations of the determinants. The idea is to find the dominating terms in the expansion of the determinant

$$\det A = \sum_\sigma \prod_i \epsilon(\sigma) A_{i,\sigma(i)}.$$

Note that we can use the above Taylor expansions inside the determinant since the number of terms in the product is at most d.

Since $f_{\mathbf{a}} \in \mathcal{V}^0$, all $t^c_{\mathbf{a}}$ are either Id^c, or an (i,j)-*substitution*. If $t^c_{\mathbf{a}} = \mathrm{Id}^c$, the summands with more than two non-diagonal terms are of the same order as the remainder term,

so that only the identity and the transpositions count in $\sum_\sigma \prod_i A_{i,\sigma(i)}$. Let $l = l(c)$, then

$$v(l) := \det(U^{\mathrm{Id}^c,\mathrm{Id}^c}(\vec{\zeta},\vec{\xi},n)) = 1 + i\frac{\xi_l}{\sqrt{n}} - \frac{1}{2n}\sum_{\substack{1\le i\le l\\ l+1\le j\le d}} \frac{|\zeta_{i,j}|^2}{\mu_i-\mu_j} + O(n^{-3/2+3\beta}\delta^{-3/2}).$$

Note that for $l = d$, we get the usual determinant of $U(\vec{\zeta},\vec{\xi},n)$ which is 1.

Consider now the case $t^c_{\mathbf{a}} \neq \mathrm{Id}^c$. Since $t^c_{\mathbf{a}}(r) \geq r$ for all r, there exists a whole column of $U^{t^c_{\mathbf{a}},\mathrm{Id}^c}$ whose entries are smaller in modulus than $O(\|\vec{\zeta}\|/\sqrt{n\delta}) = O(n^{-1/2+\beta}\delta^{-1})$. In particular if $t^c_{\mathbf{a}}$ is an (i,j)-*substitution*, then the only summand that is of this order comes from the identity. So that

$$v(i,j) := \det(U^{t^c_{\mathbf{a}},\mathrm{Id}^c}(\vec{\zeta},\vec{\xi},n)) = \frac{\zeta_{i,j}}{\sqrt{n}\sqrt{\mu_i-\mu_j}} + O(n^{-1+2\beta}\delta^{-1}). \tag{9.101}$$

Note that this approximation does not depend on $l(c)$, but only on i and j.

We now put together the estimated determinants in the product (9.88). For each $i < j$ there are $m_{i,j}$ columns of the type (i,j)-substitution. Out of the $\lambda_l - \lambda_{l+1}$ columns of length $l = l(c)$ there are $\lambda_l - \lambda_{l+1} - R_l$ of the type Id^c, with $0 \leq R_l \leq |\mathbf{m}|$.

Hence:

$$\langle f_{\mathbf{a}}|q_\lambda U(\vec{\zeta},\vec{\xi},n)^{\otimes n} f_{\mathbf{0}}\rangle = \prod_{l=1}^{d}(v(l))^{\lambda_l-\lambda_{l+1}} \prod_{1\le i<j\le d}(v(i,j))^{m_{i,j}} \prod_{l=1}^{d}(v(l))^{-R_l}. \tag{9.102}$$

Now $v(l) = 1 + O(n^{-1+2\beta}\delta^{-1})$ and $R_l \leq |\mathbf{m}| \leq n^\eta$, so the last product is $1 + O(n^{-1+2\beta+\eta}\delta^{-1})$. Similarly, since $\lambda \in \Lambda_{n,\alpha}$ we have $\lambda_l - \lambda_{l+1} = n(\mu_l - \mu_{l+1}) + O(n^\alpha)$, and we can use Lemma 9.7.12 given at the end of this section to estimate the first product as follows

$$\begin{aligned}\prod_{l=1}^{d} v(l)^{\lambda_l-\lambda_{l+1}} &= \prod_{l=1}^{d}\exp\left(i\phi_l - \frac{1}{2}\sum_{\substack{1\le i\le l\\ l+1\le j\le d}}|\zeta_{i,j}|^2\frac{\mu_l-\mu_{l+1}}{\mu_i-\mu_j}\right) r(n)\\ &= \exp\left(i\phi - \frac{\|\vec{\zeta}\|_2^2}{2}\right) r(n),\end{aligned}$$

with

$$\begin{aligned}\tilde{r}(n) &= 1 + O(n^{-1+\alpha+2\beta}\delta^{-1}, n^{-1/2+2\beta}\delta^{-1}),\\ \phi_l &= \delta_{l\neq d}\sqrt{n}(\mu_l-\mu_{l+1})\xi_l,\\ \phi &= \sqrt{n}\sum_{l=1}^{d-1}(\mu_l-\mu_{l+1})\xi_l.\end{aligned}$$

We now turn our attention to the middle product on the right side of (9.102)

$$v(i,j)^{m_{i,j}} = \left(\frac{\zeta_{i,j}}{\sqrt{n}\sqrt{\mu_i - \mu_j}}\right)^{m_{i,j}} \left(1 + O\left(n^{-1+2\beta+\eta}\delta^{-1}\right)\right),$$

where we have used that $|\mathbf{m}| \leq n^{\eta}$.

Inserting into (9.102) yields (9.95).
Note that $\langle f_{\mathbf{a}} | q_\lambda U(\vec{\zeta}, \vec{\xi}, n)^{\otimes n} f_{\mathbf{0}} \rangle = 0$ if there exist $i < j$ such that $\zeta_{i,j} = 0$ and $m_{i,j} \neq 0$.

Proof of (9.96). We may write, much like in (9.102),

$$\langle f_{\mathbf{a}} | q_\lambda U(\vec{\zeta}, \vec{\xi}, n)^{\otimes n} f_{\mathbf{0}} \rangle = \prod_{l=1}^{d} (v(l)))^{\lambda_l - \lambda_{l+1}} \prod_{\kappa} (v(\kappa))^{m_\kappa} \prod_{l=1}^{d} (v(l))^{-R_l}$$

where $0 \leq R_l \leq |\mathbf{m}| - \Gamma$ and $v(\kappa)$ is the determinant of the minor of U corresponding to having applied the column-modifier κ. We can further split the column-modifers into elementary ones $\kappa(i,j)$ and non-elementary ones κ'.

Then $\langle f_{\mathbf{a}} | q_\lambda U(\vec{\zeta}, \vec{\xi}, n)^{\otimes n} f_{\mathbf{0}} \rangle$ can be written as

$$\prod_{l=1}^{d} (v(l)))^{\lambda_l - \lambda_{l+1}} \prod_{i<j} (v(i,j))^{m_{\kappa(i,j)}} \prod_{l=1}^{d} (v(l))^{-R_l} \prod_{\kappa'} (v(\kappa'))^{m_{\kappa'}}.$$

The first three products on the right side can be treated as above. For the fourth product we give a rough upper bound based on the following observation. If the entries in the column have been modified in an admissible way, then $t^c_{\mathbf{a}}(i) = j > l(c)$ for some i, so that $|v(\kappa)| \leq C\|\vec{\zeta}\|/\sqrt{n\delta}$ for any κ, with some constant $C = C(d)$.

Thus by using the previous point

$$\left|\langle f_{\mathbf{a}} | q_\lambda U(\vec{\zeta}, \vec{\xi}, n)^{\otimes n} f_{\mathbf{0}} \rangle\right| \leq$$
$$\exp\left(-\frac{\|\vec{\zeta}\|_2^2}{2}\right) \left(\frac{C\|\vec{\zeta}\|}{\sqrt{n\delta}}\right)^{\sum_{\kappa'} m_{\kappa'}} \prod_{i<j} \left(\frac{|\zeta_{i,j}|}{\sqrt{n}\sqrt{\mu_i - \mu_j}}\right)^{m_{\kappa(i,j)}} r(9.103)$$

We obtain (9.96) by noting that the number of non-elementary modifiers is

$$\sum_{\kappa'} m_{\kappa'} = -\Gamma + \sum_{i<j} (m_{i,j} - m_{\kappa(i,j)}).$$

Proof of (9.97). Note that only admissible vectors in $\mathcal{O}_\lambda(\mathbf{m})$ can bring non-zero contributions. We shall split the sum into sub-sums using $\mathcal{O}_\lambda(\mathbf{m}) \bigcap \mathcal{V} = \bigcup_E \mathcal{V}^E(\mathbf{m})$, and compare each sub-sum against the benchmark $\mathcal{V}^0 = \mathcal{V}^{E^0}$.

From the bounds on $\vec{\zeta}$ and $\vec{z}$ we obtain $\|\vec{\zeta} + \vec{z}\| = O(n^\beta)$, so we can apply the previous points with $\vec{\zeta} + \vec{z}$ instead of $\vec{\zeta}$.

Using (9.90) and (9.95) and recalling that $\lambda \in \Lambda_{n,\alpha}$, we get:

$$\begin{aligned}
&\left\langle \sum_{f_{\mathbf{a}} \in \mathcal{V}^0} f_{\mathbf{a}} \middle| q_\lambda U(\vec{\zeta} + \vec{z}, \vec{\xi}, n)^{\otimes n} f_{\mathbf{0}} \right\rangle \\
&= \exp\left(i\phi - \frac{\|\vec{\zeta} + \vec{z}\|_2^2}{2} \right) \prod_{i<j} \frac{\left((\zeta_{i,j} + z_{i,j})\sqrt{n}\sqrt{\mu_i - \mu_j}\right)^{m_{i,j}}}{m_{i,j}!} r(n)
\end{aligned}$$

with error factor

$$r(n) = 1 + O\left(n^{-1+2\beta+\eta}\delta^{-1}, n^{-1/2+2\beta}\delta^{-1}, n^{-1+2\beta+\alpha}\delta^{-1}, n^{-1+2\eta}\delta^{-1}, n^{-1+\alpha+\eta}\delta^{-1}\right).$$

For $E \neq E^0$ we combine (9.96) and (9.91) to obtain

$$\begin{aligned}
&\left| \left\langle \sum_{f_{\mathbf{a}} \in \mathcal{V}^E} f_{\mathbf{a}} \middle| q_\lambda U(\vec{\zeta} + \vec{z}, \vec{\xi}, n) f_{\mathbf{0}} \right\rangle \right| \cdot \left| \left\langle \sum_{f_{\mathbf{a}} \in \mathcal{V}^0} f_{\mathbf{a}} \middle| q_\lambda U(\vec{\zeta} + \vec{z}, \vec{\xi}, n) f_{\mathbf{0}} \right\rangle \right|^{-1} \\
&\leq n^{-\Gamma} \prod_{i<j} \left(\frac{\lambda_i - \lambda_j}{n} \right)^{m_{\kappa(i,j)} - m_{i,j}} \frac{m_{i,j}!}{m_{\kappa(i,j)}!} \left(\frac{\|\vec{\zeta} + \vec{z}\|}{\sqrt{\delta n}} \right)^{-\Gamma} \cdot \\
&\qquad \prod_{i<j} \left(\frac{\sqrt{\delta n}|\zeta_{i,j} + z_{i,j}|}{\|\vec{\zeta} + z\|\sqrt{n}\sqrt{\mu_i - \mu_j}} \right)^{m_{\kappa(i,j)} - m_{i,j}} r(n) \\
&\leq O(n^{-\Gamma(1/2+\beta)})\delta^{-\Gamma/2} \prod_{i<j : m_{i,j} \neq 0} \left(\frac{|\zeta_{i,j} + z_{i,j}|\sqrt{\mu_i - \mu_j}}{m_{i,j}\|\vec{\zeta} + \vec{z}\|} \right)^{m_{\kappa(i,j)} - m_{i,j}} \\
&\leq O\left((2\delta^{-3/2} n^{-1/2+3\beta+2\epsilon})^\Gamma \right),
\end{aligned}$$

with $O(\cdot)$ uniform in Γ. In the second inequality we used

$$m_{i,j}!/m_{\kappa(i,j)}! \leq m_{i,j}^{m_{i,j} - m_{\kappa(i,j)}}, \qquad \sum_{i<j} (m_{\kappa(i,j)} - m_{i,j}) \geq -2\Gamma, \qquad \lambda \in \Lambda_{n,\alpha}$$

and in the third inequality we used

$$m_{i,j} \leq 2|\zeta_{i,j} + z_{i,j}|n^{\beta+\epsilon}, \qquad \frac{|\zeta_{i,j} + z_{i,j}|\sqrt{\mu_i - \mu_j}}{m_{i,j}\|\vec{\zeta} + \vec{z}\|} \leq 1.$$

Furthermore, for a given Γ, there are at most C^{Γ} different E such that $\Gamma(E) = \Gamma$, corresponding to the possible choices in the first two stages of the algorithm, where $C = C(d)$. Hence, if n is large enough, so that $2C\delta^{-3/2}n^{-1/2+3\beta+2\epsilon} < 1$, we have:

$$\left\langle \sum_{f_{\mathbf{a}} \in \mathcal{O}_{\lambda}(\mathbf{m})} f_{\mathbf{a}} \middle| q_{\lambda} U(\vec{\zeta} + z, \vec{\xi}, n) f_{\mathbf{0}} \right\rangle = \sum_{\Gamma} \left\langle \sum_{f_{\mathbf{a}} \in \mathcal{V}^{\Gamma}} f_{\mathbf{a}} \middle| q_{\lambda} U(\vec{\zeta} + z, \vec{\xi}, n) f_{\mathbf{0}} \right\rangle$$

$$= \left(1 + O(\delta^{-3/2} n^{-1/2+3\beta+2\epsilon})\right) \exp\left(i\phi - \frac{\|\vec{\zeta} + z\|_2^2}{2}\right) \prod_{i<j} \frac{\left((\vec{\zeta} + z)_{i,j}(\sqrt{n}\sqrt{\mu_i - \mu_j})\right)^{m_{i,j}}}{m_{i,j}!} r(n)$$

$$= \exp\left(i\phi - \frac{\|\vec{\zeta} + z\|_2^2}{2}\right) \prod_{i<j} \frac{\left((\vec{\zeta} + z)_{i,j}(\sqrt{n}\sqrt{\mu_i - \mu_j})\right)^{m_{i,j}}}{m_{i,j}!} r_2(n)$$

where the sum over Γ was bounded using a geometric series and

$$r_2(n) = 1 + O\left(n^{-1+2\beta+\eta}\delta^{-1}, n^{-1+\alpha+\beta}\delta^{-1}, n^{-1+2\eta}\delta^{-1}, n^{-1+\alpha+\eta}\delta^{-1}, \delta^{-3/2}n^{-1/2+3\beta+2\epsilon}\right).$$

This is exactly (9.97).

Proof of (9.98). We choose Γ^a and Γ^b satisfying the condition $\Gamma^b - \Gamma^a = |\mathbf{m}| - |\mathbf{l}|$ under which the inner products in (9.93) are non-zero. By multiplying (9.92) and (9.93), we see that:

$$\left| \left\langle \sum_{f_{\mathbf{a}} \in \mathcal{V}^{\Gamma^a}(\mathbf{l})} f_{\mathbf{a}} \middle| q_{\lambda} \sum_{f_{\mathbf{b}} \in \mathcal{V}^{\Gamma^b}(\mathbf{m})} f_{\mathbf{b}} \right\rangle \right| \leq (C|\mathbf{m}|)^{\Gamma^b} \prod_{i<j} \frac{(\lambda_i - \lambda_j)^{l_{i,j}}}{l_{i,j}!} \left(\frac{C|\mathbf{l}|^2}{n\delta^2}\right)^{\Gamma^a} \tag{9.104}$$

$$= (C|\mathbf{m}|)^{|\mathbf{m}|-|\mathbf{l}|} \prod_{i<j} \frac{(\lambda_i - \lambda_j)^{l_{i,j}}}{l_{i,j}!} \left(\frac{C|\mathbf{l}|^2|\mathbf{m}|}{n\delta^2}\right)^{\Gamma^a}.$$

It remains to sum up the upper bounds over all relevant pairs (Γ^a, Γ^b). If $n^{1-3\eta} > 2C/\delta^2$, the dominating term in the sum of bounds is that corresponding to the smallest possible Γ^a. The question is, what is the smallest possible value of Γ^a leading to non-zero inner products?

A necessary condition for $f_{\mathbf{a}}$ not to be orthogonal to $f_{\mathbf{b}}$ is that for each set S of suppressed and added values, the two vectors have the same multiplicities $m_S^{\mathbf{a}} = m_S^{\mathbf{b}}$.

The following argument provides a lower bound for $\Gamma(f_{\mathbf{a}}) + \Gamma(f_{\mathbf{b}})$. The idea is to count the minimum number of 'horizontal box shuffling' operations necessary in order to transform a Young tableau $t_{\mathbf{a}'} \in \mathcal{O}_{\lambda}(\mathbf{m})$ into the tableau $t_{\mathbf{a}}$. Since $|\mathbf{m}| \leq n^{\eta}$ and $\lambda_d \geq \delta n + O(n^{\alpha})$, the tableau $t_{\mathbf{a}'}$ can be chosen to have at most one modified box per column (thus $\Gamma(f_{\mathbf{a}'}) = 0$), and such that each of the modified columns of $t_{\mathbf{a}}$ are also modified in $t_{\mathbf{a}'}$. We also choose $t_{\mathbf{b}'}$ in a similar fashion.

Now at each step we horizontally move one elementary column modifier $\kappa(i,j)$ of $t_{\mathbf{a}'}$ (or $t_{\mathbf{b}'}$) into an already modified column, with the aim of constructing $t_{\mathbf{a}}$ (or $t_{\mathbf{b}}$).

Each such operation increases $\Gamma(f_{\mathbf{a}'}) + \Gamma(f_{\mathbf{b}'})$ by one. On the other hand the operation has the following effect on the $m_S^{\mathbf{a}'}$ (or $m_S^{\mathbf{b}'}$): the multiplicities $m_{\{(i,-),(j,+)\}}$ and m_{S_0} decrease by one, and $m_{S_0+\{(i,-),(j,+)\}}$ increases by one. Here S_0 is the signature of the column to which the box (i,j) is moved. Hence the distance $\sum_S |m_S^{\mathbf{a}'} - m_S^{\mathbf{b}'}|$ decreases by at most three. Since initially this quantity was equal to $\sum_{i<j} |l_{i,j} - m_{i,j}|$, we need at least $\sum_{i<j} |l_{i,j} - m_{i,j}|/3$ such operations before reaching our goal $m_S^{\mathbf{a}} = m_S^{\mathbf{b}}$. This means that $\Gamma(f_{\mathbf{a}}) + \Gamma(f_{\mathbf{b}}) \geq |\mathbf{l} - \mathbf{m}|/3$.

Together with $\Gamma^b - \Gamma^a = |\mathbf{m}| - |\mathbf{l}|$, this result yields $\Gamma^a \geq (|\mathbf{l} - \mathbf{m}| + 3|\mathbf{l}| - 3|\mathbf{m}|)/6$. Moreover Γ^a is non-negative.

Replacing in the above equation yields (9.98).

Proof of (9.100). Since $\mathbf{l} = \mathbf{m}$, equations (9.90) and (9.94) prove that the bound (9.104) is saturated when $\Gamma^a = 0$, up to the error factor $(1 + O(n^{-1+2\eta}/\delta))$. Hence the remainder term due to the other Γ consist in a geometric series with factor $\left(\frac{C|\mathbf{m}|^3}{n\delta^2}\right) = O(n^{1-3\eta}/\delta^2)$.

□

The only part of the proof we have still postponed is the following technical lemma:

Lemma 9.7.12. *If* $x_n = O(n^{1/2-\epsilon})$, *then*

$$\left(1 + \frac{x_n}{n}\right)^n = \exp(x_n)(1 + O(n^{-\epsilon})).$$

Proof. For simplicity we shall ignore the dependence on n and write $x = x_n$.

For any y such that $|y| \leq 1$, for any $n \in \mathbb{N}$, we have the Taylor expansion:

$$(1+y)^n = \sum_{k=1}^{\infty} \binom{n}{k} y^k.$$

Now $(n-k)^k/k! \leq \binom{n}{k} \leq n^k/k!$ for $n \geq k$. If $k \leq n^{1/2-\epsilon/2}$, then $(n-k)^k = n^k(1 + O(n^{-\epsilon}))$. If $k \geq n^{1/2-\epsilon/2}$, then $n^k/k! = O(n^{(1/2+\epsilon/2)k})$. So that if $y = x/n =$

$O(n^{-1/2-\epsilon})$,

$$\begin{aligned}(1+x/n)^n &= (1+O(n^{-\epsilon}))\sum_{k=0}^{n^{1/2-\epsilon/2}}\frac{x^k}{k!}+\sum_{k>n^{1/2-\epsilon/2}}O(n^{(1/2+\epsilon/2)k}(x/n)^k\\ &= (1+O(n^{-\epsilon}))\exp(x)+\sum_{k>n^{1/2-\epsilon/2}}(O(n^{(1/2+\epsilon/2)k}-n^k/k!)(x/n)^k\\ &= (1+O(n^{-\epsilon}))\exp(x)+O(e^{-n^{1/2-\epsilon/2}})\\ &= (1+O(n^{-\epsilon}))\exp(x),\end{aligned}$$

as $\exp(x)\geq C\exp(-n^{1/2-\epsilon}))$ for some constant $C>0$. □

9.7.3 Proof of Lemma 9.5.4 and non-orthogonality issues

Lemma 9.7.13. *Let* $(\mathbf{m},\lambda)$ *and* $(\mathbf{l},\lambda)$ *be semistandard Young tableaux with diagram* λ *and define* $|\mathbf{m}|:=\sum_{i<j}m_{ij}$ *and* $|\mathbf{l}-\mathbf{m}|:=\sum_{i<j}|l_{i,j}-m_{ij}|$.

If

$$\sum_{j>i}m_{i,j}-\sum_{j<i}m_{j,i}\neq\sum_{j>i}l_{i,j}-\sum_{j<i}l_{j,i}$$

for some $1\leq i\leq d$, *then*

$$\langle\mathbf{m},\lambda|\mathbf{l},\lambda\rangle=0.$$

Otherwise, we derive an upper bound under the following conditions. We assume that $\lambda_i-\lambda_{i+1}>\delta n$ *for all* $1\leq i\leq d-1$ *and* $\lambda_d>\delta n$, *for some* $\delta>0$. *Furthermore we assume* $|\mathbf{l}|\leq|\mathbf{m}|\leq n^\eta$ *for some* $\eta<1/3$ *and that* $Cn^{3\eta-1}/\delta^2<1$ *where* $C=C(d)$ *is a constant.*

Then:

$$|\langle\mathbf{m},\lambda|\mathbf{l},\lambda\rangle|\leq(C'n)^{-\eta(|\mathbf{m}|-|\mathbf{l}|)/4}(C'n)^{(9\eta-2)|\mathbf{m}-\mathbf{l}|/12}\delta^{(|\mathbf{m}|-|\mathbf{l}|)/2-|\mathbf{m}-\mathbf{l}|/3}(1+O(n^{-1+3\eta}/\delta)). \tag{9.105}$$

where $C'=C'(d,\eta)$ *and the constant in the remainder term depends only on* d. *The right side is of order less than* $n^{(9\eta-2)|\mathbf{m}-\mathbf{l}|/12}$ *and converges to zero for* $\eta<2/9$ *when* $n\to\infty$.

Proof. We know that $|\mathbf{m},\lambda\rangle$ is a linear combination of n-tensor product vectors in which the basis vector f_i appears exactly $\lambda_i-\sum_{j>i}m_{i,j}+\sum_{j<i}m_{j,i}$ times. As two tensor basis vectors are orthogonal if they do not have the same number of f_i in the decomposition, we get that $\langle\mathbf{m},\lambda|\mathbf{l},\lambda\rangle=0$ if $\sum_{j>i}m_{i,j}+\sum_{j<i}m_{j,i}\neq\sum_{j>i}l_{i,j}+\sum_{j<i}l_{j,i}$ for any $1\leq i\leq d$.

In the general case,

$$\langle \mathbf{m}, \lambda | \mathbf{l}, \lambda \rangle = \frac{\langle q_\lambda p_\lambda f_\mathbf{m} | q_\lambda p_\lambda f_\mathbf{l} \rangle}{\sqrt{\langle q_\lambda p_\lambda f_\mathbf{m} | q_\lambda p_\lambda f_\mathbf{m} \rangle \langle q_\lambda p_\lambda f_\mathbf{l} | q_\lambda p_\lambda f_\mathbf{l} \rangle}}. \tag{9.106}$$

We use the fact that q_λ is a projection, up to a constant factor (cf. (9.53),(9.55)), and erase the q_λ at the left of each scalar product, and we decompose $p_\lambda f_\mathbf{m}$ and $p_\lambda f_\mathbf{l}$ on orbits as in (9.85). Since the multiplicity of the elements in the orbits are the same in numerator and denominator, we end up with:

$$\langle \mathbf{m}, \lambda | \mathbf{l}, \lambda \rangle = \frac{\langle \sum_{f_\mathbf{a} \in \mathcal{O}_\lambda(\mathbf{m})} f_\mathbf{a} | q_\lambda \sum_{f_\mathbf{b} \in \mathcal{O}_\lambda(\mathbf{l})} f_\mathbf{b} \rangle}{\langle \sum_{f_\mathbf{a} \in \mathcal{O}_\lambda(\mathbf{m})} f_\mathbf{a} | q_\lambda \sum_{f_{\mathbf{a}'} \in \mathcal{O}_\lambda(\mathbf{m})} f_{\mathbf{a}'} \rangle \langle \sum_{f_\mathbf{b} \in \mathcal{O}_\lambda(\mathbf{l})} f_\mathbf{b} | q_\lambda \sum_{f_{\mathbf{b}'} \in \mathcal{O}_\lambda(\mathbf{l})} f_{\mathbf{b}'} \rangle} \tag{9.107}$$

We use (9.100) for the denominator:

$$\begin{aligned} &\left\langle \sum_{f_\mathbf{a} \in \mathcal{O}_\lambda(\mathbf{m})} f_\mathbf{a} \middle| q_\lambda \sum_{f_{\mathbf{a}'} \in \mathcal{O}_\lambda(\mathbf{m})} f_{\mathbf{a}'} \right\rangle \left\langle \sum_{f_\mathbf{b} \in \mathcal{O}_\lambda(\mathbf{l})} f_\mathbf{b} \middle| q_\lambda \sum_{f_{\mathbf{b}'} \in \mathcal{O}_\lambda(\mathbf{l})} f_{\mathbf{b}'} \right\rangle \\ &= \prod_{1 \le i < j \le d} \frac{(\lambda_i - \lambda_j)^{(m_{i,j}+l_{i,j})/2}}{\sqrt{m_{i,j}!\, l_{i,j}!}} (1 + O(n^{3\eta-1}/\delta))), \end{aligned}$$

and the numerator is bounded as in (9.98). Then, under the assumption $|\mathbf{m}| \ge |\mathbf{l}|$ we have

$$|\langle \mathbf{m}, \lambda | \mathbf{l}, \lambda \rangle| \le (C|\mathbf{m}|)^{|\mathbf{m}|-|\mathbf{l}|} \left(\frac{C|\mathbf{m}|^3}{\delta^2 n} \right)^{\Gamma_{min}} \cdot \prod_{i<j} (\lambda_i - \lambda_j)^{(l_{i,j}-m_{i,j})/2} \sqrt{\frac{m_{i,j}!}{l_{i,j}!}} \cdot \left(1 + \left(O(n^{3\eta-1}/\delta)\right)\right),$$

where $\Gamma_{min} = ((|\mathbf{l}-\mathbf{m}| + 3|\mathbf{l}| - 3|\mathbf{m}|)/6) \wedge 0$.

The factorials can be bounded as

$$\prod_{i<j} \sqrt{\frac{m_{i,j}!}{l_{i,j}!}} \le |\mathbf{m}|^{\sum (m_{i,j}-l_{i,j})_+/2} = |\mathbf{m}|^{(|\mathbf{m}-\mathbf{l}|+|\mathbf{m}|-|\mathbf{l}|)/4}.$$

Since $|\mathbf{m}| \le n^\eta$ and $Cn^{3\eta-1}/\delta^2 < 1$, we have

$$\left(\frac{C|\mathbf{m}|^3}{\delta^2 n} \right)^{\Gamma_{min}} \le \left(\frac{Cn^{3\eta-1}}{\delta^2} \right)^{(|\mathbf{l}-\mathbf{m}|+3|\mathbf{l}|-3|\mathbf{m}|)/6}.$$

Since $\lambda_i - \lambda_j > n\delta$ we have

$$\prod_{1 \le i < j \le d} (\lambda_i - \lambda_j)^{(l_{i,j}-m_{i,j})/2} \le (n\delta)^{(|\mathbf{l}|-|\mathbf{m}|)/2}.$$

The constant $C = C(d)$ can be replaced by another constant $C' = C'(d, \eta)$ such that all powers of n appear in the form $(C'n)^\gamma$. Putting the bounds together we get

$$|\langle \mathbf{m}, \lambda | \mathbf{l}, \lambda \rangle| \leq \delta^{(|\mathbf{m}|-|\mathbf{l}|)/2-|\mathbf{m}-\mathbf{l}|/3}(C'n)^{-\eta(|\mathbf{m}|-|\mathbf{l}|)/4}(C'n)^{(9\eta-2)|\mathbf{m}-\mathbf{l}|/12}(1+O(n^{-1+3\eta}/\delta))$$

□

A consequence of this lemma is the following.

Corollary 9.7.14. *Let $\eta < 2/9$ and let $(\mathbf{m}, \lambda)$ be such that $|\mathbf{m}| \leq n^\eta$. Assume as in Lemma 9.7.13 that $\lambda_i - \lambda_{i+1} > \delta n$ for all $1 \leq i \leq d-1$ and $\lambda_d > \delta n$, for some $\delta > 0$, and that $Cn^{3\eta-1}/\delta^2 < 1$ where $C = C(d)$ is a constant.*

Then there exists a constant $C'' = C''(d, \eta)$ such that

$$\sum_{\substack{|\mathbf{l}| \leq n^\eta \\ \mathbf{l} \neq \mathbf{m}}} |\langle \mathbf{m}, \lambda | \mathbf{l}, \lambda \rangle| \leq (C''n)^{(9\eta-2)/12}\delta^{-1/3}. \tag{9.108}$$

Proof. Recall that the bound (9.105) is given for $|\mathbf{m}| \geq |\mathbf{l}|$. If on the contrary $|\mathbf{l}| > |\mathbf{m}|$, we must change all the $|\mathbf{m} - \mathbf{l}|$ into $|\mathbf{l} - \mathbf{m}|$, so that these terms are always positive. Now, they are always in exponents of values less than one. We shall therefore neglect all those terms.

Hence the expression on the left side of (9.108) is bounded from above by

$$2\sum_{k \geq 1} N(k) \left[(C'n)^{(9\eta-2)/12}\delta^{-1/3}\right]^k$$

where $N(k)$ is the number of $\mathbf{l}$'s for which $|\mathbf{m} - \mathbf{l}| = k$.

Since there are $d(d-1)/2$ pairs $1 \leq i < j \leq d$, there are at most $(k+1)^{d(d-1)/2}$ different choices for the values $\{|l_{i,j} - m_{i,j}| : i < j\}$ satisfying $\sum |l_{i,j} - m_{i,j}| = k$. Moreover, there are $2^{d(d-1)/2}$ sign choices which fix $\mathbf{l} = \{l_{i,j}\}$ completely. Thus $N(k) \leq (2(k+1))^{d(d-1)/2} \leq c^k$ for some constant c which can be incorporated in the geometric series starting at $k = 1$, hence the desired estimate.

□

We use this quasi-orthogonality to build an isometry $V_\lambda : \mathcal{H}_\lambda \to \mathcal{F}$ which maps the relevant finite-dimensional vectors $|\mathbf{m}, \lambda\rangle$ 'close' to their Fock counterparts $|\mathbf{m}\rangle$. This is the aim of Lemma 9.5.4.

Lemma 9.7.15. *Let A be a contraction (i.e. $A^*A \leq \mathbf{1}$) from a finite space $\mathcal{H}$ to an infinite space $\mathcal{K}$. Then there is an $R : \mathcal{H} \to \mathcal{K}$ such that $A + R$ is an isometry and* $\mathrm{Range}(A) \perp \mathrm{Range}(R)$.

As a consequence, for any unit vector f, we have $\|Rf\|^2 = 1 - \|Af\|^2$.

Proof. As $\mathcal{K}$ is infinite-dimensional, we may consider a subspace $\mathcal{H}'$ of $\mathcal{K}$, orthogonal to $\mathrm{Range}(A)$, and the same dimension as $\mathcal{H}$, so that we can find an isomorphism I from $\mathcal{H}$ to $\mathcal{H}'$. We then take $R = I\sqrt{\mathbf{1} - A^*A}$.

□

Proof of Lemma 9.5.4. Let $A_\lambda : \mathcal{H}_\lambda \to \mathcal{F}$ be defined by

$$A_\lambda := \frac{1}{\sqrt{1 + (Cn)^{(9\eta-2)/12}/\delta^{1/3}}} \sum_{|\mathbf{l}| \leq n^\eta} |\mathbf{l}\rangle \langle \mathbf{l}, \lambda| .$$

Then,

$$\begin{aligned} A_\lambda^* A_\lambda &= \frac{1}{1 + (Cn)^{(9\eta-2)/12}/\delta^{1/3}} \sum_{|\mathbf{l}| \leq n^\eta} |\mathbf{l}, \lambda\rangle \langle \mathbf{l}, \lambda| \\ &\leq \mathbf{1}_{\mathcal{H}_\lambda}. \end{aligned}$$

where the last inequality follows from Corollary 9.7.14 and the following argument. It is enough to show that all eigenvalues of $A_\lambda^* A_\lambda$ are smaller than 1. Let $\sum_{\mathbf{m}} c_{\mathbf{m}} |\mathbf{m}, \lambda\rangle$ be an eigenvector of $A_\lambda^* A_\lambda$, and a the corresponding eigenvalue. Then by the linear independence of $|\mathbf{m}, \lambda\rangle$ we get that for each $\mathbf{l}$

$$\frac{1}{1 + (Cn)^{(9\eta-2)/12}/\delta^{1/3}} \sum_{|\mathbf{m}| \leq n^\eta} \langle \mathbf{l}, \lambda | \mathbf{m}, \lambda \rangle c_{\mathbf{m}} = a c_{\mathbf{l}}.$$

If $\mathbf{l}_0$ is an index for which $|c_{\mathbf{l}}|$ is maximum, then by taking absolute values on both sides we obtain

$$a \leq \frac{1}{1 + (Cn)^{(9\eta-2)/12}/\delta^{1/3}} \sum_{|\mathbf{m}| \leq n^\eta} |\langle \mathbf{l}, \lambda | \mathbf{m}, \lambda \rangle| \leq 1.$$

Now we may apply Lemma 9.7.15, and find an R_λ such that $A_\lambda + R_\lambda$ is an isometry, and $\mathrm{Range}(R_\lambda) \perp \mathrm{Range}(A)$, so that $\langle \mathbf{m}| R_\lambda = 0$. We define $V_\lambda := A_\lambda + R_\lambda$. Then

$$\begin{aligned} \langle \mathbf{m}| V_\lambda &= \langle \mathbf{m}| (A_\lambda + R_\lambda) \\ &= \langle \mathbf{m}| A_\lambda \\ &= \frac{1}{\sqrt{1 + (Cn)^{(9\eta-2)/12}/\delta^{1/3}}} \langle \mathbf{m}| \sum_{|\mathbf{l}| \leq n^\eta} |\mathbf{l}\rangle \langle \mathbf{l}, \lambda| \\ &= \frac{1}{\sqrt{1 + (Cn)^{(9\eta-2)/12}/\delta^{1/3}}} \langle \mathbf{m}, \lambda| . \end{aligned}$$

□

Recall By Lemma 9.7.13 we have $\langle \mathbf{m}, \lambda | \mathbf{l}, \lambda \rangle = 0$ if $m_i \neq l_i$ for some i, where m_i in the total number of i in $\mathbf{m}$ (cf. (9.125)). In particular, $|\mathbf{0}, \lambda\rangle$ is orthogonal on all other basis vectors. This means that we can choose the isometry V_λ to satisfy $V_\lambda |\mathbf{0}, \lambda\rangle = |\mathbf{0}\rangle$, and such that the relation above holds for all $0 < |\mathbf{m}| \leq n^\eta$.

9.7.4 Proof of Lemma 9.6.4 on mapping rotations into displacements

We first recall a few definitions and notations. We denote by $D^{\vec{z}}$ the displacement operation (super-operator) acting on observables in the multimode Fock space $\mathcal{F}$ as

$$D^{\vec{z}}(W(\vec{y})) := \mathrm{Ad}[W(\vec{z})]\,(W(\vec{y})) = e^{2i\sigma(\vec{y},\vec{z})}\, W(\vec{z}+\vec{y}), \qquad \vec{y}, \vec{z} \in \mathbb{C}^{d(d-1)/2}.$$

The operation acts as displacement on coherent states, in particular

$$D^{\vec{\zeta}+\vec{z}}(|\mathbf{0}\rangle\langle\mathbf{0}|) = |\vec{\zeta}+\vec{z}\rangle\langle\vec{\zeta}+\vec{z}|.$$

Similarly, on the finite dimensional space $\left(\mathbb{C}^d\right)^{\otimes n}$ we have the action (cf. (9.83))

$$\Delta^{\vec{\zeta},\vec{\xi},n}(A) = \mathrm{Ad}[U(\vec{\zeta},\vec{\xi},n)](A) := U(\vec{\zeta}/\sqrt{n}, \vec{\xi}/\sqrt{n})^{\otimes n}\, A\, U^*(\vec{\zeta}/\sqrt{n}, \vec{\xi}/\sqrt{n})^{\otimes n},$$

whose restriction to the block λ is $\Delta_\lambda^{\vec{\zeta},\vec{\xi},n} = \mathrm{Ad}[U_\lambda(\vec{\zeta},\vec{\xi},n)]$.

The isometric embedding $T_\lambda(\cdot) := V_\lambda \cdot V_\lambda^*$ and its 'adjoint' $T_\lambda^*(\cdot) := V_\lambda^* \cdot V_\lambda$ satisfy

$$T_\lambda \Delta_\lambda^{\vec{\zeta}+\vec{z},\vec{\xi},n} T_\lambda^*(|\mathbf{0}\rangle\,\langle\mathbf{0}|) = V_\lambda |\vec{\zeta}+\vec{z},\vec{\xi},\lambda\rangle\langle\vec{\zeta}+\vec{z},\vec{\xi},\lambda| V_\lambda^*$$

where $|\vec{\zeta}+\vec{z},\vec{\xi},\lambda\rangle := U_\lambda(\vec{\zeta}+\vec{z},\vec{\xi},n)\,|\mathbf{0},\lambda\rangle$ are the 'finite dimensional coherent states'.

According to Lemma 9.5.4, the coordinates of $V_\lambda|\vec{\zeta}+\vec{z},\vec{\xi},\lambda\rangle$ in the Fock basis are described by:

$$\langle \mathbf{m}|V_\lambda|\vec{\zeta}+\vec{z},\vec{\xi},\lambda\rangle = \begin{cases} \langle \mathbf{m},\lambda|U_\lambda(\vec{\zeta}+\vec{z},\vec{\xi},n)|\mathbf{0},\lambda\rangle(1+O(n^{(9\eta-2)/12}\delta^{-1/3})) \text{ if } |\mathbf{m}| \leq n^\eta; \\ \text{something not important if } |\mathbf{m}| > n^\eta. \end{cases} \tag{9.109}$$

Using the relation $\||f\rangle\langle f| - |f'\rangle\langle f'|\|_1 = 2\sqrt{1-|\langle f|f'\rangle|^2}$, which holds for unital vectors f, f', the statement of the lemma is equivalent to

$$\sup_{\|\vec{z}\|\leq n^\beta} \sup_{\vec{\zeta}\in\Theta_{n,\beta}} \sup_{\|\vec{\xi}\|\leq n^{-1/2+2\beta/\delta}} \sup_{\lambda\in\Lambda_{n,\alpha}} 1 - \left|\langle \vec{z}+\vec{\zeta}|V_\lambda|\vec{\zeta}+\vec{z},\vec{\xi},\lambda\rangle\right| = R(n)^2, \tag{9.110}$$

with $R(n)$ the original remainder term.

We shall prove formula (9.110) by decomposing these vectors in the Fock basis, that is

$$\langle\vec{\zeta}+\vec{z}|V_\lambda|\vec{\zeta}+\vec{z},\vec{\xi},\lambda\rangle=\sum_{\mathbf{m}}\langle\vec{\zeta}+\vec{z}|\mathbf{m}\rangle\langle\mathbf{m}|V_\lambda|\vec{\zeta}+\vec{z},\vec{\xi},\lambda\rangle. \tag{9.111}$$

The estimates are based on the following observations.

1) The coherent states have significant coefficients $\langle\vec{\zeta}+\vec{z}|\mathbf{m}\rangle$ only for 'small' $\mathbf{m}$'s, i.e. those in the set

$$\mathcal{M}:=\{\mathbf{m}: m_{i,j}\leq|(\vec{\zeta}+\vec{z})_{i,j}|^2n^\epsilon,\quad i<j\}. \tag{9.112}$$

In particular, since $2\beta+\epsilon<\eta$ we have $\mathcal{M}\subset\{\mathbf{m}:|\mathbf{m}|\leq n^\eta\}$.

2) The coefficients $\langle\mathbf{m}|V_\lambda|\vec{\zeta}+\vec{z},\vec{\xi},\lambda\rangle$ are uniformly close to $\exp(i\phi)\langle\vec{\zeta}+\vec{z}|\mathbf{m}\rangle$ where ϕ is a fixed real phase, in particular uniformly over $\mathbf{m}\in\mathcal{M}$.

3) If $a_\mathbf{m}$ and $b_\mathbf{m}$ are the two sets of coefficients, such that $\sum_\mathbf{m}|a_\mathbf{m}|^2=\sum_\mathbf{m}|b_\mathbf{m}|^2=1$, then

$$1-\left|\sum_{\mathbf{m}}a_\mathbf{m}b_\mathbf{m}\right|\leq 1-\left|\sum_{\mathbf{m}\in\mathcal{M}}a_\mathbf{m}b_\mathbf{m}\right|+\left|\sum_{\mathbf{m}\notin\mathcal{M}}a_\mathbf{m}b_\mathbf{m}\right|\leq 2\left(1-\left|\sum_{\mathbf{m}\in\mathcal{M}}a_\mathbf{m}b_\mathbf{m}\right|\right). \tag{9.113}$$

The precise statement in point 1) is

$$\sum_{\mathbf{m}\notin\mathcal{M}}|\langle\vec{\zeta}+\vec{z}|\mathbf{m}\rangle|^2\leq d^2n^{-\beta}. \tag{9.114}$$

Indeed, the inner products can be written as a product over the (i,j) oscillators and we have the bound

$$\sum_{\mathbf{m}\notin\mathcal{M}}|\langle\vec{\zeta}+\vec{z}|\mathbf{m}\rangle|^2\leq\sum_{i<j}\exp(-x_{i,j})\sum_{k>x_{i,j}n^\epsilon}\frac{x_{ij}^k}{k!},\qquad x_{i,j}=|(\vec{\zeta}+\vec{z})_{i,j}|^2.$$

Each of the terms in the sum is a tail of Poisson distribution and is bounded by $n^{-\epsilon n^\beta}$ if $x_{i,j}\geq 1$ and by $n^{-\beta}$ if $x_{i,j}<1$.

We turn now to point 2). From the third line of (9.109) we get

$$\begin{aligned}\langle\mathbf{m}|V_\lambda U_\lambda(\vec{\zeta}+\vec{z},\vec{\xi},n)|\mathbf{0},\lambda\rangle&=\frac{\langle y_\lambda f_\mathbf{m}|y_\lambda U(\vec{\zeta}+\vec{z},\vec{\xi},n)|f_\mathbf{0}\rangle}{\sqrt{\langle y_\lambda f_\mathbf{0}|y_\lambda f_\mathbf{0}\rangle}\sqrt{\langle y_\lambda f_\mathbf{m}|y_\lambda f_\mathbf{m}\rangle}}(1+O(n^{(9\eta-2)/12}\delta^{-1/3}))\\&=\frac{\langle p_\lambda f_\mathbf{m}|q_\lambda U(\vec{\zeta}+\vec{z},\vec{\xi},n)f_\mathbf{0}\rangle}{\sqrt{\langle p_\lambda f_\mathbf{m}|q_\lambda p_\lambda f_\mathbf{m}\rangle}}(1+O(n^{(9\eta-2)/12}\delta^{-1/3}))\end{aligned}$$

where we have used (9.55) and (9.58).

We recall that $\mathcal{O}_\lambda(\mathbf{m})$ is the orbit in $(\mathbb{C}^d)^{\otimes n}$ of $f_\mathbf{m}$ under $\mathcal{R}_\lambda$ and that we have the decomposition

$$p_\lambda f_\mathbf{m} = \sum_{f_\mathbf{a} \in \mathcal{O}_\lambda(\mathbf{m})} \frac{\#\mathcal{R}_\lambda}{\#\mathcal{O}_\lambda(\mathbf{m})} f_\mathbf{a}.$$

Then, by employing formulas (9.97) and (9.100), we can write

$$\langle \mathbf{m}|V_\lambda U_\lambda(\vec{\zeta}+\vec{z},\vec{\xi},n)|\mathbf{0},\lambda\rangle = \frac{\sum_{f_\mathbf{a}\in\mathcal{O}_\lambda(\mathbf{m})}\langle f_\mathbf{a}|q_\lambda U(\vec{\zeta}+\vec{z},\vec{\xi},n)f_\mathbf{0}\rangle}{\sqrt{\sum_{f_\mathbf{a},f_\mathbf{b}\in\mathcal{O}_\lambda(\mathbf{m})}\langle f_\mathbf{a}|q_\lambda f_\mathbf{b}\rangle}}(1+O(n^{(9\eta-2)/12}\delta^{-1/3}) \tag{9.115}$$

$$= e^{i\phi-\|\vec{\zeta}+\vec{z}\|_2^2/2}\prod_{i\le j}\frac{(\vec{\zeta}+\vec{z})_{i,j}^{m_{i,j}}}{\sqrt{m_{i,j}!}}\left(\frac{n(\mu_i-\mu_j)}{\lambda_i-\lambda_j}\right)^{m_{i,j}/2} r(n).$$

The corresponding remainder term is

$$r(n) = 1+O\big(n^{(9\eta-2)/12}\delta^{-1/3}, n^{-1+2\beta+\eta}\delta^{-1}, n^{-1/2+3\beta+2\epsilon}\delta^{-3/2}, \\ n^{-1+\alpha+2\beta}\delta^{-1}, n^{-1+\alpha+\eta}\delta^{-1}, n^{-1+3\eta}\delta^{-1}\big)$$

and the phase is:

$$\phi = \sqrt{n}\sum_{i=1}^{d-1}(\mu_i-\mu_{i+1})\xi_i.$$

Since $\lambda \in \Lambda_{n,\alpha}$ and the eigenvalues are separated by δ we have $\left(\frac{n(\mu_i-\mu_j)}{\lambda_i-\lambda_j}\right)^{m_{i,j}/2} = 1+O(n^{\alpha-1+\eta}/\delta)$ and the error can be absorbed in $r(n)$.

In conclusion, for $\mathbf{m}$ satisfying (9.112), we have:

$$\langle \mathbf{m}|V_\lambda U(\vec{\zeta}+\vec{z},\vec{\xi},n)|\mathbf{0},\lambda\rangle = \exp(i\phi)\langle \mathbf{m}|\vec{\zeta}+\vec{z}\rangle r(n).$$

Inserting this result into (9.111), and using (9.113) and (9.114), we get

$$1-\left|\langle \vec{z}+\vec{\zeta}|V_\lambda U(\vec{\zeta}+\vec{z},\vec{\xi},n)|\mathbf{0},\lambda\rangle\right| = O\left(1-r(n), \sum_{\mathbf{m}\notin\mathcal{M}}|\langle \mathbf{m}|\vec{\zeta}+\vec{z}\rangle|^2\right) = R_2(n),$$

with

$$R_2(n) = O\left(n^{(9\eta-2)/12}\delta^{-1/3}, n^{-1+2\beta+\eta}\delta^{-1}, n^{-1/2+3\beta+2\epsilon}\delta^{-3/2}, n^{-1+\alpha+2\beta}\delta^{-1}, \\ n^{-1+\alpha+\eta}\delta^{-1}, n^{-1+3\eta}\delta^{-1}, n^{-\beta}\right).$$

Through expression (9.110), noticing that $R_2(n) = R(n)^2$, we see that we have proved the lemma.

□

9.7.5 Proof of Lemma 9.6.2 on typical Young diagrams

Recall that the state $\rho^{\theta,n} := \rho_{\theta/\sqrt{n}}^{\otimes n}$ has the decomposition over 'blocks' λ given by (9.21). The probability distribution over Young diagrams $p_\lambda^{\vec{\zeta},\vec{u},n}$ depends only on the diagonal parameters $\vec{u}$ and is given by

$$p_\lambda^{\vec{\zeta},\vec{u},n} = c_n^\lambda \sum_{\mathbf{m}\in\lambda} \prod_{i=1}^{d} (\mu_i^{\vec{u},n})^{\lambda_i} \prod_{j=i+1}^{d} \left(\frac{\mu_j^{\vec{u},n}}{\mu_i^{\vec{u},n}}\right)^{m_{i,j}},$$

with

$$c_n^\lambda = \binom{n}{\lambda_1, \lambda_2, \ldots, \lambda_d} \prod_{l=1}^{d} \frac{\lambda_l! \prod_{k=l+1}^{d} (\lambda_l - \lambda_k + k - l)}{(\lambda_l + d - l)!}.$$

The above formula can be understood as follows. By invariance under rotations we can take $\vec{\zeta} = 0$ and the state is diagonal in the standard basis basis $\left(\mathbb{C}^d\right)^{\otimes n}$ formed by the vector $f_{\mathbf{a}}$. Each eigenprojector carries a weight $\prod_{i=1}^{d} (\mu^{\vec{u},n})^{m_i}$ where m_i is the multiplicity of the vector f_i in the tensor product $f_{\mathbf{a}}$. Thus, we only need to add all multiplicities over vectors that are 'inside' the block λ. Since the irreducible representation has basis $f_{\mathbf{m}}$ labelled by semistandard Young tableaux, we get a factor

$$\prod_{i=1}^{d} (\mu_i^{\vec{u},n})^{m_i} = \prod_{i=1}^{d} (\mu_i^{\vec{u},n})^{\lambda_i} \prod_{j=i+1}^{d} \left(\frac{\mu_j^{\vec{u},n}}{\mu_i^{\vec{u},n}}\right)^{m_{i,j}}.$$

The additional factor c_n^λ is the dimension of $\mathcal{K}_\lambda$, on which the state is proportional to the identity.

Recall that $\mu_i^{\vec{u},n} = \mu_i + u_i/\sqrt{n}$ for $1 \leq i \leq (d-1)$ and $\mu_d^{\vec{u},n} = \mu_d - (\sum_i u_i)/\sqrt{n}$. If $\delta \geq 2dn^{\alpha-1} \geq 2dn^{\gamma-1/2}$ then $\mu_j^{\vec{u},n}/\mu_i^{\vec{u},n} \leq 1$ for all $\|\vec{u}\| \leq n^\gamma$. Moreover $m_{i,j} \leq n$ for all (i,j), so the total number of $\mathbf{m}$'s is smaller than n^{d^2}. Thus

$$\sum_{\mathbf{m}} \prod_{i<j} (\mu^{\vec{u},n})^{\lambda_i} \left(\frac{\mu_j^{\vec{u},n}}{\mu_i^{\vec{u},n}}\right)^{m_{i,j}} \leq n^{d^2}.$$

On the other hand $\mathbf{m} = \mathbf{0}$ is always in the set of possible $\mathbf{m}$, so that

$$\sum_{\mathbf{m}} \prod_{i<j} \left(\frac{\mu_j^{\vec{u},n}}{\mu_i^{\vec{u},n}}\right)^{m_{i,j}} \geq 1.$$

One can easily verify that

$$1 \geq \prod_{l=1}^{d} \frac{\lambda_l! \prod_{k=l+1}^{d} (\lambda_l - \lambda_k + k - l)}{(\lambda_l + d - l)!} \geq \frac{1}{(n+d)^{d^2}}.$$

The remaining factor is the multinomial law. We now show that this is the dominating part. Let us write $(Y_1, \ldots, Y_d)$ for the multinomial random variable. Then we have

$$\mathbb{P}[|Y_i - n\mu_i^{\vec{u},n}| \geq x] \leq 2\exp\left(-\frac{2x^2}{n}\right). \tag{9.116}$$

Indeed each Y_i is a sum of independent Bernoulli variables $X_1, \ldots, X_n$ with $\mathbb{P}(X_k = 1) = \mu_i^{\vec{u},n}$ and $\mathbb{P}(X_k = 0) = 1 - \mu_i^{\vec{u},n}$, and by Hoeffding's inequality (van der Vaart et Wellner, J.A., 1996)

$$\mathbb{P}[|\sum_{k=1}^{n} X_k - \mathbb{E}[X_k]| \geq x] \leq 2\exp\left(-\frac{2x^2}{n}\right). \tag{9.117}$$

By definition, for any $\lambda \notin \Lambda_{n,\alpha}$ there exists an i such that $|\lambda_i - n\mu_i| \geq n^\alpha$, which implies $|\lambda_i - n\mu_i^{\vec{u},n}| \geq n^\alpha - dn^{\gamma+1/2}$. With $n^{\alpha-\gamma-1/2} > 2d$, the upper bound is simply $n^\alpha/2$ and we have

$$\begin{aligned}\sum_{\lambda \notin \Lambda n,\alpha} \|b_\lambda^{\theta,n}\|_1 = \mathbb{P}[\lambda \notin \Lambda_{n,\alpha}] &\leq n^{d^2} \sum_{i=1}^{d} \mathbb{P}[|Y_i - n\mu_i^{\vec{u},n}| \geq n^\alpha/2] \\ &\leq 2dn^{d^2}\exp(-n^{2\alpha-1}/2).\end{aligned}$$

□

9.7.6 Proof of Lemma 9.6.1 and Lemma 9.6.8

We shall use multinomials as an intermediate step. Recalling that $b_\lambda^{\theta,n} = p_\lambda^{\theta,n}\tau_\lambda^n$, we can write:

$$\begin{aligned}\left\|\mathcal{N}(\vec{u}, V_\mu) - \sum_\lambda b_\lambda^{\theta,n}\right\|_1 \leq \left\|p^{\theta,n} - M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}}\right\|_1 + \\ \left\|\mathcal{N}(\vec{u}, V_\mu) - \sum_\lambda M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}}(\lambda)\tau_\lambda^n\right\|_1,\end{aligned} \tag{9.118}$$

where $M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}}$ is the d-multinomial with coefficients $\mu_i^{\vec{u},n}$.

For background, what we really prove in this lemma is the equivalence of the following classical experiments, together with an explicit rate:

$$\begin{aligned}\mathcal{P}_n &= \left\{p^{\vec{u},n}, \|\vec{u}\| \leq n^\gamma\right\} \\ \mathcal{M}_n &= \left\{M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}}, \|\vec{u}\| \leq n^\gamma\right\} \\ \mathcal{G}_n &= \left\{\mathcal{N}(\vec{u}, V_\mu), \|\vec{u}\| \leq n^\gamma\right\}.\end{aligned}$$

Remember that $p^{\theta,n} = p^{\vec{u},n}$. We shall usually shorthand $M^{n,\vec{u}} = M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}}$.

We first bound the first term in (9.118), planning to obtain:

$$\sup_{\|\vec{u}\|\leq n^\gamma} \left\| p^{\vec{u},n} - M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}} \right\|_1 \leq C \frac{n^{-1/2+\gamma} + n^{\alpha-1}}{\delta}. \tag{9.119}$$

To show this, we rewrite:

$$\begin{aligned}
\left\| p^{\vec{u},n} - M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}} \right\|_1 &= \sum_{|\lambda|=n} |p_\lambda^{\vec{u},n} - M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}}(\lambda)| \\
&\leq \sum_{\lambda\in\Lambda_{n,\alpha}} |p_\lambda^{\vec{u},n} - M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}}(\lambda)| \\
&\quad + \sum_{\lambda\notin\Lambda_{n,\alpha}} p_\lambda^{\vec{u},n} + M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}}(\lambda).
\end{aligned}$$

Lemma 9.6.2 and (9.116) imply that for all $\|\vec{u}\| \leq n^\gamma$, and $n > (4/\delta)^{\frac{1-\alpha}{?}}$,

$$\sum_{\lambda\notin\Lambda_{n,\alpha}} p_\lambda^{\vec{u},n} + M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}}(\lambda) \leq C_1 \exp(-(C_2 n^{2\alpha-1})),$$

with C_1 and C_2 depending only on the dimension. We end the proof of (9.119) by recalling that

$$p_\lambda^{\vec{u},n} = \prod_{l=1}^d \frac{\lambda_l! \prod_{k=l+1}^d \lambda_l - \lambda_k + k - l}{(\lambda_l + d - l)!} \sum_{\mathbf{m}\in\lambda} \prod_{i<j} \left(\frac{\mu_j^{\vec{u},n}}{\mu_i^{\vec{u},n}} \right)^{m_{i,j}} M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}}(\lambda).$$

Now, for all $\|\vec{u}\| \leq n^\gamma$ and all $\lambda \in \Lambda_{n,\alpha}$, the right hand side without the multinomial is

$$\prod_{l=1}^d \prod_{k=l+1}^d \frac{n\mu_l - n\mu_k + O(n^\alpha)}{n\mu_l + O(n^\alpha)} \sum_{\mathbf{m}\in\lambda} \prod_{i<j} \left(\frac{\mu_j}{\mu_i} + O(n^{-1/2+\gamma}) \right)^{m_{i,j}}.$$

On $\Lambda_{n,\alpha}$, for $n > (4/\delta)^{\frac{1}{1-\alpha}}$, the cube $[0, n^{1/2}]^{d(d-1)/2} \subset \lambda$, so that

$$\begin{aligned}
\prod_{i<j} \frac{1 - (\frac{\mu_j}{\mu_i} + O(n^{-1/2+\gamma}))^{n^{1/2}}}{1 - \frac{\mu_j}{\mu_i} + O(n^{-1/2+\gamma})} &\leq \sum_{\mathbf{m}\in\lambda} \prod_{i<j} \left(\frac{\mu_j}{\mu_i} + O(n^{-1/2+\gamma}) \right)^{m_{i,j}} \\
&\leq \prod_{i<j} \frac{1}{1 - \frac{\mu_j}{\mu_i} + O(n^{-1/2+\gamma})}.
\end{aligned}$$

Putting together yields

$$\left| \prod_{l=1}^d \frac{\lambda_l! \prod_{k=l+1}^d \lambda_l - \lambda_k + k - l}{(\lambda_l + d - l)!} \sum_{\mathbf{m}\in\lambda} \prod_{i<j} \left(\frac{\mu_j^{\vec{u},n}}{\mu_i^{\vec{u},n}} \right)^{m_{i,j}} - 1 \right| \leq C \frac{n^{-1/2+\gamma} + n^{\alpha-1}}{\delta}.$$

We have thus proved (9.119).

We now turn our attention to the second term of (9.118). Our main tool hereon will be KMT Theorem:

Theorem 9.7.16. *(Komlós et al., 1975; Bretagnolle et Massart, 1989) Let X_i for $i \in \mathbb{N}$ be independent uniform random variables on $[0,1]$. Let F be the repartition function of this law (that is, the function $x \mapsto x$ on $[0,1]$), let F_n be the n-th empirical repartition function $F_n(t) = \frac{1}{n}\sum_{i=1}^{n} \delta_{X_i \leq t}$ and let α_n be the corresponding empirical process $\alpha_n(t) = \sqrt{n}\,(F_n(t) - F(t))$.*

Let B be a brownian bridge, that is a Gaussian stochastic process such that for $0 \leq t \leq u \leq 1$, we have $\mathbb{E}[B(t)] = 0$ and $\mathbb{E}[B(t)B(u)] = t(1-u)$.

Then we may construct these processes on the same probability space such that:

$$\mathbb{P}\left[\sup_{t\in[0,1]} \sqrt{n}\,|\alpha_n(t) - B(t)| > x + c\ln n\right] \leq K\exp(-\lambda x) \tag{9.120}$$

for all n and x, where c, K and λ are absolute positive constants.

We shall take $x = c\ln n$ below.

Now notice that the distribution of the vector $n[F_n(\mu_1^{\vec{u},n}), F_n(\mu_2^{\vec{u},n} + \mu_1^{\vec{u},n}) - F_n(\mu_1^{\vec{u},n}), \ldots, F_n(1) - F_n(1 - \mu_d^{\vec{u},n})]$ is that of the multinomial with parameters n and $\mu^{\vec{u},n}$. Now if we substract to this the vector $n\mu$ and divide by $n^{-1/2}$, as we do in our transforms τ^n and σ^n, we obtain

$$\begin{bmatrix} \alpha_n(\mu_1^{\vec{u},n}) \\ \alpha_n(\mu_2^{\vec{u},n} + \mu_1^{\vec{u},n}) - \alpha_n(\mu_1^{\vec{u},n}) \\ \vdots \\ \alpha_n(1) - \alpha_n(1 - \mu_d^{\vec{u},n}) \end{bmatrix} + \begin{bmatrix} u_1 \\ \vdots \\ u_{d-1} \\ -\sum_2^d u_i \end{bmatrix}. \tag{9.121}$$

The last part of the effect of τ_n is keeping all the components of this vector but the first, and smear out with a $(-n^{1/2}/2, n^{1/2}/2)^{d-1}$ box so that instead of a collection of peaks we have a histogram without holes between the bars.

Let us also define the Gaussian vector

$$B^{\vec{u},n} \hat{=} [B(\mu_1^{\vec{u},n}), B(\mu_2^{\vec{u},n} + \mu_1^{\vec{u},n}) - B(\mu_1^{\vec{u},n}), \ldots, B(1 - \mu_d^{\vec{u},n}) - B(\sum_{i=1}^{d-2} \mu_i^{\vec{u},n})] + [u_1, \ldots, u_{d-1}].$$

Its law is $\mathcal{N}(\vec{u}, V_{\mu^{\vec{u},n}})$, as can be easily shown with the formulas $\mathbb{E}[B(t)] = 0$ and $\mathbb{E}[B(t)B(u)] = t(1-u)$. Recall that $V_{\mu^{\vec{u},n}}$ is given by formula (9.12), with $\mu^{\vec{u},n}$ instead of μ.

To make use of Theorem 9.7.16, we must still smear out our functions. We are writing U^n for the uniform probability on $\left[\frac{f(n)}{\sqrt{n}}, \frac{f(n)}{\sqrt{n}}\right]^{d-1}$ and shall convolve. We choose later the precise $f(n)$.

Then let us write an expression where all the terms of the proof of Lemma 9.6.1 appear:

$$\begin{aligned}\left\|\mathcal{N}(\vec{u}, V_\mu) - \tau^n M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}}\right\|_1 &\leq \left\|\mathcal{N}(\vec{u}, V_\mu) - B^{\vec{u},n}\right\|_1 \\ &+ \left\|B^{\vec{u},n} - B^{\vec{u},n} \star U^n\right\|_1 \\ &+ \left\|B^{\vec{u},n} \star U^n - \tau^n M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}} \star U^n\right\|_1 \\ &+ \left\|\tau^n M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}} \star U^n - \tau^n M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}}\right\|_1 .\end{aligned} \tag{9.122}$$

Let us study the first term. We have already seen that $\left\|\mathcal{N}(\vec{u}, V_\mu) - B^{\vec{u},n}\right\|_1 = \left\|\mathcal{N}(\vec{u}, V_\mu) - \mathcal{N}(\vec{u}, V_{\mu^{\vec{u},n}})\right\|_1$Hence we must bound the distance between two Gaussians with the same mean and different variances. Since $\mu_i^{\vec{u},n} = \mu_i + u_i n^{-1/2}$ and $\|\vec{u}\|_1 \leq n^\gamma$, we have

$$\begin{aligned}\|V_\mu - V_{\mu^{\vec{u},n}}\|_1 &\leq \sum_{k,l} \left|[V_\mu]_{k,l} - [V_{\mu^{\vec{u},n}}]_{k,l}\right| \\ &\leq \sum_{1\leq i,j\leq d-1} |u_i u_j| n^{-1} + 2 * \sum_i |u_i| n^{-1/2} \left|\sum_j \mu_j\right| + \sum_i |u_i| n^{-1/2} \\ &\leq 4n^{-1/2} \sum_i |u_i| \\ &\leq 4n^{\gamma - 1/2}.\end{aligned}$$

On the other hand we can bound from above the smallest eigenvalue of V_μ. Indeed, for all $1 \leq k \leq (d-1)$, we have $[V_\mu]_{k,k} - \sum_{l\neq k} [V_\mu]_{k,l} = \mu_k(1 - \sum_{l=2}^d \mu_l) = \mu_k \mu_1 \geq \delta/d$. Hence $V_\mu \geq (\delta/d)\mathbf{1}$.

So that $\left(1 - Cn^{-1/2+\gamma}/\delta\right) V_\mu \leq V_{\mu^{\vec{u},n}} \leq \left(1 + Cn^{-1/2+\gamma}/\delta\right) V_\mu$, where C depends only on the dimension d. We end the computation of the bound for the first term of

(9.122) with:

$$\begin{aligned}
\|\mathcal{N}(\vec{u}, V_\mu) - \mathcal{N}(\vec{u}, V_{\mu^{\vec{u},n}})\|_1 &= \int \left| \frac{e^{-\frac{1}{2}x^\top V_\mu^{-1} x}}{\sqrt{(2\pi)^{d-1}\det(V_\mu)}} - \frac{e^{-\frac{1}{2}x^\top (V_{\mu^{\vec{u},n}})^{-1} x}}{\sqrt{(2\pi)^{d-1}\det(V_\mu^{\vec{u},n})}} \right| \mathrm{d}x \\
&\leq \int \frac{\exp\left(-\frac{x^\top V_\mu^{-1} x}{2(1+Cn^{-1/2+\gamma}/\delta)}\right)}{\sqrt{(2\pi(1-Cn^{-1/2+\gamma}/\delta))^{d-1}\det(V_\mu)}} \\
&\quad - \frac{\exp\left(-\frac{x^\top V_\mu^{-1} x}{2(1-Cn^{-1/2+\gamma}/\delta)}\right)}{\sqrt{(2\pi(1+Cn^{-1/2+\gamma}/\delta))^{d-1}\det(V_\mu)}} \\
&= \frac{1+Cn^{-1/2+\gamma}/\delta}{1-Cn^{-1/2+\gamma}/\delta} - \frac{1-Cn^{-1/2+\gamma}/\delta}{1+Cn^{-1/2+\gamma}/\delta} \\
&\leq C_2 n^{-1/2+\gamma}/\delta,
\end{aligned}$$

where C_2 still depends only on the dimension, as long as $Cn^{-1/2+\gamma} < \delta/2$.

The second term of (9.122) corresponds to convolving Gaussians with sharper and sharper functions. Now, we may upper bound $\|f \star g\|_1$ by $R \sup_x \|\nabla f(x)\|$ for g a probability density supported on the ball of radius R. So that

$$\left\| B^{\vec{u},n} - B^{\vec{u},n} \star U^n \right\|_1 \leq \frac{Cf(n)}{\delta\sqrt{n}},$$

where C depends only on the dimension, and where we have used $n^{\gamma-1/2} \leq \delta/2$.

The third term is the one where we use KMT theorem. Indeed, for all $\vec{u}$, for any positive x that, for all x, for all $\vec{u} \in \Xi_{n,\beta}$, using as an intermediate step the probability space $(\Omega, \mathcal{A}, q)$ on which α_n and B are built, we may write

$$\begin{aligned}
&\left\| B^{\vec{u},n} \star U^n - \tau^n M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}} \star U^n \right\|_1 \\
&\leq \int_\Omega \|B^{\vec{u},n}(\omega) \star U^n - \tau_n M^n_{\mu_1^{\vec{u},n},\ldots,\mu_d^{\vec{u},n}}(\omega) \star U^n\|_1 \, \mathrm{d}q(\omega) \\
&\leq \mathbb{P}\left[\sup_{t\in[0,1]} |\alpha_n(t) - B(t)| > \frac{x + c\ln n}{\sqrt{n}} \right] + \\
&\qquad \sup_{\|y\|_\infty \leq \frac{x+c\ln n}{\sqrt{n}}} \int_{\mathbb{R}^{d-1}} |U^n(z) - U^n(z+y)| \mathrm{d}z \\
&\leq K\exp(-\lambda x) + \left(1 - \frac{f(n) - x - c\ln n}{f(n)}\right)^{d-1}
\end{aligned}$$

We now tackle the last term of (9.122). We break it in two parts, the first being the large deviations, and the second coming explicitly from the convolution. For any ϵ,

$$\left\| \tau^n M^n_{\mu_1^{\vec{u},n},\dots,\mu_d^{\vec{u},n}} \star U^n - \tau^n M^n_{\mu_1^{\vec{u},n},\dots,\mu_d^{\vec{u},n}} \right\|_1$$
$$\leq 2 \left(\sum_{\lambda \notin \Lambda_{n,1/2+\epsilon}} M^n_{\mu_1^{\vec{u},n},\dots,\mu_d^{\vec{u},n}}(\lambda) + \sup_{\substack{\|x\| \leq n^\epsilon \\ \|x-y\|_\infty \leq f(n)/\sqrt{n}}} \left| \frac{\tau^n M^n_{\mu_1^{\vec{u},n},\dots,\mu_d^{\vec{u},n}}(x)}{\tau^n M^n_{\mu_1^{\vec{u},n},\dots,\mu_d^{\vec{u},n}}(y)} - 1 \right| \right)$$

Now, the second term can be upper bounded by

$$(1+f(n)) \sum_{j=2}^{d} \sup_{\lambda \in \Lambda_{n,1/2+\epsilon}} \left| \frac{M^n_{\mu_1^{\vec{u},n},\dots,\mu_d^{\vec{u},n}}(\lambda_1,\dots,\lambda_j,\dots,\lambda_d)}{M^n_{\mu_1^{\vec{u},n},\dots,\mu_d^{\vec{u},n}}(\lambda_1+1,\dots,\lambda_j-1,\dots,\lambda_d)} - 1 \right|$$
$$\leq (1+f(n)) \sum_{j=2}^{d} \sup_{\lambda \in \Lambda_{n,1/2+\epsilon}} \left| \frac{\lambda_1 \mu_j^{\vec{u},n}}{\lambda_j \mu_1^{\vec{u},n}} - 1 \right|$$
$$\leq (1+f(n)) C n^{-1/2+\epsilon}/\delta,$$

where we have recalled the assumption $n^{\gamma-1/2} \leq \delta/2$, and where C is a constant depending only on the dimension d.

Putting the four losses together and specifying $f(n) = n^{1/4}$ and $x = n^\epsilon$, we end up with

$$\delta(\mathcal{M}_n, \mathcal{G}_n) \leq C(n^{-1/4+\epsilon} + n^{-1/2+\gamma})/\delta$$

for $n^{-1/2+\gamma} > C\delta/2$ and C depending only on the dimension d and the universal constants c, K, λ from Theorem 9.7.16.

Adding the part (9.119), and noticing that $\alpha - 1 > \epsilon - 1/2$ for small enough ϵ, ends the proof of Lemma 9.6.1.

From here, proving Lemma 9.6.8 (that is the inverse direction) is easy enough.

Indeed, remembering that $\sigma^n \tau^n p^{\theta,n} = p^{\theta,n}$ and that σ^n is a contraction, we get

$$\left\| \sigma^n \mathcal{N}(\vec{u}, V_\mu) - p^{\vec{\zeta},\vec{u},n} \right\|_1 = \left\| \sigma^n \mathcal{N}(\vec{u}, V_\mu) - \sigma^n \tau^n p^{\vec{\zeta},\vec{u},n} \right\|_1$$
$$\leq \left\| \mathcal{N}(\vec{u}, V_\mu) - \tau^n p^{\vec{\zeta},\vec{u},n} \right\|_1.$$

So that we have the same speed and conditions as those of Lemma 9.6.1.

□

9.7.7 Proof of Lemma 9.6.3 on convergence to the thermal equilibrium state

We recall that the state ϕ on $CCR(L^2(\rho), \sigma)$ was defined in (9.47) and is the product of a classical Gaussian distribution and $d(d-1)/2$ Gaussian states $\phi_{i,j}$ of quantum harmonic oscillators, one for each pair $i < j$. $\phi_{i,j}$ are thermal equilibrium states with inverse temperature $\beta = \ln(\mu_i/\mu_j)$ (cf. (9.28)). The joint state $\phi^{\vec{0}} := \bigotimes_{i<j} \phi_{i,j}$ is then displaced to obtain $\phi^{\vec{\zeta}}$ but Lemma 9.6.3 is only concerned with $\phi^{\vec{0}}$.

It is well known that thermal equilibrium states are diagonal in the number basis and in our case

$$\phi^{\vec{0}} = \sum_{\mathbf{m} \in \mathbb{N}^{d(d-1)/2}} \prod_{i<j} \frac{\mu_i - \mu_j}{\mu_i} \left(\frac{\mu_j}{\mu_i}\right)^{m_{i,j}} |\mathbf{m}\rangle\langle\mathbf{m}|. \tag{9.123}$$

As shown in (9.60), a similar formula holds for the finite dimensional block states $\rho_\lambda^{\vec{0},\vec{u},n}$:

$$\langle \mathbf{m}, \lambda | \rho_\lambda^{\vec{0},\vec{u},n} | \mathbf{m}, \lambda \rangle = C_\lambda^{\vec{u}} \prod_{i<j}^{d} \left(\frac{\mu_j^{\vec{u},n}}{\mu_i^{\vec{u},n}}\right)^{m_{i,j}}, \tag{9.124}$$

where $C_\lambda^{\vec{u}}$ is a normalisation constant, $\mu_i^{\vec{u},n} = \mu_i + u_i/\sqrt{n}$ for $1 \leq i \leq (d-1)$ and $\mu_d^{\vec{u},n} = \mu_d - (\sum_i u_i)/\sqrt{n}$.

However there is a caveat: although $|\mathbf{m}, \lambda\rangle$ are eigenvectors of $\rho_\lambda^{\vec{0},\vec{u},n}$, they are not orthogonal to each other so we cannot directly use $|\mathbf{m}, \lambda\rangle\langle\mathbf{m}, \lambda|$ as eigenprojectors in the spectral decomposition. However, Lemma 9.5.4 gives us an estimate of the error that we incur by doing just that.

Note first that the eigenvalues of $\rho_\lambda^{\vec{0},\vec{u},n}$ are labelled by the *total* multiplicities m_i of the index i in the semistandard Young tableaux :

$$m_i := \lambda_i - \sum_{j>i} m_{i,j} + \sum_{j<i} m_{j,i}. \tag{9.125}$$

By Lemma 9.7.13 we have that $\langle \mathbf{m}, \lambda | \mathbf{l}, \lambda \rangle = 0$ if $|\mathbf{m}| \neq |\mathbf{l}|$. This allows us to split $\mathcal{H}_\lambda$ into a direct sum of orthogonal subspaces

$$\mathcal{H}_{\lambda,\eta} := \mathrm{Lin}\{|\mathbf{m}, \lambda\rangle : |\mathbf{m}| \leq n^\eta\}, \qquad \text{and} \qquad \mathcal{H}_{\lambda,\eta}^\perp := \mathrm{Lin}\{|\mathbf{l}, \lambda\rangle : |\mathbf{l}| > n^\eta\}.$$

and similarly for the Fock space $\mathcal{F} = \mathcal{F}_\eta \oplus \mathcal{F}_\eta^\perp$. Note the hidden dependence on n in the definition of the subspaces. Asymptotically, the state $\rho_\lambda^{\vec{0},\vec{u},n}$ and $\phi^{\vec{0}}$ concentrate on the 'low excitations' spaces $\mathcal{H}_{\lambda,\eta}$ and $\mathcal{F}_\eta$ with corresponding orthogonal projections $P_{\lambda,\eta}$ and P_η, respectively. More precisely,

$$\begin{aligned}\|T_\lambda(\rho_\lambda^{\vec{0},\vec{u},n}) - \phi^{\vec{0}}\|_1 &= \|T_\lambda(P_{\lambda,\eta}\rho_\lambda^{\vec{0},\vec{u},n}P_{\lambda,\eta}) - P_\eta\phi^{\vec{0}}P_\eta\|_1 + \|T_\lambda(P_{\lambda,\eta}^\perp\rho_\lambda^{\vec{0},\vec{u},n}P_{\lambda,\eta}^\perp) - P_\eta^\perp\phi^{\vec{0}}P_\eta^\perp\|_1 \\ &\leq 2\|T_\lambda(P_{\lambda,\eta}\rho_\lambda^{\vec{0},\vec{u},n}P_{\lambda,\eta}) - P_\eta\phi^{\vec{0}}P_\eta\|_1 + 2\|P_\eta^\perp\phi^{\vec{0}}P_\eta^\perp\|_1. \qquad (9.126)\end{aligned}$$

From the definition 9.32 of $\phi^{\vec{0}}$ and that of thermal states (9.27) we see that the second term on the right side of order $\max_{j<k}(\mu_k/\mu_j)^{n^\eta} = O(\exp(-\delta n^\eta))$. For the rest of the proof we shall deal with the first term on the right side.

Let us denote by $\mathcal{H}(\{m_i\}) = \mathrm{Lin}\{|\mathbf{l} : \lambda\rangle, l_i = m_i\}$ be the eigenspace of $\rho_\lambda^{\vec{0},\vec{u},n}$ and $P(\{m_i\})$ the corresponding eigenprojection . Then

$$\rho_\lambda^{\vec{0},\vec{u},n} = C_\lambda^{\vec{u}} \sum_{\{m_i\}} \prod_{i=1}^d (\mu_i^{\vec{u},n})^{m_i-\lambda_i} P(\{m_i\}).$$

As in Lemma 9.5.4 we have that for $|\mathbf{m}| \leq n^\eta$

$$P(\{m_i\}) = \frac{1}{1 + Cn^{(9\eta-2)/12}\delta^{-1/3}} \sum_{\mathbf{m}:\{m_i\}} |\mathbf{m},\lambda\rangle\langle\mathbf{m},\lambda| + E(\{m_i\})$$

where the sum runs over those $\mathbf{m}$ with total multiplicities $\{m_i\}$. The (positive) reminder has trace norm

$$\mathrm{Tr}(E(\{m_i\})) = O(n^{(9\eta-2)/12}\delta^{-1/3}) \cdot \dim(\mathcal{H}(\{m_i\})).$$

By summing over all $\{m_i\}$ satisfying $|\mathbf{m}| \leq n^\eta$ we get

$$P_{\lambda,\eta}\rho_\lambda^{\vec{0},\vec{u},n}P_{\lambda,\eta} = \frac{1}{1 + Cn^{(9\eta-2)/12}\delta^{-1/3}}\tilde{\rho}_\lambda^{\vec{0},\vec{u},n} + C_\lambda^{\vec{u}} \sum_{\{m_i\}} \prod_{i=1}^d (\mu_i^{\vec{u},n})^{m_i-\lambda_i} E(\{m_i\}),$$

where $\tilde{\rho}_\lambda^{\vec{0},\vec{u},n}$ is the approximate state

$$\tilde{\rho}_\lambda^{\vec{0},\vec{u},n} := C_\lambda^{\vec{u}} \sum_{\{m_i\}} \prod_{i=1}^d (\mu_i^{\vec{u},n})^{m_i-\lambda_i} \sum_{\mathbf{m}:\{m_i\}} |\mathbf{m},\lambda\rangle\langle\mathbf{m},\lambda|.$$

The error term has trace norm of the order

$$O(n^{(9\eta-2)/12}\delta^{-1/3}) \cdot C_\lambda^{\vec{u}} \sum_{\{m_i\}} \prod_{i=1}^d (\mu_i^{\vec{u},n})^{m_i-\lambda_i} \dim(\mathcal{H}(\{m_i\})) = O(n^{(9\eta-2)/12}\delta^{-1/3}),$$

where we have used the normalisation of the block state $\rho_\lambda^{\vec{0},\vec{u},n}$.

In conclusion

$$\|P_{\lambda,\eta}\rho_\lambda^{\vec{0},\vec{u},n}P_{\lambda,\eta} - \tilde{\rho}_\lambda^{\vec{0},\vec{u},n}\|_1 = O(n^{(9\eta-2)/12}\delta^{-1/3}). \tag{9.127}$$

The next step is to show that the block states $\tilde{\rho}_\lambda^{\vec{0},\vec{u},n}$ are mapped by T_λ close to $P_\eta\phi^{\vec{0}}P_\eta$. Using (9.60), we can write

$$T_\lambda(\tilde{\rho}_\lambda^{\vec{0},\vec{u},n}) = C_\lambda^{\vec{u}}\sum_{\mathbf{m}\in\lambda}\prod_{i<j}\left(\frac{\mu_j^{\vec{u},n}}{\mu_i^{\vec{u},n}}\right)^{m_{i,j}} T_\lambda(|\mathbf{m},\lambda\rangle\langle\mathbf{m},\lambda|). \tag{9.128}$$

If $n^{\alpha-1} \leq \delta/2$ and $\alpha > 1/2 > \eta$, we know that all $\mathbf{m}$ such that $|\mathbf{m}| \leq n^\eta$ 'fit into' λ.

Since $\mu_i^{\vec{u},n} = \mu_i + O(n^{-1/2+\gamma})$, when $|\mathbf{m}| \leq n^\eta$,

$$\left(\frac{\mu_j^{\vec{u},n}}{\mu_i^{\vec{u},n}}\right)^{m_{i,j}} = \left(\frac{\mu_j}{\mu_i}\right)^{m_{i,j}}(1 + O(n^{-1/2+\gamma+\eta}/\delta)). \tag{9.129}$$

For the normalisation constant we can write:

$$(C_\lambda^{\vec{u}})^{-1} = \sum_{|\mathbf{m}|\leq n^\eta}\prod_{i<j}\left(\frac{\mu_j^{\vec{u},n}}{\mu_i^{\vec{u},n}}\right)^{m_{i,j}} + \sum_{\mathbf{m}\in\lambda:|\mathbf{m}|\geq n^\eta}\prod_{i<j}\left(\frac{\mu_j^{\vec{u},n}}{\mu_i^{\vec{u},n}}\right)^{m_{i,j}}.$$

If $2dn^{\gamma-1/2} < \delta/2$ then the second part is less than $n^{d^2}(1-\delta/2)^{n^\eta}$ which is negligible compared to the other error terms. Hence:

$$\begin{aligned}(C_\lambda^{\vec{u}})^{-1} &= \sum_{|\mathbf{m}|\leq n^\eta}\prod_{i<j}\left(\frac{\mu_j}{\mu_i}\right)^{m_{i,j}}(1 + O(n^{-1/2+\gamma+\eta}/\delta)) \\ &= \sum_{\mathbf{m}\in\mathbb{N}^{d(d-1)/2}}\prod_{i<j}\left(\frac{\mu_j}{\mu_i}\right)^{m_{i,j}}(1 + O(n^{-1/2+\gamma+\eta}/\delta) \\ &= \prod_{i<j}\frac{\mu_i}{\mu_i - \mu_j}(1 + O(n^{-1/2+\gamma+\eta}/\delta)).\end{aligned} \tag{9.130}$$

We then recall that for unit vectors, we have $\||f\rangle\langle f| - |f'\rangle\langle f'|\|_1 = 2\sqrt{1-|\langle f|f'\rangle|^2}$. So that, using Lemma 9.5.4, we get that for $|\mathbf{m}| \leq n^\eta$

$$\|T_\lambda(|\mathbf{m},\lambda\rangle\langle\mathbf{m},\lambda|) - |\mathbf{m}\rangle\langle\mathbf{m}|\|_1 = \|V_\lambda|\mathbf{m},\lambda\rangle\langle\mathbf{m},\lambda|V_\lambda^* - |\mathbf{m}\rangle\langle\mathbf{m}|\|_1 = O(n^{(9\eta-2)/24}/\delta^{1/6}). \tag{9.131}$$

Putting the estimates (9.129), (9.130), (9.131) back into formula (9.128), we obtain $T_\lambda(\tilde{\rho}_\lambda^{\vec{0},\vec{u},n})$, so that

$$T_\lambda(\tilde{\rho}_\lambda^{\vec{0},\vec{u},n}) = \sum_{|\mathbf{m}|\leq n^\eta} \prod_{i<j} \frac{\mu_i - \mu_j}{\mu_i} \left(\frac{\mu_j}{\mu_i}\right)^{m_{i,j}} |\mathbf{m}\rangle\langle\mathbf{m}| + O(n^{-1/2+\gamma+\eta}/\delta, n^{(9\eta-2)/24}/\delta^{1/6}). \tag{9.132}$$

Comparing with (9.123), and using (9.126) and (9.127) we get the desired result.

□

9.7.8 Proof of Lemma 9.6.5 on local linearity of $SU(d)$

The key is to notice that, as we are dealing with a group, there is a r such that

$$U^{-1}(\vec{\zeta}+\vec{z},\vec{0},n)U(\vec{\zeta},\vec{0},n)U(\vec{z},\vec{0},n) = U(-\vec{\zeta}-\vec{z},\vec{0},n)U(\vec{\zeta},\vec{0},n)U(\vec{z},\vec{0},n) = U(\vec{r},\vec{s},n),$$

and similarly for the operation Δ. We shall prove below that under the condition that both $\vec{\zeta}$ and $\vec{z}$ are smaller than n^β, then $\|\vec{r}\| + \|\vec{s}\| = O(n^{-1/2+2\beta}/\delta)$. Let us call this the *domination hypothesis* for further reference.

Now, as the actions are unitary, we may rewrite the norm in Lemma as 9.6.5:

$$\begin{aligned} A &= \left\| [\Delta_\lambda^{\vec{\zeta}+\vec{z},n} - \Delta_\lambda^{\vec{\zeta},n}\Delta_\lambda^{\vec{z},n}](|\mathbf{0},\lambda\rangle\langle\mathbf{0},\lambda|) \right\|_1 \\ &= \left\| \Delta_\lambda^{-(\vec{\zeta}+\vec{z}),n}[\Delta_\lambda^{\vec{\zeta}+\vec{z},n} - \Delta_\lambda^{\vec{\zeta},n}\Delta_\lambda^{\vec{z},n}](|\mathbf{0},\lambda\rangle\langle\mathbf{0},\lambda|) \right\|_1 \\ &= \left\| [\mathrm{Id} - \Delta_\lambda^{\vec{r},\vec{s},n}](|\mathbf{0},\lambda\rangle\langle\mathbf{0},\lambda|) \right\|_1 . \end{aligned}$$

As T_λ is an isometry, we may also let it act the left and T_λ^* on the right and get:

$$\begin{aligned} A &= \left\| T_\lambda(|\mathbf{0},\lambda\rangle\langle\mathbf{0},\lambda|) - T_\lambda\Delta_\lambda^{\vec{r},\vec{s},n}T_\lambda^*(|\mathbf{0}\rangle\langle\mathbf{0}|) \right\|_1 \\ &\leq \| |\mathbf{0}\rangle\langle\mathbf{0}| - |\vec{r}\rangle\langle\vec{r}| \|_1 + \left\| |\vec{r}\rangle\langle\vec{r}| - T_\lambda\Delta_\lambda^{\vec{r},\vec{s},n}T_\lambda^*(|\mathbf{0}\rangle\langle\mathbf{0}|) \right\|_1 + \| T_\lambda(|\mathbf{0},\lambda\rangle\langle\mathbf{0},\lambda|) - |\mathbf{0}\rangle\langle\mathbf{0}| \|_1. \end{aligned}$$

By the domination hypothesis, the norm of $\vec{r}$ is smaller than $n^{-1/2+2\beta}/\delta$, hence $\langle\vec{r}|\mathbf{0}\rangle = 1 - O(n^{-1+4\beta}/\delta^2)$. Using $\| |f\rangle\langle f| - |f'\rangle\langle f'| \|_1 = 2\sqrt{1-|\langle f|f'\rangle|^2}$ we get that the first term on the right side of the inequality is $O(n^{-1/2+2\beta}/\delta)$. Notice that this is dominated by $R(n)$ given in equation (9.75) since $\eta > 2\beta$.

For the second term, we apply Lemma 9.6.4, with $\vec{z} = 0$. By the domination hypothesis, $\|\vec{s}\| \leq n^{-1/2+2\beta}/\delta$, so we may apply Lemma 9.6.4, and the remainder is given by $R(n)$ in equation (9.75).

The last term is $O(n^{(9\eta-2)/24}/\delta^{1/6})$ as shown in (9.131) which is dominated by $R(n)$.

We finish the proof of the lemma, and simultaneously that of Theorem 9.4.3, by proving the domination hypothesis. Recall that an arbitrary element in $SU(d)$ can be written in the exponential form

$$U(\vec{r},\vec{s}) := \exp\left[i\left(\sum_{i=1}^{d-1} s_i H_i + \sum_{1\le j<k\le d} \frac{\mathrm{Re}(r_{j,k})T_{j,k} + \mathrm{Im}(r_{j,k})T_{k,j}}{\sqrt{\mu_j-\mu_k}}\right)\right],$$

where $(\vec{r},\vec{s}) \in \mathbb{C}^{d(d-1)/2} \times \mathbb{R}^{d-1}$, and $T_{i,j}$, H_i are the generators of $SU(d)$ defined in (9.84). A special case of this is $U(\vec{r}) := U(\vec{r},\vec{0})$. In general, the map $(\vec{r},\vec{s}) \mapsto U(\vec{r},\vec{s})$ is not injective but becomes so if we restrict to a small enough neighbourhood $\mathcal{C}$ of the origin $(0,0) \in \mathbb{C}^{d(d-1)/2} \times \mathbb{R}^{d-1}$. On this neighbourhood it makes sense to define the inverse as a sort of 'logarithm'

$$\log U(\vec{r},\vec{s}) := (\vec{r},\vec{s}),$$

which is a C^∞ function.

By continuity of the product, if $\vec{x},\vec{y} \in \mathbb{C}^{d(d-1)/2}$ are small enough, then $U(-\vec{x}-\vec{y})U(\vec{x})U(\vec{y}) \in \mathcal{C}$. Since $\|\vec{\zeta}\| + \|\vec{z}\|/\sqrt{n} \le n^{\beta-1/2}/\delta$, we can apply this to $\vec{x} = \vec{\zeta}/\sqrt{n}, \vec{y} = \vec{z}/\sqrt{n}$ for $n > (C/\delta)^{\frac{1}{1/2-\beta}}$ with the constant C depending only on the dimension, and get

$$(\vec{r}/\sqrt{n},\vec{s}/\sqrt{n}) = f(\vec{\zeta}/\sqrt{n},\vec{z}/\sqrt{n}) := \log\left[U(-(\vec{\zeta}+\vec{z})/\sqrt{n})U(\vec{\zeta}/\sqrt{n})U(\vec{z}/\sqrt{n})\right].$$

Since f is a C^∞ function we can expand in Taylor series and it is easy to show that $f(\vec{0},\vec{0}) = (\vec{0},\vec{0})$, the first order partial derivatives are zero as well, and the second order derivatives are uniformly bounded in a neighbourhood of the origin. Thus we get

$$\vec{r} = \sqrt{n}\,O\left(\frac{\|z_{i,j}\|^2}{n(\mu_i-\mu_j)}, \frac{\|\zeta_{i,j}\|^2}{n(\mu_i-\mu_j)}\right) = O(n^{-1/2+2\beta}/\delta).$$

□

Bibliographie

Certaines des entrées BibTeX viennent de Citebase, ou de SAO/NASA Astrophysics Data System.

L. Accardi. Some trends and problems in quantum probability. In A. Frigerio L. Accardi et V. Gorini, editors, *Quantum probability and applications to the quantum theory of irreversible processes*, volume 1055 of *Lecture Notes in Mathematics, Berlin Springer Verlag*, pages 1–19, 1984.

L. Accardi et Bach, A. Central limits of squeezing operators. In Luigi Accardi et Wilhelm von Wandelfels, editors, *Quantum Probability and applications IV*, volume 1396 of *Lecture notes in mathematics*, pages 7–19. Springer, 1987.

L. Accardi et Bach, A. Quantum central limit theorem for strongly mixing random variables. *Z. W.*, pages 393–402, 1985.

A. Acin, E. Jane, et G. Vidal. Optimal estimation of quantum dynamics. *Physical Review A*, 64 :050302, 2001.

A. Acin, E. Bagan, M. Baig, Ll Masanes, et R. Munoz-Tapia. Multiple copy 2-state discrimination with individual measurements. *Physical Review A*, 71 :032338, 2005.

S. Amari. *Differential-geometrical methods in statistics.* Lecture notes in statistics. Springer Verlag, Berlin, 1985.

Erika Andersson, Stephen M. Barnett, Claire R. Gilson, et Kieran Hunter. Minimum-error discrimination between three mirror-symmetric states. *Physical Review A*, 65 :052308, 2002.

M. A. Armen, J. K. Au, J. K. Stockton, A. C. Doherty, et H. Mabuchi. Adaptive Homodyne Measurement of Optical Phase. *Phys. Rev. Lett.*, 89 :133602, 2002.

L.M. Artiles, R. Gill, et M. Guţă. An invitation to quantum tomography. *J. Royal Statist. Soc. B (Methodological)*, 67 :109–134, 2005.

Artiles, L, Gill, R., et Guţă, M. An invitation to quantum tomography. *J. Royal Statist. Soc. B (Methodological)*, 67 :109–134, 2005.

W.B. Arveson. On subalgebras of C^*-algebras. *Acta Mathematica*, 123 :141–224, 1969.

K. M. R. Audenaert, M. Nussbaum, A. Szkola, et F. Verstraete. Asymptotic Error Rates in Quantum Hypothesis Testing. arXiv :0708.4282[quant-ph].

K. M. R. Audenaert, M. Nussbaum, A. Szkola, et F. Verstraete. Asymptotic error rates in quantum hypothesis testing, 2007.

M. Audin. *Geometry.* Springer Verlag, Berlin, 2002.

E. Bagan, M. Baig, et R. Munoz-Tapia. Optimal scheme for estimating a pure qubit state via local measurements. *Phys. Rev. Lett.*, 89 :277904, 2002.

E Bagan, M Baig, et R Munoz-Tapia. Entanglement assisted alignment of reference frames using a dense covariant coding. *Physical Review A*, 69 :050303, 2004a.

E. Bagan, M. Baig, et R. Munoz-Tapia. Quantum reverse-engineering and reference frame alignment without non-local correlations. *Physical Review A*, 70 :030301, 2004b.

E. Bagan, M. Baig, R. Munoz-Tapia, et A. Rodriguez. Collective versus local measurements in a qubit mixed-state estimation. *Phys. Rev. A*, 69 :010304(R), 2004c.

E. Bagan, A. Monras, et R. Munoz-Tapia. Comprehensive analysis of quantum pure-state estimation for two-level system. *Phys. Rev. A*, 71 :062318, 2005.

E. Bagan, M. A. Ballester, R. D. Gill, A. Monras, et R. Munoz-Tapia. Optimal full estimation of qubit mixed states. *Physical Review A*, 73 :032301, 2006.

M. A. Ballester. *Estimation of Quantum States and Operations.* PhD thesis, Universiteit Utrecht, 2005a.

Manuel A. Ballester. Estimation of $SU(d)$ using entanglement. *Preprint*, 2005b. URL `http ://www.arxiv.org/abs/quant-ph/0507073`.

M. Ban, K. Kurokawa, R. Momose, et O. Hirota. Optimum measurements for discrimination among symmetric quantum states and parameter estimation. *Int. J. Theor. Phys.*, 36 :1269 – 1288, 1997.

K. Banaszek, D'Ariano, G. M., Paris, M. G. A., et Sacchi, M. F. Maximum-likelihood estimation of the density matrix. *Phys. Rev. A*, 61 :R010304, 1999.

Somshubhro Bandyopadhyay, P. Oscar Boykin, Vwani P. Roychowdhury, et Farrokh Vatan. A new proof for the existence of mutually unbiased bases. *Algorithmica*, 34(4) :512–528, 2002.

O. E. Barndorff-Nielsen et Gill, R. Fisher information in quantum statistics. *J. Phys. A*, 33 :1–10, 2000.

O. E. Barndorff-Nielsen, Gill, R., et Jupp, P. E. On quantum statistical inference (with discussion). *J. R. Statist. Soc. B*, 65 :775–816, 2003.

Stephen M. Barnett. Minimum-error discrimination between multiply symmetric states. *Phys. Rev. A*, 64(3) :030303, Aug 2001.

Stephen D. Bartlett, Terry Rudolph, et R. W. Spekkens. Classical and quantum communication without a shared reference frame. *Physical Review Letters*, 91 : 027901, 2003.

V. P. Belavkin. Generalized heisenberg uncertainty relations, and efficient measurements in quantum systems. *Theor. Math. Phys.*, 26 :213–222, 1976.

V. P. Belavkin. Optimal multiple quantum statistical hypothesis testing. *Stochastics*, 1 :315–345, 1975.

Viacheslav P. Belavkin, Giacomo Mauro D'Ariano, et Maxim Raginsky. Operational distance and fidelity for quantum channels. *Journal of Mathematical Physics*, 46 : 062106, 2005.

Charles H. Bennett, Gilles Brassard, et N. David Mermin. Quantum cryptography without bell's theorem. *Phys. Rev. Lett.*, 68(5) :557–559, Feb 1992.

Charles H. Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, et William K. Wootters. Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels. *Phys. Rev. Lett.*, 70(13) :1895–1899, Mar 1993.

J. A. Bergou, U. Herzog, et M. Hillery. Discrimination of Quantum States. In M. G. A. Paris et J. Řeháček, editors, *Quantum State Estimation*, volume 649 of *Lecture Notes in Physics, Berlin Springer Verlag*, pages 417–465, 2004.

S.N. Bernstein. On a modification of Chebyshev's inequality and of the error formula of Laplace. In *Collected works*, volume 4, 1964.

J. V. Bondar et P. Milnes. Amenability : A survey for statistical applications of hunt-stein and related conditions on groups. *Z. Wahrscheinlichkeitstheorie*, 57 : 103 – 128, 1981.

L. Bouten, Guţă, M., et Maassen, H. Stochastic schrödinger equations. *Jourrnal of Physics A*, 37 :3189–3209, 2004.

Luc Bouten, Ramon van Handel, et Matthew James. An introduction to quantum filtering, 2006. URL `http ://arxiv.org/abs/math/0601741`.

Stephen Boyd et Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004. ISBN 0521833787.

S. L. Braunstein et Caves C. M. Statistical distance and the geometry of quantum states. *Phys. Rev. Lett.*, 72 :3439–3443, 1994.

J. Bretagnolle et P. Massart. Hungarian constructions from the nonasymptotic view point. *Ann. Probab.*, 17(1) :239–256, 1989.

Dagmar Bruss. Optimal eavesdropping in quantum cryptography with six states. *Physical Review Letters*, 81 :3018, 1998.

F. Buscemi, G.M. d'Ariano, M. Keyl, P. Perinotti, et R. Werner. Clean positive operator valued measures. *J. Math. Phys*, 46 :082109, 2005.

P. Busch et P. J. Lahti. The determination of the past and the future of a physical system in quantum mechanics. *Foundations of Physics*, 19 :633–678, June 1989.

C. Butucea, M. Gu¸tă, et L. Artiles. Minimax and adaptive estimation of the Wigner function in quantum homodyne tomography with noisy data. *Annals of Statistics*, 35(2) :465–494, 2007.

V. Buzek, R. Derka, et S. Massar. Optimal quantum clocks. *Physical Review Letters*, 82 :2207, 1999.

L. Cavalier et J.-Y. Koo. Poisson intensity estimation for tomographic data using a wavelet shrinkage approach. *IEEE Trans. on Information Theory*, 48 :2794–2802, 2002.

C. M. Caves. Quantum limits on noise in linear amplifiers. *Phys. Rev. D*, 26 : 1817–1839, 1982.

A. Chefles. Quantum state discrimination. *Contemporary Physics*, 41 :401–424, June 2000.

A. Chefles et S. M. Barnett. Entanglement and unambiguous discrimination between non-orthogonal states. *Physics Letters A*, 236 :177–179, February 1997.

Anthony Chefles et Stephen M. Barnett. Optimum unambiguous discrimination between linearly independent symmetric states. *Physics Letters A*, 250 :223, 1998a.

Anthony Chefles et Stephen M. Barnett. Quantum state separation, unambiguous discrimination and exact cloning. *J.PHYS.A*, 31 :10097, 1998b.

Anthony Chefles, Richard Jozsa, et Andreas Winter. On the existence of physical transformations between sets of quantum states, 2003.

Anthony Chefles, Akira Kitagawa, Masahiro Takeoka, Masahide Sasaki, et Jason Twamley. Unambiguous discrimination among oracle operators, 2007.

Andrew M. Childs, John Preskill, et Joseph Renes. Quantum information and precision measurement. *Journal of Modern Optics*, 47 :155, 2000a.

Andrew M. Childs, John Preskill, et Joseph Renes. Quantum information and precision measurement. *Journal of Modern Optics*, 47 :155, 2000b.

G Chiribella, G M D'Ariano, P Perinotti, et M F Sacchi. Efficient use of quantum resources for the transmission of a reference frame. *Physical Review Letters*, 93 : 180503, 2004.

G. Chiribella, G. M. D'Ariano, et M. F. Sacchi. Optimal estimation of group transformations using entanglement. *Physical Review A*, 72 :042338, 2005.

Chih-Lung Chou et Li-Yi Hsu. Minimum-error discrimination between symmetric mixed quantum states. *Physical Review A*, 68 :042305, 2003.

J. I. Cirac, A. K. Ekert, et C. Macchiavello. Optimal purification of single qubits. *Phys. Rev. Lett.*, 82 :4344, 1999.

Roger B. M. Clarke, Anthony Chefles, Stephen M. Barnett, et Erling Riis. Experimental demonstration of optimal unambiguous state discrimination. *Phys. Rev. A*, 63(4) :040305, Mar 2001a.

Roger B. M. Clarke, Vivien M. Kendon, Anthony Chefles, Stephen M. Barnett, Erling Riis, et Masahide Sasaki. Experimental realization of optimal detection strategies for overcomplete states. *Physical Review A*, 64 :012303, 2001b.

C.D. Cushen et R.L. Hudson. A quantum-mechanical central limit theorem. *J. Appl. Prob.*, 8 :454–469, 1971.

Sonja Daffer et Peter L. Knight. Generating optimal states for a homodyne bell test. *Physical Review A*, 72 :032509, 2005.

Domenico D'Alessandro et Francesca Albertini. Quantum symmetries and cartan decompositions in arbitrary dimensions, 2005.

D. A. R. Dalvit, R. L. de Matos Filho, et F. Toscano. Quantum metrology at the heisenberg limit with ion traps. *New Journal of Physics*, 8 :276, 2006.

G. M. D'Ariano, Macchiavello, C., et Paris, M. G. A. Detection of the density matrix through optical homodyne tomography without filtered back projection. *Phys. Rev. A*, 50 :4298–4302, 1994.

G. M. D'Ariano, Leonhardt, U., et Paul, H. Homodyne detection of the density matrix of the radiation field. *Phys. Rev. A*, 52 :R1801–R1804, 1995.

G. M. D'Ariano, M. F. Sacchi, et J. Kahn. Minimax quantum state discrimination. *Phys. Rev. A*, 72 :032310, 2005a. URL `arXiv :quant-ph/0504048`.

G. M. D'Ariano, M. F. Sacchi, et J. Kahn. Minimax discrimination of two Pauli channels. *Phys. Rev. A*, 72 :052302, 2005b. URL `arXiv :quant-ph/0507081`.

Giacomo Mauro D'Ariano, Lorenzo Maccone, et Paoloplacido Lo Presti. Quantum calibration of measuring apparatuses. *Phys. Rev. Lett.*, 93 :250407, 2004. URL `http ://arXiv.org :quant-ph/0408116`.

E.B. Davies. On the repeated measurements of continuous observables in quantum mechanics. *J. Functional Analysis*, 6 :318–346, 1970.

S. R. Deans. *The Radon transform and some of its applications.* John Wiley & Sons, New York, 1983.

D. Dieks. Overlap and distinguishability of quantum states. *Physics Letters A*, 126 : 303–306, January 1988.

Lu-Ming Duan et Guang-Can Guo. Probabilistic cloning and identification of linearly independent quantum states. *Phys. Rev. Lett.*, 80(22) :4999–5002, Jun 1998.

F. J. Dyson. General theory of spin-wave interactions. *Phys. Rev.*, 102 :1217–1230, 1956.

H. S. Eisenberg, J. F. Hodelin, G. Khoury, et D. Bouwmeester. Multiphoton path entanglement by nonlocal bunching. *Physical Review Letters*, 94(9) :090502, 2005.

Artur K. Ekert. Quantum cryptography based on bell's theorem. *Phys. Rev. Lett.*, 67(6) :661–663, Aug 1991.

Y. C. Eldar. von Neumann measurement is optimal for detecting linearly independent mixed quantum states. *Physical Review A*, 68(5) :052303–+, November 2003.

Y. C. Eldar, A. Megretski, et G. C. Verghese. Optimal Detection of Symmetric Mixed Quantum States. *IEEE Transactions on Information Theory*, 50(6) :1198 – 1207, 2004.

Yonina C. Eldar. A semidefinite programming approach to optimal unambiguous discrimination of quantum states. *IEEE Transactions on Information Theory*, 49 :446, 2003.

F Embacher et H. Narnhofer. Strategies to measure a quantum state. *Ann. of Phys. (N.Y.)*, 311 :220, 2004.

Erdélyi. *Higher Transcendental Functions*, volume 2. McGraw-Hill, 1953.

Yuan Feng, Runyao Duan, et Zhengfeng Ji. Condition and capability of quantum state separation. *Physical Review A*, 72 :012313, 2005.

D G. Fisher, S. H. Kienle, et M. Freyberger. Quantum-state estimation by self-learning measurements. *Phys. Rev. A*, 61 :032306, 2000.

Jaromir Fiurasek et Miroslav Jezek. Optimal discrimination of mixed quantum states involving inconclusive results. *Physical Review A*, 67 :012321, 2003.

A. Fujiwara. Strong consistency and asymptotic efficiency for adaptive quantum estimation problems. *J. Phys. A*, 39 :12489–12504, 2006.

A Fujiwara et H Imai. Quantum parameter estimation of a generalized pauli channel. *Journal of Physics A : Mathematical and General*, 36(29) :8093–8103, 2003. URL `http ://stacks.iop.org/0305-4470/36/8093`.

A. Fujiwara et Nagaoka, H. Quantum fisher information and estimation for pure state models. *Phys. Lett A*, 201 :119–124, 1995.

Akio Fujiwara. Estimation of su(2) operation and dense coding : An information geometric approach. *Phys. Rev. A*, 65(1) :012316, 2001.

W. Fulton. *Young tableaux, with Applications to Representation Theory and Geometry.* Cambridge University Press, 1997.

W. Fulton et J. Harris. *Representation Theory : A First Course.* Springer Verlag, Berlin, 1991.

C. W. Gardiner et P. Zoller. *Quantum Noise.* Springer, 2004.

JM Geremia, J. K. Stockton, et H. Mabuchi. Real-time quantum feedback control of atomic spin-squeezing. *Science*, 304 :270–273, 2004.

Alexei Gilchrist, Nathan K. Langford, et Michael A. Nielsen. Distance measures to compare real and ideal quantum processes, 2004.

R. Gill. *Quantum Asymptotics*, volume 36 of *Lecture Notes-Monograph Series*, pages 255–285. IMS, 2001.

R. D. Gill. Asymptotic information bounds in quantum statistics. quant-ph/0512443, to appear in Annals of Statistics, 2005a.

R. D. Gill et S. Massar. State estimation for large ensembles. *Phys. Rev. A*, 61 : 042312, 2000.

Richard D. Gill. Asymptotic information bounds in quantum statistics. math.ST/0512443, 2005b.

Vittorio Giovannetti, Seth Lloyd, et Lorenzo Maccone. Quantum-enhanced measurements : beating the standard quantum limit. *Science*, 306 :1330, 2004.

Goodman R. et Wallach N.R. *Representations and invariants of the classical groups.* Cambridge University Press, 1998.

Lov K. Grover. A fast quantum mechanical algorithm for database search. In *STOC '96 : Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219, New York, NY, USA, 1996. ACM Press. ISBN 0-89791-785-5.

M. Guţă et A. Jenčová. Local asymptotic normality in quantum statistics. *Communications in Mathematical Physics*, 276(2) :341 – 379, 2007.

M. Guţă et J. Kahn. Local asymptotic normality for qubit states. *Phys. Rev. A*, 73 :052108, 2006. URL `arXiv :quant-ph/0512075`.

M. Guţă et J. Kahn. Local asymptotic normality for finite-dimensional systems. *Comm. Math. Phys.*, pages 79–+, March 2009. doi : 10.1007/s00220-009-0787-3. URL `arXiv :0804.3876`.

M. Guţă, B. Janssens, et J. Kahn. Optimal estimation of qubit states with continuous time measurements. *Comm. Math. Phy.*, 277(1) :127 – 160, 2008. URL `arXiv :quant-ph/0608074`.

M. Guţă. Quantum decision theory and comparison of quantum statistical experiments. in preparation.

T. Hannemann, D. Reiss, C. Balzer, W. Neuhauser, P. E. Toschek, et C. Wunderlich. Self-learning estimation of quantum states. *Phys. Rev. A*, 65 :050303–+, 2002a.

Th. Hannemann, D. Reiss, Ch. Balzer, W. Neuhauser, P. E. Toschek, et Ch. Wunderlich. Self-learning estimation of quantum states. *Phys. Rev. A*, 65 :050303(R), 2002b.

M. Hayashi. Two quantum analogues of fisher information from a large deviation viewpoint of quantum estimation. quant-ph/0202003, 2002a.

M. Hayashi. presentations at maphysto and quantop workshop on quantum measurements and quantum stochastics, aarhus, 2003, and special week on quantum statistics, isaac newton institute for mathematical sciences, cambridge, 2004.

M. Hayashi. Quantum estimation and the quantum central limit theorem. *Bulletin of the Mathematical Society of Japan*, 55 :368–391, 2003. (in Japanese ; Translated into English in quant-ph/0608198).

M. Hayashi. A linear programming approach to attainable cramér-rao type bound. In *Asymptotic theory of quantum statistical inference, Selected Papers*, 2005a.

M. Hayashi. Parallel treatment of estimation of su(2) and phase estimation. quant-ph/0407053, 2004.

M. Hayashi et K. Matsumoto. Asymptotic performance of optimal state estimation in quantum two level system. quant-ph/0411073, 2004.

M. Hayashi et K. Matsumoto. Statistical model with measurement degree of freedom and quantum physics. In Masahito Hayashi, editor, *Asymptotic theory of quantum statistical inference : selected papers*, pages 162–170. World Scientific, 2005. (English translation of a paper in Japanese published in Surikaiseki Kenkyusho Kokyuroku, vol. 35, pp. 7689-7727, 2002.).

Masahito Hayashi. *Quantum Information.* Springer-Verlag, Berlin Heidelberg, 2006.

Masahito Hayashi, editor. *Asymptotic theory of quantum statistical inference : selected papers.* World Scientific, 2005b.

Masahito Hayashi. Optimal sequence of quantum measurements in the sense of stein's lemma in quantum hypothesis testing. *MATHEMATICAL AND GENERAL*, 35 :10759, 2002b.

T. Heinonen. Optimal measurements in quantum mechanics. *Physics Letters A*, 346 :77, 2005.

C. W. Helstrom. *Quantum Detection and Estimation Theory.* Academic Press, New York, 1976.

C. W. Helstrom. Quantum detection and estimation theory. *Journal of Statistical Physics*, 1 :231–252, 1969.

U. Herzog. Optimum unambiguous discrimination of two mixed states and application to a class of similar states. *Physical Review A*, 75 :052309, 2007.

Ulrike Herzog et János A. Bergou. Minimum-error discrimination between subsets of linearly dependent quantum states. *Phys. Rev. A*, 65(5) :050305, May 2002.

Ulrike Herzog et Janos A. Bergou. Optimum unambiguous discrimination of two mixed quantum states. *Physical Review A*, 71 :050301, 2005.

W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58 :13–30, 1964.

A. S. Holevo. *Probabilistic and Statistical Aspects of Quantum Theory.* North-Holland, 1982.

A. S. Holevo. Statistical decisions in quantum theory. *Journal of Multivariate Analysis*, 3(4) :337– 394, 1973.

R. Holtz et J. Hanus. On coherent spin states. *J. Phys. A*, 7 :37, 1974.

M. Hübner. Explicit computation of the bures distance for density matrices. *Phys. Lett. A*, 163 :239–242, 1992.

R. L. Hudson et K. R. Parthasarathy. Quantum itô's formula and stochastic evolutions. *Commun. Math. Phys.*, 93 :301–323, 1984.

B. Huttner, A. Muller, J. D. Gautier, H. Zbinden, et N. Gisin. Unambiguous quantum measurement of nonorthogonal states. *Physical Review A*, 54 :3783–3789, November 1996.

Hiroshi Imai et Akio Fujiwara. Geometry of optimal estimation scheme for su(d) channels. *Journal of Physics A : Mathematical and Theoretical*, 40(16) :4391–4400, 2007. URL `http ://stacks.iop.org/1751-8121/40/4391`.

I. D. Ivanovic. How to differentiate between non-orthogonal states. *Physics Letters A*, 123 :257–259, August 1987.

I.D. Ivanovic. Geometrical description of quantum state determination. *Journal of Physics A*, 14 :3241–3245, 1981.

G. Jaeger et A. Shimony. Optimal distinction between two non-orthogonal quantum states. *Physics Letters A*, 197 :83–87, February 1995.

B. Janssens. Unifying decoherence and the heisenberg principle. arxiv.org/abs/quant-ph/ 0606093, 2006.

H. Jeffreys. An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London. Series A*, 186(1007) :453–461, 1946.

M. Jezek, J. Rehacek, et J. Fiurasek. Finding optimal strategies for minimum-error quantum-state discrimination, 2002.

Zhengfeng Ji, Hongen Cao, et Mingsheng Ying. Optimal conclusive discrimination of two states can be achieved locally. *Physical Review A*, 71 :032323, 2005.

Zhengfeng Ji, Guoming Wang, Runyao Duan, Yuan Feng, et Mingsheng Ying. Parameter estimation of quantum channels, 2006.

K. R. Jones. Fundamental limits upon the measurement of state vectors. *Phys. Rev. A*, 50 :3682, 1994.

J. Kahn. Sélection de modèles en tomographie quantique. Master's thesis, École Normale Supérieure, Université Paris-Sud, 2004.

J. Kahn. Clean positive operator valued measures for qubits and similar cases. *J. Phys. A, Math. Theor.*, 40 :4817–4832, 2007a. URL `arXiv :quant-ph/0603117`.

J. Kahn. Fast rate estimation of unitary operations in $SU(d)$. *Phys. Rev. A*, 75 : 022326, 2007b. URL `arXiv :quant-ph/0603115`.

J. Kahn. Model selection for quantum homodyne tomography. URL `arXiv :0712.2912`. Accepté par ESAIM : Probability and Statistics.

J. Kahn et D Petz. Complementary reductions for two qubits. *J. Math. Phy.*, 48 : 012107, 2007. URL `arXiv :quant-ph/0608227`.

Vladislav Kargin. On the chernoff bound for efficiency of quantum hypothesis testing. *ANNALS OF STATISTICS*, 33 :959, 2005.

M. Keyl et R. F. Werner. Estimating the spectrum of a density operator. *Phys. Rev. A*, 64 :052311, 2001.

Gen Kimura, Hajime Tanaka, et Masanao Ozawa. Solution to the mean king's problem with mutually unbiased bases for arbitrary levels. *Physical Review A*, 73 :050301, 2006.

Masahiro Kitagawa et Masahito Ueda. Squeezed spin states. *Phys. Rev. A*, 47(6) : 5138–5143, Jun 1993.

E. Knill, R. Laflamme, A. Ashikhmin, H. Barnum, L. Viola, et W. H. Zurek. Introduction to quantum error correction, 2002.

J. Komlós, P. Major, et G. Tusnády. An approximation of partial sums of independent rv-s, and the sample df. *Z. Warscheinlichkeitstheorie Verwandte*, 32 : 111–131, 1975.

K. Kraus. Complementary observables and uncertainty relations. *Phys. Rev. D*, 35 (10) :3070–3075, May 1987.

K. Kraus. *States, effects and operations.* Springer Verlag, Berlin, 1983.

J. I. Latorre, P. Pascual, et R. Tarrach. Minimal optimal generalized quantum measurements. *Phys. Rev. Lett.*, 81 :1351, 1998.

L. Le Cam. *Asymptotic Methods in Statistical Decision Theory.* Springer Verlag, New York, 1986.

L. Le Cam. Sufficiency and approximate sufficiency. *The Annals of Mathematical Statistics*, 35(4) :1419–1455, 1964.

Lucien Le Cam. Locally asymptotically normal families of distributions. Certain approximations to families of distributions and their use in the theory of estimation and testing hypotheses. *Univ. california Publ. Statist.*, 3 :37–98, 1960.

U. Leonhardt. *Measuring the Quantum State of Light.* Cambridge University Press, 1997.

U. Leonhardt, Paul, H., et D'Ariano, G. M. Tomographic reconstruction of the density matrix via pattern functions. *Phys. Rev. A*, 52 :4899–4907, 1995.

U. Leonhardt, M. Munroe, T. Kiss, Th. Richter, et M. G. Raymer. Sampling of photon statistics and density matrix using homodyne detection. *Optics Communications*, 127 :144–160, 1996.

A. I. Lvovsky et M. G. Raymer. Continuous-variable optical quantum state tomography, 2005. URL `arXiv.org :quant-ph/0511044`.

H. Mack, D. G. Fischer, et M. Freyberger. Enhanced quantum estimation via purification. *Phys. Rev. A*, 62 :042301, 2000.

H. Martens et W.M. de Muynck. Nonideal quantum measurements. *Found. Physics*, 20(3) :255–281, 1990.

S. Massar et S Popescu. Optimal extraction of information from finite quantum ensembles. *Phys. Rev. Lett.*, 74 :1259–1263, 1995.

P. Massart. *Concentration Inequalities and Model Selection.* Lecture Notes in Mathematics. Springer-Verlag, 2006. École d'été de Probabilité de Saint-Flour 2003.

K. Matsumoto. A new approach to the cramer-rao type bound of the pure state model. *J. Phys. A*, 35(13) :3111–3123, 2002.

K. Matsumoto. unpublished manuscript.

Katia Meziani. *Estimations Et Tests Non Paramétriques En Tomographie Quantique Homodyne.* PhD thesis, Université Paris VII, 2008.

Masoud Mohseni, Aephraim M. Steinberg, et Janos A. Bergou. Optical realization of optimal unambiguous discrimination for pure and mixed quantum states. *Physical Review Letters*, 93 :200403, 2004.

H. Nagaoka. On the parameter estimation problem for quantum statistical models. In M. Hayashi, editor, *Asymptotic Theory of Quantum Statistical Inference*, pages 125–132. World Scientific, 2005.

H. Nagaoka. A generalization of the simultaneous diagonalization of hermitian matrices and its relation to quantum estimation theory. *Trans. Jap. Soc. Indust. Appl. Math.*, 1 :43–56, 1991.

Hiroshi Nagaoka et Masahito Hayashi. An information-spectrum approach to classical and quantum hypothesis testing for simple hypotheses. *IEEE Transactions on Information Theory*, 53 :534, 2007.

Michael Nussbaum et Arleta Szkola. A lower bound of chernoff type for symmetric quantum hypothesis testing, 2006.

M. Ohya et Petz, D. *Quantum Entropy and its Use.* Springer Verlag, Berlin-Heidelberg, 2004.

Alexei Ourjoumtsev. *Étude théorique et expérimentale de superpositions quantiques cohérentes et d'états intriqués non-gaussiens de la lumière.* PhD thesis, Université Paris-Sud, 2007. URL `http ://tel.archives-ouvertes.fr/tel-00200715/en/`.

Masaki Owari et Masahito Hayashi. Two-way classical communication remarkably improves local distinguishability. *New Journal of Physics*, 10 :013006, 2008.

M. Ozawa. *Research Reports in Information Science A*, 74, 1980.

P.J. Lahti P. Busch et P. Mittelstaedt. *The Quantum Theory of Measurement.* Lecture Notes in Physics. Berlin Springer Verlag, 1991.

M. G. A. Paris et J. Řeháček, editors. *Quantum State Estimation*, 2004.

K. R. Parthasarathy. On Estimating the State of a Finite Level Quantum System. *ArXiv Quantum Physics e-prints*, August 2004.

Vern I. Paulsen. *Completely bounded maps and dilations.* John Wiley & Sons, Inc., New York, NY, USA, 1987. ISBN 0-470-20369-2.

A. Peres. How to differentiate between non-orthogonal states. *Physics Letters A*, 128 :19–19, March 1988.

A. Peres. *Quantum Theory : Concepts an Methods.* Kluwer Academic Press, 1993.

Asher Peres et Petra F. Scudo. Transmission of a cartesian frame by a quantum system. *Physical Review Letters*, 87 :167901, 2001.

D. Petz. *An Invitation to the Algebra of Canonical Commutation Relations.* Leuven University Press, 1990.

D. Petz. Sufficient subalgebras and the relative entropy of states of a von neumann algebra. *Commun. Math. Phys.*, 105 :123–131, 1986.

D. Petz et A. Jenčová. Sufficiency in quantum statistical inference. *Commun. Math. Phys.*, 263 :259 – 276, 2006.

D. Petz, K. M. Hangos, A. Szántó, et F. Szöllősi. State tomography for two qubits using reduced densities. *MATH.GEN.*, 39 :10901, 2006.

Dénes Petz. Complementarity in quantum systems, 2006.

Arthur O. Pittenger et Morton H. Rubin. Mutually unbiased bases, generalized spin matrices and separability. *Linear Algebra and its Applications*, 390 :255, 2004.

E. Prugorevčki. Information-theoretical aspects of quantum measurement. *International Journal of Theoretical Physics*, 16(5) :321–331, 1977.

Daowen Qiu. Minimum-error discrimination between mixed quantum states, 2007.

P. Raynal et N. Lütkenhaus. Optimal unambiguous state discrimination of two density matrices : Lower bound and class of exact solutions. *Physical Review A*, 72(2) :022342–+, August 2005.

Philippe Raynal, Norbert Lutkenhaus, et Steven J. van Enk. Reduction theorems for optimal unambiguous state discrimination of density matrices. *Physical Review A*, 68 :022308, 2003.

N.W. Rickert. Amenable groups and the fixed point property. *Trans. Amer. Math. Soc.*, 127 :221 – 232, 1967.

E. Riis et S. M. Barnett. Letter experimental demonstration of polarization discrimination at the helstrom bound. *Physial Review A*, 64 :012303, 2001.

Terry Rudolph, Robert W. Spekkens, et Peter Shipley Turner. Unambiguous discrimination of mixed states. *Physical Review A*, 68 :010301, 2003.

Massimiliano F. Sacchi. Optimal discrimination of quantum operations. *Physical Review A*, 71 :062340, 2005a.

Massimiliano F. Sacchi. Minimum error discrimination of pauli channels. *Journal of the Optical Society of America B*, 7 :S333, 2005b.

Massimiliano F. Sacchi. Entanglement can enhance the distinguishability of entanglement-breaking channels. *Physical Review A*, 72 :014305, 2005c.

Masahide Sasaki, Stephen M. Barnett, Richard Jozsa, Masao Osaki, et Osamu Hirota. Accessible information and optimal strategies for real symmetrical quantum sources. *Physical Review A*, 59 :3325, 2002.

I.V. Schensted. *A course on the application of group theory to quantum mechanics.* Neo press (Peaks Island), 1976.

S. Schiller, G. Breitenbach, S. F. Pereira, T. Müller, et J. Mlynek. Quantum statistics of the squeezed vacuum by measurement of the density matrix in the number state representation. *Phys. Rev. Lett.*, 77 :2933–2936, 1996.

J. Schwinger. Unitary Operator Bases. *Proceedings of the National Academy of Science*, 46 :570–579, April 1960.

G. A. Smith, A. Silberfarb, I. H. Deutsch, et P. S. Jessen. Efficient Quantum-State Estimation by Continuous Weak Measurement and Dynamical Control. *Phys. Rev. Lett.*, 97 :180403–+, 2006.

D. T. Smithey, Beck, M., Raymer, M. G., et Faridani, A. Measurement of the Wigner distribution and the density matrix of a light mode using optical homodyne tomography : Application to squeezed states and the vacuum. *Phys. Rev. Lett.*, 70 : 1244–1247, 1993.

Charles Stein. Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. *Proc. 3rd Berkeley Sympos. Math. Statist. Probability*, 1 :197 – 206, 1956.

W. F. Stinespring. Positive functions on c^*-algebras. *Proceedings of the American Society*, 6 :211–216, 1955.

J. K. Stockton, JM Geremia, A. C. Doherty, et H. Mabuchi. Characterizing the entanglement of symmetric multi-particle spin-1/2 systems. *Phys. Rev. A*, 67 : 022122, 2003.

H. Strasser. *Mathematical Theory of Statistics.* De Gruyter, Berlin, New York, 1985.

Xiaoming Sun, Shengyu Zhang, Yuan Feng, et Mingsheng Ying. Mathematical nature of and a family of lower bounds for the success probability of unambiguous discrimination. *Phys. Rev. A*, 65(4) :044306, Apr 2002.

E. Torgersen. *Comparison of Statistical Experiments.* Cambridge University Press, 1991.

M. A. P. Touzel, R. B. A. Adamson, et A. M. Steinberg. Optimal bounded-error strategies for projective measurements in non-orthogonal state discrimination, 2007.

A. van der Vaart. Limits of statistical experiments. unpublished manuscript.

A.W. van der Vaart. *Asymptotic Statistics.* Cambridge University Press, 1998.

A.W. van der Vaart et Wellner, J.A. *Weak Convergence and Empirical Processes.* Springer, New York, 1996.

G. Vidal, J. I. Latorre, P. Pascual, et R. Tarrach. Optimal minimal measurements of mixed states. *Phys. Rev. A*, 60 :126, 1999.

S. Virmani, M. F. Sacchi, M. B. Plenio, et D. Markham. Optimal local discrimination of two multipartite pure states. *Physics Letters A*, 288 :62, 2001.

David Vitali, Stefan Kuhr, Michel Brune, et Jean-Michel Raimond. A cavity-qed scheme for heisenberg-limited interferometry, 2006.

K. Vogel et H. Risken. Determination of quasiprobability distributions in terms of probability distributions for the rotated quadrature phase. *Phys. Rev. A*, 40 : 2847–2849, 1989.

K. Vogel et Risken, H. Determination of quasiprobability distributions in terms of probability distributions for the rotated quadrature phase. *Phys. Rev. A*, 40 : 2847–2849, 1989.

A. Wald. *Statistical Decision Functions.* John Wiley & Sons, New York, 1950.

A. Wald. Tests of statistical hypotheses concerning several parameters when the number of observations is large. *Trans. Amer. Math. Soc.*, 54 :426–482, 1943.

Jonathan Walgate, Anthony J. Short, Lucien Hardy, et Vlatko Vedral. Local distinguishability of multipartite orthogonal quantum states. *Physical Review Letters*, 85 :4972, 2000.

Guoming Wang et Mingsheng Ying. Unambiguous discrimination among quantum operations. *Physical Review A*, 73 :042301, 2006.

R. F. Werner. Optimal cloning of pure states. *Phys. Rev. A*, 58 :1827–1832, 1998.

W. K. Wootters. Statistical distance and hilbert space. *Phys. Rev. D*, 23(2) :357–362, Jan 1981.

W. K. Wootters et B. D. Fields. Optimal state-determination by mutually unbiased measurements. *Annals of Physics*, 191 :363–381, May 1989.

H. Yuen, R. Kennedy, et M. Lax. Optimum testing of multiple hypotheses in quantum detection theory. *IEEE Trans. Inform. Theory*, 21 :125– 134, 1975a.

H. Yuen, R. Kennedy, et M. Lax. Optimum testing of multiple hypotheses in quantum detection theory. *IEEE Transactions on Information Theory*, 21(2) :125–134, 1975b.

H. P. Yuen et Lax, M. Multiple-parameter quantum estimation and measurement of non-selfadjoint observables. *IEEE Trans. Inform. Theory*, 19 :740, 1973.

B. Yurke. Input states for enhancement of fermion interferometer sensitivity. *Phys. Rev. Lett.*, 56(15) :1515–1517, Apr 1986.

A. Zavatta, S. Viciani, et M. Bellini. Quantum to classical transition with single-photon-added coherent states of light. *Science*, 306 :660–662, 2004.

Jun Zhang, Jiri Vala, K. Birgitta Whaley, et Shankar Sastry. A geometric theory of non-local two-qubit operations. *Physical Review A*, 67 :042313, 2003.

Shengyu Zhang, Yuan Feng, Xiaoming Sun, et Mingsheng Ying. Upper bound for the success probability of unambiguous discrimination among quantum states. *Phys. Rev. A*, 64(6) :062103, Nov 2001.

K. Zyczkowski et H. J. Sommers. Average fidelity between random quantum states. *Phys. Rev. A*, 71 :032313, 2005.