

Transformation & Uncertainty

Some Thoughts on Quantum Probability Theory,
Quantum Statistics, and Natural Bundles

Thesis committee:

Prof. dr. E. P. van den Ban, Universiteit Utrecht

Dr. N. Datta, University of Cambridge

Prof. dr. M. Fannes, Katholieke Universiteit Leuven

Prof. dr. E. J. N. Looijenga, Universiteit Utrecht

Prof. dr. C. Wockel, Universität Hamburg

ISBN 978-90-393-5409-4

Copyright © 2010 by B. Janssens.

Transformation & Uncertainty

Some Thoughts on Quantum Probability Theory,
Quantum Statistics, and Natural Bundles

Transformatie & Onzekerheid

Ideeën over Kwantumkansrekening,
Kwantumstatistiek en Natuurlijke Bundels

(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht op gezag van de rector magnificus, prof. dr. J. C. Stoof, ingevolge het besluit van het college voor promoties in het openbaar te verdedigen op maandag 4 oktober 2010 des middags te 2.30 uur

door

Bas Janssens

geboren op 20 mei 1981 te Maastricht

Promotor: Prof. dr. R. Fernández
Copromotoren: Dr. J. W. van de Leur
Dr. J. D. M. Maassen

... and to him the meaning of an episode was not inside like a kernel but outside, enveloping the tale which brought it out only as a glow brings out a haze...

Joseph Conrad, Heart of Darkness

Contents

1	Introduction	1
2	Background	5
2.1	Quantum Mechanics	6
2.1.1	Probabilities and the Spectral Theorem	6
2.1.2	Transformation, Symmetry and Conservation Laws	9
2.2	Quantum Probability Theory	15
2.2.1	QPT on Finite Dimensional Hilbert Spaces.	15
2.2.2	QPT on General Hilbert Spaces	19
2.3	Quantum Stochastic Differential Equations	23
2.3.1	Fock Space	23
2.3.2	Weyl Operators, Fields and Momenta	24
2.3.3	The Harmonic Oscillator à la Fock	26
2.3.4	The Harmonic Oscillator à la Gauß	27
2.3.5	The Wiener Measure	27
2.3.6	Stochastic Differential Equations	30
2.3.7	Fock Space and the Wiener Measure	31
2.3.8	Quantum Stochastic Differential Equations	33
2.4	Bundles and Classical Fields	37
2.4.1	Bundles, Sections and Jets	37
2.4.2	The Euler-Lagrange Equation	39
2.4.3	Transformations and Symmetries	42
2.4.4	Infinitesimally Natural Bundles	47
3	Information Transfer Implies State Collapse	49
3.1	Introduction	49
3.2	Abstract Information Extraction	50
3.2.1	Information Extraction	51
3.2.2	State Collapse	51
3.3	Open Systems	52
3.3.1	Unitary Information Transfer and Decoherence	53
3.3.2	Unitary Information Transfer and State Collapse	57
3.3.3	Information Leakage to the Environment	59
3.4	Closed Systems	60
3.4.1	Macroscopic Systems	60

3.4.2	Examples	62
3.5	Conclusion	65
4	Unifying Decoherence and the Heisenberg Principle	69
4.1	Introduction	69
4.2	Information Transfer	70
4.2.1	General Information Transfer	71
4.2.2	Unbiased Information Transfer	73
4.3	An Operator Valued Form	74
4.3.1	A Cauchy-Schwarz Inequality	74
4.3.2	Perfect Information Transfer	76
4.4	Joint Measurement	77
4.4.1	Joint Measurement Inequality	77
4.4.2	The Beamsplitter as a Joint Measurement	77
4.4.3	Sharpness and Comparison	78
4.5	The Heisenberg Principle	79
4.5.1	A Figure of Merit	79
4.5.2	HP for the Unbiased Case	80
4.5.3	HP for General Information Transfer	82
4.5.4	Application: Resonance Fluorescence	83
4.6	Classical Coding	86
4.6.1	Classical Coding	86
4.6.2	Classical Coding Inequality	87
4.6.3	The No-Cloning method	88
4.7	Collapse of the Wave function	88
4.7.1	Collapse for Unbiased Information Transfer	89
4.7.2	Application: Perfect Qubit Measurement	91
4.7.3	Collapse for Nondestructive Measurement	91
4.7.4	Application: Nondestructive Qubit-Measurement	93
4.7.5	Sharpness of the Bounds	94
4.8	Conclusion	94
5	Optimal Pointers for Joint Measurement	97
5.1	Introduction	97
5.2	The Model	98
5.3	Quality of Information Transfer	101
5.3.1	Defining the Quality of Information Transfer	102
5.3.2	Calculating the Quality of Information Transfer	103
5.4	The Weighted Path	105
5.4.1	Restricting the Class of Pointers	106
5.4.2	Distribution of Y_t	107
5.5	Variational Calculus	108
5.5.1	Optimal Measurement of σ_x	109
5.5.2	Optimal Measurement of σ_z	111
5.6	Distribution of Pointer Variables	113
5.6.1	Calculation of h^* - and \tilde{h} -Densities	113

5.6.2	Plots of σ_x -Measurement	114
5.6.3	Plots of σ_z -Measurement	115
5.7	Discussion	116
6	Optimal Estimation of Qubit States	117
6.1	Introduction	117
6.2	State Estimation	121
6.2.1	Qubit State Estimation: the Localization Principle	123
6.3	Local Asymptotic Normality	125
6.3.1	Introduction to LAN and some Definitions	125
6.3.2	Convergence to the Gaussian Model	126
6.4	Time Evolution of the Interacting System	130
6.4.1	Quantum Stochastic Differential Equations	130
6.4.2	Solving the QSDE for the Oscillator	132
6.4.3	QSDE for Large Spin	132
6.5	The Second Stage Measurement	134
6.5.1	The Heterodyne Measurement	135
6.5.2	Energy Measurement	136
6.6	Asymptotic Optimality of the Estimator	138
6.7	Conclusions	142
A	Appendix: Proof of Theorem 6.2	143
A.1	The Map T_n	143
A.2	The Map S_n	147
B	Appendix: Proof of Theorem 6.3	148
B.1	First Estimate	148
B.2	Second Estimate	149
7	Bundles with a Lift of Infinitesimal Diffeomorphisms	153
7.1	Introduction	153
7.2	Principal Bundles as Lie Algebra Extensions	155
7.2.1	Ideals of the Lie Algebra of Vector Fields	156
7.2.2	The Splitting as a Local Map	158
7.2.3	The Splitting as a Differential Operator	159
7.3	Lie Groupoids and Algebroids of Jets	161
7.3.1	The Lie Groupoid of k -Jets	162
7.3.2	The Lie Algebroid of k -Jets	162
7.3.3	The Gauge Groupoid and its Algebroid	164
7.3.4	The Splitting as a Homomorphism of Lie Algebroids	164
7.4	The Classification Theorem	165
7.4.1	Integrating a Homomorphism of Lie Algebroids	165
7.4.2	Classification	166
7.4.3	The Bundle $\tilde{F}^{+k}(M)$	168
7.5	More General Fibre Bundles	170
7.5.1	Structured Fibre Bundles	171
7.5.2	Vector Bundles	172
7.6	Flat Connections	173

7.6.1	Lie Algebras That Do Not Contain $\mathfrak{sl}(\mathbb{R}^n)$	173
7.6.2	Lie Algebra Cohomology	175
7.7	General relativity, Gauge theory, and Spinors	176
7.7.1	General Relativity	176
7.7.2	Fermions and Spin Structures	177
7.7.3	Gauge Fields and Spin^G -structures	178
7.7.4	Induced Unitary Representations of $\widehat{\text{Diff}}^0(M)$	183
7.7.5	Some Physical Theories	188
7.8	Discussion	191
	Bibliography	192
	Samenvatting	203
	Dankwoord	213
	Curriculum Vitae	215

Chapter 1

Introduction

This thesis aims to investigate mathematically two of the most robust phenomena in modern physics. The first is the ‘uncertainty principle’ in quantum mechanics. It states that no measurement can be performed without disturbing the system. The second is the ‘equivariance principle’ in classical field theory, which states that any transformation of space-time must also transform the fields.

The word ‘robust’ should be understood in the sense of ‘independent of the dynamics of the particular model’. A physical theory can usually be separated into a kinematic and a dynamical part. The kinematic part of a theory describes the basic physical objects in terms of mathematics. The dynamical part then prescribes the precise way in which these variables evolve. Our results will largely be independent of the dynamics, making them applicable to a wide range of different models.

In the case of quantum mechanics for example, the kinematic part of the theory states that a system is described by a Hilbert space, observables by self-adjoint operators, states by vectors and transformations by unitary operators. This alone will suffice for us to prove quantitative versions of the uncertainty principle, as well as asymptotic optimality of our many-qubit state estimation scheme. The dynamics of the theory, embodied by a model-specific Hamilton operator, only come into play when we test these ideas in specific situations.

In the case of classical field theory, the kinematic part of the theory entails that space-time is modelled by a smooth manifold, fields are modelled by sections of a bundle over this manifold, and transformations are described by bundle automorphisms. This alone will suffice for us to classify the relevant bundles. The dynamical part of the theory, in the form of a model-specific Lagrangian, will have no part in this.

Although I would warmly recommend it, I am aware that the reader may not wish to read this thesis from cover to cover. Chapters 3 through 7 can be read independently, with perhaps an occasional glance at the ‘background’ chapter 2. We give a brief description of each chapter.

Chapter 2. Background We give a short introduction to quantum probability theory and classical field theory, providing background and motivation for the rest of the thesis.

Chapter 3. Information Transfer Implies State Collapse We prove that within the framework of quantum mechanics, a transfer of information from a system to the outside necessarily causes decoherence on that system. We also prove a quantitative version of this, linking the amount of decoherence to the quality of information transfer. Finally, we shed some light on the fact that coherence is very hard to observe in macroscopic systems. This chapter is an adapted version of [JM06], written together with Hans Maassen.

Chapter 4. Unifying Decoherence and the Heisenberg Principle Within the framework of quantum probability theory, we prove quantitative versions of the Heisenberg principle, the joint measurement theorem, the no-cloning theorem, and the collapse of the wave function. An important difference with chapter 3 is that all bounds will be in terms of the measurement procedure alone; they will not depend on the particular state of the measured system. This is a good thing because the state of the measured system is generally unknown—hence the need for measurement. A part of this chapter appeared earlier in [Jan08].

Chapter 5. Optimal Pointers for Joint Measurement We apply the joint measurement inequality of chapter 4 to the specific situation of a 2-level atom, i.e. a qubit, coupled to the quantized electromagnetic field. Information on σ_x and σ_z leaks into the field and affects its quadrature, which can then be probed using homodyne detection. The joint measurement inequality provides us with a sharp upper bound on the quality of any simultaneous measurement of σ_x and σ_z . Using quantum stochastic calculus, we investigate how close one can come to this bound using the above setup. Somewhat to our surprise, we found that it cannot be reached, although one can come as close as 5.6%. This chapter is an adapted version of [JB06], written together with Luc Bouten.

Chapter 6. Optimal Estimation of Qubit States We propose an asymptotically optimal scheme to estimate the state of n identically prepared qubits. The qubits are coupled to the quantized electromagnetic field for some period of time, during which continuous time measurements on the field are performed. The proof of optimality relies on the concept of local asymptotic normality for qubits: for large n , the ensemble of qubits becomes statistically equivalent to the tensor product of a classical Gaussian distribution and a Gaussian state on a quantum harmonic oscillator. The Gaussian states are well understood, which means that we can use the statistical correspondence to better understand the qubits.

In a way, this is not unlike the classical problem of determining the bias of an unfair coin that is tossed n times. In this case, the n coin tosses become statistically equivalent to a single Gaussian in the limit of large n . One should

guard against taking this analogy too far though: in the quantum case, there is a difference between the act of handing someone n identical qubits to probe, and the act of handing someone n times a single qubit to probe. At least for mixed states, a joint measurement on an ensemble of n qubits performs strictly better than the combined information of n single qubit measurements. The whole is more than the sum of its parts, so to speak. This chapter is an adapted version of [GJK08], which was written together with Mădălin Guță and Jonas Kahn.

Chapter 7. Bundles with a Lift of Infinitesimal Diffeomorphisms A bundle is called ‘natural’ if its fibres transform in a local fashion under diffeomorphisms of the base. Examples include tangent bundles, frame bundles, and most other bundles of a local geometric nature. We extend the notion of a ‘natural’ bundle to that of an ‘infinitesimally natural’ one by requiring the fibres to transform under *infinitesimal* diffeomorphisms only. We classify exactly what is gained by this extension. Bundles of spin type are never natural, but it turns out that some of them are indeed natural in the above infinitesimal sense. In classical field theory, fields are sections of a bundle over space-time. Spin and energy of a field are well-defined precisely when the corresponding bundle is infinitesimally natural. The fact that bundles that correspond to fermionic fields (such as quarks, electrons, etc.) are never natural shows that the generalization is really called for.

Samenvatting We give a non-rigorous Dutch summary of the main results, containing no formulae and as many pictures as possible. If you don’t have much experience with (this kind of) mathematics, and if you understand Dutch, then this may be a good place to start reading.

Chapter 2

Background

The chapters to come can be read independently, and are self-sufficient to a certain degree. Complete autarky however is a dream, and, when realized in a PhD thesis, probably a nightmare. From time to time then, the reader is referred elsewhere, asked to trust a statement bluntly placed before him, or simply assumed to ‘know what is going on’. This is inevitable. The flip side of the progress of science is the overspecialization that comes with it.

However, as details and applications become more and more complex, the foundations of science have a tendency to simplify with time. This makes it feasible for us to sketch the fundamentals of quantum probability theory and classical field theory in the next forty-odd pages. We aim to explain some of the more vital ingredients that will enter our reasoning, but the main goal here is to provide background and motivation for the problems that will be addressed later on.

We will start by giving a short introduction to quantum mechanics, focusing first on the way in which the theory gives rise to probability measures, and then on matters of transformation and symmetry. An even shorter introduction to quantum probability theory will follow, with an emphasis on transformations of the theory. Everything will be done first in the finite dimensional setting, and only then in the general case. The purpose of this is to prevent the elegant, yet somewhat involved functional analytic side of the theory from overshadowing its simpler algebraic counterpart.

An important tool for us will be the quantum stochastic calculus of Hudson and Parthasarathy. This is a generalization of the Itô stochastic calculus on Wiener space, which is essentially calculus with stochastic infinitesimal increments. After a primer on Fock spaces and Weyl operators, we will spend some words on the Wiener measure and Itô calculus. We will then see how the field operators in a free quantum field theory form a quantum version of the stochastic infinitesimal increments of Itô calculus, giving rise to quantum stochastic calculus.

The last part of this chapter will be a brief introduction to classical field theory from a differential geometric point of view. In a nutshell, this means

that we will consistently treat fields as sections of a fibre bundle. The emphasis is once again on transformation and symmetry. We formulate the requirement that fields should transform in a definite fashion under infinitesimal space-time transformations. The ramifications of this will become clear in chapter 7.

As already stated, we mean to sketch only the most elemental aspects of the theory. For a more complete picture, at perhaps a more responsible pace, the reader is referred to [Par92], [Fra04], [FR04], [Ban04], and [Maa04], on which we based the following account.

2.1 Quantum Mechanics

In quantum mechanics, a system is modelled by a Hilbert space \mathcal{H} . A state of the system is modelled by a ray $\mathbb{C}\psi$ in \mathcal{H} , and observables are represented by self-adjoint, possibly unbounded operators A on \mathcal{H} . Transformations of a closed system are described by unitary (or antiunitary) operators U on \mathcal{H} , and there is a distinguished continuous one-parameter group $t \mapsto U_t$ of unitary operators that describes time evolution.

Quantum mechanics can only be interpreted in terms of probabilities. There is a canonical map $(\mathbb{C}\psi, A) \mapsto \mathbb{P}_\psi$ that associates to each ray $\mathbb{C}\psi$ and self-adjoint operator A a probability measure \mathbb{P}_ψ on the spectrum of A . It is interpreted as follows.

If the system is in state $\mathbb{C}\psi$, then observation of an observable A yields an outcome $\lambda \in S$ with probability $\mathbb{P}_\psi(S)$.

Everything else follows from these postulates.

2.1.1 Probabilities and the Spectral Theorem

We will first describe the map $(\mathbb{C}\psi, A) \mapsto \mathbb{P}_\psi$ in the special case of finite dimensional Hilbert space, and then in the general case.

The finite dimensional case is simple but interesting. It is simple because a choice of orthonormal basis identifies \mathcal{H} with \mathbb{C}^n , reducing many of our problems to linear algebra. It is interesting because the results in chapter 3 and 4 are as nontrivial here as they are in the general case, and –if the reader will forgive me my patronizing remark– it might be helpful to keep the finite dimensional case in mind when reading these chapters.

Let us take the opportunity to fix some notation. All Hilbert spaces are over \mathbb{C} unless specifically mentioned otherwise, and we will take the inner product $\langle \cdot | \cdot \rangle$ to be linear on the right. The adjoint of A is denoted A^\dagger , and $\mathcal{B}(\mathcal{H})$ denotes the algebra of bounded linear operators on \mathcal{H} . For n -dimensional \mathcal{H} , this is just the algebra of $n \times n$ -matrices over \mathbb{C} .

Hilbert Spaces of Finite Dimension

An orthogonal projection E is in particular a self-adjoint operator. Because its spectrum is $\{0, 1\}$, it should be interpreted as an *event*. The outcome 1 means that the event occurs, the outcome 0 means that it does not.

If E projects onto $\mathcal{H}_1 \subset \mathcal{H}$, then we have the orthogonal decomposition $\mathcal{H} = \mathcal{H}_0 \oplus \mathcal{H}_1$. A unit vector ψ splits as $\psi = \psi_0 + \psi_1$, with $\|\psi_0\|^2 + \|\psi_1\|^2 = 1$. This leads us to define the probability distribution \mathbb{P}_ψ on $\{0, 1\}$ by

$$\mathbb{P}_\psi(0) = \|\psi_0\|^2, \quad \mathbb{P}_\psi(1) = \|\psi_1\|^2.$$

This depends only on the ray $\mathbb{C}\psi$, so that we have our map $(\mathbb{C}\psi, A) \mapsto \mathbb{P}_\psi$ in the special case that A is a projection.

We can play the same game if we split \mathcal{H} into any number of mutually orthogonal spaces. Such a decomposition determines a projection valued measure.

Definition 2.1 *Let Ω be a finite set, and \mathcal{H} a finite dimensional Hilbert space. Then a projection valued measure (or PVM) is a map $E : \Omega \rightarrow \mathcal{B}(\mathcal{H})$ such that:*

- $E(\omega)$ is an orthogonal projection for all $\omega \in \Omega$.
- If $\omega \neq \omega'$, then $E(\omega)E(\omega') = 0$.
- $\sum_{\Omega} E(\omega) = \mathbf{1}$.

If \mathcal{H}_ω is the image of $E(\omega)$, it is clear that $\mathcal{H} = \bigoplus_{\Omega} \mathcal{H}_\omega$. Accordingly, a unit vector ψ decomposes as $\psi = \sum_{\Omega} \psi_\omega$, with $\|\psi_\omega\|^2 = \langle \psi | E(\omega)\psi \rangle$ and $\sum_{\Omega} \langle \psi | E(\omega)\psi \rangle = 1$. This means that a unit vector ψ gives rise to a measure \mathbb{P}_ψ on Ω , defined by $\mathbb{P}_\psi(\omega) = \langle \psi | E(\omega)\psi \rangle$.

For any unit vector ψ , we will define the map $\rho_\psi : \mathcal{B}(\mathcal{H}) \rightarrow \mathbb{C}$, called a *pure state*, by

$$\rho_\psi(A) = \langle \psi | A\psi \rangle.$$

With this notation, we have

$$\mathbb{P}_\psi = \rho_\psi \circ E.$$

Note that $\rho_\psi = \rho_{\psi'}$ if and only if $\psi' = c\psi$ for some $c \in \mathbb{C}^\times$, so that a pure state corresponds precisely to a ray $\mathbb{C}\psi$ in the Hilbert space. We will identify the two, and think of the *pure state space* as either

$$\mathcal{P}(\mathcal{H}) = \{\rho_\psi ; \psi \in \mathcal{H}, \|\psi\| = 1\}$$

or as the space of rays

$$\mathcal{P}(\mathcal{H}) = \mathcal{H} - \{0\} / \mathbb{C}^\times.$$

We leave it up to the reader to decide whether \mathcal{P} stands for ‘pure’ or ‘projective’.

The spectral theorem says that each self-adjoint operator A gives rise to a projection valued measure on its spectrum.

Theorem 2.1 *Let A be a self-adjoint operator on a finite dimensional Hilbert space \mathcal{H} , and let $\mathbf{Spec}(A)$ be its spectrum of eigenvalues. Then there exists a unique projection valued measure $E_A : \mathbf{Spec}(A) \rightarrow \mathcal{B}(\mathcal{H})$ such that*

$$A = \sum_{\lambda \in \mathbf{Spec}(A)} \lambda E_A(\lambda).$$

The $E_A(\lambda)$ are the projections onto the eigenspaces \mathcal{H}_λ , and $\mathcal{H} = \bigoplus_{\text{Spec}(A)} \mathcal{H}_\lambda$ is the eigenspace decomposition.

The map $(\mathbb{C}\psi, A) \mapsto \mathbb{P}_\psi$ is now defined as follows.

- The ray $\mathbb{C}\psi$ is identified with the pure state $\rho_\psi : \mathcal{B}(\mathcal{H}) \rightarrow \mathbb{C}$.
- An observable A gives rise to a PVM $E_A : \text{Spec}(A) \rightarrow \mathcal{B}(\mathcal{H})$.
- The measure \mathbb{P}_ψ on $\text{Spec}(A)$ is given by $\mathbb{P}_\psi = \rho_\psi \circ E$.

We interpret $\mathbb{P}_\psi(\lambda)$ as the probability that if the system is in state $\mathbb{C}\psi$, a measurement of A yields the outcome λ .

Infinite Dimensional Hilbert Spaces

Although there do exist physically relevant systems where the Hilbert space is finite dimensional, this is not typically the case. We adapt our framework to encompass general Hilbert spaces.

We equip the algebra $\mathcal{B}(\mathcal{H})$ of bounded operators on \mathcal{H} with the *weak topology*, the coarsest topology that makes $\rho_\psi(A) = \langle \psi | A\psi \rangle$ continuous for all ψ in \mathcal{H} . We then introduce the notion of a projection valued measure for general Hilbert spaces as follows.

Definition 2.2 *Let (Ω, Σ) be a measurable space, and \mathcal{H} a Hilbert space. Then a projection valued measure or PVM is a map $E : \Sigma \rightarrow \mathcal{B}(\mathcal{H})$ with the following properties.*

- $E(S)$ is a projection for all $S \in \Sigma$.
- $E(\emptyset) = 0$, and $E(\Omega) = \mathbf{1}$.
- If $S_1 \cap S_2 = \emptyset$, then $E(S_1)E(S_2) = 0$.
- If S_i is a sequence of disjoint sets in Σ , then $E(\bigcup_{i=0}^{\infty} S_i) = \sum_{i=0}^{\infty} E(S_i)$, where the limit is in the weak topology.

Each ray $\mathbb{C}\psi$ induces a probability measure \mathbb{P}_ψ on (Ω, Σ) by applying the pure state $\rho_\psi(A) := \langle \psi | A\psi \rangle$ to E , i.e.

$$\mathbb{P}_\psi(S) = (\rho_\psi \circ E)(S).$$

If f is a measurable function on Ω , then an integral

$$I_E(f) := \int_{\Omega} f(\omega)E(d\omega) \tag{2.1}$$

can be defined by the weak limit of Riemann sums. The resulting operator $I_E(f)$ on \mathcal{H} is bounded if f is bounded, which invites us to consider I_E as a linear map $I_E : L^\infty(\Omega, \Sigma) \rightarrow \mathcal{B}(\mathcal{H})$. It is a homomorphism of $*$ -algebras, i.e. it respects multiplication $I_E(fg) = I_E(f)I_E(g)$ and adjunction $I_E(f^\dagger) = I_E(f)^\dagger$.

In particular, $I_E(f)$ is self-adjoint if f is real, and positive if f is positive. With respect to \mathbb{P}_ψ , we have $\mathbb{E}(f) = \rho_\psi(I_E(f))$.

Observables correspond to self-adjoint (i.e. symmetric, closed and possibly unbounded) operators on \mathcal{H} . The spectral theorem says that any self-adjoint operator gives rise to a projection valued measure on its spectrum, called the spectral measure.

Theorem 2.2 (spectral theorem) *Let A be a self-adjoint operator on \mathcal{H} . Then there exists a unique PVM $E_A : \Sigma(\mathbb{R}) \rightarrow \mathcal{B}(\mathcal{H})$ on the Borel sigma-algebra of \mathbb{R} such that, with λ the identity function on \mathbb{R} ,*

$$A\psi = \int_{\mathbb{R}} \lambda E(d\lambda)\psi$$

for all ψ in the domain of A . The PVM is concentrated on the spectrum of A .

Proof. See for example [Par92]. □

Any pure state ρ_ψ therefore induces a probability measure \mathbb{P}_ψ on $\mathbf{Spec}(A)$ by $\mathbb{P}_\psi = \rho_\psi \circ E_A$. Under this probability measure, the expectation of the observable A of a system in a state ψ is $\mathbb{E}(\lambda) = \rho_\psi(A)$. Its variance is $\rho_\psi(A^2) - \rho_\psi(A)^2$.

2.1.2 Transformation, Symmetry and Conservation Laws

Let us devote some attention to the way in which transformations and symmetries occur in quantum mechanics, and then give examples of relevant quantum systems. Although the proper treatment of symmetry in quantum mechanics will not play a major role in this thesis, it will help us to properly interpret the examples.

Hilbert Spaces of Finite Dimension

As before, we will first focus attention on finite dimensional Hilbert spaces, before looking at the general case. Throughout, we will denote by $U(\mathcal{H})$ the group of unitary operators on \mathcal{H} .

Transformations A unitary operator $U \in U(\mathcal{H})$ induces an automorphism $\alpha_U : A \mapsto U^\dagger A U$ of $\mathcal{B}(\mathcal{H})$. This in turn acts on $\mathcal{P}(\mathcal{H})$ by its dual, $\alpha_U^* : \rho_\psi \mapsto \rho_\psi \circ \alpha_U$. The action of the unitary operators on the pure state space is therefore governed by the group of inner automorphisms $\mathcal{P}U(\mathcal{H}) = \{\alpha_U ; U \in U(\mathcal{H})\}$.

Recall that the pure state space has the equivalent description $\mathcal{P}(\mathcal{H}) = \mathcal{H} - \{0\}/\mathbb{C}^\times$, on which α_U acts by $\mathbb{C}\psi \mapsto \mathbb{C}U\psi$. This shows that $\alpha_U = \alpha_{U'}$ if and only if $U' = cU$ for some $c \in \mathbb{C}^\times$. We therefore have the equivalent description $\mathcal{P}(U(\mathcal{H})) = U(\mathcal{H})/\mathbb{C}^\times$. Apparently, the \mathcal{P} stands for ‘projective’.

The action of an abstract Lie group G on a quantum mechanical system should therefore be described by a *projective* unitary representation, i.e. a continuous group homomorphism $\pi : G \rightarrow \mathcal{P}U(\mathcal{H})$.

Recall that a *linear* unitary representation is a continuous group homomorphism $\pi : G \rightarrow U(\mathcal{H})$. Every linear unitary representation gives rise to a projective unitary representation, but it is not true that every projective representation arises in this way.

Since the theory of linear unitary representations is well understood, we would like to reduce projective representations to linear ones. This is achieved by the following theorem, at the expense of slightly enlarging the group.

Theorem 2.3 *Let G be a connected Lie group, with universal cover \tilde{G} . Let \mathcal{H} be a finite dimensional Hilbert space. Then every projective unitary representation π of G on \mathcal{H} comes from a linear unitary representation $\tilde{\pi}$ of \tilde{G} on \mathcal{H} , in the sense that the following diagram commutes.*

$$\begin{array}{ccc} \tilde{G} & \xrightarrow{\tilde{\pi}} & U(\mathcal{H}) \\ \downarrow & & \downarrow \\ G & \xrightarrow{\pi} & \mathcal{P}U(\mathcal{H}) \end{array}$$

Proof. Since $\mathcal{H} \simeq \mathbb{C}^n$ for some finite n , we can consider π as a homomorphism $G \rightarrow \mathcal{P}(U(\mathbb{C}^n))$. The derived Lie algebra homomorphism $\text{Lie}(G) \rightarrow \mathfrak{su}(\mathbb{C}^n)$ extends to a group homomorphism $\tilde{G} \rightarrow \text{SU}(\mathbb{C}^n)$ by Lie's second theorem. \square

In what follows, we will assume that G has been replaced by the universal cover of its connected component of unity \tilde{G}_0 , and π by a linear unitary representation of \tilde{G}_0 .

Conservation Laws It is a remarkable fact of life that in every fundamental physical theory, infinitesimal symmetries correspond to conservation laws. In quantum mechanics, this correspondence is particularly straightforward.

Infinitesimal transformations are elements of the Lie algebra $\text{Lie}(G)$ of G . The following theorem describes how they give rise to observables.

Theorem 2.4 *Let π be a continuous unitary representation of G on \mathcal{H} . Then for each $X \in \text{Lie}(G)$, we have a well defined skew-symmetric operator $\dot{\pi}(X)$ defined by*

$$\dot{\pi}(X) = \pi(e^{-tX}) \frac{d}{dt} \pi(e^{tX}). \quad (2.2)$$

As $\dot{\pi}([X, Y]) = [\dot{\pi}(X), \dot{\pi}(Y)]$, this defines a Lie algebra homomorphism $\dot{\pi}$ of $\text{Lie}(G)$ into the Lie algebra $\mathfrak{su}(\mathcal{H})$ of skew-symmetric operators on \mathcal{H} .

The operator $i\dot{\pi}(X)$ is self-adjoint, and should be considered as the observable corresponding to the infinitesimal transformation X . We denote it $J(X)$. If one has to pass to the universal cover, then $\dot{\pi}$ and J are still well-defined, since the Lie algebras of G and \tilde{G}_0 are the same.

Time evolution is given by a one-parameter group of unitary transformations, i.e. a unitary representation $t \mapsto U_t$ of $(\mathbb{R}, +)$. Its Lie algebra is again \mathbb{R} , and the self-adjoint operator $J(1)$ is called the *Hamiltonian*, denoted H . The observable associated to this operator is the energy.

Equation (2.2) for the particular case of time evolution reads

$$i\frac{d}{dt}U_t = HU_t.$$

In the ‘Heisenberg picture’, where U acts by α_U , one writes $A_t = U_t^\dagger A U_t$ to obtain *Heisenberg’s equation of motion*

$$\frac{d}{dt}A_t = i[H, A_t]. \quad (2.3)$$

In the dual ‘Schrödinger picture’, where U acts by α_U^* , one writes $\psi_t = U_t \psi$ to obtain the *Schrödinger equation*

$$i\frac{d}{dt}\psi_t = H\psi_t. \quad (2.4)$$

A transformation is called a *symmetry* if it commutes with the time evolution $t \mapsto U_t$, i.e. if $[U_t, \pi(g)] = 0$ for all $t \in \mathbb{R}$. A group G is a group of symmetries if $\pi(g)$ is a symmetry for every $g \in G$. Because this implies that $\dot{\pi}(X)$ commutes with H , we immediately obtain from Heisenberg’s equation of motion the conservation law

$$\frac{d}{dt}J(X)_t = 0.$$

If X is an infinitesimal symmetry, then the probability measure on the spectrum of $J(X)$ is constant in time. For $J(X) = H$, this is called conservation of energy.

The fundamental problem in quantum mechanics is to reconstruct the time evolution $t \mapsto U_t$ from the given Hamiltonian H . This is usually done by finding the eigenspace decomposition $\mathcal{H} = \bigoplus_{\text{Spec}(H)} \mathcal{H}_E$ for H , and then computing $U_t = e^{-itH}$. This is greatly facilitated by the presence of symmetries.

The Qubit Time for an example. The internal degrees of freedom of an electron are described by the Hilbert space $\mathcal{H} = \mathbb{C}^2$. It transforms under the rotation group $\text{SO}(3, \mathbb{R})$ by a projective unitary representation.

Consider the map $\mathcal{PU}(\mathbb{C}^2) \rightarrow \text{SO}(3, \mathbb{R})$ defined as follows. In terms of the basis

$$\mathbf{1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

of $\mathcal{B}(\mathbb{C}^2)$, one readily checks that $\alpha_U(\lambda\mathbf{1}) = \lambda\mathbf{1}$, and that

$$\alpha_U\left(\sum_{i=1}^3 x_i \sigma_i\right) = \sum_{i,j=1}^3 R_{ij} x_j \sigma_i$$

for some $R \in \text{SO}(3, \mathbb{R})$. The map $\alpha_u \mapsto R$ turns out to be an isomorphism $\mathcal{PU}(\mathbb{C}^2) \xrightarrow{\sim} \text{SO}(3, \mathbb{R})$, and its inverse π is the projective representation of $\text{SO}(3, \mathbb{R})$ on \mathbb{C}^2 . If we choose the basis

$$\tau_x = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{pmatrix}, \quad \tau_y = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{pmatrix}, \quad \tau_z = \begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

of $\mathfrak{so}(3, \mathbb{R})$, then the Lie algebra homomorphism $\dot{\pi} : \mathfrak{so}(3, \mathbb{R}) \rightarrow \mathfrak{su}(\mathbb{C}^2)$ gives rise to three observables $J(\tau_x) = \frac{1}{2}\sigma_x$, $J(\tau_y) = \frac{1}{2}\sigma_y$ and $J(\tau_z) = \frac{1}{2}\sigma_z$. They should therefore be considered as the *internal angular momenta* or *spin* of the electron.

If the particle is placed in a vertically aligned magnetic field of strength B , then the effective Hamiltonian of the system is $H = B\sigma_z/2$. The group of symmetries is then *not* all of $\text{SO}(3, \mathbb{R})$, but only the subgroup $\text{SO}(2, \mathbb{R})$ of rotations around the z -axis. The only conserved quantity is therefore $H \sim \frac{1}{2}\sigma_z$. Heisenberg's equation of motion $i\frac{d}{dt}A_t = [H, A_t]$ can be solved to yield

$$t \mapsto a \left(\cos(Bt + \phi)\sigma_x + \sin(Bt + \phi)\sigma_y \right) + b\sigma_z,$$

which says that the angular momentum rotates around the z -axis with a frequency proportional to the magnetic field. In particular, the angular momentum in the z -direction is preserved, as was predicted by symmetry considerations.

The system $\mathcal{H} = \mathbb{C}^2$ is called a *qubit*, and one can use it to model any quantum system with only two (relevant) states, not just the electron. For example, in quantum optics, the interaction of an atom with a laser beam is often such that only two energy-eigenstates of the atom are affected. It is then sufficient to model this atom by a qubit \mathbb{C}^2 , rather than by its full Hilbert space. We will make extensive use of this in chapters 5 and 6.

Infinite Dimensional Hilbert Spaces

In infinite dimensional Hilbert space, the situation is similar, but slightly more involved. The only qualitative difference is that for some –but certainly not for all– Lie groups G , the universal cover \tilde{G} no longer suffices to describe all projective unitary representations, in which case a central deformation of the algebra of symmetries occurs.

Transformations A Lie group G of transformations acts on \mathcal{H} by a projective unitary representation. This is a continuous group homomorphism $G \rightarrow \mathcal{P}U(\mathcal{H})$ as before, but as we are no longer in a finite setting, we must pay some attention to matters of topology.

We equip $\mathcal{P}(\mathcal{H})$ with the *trace distance* or Kolmogorov-distance, which is the worst case difference in probability on a single event (projection) E , i.e.

$$d(\rho_\psi, \rho_\phi) = \sup_E (\rho_\psi(E) - \rho_\phi(E)).$$

The *strong topology* on $\mathcal{P}U(\mathcal{H})$ is the one defined by pointwise convergence on $\mathcal{P}(\mathcal{H})$. If we define the strong topology on $U(\mathcal{H})$ by pointwise convergence on \mathcal{H} , then the strong topology on $\mathcal{P}U(\mathcal{H})$ is the strongest one that makes $U(\mathcal{H}) \rightarrow \mathcal{P}U(\mathcal{H})$ continuous. It is with respect to these strong topologies that we require our (projective) representations to be continuous.

If the Hilbert space is infinite dimensional, one can still reduce projective unitary representations of G to linear representations of its universal cover \tilde{G} , provided that the second Lie algebra cohomology of $\text{Lie}(G)$ vanishes.

Theorem 2.5 *Let G be a connected Lie group such that $H^2(\text{Lie}(G), \mathbb{R}) = 0$, and let \tilde{G} be its universal cover. Then every projective unitary representation of G comes from a unitary representation of \tilde{G} .*

Lie algebras with vanishing second cohomology include the real semisimple Lie algebras, but also for example the Abelian Lie algebra \mathbb{R} , and the (infinite dimensional) Lie algebra of vector fields on any smooth manifold of dimension bigger than one.

If $H^2(\text{Lie}(G), \mathbb{R})$ is nonzero, then not all is lost. For every projective representation of G , there exists a central extension \hat{G} of G by the circle S^1 , and a linear unitary representation of \hat{G} that induces π . Since the Lie algebra of observables then follows $\text{Lie}(\hat{G})$ rather than $\text{Lie}(G)$, a ‘central deformation’ is said to have occurred.

Conserved Quantities Let us consider a group of transformations as a unitary representation π of G , having replaced the group by a bigger one if necessary. Then the observables corresponding to infinitesimal transformations $X \in \text{Lie}(G)$ are still precisely the derived Lie algebra elements $J(X) = i\dot{\pi}(X)$.

Recall that π is strongly continuous by definition, so that $\pi_\psi(g) = \pi(g)\psi$ is always a continuous map. A vector ψ is called *smooth* if π_ψ is smooth, and we denote the Fréchet space of smooth vectors by \mathcal{H}^∞ .

Theorem 2.6 (Stone, Gårding) *Let π be a unitary representation of a Lie group G on a Hilbert space \mathcal{H} . Then the smooth vectors \mathcal{H}^∞ lie dense in \mathcal{H} . For each $X \in \text{Lie}(G)$, the expression*

$$\dot{\pi}(X) = \exp(-tX) \frac{d}{dt} \pi(\exp(tX))$$

defines a closed, densely defined, skew-symmetric operator $\dot{\pi}(X)$. Its domain contains \mathcal{H}^∞ , and on this space, we have

$$\dot{\pi}([X, Y]) = [\dot{\pi}(X), \dot{\pi}(Y)].$$

Proof. See [Kna86, ch. III], but also [Par92, sec. 12,13]. □

In particular, the generator of the 1-parameter group $t \mapsto U_t$ of time translations is again the Hamiltonian H . In realistic models, its spectrum is usually bounded from below, but not from above.

A Point Particle on a Line Let us consider a single particle on the 1-dimensional Euclidean space \mathbb{R} . We should have a self-adjoint operator X on \mathcal{H} to represent its position. The group of Euclidean motions $(\mathbb{R}, +)$ should have a projective unitary representation π on \mathcal{H} , and its effect on X ought to be a shift, $\pi(s)^\dagger X \pi(s) = X + s\mathbf{1}$. Since $H^2(\mathbb{R}, \mathbb{R}) = 0$ and $\tilde{\mathbb{R}} = \mathbb{R}$, we may take π to be linear rather than projective. The observable $J(1) = i\dot{\pi}(1)$ is called the *momentum*, and denoted P . The above relation then differentiates to the *canonical commutation relation*

$$[X, P] = i\mathbf{1}. \tag{2.5}$$

One way to realize this is by setting $\mathcal{H} = L^2(\mathbb{R})$, $(X\psi)(x) = x\psi(x)$, and $(\pi(s)\psi)(x) = \psi(x - s)$. The associated momentum $P = -i\frac{d}{dx}$ does indeed satisfy $[X, P] = i\mathbf{1}$. A theorem of Stone and von Neumann says that this is essentially the only way to realize (the integrated version of) equation (2.5).

The PVM for X is not overly complicated; $E_X(S)$ is given by multiplication with the indicator function $\mathbf{1}_S$. That is, $(E_X(S)\psi)(x) = \mathbf{1}_S(x)\psi(x)$. This means that the probability distribution \mathbb{P}_ψ^X induced by ψ on the spectrum \mathbb{R} of X is just $\mathbb{P}_\psi^X(S) = \int_S |\psi(x)|^2 dx$. The function $|\psi|^2$ is the probability density of \mathbb{P}_ψ^X w.r.t. the Lebesgue measure.

The PVM for P can be derived from this. If F is the unitary Fourier transform on $L^2(\mathbb{R})$, then one has $P = F^\dagger(-x)F$, so that $E_P(S) = F^\dagger E_X(-S)F$. This means that $\mathbb{P}_\psi^P(S) = \int_S |F\psi|^2 dx$, and the squared Fourier transform of ψ yields the probability density of \mathbb{P}_ψ^P .

In sections 2.3.3 and 2.3.4, we will find two other ways to describe the point particle on a line.

A Point Particle in \mathbb{R}^3 Let us consider a single particle in the 3-dimensional Euclidean space \mathbb{R}^3 . We require 3 commuting self-adjoint operators X_1 , X_2 and X_3 to represent its position. The Euclidean motion group $\mathbb{R}^3 \rtimes \mathrm{O}(3)$ should have a projective unitary representation π on \mathcal{H} such that each $v \in \mathbb{R}^3$ shifts X_i to $X_i + v_i\mathbf{1}$, i.e. $\pi(v)^\dagger X_i \pi(v) = X_i + v_i\mathbf{1}$, and each rotation $R \in \mathrm{O}(3)$ rotates X_i to $\sum_{j=1}^3 R_{ij} X_j$, i.e. $\pi(R)^\dagger X_i \pi(R) = \sum_{j=1}^3 R_{ij} X_j$.

The Lie algebra $\mathbb{R}^3 \rtimes \mathfrak{so}(3)$ of the Euclidean motion group has zero second cohomology, so π induces a linear representation $\hat{\pi}$ on the level of Lie algebras. The standard basis e_x, e_y, e_z for \mathbb{R}^3 and τ_x, τ_y, τ_z for $\mathfrak{so}(3)$ gives rise to the momenta $P_k := J(e_k)$ and the angular momenta $L_k := J(\tau_k)$. The commutation relations between the P_k and L_l are then easily obtained from the Lie bracket on $\mathbb{R}^3 \rtimes \mathfrak{so}(3)$, and their commutator with the X_k follows from the requirement that the $\pi(g)$ shift or rotate the X_k . This yields the following extension of the canonical commutation relations:

$$\begin{aligned} [X_k, X_l] &= 0, & [X_k, P_l] &= i\delta_{kl}\mathbf{1}, & [P_k, P_l] &= 0 \\ [X_k, L_l] &= i\epsilon_{klm}X_m, & [L_k, L_l] &= i\epsilon_{klm}L_m, & [L_k, P_l] &= i\epsilon_{klm}P_m. \end{aligned} \quad (2.6)$$

This can be achieved by taking the Hilbert space $\mathcal{H} = L^2(\mathbb{R}^3)$, with the observables $(X_k\psi)(x) = x_k\psi(x)$ and the unitary representation $(\pi(g)\psi)(x) = \psi(g^{-1}x)$. The corresponding observables

$$P_k = -i\frac{\partial}{\partial x_k}, \quad L_k = -i\sum_{l,m} \epsilon_{klm}x_l \frac{\partial}{\partial x_m}$$

can easily be seen to obey 2.6.

Note that we have established all of this without any reference to the dynamics whatsoever. In other words, everything up to this point is independent of the particular Hamiltonian. A Hamiltonian usually takes the shape $H = \sum_{k=1}^3 P_k^2/2m + V(x)$, with V the *potential function*. This gives rise to the

Schrödinger equation

$$i \frac{d}{dt} \psi(x) = -\frac{1}{2m} \sum_{k=1}^3 \frac{\partial^2}{\partial x_k^2} \psi(x) + V(x) \psi(x).$$

It is the choice of Hamiltonian that determines which transformations are symmetries. For example, the hydrogen atom is described by the potential $V(x) = -1/\sqrt{x_1^2 + x_2^2 + x_3^2}$. In this case, the group of Euclidean motions $\mathbb{R}^3 \rtimes \mathrm{O}(3)$ is *not* a group of symmetries, but its subgroup $\mathrm{O}(3)$ of orthogonal transformations is. This means that the angular momenta L_k are conserved quantities, whereas the momenta P_k are not.

2.2 Quantum Probability Theory

Quantum probability theory (QPT) is a mathematical framework which generalizes both quantum mechanics and classical probability theory. It is ideally suited to describe *open* quantum systems, and in particular quantum measurement.

We first describe quantum probability theory on a finite dimensional Hilbert space \mathcal{H} . This allows us to focus on the algebra, postponing the functional analysis involved in the infinite dimensional case to later.

2.2.1 QPT on Finite Dimensional Hilbert Spaces.

A subalgebra \mathcal{A} of $\mathcal{B}(\mathcal{H})$ is called a $*$ -algebra if it is closed under the adjoint. In quantum probability theory, a system is modelled by a $*$ -algebra $\mathcal{A} \subset \mathcal{B}(\mathcal{H})$. The observables of the system are described by self-adjoint operators A in \mathcal{A} , and the states of the system are modelled as follows.

Definition 2.3 *Let \mathcal{H} be a finite dimensional Hilbert space, and \mathcal{A} a $*$ -subalgebra of $\mathcal{B}(\mathcal{H})$. Then a state on \mathcal{A} is a positive, normalized, linear functional $\rho : \mathcal{A} \rightarrow \mathbb{C}$. The space of states is denoted $\mathcal{S}(\mathcal{A})$.*

One should think of a state as a map that assigns to each observable A its expectation $\rho(A)$. Since the expectation should be \mathbb{R} -linear on the observables, we may as well extend it to a \mathbb{C} -linear map on all of \mathcal{A} . Positivity means that $A \geq 0$ implies $\rho(A) \geq 0$. This is reasonable; if A has nonnegative spectrum, then its expectation should also be nonnegative. Normalization, $\rho(\mathbf{1}) = 1$, means that the expectation of $\mathbf{1}$ is 1.

Quantum Mechanics

We have already seen that in our garden-variety quantum mechanics, every self-adjoint operator corresponds to an observable, so that $\mathcal{A} = \mathcal{B}(\mathcal{H})$. Any unit vector $\psi \in \mathcal{H}$ gives rise to a pure state $\rho_\psi(A) = \langle \psi | A \psi \rangle$, but it turns out that not all states are pure.

Theorem 2.7 *If \mathcal{H} is finite dimensional, then any state $\rho : \mathcal{B}(\mathcal{H}) \rightarrow \mathbb{C}$ can be written*

$$\rho(A) = \mathbf{tr}(RA)$$

for a unique positive operator R with $\mathbf{tr}(R) = 1$, called the density matrix.

Proof. Because $(A, B) \mapsto \mathbf{tr}(AB)$ is a perfect pairing, any linear functional $\rho : \mathcal{B}(\mathcal{H}) \rightarrow \mathbb{C}$ can be uniquely written as $\rho(A) = \mathbf{tr}(RA)$. If $E = |\psi\rangle\langle\psi|$ is the projection onto a 1-dimensional subspace $\mathbb{C}\psi$, then $\rho(E) = \langle\psi | R\psi\rangle$. Positivity of ρ therefore implies $R \geq 0$, and normalization corresponds to $\mathbf{tr}(R) = 1$. \square

The pure states ρ_ψ correspond to 1-dimensional projections $R = |\psi\rangle\langle\psi|$. They are the extreme points of the convex state space $\mathcal{S}(\mathcal{B}(\mathcal{H}))$.

Indeed, if we choose an orthonormal basis ψ_i of eigenvectors of R , then we can write $R = \sum_{i=1}^n p_i |\psi_i\rangle\langle\psi_i|$. Positivity implies $p_i \geq 0$, and normalization implies $\sum_{i=1}^n p_i = 1$. Any state on $\mathcal{B}(\mathcal{H})$ can therefore be written as a convex combination of vector states $\rho = \sum_i p_i \rho_{\psi_i}$, albeit not necessarily in a unique fashion.

This invites us to interpret a system in state ρ as one that is with probability p_i in state ρ_{ψ_i} . We will do so, and even elevate this to the more general *stochastic equivalence principle*, which holds for convex linear combinations of any two states ρ_1 and ρ_2 on any system \mathcal{A} .

A system in state $p\rho_1 + (1-p)\rho_2$ cannot be distinguished from one that is in state ρ_1 with probability p , and in state ρ_2 with probability $1-p$.

Classical Probability Space

Let (Ω, \mathbb{P}) be a finite probability space. A *random variable* is a map $f : \Omega \rightarrow \mathbb{C}$. It makes sense to identify two random variables f and f' if they are equal almost surely, $\mathbb{P}([f \neq f']) = 0$, and we write $L^\infty(\Omega, \mathbb{P})$ for the resulting algebra. The expectation defines a linear map $\mathbb{E} : L^\infty(\Omega, \mathbb{P}) \rightarrow \mathbb{C}$.

The Hilbert space $L^2(\Omega, \mathbb{P})$ is the same vector space, equipped with the inner product $\langle\phi | \psi\rangle = \mathbb{E}(\bar{\phi}\psi)$. Because $L^\infty(\Omega, \mathbb{P})$ acts on $L^2(\Omega, \mathbb{P})$ by multiplication, we can consider it as a commutative $*$ -subalgebra of $\mathcal{B}(L^2(\Omega, \mathbb{P}))$. This places us in the framework of quantum probability theory, with $\mathcal{A} = L^\infty(\Omega, \mathbb{P})$, $\mathcal{H} = L^2(\Omega, \mathbb{P})$, and $\rho = \mathbb{E}$. Let us describe the state space $\mathcal{S}(\mathcal{A})$.

Theorem 2.8 *Any finite probability space (Ω, \mathbb{P}) gives rise to a finite dimensional Hilbert space $L^2(\Omega, \mathbb{P})$, and a $*$ -subalgebra $L^\infty(\Omega, \mathbb{P})$ of $\mathcal{B}(L^2(\Omega, \mathbb{P}))$. If \mathbb{P}' is absolutely continuous w.r.t. \mathbb{P} , then it defines a state $\mathbb{E}' : L^\infty(\Omega, \mathbb{P}) \rightarrow \mathbb{C}$. All states arise in this way.*

Proof. The set $\{\delta_\omega ; \omega \in \Omega, \mathbb{P}(\omega) \neq 0\}$ constitutes a basis of $L^\infty(\Omega, \mathbb{P})$, so that any state ρ can be written $\rho = \sum p'_\omega \delta_\omega^*$. Positivity and normalization amount to $p'_\omega \geq 0$ and $\sum p'_\omega = 1$ respectively. The numbers p'_ω thus define a probability measure $\mathbb{P}' \ll \mathbb{P}$, and $\rho = \mathbb{E}'$ is its associated expectation. \square

Operations

We wish to describe transformations of a system $\mathcal{A} \subset \mathcal{B}(\mathcal{H})$ to a system $\mathcal{B} \subset \mathcal{B}(\mathcal{H}')$. In the context of quantum probability theory, these are called *operations*. Each operation is determined by a map $\tau : \mathcal{S}(\mathcal{A}) \rightarrow \mathcal{S}(\mathcal{B})$. We investigate which maps can reasonably be interpreted as operations.

We interpret the convex combination $p\rho_1 + (1-p)\rho_2 \in \mathcal{S}(\mathcal{A})$ as a system that is in state ρ_1 with probability p , and in state ρ_2 with probability $(1-p)$. This means that after the operation, the system \mathcal{B} must be described as being in state $\tau(\rho_1)$ with probability p , and in state $\tau(\rho_2)$ with probability $1-p$. Thus

$$\tau(p\rho_1 + (1-p)\rho_2) = p\tau(\rho_1) + (1-p)\tau(\rho_2).$$

This means that τ extends to a *linear* map $\tau : \mathcal{A}^* \rightarrow \mathcal{B}^*$. We can consider it as the dual $\tau = T^*$ of a linear map $T : \mathcal{B} \rightarrow \mathcal{A}$, and will do so from now on. The requirement that T^* map states to states implies that $T(B) \geq 0$ if $B \geq 0$ (it preserves positivity), and that $T(\mathbf{1}) = \mathbf{1}$ (normalization).

Definition 2.4 *Let \mathcal{A} and \mathcal{B} be $*$ -subalgebras of $\mathcal{B}(\mathcal{H})$ and $\mathcal{B}(\mathcal{H}')$ respectively. A map $T : \mathcal{B} \rightarrow \mathcal{A}$ is called positive if it is linear, if $B \geq 0$ implies $T(B) \geq 0$, and if $T(\mathbf{1}) = \mathbf{1}$.*

One would be tempted to interpret each positive map as an operation, but this is *not* what one should do. There is one more requirement that we must impose on T in order for it to qualify as an operation.

If $T : \mathcal{B} \rightarrow \mathcal{A}$ describes an operation from \mathcal{A} to \mathcal{B} , then the map $T \otimes \text{id}_n : \mathcal{B} \otimes \mathcal{B}(\mathbb{C}^n) \rightarrow \mathcal{A} \otimes \mathcal{B}(\mathbb{C}^n)$ describes the act of transforming \mathcal{A} to \mathcal{B} , while doing absolutely nothing on the system $\mathcal{B}(\mathbb{C}^n)$. Surely this should map states on $\mathcal{A} \otimes \mathcal{B}(\mathbb{C}^n)$ to states on $\mathcal{B} \otimes \mathcal{B}(\mathbb{C}^n)$, and we therefore require that $T \otimes \text{id}_n$ be a positive map for all $n \in \mathbb{N}$.

Definition 2.5 *A map $T : \mathcal{B} \rightarrow \mathcal{A}$ is called completely positive if*

$$T \otimes \text{id}_n : \mathcal{B} \otimes \mathcal{B}(\mathbb{C}^n) \rightarrow \mathcal{A} \otimes \mathcal{B}(\mathbb{C}^n)$$

is positive for all $n \in \mathbb{N}$.

Perhaps surprisingly, this is not automatic. There exist maps which are positive, but not completely positive. We shall model operations from \mathcal{A} to \mathcal{B} by completely positive maps $T : \mathcal{B} \rightarrow \mathcal{A}$. (CP-maps for short.)

Operations in Classical Probability Let (Ω, \mathbb{P}) and (Ω', \mathbb{P}') be finite probability spaces. Then a *transition kernel* is a positive function P on $\Omega' \times \Omega$ such that $\sum_{\Omega'} P(\omega', \omega) \mathbb{P}'(\omega') = 1$ for all $\omega \in \Omega$. One should think of $P(\omega', \omega) \mathbb{P}'(\omega')$ as the probability that a system in state ω makes a transition to ω' .

An operation $T : L^\infty(\Omega', \mathbb{P}') \rightarrow L^\infty(\Omega, \mathbb{P})$ is related to such a kernel. Indeed, if we write $Tf(\omega) = \sum_{\Omega'} P(\omega', \omega) f(\omega') \mathbb{P}'(\omega')$, then positivity of T entails $P(\omega', \omega) \geq 0$, and normalization $T(\mathbf{1}) = \mathbf{1}$ translates to $\sum_{\Omega'} P(\omega', \omega) \mathbb{P}'(\omega') = 1$. An operation is precisely an equivalence class of kernels, if we deem two kernels to be equivalent when their difference is nonzero only on those (ω', ω) with either $\mathbb{P}(\omega) = 0$ or $\mathbb{P}'(\omega') = 0$.

Operations from Quantum to Classical Systems A completely positive map $T : L^\infty(\Omega, \mathbb{P}) \rightarrow \mathcal{B}(\mathcal{H})$ is determined by its values on δ_ω , i.e. by the map $E : \Omega \rightarrow \mathcal{B}(\mathcal{H})$ defined as $E(\omega) = T(\delta_\omega)$.

Definition 2.6 A positive operator valued measure or POVM is a map $E : \Omega \rightarrow \mathcal{B}(\mathcal{H})$ that satisfies

- $E(\omega) \geq 0$ for all $\omega \in \Omega$
- $\sum_{\Omega} E(\omega) = \mathbf{1}$.

The map $E(\omega) = T(\delta_\omega)$ is a POVM; positivity of T implies the first requirement, and normalization the second. Conversely, any POVM that is measurable w.r.t. \mathbb{P} determines a CP-map $T : L^\infty(\Omega, \mathbb{P}) \rightarrow \mathcal{B}(\mathcal{H})$. A CP-map is a *-homomorphism if and only if its POVM is projection valued.

Operations between Quantum Systems Any *-homomorphism of algebras is a CP-map. In particular, the transformations $\alpha_U(A) = U^\dagger A U$ of $\mathcal{B}(\mathcal{H})$ are operations in the sense of quantum probability.

Another type of operation is considering a subsystem as part of a bigger whole. If \mathcal{H} is a Hilbert subspace of \mathcal{K} , then it comes with the inclusion $V : \mathcal{H} \hookrightarrow \mathcal{K}$, and the adjoint projection $V^\dagger : \mathcal{K} \rightarrow \mathcal{H}$. The CP-map $T : \mathcal{B}(\mathcal{K}) \rightarrow \mathcal{B}(\mathcal{H})$ corresponding to this inclusion is given by $T(A) = V^\dagger A V$.

According to the following classification theorem, any CP-map is a combination of a *-homomorphism and an inclusion.

Theorem 2.9 (Stinespring) Let \mathcal{H} and \mathcal{H}' be Hilbert spaces of finite dimension, let \mathcal{A} be a *-subalgebra of $\mathcal{B}(\mathcal{H}')$, and let T be a CP-map $\mathcal{A} \rightarrow \mathcal{B}(\mathcal{H})$. Then there exists a finite dimensional Hilbert space \mathcal{K} , an inclusion $V : \mathcal{H} \hookrightarrow \mathcal{K}$ and a *-homomorphism $\pi : \mathcal{A} \rightarrow \mathcal{B}(\mathcal{K})$ such that $T(A) = V^\dagger \pi(A) V$.

$$\begin{array}{ccc} \mathcal{B}(\mathcal{K}) & & \\ \pi \uparrow & \searrow V^\dagger \cdot V & \\ \mathcal{A} & \xrightarrow{T} & \mathcal{B}(\mathcal{H}) \end{array}$$

Proof. Let $\mathcal{K}_0 = \mathcal{A} \otimes \mathcal{H}$. Then the inclusion $V_0 : \mathcal{H} \rightarrow \mathcal{K}_0$ is given by $\psi \mapsto \mathbf{1} \otimes \psi$, and if $A \in \mathcal{A}$, then $\pi_0(A) : \mathcal{K}_0 \rightarrow \mathcal{K}_0$ is defined by $B \otimes \psi \mapsto AB \otimes \psi$. We equip \mathcal{K}_0 with the sesquilinear form defined by $\langle B \otimes \psi | B' \otimes \psi' \rangle = \langle \psi | T(B^\dagger B') \psi' \rangle_{\mathcal{H}}$.

Let $F = \sum_{i=1}^k B_i \otimes \psi_i$ be an arbitrary element of \mathcal{K}_0 . In terms of the map $T \otimes \text{id}_k : \mathcal{A} \otimes \mathcal{B}(\mathbb{C}^k) \rightarrow \mathcal{B}(\mathcal{H} \otimes \mathbb{C}^k)$, we can write

$$\langle F | F \rangle = \left\langle \Psi | T \otimes \text{id}_k \left(\underline{B}^\dagger \underline{B} \right) \Psi \right\rangle, \quad (2.7)$$

with $\Psi = \sum_{i=1}^k \psi_i \otimes e_i$ in $\mathcal{H} \otimes \mathbb{C}^k$ and $\underline{B} = \sum_{i=1}^k B_i \otimes |e_1\rangle\langle e_i|$ in $\mathcal{A} \otimes \mathcal{B}(\mathbb{C}^k)$. Because $T \otimes \text{id}_k$ is a positive map, we conclude that $\langle F | F \rangle \geq 0$. Our sesquilinear form is positive semidefinite.

We can therefore define the Hilbert space $\mathcal{K} = \mathcal{K}_0/\mathcal{N}$, with \mathcal{N} the kernel of $\langle \cdot | \cdot \rangle$. The map V_0 is an isometry w.r.t. $\langle \cdot | \cdot \rangle$, and therefore induces an inclusion $V : \mathcal{H} \hookrightarrow \mathcal{K}$ of Hilbert spaces. Its adjoint projection is given by $V^\dagger(B \otimes \psi) = T(B)\psi$.

The map $\pi_0 : \mathcal{A} \rightarrow \mathcal{B}(\mathcal{K}_0)$ is a contraction. Indeed, if we replace F by $\pi_0(A)F$ in equation (2.7), then on the right hand side, $\underline{B}^\dagger \underline{B}$ changes into the expression $\underline{B}^\dagger(A^\dagger A \otimes \mathbf{1})\underline{B}$. Since this is dominated by $\|A\|^2 \underline{B}^\dagger \underline{B}$, we have $\langle \pi_0(A)F | \pi_0(A)F \rangle \leq \|A\|^2 \langle F | F \rangle$, and the operator norm of $\pi_0(A)$ does not exceed that of A .

In particular, \mathcal{N} is invariant under $\pi_0(A)$, so that we have a map $\pi(A) : \mathcal{K} \rightarrow \mathcal{K}$. It is then straightforward to check that π is a $*$ -homomorphism, and that $V^\dagger \pi(A)V = T(A)$. \square

2.2.2 QPT on General Hilbert Spaces

We define quantum probability theory in the infinite dimensional setting, which requires a little bit of functional analysis.

Von Neumann Algebras

We have seen that in quantum mechanics, all the structure of a system is encoded in its algebra $\mathcal{B}(\mathcal{H})$ of bounded operators. A subalgebra of $\mathcal{B}(\mathcal{H})$ is called a *von Neumann algebra* if it is closed not only under the adjoint $A \mapsto A^\dagger$, but also in the weak topology. In quantum probability theory, a system is modelled by a von Neumann algebra \mathcal{A} .

A self-adjoint, possibly unbounded operator A on \mathcal{H} is said to be *affiliated* to \mathcal{A} if all its spectral projections are in \mathcal{A} . Observables of the system are modelled by self-adjoint operators affiliated to \mathcal{A} .

If \mathcal{A} and \mathcal{B} are von Neumann algebras, then we will denote by $\mathcal{A} \otimes \mathcal{B}$ the von Neumann algebra generated by the algebraic tensor product.

Normal States

A state is again described by a normalized positive linear map $\rho : \mathcal{A} \rightarrow \mathbb{C}$, but it is wise to impose a continuity condition. If A is affiliated to \mathcal{A} , then we would like its spectral measure $E : \mathbf{Spec}(A) \rightarrow \mathcal{A}$ to induce a probability measure $\mathbb{P}_\rho = \rho \circ E$ on $\mathbf{Spec}(A)$. But although E is σ -additive, i.e. $\sum_{\mathbb{N}} E(S_i) = E(\cup_{\mathbb{N}} S_i)$ in the weak topology, this need not hold for $\mathbb{P}_\rho = \rho \circ E$, unless we impose the following continuity requirement.

Definition 2.7 *A state $\rho : \mathcal{A} \rightarrow \mathbb{C}$ is called normal if it is weakly continuous on the unit ball $\{A \in \mathcal{A}; \|A\| \leq 1\}$.*

A normal state $\rho \in \mathcal{S}(\mathcal{A})$ induces a probability measure \mathbb{P}_ρ on the spectrum of any self-adjoint operator affiliated to \mathcal{A} , so that we have our beloved map $(\rho, A) \mapsto \mathbb{P}_\rho$. With respect to \mathbb{P}_ρ , we have of course $\mathbb{E}(\lambda) = \rho(A)$.

Normal States in Quantum Mechanics The normal states of an ordinary quantum system $\mathcal{B}(\mathcal{H})$ can again be described by a density matrix.

Theorem 2.10 *Let \mathcal{H} be a separable Hilbert space. Then any normal state ρ on $\mathcal{B}(\mathcal{H})$ can be written*

$$\rho(A) = \mathbf{tr}(RA)$$

for a unique positive trace-class operator R on \mathcal{H} with $\mathbf{tr}(R) = 1$.

Proof. See for example chapter 7 of [KR86]. □

A positive normalized trace-class operator admits a basis of eigenvectors ψ_i with nonnegative eigenvalues p_i summing to 1. In other words, any state can be written as a countably infinite convex combination of pure states, $\rho = \sum_{i=1}^{\infty} p_i \rho_{\psi_i}$.

Normal States in Classical Probability A classical probability space, denoted $(\Omega, \Sigma, \mathbb{P})$, gives rise to the algebra $L^\infty(\Omega, \Sigma, \mathbb{P})$ of essentially bounded random variables on Ω , i.e. bounded measurable functions up to the equivalence \sim , with $f \sim g$ if $f = g$ almost surely.

Since $L^\infty(\Omega, \Sigma, \mathbb{P})$ acts on the Hilbert space $L^2(\Omega, \Sigma, \mathbb{P})$ by multiplication, we can regard it as a commutative $*$ -subalgebra of $\mathcal{B}(L^2(\Omega, \Sigma, \mathbb{P}))$, which turns out to be weakly closed.

Theorem 2.11 *The algebra $L^\infty(\Omega, \Sigma, \mathbb{P})$ is a commutative von Neumann subalgebra of $\mathcal{B}(L^2(\Omega, \Sigma, \mathbb{P}))$. If \mathbb{P}' is absolutely continuous w.r.t. \mathbb{P} , then the expectation $\mathbb{E}'(f) = \int_{\Omega} f(\omega) \mathbb{P}'(d\omega)$ constitutes a normal state on $L^\infty(\Omega, \Sigma, \mathbb{P})$, and every normal state arises in this way.*

Proof. See e.g. [Maa03] and the references therein. □

In short, a classical probability space $(\Omega, \Sigma, \mathbb{P})$ gives rise to a commutative von Neumann algebra, and a normal state corresponds to a probability measure $\mathbb{P}' \ll \mathbb{P}$. Essentially every¹ commutative von Neumann algebra arises in this way.

Operations

The definition of a completely positive map remains unaltered. A linear map $T : \mathcal{B} \rightarrow \mathcal{A}$ between von Neumann algebras is called positive if $T(\mathbf{1}) = \mathbf{1}$ and if $B \geq 0$ implies $T(B) \geq 0$. It is called completely positive if $T \otimes \text{id}_n : \mathcal{B} \otimes \mathcal{B}(\mathbb{C}^n) \rightarrow \mathcal{A} \otimes \mathcal{B}(\mathbb{C}^n)$ is positive for all $n \in \mathbb{N}$. This already implies that $T \otimes \text{id} : \mathcal{B} \otimes \mathcal{B}(\mathcal{H}) \rightarrow \mathcal{A} \otimes \mathcal{B}(\mathcal{H})$ is positive for arbitrary \mathcal{H} , so that we need not add this as an extra requirement.

We do however require T^* to map normal states to normal states. We therefore model operations by *weakly continuous* CP-maps.

¹To be more precise, a state ρ is called *faithful* if $A \geq 0$ and $\rho(A) = 0$ implies $A = 0$. Any commutative von Neumann algebra \mathcal{A} with a faithful normal state ρ is isomorphic to $L^\infty(\Omega, \Sigma, \mathbb{P})$, and the state ρ corresponds to the expectation w.r.t. \mathbb{P} .

Operations in Classical Probability Let $(\Omega, \Sigma, \mathbb{P})$ and $(\Omega', \Sigma', \mathbb{P}')$ be two probability spaces. A *transition kernel* is a positive measurable function P on $\Omega' \times \Omega$ such that $\int_{\Omega'} P(\omega', \omega) \mathbb{P}'(d\omega') = 1$ for all $\omega \in \Omega$. We will identify two transition kernels if they are equal almost surely w.r.t. $\mathbb{P}' \times \mathbb{P}$. One should think of $P(\omega', \omega) \mathbb{P}'(d\omega')$ as the probability that the system ends up in $d\omega'$, provided that it starts in ω .

Completely positive maps between classical probability spaces correspond precisely to transition kernels.

Theorem 2.12 *A transition kernel P gives rise to the weakly continuous CP-map $T_P : L^\infty(\Omega', \Sigma', \mathbb{P}') \rightarrow L^\infty(\Omega, \Sigma, \mathbb{P})$ by $T_P(f)(\omega) = \int_{\Omega'} f(\omega') p(\omega, \omega') \mathbb{P}'(d\omega')$. Every weakly continuous CP-map $T : L^\infty(\Omega', \Sigma', \mathbb{P}') \rightarrow L^\infty(\Omega, \Sigma, \mathbb{P})$ is of this form.*

Proof. We prove the second statement. The map T gives rise to the normal state ρ on $L^\infty(\Omega', \Sigma', \mathbb{P}') \otimes L^\infty(\Omega, \Sigma, \mathbb{P})$ defined by $\rho(f' \otimes f) = \mathbb{E}_{\mathbb{P}}(fT(f'))$. According to theorem 2.11, this corresponds to a probability measure on $\Omega' \times \Omega$ which is absolutely continuous w.r.t. $\mathbb{P}' \times \mathbb{P}$. Its Radon-Nikodym derivative, well defined up to equivalence, is a transition kernel P . The requirement $\int_{\Omega'} P(\omega', \omega) \mathbb{P}'(d\omega') = 1$ corresponds to the fact that \mathbb{P} is the marginal probability distribution on Ω , as is clear from $\rho(\mathbf{1} \otimes f) = \mathbb{E}_{\mathbb{P}}(f)$. \square

Operations from Quantum to Classical Systems We have seen that a projection valued measure $E : \Sigma \rightarrow \mathcal{B}(\mathcal{H})$ gives rise to a *-algebra homomorphism $I_E : L^\infty(\Omega, \Sigma, \mathbb{P}) \rightarrow \mathcal{B}(\mathcal{H})$, by way of

$$I_E(f) := \int_{\Omega} f(\omega) E(d\omega). \quad (2.8)$$

If $E \ll \mathbb{P}$, this factors through a weakly continuous homomorphism $I_E : L^\infty(\Omega, \Sigma, \mathbb{P}) \rightarrow \mathcal{B}(\mathcal{H})$ of von Neumann algebras. Conversely, any weakly continuous homomorphism $T : L^\infty(\Omega, \Sigma, \mathbb{P}) \rightarrow \mathcal{B}(\mathcal{H})$ of von Neumann algebras gives rise to a PVM by way of $E(S) = T(\mathbf{1}_S)$.

In order to capture all CP-maps, not just the homomorphisms, we introduce the notion of a Positive Operator Valued Measure or POVM (cf. def. 2.6).

Definition 2.8 *Let (Ω, Σ) be a measurable space, and \mathcal{H} a Hilbert space. Then a POVM is a map $E : \Sigma \rightarrow \mathcal{B}(\mathcal{H})$ with the following properties.*

- $E(S) \geq 0$ for all $S \in \Sigma$.
- $E(\Omega) = \mathbf{1}$.
- If S_i is a sequence of disjoint sets in Σ , then $\sum_{i=1}^{\infty} E(S_i) = E(\bigcup_{i=1}^{\infty} S_i)$ in the weak topology.

A weakly continuous CP-map $T : L^\infty(\Omega, \Sigma, \mathbb{P}) \rightarrow \mathcal{B}(\mathcal{H})$ corresponds to a POVM $E \ll \mathbb{P}$ by $T(f) = \int_{\Omega} f(\omega) E(d\omega)$. The POVM is a PVM (cf. def. 2.2) if and only if T is a homomorphism.

Operations in Quantum Systems Although every homomorphism of von Neumann algebras is a CP-map, it is certainly not true that every CP-map is a homomorphism. We do however have the following theorem.

Theorem 2.13 *Every CP-map $T : \mathcal{B} \rightarrow \mathcal{A}$ with a completely positive inverse is an isomorphism.*

See page 76 for the proof. The group of invertible operations of a system \mathcal{A} is thus simply $\text{Aut}^{\text{ct}}(\mathcal{A})$, the group of weakly continuous automorphisms. Its action on the state space $\mathcal{S}(\mathcal{A})$ equips it with the topology of pointwise convergence in trace distance. (See page 12.) From the point of view of quantum probability theory, the action of a Lie group G should therefore be described by a continuous homomorphism $\pi : G \rightarrow \text{Aut}^{\text{ct}}(\mathcal{A})$. The following theorem shows that $\text{Aut}^{\text{ct}}(\mathcal{B}(\mathcal{H})) = \mathcal{P}U(\mathcal{H})$.

Theorem 2.14 (Wigner) *Let \mathcal{H} be a separable Hilbert space. Then every weakly continuous automorphism α of $\mathcal{B}(\mathcal{H})$ is inner, and can be written $\alpha(A) = U^\dagger A U$ for some unitary U .*

Proof. Because α^* maps normal states to normal states, we can consider it as a continuous trace-preserving \mathbb{C} -linear map on the trace-class operators $\mathcal{T}(\mathcal{H})$, cf. theorem 2.10. Because α^* maps R to $\alpha^{-1}(R)$, it restricts to a map on the pure state space $\mathcal{P}(\mathcal{H})$, and α is determined by this restriction because of continuity. The map on $\mathcal{P}(\mathcal{H})$ preserves the pairing $(\mathbb{C}\phi, \mathbb{C}\psi) = |\langle \phi | \psi \rangle| / \|\phi\| \|\psi\|$ because it preserves the trace. According to Wigner's classification of the automorphisms of $\mathcal{P}(\mathcal{H})$, any such map must be given by $\alpha^*(|\psi\rangle\langle\psi|) = U|\psi\rangle\langle\psi|U^\dagger$ with U either unitary or anti-unitary. The map $U \cdot U^\dagger$ can only be extended to a \mathbb{C} -linear map on $\mathcal{T}(\mathcal{H})$ if U is unitary, or of course if \mathcal{H} is one-dimensional, in which case the theorem is trivially true. \square

An action of G by weakly continuous CP-maps on $\mathcal{B}(\mathcal{H})$ is therefore precisely a projective unitary representation. In particular, *invertible* time evolution on a *closed* quantum system is still described by a Hamiltonian, and we do not gain anything new.

Non-invertible time evolution on an *open* quantum system however should be described by a continuous homomorphism of the semigroup $(\mathbb{R}^{\geq 0}, +)$ into the semigroup of weakly continuous CP-maps on $\mathcal{B}(\mathcal{H})$, which will in general *not* be automorphic. CP-maps are classified by the following theorem.

Theorem 2.15 (Stinespring) *Let \mathcal{A} be a von Neumann algebra, \mathcal{H} a Hilbert space, and let T be a CP-map $\mathcal{A} \rightarrow \mathcal{B}(\mathcal{H})$. Then there exists a Hilbert space \mathcal{K} , an inclusion $V : \mathcal{H} \hookrightarrow \mathcal{K}$ and a $*$ -homomorphism $\pi : \mathcal{A} \rightarrow \mathcal{B}(\mathcal{K})$ such that $T(A) = V^\dagger \pi(A) V$.*

$$\begin{array}{ccc} \mathcal{B}(\mathcal{K}) & & \\ \pi \uparrow & \searrow V^\dagger \cdot V & \\ \mathcal{A} & \xrightarrow{T} & \mathcal{B}(\mathcal{H}) \end{array}$$

If we impose that \mathcal{K} be the closure of $\pi(\mathcal{A})\mathcal{H}$, then (π, V, \mathcal{K}) is unique up to unitary transformation, and T is weakly continuous if and only if π is.

Proof. See Stinespring's paper [Sti55]. \square

Summarizing

We arrive at the following picture of quantum probability theory. A system is modelled by a von Neumann algebra \mathcal{A} of bounded operators on a Hilbert space \mathcal{H} . An observable is modelled by a self-adjoint operator affiliated to \mathcal{A} . A state on this system is described by a normalized positive linear map $\rho : \mathcal{A} \rightarrow \mathbb{C}$ that is weakly continuous on the unit ball, and an operation from \mathcal{A} to \mathcal{B} is described by a weakly continuous completely positive map $T : \mathcal{B} \rightarrow \mathcal{A}$.

2.3 Quantum Stochastic Differential Equations

This concludes our description of quantum probability spaces in general. We will now focus on a single model that will be of particular relevance to this thesis: the second quantized electromagnetic field. Its Hilbert space $\mathcal{F}(L^2(\mathbb{R}^+))$ is the Fock space of $L^2(\mathbb{R}^+)$.

2.3.1 Fock Space

If \mathcal{H} is a Hilbert space, then so is its n -fold symmetric tensor power $\mathcal{H}^{\otimes_s n}$. We set $\mathcal{H}^{\otimes_s 0} := \mathbb{C}$, and define the Fock space over \mathcal{H} as follows.

Definition 2.9 *The (symmetric) Fock space $\mathcal{F}(\mathcal{H})$ over a Hilbert space \mathcal{H} is defined as*

$$\mathcal{F}(\mathcal{H}) = \bigoplus_{n=0}^{\infty} \mathcal{H}^{\otimes_s n},$$

the Hilbert space of sequences $n \mapsto v_n$ with $v_n \in \mathcal{H}^{\otimes_s n}$ and $\sum_n \langle v_n | v_n \rangle < \infty$.

In many respects, taking \mathcal{H} into $\mathcal{F}(\mathcal{H})$ can be seen as exponentiating a Hilbert space. For example, the set of *exponential vectors* $\{e(v); v \in \mathcal{H}\}$, with

$$e(v) := \bigoplus_{n=0}^{\infty} \frac{1}{\sqrt{n!}} v^{\otimes n},$$

is linearly independent and dense in $\mathcal{F}(\mathcal{H})$ (see [Par92]). They are easily seen to satisfy

$$\langle e(v) | e(w) \rangle_{\mathcal{F}(\mathcal{H})} = \exp(\langle v | w \rangle_{\mathcal{H}}).$$

We will denote the vectors $v^{\otimes n}$ by $|v\rangle_n$. The normalized exponential vectors are called *coherent states*. We denote them by $|v\rangle := e^{-\frac{1}{2}\langle v | v \rangle} e(v)$.

Proposition 2.16 *There is a natural isomorphism of Hilbert spaces*

$$\mathcal{F}(\mathcal{H}_1 \oplus \mathcal{H}_2) \simeq \mathcal{F}(\mathcal{H}_1) \otimes \mathcal{F}(\mathcal{H}_2),$$

under which $e(v_1 \oplus v_2) \simeq e(v_1) \otimes e(v_2)$.

Proof. Since the exponential vectors are dense and linearly independent, the expression $e(v_1 \oplus v_2) \mapsto e(v_1) \otimes e(v_2)$ yields a densely defined linear map. Since both $\langle e(v_1 \oplus v_2) | e(w_1 \oplus w_2) \rangle$ and $\langle e(v_1) \otimes e(v_2) | e(w_1) \otimes e(w_2) \rangle$ equal $\exp(\langle v_1 | w_1 \rangle + \langle v_2 | w_2 \rangle)$, it extends to a unitary isomorphism. \square

By the same token, it is clear that an isometry $U : \mathcal{H}_1 \rightarrow \mathcal{H}_2$ induces the isometry $\mathcal{F}(U) : \mathcal{F}(\mathcal{H}_1) \rightarrow \mathcal{F}(\mathcal{H}_2)$ defined by $\mathcal{F}(U)e(v) = e(Uv)$. Because $\mathcal{F}(U' \circ U) = \mathcal{F}(U') \circ \mathcal{F}(U)$, one can view \mathcal{F} as a functor from the category of Hilbert spaces to itself, taking direct sums into tensor products². In particular, $\mathcal{F}(\mathcal{H})$ carries a unitary representation of $U(\mathcal{H})$.

2.3.2 Weyl Operators, Fields and Momenta

The unitary representation $U \mapsto \mathcal{F}(U)$ of $U(\mathcal{H})$ extends to a *projective* unitary representation of the group $U(\mathcal{H}) \ltimes \mathcal{H}$ with product $(U, v)(U', v') = (UU', Uv' + v)$.

Proposition 2.17 *The Weyl operators $W(U, v)$, defined by*

$$W(U, v)e(w) = \exp\left(-\frac{1}{2}\|v\|^2 - \langle v | Uw \rangle\right)e(Uw + v), \quad (2.9)$$

constitute a projective unitary representation of $U(\mathcal{H}) \ltimes \mathcal{H}$, satisfying

$$W(U, v)W(U', v') = \exp(-i\text{Im}\langle v | Uv' \rangle)W(UU', Uv' + v). \quad (2.10)$$

Proof. Since $\langle W(U, v)ew | W(U, v)e(w') \rangle = \langle ew | ew' \rangle$, equation (2.9) does indeed define a unitary operator. Equation (2.10) is then readily verified on the dense set of exponential vectors. Continuity essentially follows from continuity of the map $v \mapsto e(v)$, see [Par92] for details. \square

We will denote $W(\mathbf{1}, v)$ by $W(v)$, and designate $W(U, 0)$ by its old name $\mathcal{F}(U)$. If \mathcal{H} is finite dimensional, we find a projective Lie algebra homomorphism $\dot{W} : \mathfrak{su}(\mathcal{H}) \ltimes \mathcal{H} \rightarrow \mathfrak{su}(\mathcal{F}(\mathcal{H}))$, cf. section 2.1.2. Rather than making this precise for infinite dimensional \mathcal{H} , we will restrict attention to Lie subgroups of $U(\mathcal{H}) \ltimes \mathcal{H}$ of dimension 1 or 2.

Associated Observables

Every self-adjoint operator H gives rise to a one-parameter Lie group in $U(\mathcal{H}) \ltimes \mathcal{H}$, as does every vector $v \in \mathcal{H}$. The restriction of W to these one-parameter groups is a linear unitary representation. We can therefore define the skew-symmetric closed operators

$$\dot{\mathcal{F}}(H) = \mathcal{F}(e^{itH}) \frac{d}{dt} \mathcal{F}(e^{-itH}) \quad (2.11)$$

and

$$\dot{W}(v) = W(-tv) \frac{d}{dt} W(tv). \quad (2.12)$$

The observable $\lambda(H) := i\dot{\mathcal{F}}(H)$ is called the differential second quantization of H . We will not be needing commutation relations involving it. (See however [Par92].) The operators $J(v) = i\dot{W}(v)$ are called fields and momenta.

²In the words of E. Nelson: ‘Quantization is a mystery, but second quantization is a functor.’

They satisfy the *canonical commutation relation*

$$[J(u), J(v)] = 2i \operatorname{Im}\langle u | v \rangle \mathbf{1}, \quad (2.13)$$

as can be seen by differentiating (2.10). Since every exponential vector is in the domain of every product $J(v_1) \dots J(v_n)$, the l.h.s. of equation (2.13) is densely defined.

Fields and Momenta

As is the case in Hamiltonian mechanics, the division of the $J(v)$ into fields and momenta is to a certain extent arbitrary. If we choose a completely real subspace $\mathcal{H}_{\mathbb{R}}$, then we define for $f \in \mathcal{H}_{\mathbb{R}}$ the *fields* to be

$$\Phi(f) = -J(f),$$

and their *conjugate momenta*

$$\Pi(f) = -J(if).$$

They satisfy the commutation relations

$$[\Phi(f), \Phi(g)] = 0, \quad [\Phi(f), \Pi(g)] = 2i\langle f | g \rangle, \quad [\Pi(f), \Pi(g)] = 0.$$

Creation and Annihilation Operators

Even without a completely real subspace, we can introduce for any $u \in \mathcal{H}$ the creation and annihilation operators

$$a^\dagger(u) = -\frac{1}{2}(J(iu) + iJ(u)), \quad a(u) = -\frac{1}{2}(J(iu) - iJ(u)), \quad (2.14)$$

respectively, satisfying

$$[a(u), a(v)] = 0, \quad [a(u), a^\dagger(v)] = \langle u | v \rangle \mathbf{1}, \quad [a^\dagger(u), a^\dagger(v)] = 0.$$

By sandwiching (2.9) between exponential vectors and applying equation (2.12), we obtain:

$$\langle ev | a(u)e(v') \rangle = \langle u | v' \rangle \langle ev | e(v') \rangle, \quad (2.15)$$

$$\langle ev | a^\dagger(u)e(v') \rangle = \langle v | u \rangle \langle ev | e(v') \rangle. \quad (2.16)$$

We will also have use for the second order correlation functions

$$\langle ev | a^\dagger(u_1)a^\dagger(u_2)e(v') \rangle = \langle v | u_1 \rangle \langle v | u_2 \rangle \langle ev | e(v') \rangle \quad (2.17)$$

$$\langle ev | a(u_1)a(u_2)e(v') \rangle = \langle u_1 | v' \rangle \langle u_2 | v' \rangle \langle ev | e(v') \rangle \quad (2.18)$$

$$\langle ev | a^\dagger(u_1)a(u_2)e(v') \rangle = \langle v | u_1 \rangle \langle u_2 | v' \rangle \langle ev | e(v') \rangle \quad (2.19)$$

$$\begin{aligned} \langle ev | a(u_1)a^\dagger(u_2)e(v') \rangle &= \langle u_1 | v' \rangle \langle v | u_2 \rangle \langle ev | e(v') \rangle \\ &+ \langle u_1 | u_2 \rangle \langle ev | e(v') \rangle. \end{aligned} \quad (2.20)$$

The seemingly innocuous hiccup at (2.20), essentially due to the fact that a and a^\dagger do not commute, is what will eventually make the quantum stochastic calculus of section 2.3.8 work.

2.3.3 The Harmonic Oscillator à la Fock

It is instructive to see how all of this works out in the simple case $\mathcal{F}(\mathbb{C})$. This describes the harmonic oscillator, i.e. a particle moving in a potential $V(x) \sim x^2$.

The n -fold symmetric tensor product $\mathbb{C}^{\otimes_s n}$ of \mathbb{C} is one dimensional, and we choose the unit basis element $|1\rangle_n = 1 \otimes \dots \otimes 1$. We denote it $|n\rangle := |1\rangle_n$ for brevity. These combine to form the *particle basis* $\{|n\rangle; n \in \mathbb{N}\}$ of the Fock space

$$\mathcal{F}(\mathbb{C}) = \bigoplus_{n=0}^{\infty} \mathbb{C}^{\otimes_s n}.$$

In terms of this particle basis, we have the exponential vectors

$$e(z) = \sum_{n=0}^{\infty} \frac{z^n}{\sqrt{n!}} |n\rangle.$$

Conversely, we can recover the particle basis from the exponential vectors by

$$|n\rangle = \frac{1}{\sqrt{n!}} \left. \frac{d^n}{dz^n} \right|_{z=0} e(z). \quad (2.21)$$

Let us investigate what the observables $\lambda(h) = i\dot{\mathcal{F}}(h)$ and $J(v) = i\dot{W}(\xi)$ look like in the particle basis. (They combine to form a projective representation of the three-dimensional real Lie algebra $\mathfrak{su}(\mathbb{C}) \times \mathbb{C}$.)

The 1-dimensional unitary group $U(\mathbb{C})$ acts on the exponential vectors by $W(e^{-ih})e(z) = e(e^{-ih}z)$. From equation (2.21), we see that $W(e^{-ith})|n\rangle = e^{-inth}|n\rangle$, so that

$$\lambda(h)|n\rangle = nh|n\rangle. \quad (2.22)$$

We denote the second quantized Hamiltonian by $H = \lambda(1)$. The particle basis is a basis of eigenvectors of H .

Using equation (2.21) and the fact that $J(v) = i \frac{d}{dt} \Big|_0 W(tv)$, a short calculation starting from

$$W(tv)e(z) = \exp(-\frac{1}{2}t^2|v|^2 - t\bar{v}z)e(z + tv)$$

yields

$$J(v)|n\rangle = iv\sqrt{n+1}|n+1\rangle - i\bar{v}\sqrt{n}|n-1\rangle$$

for $n \geq 1$, and $J(v)|0\rangle = iv|1\rangle$. From equation (2.14), we then have

$$a(1)|n\rangle = \sqrt{n}|n-1\rangle, \quad a^\dagger(1)|n\rangle = \sqrt{n+1}|n+1\rangle. \quad (2.23)$$

From equation 2.13, we see that $[J(1), J(i)] = 2i$. If we define $X = -J(1)$ and $P = -J(i)/2$, we have the canonical commutation relation $[X, P] = i$. The creation and annihilation operators can be expressed as

$$a(1) = P - \frac{1}{2}iX, \quad a^\dagger(1) = P + \frac{1}{2}iX$$

with $[a(1), a^\dagger(1)] = \mathbf{1}$.

We see from equations (2.22) and (2.23) that $\lambda(1) = a(1)a^\dagger(1) - \mathbf{1}$, so that

$$H = P^2 + \frac{1}{4}X^2 - \frac{1}{2}\mathbf{1}.$$

This means that $\mathcal{F}(\mathbb{C})$ serves as a Hilbert space for the harmonic oscillator in which H is diagonal w.r.t. the particle basis. The Schrödinger equation $i\frac{d}{dt}\psi(t) = H\psi(t)$ is therefore trivially solved, the fundamental solutions are $\psi(t) = e^{-int}|n\rangle$.

The fact that H is diagonal means that X and P are not, because they do not commute with H . We will now consider a representation in which X is diagonal, but H and P are not.

2.3.4 The Harmonic Oscillator à la Gauß

Consider the Hilbert space $L^2(\mathbb{R}, \mathbb{P}_G)$, with \mathbb{P}_G is the Gaussian probability measure

$$\mathbb{P}_G(dx) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2} dx.$$

It has the characteristic function $\chi_G(z) = \mathbb{E}(e^{izx}) = e^{-\frac{1}{2}z^2}$. Because the functions $e_z(x) = e^{izx}$ are in $L^2(\mathbb{R}, \mathbb{P}_G)$ not only for $z \in \mathbb{R}$, but for all $z \in \mathbb{C}$, we may consider $\chi_G(z)$ as a holomorphic function.

The inner product $\langle e_z | e_w \rangle = \mathbb{E}(e^{i(w-\bar{z})x})$ equals $\chi_G(w - \bar{z}) = e^{-\frac{1}{2}(w-\bar{z})^2}$. The linear span of the e_z is dense in $L^2(\mathbb{R}, \mathbb{P}_G)$, and if we define $E_z = e^{\frac{1}{2}z^2} e_z$, then $\langle E_z | E_w \rangle = e^{\bar{z}w}$.

Recall that similarly, the exponential vectors $e(z)$ span $\mathcal{F}(\mathbb{C})$, and that $\langle e(z) | e(w) \rangle = e^{\bar{z}w}$. The map $e(z) \mapsto E_z$ therefore yields an isomorphism $\gamma : \mathcal{F}(\mathbb{C}) \xrightarrow{\sim} L^2(\mathbb{R}, \mathbb{P}_G)$. We can now transport the Weyl operators to $L^2(\mathbb{R}, \mathbb{P}_G)$, as well as X , P and H .

On exponential vectors, the Weyl operators are given by $\gamma W(v)\gamma^{-1}E_z(x) = e^{\frac{1}{2}v(v-\bar{v})} e^{ivx} E_z(x+v-\bar{v})$. The E_z are dense, so we must have $\gamma W(v)\gamma^{-1}\psi(x) = e^{\frac{1}{2}v(v-\bar{v})} e^{ivx}\psi(x+v-\bar{v})$ for all $\psi \in L^2(\mathbb{R}, \mathbb{P}_G)$. The corresponding observables are then obtained from $J(v) = i\frac{d}{dt}|_0\gamma W(tv)\gamma^{-1}$. We get $J(1) = -x$ and $J(i) = i(2\frac{d}{dx} - x)$, so that $X = x$, $P = -i\frac{d}{dx} + \frac{1}{2}ix$ and $H = -\frac{d^2}{dx^2} + x\frac{d}{dx}$. We conclude that indeed, in the Gaussian picture, X is diagonal.

Although H is not diagonal here, we can use γ to transport its basis $|n\rangle$ of eigenvectors to $L^2(\mathbb{R}, \mathbb{P}_G)$. The particle basis then transforms into the Hermite polynomials

$$\psi_n(x) = \frac{1}{\sqrt{n!}} \frac{d^n}{dz^n} \Big|_0 e^{\frac{1}{2}z^2 + izx},$$

which apparently satisfy $(-\frac{d^2}{dx^2} + x\frac{d}{dx})\psi_n = n\psi_n$. Schrödinger's equation $i\frac{d}{dt}\psi = (-\frac{d^2}{dx^2} + x\frac{d}{dx})\psi$ therefore has the fundamental solutions $e^{-int}\psi_n(x)$.

2.3.5 The Wiener Measure

We wish to obtain a Gaussian representation of the Fock space $\mathcal{F}(L^2(\mathbb{R}^+))$, in which the fields $\Phi(f)$ for all real functions f are *simultaneously* diagonal-

ized. The germane probability distribution turns out to be the so-called Wiener measure, obtained from the random walk in the limit of small step size.

The Random Walk

The random walk is the following stochastic process. Divide the time line \mathbb{R}^+ into equal intervals of length Δt . At the start of each time interval, flip a coin to step either left or right by an amount $\sqrt{\Delta t}$. If we let Δt tend to zero, then according to the central limit theorem, the position B_t at time t will have a Gaussian distribution with $\mu = 0$ and $\sigma = \sqrt{t}$.

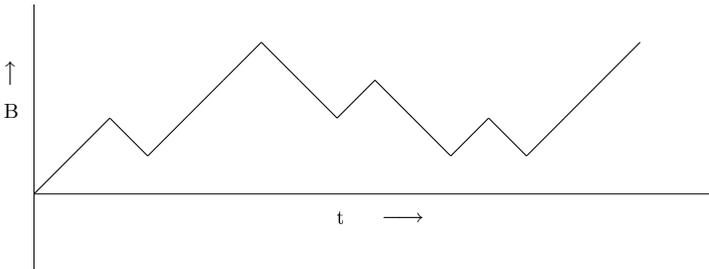


FIG. 1: Random walk, produced with a 0,05€ random bit generator.

For each finite step size, the random walk places a probability measure $\mathbb{P}_{\Delta t}$ on the set of possible paths starting at 0. The Wiener measure \mathbb{P}_W is its limit as Δt tends to zero. It lives on $C_S(\mathbb{R}^+)$, the space of continuous paths $B : \mathbb{R}^+ \rightarrow \mathbb{R}$ that start at 0 and grow at most polynomially. We briefly sketch its construction.

Measures on the Space of Tempered Distributions

Because the increments ΔB_t rather than the endpoints B_T are independent, it is natural to view B_T as the ‘noise’ $N_t = \frac{d}{dt} B_t$ integrated from 0 to T , and then consider \mathbb{P}_W as a probability measure on the set $\Omega = \{N = \frac{d}{dt} B; B \in C_S(\mathbb{R}^+)\}$ of noises.

Unfortunately, not every continuous function is differentiable. Because one cannot expect the erratic ($\Delta B_t \sim \sqrt{\Delta t}$) paths to be differentiable, we cannot solve this problem by restricting the space of paths. We therefore define the noise N_t as the tempered distribution

$$N(\phi) = - \int_0^\infty B_t \frac{d}{dt} \phi dt,$$

a continuous linear functional on the space $C_S^\infty(\mathbb{R}^+)$ of rapidly decreasing functions vanishing at 0. Using highly illegal partial integration, we write this symbolically as $N(\phi) = \int_0^\infty \phi(t) dB_t$.

Having realized our noises as a subset $\Omega \subset \mathcal{D}_S(\mathbb{R}^+)$ of the tempered distributions, we can (and will) consider the Wiener measure as a measure on $\mathcal{D}_S(\mathbb{R}^+)$ with $\mathbb{P}_W(\Omega) = 1$.

Now for any probability measure \mathbb{P} on $\mathcal{D}_S(\mathbb{R}^+)$, we define the characteristic function $\chi : C_S^\infty(\mathbb{R}^+) \rightarrow \mathbb{C}$ to be

$$\chi(\phi) = \int_{\mathcal{D}_S(\mathbb{R}^+)} \exp(iF(\phi)) \mathbb{P}(dF).$$

According to the following theorem, a probability measure is completely characterized by its characteristic function. (Hence the name.)

Theorem 2.18 (Bochner-Minlos) *A functional $\chi : C_S^\infty(\mathbb{R}^+) \rightarrow \mathbb{C}$ is the characteristic function of a probability measure on the space $\mathcal{D}_S(\mathbb{R}^+)$ of tempered distributions if and only if it is continuous, if $\chi(0) = 1$ and if it is positive semidefinite, i.e. $\sum_{i,j} \bar{z}_i z_j \chi(\phi_i - \phi_j) \geq 0$ for all $z_1 \dots z_m \in \mathbb{C}$, $\phi_1 \dots \phi_m \in C_S^\infty(\mathbb{R}^+)$.*

Proof. Any characteristic function is continuous, and $\chi(0) = \mathbb{P}(\mathcal{D}_S(\mathbb{R}^+))$ is equal to 1. It is also positive semidefinite, as

$$\sum_{i,j} \bar{z}_i z_j \int_{\mathcal{D}_S(\mathbb{R}^+)} \exp(iF(\phi_i - \phi_j)) \mathbb{P}(dF) = \int_{\mathcal{D}_S(\mathbb{R}^+)} \left| \sum_i z_i \exp(iF(\phi_i)) \right|^2 \mathbb{P}(dF)$$

is certainly nonnegative. For the converse statement, which is of course the hard part, see e.g. [GV64, p. 322]. \square

Our strategy is now to calculate the characteristic function $\chi_{\Delta t}$ of the random walk process with step size Δt , and show that it tends to a limit χ_W as Δt tends to zero. We then use theorem 2.18 to define \mathbb{P}_W as the unique measure characterized by χ_W .

The Wiener Measure

Calculating $\chi_{\Delta t}$ is rather straightforward. If we think of $\omega \in \{\pm 1\}^{\mathbb{N}}$ as an infinite coin toss, then $B_t(\omega) = \int_0^t N_s(\omega) ds$, with $N(\omega) = \sum_{k=0}^{\infty} \frac{\omega_k}{\sqrt{\Delta t}} \mathbf{1}_{[k\Delta t, (k+1)\Delta t)}$. As a distribution, this reads

$$N(\phi) = \sum_{k=0}^{\infty} \frac{\omega_k}{\sqrt{\Delta t}} \int_{k\Delta t}^{(k+1)\Delta t} \phi(s) ds.$$

Then

$$\begin{aligned} \chi_{\Delta t}(\phi) &= \lim_{n \rightarrow \infty} 2^{-n} \sum_{\omega \in \{\pm 1\}^n} \exp\left(\sum_{k=0}^{n-1} \frac{i\omega_k}{\sqrt{\Delta t}} \int_{k\Delta t}^{(k+1)\Delta t} \phi(s) ds \right) \\ &= \lim_{n \rightarrow \infty} \prod_{k=0}^{n-1} \cos\left(\frac{1}{\sqrt{\Delta t}} \int_{k\Delta t}^{(k+1)\Delta t} \phi(s) ds \right). \end{aligned}$$

This means that as the step size Δt tends to zero, the characteristic function of the random walk process approaches the limit $\chi_{\Delta t}(\phi) \rightarrow \exp(-\frac{1}{2} \int_0^\infty \phi^2(s) ds)$. The function $\chi_W(\phi) = \exp(-\frac{1}{2} \int_0^\infty \phi^2(s) ds)$ is indeed continuous and positive semidefinite, so that we may finally define the Wiener measure.

Definition 2.10 *The Wiener measure \mathbb{P}_W is the unique probability measure \mathbb{P}_W on $\mathcal{D}_S(\mathbb{R}^+)$ with characteristic function $\chi_W(\phi) = \exp(-\frac{1}{2} \int_0^\infty \phi^2(s) ds)$.*

One can show that the Wiener measure concentrates on the set of noises, $\mathbb{P}_W(\Omega) = 1$. It therefore induces a probability measure on $C_S(\mathbb{R}^+)$, which we denote again by \mathbb{P}_W . It is the unique one with the following properties.

- Every increment $B_{t_2} - B_{t_1}$ has a Gaussian distribution with $\mu = 0$ and $\sigma = \sqrt{t_2 - t_1}$.
- The increments are independent, i.e. for any sequence $t_1 \leq \dots \leq t_k$ of times, the steps $B_{t_{i+1}} - B_{t_i}$ are mutually independent.
- The increments are stationary, i.e. the joint probability distribution of the steps $(B_{t_1} - B_{s_1}), \dots, (B_{t_k} - B_{s_k})$ does not change if all t_i and s_i are shifted by the same amount.

2.3.6 Stochastic Differential Equations

Consider B_s as a random variable on the space $\Omega \subset \mathcal{D}_S(\mathbb{R}^+)$ of noises. Let Σ_t be the σ -algebra generated by $\{B_s; 0 \leq s \leq t\}$. The σ -algebras Σ_s with $s \in \mathbb{R}^+$ filter Ω , in the sense that Σ_t is finer than Σ_s if $t > s$.

A *process* is a measurable map $Y : \Omega \times \mathbb{R}^+ \rightarrow \mathbb{R}$. A process is called *adapted* if Y_t is measurable w.r.t. Σ_t . One should think of an adapted process as one in which Y_t depends only on the paths B_s up to time t . We denote by $L^2(\Omega \times \mathbb{R}^+, \Sigma_t)$ the Hilbert space of adapted processes that are square integrable w.r.t. the product of \mathbb{P}_W and the Lebesgue measure.

An adapted step function is an adapted process of the form

$$Y_t(\omega) = \sum_{k=0}^{\infty} c_k(\omega) \mathbf{1}_{[t_k, t_{k+1})}(t)$$

for some increasing sequence t_k with $t_0 = 0$ and $t_k \rightarrow \infty$. Adaptedness implies that the c_k are measurable w.r.t. Σ_{t_k} . If Y is an adapted step function, and $T = t_n$, then we define the stochastic integral

$$\int_0^T Y_t dB_t = \sum_{k=0}^{n-1} c_k(\omega) (B_{t_{k+1}}(\omega) - B_{t_k}(\omega)).$$

If T is not one of the t_k , we simply make a subdivision to fix this.

Theorem 2.19 (Itô) *The adapted step functions are dense in $L^2(\Omega \times \mathbb{R}^+, \Sigma_t)$, and the map $Y \mapsto \int_0^\infty Y_s dB_s$ on step functions extends to an isometry $I : L^2(\Omega \times \mathbb{R}^+, \Sigma_t) \rightarrow L^2(\Omega, \mathbb{P}_W)$. We denote $I(Y \mathbf{1}_{[0,t]})$ by $\int_0^t Y_s dB_s$. The process $Z_t = \int_0^t Y_s dB_s$ has a version with $t \mapsto Z_t(\omega)$ continuous a.s.*

Now that we know how to integrate adapted processes, we can define the *stochastic differential equation*

$$dX_t = A(X_t, t)dt + S(X_t, t)dB_t \tag{2.24}$$

to be just shorthand for the integral equation

$$X_T = X_0 + \int_0^T A(X_t, t)dt + \int_0^T S(X_t, t)dB_t.$$

If A and S both satisfy $|f(x, t)| \leq C(1 + |x|)$ and $|f(x, t) - f(x', t)| \leq D|x - x'|$ uniformly for $t \in [0, T]$, then a Picard iteration scheme shows that equation (2.24) has a unique continuous adapted solution X_t for any given X_0 .

Theorem 2.20 *If X_t and X'_t satisfy $dX_t = A(X_t, t)dt + S(X_t, t)dB_t$ and $dX'_t = A'(X'_t, t)dt + S'(X'_t, t)dB_t$, then $X_t X'_t$ satisfies $d(X_t X'_t) = X_t dX'_t + dX_t X'_t + dX_t dX'_t$, where $dX_t dX'_t$ should be evaluated according to the following Itô table.*

\times	dt	dB_t
dt	0	0
dB_t	0	dt

In other words,

$$d(X_t X'_t) = (X_t A'_t + A_t X'_t + S_t S'_t) dt + (X_t S'_t + S X'_t) dB_t.$$

This is what ultimately makes stochastic differential equations tractable. We can just do calculus, as long as we keep track of the second order terms.

For example, let us solve the stochastic differential equation

$$dX_t = X_t dB_t \tag{2.25}$$

with initial condition $X_0 = 1$. One would expect exponential behaviour. For any C^2 function f , we have $df(X_t) = f'(X_t)dX_t + \frac{1}{2}f''(X_t)(dX_t)^2$. This means that $\exp(B_t)$ will not quite be a solution; $d\exp(B_t) = \exp(B_t)dB_t + \frac{1}{2}\exp(B_t)dt$. We correct this, and try $X_t = \exp(B_t - \frac{1}{2}t)$ instead. This does the trick; $d\exp(B_t - \frac{1}{2}t) = (dB_t - \frac{1}{2}dt)\exp(B_t - \frac{1}{2}t) + \frac{1}{2}\exp(B_t - \frac{1}{2}t)dt = \exp(B_t - \frac{1}{2}t)dB_t$, so that $X_t = \exp(B_t - \frac{1}{2}t)$ is the unique solution of equation (2.25) with $X_0 = 1$.

2.3.7 Fock Space and the Wiener Measure

We can now realize the fields $\Phi(f)$ as multiplication operators on the Hilbert space $L^2(C_S(\mathbb{R}^+), \mathbb{P}_W)$. The construction of the required isomorphism $\gamma : \mathcal{F}(L^2(\mathbb{R}^+)) \xrightarrow{\sim} L^2(C_S(\mathbb{R}^+), \mathbb{P}_W)$ is analogous to that of the homonymous map of section 2.3.4.

The characteristic function $\chi_W : C_S^\infty(\mathbb{R}^+) \rightarrow \mathbb{C}$ of the Wiener measure extends to a holomorphic map $\chi_W : C_S^\infty(\mathbb{R}^+, \mathbb{C}) \rightarrow \mathbb{C}$, given by the same formula $\chi_W(\phi) = \exp(-\frac{1}{2}\int_0^\infty \phi^2(s)ds)$. The functions $e_\phi : F \mapsto \exp(iF(\phi))$ therefore have finite expectation $\chi_W(\phi) = \mathbb{E}(e_\phi)$ for all $\phi \in C_S^\infty(\mathbb{R}^+, \mathbb{C})$. They are even square integrable. Indeed,

$$\begin{aligned} \langle e_\phi | e_\psi \rangle &= \int_{\mathcal{D}_S(\mathbb{R}^+)} \exp(-iF(\bar{\phi})) \exp(iF(\psi)) \mathbb{P}_W(dF) \\ &= \int_{\mathcal{D}_S(\mathbb{R}^+)} \exp(iF(\psi - \bar{\phi})), \end{aligned}$$

so that $\langle e_\phi | e_\psi \rangle = \chi_W(\psi - \bar{\phi})$, and in particular $\langle e_\phi | e_\phi \rangle < \infty$. Moreover, in terms of the renormalized functions

$$E_\phi(F) = \exp\left(\frac{1}{2} \int_0^\infty \phi^2(t) dt\right) \exp(iF(\phi)), \quad (2.26)$$

we have the coveted relation $\langle E_\phi | E_\psi \rangle = \exp(\langle \phi | \psi \rangle)$. Since the E_ϕ are dense in $L^2(\mathcal{D}_S(\mathbb{R}^+), \mathbb{P}_W)$, the map $e(\phi) \mapsto E_\phi$ extends to an isomorphism. It is this map, considered as an isomorphism $\gamma : \mathcal{F}(L^2(\mathbb{R}^+)) \xrightarrow{\sim} L^2(C_S(\mathbb{R}^+), \mathbb{P}_W)$, that realizes the fields $\Phi(f)$ as multiplication operators.

Theorem 2.21 *The isomorphism $\gamma : \mathcal{F}(L^2(\mathbb{R}^+)) \xrightarrow{\sim} L^2(C_S(\mathbb{R}^+), \mathbb{P}_W)$ sends $e(\phi)$ to*

$$\gamma e(\phi) : B \mapsto \exp\left(\frac{1}{2} \int_0^\infty \phi(t)^2 dt\right) \exp\left(i \int_0^\infty \phi(t) dB_t\right) \quad (2.27)$$

for all $\phi \in L^2(\mathbb{R}^+, \mathbb{C})$, and $\Phi(f)$ to

$$\gamma \Phi(f) \gamma^{-1} : L(B) \mapsto \left(\int_0^\infty f dB_t\right) L(B) \quad (2.28)$$

for all $f \in L^2(\mathbb{R}^+, \mathbb{R})$. Moreover, the probability measure induced on $C_S(\mathbb{R}^+)$ by the vacuum state $e(0)$ is precisely the Wiener measure \mathbb{P}_W .

Proof. If we consider the Wiener measure to be based on $C_S(\mathbb{R}^+)$, then equation (2.26) and the definition of γ yield (2.27). Using equation (2.9), one calculates that

$$\gamma W(v) \gamma^{-1} E_\phi(B) = \exp\left(i \int_0^\infty v(dB + \text{Im}(v) ds)\right) E_\phi(B_t + 2 \int_0^t \text{Im}(v) ds).$$

Because the E_ϕ are dense, this continues to hold if one replaces E_ϕ by any $L \in L^2(C_S(\mathbb{R}^+), \mathbb{P}_W)$. Differentiation then yields³ equation (2.28).

The vacuum state $e(0)$ induces a probability measure $\mathbb{P}_{(f_1, \dots, f_n)}$ on the joint spectrum \mathbb{R}^n of $(\Phi(f_1), \dots, \Phi(f_n))$. We show that this coincides with the joint probability distribution $\mathbb{P}_{(f_1, \dots, f_n)}^W$ of $(\int_0^\infty f_1(t) dB_t, \dots, \int_0^\infty f_n(t) dB_t)$ under the Wiener measure. This suffices to prove the theorem, because the cylinder sets $\{F; (F(f_1), \dots, F(f_n)) \in S\}$ with open $S \subset \mathbb{R}^n$ generate the Borel σ -algebra of $\mathcal{D}_S(\mathbb{R}^+)$.

The characteristic function of $\mathbb{P}_{(f_1, \dots, f_n)}$ is

$$\chi(f_1, \dots, f_n; k_1, \dots, k_n) = \langle e(0) | \prod_{r=1}^n \exp(i\Phi(k_r f_r)) e(0) \rangle.$$

Because $\exp(i\Phi(kf)) = W(kf)$, and $\prod_i W(k_i f_i) = W(\sum_i k_i f_i)$, we see that $\chi(f_1, \dots, f_n; k_1, \dots, k_n) = \chi(\sum_i k_i f_i)$ with $\chi(f) = \langle e(0) | W(f) e(0) \rangle$. It equals $\chi(f) = \exp(-\frac{1}{2} \langle f | f \rangle)$ by equation (2.9).

³The conjugate momenta are given by $\Pi(f)L(B) = \left(i \int_0^\infty f dB_t - 2i \frac{\partial}{\partial \bar{f}}\right) L(B)$, but we will not need them here.

Similarly, $\mathbb{P}_{(f_1, \dots, f_n)}^W$ has characteristic function

$$\chi_W(f_1, \dots, f_n; k_1, \dots, k_n) = \int \prod_{r=1}^n \exp(ik_r \int_0^\infty f_r(t) dB_t) \mathbb{P}_W(dB).$$

This is equal to $\chi_W(\sum_i k_i f_i)$, with $\chi_W(f) = \int \exp(i \int_0^\infty f(t) dB_t) \mathbb{P}_W(dB)$. We have $\chi_W(f) = \exp(-\frac{1}{2} \langle f | f \rangle)$ by definition 2.10. The probability measures $\mathbb{P}_{(f_1, \dots, f_n)}$ and $\mathbb{P}_{(f_1, \dots, f_n)}^W$ now agree because their characteristic functions do. \square

2.3.8 Quantum Stochastic Differential Equations

We have defined stochastic calculus in terms of the random variables B_t . According to theorem (2.21), they correspond to the field operators $\Phi(\mathbf{1}_{[0,t]}) = i(a(\mathbf{1}_{[0,t]}) - a^\dagger(\mathbf{1}_{[0,t]}))$ through the isomorphism γ .

However, we may as well choose to diagonalize⁴ the conjugate momenta $\Pi(\mathbf{1}_{[0,t]}) = a(\mathbf{1}_{[0,t]}) + a^\dagger(\mathbf{1}_{[0,t]})$, and integrate w.r.t. those instead. We then have two *noncommuting* notions of stochastic integration living on the same Fock space.

This train of thought is brought to its logical conclusion by the theory of quantum stochastic integration, due to Hudson and Parthasarathy.

Adapted Processes

We take $\mathcal{H} = L^2(\mathbb{R}^+)$, and write $\mathcal{H} = \mathcal{H}_t \oplus \mathcal{H}^t$ with $\mathcal{H}_t = L^2([0, t])$ and $\mathcal{H}^t = L^2([t, \infty))$. If we write $\mathcal{F} = \mathcal{F}(\mathcal{H})$, $\mathcal{F}_t = \mathcal{F}(\mathcal{H}_t)$ and $\mathcal{F}^t = \mathcal{F}(\mathcal{H}^t)$, then this translates into $\mathcal{F} = \mathcal{F}_t \otimes \mathcal{F}^t$.

We take S to be the dense set of locally bounded functions in \mathcal{H} . The set $e(S)$ of exponential vectors then spans \mathcal{F} .

We choose an auxiliary Hilbert space \mathcal{K} . One may think of \mathcal{F} as the Hilbert space of the electromagnetic field, and of \mathcal{K} as the Hilbert space of some system (an atom, say) in interaction with this field. Our dynamics will take place in $\mathcal{K} \otimes \mathcal{F}$, the Hilbert space of the combined system.

The von Neumann algebra $\mathcal{N} = B(\mathcal{K} \otimes \mathcal{F})$ of observables possesses the natural filtration $\{\mathcal{N}_t; t \in \mathbb{R}^+\}$, with $\mathcal{N}_t = B(\mathcal{H} \otimes \mathcal{F}_t) \otimes \mathbf{1}_{\mathcal{F}^t}$ the von Neumann algebra of observables up to time t . We denote its commutant by \mathcal{N}^t . It is the algebra of future events, $\mathcal{N}^t = \mathbf{1}_{\mathcal{K} \otimes \mathcal{F}_t} \otimes B(\mathcal{F}^t)$.

Definition 2.11 *An adapted process based on S is a family $\{F_t; t \in \mathbb{R}^+\}$ of operators such that F_t is affiliated to \mathcal{N}_t , and $\mathcal{K} \otimes e(S)$ is in the domain of each member.*

The operator F_t does not ‘know’ what happens after time t . The common core is a convenient and not overly restrictive technical requirement.

⁴This is easily done by composing $\gamma : \mathcal{F}(L^2(\mathbb{R}^+)) \xrightarrow{\sim} L^2(C_S(\mathbb{R}^+), \mathbb{P}_W)$ with the ‘Fourier transform’ $\mathcal{F}(i) \in U(\mathcal{F}(L^2(\mathbb{R}^+)))$.

Definition 2.12 *An adapted process F_t is called simple if there is an increasing sequence t_k with $t_0 = 0$ and $t_k \rightarrow \infty$ such that $F_t = \sum_{k=0}^{\infty} F_k \mathbf{1}_{[t_k, t_{k+1})}(t)$. It is called continuous if the map $\gamma : t \mapsto F_t u \otimes e(f)$ is continuous for all $u \in \mathcal{K}$, $f \in S$, and it is called locally square integrable if γ is measurable, and $\int_0^t \|\gamma(s)\|^2 ds$ is finite for each $t \in \mathbb{R}^+$.*

The locally square integrable processes are sufficiently general to cater to our needs, and the simple ones are, as the name suggests, easy to handle. Luckily, locally square integrable processes can be approximated by simple ones.

Proposition 2.22 (Muthuramalingam, [HP84]) *Any locally square integrable process F_t can be approximated by a sequence of simple processes F_t^n , in the sense that for each $t \in \mathbb{R}^+$, $u \in \mathcal{K}$ and $f \in S$, we have*

$$\lim_{n \rightarrow \infty} \int_0^t \|(F_s - F_s^n)u \otimes e(f)\|^2 ds = 0.$$

Proof. One first shows that the simple processes approximate continuous ones. From this, one sees that it suffices to approximate F_t by the convolution processes $F_t^k = (F * k\mathbf{1}_{[0, 1/k]})(t)$, which are continuous. See [HP84] for details. \square

Quantum Stochastic Integration

Stochastic integration is integration against the distinguished Wiener process B_t . In the noncommutative version, there are two separate processes which take its role: the annihilation process $A(t) = a(\mathbf{1}_{[0, t]})$ and the creation process $A^\dagger(t) = a^\dagger(\mathbf{1}_{[0, t]})$.

Definition 2.13 *Let $F_t = \sum_{k=0}^{\infty} F_k \mathbf{1}_{[t_k, t_{k+1})}(t)$ be a simple process, and let $T = t_n$. (Refine the sequence t_k by inserting T if necessary.) We then define the stochastic integral*

$$\int_0^t F_t dA_t = \sum_{k=0}^{n-1} F_k (A(t_{k+1}) - A(t_k)).$$

The expression $\int_0^t F_t dA_t^\dagger$ is defined in a similar fashion, with A^\dagger in stead of A , and of course we have $\int_0^t F_t dt = \sum_{k=0}^{n-1} F_k (t_{k+1} - t_k)$ as usual.

Note that F_k is affiliated to \mathcal{N}_{t_k} , acting essentially on $\mathcal{K} \otimes \mathcal{F}_{t_k}$. On the other hand, $A(t_{k+1}) - A(t_k) = a(\mathbf{1}_{[t_k, t_{k+1})})$ is affiliated to its commutant \mathcal{N}'_{t_k} , acting essentially on \mathcal{F}^{t_k} . In particular, F_t commutes with ΔA_{t_k} and $\Delta A_{t_k}^\dagger$. We have used (and will use from now on) the notation Δf_t for $f(t_{k+1}) - f(t_k)$.

We can now define the stochastic integral of any locally square integrable process F_t by approximating it with simple processes. If $F_t^n \rightarrow F_t$ in the sense of proposition 2.22, then one can show, (and this is not entirely trivial [HP84]), that

$$\lim_{n \rightarrow \infty} \int_0^t F_t^n dA_t$$

converges on vectors of the type $u \otimes e(f)$. We will take this to be our stochastic integral $\int_0^t F_t dA_t$ (and similarly for $\int_0^t F_t dA_t^\dagger$). We introduce the ‘stochastic differential equation’ $dM_t = F_t dA_t + G_t dA_t^\dagger + H_t dt$ as shorthand notation for the integral equation $M_t = M_0 + \int_0^t F_s dA_s + \int_0^t G_s dA_s^\dagger + \int_0^t H_s ds$.

A Quantum Version of Itô’s Formula

We can now prove a noncommutative version of Itô’s formula for stochastic integrals of continuous processes.

Theorem 2.23 (Hudson, Parthasarathy) *Let M_t and M'_t be integrals of continuous processes, i.e. $dM_t = F_t dA_t + G_t dA_t^\dagger + H_t dt$ and $dM'_t = F'_t dA_t + G'_t dA_t^\dagger + H'_t dt$. Suppose that the products between the operators M_t, F_t, G_t, H_t and M'_s, F'_s, G'_s, H'_s are well defined on $\mathcal{K} \otimes e(S)$. Then the adapted process $t \mapsto M_t M'_t$ satisfies $d(M_t M'_t) = d(M_t) M'_t + M_t d(M'_t) + dM_t dM'_t$, where $dM_t dM'_t$ should be evaluated according to the following quantum Itô table.*

\times	dt	dA_t	dA_t^\dagger
dt	0	0	0
dA_t	0	0	dt
dA_t^\dagger	0	0	0

In other words,

$$d(M_t M'_t) = (F_t M'_t + M_t F'_t) dA_t + (G_t M'_t + M_t G'_t) dA_t^\dagger + (H_t M'_t + M_t H'_t + F_t G'_t) dt.$$

Proof. We expand $\Delta(M_t M'_t) = (\Delta M_t) M'_t + M_t (\Delta M'_t) + \Delta M_t \Delta M'_t$, and focus on the last term. In the limit $\Delta t \rightarrow 0$, this equals

$$\begin{aligned} \Delta M_t \Delta M'_t &= F_t F'_t (\Delta A_t)^2 + G_t G'_t (\Delta A_t^\dagger)^2 + H_t H'_t (\Delta t)^2 \\ &\quad + F_t G'_t \Delta A_t \Delta A_t^\dagger + (F_t H'_t + H_t F'_t) \Delta A_t \Delta t \\ &\quad + G_t F'_t \Delta A_t^\dagger \Delta A_t + (G_t H'_t + H_t G'_t) \Delta A_t^\dagger \Delta t. \end{aligned}$$

Each one of these 7 terms is a product of a term in \mathcal{N}_t and one in \mathcal{N}^t , so that we can split their contributions. For example, take $F_t G'_t \Delta A_t \Delta A_t^\dagger$. Consider it as $F_t G'_t \otimes \mathbf{1}_{\mathcal{F}^t} \times \mathbf{1}_{\mathcal{K} \otimes \mathcal{F}_t} \otimes \Delta A_t \Delta A_t^\dagger$, and split $u \otimes e(h) = (u \otimes e(h \mathbf{1}_{[0,t]})) \otimes e(h \mathbf{1}_{[t,\infty]})$ accordingly. We see that

$$\langle u \otimes e(h) | F_t G'_t \Delta A_t \Delta A_t^\dagger u' \otimes e(h') \rangle$$

equals

$$\langle u \otimes e(h \mathbf{1}_{[0,t]}) | F_t G'_t u' \otimes e(h' \mathbf{1}_{[0,t]}) \rangle$$

times

$$\langle e(h \mathbf{1}_{[t,\infty]}) | a(\mathbf{1}_{[t_k, t_{k+1}]}) a^\dagger(\mathbf{1}_{[t_k, t_{k+1}]}) e(h' \mathbf{1}_{[t,\infty]}) \rangle.$$

A look at equation (2.20) reveals that this second factor equals

$$\left\{ \left(\int_{t_k}^{t_{k+1}} h'(s) ds \right) \left(\int_{t_k}^{t_{k+1}} \bar{h}(s) ds \right) + \Delta t \right\} \langle e(h \mathbf{1}_{[t,\infty]}) | e(h' \mathbf{1}_{[t,\infty]}) \rangle.$$

As h' and h are locally bounded, the first term is quadratic in Δt . If we take $\Delta t \rightarrow 0$, we are left with $\langle u \otimes e(h) | F_t G_t' u' \otimes e(h') \rangle \Delta t$. The other 6 terms are of second order, as equations (2.17), (2.18), (2.19) lack the extra $\langle u_1 | u_2 \rangle$ term of equation (2.20). \square

Quantum Stochastic Differential Equations

Let V_1, V_2 and V_3 be bounded operators on \mathcal{K} . We can then define the quantum stochastic differential equation

$$dM_t = M_t(V_1 dA_t + V_2 dA_t^\dagger + V_3 dt) \quad (2.29)$$

as shorthand for the quantum integral equation

$$M_t = M_0 + \int_0^t M_s(V_1 dA_s + V_2 dA_s^\dagger + V_3 ds).$$

Theorem 2.24 *Equation (2.29) has a unique continuous solution with $M_0 = \mathbf{1}$. It is unitary if and only if (V_1, V_2, V_3) takes the shape $(V^\dagger, -V, iH - \frac{1}{2}V^\dagger V)$.*

Proof. The existence of a unique solution can be shown by means of a Picard-type iteration procedure, see [HP84]. The solution is unitary, $M_t^\dagger M_t = \mathbf{1}$, if and only if $d(M_t^\dagger M_t) = 0$. Using theorem 2.23, we calculate this by adding $M_t^\dagger dM_t = V_1 dA_t + V_2 dA_t^\dagger + V_3 dt$, $d(M_t^\dagger)M_t = V_1^\dagger dA_t^\dagger + V_2^\dagger dA_t + V_3^\dagger dt$ and $d(M_t^\dagger)d(M_t) = V_2^\dagger V_2 dt$. The result is identically zero if and only if $V_1 + V_2^\dagger = 0$ and $V_3 + V_3^\dagger + V_2^\dagger V_2 = 0$. \square

If we think of $\mathcal{F}(L^2(\mathbb{R}^+))$ as the Hilbert space of the electromagnetic field, and of \mathcal{K} as the Hilbert space of an atom coupled to this field, then time evolution is described by the Schrödinger equation

$$i \frac{d}{dt} U^\lambda(t) = (H_S \otimes \mathbf{1} + \mathbf{1} \otimes H_F + \lambda H_I) U^\lambda(t)$$

on $\mathcal{K} \otimes \mathcal{F}(L^2(\mathbb{R}^+))$. In this expression, H_S is the Hamiltonian of the system \mathcal{K} , H_F is the free field Hamiltonian on $\mathcal{F}(L^2(\mathbb{R}^+))$, and $H_I = i(D \otimes a^\dagger(g) - D^\dagger \otimes a(g))$ is the interaction Hamiltonian on $\mathcal{K} \otimes \mathcal{F}(L^2(\mathbb{R}^+))$ with coupling constant λ .

Since $U^0(t)$ can often be easily solved, it suffices to find the interaction picture time evolution $\hat{U}^\lambda(t) := U^\lambda(t)U^0(-t)$. The ‘van Hove’ limit of $U^\lambda(t/\lambda^2)$ for $\lambda \downarrow 0$ exists under suitable conditions. According to a theorem of Accardi, Frigerio and Lu, see [AFL90], it is given by a unitary quantum stochastic differential equation of the form

$$dU_t = U_t(V^\dagger dA_t - V dA_t^\dagger - \frac{1}{2}V^\dagger V dt).$$

This equation will play a pivotal role in chapters 5 and 6.

2.4 Bundles and Classical Fields

We have reviewed the foundations of quantum mechanics, and scratched the surface of a tiny part of quantum field theory. We will now take a step back, and submit the basic framework of classical field theory to a closer look.

In classical field theory, a system is modelled by a smooth fibre bundle $\pi : F \rightarrow M$, and fields are described by sections of π . The base manifold M represents space-time. The dynamics are prescribed by a Lagrangian density $\mathcal{L} : J^1F \rightarrow \wedge^n T^*M$, with n the dimension of M . Transformations of the theory correspond to automorphisms of F .

2.4.1 Bundles, Sections and Jets

Intuitively speaking, a fibre bundle over M with fibre F_0 is a smoothly varying family $\{F_x; x \in M\}$ of copies of F_0 . A section is a function $\phi : M \rightarrow F$ such that its value on $x \in M$ lies in the fibre F_x over x . The k -jet $j_x^k \phi$ of ϕ at x is its value $\phi(x)$ at x , together with its derivatives at x up to order k . Let us make this a bit more precise.

Fibre Bundles

If F, F', M and M' are smooth manifolds, then an *isomorphism* between the maps $\pi : F \rightarrow M$ and $\pi' : F' \rightarrow M'$ is by definition a pair (ϕ, ϕ_M) of diffeomorphisms $\phi : F \rightarrow F'$ and $\phi_M : M \rightarrow M'$ such that $\pi' \circ \phi = \phi_M \circ \pi$. Since ϕ_M is completely determined by ϕ , we will suppress it in the notation, and write ϕ for (ϕ, ϕ_M) . If $M = M'$, then ϕ is called *vertical* if $\phi_M = \text{id}_M$.

Definition 2.14 *Let F, M and F_0 be smooth manifolds. Then a fibre bundle $\pi : F \rightarrow M$ with fibre F_0 is a smooth map with the property that each $x \in M$ has an open neighbourhood U such that the restriction $\pi : \pi^{-1}(U) \rightarrow U$ is vertically isomorphic to the projection $\pi' : U \times F_0 \rightarrow U$ on the first factor.*

In short, F is locally the product of the base by F_0 . A vertical isomorphism $\phi : \pi^{-1}(U) \rightarrow U \times F_0$ is called a *local trivialization*. The manifold $F_x := \pi^{-1}(\{x\})$ is called the *fibre* over x . Each local trivialization restricts to a diffeomorphism $\phi|_{F_x} : F_x \xrightarrow{\sim} \{x\} \times F_0$, so that in particular, any fibre F_x is diffeomorphic to F_0 .

For any fibre bundle $F \rightarrow M$, we can produce *adapted co-ordinates* as follows. Choose co-ordinates $x^\mu : M \supset U \rightarrow \mathbb{R}$ on M , co-ordinates $v^a : F_0 \supset V \rightarrow \mathbb{R}$ on F_0 , and local trivializations $\phi_U : \pi^{-1}(U) \rightarrow U \times F_0$. The adapted co-ordinates on F are then $(x^\mu \circ \phi_U, v^a \circ \phi_U)$. We will simply denote them (x^μ, v^a) . They have the convenient property that $x^\mu(f) = x^\mu(f')$ for all μ if and only if f and f' are in the same fibre.

Finally, if $F \rightarrow M$ is any fibre bundle, and $f : N \rightarrow M$ is a smooth map, then we can define the pullback bundle $f^*F \rightarrow N$ as the manifold $f^*F = \{(n, f) \in N \times F; f(n) = \pi(f)\}$, with projection $(n, f) \mapsto n$. The idea is that the map f allows one to regard the fibre over $f(n) \in M$ as a fibre over $n \in N$, i.e. $f^*F_n = F_{f(n)}$.

Vector Bundles and Principal Fibre Bundles

It is useful to study fibre bundles with additional structure on the fibres. A fibre bundle $\pi : E \rightarrow M$ is called a *vector bundle* if E_0 and each of the fibres E_x are vector spaces, and if one can choose the local trivializations ϕ over U in such a way that for any $x \in M$, the map $\phi|_{E_x} : E_x \xrightarrow{\sim} E_0$ is an isomorphism of vector spaces. Similarly, a fibre bundle $\pi : P \rightarrow M$ is called a *principal fibre bundle* with Lie group G if P_0 and the fibres P_x carry a free transitive right action of G , and if the local trivializations can be chosen to be G -equivariant.

The tangent bundle is a natural example of a vector bundle. The *tangent space* $T_x M$ of M at $x \in M$ is by definition the space of vectors tangent to M at x . If M is n -dimensional, then $T_x M$ is a vector space isomorphic to \mathbb{R}^n . This makes the *tangent bundle* $TM = \cup_{x \in M} T_x M$ into a vector bundle, with projection $\pi : TM \rightarrow M$ assigning to each tangent vector $\xi_x \in T_x M$ its base point $\pi(\xi_x) = x$.

Starting from the tangent bundle, one can construct many other interesting fibre bundles. If $\pi : F \rightarrow M$ is a any fibre bundle, then the kernel $T_f^V F$ of the pushforward $\pi_* : T_f F \rightarrow T_{\pi(f)} M$ is a linear subspace of $T_f F$, and we obtain the vertical tangent bundle $T^V F \rightarrow F$, a vector subbundle of $TF \rightarrow F$. One can think of it as the union of all the separate tangent spaces $T(F_x)$.

The *cotangent bundle* $\pi : T^* M \rightarrow M$ is the vector bundle defined by letting $T_x^* M$ be the dual of $T_x M$, i.e. the vector space of linear maps $T_x M \rightarrow \mathbb{R}$.

Similarly, a *volume form* on the vector space $T_x M$ is an alternating multilinear map $T_x M \times \dots \times T_x M \rightarrow \mathbb{R}$ on n copies of $T_x M$. We denote the 1-dimensional space of volume forms on $T_x M$ by $\wedge^n T_x^* M$. Their union $\wedge^n T^* M = \cup_{x \in M} \wedge^n T_x^* M$ is a line bundle; a vector bundle over M with 1-dimensional fibre.

The *frame bundle* $F(M) \rightarrow M$ is a natural example of a principal fibre bundle. A frame at x is by definition a linear isomorphism $f : \mathbb{R}^n \rightarrow T_x(M)$. If we fix the standard basis (e_1, \dots, e_n) of \mathbb{R}^n , then frames f at x correspond precisely to bases $(f(e_1), \dots, f(e_n))$ of $T_x(M)$. The space of frames at x , denoted $F_x(M)$, has a free transitive right action $g : f \mapsto f \circ g$ of $\text{GL}(\mathbb{R}^n)$, making $F(M) = \cup_{x \in M} F_x(M)$ into a principal fibre bundle.

Sections

A *section* of $\pi : F \rightarrow M$ is a smooth map $\phi : M \rightarrow F$ such that $\pi \circ \phi = \text{id}_M$, and we denote the space of sections by $\Gamma(F)$.

If F happens to be the trivial bundle $M \times F_0$, then sections correspond precisely to functions $M \rightarrow F_0$. One way to view sections and bundles is as follows. Take the base manifold M , and cut it into parts U_i that overlap only at their boundaries, and that are small enough to fit into a trivializing neighbourhood. A section ϕ of F is then a collection of functions $\phi_i : U_i \rightarrow F_0$, and the bundle $F \rightarrow M$ encodes the boundary conditions that ϕ_i and ϕ_j have to satisfy on $\partial U_i \cap \partial U_j$.

A section of the tangent bundle $TM \rightarrow M$ is precisely a smooth vector field on M , so that the commutator bracket $[\xi, \chi]$ of vector fields endows $\Gamma(TM)$ with a structure of Lie algebra. A section of $T^* M$ is called a 1-form, and a

section of $\wedge^n T^*M$ a volume form.

The tangent bundle $TS^2 \rightarrow S^2$ of the 2-sphere is a nice nice example of a vector bundle that is not trivial. Every vector field on S^2 must have at least one zero,⁵ showing that a section of $TS^2 \rightarrow S^2$ is not quite the same as a function $S^2 \rightarrow \mathbb{R}^2$.

Jet Bundles

Intuitively speaking, the k -jet of a section ϕ at x is its value $\phi(x)$, together with its derivatives at x up to order k .

To put this in a more formal setting, we introduce the relation \sim on the local sections around x , where $\phi \sim \phi'$ if ϕ and ϕ' have the same value at x , and if their derivatives agree up to order k . The k -jet of ϕ at x , denoted $j_x^k \phi$, is now simply the equivalence class of ϕ modulo \sim .

We denote the set of all k -jets at x by $J_x^k F$, and we define $J^k F := \cup_{x \in M} J_x^k F$ to be the *jet bundle*. It is a fibre bundle in two different ways. The map $\pi_k : J^k F \rightarrow M$ defined by $j_x^k(\phi) \mapsto x$ makes $J^k F$ into a bundle over M , and the map $\pi_{k,0} : J^k F \rightarrow F$ that maps $j_x^k(\phi)$ to $\phi(x)$ makes it into a bundle over F . We have $\pi \circ \pi_{k,0} = \pi_k$. Any section ϕ of $\pi : F \rightarrow M$ yields a section $j^1 \phi$ of $\pi_k : J^k F \rightarrow M$, but not all sections of $J^k F$ are of this form.

Adapted co-ordinates $(\tilde{x}^\mu, \tilde{v}^a)$ on F give rise to the co-ordinates (x^μ, v^a, v_μ^a) on $J^1 F$ by $x^\mu(j_x^1 \phi) = \tilde{x}^\mu(x)$, $v^a(j_x^1 \phi) = \tilde{v}^a(\phi(x))$ and $v_\mu^a(j_x^1 \phi) = \partial_\mu \tilde{v}^a(\phi(x))$. If ϕ is a section of F , then we denote $\phi^a := v^a \circ \phi$, so that $\partial_\mu \phi^a = v_\mu^a(j^1 \phi)$.

2.4.2 The Euler-Lagrange Equation

Having described its basic ingredients, we now have a closer look at classical field theory itself. A physical system is described by a fibre bundle $\pi : F \rightarrow M$, the base manifold M representing space-time. Fields are modelled by sections of π , and the dynamics are prescribed by a bundle map $\mathcal{L} : J^1 F \rightarrow \wedge^n(T^*M)$ called⁶ the *Lagrangian density*, or Lagrangian for short. It singles out a class of preferred sections by means of the *action principle*. These are considered to be the fields of physical relevance.

The Action Principle

Let U be an open submanifold of M with smooth boundary and compact closure. We define the *action functional* $S_U : \Gamma(F) \rightarrow \mathbb{R}$ by $S_U(\phi) = \int_U j^1 \phi^* \mathcal{L}$. The action principle then states that the physical fields are precisely the critical points of the action S_U under infinitesimal variations with support inside U .

An infinitesimal variation of a section $\phi \in \Gamma(F)$ is precisely a section $\delta \phi$ of $\phi^* T^V F$, the pullback of the vertical tangent bundle of F along ϕ . Indeed, if ϕ_ε

⁵The theorem of P. Bohl, J. Hadamard and L. E. J. Brouwer that every (continuous) vector field on a sphere of even dimension must have at least one zero is known colloquially as the ‘hairy ball theorem’: one cannot comb the hair on a $2n$ -sphere without creating a crown.

⁶Generalization to Lagrangian densities on $J^k F$ with $k > 1$ is straightforward, and the Euler-Lagrange equation then becomes a PDE of order $k + 1$. We focus on the case $k = 1$ for notational convenience, but also because most fundamental field equations are of order two.

is a one-parameter family of sections of F with $\phi_0 = \phi$, then $\delta\phi(x) = \frac{d}{d\varepsilon}|_0\phi_\varepsilon$ yields a vertical vector in $T_{\phi(x)}F$ for each $x \in M$.

If we choose adapted co-ordinates (x^μ, v^a, v_μ^a) , then we can write the Lagrangian density as $\mathcal{L} = L(x^\mu, v^a, v_\mu^a)\text{vol}$, with $\text{vol} = dx^1 \wedge \dots \wedge dx^n$ the volume form assigning unit volume to the frame $(\frac{\partial}{\partial x^1}, \dots, \frac{\partial}{\partial x^n})$, and L the so-called Lagrangian function. The action then reads $S_U(\phi) = \int_U L(x^\mu, \phi^a, \partial_\mu \phi^a)\text{vol}$.

In order to avoid a jungle of sigmas, we will adopt the convention that repeated indices indicate summation. We also use the shorthand notation $\frac{\partial L}{\partial \phi^a}$ for $x \mapsto \frac{\partial L}{\partial v^a}(x^\mu(x), \phi^a(x), \partial_\mu \phi^a(x))$, and similarly, we write $\frac{\partial L}{\partial(\partial_\mu \phi^a)}$ for the function $x \mapsto \frac{\partial L}{\partial v_\mu^a}(x^\mu(x), \phi^a(x), \partial_\mu \phi^a(x))$.

The following theorem shows that critical sections correspond to solutions of a second order P.D.E. called the *Euler-Lagrange equation*.

Theorem 2.25 (Euler-Lagrange) *A section ϕ of F is a stationary point of the action functionals S_U under variations $\delta\phi$ with support strictly contained in U if and only if it satisfies the second order P.D.E.*

$$[\mathcal{L}]^\phi = 0, \quad (2.30)$$

called the *Euler-Lagrange equation*. The map $[\mathcal{L}]^\phi : \phi^*T^V F \rightarrow \wedge^n T^*M$ of vector bundles is well defined by

$$[\mathcal{L}]^\phi(\delta\phi) = \left(\left(\partial_\mu \frac{\partial L}{\partial(\partial_\mu \phi^a)} \right) - \frac{\partial L}{\partial \phi^a} \right) \text{vol} \cdot \delta\phi^a. \quad (2.31)$$

The r.h.s. of (2.31) does not depend on the choice of co-ordinates, whether ϕ is critical or not.

Proof. Since $\frac{\partial L}{\partial \phi^a}(x)$ and $\frac{\partial L}{\partial(\partial_\mu \phi^a)}(x)$ only depend on ϕ through $j_x^1\phi$, we see from (2.31) that $[\mathcal{L}]^\phi(x)$ depends only on ϕ through $j_x^2\phi$, so that (2.30) is indeed a P.D.E. of order 2.

The variation of S_D along $\delta\phi = \delta\phi^a \partial_a$ reads

$$\delta S_U(\phi) = \int_U \left(\frac{\partial L}{\partial \phi^a} \delta\phi^a + \frac{\partial L}{\partial(\partial_\mu \phi^a)} \delta\partial_\mu \phi^a \right) \text{vol}.$$

Note that $\delta\partial_\mu \phi^a = \partial_\mu \delta\phi^a$, because $\frac{d}{d\varepsilon}|_0 \partial_\mu \phi_\varepsilon^a = \partial_\mu \frac{d}{d\varepsilon}|_0 \phi_\varepsilon^a$. We can then write

$$\frac{\partial L}{\partial(\partial_\mu \phi^a)} \delta\partial_\mu \phi^a \text{vol} = d \left(\frac{\partial L}{\partial(\partial_\mu \phi^a)} \delta\phi^a i_{\partial_\mu} \text{vol} \right) - \left(\partial_\mu \frac{\partial L}{\partial(\partial_\mu \phi^a)} \right) \delta\phi^a \text{vol}.$$

If we now define $[\mathcal{L}]^\phi : \phi^*T^V F \rightarrow \wedge^n T^*M$ by

$$[\mathcal{L}]^\phi(\delta\phi) = \left(\left(\partial_\mu \frac{\partial L}{\partial(\partial_\mu \phi^a)} \right) - \frac{\partial L}{\partial \phi^a} \right) \text{vol} \cdot \delta\phi^a \quad (2.32)$$

and $\{\mathcal{L}\}^\phi : \phi^*T^V F \rightarrow \wedge^{n-1} T^*M$ by

$$\{\mathcal{L}\}^\phi(\delta\phi) = \frac{\partial L}{\partial(\partial_\mu \phi^a)} i_{\partial_\mu} \text{vol} \cdot \delta\phi^a, \quad (2.33)$$

then we have

$$\delta S_U(\phi) = - \int_U [\mathcal{L}]^\phi(\delta\phi) + \int_{\partial U} \{\mathcal{L}\}^\phi(\delta\phi). \quad (2.34)$$

It can be seen from (2.34) that $[\mathcal{L}]^\phi$ and $\{\mathcal{L}\}^\phi$ do not depend on our choice of co-ordinates, a fact not at all obvious from equations (2.31) and (2.33).

The argument runs as follows. Let $[\mathcal{L}]_A^\phi$ and $[\mathcal{L}]_B^\phi$ be two versions of (2.31), defined w.r.t different co-ordinate systems. Let $\delta\phi$ be any section of $\phi^*T^V(F)$ that vanishes on ∂U . Then the second term on the r.h.s. of (2.34) vanishes, and we have both $\delta S_U(\phi) = - \int_U [\mathcal{L}]_A^\phi(\delta\phi)$ and $\delta S_U(\phi) = - \int_U [\mathcal{L}]_B^\phi(\delta\phi)$. This means that $\int_U ([\mathcal{L}]_A^\phi - [\mathcal{L}]_B^\phi)(\delta\phi) = 0$ for all infinitesimal variations vanishing at ∂U , so that $[\mathcal{L}]_A^\phi = [\mathcal{L}]_B^\phi$ holds over the interior of U . But since U was arbitrary, equality must hold over all of M .

We can now show that also $\{\mathcal{L}\}^\phi$ does not depend on the choice of co-ordinates. Let $\delta\phi$ be any section of $\phi^*T^V(F)$, not necessarily vanishing at ∂U . According to equation (2.34), both $\int_{\partial U} \{\mathcal{L}\}_A^\phi(\delta\phi)$ and $\int_{\partial U} \{\mathcal{L}\}_B^\phi(\delta\phi)$ are equal to $\delta S_U(\phi) + \int_U [\mathcal{L}]^\phi(\delta\phi)$. (We just established that $[\mathcal{L}]^\phi(\delta\phi)$ is the same for A and B .) Once again, $\int_{\partial U} \{\mathcal{L}\}_A^\phi(\delta\phi) - \{\mathcal{L}\}_B^\phi(\delta\phi) = 0$ for arbitrary $\delta\phi$ and U implies that $\{\mathcal{L}\}_A^\phi = \{\mathcal{L}\}_B^\phi$ must hold over all of M .

It is clear from (2.34) that S_U is stationary at ϕ under variations $\delta\phi$ with support inside U if and only if $\int_U [\mathcal{L}]^\phi(\delta\phi)$ is zero for all such variations. This is the case if and only if $[\mathcal{L}]^\phi$ is identically zero. \square

Example: the Klein-Gordon Equation

One of the simplest examples of an Euler-Lagrange equation the following. Let $M = \mathbb{R}^4$, and equip it with the Minkowski metric η of signature $(-, +, +, +)$. In co-ordinates (x^0, x^1, x^2, x^3) , this reads $\eta(\xi, \chi) = \eta_{\mu\nu} \xi^\mu \chi^\nu$, with $\eta_{\mu\nu}$ the diagonal matrix $\text{diag}(-1, 1, 1, 1)$. We denote $\xi_\mu = \eta_{\mu\nu} \xi^\nu$, so that $\xi^\mu = \eta^{\mu\nu} \xi_\nu$.

Let F be the trivial bundle $\mathbb{R}^4 \times \mathbb{C} \rightarrow \mathbb{R}^4$, which comes with co-ordinates $(x^\mu; v^a)$. The index μ runs from 0 to 3, and a runs from 0 to 1. (We write $z \in \mathbb{C}$ as $z = v^0 + iv^1$.) We define the Lagrangian

$$\mathcal{L}(j^1\phi) = \left(\frac{1}{2m} \partial^\mu \phi_a \partial_\mu \phi^a - \frac{m}{2} \phi_a \phi^a \right) \text{vol}.$$

We have $\frac{\partial \mathcal{L}}{\partial \phi^a} = m\phi_a$ and $\frac{\partial \mathcal{L}}{\partial (\partial_\mu \phi^a)} = m^{-1} \partial^\mu \phi_a$, so that the Euler-Lagrange equation in this case is the *Klein Gordon* equation

$$\partial^\mu \partial_\mu \phi = m^2 \phi.$$

It allows for plane wave solutions $\phi(x) = \exp(k_\mu x^\mu)$ with $k^\mu k_\mu = m^2$.

Note that certainly not all Euler-Lagrange equations are linear. For instance, if F is the bundle $\mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}^2$, η is of signature $(+, -)$, and $\mathcal{L}(j^1\phi) = \frac{1}{2} \partial_\mu \phi \partial^\mu \phi - (1 - \cos(\phi))$, then the resulting Euler-Lagrange equation $\partial_\mu \partial^\mu \phi = -\sin(\phi)$, rather childishly called the sine-Gordon equation, is a nonlinear P.D.E. with some bearing on surfaces of constant negative curvature.

Example: Electromagnetism

The most classical of all field theories is electromagnetism. The base manifold M is again \mathbb{R}^4 with the Minkowski metric, but the bundle F is now the cotangent bundle $T^*M \rightarrow M$. This is a trivial vector bundle for $M = \mathbb{R}^4$. If dx^ν is the dual of ∂_{x^ν} in T_x^*M , then a global trivialization is given by describing $p_\nu dx^\nu \in T_x^*M$ with the co-ordinates (x^μ, p_ν) .

The basic field in electromagnetism is a 1-form $A \in \Gamma(T^*M)$ called the electromagnetic four-potential. We set $p_\nu(A(x)) = A_\nu(x)$, and write $A = A_\mu(x)dx^\mu$. We define the antisymmetric tensor $F_{\mu\nu} := \partial_\mu A_\nu - \partial_\nu A_\mu$, so that $F = 2dA$, with $F = F_{\mu\nu}dx^\mu \wedge dx^\nu$. The (vacuum) EM-Lagrangian is then

$$\mathcal{L} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu}\text{vol}.$$

We have $\frac{\partial \mathcal{L}}{\partial \partial_\mu A_\nu} = -F^{\mu\nu}$, and $\frac{\partial \mathcal{L}}{\partial A_\nu} = 0$. The Euler-Lagrange equation is then the (vacuum) *Maxwell equation*

$$\partial_\mu F^{\mu\nu} = 0.$$

The 3 components $E_i = F_{i0}$ are identified with the electric field, and the 3 components $B_i = \frac{1}{2}\epsilon_{ijk}F_{jk}$ with the magnetic field. Light is described by the plane wave solutions $A_\mu(x) = a_\mu \exp(ik_\mu x^\mu)$ with $k_\mu k^\mu = 0$, $a_\mu k^\mu = 0$ and, at least in the Coulomb gauge, $a_0 = 0$.

2.4.3 Transformations and Symmetries

Transformations of the system correspond to automorphisms of the fibre bundle. An automorphism of $\pi : F \rightarrow M$ is a pair (α, α_M) with α a diffeomorphism of the total space F , α_M a diffeomorphism of the base M , and $\pi \circ \alpha = \alpha_M \circ \pi$. If F is a vector bundle, we additionally require that α be linear on the fibres, and if F is a principal fibre bundle, we require α to be equivariant w.r.t. the group action. Because α_M is completely determined by α , we drop it from the notation, and write α instead of (α, α_M) . An automorphism is called *vertical* if $\alpha_M = \text{id}$.

The automorphisms of F form the group $\text{Aut}(F)$, and the vertical automorphisms $\text{Aut}^V(F)$ form a normal subgroup. If we denote the diffeomorphism group of M by $\text{Diff}(M)$, then $\alpha \mapsto \alpha_M$ is a homomorphism from $\text{Aut}(F)$ to $\text{Diff}(M)$ with kernel $\text{Aut}^V(F)$. Its image $\text{Diff}^F(M)$, the group of liftable diffeomorphisms, is an open subgroup of $\text{Diff}(M)$. We have the exact sequence

$$1 \rightarrow \text{Aut}^V(F) \rightarrow \text{Aut}(F) \rightarrow \text{Diff}^F(M) \rightarrow 1. \quad (2.35)$$

The group $\text{Aut}(F)$ of transformations acts on the space $\Gamma(F)$ of fields by $\phi \mapsto \alpha \circ \phi \circ \alpha_M^{-1}$. Note however that without any additional structure, the diffeomorphism group $\text{Diff}(M)$ does *not* have a natural action on the fields $\Gamma(F)$. The extra structure that is required to have $\text{Diff}^F(M)$ act on the fields $\Gamma(F)$ is precisely a homomorphism $\Sigma : \text{Diff}^F(M) \rightarrow \text{Aut}(F)$ splitting (2.35).

An automorphism α of F acts on J^1F by $j^1\alpha : j_x^1\phi \mapsto j_x^1(\alpha \circ \phi \circ \alpha_M^{-1})$. This means that α takes the Lagrangian $\mathcal{L} : J^1F \rightarrow \wedge^n T^*M$ into $\alpha_M^* \circ \mathcal{L} \circ j^1\alpha$.

Definition 2.15 *An automorphism α of F is called a symmetry of the Lagrangian $\mathcal{L} : J^1F \rightarrow \wedge^n T^*M$ if $\mathcal{L} = \alpha_M^* \circ \mathcal{L} \circ j^1\alpha$.*

This is a sensible definition. If ϕ is a local section of F over U , then $\alpha \circ \phi \circ \alpha_M^{-1}$ is a local section over $\alpha(U)$. One readily checks that if α is a symmetry of \mathcal{L} , then $S_U(\phi) = S_{\alpha(U)}(\alpha \circ \phi \circ \alpha_M^{-1})$. In particular, $\alpha \circ \phi \circ \alpha_M^{-1}$ is a critical section if and only if ϕ is, and the action of α on $\Gamma(F)$ restricts to the solution space $\Gamma_{EL}(F)$ of the Euler-Lagrange equation.

Infinitesimal Symmetries

It will be expedient to look at this picture from the infinitesimal point of view. The Lie algebra of $\text{Diff}(M)$ is, at least morally speaking, the Lie algebra $\Gamma(TM)$ of vector fields on M . The group $\text{Aut}(F)$ corresponds to the Lie algebra $\Gamma^P(TF)$ of projectable vector fields on F , i.e. vector fields $\xi \in \Gamma(TF)$ such that $\pi_*\xi_f = \pi_*\xi_{f'}$ whenever $\pi(f) = \pi(f')$. The pushforward $\pi_* : \Gamma^P(TF) \rightarrow \Gamma(TM)$ is not only well defined, but also a homomorphism of Lie algebras. Its kernel, the ideal $\Gamma^V(TF)$ of vertical vector fields, corresponds to the group $\text{Aut}^V(F)$, and we have the exact sequence of Lie algebras

$$0 \longrightarrow \Gamma^V(TF) \longrightarrow \Gamma^P(TF) \xrightarrow{\pi_*} \Gamma(TM) \longrightarrow 0. \quad (2.36)$$

The infinitesimal variation $\delta\phi \in \Gamma(\phi^*T^V F)$ of ϕ under $\xi \in \Gamma^P(TF)$ is given by $\delta\phi_x = \frac{d}{d\varepsilon}|_0 \alpha_\varepsilon \circ \phi \circ \alpha_{\varepsilon, M}^{-1}(x)$, with $\alpha_\varepsilon = \exp(\varepsilon\xi)$ and $\alpha_{\varepsilon, M} = \exp(\varepsilon\pi_*\xi)$ the induced flows on F and M respectively. It is readily seen to be $\delta\phi_x = d^v\phi(\xi_{\phi(x)})$, with $d^v\phi(\xi) = \xi - \phi_*\pi_*\xi$ the vertical projection $T_{\phi(x)}F \rightarrow T_{\phi(x)}^V F$ along the image of $\phi_* : T_xM \rightarrow T_{\phi(x)}F$.

In order to obtain the infinitesimal version of definition 2.15, we require that $\frac{d}{d\varepsilon}|_0 \alpha_{\varepsilon, M}^* \mathcal{L} \circ j^1\alpha_\varepsilon$ vanish. Variation of $\alpha_{\varepsilon, M}^*$ gives rise to a term $L_{\pi_*\xi}\mathcal{L}(j^1(\phi)) = d(i_{\pi_*\xi}\mathcal{L}(j^1\phi))$. Variation of \mathcal{L} through the variation $\frac{d}{d\varepsilon}\alpha_\varepsilon \circ \phi \circ \alpha_{\varepsilon, M}^{-1}$ of the field gives rise to an additional term $-[\mathcal{L}]^\phi(\delta\phi) + d(\{\mathcal{L}\}^\phi(\delta\phi))$, cf. equation (2.34), with $\delta\phi_x = d^v\phi(\xi_{\phi(x)})$. This leads us to the following definition of infinitesimal symmetries.

Definition 2.16 *A projectable vector field $\xi \in \Gamma^P(TF)$ is called an infinitesimal symmetry of the Lagrangian $\mathcal{L} : J^1F \rightarrow \wedge^n T^*M$ if*

$$-[\mathcal{L}]^\phi(d^v\phi(\xi)) + d(i_{\pi_*\xi}\mathcal{L} + \{\mathcal{L}\}^\phi(d^v\phi(\xi))) = 0.$$

From this definition, we immediately see that infinitesimal symmetries give rise to conserved quantities. If we define the *Noether current* associated to the infinitesimal symmetry ξ and the field ϕ to be the $(n-1)$ -form

$$J^\phi(\xi) = \{\mathcal{L}\}^\phi(d^v\phi(\xi)) + i_{\pi_*\xi}\mathcal{L}(j^1(\phi)), \quad (2.37)$$

then the following theorem follows straight from the definitions.

Theorem 2.26 (Noether) *If ϕ is a critical point of the action, and ξ is an infinitesimal symmetry, then $dJ^\phi(\xi) = 0$.*

Proof. We have $\delta\mathcal{L} = -[\mathcal{L}]^\phi(d^v\phi(\xi)) + dJ^\phi(\xi)$. Now $\delta\mathcal{L} = 0$ because ξ is an infinitesimal symmetry, and $[\mathcal{L}]^\phi = 0$ by the Euler-Lagrange equation. \square

The rationale behind considering closed $(n-1)$ -forms as conserved quantities is the following. Let t be a (local) time co-ordinate, and let Σ be a compact $(n-1)$ -dimensional submanifold with boundary of the (local) time-slice $t = t_0$. If we interpret $Q_\Sigma = \int_\Sigma J^\phi(\xi)$ as the ‘charge’ of Σ , then $\frac{d}{dt}Q_\Sigma = \int_\Sigma \mathcal{L}_{\partial_t} J^\phi(\xi)$, with \mathcal{L}_{∂_t} the Lie derivative along ∂_t . Now $\mathcal{L}_{\partial_t} J^\phi(\xi) = di_{\partial_t} J^\phi(\xi)$ because $dJ^\phi(\xi) = 0$, and $\int_\Sigma di_{\partial_t} J^\phi(\xi) = \int_{\partial\Sigma} i_{\partial_t} J^\phi(\xi)$ by Gauß’ theorem. If we now interpret $\Phi_{\partial\Sigma} = \int_{\partial\Sigma} i_{\partial_t} J^\phi(\xi)$ as the flux through $\partial\Sigma$, then $\frac{d}{dt}Q_\Sigma = \Phi_{\partial\Sigma}$ says that the rate of change of the charge of Σ is equal to the flux through its boundary.

Strictly speaking then, Q_Σ is not a conserved quantity at all; we merely know that its rate of change can be attributed to a flux. However, if the current goes to zero fast enough at spatial infinity, then the total ‘charge’ of the time-slice $t = t_0$ is conserved.

Example: the Current of a Klein-Gordon Field

The Klein-Gordon Lagrangian allows the one-parameter group of vertical rotations $\alpha_\theta : (x^\mu; v^0, v^1) \mapsto (x^\mu; \cos(\theta)v^0 + \sin(\theta)v^1, -\sin(\theta)v^0 + \cos(\theta)v^1)$ as a group of symmetries. The corresponding infinitesimal symmetry is the vector field $\xi = v^1 \frac{d}{dv^0} - v^0 \frac{d}{dv^1}$. We calculate the corresponding current from equation (2.37). Because ξ is vertical, $\pi_*\xi = 0$ and $d^v\phi(\xi) = \xi$. Since $\{\mathcal{L}_\alpha^\phi = \frac{1}{m}\partial^\mu\phi_\alpha i_{\partial_\mu}\text{vol}, \text{ we get } J^\phi(\xi) = \frac{1}{m}(\phi^0\partial^\mu\phi^1 - \phi^1\partial^\mu\phi^0)i_{\partial_\mu}\text{vol}$. It is not hard to check ‘by hand’ that $\partial_\mu J^\mu(\xi) = 0$ for solutions of the Klein-Gordon equation, as predicted by Emmy Noether.

Transformations and Symmetries of Space-Time

We would like to associate conserved quantities to infinitesimal symmetries of space-time. However, in order to even define what a space-time symmetry is, we need to know how $\text{Diff}(M)$ acts on the fields. In other words, we need a homomorphism $\Sigma : \text{Diff}^F(M) \rightarrow \text{Aut}(F)$ that splits (2.35). A transformation $\alpha_M \in \text{Diff}(M)$ is then called a space-time symmetry of \mathcal{L} if $\Sigma(\alpha_M)$ is a symmetry in the sense of definition 2.15.

An infinitesimal transformation of space-time is a vector field $\xi_M \in \Gamma(TM)$. In order to describe its action on the fields, we need a splitting $\sigma : \Gamma(TM) \rightarrow \Gamma^P(TF)$ of the exact sequence (2.36) of Lie-algebras. A vector field $\xi_M \in \Gamma(TM)$ is called an infinitesimal space-time symmetry of \mathcal{L} if $\sigma(\xi_M)$ is a symmetry in the sense of definition 2.16.

This Lie algebra homomorphism $\sigma : \Gamma(TM) \rightarrow \Gamma^P(TF)$ will be the object of our attention in chapter 7. We will argue there that it is of more fundamental importance than its counterpart $\Sigma : \text{Diff}(M) \rightarrow \text{Aut}(F)$ at the level of groups.

Although every flat connection gives rise to a splitting of (2.36), it is certainly not true that every splitting arises in this way. However, in chapter 7, proposition 7.7 and corollary 7.17, we will prove that under reasonable assumptions⁷, σ must be a differential operator of finite order k .

⁷Namely, we assume that the fibre bundle F is associated to a principal fibre bundle, and that the infinitesimal transformations respect this association.

This means that the map of sections $\sigma : \Gamma(TM) \rightarrow \Gamma^P(TF)$ comes from a map $\nabla : \pi^* J^k TM \rightarrow TF$ of bundles over F . In local co-ordinates, there then exist functions $C_\tau^{a;\rho_1 \dots \rho_j}(x^\mu, v^a)$ such that

$$\sigma(\xi^\mu \partial_\mu) = \xi^\mu \partial_\mu + (C_\tau^{a;\xi^\tau} + C_\tau^{a;\rho} \partial_\rho \xi^\tau + \dots + C_\tau^{a;\rho_1 \dots \rho_k} \partial_{\rho_1} \dots \partial_{\rho_k} \xi^\tau) \partial_a.$$

Note that the case $k = 0$ corresponds to a flat connection.

The Stress-Energy-Momentum Tensor

Once we have a splitting σ , it is tempting to define the associated conserved current as $J^\phi(\sigma(\xi_M))$. Appealing as this may seem, it is *not* the proper course of action in the presence of a sufficiently large group of space-time symmetries, as noted first in [Noe18].

Instead, one can use σ in order to construct a *Stress-Energy-Momentum (SEM) tensor* as follows. Choose adapted co-ordinates, and expand

$$\delta \mathcal{L} = \left(\partial_\mu(\xi^\mu L) + \frac{\partial L}{\partial \phi^a} \delta \phi^a + \frac{\partial L}{\partial \partial_\mu \phi^a} \partial_\mu \delta \phi^a \right) \text{vol}$$

into

$$\delta \mathcal{L} = T_\tau \xi^\tau \text{vol} + T^\rho{}_\tau \partial_\rho \xi^\tau \text{vol} + \dots + T_\tau^{\rho_1 \dots \rho_{k+1}} \partial_{\rho_1} \dots \partial_{\rho_{k+1}} \xi^\tau \text{vol}, \quad (2.38)$$

using $\delta \phi = d^v \phi(\sigma(\xi_M))$, i.e.

$$\delta \phi^a = (C_\tau^a - \partial_\tau \phi^a) \xi^\tau + C_\tau^{a;\rho} \partial_\rho \xi^\tau + \dots + C_\tau^{a;\rho_1 \dots \rho_k} \partial_{\rho_1} \dots \partial_{\rho_k} \xi^\tau.$$

The tensor density $T^\rho{}_\tau(\phi) \text{vol}$ is then the desired SEM-tensor. If Σ is a compact $(n-1)$ -dimensional submanifold with boundary of the equal-time slice $t = t_0$, then for τ equal to zero, the quantity $P_\tau = \int_\Sigma T^\rho{}_\tau(\phi) i_{\partial_\rho} \text{vol}$ is interpreted as the energy of the fields in Σ . For τ equal to 1 through $n-1$, it is interpreted as the momentum.

The problem with this definition is of course that the tensor density $T^\rho{}_\tau(\phi) \text{vol}$ may well depend on the choice of co-ordinates. However, if \mathcal{L} is diffeomorphism invariant, then $\delta \mathcal{L} = 0$ in equation (2.38) for any choice of $\xi \in \Gamma(TM)$, forcing all the $T_\tau^{\rho_1 \dots \rho_s}(\phi)$ to vanish identically.

Example: the SEM-Tensor of Electromagnetism

The tangent bundle comes with a natural lift of $\text{Diff}(M)$. For any diffeomorphism α , the pushforward α_* maps $T_x M$ to $T_{\alpha(x)} M$, and is therefore an automorphism of TM covering α . The lift $\Sigma(\alpha) = \alpha_*$ is a homomorphism of groups.

Similarly, the exact sequence (2.35) for the cotangent bundle T^*M is split by the inverse pullback $\Sigma(\alpha) = (\alpha^{-1})^*$, considered as a map $\text{Diff}(M) \rightarrow \text{Aut}(T^*M)$. Its infinitesimal version is the Lie algebra homomorphism $\sigma(\xi) = \frac{d}{d\varepsilon} \big|_0 \exp(-\varepsilon \xi)^*$, which reads $\sigma(\xi^\mu \partial_\mu) = \xi^\mu \partial_\mu - p_\sigma \partial_\nu \xi^\sigma \frac{\partial}{\partial p_\nu}$ in local co-ordinates.

The infinitesimal variation $\delta A \in \Gamma(A^*(T^V T^* M))$ of $A \in \Gamma(T^* M)$ along $\sigma(\xi)$ reads $\delta A = d^v A(\sigma(\xi)) = \sigma(\xi)_{A(x)} - A_* \xi$. In local co-ordinates, this is $\delta A = (\xi^\mu \partial_\mu - A_\sigma \partial_\nu \xi^\sigma \frac{\partial}{\partial p_\nu}) - (\xi^\mu \partial_\mu + \xi^\sigma \partial_\sigma A_\nu \frac{\partial}{\partial p_\nu})$, i.e. $\delta A = -(A_\sigma \partial_\nu \xi^\sigma + \xi^\sigma \partial_\sigma A_\nu) \frac{\partial}{\partial p_\nu}$. If we identify $A^*(T^V T^* M)$ with $T^* M$, which we can do because $T^* M$ happens to be a vector bundle, then we simply have $\delta A = -\mathcal{L}_\xi A$.

Now that we know that action of $\Gamma(TM)$ on $\Gamma(T^* M)$ is the Lie derivative – What else could it have been? – calculating the SEM-tensor is straightforward. We expand $\delta \mathcal{L} = -(F^{\mu\nu} \partial_\mu \delta A_\nu - \frac{1}{4} \partial_\tau (F^{\mu\nu} F_{\mu\nu} \xi^\tau)) \text{vol}$ into derivatives of ξ .

$$\begin{aligned} \delta \mathcal{L} &= (F^{\mu\nu} \partial_\mu (\xi^\tau \partial_\tau A_\nu + A_\tau \partial_\nu \xi^\tau) - \frac{1}{4} \partial_\tau (F^{\mu\nu} F_{\mu\nu} \xi^\tau)) \text{vol} \\ &= (F^{\mu\nu} \partial_\mu \partial_\tau A_\nu - \frac{1}{4} \partial_\tau (F^{\mu\nu} F_{\mu\nu})) \xi^\tau \text{vol} \\ &\quad + (F^{\mu\nu} (\partial_\mu \xi^\tau \partial_\tau A_\nu + \partial_\mu A_\tau \partial_\nu \xi^\tau) - \frac{1}{4} F^{\mu\nu} F_{\mu\nu} \partial_\tau \xi^\tau) \text{vol} \\ &\quad + A_\tau F^{\mu\nu} \partial_\mu \partial_\nu \xi^\tau \text{vol} \\ &= -\frac{1}{4} \partial_\tau (F^{\mu\nu} F_{\mu\nu}) \xi^\tau \text{vol} \\ &\quad + (F^{\mu\nu} F_{\tau\nu} - \frac{1}{4} F^{\mu\nu} F_{\mu\nu} \delta^\mu{}_\tau) \partial_\mu \xi^\tau \text{vol} \end{aligned}$$

In the last step, we used that $F_{\mu\nu}$ is antisymmetric, and that $F^{\mu\nu} \partial_\mu A_\tau \partial_\nu \xi^\tau = -F^{\mu\nu} \partial_\nu A_\tau \partial_\mu \xi^\tau$. According to our definition, the SEM-tensor of electromagnetism is

$$T^\mu{}_\tau(A) = F^{\mu\nu} F_{\tau\nu} - \frac{1}{4} F^{\mu\nu} F_{\mu\nu} \delta^\mu{}_\tau. \quad (2.39)$$

In particular, the energy density T^{00} of the EM-field is $\frac{1}{2}(|E|^2 + |B|^2) \text{vol}$.

Remark The above tensor density is co-ordinate independent, which, in our limited setup, is a small miracle. As ever so often, the miracle disappears if we broaden our view a little. The transformation behaviour $\Sigma : \text{Diff}(M) \rightarrow \text{Aut}(F)$ of the EM-field allows us to couple it to a metric $g^{\mu\nu}$, resulting in a Lagrangian that equals the one above for $g = \eta$, but is $\text{Diff}(M)$ -invariant because the transformation of the metric cancels that of the fields. The SEM-tensor then vanishes identically. It is the sum of a term resulting from δA , and an equal term of opposite sign resulting from δg . The former term was calculated above, and the latter term can be seen to represent a bona fide tensor density from general principles. See e.g. [FR04] for details.

Remark Our definition of the SEM-tensor differs from the one in [GM92] and [FR04] by a term $[\mathcal{L}]_a C_\tau^{\alpha;\mu}$ which vanishes on shell, i.e. for solutions of the Euler-Lagrange equations. Consequently, the above expression for the SEM-tensor is correct on shell as well as off shell, whereas the derivation in [GM92] is valid only on shell. The term $[\mathcal{L}]_a C_\tau^{\alpha;\mu}$ does appear in Emily Noether's seminal paper [Noe18], but her aim is slightly different from ours. She expresses $d(J^\mu(\xi) - [\mathcal{L}]_a C_\tau^{\alpha;\mu} \xi^\tau \text{vol})$ in terms of $[\mathcal{L}]^\phi$ and its derivatives, showing that the Euler-Lagrange equations are linearly dependent in the presence of symmetry.

2.4.4 Infinitesimally Natural Bundles

The point of deriving (2.39) is not to acquaint the reader of the precise form of the electromagnetic SEM-tensor. It will not occur in this thesis again, and could be found in a myriad of introductory books on electrodynamics if it did. The point is to stress the crucial role that the splitting $\sigma : \Gamma(TM) \rightarrow \Gamma^P(TF)$ of (2.36) plays in its derivation.

Note that in order to derive equation (2.39), we needed to use the nontrivial splitting derived from the pullback, even though $T^*\mathbb{R}^4 \rightarrow \mathbb{R}^4$ is trivial as a vector bundle! Had we used the trivial splitting $\hat{\sigma}(\xi^\mu \partial_\mu) = \xi^\mu \partial_\mu$ that comes from the global trivialization of T^*M , we would have obtained the *incorrect* SEM-tensor

$$T^\mu{}_\tau = F^{\mu\nu} \partial_\tau A_\nu - \frac{1}{4} F^{\mu\nu} F_{\mu\nu} \delta^\mu{}_\tau,$$

which differs from the correct one by a term $d(F^{\mu\nu} A_\tau i_{\partial_\nu} \text{vol}) - (\partial_\nu F^{\mu\nu}) A_\tau \text{vol}$. This would increase the electromagnetic energy of a 3-volume $\Sigma \subset \mathbb{R}^3$ by a boundary term $\int_{\partial\Sigma} A_0 \vec{E} \cdot d\vec{S}$, which is accessible to experiment. It is known to be incorrect.

We conclude that a classical field theory is not fully specified by a fibre bundle $F \rightarrow M$ and a Lagrangian density $\mathcal{L} : J^1F \rightarrow \wedge^n T^*M$. In order to derive the correct SEM-tensor, one has to specify the reaction of the fields to infinitesimal space-time transformations, that is one has to specify a Lie algebra homomorphism $\sigma : \Gamma(TM) \rightarrow \Gamma^P(TF)$ splitting the exact sequence (2.36).

Even if one has no interest in the SEM-tensor whatsoever, this is a very natural thing to ask for. It is hard to imagine fields on space-time without a well defined transformation behaviour under infinitesimal space-time transformations. For one thing, such fields would have ill-defined spin.

We therefore consider the splitting $\sigma : \Gamma(TM) \rightarrow \Gamma^P(TF)$ as part of the input data of a classical field theory, on equal footing with the bundle $\pi : F \rightarrow M$ and the Lagrangian $\mathcal{L} : J^1F \rightarrow \wedge^n T^*M$. This poses an obvious restriction on the fibre bundle $F \rightarrow M$, namely that it has to admit at least one such splitting. We will call such fibre bundles *infinitesimally natural*. In chapter 7, we will classify these infinitesimally natural bundles—or at least the ones with sufficiently structured fibres.

Chapter 3

Information Transfer Implies State Collapse

We attempt to clarify certain puzzles concerning state collapse and decoherence. In open quantum systems decoherence is shown to be a necessary consequence of the transfer of information to the outside; we prove an upper bound for the amount of coherence which can survive such a transfer. We claim that in large closed systems decoherence has never been observed, but we will show that it is usually harmless to assume its occurrence. An independent postulate of state collapse over and above Schrödinger's equation and the probability interpretation of quantum states, is shown to be redundant.

3.1 Introduction

In its most basic formulation, quantum theory encodes the preparation of a system in a pure quantum state, a unit vector ψ in a Hilbert space \mathcal{H} . Observables are modelled by (say, nondegenerate) self-adjoint operators on \mathcal{H} . The expectation value of an observable A in a state ψ is given by $\langle \psi | A \psi \rangle$. If a is an eigenvalue of A and ψ_a a unit eigenvector, and information concerning A is somehow extracted from the system, then the probability for the value a to be observed is $|\langle \psi_a | \psi \rangle|^2$. If this observation is indeed made, then the subsequent behaviour of the system is predicted using the pure state ψ_a . This is called *state collapse*. It follows that, if the information extraction has taken place but the information on the value of A is disregarded, then the subsequent behaviour can be described optimally using a mixture of eigenstates. This is called *decoherence*. In this chapter we substantiate the following claim concerning decoherence and state collapse.

Decoherence is only observed in open systems, where it is a necessary consequence of the transfer of information to the outside.

So the observed occurrence of decoherence does not contradict the unitary time evolution postulated by quantum mechanics, since open systems do not

evolve unitarily. Decoherence can be explained in quantum theory by embedding the quantum system into a larger, closed whole, which in itself evolves unitarily. This is well-known (see e.g. [Neu32]). We add the observation that decoherence is not only a *possibility* for an open system, but a *necessary consequence* of the leakage of information out of the system. We prove an inequality relating the decoherence between two pure states to the degree in which a decision between the two is possible by a measurement outside. This is the content of theorem 3.3 in section 3.3.

Also, we have claimed that one *does not* actually observe decoherence in closed macroscopic systems. First of all, most of the systems that are ever observed are actually open, since it is extremely difficult to shield large systems from interaction. But more to the point, the difference between coherence and decoherence can only be seen by measuring some highly exotic ‘stray observables’ which are almost always forbiddingly hard to observe. And indeed, in those rare cases where experimenters have succeeded in measuring them, ordinary unitary evolution was found, not decoherence. (See [AN⁺99], [FP⁺00], [WH⁺00].)

We illustrate the latter point in section 3.4, where we show that the measurement of two classes of observables *can not* reveal the difference between coherence and decoherence: a class of microscopic observables and a large class of macroscopic observables. Take as an example a volume of gas. Microscopic observables such as the position of one particular atom in a gas, only relate to a small fraction of the system. Macroscopic observables like the center of mass of the gas, are the average over a large number of microscopic observables. Belonging neither to the macroscopic nor to the microscopic class, the ‘stray observables’ referred to above describe detailed correlations between large numbers of atoms in the gas. This kind of information is experimentally almost inaccessible.

Driving home our point concerning decoherence in closed systems: coherent superpositions of macroscopically distinguishable states are not the strange monsters produced by a quantum theory applied outside its domain. They are, on the contrary, everyday occurrences which, however, *can not* be distinguished from the more classical incoherent superpositions in practice, and can therefore always be regarded as such.

3.2 Abstract Information Extraction

Quantum phenomena are inherently stochastic. This means that, if quantum systems are prepared in identical ways, then nevertheless different events may be observed. A quantum state describes an ensemble of physical systems, e.g. a beam of particles, and is modelled by a normalized trace-class operator ρ on the Hilbert space. The expectation value of an observable A in the state ρ is then $\text{tr}(\rho A)$.

An information extraction or measurement on a quantum state is to be considered as the partition of such an ensemble into subensembles, each subensemble corresponding to a measurement outcome. Let us, in the present section,

not wonder *how* the splitting of ensembles can be described by quantum theory, but let us see what such an information extraction, if it can be done, will entail for the subsequent behaviour of the subensembles. Note that this process may serve as part of the preparation for further experiments on the system, so that it must again lead to a state.

3.2.1 Information Extraction

For simplicity let us assume that only two outcomes can occur, labelled 0 and 1, say with probabilities p_0 and p_1 . The ensemble is then split in two parts, described by their respective states ρ_0 and ρ_1 . The map

$$M : \rho \mapsto p_0\rho_0 \oplus p_1\rho_1 \tag{3.1}$$

must be normalized, affine and positive. Indeed, normalization is the property that $p_0 + p_1 = 1$, and positivity is the requirement that states must be mapped to states. The affine property entails that for all states ρ and ϑ on the original system, and for all $\lambda \in [0, 1]$,

$$M(\lambda\rho + (1 - \lambda)\vartheta) = \lambda M(\rho) + (1 - \lambda)M(\vartheta) .$$

This follows from the physical principle that a system which is prepared in the state ρ with probability λ and in the state ϑ with probability $1 - \lambda$, say by tossing a coin, can not be distinguished from a physical system in the state $\lambda\rho + (1 - \lambda)\vartheta$. We emphasize that indeed this is a *physical* principle, not a matter of definitions. It states, for instance, that a bundle of particles having 50% spin up and 50% spin down can not be distinguished from a bundle having 50% spin left and 50% spin right. This is a falsifiable statement.

3.2.2 State Collapse

The above elementary observations are sufficient to prove that information extraction implies state collapse. If M distinguishes perfectly between the pure states ψ_0 and ψ_1 , then of course $p_0 = 1$ in case $\rho = |\psi_0\rangle\langle\psi_0|$, and $p_1 = 1$ if $\rho = |\psi_1\rangle\langle\psi_1|$.

Proposition 3.1 *Let $\mathcal{T}(\mathcal{H})$ denote the space of trace class operators on a Hilbert space \mathcal{H} , and let the map $M : \mathcal{T}(\mathcal{H}) \rightarrow \mathcal{T}(\mathcal{H}) \oplus \mathcal{T}(\mathcal{H}) : \rho \mapsto M_0(\rho) \oplus M_1(\rho)$ be the linear extension of some normalized, affine and positive map on the states. Suppose that unit vectors ψ_0 and ψ_1 exist such that*

$$\begin{aligned} M(|\psi_0\rangle\langle\psi_0|) &= M_0(|\psi_0\rangle\langle\psi_0|) \oplus 0 \\ M(|\psi_1\rangle\langle\psi_1|) &= 0 \oplus M_1(|\psi_1\rangle\langle\psi_1|) . \end{aligned}$$

Then we have $M(|\psi_0\rangle\langle\psi_1|) = M(|\psi_1\rangle\langle\psi_0|) = 0$.

Proof. The positivity of M yields $M(|\varepsilon e^{i\varphi}\psi_0 + \psi_1\rangle\langle\varepsilon e^{i\varphi}\psi_0 + \psi_1|) \geq 0$ as an operator inequality. In particular, the 0-th component must be positive. As

$M_0(|\psi_1\rangle\langle\psi_1|) = 0$, it follows that for all $\varepsilon, \phi \in \mathbb{R}$, we have

$$\varepsilon^2 M_0(|\psi_0\rangle\langle\psi_0|) + \varepsilon (e^{i\varphi} M_0(|\psi_0\rangle\langle\psi_1|) + e^{-i\varphi} M_0(|\psi_1\rangle\langle\psi_0|)) \geq 0.$$

Taking the limit $\varepsilon \downarrow 0$ yields $(e^{i\varphi} M_0(|\psi_0\rangle\langle\psi_1|) + e^{-i\varphi} M_0(|\psi_1\rangle\langle\psi_0|)) \geq 0$ for all $\varphi \in \mathbb{R}$. In particular for $\varphi = 0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}$, implying that both $M_0(|\psi_0\rangle\langle\psi_1|)$ and $M_0(|\psi_1\rangle\langle\psi_0|)$ are equal to zero. Exchanging the roles of ψ_0 and ψ_1 in the argument above shows that also $M_1(|\psi_0\rangle\langle\psi_1|)$ and $M_1(|\psi_1\rangle\langle\psi_0|)$ vanish, proving the proposition. \square

We may draw two conclusions from proposition 3.1. The first is that, for all $|\psi\rangle = \alpha_0|\psi_0\rangle + \alpha_1|\psi_1\rangle$, we have

$$(M_0 + M_1)(|\psi\rangle\langle\psi|) = (M_0 + M_1)(|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|). \quad (3.2)$$

In words: for the prediction of events *after* the splitting of the ensemble in two, it no longer matters whether *before* the splitting the system was in the pure state $|\alpha_0\psi_0 + \alpha_1\psi_1\rangle\langle\alpha_0\psi_0 + \alpha_1\psi_1|$ or in the mixed state $|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|$. This phenomenon, which is a direct consequence of the structure (3.1) of the measurement process, we will call *decoherence*.

The second conclusion from proposition 3.1 is the following. For all $|\psi\rangle = \alpha_0|\psi_0\rangle + \alpha_1|\psi_1\rangle$, we have

$$M(|\psi\rangle\langle\psi|) = |\alpha_0|^2 M_0(|\psi_0\rangle\langle\psi_0|) \oplus |\alpha_1|^2 M_1(|\psi_1\rangle\langle\psi_1|). \quad (3.3)$$

In words: if an ensemble is split in two parts, then the ‘0-ensemble’ will further behave as if the system had been in state ψ_0 instead of ψ prior to splitting, and the ‘1-ensemble’ as if it had been in state ψ_1 instead of ψ . This phenomenon will be called *collapse*.

Throughout this chapter, we will maintain a sharp distinction between the collapse $M : \mathcal{T}(\mathcal{H}) \rightarrow \mathcal{T}(\mathcal{H}) \oplus \mathcal{T}(\mathcal{H})$ and the decoherence $(M_0 + M_1) : \mathcal{T}(\mathcal{H}) \rightarrow \mathcal{T}(\mathcal{H})$. The former represents the splitting of an ensemble in two parts by means of measurement, whereas the latter represents the splitting and subsequent recombination of this ensemble.

3.3 Open Systems

A decoherence-mapping $(M_0 + M_1) : \mathcal{T}(\mathcal{H}) \rightarrow \mathcal{T}(\mathcal{H})$ maps the pure (vector) state $|\alpha_0\psi_0 + \alpha_1\psi_1\rangle\langle\alpha_0\psi_0 + \alpha_1\psi_1|$ and the mixed state $|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|$ to the same final state. Since unitary maps preserve purity, there can not exist a unitary map $U : \mathcal{H} \rightarrow \mathcal{H}$ such that for all $\rho \in \mathcal{T}(\mathcal{H})$:

$$(M_0 + M_1)(\rho) = U\rho U^*.$$

However, according to Schrödinger’s equation the development of a closed quantum system is given by a unitary operator. We conclude that the decoherence (3.2) is impossible in a closed system. On the other hand decoherence is a well known and experimentally confirmed phenomenon.

We will therefore consider open systems, i.e. quantum systems which do not obey the Schrödinger equation, but are part of a larger system which does. It has often been pointed out (e.g. [Neu32], [Zur82]) that decoherence can well occur in this situation, provided that states are only evaluated on the observables of the smaller system. We are more ambitious here: we shall prove that this form of ‘local’ decoherence is not just a *possible*, but an *unavoidable* consequence of information-transfer out of the open system.

3.3.1 Unitary Information Transfer and Decoherence

We assume that the open system has Hilbert space \mathcal{H} , and that its algebra of observables is given by $B(\mathcal{H})$, the bounded operators on \mathcal{H} . We may then assume that the larger system has Hilbert space $\mathcal{K} \otimes \mathcal{H}$, since the only way to represent $B(\mathcal{H})$ on a Hilbert space is in the form $A \mapsto \mathbf{1} \otimes A$ [Tak79]. We may think¹ of $B(\mathcal{K})$ as the observable algebra of some ancillary system in contact with our open quantum system. In this context, \mathcal{H} will be referred to as the ‘open system’, \mathcal{K} as the ‘ancilla’ and $\mathcal{K} \otimes \mathcal{H}$ as the ‘closed system’.

We couple the system to the ancilla during a finite time interval $[0, t]$. Let $\tau \in \mathcal{T}(\mathcal{K})$ denote the state of the ancilla at time 0, and $\rho \in \mathcal{T}(\mathcal{H})$ that of the small system. The effect of the interaction is described by a unitary operator $U : \mathcal{K} \otimes \mathcal{H} \rightarrow \mathcal{K} \otimes \mathcal{H}$, and the state of the pair at time t is given by the density operator $U(\tau \otimes \rho)U^* \in \mathcal{T}(\mathcal{K} \otimes \mathcal{H})$. For convenience, we will define the information transfer map $T : \mathcal{T}(\mathcal{H}) \rightarrow \mathcal{T}(\mathcal{K} \otimes \mathcal{H})$ by $T(\rho) := U(\tau \otimes \rho)U^*$.

Decoherence

In the above setup, we are interested in distinguishing whether the open system \mathcal{H} was in state $|\psi_0\rangle$ or $|\psi_1\rangle$ at time 0. This can be done if there exists a ‘pointer observable’ $B \otimes \mathbf{1}$ in the ancilla $B(\mathcal{K})$ which takes average value b_0 in state $T(|\psi_0\rangle\langle\psi_0|)$ and b_1 in state $T(|\psi_1\rangle\langle\psi_1|)$. By looking only at the ancilla \mathcal{K} at time t , we are then able to gain information on the state of the open system \mathcal{H} at time 0. We say that information is *transferred* from \mathcal{H} to \mathcal{K} .

Under these circumstances, we wish to prove that decoherence occurs on the open system. We prepare the ground by proving the following lemma.

Lemma 3.2 . *Let ϑ_0, ϑ_1 be unit vectors in a Hilbert space \mathcal{L} , and let A and B be bounded self-adjoint operators on \mathcal{L} satisfying $\|[A, B]\| \leq \delta \|A\| \cdot \|B\|$. For $j = 0$ or 1, let $b_j := \langle \theta_j | B \theta_j \rangle$ denote the expectation and $\sigma_j^2 := \langle \theta_j | B^2 \theta_j \rangle - \langle \theta_j | B \theta_j \rangle^2$ the variance of B in the state ϑ_j . Then, if $b_0 \neq b_1$,*

$$|\langle \vartheta_0 | A \vartheta_1 \rangle| \leq \frac{\delta \|B\| + \sigma_0 + \sigma_1}{|b_0 - b_1|} \|A\| .$$

¹Sometimes it may happen, as for instance in fermionic systems, that the observables of the ancilla do not all commute with those of the open system. Also the observable algebra on \mathcal{K} may be smaller than $B(\mathcal{K})$, but we will neglect these complications here.

Proof. We write $|(b_0 - b_1)\langle \vartheta_0 | A \vartheta_1 \rangle|$ as $|\langle \vartheta_0 | (A(B - b_1) - (B - b_0)A + [B, A]) \vartheta_1 \rangle|$. Since $\|(B - b_j)\vartheta_j\|^2 = \langle \vartheta_j | (B - b_j)^2 \vartheta_j \rangle = \sigma_j^2$, we can then use the Cauchy-Schwarz inequality to show that

$$|(b_0 - b_1)\langle \vartheta_0 | A \vartheta_1 \rangle| \leq \|A\|(\sigma_1 + \sigma_0) + \delta \|A\| \cdot \|B\|.$$

□

Note that, for $\delta = \sigma_0 = \sigma_1 = 0$, lemma 3.2 merely states that commuting operators respect each other's eigenspaces. We proceed to prove that information transfer causes decoherence on the open system. (See [Jan04].)

Theorem 3.3 *Let ψ_0 and ψ_1 be mutually orthogonal unit vectors in a Hilbert space \mathcal{H} , and let $\tau \in \mathcal{T}(\mathcal{K})$ be a state on a Hilbert space \mathcal{K} . Let $U : \mathcal{K} \otimes \mathcal{H} \rightarrow \mathcal{K} \otimes \mathcal{H}$ be unitary and define $T : \mathcal{T}(\mathcal{H}) \rightarrow \mathcal{T}(\mathcal{K} \otimes \mathcal{H})$ by $T(\rho) = U(\tau \otimes \rho)U^*$. Let B be a bounded self-adjoint operator on $\mathcal{K} \otimes \mathcal{H}$, and denote by b_j and σ_j^2 its expected value and variance in the state $T(|\psi_j\rangle\langle\psi_j|)$ for $j = 0, 1$. Suppose that $b_0 \neq b_1$. Then for all $\psi = \alpha_0\psi_0 + \alpha_1\psi_1$ with $|\alpha_0|^2 + |\alpha_1|^2 = 1$ and for all bounded self-adjoint operators A on $\mathcal{K} \otimes \mathcal{H}$ such that $\| [A, B] \| \leq \delta \|A\| \cdot \|B\|$, we have*

$$\left| \mathbf{tr}(T(|\psi\rangle\langle\psi|)A) - \mathbf{tr}\left(T(|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|)A\right) \right| \leq \frac{\delta \|B\| + \sigma_0 + \sigma_1}{|b_0 - b_1|} \|A\|. \quad (3.4)$$

Proof. First, we prove (3.4) in the special case that $\tau = |\varphi\rangle\langle\varphi|$ for some vector $\varphi \in \mathcal{K}$. We introduce the notation $\theta_j := U(\varphi \otimes \psi_j)$. Recall that the expectation of B is given by $b_j = \mathbf{tr}(T(|\psi_j\rangle\langle\psi_j|)B)$, and its variance by $\sigma_j^2 = \mathbf{tr}(T(|\psi_j\rangle\langle\psi_j|)B^2) - \mathbf{tr}^2(T(|\psi_j\rangle\langle\psi_j|)B)$. In terms of θ_j , this reduces to $b_j = \langle \theta_j | B \theta_j \rangle$ and $\sigma_j^2 = \langle \theta_j | B^2 \theta_j \rangle - \langle \theta_j | B \theta_j \rangle^2$. Similarly, the l.h.s. of (3.4) equals $|\overline{\alpha_0}\alpha_1\langle \theta_0 | A \theta_1 \rangle + \alpha_0\overline{\alpha_1}\langle \theta_1 | A \theta_0 \rangle|$, a quantity bounded by $|\langle \theta_0 | A \theta_1 \rangle|$ since $2|\alpha_0| \cdot |\alpha_1| \leq 1$. Formula (3.4) is then a direct application of lemma 3.2.

To reduce the general case to the case above, we note that a non-pure state τ can always be represented as a vector state. Explicitly, suppose that τ decomposes as $\tau = \sum_{i \in \mathbb{N}} |\beta_i|^2 |\varphi_i\rangle\langle\varphi_i|$. Then define the Hilbert space $\tilde{\mathcal{K}} := \bigoplus_{i \in \mathbb{N}} \mathcal{K}_i$, where each \mathcal{K}_i is a copy of \mathcal{K} . Now since $(\bigoplus_{i \in \mathbb{N}} \mathcal{K}_i) \otimes \mathcal{H} \cong \bigoplus_{i \in \mathbb{N}} (\mathcal{K}_i \otimes \mathcal{H})$, we may define, for each $X \in B(\mathcal{K} \otimes \mathcal{H})$, the operator $\tilde{X} \in B(\tilde{\mathcal{K}} \otimes \mathcal{H})$ by diagonal action on the components of the sum, i.e. $\tilde{X}(\bigoplus_{i \in \mathbb{N}} (k_i \otimes h_i)) := \bigoplus_{i \in \mathbb{N}} X(k_i \otimes h_i)$. If we now define the vector $\tilde{\varphi} \in \tilde{\mathcal{K}}$ by $\tilde{\varphi} = \bigoplus_i \beta_i \varphi_i$, then we have for all $X \in \mathcal{K} \otimes \mathcal{H}$ and $\chi \in \mathcal{H}$:

$$\begin{aligned} \mathbf{tr}(\tilde{U}(|\tilde{\varphi}\rangle\langle\tilde{\varphi}| \otimes |\chi\rangle\langle\chi|)\tilde{U}^* \tilde{X}) &= \langle \bigoplus_{i \in \mathbb{N}} (\beta_i \varphi_i \otimes \chi) | \tilde{U}^* \tilde{X} \tilde{U} \bigoplus_{j \in \mathbb{N}} (\beta_j \varphi_j \otimes \chi) \rangle_{\tilde{\mathcal{K}} \otimes \mathcal{H}} \\ &= \langle \bigoplus_{i \in \mathbb{N}} (\beta_i \varphi_i \otimes \chi) | \bigoplus_{j \in \mathbb{N}} U^* X U (\beta_j \varphi_j \otimes \chi) \rangle_{\tilde{\mathcal{K}} \otimes \mathcal{H}} \\ &= \sum_{i \in \mathbb{N}} |\beta_i|^2 \langle (\varphi_i \otimes \chi) | U^* X U (\varphi_i \otimes \chi) \rangle_{\mathcal{K} \otimes \mathcal{H}} \\ &= \sum_{i \in \mathbb{N}} |\beta_i|^2 \mathbf{tr}(U(|\varphi_i\rangle\langle\varphi_i| \otimes |\chi\rangle\langle\chi|)U^* X) \\ &= \mathbf{tr}(U(\tau \otimes |\chi\rangle\langle\chi|)U^* X) \end{aligned}$$

The second step is due to the diagonal action of the operators on $\tilde{\mathcal{K}} \otimes \mathcal{H}$. The problem is now reduced to the vector-case by applying the above to $\chi = \psi$, $\chi = \psi_0$ or $\chi = \psi_1$ and on the other hand $X = A$, $X = B$ or $X = B^2$. \square

The backbone of theorem 3.3 is formed by the special case that $\sigma_0 = \sigma_1 = 0$, $[A, B] = 0$ and $\tau = |\phi\rangle\langle\phi|$, which allows for a short and transparent proof.

In order to arrive at a physical interpretation of theorem 3.3, we focus on the case $B = \tilde{B} \otimes \mathbf{1}$, when information is transferred from \mathcal{H} to \mathcal{K} . Indeed, examining \mathcal{K} at time t yields information about \mathcal{H} at time 0.

Quality of Information Transfer

A small ratio $\frac{\sigma_0 + \sigma_1}{|b_0 - b_1|}$ indicates a good quality of information transfer. The ratio equals 0 in the perfect case, when $\sigma_0 = \sigma_1 = 0$. Thus $\tilde{B} \otimes \mathbf{1}$ takes a definite value of either b_0 or b_1 , depending on whether the initial state of \mathcal{H} was $|\psi_0\rangle$ or $|\psi_1\rangle$. In this case, one can infer the initial state of \mathcal{H} with certainty by inspecting only the ancilla \mathcal{K} . More generally, it is still possible to reliably determine from the ancilla \mathcal{K} whether the open system \mathcal{H} was initially in state $|\psi_0\rangle$ or $|\psi_1\rangle$ as long as the standard deviations are small compared to the difference in mean, $\sigma_0, \sigma_1 \ll |b_0 - b_1|$.

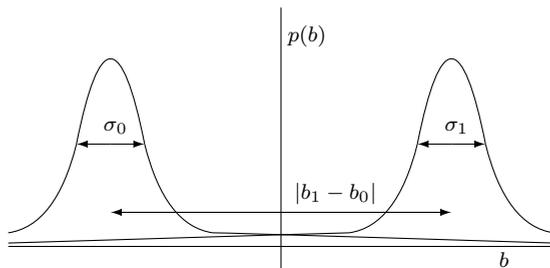


FIG. 2: Probability densities p of B for input $|\psi_0\rangle$ and $|\psi_1\rangle$

As the ratio increases, the restriction (3.4) gets less severe, reaching triviality at $\sigma_0 + \sigma_1 = 2|b_0 - b_1|$.

Decoherence on the Commutant of the Pointer

Assume perfect information transfer, i.e. $\sigma_0 = \sigma_1 = 0$. If $[A, B] = 0$, then theorem 3.3 says that coherent and mixed initial states yield identical distributions of A at time t . In order to distinguish, at time t , whether or not \mathcal{H} was in a pure state at time 0, we will have to use observables A which do not commute with B . But then A and B cannot be observed simultaneously. Summarizing:

At time t , it is possible to distinguish whether \mathcal{H} was in state ψ_0 or ψ_1 at time 0. It is also possible to distinguish whether \mathcal{H} was in state ψ or $|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|$ at time 0. But it is not possible to do both.

We emphasize that this holds even when one has all observables of the entire closed system $\mathcal{K} \otimes \mathcal{H}$ at one's disposal.

Decoherence on the Open System

We consider the final state of the open system \mathcal{H} , obtained from the final state of the closed system $\mathcal{K} \otimes \mathcal{H}$ by tracing out the degrees of freedom of the ancilla \mathcal{K} : an initial state $\rho \in S(\mathcal{H})$ yields final state $\mathbf{tr}_{\mathcal{K}}(T(\rho)) \in S(\mathcal{H})$.

Suppose that information is transferred to a pointer $B = \tilde{B} \otimes \mathbf{1}$ in the ancilla \mathcal{K} with perfect quality, that is $\sigma_0 = \sigma_1 = 0$. Since $[\mathbf{1} \otimes \tilde{A}, \tilde{B} \otimes \mathbf{1}] = 0$, we see from theorem 3.3 that $\mathbf{tr}(T(|\psi\rangle\langle\psi|)(\mathbf{1} \otimes \tilde{A}))$ must equal $\mathbf{tr}(T(|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|)(\mathbf{1} \otimes \tilde{A}))$ for all $\tilde{A} \in B(\mathcal{H})$, or equivalently

$$\mathbf{tr}_{\mathcal{K}}(T(|\psi\rangle\langle\psi|)) = \mathbf{tr}_{\mathcal{K}}\left(T(|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|)\right). \quad (3.5)$$

In words:

Suppose that at time t , by making a hypothetical measurement of \tilde{B} on the ancilla, it would be possible to distinguish perfectly whether the open system had been in state ψ_0 or ψ_1 at time 0. Then, by looking only at the observables of the open system, it is not possible to distinguish whether \mathcal{H} had been in the pure state $\psi = \alpha_0\psi_0 + \alpha_1\psi_1$ or the collapsed state $|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|$ at time 0.

This statement holds true, regardless whether \tilde{B} is actually measured or not. (So we do not assume here that such a measurement is physically possible.) We have shown that the map $M_0 + M_1 = \mathbf{tr}_{\mathcal{K}} \circ T$, with $T : \mathcal{T}(\mathcal{H}) \rightarrow \mathcal{T}(\mathcal{K} \otimes \mathcal{H})$ the information-transfer operation defined by $T(\rho) := U(\tau \otimes \rho)U^*$, constitutes a physical realization of the abstract decoherence mapping $(M_0 + M_1)$ of section 3.2.

All in all, we have proven that decoherence is an unavoidable consequence of information transfer out of an open system.

Example

The simplest possible example of unitary information transfer is the following. Let $\mathcal{K} \sim \mathcal{H} \sim \mathbb{C}^2$ be the Hilbert space of a qubit; let $\psi_0 = (1, 0)$ and $\psi_1 = (0, 1)$ be the ‘computational basis’, and let $U : \mathbb{C}^2 \otimes \mathbb{C}^2 \rightarrow \mathbb{C}^2 \otimes \mathbb{C}^2$ be the ‘controlled-not gate’. Explicitly, U is defined by $U|\psi_1 \otimes \psi_1\rangle = |\psi_0 \otimes \psi_1\rangle$, $U|\psi_0 \otimes \psi_1\rangle = |\psi_1 \otimes \psi_1\rangle$, $U|\psi_1 \otimes \psi_0\rangle = |\psi_1 \otimes \psi_0\rangle$, and $U|\psi_0 \otimes \psi_0\rangle = |\psi_0 \otimes \psi_0\rangle$. That is, it flips the first qubit whenever the second qubit is set to 1. Let τ be the 0 state of the first qubit.

Since the initial state of the second qubit can be read off from the first, this situation satisfies the hypotheses of theorem 3.3 with $B = \sigma_z \otimes \mathbf{1}$ and $\sigma_0 = \sigma_1 = 0$. We verify equation (3.5). For any state $|\psi\rangle = \alpha_0|\psi_0\rangle + \alpha_1|\psi_1\rangle$:

$$\begin{aligned} U|\psi_0 \otimes \psi\rangle &= \alpha_0|\psi_0 \otimes \psi_0\rangle + \alpha_1|\psi_1 \otimes \psi_1\rangle := |\theta\rangle; \\ \mathbf{tr}_{\mathcal{K}}(|\theta\rangle\langle\theta|) &= |\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|. \end{aligned}$$

Thus we have $\mathbf{tr}_{\mathcal{K}}(T(|\psi\rangle\langle\psi|)) = |\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|$. This agrees with equation (3.5), as one can easily check that $\mathbf{tr}_{\mathcal{K}}(T(|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|))$ equals $|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|$ as well.

3.3.2 Unitary Information Transfer and State Collapse

We have derived that, in the context of information transfer to an ancillary system, the initial states $|\psi\rangle\langle\psi|$ and $|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|$ lead to the same final state. This is decoherence.

State collapse is a much stronger statement: if outcome ‘0’ is observed, then the system will further behave as if its initial state had been ψ_0 instead of ψ . Similarly, if outcome ‘1’ is observed, then the system will behave as if its initial state had been ψ_1 . Now suppose that we ignore the outcome. Since ‘0’ happens with probability $|\alpha_0|^2$ and ‘1’ with probability $|\alpha_1|^2$, the system will behave as if its initial state had been $|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|$. We see that collapse implies decoherence.

The converse does not hold however: imagine a Stern-Gerlach experiment, in which a beam of particles in a σ_x -eigenstate is split in two according to spin in the z -direction. State collapse is the statement that one beam consists of particles with positive spin, the other of particles with negative spin and that both beams have equal intensity. Decoherence is the statement that both outgoing beams *together* consist for 50% of positive-spin particles and for 50% of negative-spin particles. The former statement is strictly stronger than the latter, and deserves separate investigation.

We will therefore answer the following question: suppose that we transfer information to an ancilla \mathcal{K} , and then separate \mathcal{K} from \mathcal{H} , dividing \mathcal{H} into subensembles according to outcome. What states do we use to describe these subensembles?

Joint Probability Distributions

A special case of an observable is an *event* p , which in quantum mechanics is represented by a projection P . The relative frequency of occurrence of p is given by $\mathbb{P}(p = 1) = \mathbf{tr}(\rho P)$.

The projection $\mathbf{1} - P$ is interpreted as ‘not p ’. Furthermore, if a projection Q corresponding to an observable q commutes with P , then PQ is again a projection. According to quantum mechanics, p and q can then be observed simultaneously, and the projection PQ is interpreted as the event ‘ p and q are both observed’.

A state ρ therefore induces a joint probability distribution on p and q :

$$\begin{aligned} \mathbf{tr}(\rho PQ) &= \mathbb{P}(p = 1, q = 1) \quad , \quad \mathbb{P}(p = 0, q = 1) = \mathbf{tr}(\rho(\mathbf{1} - P)Q) \\ \mathbf{tr}(\rho P(\mathbf{1} - Q)) &= \mathbb{P}(p = 1, q = 0) \quad , \quad \mathbb{P}(p = 0, q = 0) = \mathbf{tr}(\rho(\mathbf{1} - P)(\mathbf{1} - Q)) \end{aligned}$$

Particularly relevant is the case in which ρ is a state on a combined space $\mathcal{K} \otimes \mathcal{H}$, and the projections are of the form $Q \otimes \mathbf{1}$ and $\mathbf{1} \otimes P$. (The commuting projections are properties of different systems.) We then see that $\mathbb{P}(p = 1, q = 1)$ equals $\mathbf{tr}((\mathbf{1} \otimes P)(Q \otimes \mathbf{1})\rho)$, which in turn is equal to $\mathbf{tr}(P \mathbf{tr}_{\mathcal{K}}((Q \otimes \mathbf{1})\rho))$. This holds for all projections P on \mathcal{H} , so that the normalized version of $\mathbf{tr}_{\mathcal{K}}((Q \otimes \mathbf{1})\rho) \in \mathcal{T}(\mathcal{H})$ must be interpreted as the state of \mathcal{H} , given that $q = 1$. Similarly, the normalized version of $\mathbf{tr}_{\mathcal{K}}((\mathbf{1} - Q) \otimes \mathbf{1})\rho \in \mathcal{T}(\mathcal{H})$ is the state of \mathcal{H} , given that $q = 0$ is observed.

Collapse

Let $T : \rho \mapsto U(\tau \otimes \rho)U^*$ from $\mathcal{T}(\mathcal{H})$ to $\mathcal{T}(\mathcal{K} \otimes \mathcal{H})$ be an information transfer from \mathcal{H} to a pointer-projection $Q \in B(\mathcal{K})$. That is, $\mathbf{tr}((Q \otimes \mathbf{1})T(|\psi_0\rangle\langle\psi_0|)) = 0$ and $\mathbf{tr}((Q \otimes \mathbf{1})T(|\psi_1\rangle\langle\psi_1|)) = 1$, so that at time t , one can see from \mathcal{K} whether \mathcal{H} was in state ψ_0 or ψ_1 at time 0.

Since $Q \otimes \mathbf{1}$ commutes with all of $\mathbf{1} \otimes B(\mathcal{H})$, it is possible to separate \mathcal{H} from \mathcal{K} , and divide \mathcal{H} into subensembles according to the outcome of Q . This is done as follows: with any measurement on \mathcal{H} , a simultaneous measurement of Q on \mathcal{K} is performed to determine in which ensemble this particular system should fall. It follows from the above that the 1-ensemble should be described by the normalized version of $M_1(\rho) := \mathbf{tr}_{\mathcal{K}}((Q \otimes \mathbf{1})T(\rho))$, and the 0-ensemble by the normalized version of $M_0(\rho) := \mathbf{tr}_{\mathcal{K}}((\mathbf{1} - Q \otimes \mathbf{1})T(\rho))$. Since Q commutes with $B(\mathcal{H})$, this is just conditioning on a classical probability space at time t . We have arrived at an interpretation of the map $M(\rho) := M_0(\rho) \oplus M_1(\rho)$ of section 3.2.

We will now prove that M takes the form $M(|\psi\rangle\langle\psi|) = |\alpha_0|^2 \mathbf{tr}_{\mathcal{K}}T(|\psi_0\rangle\langle\psi_0|) \oplus |\alpha_1|^2 \mathbf{tr}_{\mathcal{K}}T(|\psi_1\rangle\langle\psi_1|)$. This is a strong physical statement. For instance, any spin-system $\alpha_0|\psi_0\rangle + \alpha_1|\psi_1\rangle$ that is found to have spin 1 in the z -direction may subsequently be treated as if it had been in state ψ_1 at time 0. This is nontrivial: a priori, it is perfectly conceivable that the different initial states ψ_0 and ψ_1 result in different final states, even though they yield the same Q -output.

One could alternatively, (and more traditionally), arrive at the ‘collapse of the wavefunction’ $M(|\psi\rangle\langle\psi|) = |\alpha_0|^2 \mathbf{tr}_{\mathcal{K}}T(|\psi_0\rangle\langle\psi_0|) \oplus |\alpha_1|^2 \mathbf{tr}_{\mathcal{K}}T(|\psi_1\rangle\langle\psi_1|)$ by assuming that, at time 0, the quantum system makes either the jump $|\psi\rangle\langle\psi| \mapsto |\psi_0\rangle\langle\psi_0|$ or the jump $|\psi\rangle\langle\psi| \mapsto |\psi_1\rangle\langle\psi_1|$. Since we arrive at the same conclusion, namely the above ‘collapse of the wavefunction’, using only open systems, unitary transformations and the probabilistic interpretation of quantum mechanics, such an assumption of ‘jumps’ at time 0 is made redundant.

Proposition 3.4 *Let $T : \rho \mapsto U(\tau \otimes \rho)U^*$ from $\mathcal{T}(\mathcal{H})$ to $\mathcal{T}(\mathcal{K} \otimes \mathcal{H})$ satisfy $\mathbf{tr}((Q \otimes \mathbf{1})T(|\psi_0\rangle\langle\psi_0|)) = 0$ and $\mathbf{tr}((Q \otimes \mathbf{1})T(|\psi_1\rangle\langle\psi_1|)) = 1$ for some ‘pointer-projection’ Q on \mathcal{K} . Define a map $M : \mathcal{T}(\mathcal{H}) \mapsto \mathcal{T}(\mathcal{H}) \oplus \mathcal{T}(\mathcal{H})$ by $M(\rho) := \mathbf{tr}_{\mathcal{K}}((\mathbf{1} - Q \otimes \mathbf{1})T(\rho)) \oplus \mathbf{tr}_{\mathcal{K}}((Q \otimes \mathbf{1})T(\rho))$. Then for $\psi = \alpha_0\psi_0 + \alpha_1\psi_1$ we have $M(|\psi\rangle\langle\psi|) = |\alpha_0|^2 \mathbf{tr}_{\mathcal{K}}T(|\psi_0\rangle\langle\psi_0|) \oplus |\alpha_1|^2 \mathbf{tr}_{\mathcal{K}}T(|\psi_1\rangle\langle\psi_1|)$.*

This can be seen almost directly from proposition 3.1:

Proof. Since $M_1(|\psi_0\rangle\langle\psi_0|) \geq 0$ is a positive operator, we may conclude from $\mathbf{tr}(M_1(|\psi_0\rangle\langle\psi_0|)) = 0$ that $M_1(|\psi_0\rangle\langle\psi_0|)$ itself must vanish. Similarly, reversing the role of 0 and 1, we see that $M_0(|\psi_1\rangle\langle\psi_1|) = 0$. From proposition 3.1, we then obtain $M(|\psi\rangle\langle\psi|) = |\alpha_0|^2 M_0(|\psi_0\rangle\langle\psi_0|) \oplus |\alpha_1|^2 M_1(|\psi_1\rangle\langle\psi_1|)$. The proof is completed by noting from $\mathbf{tr}_{\mathcal{K}}((\mathbf{1} - Q \otimes \mathbf{1})T(|\psi_1\rangle\langle\psi_1|)) = 0$ that $\mathbf{tr}_{\mathcal{K}}(T(|\psi_1\rangle\langle\psi_1|)) = \mathbf{tr}_{\mathcal{K}}((Q \otimes \mathbf{1})T(|\psi_1\rangle\langle\psi_1|)) = M_1(|\psi_1\rangle\langle\psi_1|)$, and similarly, again reversing the role of 0 and 1, that $\mathbf{tr}_{\mathcal{K}}(T(|\psi_0\rangle\langle\psi_0|)) = M_0(|\psi_0\rangle\langle\psi_0|)$. \square

We summarize:

Consider an ensemble of systems of type \mathcal{H} in state ψ . Suppose that information is transferred to a pointer-projection Q on an ancillary system \mathcal{K} . Subsequently, the ensemble is divided into two subensembles according to outcome. Then all observations on \mathcal{H} made afterwards, conditioned on the observation that the measurement outcome was 0, will be as if the system had originally been in the collapsed state ψ_0 instead of ψ . No independent ‘collapse postulate’ is needed to arrive at this conclusion.

Example

In the simple model of information transfer introduced in Section 3.3.1, we will now demonstrate why repeated spin-measurements yield identical outcomes.

The probed system is once again a single spin $\mathcal{H} = \mathbb{C}^2$, whereas the ancillary system now consists of two spins, $\mathcal{K} = \mathbb{C}^2 \otimes \mathbb{C}^2$ in initial state $|\psi_0 \otimes \psi_0\rangle$. Repeated information-transfer, first to pointer $\sigma_{z,1}$ and then to $\sigma_{z,2}$, is then represented by the unitary $U := U_2 U_1$ on $\mathcal{K} \otimes \mathcal{H}$. In this expression, U_1 is the controlled not-gate flipping the first qubit of \mathcal{K} if \mathcal{H} is set to 1, and U_2 flips the second qubit of \mathcal{K} if \mathcal{H} is set to 1.

Since $U|\psi_0 \otimes \psi_0 \otimes (\alpha_0 \psi_0 + \alpha_1 \psi_1)\rangle = |\alpha_0 \psi_0 \otimes \psi_0 \otimes \psi_0\rangle + |\alpha_1 \psi_1 \otimes \psi_1 \otimes \psi_1\rangle$, we can explicitly calculate the joint probability distribution on the two pointers $\sigma_{z,1}$ and $\sigma_{z,2}$ in the final state:

$$\begin{aligned} \mathbb{P}(s_{z,1} = 1, s_{z,2} = 1) &= |\alpha_1|^2 & , & & 0 &= \mathbb{P}(s_{z,1} = 1, s_{z,2} = -1) \\ \mathbb{P}(s_{z,1} = -1, s_{z,2} = 1) &= 0 & , & & |\alpha_0|^2 &= \mathbb{P}(s_{z,1} = -1, s_{z,2} = -1) \end{aligned}$$

In particular, we see that if the first outcome is 1 (which happens with probability $|\alpha_1|^2$), then so is the second. proposition 3.4 shows that this is the general situation, independent of the (rather simplistic) details of this particular model.

3.3.3 Information Leakage to the Environment

On closed systems decoherence does not occur, because unitary time evolution preserves the purity of states. However, macroscopic systems are almost never closed.

Imagine, for example, that $\mathcal{H} = \mathbb{C}^2$ represents a two-level atom, and \mathcal{K} some large measuring device. Information about the energy $\mathbf{1} \otimes \sigma_z$ of the atom is transferred to the apparatus, where it is stored as the position $\vec{B} \otimes \mathbf{1}$ of a pointer. Then as soon as information on the pointer-position $\vec{B} \otimes \mathbf{1}$ leaves the system, collapse on the combined atom-apparatus system takes place. For example, a ray of light may reflect on the pointer, revealing its position to the outside world. (See [JZ85].) It is of course immaterial whether or not someone is actually *looking* at the photons. If even the smallest speck of light were to fall on the pointer, the information about the pointer position would already be encoded in the light, causing full collapse on the atom-apparatus system. (See [Zur82] for an example.)

The quality of this information transfer will not be perfect. If a macroscopic system is interacting normally with the outside world, (the occasional photon happens to scatter on it, for instance), then a number of macroscopic observables X will leak information continually, with a macroscopic uncertainty σ . This enables us to apply theorem 3.3. It says that all coherences between eigenstates ψ_{x_1} and ψ_{x_2} of macroscopic observables X are continually vanishing on the macroscopic system \mathcal{L} , provided that their eigenvalues x_1 and x_2 satisfy the inequality $|x_1 - x_2| \gg 2\sigma$. (The pointer, e.g. a beam of light, is outside the system, so that $\delta = 0$.)

Take for example a collection of N spins, with Hilbert space $\mathcal{L} = \bigotimes_{i=1}^N \mathbb{C}^2$. Suppose that the average spin-observables $S_\alpha = \frac{1}{N} \sum_{i=1}^N \sigma_\alpha^i$ (for $\alpha = x, y, z$) are continually being measured with an accuracy² $N^{-1/2} \ll \sigma \ll 1$. Then between macroscopically different eigenstates of S_α , i.e. states for which the eigenvalues satisfy $|s_\alpha - s'_\alpha| \gg \sigma$, coherences are constantly disappearing. However, the information leakage need not have any effect on states which only differ on a microscopic scale. Take for instance $\rho \otimes |+\rangle\langle +|$ and $\rho \otimes |-\rangle\langle -|$, with ρ an arbitrary state on $N - 1$ spins. Indeed, $|s_\alpha - s'_\alpha| \leq 2/N \ll \sigma$, so theorem 3.3 is vacuous in this case: no decoherence occurs.

We see how the variance σ^2 produces a smooth boundary between the macroscopic and the microscopic world: macroscopically distinguishable states (involving S_α -differences $\gg \sigma$) continually suffer from loss of coherence, while states that only differ microscopically (involving S_α -differences $\ll \sigma$) are unaffected.

In case of a system monitored by a macroscopic measurement apparatus, we are interested in coherence between eigenstates of the macroscopic pointer. By definition, these eigenstates are macroscopically distinguishable. We may then give the following answer to the question why it is so hard, in practice, to witness coherence:

If information leaks from the pointer into the outside world, decoherence takes place on the combination of system and measurement apparatus. In practice, macroscopic pointers constantly leak information.

3.4 Closed Systems

Closed systems evolve according to unitary time evolution, so that coherence which is present initially will still be there at later times. Yet on macroscopic systems, coherent superpositions are almost never observed. Why is this so?

3.4.1 Macroscopic Systems

Because of the direct link that it provides between the scale of a system on the one hand, and on the other hand the difficulties in witnessing coherence, we

²Since $[S_x, S_y] \neq 0$, they cannot be simultaneously measured with complete accuracy, see e.g. [Wer01]. However, this problem disappears if the accuracy satisfies $\sigma^2 \geq \frac{1}{2} \|[S_x, S_y]\| = \frac{1}{N}$, see [Jan04]. For large N , (typically $N \sim 6 \times 10^{23}$), this allows for extremely accurate measurement.

feel that the following line of reasoning, essentially due to Hepp [Hep72], is the most important mechanism hiding coherence.

Let us first define what we mean by macroscopic and microscopic observables. We consider a system consisting of N distinct subsystems, i.e. $\mathcal{K} = \bigotimes_{i=1}^N \mathcal{K}_i$. If one thinks of \mathcal{K}_i as the atoms out of which a macroscopic system \mathcal{K} is constructed, N may well be in the order of 10^{23} .

We will define the *microscopic* observables to be the ones that refer only to one particular subsystem \mathcal{K}_i :

Definition 3.1 *An observable $X \in B(\mathcal{K})$ is called microscopic if it is of the form $X = \mathbf{1} \otimes \dots \otimes \mathbf{1} \otimes X_i \otimes \mathbf{1} \otimes \dots \otimes \mathbf{1}$ for some $i \in \{1, 2, \dots, N\}$ and some $X_i \in B(\mathcal{K}_i)$.*

In this situation we will identify $X_i \in B(\mathcal{K}_i)$ with $X \in B(\mathcal{K})$. We take macroscopic observables to be averages of microscopic observables ‘of the same size’:

Definition 3.2 *An observable $Y \in B(\mathcal{K})$ is called macroscopic if it is of the form $Y = \frac{1}{N} \sum_{i=1}^N Y_i$, with $Y_i \in B(\mathcal{K}_i)$ such that $\|Y_i\| \leq \|Y\|$.*

We will only use the term ‘macroscopic’ in this narrow sense from here on, even though there do exist observables which are called ‘macroscopic’ in daily life, but do not fall under the above definition.

Now suppose that we transfer information from a system \mathcal{H} to a macroscopic system $\mathcal{K} = \bigotimes_{i=1}^N \mathcal{K}_i$, using a macroscopic pointer $\tilde{B} \in B(\mathcal{K})$. As explained before, we then have a map $T : \mathcal{T}(\mathcal{H}) \rightarrow \mathcal{T}(\mathcal{K} \otimes \mathcal{H})$ such that the pointer $\tilde{B} \otimes \mathbf{1}$ has different expectation values b_0 and b_1 in the states $T(|\psi_0\rangle\langle\psi_0|)$ and $T(|\psi_1\rangle\langle\psi_1|)$.

Since the pointer \tilde{B} is macroscopic, it is unrealistic to require $T(|\psi_0\rangle\langle\psi_0|)$ and $T(|\psi_1\rangle\langle\psi_1|)$ to be eigenstates of \tilde{B} . Instead, we will require their standard deviations in \tilde{B} to be negligible compared to their difference in mean, that is $\sigma_0 \ll |b_0 - b_1|$ and $\sigma_1 \ll |b_0 - b_1|$.

After this information transfer, we try to distinguish whether the system \mathcal{H} had initially been in the coherent state $\alpha_0|\psi_0\rangle + \alpha_1|\psi_1\rangle$ or in the incoherent mixture $|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|$. We have already shown that this cannot be done by measuring observables in $\mathbf{1} \otimes B(\mathcal{H})$. The following adaptation of theorem 3.3 shows that it is also impossible to do this by measuring macroscopic or microscopic observables on the closed system $\mathcal{K} \otimes \mathcal{H}$.

Corollary 3.5 *Let ψ_0 and ψ_1 be orthogonal unit vectors in a Hilbert space \mathcal{H} and let $\tau \in \mathcal{T}(\mathcal{K})$ be a state on the Hilbert space $\mathcal{K} = \bigotimes_{i=1}^N \mathcal{K}_i$. Let $U : \mathcal{K} \otimes \mathcal{H} \rightarrow \mathcal{K} \otimes \mathcal{H}$ be unitary and define $T : \mathcal{T}(\mathcal{H}) \rightarrow \mathcal{T}(\mathcal{K} \otimes \mathcal{H})$ by $T(\rho) = U(\tau \otimes \rho)U^*$. Let \tilde{B} be a macroscopic observable in $B(\mathcal{K})$, and define $B := \tilde{B} \otimes \mathbf{1}$. Denote by b_j and σ_j^2 its expected value and variance in the state $T(|\psi_j\rangle\langle\psi_j|)$ for $j = 0, 1$. Suppose that $b_0 \neq b_1$. Then for all $\psi = \alpha_0\psi_0 + \alpha_1\psi_1$ with $|\alpha_0|^2 + |\alpha_1|^2 = 1$ and for all microscopic and macroscopic observables $A \in B(\mathcal{K} \otimes \mathcal{H})$, we have*

$$\left| \text{tr}(T(|\psi\rangle\langle\psi|)A) - \text{tr}\left(T(|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|)A\right) \right| \leq \frac{2}{|b_0 - b_1|} (\|B\| + \sigma_0 + \sigma_1) \|A\|.$$

Proof. If A is microscopic, we have the inequality

$$\|[A, B]\| = \left\| \left[A_i, \frac{1}{N} \sum_{j=1}^N B_j \right] \right\| = \frac{1}{N} \|[A_i, B_i]\| \leq \frac{2\|A\|\|B\|}{N}.$$

If A is macroscopic, we have

$$\|[A, B]\| = \left\| \left[\frac{1}{N+1} \sum_{i=0}^N A_i, \frac{1}{N} \sum_{j=1}^N B_j \right] \right\| = \frac{1}{N(N+1)} \sum_{i=1}^N \|[A_i, B_i]\| \leq \frac{2\|A\|\|B\|}{N}.$$

Either way, we can now apply theorem 3.3. \square

3.4.2 Examples

In order to illustrate the above, we discuss four examples of information transfer to a macroscopic system.

The Finite Spin-Chain

We study a single spin $\mathcal{H} = \mathbb{C}^2$ in interaction with a large but finite spin-chain $\mathcal{K} = \bigotimes_{i=1}^N \mathbb{C}^2$, the latter acting as a measurement apparatus. Once again, let $\psi_0 = (1, 0)$ and $\psi_1 = (0, 1)$ be the ‘computational basis’. Initially, all spins in the spin-chain are down: $\tau = |\psi_0 \otimes \dots \otimes \psi_0\rangle \langle \psi_0 \otimes \dots \otimes \psi_0|$. Let $U_i : \mathcal{K} \otimes \mathcal{H} \rightarrow \mathcal{K} \otimes \mathcal{H}$ be the ‘controlled-not gate’, which flips spin number i in the chain whenever the single qubit is set to 1. (We define $U_j = \mathbf{1}$ for $j \notin \{1, 2, \dots, N\}$.)

$$U_i = \mathbf{1} \otimes P_- + \sigma_{x,i} \otimes P_+ \quad \text{with} \quad P_+ = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad P_- = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}.$$

In discrete time $n \in \mathbb{Z}$, the unitary evolution is given by $n \mapsto U_n U_{n-1} \dots U_2 U_1$. (See [Hep72].) This represents a single spin flying over a spin-chain from 1 to N , interacting with spin n at time n .

Obviously, we have $U_N |\psi_0 \otimes \dots \otimes \psi_0\rangle \otimes |\psi_0\rangle = |\psi_0 \otimes \dots \otimes \psi_0\rangle \otimes |\psi_0\rangle$ and $U_N |\psi_0 \otimes \dots \otimes \psi_0\rangle \otimes |\psi_1\rangle = |\psi_1 \otimes \dots \otimes \psi_1\rangle \otimes |\psi_1\rangle$. We consider the average spin of the spin-chain as pointer, $B = \frac{1}{N} \sum_{i=1}^N \sigma_{z,i}$. This makes the map $T : \rho \mapsto U_N \tau \otimes \rho U_N^*$ an information transfer to a macroscopic system. Applying Corollary 3.5 with $b_0 = -1$, $b_1 = 1$ and $\sigma_0 = \sigma_1 = 0$, we see that

$$\left| \text{tr}(T(|\alpha_0 \psi_0 + \alpha_1 \psi_1\rangle \langle \alpha_0 \psi_0 + \alpha_1 \psi_1|)A) - \text{tr}(T(|\alpha_0|^2 |\psi_0\rangle \langle \psi_0| + |\alpha_1|^2 |\psi_1\rangle \langle \psi_1|)A) \right|$$

is bounded by $\frac{1}{N} \|A\|$ for all microscopic and macroscopic $A \in B(\mathcal{K} \otimes \mathcal{H})$. It is not hard to see that in this particular model, the estimated quantity is even identically zero. Indeed, $\langle \psi_0 \otimes \dots \otimes \psi_0 | X_i \psi_1 \otimes \dots \otimes \psi_1 \rangle = \langle \psi_0 | \psi_1 \rangle^{N-1} \langle \psi_0 | X_i \psi_1 \rangle = 0$ for all microscopic X_i .

Of course coherence *can* be detected on the closed system $\mathcal{K} \otimes \mathcal{H}$, but only using observables that are neither macroscopic nor microscopic, such as the ‘product of all spins’ $\sigma_x \otimes \dots \otimes \sigma_x$.

Finite Spin-Chain at Nonzero Temperature

A more realistic initial state for the spin-chain is the nonzero-temperature state $\tau_\beta = \frac{e^{-\beta H}}{\text{tr} e^{-\beta H}}$. For the spin-chain Hamiltonian we will take $H = \sum_i \sigma_{z,i} = NB$, so that τ_β becomes the tensor product of N copies of the \mathbb{C}^2 -state

$$\hat{\tau}_\beta = \frac{1}{e^\beta + e^{-\beta}} \begin{pmatrix} e^{-\beta} & 0 \\ 0 & e^\beta \end{pmatrix}.$$

With the same time evolution as before, we have $T(|\psi_0\rangle\langle\psi_0|) = |\psi_0\rangle\langle\psi_0| \otimes \rho_\beta$ and $T(|\psi_1\rangle\langle\psi_1|) = |\psi_1\rangle\langle\psi_1| \otimes \rho_{-\beta}$. Again we choose the mean energy B as our pointer. A brief calculation shows that $\text{tr}(B\tau_\beta) = \frac{e^{-\beta} - e^\beta}{e^\beta + e^{-\beta}} =: \varepsilon(\beta)$ and that $\text{tr}(B^2\tau_\beta) - \text{tr}(B\rho_\beta)^2 = \frac{1}{N}(1 - \varepsilon^2(\beta))$. Corollary 3.5 now gives us, for microscopic and macroscopic A ,

$$\begin{aligned} & \left| \text{tr}(T(|\alpha_0\psi_0 + \alpha_1\psi_1\rangle\langle\alpha_0\psi_0 + \alpha_1\psi_1|)A) - \text{tr}\left(T(|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|)A\right) \right| \\ & \leq \left(\frac{1}{\varepsilon(\beta)N} + \frac{\sqrt{1 - \varepsilon^2(\beta)}}{\varepsilon(\beta)\sqrt{N}} \right) \|A\|. \end{aligned}$$

For large N , we see that the term $\sim \frac{1}{N}$ due to the fact that $[A, B] \neq 0$ is dominated by the thermodynamical fluctuations, which of course go as $\sim \frac{1}{\sqrt{N}}$. In statistical physics, it is standard practice to neglect even the latter.

Energy as a Pointer

Hamiltonians often fail to be macroscopic in our narrow sense of the word, since they are generically unbounded and contain interaction terms. However, this does not imply failure of our scheme to estimate coherence.

For example, consider an N -particle system with Hilbert space $\mathcal{K} = \bigotimes_{i=1}^N \mathcal{K}_i$ and Hamiltonian $H = \sum_{i=1}^N \frac{p_i^2}{2m_i} + V(x_1, x_2, \dots, x_N)$. Information is transferred from \mathcal{H} to \mathcal{K} with H as pointer, that is the two states $\text{tr}_{\mathcal{H}}(T(|\psi_0\rangle\langle\psi_0|))$ and $\text{tr}_{\mathcal{H}}(T(|\psi_1\rangle\langle\psi_1|))$ have different energies E and E' . Without loss of generality, assume that they are vector states: $\text{tr}_{\mathcal{H}}(T(|\psi_0\rangle\langle\psi_0|)) = |\psi\rangle\langle\psi|$ and $\text{tr}_{\mathcal{H}}(T(|\psi_1\rangle\langle\psi_1|)) = |\psi'\rangle\langle\psi'|$. (Density matrices can always be represented as vectors on a different Hilbert space, cf. the proof of theorem 3.3.)

We thus have two vector states $|\psi\rangle$ and $|\psi'\rangle$ with two different energies, $E := \langle\psi|H\psi\rangle$ and $E' := \langle\psi'|H\psi'\rangle$. We estimate the coherence between $|\psi\rangle$ and $|\psi'\rangle$ on x_n , the position of particle n .

$$\begin{aligned} (E - E')\langle\psi|x_n\psi'\rangle &= \langle E\psi|x_n\psi'\rangle - \langle x_n\psi|E'\psi'\rangle \\ &= \langle H\psi - (H - E)\psi|x_n\psi'\rangle - \langle x_n\psi|H\psi' - (H - E')\psi'\rangle \\ &= \langle [H, x_n]\psi|\psi'\rangle - \langle (H - E)\psi|x_n\psi'\rangle + \langle x_n\psi|(H - E')\psi'\rangle \end{aligned}$$

Now since $[H, x_n] = \frac{1}{2m_n}[p_n^2, x_n] = \frac{-i\hbar p_n}{m_n}$, we can apply the Cauchy-Schwarz inequality in each term to obtain

$$|E - E'| |\langle \psi | x_n \psi' \rangle| \leq \frac{\hbar}{m_n} \sqrt{\langle \psi | p_n^2 \psi \rangle} + \sqrt{\langle \psi | x_n^2 \psi \rangle} \sqrt{\langle \psi' | (H - E')^2 \psi' \rangle} \\ + \sqrt{\langle \psi' | x_n^2 \psi' \rangle} \sqrt{\langle \psi | (H - E)^2 \psi \rangle}.$$

If we define the characteristic speed $V_n := \sqrt{\langle \psi | (p_n/m_n)^2 \psi \rangle}$, the characteristic positions $X_n := \sqrt{\langle \psi | x_n^2 \psi \rangle}$ and $X'_n := \sqrt{\langle \psi' | x_n^2 \psi' \rangle}$, and the standard deviations $\sigma := \sqrt{\langle \psi | (H - E)^2 \psi \rangle}$ and $\sigma' := \sqrt{\langle \psi' | (H - E')^2 \psi' \rangle}$, we obtain

$$|\langle \psi | x_n \psi' \rangle| \leq \frac{\hbar V_n + \sigma X'_n + \sigma' X_n}{|E - E'|}.$$

As such, this doesn't tell us very much. We will have to make some physically plausible assumptions on the state of the system in order to obtain results. First, we assume that the system is encased in an $L \times L \times L$ box so that $X_n, X'_n \leq L$. Also, we assume $V_n < c$. This yields $|\langle \psi | x_n \psi' \rangle| \leq \frac{\hbar c + L(\sigma + \sigma')}{|E - E'|}$. Secondly, we assume that scaling the system in any meaningful way will produce $|E - E'| \sim N$ and $\sigma + \sigma' \sim \sqrt{N}$, so that the coherence on x_n approaches zero as $\sim \frac{1}{\sqrt{N}}$. Notice the almost thermodynamic lack of detail required for this estimate.

Schrödinger's Cat

Let us finally analyze the rather drastic extraction of information from a radioactive particle that has become known³ as 'Schrödinger's cat'. (See [Sch35].) The experiment is performed as follows. We are interested in a radioactive particle. Is it in a decayed state ψ_0 or in a non-decayed state ψ_1 ?

In order to determine this, we set up the following experiment. A Geiger counter is placed next to the radioactive particle. If the particle decays, then the Geiger counter clicks. A mechanism then releases a hammer, which smashes a vial of hydrocyanic acid, killing a cat. All of this happens in a closed box no higher than $1m$, and completely impenetrable to information. A measurement of the atom is done as follows: first, place it inside the box. Then wait for a period of time that is long compared to the decay time of the atom. Finally, open the box, and inspect whether the cat has dropped dead or is still standing upright.

The atom is described by a Hilbert space \mathcal{H} , the combination of Geiger counter, mechanism, hammer, vial and cat by a Hilbert space \mathcal{K} . Initially, the latter is prepared in a state $|\theta\rangle$. As a pointer, we take the center of mass of the cat, $Z := \frac{1}{N} \sum_{i=1}^N z_i$. In this expression, N is the amount of atoms out of which the cat is constructed, and z_i is the z -component of particle number i . (It is a harmless assumption that all atoms in the cat have the same mass.) Since the box only measures $1m$ in height, we may take $\|Z\| = 1$. The unitary evolution

³Actually, Schrödinger's proposal was slightly different. In the original thought experiment, death of the cat was correlated with decay of the atom at time t instead of 0, which wouldn't make it an information transfer in our sense of the word.

$U \in B(\mathcal{K} \otimes \mathcal{H})$ then produces $U|\psi_0 \otimes \theta\rangle := |\gamma_0\rangle$ and $U|\psi_1 \otimes \theta\rangle := |\gamma_1\rangle$, which are eigenstates⁴ of Z with different eigenvalues.

Suppose that, initially, the atom is either in the decayed state ψ_0 with probability $|\alpha_0|^2$ or in the non-decayed state ψ_1 with probability $|\alpha_1|^2$. That is, the initial state is the incoherent mixture $|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|$. By linearity, the final state is then the incoherent state $|\alpha_0|^2|\gamma_0\rangle\langle\gamma_0| + |\alpha_1|^2|\gamma_1\rangle\langle\gamma_1|$.

On the other hand, if the atom starts out in the coherent superposition $\alpha_0|\psi_0\rangle + \alpha_1|\psi_1\rangle$, then the combined system ends up in the coherent state $U(\alpha_0\psi_0 + \alpha_1\psi_1) \otimes \theta = \alpha_0|\gamma_0\rangle + \alpha_1|\gamma_1\rangle$.

The question is now this: why do we not notice the difference between these two situations if we open the box? First of all, according to theorem 3.3 (and the observations following it in section 3.3.1), it is impossible to detect coherence between γ_0 and γ_1 *and* ascertain the position of the cat. Upon opening the black box, we must make a choice.

Secondly, according to the discussion in section 3.3.3, the coherences between the macroscopically different states γ_0 and γ_1 are extremely volatile. Any speck of light falling on the cat will reveal its position with reasonable accuracy, causing the coherence to disappear according to theorem 3.3.

Yet even if we were able to open the box without any information on the position of the cat leaking out, even then would we be unable to detect coherence between γ_0 and γ_1 . Apply Corollary 3.5 to the transfer of information from atom to cat. We have $\sigma_0 = \sigma_1 = 0$, and with pointer Z we have $\|Z\| = 1$ (the height of the box is 1 m) and $z_1 - z_0 = 0.1$ (the difference between a cat that is standing up and one that has dropped dead is 10 cm). We then obtain for all macroscopic and microscopic A :

$$|\langle\alpha_0\gamma_0 + \alpha_1\gamma_1 | A\alpha_0\gamma_0 + \alpha_1\gamma_1\rangle - (|\alpha_0|^2\langle\gamma_0 | A\gamma_0\rangle + |\alpha_1|^2\langle\gamma_1 | A\gamma_1\rangle)| \leq \frac{20}{N}\|A\|.$$

On the subset of observables we are normally able to measure, the distinction between coherent and incoherent mixtures practically vanishes for $N \sim 10^{23}$. For all practical intents and purposes, it is completely harmless to assume that the final state of the cat is $|\alpha_0|^2|\gamma_0\rangle\langle\gamma_0| + |\alpha_1|^2|\gamma_1\rangle\langle\gamma_1|$ instead of $\alpha_0|\gamma_0\rangle + \alpha_1|\gamma_1\rangle$. But it would be false to state that the former has actually been observed.

3.5 Conclusion

In open systems, we have proven that decoherence is a necessary consequence of information transfer to the outside. More in detail, we have reached the following conclusions:

- Suppose that an open system \mathcal{H} interacts with an ancillary system \mathcal{K} in such a way, that it is possible, in principle, to determine from \mathcal{K} whether \mathcal{H} had been in state ψ_0 or ψ_1 before the interaction. If \mathcal{H} started out in a coherent state $\alpha_0|\psi_0\rangle + \alpha_1|\psi_1\rangle$, then it will behave after the information

⁴As discussed before, it would be more realistic to allow for a nonzero variance $0 < \sigma_j \ll 1$ instead of requiring θ_j to be eigenstates of Z . We use $\sigma_j = 0$ for clarity, leaving the argument essentially unchanged.

transfer as if it had started out in the incoherent mixture $|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|$ instead. This is called ‘decoherence’.

- Suppose again that the information whether \mathcal{H} was in state ψ_0 or ψ_1 is transported to an ancillary system \mathcal{K} . This is done with an ensemble of \mathcal{H} -systems described by the state $\alpha_0|\psi_0\rangle + \alpha_1|\psi_1\rangle$. The ensemble is then split into subensembles, according to outcome. The ‘0-ensemble’ then behaves as if it had been in state ψ_0 at the beginning of the procedure, and the ‘1-ensemble’ as if it had started in state ψ_1 . This is called ‘state collapse’.
- These results were obtained entirely within the framework of traditional quantum mechanics and unitary time evolution on a larger, closed system containing \mathcal{H} . No ‘reduction-postulate’ is needed. From proposition 3.1, we see that any information extraction causes collapse, quite independent of its particular mechanism.
- On the closed system containing the smaller, open one no decoherence occurs in principle. In practice however, closed systems are very hard to achieve. We have argued that information transfer from a macroscopic observable A , performed with macroscopic precision σ , causes decoherence between eigenstates of A if their values satisfy $\sigma \ll |a_1 - a_0|$. Since information on macroscopic observables tends to leak out, coherence between macroscopically different states tends to vanish.

Still, even if the combined system $\mathcal{K} \otimes \mathcal{H}$ is considered perfectly closed, there are some results to be obtained. Again, we investigated the case that a system \mathcal{H} interacts unitarily with a system \mathcal{K} in such a way that the information whether \mathcal{H} was in state ψ_0 or ψ_1 can be read off from a pointer in \mathcal{K} . We have reached the following conclusions concerning the closed system $\mathcal{K} \otimes \mathcal{H}$:

- Using only observables on the closed system that commute with the pointer, it is impossible to detect whether \mathcal{H} had started out in state $\alpha_0|\psi_0\rangle + \alpha_1|\psi_1\rangle$ or $|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|$. Physically, this means that it is impossible to distinguish between coherent and incoherent initial states while at the same time distinguishing between ψ_0 and ψ_1 .
- Suppose that the closed system $\mathcal{K} \otimes \mathcal{H}$ is macroscopic, and that one has access to its macroscopic and microscopic observables only. Then it is almost impossible to distinguish whether \mathcal{H} had started out in state $\alpha_0|\psi_0\rangle + \alpha_1|\psi_1\rangle$ or $|\alpha_0|^2|\psi_0\rangle\langle\psi_0| + |\alpha_1|^2|\psi_1\rangle\langle\psi_1|$. We have obtained upper bounds on the coherences $\langle\psi_0|A\psi_1\rangle$, evaluated on microscopic or macroscopic A . Assuming perfect information transfer ($\sigma_0 = \sigma_1 = 0$), they approach zero as $\sim \frac{1}{N}$, where N is the size of the system.

In short: no decoherence ever occurs on perfectly closed systems, even if they are macroscopic. It is just very hard to distinguish coherent from incoherent states, creating the false impression that it does.

The link between decoherence and macroscopic systems was brought forward by Klaus Hepp in his fundamental paper [Hep72], where he considered infinite closed systems, displaying decoherence in infinite time. In infinite systems, the microscopic observables form a non-commutative C^* -algebra \mathcal{A} . Its weak closure \mathcal{A}'' is considered as the (von Neumann-)algebra of all observables. The macroscopic observables form a commutative algebra \mathcal{C} which is contained in the centre of \mathcal{A}'' , i.e. $\mathcal{C} \subset \mathcal{Z} = \{Z \in \mathcal{A}'' \mid [Z, A] = 0 \ \forall A \in \mathcal{A}''\}$, yet is almost disjoint from the microscopic observables: $\mathcal{C} \cap \mathcal{A} = \mathbb{C}\mathbf{1}$. Transfer of information to a macroscopic observable therefore implies perfect decoherence on all microscopic and macroscopic observables (cf. section 3.3.1).

Unfortunately, this transfer cannot be done by any automorphic time evolution, since the macroscopic observables are central. Hepp proposed information transfer by a $t \rightarrow \infty$ limit of automorphisms. He was able to show that this causes decoherence in the weak-operator sense. That is, on each fixed microscopic observable, the coherence becomes arbitrarily small for sufficiently large time t .

The paper was criticized by John Bell a few years later [Bel75], on the grounds that, for each fixed time t , there are observables to be found on which coherence is not small. Since Bell was of the opinion that a ‘wave packet reduction’, even on closed systems, ‘takes over from the Schrödinger equation’, this was not to his satisfaction. He did agree however that these observables would become arbitrarily difficult to observe in practice for large t .

By considering large but finite closed systems subject to unitary time evolution, we hope to clarify the role that macroscopic systems play in making us mistake coherent superpositions for classical mixtures. It seems striking that the same, simple mathematics can also be used to understand why open systems do undergo decoherence as soon as they lose information.

Chapter 4

Unifying Decoherence and the Heisenberg Principle

The laws of quantum mechanics impose severe restrictions on the quality of measurement. In this chapter, we investigate some of these restrictions from a quantum probabilistic point of view. We introduce state independent figures of merit for measurement performance, and then use them to prove sharp bounds.

4.1 Introduction

In quantum probability theory, an observable is modelled by a self-adjoint operator A in a von Neumann algebra \mathcal{A} , and a state by a positive, normalized, linear map $\rho : \mathcal{A} \rightarrow \mathbb{C}$. As discussed in chapter 2, a (normal) state induces a probability measure on the spectrum of A . It is the objective of a *quantum measurement* to portray this probability measure as faithfully as possible.

According to the uncertainty relation of Heisenberg, Kennard and Robertson

$$\sigma_X \sigma_Y \geq \frac{1}{2} |\rho([X, Y])|,$$

(see [Hei27, Ken27, Rob29]), there is an inherent variance in the quantum state. Furthermore, quantum theory puts severe restrictions on the performance of measurement. These restrictions, which come on top of the measurement restrictions implied by the above uncertainty relation, fall into four distinct classes.

- I The impossibility of perfect *joint measurement*. It is not possible to perform a simultaneous measurement of two noncommuting observables in such a way that both measurements have perfect quality.
- II The *Heisenberg principle*, (see [Hei27]). This states that quantum information cannot be extracted from a system without disturbing that system.
- III The impossibility of *classical coding*. One cannot extract classical information from a quantum system and then use this information to reconstruct the quantum state.

IV *The collapse of the wave function.* When information is extracted from a quantum system, a so-called decoherence is experimentally known to occur on this system.

We aim to find a quantitative description for each of these fundamental measurement restrictions.

The subject of uncertainty relations in quantum measurement is already endowed with an extensive literature. For example, the Heisenberg principle and the impossibility of joint measurement are quantitatively illustrated in [AK65, Oza03, Ish91, Hal04]. However, the inequalities in these papers depend on the state ρ , which somewhat limits their practical use. Indeed, the bound on the measurement quality can only be calculated if the state ρ is known, in which case there is no need for a measurement in the first place.

In section 4.2, we introduce two state independent figures of merit: the *maximal added variance* and the *measurement infidelity*. These will lead us quite naturally to state independent bounds on the performance of measurement.

In section 4.3, we introduce our main tool in the proof of these inequalities: on operator valued Cauchy-Schwarz inequality.

In section 4.4, we prove a sharp, state independent bound on the performance of jointly unbiased measurement. This generalizes the impossibility of perfect joint measurement.

In section 4.5, we prove a sharp, state independent bound on the performance of a measurement in terms of the maximal disturbance that it causes. This generalizes the Heisenberg principle.

In section 4.6, using a result of Keyl and Werner, we prove a sharp, state independent bound on the performance of classical coding.

In section 4.7, we investigate the collapse of the wave function. In contrast with the Heisenberg principle and its abundance of inequalities, this phenomenon has mainly been investigated in specific examples (see e.g. [Hep72, Zur82, JZ85]). Although there are some bounds on the remaining coherence in terms of the measurement quality (for example [Sew05], but also theorem 3.3 in the previous chapter), a sharp, information-theoretic inequality does not yet appear to exist. We will provide such an inequality, and prove a sharp upper bound on the amount of coherence that can survive information transfer.

First of all, this generalizes the collapse of the wave function, and shows that decoherence is a mathematical consequence of information extraction. But it also shows that no information can be extracted if all coherence is left perfectly intact, which is precisely the Heisenberg principle. It is therefore a *unified* description of both the Heisenberg principle and the collapse of the wave function.

4.2 Information Transfer

We wish to model information transfer between quantum systems. Recall from chapter 2 that a quantum system is modelled by a *quantum probability space*, a pair (\mathcal{A}, ρ) of a von Neumann-algebra \mathcal{A} and a normal state $\rho \in \mathcal{S}(\mathcal{A})$. Information transfer then proceeds as follows. We are interested in the distribution

of the observable $A \in \mathcal{A}$, with the system \mathcal{A} in some unknown state ρ . We perform the operation $T^* : \mathcal{S}(\mathcal{A}) \rightarrow \mathcal{S}(\mathcal{B})$, and then observe the ‘pointer’ B in the system \mathcal{B} in order to obtain information on A .

4.2.1 General Information Transfer

We take the point of view (see [Hal04]) that *any* CP-map $T : \mathcal{B} \rightarrow \mathcal{A}$ is an information transfer from *any* observable $A \in \mathcal{A}$ to *any* pointer $B \in \mathcal{B}$. Undeniably though, for the purpose of transferring information, some operations are better than others. What does it mean for a CP-map T to be ‘good’ at transferring information?

Perfect Information Transfer

The initial state ρ defines a probability measure \mathbb{P}_i on the spectrum of A by $\mathbb{P}_i(S) := \rho(\mathbf{1}_S(A))$, where $\mathbf{1}_S(A)$ denotes the spectral projection of A associated to the set S . Similarly, the final state $T^*\rho$ defines a probability distribution \mathbb{P}_f on the spectrum of B .

The object of information transfer is to gain information on \mathbb{P}_i , the probability measure of the observable A *before* the operation, by looking at the probability measure \mathbb{P}_f of the pointer B *afterwards*. We will therefore consider an information transfer to be perfect if $\mathbb{P}_i = \mathbb{P}_f$ for any initial state ρ . Translated to algebra, this reads as follows.

Definition 4.1 *An information transfer $T : \mathcal{B} \rightarrow \mathcal{A}$ from $A \in \mathcal{A}$ to $B \in \mathcal{B}$ is called perfect if $T(B) = A$ and if the restriction of T to B'' , the von Neumann algebra generated by B , is an isomorphism $B'' \rightarrow A''$ of commutative non Neumann algebras.*

Indeed, $\mathbb{P}_i = \mathbb{P}_f$ if and only if $\mathbb{E}_i(g) = \mathbb{E}_f(g)$ for any measurable function g . In other words, $T^*\rho(g(B)) = \rho(g(A))$ for all $\rho \in \mathcal{S}(\mathcal{A})$ and $g \in L^\infty(\mathbb{R})$. But if $\rho(T(g(B)))$ equals $\rho(g(A))$ for any normal state ρ , we must have $T(g(B)) = g(A)$ and vice versa. Since $X'' = \{g(X) | g \in L^\infty(\mathbb{R})\}$, the two notions of perfection coincide.

Measurement Infidelity

Not unlike a man, an information transfer which is interesting is not necessarily perfect. The aim of information transfer is to infer \mathbb{P}_i from \mathbb{P}_f , so surely T is *good* if \mathbb{P}_i and \mathbb{P}_f are *close*.

Recall that both \mathbb{P}_i and \mathbb{P}_f depend on the state ρ . If we choose a metric d on the space of probability measures, then we can define the *measurement infidelity* δ to be the worst case distance between the input and the output measure,

$$\delta := \sup_{\rho \in \mathcal{S}(\mathcal{A})} d(\mathbb{P}_i, \mathbb{P}_f).$$

We consider T to be a good information transfer from A to B if δ is small. Clearly, T is perfect if and only if $\delta = 0$.

As a metric on the space of probability measures, we employ the *trace distance* (a.k.a. variational distance or Kolmogorov distance)

$$d(\mathbb{P}_f, \mathbb{P}_i) := \sup_{S \subset \mathbb{R}} \{|\mathbb{P}_i(S) - \mathbb{P}_f(S)|\}.$$

It is the difference between the probability that the event S occurs in the distribution \mathbb{P}_i and the probability that it occurs in the distribution \mathbb{P}_f , for the worst case Borel set $S \subset \mathbb{R}$.

Let us expand this definition, and write it in a more algebraic fashion. We have $\mathbb{P}_i(S) = \rho(\mathbf{1}_S(A))$ and $\mathbb{P}_f(S) = \rho(T(\mathbf{1}_S(B)))$, so that $\sup_{\rho} d(\mathbb{P}_i, \mathbb{P}_f)$ can be written $\sup_{\rho, S} |\rho(\mathbf{1}_S(A) - T(\mathbf{1}_S(B)))|$. As $\sup_{\rho} |\rho(X)| = \|X\|$, this is just $\sup_S \|\mathbf{1}_S(A) - T(\mathbf{1}_S(B))\|$, and we arrive at the following less intuitive yet cleaner expression.

Definition 4.2 *Let $T : \mathcal{B} \rightarrow \mathcal{A}$ be a CP-map. Its measurement infidelity δ in transferring information from the observable A to the pointer B is defined as $\delta := \sup_S \|\mathbf{1}_S(A) - T(\mathbf{1}_S(B))\|$, where S runs over the Borel subsets of \mathbb{R} .*

The following is a standard example of a perfect information transfer.

Example: von Neumann Qubit Measurement

Let $\mathcal{A} = M_2(\mathbb{C})$ and $A = \sigma_z$. Let $\Omega = \{+1, -1\}$ be the spectrum of σ_z . Denote by $L^\infty(\Omega)$ the (commutative von Neumann) algebra of \mathbb{C} -valued random variables on Ω . (This is just \mathbb{C}^2 with pointwise multiplication.) A state $L^\infty(\Omega) \rightarrow \mathbb{C}$ is precisely the expectation \mathbb{E} with respect to some probability measure \mathbb{P} on Ω . Define the probability measures \mathbb{P}_\pm to assign probability 1 to ± 1 , so that $\mathbb{E}_\pm(f) = f(\pm 1)$. We take $\mathcal{B} = M_2(\mathbb{C}) \otimes L^\infty(\Omega)$ and $B = \mathbf{1} \otimes (\delta_{+1} - \delta_{-1})$.

The von Neumann-measurement $T : M_2(\mathbb{C}) \otimes L^\infty(\Omega) \rightarrow M_2(\mathbb{C})$ is then defined as $T(X \otimes f) := f(+1)P_+XP_+ + f(-1)P_-XP_-$, with $P_+ = |\uparrow\rangle\langle\uparrow|$ and $P_- = |\downarrow\rangle\langle\downarrow|$. The dual $T^* : \mathcal{S}(M_2(\mathbb{C})) \rightarrow \mathcal{S}(M_2(\mathbb{C})) \otimes \mathcal{S}(L^\infty(\Omega))$ reads $T^*\rho = \rho(P_+)|\uparrow\rangle\langle\uparrow| \otimes \mathbb{P}_+ + \rho(P_-)|\downarrow\rangle\langle\downarrow| \otimes \mathbb{P}_-$.

In words: with probability $\rho(P_+)$, the output $+1$ occurs and the qubit is left in state $|\uparrow\rangle$. With probability $\rho(P_-)$, the output -1 occurs, leaving the qubit in state $|\downarrow\rangle$. The von Neumann-measurement T is a perfect information transfer from $\sigma_z \in M_2(\mathbb{C})$ to $\mathbf{1} \otimes (\delta_{+1} - \delta_{-1}) \in M_2(\mathbb{C}) \otimes L^\infty(\Omega)$.

POVMs

Quantum measurements are often (e.g. [Hol82, Hal04]) modelled by Positive Operator Valued Measures or POVMs. From a CP-map $\mathcal{B} \rightarrow \mathcal{A}$ and a pointer $B \in \mathcal{B}$, we may distill the POVM $\mu : \Sigma(\mathbf{Spec}(B)) \rightarrow \mathcal{A}$ by way of $\mu(S) := T(\mathbf{1}_S(B))$. It will only be projection valued if the information transfer is perfect.

For instance, the above von Neumann qubit measurement gives rise to the POVM $\mu : \Omega \rightarrow M_2$ defined by $\mu(+1) = P_+$ and $\mu(-1) = P_-$.

Note that from the POVM, the CP-map can be reconstructed only on B'' . A CP-map can thus be seen as an extension of a POVM that keeps track of the system output as well as the measurement output. Since we will be interested

in disturbance of the system, it is imperative that we consider the full CP-map rather than merely its POVM.

4.2.2 Unbiased Information Transfer

Not unlike a man, it does an information transfer credit to be *unbiased*. This means that the *expectation* of the observable A before the operation is equal to the *expectation* of the pointer B afterwards, for any initial state ρ . In other words, if $T^*\rho(B) = \rho(A)$ for any ρ . This is equivalent to $\rho(T(B) - A) = 0$ for any ρ , and thus to $T(B) = A$.

Definition 4.3 A CP-map $T : \mathcal{B} \rightarrow \mathcal{A}$ is called an unbiased information transfer from $A \in \mathcal{A}$ to $B \in \mathcal{B}$ if $T(B) = A$.

Any perfect information transfer is unbiased, but the converse does not hold. Schematically, we have

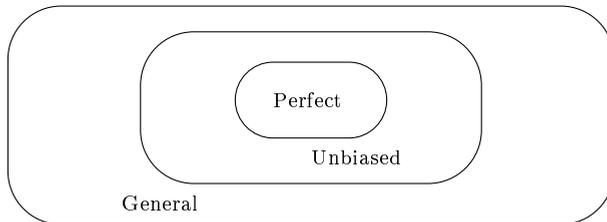


FIG. 3: Kinds of information transfer.

Maximal Added Variance

For unbiased information transfer, there is a natural figure of merit other than the measurement infidelity. The variance of an observable X is defined as $\mathbb{V}\text{ar}(X, \rho) = \mathbb{E}_\rho(X^2) - \mathbb{E}_\rho(X)^2$, or equivalently $\rho(X^2) - \rho(X)^2$. An unbiased information transfer is good if the variance $\mathbb{V}\text{ar}(B, T^*(\rho))$ of the pointer in the final state is as low as possible.

The output variance can be split in two parts. One part $\mathbb{V}\text{ar}(A, \rho)$ is the variance of the input, which is intrinsic to the quantum state ρ . The other part $\mathbb{V}\text{ar}(B, T^*(\rho)) - \mathbb{V}\text{ar}(A, \rho) \geq 0$ is *added* by the measurement procedure.

It would be unfair to hold the information transfer responsible for uncertainty that was already there to begin with, so we will take this added variance to measure its performance. The *maximal added variance*,

$$\Sigma^2 := \sup_{\rho \in \mathcal{S}(\mathcal{A})} \mathbb{V}\text{ar}(B, T^*(\rho)) - \mathbb{V}\text{ar}(A, \rho),$$

will be our figure of (de)merit.

For example, in the case of perfect information transfer from A to B , we have that $\mathbb{V}\text{ar}(B, T^*(\rho)) = \mathbb{V}\text{ar}(A, \rho)$; the maximal added variance is 0. There is uncertainty in the measurement outcome, but all uncertainty ‘comes from’ the quantum state, and none is added by the measurement procedure.

To obtain a cleaner definition, we expand $\text{Var}(B, T^*(\rho)) - \text{Var}(A, \rho)$ into $(\rho(T(B^2)) - \rho(T(B))^2) - (\rho(A^2) - \rho(A)^2)$. Using the fact that $T(B) = A$, we obtain $\Sigma^2 = \sup_\rho \rho(T(B^2) - T(B)^2)$, and thus $\Sigma^2 = \|T(B^2) - T(B)^2\|$.

Definition 4.4 *The maximal added variance of an unbiased information transfer $T : \mathcal{B} \rightarrow \mathcal{A}$ from an observable $A = T(B)$ to a pointer B is defined as*

$$\Sigma^2 := \|T(B^2) - T(B)^2\|. \quad (4.1)$$

4.3 An Operator Valued Form

Inspired by formula (4.1), we define an operator valued form which allows us to consider Σ as the ‘length’ of the pointer; $\Sigma = \sqrt{\|(B, B)\|}$.

Definition 4.5 *Let $T : \mathcal{B} \rightarrow \mathcal{A}$ be a CP-map. We define the operator-valued sesquilinear form $(\cdot, \cdot)_T : \mathcal{B} \times \mathcal{B} \rightarrow \mathcal{A}$ by*

$$(X, Y)_T := T(X^\dagger Y) - T(X)^\dagger T(Y). \quad (4.2)$$

If there is no danger of confusion, we will drop the subscript T .

As we will see, the form $(\cdot, \cdot)_T$ is sesquilinear and positive semidefinite. Its most useful feature however is that it satisfies an operator valued Cauchy-Schwarz inequality.

4.3.1 A Cauchy-Schwarz Inequality

Before proving the operator valued Cauchy-Schwarz inequality, let us first check that $(\cdot, \cdot)_T$ is indeed sesquilinear and positive semidefinite.

Proposition 4.1 *The form $(\cdot, \cdot)_T$ from definition 4.5 is linear in the second argument, it satisfies $(X, Y)^\dagger = (Y, X)$, and $(X, X) \geq 0$ for all $X \in \mathcal{B}$.*

Proof. Linearity in the second argument follows from linearity of T , and the fact that $(X, Y)^\dagger = (Y, X)$ is a consequence of $T(X^\dagger) = T(X)^\dagger$. The following proof of positivity is taken from [Maa04].

As T is completely positive, $T \otimes \text{id} : \mathcal{B} \otimes M_2(\mathbb{C}) \rightarrow \mathcal{A} \otimes M_2(\mathbb{C})$ is positive, so that

$$T \otimes \text{id} \left(\begin{pmatrix} X & \mathbf{1} \\ 0 & 0 \end{pmatrix}^\dagger \begin{pmatrix} X & \mathbf{1} \\ 0 & 0 \end{pmatrix} \right) = \begin{pmatrix} T(X^\dagger X) & T(X^\dagger) \\ T(X) & \mathbf{1} \end{pmatrix}$$

is a positive matrix. Using the fact that $A \geq 0$ implies $B^\dagger A B \geq 0$, we find that

$$\begin{pmatrix} \mathbf{1} & 0 \\ -T(X) & 0 \end{pmatrix}^\dagger \begin{pmatrix} T(X^\dagger X) & T(X^\dagger) \\ T(X) & \mathbf{1} \end{pmatrix} \begin{pmatrix} \mathbf{1} & 0 \\ -T(X) & 0 \end{pmatrix}$$

is also positive, and hence $T(X^\dagger X) - T(X)^\dagger T(X) \geq 0$ as required. \square

We are now ready to prove the operator valued Cauchy-Schwarz inequality. Simple as it may be, it will be crucial to the remainder of this chapter. All our bounds on the performance of information transfer will depend upon it.

Lemma 4.2 (Cauchy-Schwarz) *Let $T : \mathcal{B} \rightarrow \mathcal{A}$ be a CP-map, and define $(X, Y) := T(X^\dagger Y) - T(X)^\dagger T(Y)$. Then for all $X, Y \in \mathcal{B}$:*

$$(X, Y)(Y, X) \leq \|(Y, Y)\|(X, X). \quad (4.3)$$

Proof. If $T : \mathcal{B} \rightarrow \mathcal{A}$ is completely positive¹, then so is $T \otimes \text{id} : \mathcal{B} \otimes M_2(\mathbb{C}) \rightarrow \mathcal{A} \otimes M_2(\mathbb{C})$. According to proposition 4.1,

$$\left(\begin{pmatrix} X & Y \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} X & Y \\ 0 & 0 \end{pmatrix} \right)_{T \otimes \text{id}} = \begin{pmatrix} (X, X)_T & (X, Y)_T \\ (Y, X)_T & (Y, Y)_T \end{pmatrix}$$

is a positive matrix, so that also

$$\begin{pmatrix} \mathbf{1} & 0 \\ -(Y, X)_T & 0 \end{pmatrix}^\dagger \begin{pmatrix} (X, X)_T & (X, Y)_T \\ (Y, X)_T & (Y, Y)_T \end{pmatrix} \begin{pmatrix} \mathbf{1} & 0 \\ -(Y, X)_T & 0 \end{pmatrix} \quad (4.4)$$

is positive, and therefore $(X, X) - 2(X, Y)(Y, X) + (X, Y)(Y, Y)(Y, X) \geq 0$. Rewrite this as $(X, Y)(Y, X) \leq (X, X) + (X, Y)((Y, Y) - \mathbf{1})(Y, X)$, and substitute λY for Y . If λ is such that

$$|\lambda|^2 \|(Y, Y)\| \leq 1, \quad (4.5)$$

then $|\lambda|^2(X, Y)(|\lambda|^2(Y, Y) - \mathbf{1})(Y, X)$ is negative, so that

$$|\lambda|^2(X, Y)(Y, X) \leq (X, X). \quad (4.6)$$

If $(Y, Y) = 0$, then (4.5) holds for arbitrary λ , so that (4.6) forces $(X, Y)(Y, X)$ to be zero. If $(Y, Y) \neq 0$, then taking $|\lambda|^2 = \|(Y, Y)\|^{-1}$ in (4.6) yields the inequality $(X, Y)(Y, X) \leq \|(Y, Y)\|(X, X)$. Either way, we have (4.3), which was to be demonstrated. \square

Remark The special case that $(X, X) = 0$ implies $(X, Y) = 0$ is due to Reinhard Werner, see e.g. [Wer01].

We will have occasion for the following corollary.

Corollary 4.3 *Let $C : \mathcal{A} \rightarrow \mathcal{N}$ and $T : \mathcal{B} \rightarrow \mathcal{A}$ be completely positive maps. Define $(\cdot, \cdot)_{C, T} : \mathcal{B} \times \mathcal{B} \rightarrow \mathcal{N}$ by $(X, Y)_{C, T} = C((X, Y)_T)$. Then*

$$(X, Y)_{C, T}(Y, X)_{C, T} \leq \|(Y, Y)_{C, T}\|(X, X)_{C, T}.$$

Proof. This is a minor adaptation of the proof of lemma 4.2. The map $C \otimes \text{id} : \mathcal{A} \otimes M_2(\mathbb{C}) \rightarrow \mathcal{N} \otimes M_2(\mathbb{C})$ is positive. If we apply it to each of the three factors in equation (4.4) separately, the result will still be positive. This has the effect of replacing every $(\cdot, \cdot)_T$ that occurs in the entries by $(\cdot, \cdot)_{C, T}$. The rest of the proof then goes through uninterrupted. \square

¹We state the result for completely positive maps, but it is clear from the proof that 4-positivity suffices.

4.3.2 Perfect Information Transfer

If an information transfer is perfect, then its maximal added variance $\Sigma^2 = \|(B, B)\|$ will be zero. Indeed, the initial and final probability measures are identical, so that no variance is added. Using our newly acquired Cauchy-Schwarz inequality, we will show that the converse is also true.

Theorem 4.4 *Let $T : \mathcal{B} \rightarrow \mathcal{A}$ be a CP-map, let $B \in \mathcal{B}$ be Hermitean. Then among the statements*

- 1 $(B, B) = 0$.
- 2 *The restriction of T to B'' , the von Neumann algebra generated by B , is a *-homomorphism $B'' \rightarrow T(B)''$.*
- 3 $(f(B), f(B)) = 0$ for all measurable functions f on the spectrum of B .
- 4 T maps the relative commutant $B' = \{X \in \mathcal{A}; [X, B] = 0\}$ into $T(B)'$.

the following relations hold: (1) \Leftrightarrow (2) \Leftrightarrow (3) \Rightarrow (4).

Proof. First, we prove (1) \Rightarrow (2). By Cauchy-Schwarz (lemma 4.2), we have $T(B^n) - T(B)T(B^{n-1}) \leq \|(B, B)\|(B^{n-1}, B^{n-1}) = 0$. Induction on n shows that $T(B^n) = T(B)^n$, so we have $T(f(B)) = f(T(B))$ for all polynomials f by linearity. Thus T is a *-homomorphism from the algebra of polynomials on the spectrum of B to that on $T(B)$. Since T is weakly continuous, this statement extends to the algebras of measurable functions on the spectra of B and $T(B)$, isomorphic to B'' and $T(B)''$ respectively. For (2) \Rightarrow (3), note that $T(f(B)^2) = T(f(B))^2$. For (3) \Rightarrow (1), one simply takes $f(x) = x$. Finally we prove the implication (1) \Rightarrow (4). Suppose that $[A, B] = 0$. Then $[T(B), T(A)] = T([A, B]) - [T(A), T(B)] = (A^\dagger, B) - (B^\dagger, A)$. (B is Hermitean.) By Cauchy-Schwarz, the last term equals zero if (B, B) does. \square

So the maximal added variance Σ^2 is zero if and only if T is a perfect information transfer, i.e. a *-homomorphism $B'' \rightarrow \mathcal{A}$.

We can now give an easy proof of theorem 2.13 in chapter 2.

Theorem 4.5 *Every CP-map $T : \mathcal{B} \rightarrow \mathcal{A}$ with a completely positive inverse is an isomorphism.*

Proof. Let S be the inverse of T . According to proposition 4.1, $(A, A)_S \geq 0$, so that also $T((A, A)_S) \geq 0$. Expanding this as $T(S(A^\dagger A) - S(A^\dagger)S(A)) \geq 0$ and using the fact that T is the inverse of S , we see that $A^\dagger A \geq T(S(A^\dagger)S(A))$. On the other hand, $T(S(A^\dagger)S(A)) - T(S(A^\dagger))T(S(A)) \geq 0$ implies $A^\dagger A \leq T(S(A^\dagger)S(A))$, so that $A^\dagger A = T(S(A^\dagger)S(A))$ for all $A \in \mathcal{A}$. For $A = T(B)$, this reads $T(B)^\dagger T(B) = T(B^\dagger B)$ for all $B \in \mathcal{B}$, or $(B, B)_T = 0$. The Cauchy-Schwarz inequality then implies that also $(B'^\dagger, B)_T = 0$ for all $B', B \in \mathcal{B}$, so that $T(B'B) = T(B')T(B)$. \square

4.4 Joint Measurement

In a jointly unbiased measurement, information on two observables A and \tilde{A} is transferred to two *commuting* pointers B and \tilde{B} . Lemma 4.2 will provide a sharp bound on its performance.

4.4.1 Joint Measurement Inequality

If A and \tilde{A} do not commute, then it is not possible for both information transfers to be perfect. (See [Neu32], [Wer01].) The following inequality states that the degree of imperfection is determined by the amount of noncommutativity.

Theorem 4.6 *Let $T : \mathcal{B} \rightarrow \mathcal{A}$ be a CP-map, let B, \tilde{B} be commuting Hermitean observables in \mathcal{B} , and define $A := T(B)$, $\tilde{A} := T(\tilde{B})$, $\Sigma_B^2 := \|(B, B)\|$ and $\Sigma_{\tilde{B}}^2 := \|(\tilde{B}, \tilde{B})\|$. Then*

$$\Sigma_B \Sigma_{\tilde{B}} \geq \frac{1}{2} \|[A, \tilde{A}]\|. \quad (4.7)$$

Proof. Since $[B, \tilde{B}] = 0$, we have $[\tilde{A}, A] = T([B, \tilde{B}]) - [T(B), T(\tilde{B})] = (B, \tilde{B}) - (\tilde{B}, B)$. By Cauchy-Schwarz, the latter is at most $2\Sigma_B \Sigma_{\tilde{B}}$ in norm. \square

The following application will show that this bound is sharp.

4.4.2 The Beamsplitter as a Joint Measurement

A *beamsplitter* is a device which takes two beams of light as input. A certain fraction of each incident beam is refracted and the rest is reflected, in such a way that the refracted part of the first beam coincides with the reflected part of the second and vice versa.

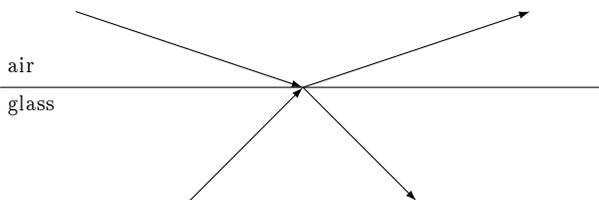


FIG. 4: Beamsplitter.

We will show that the beamsplitter serves as an optimal joint unbiased measurement.

The Maximal Added Variance of the Beamsplitter

In cavity QED, a single mode in the field is described by a Hilbert space \mathcal{H} of a harmonic oscillator, with creation and annihilation operators a^\dagger and a satisfying $[a, a^\dagger] = 1$, as well as $x = \frac{a+a^\dagger}{\sqrt{2}}$ and $p = \frac{a-a^\dagger}{\sqrt{2}i}$. The coherent states $|\alpha\rangle = e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle$ are dense in \mathcal{H} , and satisfy $a|\alpha\rangle = \alpha|\alpha\rangle$.

Quantummechanically, a beamsplitter is described by the unitary operator $U = \exp(\theta(a^\dagger \otimes a - a \otimes a^\dagger))$ on $\mathcal{H} \otimes \mathcal{H}$. In terms of the coherent vectors, we have $U|\alpha\rangle \otimes |\beta\rangle = |\alpha \cos(\theta) + \beta \sin(\theta)\rangle \otimes |-\alpha \sin(\theta) + \beta \cos(\theta)\rangle$. Note that $U^\dagger a \otimes \mathbf{1} U = \cos(\theta)a \otimes \mathbf{1} + \sin(\theta)\mathbf{1} \otimes a$ and that $U^\dagger \mathbf{1} \otimes a U = -\sin(\theta)a \otimes \mathbf{1} + \cos(\theta)\mathbf{1} \otimes a$. (This can be seen by sandwiching both sides between coherent vectors.) Since the map $Y \mapsto U^\dagger Y U$ respects $+$, \cdot and \dagger , we readily calculate

$$\begin{aligned} U^\dagger x \otimes \mathbf{1} U &= \cos(\theta)x \otimes \mathbf{1} + \sin(\theta)\mathbf{1} \otimes x, \\ U^\dagger x^2 \otimes \mathbf{1} U &= \cos^2(\theta)x^2 \otimes \mathbf{1} + 2\sin(\theta)\cos(\theta)x \otimes x + \sin^2(\theta)\mathbf{1} \otimes x^2, \\ U^\dagger \mathbf{1} \otimes p U &= -\sin(\theta)p \otimes \mathbf{1} + \cos(\theta)\mathbf{1} \otimes p, \\ U^\dagger \mathbf{1} \otimes p^2 U &= \sin^2(\theta)p^2 \otimes \mathbf{1} - 2\cos(\theta)\sin(\theta)p \otimes p + \cos^2(\theta)\mathbf{1} \otimes p^2. \end{aligned}$$

Let us identify normal states ρ on $\mathcal{B}(\mathcal{H})$ with normalized trace class operators $R \in \mathcal{T}(\mathcal{H})$ by way of $\rho(X) = \text{tr}(RX)$. We are now interested in the map $R \mapsto UR \otimes |0\rangle\langle 0|U^\dagger$, from $\mathcal{S}(\mathcal{H})$ to $\mathcal{S}(\mathcal{H}) \otimes \mathcal{S}(\mathcal{H})$. In other words, we feed the beamsplitter only one beam of light in a state ρ , the other input being the vacuum. The dual of this is the CP-map $T : \mathcal{B}(\mathcal{H}) \otimes \mathcal{B}(\mathcal{H}) \rightarrow \mathcal{B}(\mathcal{H})$ defined by $T(Y) := id \otimes \Phi(U^\dagger Y U)$, with Φ the vacuum state $\Phi(X) = \langle 0|X|0\rangle$.

Take for instance the pointer $B = \cos^{-1}(\theta)x \otimes \mathbf{1}$. Then $T(B) = x\langle 0|\mathbf{1}|0\rangle + \tan(\theta)\mathbf{1}\langle 0|x|0\rangle = x$. Starting with the pointer $\tilde{B} = -\sin^{-1}(\theta)\mathbf{1} \otimes p$ instead, we end up with $T(\tilde{B}) = p$. Apparently, splitting a beam of light in two parts, measuring $x \otimes \mathbf{1}$ in the first beam and $\mathbf{1} \otimes p$ in the second, and then compensating for the loss of intensity provides a *simultaneous unbiased measurement* of x and p in the original beam. Since $[x, p] = i$, we must² have $\Sigma_B \Sigma_{\tilde{B}} \geq \frac{1}{2}$.

We now calculate Σ_B and $\Sigma_{\tilde{B}}$ explicitly. From $\langle 0|x^2|0\rangle = \frac{1}{2}$, we see that $T(B^2) = x^2 + \frac{1}{2}\tan^2(\theta)\mathbf{1}$. Thus $\Sigma_B^2 = \|(B, B)\| = \frac{1}{2}\tan^2(\theta)$. Similarly $\Sigma_{\tilde{B}}^2 = \frac{1}{2}\tan^{-2}(\theta)$. We see that $\Sigma_B \Sigma_{\tilde{B}} = \frac{1}{2}$, so that the beamsplitter is indeed an optimal jointly unbiased measurement.

4.4.3 Sharpness and Comparison

The beamsplitter is an optimal joint measurement in the sense that it minimizes the product $\Sigma_B^2 \Sigma_{\tilde{B}}^2$ of the maximal added variances. By scaling B , optimal joint measurements can be found for arbitrary values of Σ_B and $\Sigma_{\tilde{B}}$, which shows the bound in theorem 4.6 to be sharp.

The beamsplitter also performs well with other figures of merit. For example, if the quality of joint measurement is judged by the state-dependent cost $R(T) := \text{Var}(B, T^*(\rho)) + \text{Var}(\tilde{B}, T^*(\rho))$, then at least for Gaussian ρ , the optimal measurement is the above beamsplitter with $\theta = \pi/4$. (See [Hol82].)

Choosing a figure of merit is –to a degree– a moral decision. Is it fair to prefer the angle $\theta = \pi/4$ over the others, or should we deem all angles to be of equal value? The former point of view leads to the figure of merit $R(T)$, the latter to $\Sigma_B \Sigma_{\tilde{B}}$.

We will take the second point of view, and propose the product of the two maximal added variances as a figure of merit for joint unbiased measurement.

²We neglect the technical complication of x and p being unbounded operators.

In chapter 5, we will show that homodyne detection of the spontaneous decay of a two-level atom constitutes a joint measurement with $\Sigma_B \Sigma_{\tilde{B}} = 1.056$, slightly above the bound $\Sigma_B \Sigma_{\tilde{B}} \geq 1$ provided by theorem 4.6.

4.5 The Heisenberg Principle

The Heisenberg Principle may be stated as follows:

If all states are left intact, then no quantum information can be extracted from a system.

This alludes to information transfer from an initial system \mathcal{A} to a final system consisting of two parts: the system \mathcal{A} and an ancilla \mathcal{B} , containing the pointer B . We thus have an information transfer $T : \mathcal{A} \otimes \mathcal{B} \rightarrow \mathcal{A}$ from A to $\mathbf{1} \otimes B$.

An initial state $\rho \in \mathcal{S}(\mathcal{A})$ gives rise to a final state $T^* \rho \in \mathcal{S}(\mathcal{A} \otimes \mathcal{B})$. Restricting this final state to the system \mathcal{A} (i.e. taking the partial trace over \mathcal{B}) yields a ‘residual’ state $R^* \rho \in \mathcal{S}(\mathcal{A})$, whereas taking the partial trace over \mathcal{A} yields the final state $Q^* \rho \in \mathcal{S}(\mathcal{B})$ of the ancilla. We define the CP-maps $R : \mathcal{A} \rightarrow \mathcal{A}$ by $R(A) := T(A \otimes \mathbf{1})$ and $Q : \mathcal{B} \rightarrow \mathcal{B}$ by $Q(B) := T(\mathbf{1} \otimes B)$. The map R describes what happens to \mathcal{A} if we forget about the ancilla \mathcal{B} , and Q describes the ancilla, neglecting the original system \mathcal{A} .

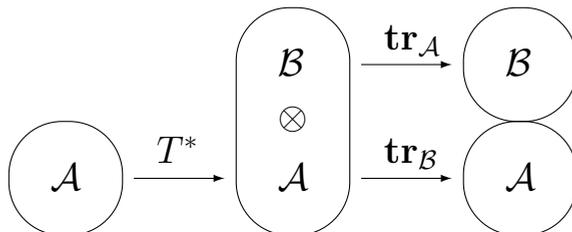


FIG. 5: Constructing R^* and Q^* from T^* .

We wish to find a quantitative version of the Heisenberg principle, i.e. we want to relate the imperfection of the extracted quantum-information to the amount of state disturbance.

4.5.1 A Figure of Merit

We have already established a figure of merit for the quality of information transfer. For general information transfer, we use the measurement infidelity δ of definition 4.2, and for unbiased information transfer the maximal added variance Σ^2 of definition 4.4. We must now find a figure of merit to quantify how well states on \mathcal{A} are preserved by R^* .

For any initial state $\rho \in \mathcal{S}(\mathcal{A})$, we want the final, residual state $R^* \rho$ to be as close to ρ as possible. We therefore define the *maximal disturbance* Δ to be the maximum trace distance between the initial and the final state,

$$\Delta := \sup_{\rho \in \mathcal{S}(\mathcal{A})} d(\rho, R^* \rho). \quad (4.8)$$

The trace distance or Kolmogorov distance between two states ρ and σ is defined as the worst case difference in probability that an event P occurs,

$$d(\rho, \sigma) := \sup_P (\rho(P) - \sigma(P)),$$

where P runs over the projections (or ‘events’) in \mathcal{A} . It is shown in [NC00] that one may equivalently have P run over all positive operators $0 \leq P \leq \mathbf{1}$, rather than just the projections.

The name ‘trace distance’ stems from the fact that if $\rho = \mathbf{tr}(R \cdot)$ and $\sigma = \mathbf{tr}(S \cdot)$ for trace class operators R and S , then $d(\rho, \sigma) = \frac{1}{2} \mathbf{tr}(|R - S|)$.

In order to arrive at a cleaner definition, replace $R^*(\rho)(P)$ in

$$\Delta = \sup\{ |R^*(\rho)(P) - \rho(P)| ; \rho \in \mathcal{S}(\mathcal{A}), 0 \leq P \leq \mathbf{1} \}$$

by $\rho(R(P))$ in order to obtain $\Delta = \sup_{0 \leq P \leq \mathbf{1}} \|R(P) - P\|$. Summarizing, we arrive at the following quantification of state disturbance.

Definition 4.6 *The maximal disturbance Δ of a map $R : \mathcal{A} \rightarrow \mathcal{A}$ is given by*

$$\Delta := \sup\{ \|R(P) - P\| ; P \in \mathcal{A}, 0 \leq P \leq \mathbf{1} \}.$$

One has $\Delta = 0$ if and only if the operation R leaves all states of the original system perfectly in place, or equivalently iff $R = \text{id}$.

4.5.2 HP for the Unbiased Case

We first turn our attention to unbiased information transfer. The imperfection of the transfer is then captured in the maximal added variance Σ^2 .

The Heisenberg principle only holds for quantum-information. Classical observables are contained in the centre $\mathcal{Z} = \{A \in \mathcal{A}; [A, X] = 0 \ \forall X \in \mathcal{A}\}$, whereas quantum observables are not. The degree in which an observable A is ‘quantum’ is given by its distance to the centre $d(A, \mathcal{Z}) = \inf_{Z \in \mathcal{Z}} \|A - Z\|$. In the following, we will take the algebra of observables to be $\mathcal{A} = B(\mathcal{H})$ for some Hilbert space \mathcal{H} . The centre is then simply $\mathbb{C}\mathbf{1}$.

Theorem 4.7 *Let $T : B(\mathcal{H}) \otimes \mathcal{B} \rightarrow B(\mathcal{H})$ be a CP-map, let $B \in \mathcal{B}$ be Hermitian. Define $A := T(\mathbf{1} \otimes B)$, and $\Sigma^2 := \|(\mathbf{1} \otimes B, \mathbf{1} \otimes B)\|$. Furthermore, define $\Delta := \sup_P \{\|R(P) - P\|\}$, with R the restriction of T to $B(\mathcal{H}) \otimes \mathbf{1}$. Then*

$$\Sigma \geq d(A, \mathcal{Z}) \frac{\frac{1}{2} - \Delta}{\sqrt{\Delta(1 - \Delta)}}. \quad (4.9)$$

This bound is sharp in the sense that for all $\Delta \in [0, \frac{1}{2}]$, there exist T and A for which (4.9) attains equality.

We prepare the ground for theorem 4.7 by means of the following lemma.

Lemma 4.8 *Let P be a projection, and let $0 \leq X \leq \mathbf{1}$ be such that $\|P - X\| \leq \Delta$. Then $\|X(\mathbf{1} - X)\| \leq \Delta(1 - \Delta)$.*

Proof. Let x be in $\mathbf{Spec}(X)$, the spectrum of X . As $0 \leq X \leq 1$, we have $0 \leq x \leq 1$. Without loss of generality, assume that there exists a normalized eigenvector ψ such that $X\psi = x\psi$. (If this is not the case, one may complete the proof using approximate eigenvectors.) Decompose ψ over the eigenspaces of P , i.e. write $\psi = \chi_1 + \chi_0$, with χ_1 in the image and χ_0 in the kernel of P . Then $(X - P)\psi = (x - 1)\chi_1 + x\chi_0$. Since $\|\chi_1\|^2 + \|\chi_0\|^2 = 1$, the inequality $\Delta^2 \geq \|(X - P)\psi\|^2 = (x - 1)^2\|\chi_1\|^2 + x^2\|\chi_0\|^2$ implies that either $(1 - x)^2 \leq \Delta^2$ or $x^2 \leq \Delta^2$. We conclude that $\mathbf{Spec}(X) \subseteq [0, \Delta] \cup [(1 - \Delta), 1]$. This implies that $\mathbf{Spec}(X - X^2) \subseteq [0, \Delta(1 - \Delta)]$, as desired. \square

We proceed with the proof of theorem 4.7.

Proof. We may assume $\Delta < \frac{1}{2}$, as inequality (4.9) is trivially satisfied otherwise. Let $x := \sup(\mathbf{Spec}(A))$ and $y := \inf(\mathbf{Spec}(A))$, so that $d(A, \mathcal{Z}) = (x - y)/2$. Without loss of generality, assume that there exist normalized eigenvectors ψ_x and ψ_y satisfying $A\psi_x = x\psi_x$ and $A\psi_y = y\psi_y$. (If this is not the case, one uses approximate eigenvectors.) Define $\psi_- := (\psi_x + \psi_y)/\sqrt{2}$, and $\tilde{B} := |\psi_- \rangle \langle \psi_-|$.

On the one hand, we have $\|R(\tilde{B}) - \tilde{B}\| \leq \Delta$, so that $\|[A, R(\tilde{B}) - \tilde{B}]\| \leq 2\Delta d(A, \mathcal{Z})$. On the other hand, we have $\|[A, \tilde{B}]\| = d(A, \mathcal{Z})$. We then use the triangle inequality to see that

$$\begin{aligned} \| [T(\mathbf{1} \otimes B), T(\tilde{B} \otimes \mathbf{1})] \| &= \| [A, R(\tilde{B})] \| \\ &= \| [A, \tilde{B}] + [A, R(\tilde{B}) - \tilde{B}] \| \\ &\geq d(A, \mathcal{Z})(1 - 2\Delta). \end{aligned}$$

This brings us in a position to apply theorem 4.6 to the commuting pointers $\tilde{B} \otimes \mathbf{1}$ and $\mathbf{1} \otimes B$, yielding

$$2\Sigma \sqrt{\|(\tilde{B} \otimes \mathbf{1}, \tilde{B} \otimes \mathbf{1})_T\|} \geq d(A, \mathcal{Z})(1 - 2\Delta). \quad (4.10)$$

In order to estimate $\|(\tilde{B} \otimes \mathbf{1}, \tilde{B} \otimes \mathbf{1})_T\| = \|(\tilde{B}, \tilde{B})_R\|$, we use the fact that $\tilde{B}^2 = \tilde{B}$ to write $(\tilde{B}, \tilde{B})_R = R(\tilde{B}) - R(\tilde{B})^2$. Lemma 4.8 with $X = R(\tilde{B})$ and $P = \tilde{B}$ then yields $\|R(\tilde{B}) - R(\tilde{B})^2\| \leq \Delta(1 - \Delta)$. Inserting

$$\sqrt{\Delta(1 - \Delta)} \geq \sqrt{\|(\tilde{B} \otimes \mathbf{1}, \tilde{B} \otimes \mathbf{1})_T\|}$$

into inequality (4.10) yields $2\Sigma \sqrt{\Delta(1 - \Delta)} \geq d(A, \mathcal{Z})(1 - 2\Delta)$, which was to be demonstrated. For sharpness of the bound, see section 4.7.5. \square

In the case of no disturbance, $\Delta = 0$, we see that $\Sigma \rightarrow \infty$. No information transfer from \mathcal{A} is allowed if all states on \mathcal{A} are left intact. This is Werner's (see [Wer01]) formulation of the Heisenberg principle.

In the opposite case of perfect information transfer, $\Sigma = 0$, inequality 4.9 shows that Δ must equal at least one half. We shall see in section 4.7 that this corresponds with a so-called 'collapse of the wave function'.

These two extreme situations are connected by theorem 4.7 in a continuous fashion, as indicated in the graph below.

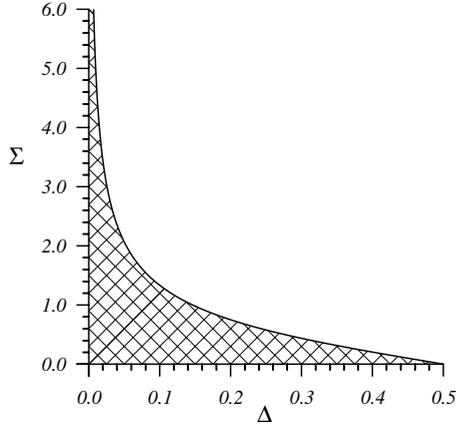


FIG. 6: The combinations (Δ, Σ) below the curve are forbidden, those above are allowed. (With $d(A, \mathcal{Z}) = 1$.)

The upper left corner of the curve illustrates the Heisenberg principle, whereas in the the lower right corner, we can see the collapse of the wave function.

4.5.3 HP for General Information Transfer

We now prove a version of the Heisenberg Principle for general information transfer. We formulate it as a trade-off between the maximal disturbance Δ and the measurement infidelity δ .

Corollary 4.9 *Let $T : B(\mathcal{H}) \otimes \mathcal{B} \rightarrow B(\mathcal{H})$ be a CP-map, let $A \in B(\mathcal{H})$ and $B \in \mathcal{B}$ be Hermitean, $A \notin \mathcal{Z} = \mathbb{C}\mathbf{1}$. Define $\Delta := \sup_P \{\|R(P) - P\|\}$, with R the restriction of T to $B(\mathcal{H}) \otimes \mathbf{1}$. Define $\delta := \sup_S \{\|T(\mathbf{1} \otimes \mathbf{1}_S(B)) - \mathbf{1}_S(A)\|\}$. Then, for δ and Δ in $[0, \frac{1}{2}]$, we have*

$$\left(\frac{1}{2} - \delta\right)^2 + \left(\frac{1}{2} - \Delta\right)^2 \leq \frac{1}{4}. \quad (4.11)$$

This bound is sharp in the sense that for all $\Delta \in [0, \frac{1}{2}]$, there exists a T for which (4.11) attains equality.

Proof. Choose a nontrivial subset S of $\mathbf{Spec}(A)$ and put $P := \mathbf{1} \otimes \mathbf{1}_S(B)$. Since $\|T(P) - \mathbf{1}_S(A)\| \leq \delta$, lemma 4.8 yields $\Sigma^2 = \|T(P) - T(P)^2\| \leq \delta(1 - \delta)$. From the proof of the lemma, we see that $\mathbf{Spec}(T(P))$ is contained in $[0, \delta] \cup [1 - \delta, 1]$. Since $\mathbf{Spec}(T(P))$ contains points in both $[0, \delta]$ and $[1 - \delta, 1]$, we have $d(T(P), \mathcal{Z}) \geq \frac{1}{2} - \delta$. Then apply theorem 4.7 to the pointer P to obtain $\sqrt{\delta(1 - \delta)} \geq (\frac{1}{2} - \delta)(\frac{1}{2} - \Delta) / \sqrt{\Delta(1 - \Delta)}$, or equivalently $(\frac{1}{2} - \delta)^2 + (\frac{1}{2} - \Delta)^2 \leq \frac{1}{4}$. For sharpness, see section 4.7.5. \square

A measurement which does not disturb any state ($\Delta = 0$) cannot yield information ($\delta \geq \frac{1}{2}$). This is the Heisenberg principle. On the other hand, perfect information ($\delta = 0$) implies full disturbance ($\Delta \geq \frac{1}{2}$), corresponding to the collapse of the wave function. Both extremes are connected in a continuous fashion, as depicted below.

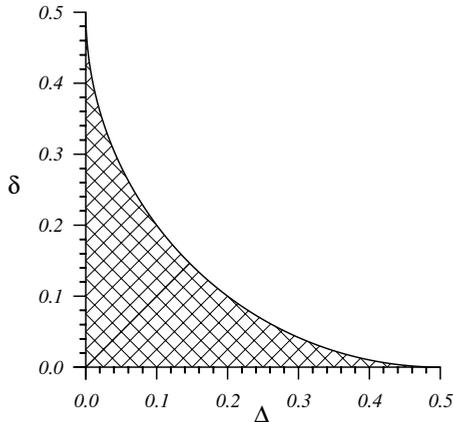


FIG. 7: The combinations (Δ, δ) below the curve are forbidden, those above are allowed.

4.5.4 Application: Resonance Fluorescence

Corollary 4.9 may be used to determine the minimum amount of disturbance if the quality of the measurement is known. Alternatively, if the system is only mildly disturbed, one may find a bound on the attainable measurement quality. Let us concentrate on the latter option.

We investigate the radiation emission of a laser-driven two-level atom. The emitted EM-radiation yields information on the atom. A two-level atom (i.e. a qubit) only has three independent observables: σ_x , σ_y and σ_z . There are various ways to probe the EM field: photon counting, homodyne detection, heterodyne detection, et cetera. For a strong ($\Omega \gg 1$) resonant ($\omega_{\text{laser}} = \omega_{\text{atom}}$) laser, we will use corollary 4.9 to prove that *any* EM-measurement of σ_x , σ_y or σ_z will have a measurement infidelity of at least

$$\delta \geq \frac{1}{2} - \frac{1}{2} \sqrt{1 - e^{-\frac{3}{2}\lambda^2 t}},$$

with λ the coupling constant. For a measurement with two outcomes, δ is the maximal probability of getting the wrong outcome.

Unitary Evolution on the Closed System

The atom is modelled by the Hilbert-space \mathbb{C}^2 (only two energy-levels are deemed relevant). In the field, we discern a *forward* and a *side* channel, each described by a bosonic Fock-space \mathcal{F} . The laser is put on the forward channel,

which is thus initially in the state Φ_Ω , the coherent state with frequency ω and strength Ω . (The field strength is parametrized by the frequency of the induced Rabi-oscillations). The side channel starts in the vacuum state Φ_0 .

If we permit ourselves the luxury of identifying a state with its trace class operator, then the state at time t is given by

$$T_t^*(\rho) = U(t)(\rho \otimes \Phi_\Omega \otimes \Phi_0)U^\dagger(t),$$

with time evolution

$$\frac{d}{dt}U_t = -i(H_S + H_F + \lambda H_I)U_t.$$

$H_S \in B(\mathbb{C}^2)$ is the Hamiltonian of the two-level atom, $H_F \in B(\mathcal{F} \otimes \mathcal{F})$ that of the field and $\lambda H_I \in B(\mathbb{C}^2) \otimes B(\mathcal{F} \otimes \mathcal{F})$ is the interaction-Hamiltonian. Define the interaction-picture time evolution by

$$\hat{T}_t^*(\rho) := U_1(t)^\dagger U_2(t)^\dagger T_t^*(\rho) U_2(t) U_1(t),$$

where $U_1(t) := e^{-iH_S t}$ and $U_2(t) := e^{-iH_F t}$ are the 1-parameter groups of unitaries for the ‘unperturbed’ time evolution.

We now investigate \hat{T}_t instead of T_t . We are looking for a bound on the measurement infidelity $\delta = \sup_S \{\|T(\mathbf{1} \otimes \mathbf{1}_S(B)) - \mathbf{1}_S(A)\|\}$ of T , but if we move to the interaction picture and define $\hat{B} := U_2^\dagger B U_2$, then $\hat{T}(\mathbf{1}_S(\hat{B})) = T(\mathbf{1}_S(B))$, so that $\hat{\delta} = \sup_S \{\|\hat{T}(\mathbf{1} \otimes \mathbf{1}_S(\hat{B})) - \mathbf{1}_S(A)\|\} = \delta$. If we find the interaction-picture disturbance $\hat{\Delta}$, then corollary 4.9 will yield a bound on $\hat{\delta}$, and therefore also on δ .

In the weak coupling limit $\lambda \downarrow 0$, the interaction picture time evolution \hat{T}_t is given by $\hat{T}_t^*(\rho) = \hat{U}(t/\lambda^2)(\rho \otimes \Phi_\Omega \otimes \Phi_0)\hat{U}^\dagger(t/\lambda^2)$, where the evolution of the unitary cocycle $t \mapsto \hat{U}_t$ is described (see [AFL90]) by a Quantum Stochastic Differential Equation or QSDE. Explicitly calculating the maximal added variances Σ^2 by solving the QSDE is in general rather nontrivial, if indeed possible at all.

In chapter 5, we will do this for the case of spontaneous decay, i.e. $\Omega = 0$, with the map \hat{T}_t restricted to the commutative algebra of homodyne measurement results. But here we take an other approach.

Master Equation for the Open System

Fortunately, in contrast to the somewhat complicated time evolution \hat{T}_t of the combined system, the evolution *restricted to the two-level system* is both well-known and uncomplicated. If we use λ^{-2} as a unit of time, then the restricted evolution $\hat{R}_t^*(\rho) := \mathbf{tr}_{\mathcal{F} \otimes \mathcal{F}} \hat{T}_t^*(\rho)$ of the two-level system is known (see [BGM04]) to satisfy the Master equation

$$\frac{d}{dt} \hat{R}_t^*(\rho) = L(\hat{R}_t^*(\rho)), \quad (4.12)$$

with Liouvillian $L(\rho) := \frac{1}{2}i\Omega[e^{-i(\omega-E)t}V + e^{i(\omega-E)t}V^\dagger, \rho] - \frac{1}{2}\{V^\dagger V, \rho\} + V\rho V^\dagger$. In this expression, E is the energy-spacing of the two-level atom and $V^\dagger = \sigma_+$,

$V = \sigma_-$ are its raising and lowering operators. In the case $\omega = E$ of *resonance fluorescence*, we obtain

$$L(\rho) = \frac{1}{2}i\Omega[V + V^\dagger, \rho] - \frac{1}{2}\{V^\dagger V, \rho\} + V\rho V^\dagger.$$

If we parametrize a state by its Bloch-vector $\hat{R}_t^*(\rho) = \frac{1}{2}(\mathbf{1} + x\sigma_x + y\sigma_y + z\sigma_z)$, then equation 4.12 is simply the following differential equation on \mathbb{R}^3 :

$$\frac{d}{dt} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} -\frac{1}{2} & 0 & 0 \\ 0 & -\frac{1}{2} & \Omega \\ 0 & -\Omega & -1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} - \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

This can be solved explicitly. For $\Omega \gg 1$, the solution approaches

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} e^{-\frac{1}{2}t} & 0 & 0 \\ 0 & e^{-\frac{3}{4}t} \cos(\Omega t) & e^{-\frac{3}{4}t} \sin(\Omega t) \\ 0 & -e^{-\frac{3}{4}t} \sin(\Omega t) & e^{-\frac{3}{4}t} \cos(\Omega t) \end{pmatrix} \begin{pmatrix} x_0 \\ y_0 \\ z_0 \end{pmatrix}.$$

If we move to the interaction picture once more to counteract the Rabi oscillations, i.e. with $U_1(t) = e^{\frac{1}{2}\Omega t \sigma_x}$ and $U_2 = \mathbf{1}$, we see that the time evolution is transformed to

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} e^{-\frac{1}{2}t} & 0 & 0 \\ 0 & e^{-\frac{3}{4}t} & 0 \\ 0 & 0 & e^{-\frac{3}{4}t} \end{pmatrix} \begin{pmatrix} x_0 \\ y_0 \\ z_0 \end{pmatrix}.$$

Since the trace distance $D(\rho, \tau)$ is exactly half the Euclidean distance between the Bloch vectors of ρ and τ , (see [NC00]), we see that $\Delta = \frac{1}{2}(1 - e^{-\frac{3}{4}t})$. For any measurement of σ_x , σ_y or σ_z , we therefore have $\delta \geq \frac{1}{2} - \frac{1}{2}\sqrt{1 - e^{-\frac{3}{4}t}}$ by corollary 4.9 (remember that t is in units of λ^{-2}).

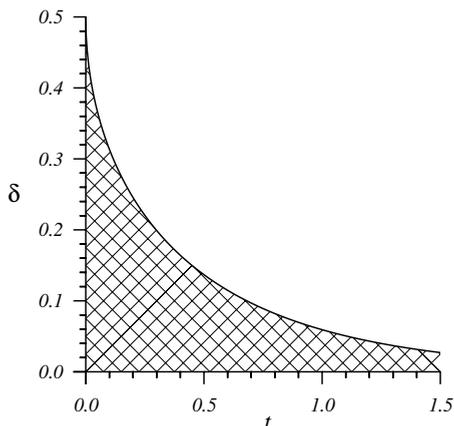


FIG. 8: Lower bound on δ in terms of t (in units of λ^{-2}).

4.6 Classical Coding

In classical coding, a single quantum state is encoded into classical information. Decoding this classical information in order to regain the original quantum state is known to be impossible. However, one can attempt to construct a state which comes as close as possible. We give bounds on the smallest possible trace distance between the original and the decoded state which can be reached.

We start with an approach based on the operator-valued Cauchy-Schwarz inequality. It will yield a bound, but in contrast to the case of joint measurement and the Heisenberg principle, it will not be sharp. We will then prove a sharp bound, starting from Keyl and Werner's no-cloning theorem [KW99].

4.6.1 Classical Coding

We are interested in operations that take quantum states into classical probability measures and vice versa. As we have seen in chapter 2, a classical probability space $(\Omega, \Sigma, \mathbb{P})$ corresponds to a *commutative* quantum probability space (\mathcal{A}, ρ) . The algebra is given by the random variables 'up to measure zero', $\mathcal{A} := L^\infty(\Omega, \Sigma, \mathbb{P})$, and the state is simply the expectation $\rho(f) = \int_\Omega f(\omega) \mathbb{P}(d\omega)$.

An operation that maps quantum states in $\mathcal{S}(B(\mathcal{H}))$ to probability measures in $\mathcal{S}(\mathcal{A})$ must therefore be the dual of a CP-map $C : \mathcal{A} \rightarrow B(\mathcal{H})$, with $\mathcal{A} := L^\infty(\Omega, \Sigma, \mathbb{P})$. According to section 4.2.1, such a map is precisely a POVM $\Sigma \rightarrow B(\mathcal{H})$. Similarly, an operation that maps classical probability distributions into quantum states must be the dual of a CP-map $D : B(\mathcal{H}) \rightarrow \mathcal{A}$.

Classical Coding

By *classical coding*, we mean the following procedure. First, classical information is extracted from a quantum system. This is described by the coding map $C : \mathcal{A} \rightarrow B(\mathcal{H})$, with $\mathcal{A} = L^\infty(\Omega, \Sigma)$. Then, on the basis of this classical information, the original state is reconstructed as well as possible by means of some decoding procedure. This is described by a CP-map $D : B(\mathcal{H}) \rightarrow \mathcal{A}$. In the dual picture, this gives

$$\mathcal{S}(B(\mathcal{H})) \xrightarrow{C^*} \mathcal{S}(\mathcal{A}) \xrightarrow{D^*} \mathcal{S}(B(\mathcal{H})).$$

The coding procedure is flawless iff every state is reconstructed perfectly, i.e. iff $C \circ D : B(\mathcal{H}) \rightarrow B(\mathcal{H})$ is the identity.

It is well known [Wer01, Maa04] that perfect classical coding is impossible. So let us investigate how close we can come to perfection.

A Figure of Merit

To define a figure of merit for classical coding, we proceed in close analogy to the maximal disturbance Δ .

For any initial state $\rho \in \mathcal{S}(B(\mathcal{H}))$, we want the final, reconstructed state $D^*C^*\rho$ to be as close to ρ as possible. We therefore define the *coding imperfection* Δ to be the maximum trace distance between the initial state and its

reconstruction,

$$\Delta := \sup_{\rho \in \mathcal{S}(\mathcal{B}(\mathcal{H}))} d(\rho, D^*C^*\rho).$$

Comparing this with equation (4.8), we see that this is precisely the maximal disturbance for the map $R^* = D^*C^*$. We can follow the reasoning leading up to definition 4.6 word by word to arrive at the following reformulation.

Definition 4.7 *A classical coding procedure consists of a commutative von Neumann algebra \mathcal{A} , and a pair of CP-maps $C : \mathcal{A} \rightarrow \mathcal{B}(\mathcal{H})$ and $D : \mathcal{B}(\mathcal{H}) \rightarrow \mathcal{A}$. Its coding imperfection is defined as*

$$\Delta = \sup_{0 \leq B \leq \mathbf{1}} \|B - CD(B)\|.$$

The coding imperfection quantifies the quality of the coding procedure: the better the procedure, the smaller the imperfection. One would only have $\Delta = 0$ if all states could be encoded perfectly. We now investigate how close to zero Δ can come.

4.6.2 Classical Coding Inequality

We use the Cauchy-Schwarz inequality to obtain a the bound on Δ for any classical coding procedure. (The proof is inspired by the ‘impossibility of classical coding’ version in [Maa04].)

Theorem 4.10 *Any classical coding procedure has coding imperfection*

$$\Delta \geq (3 - \sqrt{5})/4.$$

Proof. Take two orthogonal vectors $\psi, \phi \in \mathcal{H}$, and define X to be the projection on ψ , and Y the projection on $(\psi + \phi)/\sqrt{2}$. We have $\|[X, Y]\| = 1/2$. Since \mathcal{A} is Abelian, we have $D(X)D(Y) = D(Y)D(X)$, and we can write

$$[X, Y] = [X, Y] - CD([X, Y]) + \tag{4.13}$$

$$C(D(XY) - D(X)D(Y)) - \tag{4.14}$$

$$C(D(YX) - D(Y)D(X)).$$

We will bound the r.h.s. in terms of Δ . Remembering that the l.h.s. is at least $1/2$ in norm will then yield a minimum value on Δ .

We start with (4.13). Like any antihermitean operator, $[X, Y]$ can be written as $[X, Y] = i(A_+ - A_-)$, with $0 \leq A_{\pm} \leq \|[X, Y]\|I = \frac{1}{2}I$. Therefore, we have $\|[X, Y] - CD([X, Y])\| \leq \Delta\|A_+\| + \Delta\|A_-\| \leq \Delta$.

We then proceed with (4.14). Consider the positive semidefinite sesquilinear form $(X, Y)_{C,D} := C(D(X^\dagger Y) - D(X)^\dagger D(Y))$, in terms of which the expression (4.14) equals $(X, Y)_{C,D} - (Y, X)_{C,D}$. According to corollary 4.3, we then have

$$\|(X, Y)_{C,D} - (Y, X)_{C,D}\| \leq 2\sqrt{\|(X, X)\| \|(Y, Y)\|}.$$

Now $(X, X)_{C,D} = (X, X)_{CD} - (D(X), D(X))_C \leq (X, X)_{CD}$, which in turn can be written $(X, X)_{CD} = CD(X)(I - CD(X))$ because $X^\dagger X = X$.

As X is a projection, and $\|X - CD(X)\| \leq \Delta$, we can apply lemma 4.8 to obtain $\|CD(X) - CD(X)^2\| \leq \Delta(1 - \Delta)$, or equivalently $\|(X, X)_{CD}\| \leq \Delta(1 - \Delta)$. Of course, we similarly have $\|(Y, Y)_{CD}\| \leq \Delta(1 - \Delta)$. All in all, we can bound (4.14) as $\|(X, Y)_{C,D} - (Y, X)_{C,D}\| \leq 2\Delta(1 - \Delta)$.

We conclude that $1/2 = \|[X, Y]\| \leq \Delta + 2\Delta(1 - \Delta)$, or $(\Delta - 3/4)^2 \leq 5/16$. Thus $\Delta \geq (3 - \sqrt{5})/4$, which was to be proven. \square

The bound $\Delta \geq (3 - \sqrt{5})/4$ is unfortunately not the best possible one. However, the good news is that the best bound follows easily from the no-cloning theorem of Keyl and Werner [KW99].

4.6.3 The No-Cloning method

Each classical coding procedure yields a ‘cloning machine’. All one has to do is to decode the classical information an arbitrary amount M of times, rather than just once. A bound for cloning will therefore induce a bound for classical coding.

Theorem 4.11 *Any classical coding procedure on a Hilbert space \mathcal{H} of dimension d has coding imperfection*

$$\Delta \geq \frac{d-1}{d+1}.$$

Proof. We set $\mathcal{H} = \mathbb{C}^d$, and $\mathcal{A} = L^\infty(\Omega)$ with $\Omega = \{1, 2, \dots, n\}$. The ‘diagonal’ map $\Omega \rightarrow \Omega^M : i \mapsto (i, \dots, i)$ induces the ‘classical cloning’ map $K : \mathcal{A}^{\otimes M} \rightarrow \mathcal{A}$, i.e. $(Kf)(i) := f(i, \dots, i)$. (Note that for e.g. $\Omega = \mathbb{R}$, a cloning map poses difficulties.)

The composition $T := C \circ K \circ D^{\otimes M}$, mapping $B(\mathbb{C}^d)^{\otimes M}$ to $B(\mathbb{C}^d)$, is a so-called $1 \rightarrow M$ cloner. By construction of the operation T , it is clear that $T(I \otimes \dots \otimes B \otimes \dots \otimes I) = CD(B)$. The main result of [KW99] then says that $\sup_{0 \leq B \leq I} \|T(I \otimes \dots \otimes B \otimes \dots \otimes I) - B\| \geq \frac{(M-1)d-1}{M} \frac{d-1}{d+1}$. Since $M \in \mathbb{N}^+$ was arbitrary, this implies $\Delta \geq \frac{d-1}{d+1}$. \square

The number $(3 - \sqrt{5})/4 \approx 0.19$, obtained from the Cauchy-Schwarz inequality, is inferior to the $\frac{1}{3} \approx 0.33$ from the no-cloning theorem in the best possible case $d = 2$, in usefulness as well as in the standard order on \mathbb{R} . Theorem 4.11 is thus the better result.

4.7 Collapse of the Wave function

Recall that the Heisenberg principle states that if all states are left intact, then no information can be extracted from a system. The ‘collapse of the wave function’ may be seen as the flip side of the Heisenberg principle. It states that if information is extracted from a system, then its states undergo a very specific kind of perturbation, called decoherence.

4.7.1 Collapse for Unbiased Information Transfer

We start out by investigating unbiased information transfer. We prove a sharp upper bound on the amount of remaining coherence in terms of the measurement quality.

Theorem 4.12 *Let $T : \mathcal{A} \otimes \mathcal{B} \rightarrow \mathcal{A}$ be a CP-map. Let $B \in \mathcal{B}$ be Hermitean, and consider T as an unbiased information transfer from $A = T(\mathbf{1} \otimes B)$ to $\mathbf{1} \otimes B$ with maximal added variance $\Sigma^2 = \|(\mathbf{1} \otimes B, \mathbf{1} \otimes B)\|$. Suppose that ψ_x and ψ_y are unit eigenvectors of A with different eigenvalues x and y respectively. Define, for $|\alpha|^2 + |\beta|^2 = 1$, the coherent and the incoherent state*

$$\rho_{\text{coh}} := |\alpha\psi_x + \beta\psi_y\rangle\langle\alpha\psi_x + \beta\psi_y|, \quad \rho_{\text{inc}} := |\alpha|^2|\psi_x\rangle\langle\psi_x| + |\beta|^2|\psi_y\rangle\langle\psi_y|.$$

If $R : \mathcal{A} \rightarrow \mathcal{A}$ is the restriction of T to $\mathcal{A} \otimes \mathbf{1}$, then we have the following bound in trace distance:

$$d\left(R^*(\rho_{\text{coh}}), R^*(\rho_{\text{inc}})\right) \leq \frac{\Sigma/|x-y|}{\sqrt{1+4(\Sigma/|x-y|)^2}}. \quad (4.15)$$

This bound is sharp in the sense that for all values of $\Sigma/|x-y|$, there exists a CP-map T for which (4.15) attains equality.

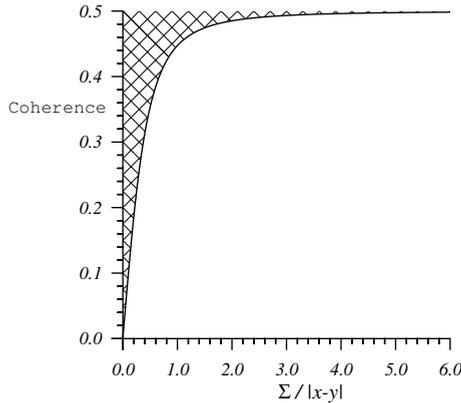


FIG. 9: Bound on the coherence as a function of $\Sigma/|x-y|$. Only points below the curve are allowed.

Consider the ideal case of perfect ($\Sigma = 0$) information transfer. Suppose that the system \mathcal{A} is initially in the coherent state $\rho_{\text{coh}} = |\alpha\psi_x + \beta\psi_y\rangle\langle\alpha\psi_x + \beta\psi_y|$. Then theorem 4.12 says that, after the information transfer to the ancilla \mathcal{B} , the system \mathcal{A} cannot be distinguished from one that started out in the incoherent state $\rho_{\text{inc}} = |\alpha|^2|\psi_x\rangle\langle\psi_x| + |\beta|^2|\psi_y\rangle\langle\psi_y|$ instead. As far as the behaviour of \mathcal{A} after the information transfer is concerned, it is therefore completely harmless to assume that a collapse

$$|\alpha\psi_x + \beta\psi_y\rangle\langle\alpha\psi_x + \beta\psi_y| \mapsto |\alpha|^2|\psi_x\rangle\langle\psi_x| + |\beta|^2|\psi_y\rangle\langle\psi_y|$$

will have occurred *at the start* of the procedure.

Proof. The l.h.s. of equation (4.15) can be written

$$\sup\{\bar{\alpha}\beta\langle\psi_x|R(P)\psi_y\rangle + \text{c.c.}; P \in \mathcal{A}, P^\dagger = P^2 = P\}.$$

Furthermore, since $2|\alpha||\beta| \leq 1$, it suffices to bound the ‘coherence’ $\langle\psi_x|R(P)\psi_y\rangle$ on all projections P . Now $(x-y)\langle\psi_x|R(P)\psi_y\rangle = \langle\psi_x|[A, R(P)]\psi_y\rangle$, and furthermore $[A, R(P)] = (P \otimes \mathbf{1}, \mathbf{1} \otimes B) - (\mathbf{1} \otimes B, P \otimes \mathbf{1})$. Therefore

$$(x-y)\langle\psi_x|R(P)\psi_y\rangle = \langle\psi_x|(P \otimes \mathbf{1}, \mathbf{1} \otimes B)\psi_y\rangle - \langle\psi_x|(\mathbf{1} \otimes B, P \otimes \mathbf{1})\psi_y\rangle, \quad (4.16)$$

and we will bound these last two terms.

By Stinespring’s theorem [Sti55, Tak79], we may assume without loss of generality that T is of the form $T(X) = V^\dagger X V$ for some contraction V . We then obtain $(X, Y) = V^\dagger X^\dagger (\mathbf{1} - V V^\dagger) Y V$. Defining $g(X) := \sqrt{\mathbf{1} - V V^\dagger} X V$, we have $(X, Y) = g(X)^\dagger g(Y)$, and therefore $\|g(\mathbf{1} \otimes B)\| = \Sigma$. This leads to the estimate

$$\begin{aligned} \langle\psi_x|(P \otimes \mathbf{1}, \mathbf{1} \otimes B)\psi_y\rangle &= \langle g(P \otimes \mathbf{1})\psi_x | g(\mathbf{1} \otimes B)\psi_y \rangle \\ &\leq \|g(P \otimes \mathbf{1})\psi_x\| \|g(\mathbf{1} \otimes B)\psi_y\| \\ &\leq \Sigma \sqrt{\langle\psi_x|(P \otimes \mathbf{1}, P \otimes \mathbf{1})\psi_x\rangle}. \end{aligned}$$

We will bound $\langle\psi_x|(T(P^2 \otimes \mathbf{1}) - T(P \otimes \mathbf{1})^2)\psi_x\rangle = \langle\psi_x|(R(P) - R(P)^2)\psi_x\rangle$ in terms of the coherence. For brevity, let us denote $\langle\psi_x|X\psi_x\rangle$ by X_{xx} . Since $\psi_x \perp \psi_y$, we have $(R(P)^2)_{xx} \geq |R(P)_{xx}|^2 + |R(P)_{xy}|^2$, so that

$$(R(P) - R(P)^2)_{xx} \leq R(P)_{xx}(1 - R(P)_{xx}) - |R(P)_{xy}|^2.$$

Since $x(1-x) \leq \frac{1}{4}$ for all $x \in \mathbb{R}$, this is at most $\frac{1}{4} - |R(P)_{xy}|^2$.

All in all, we have obtained

$$(P \otimes \mathbf{1}, \mathbf{1} \otimes B)_{xy} \leq \Sigma \sqrt{\frac{1}{4} - |R(P)_{xy}|^2},$$

and of course the same for $x \leftrightarrow y$. Plugging this into equation 4.16 yields

$$|x-y||R(P)_{xy}| \leq 2\Sigma \sqrt{\frac{1}{4} - |R(P)_{xy}|^2},$$

or equivalently the desired result $|R(P)_{xy}| \leq \frac{\Sigma|x-y|}{\sqrt{1+4(\Sigma/|x-y|)^2}}$. For sharpness, see section 4.7.5. \square

We have seen that if $\Sigma = 0$, then theorem 4.12 is precisely the collapse of the wave function. Now consider the other extreme of a measurement which leaves all states intact, i.e. $R^*\rho = \rho$ for all ρ . Then there exist states for which the l.h.s. of equation (4.15) equals $\frac{1}{2}$, forcing $\Sigma \rightarrow \infty$; no information can be obtained. This is Werner’s formulation of the Heisenberg principle.

Theorem 4.12 thus unifies the Heisenberg principle and the collapse of the wave function. For $\Sigma = 0$ we have a full decoherence, whereas if all states are left intact, we have $\Sigma \rightarrow \infty$. For all intermediate cases, the bound 4.15 on the remaining coherence is an increasing function of $\Sigma/|x-y|$.

This agrees with physical intuition: decoherence between ψ_x and ψ_y is expected to occur in case the information transfer is able to distinguish between the two. This is the case if the added variance is small w.r.t. the differences in mean.

4.7.2 Application: Perfect Qubit Measurement

In section 4.2, we have encountered the von Neumann Qubit measurement. Now consider any perfect measurement T of σ_z with pointer $\mathbf{1} \otimes (\delta_+ - \delta_-)$ which leaves $|\uparrow\rangle\langle\uparrow|$ and $|\downarrow\rangle\langle\downarrow|$ in place, that is $R^*(|\uparrow\rangle\langle\uparrow|) = |\uparrow\rangle\langle\uparrow|$ and $R^*(|\downarrow\rangle\langle\downarrow|) = |\downarrow\rangle\langle\downarrow|$. (Such a measurement is often called *nondestructive*.) Theorem 4.12 then reads $R^*(|\alpha\uparrow + \beta\downarrow\rangle\langle\alpha\uparrow + \beta\downarrow|) = |\alpha|^2|\uparrow\rangle\langle\uparrow| + |\beta|^2|\downarrow\rangle\langle\downarrow|$, as illustrated below.

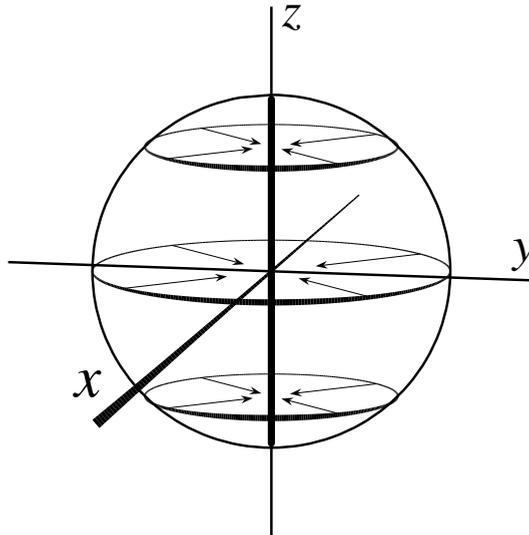


FIG. 10: Collapse on the Bloch-sphere for perfect measurement.

Incidentally, the trace distance between the centre of the Bloch sphere and its surface is $\frac{1}{2}$, so that we read off $\Delta = \sup\{D(R^*(\rho), \rho); \rho \in \mathcal{S}(M_2)\} = \frac{1}{2}$. This was predicted by theorem 4.7.

4.7.3 Collapse for Nondestructive Measurement

We will prove a sharp bound on the remaining coherence in general information transfer. The word ‘general’ in the previous sentence means that the measurement need not be unbiased, but we will focus attention on nondestructive measurements.

A measurement of A is called ‘nondestructive’ (or ‘conserving’ or ‘quantum nondemolition’) if it leaves the eigenstates of A intact, so that repetition of the measurement will yield the same result. For example, the measurement in section 4.2.1 is nondestructive, the one in section 4.5.4 is destructive.

Restricting attention to the nondestructive case is certainly not uncommon in quantum measurement theory (see [Per93]). In all fairness though, we make the nondemolition assumption out of technical convenience, not because we believe it to be of fundamental importance.

Corollary 4.13 *Let $T : B(\mathcal{H}) \otimes \mathcal{B} \rightarrow B(\mathcal{H})$ be a CP-map, and let $A \in B(\mathcal{H})$ and $B \in \mathcal{B}$ be Hermitean. Suppose that A has an orthonormal basis of eigenvectors ψ_i with eigenvalue a_i . If $a_i \neq a_j$, and $|\alpha|^2 + |\beta|^2 = 1$, define the coherent and the incoherent state*

$$\rho_{\text{coh}} := |\alpha\psi_i + \beta\psi_j\rangle\langle\alpha\psi_i + \beta\psi_j|, \quad \rho_{\text{inc}} := |\alpha|^2|\psi_i\rangle\langle\psi_i| + |\beta|^2|\psi_j\rangle\langle\psi_j|.$$

Suppose that T is nondestructive, i.e. $R^(|\psi_i\rangle\langle\psi_i|) = |\psi_i\rangle\langle\psi_i|$ for all ψ_i , with R the restriction of T to $B(\mathcal{H}) \otimes \mathbf{1}$. Then if the measurement infidelity $\delta := \sup_S \{\|T(\mathbf{1} \otimes \mathbf{1}_S(B)) - \mathbf{1}_S(A)\|\}$ satisfies $\delta \in [0, \frac{1}{2}]$, then*

$$d\left(R^*(\rho_{\text{coh}}), \rho_{\text{inc}}\right) \leq \sqrt{\delta(1-\delta)}. \quad (4.17)$$

This bound is sharp in the sense that for all $\delta \in [0, \frac{1}{2}]$, there exists a CP-map T for which (4.17) attains equality.

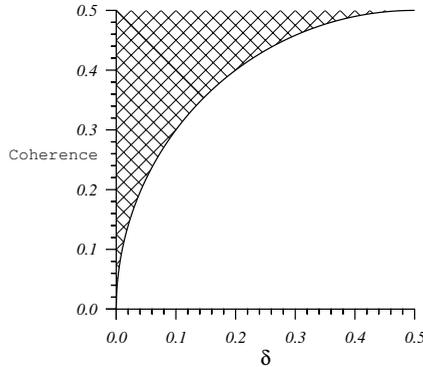


FIG. 11: Bound on the coherence in terms of δ .
Only points below the curve are allowed.

Proof. For sharpness, see section 4.7.5. Choose a set S such that $a_i \in S$ and $a_j \notin S$. T is an unbiased measurement of $T(\mathbf{1} \otimes \mathbf{1}_S(B))$ with pointer $\mathbf{1}_S(B)$ and maximal added variance $\Sigma^2 \leq \delta(1-\delta)$ (cf. the proof of corollary 4.9). We will prove that ψ_i and ψ_j are eigenvectors of $T(\mathbf{1} \otimes \mathbf{1}_S(B))$ with eigenvalues x and y which differ by at least $1-2\delta$.

Define $P_i := |\psi_i\rangle\langle\psi_i|$. Since T is nondestructive, we have $\langle\psi_j | R(P_i)\psi_j\rangle = \langle\psi_j | P_i\psi_j\rangle$ for all j . Apparently, $R(P_i)$ has only one nonzero diagonal element, a 1 at position (i, i) . Since $R(P_i) \geq 0$, this implies $R(P_i) = P_i$.

From $(P_i \otimes \mathbf{1}, P_i \otimes \mathbf{1}) = 0$, we obtain $(P_i \otimes \mathbf{1}, \mathbf{1} \otimes \mathbf{1}_S(B)) = 0$ by Cauchy-Schwarz. Since $P_i = T(P_i \otimes \mathbf{1})$, we have $[T(\mathbf{1} \otimes \mathbf{1}_S(B)), P_i] = (P_i \otimes \mathbf{1}, \mathbf{1} \otimes \mathbf{1}_S(B)) - (\mathbf{1} \otimes \mathbf{1}_S(B), P_i \otimes \mathbf{1}) = 0$. Therefore ψ_i is an eigenvector of $T(\mathbf{1} \otimes \mathbf{1}_S(B))$, with eigenvalue x , say. By a similar reasoning, ψ_j is also an eigenvector, denote its eigenvalue by y .

Since $\|T(\mathbf{1} \otimes \mathbf{1}_S(B)) - \mathbf{1}_S(A)\| \leq \delta$, we have in particular $\|(T(\mathbf{1} \otimes \mathbf{1}_S(B)) - \mathbf{1}_S(A))\psi_i\| = |x-1| \leq \delta$ and $\|(T(\mathbf{1} \otimes \mathbf{1}_S(B)) - \mathbf{1}_S(A))\psi_j\| = |y| \leq \delta$, so that $|x-y| \geq 1-2\delta$. We can now apply Theorem 4.12.

On the l.h.s. of the bound (4.15), we may substitute $R^*(\rho_{\text{inc}}) = \rho_{\text{inc}}$ on account of T being nondestructive. On the r.h.s., we substitute $\Sigma = \sqrt{\delta(1-\delta)}$ and $|x-y| = (1-2\delta)$. Strikingly enough, this yields the bound

$$(\sqrt{\delta(1-\delta)}/(1-2\delta))/\sqrt{1+4(\delta(1-\delta)/(1-2\delta)^2)} = \sqrt{\delta(1-\delta)} \quad \square$$

For perfect measurement ($\delta = 0$), this yields $R^*(|\alpha\psi_i + \beta\psi_j\rangle\langle\alpha\psi_i + \beta\psi_j|) = |\alpha|^2|\psi_i\rangle\langle\psi_i| + |\beta|^2|\psi_j\rangle\langle\psi_j|$; all coherence between ψ_i and ψ_j must vanish. This collapse of the wave function is illustrated in the lower left corner of fig. 11. On the other hand, if all states are left intact so that $R^* = \text{id}$, then we must have $\delta = \frac{1}{2}$; no information can be gained. This is illustrated in the upper right corner of fig.11. Corollary 4.13 is a unified description of the Heisenberg principle and the collapse of the wave function.

4.7.4 Application: Nondestructive Qubit-Measurement

In quantum information theory, a σ_z -measurement is often taken to yield output $+1$ or -1 , according to whether the input was $|\uparrow\rangle$ or $|\downarrow\rangle$. It is nondestructive if it leaves the states $|\uparrow\rangle$ and $|\downarrow\rangle$ intact, yet it is only unbiased if it is perfect. Corollary 4.13 shows that in the nondestructive case, the Bloch-sphere collapses to the cigar-shaped region depicted below:

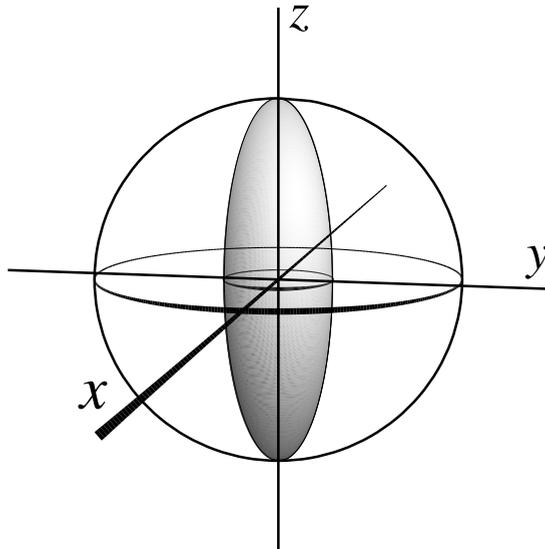


FIG. 12: Collapse on the Bloch-sphere with $\delta = 0.01$.

Current single-qubit readout technology is just now moving into the regime where the bound (4.17) becomes significant. in [LD⁺06], a nondestructive measurement of a SQUID-qubit was described, with experimentally determined measurement infidelity $\delta = 0.13$. The bound then reads $d(R^*\rho_{\text{coh}}, \rho_{\text{inc}}) \leq 0.336$.

4.7.5 Sharpness of the Bounds

We have yet to prove sharpness of the bounds. As it turns out, the following CP-map will do the trick for all the bounds in a single stroke.

Let $\Omega = \{+1, -1\}$, and $\mathcal{B} = L^\infty(\Omega)$. Let

$$V_+ := \begin{pmatrix} \sqrt{1-p} & 0 \\ 0 & \sqrt{p} \end{pmatrix}, \quad V_- := \begin{pmatrix} \sqrt{p} & 0 \\ 0 & \sqrt{1-p} \end{pmatrix},$$

and define $T : M_2(\mathbb{C}) \otimes \mathcal{B} \rightarrow M_2(\mathbb{C})$ by

$$T(X \otimes f) := f(+1)V_+XV_+ + f(-1)V_-XV_-.$$

For $p = 0$, this is the von Neumann-measurement. As a measurement of σ_z with pointer $B := (\delta_+ - \delta_-)/(1 - 2p)$, we have $\delta = p$. This yields bounds on the disturbance and on the coherence. Corollary 4.9 and theorem 4.7 yield $\Delta \geq \frac{1}{2} - \sqrt{p(1-p)}$, corollary 4.13 and theorem 4.12 yield $d(R^*(\rho_{\text{coh}}), \rho_{\text{inc}}) \leq \sqrt{p(1-p)}$. We now explicitly calculate the restriction of T to M_2 , and find

$$R^*(\rho) = \begin{pmatrix} \rho_{11} & 2\sqrt{p(1-p)}\rho_{12} \\ 2\sqrt{p(1-p)}\rho_{21} & \rho_{22} \end{pmatrix}.$$

The maximal remaining coherence occurs for $\alpha = \beta = 1/\sqrt{2}$, for which it equals $\sqrt{p(1-p)}$. The maximal disturbance equals $\Delta = \frac{1}{2} - \sqrt{p(1-p)}$. This shows the bounds (4.9), (4.11), (4.15) and (4.17) to be sharp.

4.8 Conclusion

Our expedition into the heart of darkness has brought us the following spoils.

- I Theorem 4.6 provides a sharp, state independent bound on the performance of unbiased joint measurement of noncommuting observables. In the case of perfect ($\Sigma = 0$) measurement of one observable, it implies that no information whatsoever ($\Sigma' = \infty$) can be gained on the other.
- II Theorem 4.7 (for unbiased information transfer) and corollary 4.9 (for general information transfer) provide sharp, state independent bounds on the performance of a measurement in terms of the maximal disturbance that it causes. In the case of zero disturbance, when all states are left intact, it follows that no information can be obtained. This is the Heisenberg principle.
- III Theorem 4.11, based on a result of Keyl and Werner, provides a sharp, state independent bound on the quality of classical coding. In contrast to I, II and IV, the result based on the Cauchy-Schwarz inequality is not optimal in this case.

IV Theorem 4.12 (for unbiased information transfer) and corollary 4.13 (for nondestructive information transfer) provide a sharp upper bound on the amount of coherence which can survive information transfer. For perfect information transfer, all coherence vanishes. This clearly proves that decoherence on a system is a mathematical consequence of information transfer out of this system. If, on the other hand, all states are left intact, then it follows that no information can be obtained. This is the Heisenberg principle. Theorem 4.12 and corollary 4.13 connect these two extremes in a continuous fashion; they form a unified description of the Heisenberg principle and the collapse of the wave function.

All these restrictions follow directly from the formalism of quantum probability theory as set forth in chapter 2.

Chapter 5

Optimal Pointers for Joint Measurement

Having concluded our investigation into the theoretical restrictions of quantum measurement in general, we turn our attention towards a particular example of a joint measurement procedure. In this chapter, we study a qubit in interaction with the electromagnetic field. By means of homodyne detection, the field-quadrature $A_t + A_t^\dagger$ is observed continuously in time. Due to the interaction, information about the initial state of the qubit is transferred into the field, thus influencing the homodyne measurement results. We construct random variables (pointers) on the probability space of homodyne measurement outcomes having distributions close to the initial distributions of σ_x and σ_z . Using variational calculus, we find the pointers that are optimal. These optimal pointers are very close to hitting the bound imposed by the joint measurement inequality, theorem 4.6 in chapter 4.

5.1 Introduction

The implementation of quantum filtering and control [Blv88] in recent experiments [AS⁺02], [GSM04] has brought new interest to the field of continuous time measurement of quantum systems [Dav69], [Dav76], [Blv88], [Blv92b], [Car93], [WM93], [BGM04]. In particular, homodyne detection has played a considerable role in this development [Car93]. In this chapter, we aim to gain insight into the transfer of information about the initial state of a qubit from this qubit, a two-level atom, to the homodyne photocurrent, which is observed in actual experiments. Our goal is to perform a joint measurement of two noncommuting observables in the initial system. In order to achieve this, we construct random variables (pointers) on the space of possible homodyne measurement results, having distributions close (in a sense to be defined) to the distributions of these observables in the initial state.

The problem of joint measurement of noncommuting observables has been studied by several authors before, see [Neu32], [Dav76], [Hol82] and the references therein. As a measure for the quality of an unbiased measurement, we use the *maximal added variance*, i.e. the difference between the variance of the pointer in the final state and the variance of the observable in the initial state, evaluated in the worst case initial state. (See page 73.) Recall that according to theorem 4.6 in chapter 4, these maximal added variances for two pointers, corresponding to two noncommuting observables of the initial system, satisfy the Heisenberg-like relation (4.7) that bounds how well their joint measurement can be performed.

The present chapter concentrates on the example of a qubit coupled to the quantized electromagnetic field. We study this system in the weak coupling limit [Gou05], i.e. the interaction between qubit and field is governed by a quantum stochastic differential equation in the sense of Hudson and Parthasarathy [HP84]. In the electromagnetic field we perform a homodyne detection experiment. Its integrated photocurrent is the measurement result for measurement of the field-quadrature $A_t + A_t^\dagger$ continuously in time. Using the characteristic functions introduced by Barchielli and Lupieri [BL85], we find the probability density for these measurement results. In this density the x - and z -component of the Bloch vector of the initial state appear, indicating that homodyne detection is in fact a joint measurement of σ_x and σ_z in the initial state.

Our goal is to construct random variables (pointers) on the probability space of homodyne measurement results having distributions as close as possible to those of the observables σ_x and σ_z in the initial state of the qubit. ‘As close as possible’ is taken to mean that the pointer must give an unbiased estimate of the observable, with its maximal added variance as low as possible. Using an argument due to Wiseman [Wis96], we first show that optimal random variables will only depend on the endpoint of a weighted path of the integrated photocurrent. Allowed to restrict our attention to this smaller class of pointers, we are able to use standard variational calculus to obtain the optimal random variables. They do not achieve the bound imposed by the joint measurement theorem 4.6, but will be off by less than 5.6%.

The remainder of the chapter is organized as follows. In section 5.2 we introduce the model of the qubit coupled to the field in the weak coupling limit. In section 5.3, we briefly recall the definition of maximal added variance, as well as the Heisenberg-like relation (4.7) for joint measurement. In section 5.4 we calculate the characteristic function of Barchielli and Lupieri for the homodyne detection experiment. Section 5.5 deals with the variational calculus to find the optimal pointers. In section 5.6 we calculate the densities of the optimal pointers and then capture our main results graphically. In the last section we discuss our results.

5.2 The Model

We consider a two-level atom, i.e. a qubit, in interaction with the quantized electromagnetic field. The qubit is described by \mathbb{C}^2 and the electromagnetic

field by the *symmetric Fock space* \mathcal{F} over the Hilbert space of quadratically integrable functions $L^2(\mathbb{R})$ (space of one-photon wave functions), i.e.

$$\mathcal{F} := \mathbb{C} \oplus \bigoplus_{k=1}^{\infty} L^2(\mathbb{R})^{\otimes_s k}.$$

With the Fock space \mathcal{F} we can describe superpositions of field-states with different numbers of photons. The joint system of qubit and field together is described by the Hilbert space $\mathbb{C}^2 \otimes \mathcal{F}$.

The interaction between the qubit and the electromagnetic field is studied in the weak coupling limit [Gou06], [Gou05], [AFL90]. This means that in the interaction picture the unitary dynamics of the qubit and the field together is given by a quantum stochastic differential equation (QSDE) in the sense of Hudson and Parthasarathy [HP84]

$$dU_t = \left\{ \sigma_- dA_t^\dagger - \sigma_+ dA_t - \frac{1}{2} \sigma_+ \sigma_- dt \right\} U_t, \quad (5.1)$$

with

$$\sigma_- = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, \quad \sigma_+ = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix},$$

and $U_0 = \mathbf{1}$. The operators σ_- and σ_+ are the annihilator and creator on the two-level system. The field annihilation and creation processes are denoted A_t and A_t^\dagger , respectively. Keep in mind that the evolution U_t acts nontrivially on the combined system $\mathbb{C}^2 \otimes \mathcal{F}$, whereas σ_\pm and A_t are understood to designate the single system-operators $\sigma_\pm \otimes \mathbf{1}$ and $\mathbf{1} \otimes A_t$. Throughout this chapter we will remain in the interaction picture. Equation (5.1) should be understood as a shorthand for the integral equation

$$U_t = \mathbf{1} + \int_0^t \sigma_- U_\tau dA_\tau^\dagger - \int_0^t \sigma_+ U_\tau dA_\tau - \frac{1}{2} \int_0^t \sigma_+ \sigma_- U_\tau d\tau,$$

where the integrals on the right-hand side are stochastic integrals in the sense of Hudson and Parthasarathy [HP84], and Picard iterations can be used to show that (5.1) has a unique solution. The value of these integrals does not lie so much in their actual definition (see section 2.3.8), but in the Itô rule satisfied by them, allowing for easy calculations. The following is a slight extension of theorem 2.23.

Theorem 5.1 (Quantum Itô rule [HP84], [Par92]) *Let X_t and Y_t be quantum stochastic integrals of the form*

$$\begin{aligned} dX_t &= C_t dA_t + D_t dA_t^\dagger + E_t dt \\ dY_t &= F_t dA_t + G_t dA_t^\dagger + H_t dt \end{aligned}$$

for some stochastically integrable processes C_t, D_t, E_t, F_t, G_t and H_t (see [HP84], [Par92] for definitions). Suppose that X_t has an adjoint process, and that all the necessary products exist and are stochastically integrable. Then the process $X_t Y_t$ satisfies the relation

$$d(X_t Y_t) = X_t dY_t + (dX_t) Y_t + dX_t dY_t,$$

where $dX_t dY_t$ should be evaluated according to the quantum Itô table:

	dA_t	dA_t^\dagger	dt
dA_t	0	dt	0
dA_t^\dagger	0	0	0
dt	0	0	0

i.e. $dX_t dY_t = C_t G_t dt$.

As a corollary we have that, for any $f \in C^2(\mathbb{R})$, the process $f(X_t)$ satisfies $d(f(X_t)) = f'(X_t)dX_t + \frac{1}{2}f''(X_t)(dX_t)^2$, where $(dX_t)^2$ should be evaluated according to the quantum Itô table.

Let us return to equation (5.1). In order to illustrate how the quantum Itô rule will be used, we calculate the time evolution on the qubit explicitly. We choose an orthonormal basis of \mathbb{C}^2 , and identify the algebra $\mathcal{B}(\mathbb{C}^2)$ of qubit-observables with $M_2(\mathbb{C})$, the algebra of 2×2 -matrices. The algebra of observables in the field is given by $\mathcal{B}(\mathcal{F})$, the bounded operators on \mathcal{F} . If $\text{id} : M_2(\mathbb{C}) \rightarrow M_2(\mathbb{C})$ is the identity map and $\phi : \mathcal{B}(\mathcal{F}) \rightarrow \mathbb{C}$ is the expectation with respect to the vacuum state $\Omega := 1 \oplus 0 \oplus 0 \oplus \dots \in \mathcal{F}$ (i.e. $\phi(Y) := \langle \Omega | Y \Omega \rangle$), then time evolution on the qubit $T_t : M_2(\mathbb{C}) \rightarrow M_2(\mathbb{C})$ is given by $T_t(X) := \text{id} \otimes \phi(U_t^\dagger X \otimes \mathbf{1}U_t)$. On the combined system, the full time evolution $j_t : M_2(\mathbb{C}) \otimes \mathcal{B}(\mathcal{F}) \rightarrow M_2(\mathbb{C}) \otimes \mathcal{B}(\mathcal{F})$ is given by $j_t(W) := U_t^\dagger W U_t$. In a diagram this reads

$$\begin{array}{ccc}
 M_2(\mathbb{C}) & \xrightarrow{T_t} & M_2(\mathbb{C}) \\
 \text{id} \otimes \mathbf{1} \downarrow & & \uparrow \text{id} \otimes \phi \\
 M_2(\mathbb{C}) \otimes \mathcal{B}(\mathcal{F}) & \xrightarrow{j_t} & M_2(\mathbb{C}) \otimes \mathcal{B}(\mathcal{F}).
 \end{array} \tag{5.2}$$

In the Schrödinger picture the arrows would be reversed. A qubit-state ρ would be extended with the vacuum to $\rho \otimes \phi$, time evolution would change it into $\rho \otimes \phi(U_t^\dagger \cdot U_t)$, and in the last step the partial trace over the field would be taken, resulting in the state $\rho \circ T_t$.

Using the Itô rule we can derive a (matrix-valued) differential equation for $T_t(X)$ as follows. We start by using the Itô rule in order to write

$$\begin{aligned}
 dT_t(X) &= \text{id} \otimes \phi(d(U_t^\dagger X \otimes \mathbf{1}U_t)) \\
 &= \text{id} \otimes \phi((dU_t^\dagger)X \otimes \mathbf{1}U_t + U_t^\dagger X \otimes \mathbf{1}(dU_t) + (dU_t^\dagger)X \otimes \mathbf{1}(dU_t)).
 \end{aligned} \tag{5.3}$$

We then use the fact that stochastic integrals with respect to dA_t and dA_t^\dagger vanish with respect to the vacuum expectation, leaving us only with the dt terms. The QSDE for U_t^\dagger easily follows from (5.1);

$$dU_t^\dagger = U_t^\dagger \left\{ \sigma_+ dA_t - \sigma_- dA_t^\dagger - \frac{1}{2} \sigma_+ \sigma_- dt \right\}, \quad U_0^\dagger = \mathbf{1}.$$

The first term in equation 5.3 can then be written

$$\begin{aligned} \text{id} \otimes \phi(dU_t^\dagger X \otimes \mathbf{1}U_t) &= \text{id} \otimes \phi(U_t^\dagger \{ \sigma_+ dA_t - \sigma_- dA_t^\dagger - \frac{1}{2} \sigma_+ \sigma_- dt \} X \otimes \mathbf{1}U_t) \\ &= \text{id} \otimes \phi(U_t^\dagger \{ -\frac{1}{2} \sigma_+ \sigma_- dt \} X \otimes \mathbf{1}U_t) \\ &= T_t(-\frac{1}{2} \sigma_+ \sigma_- X) dt. \end{aligned}$$

Similarly, the second term yields

$$\text{id} \otimes \phi(U_t^\dagger X \otimes \mathbf{1}dU_t) = T_t(-\frac{1}{2} X \sigma_+ \sigma_-) dt.$$

For the third term, we use the quantum Itô table to see that

$$U_t^\dagger \{ \sigma_+ dA_t - \sigma_- dA_t^\dagger - \frac{1}{2} \sigma_+ \sigma_- dt \} X \otimes \mathbf{1} \{ \sigma_- dA_t^\dagger - \sigma_+ dA_t - \frac{1}{2} \sigma_+ \sigma_- dt \} U_t$$

equals $U_t^\dagger (\sigma_+ X \sigma_- \otimes \mathbf{1}) U_t dt$, so that

$$\text{id} \otimes \phi(dU_t^\dagger X \otimes \mathbf{1} dU_t) = T_t(\sigma_+ X \sigma_-) dt.$$

Combining these three terms, equation (5.3) becomes

$$\frac{d}{dt} T_t(X) = T_t(L(X)), \quad (5.4)$$

with L is the Lindblad generator

$$L(X) := -\frac{1}{2}(\sigma_+ \sigma_- X + X \sigma_+ \sigma_-) + \sigma_+ X \sigma_-.$$

The matrix-valued differential equation (5.4) with initial condition $T_0(X) = X$ is solved by $T_t(X) = \exp(tL)(X)$, which is exactly the time evolution of a two-level system spontaneously decaying to the ground state, as it should be. (Compare this to the discussion in section 4.5.4, with zero field strength Ω .) Although the arguments above are completely standard (cf. [HP84]), they do illustrate nicely and briefly some of the techniques used also in following sections.

5.3 Quality of Information Transfer

Now suppose we do a homodyne detection experiment, enabling us to measure the observables $A_t + A_t^\dagger$ in the field continuously in time [Bac90]. If initially the qubit is in state ρ , then at time t the qubit and field together are in a state ρ^t on $M_2(\mathbb{C}) \otimes \mathcal{B}(\mathcal{F})$ given by $\rho^t(W) := \rho \otimes \phi(U_t^\dagger W U_t) = \rho(\text{id} \otimes \phi(U_t^\dagger W U_t))$. We find $\rho^t(\mathbf{1} \otimes (A_t + A_t^\dagger))$ by solving the differential equation

$$d\left(\text{id} \otimes \phi(U_t^\dagger \mathbf{1} \otimes (A_t + A_t^\dagger) U_t)\right) = \text{id} \otimes \phi\left(d(U_t^\dagger \mathbf{1} \otimes (A_t + A_t^\dagger) U_t)\right) \quad (5.5)$$

$$= \text{id} \otimes \phi\left(U_t^\dagger (\sigma_- + \sigma_+) \otimes \mathbf{1} U_t\right) dt \quad (5.6)$$

$$= \exp(tL)(\sigma_- + \sigma_+) dt \quad (5.7)$$

$$= e^{-\frac{t}{2} \sigma_x} dt. \quad (5.8)$$

In order to go from (5.5) to (5.6), we used once again the fact that the ground state kills all terms with dA_t and dA_t^\dagger . Using the quantum Itô table to isolate the terms with dt , we then expand $d(U_t^\dagger \mathbf{1} \otimes (A_t + A_t^\dagger)U_t)$ into six terms, three of first order and three of second order. The dt -terms coming from $dU_t^\dagger \mathbf{1} \otimes (A_t + A_t^\dagger)U_t$, $U_t^\dagger \mathbf{1} \otimes (A_t + A_t^\dagger)dU_t$ and $dU_t^\dagger \mathbf{1} \otimes (A_t + A_t^\dagger)dU_t$ cancel each other out, and obviously $U_t^\dagger \mathbf{1} \otimes (dA_t + dA_t^\dagger)U_t$ lacks terms with dt . We are left with $dU_t^\dagger \mathbf{1} \otimes (dA_t + dA_t^\dagger)U_t$ and $U_t^\dagger \mathbf{1} \otimes (dA_t + dA_t^\dagger)dU_t$, which contribute $U_t^\dagger((\sigma_+ + \sigma_-) \otimes \mathbf{1})U_t$.

This shows that regardless the initial state ρ of the qubit, the expectation of $(A_t + A_t^\dagger)$ in the final state ρ^t will equal the expectation of $(2 - 2e^{-\frac{t}{2}})\sigma_x$ in the initial state ρ .

5.3.1 Defining the Quality of Information Transfer

The process at hand is thus a transfer of information about σ_x to the ‘pointer’ $(A_t + A_t^\dagger)$, which can be read off by means of homodyne detection. We recall definition 4.3 of *unbiased measurement* in the particular case of a qubit measurement by means of the electromagnetic field.

Definition 5.1 (Unbiased Measurement) *Let X be an observable of the qubit, i.e. a self-adjoint element of $M_2(\mathbb{C})$, and let Y be an observable of the field, i.e. a self-adjoint operator in (or affiliated to) $\mathcal{B}(\mathcal{F})$. Then an unbiased measurement M of X with pointer Y is by definition a completely positive map $M : \mathcal{B}(\mathcal{F}) \rightarrow M_2(\mathbb{C})$ such that $M(Y) = X$.*

Needless to say, for each fixed point t in time, the map $M : \mathcal{B}(\mathcal{F}) \rightarrow M_2(\mathbb{C})$ defined by $M(B) := \text{id} \otimes \phi(U_t^\dagger \mathbf{1} \otimes BU_t)$ is a measurement of σ_x with pointer $Y = (2 - 2e^{-\frac{t}{2}})^{-1}(A_t + A_t^\dagger)$. This means that, after the measurement procedure of coupling to the field in the vacuum state and allowing for interaction with the qubit for t time units, the distribution of the measurement results of the pointer Y has inherited the expectation of σ_x , regardless of the initial state ρ . However, we are more ambitious and would like its distribution as a whole to resemble that of σ_x . We recall from section 4.2.2 that the performance of an unbiased measurement is gauged by its *maximal added variance* Σ^2 . We will call its square root Σ the *quality* of measurement.

Definition 5.2 (Quality) *Let $M : \mathcal{B}(\mathcal{F}) \rightarrow M_2(\mathbb{C})$ be an unbiased measurement of X with pointer Y . Then its quality Σ is defined by*

$$\Sigma^2 := \sup \left\{ \left| \text{Var}_{\rho \circ M}(Y) - \text{Var}_\rho(X) \right| \mid \rho \in \mathcal{S}(M_2(\mathbb{C})) \right\},$$

where $\mathcal{S}(M_2(\mathbb{C}))$ denotes the state space of $M_2(\mathbb{C})$ (i.e. all positive normalized linear functionals on $M_2(\mathbb{C})$).

This means that Σ^2 is the variance added to the initial distribution of X by the measurement procedure M for the worst case initial state ρ . As we have seen,

$$\text{Var}_{\rho \circ M}(Y) - \text{Var}_\rho(X) = \rho(M(Y^2) - M(Y)^2),$$

which implies that $\Sigma^2 = \|M(Y^2) - M(Y)^2\|$, where $X \mapsto \|X\|$ denotes the operator norm on $M_2(\mathbb{C})$. In particular this shows that Σ^2 is positive, as one might expect. It follows from theorem 4.4 that Σ equals zero if and only if the measurement procedure M exactly carries over the distribution of X to Y . In short, Σ is a suitable measure for how well M transfers information about X to the pointer Y .

5.3.2 Calculating the Quality of Information Transfer

Let us return to the example at hand, i.e. $M(B) = \text{id} \otimes \phi(U_t^\dagger \mathbf{1} \otimes BU_t)$, with field-observable $Y = (2 - 2e^{-\frac{t}{2}})^{-1}(A_t + A_t^\dagger)$ as a pointer for σ_x . Let us calculate its quality, which amounts to evaluating $M(Y^2) = (2 - 2e^{-\frac{t}{2}})^{-2}M((A_t + A_t^\dagger)^2)$. To this aim, we will first introduce some ideas which will be of use to us in later calculations as well.

Definition 5.3 *Let f and h be real valued functions, h twice differentiable. Let Y_t be given by $dY_t = f(t)(dA_t + dA_t^\dagger)$, $Y_0 = 0$. For $X \in M_2(\mathbb{C})$ we define*

$$F_h(X, t) := \text{id} \otimes \phi(U_t^\dagger X \otimes h(Y_t)U_t).$$

When no confusion can arise we shall shorten $F_h(X, t)$ to $F_h(X)$.

The homodyne detection experiment has given us a measurement result (the integrated photocurrent) which is just the path of measurement results of $A_t + A_t^\dagger$ continuously in time. Given this result, we post-process it by weighting the increments of the path with the function $f(t)$ and letting $h(y)$ act on the result. The following lemma will considerably shorten calculations.

Lemma 5.2

$$\frac{dF_h(X)}{dt} = F_h(L(X)) + f(t)F_{h'}(\sigma_+X + X\sigma_-) + \frac{1}{2}f(t)^2F_{h''}(X)$$

Proof. Just as we did in equation (5.5), we use that the vacuum expectation kills all terms with dA_t and dA_t^* . We then expand $d(U_t^\dagger X \otimes h(Y_t)U_t)$ into three first order and three second order terms. Using theorem 5.1 we see that after the vacuum expectation, the terms $(dU_t^\dagger)X \otimes h(Y_t)U_t$, $U_t^\dagger X \otimes h(Y_t)dU_t$ and $(dU_t^\dagger)X \otimes h(Y_t)dU_t$ make up $F_h(L(X))dt$. From

$$dh(Y_t) = h'(Y_t)f(t)(dA_t + dA_t^\dagger) + \frac{1}{2}h''(Y_t)f(t)^2dt,$$

we find that, after taking vacuum expectations, the terms $d(U_t^\dagger)d(X \otimes h(Y_t))U_t$ and $U_t^\dagger d(X \otimes h(Y_t))dU_t$ make up the second term $f(t)F_{h'}(\sigma_+X + X\sigma_-)dt$, and $U_t^\dagger d(X \otimes h(Y_t))U_t$ provides the last term $\frac{1}{2}f(t)^2F_{h''}(X)dt$. \square

We are now well-equipped to calculate $M((A_t + A_t^\dagger)^2)$. Choose $f(t) = 1$ and $h(x) = x^2$. (The maps $x \mapsto x^n$ will be denoted \mathfrak{X}^n hereafter.) Then we have $M((A_t + A_t^\dagger)^2) = F_{\mathfrak{X}^2}(\mathbf{1})$ and by lemma 5.2

$$\frac{dF_{\mathfrak{X}^2}(\mathbf{1})}{dt} = 2F_{\mathfrak{X}}(\sigma_- + \sigma_+) + F_1(\mathbf{1}) = 2F_{\mathfrak{X}}(\sigma_- + \sigma_+) + \mathbf{1}, \quad F_{\mathfrak{X}^2}(\mathbf{1}, 0) = 0. \quad (5.9)$$

Applying lemma 5.2 to $F_{\mathfrak{X}}(\sigma_- + \sigma_+)$, we obtain

$$\frac{dF_{\mathfrak{X}}(\sigma_- + \sigma_+)}{dt} = -\frac{1}{2}F_{\mathfrak{X}}(\sigma_- + \sigma_+) + 2F_1(\sigma_+\sigma_-), \quad F_{\mathfrak{X}}(\sigma_- + \sigma_+, 0) = 0. \quad (5.10)$$

Finally, $F_1(\sigma_+\sigma_-)$ satisfies

$$\frac{dF_1(\sigma_+\sigma_-)}{dt} = -F_1(\sigma_+\sigma_-), \quad F_1(\sigma_+\sigma_-, 0) = \sigma_+\sigma_-. \quad (5.11)$$

Solving (5.11), (5.10) and (5.9) successively leads first to $F_1(\sigma_+\sigma_-) = e^{-t}\sigma_+\sigma_-$, then to $F_{\mathfrak{X}}(\sigma_- + \sigma_+) = 4(e^{-\frac{t}{2}} - e^{-t})\sigma_+\sigma_-$ and finally to $F_{\mathfrak{X}^2}(\mathbf{1}) = 8(e^{-\frac{t}{2}} - 1)^2\sigma_+\sigma_- + tI$. Consequently, the quality of the measurement M of σ_x with pointer $Y = (2 - 2e^{-\frac{t}{2}})^{-1}(A_t + A_t^\dagger)$ is given by

$$\begin{aligned} \Sigma^2 &= \|M(Y^2) - M(Y)^2\| \\ &= \left\| \frac{8(e^{-\frac{t}{2}} - 1)^2\sigma_+\sigma_- + tI}{(2 - 2e^{-\frac{t}{2}})^2} - \mathbf{1} \right\| \\ &= \left\| 2\sigma_+\sigma_- + \left(\frac{t}{(2 - 2e^{-\frac{t}{2}})^2} - 1 \right) \mathbf{1} \right\| \\ &= \frac{t}{(2 - 2e^{-\frac{t}{2}})^2} + 1. \end{aligned}$$

This expression takes its minimal value 2.228 at $t = 2.513$, leading to a quality $\Sigma = 1.493$.

The calculation above has an interesting side product. The observable $M((A_t + A_t^\dagger)^2)$ depends linearly on σ_z , indicating that in addition to information on σ_x , also information on σ_z in the initial qubit-state ends up in the measurement outcome. Indeed, if we use as a pointer

$$\tilde{Y} := \frac{(A_t + A_t^\dagger)^2 - tI}{4(e^{-\frac{t}{2}} - 1)^2} - \mathbf{1}, \quad (5.12)$$

then we have $M(\tilde{Y}) = \sigma_z$, so that M is also a measurement of σ_z with pointer \tilde{Y} .

Note that the pointers Y and \tilde{Y} commute, i.e. measuring $A_t + A_t^\dagger$ via the homodyne detection scheme is an indirect *joint measurement* of σ_x and σ_z . If we would also like to gain some information about σ_y , we could for example sweep the measured quadrature through $[0, 2\pi)$ in time by measuring $e^{i\omega t}A_t + e^{-i\omega t}A_t^\dagger$ instead. We will however restrict ourselves to continuous time measurement of $A_t + A_t^\dagger$, as additional information on σ_y would deteriorate the quality of σ_x - and/or σ_z -measurement. We formulate the joint measurement theorem 4.7 in the particular case of qubit measurement by means of an electromagnetic field.

Theorem 5.3 (Joint Measurement) *Let $M : \mathcal{B}(\mathcal{F}) \rightarrow M_2(\mathbb{C})$ be an unbiased measurement of self-adjoint observables $X \in M_2(\mathbb{C})$ and $\tilde{X} \in M_2(\mathbb{C})$ with self-adjoint commuting pointers Y and \tilde{Y} in (or affiliated to) $\mathcal{B}(\mathcal{F})$, respectively. Then for their corresponding qualities Σ and $\tilde{\Sigma}$ the following relation holds*

$$2\Sigma\tilde{\Sigma} \geq \|[X, \tilde{X}]\|.$$

Denote by $\tilde{\Sigma}$ the quality of the σ_z measurement with the pointer \tilde{Y} defined in (5.12). Since $[\sigma_x, \sigma_z] = -2i\sigma_y$, the qualities Σ and $\tilde{\Sigma}$ (corresponding to the pointers Y and \tilde{Y} , respectively) satisfy the inequality

$$\Sigma\tilde{\Sigma} \geq 1. \quad (5.13)$$

Using similar techniques as before, that is recursively calculating $F_{\mathfrak{X}^4}(\mathbf{1})$ via lemma 5.2, we find

$$\tilde{\Sigma}^2 = \frac{t^2}{8(e^{-\frac{t}{2}} - 1)^4} + \frac{2t - 4(e^{-\frac{t}{2}} - 1)^2}{(e^{-\frac{t}{2}} - 1)^2}.$$

This expression takes its minimal value 8.836 at $t = 2.513$. This leads to a quality $\tilde{\Sigma} = 2.973$, which means that $\Sigma\tilde{\Sigma} = 4.437$, i.e. we are far removed from hitting the bound 1 in (5.13). However, there is still some room for manoeuvring by post-processing of the homodyne measurement data.

5.4 The Weighted Path

Let us presently return to our homodyne detection experiment. We observe $A_\tau + A_\tau^\dagger$ continuously in time, i.e. the result of our measurement is a path ω of measurement results ω_τ (the photocurrent integrated up to time τ) for $A_\tau + A_\tau^\dagger$. This means that we have a space Ω of all possible measurement paths and that we can identify an operator $A_\tau + A_\tau^\dagger$ with the map from Ω to \mathbb{R} mapping a measurement path $\omega \in \Omega$ to the measurement result ω_τ at time τ . That is, we have simultaneously diagonalized the family of commuting operators $\{A_\tau + A_\tau^\dagger \mid \tau \geq 0\}$ and viewed them as random variables on the spectrum Ω . The spectral projectors of the operators $\{A_\tau + A_\tau^\dagger \mid 0 \leq \tau \leq t\}$ endow Ω with a filtration of σ -algebras $\{\Sigma_t; t \in \mathbb{R}^+\}$. Furthermore, the states ρ^τ , defined by $\rho^\tau(W) := \rho \otimes \phi(U_\tau^\dagger W U_\tau)$ provide a family of consistent measures \mathbb{P}_τ on (Ω, Σ_τ) , turning it into the probability space $(\Omega, \Sigma_t, \mathbb{P})$. (See e.g. [Bou04].)

We aim to find random variables on $(\Omega, \Sigma_t, \mathbb{P})$ having distributions resembling those of σ_x and σ_z in the initial state ρ . In the previous section we used the random variables

$$Y(\omega) = \frac{\omega_\tau}{2 - 2e^{-\frac{\tau}{2}}} \quad \text{and} \quad \tilde{Y}(\omega) = \frac{\omega_\tau^2 - \tau}{4(e^{-\frac{\tau}{2}} - 1)^2} - 1, \quad \tau = 2.513 \quad (5.14)$$

for σ_x and σ_z , respectively. Our next goal is to find the optimal random variables, in the sense of the previously defined quality.

5.4.1 Restricting the Class of Pointers

In our specific example, M is given by $M(B) = \text{id} \otimes \phi(U_\tau^\dagger \mathbf{1} \otimes BU_\tau)$. Note that stochastic integrals with respect to the annihilator A_τ acting on the vacuum vector Ω are zero. Therefore, we can modify U_τ to Z_τ , given by

$$dZ_\tau = \left\{ \sigma_- (dA_\tau^\dagger + dA_\tau) - \frac{1}{2} \sigma_+ \sigma_- d\tau \right\} Z_\tau, \quad Z_0 = \mathbf{1},$$

without affecting M [Blv92a]. Therefore, for all $B \in \mathcal{B}(\mathcal{F})$, we have $M(B) = \text{id} \otimes \phi(U_\tau^\dagger \mathbf{1} \otimes BU_\tau) = \text{id} \otimes \phi(Z_\tau^\dagger \mathbf{1} \otimes BZ_\tau)$. The solution Z_t can readily be found, it is given by

$$Z_t = \begin{pmatrix} e^{-\frac{1}{2}t} & 0 \\ \int_0^t e^{-\frac{1}{2}\tau} (dA_\tau + dA_\tau^\dagger) & 1 \end{pmatrix}.$$

Note that Z_t , as a matrix valued function of the measurement path, is an element of $M_2(\mathbb{C}) \otimes \mathcal{C}_t$, where \mathcal{C}_t is the commutative von Neumann algebra generated by $A_\tau + A_\tau^\dagger$, $0 \leq \tau \leq t$. Moreover we see that Z_t is not a function of all the $(A_\tau + A_\tau^\dagger)$'s separately, it is only a function of the endpoint of the weighted path $Y_t = \int_0^t e^{-\frac{1}{2}\tau} (dA_\tau + dA_\tau^\dagger)$ [Wis96]. Therefore if we define $\mathcal{S}_t \subset \mathcal{C}_t$ to be the commutative von Neumann algebra generated by Y_t , then we even have $Z_t \in M_2(\mathbb{C}) \otimes \mathcal{S}_t$.

Denote by $C \mapsto \mathbb{E}[C | \mathcal{S}_t]$ the unique classical conditional expectation from \mathcal{C}_t onto \mathcal{S}_t that leaves ϕ invariant, i.e. $\phi(\mathbb{E}[C | \mathcal{S}_t]) = \phi(C)$ for all $C \in \mathcal{C}_t$. We can extend $\mathbb{E}[\cdot | \mathcal{S}_t]$ by tensoring it with the identity map on the 2×2 matrices to obtain a map $\text{id} \otimes \mathbb{E}[\cdot | \mathcal{S}_t]$ from $M_2(\mathbb{C}) \otimes \mathcal{C}_t$ onto $M_2(\mathbb{C}) \otimes \mathcal{S}_t$. From the positivity of $\mathbb{E}[\cdot | \mathcal{S}_t]$ as a map between commutative algebras, it follows that $\text{id} \otimes \mathbb{E}[\cdot | \mathcal{S}_t]$ is completely positive. Since $\mathbb{E}[\cdot | \mathcal{S}_t]$ satisfies $\mathbb{E}[CS | \mathcal{S}_t] = \mathbb{E}[C | \mathcal{S}_t]S$ for all $C \in \mathcal{C}_t$ and $S \in \mathcal{S}_t$, we find that $\text{id} \otimes \mathbb{E}[\cdot | \mathcal{S}_t]$ satisfies the *module property*, i.e.

$$\text{id} \otimes \mathbb{E}[\cdot | \mathcal{S}_t](A_1 B A_2) = A_1 \left(\text{id} \otimes \mathbb{E}[\cdot | \mathcal{S}_t](B) \right) A_2,$$

for all $A_1, A_2 \in M_2(\mathbb{C}) \otimes \mathcal{S}_t$ and $B \in M_2(\mathbb{C}) \otimes \mathcal{C}_t$. Moreover, if ρ is a state on $M_2(\mathbb{C})$, then it follows from the invariance of ϕ under $\mathbb{E}[\cdot | \mathcal{S}_t]$ that $\text{id} \otimes \mathbb{E}[\cdot | \mathcal{S}_t]$ leaves $\rho \otimes \phi$ invariant. We conclude that, given ρ on $M_2(\mathbb{C})$, the map $\text{id} \otimes \mathbb{E}[\cdot | \mathcal{S}_t]$ from $M_2(\mathbb{C}) \otimes \mathcal{C}_t$ onto $M_2(\mathbb{C}) \otimes \mathcal{S}_t$ is the unique conditional expectation in the noncommutative sense of [Tak71] that leaves $\rho \otimes \phi$ invariant. We will use the shorthand $\mathbb{E}_{\mathcal{S}_t}$ for $\text{id} \otimes \mathbb{E}[\cdot | \mathcal{S}_t]$ in the following.

Lemma 5.4 *Let $C \in \mathcal{C}_t$ be a pointer with quality Σ_C such that $M(C) = X$. Then $\tilde{C} := \mathbb{E}[C | \mathcal{S}_t]$ is also a pointer with $M(\tilde{C}) = X$, and its quality is at least as good; $\Sigma_{\tilde{C}} \leq \Sigma_C$.*

Proof. Note that for all states ρ on $M_2(\mathbb{C})$ we have

$$\begin{aligned}
\rho(M(\tilde{C})) &= \rho \otimes \phi(Z_t^\dagger \mathbf{1} \otimes \tilde{C} Z_t) \\
&= \rho \otimes \phi\left(Z_t^\dagger \mathbb{E}_{\mathcal{S}_t}(\mathbf{1} \otimes C) Z_t\right) \\
&= \rho \otimes \phi\left(\mathbb{E}_{\mathcal{S}_t}(Z_t^\dagger \mathbf{1} \otimes C Z_t)\right) \\
&= \rho \otimes \phi(Z_t^\dagger \mathbf{1} \otimes C Z_t) \\
&= \rho(M(C)).
\end{aligned}$$

In the third step, we used the module property and the fact that Z_t is an element of $M_2(\mathbb{C}) \otimes \mathcal{S}_t$, and in the fourth step, we used invariance of $\rho \otimes \phi$. Since this holds for all states ρ on $M_2(\mathbb{C})$, we conclude that $M(\tilde{C}) = M(C) = X$.

As for the variance, we note first that the conditional expectation $\mathbb{E}_{\mathcal{S}_t}$ is a completely positive identity preserving map. Therefore, for all self-adjoint $C \in \mathcal{C}_t$, we have

$$\mathbb{E}_{\mathcal{S}_t}(\mathbf{1} \otimes C^2) \geq \left(\mathbb{E}_{\mathcal{S}_t}(\mathbf{1} \otimes C)\right)^2$$

by proposition 4.1.

We can now apply the same strategy as before. For all states ρ on $M_2(\mathbb{C})$ we have

$$\begin{aligned}
\rho(M(C^2)) &= \rho \otimes \phi(Z_t^\dagger \mathbf{1} \otimes C^2 Z_t) \\
&= \rho \otimes \phi\left(\mathbb{E}_{\mathcal{S}_t}(Z_t^\dagger \mathbf{1} \otimes C^2 Z_t)\right) \\
&= \rho \otimes \phi\left(Z_t^\dagger \mathbb{E}_{\mathcal{S}_t}(\mathbf{1} \otimes C^2) Z_t\right) \\
&\geq \rho \otimes \phi\left(Z_t^\dagger \left(\mathbb{E}_{\mathcal{S}_t}(\mathbf{1} \otimes C)\right)^2 Z_t\right) \\
&= \rho(M(\tilde{C}^2)).
\end{aligned}$$

Thus $M(C^2) \geq M(\tilde{C}^2)$, so that $M(C^2) - M(C)^2 \geq M(\tilde{C}^2) - M(\tilde{C})^2$, and in particular $\Sigma_C^2 \geq \Sigma_{\tilde{C}}^2$. \square

This has a very useful consequence: if we are looking for pointers that record, say, σ_x or σ_z in an optimal fashion, then it suffices to examine only pointers in \mathcal{S}_t . Instead of sifting through the collection of all random variables on the measurement outcomes, we are thus allowed to confine the scope of our search to the rather transparent collection of measurable functions of Y_t . In the following, we will look at such pointers $h_t(Y_t)$. We will usually drop the subscript t on h to make the notation lighter.

5.4.2 Distribution of Y_t

At this point we are interested in the probability distribution of the random variable Y_t . Its characteristic function [BL85] is given by

$$E(k) := \mathbb{E}_{\rho^s}[\exp(-ikY_t)] = \rho \otimes \phi\left(U_t^\dagger \mathbf{1} \otimes \exp(-ikY_t) U_t\right) = \rho\left(F_{\exp(-ik\hat{x})}(\mathbf{1})\right),$$

so that we need only calculate $F_{\exp(-ik\mathfrak{X})}(\mathbf{1})$. For notational convenience we will replace the subscript $\exp(-ik\mathfrak{X})$ by k in the following. Using lemma 5.2, we find the following system of matrix valued differential equations:

$$\begin{aligned}\frac{d}{dt}F_k(\mathbf{1}) &= -ik e^{-\frac{t}{2}}F_k(\sigma_x) - \frac{k^2 e^{-t}}{2}F_k(\mathbf{1}), \\ \frac{d}{dt}F_k(\sigma_x) &= -\frac{1}{2}F_k(\sigma_x) - 2ik e^{-\frac{t}{2}}F_k(\sigma_+\sigma_-) - \frac{k^2 e^{-t}}{2}F_k(\sigma_x), \\ \frac{d}{dt}F_k(\sigma_+\sigma_-) &= -F_k(\sigma_+\sigma_-) - \frac{k^2 e^{-t}}{2}F_k(\sigma_+\sigma_-),\end{aligned}$$

with initial values $F_k(\mathbf{1}, 0) = \mathbf{1}$, $F_k(\sigma_x, 0) = \sigma_x$ and $F_k(\sigma_+\sigma_-, 0) = \sigma_+\sigma_-$. (Recall that $\sigma_x = \sigma_+ + \sigma_-$.) Solving this system leads to

$$F_k(\mathbf{1}) = e^{-\frac{k^2(1-e^{-t})}{2}} \left(\mathbf{1} - ik(1-e^{-t})\sigma_x - k^2(1-e^{-t})^2\sigma_+\sigma_- \right).$$

We define the Fourier transform to be $\mathcal{F}(f)(x) := \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f(k)e^{ikx} dk$. Then the probability density of Y_t with respect to the Lebesgue measure is given by $\frac{1}{\sqrt{2\pi}}\mathcal{F}(E)(x) = \frac{1}{\sqrt{2\pi}}\rho(\mathcal{F}(F_k(\mathbf{1}))(x))$. Defining $p(x) := \frac{1}{\sqrt{2\pi}}\mathcal{F}(F_k(\mathbf{1}))(x)$, we can write

$$p(x) = \frac{e^{-\frac{1}{2}\frac{x^2}{1-e^{-t}}}}{\sqrt{2\pi(1-e^{-t})}} \left(\mathbf{1} + x\sigma_x + (x^2 - 1 + e^{-t})\sigma_+\sigma_- \right),$$

i.e. Y_t is distributed according to a Gaussian perturbed by the matrix elements of the initial state $\rho(\sigma_x)$ and $\rho(\sigma_+\sigma_-) = \frac{1}{2}\rho(\sigma_z) + \frac{1}{2}$. No information about ρ on σ_y enters the distribution though. To gain information about σ_y we would have to change our continuous time measurement setup, as we discussed before. If we absorb a constant $(1-e^{-t})^{-\frac{1}{2}}$ in the definition of Y_t , i.e. if we redefine $Y_t := (1-e^{-t})^{-\frac{1}{2}} \int_0^t e^{-\frac{\tau}{2}}(dA_\tau + dA_\tau^\dagger)$, then its density becomes

$$p(y) = \frac{e^{-\frac{y^2}{2}}}{\sqrt{2\pi}} \left(\mathbf{1} + \beta_t y \sigma_x + \beta_t^2 (y^2 - 1) \sigma_+\sigma_- \right), \quad (5.15)$$

where $\beta_t := \sqrt{1-e^{-t}}$.

5.5 Variational Calculus

In lemma 5.4, we have shown that it suffices to consider only random variables of the form $h(Y_t)$ for some measurable h . In equation (5.15), we have captured the probability distribution of Y_t . All that remains now is to calculate the optimal h , which can be done with variational calculus.

5.5.1 Optimal Measurement of σ_x

We seek the function h^* for which the quality Σ of the pointer $h^*(Y_t)$ for σ_x -measurement is optimal. In other words, we need

$$\Sigma^2 := \left\| \int_{-\infty}^{\infty} h^2(y)p(y)dy - \left(\int_{-\infty}^{\infty} h(y)p(y)dy \right)^2 \right\| := \left\| \begin{pmatrix} d_1 & 0 \\ 0 & d_2 \end{pmatrix} \right\| \quad (5.16)$$

to be minimal under the restriction $\int_{-\infty}^{\infty} h(y)p(y)dy = \sigma_x$.

Now Σ^2 is the norm of a diagonal 2×2 -matrix with entries d_1 and d_2 . Both depend smoothly on h , but $\Sigma^2 = \max\{d_1, d_2\}$ does not. There are three possibilities:

- I) $\Sigma^2 = d_1$ in some open neighbourhood of h^* . To find these h^* , we must minimize the smooth functional d_1 and then check whether $d_1 < d_2$.
- II) $\Sigma^2 = d_2$ in some open neighbourhood of h^* . To find these h^* , we must minimize d_2 and check whether $d_2 < d_1$.
- III) $d_1 = d_2$ for h^* . To find these h^* , we must minimize d_1 subject to the condition $d_1 = d_2$.

In principle, we need three different functionals Λ_1 , Λ_2 and Λ_3 for these three distinct cases. However, it turns out that we can make due with the following functional

$$\begin{aligned} \Lambda(h, \kappa, \gamma_1, \gamma_2, \gamma_3) := & \left(\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} h^2(y)e^{-\frac{1}{2}y^2} dy - 1 \right) + \\ & \kappa \left(\frac{\beta_t^2}{\sqrt{2\pi}} \int_{-\infty}^{\infty} h^2(y)(y^2 - 1)e^{-\frac{1}{2}y^2} dy \right) + \\ & \gamma_1 \left(\int_{-\infty}^{\infty} h(y)e^{-\frac{1}{2}y^2} dy \right) + \\ & \gamma_2 \left(\int_{-\infty}^{\infty} h(y)(y^2 - 1)e^{-\frac{1}{2}y^2} dy \right) + \\ & \gamma_3 \left(\frac{\beta_t}{\sqrt{2\pi}} \int_{-\infty}^{\infty} h(y)ye^{-\frac{1}{2}y^2} dy - 1 \right). \end{aligned} \quad (5.17)$$

The constants γ_1, γ_2 and γ_3 are the Lagrange multipliers enforcing the identity $\int_{-\infty}^{\infty} h(y)p(y)dy = \sigma_x$. These are needed in all cases: Λ_1 , Λ_2 and Λ_3 . One can readily check that setting $\kappa = 0$ in Λ yields Λ_1 , setting $\kappa = 1$ yields Λ_2 and considering κ as a free Lagrange multiplier forces $d_1 = d_2$, so that one has $\Lambda_3 = \Lambda$.

All three cases lead to similar optimality conditions. The requirement that the optimal solution be stable under first order variations yields h^* satisfying either

$$h^*(x) = \frac{C_1 x + C_2}{x^2 + \varepsilon} + C_3 \quad (5.18)$$

or

$$h^*(x) = C_4 x^2 + C_5 x + C_6 \quad (5.19)$$

for some real constants $C_1, C_2, C_3, C_4, C_5, C_6$ and ε depending on $\kappa, \gamma_1, \gamma_2, \gamma_3$.

Suppose that h^* takes the form (5.19). The constraint $\int_{-\infty}^{\infty} h^*(y)p(y)dy = \sigma_x$ will then force $C_4 = C_6 = 0$ and $C_5 = \beta_t^{-1}$, so that $h^*(y) = \beta_t^{-1}y$. The random variable we are investigating is simply the observed path, weighted by the function $f(\tau) = \beta_t^{-1}e^{-\tau/2}$, with t the final time of measurement. Since all the integrals we encounter are Gaussian moments, we can readily compute $M(h^{*2}(Y_t)) = \int_{-\infty}^{\infty} h^{*2}(y)p(y)dy$ to be $2\sigma_+\sigma_- + \beta_t^{-2}\mathbf{1}$. Thus

$$\Sigma^2 = \|(2\sigma_+\sigma_- + \beta_t^{-2}\mathbf{1}) - (\sigma_- + \sigma_+)^2\| = 1 + \beta_t^{-2}.$$

For $t \rightarrow \infty$, this amounts to $\Sigma \rightarrow \sqrt{2}$. Already, we have improved on the naive result $\Sigma = 1.493$ obtained previously.

We proceed with the more involved case (5.18), which will provide us with the optimal solution. Before we continue with the constants C_1, C_2, C_3 and ε however, we calculate some integrals for later use.

Definition 5.4 Define the error function $\operatorname{erf}(x)$ and integrals $I(\varepsilon)$ and $J(\varepsilon)$ by

$$\operatorname{erf}(x) := \frac{2}{\sqrt{\pi}} \int_0^x e^{-u^2} du, \quad I(\varepsilon) := \int_{-\infty}^{\infty} \frac{e^{-\frac{x^2}{2}}}{x^2 + \varepsilon} dx, \quad J(\varepsilon) := \int_{-\infty}^{\infty} \frac{e^{-\frac{x^2}{2}}}{(x^2 + \varepsilon)^2} dx.$$

Lemma 5.5

$$J(\varepsilon) = \frac{\sqrt{2\pi} + (1 - \varepsilon)I(\varepsilon)}{2\varepsilon} \quad \text{and} \quad I(\varepsilon) = \pi\sqrt{\frac{e^\varepsilon}{\varepsilon}} \left(1 - \operatorname{erf}\left(\sqrt{\frac{\varepsilon}{2}}\right)\right).$$

Proof. Since the Fourier transform of $e^{-\sqrt{\varepsilon}|k|}$ is equal to $\sqrt{\frac{2\varepsilon}{\pi}} \frac{1}{x^2 + \varepsilon}$, we find

$$\begin{aligned} I(\varepsilon) &= \sqrt{\frac{\pi}{2\varepsilon}} \int_{-\infty}^{\infty} \mathcal{F}(e^{-\sqrt{\varepsilon}|k|}) \mathcal{F}(e^{-\frac{k^2}{2}}) dx \\ &= \sqrt{\frac{\pi}{2\varepsilon}} \int_{-\infty}^{\infty} e^{-\sqrt{\varepsilon}|k|} e^{-\frac{k^2}{2}} dk, \end{aligned}$$

in view of the unitarity of the Fourier transform. Then

$$\begin{aligned} I(\varepsilon) &= \sqrt{\frac{2\pi}{\varepsilon}} \int_0^{\infty} e^{-\sqrt{\varepsilon}k} e^{-\frac{k^2}{2}} dk \\ &= \sqrt{\frac{2\pi}{\varepsilon}} e^{\frac{1}{2}\varepsilon} \int_{\sqrt{\varepsilon}}^{\infty} e^{-\frac{u^2}{2}} du \\ &= \pi\sqrt{\frac{e^\varepsilon}{\varepsilon}} \left(1 - \operatorname{erf}\left(\sqrt{\frac{\varepsilon}{2}}\right)\right). \end{aligned}$$

The expression for J then follows from

$$\begin{aligned} 0 &= \frac{xe^{-\frac{x^2}{2}}}{x^2 + \varepsilon} \Big|_{-\infty}^{\infty} = \int_{-\infty}^{\infty} \frac{d}{dx} \left(\frac{xe^{-\frac{x^2}{2}}}{x^2 + \varepsilon} \right) dx = \int_{-\infty}^{\infty} \left(\frac{1-x^2}{x^2 + \varepsilon} - \frac{2x^2}{(x^2 + \varepsilon)^2} \right) e^{-\frac{x^2}{2}} dx \\ &= \int_{-\infty}^{\infty} \left(-1 + \frac{\varepsilon - 1}{x^2 + \varepsilon} + \frac{2\varepsilon}{(x^2 + \varepsilon)^2} \right) e^{-\frac{x^2}{2}} dx = -\sqrt{2\pi} + (\varepsilon - 1)I(\varepsilon) + 2\varepsilon J(\varepsilon). \end{aligned}$$

□

The condition $\int_{-\infty}^{\infty} h^*(y)p(y)dy = \sigma_- + \sigma_+ = \sigma_x$ implies

$$C_1 = \frac{\sqrt{2\pi}}{\beta_t(\sqrt{2\pi} - \varepsilon I(\varepsilon))}, \quad C_2 = C_3 = 0$$

which fixes C_1 as a function of ε . The next step is to express d_1 and d_2 in terms of ε :

$$\begin{aligned} d_2 &= \frac{C_1^2}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \frac{y^2}{(y^2 + \varepsilon)^2} e^{-\frac{y^2}{2}} dy - 1 \\ &= \frac{C_1^2}{\sqrt{2\pi}} \left(I(\varepsilon) - \varepsilon J(\varepsilon) \right) - 1, \\ d_1 &= \frac{C_1^2 \beta_t^2}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \frac{y^2(y^2 - 1)}{(y^2 + \varepsilon)^2} e^{-\frac{y^2}{2}} dy + d_2 \\ &= \frac{C_1^2 \beta_t^2}{\sqrt{2\pi}} \left(\sqrt{2\pi} - (1 + 2\varepsilon)I(\varepsilon) + \varepsilon(1 + \varepsilon)J(\varepsilon) \right) + d_2. \end{aligned}$$

First, we use lemma 5.5 to express the above in terms of elementary functions and the error function. Then, using Maple, we find that $\varepsilon \mapsto \max\{d_1, d_2\}$ has a unique minimum at $\varepsilon = 0.605$, for which $d_1 = d_2 = 0.470$. This leads to a C_1 that equals 2.359, and to a quality of

$$\Sigma = \sqrt{\max\{d_1, d_2\}} = 0.685.$$

5.5.2 Optimal Measurement of σ_z

For optimal σ_z -measurement, we can run the same program. We search for the function \tilde{h} that optimizes the quality $\tilde{\Sigma}$, under the restriction that $\tilde{h}(Y_t)$ be a pointer for σ_z -measurement. That is, we search for a function \tilde{h} minimizing the functional of equation (5.16), but now under the restriction $\int_{-\infty}^{\infty} h(y)p(y)dy = \sigma_z$. Again there are three cases of interest, $d_1 = d_2$, $d_1 > d_2$ and $d_2 > d_1$, which we can treat simultaneously by introducing, analogous to equation (5.17), the functional

$$\begin{aligned} \tilde{\Lambda}(h, \kappa, \gamma_1, \gamma_2, \gamma_3) &:= \left(\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} h^2(y) e^{-\frac{1}{2}y^2} dy - 1 \right) \\ &+ \kappa \left(\frac{\beta_t^2}{\sqrt{2\pi}} \int_{-\infty}^{\infty} h^2(y)(y^2 - 1) e^{-\frac{1}{2}y^2} dy \right) \\ &+ \gamma_1 \left(\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} h(y) e^{-\frac{1}{2}y^2} dy + 1 \right) \\ &+ \gamma_2 \left(\frac{\beta_t^2}{\sqrt{2\pi}} \int_{-\infty}^{\infty} h(y)(y^2 - 1) e^{-\frac{1}{2}y^2} dy - 2 \right) \\ &+ \gamma_3 \left(\int_{-\infty}^{\infty} h(y) y e^{-\frac{1}{2}y^2} dy \right). \end{aligned}$$

Indeed, γ_1, γ_2 and γ_3 are the Lagrange multipliers enforcing the restriction $\int_{-\infty}^{\infty} h(y)p(y)dy = \sigma_z$. Again, the functional Σ^2 of equation (5.16) depends

nondifferentiably on h when $d_1 = d_2$. We then have to search for the optimum among the points of nondifferentiability, in which case κ is the Lagrange multiplier confining us to these points. If $d_1 > d_2$ then $\kappa = 1$ and if $d_2 > d_1$ then $\kappa = 0$. Summarizing, wherever Λ takes its minimal value, optimality implies $\frac{\delta \tilde{\Lambda}}{\delta \tilde{h}}(\tilde{h}, \kappa, \gamma_1, \gamma_2, \gamma_3) = 0$ for some $\kappa, \gamma_1, \gamma_2$ and γ_3 . Performing the functional derivative yields either

$$\tilde{h}(x) = \frac{D_1 x + D_2}{x^2 + \delta} + D_3 \quad (5.20)$$

or

$$\tilde{h}(x) = D_4 x^2 + D_5 x + D_6 \quad (5.21)$$

for some (time-dependent) constants $D_1, D_2, D_3, D_4, D_5, D_6$ and δ depending on $\kappa, \gamma_1, \gamma_2$ and γ_3 .

Again, we begin with the least demanding case (5.21), resulting from $\kappa = 0$. The condition $\int_{-\infty}^{\infty} \tilde{h}(y)p(y)dy = \sigma_z$ implies $D_5 = 0$, $D_4 = \beta_t^{-2}$ and $D_6 = -1 - \beta_t^2$. For $t \rightarrow \infty$, this leads to

$$\begin{aligned} \Sigma^2 &= \|M(\tilde{h}^2(Y_t)) - M(\tilde{h}(Y_t))^2\| \\ &= \|(4\sigma_+ \sigma_- + 3\mathbf{1}) - \mathbf{1}\| \\ &= 6, \end{aligned}$$

so that $\Sigma \rightarrow \sqrt{6}$.

This improves the result $\tilde{\Sigma} = 2.973$ obtained previously, but once again the ultimate bound will be reached in the more arduous case (5.20). There, the condition $\int_{-\infty}^{\infty} \tilde{h}(y)p(y)dy = \sigma_z$ implies

$$\begin{aligned} D_3 &= -\frac{\sqrt{2\pi} + I(\delta)D_2}{\sqrt{2\pi}}, \\ D_2 &= \frac{2\sqrt{2\pi}}{\beta_t^2(\sqrt{2\pi} - (1 + \delta)I(\delta))}, \\ 0 &= D_1(\sqrt{2\pi} - \delta I(\delta)). \end{aligned}$$

This leads to expressions for d_1 and d_2 as a function of δ . Using lemma 5.5 and Maple once more, we find that the function $\delta \mapsto \max\{d_1, d_2\}$ has a unique minimum at $\delta = 2.701$, for which $d_1 = d_2 = 2.373$. This leads to a quality of

$$\tilde{\Sigma} = \sqrt{\max\{d_1, d_2\}} = 1.540,$$

attained for $D_1 = 0$, $D_2 = -21.649$ and $D_3 = 5.391$. For the joint measurement this leads to

$$\Sigma \tilde{\Sigma} = 1.056.$$

Although we did not achieve the bound of 1 provided by theorem 5.3, we have come as close as the measurement setup allows. We conclude that, using the setup investigated in this chapter, no simultaneous measurement of σ_x and σ_z will be able to approach the quantum bound by more than 5.6 %. Furthermore, we have identified the unique pointers for this optimal measurement in equations (5.18) and (5.20).

5.6 Distribution of Pointer Variables

We have designed pointers $h^*(Y_t)$ and $\tilde{h}(Y_t)$ in such a way that their distributions in the final state best resemble the distributions of σ_x and σ_z in the initial state. We will now calculate and plot these final densities.

5.6.1 Calculation of h^* - and \tilde{h} -Densities

Let ρ be the initial state of the qubit and let it be parameterized by its Bloch vector (P_x, P_y, P_z) . By equation (5.15), the density $q(y)$ of Y_t is given by

$$q(y) = \rho(p(y)) = \frac{e^{-\frac{1}{2}y^2}}{\sqrt{2\pi}} \left(1 + \beta_t y P_x + \beta_t^2 (y^2 - 1) \frac{P_z + 1}{2} \right). \quad (5.22)$$

We are interested in the the distributions $r(x)$ and $s(x)$ of $h^*(Y_t)$ and $\tilde{h}(Y_t)$ respectively. Let us start with h^* . From equation (5.18), we first calculate the points y where $h^*(y) = x$ for some fixed value of x .

$$y_{\pm} = \frac{C_1 \pm \sqrt{C_1^2 - 4x^2\varepsilon}}{2x}$$

By the Frobenius-Peron equation (see e.g. [Ott93]), $r(x)$ is given by

$$r(x) = \sum_{+,-} \frac{q(y_{\pm})}{|h^{*'}(y_{\pm})|},$$

which leads immediately to

$$r(x) = \sum_{+,-} \frac{(y_{\pm}^2 + \varepsilon)^2 \left(1 + \beta_t y_{\pm} P_x + \beta_t^2 (y_{\pm}^2 - 1) \frac{P_z + 1}{2} \right) e^{-\frac{1}{2}y_{\pm}^2}}{C_1 |y_{\pm}^2 - \varepsilon| \sqrt{2\pi}}, \quad (5.23)$$

where it is understood that $r(x) \neq 0$ only for $x \in [-\frac{C_1}{2\sqrt{\varepsilon}}, \frac{C_1}{2\sqrt{\varepsilon}}]$. We run a similar analysis for $s(x)$. The points y in which $\tilde{h}(y) = x$ are given by

$$y_{\pm} = \pm \sqrt{\frac{(x - D_3)\delta - D_2}{D_3 - x}}.$$

This leads to

$$s(x) = \sum_{+,-} \frac{(y_{\pm}^2 + \delta)^2 \left(1 + \beta_t y_{\pm} P_x + \beta_t^2 (y_{\pm}^2 - 1) \frac{P_z + 1}{2} \right) e^{-\frac{1}{2}y_{\pm}^2}}{2|D_2 y_{\pm}| \sqrt{2\pi}}, \quad (5.24)$$

with $s(x) \neq 0$ only for $x \in [D_3 + \frac{D_2}{\delta}, D_3]$. We proceed with a graphical illustration of the results obtained so far.

5.6.2 Plots of σ_x -Measurement

According to formula 5.22, the distribution of the endpoint of the weighted path depends on the input qubit-state. For instance, the negative σ_x -eigenstate, the tracial state and the positive σ_x -eigenstate lead to the distributions below:

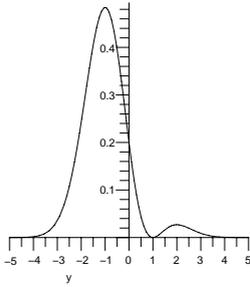


FIG.13: Probability density of the endpoint of the weighted path for the input state $|\leftarrow\rangle$.

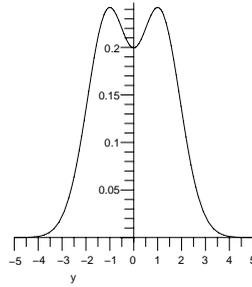


FIG.14: Probability density of the endpoint of the weighted path for the tracial input state.

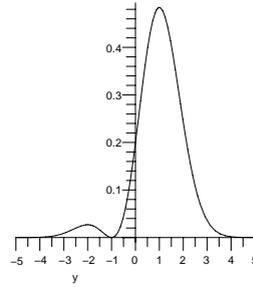


FIG.15: Probability density of the endpoint of the weighted path for the input state $|\rightarrow\rangle$.

In order to estimate σ_x , we use the pointer of σ_x given by

$$h^*(x) = \frac{C_1 x}{x^2 + \varepsilon}, \tag{5.25}$$

with $\varepsilon = 0.605$ and $C_1 = 2.359$. It is illustrated to the right. In formula (5.23), we have calculated the probability distribution of this pointer under the distribution of the endpoint of the weighted path. The distribution of the weighted path is illustrated above, that of the pointer below.

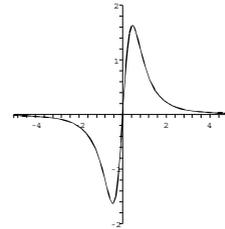


FIG.16: Pointer for σ_x

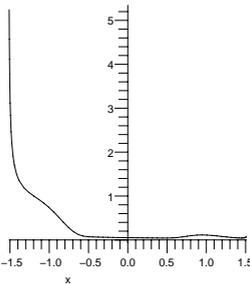


FIG.17: Probability density of the σ_x -pointer for the input state $|\leftarrow\rangle$.

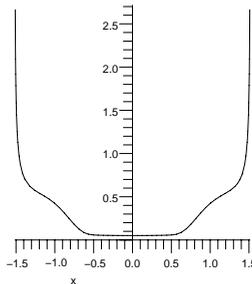


FIG.18: Probability density of the σ_x -pointer for the tracial input state.

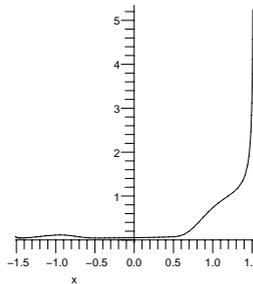


FIG.19: Probability density of the σ_x -pointer for the input state $|\rightarrow\rangle$.

5.6.3 Plots of σ_z -Measurement

We repeat this for the σ_z -pointer. By formula 5.22, the positive σ_z -eigenstate, the tracial state and the negative σ_z -eigenstate lead to the distributions of the endpoint of the weighted path that are shown below:

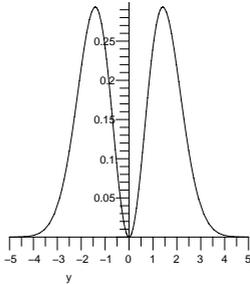


FIG. 20: Probability density of the endpoint of the weighted path for the input state $|\uparrow\rangle$.

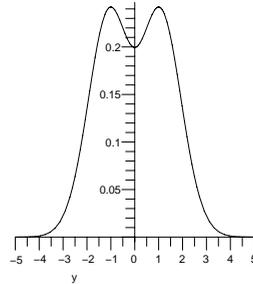


FIG. 21: Probability density of the endpoint of the weighted path for the tracial input state.

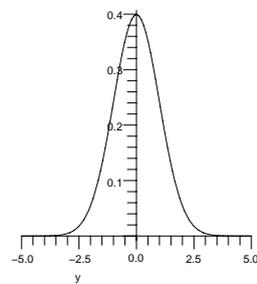


FIG. 22: Probability density of the endpoint of the weighted path for the input state $|\downarrow\rangle$.

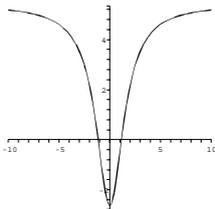


FIG. 23: Pointer for σ_z .

In order to estimate σ_z , we use the pointer of σ_z illustrated here to the left. It is given by

$$\tilde{h}(x) = \frac{D_2}{x^2 + \delta} + D_3, \quad (5.26)$$

with $\delta = 2.701$, $D_2 = -21.649$ and $D_3 = 5.391$. From formula (5.24), we read off the probability distributions of this pointer under the distributions of the endpoint of the weighted path. The probability distributions of the weighted path are illustrated above, those of the pointer below.

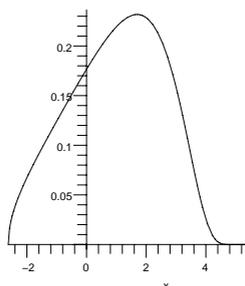


FIG. 24: Probability density of the σ_z -pointer for the input state $|\uparrow\rangle$.

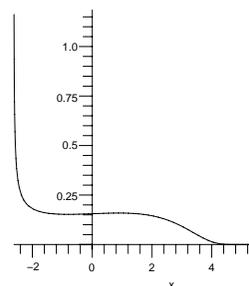


FIG. 25: Probability density of the σ_z -pointer for the tracial input state.

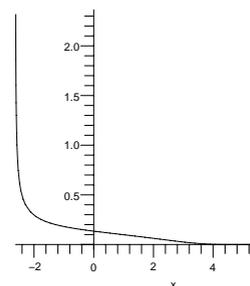


FIG. 26: Probability density of the σ_z -pointer for the input state $|\downarrow\rangle$.

5.7 Discussion

In this chapter, we have investigated homodyne detection of spontaneous decay of a two-level atom into the electromagnetic field. We have seen how the photocurrent, besides carrying information on σ_x (which is immediate from the innovations term in the filtering equation), also carries information on σ_z . Homodyne detection can thus be viewed as a joint measurement of the non-commuting observables σ_x and σ_z in the initial state of the qubit, and we have identified the optimal pointers for this procedure in equations (5.25) and (5.26).

One particular feature of the pointers we constructed might seem counter-intuitive at first: they yield values outside $[-1, 1]$ with nonzero probability. This is a direct result of our requirement that the measurement be unbiased. Suppose, for example, that the input state is $|\uparrow\rangle$, so that σ_z has value 1. Since the photocurrent carries information on σ_x as well, its information on σ_z is certainly flawed, and will yield estimates $\sigma_z < 1$ at least some of the time. Unbiasedness then implies that also estimates $\sigma_z > 1$ must occur.

On the other hand, an unbiased measurement will yield *on average* the ‘true’ value of σ_z for any possible input state. (Not just for the 3 possibilities sketched on page 115.) In repeated experiments, optimality of our pointers guarantees fast convergence to these averages.

Theorem 5.3 provides a theoretical bound for the quality of joint measurement of σ_x and σ_z . No conceivable measurement procedure can ever achieve $\Sigma\tilde{\Sigma} < 1$. It is now clear that this bound cannot be met by way of homodyne detection: a small part of the information extracted from the atom is simply lost in this particular procedure. Constructing the optimal pointers on the photocurrent does yield $\Sigma\tilde{\Sigma} = 1.056$ however, a figure much closer to the bound than the 4.437 provided by the naïve choice of (5.14).

Chapter 6

Optimal Estimation of Qubit States

In this chapter, we propose an adaptive, two step strategy for the estimation of mixed qubit states. We show that the strategy is optimal in a local minimax sense for the trace norm distance as well as other locally quadratic figures of merit. Local minimax optimality means that given n identical qubits, there exists no estimator which can perform better than the proposed estimator on a neighbourhood of size $n^{-1/2}$ of an arbitrary state. In particular, it is asymptotically Bayesian optimal for a large class of prior distributions.

We present a physical implementation of the optimal estimation strategy based on continuous time measurements in a field that couples with the qubits, not unlike the setup in the previous chapter.

The crucial ingredient of the result is the concept of local asymptotic normality (or LAN) for qubits. This means that, for large n , the statistical model described by n identically prepared qubits is locally equivalent to a model with only a classical Gaussian distribution and a Gaussian state of a quantum harmonic oscillator.

The term ‘local’ refers to a shrinking neighbourhood around a fixed state ρ_0 . An essential result is that the neighbourhood radius can be chosen arbitrarily close to $n^{-1/4}$. This allows us to use a two step procedure by which we first localize the state within a smaller neighbourhood of radius $n^{-1/2+\epsilon}$, and then use LAN to perform optimal estimation.

6.1 Introduction

State estimation is a central topic in quantum statistical inference [Hol82, Hst76, BG⁺03, Hay05]. 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*

$$R_\rho(M, \hat{\rho}) = \mathbb{E} (d(\hat{\rho}(X), \rho)^2), \quad (6.1)$$

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}_ρ^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)

$$R_\pi(M, \hat{\rho}) = \int \pi(d\rho) R_\rho(M, \hat{\rho}), \quad (6.2)$$

or by considering a maximum risk (pointwise or minimax setup)

$$R_{\max}(M, \hat{\rho}) = \sup_\rho R_\rho(M, \hat{\rho}). \quad (6.3)$$

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 (cf. chapter 4) 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 [Vaa98] of estimating a parameter θ from independent identically distributed random variables X_1, \dots, X_n with distribution \mathbb{P}_θ , and some of the methods developed in this paper 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 6.2.

In this paper we give an asymptotically optimal measurement strategy for qubit states that is based on the technique of *local asymptotic normality* introduced in [GK06, GJ07]. The technique is a quantum generalization of Le Cam’s classical statistical result [LC86], and builds on previous work of Hayashi and Matsumoto [Hay04, HM06]. 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 [BNG00], and was subsequently used in [GM00, Hay02, HM05]. The aim there is similar: one wants to first localize the state, and then perform a suitably tailored measurement which performs optimally around a given state. A different adaptive technique was proposed independently by Nagaoka [Nag05] and further developed in [Fuj06].

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 6.1) 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 neighbourhood (up to a unitary rotation), as illustrated in the figure below using the Bloch sphere representation of qubit states.

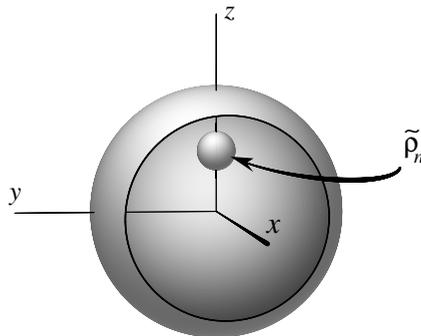


FIG. 27: After the first measurement stage, the state ρ lies in a small ball centered at $\tilde{\rho}_n$.

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 than $\epsilon > 0$ implying that the two stage procedure as a whole is asymptotically optimal for any state, as depicted in the figure below.

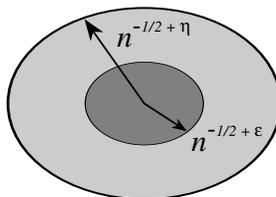


FIG. 28: The smaller domain is the localization region of the first step. The second stage estimator is optimal for all states in the bigger domain.

The optimality of the second stage relies heavily on the principle of *local asymptotic normality* or LAN, see [Vaa98], 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 was the case in [GK06].

Let ρ_0 be a fixed state. We parameterize the neighbouring 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 in the limit $n \rightarrow \infty$ to a Gaussian state of the form $N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}$, in a sense explained in theorem 6.2. 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} [Hol82, YL73].

The idea of approximating a sequence of statistical experiments by a Gaussian one goes back to Wald [Wal43], and was subsequently developed by Le Cam [LC86] 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 the Japanese paper [Hay03], as well as [Hay04] and were subsequently developed in [HM06]. In theorem 6.2 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 in [GJ07] 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 6.2.

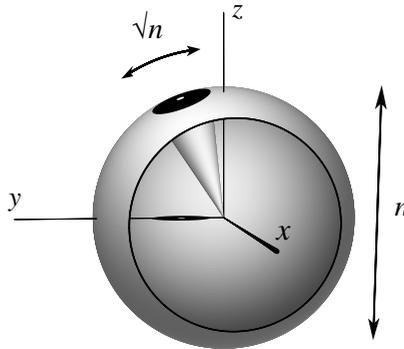


FIG. 29: Total spin representation of the state of $n \gg 1$ spins. The quantum fluctuations of the x and y spin directions coincide with those of a coherent state of a harmonic oscillator.

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 [HH74]. Indeed, it has been known since Dyson [Dys56] 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 the figure above, 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 [GSM04].

In theorem 6.3 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 paper is theorem 6.4 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 [Vaa98]. Also, its asymptotic risk is calculated explicitly.

The paper is structured as follows: in section 6.2, we show that the first stage of the measurement sufficiently localizes the state. In section 6.3, we prove that LAN holds with radius of validity $n^{-1/2+\eta}$, and we bound its rate of convergence. Sections 6.4 and 6.5 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 6.6, 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.

6.2 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 [Hst76] and Holevo [Hol82] and the more recent review paper [BG⁺03]. The collection [Hay05] 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 self-adjoint 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 $\mathbf{tr}(\rho a)$. A measurement M with outcomes in a measurable space (\mathcal{X}, Σ) is completely determined by a σ -additive

collection of positive self-adjoint operators $M(A)$ on \mathcal{H} , where A is an event in Σ . This collection is called a positive operator valued measure (cf. def. 2.6). The distribution of the results X when the system is in state ρ is given by $P_\rho(A) = \text{tr}(\rho M(A))$.

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 than 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 [VR89]. This allows the estimation with arbitrary precision [DLP95, LPA95, LM⁺96] of the whole density matrix of a monochromatic beam of light by repeatedly measuring a sufficiently large number of identically prepared beams [SB⁺93, SB⁺96, ZVB04]. In [AGG05, BGA07] it is 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 parameterizing 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 [Mat02]. However for multi-dimensional families of mixed states this is no longer the case and joint measurements perform strictly better than separate ones [GM00].

In the Bayesian setup, one optimizes $R_\pi(M_n, \hat{\rho}_n)$ for some prior distribution π . We refer to [Jon94, MP95, LPT98, FKF00, HR⁺02, BMM02, EN04, BM⁺05] for the pure state case, and to [CEM99, VL⁺99, MFF00, KW01, BB⁺04, ZS05], and [BB⁺06] 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 in [BB⁺06, HM06] 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 [Hay02, HM05, GM00, BNG00, FN95, Mat02, BG⁺03, HM06] 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 (6.3)), 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 [Gil08], 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.

6.2.1 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 $P_\rho(dx) := \mathbf{tr}(\rho M(dx))$ 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 (see p. 80)

$$\|\rho_1 - \rho_2\|_1 := \mathbf{tr}(|\rho_1 - \rho_2|),$$

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

$$\limsup_{n \rightarrow \infty} \sup_{\|\rho - \rho_0\|_1 \leq n^{-1/2+\epsilon}} nR_\rho(M_n, \hat{\rho}_n) \leq \limsup_{n \rightarrow \infty} \sup_{\|\rho - \rho_0\|_1 \leq n^{-1/2+\epsilon}} nR_\rho(N_n, \check{\rho}_n),$$

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 than $\hat{\rho}_n$ 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 paper 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 neighbourhood 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 neighbourhood. The second step ensures optimality and requires more sophisticated techniques inspired by the theory of local asymptotic normality for qubit states [GK06]. 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 neighbourhood of a known state. This may serve also as an a posteriori justification of the definition of local minimax optimality.

Lemma 6.1 *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*

$$\tilde{\rho}_n = \frac{1}{2}(\mathbf{1} + \tilde{\mathbf{r}}\sigma), \quad \text{if } |\tilde{\mathbf{r}}| \leq 1,$$

where $\tilde{\mathbf{r}} = (\tilde{r}_x, \tilde{r}_y, \tilde{r}_z)$ is the vector average of the measured components. If $|\tilde{\mathbf{r}}| > 1$ 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

$$\mathbb{P}(\|\tilde{\rho}_n - \rho\|_1^2 > 3n^{2\epsilon-1}) \leq 6 \exp(-\frac{1}{2}\tilde{n}n^{2\epsilon-1}), \quad \forall \rho.$$

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

$$\mathbb{E}(\|\tilde{\rho}_n - \rho\|_1^2 \chi_E) \leq 24 \exp(-\frac{1}{2}n^{2\epsilon-\kappa}) = \mathcal{O}(1).$$

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 [VW96] then states that for all real numbers $c > 0$, 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

$$\mathbb{P}\left(\sum_1^3 |\tilde{r}_i - r_i|^2 > 3n^{2\epsilon-1}\right) \leq 6 \exp(-\frac{1}{2}\tilde{n}n^{2\epsilon-1}) \quad \forall \rho,$$

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 $\|\rho - \tilde{\rho}_n\|_1^2 \leq 4$. \square

In the second step of the measurement procedure we rotate the remaining $n - \tilde{n}$ qubits such that after rotation the vector $\tilde{\mathbf{r}}$ 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 6.4 and 6.5, 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 6.1 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 neighbourhood 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 neighbourhood of order $n^{-1/2+\epsilon}$ of the true state ρ , as shown in figure 28. In the next section we will show that for such neighbourhoods, 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.

6.3 Local Asymptotic Normality

The optimality of the second stage of the measurement relies on the concept of local asymptotic normality [Vaa98, GK06]. 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 $\mathcal{O}(n^{-1/4+\eta+\epsilon})$ for arbitrarily small ϵ .

6.3.1 Introduction to LAN and some Definitions

Let ρ_0 be a fixed state, which by rotational symmetry can be chosen of the form

$$\rho_0 = \begin{pmatrix} \mu & 0 \\ 0 & 1 - \mu \end{pmatrix}, \quad (6.4)$$

for a given $\frac{1}{2} < \mu < 1$. We parameterize the neighbouring 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

$$\rho_{\mathbf{v}} := U(\mathbf{v}) \begin{pmatrix} \mu + v_z & 0 \\ 0 & 1 - \mu - v_z \end{pmatrix} U(\mathbf{v})^\dagger, \quad (6.5)$$

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 in [BB⁺06] and in [HM06], the latter coming pretty close to our formulation of local asymptotic normality.

The precise meaning of the convergence is given in theorem 6.2 below. In short, there exist quantum channels (completely positive maps) 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 [Hol82, YL73], 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 qubit state $\hat{\rho}_n := \rho_{\hat{\mathbf{u}}/\sqrt{n}}$. The precise result is formulated in theorem 6.4

6.3.2 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

$$N^{\mathbf{u}} := N(u_z, \mu(1 - \mu)), \quad (6.6)$$

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 (see sec. 2.3.1) as

$$\mathcal{F}(\mathfrak{h}) := \bigoplus_{n=0}^{\infty} \mathfrak{h} \otimes_s \cdots \otimes_s \mathfrak{h}, \quad (6.7)$$

with \otimes_s denoting the symmetric tensor product. Thus $\mathcal{F}(\mathbb{C})$ is the simplest example of a Fock space. Let

$$\phi := (1 - p) \sum_{k=0}^{\infty} p^k |k\rangle \langle k| \quad (6.8)$$

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 $D(\alpha) := \exp(\alpha a^\dagger - \bar{\alpha} a)$ is the displacement operator (i.e. the Weyl operator $W(\alpha)$, cf. sec. 2.3.3), 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^\dagger and a are the creation and annihilation operators on $\mathcal{F}(\mathbb{C})$, satisfying $[a, a^\dagger] = \mathbf{1}$. The family $\phi^{\mathbf{u}}$ of states in which we are interested is given by

$$\phi^{\mathbf{u}} := \phi(\sqrt{2\mu - 1}\alpha_{\mathbf{u}}), \quad \mathbf{u} \in \mathbb{R}^3, \quad (6.9)$$

with $\alpha_{\mathbf{u}} := -u_y + iu_x$. 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

$$\rho_n^{\mathbf{u}} := \rho_{\mathbf{u}/\sqrt{n}}^{\otimes n}, \quad (6.10)$$

is asymptotically identical to that contained in the couple $(N^{\mathbf{u}}, \phi^{\mathbf{u}})$. More precisely:

Theorem 6.2 *Let $\rho_n^{\mathbf{u}}$ be the family of states (6.5) on the Hilbert space $(\mathbb{C}^2)^{\otimes n}$, let $N^{\mathbf{u}}$ be the family (6.6) of Gaussian distributions, and let $\phi^{\mathbf{u}}$ be the family (6.9) of displaced thermal equilibrium states of a quantum oscillator. Then for each n there exist quantum channels (trace preserving CP maps)*

$$\begin{aligned} T_n &: \mathcal{T}((\mathbb{C}^2)^{\otimes n}) \rightarrow L^1(\mathbb{R}) \otimes \mathcal{T}(\mathcal{F}(\mathbb{C})), \\ S_n &: L^1(\mathbb{R}) \otimes \mathcal{T}(\mathcal{F}(\mathbb{C})) \rightarrow \mathcal{T}((\mathbb{C}^2)^{\otimes n}) \end{aligned}$$

with $\mathcal{T}(\mathcal{H})$ the trace-class operators on \mathcal{H} , such that, for any $0 \leq \eta < 1/4$ and any $\epsilon > 0$,

$$\sup_{\|\mathbf{u}\| \leq n^\eta} \|N^{\mathbf{u}} \otimes \phi^{\mathbf{u}} - T_n(\rho_n^{\mathbf{u}})\|_1 = \mathcal{O}(n^{-1/4+\eta+\epsilon}), \quad (6.11)$$

$$\sup_{\|\mathbf{u}\| \leq n^\eta} \|\rho_n^{\mathbf{u}} - S_n(N^{\mathbf{u}} \otimes \phi^{\mathbf{u}})\|_1 = \mathcal{O}(n^{-1/4+\eta+\epsilon}). \quad (6.12)$$

Moreover, for each $\epsilon_2 > 0$ there exists a function $f(n)$ of order $\mathcal{O}(n^{-1/4+\eta+\epsilon})$ such that the above convergence rates are bounded by $f(n)$, with f independent of ρ^0 as long as $|\frac{1}{2} - \mu| > \epsilon_2$.

Remark Note that the equations (6.11) and (6.12) imply that the expressions on the left side converge to zero as $n \rightarrow \infty$. Following the classical terminology of Le Cam [LC86], we will call this type of result *strong convergence* of quantum statistical models (experiments). Another local asymptotic normality result has been derived in [GJ07] 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. A similar relation is conjectured to hold in the quantum set-up, but for the moment this has been shown only under additional assumptions [GJ07].

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 [GJ07] holds for the larger class of finite dimensional states while the strong convergence has more direct consequences as it is shown in this paper 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 [Wer98] and quantum amplification [Cav82], 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 [GK06]. It is also shown that a similar result holds uniformly over $\|\mathbf{u}\| < C$ for any fixed finite constant C . In [GJ07], it is shown that such maps also exist 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 neighbourhoods $\|\mathbf{u}\| \leq n^\eta$ 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 neighbourhood $\|\mathbf{u}\| \leq n^\eta$ for which local asymptotic normality still holds (see figure 28).

Proof. Following [GK06] we will first indicate how the channels T_n are constructed. The technical details of the proof can be found in appendix 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

$$\tilde{\pi}_n(\tau) : v_1 \otimes \cdots \otimes v_n \mapsto v_{\tau^{-1}(1)} \otimes \cdots \otimes v_{\tau^{-1}(n)}, \quad \tau \in S(n).$$

As $[\pi_n(u), \tilde{\pi}_n(\tau)] = 0$ for all $u \in SU(2), \tau \in S(n)$, we have the decomposition

$$(\mathbb{C}^2)^{\otimes n} = \bigoplus_{j=0, 1/2}^{n/2} \mathcal{H}_j \otimes \mathcal{H}_n^j. \quad (6.13)$$

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

$$\rho_n^{\mathbf{u}} = \bigoplus_{j=0,1/2}^{n/2} p_{n,\mathbf{u}}(j) \rho_{j,n}^{\mathbf{u}} \otimes \frac{\mathbf{1}}{n_j}. \quad (6.14)$$

The probability distribution $p_{n,\mathbf{u}}(j)$ is given by [BB⁺06]:

$$p_{n,\mathbf{u}}(j) := \frac{n_j}{2\mu_{\mathbf{u}} - 1} (1 - \mu_{\mathbf{u}})^{\frac{n}{2}-j} \mu_{\mathbf{u}}^{\frac{n}{2}+j+1} (1 - p_{\mathbf{u}}^{2j+1}), \quad (6.15)$$

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

$$p_{n,\mathbf{u}}(j) := B_{n,\mu_{\mathbf{u}}}(n/2 + j) \times K(j, n, \mu, \mathbf{u}), \quad (6.16)$$

where

$$B_{n,\nu}(k) := \binom{n}{k} \nu^k (1 - \nu)^{n-k}, \quad k = 0, \dots, n$$

is a binomial distribution, and the factor $K(j, n, \mu, \mathbf{u})$ is given by

$$K(j, n, \mu, \mathbf{u}) := (1 - p_{\mathbf{u}}^{2j+1}) \frac{n + (2(j - j_n - \sqrt{n}u_z) + 1)/(2\mu_{\mathbf{u}} - 1)}{n + (j - j_n - \sqrt{n}u_z + 1)/\mu_{\mathbf{u}}},$$

with $j_n := n(\mu - 1/2)$. Now $K(j, n, \mu, \mathbf{u}) = 1 + \mathcal{O}(n^{-1/2+\epsilon})$ 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 [Vaa98] 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 converge to the Gaussian distribution in total variation norm.

The measurement “which block”, corresponding to the decomposition (6.14), 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 “smoothing” with a Gaussian kernel of mean $g_n(j)$ and variance $1/(2\sqrt{n})$, i.e. with $\tau_{n,j} := (n^{1/4}/\sqrt{\pi}) \exp(-\sqrt{n}(x - g_n(j))^2)$.

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 [GK06] and consists of embedding each block state $\rho_{j,n}^{\mathbf{u}}$ into the state space of the oscillator by means of an isometry $V_j : \mathcal{H}_j \rightarrow \mathcal{F}(\mathbb{C})$,

$$V_j : |j, m\rangle \mapsto |j - m\rangle,$$

where $\{|j, m\rangle : m = -j, \dots, j\}$ is the eigenbasis of the total spin component $L_z := \sum_i \sigma_z^{(i)}$, cf. equation (5.1) of [GK06]. Then the action of the channel T_n is

$$T_n : \bigoplus_j p_{n,\mathbf{u}}(j) \rho_{j,n}^{\mathbf{u}} \otimes \frac{\mathbf{1}}{n_j} \mapsto \sum_j p_{n,\mathbf{u}}(j) \tau_{n,j} \otimes V_j \rho_{j,n}^{\mathbf{u}} V_j^\dagger.$$

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 (6.13) (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 A. \square

6.4 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 6.5.2, 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 6.5.

6.4.1 Quantum Stochastic Differential Equations

In the weak coupling limit, [AFL90, GZ04], the joint evolution of the qubits and field can be described mathematically by quantum stochastic differential equations (QSDE). See section 2.3 for a short introduction, and [HP84, Par92] for a thorough account. 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 (6.7). An important linearly complete set in $\mathcal{F}(L^2(\mathbb{R}))$ is that of the exponential vectors

$$e(f) := \bigoplus_{n=0}^{\infty} \frac{1}{\sqrt{n!}} f^{\otimes n} := \bigoplus_{n=0}^{\infty} \frac{1}{\sqrt{n!}} |f\rangle_n, \quad f \in L^2(\mathbb{R}), \quad (6.17)$$

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 $A_t^\dagger := a^\dagger(\chi_{[0,t]})$ 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^\dagger play the role of non-commuting integrators in quantum stochastic differential equations, comparable to the role of Brownian motion in classical stochastic calculus.

We now consider the joint unitary evolution for qubits and field defined by the quantum stochastic differential equation [HP84, BGM04]:

$$dU_n(t) = (a_n dA_t^\dagger - a_n^\dagger dA_t - \frac{1}{2} a_n^\dagger a_n dt) U_n(t),$$

where $U_n(t)$ is a unitary operator on $(\mathbb{C}^2)^{\otimes n} \otimes \mathcal{F}(L^2(\mathbb{R}))$, and

$$a_n := \frac{1}{\sqrt{2j_n}} \sum_{k=1}^n \sigma_+^{(k)}, \quad \sigma_+^{(k)} := \mathbf{1} \otimes \cdots \otimes (\sigma_x + i\sigma_y)/2 \otimes \cdots \otimes \mathbf{1},$$

with $j_n := (\mu - 1/2)n$. 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 (6.13), 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

$$U_{j,n}(t) : \mathcal{H}_j \otimes \mathcal{F}(L^2(\mathbb{R})) \rightarrow \mathcal{H}_j \otimes \mathcal{F}(L^2(\mathbb{R})),$$

is the restricted cocycle

$$dU_{j,n}(t) = (a_j dA_t^\dagger - a_j^\dagger dA_t - \frac{1}{2} a_j^\dagger a_j dt) U_{j,n}(t), \quad (6.18)$$

with a_j acting on the basis $|j, m\rangle$ of \mathcal{H}_j as

$$\begin{aligned} a_j |j, m\rangle &= \sqrt{j-m} \sqrt{(j+m+1)/2j_n} |j, m+1\rangle, \\ a_j^\dagger |j, m\rangle &= \sqrt{j-m+1} \sqrt{j+m/2j_n} |j, m-1\rangle. \end{aligned}$$

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 (6.14) 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 $n \rightarrow \infty$, 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.

6.4.2 Solving the QSDE for the Oscillator

Let a^\dagger 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

$$dU(t) = (adA_t^\dagger - a^\dagger dA_t - \frac{1}{2}a^\dagger adt)U(t).$$

We choose the initial (unnormalized) 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 is given by $f_t(s) := f(s)\chi_{[0,t]}(s)$ 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

$$dI(t) = (\bar{\beta} \frac{d}{dt} \alpha(t) + \bar{g}(t)f(t)) I(t)dt. \quad (6.19)$$

We now calculate $\frac{d}{dt}I(t)$ with the help of the QSDE. Since the annihilator acts as $A_t e(f) = \langle \chi_{[0,t]} | f \rangle e(f)$, 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

$$dI(t) = \langle e(\beta) \otimes e(g) | (adA_t^\dagger - a^\dagger dA_t - \frac{1}{2}a^\dagger adt)\psi(t) \rangle,$$

so that

$$dI(t) = (\bar{g}(t)\alpha(t) - \frac{1}{2}\bar{\beta}\alpha(t))I(t)dt. \quad (6.20)$$

Equating (6.19) with (6.20) for all t , all β , and all 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, $\psi(t) = e(\mathbf{z}e^{-\frac{1}{2}t}) \otimes e(\mathbf{z}e^{-\frac{1}{2}t}\chi_{[0,t]}(s))$. For later use, we denote by $\psi_{\mathbf{z}}(t) := U(t)|\mathbf{z}\rangle \otimes |\Omega\rangle = e^{-|\mathbf{z}|^2/2}U(t)e(\mathbf{z}) \otimes |\Omega\rangle$ the *normalized* solution of the QSDE.

6.4.3 QSDE for Large Spin

We consider now the unitary evolution for qubits and field:

$$dU_n(t) = (a_n dA_t^\dagger - a_n^\dagger dA_t - \frac{1}{2}a_n^\dagger a_n dt)U_n(t).$$

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 can write $\psi(t+dt) - \xi(t+dt)$ as

$$\psi(t+dt) - U_{t+dt}^t \xi(t) + U_{t+dt}^t \xi(t) - \xi(t) + \xi(t) - \xi(t+dt).$$

This equals

$$U_{t+dt}^t [\psi(t) - \xi(t)] + [U(t+dt) - U(t)]U(t)^{-1}\xi(t)[\xi(t) - \xi(t+dt)],$$

so that

$$\psi(t+dt) - \xi(t+dt) = U_{t+dt}^t [\psi(t) - \xi(t)] + G_{dt}\xi(t) - d\xi(t).$$

By taking norms we finally obtain

$$d\|\psi(t) - \xi(t)\| \leq \|G_{dt}\xi(t) - d\xi(t)\|. \quad (6.21)$$

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 (6.18)) for a “typical” $j \in \mathcal{J}_n$ (see equation (6.39)). 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 6.3 *Let $\rho_{j,n}^{\mathbf{u}}(t) = U_{j,n}(t) [\rho_{j,n}^{\mathbf{u}} \otimes \Phi] U_{j,n}^\dagger(t)$ be the j -th block of the state of qubits and field at time t . Let $\phi^{\mathbf{u}}(t) := U(t) [\phi^{\mathbf{u}} \otimes \Phi] U(t)^\dagger$ be the joint state of the oscillator and field at time t . For any $\eta < 1/6$, for any $\epsilon > 0$,*

$$\sup_{j \in \mathcal{J}_n} \sup_{\|\mathbf{u}\| \leq n^\eta} \sup_t \|\rho_{j,n}^{\mathbf{u}}(t) - \phi^{\mathbf{u}}(t)\|_1 = \mathcal{O}(n^{-1/4+\eta+\epsilon}, n^{-1/2+3\eta+\epsilon}). \quad (6.22)$$

Proof. From the proof of the local asymptotic normality theorem 6.2 we know that the initial states of the two unitary evolutions are asymptotically close to each other

$$\sup_{j \in \mathcal{J}_n} \sup_{\|\mathbf{u}\| \leq n^\eta} \|\rho_{j,n}^{\mathbf{u}} - \phi^{\mathbf{u}}\|_1 = \mathcal{O}(n^{-1/4+\eta+\epsilon}). \quad (6.23)$$

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 $\rho_{j,n}^{\mathbf{u}}$:

$$\sup_{j \in \mathcal{J}_n} \sup_{\|\mathbf{u}\| \leq n^\eta} \|\tilde{\rho}_{j,n}^{\mathbf{u}} - \phi^{\mathbf{u}}\|_1 = \mathcal{O}(e^{-n^\epsilon}). \quad (6.24)$$

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

$$\sup_{j \in \mathcal{J}_n} \sup_{\|\mathbf{u}\| \leq n^\eta} \sup_t \|\tilde{\rho}_{j,n}^{\mathbf{u}}(t) - \phi^{\mathbf{u}}(t)\|_1 = \mathcal{O}(n^{-1/4+\eta+\epsilon}, n^{-1/2+3\eta+\epsilon}). \quad (6.25)$$

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 (6.23), (6.24) and (6.25), and using the triangle inequality, we get

$$\sup_{j \in \mathcal{J}_n} \sup_{\|\mathbf{u}\| \leq n^\eta} \sup_t \|\rho_{j,n}^{\mathbf{u}}(t) - \phi^{\mathbf{u}}(t)\|_1 = \mathcal{O}(n^{-1/4+\eta+\epsilon}, n^{-1/2+3\eta+\epsilon}).$$

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 .

$$\begin{array}{ccccc} \mathcal{S}(\mathcal{H}_j) & \xrightarrow{\text{Id}_j \otimes \Phi} & \mathcal{S}(\mathcal{H}_j \otimes \mathcal{F}) & \xrightarrow{U_{j,n}(t)} & \mathcal{S}(\mathcal{H}_j \otimes \mathcal{F}) \\ \downarrow V_j \cdot V_j^\dagger & & \downarrow & & \downarrow \\ \mathcal{S}(\mathcal{F}(\mathbb{C})) & \xrightarrow{\text{Id} \otimes \Phi} & \mathcal{S}(\mathcal{F}(\mathbb{C}) \otimes \mathcal{F}) & \xrightarrow{U(t)} & \mathcal{S}(\mathcal{F}(\mathbb{C}) \otimes \mathcal{F}) \end{array}$$

For the rest of the proof, we refer to appendix B. \square

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.

6.5 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 neighbourhood such that the machinery of local asymptotic normality of theorem 6.2 can be applied.

In theorem 6.3 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

$$\begin{aligned} \rho_n^{\mathbf{u}}(t) & := \bigoplus_{j=0,1/2}^{n/2} p_{n,\mathbf{u}}(j) \frac{1}{2\pi s^2} \int_{\mathbb{C}} d\mathbf{z} e^{-|\mathbf{z} - \sqrt{2\mu-1}\alpha_{\mathbf{u}}|^2/2s^2} \exp(-|\mathbf{z}|^2) \times \\ & |e(\mathbf{z}e^{-t/2})_j\rangle \langle e(\mathbf{z}e^{-t/2})_j| \otimes |e(\mathbf{z}e^{-u/2}\chi_{[0,t]}(u))\rangle \langle e(\mathbf{z}e^{-u/2}\chi_{[0,t]}(u))| \\ & + \mathcal{O}(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon}), \end{aligned}$$

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

$$\exp(-|\mathbf{z}|^2) \langle e(\mathbf{z}e^{-t/2})_j | j, j \rangle \langle e(\mathbf{z}e^{-u/2} \chi_{[0,t]}(u)) | e(\mathbf{z}e^{-u/2}) \rangle = \exp(-|\mathbf{z}|^2 e^{-t}). \quad (6.26)$$

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 $\rho_n^{\mathbf{u}}(t) = \rho_n^{\mathbf{u}}(\infty) + \mathcal{O}(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon})$ with

$$\rho_n^{\mathbf{u}}(\infty) := \bigoplus_{j=0,1/2}^{n/2} p_{n,\mathbf{u}}(j) |j, j\rangle \langle j, j| \otimes \left[\frac{1}{2\pi s^2} \int_{\mathbb{C}} d\mathbf{z} e^{-|\mathbf{z}-\sqrt{2\mu-1}\alpha_{\mathbf{u}}|^2/2s^2} |e(\mathbf{z}e^{-u/2})\rangle \langle e(\mathbf{z}e^{-u/2})| \exp(-|\mathbf{z}|^2) \right]. \quad (6.27)$$

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 6.2. 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 [Hol82, YL73]; for the latter, we perform a “which block” measurement. These measurements are described in the next two subsections.

6.5.1 The Heterodyne Measurement

A heterodyne measurement is a “joint measurement” of the two quadratures $\mathbf{Q} := (a + a^\dagger)/\sqrt{2}$ and $\mathbf{P} := -i(a - a^\dagger)/\sqrt{2}$ 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 $\mathcal{O}(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon})$. From (6.9) 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

$$N(u_x, \mu/(2(2\mu-1)^2)) \times N(u_y, \mu/(2(2\mu-1)^2)). \quad (6.28)$$

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) :

$$\hat{u}_i = \begin{cases} \tilde{u}_i & \text{if } |\tilde{u}_i| \leq 3n^\eta \\ 0 & \text{if } |\tilde{u}_i| > 3n^\eta \end{cases} \quad (6.29)$$

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.

6.5.2 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 (6.27) (for large time) over the field

$$\tau_n^{\mathbf{u}} = \sum_{j=0,1/2}^{n/2} p_{n,\mathbf{u}}(j) |j, j\rangle \langle j, j| + \mathcal{O}(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon}).$$

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 $\mathcal{O}(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon})$). We couple the atoms to a new field (in the vacuum state $|\Omega\rangle$) by means of the interaction

$$dU_t = \{J_z(dA_t^\dagger - dA_t) - \frac{1}{2}J_z^2 dt\}U_t,$$

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 $B_t = (A_t^\dagger - A_t)/i$, the solution is easily seen to be

$$U_t = \exp(J_z \otimes (A_t^\dagger - A_t)).$$

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

$$d \exp(iJ_z \otimes B_t) = \{iJ_z dB_t - \frac{1}{2}J_z^2 dt\} \exp(iJ_z \otimes B_t).$$

For an initial state $|j, m\rangle \otimes |\Omega\rangle$, this evolution gives rise to the final state

$$\begin{aligned} U_t |j, m\rangle \otimes \Omega &= |j, m\rangle \otimes \exp((m/\sqrt{n})(A_t^\dagger - A_t))\Omega \\ &= |j, m\rangle \otimes |(m/\sqrt{n})\chi_{[0,t]} \rangle, \end{aligned}$$

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

$$\begin{aligned} U_t \tau_n^{\mathbf{u}} \otimes \Phi U_t^\dagger &= \sum_{j=0,1/2}^{n/2} p_{n,\mathbf{u}}(j) |j, j\rangle\langle j, j| \otimes |j/\sqrt{n}\chi_{[0,t]}\rangle\langle j/\sqrt{n}\chi_{[0,t]}| \\ &\quad + \mathcal{O}(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon}). \end{aligned}$$

The final state of the field results from a partial trace over the atoms; it is given by

$$\sum_{j=0,1/2}^{n/2} p_{n,\mathbf{u}}(j) |(j/\sqrt{n})\chi_{[0,t]}\rangle\langle (j/\sqrt{n})\chi_{[0,t]}| + \mathcal{O}(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon}). \quad (6.30)$$

We now perform a homodyne measurement on the field, which amounts to a direct measurement of $(A_t + A_t^\dagger)/2t$. 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, $\mathbb{E}((A_t + A_t^\dagger)/2t) = j/\sqrt{n}$, whereas $\mathbb{V}\text{ar}((A_t + A_t^\dagger)/2t) = 1/(4t)$. Thus the probability distribution $p_{n,\mathbf{u}}$ is reproduced up to order $\mathcal{O}(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon})$ in L^1 -distance.

The following is a reminder from the proof of theorem 6.2. 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 $\mathcal{O}(n^{\eta-1/2}) + \mathcal{O}(n^{\epsilon-1/2})$. 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

$$N(u_z, \mu(1 - \mu)) + \mathcal{O}(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon}) + \mathcal{O}(n^{\eta-1/2}) + \mathcal{O}(n^{\epsilon-1/2}). \quad (6.31)$$

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

$$\hat{u}_z = \begin{cases} g_n & \text{if } |g_n| \leq 3n^\eta \\ 0 & \text{if } |g_n| > 3n^\eta. \end{cases} \quad (6.32)$$

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 $|\hat{u}_z - u_z| \leq |\tilde{u}_z - u_z|$.

6.6 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 6.4 *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 $\hat{\rho}_n$, the following local asymptotic minimax result holds for any $0 < \epsilon < 1/12$:*

$$\limsup_{n \rightarrow \infty} \sup_{\|\rho - \rho_0\|_1 \leq n^{-1/2+\epsilon}} nR(\rho, \hat{\rho}_n) \leq \limsup_{n \rightarrow \infty} \sup_{\|\rho - \rho_0\|_1 \leq n^{-1/2+\epsilon}} nR(\rho, \hat{\rho}_n). \quad (6.33)$$

Let $(\mu_0, 1 - \mu_0)$ be the eigenvalues of ρ_0 with $\mu_0 > 1/2$. Then the local asymptotic minimax risk is

$$\limsup_{n \rightarrow \infty} \sup_{\|\rho - \rho_0\|_1 \leq n^{-1/2+\epsilon}} nR(\rho, \hat{\rho}_n) = R_{\text{minimax}}(\mu_0) = 8\mu_0 - 4\mu_0^2. \quad (6.34)$$

Proof. 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 $\|\rho - \sigma\|_1^2 \leq 4$. Thus

$$R_1 = 4\mathbb{P}(\|\rho - \tilde{\rho}_n\|_1 \geq n^{-1/2+\epsilon}).$$

According to lemma 6.1 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 $\rho^{\otimes n} = \rho_{\mathbf{u}}^{\mathbf{u}}$ with \mathbf{u} the local parameter around $\tilde{\rho}_n$.

We get the following inequalities for the second term in the risk.

$$\begin{aligned}
\mathbb{E}(\|\rho - \hat{\rho}_n\|_1^2 \chi_E) &\leq \mathbb{E} \left[\|\hat{\rho}_n - \rho\|_1^2 \mid \|\tilde{\rho}_n - \rho\|_1 \leq n^{-1/2+\epsilon} \right] \\
&\leq \sup_{\|\rho - \rho_0\| < n^{-1/2+\epsilon}} \mathbb{E} \left[\|\hat{\rho}_n - \rho\|_1^2 \mid \tilde{\rho}_n = \rho_0 \right] \\
&\leq \sup_{\|\rho - \rho_0\| < n^{-1/2+\epsilon}} \mathbb{E}_{\rho_n^{\mathbf{u}}(\infty)} \left[\|\hat{\rho}_n - \rho\|_1^2 \mid \tilde{\rho}_n = \rho_0 \right] \\
&\quad + \sup_{\|\rho - \rho_0\| < n^{-1/2+\epsilon}} \|\rho_n^{\mathbf{u}}(t) - \rho_n^{\mathbf{u}}(\infty)\|_1 \sup_{\hat{\mathbf{u}}_n} \|\hat{\rho}_n - \rho\|_1^2 \\
&\leq \sup_{\|\rho - \rho_0\| < n^{-1/2+\epsilon}} \mathbb{E}_{\rho_n^{\mathbf{u}}(\infty)} \left[\|\hat{\rho}_n - \rho\|_1^2 \mid \tilde{\rho}_n = \rho_0 \right] \\
&\quad + cn^{-1+2\eta} \sup_{\|\rho - \rho_0\| < n^{-1/2+\epsilon}} \|\rho_n^{\mathbf{u}}(t) - \rho_n^{\mathbf{u}}(\infty)\|_1 \\
&= R_2. \tag{6.35}
\end{aligned}$$

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 (6.29) and (6.32) 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 $\rho^0 = \tilde{\rho}$ 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 $\mathcal{O}(n^{-5/4+3\eta+\delta}, n^{-3/2+5\eta+\delta})$ where $\delta > 0$ can be chosen arbitrarily small. Indeed, in the end of section 6.4, we have proven that after time $t \geq \ln n$, we have $\|\rho_n^{\mathbf{u}}(t) - \rho_n^{\mathbf{u}}(\infty)\|_1 = \mathcal{O}(n^{-1/4+\eta+\delta}, n^{-1/2+3\eta+\delta})$. 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 parameterization around $\rho_0 = \tilde{\rho}$ as $\rho_{\mathbf{u}_n/\sqrt{n}}$. We have

$$\begin{aligned}
\|\hat{\rho}_n - \rho\|_1^2 &= \|\rho_{\mathbf{u}/\sqrt{n}} - \rho_{\hat{\mathbf{u}}_n/\sqrt{n}}\|_1^2 \\
&= \frac{(u_z - \hat{u}_z)^2 + (2\mu - 1)^2((u_x - \hat{u}_x)^2 + (u_y - \hat{u}_y)^2)}{n} \\
&\quad + \mathcal{O}(\|\mathbf{u} - \hat{\mathbf{u}}_n\|^3 n^{-3/2}). \tag{6.36}
\end{aligned}$$

The remainder term $\mathcal{O}(\|\mathbf{u} - \hat{\mathbf{u}}_n\|^3 n^{-3/2})$ is negligible. It is $\mathcal{O}(n^{3\eta-3/2})$ 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 (6.29) and (6.32).

We now evaluate $\mathbb{E}_{\rho_n^{\mathbf{u}}(\infty)} [d(\mathbf{u}, \hat{\mathbf{u}}_n)^2]$, where $d(\mathbf{u}, \hat{\mathbf{u}}_n)^2$ is defined as

$$d(\mathbf{u}, \mathbf{v})^2 := 4 [(u_z - v_z)^2 + (2\mu - 1)^2((u_x - v_x)^2 + (u_y - v_y)^2)]. \quad (6.37)$$

Note that the risk of $\hat{\mathbf{u}}_n$ is smaller than that of $\tilde{\mathbf{u}}_n$ (see discussion below (6.29) and (6.32)). Under the law $\mathbb{P}_{\rho_n^{\mathbf{u}}(\infty)}$ the estimator $\tilde{\mathbf{u}}_n$ has a Gaussian distribution as shown in (6.28) and (6.31) with fixed and known variance and unknown expectation. In statistics this type of model is known as a Gaussian shift experiment [Vaa98]. Using (6.28) and (6.31), 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 (6.36), we obtain (6.34).

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 $\hat{\rho}_n$ we have

$$R_0 = \limsup_{n \rightarrow \infty} \sup_{\|\rho - \rho_0\|_1 \leq n^{-1/2+\epsilon}} nR(\rho, \hat{\rho}_n) \geq 8\mu_0 - 4\mu_0^2.$$

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 [GK06]. 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 [Vaa98].

We now end the proof by using the other direction of LAN. Suppose that there exists a better sequence of estimators $\hat{\rho}_n$ 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 $\hat{\mathbf{u}}$ 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 transmittivity of the beamsplitter ($r^2 + t^2 = 1$), then the transmitted beam has state $\phi_{\text{tr}}^{\mathbf{u}} = \phi^{t\mathbf{u}}$ and the reflected one $\phi_{\text{ref}}^{\mathbf{u}} = \phi^{r\mathbf{u}}$. 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 6.1. 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 6.2 and by (6.12) (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 $\hat{\rho}_n$ 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 (6.29) and (6.32). The risk of this estimation procedure for $\phi^{\mathbf{u}}$ is then bounded from above by the sum of three terms: the risk $nR_\rho(\hat{\rho}_n)/t^2$ coming from the qubit estimation, the error contribution from the map S_n which is $a^2n^{-1/4+\epsilon+\delta}$, and the localization risk contribution $\tilde{\epsilon}$. This risk bound uses the same technique as the third inequality of (6.35). The second contribution can be made arbitrarily small by choosing n large enough, for $\epsilon < 1/4$. From our assumption we have $R_0 < R_{\text{minimax}}(\mu_0)$ and we can choose t close to one such that $R_0/t^2 < R_{\text{minimax}}(\mu_0)$ and further choose $\tilde{\epsilon}$ such that $R_0/t^2 + \tilde{\epsilon} < R_{\text{minimax}}(\mu_0)$.

In conclusion, we get that the risk for estimating \mathbf{u} is asymptotically smaller than the risk of the heterodyne measurement combined with observing the classical part which is known to be minimax [GK06]. Hence no such sequence $\hat{\rho}_n$ exists, and $\hat{\rho}_n$ is optimal. \square

Remark In theorem 6.33, 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.

$$d^2(\rho_{\hat{\mathbf{u}}}, \rho_{\mathbf{u}}) = \sum_{\alpha, \beta=x, y, z} \gamma_{\alpha\beta}(u_\alpha - \hat{u}_\alpha)(u_\beta - \hat{u}_\beta) + \mathcal{O}(\|\mathbf{u} - \hat{\mathbf{u}}\|^3).$$

For instance, one may choose $d^2(\hat{\rho}, \rho) = 1 - F^2(\hat{\rho}, \rho)$ with the fidelity $F(\hat{\rho}, \rho) := \text{tr}(\sqrt{\sqrt{\hat{\rho}}\rho\sqrt{\hat{\rho}}})$. 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

$$\limsup_{n \rightarrow \infty} \sup_{\|\rho - \rho_0\|_1 \leq n^{-1/2+\epsilon}} nR_F(\rho, \hat{\rho}_n) = \mu_0 + 1/4, \quad (6.38)$$

with the same asymptotically optimal $\hat{\rho}$. The asymptotic rate $R_F \sim \frac{4\mu_0+1}{4n}$ was found earlier in [BB⁺06], using different methods.

6.7 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 [BB⁺06, HM06], 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 6.2, 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 6.3, 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, in the limit $n \rightarrow \infty$, 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 neighbourhoods 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 6.4 and 6.5 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 6.6, 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 $\mu_0 > \frac{1}{2}$ an eigenvalue of ρ .

More precisely, we have

$$\limsup_{n \rightarrow \infty} \sup_{\|\rho - \rho_0\|_1 \leq n^{-1/2+\epsilon}} nR(\rho, \hat{\rho}_n) = 8\mu_0 - 4\mu_0^2.$$

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 $\|\hat{\rho} - \rho\|_1 := \text{tr}(|\hat{\rho} - \rho|)$. 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 6.2 does not apply in this case. The effect of the singularity is that the family of states (6.9) collapses to a single degenerate state of infinite temperature. However this phenomenon is only due to our particular parameterization, which was chosen for its convenience in describing the local neighbourhoods around arbitrary states, with the exception of the totally mixed state. Had we chosen a different parameterization, 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 (6.18)). 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.

A Appendix: Proof of Theorem 6.2

Here we give the technical details of the proof of local asymptotic normality with “slowly growing” local neighbourhoods $\|\mathbf{u}\| \leq n^\eta$, with $\eta < 1/4$. We start with the map T_n .

A.1 The Map T_n

Let us define, for $0 < \epsilon < (1/4 - \eta)$ the interval

$$\mathcal{J}_n = \left\{ j; (\mu - 1/2)n - n^{1/2+\epsilon} \leq j \leq (\mu - 1/2)n + n^{1/2+\epsilon} \right\}. \quad (6.39)$$

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

$\mu - 1/2 \geq \epsilon_2$:

$$\lim_{n \rightarrow \infty} p_{n,\mathbf{u}}(\mathcal{J}_n) = 1 - \mathcal{O}(n^{-1/2+\epsilon}). \quad (6.40)$$

This is a consequence of Hoeffding's inequality applied to the binomial distribution, and recalling that $p_{n,\mathbf{u}}(j) = B(n/2 + j)(1 + \mathcal{O}(n^{-1/2+\epsilon}))$ 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

$$3 \sum_{j \notin \mathcal{J}_n} p_{n,j}^{\mathbf{u}} + \left\| N^{\mathbf{u}} - \sum_{j \in \mathcal{J}_n} p_{n,\mathbf{u}}(j) \tau_{n,j} \right\|_1 + \sup_{j \in \mathcal{J}_n} \|V_j \rho_{j,n}^{\mathbf{u}} V_j^\dagger - \phi^{\mathbf{u}}\|_1. \quad (6.41)$$

The first two terms are ‘‘classical’’ and converge to zero uniformly over $\|\mathbf{u}\| \leq n^\eta$. For the first term, this is (6.40), while the second term converges uniformly on $\mu - 1/2 \geq \epsilon_2$ at rate $n^{\eta-1/2}$ [KG09]. The third term can be analyzed as in proposition 5.1 of [GK06]:

$$\left\| V_j \rho_{n,j}^{\mathbf{u}} V_j^\dagger - \phi^{\mathbf{u}} \right\|_1 \leq \left\| \rho_{n,j}^{\mathbf{u}} - V_j^\dagger \phi^{\mathbf{u}} V_j \right\|_1 + \|\phi^{\mathbf{u}} - P_j \phi^{\mathbf{u}} P_j\|_1, \quad (6.42)$$

where $P_j := V_j V_j^\dagger$ 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

$$\phi^{\mathbf{u}} = \frac{1}{\sqrt{2\pi s^2}} \int e^{-|\mathbf{z} - \sqrt{2\mu-1}\alpha_{\mathbf{u}}|^2/2s^2} (|\mathbf{z}\rangle\langle\mathbf{z}|) d^2\mathbf{z}, \quad (6.43)$$

where $s^2 := (1 - \mu)/(4\mu - 2)$.

The second term on the right side of (6.42) 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

$$\int e^{-|\mathbf{z}|^2/2s^2} \|P_j^\perp|\mathbf{z} + \sqrt{2\mu-1}\alpha_{\mathbf{u}}\rangle\| d^2\mathbf{z}. \quad (6.44)$$

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 $\mathcal{O}(e^{-n^{2(\eta+\delta)}/(2s^2)})$. 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 $\|P_j^\perp|\mathbf{z}\rangle\| \leq |\mathbf{z}|^j/\sqrt{j!} = \mathcal{O}(e^{-n(1/2-\eta-2\delta)})$ 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:

$$\begin{aligned} \left\| \rho_{n,j}^{\mathbf{u}} - V_j^\dagger \phi^{\mathbf{u}} V_j \right\|_1 &= \left\| \text{Ad} \left[U_j \left(\frac{\mathbf{u}}{\sqrt{n}} \right) \right] (\rho_{n,j}^{\mathbf{0}}) - V_j^\dagger \phi^{\mathbf{u}} V_j \right\|_1 \leq \\ &\left\| \rho_{n,j}^{\mathbf{0}} - V_j^\dagger \phi^{\mathbf{0}} V_j \right\|_1 + \left\| \text{Ad} \left[U_j \left(\frac{\mathbf{u}}{\sqrt{n}} \right) \right] (V_j^\dagger \phi^{\mathbf{0}} V_j) - V_j^\dagger \phi^{\mathbf{u}} V_j \right\|_1. \end{aligned} \quad (6.45)$$

The first term on the right side does not depend on \mathbf{u} . From the proof of lemma 5.4 of [GK06], we know that

$$\left\| \rho_{n,j}^{\mathbf{0}} - V_j^\dagger \phi^{\mathbf{0}} V_j \right\|_1 \leq \left(\frac{p^{2j+1}}{1-p^{2j+1}} + p^{2j+1} \right)$$

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 (6.45) can be bounded again by a Gaussian integral

$$\frac{1}{\sqrt{2\pi s^2}} \int e^{-|\mathbf{z}|^2/2s^2} \|\Delta(\mathbf{u}, \mathbf{z}, j)\|_1 d^2 \mathbf{z}, \quad (6.46)$$

where the operator $\Delta(\mathbf{u}, \mathbf{z}, j)$ is given by

$$\begin{aligned} \Delta(\mathbf{u}, \mathbf{z}, j) &:= \text{Ad} \left[U_j \left(\frac{\mathbf{u}}{\sqrt{n}} \right) \right] \left(V_j^\dagger |\mathbf{z}\rangle \langle \mathbf{z}| V_j \right) \\ &\quad - V_j^\dagger |\mathbf{z} + \sqrt{2\mu - 1}\alpha_{\mathbf{u}}\rangle \langle \mathbf{z} + \sqrt{2\mu - 1}\alpha_{\mathbf{u}}| V_j. \end{aligned}$$

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 the vector $\tilde{\mathbf{z}} \in \mathbb{R}^2$ be such that $\alpha_{\tilde{\mathbf{z}}} = \mathbf{z}/\sqrt{2\mu - 1}$, and introduce the notation $\psi(n, j, \mathbf{v}) = V_j U_j(\mathbf{v}/\sqrt{n})|j, j\rangle$. Then, up to a $\sqrt{2}$ factor, $\|\Delta(\mathbf{u}, \mathbf{z}, j)\|_1$ is bounded from above by

$$\begin{aligned} &\left\| \psi(n, j, \tilde{\mathbf{z}}) - |\mathbf{z}\rangle \right\| + \\ &\left\| \psi(n, j, \mathbf{u} + \tilde{\mathbf{z}}) - |\mathbf{z} + \sqrt{2\mu - 1}\alpha_{\mathbf{u}}\rangle \right\| + \\ &\left\| U_j \left(\frac{\mathbf{u}}{\sqrt{n}} \right) U_j \left(\frac{\tilde{\mathbf{z}}}{\sqrt{n}} \right) |jj\rangle - U_j \left(\frac{\mathbf{u} + \tilde{\mathbf{z}}}{\sqrt{n}} \right) |jj\rangle \right\|. \end{aligned} \quad (6.47)$$

This is obtained by adding and subtracting the states $|\psi(n, j, \tilde{\mathbf{z}})\rangle \langle \psi(n, j, \tilde{\mathbf{z}})|$ 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 first two terms are similar, we want to dominate them uniformly. We replace $\mathbf{u} + \tilde{\mathbf{z}}$ by $\tilde{\mathbf{z}}$ with $|\mathbf{z}| \leq 2n^{\eta+\delta}$, and then write:

$$\begin{aligned} \|\psi(n, j, \tilde{\mathbf{z}}) - |\mathbf{z}\rangle\|^2 &= \sum_{k=0}^{\infty} |\langle k|\psi(n, j, \tilde{\mathbf{z}})\rangle - \langle k|\mathbf{z}\rangle|^2 \\ &\leq \sum_{k=0}^{r-1} |\langle k|\psi(n, j, \tilde{\mathbf{z}})\rangle - \langle k|\mathbf{z}\rangle|^2 \\ &\quad + 2 \sum_{k=r}^{\infty} (|\langle k|\psi(n, j, \tilde{\mathbf{z}})\rangle|^2 + |\langle k|\mathbf{z}\rangle|^2). \end{aligned} \quad (6.48)$$

If $\mathbf{z} = |\mathbf{z}|e^{i\theta}$ then we have [HM06]

$$\begin{aligned}\langle k|\psi(n, j, \tilde{\mathbf{z}})\rangle &= \sqrt{\binom{2j}{k}} (\sin(|\mathbf{z}|/\sqrt{n})e^{i\theta})^k (\cos(|\mathbf{z}|/\sqrt{n}))^{2j-k}, \\ \langle k|\mathbf{z}\rangle &= \exp\left(-\frac{(2\mu-1)|\mathbf{z}|^2}{2}\right) \frac{(e^{i\theta}|\mathbf{z}|\sqrt{2\mu-1})^k}{\sqrt{k!}}.\end{aligned}$$

In (6.48), we choose $r = n^{2\eta+\epsilon_3}$ with ϵ_3 such that it satisfies the conditions $2\delta+2\eta+\epsilon < 2\eta+\epsilon_3+\epsilon < 1/2$ and $\eta+\epsilon_3 < 1/4$. The tail sums are then of the order

$$\begin{aligned}\sum_{k=r}^{\infty} |\langle k|\mathbf{z}\rangle|^2 &\leq \frac{|\mathbf{z}|^{2r}}{r!} \leq \frac{(2n^{(\eta+\delta)})2n^{2\eta+\epsilon_3}}{(n^{2\eta+\epsilon_3})!} \\ &= \mathcal{O}(\exp(-n^{2\eta+\epsilon_3})), \\ \sum_{k=r}^{\infty} |\langle k|\psi(n, j, \tilde{\mathbf{z}})\rangle|^2 &\leq \sum_{k=r}^j \left(\frac{|\mathbf{z}|^2}{n}\right)^k \frac{(2j)!}{(2j-k)!k!} \leq n \frac{|\mathbf{z}|^{2r}}{r!} \\ &= \mathcal{O}(\exp(-n^{2\eta+\epsilon_3})).\end{aligned}$$

For the finite sums we use the following estimates which are uniform over all $|\mathbf{z}| \leq 2n^{\eta+\delta}$, $k \leq r$, $j \in \mathcal{J}_n$:

$$\begin{aligned}\sqrt{\binom{2j}{k}} &= \frac{((2\mu-1)n)^{k/2}}{\sqrt{k!}} (1 + \mathcal{O}(n^{-1/2+\epsilon+2\eta+\epsilon_3})), \\ (\sin(|\mathbf{z}|/\sqrt{n}))^k &= (|\mathbf{z}|/\sqrt{n})^k (1 + \mathcal{O}(n^{4\eta+\epsilon_3+2\delta-1})), \\ (\cos(|\mathbf{z}|/\sqrt{n}))^{2j-k} &= \exp\left(-\frac{(2\mu-1)|\mathbf{z}|^2}{2}\right) (1 + \mathcal{O}(n^{2\eta-1/2+\epsilon+2\delta})),\end{aligned}$$

where we have used on the last line that $(1+x/n)^n = \exp(x)(1 + \mathcal{O}(n^{-1/2}x))$ for $x \leq n^{1/2-\epsilon_4}$ (cf. [KG09]). This is enough to show that the finite sum converges uniformly to zero at rate $\mathcal{O}(n^{2\eta-1/2+\epsilon+\epsilon_3})$ (at least if ϵ_3 is small enough). The first two terms in (6.47) decrease as the square root of this, that is $\mathcal{O}(n^{\eta-1/4+\epsilon/2+\epsilon_3/2})$.

Notice that the error 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 now pass to the third term of (6.47). 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 self-adjoint elements of $M(\mathbb{C}^2)$, then

$$e^{-i(X_1+X_2)} e^{iX_1} e^{iX_2} e^{[X_1, X_2]/2} = \mathbf{1} + \mathcal{O}(X_{i_1} X_{i_2} X_{i_3}), \quad (6.49)$$

where the $\mathcal{O}(\cdot)$ 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 $\mathbf{z}, \mathbf{u} \leq n^{\eta+\epsilon_3}$:

$$\begin{aligned} U(\beta) &:= U\left(-\frac{\mathbf{u}+\mathbf{v}}{\sqrt{n}}\right)U\left(\frac{\mathbf{u}}{\sqrt{n}}\right)U\left(\frac{\mathbf{v}}{\sqrt{n}}\right)\exp(i(u_x v_y - u_y v_x)\sigma_z/n) \\ &= \begin{bmatrix} 1 + \mathcal{O}(n^{-2+4\eta+4\epsilon_3}) & \mathcal{O}(n^{-3/2+3\eta+3\epsilon_3}) \\ \mathcal{O}(n^{-3/2+3\eta+3\epsilon_3}) & 1 + \mathcal{O}(n^{-2+4\eta+4\epsilon_3}) \end{bmatrix}. \end{aligned} \quad (6.50)$$

Finally, using the fact that $|j, j\rangle$ is an eigenvector of L_z , the third term in (6.47) can be written as

$$\| |j, j\rangle\langle j, j| - U_j(\beta)|j, j\rangle\langle j, j|U_j(\beta)^\dagger \|$$

and both states are pure, so it suffices to show that the scalar product converges to one uniformly. Using (6.50) and the expression of $\langle j|U_j(\beta)|j\rangle$ [HM06] we get, as $j \leq n$,

$$\langle j, j|U_j(\beta)|j, j\rangle = [U(\beta)_{1,1}]^j = 1 + \mathcal{O}(n^{-1+4\eta+4\epsilon_3}),$$

which implies that the third term in (6.47) is of order $\mathcal{O}(n^{-1+4\eta+4\epsilon_3})$. By choosing ϵ_3 and ϵ small enough, we obtain that all terms used in bounding (6.46) are uniformly $\mathcal{O}(n^{-1/4+\eta+\epsilon})$ for any $\epsilon > 0$.

This ends the proof of convergence (6.11) from the n qubit state to the oscillator.

A.2 The Map S_n

The opposite direction (6.12) does not require much additional estimation, so we 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 $\rho_n^{\mathbf{u}}$ by means of a completely positive map S_n , or at least come as close as possible.

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

$$\sup_{\|\mathbf{u}\| \leq n^\eta} \|q_n^{\mathbf{u}} - p_n^{\mathbf{u}}\|_1 = \mathcal{O}(n^{\eta-1/2}). \quad (6.51)$$

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 $\text{SU}(2)$:

$$\tau_{n,j}^{\mathbf{u}} := \left(V_j^\dagger \phi^{\mathbf{u}} V_j + \text{tr}(P_j^\perp \phi^{\mathbf{u}}) \mathbf{1} \right) \otimes \frac{\mathbf{1}}{n_j}.$$

The transformation $\phi^{\mathbf{u}} \mapsto \tau_{n,j}^{\mathbf{u}}$ is trace preserving and completely positive [GK06].

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 $\mathcal{O}(n^{\eta-1/2})$ (see (6.51)), the second term is $\mathcal{O}(n^{\epsilon-1/2})$ (see (6.40)). As for the third term, we use the triangle inequality to write, for $j \in \mathcal{J}_n$,

$$\|\tau_{n,j}^{\mathbf{u}} - \rho_{n,j}^{\mathbf{u}}\|_1 \leq \|\tau_{n,j}^{\mathbf{u}} - V_j^\dagger \phi^{\mathbf{u}} V_j^\dagger\|_1 + \|V_j^\dagger \phi^{\mathbf{u}} V_j^\dagger - \rho_{n,j}^{\mathbf{u}}\|_1.$$

The first term is $\mathcal{O}(e^{-n(1/2-\eta-2\delta)})$, according to the discussion following equation (6.44). The second term on the right is $\mathcal{O}(n^{-1/4+\eta+\epsilon})$ according to equations (6.45) through (6.50).

Summarizing, we have $\|S_n(N^{\mathbf{u}} \otimes \phi^{\mathbf{u}}) - \rho_n^{\mathbf{u}}\|_1 = \mathcal{O}(n^{-1/4+\eta+\epsilon})$, which establishes the proof in the inverse direction. \square

B Appendix: Proof of Theorem 6.3

We split the proof of theorem 6.3 in two steps.

B.1 First Estimate

We build up the state $\tilde{\rho}_{j,n}^{\mathbf{u}}$ 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

$$|\psi_{\mathbf{z},j}^n\rangle := \frac{1}{\|P_{\tilde{m}}|\mathbf{z}\rangle\|} \sum_{m=0}^{\tilde{m}} \frac{|\mathbf{z}|^m}{\sqrt{m!}} |m\rangle, \quad (6.52)$$

We mix the above states to obtain

$$\tilde{\rho}_{j,n}^{\mathbf{u}} := \frac{1}{\sqrt{2\pi s^2}} \int e^{-|\mathbf{z} - \sqrt{2\mu-1}\alpha_{\mathbf{u}}|^2/2s^2} (|\psi_{\mathbf{z},j}^n\rangle\langle\psi_{\mathbf{z},j}^n|) d^2\mathbf{z}.$$

Recall that $s^2 = (1-\mu)(4\mu-2)$, and

$$\phi^{\mathbf{u}} = \frac{1}{\sqrt{2\pi s^2}} \int e^{-|\mathbf{z} - \sqrt{2\mu-1}\alpha_{\mathbf{u}}|^2/2s^2} (|\mathbf{z}\rangle\langle\mathbf{z}|) d^2\mathbf{z}.$$

From the definition of $|\psi_{\mathbf{z},j}^n\rangle$ we have

$$\| |\psi_{\mathbf{z},j}^n\rangle - |\mathbf{z}\rangle \| \leq \sqrt{2} \frac{|\mathbf{z}|^{\tilde{m}}}{\sqrt{\tilde{m}!}} \wedge 2, \quad (6.53)$$

which implies that $\|\tilde{\rho}_{j,n}^{\mathbf{u}} - \phi^{\mathbf{u}}\|_1 = \mathcal{O}(e^{-n^{2(\eta+\epsilon)}})$ for any $\epsilon > 0$, for any $\gamma \geq 2(\eta+\epsilon)$. Indeed, we bound it by

$$\frac{\sqrt{2}}{\sqrt{\pi s^2}} \int e^{-|\mathbf{z}|^2/2s^2} \left(\frac{|\mathbf{z} + \sqrt{2\mu - 1}\alpha_{\mathbf{u}}|^{\tilde{m}}}{\sqrt{\tilde{m}!}} \wedge \sqrt{2} \right) d^2\mathbf{z},$$

and 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 $\mathcal{O}(e^{-n^{2(\eta+\epsilon)}})$. The integral over the disk $|\mathbf{z}| \leq n^{\eta+\epsilon}$ is bounded by the supremum of (6.53) since the Gaussian integrates to one, and is $\mathcal{O}(e^{-(\gamma/2 - \eta - \epsilon)n^\gamma})$. 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 $\epsilon_2 > 0$.

B.2 Second Estimate

We now compare the evolved qubits state $\tilde{\rho}_{j,n}^{\mathbf{u}}(t)$ and the evolved oscillator state $\phi^{\mathbf{u}}(t)$. Let $|\psi_{m,j}^n(t)\rangle = U_{j,n}(t) |m\rangle \otimes |\Omega\rangle$ 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 $|\psi_{m,j}^n(t)\rangle$,

$$|\xi_{m,j}^n(t)\rangle := \sum_{i=0}^m c_n(m, i) \alpha_i(t) |m-i\rangle \otimes |e^{-1/2u} \chi_{[0,t]}(u)\rangle_i, \quad (6.54)$$

where $\alpha_i(t) = \exp((-m+i)t/2)$, and we define the numbers $c_n(m, i)$ inductively by $c_n(m, 0) := 1$ and

$$c_n(m, i) := c_n(m, i-1) \sqrt{\frac{2j-m+i}{2j_n}} \sqrt{\frac{m-i+1}{i}}.$$

The n -particle vector $|f\rangle_n := f^{\otimes n}$ was defined in (6.17). Note that in particular, for $\mu - 1/2 > \epsilon_2$ and $j \in \mathcal{J}_n$, we have

$$c_n(m, i) \leq \sqrt{\binom{m}{i}} \left(1 + \frac{2}{\epsilon_2} n^{-1/2+\epsilon}\right)^i.$$

We now apply the estimate (6.21). By direct computations, we get

$$\begin{aligned} d|\xi_{m,j}^n(t)\rangle &= -\frac{1}{2} \sum_{i=0}^m c_n(m, i) \alpha_i(t) (m-i) |m-i\rangle \otimes |e^{-1/2u} \chi_{[0,t]}(u)\rangle_i dt \\ &\quad + \sum_{i=1}^m c_n(m, i) \alpha_{i-1}(t) |m-i\rangle \otimes |e^{-1/2u} \chi_{[0,t]}(u)\rangle_{i-1} \otimes_s |\chi_{[t, t+dt]}\rangle, \end{aligned} \quad (6.55)$$

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 obtain

$$\begin{aligned} G_{dt} |\xi_{m,j}^n(t)\rangle &= \sum_{i=0}^m c_n(m, i) \alpha_i(t) \sqrt{\frac{(m-i)(2j-m+i+1)}{2j_n(i+1)}} \times \\ &\quad |m-i-1\rangle \otimes |e^{-1/2u} \chi_{[0,t]}(u)\rangle_i \otimes_s |\chi_{[t,t+dt]}\rangle \\ &- \frac{1}{2} \sum_{i=0}^m c_n(m, i) \alpha_i(t) (m-i) \frac{2j-m+i+1}{2j_n} |m-i\rangle \otimes |e^{-1/2u} \chi_{[0,t]}(u)\rangle_i dt. \end{aligned}$$

In the first term of the right hand side of the above equation, 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 (6.55). Thus

$$\begin{aligned} G_{dt} |\xi_{m,j}^n(t)\rangle - d|\xi_{m,j}^n(t)\rangle &= \\ \frac{1}{2} \sum_{i=0}^{m-1} c_n(m, i) \alpha_i(t) (m-i) \frac{2(j_n-j)+m-i-1}{2j_n} |m-i\rangle &\otimes |e^{-1/2u} \chi_{[0,t]}(u)\rangle_i dt. \end{aligned}$$

Then, using $c_n(m, i) \leq \sqrt{\binom{m}{i} (1 + (2/\epsilon_2) n^{-1/2+\epsilon})^i}$, we find that the expression $\|G_{dt} \xi_{m,j}^n(t) - d\xi_{m,j}^n(t)\|$ is bounded from above by

$$\frac{1}{2} \left[\sum_{i=0}^{m-1} \binom{m}{i} ((1 + n^{-1/2+\epsilon})(1 - e^{-t}))^i e^{-(m-i)t} \left(\frac{(2(j_n-j)+m-i-1)(m-i)}{2j_n} \right)^2 \right]^{1/2} dt.$$

We have

$$\frac{(2(j_n-j)+m-i-1)(m-i)}{2j_n} = \mathcal{O}(m(n^{-1/2+\epsilon} + n^{-1}m)).$$

Inside the sum, we recognize the binomial terms with the m^{th} term missing. Thus the sum is

$$\begin{aligned} &\left(1 + n^{-1/2+\epsilon} - e^{-t} n^{-1/2+\epsilon}\right)^m - \left((1 - e^{-t})(1 + n^{-1/2+\epsilon})\right)^m \\ &\leq (1 + n^{-1/2+\epsilon})^m (1 - (1 - e^{-t})^m) \leq (1 + n^{-1/2+\epsilon})^m m e^{-t}. \end{aligned}$$

Then there exists a constant C (independent of μ if $\mu - 1/2 \geq \epsilon_2$) such that

$$\|G_{dt} \xi_{m,j}^n(t) - d\xi_{m,j}^n(t)\| \leq \frac{C}{2} e^{-t/2} m^{3/2} (n^{-1/2+\epsilon} + mn^{-1}) \left(1 + \frac{2}{\epsilon_2} n^{-1/2+\epsilon}\right)^{m/2}$$

By integrating over t , we finally obtain

$$\|\psi_{m,j}^n(t) - \xi_{m,j}^n(t)\| \leq C m^{3/2} (n^{-1/2+\epsilon} + mn^{-1}) \left(1 + \frac{2}{\epsilon_2} n^{-1/2+\epsilon}\right)^{m/2}. \quad (6.56)$$

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 on γ are

$$2\eta + 2\epsilon < \gamma < 1/3 - 2\epsilon/3.$$

Now consider the vector $|\psi_{\mathbf{z},j}^n\rangle$ as defined in (6.52), and introduce the notation $|\psi_{\mathbf{z},j}^n(t)\rangle = U_{j,n}(t)|\psi_{\mathbf{z},j}^n\rangle \otimes |\Omega\rangle$. Then based on (6.54), we choose the approximate solution

$$|\xi_{\mathbf{z},j}^n(t)\rangle = e^{-|\mathbf{z}|^2/2} \sum_{m=0}^{\tilde{m}} \frac{|\mathbf{z}|^m}{\sqrt{m!}} \sum_{i=0}^m c_n(m,i) \alpha_i(t) |m-i\rangle \otimes |e^{-1/2u} \chi_{[0,t]}(u)\rangle_i.$$

Note that the vectors $|\psi_{k,j}^n(t)\rangle$ and $|\xi_{k,j}^n(t)\rangle$ live in the “ k -particle” subspace of $\mathcal{H}_j \otimes \mathcal{F}(L^2(\mathbb{R}))$ and thus are orthogonal to all vectors $|\psi_{p,j}^n(t)\rangle$ and $|\xi_{p,j}^n(t)\rangle$ with $p \neq k$. By (6.56), the error is

$$\begin{aligned} & \left\| |\psi_{\mathbf{z},j}^n(t)\rangle - |\xi_{\mathbf{z},j}^n(t)\rangle \right\| \leq \\ & C e^{-|\mathbf{z}|^2/2} \left(\sum_{m=0}^{\tilde{m}} \frac{|\mathbf{z}|^{2m}}{m!} m^3 (n^{-1/2+\epsilon} + mn^{-1})^2 \left(1 + \frac{2}{\epsilon_2} n^{-1/2+\epsilon}\right)^m \right)^{1/2} + \frac{|\mathbf{z}|^{2\tilde{m}}}{\tilde{m}!} \leq \\ & C \tilde{m}^{3/2} (n^{-1/2+\epsilon} + \tilde{m}n^{-1}) \left(1 + \frac{2}{\epsilon_2} n^{-1/2+\epsilon}\right)^{\tilde{m}/2} + \frac{|\mathbf{z}|^{2\tilde{m}}}{\tilde{m}!}. \end{aligned} \tag{6.57}$$

We now compare the approximate solution $|\xi_{\mathbf{z},j}^n(t)\rangle$ with the “limit” solution $|\psi_{\mathbf{z}}(t)\rangle$ for the oscillator coupled with the field as described in section 6.4.2. We can write

$$|\psi_{\mathbf{z}}(t)\rangle = e^{-|\mathbf{z}|^2/2} \sum_{m=0}^{\infty} \frac{|\mathbf{z}|^m}{\sqrt{m!}} \sum_{i=0}^m \sqrt{\binom{m}{i}} e^{-(m-i)t/2} |m-i\rangle \otimes |e^{-1/2u} \chi_{[0,t]}(u)\rangle_i.$$

Then

$$\begin{aligned} & \left\| |\xi_{\mathbf{z},j}^n(t)\rangle - |\psi_{\mathbf{z}}(t)\rangle \right\|^2 = e^{-|\mathbf{z}|^2} \sum_{m=\tilde{m}}^{\infty} \frac{|\mathbf{z}|^{2m}}{m!} \\ & + e^{-|\mathbf{z}|^2} \sum_{m=0}^{\tilde{m}} \frac{|\mathbf{z}|^{2m}}{m!} \sum_{i=0}^m e^{-(m-i)t} \left| c_n(m,i) - \sqrt{\binom{m}{i}} \right|^2 (1 - e^{-t})^i. \end{aligned}$$

Now

$$\begin{aligned} \left| c_n(m,i) - \sqrt{\binom{m}{i}} \right|^2 & \leq |c_n(m,i)^2 - \binom{m}{i}| \\ & \leq \binom{m}{i} \left| 1 - \prod_{p=1}^i \left(1 + \frac{2(j-j_n)-m+p}{2j_n} \right) \right| \\ & \leq C_2 \binom{m}{i} m n^{-1/2+\epsilon}, \end{aligned}$$

where C_2 does not depend on μ as long as $\mu - 1/2 \geq \epsilon_2$ (recall that the dependence on μ is hidden in $j_n = (2\mu - 1)n$). Thus

$$\begin{aligned} \|\xi_{\mathbf{z},j}^n(t) - \psi_{\mathbf{z}}(t)\|^2 &\leq C_2 n^{-1/2+\epsilon} e^{-|\mathbf{z}|^2} \sum_{m=0}^{\tilde{m}} \frac{m|\mathbf{z}|^{2m}}{m!} + \frac{|\mathbf{z}|^{2\tilde{m}}}{\tilde{m}!} \\ &\leq C_2 n^{-1/2+\epsilon} |\mathbf{z}|^2 + \frac{|\mathbf{z}|^{2\tilde{m}}}{\tilde{m}!}. \end{aligned} \quad (6.58)$$

From (6.57) and (6.58) we get

$$\begin{aligned} \|\psi_{\mathbf{z},j}^n(t) - \psi_{\mathbf{z}}(t)\| &\leq 2 \wedge \left[C\tilde{m}^{3/2}(n^{-1/2+\epsilon} + \tilde{m}n^{-1}) \left(1 + \frac{2}{\epsilon_2} n^{-1/2+\epsilon}\right)^{\tilde{m}/2} \right. \\ &\quad \left. + \frac{|\mathbf{z}|^{2\tilde{m}}}{\tilde{m}!} + \left[C_2 n^{-1/2+\epsilon} |\mathbf{z}|^2 + \frac{|\mathbf{z}|^{2\tilde{m}}}{\tilde{m}!} \right]^{1/2} \right], \end{aligned}$$

and we denote the r.h.s. by $E(\tilde{m}, n, \mathbf{z})$. 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

$$\tilde{\rho}_{j,n}^{\mathbf{u}} := \frac{1}{\sqrt{2\pi s^2}} \int e^{-|\mathbf{z} - \sqrt{2\mu-1}\alpha_{\mathbf{u}}|^2/2s^2} (|\psi_{\mathbf{z},j}^n\rangle\langle\psi_{\mathbf{z},j}^n|) d^2\mathbf{z}.$$

If we define the evolved states

$$\tilde{\rho}_{j,n}^{\mathbf{u}}(t) := U_{j,n}(t)\tilde{\rho}_{j,n}^{\mathbf{u}}U_{j,n}(t)^\dagger \quad \text{and} \quad \phi^{\mathbf{u}}(t) := U(t)\phi^{\mathbf{u}}U(t)^\dagger,$$

then $\sup_{j \in \mathcal{J}_n} \sup_{\|\mathbf{u}\| \leq n^\eta} \|\tilde{\rho}_{j,n}^{\mathbf{u}}(t) - \phi^{\mathbf{u}}(t)\|_1$ is bounded by

$$\sup_{\|\mathbf{u}\| \leq n^\eta} \frac{1}{\sqrt{\pi s^2}} \int e^{-|\mathbf{z} - \sqrt{2\mu-1}\alpha_{\mathbf{u}}|^2/2s^2} E(\tilde{m}, n, \mathbf{z}) d^2\mathbf{z}.$$

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 $e^{-n^{\eta+\epsilon}}$. The inner part is then dominated by $\sup_{|\mathbf{z}| \leq n^{\eta+\epsilon}} E(\tilde{m}, n, \mathbf{z})$. Now on the one hand, we want (6.57) to be small, so that \tilde{m} cannot be too big. 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

$$E(\tilde{m}, n, \mathbf{z}) = \mathcal{O}(n^{\eta-1/4+\epsilon}, n^{3\eta-1/2+\epsilon}),$$

for any small enough $\epsilon > 0$. Hence we obtain (6.22). \square

Chapter 7

Bundles with a Lift of Infinitesimal Diffeomorphisms

This chapter is not about algebra and quantum mechanics, but about differential geometry and classical field theory. We slightly extend the notion of a natural fibre bundle by requiring diffeomorphisms of the base to lift to automorphisms of the bundle only infinitesimally, i.e. at the level of the Lie algebra of vector fields. Spin structures are natural only in this extended sense. We classify the fibre bundles with this property, under the additional assumption of a finite dimensional structure group. This includes all spin structures, but only some Spin^c and SpinG -structures. This classification links the gauge group G to the topology of space-time.

Remark Many of the results in this chapter, including the main theorem 7.14, can be found in Lecomte's paper [Lec85]. Our results were obtained independently.

7.1 Introduction

All bundles are equal, but some are more natural than others. For example, the naturality of tangent bundles, frame bundles and jet bundles is beyond all question. A bundle is called 'natural' if diffeomorphisms of the base lift to automorphisms of the bundle in a local fashion.

Let us phrase this more carefully. Any principal fibre bundle $\pi : P \rightarrow M$ determines an exact sequence of groups

$$1 \rightarrow \Gamma_c(\text{Ad}(P)) \rightarrow \text{Aut}_c(P) \rightarrow \text{Diff}_c^P(M) \rightarrow 1. \quad (7.1)$$

In this expression, $\text{Aut}_c(P)$ is the group of automorphisms of P that are trivial outside a compact subset of M . Its image $\text{Diff}_c^P(M)$, the group of compactly supported liftable diffeomorphisms of M , is an open subgroup of $\text{Diff}_c(M)$. In particular, it is big enough to contain its connected component of unity. We have identified the gauge group of vertical automorphisms with the group

of sections of the adjoint bundle $\text{Ad}(P)$, and $\Gamma_c(\text{Ad}(P))$ is the subgroup of compactly supported ones.

Definition 7.1 *A ‘natural’ principal fibre bundle is a principal fibre bundle P for which all diffeomorphisms are liftable and (7.1) is split, together with a distinguished splitting homomorphism $\Sigma : \text{Diff}_c(M) \rightarrow \text{Aut}_c(P)$. Moreover, the splitting homomorphism Σ is required to be local in the sense that it should lift each diffeomorphism $\phi : U_1 \rightarrow U_2$ between open subsets of M to a bundle isomorphism $\Sigma(\phi) : \pi^{-1}(U_1) \rightarrow \pi^{-1}(U_2)$ in a functorial fashion.*

Natural bundles have been classified. A theorem of Salviooli, Palais, Terng, Epstein and Thurston [Sal72, PT77, ET79] states that any natural fibre bundle is associated to the k^{th} order frame bundle $F^k(M)$.

In this chapter, we seek to extend the notion of a natural fibre bundle in two separate ways. First of all, we do not require locality, but prove it. And secondly, we only require diffeomorphisms of the base to lift to automorphisms of the bundle infinitesimally, i.e. at the level of Lie algebras. We will call the principal fibre bundles which are natural in this extended sense ‘infinitesimally natural’, and our main objective will be their classification.

Let us formulate this more precisely. For any principal fibre bundle, the sequence of groups (7.1) gives rise to the exact sequence of Lie algebras

$$0 \rightarrow \Gamma_c(\text{ad}(P)) \rightarrow \Gamma_c(TP)^G \rightarrow \Gamma_c(TM) \rightarrow 0, \quad (7.2)$$

where ‘ c ’ stands for ‘0 outside a compact subset of M ’.

Definition 7.2 *An ‘infinitesimally natural’ principal fibre bundle is a principal fibre bundle P , together with a distinguished Lie algebras homomorphism $\sigma : \Gamma_c(TM) \rightarrow \Gamma_c(TP)^G$ that splits the exact sequence of Lie algebras (7.2).*

We emphasize that we will not require σ to be local, continuous or $C^\infty(M)$ -linear, and it need not come from a map of bundles. We only require σ to be a homomorphism of Lie algebras.

The price we pay for this level of generality is that we must confine ourselves to fibre bundles with a finite dimensional structure group. This allows us to first study principal fibre bundles, and then generalize the results to more general fibre bundles with a finite dimensional structure group, such as vector bundles.

There are three main reasons for wanting to extend the notion of a natural bundle. First of all, natural bundles are meant to describe geometric objects, but not all geometry is local. One might feel that the locality requirement on Σ is therefore inappropriate. The universal cover $\tilde{M} \rightarrow M$ for instance is not a natural bundle, although its ties with the global topology of M are unmistakable. It does however allow for a lift of infinitesimal diffeomorphisms, and is therefore infinitesimally natural.

The second reason comes from field theory. Fields are described by sections of a fibre bundle F over space-time M , where F is associated to some principal fibre bundle P . Fermions, for example, are sections of a spinor bundle F , associated to a spin structure P . As spinors transform according to a projective

representation of the Lorentz group, they acquire a minus sign upon a full rotation. This means that there is no hope of lifting global diffeomorphisms, so that spinor bundles cannot be natural. One might feel however that it would be unfair to discard them as unnatural. After all, they do occur in nature, or at least in our mathematical description of it. As one can lift diffeomorphisms infinitesimally, i.e. at the level of vector fields, a spinor bundle does constitute an infinitesimally natural bundle.

The third reason is that in field theory, one would like to construct a stress-energy-momentum tensor, corresponding to infinitesimal transformations of space-time M . Unfortunately, $\text{Diff}(M)$ does not act on the fields $\Gamma(E)$ directly. Since $\text{Aut}(P)$ does, we need an infinitesimal lift of $\text{Diff}(M)$ into $\text{Aut}(P)$, i.e. a Lie algebra homomorphism σ that splits (7.2). This means that an infinitesimally natural bundle is precisely what one needs in order to construct a SEM-tensor from Noether's principle. (See chapter 2, or e.g. [GM92, FR04].) The splitting σ is to be interpreted as the transformation behaviour of the fields under infinitesimal space-time transformations.

The outline of the chapter is as follows. Sections 7.2, 7.3 and 7.4 are devoted to the classification of infinitesimally natural principal fibre bundles. The central result is theorem 7.14, which states the following.

Any infinitesimally natural principal fibre bundle is associated to the universal cover of the connected component of the k^{th} order frame bundle $\tilde{F}^{+k}(M)$.

We extend this to fibre bundles with a finite dimensional structure group in section 7.5, with special attention for vector bundles. In section 7.6, we seek conditions under which a splitting of (7.2) gives rise to a flat connection.

Finally, in section 7.7, we study spin structures in the presence of a gauge field. Ordinary spin structures are infinitesimally natural, as they carry a canonical splitting of (7.2). In the presence of a gauge group G however, one should consider Spin^G -structures instead (see [HP78], [AI80]). For example, if $G = U(1)$, the appropriate bundles are Spin^c -structures. In contrast to ordinary spin structures, not all Spin^G -structures are infinitesimally natural. We show that for a compact gauge group, infinitesimally natural Spin^G -structures correspond precisely to homomorphisms $\pi_1(F(M)) \rightarrow G$ that are injective on $\pi_1(\text{GL}(\mathbb{R}^n))$.

In a sense, this shows that complicated manifolds call for complicated gauge groups. Some manifolds, such as CP^2 , do not allow for any infinitesimally natural Spin^G -structure at all. If one feels that fermions should have a well defined transformation behaviour under infinitesimal space-time transformations, then these manifolds are disqualified as models for space-time.

7.2 Principal Bundles as Lie Algebra Extensions

We seek to classify infinitesimally natural principal fibre bundles. It is to this end that we study Lie algebra homomorphisms σ that split (7.2). In this section, we will prove that σ must be a differential operator of finite order.

The first step, to be taken in section 7.2.1, is to show that maximal ideals in $\Gamma_c(TM)$ correspond precisely to points in M . Using this, we will see in section 7.2.2 that σ must be a local map. Once again, we stress that locality of σ will be a theorem, not an assumption. We will then prove, in section 7.2.3, that σ is in fact a differential operator.

Before we proceed, let us have a closer look at the exact sequence of Lie algebras (7.2), derived from a smooth principal G -bundle $\pi : P \rightarrow M$. The last term, $\Gamma_c(TM)$, is the Lie algebra of smooth, compactly supported vector fields on M . The middle term, $\Gamma_c(TP)^G$, is the Lie algebra of G -invariant vector fields v on P such that $\pi(\text{Supp}(v))$ is compact. The pushforward defines a Lie algebra homomorphism $\pi_* : \Gamma_c(TP)^G \rightarrow \Gamma_c(TM)$ because of G -invariance, and its kernel $\Gamma_c(TP)_v^G$ is the ideal of vertical vector fields. It can be identified with $\Gamma_c(\text{ad}(P))$, the Lie algebra of compactly supported sections of the adjoint bundle $\text{ad}(P) := P \times_G \mathfrak{g}$, where \mathfrak{g} is the Lie algebra of G .

7.2.1 Ideals of the Lie Algebra of Vector Fields

The following lemma, due to Shanks and Pursell [SP54], constitutes the linchpin of the proof. It identifies the maximal ideals of $\Gamma_c(TM)$ as points in M . The proof is taken from [SP54], with perhaps some minor clarifications.

Lemma 7.1 *Let $\Gamma_c(TM)$ be the Lie algebra of smooth compactly supported vector fields on M . Then the maximal ideals of $\Gamma_c(TM)$ are labelled by points q in M , I_q being the ideal of vector fields in $\Gamma_c(TM)$ which are zero and flat at q . That is, $I_q = \{v \in \Gamma_c(TM) \mid v(q) = 0 \text{ and } (\text{ad}(w_{i_1}) \dots \text{ad}(w_{i_n})v)(q) = 0 \forall w_{i_1}, \dots, w_{i_n}, \in \Gamma_c(TM)\}$.*

Proof. Let $I \subseteq \Gamma_c(TM)$ be an ideal. Suppose that there exists a point $q \in M$ such that $v(q) = 0$ for all $v \in I$. Then $I \subseteq I_q$. Indeed, for all $w_{i_1}, \dots, w_{i_n} \in \Gamma_c(TM)$, one has $(\text{ad}(w_{i_1}) \dots \text{ad}(w_{i_n})v)(q) = 0$ because the l.h.s. is in I .

Now suppose that I does not have such a point. We will prove that this implies $I = \Gamma_c(TM)$. From this, the lemma will follow. If I is a maximal ideal, then surely $I \neq \Gamma_c(TM)$, so we must have $I \subseteq I_{\bar{q}}$. But it is clear from its definition that $I_{\bar{q}}$ is an ideal, so, I being maximal, we must have $I = I_{\bar{q}}$. Conversely, the ideals I_q are all maximal, because any enveloping ideal must either be all of $\Gamma_c(TM)$, or else be contained some $I_{\bar{q}}$, which cannot be unless $\bar{q} = q$.

We are therefore left to prove the following: suppose that I is an ideal such that for all $q \in M$, there exists a $v \in I$ with $v(q) \neq 0$. Then $I = \Gamma_c(TM)$.

We will shortly prove the following statement: for all $q \in M$, there exists a neighbourhood U_q such that all $w \in \Gamma_c(TM)$ with $\text{Supp}(w) \subset U_q$ can be written as $w = [v, u]$ with $v \in I$ and $u \in \Gamma_c(TM)$.

Such w are therefore elements of I . Membership of I then extends to arbitrary $w \in \Gamma_c(TM)$ by covering its support with a finite amount of neighbourhoods $U_{q_1} \dots U_{q_N}$, for which we construct a partition of unity $\sum_{s=1}^N f_s = 1$, $\text{Supp}(f_s) \subseteq U_{q_s}$. We may then write $w = \sum_{s=1}^N f_s w = \sum_{s=1}^N [v_s, u_s] \in I$. This concludes the proof modulo lemma 7.2, which is next in line. \square

Although one can easily find $u \in \Gamma_c(TM)$ and $v \in I$ such that $w = [u, v]$ locally, it is not always clear how to extend these to global vector fields while simultaneously satisfying $w = [u, v]$. For instance, with M the circle S^1 and $v = w = \partial_\theta$, the solution $u = \theta \partial_\theta$ does not globally exist. We need to do some work in order to define a proper cutoff procedure.

Lemma 7.2 *Let I be a maximal ideal of $\Gamma_c(TM)$ containing a vector field which does not vanish at $q \in M$. Then q has a neighbourhood U_q such that all $w \in \Gamma_c(TM)$ with $\text{Supp}(w) \subset U_q$ can be written $w = [v, u]$ with $v \in I$ and $u \in \Gamma_c(TM)$.*

Proof. Choose $v \in I$ with $v(q) \neq 0$. There exist local co-ordinates x_1, \dots, x_n and an open neighbourhood W of q such that $v|_W = \partial_1$. Choose W to be a block centered around q , and nest two smaller blocks (in local co-ordinates) inside, so that $W \supset V \supset U$.

We take $W = (-\varepsilon, \varepsilon)^n$, $V = (-\frac{2}{3}\varepsilon, \frac{2}{3}\varepsilon)^n$, and $U = (-\frac{1}{3}\varepsilon, \frac{1}{3}\varepsilon)^n$. Choose a smooth function g on M such that $g|_U = x_1$ and $g|_{M \setminus V} = 0$. Require also that $\partial_i g(x) = 0$ for $i \neq 1$ and $(x_2, \dots, x_n) \in (-\frac{1}{3}\varepsilon, \frac{1}{3}\varepsilon)^{n-1}$. Define $h := \partial_1 g$, and set $\tilde{v} := [v, g\partial_1] = h\partial_1$. Then $\tilde{v}|_U = \partial_1$, $\tilde{v}|_{M \setminus V} = 0$, and most importantly $\tilde{v} \in I$.

Now let $w \in \Gamma_c(TM)$ with $\text{Supp}(w) \subset U$. We will find a global vector field $u \in \Gamma_c(TM)$ that realizes $[\tilde{v}, u] = w$.

For $i \neq 1$, the i^{th} component of the above reads $h\partial_1 u^i = w^i$. In the region $x_1 \leq \frac{1}{3}\varepsilon$, we set $u^i(x_1, \dots, x_n) = \int_{-\infty}^{x_1} w^i(t, x_2, \dots, x_n) dt$. On U , where $\tilde{v} = \partial_1$, we then have $\partial_1 u^i = w^i$ and therefore $h\partial_1 u^i = w^i$. This is obviously also correct for points outside U with $x_1 \leq \frac{1}{3}\varepsilon$, as both u^i and w^i are zero.

For $\frac{1}{3}\varepsilon \leq x_1 \leq \frac{2}{3}\varepsilon$, where $w = 0$, we let $u^i(x_1, x_2, \dots, x_n)$ be a constant function of x_1 , so that $u^i(x_1, x_2, \dots, x_n) = u^i(\frac{1}{3}\varepsilon, x_2, \dots, x_n)$. We then have $\partial_1 u^i(x_1, \dots, x_n) = 0$, guaranteeing $h\partial_1 u^i = w^i$. Note that this does not effect the smoothness of u^i .

Finally, for $\frac{2}{3}\varepsilon \leq x_1 \leq \varepsilon$, let u^i tend to zero, and let u^i be zero for $x_1 \geq \varepsilon$. This can be done in such a way that u^i remains smooth. Since both h and w are zero, we have $h\partial_1 u^i = w^i$ on all of M .

The case $i = 1$ is handled similarly, the only difference being that the first component of $[\tilde{v}, u] = w$ is now $h\partial_1 u^1 - \sum_j u^j \partial_j h = w^1$, the term $u^1 \partial_1 h$ of which cannot be dispensed with. We have arranged for $u^i(x)$, with $i \neq 0$, to be zero if $(x_2, \dots, x_n) \notin (-\frac{1}{3}\varepsilon, \frac{1}{3}\varepsilon)^{n-1}$, so that $\partial_i h$ equals zero if u^i does not. This leaves us with $h\partial_1 u^1 - u^1 \partial_1 h = w^1$.

For $x_1 \leq \frac{1}{3}\varepsilon$, one once again sets $u^1(x_1, \dots, x_n) = \int_{-\infty}^{x_1} w^1(t, x_2, \dots, x_n) dt$. For $\frac{1}{3}\varepsilon \leq x_1 \leq \frac{2}{3}\varepsilon$ however, one now has to define $u^1(x_1, \dots, x_n) = h(x_1, \dots, x_n) u^1(\frac{1}{3}\varepsilon, x_2, \dots, x_n)$ in order for $h\partial_1 u^1 - u^1 \partial_1 h = w^1$ to hold. Since this renders u^1 zero on the boundary of V , one is then free to define u^i to be zero on $M \setminus V$.

Thus, if $\text{Supp}(w) \subset U$, we see that $w = [\tilde{v}, u]$ with $\tilde{v} \in I$ and $u \in \Gamma_c(TM)$. This concludes the proof lemma 7.2, and thereby that of lemma 7.1. \square

A maximal subalgebra A of a Lie algebra L is either self-normalizing or ideal. Indeed it is contained in its normalizer, which therefore equals either A or L . A theorem of Barnes [Bar62] states that a finite-dimensional Lie algebra is

nilpotent if and only if¹ every maximal subalgebra is an ideal. On the other extreme:

Proposition 7.3 *Let L be a Lie algebra over a field \mathbf{K} , and let \mathcal{S} be the set of subspaces $A \subset L$ such that A is both an ideal and a maximal subalgebra. Then*

$$[L, L] = \bigcap_{A \in \mathcal{S}} A.$$

In particular, since the r.h.s reads ‘ L ’ if $\mathcal{S} = \emptyset$, L is perfect ($[L, L] = L$) if and only if every maximal subalgebra is self-normalizing.

Proof. Let $X \notin [L, L]$. Then choose $[L, L] \subseteq A \subsetneq L$ where A has codimension 1 in L , and $X \notin A$. A is an ideal maximal subalgebra, which does not contain X . Thus $\bigcap_{A \in \mathcal{S}} A \subseteq [L, L]$.

Let A be an ideal maximal subalgebra, and $X \notin A$. Then $A + \mathbf{K}X$ is a subalgebra strictly containing A , so that it must equal L . Thus $[L, L] = [A + \mathbf{K}X, A + \mathbf{K}X] \subseteq A$, whence $[L, L] \subseteq \bigcap_{A \in \mathcal{S}} A$. \square

As a corollary, we have the following well known statement:

Corollary 7.4 *The Lie algebra $\Gamma_c(TM)$ is perfect;*

$$[\Gamma_c(TM), \Gamma_c(TM)] = \Gamma_c(TM).$$

Proof. According to lemma 7.1, the maximal ideals are precisely the ideals I_q of vector fields in $\Gamma_c(TM)$ which are zero and flat at q . I_q is strictly contained in the subalgebra A_q of vector fields which are zero at q , so that no ideal is a maximal subalgebra. So every maximal subalgebra is self-normalizing. \square

7.2.2 The Splitting as a Local Map

With the main technical obstacles out of the way, we turn our attention to the sequence (7.2). We now prove that σ is a local map.

Lemma 7.5 *Let $P \rightarrow M$ be a principal G -bundle over M , with G any Lie group. Let $\sigma : \Gamma_c(TM) \rightarrow \Gamma_c(TP)^G$ be a Lie algebra homomorphism splitting the exact sequence of Lie-algebras*

$$0 \rightarrow \Gamma_c(TP)_v^G \rightarrow \Gamma_c(TP)^G \rightarrow \Gamma_c(TM) \rightarrow 0. \quad (7.3)$$

Then σ is local in the sense that $\pi(\text{Supp}(\sigma(v))) \subseteq \text{Supp}(v)$.

Proof. Any principal fibre bundle possesses an equivariant connection 1-form $\omega \in \Omega^1(P, \mathfrak{g})$, which enables one to lift vector fields. More precisely, the lifting map $\gamma : \Gamma_c(TM) \rightarrow \Gamma_c(TP)^G$ is defined by the requirement that $\omega_p(\gamma(v_{\pi(p)}))$ equal zero for all $p \in P$, and that $\pi_* \circ \gamma$ be the identity. The lifting map γ splits

¹Actually, the ‘only if’ part in Barnes’ theorem is not written down in [Bar62], but this is immediately clear from the proof of Engel’s theorem. (See e.g. [Hum72]).

the exact sequence (7.3) as a sequence of $C^\infty(M)$ -modules, but generally *not* as a sequence of Lie algebras, since it need not be a homomorphism.

Define $f := \sigma - \gamma$. Then $\pi_* \circ f = 0$, so that f is a map $\Gamma_c(TM) \rightarrow \Gamma_c(TP)_v^G$. Since γ is local by definition, the lemma will follow if we show f to be local.

The homomorphism property of σ , $[\sigma(v), \sigma(w)] - \sigma([v, w]) = 0$, translates to

$$\begin{aligned} f([v, w]) - [f(v), f(w)] = \\ [\gamma(v), f(w)] - [\gamma(w), f(v)] + [\gamma(v), \gamma(w)] - \gamma([v, w]). \end{aligned} \quad (7.4)$$

The G -action effects a Lie algebra isomorphism $\Gamma_c(TP)_v^G \simeq \Gamma_c(\text{ad}(P))$. From that perspective, we have for $s \in \Gamma_c(\text{ad}(P))$ that $[\gamma(v), s] = \nabla_v s$, the covariant derivative along v . The curvature of the connection is then given by $R(v, w) = [\gamma(v), \gamma(w)] - \gamma([v, w])$. We rewrite equation (7.4) as

$$f([v, w]) - [f(v), f(w)] = \nabla_v f(w) - \nabla_w f(v) + R(v, w). \quad (7.5)$$

Pick $m \in M$, and identify the fibre of $\text{ad}(P)$ over m with the structure Lie algebra \mathfrak{g} . The restriction $r_m : \Gamma_c(\text{ad}(P)) \rightarrow \mathfrak{g}$ is a homomorphism of Lie algebras. Now define $A_m := \{v \in \Gamma_c(TM) \mid v(m) = 0\}$ to be the maximal subalgebra of vector fields that vanish at m , and consider the map $\tilde{f}_m : A_m \rightarrow \mathfrak{g}$ defined by $\tilde{f}_m = r_m \circ f|_{A_m}$. Equation (7.5) transforms into

$$\tilde{f}_m([v, w]) - [\tilde{f}_m(v), \tilde{f}_m(w)] = \nabla_v(\tilde{f}_m(w))_m - \nabla_w(\tilde{f}_m(v))_m + R(v, w)_m, \quad (7.6)$$

the r.h.s. of which vanishes because ∇ and R are $C^\infty(M)$ -linear in v and w , and $v(m) = w(m) = 0$. This means that \tilde{f}_m is a Lie algebra homomorphism.

Let us restrict \tilde{f}_m even further to the ideal $\Gamma_c(T(M - \{m\})) \subset A_m$ (which is non-maximal), and note that $\hat{f}_m : \Gamma_c(T(M - \{m\})) \rightarrow \mathfrak{g}$ is a homomorphism. Its kernel $\text{Ker}(\hat{f}_m)$ is therefore an ideal in $\Gamma_c(T(M - \{m\}))$, and one of finite codimension at that. Indeed,

$$\Gamma_c(T(M - \{m\}))/\text{Ker}(\hat{f}_m) \simeq \text{Im}(\hat{f}_m) \subseteq \mathfrak{g}.$$

According to lemma 7.1 however, all proper ideals are of infinite codimension, forcing $\text{Ker}(\hat{f}_m) = \Gamma_c(T(M - \{m\}))$. But the vanishing of \hat{f}_m for all $m \in M$ is tantamount to locality of f in the sense that $\text{Supp}(f(v)) \subseteq \text{Supp}(v)$. \square

We gather some notation that we introduced in the course of the proof. If $\sigma : \Gamma_c(TM) \rightarrow \Gamma_c(TP)^G$ is a Lie algebra homomorphism splitting π_* and $\gamma : \Gamma_c(TM) \rightarrow \Gamma_c(TP)^G$ is the lift induced by an equivariant connection, then we define $f := \sigma - \gamma$. It maps to the vertical vector fields, identified with $\Gamma(\text{ad}(P))$. We define A_m to be the maximal subalgebra of vector fields that vanish at m , and $\tilde{f}_m : A_m \rightarrow \mathfrak{g}$ to be the restriction of f to A_m , followed by the map $\Gamma(\text{ad}(P)) \rightarrow \mathfrak{g}$ which picks out the fibre over m and identifies it with \mathfrak{g} .

7.2.3 The Splitting as a Differential Operator

In this section, we will prove that σ is a differential operator of finite order. Since f is local, it defines a map from the sheaf of smooth sections of TM to

the sheaf of smooth sections of $\text{ad}(P)$. An elegant theorem of Peetre ([Pee60]) then says that f , and therefore σ , must be a differential operator of locally finite order. All we need to do is find a global bound on the order, which will occupy us for the remainder of the section.

The fact that f is a differential operator of locally finite order means that for each $m \in M$, there exists an r such that $H_m^r = \{v \in A_m \mid j_m^r(v) = 0\}$ is contained in $\ker(\tilde{f}_m)$, where $j_m^r(v)$ is the r -jet of v at m . Consequently, \tilde{f}_m factors through the jet Lie algebra $J_m^{r,0}(TM) := A_m/H_m^r$. Since we do not know r , we define

$$\text{Vec}_n := \varprojlim A_m/H_m^r,$$

and remark that \tilde{f}_m induces a homomorphism $\text{Vec}_n \rightarrow \mathfrak{g}$. The Lie algebra Vec_n depends on M only through its dimension n .

Local co-ordinates provide one with a basis $x^{\vec{\alpha}}\partial_i$, where $x^{\vec{\alpha}}$ is shorthand for $x_1^{\alpha_1} \dots x_n^{\alpha_n}$. With $\text{Vec}_n^k = \text{Span}\{x^{\vec{\alpha}}\partial_i \mid |\vec{\alpha}| = k+1, i = 1 \dots n\}$, one may write

$$\text{Vec}_n = \bigoplus_{k=0}^{\infty} \text{Vec}_n^k.$$

That is, each element of Vec_n can be uniquely written as a *finite* sum of homogeneous vector fields. Note that Vec_n^k is the k -eigenspace of the Euler vector field $E := \sum_{i=1}^n x^i \partial_i$. If I is an ideal containing $v = \sum_{k=0}^N v_k$, then $\text{ad}(E)^j v = \sum_{k=0}^N k^j v_k \in I$ for all j , so that $v_k \in I$. Thus any ideal splits into homogeneous components

$$I = \bigoplus_{k=0}^{\infty} I^k$$

with $I^k = I \cap \text{Vec}_n^k$. This renders the ideal structure of Vec_n more or less tractable, so that we may prove the following bound on the order of \tilde{f}_m .

Lemma 7.6 *The order of the differential operator σ is at most $\dim(\mathfrak{g})$ unless $\dim(M) = 1$ and $\dim(\mathfrak{g}) = 2$, in which case the order is at most 3.*

Proof. We closely follow Epstein and Thurston, [ET79]. One checks by hand that the only ideals of $\text{Vec}_1 = \text{Span}\{x^k \partial \mid k \geq 1\}$ are $\text{Span}\{x^2 \partial, x^k \partial \mid k \geq 4\}$ and $\text{Span}\{x^k \partial \mid k \geq N\}$ with $N \geq 1$.

Consider Vec_1 as a subalgebra of Vec_n , define K to be the kernel of \tilde{f}_m , and let $K_1 := K \cap \text{Vec}_1$. We then have injective homomorphisms

$$\text{Vec}_1/K_1 \hookrightarrow \text{Vec}_n/K \hookrightarrow \mathfrak{g},$$

so that $\dim(\text{Vec}_1/K_1) \leq \dim(\mathfrak{g})$. As K_1 is an ideal, it must be of the shape mentioned above. This leads us to conclude that $x_1^k \partial_1 \in K$ for all $k > \dim(\mathfrak{g})$ unless $\dim(\mathfrak{g}) = 2$, in which case $x_1^k \partial_1 \in K$ for all $k > 3$, and $x_1^2 \partial_1 \in K$.

The following short calculation shows that if $\dim(\mathfrak{g}) = 2$ and $\dim(M) > 1$, then also $x_1^3 \partial_1 \in K$. As K contains $x_1^2 \partial_1$, it also contains $[x_1^2 \partial_1, x_1 \partial_2] = x_1^2 \partial_2$, and thus $[x_1^2 \partial_2, x_1 x_2 \partial_1] = x_1^3 \partial_1 - 2x_1^2 x_2 \partial_2$. But by bracketing with $x_1^2 \partial_2$ and $x_2 \partial_1$ respectively, we see that $x_1^3 \partial_1 - 3x_1^2 x_2 \partial_2$ is in K , ergo $x_1^3 \partial_1 \in K$.

The next step is to show that if $x_1^s \partial_1 \in K$, then K also contains all $x^{\vec{\alpha}} \partial_i$ with $|\vec{\alpha}| = s$. First of all, we remain in K if we repeatedly apply $\text{ad}(x_i \partial_1)$ to $x_1^s \partial_1$, to the effect of replacing x_1 by x_i up to a nonzero factor. This shows that $x^{\vec{\alpha}} \partial_1 \in K$. Then the relation $x^{\vec{\alpha}} \partial_i = [x^{\vec{\alpha}} \partial_1, x_1 \partial_i] + x_1 \partial_i x^{\vec{\alpha}} \partial_1$ transfers membership of K from right to left.

In the generic case $\dim(\mathfrak{g}) \neq 2$, $\dim(M) \neq 1$, we may conclude that the order of σ is at most $\dim(\mathfrak{g})$, because $H_m^{\dim(\mathfrak{g})} \subset K$. In the exceptional case that $\dim(\mathfrak{g}) = 2$ and $\dim(M) = 1$, the order of σ is at most 3. \square

In particular, σ is a differential operator of finite rather than locally finite order. Let us summarize our progress so far.

Proposition 7.7 *Let P be an infinitesimally natural principal G -bundle. Then $\sigma : \Gamma_c(TM) \rightarrow \Gamma_c(TP)^G$ factors through the bundle of k -jets, where $k = 3$ if $\dim(M) = 1$, $\dim(\mathfrak{g}) = 2$ and $k = \dim(\mathfrak{g})$ otherwise. If we identify equivariant vector fields on P with sections of the Atiyah bundle TP/G , we can therefore define a bundle map $\nabla : J^k(TM) \rightarrow TP/G$ by $\nabla(j_m^k(v)) := \sigma(v)_m$. It makes the following diagram commute.*

$$\begin{array}{ccc} \Gamma(TM) & \xrightarrow{\sigma} & \Gamma(TP/G) \\ & \searrow & \nearrow \nabla \\ & & \Gamma(J^k(TM)) \end{array}$$

The point is that although σ is defined only on sections, ∇ comes from a veritable bundle map $J^k(TM) \rightarrow TP/G$. Note that although σ was only defined on $\Gamma_c(TM)$, it extends to $\Gamma(TM)$ by locality.

7.3 Lie Groupoids and Algebroids of Jets

The bundles $J^k(TM)$ and TP/G are Lie algebroids, and it will be essential for us to prove that $\nabla : J^k(TM) \rightarrow TP/G$ is a homomorphism of Lie algebroids. In order to do this, we will first have a closer look at $J^k(TM)$ and TP/G , and at their corresponding Lie groupoids.

Let us first set some notation. The jet group $G_{0,0}^k(\mathbb{R}^n)$ is the group of k -jets of diffeomorphisms of \mathbb{R}^n that fix 0. It is the semi-direct product of $\text{GL}(\mathbb{R}^n)$ and the connected, simply connected, unipotent Lie group of k -jets that equal the identity to first order.

The subgroup $G_{0,0}^{+k}(\mathbb{R}^n)$ of orientation preserving k -jets is connected, but not simply connected. As $G_{0,0}^{+k}(\mathbb{R}^n)$ retracts to $\text{SO}(\mathbb{R}^n)$, its homotopy group is isomorphic to $\{1\}$ if $n = 1$, to \mathbb{Z} if $n = 2$, and to $\mathbb{Z}/2\mathbb{Z}$ if $n > 2$. For brevity, we introduce the following notation.

Definition 7.3 *If $k > 0$, we denote $\pi_1(G_{0,0}^{+k}(\mathbb{R}^n))$ by Z .*

Thus for $n > 2$, the universal cover $\tilde{G}_{0,0}^{+k}(\mathbb{R}^n) \rightarrow G_{0,0}^{+k}(\mathbb{R}^n)$ is $2 : 1$, and restricts to the spin group over $\text{SO}(\mathbb{R}^n)$.

7.3.1 The Lie Groupoid of k -Jets

In this section, we define the Lie groupoid $G^k(M)$ of k -jets, its maximal source-connected Lie subgroupoid $G^{+k}(M)$, and the k^{th} order frame bundle $F^k(M)$.

Denote by $G_{m',m}^k(M)$ the manifold of k -jets at m of diffeomorphisms of M which map m to m' , and denote by $G^k(M) = \cup_{M \times M} G_{m',m}^k(M)$ the groupoid of k -jets. If $j_m^k(\alpha)$ is a k -jet at m of a diffeomorphism α , then its source is $s(j_m^k(\alpha)) = m$, its target is $t(j_m^k(\alpha)) = \alpha(m)$, and multiplication is given by concatenation.

Denote by $G_{*,m}^k(M)$ the manifold $s^{-1}(m)$ of k -jets with source m . The target map $t : G_{*,m}^k(M) \rightarrow M$ endows it with a structure of principal fibre bundle, the structure group $G_{m,m}^k(M) \simeq G_{0,0}^k(\mathbb{R}^n)$ acting freely and transitively on the right. As $G^1(M)_{*,m}$ is isomorphic to the frame bundle $F(M)$, one calls $G^k(M)_{*,m}$ the k^{th} order frame bundle, sometimes denoted $F^k(M)$.

Lemma 7.8 *Let M be connected and let $k \geq 1$. Then $G^k(M)$ is source-connected if and only if M is not orientable.*

Proof. We may as well consider $k = 1$, because the fibres of $G^k(M) \rightarrow G^1(M)$ can be contracted. Each source fibre $G^1(M)_{*,m}$ of $G^1(M)$ is isomorphic to the frame bundle. By definition, M is oriented precisely when the frames can be grouped into positively and negatively oriented ones. \square

Definition 7.4 *We define $G^{+k}(M)$ to be the maximal source-connected Lie subgroupoid of $G^k(M)$, and denote its source fibre by $F^{+k}(M)$.*

In the light of the previous lemma, this means that $G^{+k}(M)$ is the Lie groupoid of k -jets of orientation preserving diffeomorphisms if M is orientable, and simply $G^k(M)$ if M is not.

Note that the map $D : \text{Diff}(M) \rightarrow \text{Diff}(G^k(M))$ defined by $D\alpha : j_m^k(\gamma) \mapsto j_m^k(\alpha\gamma)$ is a homomorphism of groups. We will call it the k^{th} order derivative. Because $D\alpha$ is source-preserving and right invariant, it defines a homomorphism $\text{Diff}(M) \rightarrow \text{Aut}_{G_{*,m}^k(M)}^{G_{m,m}^k(M)}(G_{*,m}^k(M))$, splitting the exact sequence of groups (7.1). This makes $G_{*,m}^k(M) = F^k(M)$ into a natural bundle. Note that $F^{+k}(M)$ is infinitesimally natural.

7.3.2 The Lie Algebroid of k -Jets

The bundle $J^k(TM)$ possesses a structure of Lie algebroid, induced by the Lie groupoid $G^k(M)$. We now describe the Lie bracket on $\Gamma(J^k(TM))$ explicitly. Later, in section 7.3.4, we will use this to show that ∇ is a Lie algebroid homomorphism.

The Lie algebroid of $G^k(M)$ is a vector bundle $A \rightarrow M$. Its fibre A_m is by definition the subspace of the tangent space of $G^k(M)$ at $j_m^k(\text{id})$ which is annihilated by ds . Sections of A therefore correspond to right-invariant vector fields on $G^k(M)$ parallel to the source fibres.

Each curve in $G_{*,m}^k(M)$ through $j_m^k(\text{id})$ takes the shape $c(t) = j_m^k(\alpha_t)$ with $\alpha_0 = \text{id}$, so that its tangent vector $a \in A_m$ takes the shape $a = j_m^k(v)$, with $v = \partial_t|_0 \alpha_t$. This shows that $A \simeq J^k(TM)$.

The anchor $dt : J^k(TM) \rightarrow TM$ is easily seen to be the canonical projection, so we shall denote it by π . The Lie bracket on $\Gamma(J^k(TM))$ however, which is defined as the restriction of the commutator bracket on $\Gamma(TG^k(M))$ to the right invariant source preserving vector fields, perhaps deserves some comment.

Define $J^{k,0}(TM)$ to be the kernel of π , and consider the exact sequence of Lie algebras

$$0 \rightarrow \Gamma(J^{k,0}(TM)) \rightarrow \Gamma(J^k(TM)) \xrightarrow{\pi} \Gamma(TM) \rightarrow 0.$$

It is split by $j^k : \Gamma(TM) \rightarrow \Gamma(J^k(TM))$, the infinitesimal version of the k^{th} order derivative. (The sequence of Lie algebroids of course does not split, as differentiation is not linear over $C^\infty(M)$.)

We will describe the Lie bracket on $\Gamma(J^k(TM))$ by giving it on $\Gamma(TM)$ and $\Gamma(J^{k,0}(TM))$ separately, and then giving the action of $\Gamma(TM)$ on $\Gamma(J^{k,0}(TM))$.

Proposition 7.9 *Let u and u' be sections of TM , and let $\tau : m \mapsto j_m^k(v_m)$ and $\tau' : m \mapsto j_m^k(v'_m)$ be sections of $J^{k,0}(TM)$. Then*

$$\begin{aligned} [j^k(u), j^k(u')]_m &= j_m^k([u, u']), \\ [\tau, \tau']_m &= j_m^k([v_m, v'_m]), \\ [j^k(u), \tau]_m &= j_m^k([u, v_m]) + j_m^k(d_u|_m(x \mapsto v_x)), \end{aligned}$$

where $d_u|_m(x \mapsto v_x)$ is the ordinary derivative at m along u of a map from M to $\Gamma(TM)$. Although both terms on the right hand side depend on the choice of $m \mapsto v_m$, their sum does not.

Proof. The first equality is clear, as j^k is a homomorphism of Lie algebras. The second equality can be seen as follows. Consider the bundle of groups $G^k(M)_{*,*} := \{j_m^k(\alpha) \in G^k(M) \mid \alpha(m) = m\}$, with bundle map $s = t$. Its sections $\Gamma(G^k(M)_{*,*})$ form a group under pointwise multiplication, the Lie algebra of which is $\Gamma(J^{k,0}(TM))$, with the pointwise bracket. As $\Gamma(G^k(M)_{*,*})$ acts from the left on $G^k(M)$ by $j_m^k(\gamma) \mapsto j_m^k(\alpha_m) \circ j_m^k(\gamma)$, respecting both the source map and right multiplication, the inclusion $\Gamma(J^{k,0}(TM)) \rightarrow \Gamma(J^k(TM))$ is a homomorphism of Lie algebras. This proves the second line.

To verify the third line, we must choose a smooth map $x \mapsto v_x$ from M to $\Gamma(TM)$ such that $\tau_x = j_x^k(v_x)$ in a neighbourhood of m . Each v_x necessarily has a zero at x . If we denote $m(s) := \exp(su)m$, then the bracket $[j^k(u), \tau]_m$ is by definition² minus the mixed second derivative along s and t at 0 of the groupoid commutator

$$j_{m(s)}^k(\exp(-su))j_{m(s)}^k(\exp(-tv_{m(s)}))j_m^k(\exp(su))j_m^k(\exp(tv_m)),$$

which is just $j_m^k(\exp(-su)\exp(-tv_{m(s)})\exp(su)\exp(tv_m))$. The terms not involving derivatives of $s \mapsto m(s)$ yield $j_m^k([u, v_m])$, and the terms which do provide the extra $j_m^k(d_u|_m(x \mapsto v_x))$. \square

²The Lie algebra of the diffeomorphism group is the Lie algebra of vector fields, but the exponential map is given by $v \mapsto \exp(-v)$, where \exp denotes the unit flow along v . This is why the groupoid commutator might seem odd at first sight.

7.3.3 The Gauge Groupoid and its Algebroid

Given a principal G -bundle $\pi : P \rightarrow M$, one can define the gauge groupoid $(P \times P)/G$, that is the pair groupoid modded out by the diagonal action. Source and target come from projection on the second and first term respectively, $M \hookrightarrow (P \times P)/G$ as $\text{id}_{\pi(p)} = [(p, p)]$, and multiplication is well defined by $[(r, q)] \circ [(q, p)] = [(r, p)]$. An element $[(q, p)]$ corresponds precisely to a G -equivariant diffeomorphism $\pi^{-1}(p) \rightarrow \pi^{-1}(q)$, and the product to concatenation of maps.

Its Lie algebroid TP/G is sometimes called the Atiyah algebroid. Indeed, the subspace of $T_{\text{id}_m}((P \times P)/G)$ which annihilates ds is canonically $(TP/G)_m$. A section of TP/G can be identified with a G -equivariant section of TP , endowing $\Gamma(TP/G)$ with the Lie bracket that comes from $\Gamma(TP)^G$.

7.3.4 The Splitting as a Homomorphism of Lie Algebroids

The point of considering the Lie algebroid structure of $J^k(TM)$ was of course to prove the following.

Lemma 7.10 *The map $\nabla : J^k(TM) \rightarrow TP/G$ is a homomorphism of Lie algebroids.*

Proof. The fact that ∇ respects the anchor is immediate. As $\pi_* \circ \nabla(j_m^k(v_m)) = \pi_* \circ \sigma(v_m)(m) = v_m(m)$, it equals $\pi \circ j^k(v_m)$ in the point m .

We now show that $\nabla : \Gamma(J^k(TM)) \rightarrow \Gamma(TP/G)$ is a homomorphism of Lie algebras. First of all, the restriction of ∇ to $j^k(\Gamma(TM))$ is a homomorphism. Indeed, $[\nabla(j^k(v)), \nabla(j^k(w))] = [\sigma(v), \sigma(w)]$, which equals $\sigma([v, w])$ because σ is a homomorphism. This in turn is just $\nabla(j^k([v, w]))$, so that $[\nabla(j^k(v)), \nabla(j^k(w))] = \nabla([j^k(v), j^k(w)])$.

Secondly, its restriction to $\Gamma(J^{k,0}(TM))$ is a homomorphism. If $\tau_x = j_x^k(v_x)$ and $v_x = j_x^k(w_x)$ are sections of $J^{k,0}(TM)$, then $\nabla\tau$ and ∇v are in the kernel of the anchor. This implies that their commutator at a certain point m depends only on their values at m , not on their derivatives. To find the commutator at m , we may therefore replace $j_x^k(v_x)$ by $j_x^k(v_m)$ and likewise $j_x^k(w_x)$ by $j_x^k(w_m)$. We then see that $[\nabla j_x^k(v_x), \nabla j_x^k(w_x)]_m = [\nabla j_x^k(v_m), \nabla j_x^k(w_m)]_m$. We already know that this is $\nabla(j_x^k([v_m, w_m]))_m$, so that $[\nabla(\tau), \nabla(v)]_m = \nabla([\tau, v])_m$.

The last step is to show that ∇ respects the bracket between $j^k(\Gamma(TM))$ and $\Gamma(J^{k,0}(TM))$. Again, let $j^k(v)$ be an element of the former and $\tau_x = j_x^k(w_x)$ of the latter. Considered as an equivariant vector field on P , the vertical vector field $\nabla(\tau)$ takes the value $\sigma_p(w_{\pi(p)})$ at $p \in P$. Then $[\nabla(j^k(v)), \nabla(\tau)]$ is the Lie derivative along $\sigma(v)$ of the vertical vector field $\sigma_p(w_{\pi(p)})$. Differentiating along $\sigma(v)$ is done by considering $(p, p') \mapsto \sigma_p(w_{\pi(p')})$, differentiating w.r.t. p and p' separately, and then putting $p = p'$. This results in $[\sigma(v), \sigma_p(w_{\pi(p)})]_{p_0} = [\sigma(v), \sigma(w_{\pi(p_0)})]_{p_0} + \sigma(d_v|_{\pi(p_0)}(x \mapsto w_x))$, which is in turn the same as $\sigma_{p_0}([v, w_{\pi(p_0)}] + d_v|_{\pi(p_0)}(m \mapsto w_m))$, so that $[\nabla(j^k(v)), \nabla(\tau)] = \nabla([j^k(v), \tau])$ as required. Therefore ∇ must be a homomorphism on all of $\Gamma(J^k(TM))$. \square

Definition 7.5 *A connection ∇ of a Lie algebroid A on a vector bundle E is by definition a bundle map of A into $\text{DO}^1(E)$, the first order differential operators on E , which respects the anchor. If moreover it is a morphism of Lie algebroids, then the connection is called flat. A flat connection of A on E is also called a representation of A on E .*

This explains our notation for the map ∇ induced by σ . Given a representation V of G , one may form the associated vector bundle $E := P \times_G V$. The map ∇ then defines a Lie algebroid homomorphism of $\Gamma(J^k(TM))$ into the Lie algebroid of first order differential operators on E . (Simply consider a section of E as a G -equivariant function $P \rightarrow V$, and let $\Gamma^G(TP)$ act by Lie derivative.) By definition, this is a flat connection, or equivalently a Lie algebroid representation.

7.4 The Classification Theorem

We use the fact that $\nabla : \Gamma(J^k(TM)) \rightarrow \Gamma(TP/G)$ is a homomorphism of Lie algebroids to find a corresponding homomorphism of Lie groupoids. This will give us the desired classification of infinitesimally natural principal fibre bundles.

7.4.1 Integrating a Homomorphism of Lie Algebroids

The following theorem states that homomorphisms of Lie algebroids induce homomorphisms of Lie groupoids if the initial groupoid is source-simply connected.

Theorem 7.11 (Lie II for algebroids) *Let G and H be Lie groupoids, with corresponding Lie algebroids A and B respectively. Let $\nabla : A \rightarrow B$ be a homomorphism of Lie algebroids. If G is source-simply connected, then there exists a unique homomorphism $G \rightarrow H$ of Lie groupoids which integrates ∇ .*

Remark The result was probably announced first in [Pra66], and proofs have appeared e.g. in [MX00] and [MM02]. We follow [CF03], which the reader may consult for details.

Sketch of proof. The idea is that ∇ allows one to lift piecewise smooth paths of constant source in G to piecewise smooth paths of constant source in H . Source-preserving piecewise smooth homotopies in G of course do not affect the endpoint of the path in H , so that, if G is source-simply connected, one obtains a map $G \rightarrow H$ by identifying elements g of G with equivalence classes of source preserving paths from $\text{id}_{s(g)}$ to g . One checks that this is the unique homomorphism of Lie groupoids integrating ∇ . \square

Unfortunately, $G^k(M)$ is not always source connected, let alone source-simply connected. Recall that $G^{+k}(M)$ is the maximal source-connected Lie subgroupoid of $G^k(M)$, and therefore has the same Lie algebroid $J^k(TM)$.

We define $\tilde{G}^{+k}(M)$ to be the set of piecewise smooth, source preserving paths in $G^{+k}(M)$ beginning at an identity, modulo piecewise smooth, source preserving homotopies. It is a smooth manifold because $G^{+k}(M)$ is, and a Lie

groupoid under the unique structure making the projection on the endpoint $\tilde{G}^{+k}(M) \rightarrow G^{+k}(M)$ into a morphism of groupoids. Explicitly, the multiplication is given as follows. If $g(t)$ a path from id_m to $g(1)_{m'm}$, and $h(t)$ a path from $\text{id}_{m'}$ to $h(1)_{m''m'}$, then the product $[h] \circ [g]$ is $[(h \cdot g(1)) * g]$, where the dot denotes groupoid multiplication and the star concatenation of paths. The proof of associativity is the usual one.

Note that the source fibre $\tilde{G}^{+k}(M)_{*,m}$ is precisely the universal cover of the connected component of the k^{th} order frame bundle $G^{+k}(M)_{*,m} = F^{+k}(M)$. In order to cut down on the subscripts, we introduce new notation for $\tilde{G}^{+k}(M)_{*,m}$ and its structure group $\tilde{G}^{+k}(M)_{m,m}$.

Definition 7.6 *We denote the universal cover of the connected component of the k^{th} order frame bundle by $\tilde{F}^{+k}(M)$, and its structure group by $G(k, M)$.*

It is an infinitesimally natural bundle because $F^{+k}(M)$ is. Note that $G(k, M)$ is not the universal cover of $G_{m,m}^{+k}(M)$, but rather its extension by $\pi_1(F^k(M))$. As $\pi_1(F^k(M)) = \pi_1(F(M))$, we have the exact sequence of groups

$$1 \rightarrow \pi_1(F(M)) \rightarrow G(k, M) \rightarrow G_{m,m}^{+k}(M) \rightarrow 1.$$

The group $G_{m,m}^{+k}(M)$ in turn is isomorphic to $G_{0,0}^{+k}(\mathbb{R}^n)$ if M is orientable, and to $G_{0,0}^k(\mathbb{R}^n)$ if it is not.

7.4.2 Classification

Now that we've found a source-simply connected Lie groupoid with $J^k(TM)$ as Lie algebroid, we can finally apply Lie's second theorem for algebroids to obtain the following.

Proposition 7.12 *If σ splits the exact sequence of Lie algebras (7.2), then it induces a morphism of groupoids $\exp \nabla : \tilde{G}^{+k} \rightarrow (P \times P)/G$ such that the following diagram commutes, with \exp_m the flow along a vector field starting at id_m .*

$$\begin{array}{ccccc}
 \Gamma(TM) & \xrightarrow{\sigma} & \Gamma(TP/G) & \xrightarrow{\exp_m} & (P \times P)/G \\
 & \searrow^{j^k} & \nearrow^{\nabla} & & \nearrow^{\exp \nabla} \\
 & & \Gamma(J^k(TM)) & \xrightarrow{\exp_m} & \tilde{G}^{+k}(M)
 \end{array}$$

Proof. As $\tilde{G}^{+k}(M)$ is a source-simply connected Lie groupoid with $J^k(TM)$ as Lie algebroid, we can apply Lie's second theorem for algebroids. \square

It is perhaps worth while to formulate this for general transitive Lie groupoids, as it clarifies the link with the recent work of Grabowski, Kotov and Poncin [GKP09]. For the Atiyah algebroid A of a principal fibre bundle with a connected, reductive structure group, they classify Lie algebra isomorphisms of $\Gamma(A)$ in terms of the Lie algebroid isomorphisms of A .

Proposition 7.13 *Let $\mathcal{G} \rightrightarrows M$ be a transitive Lie groupoid, with Lie algebroid A . The kernel of the anchor K is then a bundle of Lie algebras with fixed dimension d . Suppose that the sequence*

$$0 \rightarrow \Gamma(K) \rightarrow \Gamma(A) \rightarrow \Gamma(TM) \rightarrow 0,$$

with K the kernel of the anchor, splits as a sequence of Lie algebras. Then this splitting is induced by a morphism of Lie algebroids $\nabla : J^k(TM) \rightarrow A$, and there is a corresponding morphism of Lie groupoids $\tilde{G}^{+k}(M) \rightarrow \mathcal{G}$. The number k is at most 3 if d is 2 and $\dim(M) = 1$, and at most d otherwise.

Proof. (Or rather a flimsy sketch thereof.) Analogous to the case of the gauge groupoid. \square

We have paved the way for a classification of infinitesimally natural principal fibre bundles.

Theorem 7.14 *Let $\pi : P \rightarrow M$ be an infinitesimally natural principal G -bundle with splitting σ of (7.2). Then there exists a group homomorphism $\rho : G(k, M) \rightarrow G$ such that the bundle P is associated to $\tilde{F}^{+k}(M)$ through ρ , i.e.*

$$P \simeq \tilde{F}^{+k}(M) \times_{\rho} G.$$

Moreover, σ is induced by the canonical one for $\tilde{F}^{+k}(M)$.

Proof. Fix a base point m on M . The map $\exp \nabla$ yields a homomorphism of groups $\rho : \tilde{G}_{m,m}^{+k}(M) \rightarrow ((P \times P)/G)_{m,m}$, the latter isomorphic to G , the former to $G(k, M)$.

The map $\tilde{G}_{*,m}^{+k}(M) \times_{\rho} ((P \times P)/G)_{m,m} \rightarrow ((P \times P)/G)_{*,m}$ which is given by $(g_{m',m}, p_{m,m}) \mapsto (\exp \nabla(g_{m',m})) \cdot p_{m,m}$ is well defined and injective because two pairs share the same image if and only if they are equivalent modulo $\tilde{G}_{m,m}^{+k}(M)$. It is also surjective and G -equivariant, and hence an isomorphism of principal G -bundles. As $((P \times P)/G)_{*,m} \simeq P$ and $\tilde{G}_{*,m}^{+k}(M) \times_{\rho} ((P \times P)/G)_{m,m} \simeq \tilde{F}^{+k}(M) \times_{\rho} G$, the equivalence is proven. The remark on σ follows from the construction. \square

This classifies the infinitesimally natural principal fibre bundles. They are all associated (via a group homomorphism) to the bundle $\tilde{G}_{*,m}^{+k} = \tilde{F}^{+k}(M)$.

The classification of natural principal fibre bundles is now an easy corollary. The following well known result ([PT 77], [Ter 78]) states that they are precisely the ones associated to $G_{*,m}^k(M) = F^k(M)$.

Corollary 7.15 *Let $\pi : P \rightarrow M$ be a natural principal G -bundle with local splitting Σ of (7.1). Then P is associated to $F^k(M)$. That is, there exists a homomorphism $\rho : G_{0,0}^k(\mathbb{R}^n) \rightarrow G$ such that*

$$P \simeq F^k(M) \times_{\rho} G.$$

Moreover, Σ is induced by the canonical one for $F^k(M)$.

Proof. As the homomorphism $\Sigma : \text{Diff}_c(M) \rightarrow \text{Aut}_c(P)$ is local, it induces a homomorphism of groupoids $\Sigma : \text{Germ}(M) \rightarrow (P \times P)/G$, with $\text{Germ}(M)$ the groupoid of germs of diffeomorphisms of M . We need but show that Σ factors through $j^k : \text{Germ}(M) \rightarrow G^k(M)$ for some $k > 0$, cf. the proof of theorem 7.14.

The Lie algebra homomorphism $\sigma : \Gamma_c(TM) \rightarrow TP/G$ defined by $\sigma(v) := \partial_t|_0 \Sigma(\exp(tv))$ is local by assumption, and according to proposition 7.7 it factors through the k -jets for some $k > 0$. It suffices to show that $\Sigma(\phi)_{m,m} = \text{id}_{m,m}$ for any $\phi \in \text{Germ}_{m,m}(M)$ that agrees with the identity to k^{th} order at m .

In local co-ordinates $\{x^i\}$, we write $\phi^i(x) = x^i + v^i(x)$, where $v : \mathbb{R}^n \rightarrow \mathbb{R}^n$ vanishes to k^{th} order. We define the one parameter family of germs of diffeomorphisms $\phi_t^i(x) := x^i + tv^i(x)$. Then $\partial_t|_\tau \Sigma(\phi_\tau)^{-1} \Sigma(\phi_t)_m = 0$, as it equals $\sigma_m(\partial_t|_\tau \phi_\tau^{-1} \phi_t)$, the image of a vector field that vanishes to order k at m . Therefore $t \mapsto \Sigma(\phi_t)_{m,m}$ is constant, and $\Sigma(\phi)_{m,m} = \text{id}_{m,m}$ as required. \square

To summarize: natural principal fibre bundles are associated to a higher frame bundle, whereas infinitesimally natural principal fibre bundles are associated to the universal cover of a higher frame bundle.

7.4.3 The Bundle $\tilde{F}^{+k}(M)$

The above considerations prompt a few remarks on the universal cover of the connected component of the frame bundle $\tilde{F}^{+k}(M)$, and on its (disconnected) structure group $G(k, M)$. Recall that they are just the source fibre $\tilde{G}_{*,m}^{+k}(M)$ and isotropy group $\tilde{G}_{m,m}^{+k}(M)$ of $\tilde{G}^{+k}(M)$.

General Manifolds

If $\pi_1(M)$ is the homotopy groupoid of M , define the homomorphism of groupoids $\text{Pr} : \pi_1(M)_{m',m} \rightarrow \pi_0(G_{m',m}(M))$ by lifting a path in M to a path in $G^k(M)$ with fixed source, and taking the connected component of its end point. It makes

$$\begin{array}{ccccc}
 & & \pi_1(M)_{m',m} & & \\
 & \nearrow & & \searrow & \\
 \tilde{G}_{m',m}^{+k}(M) & & & & \pi_0(G_{m',m}^k(M)) \\
 & \searrow & & \nearrow & \\
 & & G_{m',m}^k(M) & &
 \end{array}$$

into a commutative diagram.

Define $(G^k(M) \times \pi_1(M))^{\text{Pr}}$ to be the groupoid of pairs $(g, [f])$ such that $\pi_0(g) = \text{Pr}([f])$. If M is orientable, this is simply $G^{+k}(M) \times \pi_1(M)$. The map of groupoids $\tilde{G}^{+k}(M) \rightarrow (G^k(M) \times \pi_1(M))^{\text{Pr}}$ is well defined and surjective. It restricts to a covering map of principal fibre bundles

$$\tau : \tilde{G}_{*,m}^{+k}(M) \rightarrow (G^k(M) \times \pi_1(M))_{*,m}^{\text{Pr}}. \tag{7.7}$$

The kernel of the corresponding cover of groups is precisely $i_* \pi_1(G_{m,m}^{+k}(M))$, with $i : G_{m,m}^{+k}(M) \rightarrow G_{*,m}^{+k}(M)$ the inclusion. Note that i_* has a nonzero kernel precisely when a vertical loop is contractible in $G_{*,m}^{+k}(M)$, but not by a

homotopy which stays inside the fibre. Denoting $\pi_1(G_{m,m}^{+k}(M))$ by Z , we obtain the exact sequence

$$1 \rightarrow Z/\text{Ker}(i_*) \xrightarrow{i_*} \tilde{G}_{m,m}^{+k}(M) \xrightarrow{\tau} (G^k(M) \times \pi_1(M))_{m,m}^{\text{Pr}} \rightarrow 1. \quad (7.8)$$

A moment's thought reveals that this extension is central: if $g(t)$ is a path in $G_{m,m}^{+k}(M)$ and $h(t)$ one in $G_{*,m}^{+k}(M)$, then both $h * (i \circ g)$ and $(i \circ g) \cdot h(1) * h$ can be homotoped into $t \mapsto h(t)g(t)$.

We may as well restrict attention to the case $k = 1$, in which $G_{*,m}^1(M)$ is the frame bundle $F(M)$. Indeed, as $G_{*,m}^{+k}(M) \rightarrow G_{*,m}^{1,+}(M)$ has contractible fibres, $\tilde{G}_{*,m}^{+k}(M)$ is just the pullback of $G_{*,m}^{+k}(M)$ along $\tilde{G}_{*,m}^{1,+}(M) \rightarrow G_{*,m}^{1,+}(M)$.

Orientable Manifolds

For orientable manifolds, the situation simplifies. If we identify the connected component of $G_{m,m}^1(M)$ with $\text{GL}^+(\mathbb{R}^n)$, we obtain a homomorphism i_* of $\widetilde{\text{GL}}^+(\mathbb{R}^n)$ into $\tilde{G}_{m,m}^{1,+}(M)$. There is a second homomorphism $\pi_1(F(M)) \rightarrow \tilde{G}_{m,m}^{1,+}(M)$. Their images intersect in $Z/\text{Ker}(i_*)$, and commute by an argument similar to the one on centrality of (7.8). If we define $(\widetilde{\text{GL}}^+(\mathbb{R}^n) \times \pi_1(F(M)))_Z$ to be the quotient of $\widetilde{\text{GL}}^+(\mathbb{R}^n) \times \pi_1(F(M))$ by the equivalence $(gz, h) \sim (g, zh)$, we can regard it as a subgroup of $\tilde{G}_{m,m}^{1,+}(M)$. Note that if $\text{Ker}(i_*)$ is nonzero, the above equivalence relation sets it to 1.

If M is orientable, we may restrict our attention to $F^+(M)$, which has connected fibres. Any path in $F^+(M)$ which starts and ends in the same fibre can therefore be obtained by combining a closed loop with a path in $\text{GL}^+(\mathbb{R}^n)$. For orientable manifolds, we thus have $\tilde{G}_{m,m}^{1,+}(M) \simeq (\widetilde{\text{GL}}^+(\mathbb{R}^n) \times \pi_1(F^+(M)))_Z$, and in the same vein

$$G(k, M) \simeq (G(k, \mathbb{R}^n) \times \pi_1(F^+(M)))_Z. \quad (7.9)$$

Spin Manifolds

Let M be an orientable manifold, equipped with a pseudo-Riemannian metric g of signature $\eta \in \text{Bil}(\mathbb{R}^n)$. Then $OF_g^+ := \{f \in F^+(M) \mid f^*g = \eta\}$ is the bundle of positively oriented orthogonal frames. A spin structure is then by definition an $\widetilde{\text{SO}}(\eta)$ -bundle³ Q over M , plus a map $u : Q \rightarrow OF_g^+$ such that the following diagram commutes, with κ the canonical homomorphism $\widetilde{\text{SO}}(\eta) \rightarrow \text{SO}(\eta)$. A manifold is called spin if it admits a spin structure.

³There is a subtlety here. Suppose η has indefinite signature, say $(3, 1)$. The group $\text{SO}(3, 1)$ has 2 connected components, so that a universal cover does not exist. As it is a subgroup of the simply connected group $\text{GL}^+(\mathbb{R}^4)$, we simply define $\widetilde{\text{SO}}(3, 1)$ to be $\kappa^{-1}(\text{SO}(3, 1))$ with $\kappa : \widetilde{\text{GL}}^+(\mathbb{R}^4) \rightarrow \text{GL}^+(\mathbb{R}^4)$ the covering map. Thus $\widetilde{\text{SO}}(3, 1)$ is, perhaps surprisingly, not isomorphic to the 2-component spin group $\text{Spin}(3, 1)$. Indeed, if T is time inversion and P is the inversion of 3 space co-ordinates, then $(PT)^2 = \mathbf{1}$ in $\widetilde{\text{SO}}(3, 1)$, as opposed to $(PT)^2 = -\mathbf{1}$ in $\text{Spin}(3, 1)$. Therefore $\pi^{-1}(\pm\mathbf{1}) \simeq \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ in $\widetilde{\text{SO}}(3, 1)$, whereas $\pi^{-1}(\pm\mathbf{1}) = \mathbb{Z}/4\mathbb{Z}$ in $\text{Spin}(3, 1)$ (see [BW⁺01]).

Of course the connected component of unity of $\widetilde{\text{SO}}(3, 1)$ and that of $\text{Spin}(3, 1)$ are both isomorphic to $\text{SL}(\mathbb{C}^2)$, so that none of this is relevant if M is both orientable and time-orientable, i.e. if the structure group of the frame bundle reduces to $\text{SO}^1(3, 1)$.

$$\begin{array}{ccc}
 \widetilde{\mathrm{SO}}(\eta) & \xrightarrow{\kappa} & \mathrm{SO}(\eta) \\
 \widehat{Q} & \xrightarrow{u} & \widehat{OF}_g^+ \\
 & \searrow & \swarrow \\
 & M &
 \end{array}$$

Define $\widehat{Q} := Q \times_{\widetilde{\mathrm{SO}}(\eta)} \widetilde{\mathrm{GL}}^+(\mathbb{R}^n)$, and let us again denote the induced map $\widehat{Q} \rightarrow F^+(M)$ by u . As any cover of $F^+(M)$ by a $\widetilde{\mathrm{GL}}^+(\mathbb{R}^n)$ -bundle can be obtained in this way, there is a 1:1 correspondence between spin covers of $\widehat{OF}_g^+(M)$ and $F^+(M)$. In particular, whether or not M is spin does not depend on the metric.

The Serre spectral sequence gives rise to the exact sequence

$$1 \rightarrow Z/\mathrm{Ker}(i_*) \rightarrow \pi_1(F^+(M)) \rightarrow \pi_1(M) \rightarrow 1. \tag{7.10}$$

The following proposition is well known.

Proposition 7.16 *A spin structure exists if and only if $i_* : Z \rightarrow \pi_1(F^+(M))$ is injective and (7.10) splits as a sequence of groups. If spin structures exist, then equivalence classes of spin covers correspond to splittings of (7.10).*

Proof. This will follow from theorem 7.23 later on, but see e.g. [Mor01] for an independent proof. Our criterion for M to be spin is equivalent to the vanishing of the second Stiefel-Whitney class, see e.g. [LM94]. \square

Remark In terms of group cohomology, one can consider the sequence (7.10) as an element $[\omega] \in H^2(\pi_1(M), Z/\mathrm{Ker}(i_*))$. Spin bundles exist if and only if both $\mathrm{Ker}(i_*)$ and $[\omega]$ are trivial, in which case they are indexed by $H^1(\pi_1(M), Z)$.

If a spin structure exists, then \widetilde{F}^+ is simply the pullback along the universal cover $\widetilde{M} \rightarrow M$ of $\widehat{Q} \rightarrow M$. The picture then becomes

$$\begin{array}{ccccc}
 \widetilde{F}^{+k}(M) & \longrightarrow & u_*(F^{+k}(M)) & \longrightarrow & F^{+k}(M) \\
 \downarrow & & \downarrow & & \downarrow \\
 \widetilde{F}^+(M) & \longrightarrow & \widehat{Q} & \xrightarrow{u} & F^+(M) \\
 \downarrow & & \downarrow & \swarrow & \\
 \widetilde{M} & \longrightarrow & M & &
 \end{array}$$

with each of the three squares a pullback square.

7.5 More General Fibre Bundles

In this section, we will prove a version of theorem 7.14 for fibre bundles which are not principal. It would however be overly optimistic to expect an analogue

of theorem 7.14 to hold for arbitrary smooth fibre bundles, so we will restrict ourselves to those bundles that carry a sufficiently rigid structure on their fibres.

7.5.1 Structured Fibre Bundles

We start by making this statement more precise.

Definition 7.7 *Let \mathbf{C} be a subcategory of the category of smooth manifolds such that the group of automorphisms of each object of \mathbf{C} is a finite dimensional Lie group. Then a ‘structured fibre bundle’ with structure \mathbf{C} and fibre $F_0 \in \text{ob}(\mathbf{C})$ is by definition a smooth fibre bundle $\pi : F \rightarrow M$ where the fibres are objects of \mathbf{C} . We also require each point to possess a neighbourhood U and a local trivialization $\phi : \pi^{-1}(U) \rightarrow F_0 \times U$ such that ϕ restricted to a single fibre is a \mathbf{C} -isomorphism $\pi^{-1}(m) \rightarrow F_0$.*

For example, a structured fibre bundle in the category of finite dimensional vector spaces is a vector bundle.

If $\pi : F \rightarrow M$ is any smooth fibre bundle, then an automorphism of π is by definition a diffeomorphism α of F such that $\pi(f) = \pi(f')$ implies $\pi(\alpha(f)) = \pi(\alpha(f'))$. It is called vertical if it maps each fibre to itself.

Definition 7.8 *We define an automorphism of a structured fibre bundle to be an automorphism of the smooth fibre bundle such that its restriction to each single fibre is an isomorphism in \mathbf{C} .*

One can then construct a sequence of groups

$$1 \rightarrow \text{Aut}_c^{\mathbf{C}}(F)^V \rightarrow \text{Aut}_c^{\mathbf{C}}(F) \rightarrow \text{Diff}_c^F(M) \rightarrow 1 \quad (7.11)$$

and its corresponding exact sequence of Lie algebras

$$0 \rightarrow \Gamma_c^{\mathbf{C}}(TF)^V \rightarrow \Gamma_c^{\mathbf{C}}(TF)^P \rightarrow \Gamma_c(TM) \rightarrow 0, \quad (7.12)$$

where ‘V’ is for vertical, ‘P’ for projectable, and c again stands for ‘0 outside a compact subset of M ’.

The proof of the following corollary of theorem 7.14 is now a formality.

Corollary 7.17 *Let $\pi : F \rightarrow M$ be a structured fibre bundle with fibre F_0 such that (7.12) splits as a sequence of Lie algebras. Then there exists an action ρ of $G(k, M)$ by \mathbf{C} -automorphisms on a single fibre F_m such that*

$$F = \tilde{F}^{+k}(M) \times_{\rho} F_0.$$

Proof. Construct the principal $\text{Aut}^{\mathbf{C}}(F_0)$ -bundle $\pi : P \rightarrow M$, the fibre over x of which is precisely the set of \mathbf{C} -isomorphisms $\phi : F_m \rightarrow F_x$. Then there is a natural isomorphism $\text{Aut}_c^{\mathbf{C}}(F) \simeq \text{Aut}_c(P)^G$ under which the vertical subgroups of the two correspond, so that the exact sequence (7.11) is isomorphic to (7.1), and therefore (7.12) to (7.2). As $F = P \times_{\text{Aut}(F_m)} F_m$, we can now apply theorem 7.14 to P in order to substantiate our claim. \square

7.5.2 Vector Bundles

We specialize to the case of vector bundles. These are precisely structured fibre bundles in the category of finite dimensional vector spaces.

The exact sequence of Lie algebras (7.12) for a vector bundle E with fibre V is then

$$1 \rightarrow \mathrm{DO}_c^0(E) \rightarrow \mathrm{DO}_c^1(E) \rightarrow \Gamma_c(TM) \rightarrow 0, \quad (7.13)$$

where $\mathrm{DO}_c^1(E)$ is the Lie algebra of compactly supported 1st order differential operators on $\Gamma(E)$, and $\mathrm{DO}_c^0(E)$ the ideal of 0th order ones, that is to say $\mathrm{DO}_c^0(E) \simeq \Gamma_c(E \otimes E^*)$.

Corollary 7.17 then says that (7.13) splits as a sequence of Lie algebras if and only if there is a representation ρ of $G(k, M)$ on V such that $E \simeq \tilde{F}^{+k}(M) \times_\rho V$.

But thanks to the fact that all finite dimensional representations of the universal cover of $\mathrm{GL}^+(\mathbb{R}^n)$ factor through $\mathrm{GL}^+(\mathbb{R}^n)$ itself, we can even say something slightly stronger.

Proposition 7.18 *Let $E \rightarrow M$ be a vector bundle for which (7.13) splits as a sequence of Lie algebras. Then there exists a representation ρ of the group $(G^k \times \pi_1(M))_{*,m}^{\mathrm{Pr}}$ on V such that*

$$E \simeq (G^k(M) \times \pi_1(M))_{*,m}^{\mathrm{Pr}} \times_\rho V.$$

Remark If M is orientable, this reads $E \simeq \pi^*F^{+k}(M) \times_\rho V$. In this expression, $\pi^*F^{+k}(M)$ is the pullback of $F^{+k}(M)$ along $\pi: \tilde{M} \rightarrow M$, considered as a principal $G_{0,0}^{+k}(\mathbb{R}^n) \times \pi_1(M)$ -bundle over M .

Proof. Consider the restriction of the map τ in equation (7.7) to the group $\tilde{G}_{m,m}^{+k}(M)$. In order to prove the proposition, we need but show that its kernel Z acts trivially on V . For $k = 0$, this is clear.

If k is at least 1, the homomorphism $\tilde{\mathrm{GL}}^+(\mathbb{R}^n) \rightarrow \tilde{G}_{m,m}^{+k}(M)$ makes V into a finite dimensional representation space for $\tilde{\mathrm{GL}}^+(\mathbb{R}^n)$. But it is known (see [Kna96, p. 311]) that all finite dimensional representations of its cover factor through $\mathrm{GL}^+(\mathbb{R}^n)$ itself. This implies that the subgroup Z which covers the identity must act trivially on V , and we may consider

$$\tilde{G}_{*,m}^{+k}(M)/Z \simeq (G^k(M) \times \pi_1(M))_{*,m}^{\mathrm{Pr}}$$

to be the underlying bundle, as announced. \square

This reduces the problem of classifying vector bundles with split sequence (7.13) to the representation theory of $(G^k \times \pi_1(M))_{*,m}^{\mathrm{Pr}}$.

The above extends a result [Ter78] of Terng, in which she classifies vector bundles which allow for a local splitting of the sequence of groups (7.11). It is an extension first of all in the sense that we prove, rather than assume, that the splitting is local. Secondly, we have shown that in classifying vector bundles with split sequence (7.13) of Lie algebras rather than groups, one encounters only slightly more. Intuitively speaking, the extra bit is the representation theory of $\pi_1(M)$. We refer to [Ter78] for a thorough exposition of the representation theory of $G_{0,0}^k(\mathbb{R}^n)$.

7.6 Flat Connections

Having concluded our classification of bundles in which (7.2) is split as a sequence of Lie algebras, the time has come to apply our newly acquired knowledge.

In this section, we will investigate splittings that come from a flat equivariant connection on a principal G -bundle $P \rightarrow M$. We will prove that if the Lie algebra \mathfrak{g} of G does not contain $\mathfrak{sl}(\mathbb{R}^n)$ as a subalgebra, then the sequence of Lie algebras (7.2) splits if and only if P admits a flat equivariant connection. In other words, the sequence (7.2) then splits as a sequence of Lie algebras if and only if it splits as a sequence of Lie algebras and $C^\infty(M)$ -modules.

Note that this is certainly not the case for general groups G . The frame bundle for example always allows for a splitting of (7.2), but usually not for a flat connection.

7.6.1 Lie Algebras That Do Not Contain $\mathfrak{sl}(\mathbb{R}^n)$

Although lemma 7.6 exhibits σ as a differential operator of finite order, the bound on the order is certainly not optimal. With full knowledge of the Lie algebras at hand, sharper restrictions can be put on the kernel of σ . In particular, if \mathfrak{g} does not contain $\mathfrak{sl}(\mathbb{R}^n)$, there is only a single relevant ideal, and σ is of order at most 1. For notation, see section 7.2.2.

Lemma 7.19 *Let $n = 1$, and let \mathfrak{g} be such that it does not contain two nonzero elements such that $[X, Y] = Y$. Or let $n \geq 2$, and let \mathfrak{g} be such that it does not admit $\mathfrak{sl}(\mathbb{R}^n)$ as a subalgebra. Then the kernel of the homomorphism $\check{f}_m : \text{Vec}_n \rightarrow \mathfrak{g}$ contains $\{v \in \text{Vec}_n \mid \text{Div}_m(v) = 0\}$.*

Proof. We start with the case $n = 1$. Again, we note that the only ideals of $\text{Vec}_1 = \text{Span}\{x^k \partial_x \mid k \geq 1\}$ are $\text{Span}\{x^2 \partial_x, x^k \partial_x \mid k \geq 4\}$, and for each $N \geq 1$ an ideal $\text{Span}\{x^k \partial_x \mid k \geq N\}$. The corresponding quotients all contain two elements X and Y with $[X, Y] = Y$, except the ideals corresponding to $N = 1, 2$. This means that also \mathfrak{g} , containing $\text{Vec}_1 / \ker(\check{f}_m)$ as the image of \check{f}_m , will possess X and Y such that $[X, Y] = Y$ unless the kernel of \check{f}_m contains the ideal $\text{Span}\{x^k \partial_x \mid k \geq 2\} = \{v \in \text{Vec}_1 \mid \text{Div}_m(v) = 0\}$.

Now for $n \geq 2$. Under the identification $\text{Vec}_n^0 \simeq \mathfrak{gl}(\mathbb{R}^n)$ given by $x_i \partial_j \mapsto e_{ij}$, the Euler vector field is the identity $\mathbf{1}$ and Div_m becomes the trace. As $\ker(\check{f}_m)^0$ is an ideal in $\mathfrak{gl}(\mathbb{R}^n)$, it can be either 0 , $\mathbb{R}\mathbf{1}$, $\mathfrak{sl}(\mathbb{R}^n)$ or $\mathbb{R}\mathbf{1} \oplus \mathfrak{sl}(\mathbb{R}^n)$. In the former two cases, $\text{Im}(\check{f}_m) \simeq \text{Vec}_n / \ker(\check{f}_m)$, and hence \mathfrak{g} , would contain $\mathfrak{sl}(\mathbb{R}^n)$ as a subalgebra, contradicting the hypothesis. Hence $\mathfrak{sl}(\mathbb{R}^n) \subseteq \ker(\check{f}_m)$. If we now show that $[\text{Vec}_n, \mathfrak{sl}(\mathbb{R}^n)] = \mathfrak{sl}(\mathbb{R}^n) \bigoplus_{k=1}^{\infty} \text{Vec}_n^k$, the proof will be complete.

Let $i \neq j$. We then have $[x_i \partial_j, x_j x_i^\alpha \partial_j] = (\alpha_j + 1)x_i x_i^\alpha \partial_j$, showing that $x_i x_i^\alpha \partial_j \in [\text{Vec}_n, \mathfrak{sl}(\mathbb{R}^n)]$. The only basis elements not of this shape are of the form $x_j^k \partial_j$. But $[x_j \partial_i, x_i x_j^{k-1} \partial_j] = x_j^k \partial_j - x_i x_j^{k-1} \partial_i$. If $k \geq 2$, the latter part was just shown to be in $[\text{Vec}_n, \mathfrak{sl}(\mathbb{R}^n)]$, so that also $x_j^k \partial_j \in [\text{Vec}_n, \mathfrak{sl}(\mathbb{R}^n)]$. If $k = 1$, the elements $x_j \partial_j - x_i \partial_i$ join $x_i \partial_j$ to form a basis of $\mathfrak{sl}(\mathbb{R}^n)$. \square

This rather limits the possibilities. Not only can we restrict to first order, but also the Lie algebroid map $\nabla : J^1(TM) \rightarrow TP/G$ vanishes on the trace-zero jets $K_m = \{j_m^1(v) \in J^1(TM) \mid v(m) = 0 \text{ and } \mathbf{tr}(v) = 0\}$, so that it factors through the ‘trace Lie algebroid’ $\mathbf{tr}_m(M) := J_m^1(TM)/K_m$.

This in turn is the Lie algebroid of the ‘determinant groupoid’ $\text{Det}(M)$. An element $[\alpha]_{m',m}$ of $\text{Det}(M)_{m',m}$ is by definition an equivalence class of diffeomorphisms mapping m to m' , with $\alpha \sim \beta$ if and only if $\text{Det}(\beta^{-1}\alpha) = 1$.

As $[\alpha]_{m',m}$ identifies $\wedge^n(T_m^*M)$ with $\wedge^n(T_{m'}^*M)$, the source fibre $\text{Det}(M)_{*,m}$ is isomorphic to the determinant line bundle $\wedge^n(T^*M) \rightarrow M$. Its⁴ connected component $\wedge^{n,+}(T^*M)$ is the bundle of positive top forms if M is orientable, and the whole bundle otherwise.

Its universal covering space is the bundle $\wedge^{n,+}(T^*\tilde{M})$ of positive top forms on \tilde{M} . Indeed, \tilde{M} is always orientable, regardless of whether or not M is. This means that $\wedge^n(T^*\tilde{M})$ is a trivial bundle, and that its connected component $\wedge^{n,+}(T^*\tilde{M}) \simeq \tilde{M} \times \mathbb{R}^+$ is simply connected. The covering map is induced by the map $\tilde{M} \rightarrow M$. This leads to the following version of theorem 7.14.

Proposition 7.20 *Let P be a principal G -bundle over an n -dimensional manifold M . Let G be such that its Lie algebra \mathfrak{g} does not contain $\mathfrak{sl}(\mathbb{R}^n)$ if $n > 1$, or $[X, Y] = Y$ if $n = 1$. Then there is a homomorphism $\pi_1(M) \times \mathbb{R}^+ \rightarrow G$ associating P to the principal $\pi_1(M) \times \mathbb{R}^+$ -bundle $\wedge^{n,+}(T^*\tilde{M}) \rightarrow M$.*

$$P \simeq \wedge^{n,+}(T^*\tilde{M}) \times_{\pi_1(M) \times \mathbb{R}^+} G.$$

We may even classify the possible splittings.

Corollary 7.21 *Under the hypotheses of proposition 7.20, any Lie-algebra homomorphism $\sigma : \Gamma_c(TM) \rightarrow \Gamma_c(TP)^G$ which splits the sequence of Lie algebras (7.2) can be written*

$$\sigma = \nabla^\mu + \Lambda \text{Div}_\mu, \quad (7.14)$$

where ∇^μ is a flat equivariant connection on P , and Λ a section of $\text{ad}(P)$ which is constant w.r.t. the connection induced on $\text{ad}(P)$ by ∇^μ .

Remark In particular, this shows that there exists a flat connection which splits (7.2), even though most splittings are not flat connections.

Proof. First, we prove the case $P = \wedge^{n,+}(T^*\tilde{M})$. Pick a nonzero (pseudo-) density μ on M . This induces an honest density $\tilde{\mu}$ on \tilde{M} , which in turn identifies $\wedge^{n,+}(T^*\tilde{M})$ with $\tilde{M} \times \mathbb{R}^+$. The local trivializations of $\wedge^{n,+}(T^*\tilde{M}) \rightarrow \tilde{M}$ and $\tilde{M} \rightarrow M$ combine to locally trivialize $\wedge^{n,+}(T^*\tilde{M}) \rightarrow M$. This yields a flat equivariant connection ∇^μ on $\wedge^{n,+}(T^*\tilde{M})$, which annihilates $\tilde{\mu}$.

The splitting σ is uniquely determined by the action of $\sigma(v)$ on local sections $\tilde{\nu}$, which reads $\sigma(v)(\tilde{\nu}) = \pi^* \circ \mathcal{L}_v \circ \pi^{*-1} \tilde{\nu}$, where π is the map from \tilde{M} to M . If we define the divergence w.r.t. μ by the requirement that the Lie derivative

⁴An isomorphism $\wedge^n(T^*M) \simeq \text{Det}(M)_{*,m}$ is only given after a choice of $\lambda_0 \in \wedge^n(T_m^*M)$. This determines the connected component.

$\mathcal{L}_v\mu$ equal $\text{Div}_\mu(v)\mu$, then we have

$$\begin{aligned}\sigma(v)(f\tilde{\mu}) &= \pi^*(\mathcal{L}_v(f\mu)) \\ &= \pi^*(v(f)\mu + \text{Div}_\mu(v)f\mu) \\ &= v(f)\tilde{\mu} + \text{Div}_\mu(v)f\tilde{\mu} \\ &= \nabla_v^\mu(f\tilde{\mu}) + \text{Div}_\mu(v)f\tilde{\mu}.\end{aligned}$$

This shows that $\sigma(v) = \nabla_v^\mu + \Lambda\text{Div}_\mu(v)$, with $\Lambda = \partial_r$, the equivariant vertical vector field defined by the action of \mathbb{R}^+ on $\wedge^n(T^*\tilde{M})$. (Equivariant vertical vector fields on P correspond to sections of $\text{ad}(P)$.)

The general case follows by proposition 7.20. \square

7.6.2 Lie Algebra Cohomology

If we specialize to the case of a trivial bundle over an Abelian group G , we find ourselves in the realm of Lie algebra cohomology. The continuous cohomology of the Lie algebra of vector fields with values in the functions has already been unravelled in all degrees [Fuk86]. Corollary 7.21 describes this cohomology only in degree 1, but now with all cocycles rather than just the continuous ones.

Corollary 7.22 *Let H_{LA} denote Lie algebra cohomology and H_{dR} de Rham cohomology. Let \mathfrak{g} be Abelian, and consider the representation $C_c^\infty(M, \mathfrak{g})$ of $\Gamma_c(TM)$ where a vector field v acts by the Lie derivative \mathcal{L}_v . Then*

$$H_{LA}^1(\Gamma_c(TM), C_c^\infty(M, \mathfrak{g})) \simeq H_{dR}^1(M, \mathfrak{g}) \oplus \mathfrak{g}.$$

Proof. Consider the trivial bundle $M \times G \rightarrow M$ over an Abelian Lie group G , which comes equipped with a flat connection ∇^0 , which acts as Lie derivative. Note that Abelian \mathfrak{g} certainly satisfy the conditions of propositions 7.20 and 7.21. View $\Gamma_c(\text{ad}(P)) \simeq C_c^\infty(M, \mathfrak{g})$ as a representation of $\Gamma_c(TM)$, and consider its Lie algebra cohomology. An n -cochain is an alternating linear map $\Gamma_c(TM)^n \rightarrow C_c^\infty(M, \mathfrak{g})$. For $f^1 \in C^1$, closure $\delta f^1 = 0$ amounts to

$$\mathcal{L}_v f^1(w) - \mathcal{L}_w f^1(v) - f^1([v, w]) = 0.$$

Due to this cocycle condition, $\sigma = \nabla^0 + f^1$ is once again a Lie algebra homomorphism splitting π_* . According to corollary 7.21, it must therefore take the shape $\sigma = \nabla^\mu + \Lambda\text{Div}_\mu$, where $\Lambda \in \mathfrak{g}$ is constant. One can write $\nabla^\mu = \nabla^0 + \omega^1$ for some closed 1-form ω^1 , so that $f^1 = \omega^1 + \Lambda\text{Div}_\mu$. This classifies the closed 1-cocycles.

Exact 1-cocycles satisfy $f^1(v) = \delta f^0(v) = \mathcal{L}_v f^0 = df^0(v)$, with f^0 a 0-cocycle, that is an element of $C^\infty(M, \mathfrak{g})$.

Note that a change of density $\mu' = e^h\mu$ alters f^1 by a mere coboundary Λdh , so that the choice of μ is immaterial. The class of $\omega^1 + \Lambda\text{Div}_\mu$ modulo δC^0 is therefore determined by $[\omega^1] \in H_{dR}^1(M, \mathfrak{g})$ and $\Lambda \in \mathfrak{g}$. \square

Continuity turns out to be implied by the closedness-condition. A similar situation was encountered by Takens in [Tak73], when proving that all derivations of $\Gamma_c(TM)$ are inner, i.e. $H_{LA}^1(\Gamma_c(TM), \Gamma_c(TM)) = 0$.

7.7 General relativity, Gauge theory, and Spinors

In this section, we briefly reflect on the relationship between (generalized) spin structures and principal fibre bundles with a splitting of (7.2).

We formulate general relativity and gauge theory, including the fermionic fields, in terms of principal fibre bundles over the manifold M which describes space-time. For convenience, we will take M to be a smooth and orientable manifold of dimension at least 2.

7.7.1 General Relativity

The fundamental degrees of freedom in general relativity are a pseudo-Riemannian metric g on space-time M , and a connection ∇ on TM .

We identify the metric with a section of $F^+(M)/\text{SO}(\eta)$ by associating to $g_x : T_x M \times T_x M \rightarrow \mathbb{R}$ the coset of all frames $f : \mathbb{R}^n \rightarrow T_x M$ such that $f^*g = \eta$. We identify the connection on TM with an equivariant connection on $F^+(M)$, that is a section of $J^1(F^+(M))/\text{GL}^+(\mathbb{R}^n) \rightarrow M$. (Its value at x is the class of a 1-jet at x of a section $\phi \in \Gamma(F^+)$ with $\nabla_x \phi = 0$ [Sar09].) These two, the metric g and the connection ∇ , are conveniently combined into a single section of the fibre bundle $J^1(F^+(M))/\text{SO}(\eta)$.

The dynamics of the theory are then governed by the Einstein-Hilbert action $S_{EH} : \Gamma(J^1(F^+(M))/\text{SO}(\eta)) \rightarrow \mathbb{R}$, defined in terms of the Ricci scalar R by

$$S_{EH}(g, \nabla) = \int_M R(g, \nabla) \sqrt{|g|} dx_0 \dots dx_n.$$

In the Einstein-Hilbert approach, the connection is constrained to equal the Levi-Civita connection, and only the metric is varied. In the Palatini approach [Ash91], the connection varies independently, and the fact that ∇ is the Levi-Civita connection of the metric is a consequence of the field equations.

Either way, it is clear that the fields transform in a natural fashion under diffeomorphisms of M . Any orientation-preserving diffeomorphism α of M lifts to an automorphism $\Sigma(\alpha)$ of $J^1(F^+(M))/\text{SO}(\eta)$, defined by $\Sigma(\alpha)(j_m^1(\phi)) = j_m^1(\alpha_* \circ \phi \circ \alpha^{-1})$. This splits the sequence of groups

$$1 \rightarrow \text{Aut}^V(J^1(F^+(M))/\text{SO}(\eta)) \rightarrow \text{Aut}(J^1(F^+(M))/\text{SO}(\eta)) \rightarrow \text{Diff}^+(M) \rightarrow 1.$$

This splitting is central to the theory of general relativity. The requirement that the action be invariant under co-ordinate transformations, $S_{EH}(\Sigma(\alpha)\phi) = S_{EH}(\phi)$, cannot even be formulated without providing Σ explicitly.

Note that as the above sequence of groups splits, so does the corresponding sequence of Lie algebras of vector fields.

7.7.2 Fermions and Spin Structures

We wish to describe fermions. As these are known to transform under Lorentz transformations by a projective representation rather than a linear one, we must extend our framework.

The Bundle

Suppose that we have a spin structure Q w.r.t a background metric g . Let $u : \hat{Q} \rightarrow F^+(M)$ be as in section 7.4.3, and let V be a unitary spinor representation⁵ of $\widetilde{SO}(\eta)$. Then one has the composite bundle

$$\hat{Q} \times_{\widetilde{SO}(\eta)} V \rightarrow F^+(M)/SO(\eta) \rightarrow M.$$

A section $\tau : M \rightarrow \hat{Q} \times_{\widetilde{SO}(\eta)} V$ can then be interpreted as a metric g along with a fermionic field ψ . Consider g as the induced section of $F^+(M)/SO(\eta)$. Use g to construct the spinor bundle $u^{-1}(OF_g^+) \times_{\widetilde{SO}(\eta)} V$, and obtain a section ψ by simply restricting the image of τ .

In the same vein, we will describe physical fields by sections of the fibre bundle $J^1(\hat{Q}) \times_{\widetilde{SO}(\eta)} V$. This is equivalent to providing a triple of sections: one of $F^+(M)/SO(\eta)$, one of $J^1(F^+(M))/GL^+(\mathbb{R}^n)$, and one of $u^{-1}(OF_g^+) \times_{\widetilde{SO}(\eta)} V$. These correspond to the metric $g_{\mu\nu}$, the Levi Civita-connection $\Gamma_{\mu\beta}^\alpha$, and the fermionic field ψ^i respectively.

Transformation Behaviour

Let us investigate its transformation behaviour. As a spinor changes sign under a 2π -rotation, there is no hope of finding an interesting homomorphism of groups $\text{Diff}^+(M) \rightarrow \text{Aut}(J^1(\hat{Q}) \times_{\widetilde{SO}(\eta)} V)$. There is however a canonical homomorphism of Lie algebras.

Because $u : \hat{Q} \rightarrow F^+(M)$ has discrete fibres, it has a unique flat equivariant connection ∇ . This means that the exact sequence of Lie algebras

$$0 \rightarrow \Gamma_c(\text{ad}(\hat{Q})) \rightarrow \Gamma_c(T\hat{Q}) \xrightarrow{\widetilde{GL}^+(\mathbb{R}^n)} \Gamma_c(TM) \rightarrow 0$$

is split by $\sigma := \nabla \circ D$, with $D : \Gamma_c(TM) \rightarrow \Gamma_c(TF^+(M))^{\text{GL}^+(\mathbb{R}^n)}$ the first order derivative $D(v)_f = \partial/\partial t|_0 \exp(tv)_* \circ f$. This induces a splitting for $J^1(\hat{Q})$ by prolongation (see e.g. [FR04]), and consequently also one for $J^1(\hat{Q}) \times_{\widetilde{SO}(\eta)} V$. Note that $u_* \circ \sigma$ equals D .

We see that \hat{Q} is an infinitesimally natural bundle, and that the canonical splitting σ of (7.2) does not come from a splitting of groups. We would like to emphasize that even if a splitting at the level of groups does exist, it will not be physically relevant, because it cannot reproduce the minus sign under a full rotation that one knows and loves in fermions.

⁵The indefinite article is appropriate since there is a choice involved here. The connected component of $\mathbf{1}$ of $\widetilde{SO}(3, 1)$ is $\text{Spin}^1(3, 1) \simeq \text{SL}^2(\mathbb{C})$. A spinor representation for the connected component can then be unambiguously derived from a Clifford algebra representation [Her74]. But as $\widetilde{SO}(3, 1)$ is not isomorphic to $\text{Spin}(3, 1)$, the action of the order 2 central elements covering PT will have to be specified 'by hand'. Again, this is not relevant if M is time-orientable as well as orientable.

Take for example the spin structure $Q = \mathbb{R}^n \times \widetilde{SO}(\eta)$ over \mathbb{R}^n , and lift $\alpha \in \text{Diff}^+(M)$ to $\text{Aut}(Q)$ by $\Sigma(\alpha)(m, q) = (\alpha(m), q)$. This the wrong thing to do: if we restrict attention to $SO(\eta) < \text{Diff}^+(M)$, we see that sections of $Q \times_{\widetilde{SO}(\eta)} V$ transform under the trivial representation of the Lorentz group. Our fermions are Lorentz scalars rather than spin-1/2 particles. This is a very real problem: using the wrong splitting will generally result in an incorrect energy-momentum tensor [GM92].

We conclude that not only the bundle Q and the covering map $u : Q \rightarrow OF_g$ are relevant, but also the splitting $\sigma : \Gamma_c(TM) \rightarrow \Gamma_c(T\hat{Q})^{\widetilde{GL}^+(\mathbb{R}^n)}$. It must satisfy $u_* \circ \sigma = D$ in order for the metric $g \in \Gamma(TF^+(M)/SO(\eta))$ to transform properly. Luckily, such a σ is naturally associated to any ordinary spin structure Q . For generalized spin structures however, this is no longer so.

7.7.3 Gauge Fields and Spin^G-structures

In the presence of gauge fields, the topological conditions on M in order to support a spin structure are more relaxed. Intuitively, this is because the gauge group G can absorb some of the indeterminacy that stems from the 2:1 cover of the Lorentz group.

Generalized Spin Structures

This is made more rigorous by the notion of a generalized spin structure or Spin^G-structure [AI80].

Definition 7.9 *Let G be a Lie group with a central subgroup isomorphic to Z . Let⁶ $\text{Spin}^G := \widetilde{SO}(\eta) \times_Z G$. A Spin^G-structure is a Spin^G-bundle Q over M , together with a map $u : Q \rightarrow OF_g^+$ that makes*

$$\begin{array}{ccc}
 \text{Spin}^G & \xrightarrow{\kappa} & \text{SO}(\eta) \\
 \widehat{Q} & \xrightarrow{u} & \widehat{OF}_g^+ \\
 & \searrow & \swarrow \\
 & M &
 \end{array}$$

commute. We again denote the map $(x, g) \mapsto \kappa(x)$ by κ .

This gives rise to the principal $\widetilde{GL}(\mathbb{R}^n) \times_Z G$ -bundle $\hat{Q} := Q \times_{\widetilde{SO}(\eta)} \widetilde{GL}(\mathbb{R}^n)$. If $G = Z$, we recover the notion of a spin structure. Apart from spin structures, the best known examples of Spin^G-structures are probably the spin^c-structures, which are precisely Spin^G-structures for $G = U(1)$.

Let V be a representation of Spin^G. The ‘physical bundle’ is then the fibre $J^1(\hat{Q}) \times_{\text{Spin}^G} V$. A single section of $J^1(\hat{Q}) \times_{\text{Spin}^G} V$ represents a metric, a Levi Civita-connection, a gauge field and a fermionic field.

⁶This notation is convenient but slightly misleading. Be ware that if η is of signature $+++ -$, then Spin^Z is isomorphic to $\widetilde{SO}(\eta)$, not to Spin(3, 1).

The classification theorem for infinitesimally natural Spin^G -structures will take the following form.

Theorem 7.23 *Let M be an orientable smooth manifold of dimension $n \geq 2$. Let G be a Lie group with a central subgroup Z isomorphic to $\pi_1(\text{GL}(\mathbb{R}^n))$, and let $\text{Lie}(G)$ be such that it does not contain $\mathfrak{sl}(\mathbb{R}^n)$. Finally, let (Q, u) be an infinitesimally natural Spin^G -structure over M . Then $i_* : Z \rightarrow \pi_1(F^+(M))$ is injective, and there exists a homomorphism $\tau : \pi_1(F^+(M)) \rightarrow G$ which fixes Z , and makes (Q, u) isomorphic to the bundle*

$$\widetilde{OF}_g^+ \times_\tau G,$$

with u the natural projection map onto OF_g^+ .

Remark We have already seen that an ordinary spin structure is always infinitesimally natural, so that proposition 7.16 is just the special case $G = Z$.

Theorem 7.23 rather simplifies the data needed to construct the relevant bundle $J^1(\hat{Q}) \times_{\text{Spin}^G} V$. Indeed, it suffices to have a triple $(M, G, (\rho, V))$ of:

- An orientable manifold M such that $i_* : \pi_1(\text{SO}(\eta)) \rightarrow \pi_1(F(M))$ is injective. Its image $Z = \{\pm 1\}$ is then central in $\pi_1(F(M))$.
- A representation ρ of $\widetilde{\text{SO}}(\eta) \times_Z \pi_1(F(M))$ on V which is unitary when restricted to $\pi_1(F(M))$, and faithful on Z .
- A subgroup $G < U(V)$ that commutes with the image of $\widetilde{\text{SO}}(\eta)$ under ρ , and contains the image of $\pi_1(F(M))$.

We can then form $\hat{Q} = \widetilde{F}^+(M) \times_{\pi_1(F)} G$, with structure group $\text{Spin}^G = \widetilde{\text{SO}}(\eta) \times_Z G$. The relevant bundle is $J^1(\hat{Q}) \times_{\text{Spin}^G} V$, a single section providing the metric, Levi-Civita connection, gauge field and fermions.

Note that the ‘gauge bundle’ is the trivial bundle $P = M \times G/Z$, and that the bundle describing fermions and metric is $\widetilde{F}^+(M) \times_H V$, where H is defined as $H := \widetilde{\text{SO}}(\eta) \times_Z \pi_1(F)$.

According to theorem 7.23, the above setting exhausts the possibilities— at least under the (natural) assumptions that Q is infinitesimally natural, and that V is a faithful unitary representation for G . Indeed, G must then be compact, so that its Lie-algebra cannot allow $\mathfrak{sl}(\mathbb{R}^n)$ as a subalgebra.

Two Lemmas

We set out to prove theorem 7.23. We start with two lemmas designed to explicate the homomorphism $\rho : G(k, M) \rightarrow \widetilde{\text{GL}}(\mathbb{R}^n) \times_Z G$.

Lemma 7.24 *Denote the natural projection map $\widetilde{F}^{+k}(M) \rightarrow F^+(M)$ by π , and write $\nu : G(k, M) \rightarrow \text{GL}^+(\mathbb{R}^n)$ for the corresponding homomorphism of groups. Then each infinitesimally natural Spin^G -structure is isomorphic to one for which $\kappa \circ \rho = \nu$ and $u \circ \exp \nabla = \pi$.*

Proof. Denote $u \circ \exp \nabla$ by γ . Pick $\tilde{f} \in \widetilde{F}_x^{+k}(M)$, and let $\pi(\tilde{f}) = f$. There exists a $c \in \mathrm{GL}^+(\mathbb{R}^n)$ such that $\gamma(\tilde{f}) = fc$. Equivariance of γ then implies $\gamma(\tilde{f}\tilde{y}) = fc\kappa(\rho(\tilde{y}))$ for all $\tilde{y} \in G(k, M)$.

We required $\gamma_* \circ \tilde{D} = D$, but clearly we also have $\pi_* \circ \tilde{D} = D$. Therefore, if $\exp(\tilde{D}(v))\tilde{f} = \tilde{f}\tilde{y}$ and $\exp(D(v))f = fy$, we must have both $\nu(\tilde{y}) = y$ and $\kappa(\rho(\tilde{y})) = c^{-1}yc$.

We may obtain all \tilde{y} in $G(k, M)^0$, the connected component of $G(k, M)$, by choosing an appropriate v , so that we have $\kappa(\rho(\tilde{y})) = c^{-1}\nu(\tilde{y})c$ for all $\tilde{y} \in G(k, M)^0$. Although c depends on \tilde{f} a priori, it turns out to be constant up to scaling. Indeed, as both κ and ρ are constant, so is the adjoint action of c . Moreover, $\tilde{f} \mapsto c(\tilde{f})$ is invariant under $G(k, M)^0$, making it a function on $\widetilde{F}^{+k}(M)/G(k, M)^0 \simeq \tilde{M}$. All in all, we have established that there exist $c_0 \in \mathrm{SL}(\mathbb{R}^n)$ and $h : \tilde{M} \rightarrow \mathbb{R}^+$ such that $c(\tilde{f}) = c_0h(\tilde{x})$, with \tilde{x} the projection of \tilde{f} to \tilde{M} . We may write $u \circ \exp \nabla(\tilde{f}) = \pi(f)c_0h(\tilde{x})$.

We show that we may as well take c and h to be 1. Pick a $\tilde{c} \in \widetilde{\mathrm{SL}}(\mathbb{R}^n)$ which covers c , and construct the bundle $\hat{Q}_c := \hat{Q} \times_{\mathrm{Ad}(\tilde{c})} \widetilde{\mathrm{GL}}(\mathbb{R}^n)$. It is isomorphic to \hat{Q} , with isomorphism $\hat{Q}_c \rightarrow \hat{Q}$ given by $[q, y] \mapsto q\tilde{c}^{-1}h^{-1}(\tilde{x})y$.

If we simply pull back the covering map on Q , we obtain $u_c : \hat{Q}_c \rightarrow F^+(M)$ given by $[q, y] \mapsto u(q\tilde{c}^{-1}h^{-1}(\tilde{x})y)$. This makes $Q_c := u_c^{-1}(OF_g)$ into a Spin^G -structure isomorphic to Q , but with the desired properties. \square

Recall from (7.9) that $G(k, M) \simeq G(k, \mathbb{R}^n) \times_Z \pi_1(F^+(M))$, and that we have $G(k, \mathbb{R}^n) \simeq \widetilde{\mathrm{GL}}^+(\mathbb{R}^n) \ltimes G^{>1}$, with $G^{>1}$ the subgroup of k -jets that are the identity to first order. We unravel ρ , considering it as a map

$$\widetilde{\mathrm{SL}}(\mathbb{R}^n) \times \mathbb{R}^+ \ltimes G^{>1} \times_Z \pi_1(F^+(M)) \rightarrow \widetilde{\mathrm{SL}}(\mathbb{R}^n) \times \mathbb{R}^+ \times_Z G.$$

Lemma 7.25 *Under the assumptions of theorem 7.23, the map ρ is completely determined by a homomorphism $\bar{\rho}_1 : \pi_1(F^+(M)) \rightarrow G$ fixing Z , a homomorphism $\bar{\rho}_2 : \pi_1(F^+(M)) \rightarrow \mathbb{R}^+$, and an element Λ of $\mathrm{Lie}(G)$ which commutes with $\mathrm{Im}(\bar{\rho}_1)$. We have*

$$\rho : (\tilde{x}, e^t, g, [p]) \mapsto (\tilde{x}, e^t\bar{\rho}_2([p]), e^{t\Lambda}\bar{\rho}_1([p])).$$

Proof. Consider $\dot{\rho}$ as a Lie algebra homomorphism from $\mathfrak{sl}(\mathbb{R}^n) \times \mathbb{R} \ltimes \mathfrak{g}^{>1}$ to $\mathfrak{sl}(\mathbb{R}^n) \times \mathbb{R} \times \mathrm{Lie}(G)$, and let $\dot{\rho}_{ij}$ be its (i, j) component for $i, j \in \{1, 2, 3\}$.

Because $\mathfrak{sl}(\mathbb{R}^n)$ is simple and not contained in $\mathrm{Lie}(G)$, we must have $\dot{\rho}_{13} = 0$. Due to the previous lemma, $\dot{\rho}_{12} = 0$ and $\dot{\rho}_{11} = \mathrm{id}$. As $\widetilde{\mathrm{SL}}(\mathbb{R}^n)$ is simply connected, we must have $(\tilde{x}, 1, 1, 1) \mapsto (\tilde{x}, 1, 1)$. In particular, this forces the map $i_* : Z \rightarrow \pi_1(F^+(M))$ to be injective.

Again due to the previous lemma, $\dot{\rho}_{21} = 0$ and $\dot{\rho}_{22} = \mathrm{id}$. Define the ‘scaling element’ $\Lambda := \dot{\rho}_{23}(1)$. Then $(1, e^t, 1, 1) \mapsto (1, e^t, e^{t\Lambda})$. The image of $\pi_1(F^+(M))$ must commute with $(\tilde{x}, e^t, e^{t\Lambda})$, so that $\bar{\rho} : \pi_1(F^+(M)) \rightarrow \mathbb{R}^+ \times G$ is well defined, and Λ commutes with its image.

We now show that $\dot{\rho}(\mathfrak{g}^{>1}) = 0$. First of all, as $[\mathfrak{sl}(\mathbb{R}^n), \mathfrak{sl}(\mathbb{R}^n) + \mathfrak{g}^{>1}]$ equals $\mathfrak{sl}(\mathbb{R}^n) + \mathfrak{g}^{>1}$ if n is at least 2 (see Lemma 7.19), we must have

$$\dot{\rho}(\mathfrak{g}^{>1}) \subset [\dot{\rho}(\mathfrak{sl}(\mathbb{R}^n)), \dot{\rho}(\mathfrak{sl}(\mathbb{R}^n) + \mathfrak{g}^{>1})] \subset \mathfrak{sl}(\mathbb{R}^n) \oplus 0 \oplus 0.$$

But on the other hand, as $[\mathbb{R}, \mathfrak{g}^{>1}] = \mathfrak{g}^{>1}$, (recall that \mathbb{R} represents the Euler vector field), we have

$$\dot{\rho}(\mathfrak{g}^{>1}) = [\dot{\rho}(\mathbb{R}), \dot{\rho}(\mathfrak{g}^{>1})] \subset 0 \oplus 0 \oplus \text{Lie}(G).$$

The intersection being zero, we have $\dot{\rho}(\mathfrak{g}^{>1}) = 0$. But then $\rho(G^{>1}) = 1$, because $G^{>1}$ is simply connected. \square

Proof of Theorem 7.23

Clearly, \mathbb{R}^+ cannot cause any topological obstruction, so we should be able to eliminate both $\bar{\rho}_2$ and Λ from the story. The former is easy. recall that we seek Q , not \hat{Q} . As Q is a subbundle of \hat{Q}/\mathbb{R}^+ , we will focus on the latter from now on, allowing us to simply disregard $\bar{\rho}_2$.

In order to remove Λ , we choose a volume form λ on M . This endows each frame $f \in F_x(M)$ with a volume $\text{vol}_\lambda(f)$. A frame has volume 1 precisely when it is the jet of a diffeomorphism which preserves λ . Denote by $F^\lambda(M)$ the $\text{SL}(\mathbb{R}^n)$ -bundle of frames with volume 1. Its universal cover $\tilde{F}^\lambda(M)$ is its inverse image under π , a principal $\widetilde{\text{SL}}(\mathbb{R}^n) \times_Z \pi_1(F^+(M))$ -bundle.

Define the isomorphism

$$\hat{Q} \rightarrow \tilde{F}^\lambda(M) \times_{\pi_1(F^+(M))} (G \times \mathbb{R}^+)$$

by

$$(\tilde{f}, g) \mapsto \left(\tilde{f} \text{vol}_\lambda^{-1}(\pi(\tilde{f})), \exp(\log(\text{vol}_\lambda(\pi(\tilde{f}))) \Lambda)g, \text{vol}_\lambda(\pi(\tilde{f})) \right)$$

where we consider \hat{Q} as $(\tilde{F}^{+k}(M) \times_{\bar{\rho}_1 \bar{\rho}_2} G)$. One can see that it is well defined, and that it intertwines the natural maps to F/\mathbb{R}^+ .

This shows that \hat{Q}/\mathbb{R}^+ , and therefore the spin structure (Q, u) , is completely determined by the homomorphism $\bar{\rho}_1 : \pi_1(F^+(M)) \rightarrow G$. We denote it by τ from now on.

If we choose λ to be the volume form induced by the metric g , then OF_g^+ is a subbundle of $F_\lambda(M)$. Since $\pi_1(F^+(M)) = \pi_1(OF_g^+)$, we have $Q \simeq \widetilde{OF}_g^+ \times_\tau G$. The spin map u is simply the projection $\widetilde{OF}_g^+ \rightarrow OF_g^+$. The principal G -bundle $Q \rightarrow OF_g$ has a flat equivariant connection induced by the one on $\widetilde{OF}_g^+ \rightarrow OF_g^+$. This concludes the proof of theorem 7.23. \square

Tracking back through the isomorphisms, we can formulate the following.

Corollary 7.26 *Under the assumptions of theorem 7.23, there exists an element $\Lambda \in \text{Lie}(G)$ which commutes with the image of τ , such that the splitting $\sigma : \Gamma(TM) \rightarrow \Gamma(T\hat{Q}/\mathbb{R}^+) \widetilde{\text{SL}}(\mathbb{R}^n) \times_Z G$ is given by*

$$\sigma(v) = \nabla \circ \bar{D}(v) + \text{Div}_\lambda(v)\Lambda,$$

where λ is the volume form induced by g , \bar{D} is the natural lift from $\Gamma(TM)$ to $\Gamma(\tilde{F}^+(M)/\mathbb{R}^+)$, and we have identified Λ with the vector field on \hat{Q}/\mathbb{R}^+ induced by the action of the Lie algebra element.

It is clear that two different homomorphisms τ_1 and τ_2 yield isomorphic Spin^G -structures if one is obtained from the other by conjugation within G .

7.7.4 Induced Unitary Representations of $\widetilde{\text{Diff}}^0(M)$

Before we explore the ramifications of theorem 7.23 on some theories of physical interest, we pause for a second to examine the representations of $\widetilde{\text{Diff}}^0(M)$ that live on infinitesimally natural bundles of the type described on page 180.

Recall that a metric is precisely a section of the bundle of symmetric spaces $\pi : R(M) \rightarrow M$, with $R(M) := F^+(M)/\text{SO}(\eta)$. The bundle describing fermions and metric is $E(V) := \widetilde{F}^+(M) \times_{\text{Spin}^Z} V$, with V a representation of $\text{Spin}^Z = \widetilde{\text{SO}}(\eta) \times_Z \pi_1(F(M))$ of spinor type. The projection $E(V) \rightarrow R(M)$ makes it into a vector bundle over $R(M)$, and the projection $E(V) \rightarrow M$ makes it into a fibre bundle, unfortunately not a vector bundle, over M .

If we fix M to be \mathbb{R}^n , and g to be the Minkowski metric η , then the action of the connected Poincaré group $\text{SO}^0(\eta) \ltimes \mathbb{R}^n$ on $OF_\eta^+(M)$ induces an action (by automorphisms) of its cover on $\widetilde{OF}_\eta^+(M)$. This in turn induces a representation on the sections of the associated vector bundle $\widetilde{OF}_\eta^+(M) \times_{\widetilde{\text{SO}}^0(\eta)} V$, the space of spinors. The Minkowski metric gives rise to a unique (up to a scalar) Poincaré-invariant volume form vol_η on M , which we use to turn the nondegenerate Lorentz-invariant form on V into the nondegenerate Poincaré-invariant form $\langle \psi | \psi' \rangle = \int_M (\psi, \psi') \text{vol}_\eta$ on the space of sections. The group action commutes with the Dirac operator, and the representation therefore decomposes according to mass.

We would like to do something similar for the diffeomorphism group instead of the Poincaré group. The action of $\text{Diff}^0(M)$ on $F^+(M)$ induces an action by automorphisms of $\widetilde{\text{Diff}}^0(M)$ on $\widetilde{F}^+(M)$, and therefore also on $R(M)$ and $E(V)$. Unfortunately, the sections $\Gamma(E(V) \rightarrow M)$ of physical relevance do not form a representation, or even a vector space for that matter.

We do however have a representation of $\widetilde{\text{Diff}}^0(M)$ on $\Gamma(E(V) \rightarrow R(M))$. We will show that $R(M)$ carries a unique (up to a scalar) $\text{Diff}^0(M)$ -invariant volume form⁷ vol , so that $\Gamma(E(V) \rightarrow R(M))$ has a nondegenerate $\widetilde{\text{Diff}}^0(M)$ -invariant form $\langle \phi | \psi \rangle = \int_{R(M)} (\phi, \psi) \text{vol}$.

An Invariant Volume

The only special asset of $\text{SO}(\eta)$ that we will use is the fact that it is an open subgroup of the fixed point set of the involution $g \mapsto \eta(g^t)^{-1}\eta$ of $\text{SL}(\mathbb{R}^n)$, so we may as well work in that setting.

Proposition 7.27 *Let M be a connected n -dimensional smooth manifold, and let $H < \text{SL}(\mathbb{R}^n)$ be an open subgroup of the fixed point set of an involution. Then the bundle $R(M) = F^+(M)/H$ of symmetric spaces admits a diffeomorphism-invariant volume form vol . It is everywhere nonzero, and unique up to \mathbb{R}^\times .*

Proof. We write \mathfrak{h} for the Lie algebra of H . The decomposition $\mathfrak{gl}(\mathbb{R}^n) = \mathfrak{h} \oplus \mathfrak{m}$ of $\mathfrak{gl}(\mathbb{R}^n)$ into a positive and a negative eigenspace of the Lie algebra involution

⁷It would be tempting to interpret the volume form $(\psi, \psi) \text{vol}$ on $R(M)$ as the probability density of observing a fermion at a certain point of space-time and with a certain value of the metric, but as far as I am aware, this is not what physicists do.

is invariant under the adjoint action of \mathfrak{h} . In particular, \mathfrak{m} is an \mathfrak{h} -representation. We construct a canonical map

$$\wedge^n(\mathbb{R}^{n*}) \otimes \wedge^k(\mathfrak{gl}(\mathbb{R}^n)^*)^{\mathfrak{h}} \xrightarrow{\sim} \Omega^{n+k}(R(M))^{\text{Diff}^0(M)},$$

with $k = \dim(\mathfrak{m})$. Choose $\lambda \in \wedge^n(\mathbb{R}^{n*})$ and $\mu \in \wedge^k(\mathfrak{gl}(\mathbb{R}^n)^*)^{\mathfrak{h}}$. The former is a volume form on \mathbb{R}^n , the latter can be thought of as a volume form on $\mathfrak{gl}(\mathbb{R}^n)/\mathfrak{h} \simeq \mathfrak{m}$. We construct the volume form on $R(M)$ over each point $[f] \in R(M)$ separately, using the isomorphism $T_f F^+(M)/i_f(\mathfrak{h}) \xrightarrow{\sim} T_{[f]}R(M)$ induced by the quotient map $F^+(M) \rightarrow R(M)$.

Each tangent space $T_f F^+(M)$ fits into the exact sequence of vector spaces

$$0 \longrightarrow \mathfrak{gl}(\mathbb{R}^n) \xrightarrow{i_f} T_f F^+(M) \xrightarrow{\theta_f} \mathbb{R}^n \longrightarrow 0,$$

where $i_f(\xi) := \frac{d}{dt}|_0 f \circ e^{t\xi}$, and $\theta_f = f^{-1} \circ \pi_*$.

The n -form $\theta_f^* \lambda$ annihilates the image of i_f , so in particular, it defines an n -form on $T_{[f]}R(M)$. This is in fact independent of the choice of f representing $[f]$. The two identifications $T_{[f]}R(M) \simeq T_f F^+(M)/i_f(\mathfrak{h})$ and $T_{[f]}R(M) \simeq T_{fh} F^+(M)/i_{fh}(\mathfrak{h})$ are intertwined by R_{h*} , the pushforward of right multiplication by h . We have $\theta_{fh} \circ R_{h*} = h^{-1} f^{-1} \pi_* R_{h*} = h^{-1} f^{-1} \pi_* = h^{-1} \circ \theta_f$, so that $(R_{h*})^* \theta_{fh}^* \lambda = \theta_f^* h^{-1*} \lambda$. Because $h^{-1*} \lambda = \lambda$, ($h \in H < \text{SL}(\mathbb{R}^n)$ is volume preserving), we even have $(R_{h*})^* \theta_{fh}^* \lambda = \theta_f^* \lambda$. Since the two forms $\theta_f^* \lambda$ and $\theta_{fh}^* \lambda$ are related by R_{h*} , they define the same n -form on $T_{[f]}R(M)$.

In order to transport μ to $T_f F^+(M)$, we need to choose a splitting $s_f : T_f F^+(M) \rightarrow \mathfrak{gl}(\mathbb{R}^n)$ such that $s_f \circ i_f = \text{id}$. Consider the $n+k$ -form $s_f^* \mu \wedge \theta_f^* \lambda$. Although $s_f^* \mu$ depends on the choice of splitting, the form $s_f^* \mu \wedge \theta_f^* \lambda$ does not. Indeed, since $(s_f - s'_f) \circ i_f = 0$ for every two sections s_f and s'_f , we have $i_f^*(s_f^* \mu - s'^*_f \mu) = 0$, which implies that $(s_f^* \mu - s'^*_f \mu) \wedge \theta_f^* \lambda = 0$.

It is also independent of the choice of f . Since $R_{h*} \circ i_f(\xi) = \frac{d}{dt}|_0 f e^{t\xi} h = \frac{d}{dt}|_0 f h (h^{-1} e^{t\xi} h) = i_{fh} \circ \text{Ad}_{h^{-1}}(\xi)$, we obtain $s_f = \text{Ad}_h \circ s_{fh} \circ R_{h*}$, so that $s_f^* \mu = (R_{h*})^* s_{fh}^* \text{Ad}_h^* \mu$.

Note that Ad_h respects the nondegenerate bilinear form $\kappa(\xi, \chi) = \text{tr}(\xi\chi)$ on $\mathfrak{gl}(\mathbb{R}^n)$, and hence its restriction to \mathfrak{m} . Since $\text{Ad}_h : \mathfrak{m} \rightarrow \mathfrak{m}$ is orthogonal w.r.t. κ , it certainly preserves the volume form μ .

Now since $\text{Ad}_h^* \mu = \mu$, we have $s_f^* \mu = (R_{h*})^* s_{fh}^* \mu$, as required. This means that $\text{vol}_{[f]} = s_f^* \mu \wedge \theta_f^* \lambda$ is a well-defined $(n+k)$ -form on $T_{[f]}R(M)$ for any $[f]$. Since our construction varies smoothly with $[f]$, they combine into a smooth volume form vol on $R(M)$, which clearly depends only on $\lambda \otimes \mu$.

A diffeomorphism $\alpha \in \text{Diff}^0(M)$ induces an automorphism α_* of $F^+(M)$ by pushforward, and therefore also one of $R(M)$. The diffeomorphism invariance of the volume form is a straightforward consequence of the diffeomorphism invariance of $f \mapsto \theta_f$ and $f \mapsto i_f$. Indeed, invariance of the volume, $(\alpha_*)^* \text{vol} = \text{vol}$, is equivalent to $(\alpha_{**})^*(s_{\alpha_* f}^* \mu \wedge \theta_{\alpha_* f}^* \lambda) = s_f^* \mu \wedge \theta_f^* \lambda$. This is true because $\theta_{\alpha_* f} \circ \alpha_{**} = \theta_f$ and $s_{\alpha_* f} \circ \alpha_{**} = s_f$. The first can be seen by expanding the definitions and using $\alpha^{-1} \circ \pi \circ \alpha_* = \pi$, the second is a consequence of $\alpha_{**} \circ i_f(\xi) = \frac{d}{dt} \alpha_* f e^{t\xi} = i_{\alpha_* f}(\xi)$.

If M is connected, then $\text{Diff}^0(M)$ acts transitively on $R(M)$, so that the invariant volume is determined by its value in a single point. It is therefore unique up to a real scalar. The fact that it is nonzero is clear from the construction. \square

Remark The space $R(M)$ has more structure than just a canonical volume. In fact, its frame bundle has a canonical $\text{Diff}^0(M)$ -invariant subbundle with structure group $H \times \text{Lin}(\mathbb{R}^n, \mathfrak{m})$, with multiplication given by $(h, v)(h', v') = (hh', \text{Ad}_h \circ v' + v \circ h')$. This implies the existence of an invariant volume, because the structure group is contained in $\text{SL}(\mathbb{R}^n \times \mathfrak{m})$.

Now let \tilde{H} be the preimage of H under $G(1, M) \rightarrow \text{GL}^+(\mathbb{R}^n)$, an extension of H by $\pi_1(F(M))$. (Cf. def. 7.6.) If V is a representation of \tilde{H} , we construct the vector bundle $E(V) \rightarrow R(M)$ by $E(V) := \tilde{F}^+(M) \times_{\tilde{H}} V$.

The action of $\text{Diff}^0(M)$ on $F^+(M)$ induces an action by automorphisms of $\tilde{\text{Diff}}^0(M)$ on $\tilde{F}^+(M)$, and therefore also on $R(M)$ and $E(V)$. (It is just the integrated version of the natural lift of $\Gamma(TM)$.) This makes $\Gamma(E(V))$, the space of smooth sections of $E(V) \rightarrow R(M)$, into a representation of $\tilde{\text{Diff}}^0(M)$.

If V has an \tilde{H} -invariant (bilinear or sesquilinear) form (\cdot, \cdot) , then this induces a form on each fibre of $E(V) \rightarrow R(M)$. This defines a form $\langle \phi | \psi \rangle = \int_{R(M)} (\phi, \psi) \text{vol}$ on $\Gamma(E(V))$. Because vol is everywhere nonzero, the form $\langle \cdot | \cdot \rangle$ is positive definite and/or nondegenerate if and only if (\cdot, \cdot) is. Moreover, the form $\langle \cdot | \cdot \rangle$ is $\tilde{\text{Diff}}^0(M)$ -invariant because the volume is. It is perhaps worth while to summarize this in a corollary.

Corollary 7.28 *Let $H < \text{SL}(\mathbb{R}^n)$ be an open subgroup of the fixed point set of an involution. Let \tilde{H} be its preimage in $G(1, M)$, and let V be a representation of \tilde{H} with invariant form (\cdot, \cdot) . Then the induced representation $\Gamma(E(V))$ of $\tilde{\text{Diff}}^0(M)$ has an invariant form $\langle \phi | \psi \rangle = \int_{R(M)} (\phi, \psi) \text{vol}$, with vol the Diff^0 -invariant volume on $R(M)$. It is nondegenerate and/or positive definite if and only if (\cdot, \cdot) is. In particular, any unitary representation of \tilde{H} induces a unitary representation of $\tilde{\text{Diff}}^0(M)$.*

In the case that V is a spinor type representation of $\widetilde{\text{SO}}(\eta) \times_{\mathbb{Z}} \pi_1(F)$, the induced representation on $\Gamma(E(V))$ is only unitary if η is positive definite. On the other hand, an invariant sesquilinear form exists for any signature.

Decomposition into Irreducibles

These induced representations will in general not be irreducible. We sketch the decomposition into irreps, in the particular case that M is a parallelizable, compact, connected manifold of dimension ≥ 2 , and V is a unitary representation of \tilde{H} that factors through H . (So that the induced representation will factor through $\text{Diff}^0(M)$.) The problem then essentially reduces to harmonic analysis on the symmetric space $\Delta = \text{SL}(\mathbb{R}^n)/H$.

We fix once and for all a global section $f : M \rightarrow F^+(M)$ of the frame bundle. It identifies $F^+(M)$ with $M \times \text{GL}^+(\mathbb{R}^n)$, by way of $(x, g) \mapsto f_x g$. If

$\alpha_* : T_x M \rightarrow T_{\alpha(x)} M$ is the pushforward, then we define $D_x \alpha$ to be the matrix of α_* w.r.t. the frames at x and $\alpha(x)$, i.e. $D_x \alpha := f_{\alpha(x)}^{-1} \circ \alpha_* \circ f_x$. The action of α on $M \times \mathrm{GL}^+(\mathbb{R}^n)$ is then $\alpha : (x, g) \mapsto (\alpha(x), D_x \alpha \cdot g)$.

We identify $\mathbb{R}^+ \times \mathrm{SL}(\mathbb{R}^n)$ with $\mathrm{GL}^+(\mathbb{R}^n)$ by $(\lambda, g) \mapsto \sqrt[n]{\lambda} g$. With $\Delta = \mathrm{SL}(\mathbb{R}^n)/H$, we then have $R(M) = M \times \mathbb{R}^+ \times \Delta$. The action of α on $R(M)$ is $\alpha : (x, \lambda, \omega) \mapsto (\alpha(x), \det(D_x \alpha), \check{D}_x \alpha \cdot \omega)$, with $\check{D}_x \alpha := D_x \alpha / \sqrt[n]{\det(D_x \alpha)}$.

If vol_M is the volume on M defined by the frame, and vol_Δ is the unique (up to \mathbb{R}^\times) $\mathrm{SL}(\mathbb{R}^n)$ -invariant volume on Δ , then it is not hard to see that the unique (up to \mathbb{R}^\times) Diff^0 -invariant volume on $R(M)$ is $\mathrm{vol} = \mathrm{vol}_M \wedge d\lambda/\lambda^2 \wedge \mathrm{vol}_\Delta$.

The bundle $E(V)$ is simply $M \times \mathbb{R}^+ \times E_\Delta(V)$, with $E_\Delta(V)$ the vector bundle $\mathrm{SL}(\mathbb{R}^n) \times_H V$ over Δ . Its space \mathcal{H} of L^2 -sections constitutes a unitary representation of $\mathrm{SL}(\mathbb{R}^n)$, the induced representation of V .

We identify the space of L^2 -sections of $E(V)$ with $L^2(M, L^2(\mathbb{R}^+) \otimes \mathcal{H})$, where \mathbb{R}^+ is equipped with $d\lambda/\lambda^2$, and M with vol_M . The representation of $\mathrm{Diff}^0(M)$ is then given by

$$(\alpha\psi)(x) = \rho(\det(D_x \alpha^{-1}))^{-1} \otimes \pi(\check{D}_x \alpha^{-1})^{-1} \psi(\alpha^{-1}(x)),$$

with ρ the nonunitary \mathbb{R}^+ -representation $(\rho(c)f)(\lambda) = f(c^{-1}\lambda)$ on the Hilbert space $L^2(\mathbb{R}^+, d\lambda/\lambda^2)$, and π the unitary $\mathrm{SL}(\mathbb{R}^n)$ -representation on \mathcal{H} .

We decompose ρ into irreducible representations. The unitary transformation $W : L^2(\mathbb{R}^+, d\lambda/\lambda^2) \rightarrow L^2(\mathbb{R}, dt)$ defined by $(Wf)(t) = e^{-t/2} f(e^t)$ takes ρ into $\rho'(c) = W\rho W^{-1}(c)$, with $(\rho'f')(t) = c^{-1/2} f'(t - \log(c))$. The Fourier transform $\mathcal{F} : L^2(\mathbb{R}, dt) \rightarrow L^2(\mathbb{R}, dk)$ then takes ρ' into $\hat{\rho}(c) = \mathcal{F}\rho'(c)\mathcal{F}^{-1}$, with $(\hat{\rho}(c)\hat{f})(k) = c^{-1/2-ik} \hat{f}(k)$. We conclude that the overall transformation $\mathcal{F} \circ W : L^2(\mathbb{R}^+, d\lambda/\lambda^2) \rightarrow L^2(\mathbb{R}, dk)$, given by

$$\hat{f}(k) = \frac{1}{\sqrt{2\pi}} \int_0^\infty \lambda^{1/2-ik} f(\lambda) d\lambda/\lambda^2,$$

identifies ρ with the representation $(\hat{\rho}(c)\hat{f})(k) = c^{-1/2-ik} \hat{f}(k)$ on $L^2(\mathbb{R}, dk)$. We write $L^2(\mathbb{R}^+, d\lambda/\lambda^2) \simeq \int_{\mathbb{R}}^\oplus \mathbb{C}_k dk$ and $\rho \simeq \int_{\mathbb{R}}^\oplus \rho_k dk$, with (ρ_k, \mathbb{C}_k) the one-dimensional representation $\rho_k(c)z = c^{-\frac{1}{2}-ik} z$. (See e.g. [Dix69, ch. II] on direct integrals.)

Consequently, the $\mathrm{Diff}^0(M)$ -representation $L^2(M \times \mathbb{R}^+, \mathcal{H})$, with measure $\mathrm{vol}_M \wedge d\lambda/\lambda^2$, is unitarily equivalent to $L^2(M \times \mathbb{R}, \mathcal{H})$, with measure $\mathrm{vol}_M \wedge dk$ and representation

$$(\alpha\psi)(x, k) = \det(D_x \alpha^{-1})^{1/2+ik} \pi(\check{D}_x \alpha^{-1})^{-1} \psi(\alpha^{-1}(x), k). \quad (7.15)$$

We thus have a decomposition $L^2(M, L^2(\mathbb{R}^+) \otimes \mathcal{H}) \simeq \int_{\mathbb{R}}^\oplus L^2(M, \mathbb{C}_k \otimes \mathcal{H}) dk$ into subrepresentations, where the unitary $\mathrm{Diff}^0(M)$ -representation on the Hilbert space $L^2(M, \mathbb{C}_k \otimes \mathcal{H})$ is given by equation (7.15), but now with fixed k .

In general, the representations $L^2(M, \mathbb{C}_k \otimes \mathcal{H})$ will still be reducible. Now suppose that we can decompose \mathcal{H} into unitary irreps as $\mathcal{H} = \int_B^\oplus \mathcal{H}_\beta \mu(d\beta)$ for some measure space (B, μ) , with $\pi = \int_B^\oplus \pi_\beta \mu(d\beta)$. Then

$$L^2(M, L^2(\mathbb{R}^+) \otimes \mathcal{H}) \simeq \int_{\mathbb{R} \times B}^\oplus L^2(M, \mathbb{C}_k \otimes \mathcal{H}_\beta) dk \mu(d\beta), \quad (7.16)$$

with the representation on $L^2(M, \mathbb{C}_k \otimes \mathcal{H}_\beta)$ given by

$$(\alpha\psi)(x) = \det(D_x\alpha^{-1})^{1/2+ik} \pi_\beta(\check{D}_x\alpha^{-1})^{-1}\psi(\alpha^{-1}(x)). \quad (7.17)$$

According to a theorem stated (but not proven) in [Is72], these representations of $\text{Diff}^0(M)$ are irreducible, at least under the assumption that M is a compact manifold of dimension at least 2.

This means that the problem of decomposing the L^2 -closure of $\Gamma(E(V))$ into irreducible representations of $\text{Diff}^0(M)$ reduces to that of decomposing the L^2 -closure of $\Gamma(E_\Delta(V))$ into irreps of $\text{SL}(\mathbb{R}^2)$. Although this is not an easy problem, it is one that has been studied in considerable detail.

Example: The Hyperbolic Plane

We explicitly carry out the decomposition in the special case that M is the 2-torus T^2 with global frame $(\frac{d}{dx}, \frac{d}{dy})$, the subgroup is H is $\text{SO}(\mathbb{R}^2)$, and $V = \mathbb{C}$ is the trivial representation.

We identify the symmetric space $\text{SL}(\mathbb{R}^2)/\text{SO}(\mathbb{R}^2)$ with the open unit disk $\Delta = \{z \in \mathbb{C}; |z| < 1\}$. It comes equipped with the $\text{SL}(\mathbb{R}^2)$ -action $g \cdot z = A(g)z + B(g)/(\bar{B}(g)z + \bar{A}(g))$, where $A(g)$ and $B(g)$ are defined by

$$\begin{pmatrix} A & B \\ \bar{B} & \bar{A} \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1 & -i \\ 1 & i \end{pmatrix} g \begin{pmatrix} 1 & 1 \\ i & -i \end{pmatrix}. \quad (7.18)$$

(This is the identification of $\text{SU}(1, 1)$ with $\text{SL}(\mathbb{R}^2)$ that comes from the complex isomorphism of Δ with the upper half plane.) The $\text{SL}(\mathbb{R}^2)$ -invariant volume on Δ is

$$\text{vol}_\Delta = \frac{idz \wedge d\bar{z}}{2(1 - |z|^2)^2}.$$

Since $E_\Delta(V)$ is the trivial bundle $\Delta \times \mathbb{C}$, its space of L^2 -sections is $L^2(\Delta, \text{vol}_\Delta)$. It decomposes into the spherical principal series representations π_β , which we will define shortly.

The unit circle $S^1 = \{\omega \in \mathbb{C}; |\omega| = 1\}$, equipped with the rotation invariant measure $d\omega/(2\pi i\omega)$, carries an $\text{SL}(\mathbb{R}^2)$ -action $\omega \mapsto g \cdot \omega$ because it is the boundary of Δ . Define the Fourier transform $\mathcal{F} : L^2(\Delta) \xrightarrow{\sim} L^2(\mathbb{R}^+ \times S^1)$ by

$$(\mathcal{F}f)(\beta, \omega) = \int_\Delta f(z) \left(\frac{1 - |z|^2}{|\omega - z|^2} \right)^{(1-i\beta)/2} \text{vol}_\Delta.$$

If we equip \mathbb{R}^+ with the measure $\mu(d\beta) = (2\pi)^{-1} \beta \tanh(\frac{1}{2}\pi\beta) d\beta$, then it is a unitary isomorphism [Hel84, p. 33]. By direct calculation, one verifies that

$$\left(\mathcal{F}\pi(g^{-1})\mathcal{F}^{-1}f \right) (\beta, \omega) = |\bar{B}(g)\omega + \bar{A}(g)|^{i\beta-1} \hat{f}(\beta, g \cdot \omega),$$

with $A(g)$ and $B(g)$ as in equation (7.18). In other words, the Fourier transform identifies the representation of $\text{SL}(\mathbb{R}^2)$ on $L^2(\Delta)$ with the direct integral $\int_{\mathbb{R}^+}^{\oplus} \pi_\beta \mu(d\beta)$, where π_β is the spherical principal series representation of $\text{SL}(\mathbb{R}^2)$ on $L^2(S^1)$ defined by $(\pi_\beta(g^{-1})f)(\omega) = |\bar{B}(g)\omega + \bar{A}(g)|^{i\beta-1} f(g \cdot \omega)$.

In this particular case, equation (7.16) means that the Diff^0 -representation on the L^2 -sections of $E(V)$ is unitarily equivalent to $L^2(T^2 \times S^1 \times \mathbb{R}^+ \times \mathbb{R}^+)$, with the measure $d\phi \wedge d\theta \wedge d\omega / (2\pi i\omega) \wedge dk \wedge (2\pi)^{-1} \beta \tanh(\pi\beta/2) d\beta$, and the representation

$$(\pi(\alpha)\psi)(x, \omega, k, \beta) = c_{k,\beta}(D_x\alpha^{-1})\psi(\alpha^{-1}(x), D_x\alpha^{-1} \cdot \omega, k, \beta),$$

with the cocycle (A and B as in (7.18))

$$c_{k,\beta}(D_x\alpha^{-1}) = \det(D_x\alpha^{-1})^{1/2+ik} |\bar{B}(D_x\alpha^{-1})\omega + \bar{A}(D_x\alpha^{-1})|^{i\beta-1}.$$

(There is no need for the projection $\check{D}_x\alpha^{-1}$, as the scalars act trivially on Δ .)

The unitary representations of $\text{Diff}^0(T^2)$ that one gets when fixing k and β are irreducible according to [Ism72].

7.7.5 Some Physical Theories

If we accept that any physically relevant Spin^G -structure must be infinitesimally natural, and that G should be the gauge group of the theory, then theorem 7.23 provides a link between the spectrum of elementary particles and the global topology of space-time.

The fact that the mere existence of a generalized spin structure may place restrictions on the space-time manifold was recognized by Hawking and Pope [HP78]. Generalized spin structures were classified [AI80], and it was found that if the Lie group contains $\text{SU}(2)$, then ‘universal spin structures’ exist [BFF78], irrespective of the topology of M . In particular, there are no topological obstructions to the existence of a Spin^G -structure as soon as $\text{SU}(2) < G$.

But according to theorem 7.23, this changes if one requires the Spin^G -structure to be infinitesimally natural. Universal spin structures then exist only for certain noncompact groups. For compact G , the requirement that there exist a homomorphism $\pi_1(F^+(M)) \rightarrow G$ fixing Z provides an obstruction on the space-time manifold M in terms of the group of internal symmetries G . Let us see what this means in some specific cases.

Weyl Spinors

Consider a single massless charged Weyl spinor coupled to a $U(1)$ gauge field. For simplicity, let us assume that M is 4-dimensional, oriented, and time-oriented, so that $Z = \mathbb{Z}/2\mathbb{Z}$ and we may use $\text{SL}(\mathbb{C}^2)$ instead of $\widetilde{\text{SO}}(\eta)$.

This means that $G = U(1)$ and $V = \mathbb{C}^2 \otimes \mathbb{C}_q$, the two-dimensional defining representation of $\text{SL}(\mathbb{C}^2)$ tensored with the one dimensional defining representation of $U(1)$.

Note that the representation descends to $\text{Spin}^c = \text{SL}(\mathbb{C}^2) \times_Z U(1)$. (Z is just ± 1 in $U(1)$.) This means that upon choosing a Spin^c -structure Q , the configuration space is $\Gamma(J^1(\hat{Q}) \times_{\text{Spin}^c} V)$.

Let us now impose the requirement that Q be infinitesimally natural. Theorem 7.23 tells us that there must then be a homomorphism $\pi_1(F^+(M)) \rightarrow U(1)$ sending $\mathbb{Z}/2\mathbb{Z}$ to ± 1 . If $\pi_1(M)$ is finitely generated, then the image of $\pi_1(F^+(M))$ in $U(1)$ is a finitely generated subgroup containing ± 1 . It

must be isomorphic to $\mathbb{Z}^n \times (\mathbb{Z}/2m\mathbb{Z})$ for some $n, m \in \mathbb{N}$. If we then send \mathbb{Z}^n to 1, we obtain a homomorphism $\pi_1(F^+(M)) \rightarrow \mathbb{Z}/2m\mathbb{Z}$. The sequence $1 \rightarrow \mathbb{Z}/2\mathbb{Z} \rightarrow \mathbb{Z}/2m\mathbb{Z} \rightarrow \mathbb{Z}/m\mathbb{Z} \rightarrow 1$ splits precisely when m is odd.

We conclude that an infinitesimally natural Spin^c -structure exists on M if and only if there is a surjective homomorphism $\pi_1(F^+(M)) \rightarrow \mathbb{Z}/2m\mathbb{Z}$ which preserves $\mathbb{Z}/2\mathbb{Z}$. Only if m is odd does this give rise to a spin structure.

Dirac Spinors

Next, consider the case of a Dirac spinor. That is, $V = \mathbb{C}^4 \otimes \mathbb{C}_q$, with \mathbb{C}^4 the representation of $\text{Cl}(\mathbb{C}^4)$ which splits into two identical irreps $\mathbb{C}^2 \oplus \mathbb{C}^2$ under $\text{SL}(\mathbb{C}^2)$, the left handed and right handed spinors.

The fact that V is reducible under Spin^c makes us re-examine our assumption that the group G in theorem 7.23 should be the gauge group $U(1)$. Indeed, the unitary commutant of $\text{SL}(\mathbb{C}^2)$ in V is $U(2)$ rather than $U(1)$. If we take any discrete subgroup $H < U(2)$ and form the group $U(1)_H$ generated by H and $U(1)$, can we take Q to be a Spin^G -structure with structure group $U(1)_H$?

As far as only the kinematics is concerned, the answer is yes. The generic fibre of $J^1(\hat{Q}) \times_{\text{Spin}^G} V$ is the same for $G = U(1)$ as it is for $G = U(1)_H$, so adding H will not change the space of local sections.

But if we take into account the dynamics, the answer becomes: ‘only if H is a group of symmetries of the Lagrangian’. The reason is that Lagrangians are usually defined in local co-ordinates, yielding a local action S_U for each co-ordinate patch $U \subset M$. However, as the dynamics of the theory should be governed by a global action functional S_M , it is necessary for S_U and S_V to agree on $U \cap V$. This means that if H is part of the structure group of the bundle, then it must leave the Lagrangian invariant. If H is a global symmetry, then the transition functions must be constant. This is automatic if H is discrete.

For instance, in the case of a massive Dirac fermion, the subgroup of $U(2)$ which preserves the Lagrangian is precisely the diagonal $U(1)$. This means that the relevant Spin^G -structures are precisely the Spin^c -structures classified above.

For massless Dirac spinors, the left and right Weyl spinors decouple, so that the relevant symmetry group is $U_L(1) \times U_R(1)$. Although the requirement on a manifold to carry a Spin^G -structure does not change, this does give us more Spin^G -structures for the same manifold.

This illustrates that we may enlarge the gauge group G by any group of discrete symmetries of the Lagrangian in order to obtain Spin^G -structures. In particular, this means that infinitesimally natural Spin^c -structures are allowed even for uncharged Weyl spinors. (The image of $\pi_1(F^+(M))$ is automatically discrete.)

The Standard Model

In the the standard model of elementary particle physics, the gauge group G is $(\text{SU}(3) \times \text{SU}(2)_L \times \text{U}(1)_Y)/N$, with N the cyclic subgroup of order 6 generated by $(e^{2\pi i/3}\mathbf{1}, -\mathbf{1}, e^{2\pi i/6})$. It is isomorphic to $S(U(3) \times U(2))$, a subgroup of $\text{SU}(5)$, and it has a unique central subgroup of order 2 generated by $\text{diag}(1, 1, 1, -1, -1)$.

The fermion representation V for a single family can be conveniently described (see e.g. [BH10]) by $\mathbb{C}^2 \otimes \wedge^* \mathbb{C}^5$, the tensor product of the defining representation of $\mathrm{SL}(\mathbb{C}^2)$ and the exterior algebra of the defining representation of $\mathrm{SU}(5)$. Under $\mathrm{SL}(\mathbb{C}^2) \times S(U(3) \times U(2))$, this decomposes into 12 irreps corresponding to left and right handed electrons, neutrinos, up and down quarks and their antiparticles.

Unfortunately, $\mathrm{diag}(1, 1, 1, -1, -1) \in G$ acts by $+1$ on right-handed fermions, whereas $-1 \in \mathrm{SL}(\mathbb{C}^2)$ acts by -1 . This means that V does not define a representation of $\mathrm{SL}(\mathbb{C}^2) \times_Z S(U(3) \times U(2))$ if one were to identify the central order 2 elements on both sides.

As the gauge group alone is of no use when trying to find a Spin^G -structure, one has to involve the group of global symmetries of the standard model Lagrangian. It contains the gauge group G , but also (at least on the classical level) the global $U(1)_B \times U(1)_L$ -symmetries that rotate quarks and leptons independently. (These are connected to baryon and lepton number.)

We conclude that the only infinitesimally natural Spin^G -structures relevant to the standard model are the ones associated to homomorphisms

$$\pi_1(F^+(M)) \rightarrow \hat{G} \tag{7.19}$$

that preserve Z , the subgroup of $U(1)_B \times U(1)_L$ generated by $(-1, -1)$. In this expression, \hat{G} is the group of global symmetries of the standard model Lagrangian, which at least contains $S(U(3) \times U(2)) \times U(1)_B \times U(1)_L$.

Infinitesimally Natural Spin^G -Structures for the Standard Model

It seems that any manifold which possesses an infinitesimally natural Spin^G -structure for the standard model automatically permits an infinitesimally natural Spin^c -structure. (This is at least the case if the group \hat{G} of global symmetries is $S(U(3) \times U(2)) \times U(1)_B \times U(1)_L$.) On the other hand, there do exist Spin^G -structures for the standard model which are not Spin^c . We construct an example.

Consider de Sitter space $H = \{\vec{x} \in \mathbb{R}^5 \mid -x_0^2 + x_1^2 + x_2^2 + x_3^2 + x_4^2 = 1\}$, which has a pseudo-Riemannian metric g with constant curvature induced by the Minkowski metric in the ambient \mathbb{R}^5 . Its group of orientation preserving isometries is $\mathrm{SO}(1, 4)$, and $H \simeq \mathrm{SO}(1, 4)/\mathrm{SO}(1, 3)$. Denote by $OF_g^{+\uparrow}(H)$ the bundle of orthogonal frames with positive orientation and time-orientation. By viewing $OF_g^{+\uparrow}(H)$ as a submanifold of $\mathbb{R}^5 \times \mathrm{SO}(1, 4)^0$, one can see that $\mathrm{SO}(1, 4)^0$ acts freely and transitively by $x : f \mapsto x_*f$. Therefore $OF_g^{+\uparrow}(H)$ is diffeomorphic to $\mathrm{SO}(1, 4)^0$.

Now let $\Gamma < \mathrm{SO}(4)$ be a discrete group which acts freely, isometrically and properly discontinuously on S^3 . The manifold $\Gamma \backslash S^3$ is called a spherical space form. (See [Wol67] for a classification.) As Γ includes into $\mathrm{SO}(1, 4)^0$, it acts on H , making $M = \Gamma \backslash H$ into a pseudo-Riemannian manifold with constant curvature.

We immediately see that $\pi_1(M) = \Gamma$, because H is simply connected. We calculate the homotopy group of the frame bundle. Because $OF_g^{+\uparrow}(M)$ is just $\Gamma \backslash OF_g^{+\uparrow}(H)$, it is isomorphic to $\Gamma \backslash \mathrm{SO}(1, 4)^0$. Going to the universal cover,

we see that $OF_g^+(M) = \tilde{\Gamma} \backslash \widetilde{SO}(1,4)^0$. As $\Gamma < SO(4)$, we may consider $\tilde{\Gamma}$ to be the preimage of Γ in $\text{Spin}(4)$. As the universal cover is simply connected, it is now clear that $\pi_1(OF_g^+(M)) = \tilde{\Gamma}$. We get for free a homomorphism $\tilde{\Gamma} \rightarrow \text{Spin}(4) \simeq \text{SU}(2)_l \times \text{SU}(2)_r$, which maps the noncontractible loop in the fibre to $(-1, -1)$.

Triggered by the WMAP-data on cosmic background radiation, there has been some interest in the case where Γ is I^* , the binary icosahedral group [LW⁺03]. Some evidence against [KC⁺07] as well as in favour of [RB⁺08] this hypothesis appears to have been found. We carefully avoid choosing sides in the debate, and simply point out that $I^* \backslash H$ gives rise to an interesting infinitesimally natural Spin^G -structure.

Under the identification $\text{Spin}(4) \simeq \text{SU}(2)_l \times \text{SU}(2)_r$, we see that $\Gamma = I^* \times \mathbf{1}$ lives only in $\text{SU}(2)_l$, so that $\tilde{\Gamma}$ is the direct product of $I^* \times \mathbf{1}$ and the $\mathbb{Z}/2\mathbb{Z}$ generated by $(-1, -1)$. One can therefore define a homomorphism (7.19) by identifying $\text{SU}(2)_l$ with $\text{SU}(2)_L < G$, and mapping $(-1, -1)$ to $(-1, -1) \in U(1)_B \times U(1)_L$. This yields an infinitesimally natural Spin^G -structure which uses the noncommutativity of the gauge group in an essential fashion. Note however that ‘ordinary’ Spin^c -structures also exist on M .

Extensions of the Standard Model

The fact that $S(U(3) \times U(2))$ does not contribute to the obstruction of finding Spin^G -structures on M is due to the fact that it never acts by -1 on V . This is not true for some GUT-type extensions of the standard model, such as the Pati-Salam $\text{SU}(2)_L \times \text{SU}(2)_R \times \text{SU}(4)$ model and anything which extends it, for example $\text{Spin}(10)$.

If N is the group of order 2 generated by $(-1, -1, -1)$, then infinitesimally natural Spin^G -structures in the Pati-Salam model correspond, neglecting global symmetries, to homomorphisms $\pi_1(F^+(M)) \rightarrow \text{SU}(2)_L \times \text{SU}(2)_R \times \text{SU}(4)/N$ which take Z to $\langle(-1, -1, \mathbf{1})\rangle$. This has the rather intriguing consequence that there may well exist space-time manifolds M which are compatible with the Pati-Salam model, but not with the standard model. A manifold M would have this property if the smallest quotient of $\pi_1(F(M))$ containing Z is a nonabelian subgroup of $\text{SU}(2)_L \times \text{SU}(2)_R \times \text{SU}(4)$ containing $(-1, -1, \mathbf{1})$.

7.8 Discussion

A natural bundle is one in which diffeomorphisms of the base lift to automorphisms of the bundle in a local fashion. In corollary 7.15, we have rederived the well known [PT77] result that any natural principal fibre bundle is associated to the k^{th} order frame bundle $F^k(M)$.

Bundles associated to spin structures are almost never natural, but nonetheless indispensable to physics. We would therefore like to extend the notion of ‘naturalness’. One way to do this is to introduce ‘gauge natural bundles’ [Eck81] [KMS93], which do not transform according to diffeomorphisms of the base M , but according to automorphisms of a bundle over M which has to be specified.

We here propose to accommodate spin structures in a different fashion. Although the Lorentz group does not act on a spin structure, its associated Lie algebra does. Analogously, we define an ‘infinitesimally natural’ bundle over M to be a bundle in which it is not the group of diffeomorphisms of M that lifts, but its Lie algebra of vector fields. This has the advantage that the link between the base and the fibres is not lost, so that a stress-energy-momentum tensor can be constructed from Noether’s principle [Noe18].

We assume only that the lift is a homomorphism of Lie algebras. However, the careful analysis by Gotay and Marsden [GM92] reveals that in order to define a SEM-tensor from Noether’s principle, one needs a lift that is given by a differential operator. Assuming only that the lift is a homomorphism, we prove in proposition 7.7 that it is actually a differential operator. This shows that one can construct a SEM-tensor from Noether’s principle precisely when the bundle is infinitesimally natural. We therefore propose to describe fields on space-time by sections of infinitesimally natural bundles.

We have classified the infinitesimally natural principal fibre bundles. According to theorem 7.14, they are associated to the universal cover $\tilde{F}^{+k}(M)$ of the k^{th} order frame bundle, and their transformation behaviour is governed by the disconnected group $G(k, M)$. This group depends on the manifold. It generalizes the Pin group in the sense that it regulates parity, time reversal and Lorentz transformations.

Theorem 7.14 is originally due to Lecomte [Lec85]. Although our proof was obtained independently, it does rather resemble Lecomte’s. We pause to highlight a few differences. First of all, Lecomte does not work in the compactly supported setting: he bases his version of lemma 7.5 on results of Amemiya [Ame75] rather than on Pursell and Shanks’ lemma 7.2.1. Secondly, our uniform bound, in proposition 7.7, is different from Lecomte’s. But the most significant difference is probably that Lecomte bases his proof on integration of foliations rather than on integration of Lie algebroids, which probably makes our proof easier to generalize. Strikingly enough, Lecomte’s work does not appear to be motivated by physics; he mentions spin structures only briefly.

The consistent description of fermions in the presence of a gauge group G requires a Spin^G -structure rather than a spin structure [HP78], [AI80]. All spin structures are infinitesimally natural, but some Spin^G -structures are not. As described in theorem 7.23, infinitesimally natural Spin^G -structures for compact G correspond to homomorphisms $\pi_1(F^+(M)) \rightarrow G$ that are injective on $\pi_1(\text{GL}^+(\mathbb{R}^n))$.

The existence of infinitesimally natural Spin^G -structures therefore provides an obstruction on space-time M in terms of the symmetry group G . Some manifolds are even excluded for any (compact) choice of G . For example, $\mathbb{C}P^2$ does not admit any infinitesimally natural Spin^G -structure because its frame bundle is simply connected. In our eyes, this disqualifies it as a model for space-time.

If all global symmetries are gauged, then theorem 7.23 constitutes the “connexion between the topology of space-time and the spectrum of elementary particles” alluded to in [HP78].

Bibliography

- [AFL90] L. Accardi, L. Frigerio, and Y. G. Lu. The weak coupling limit as a quantum functional central limit. *Commun. Math. Phys.*, 131:537–570, 1990.
- [Ame75] I. Amemiya. Lie algebra of vector fields and complex structure. *J. Math. Soc. Japan*, 72(4):545–549, 1975.
- [AGG05] L. Artiles, R. D. Gill, and M. Guță. An invitation to quantum tomography. *J. Royal Statist. Soc. B (Methodological)*, 67:109–134, 2005.
- [AI80] S. J. Avis and C. J. Isham. Generalized spin structures on four dimensional space-times. *Comm. Math. Phys.*, 72:103–118, 1980.
- [AK65] E. Arthurs and J. Kelly. On simultaneous measurement on a pair of conjugate observables. *Bell. Syst. Tech. J.*, 44:725, 1965.
- [AN⁺99] M. Arndt, O. Nairz, J. Vos-Andreae, C. Keller, G. van der Zouw, and A. Zeilinger. Wave–particle duality of C_{60} –molecules. *Nature*, 401:680–682, 1999.
- [AS⁺02] M. Armen, J. Stockton, A. Doherty, and H. Mabuchi. Adaptive homodyne measurement of optical phase. *Phys. Rev. Lett.*, 89:133602, 2002.
- [Ash91] A. Ashtekar. Lectures on non-perturbative canonical gravity. In *Advanced Series in Astrophysics and Cosmology*, volume 6. World Scientific, 1991.
- [Bac90] A. Barchielli. Direct and heterodyne detection and other applications of quantum stochastic calculus to quantum optics. *Quantum Opt*, 2:423–441, 1990.
- [Ban04] E. P. van den Ban. Applications of representation theory in classical quantum mechanics. Lecture notes for the MRI Spring School Lie groups in Analysis, Geometry and Mechanics, 2004.
- [Bar62] D. W. Barnes. Nilpotency of Lie algebras. *Math. Zeitschr.*, 79:237–238, 1962.

- [BB⁺06] E. Bagan, M. A. Ballester, R. D. Gill, A. Monras, and R. Muñoz-Tapia. Optimal full estimation of qubit mixed states. *Phys. Rev. A*, 73:032301, 2006.
- [BB⁺04] E. Bagan, M. Baig, R. Muñoz-Tapia, and A. Rodriguez. Collective versus local measurements in a qubit mixed-state estimation. *Phys. Rev. A*, 69:010304, 2004.
- [Bel75] J. Bell. On wave packet reduction in the Coleman–Hepp model. *Helv. Phys. Acta*, 48:93–98, 1975.
- [BFF78] A. Back, P. G. O. Freund, and M. Forger. New gravitational instantons and universal spin structures. *Phys. Lett.*, 77B:181–184, 1978.
- [BGA07] C. Butucea, M. Guță, and L. Artiles. Minimax and adaptive estimation of the Wigner function in quantum homodyne tomography with noisy data. *Ann. Stat.*, 35(2):465–494, 2007.
- [BGM04] L. Bouten, M. Guță, and H. Maassen. Stochastic Schrödinger equations. *J. Phys. A*, 37:3189–3209, 2004.
- [BH10] J. Baez and J. Huerta. The algebra of grand unified theories. *Bull. Am. Math. Soc.*, 47(3):483–552, 2010.
- [BL85] A. Barchielli and G. Lupieri. Quantum stochastic calculus, operation valued stochastic processes, and continual measurements in quantum mechanics. *J. Math. Phys.*, 26:2222–2230, 1985.
- [Blv88] V. Belavkin. Nondemolition stochastic calculus in Fock space and nonlinear filtering and control in quantum systems. In *Proceedings XXIV Karpacz winter school, stochastic methods in mathematics and physics*, pages 310–324, 1988. World Scientific.
- [Blv92a] V. Belavkin. Quantum continual measurements and a posteriori collapse on CCR. *Commun. Math. Phys.*, 146:611–635, 1992.
- [Blv92b] V. Belavkin. Quantum stochastic calculus and quantum nonlinear filtering. *Journal of Multivariate Analysis*, 42:171–201, 1992.
- [BMM02] E. Bagan, M. Baig, and R. Muñoz-Tapia. Optimal scheme for estimating a pure qubit state via local measurements. *Phys. Rev. Lett.*, 89:277904, 2002.
- [BM⁺05] E. Bagan, A. Monras, and R. Muñoz-Tapia. Comprehensive analysis of quantum pure-state estimation for two-level systems. *Phys. Rev. A*, 71:062318, 2005.
- [BNG00] O. E. Barndorff-Nielsen and R. D. Gill. Fisher information in quantum statistics. *J. Phys. A.*, 33:1–10, 2000.

- [BG⁺03] O. E. Barndorff-Nielsen, R. D. Gill, and P. E. Jupp. On quantum statistical inference (with discussion). *J. R. Statist. Soc. B*, 65:775–816, 2003.
- [Bou04] L. Bouten. *Filtering and Control in Quantum Optics*. PhD thesis, Radboud Universiteit Nijmegen, 2004.
- [BW⁺01] M. Berg, C. DeWitt-Morette, S. Gwo, and E. Kramer. The pin groups in physics: C, P and T. *Rev. Math. Phys.*, 13:953–1034, 2001.
- [Car93] H. Carmichael. An open systems approach to quantum optics. In *Lecture notes in Physics*, volume 18. Springer-Verlag, 1993.
- [Cav82] C. M. Caves. Quantum limits on noise in linear amplifiers. *Phys. Rev. D*, 26:1817–1839, 1982.
- [CEM99] J. I. Cirac, A. K. Ekert, and C. Macchiavello. Optimal purification of single qubits. *Phys. Rev. Lett.*, 82:4344, 1999.
- [CF03] M. Crainic and R. L. Fernandes. Integrability of Lie brackets. *Ann. of Math.*, 157(2):575–620, 2003.
- [Dav69] E. Davies. Quantum stochastic processes. *Commun. Math. Phys.*, 15:277–304, 1969.
- [Dav76] E. Davies. *Quantum Theory of Open Systems*. Academic Press, 1976.
- [Dix69] J. Dixmier. *les algèbres d’opérateurs dans l’espace Hilbertien (algèbres de von Neumann)*. Gauthier-Villars, 1969.
- [DLP95] G. M. D’Ariano, U. Leonhardt, and H. Paul. Homodyne detection of the density matrix of the radiation field. *Phys. Rev. A*, 52:R1801–R1804, 1995.
- [Dys56] F. J. Dyson. General theory of spin-wave interactions. *Phys. Rev.*, 102:1217–1230, 1956.
- [Eck81] D. J. Eck. Gauge-natural bundles and generalized gauge theories. In *Memoirs of the AMS*, volume 33, No. 247. AMS, 1981.
- [EN04] F. Embacher and H. Narnhofer. Strategies to measure a quantum state. *Ann. Phys.*, 311:220, 2004.
- [ET79] D. B. A. Epstein and W. P. Thurston. Transformation groups and natural bundles. *Proc. London Math. Soc.*, 38(3):219–236, 1979.
- [FKF00] D. G. Fisher, S. H. Kienle, and M. Freyberger. Quantum-state estimation by self-learning measurements. *Phys. Rev. A*, 61:032306, 2000.
- [FN95] A. Fujiwara and H. Nagaoka. Quantum Fisher metric and estimation for pure state models. *Physics Letters A*, 201:119–124, 1995.

- [FP⁺00] J. Friedman, V. Patel, W. Chen, S. K. Tolpygo, and J. E. Lukens. Quantum superposition of distinct macroscopic states. *Nature*, 406:43–46, 2000.
- [FR04] M. Forger and H Römer. Currents and the energy-momentum tensor in classical field theory: A fresh look at an old problem. *Ann. Phys.*, 309:306–389, 2004.
- [Fra04] Th. Frankel. *The Geometry of Physics*. Cambridge University Press, 2004.
- [Fuj06] A. Fujiwara. Strong consistency and asymptotic efficiency for adaptive quantum estimation problems. *J. Phys. A*, 39:12489–12504, 2006.
- [Fuk86] D. B. Fuks. *Cohomology of Infinite Dimensional Lie Algebras*. Plenum, 1986.
- [Gil08] R. D. Gill. Conciliation of Bayes and pointwise quantum state estimation. In V. Belavkin, editor, *Quantum Stochastics & Information: Statistics, Filtering & Control*, pages 239–261, 2008. World Scientific.
- [GJ07] M. Guță and A. Jenčová. Local asymptotic normality in quantum statistics. *Commun. Math. Phys.*, 276:341–379, 2007.
- [GJK08] M. Guță, B. Janssens, and J. Kahn. Optimal estimation of qubit states with continuous time measurements. *Comm. Math. Phys.*, 277:127–160, 2008.
- [GK06] M. Guță and J. Kahn. Local asymptotic normality for qubit states. *Phys. Rev. A*, 73:052108, 2006.
- [GKP09] J. Grabowski, A. Kotov, and N. Poncin. Geometric structures encoded in the Lie structure of an Atiyah algebroid, 2009.
- [GM92] M. J. Gotay and J. E. Marsden. Stress-energy-momentum tensors and the Belinfante-Rosenfeld formula. *Contemp. Math.*, 132:367–392, 1992.
- [GM00] R. D. Gill and S. Massar. State estimation for large ensembles. *Phys. Rev. A.*, 61:042312, 2000.
- [Gou05] J. Gough. Quantum flows as Markovian limit of emission, absorption and scattering interactions. *Commun. Math. Phys.*, 254:489–512, 2005.
- [Gou06] J. Gough. Feynman diagrams and the quantum stochastic calculus. *Banach Center Publ.*, 73:187–198, 2006.
- [GSM04] J. Geremia, J. Stockton, and H. Mabuchi. Real-time quantum feedback control of atomic spin-squeezing. *Science*, 304:270–273, 2004.
- [GV64] I. M. Gelfand and N. J. Vilenkin. *Verallgemeinerte Funktionen (Distributionen) IV*. VEB Deutscher Verlag der Wissenschaften, 1964.

- [GZ04] C. W. Gardiner and P. Zoller. *Quantum Noise*. Springer, 2004.
- [Hal04] M. Hall. Prior information: How to circumvent the standard joint-measurement uncertainty relation. *Phys. Rev. A*, 69:052113, 2004.
- [Hay02] M. Hayashi. Two quantum analogues of Fisher information from a large deviation viewpoint of quantum estimation. *J. Phys. A: Math. Gen.*, 35:7689–7727, 2002.
- [Hay03] M. Hayashi. Quantum estimation and the quantum central limit theorem. *Bulletin of the Mathematical Society of Japan*, 55:368–391, 2003. Japanese; Translated into English in quant-ph/0608198.
- [Hay04] 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.
- [Hay05] M. Hayashi, editor. *Asymptotic Theory of Quantum Statistical Inference: Selected Papers*. World Scientific, 2005.
- [Hei27] W. Heisenberg. Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik. *Z. Phys.*, 43:172–198, 1927.
- [Hel84] S. Helgason. *Groups and Geometric Analysis: Integral Geometry, Invariant Differential Operators, and Spherical Functions*. Academic Press, 1984.
- [Hep72] K. Hepp. Quantum theory of measurement and macroscopic observables. *Helv. Phys. Acta*, 45:237–248, 1972.
- [Her74] R. Hermann. Spinors, Clifford and Cayley algebras. In *Interdisciplinary mathematics*, volume VII. MATH SCI Press, Brookline, 1974.
- [HH74] R. Holtz and J. Hanus. On coherent spin states. *J. Phys. A*, 7:37, 1974.
- [HM05] M. Hayashi and K. Matsumoto. Statistical model with measurement degree of freedom and quantum physics. In M. Hayashi, editor, *Asymptotic Theory of Quantum Statistical Inference: Selected Papers*, pages 7689–7727. World Scientific, 2005. English translation of a paper in Japanese published in Surikaiseki Kenkyusho Kokyuroku, vol. 35, pp. 7689-7727, 2002.
- [HM06] M. Hayashi and K. Matsumoto. Asymptotic performance of optimal state estimation in quantum two level system, 2006.
- [Hol82] A. S. Holevo. *Probabilistic and Statistical Aspects of Quantum Theory*. North Holland Publishing Company, 1982.
- [HP78] S. W. Hawking and C. N. Pope. Generalized spin structures in quantum gravity. *Phys. Lett.*, 73B:42–44, 1978.

- [HP84] R. Hudson and K. Parthasarathy. Quantum Itô's formula and stochastic evolutions. *Commun. Math. Phys.*, 93:301–323, 1984.
- [HR⁺02] T. Hannemann, D. Reiss, C. Balzer, W. Neuhauser, P. E. Toschek, and C. Wunderlich. Self-learning estimation of quantum states. *Phys. Rev. A*, 65:050303, 2002.
- [Hst76] C. W. Helstrom. *Quantum Detection and Estimation Theory*. Academic Press, 1976.
- [Hum72] J. E. Humphreys. *Introduction to Lie Algebras and Representation Theory*. Springer-Verlag, 1972.
- [Ish91] S. Ishikawa. Uncertainty relations in simultaneous measurements for arbitrary observables. *Rep. Math. Phys.*, 29:257–273, 1991.
- [Ism72] R. S. Ismagilov. On unitary representations of the group of diffeomorphisms of a compact manifold. *Funktsional'nyi Analiz i Ego Prilozheniya*, 6(1):79–80, 1972.
- [Jan04] B. Janssens. Quantum measurement, a coherent description. Master's thesis, Radboud Universiteit Nijmegen, 2004.
- [Jan08] B. Janssens. Classical coding and the cauchy-schwarz inequality. In V. P. Belavkin and M. Guță, editors, *Quantum Stochastics & Information: Statistics, Filtering & Control*, pages 375–382, 2008. World Scientific.
- [JB06] B. Janssens and L. Bouten. Optimal pointers for joint measurement of σ_x and σ_z via homodyne detection. *J. Phys. A*, 39:2773–2790, 2006.
- [JM06] B. Janssens and H. Maassen. Information transfer implies state collapse. *J. Phys. A*, 39:9845–9860, 2006.
- [Jon94] K. R. Jones. Fundamental limits upon the measurement of state vectors. *Phys. Rev. A*, 50:3682, 1994.
- [JZ85] E. Joos and H. Zeh. The emergence of classical properties through interaction with the environment. *Z. Phys. B*, 59:223–243, 1985.
- [KC⁺07] J. S. Key, N. J. Cornish, D. N. Spergel, and G. D. Starkman. Extending the WMAP bound on the size of the universe. *Phys. Rev.*, 75D:084034, 2007.
- [Ken27] E. Kennard. Zur Quantenmechanik einfacher Bewegungstypen. *Z. Phys.*, 44:326–325, 1927.
- [KG09] J. Kahn and M. Guță. Local asymptotic normality for finite dimensional quantum systems. *Comm. Math. Phys.*, 289(2):597–652, 2009.
- [KMS93] I. Kolář, P. W. Michor, and J. Slovák. *Natural Operations in Differential Geometry*. Springer-Verlag, 1993.

- [Kna86] A. W. Knap. *Representation Theory of Semisimple Groups: An Overview Based on Examples*. Princeton University Press, 1986.
- [Kna96] A. W. Knap. *Lie Groups Beyond an Introduction*. Birkhäuser, Boston, first edition, 1996.
- [KR86] R. V. Kadison and J. R. Ringrose. *Fundamentals of the Theory of Operator Algebras*, volume II. Academic Press, 1986.
- [KW99] M. Keyl and R. Werner. Optimal cloning of pure states, judging single clones. *J. Math. Phys.*, 40:3283–3299, 1999.
- [KW01] M. Keyl and R. F. Werner. Estimating the spectrum of a density operator. *Phys. Rev. A*, 64:052311, 2001.
- [LC86] L. Le Cam. *Asymptotic Methods in Statistical Decision Theory*. Springer Verlag, 1986.
- [LD⁺06] A. Lupaşcu, E. F. C. Driessen, L. Roschier, C. J. P. M. Harmans, and J. E. Mooij. High-contrast dispersive readout of a superconducting flux qubit using a nonlinear resonator. *Phys. Rev. Lett.*, 96:127003, 2006.
- [Lec85] P. B. A. Lecomte. Sur la suite exacte canonique associée à un fibré principal. *Bulletin de la S. M. F.*, 113:259–271, 1985.
- [LM94] H. B. Lawson and M.-L. Michelsohn. *Spin geometry*. Princeton University Press, second edition, 1994.
- [LM⁺96] U. Leonhardt, M. Munroe, T. Kiss, T. Richter, and M. G. Raymer. Sampling of photon statistics and density matrix using homodyne detection. *Optics Communications*, 127:144–160, 1996.
- [LPA95] U. Leonhardt, H. Paul, and G. M. D’Ariano. Tomographic reconstruction of the density matrix via pattern functions. *Phys. Rev. A*, 52:4899–4907, 1995.
- [LPT98] J. I. Latorre, P. Pascual, and R. Tarrach. Minimal optimal generalized quantum measurements. *Phys. Rev. Lett.*, 81:1351, 1998.
- [LW⁺03] J. Luminet, J. R. Weeks, A. Riazuelo, R. Lehoucq, and J. Uzan. Dodecahedral space topology as an explanation for weak wide-angle temperature correlations in the cosmic microwave background. *Nature*, 425:593–595, 2003.
- [Maa03] H. Maassen. Quantum probability applied to the damped harmonic oscillator. In S. Attal and J. M. Lindsay, editors, *Quantum Probability Communications XII*, pages 23–58. World Scientific, 2003.
- [Maa04] H. Maassen. Quantum probability, quantum information and quantum computing, 2004. Lecture notes, University of Nijmegen.

- [Mat02] K. Matsumoto. A new approach to the Cramer-Rao type bound of the pure state model. *J. Phys. A*, 35(13):3111–3123, 2002.
- [MFF00] H. Mack, D. G. Fischer, and M. Freyberger. Enhanced quantum estimation via purification. *Phys. Rev. A*, 62:042301, 2000.
- [MM02] I. Moerdijk and J. Mrčun. On integrability of infinitesimal actions. *Amer. J. Math.*, 124:567–593, 2002.
- [Mor01] S. Morrison. Classifying spinor structures. Master’s thesis, University of New South Wales, 2001.
- [MP95] S. Massar and S. Popescu. Optimal extraction of information from finite quantum ensembles. *Phys. Rev. Lett.*, 74:1259–1263, 1995.
- [MX00] K. C. H. Mackenzie and P. Xu. Integration of Lie bialgebroids. *Topology*, 39:445–467, 2000.
- [Nag05] 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.
- [NC00] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [Neu32] J. von Neumann. *Mathematische Grundlagen der Quantenmechanik*. Springer-Verlag, 1932.
- [Noe18] A. E. Noether. Invariante Variationsprobleme. *Nachr. d. König. Gesellsch. d. Wiss. zu Göttingen*, pages 235–257, 1918.
- [Ott93] E. Ott. *Chaos in Dynamical Systems*. Cambridge University Press, 1993.
- [Oza03] M. Ozawa. Universally valid reformulation of the Heisenberg uncertainty principle on noise and disturbance in measurement. *Phys. Rev. A*, 67:042105, 2003.
- [Par92] K. Parthasarathy. *An Introduction to Quantum Stochastic Calculus*. Birkhäuser, 1992.
- [Pee60] J. Peetre. Réctification (*sic*) à l’article «une caractérisation abstraite des opérateurs différentiels». *Math. Scand.*, 8:116–120, 1960.
- [Per93] A. Peres. *Quantum Theory: Concepts and Methods*. Kluwer Academic Publishers, 1993.
- [Pra66] J. Pradines. Théorie de Lie pour les groupoides différentiables, relation entre propriétés locales et globales. *Comptes Rendus Acad. Sci. Paris A*, 263:907–910, 1966.

- [PT77] R. S. Palais and C. L. Terng. Natural bundles have finite order. *Topology*, 16:271–277, 1977.
- [RB⁺08] B. F. Roukema, Z. Bulinski, A. Szaniewska, and N. E. Gaudin. Optimal phase of the generalised Poincaré dodecahedral space hypothesis implied by the spatial cross-correlation function of the WMAP sky maps. *A & A*, 486:55–72, 2008.
- [Rob29] H. Robertson. The uncertainty principle. *Phys. Rev.*, 34:163–164, 1929.
- [Sal72] S. E. Salvioli. On the theory of geometric objects. *J. Diff. Geom.*, 7:257–278, 1972.
- [Sar09] G. Sardanashvily. Fibre bundles, jet manifolds and Lagrangian theory. Lectures for theoreticians, 2009.
- [SB⁺96] S. Schiller, G. Breitenbach, S. F. Pereira, T. Müller, and 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.
- [SB⁺93] D. T. Smithey, M. Beck, M. G. Raymer, and A. Faridani. 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.
- [Sch35] E. Schrödinger. Die gegenwärtige Situation in der Quantenmechanik. *Naturwissenschaften*, 23:807–812, 823–828, 844–849, 1935.
- [Sew05] G. Sewell. On the mathematical structure of quantum measurement theory. *Rep. Math. Phys.*, 56:271–290, 2005.
- [SP54] M. E. Shanks and L. E. Pursell. The Lie algebra of a smooth manifold. *Proceedings of the American Mathematical Society*, 5(3):468–472, 1954.
- [Sti55] W. F. Stinespring. Positive functions on C*-algebras. *Proc. Amer. Math. Soc*, 6:211–216, 1955.
- [Tak71] M. Takesaki. Conditional expectations in von Neumann algebras. *J. Funct. Anal.*, 9:306–321, 1971.
- [Tak73] F. Takens. Derivations of vector fields. *Comp. Math.*, 26:151–158, 1973.
- [Tak79] M. Takesaki. *Theory of Operator Algebras I*. Springer-Verlag, 1979.
- [Ter78] C. L. Terng. Natural vector bundles and natural differential operators. *Am. J. Math.*, 100(4):775–828, 1978.

- [Vaa98] A. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.
- [VW96] A. van der Vaart and J. Wellner. *Weak Convergence and Empirical Processes*. Springer, 1996.
- [WH⁺00] C. H. van der Wal, A. C. J. ter Haar, F. K. Wilhelm, R. N. Schouten, C. J. P. M. Harmans, T. P. Orlando, S. Lloyd, and J. E. Mooij. Quantum superposition of macroscopic persistent-current states. *Science*, 290:773–777, 2000.
- [VL⁺99] G. Vidal, J. I. Latorre, P. Pascual, and R. Tarrach. Optimal minimal measurements of mixed states. *Phys. Rev. A*, 60:126, 1999.
- [VR89] K. Vogel and H. Risken. Determination of quasiprobability distributions in terms of probability distributions for the rotated quadrature phase. *Phys. Rev. A*, 40:2847–2849, 1989.
- [Wal43] A. Wald. Tests of statistical hypotheses concerning several parameters when the number of observations is large. *Trans. Amer. Math. Soc.*, 54:426–482, 1943.
- [Wer98] R. Werner. Optimal cloning of pure states. *Phys. Rev. A*, 58:1827–1832, 1998.
- [Wer01] R. Werner. Quantum information theory | an invitation. In *Quantum Information | an Introduction to Basic Theoretical Concepts and Experiments, STMP*, volume 173, pages 14–57. Springer-Verlag, 2001.
- [Wis96] H. Wiseman. Quantum trajectories and quantum measurement theory. *Quantum Semiclass.*, 8:205–222, 1996.
- [WM93] H. Wiseman and G. Milburn. Quantum theory of field-quadrature measurements. *Phys. Rev. A*, 47:642–662, 1993.
- [Wol67] J. A. Wolf. *Spaces of constant curvature*. McGraw-Hill, 1967.
- [YL73] H. P. Yuen and M. Lax. Multiple-parameter quantum estimation and measurement of non-selfadjoint observables. *IEEE Trans. Inform. Theory*, 19:740, 1973.
- [ZS05] K. Życzkowski and H. J. Sommers. Average fidelity between random quantum states. *Phys. Rev. A*, 71:032313, 2005.
- [Zur82] W. Zurek. Environment-induced superselection rules. *Phys. Rev. D*, 26:1862–1880, 1982.
- [ZVB04] A. Zavatta, S. Viciani, and M. Bellini. Quantum to classical transition with single-photon-added coherent states of light. *Science*, 306:660–662, 2004.

Samenvatting

Het meetproces in de kwantummechanica is een wat bizarre aangelegenheid. Iets wat in het dagelijks leven volstrekt vanzelfsprekend is, zoals het gelijktijdig meten van twee verschillende dingen, kan op kleine schaal opeens volstrekt onmogelijk blijken.

In hoofdstuk 3 en 4 is een aantal van dit soort fundamentele beperkingen zo scherp mogelijk geformuleerd. Je kunt immers moeilijk je grenzen opzoeken, als je niet weet waar die liggen. Vervolgens is voor een tweetal meetmethoden vastgesteld hoeveel verbetering er theoretisch nog mogelijk is. Het zijn allebei methoden om de inwendige toestand van een atoom vast te stellen aan de hand van het licht dat het uitstraalt. De methode van hoofdstuk 6 blijkt optimaal te zijn, die van hoofdstuk 5 laat nog een klein beetje ruimte voor verbetering.

Hoofdstuk 7 gaat over een ander deel van het onderzoek. Alle velden die in de natuur voorkomen (elektrische velden, magnetische velden, . . .) bewegen mee met kleine vervormingen van de tijdruimte. Het belangrijkste resultaat in dit deel van het proefschrift is een classificatie van alle ‘meebewegende’ velden die theoretisch mogelijk zijn. Verrassend genoeg zijn dat er in zekere zin tamelijk weinig, en bovendien hangen de mogelijkheden sterk af van de vorm van de tijdruimte.

In de rest van deze samenvatting wil ik graag *zonder wiskundige formules* een preciezer beeld geven van de inhoud van dit proefschrift. Ik beperk me daarbij tot de wiskundige ideeën in hoofdstuk 7, omdat die het best uit te drukken zijn in woorden en vooral ook plaatjes.

Elektrische Velden en de Raakbundel

Een elektrisch veld is iets wat trekt aan geladen deeltjes. Hoe sterker het elektrisch veld, hoe harder de deeltjes in de richting van dat veld getrokken worden.

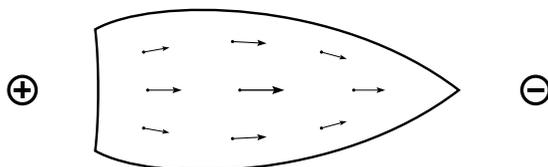


FIG. 30: Een elektrisch veld zet geladen deeltjes in beweging.

Dit is bijzonder handig, bijvoorbeeld als je een broek wilt strijken. Je neemt dan een dunne metalen plaat, die vol zit met geladen deeltjes (elektronen). Plaats deze bovenop de te strijken broek. Zorg vervolgens voor een elektrisch veld in de metalen plaat, bijvoorbeeld m.b.v. een stopcontact. De geladen deeltjes worden in de richting van het veld door de plaat getrokken, die door alle wrijving opwarmt. Door vervolgens de plaat zachtjes heen en weer over de broek te bewegen kan men vouwen en kreukels verwijderen, of deze er juist in strijken (in het geval van de auteur vooral dat laatste).

De Ruimte van Punten

Laten we proberen om het elektrisch veld wiskundig te beschrijven. Stap één is een beschrijving van de ruimte M van punten waar het veld kan leven. Dit vraagt om een zekere mate van algemeenheid. Wat M precies is, hangt immers af van de situatie.

In een dunne metalen plaat kan ieder punt bijvoorbeeld beschreven worden door 2 coördinaten, lengte (x) en breedte (y). De ruimte M heet dan 2-dimensionaal. Maar elektrische velden leven ook in de aarde, in de lucht en zelfs in de lege ruimte. Je hebt daar 3 coördinaten nodig om een punt te beschrijven: lengte (x), breedte (y) en hoogte (z). De ruimte M is dus 3-dimensionaal. Wonderlijk genoeg houdt het hier niet op. Volgens Einstein zijn tijd (1-dimensionaal) en ruimte (3-dimensionaal) ‘van hetzelfde spul gemaakt’. Er zijn dus 4 coördinaten nodig om een punt in de tijdruimte te beschrijven: lengte (x), breedte (y), hoogte (z) en tijd (t). Met andere woorden: de tijdruimte M is 4-dimensionaal.

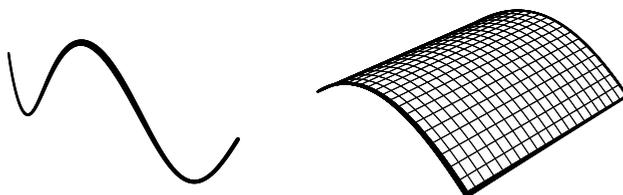


FIG. 31: Gladde variëteiten van dimensie 1 (links) en 2 (rechts).

Om dit alles in één klap te vangen, beschrijven we de ruimte M van punten als een *gladde variëteit* (smooth manifold). Gladde variëteiten heb je in elke gewenste dimensie. Een gladde variëteit van dimensie 1 is niets anders dan een (kromme) lijn. Een gladde variëteit van dimensie 2 is een (gekromd) oppervlak.

Zoals gezegd is onze tijdruimte M een gladde variëteit van dimensie 4. Dit is lastig te tekenen, zeker aangezien de tijdruimte ook nog eens gekromd is. Maar wiskundig gezien zijn 4 dimensies niet veel moeilijker dan 1 of 2. Bij woorden als ‘gladde variëteit’, ‘ruimte’ of ‘tijdruimte’ kan men dus gerust figuur 31 in het achterhoofd houden.

De Raakbundel

Stap één was het beschrijven van de ruimte M van punten. Stap twee is het veld in een vast punt⁸ x van M . Het veld geeft aan hoe hard er aan een geladen deeltje wordt getrokken. Hoe sterker het veld, hoe groter de trekkracht. Maar daarnaast bepaalt het veld ook in welke richting er wordt getrokken. Het veld in x is dus iets wat naast grootte ook richting heeft: een pijl of *vector*. Hij moet natuurlijk raken aan M , omdat de deeltjes anders de ruimte zouden verlaten. Kortom: het elektrisch veld in een punt x beschrijven we wiskundig als een vector die aangrijpt in x en raakt aan M .

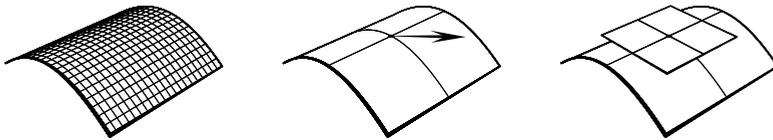


FIG. 32: Links een gladde variëteit M van dimensie 2. In het midden een raakvector die aangrijpt in het punt x . Rechts het raakvlak in dit punt. Alle vectoren die raken aan M en aangrijpen in x hebben hun kop in dit vlak.

Voor de derde en laatste stap in de beschrijving van het veld moeten we ons realiseren dat het elektrisch veld niet louter leeft in één punt. Voor ieder mogelijk punt in de ruimte bepaalt het hoe hard er aan geladen deeltjes getrokken wordt. We beschrijven het veld dan ook als een hele familie vectoren, één raakvector die aangrijpt in x voor elk punt x van de ruimte M .

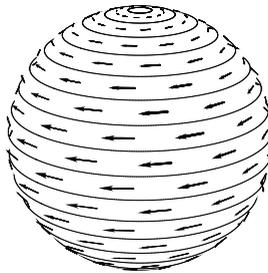


FIG. 33: De bolschil is een 2-dimensionale gladde variëteit. Het elektrisch veld hierop wordt beschreven door een snede van zijn raakbundel. Dit is iets wat in ieder punt x één raakvector uitkiest: de waarde van het veld in x .

Alles bij elkaar geeft dit een wiskundige beschrijving van het elektrisch veld. De ruimte waarop het veld leeft, beschrijven we met een gladde variëteit M .

⁸We schrijven kortweg x in plaats van (x, y) , (x, y, z) of (x, y, z, t) . Op een heel proefschrift scheelt dat alras een bladzij.

Bij ieder punt x van M hoort een *raakvlak* (tangent space) T_xM , de ruimte van alle vectoren die raken aan M en aangrijpen in x (zie figuur 32). Het raakvlak in x beschrijft precies alle mogelijke waarden die het elektrisch veld in x kan aannemen. We definiëren de *raakbundel* (tangent bundle) als de familie van alle mogelijke raakvlakken T_xM , waar x loopt over de punten van M .

We beschrijven het elektrisch veld als een *sede* van de raakbundel. Dit is iets wat in ieder raakvlak T_xM precies één vector uitkiest: de waarde van het veld in x .

Algemene Velden en Vezelbundels

Het elektrisch veld is niet het enige veld. Zo bestaat er bijvoorbeeld een drukveld, dat als het ware bijhoudt in hoeverre een rubber balletje zou worden samengedrukt als het de moed had op een bepaalde plek in de ruimte zijn bol gezicht te vertonen. Een trilling in het drukveld ervaren wij als geluid. Ook het elektrisch veld kan trillen: dit doet het samen met zijn tweelingbroer, het magnetisch veld. Rond 1864 ontdekte J.C. Maxwell⁹ dat licht precies bestaat uit dit soort gecombineerde elektromagnetisch trillingen.

Maar ook in de moderne natuurkunde spelen velden een hoofdrol. Het hele idee van een ‘fundamenteel deeltje’ is daar zelfs vervangen door dat van een ‘fundamenteel veld’! Deeltjes zijn dan niets anders dan trillingen in dit veld.

Het loont dus de moeite om de voorgaande bespiegelingen iets algemener te trekken, en een wiskundige beschrijving te geven voor velden in het algemeen. Dit hebben we gedaan in hoofdstuk 2.4. Een sleutelrol is daarbij weggelegd voor het begrip *vezelbundel* (fibre bundle), dat de raakbundel generaliseert.

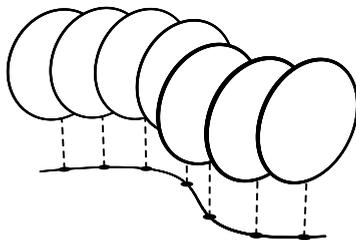


FIG. 34: Schets van een vezelbundel. De basisruimte M is de kromme beneden. Bij ieder punt x op de kromme hoort een vezel F_x . Alle vezels hebben hier de vorm van een schijf.

Een vezelbundel over een gladde variëteit M is per definitie een familie van *vezels*, één vezel F_x bij ieder punt x van M . De vezel bij het punt x stelt de verzameling voor van waarden die het veld mag aannemen in x . Die vezels

⁹Hij publiceerde zijn bevindingen in "A Dynamical Theory of the Electromagnetic Field", *Phil. Trans. R. Soc. Lond.* **155**, blz. 459–512, 1865. Maxwell is veel geprezen voor zijn ‘ontdekking van het licht’, iets wat voor een man met die initialen nochtans geen grote opgaaf zal zijn geweest.

mogen van alles zijn, maar we eisen wel dat alle vezels er ‘hetzelfde uitzien’ (in jargon: isomorf zijn).

Een belangrijk voorbeeld van een vezelbundel is de raakbundel. De vezel F_x bij een punt x wordt hier gevormd door het raakvlak T_xM , de ruimte van alle vectoren die raken aan M en aangrijpen in het punt x . Alle vezels zien er hetzelfde uit, in dit geval als vlakken.

In het algemeen beschrijven we een veld als een *snede* (section) van de vezelbundel. Een snede is iets wat in iedere vezel F_x precies één waarde uitkiest.

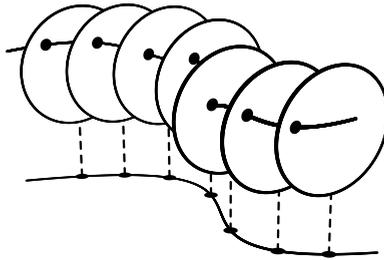


FIG. 35: Schets van een snede. In iedere vezel is één waarde aangewezen.

De wiskundige beschrijving van een veld werkt dus als volgt. De ruimte waarop het veld leeft, beschrijven we met een gladde variëteit M . Het soort veld (bijvoorbeeld: het elektrisch veld) wordt bepaald door een vezelbundel (bijvoorbeeld: de raakbundel). Bij ieder punt x van M hoort dan een vezel F_x van waarden die het veld daar aan kan nemen (voor de raakbundel is dit het raakvlak T_xM). Het veld zelf wordt beschreven door een snede van de vezelbundel. Die wijst in iedere vezel F_x één waarde aan: de waarde van het veld in x . (Een snede van de raakbundel wijst bij ieder punt dus één raakvector aan.)

Natuurlijke Bundels

Alle vezelbundels zijn gelijk, maar sommige vezelbundels zijn gelijkere dan andere. Een vezelbundel heet *natuurlijk* als iedere vervorming (jargon: diffeomorfisme) van de ruimte M een welbepaald effect heeft op zijn vezels. (Zie definitie 7.1 op blz. 154 voor een preciezere formulering.) Je kunt de natuurlijke bundels zien als een soort ‘elite’ onder de vezelbundels.

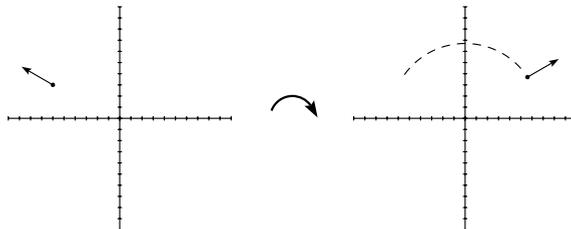


FIG. 36: Een raakvector draait mee als je het vlak draait.

Elke raakbundel is natuurlijk. Bekijk bijvoorbeeld het platte vlak M , met een raakvector die aangrijpt in een punt x van M . Vervolgens draaien we het vlak. Wat gebeurt er met de vector? Het antwoord is duidelijk: *de vector beweegt mee als je M roteert*, zie figuur 36. De raakbundel heet ‘natuurlijk’ omdat elke vervorming van de ruimte M een welbepaald effect heeft op de raakvectoren.

Dit heeft belangrijke gevolgen voor het elektrisch veld, dat wordt beschreven door een snede van de raakbundel. Bij een vervorming van de ruimte M is er een welbepaald effect op de raakvectoren in *ieder* afzonderlijk punt van M . Omdat een elektrisch veld niets anders is dan een manier om aan *alle* punten van M een raakvector in dat punt toe te kennen, moeten we concluderen dat *het gehele veld meeverandert* als M wordt vervormd. Zie figuur 37.

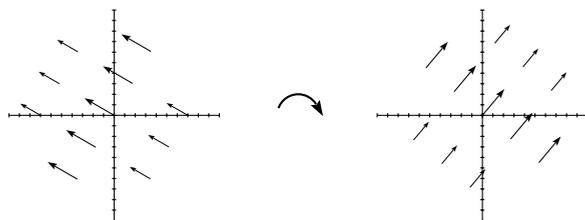


FIG. 37: Ook het veld draait mee als je het vlak draait.

Aan het einde van hoofdstuk 2.4 wordt uitgelegd dat het meebewegen van het veld bij kleine vervormingen van M noodzakelijk is om de *energie* (jargon: de ‘stress-energy-momentum tensor’) van een veld te definiëren.

Het feit dat elektrische en magnetische velden energie kunnen bevatten, hoeft niet te verbazen. Licht is immers niets anders dan een golf in het elektromagnetisch veld. Iedereen die wel eens roodverbrand onder de zonnepanels vandaan is gekomen, heeft het energiedragend vermogen van het elektromagnetisch veld dus experimenteel vastgesteld, om niet te zeggen aan den lijve ondervonden.

Energietransport door het elektromagnetisch veld heeft trouwens ook bijverschijnselen die prettiger zijn dan eerstegraads brandwonden. Leven op aarde, om eens wat te noemen. De energie die het leven hier mogelijk maakt, komt direct of indirect van de zon. Maar als het elektromagnetisch veld geen energie zou dragen, dan zou al die zonne-energie vrolijk op de zon achterblijven. De aarde zou dan koud en levenloos achterblijven.

Enfin, elektrische velden kunnen dus, samen met magnetische velden, energie overdragen. Dit komt doordat de raakbundel natuurlijk is. Hoe zit dat met andere velden?

We weten uit de natuurkunde dat velden eigenlijk altijd in staat zijn energie te dragen. Het ligt dus voor de hand om velden te beschrijven met *natuurlijke* vezelbundels. Het eerste voorbeeld van een natuurlijke vezelbundel hebben we dus al gezien: de raakbundel.

Natuurlijke bundels zijn witte raven onder de vezelbundels. Een algemene vezelbundel is zomaar een familie vezels F_x , zonder enige verdere structuur. In het algemeen heb je geen flauw idee hoe de vezels zouden moeten meeveranderen

als de ruimte M vervormt, en is het zelfs nog maar de vraag of dit hoe dan ook wel op een consistente manier kan. Als een vezelbundel daarentegen natuurlijk is, heeft elke vezel F_x als het ware een band met zijn bijbehorend punt x . Als de ruimte M vervormt, ‘voelen’ de vezels F_x hoe hun eigen punt beweegt, en bewegen mee, vgl. figuur 36 en 37.

Infinitiesimaal Natuurlijke Bundels

Een bundel verdient dus de eretitel ‘natuurlijk’ als elke vervorming van de ruimte M een welbepaald effect heeft op de vezels F_x . Die eis is soms een beetje al te zwaar. We noemen een bundel daarom *infinitiesimaal natuurlijk* als elke *kleine* vervorming van M een welbepaald effect heeft op de vezels F_x . (Zie ook definitie 7.2 op blz. 154.) Dat is een soort troostprijs. Als je een bundel bent, is infinitesimaal natuurlijk zijn niet zo goed als écht natuurlijk zijn, maar het is toch al heel wat. Voor het definiëren van energie zijn bijvoorbeeld alleen *kleine* vervormingen nodig. Met andere woorden, een veld kan energie dragen precies dan als hij hoort bij een infinitesimaal natuurlijke bundel.

De zogenaamde *spinorbundels* zijn infinitesimaal natuurlijke vezelbundels bij uitstek. De spinorbundel bij het platte vlak M is precies de raakbundel, behalve dat vectoren *op halve snelheid* met de ruimte meedraaien, zie figuur 38. Zo een halfslachtige vector heet ook wel een *spinor*.

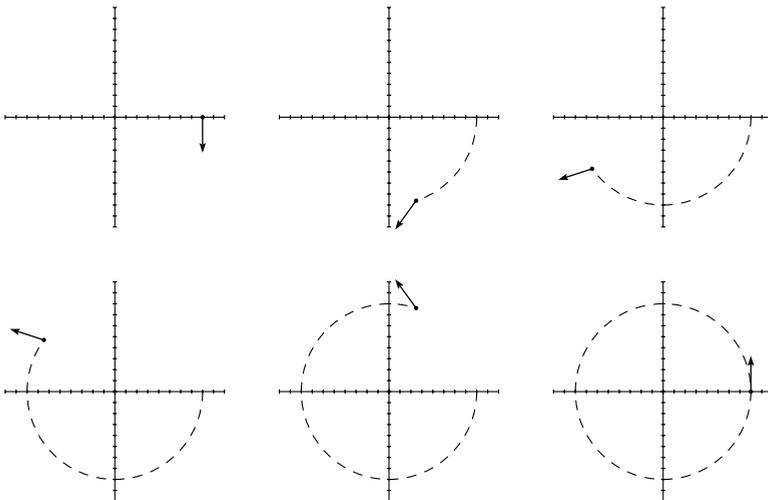


FIG. 38: Een spinor draait mee met de ruimte, maar half zo snel.

Als de ruimte 360° om haar as draait, is er netto helemaal niets gebeurd. Als je immers een heel rondje draait, is daarna alles weer precies zoals het geweest is. In figuur 38 valt echter te zien dat de spinor desondanks is omgeklapt, van \downarrow naar \uparrow . Dat is een beetje merkwaardig, en deze kleine ongerijmdheid komt de arme spinorbundel dan ook op verlies van het predicaat ‘natuurlijk’ te staan.

Dit soort problemen ontstaat echter alleen bij *grote* vervormingen van de ruimte. Bij *kleine* vervormingen bewegen de spinoren wél netjes mee. We strijken daarom over ons hart en reiken de troostprijs uit: spinorbundels mogen dan niet ‘natuurlijk’ zijn, ze zijn ten minste nog wel ‘infinitesimaal natuurlijk’.

Een spinorveld is een snede van de spinorbundel. Het wijst in ieder punt van de ruimte M één spinor aan, de waarde van het veld in dat punt. Spinorvelden worden veel gebruikt in de natuurkunde. Je zou zeggen dat de rare flip van \downarrow naar \uparrow onder een rotatie van 360° problemen oplevert, maar dat blijkt mee te vallen. De reden hiervoor is wat subtiel; het komt er uiteindelijk op neer dat een spinorveld nooit direct kan worden waargenomen, we zien alleen zijn kwadraten. En in het kwadraat van een spinorveld heb je twee flips, die elkaar opheffen.

Het is ook zeker niet zo dat spinorvelden alleen gebruikt worden om bizarre effecten te beschrijven in obscure hoekjes van de moderne fysica. In het ‘standaard model’, de natuurkundige theorie die al sinds de jaren '70 vrijwel algemeen geaccepteerd wordt, wordt alle huis-tuin-en-keukenmaterie beschreven door spinorvelden. Net als licht een trilling is in het elektromagnetisch veld, zijn sommige elementaire deeltjes trillingen in een spinorveld. Twee van dat soort deeltjes, de *elektronen* en de *quarks*, klonteren via een aantal tussenstappen aan een tot atomen. Atomen op hun beurt zijn weer de bouwstenen van alle materie om ons heen. Welbeschouwd is dus iedere steen, iedere appel, iedere hond, mens of neushoorn een serie trillingen in één groot spinorveld.

De Hamvraag

We hebben gezien dat verschillende soorten velden worden beschreven door verschillende soorten vezelbundels. Een vezelbundel komt echter alleen in aanmerking om een veld te beschrijven als hij *infinitesimaal natuurlijk* is. Zijn vezels moeten dan als het ware meebewegen met kleine vervormingen van de ruimte, zoals in figuur 36, 37 en 38.

Dit laat de natuur relatief weinig vezelbundels om uit te kiezen. Als we ze allemaal kunnen beschrijven, dan geeft ons dit een idee welke velden wél in de natuur kunnen voorkomen, en welke niet. De hamvraag is dus:

‘Welke vezelbundels zijn infinitesimaal natuurlijk?’

Het is natuurlijk niet zo dat, als een veld in de natuur *mag* voorkomen, hij dat ook automatisch *doet*. Het infinitesimaal natuurlijk zijn van zijn vezelbundel moet je dan ook zien als een minimumvereiste voor een veld om in de natuur voor te komen.

Hoofdstuk 7 is gewijd aan het beantwoorden van de hamvraag. Het belangrijkste resultaat is de classificatiestelling 7.14 op bladzijde 167. Zij zegt dat alle infinitesimaal natuurlijke vezelbundels afstammen van één overzichtelijke oerfamilie, de zogenaamde ‘universele overdekkingen van bundels van k -frames’.

Hoe die beesten er precies uitzien, is vrij goed bekend (zie sectie 7.4.3 voor een algemene beschrijving en figuur 39 voor een eenvoudig voorbeeld). Van belang is vooral dat het er niet al te veel zijn, en dat hun structuur sterk afhangt van de gladde variëteit M . Voor iedere M geeft ons dit een ‘grijze’ lijst van

velden die infinitesimaal natuurlijk zijn en dus in de natuur kunnen voorkomen, en een ‘zwarte’ lijst van velden die dat zeker niet kunnen.

De grap is nu dat we eigenlijk al wel een heel aardig beeld hebben van de velden die in het wild voorkomen. De laatste decennia hebben fysici (CERN, Fermilab etc.) enorme deeltjesversnellers gebouwd, waarmee een ‘witte’ lijst is opgesteld van velden die echt gemeten zijn, en die dus zeker in de natuur voorkomen.

Aan de andere kant hebben we eigenlijk geen flauw benul van de vorm van onze 4-dimensionale tijdruimte M . Zij is op grote schaal recht voorzover wij kunnen zien, maar dat is helaas niet ver genoeg om uitsluitsel te kunnen geven.

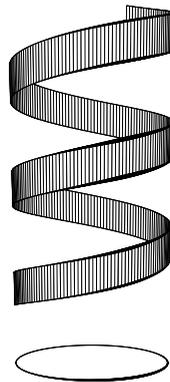


FIG. 39: Schets van de universele overdekking van de bundel van k -frames over de cirkel, voor het geval $k = 1$.

We willen dus eigenlijk stelling 7.14 ondersteboven houden! In plaats van alle mogelijke velden te bepalen voor een gegeven tijdruimte M , bepalen we nu alle mogelijke tijdruimten, gegeven de ‘witte’ lijst van velden waarvan we al weten dat ze voorkomen. Dit is de portee van stelling 7.23 op bladzijde 180.

Een aantal mogelijkheden kun je nu al direct afstrepen. De gladde 4-dimensionale variëteit M met de naam $\mathbb{C}P^2$ valt bijvoorbeeld af, omdat zijn ‘zwarte lijst’ volgens stelling 7.14 overlapt met de ‘witte lijst’ van waargenomen velden. Als het beeld dat we hebben van de elementaire velden correct en compleet is, kan de tijdruimte M dus nooit van deze vorm zijn.

Welbeschouwd is dit een wonderlijke conclusie: door goed te kijken naar de kleinst mogelijke structuren, de elementaire deeltjes waaruit alle materie is opgebouwd, krijgen we informatie over de vorm van ons heelal.

Dankwoord

Voor de inhoud van dit proefschrift draagt de auteur de ongedeelde verantwoordelijkheid, maar zeker niet de ongedeelde eer. Graag wil ik hier iedereen van harte danken die er een bijdrage aan heeft geleverd.

In de eerste plaats zijn dat natuurlijk mijn copromotoren, Hans Maassen en Johan van de Leur, die mij vertrouwd hebben gemaakt met respectievelijk de kwantumkansrekening en de Lie-theorie. De vele plezierige discussies met Hans tijdens het laatste jaar van mijn studie in Nijmegen hebben de basis gevormd voor hoofdstuk 3 en 4. De eerste daarvan is dan ook gepubliceerd onder ons beider naam. Johan heeft mij begeleid bij mijn promotietraject in Utrecht en ik wil hem graag bedanken voor alle hulp en raad die ik van hem heb gehad, al dan niet wiskundig van aard. Voor zijn enthousiaste en essentiële hulp bij hoofdstuk 7 ben ik ook Marius Crainic veel dank verschuldigd.

Verder wil ik graag mijn coauteurs bedanken voor de fijne en leerzame samenwerking. Hoofdstuk 5 is geschreven samen met Luc Bouten, en hoofdstuk 6 samen met Mădălin Guță en Jonas Kahn. Graag dank ik ook de respectievelijke thuisbases van Luc en Mădălin voor hun gastvrijheid: het Mabuchilab, toen nog aan het California Institute of Technology, en de School of Mathematical Sciences aan de University of Nottingham.

Hartelijk dank ook aan mijn promotor, Roberto Fernández, en aan de leescommissie: Erik van den Ban, Nilanjana Datta, Mark Fannes, Eduard Looijenga en Christoph Wockel.

Ten slotte wil ik graag mijn medewiskundigen bedanken voor alle interessante gesprekken, waar ik veel van geleerd heb. Bij naam wil van hen slechts diegenen noemen die ik helaas niet meer persoonlijk bedanken kan: Hanneke Janssen en Hans Duistermaat.

Curriculum Vitae

Bas Janssens werd op 20 mei 1981 geboren te Maastricht en behaalde in 1999 cum laude zijn gymnasiumdiploma aan het Stella Maris College in Meerssen. Hij studeerde natuurkunde (1999-2004) en wiskunde (2000-2005) aan de Katholieke Universiteit Nijmegen (nu Radboud Universiteit Nijmegen), waar hij onder begeleiding van Hans Maassen cum laude afstudeerde in de richting van de kwantumkansrekening. Vervolgens bracht hij korte tijd door aan het California Institute of Technology (waar hij samenwerkte met Luc Bouten) en aan de University of Nottingham (waar hij samenwerkte met Jonas Kahn en Mădălin Guță). In 2006 begon hij zijn promotieonderzoek in de richting van oneindig-dimensionale Lie-algebras aan de Universiteit Utrecht, onder begeleiding van Johan van de Leur. Dit proefschrift is niet enkel een verslag van dit promotieonderzoek, maar van al zijn bevindingen tot nu toe.

