



THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

Exploring Quantum Foundations with Single Photons

Martin Ringbauer
BSc. BSc. The University of Vienna



*A thesis submitted for the degree of Doctor of Philosophy at
The University of Queensland in 2016
School of Mathematics and Physics*

Abstract

Exploring Quantum Foundations with Single Photons

Martin Ringbauer

School of Mathematics and Physics

The University of Queensland

Quantum mechanics is our most successful physical theory and has been confirmed to extreme accuracy, yet, a century after its inception it is still unclear what it says about the nature of reality. In this thesis, I explore some of the foundational questions which are central to our understanding of quantum mechanics using single photons as an experimental platform. Three experiments form the core of the thesis, studying, respectively, the role of reality, causality, and uncertainty in quantum mechanics. These experiments shed light on decade-old questions that have previously been thought to be outside the realm of experimental physics. The results contribute to our understanding of the structure of quantum mechanics, and reveal novel aspects of phenomena that were believed to be well understood. In three further experiments I also touch upon the topics of finding physical principles behind quantum mechanics, quantum effects in extreme gravitational fields, and developing a pathway towards tests of macroscopic quantum phenomena. Starting out as a fringe discipline, the field of quantum foundations has developed into an influential area of research. It is now becoming possible to turn many of the foundational questions in quantum mechanics from topics of philosophy into topics of physics. Subjecting these questions to rigorous experimental tests the field is making progress in the quest for understanding our best physical theory.

Declaration by author

This thesis is composed of my original work, and contains no material previously published or written by another person except where due reference has been made in the text. I have clearly stated the contribution by others to jointly-authored works that I have included in my thesis.

I have clearly stated the contribution of others to my thesis as a whole, including statistical assistance, survey design, data analysis, significant technical procedures, professional editorial advice, and any other original research work used or reported in my thesis. The content of my thesis is the result of work I have carried out since the commencement of my research higher degree candidature and does not include a substantial part of work that has been submitted to qualify for the award of any other degree or diploma in any university or other tertiary institution. I have clearly stated which parts of my thesis, if any, have been submitted to qualify for another award.

I acknowledge that an electronic copy of my thesis must be lodged with the University Library and, subject to the policy and procedures of The University of Queensland, the thesis be made available for research and study in accordance with the Copyright Act 1968 unless a period of embargo has been approved by the Dean of the Graduate School.

I acknowledge that copyright of all material contained in my thesis resides with the copyright holder(s) of that material. Where appropriate I have obtained copyright permission from the copyright holder to reproduce material in this thesis.

Publications during candidature

Peer-reviewed publications

1. *Experimental Joint Quantum Measurements with Minimum Uncertainty*,
M. Ringbauer, D. Biggerstaff, M.A. Broome, A. Fedrizzi, C. Branciard, A.G. White
Physical Review Letters **112**, 020401 (2014). ([arXiv:1308.5688](#))
2. *Experimental Simulation of Closed Timelike Curves*,
M. Ringbauer, M.A. Broome, C.R. Myers, A.G. White, T.C. Ralph
Nature Communications **5**, 4145 (2014). ([arXiv:1501.05014](#))
3. *Information Causality in the Quantum and Post-Quantum Regime*,
M. Ringbauer, A. Fedrizzi, D.W. Berry, A.G. White
Scientific reports **4**, 6955 (2014). ([arXiv:1406.5034](#))
4. *Experimental characterization of quantum dynamics with initial system-environment correlations*,
M. Ringbauer, C.J. Wood, K. Modi, A. Gilchrist, A.G. White, A. Fedrizzi
Physical Review Letters **114**, 090402 (2015). ([arXiv:1410.5826](#))
5. *Measurements on the Reality of the Wavefunction*,
M. Ringbauer, B. Duffus, C. Branciard, E.G. Cavalcanti, A.G. White, A. Fedrizzi
Nature Physics **11**, 249-254 (2015). ([arXiv:1412.6213](#))

Online e-print publications

6. *Experimental Test of Nonlocal Causality*,
M. Ringbauer, C. Giarmatzi, R. Chaves, F. Costa, A.G. White, A. Fedrizzi
Science Advances **2**, e1600162-e1600162 (2016). ([arXiv:1602.02767](#))
7. *Generation of Mechanical Interference Fringes by Multi-Photon Quantum Measurement*,
M. Ringbauer, T.J. Weinhold, A.G. White, M.R. Vanner
([arXiv:1602.05955](#))

Publications included in this thesis

Information Causality in the Quantum and Post-Quantum Regime,

M. Ringbauer, A. Fedrizzi, D.W. Berry, A.G. White

Scientific reports **4**, 6955 (2014).

Reproduced in Appendix A.

Contributor	Statement of contribution
M. Ringbauer (Candidate)	Design and construction of experiment (70%) Data acquisition (100%) Data analysis and interpretation (65%) Complete first draft of manuscript (40%) Final draft of manuscript (25%) Referee replies and final manuscript revision (30%) Initial development of concept (10%)
A. Fedrizzi	Design and construction of experiment (30%) Data analysis and interpretation (35%) Complete first draft of manuscript (40%) Final draft of manuscript (25%) Referee replies and final manuscript revision (30%) Initial development of concept (30%)
D.W. Berry	Theoretical model (100%) Complete first draft of manuscript (20%) Final draft of manuscript (25%) Referee replies and final manuscript revision (20%) Initial development of concept (50%)
A.G. White	Data analysis and interpretation (10%) Final draft of manuscript (25%) Referee replies and final manuscript revision (20%) Initial development of concept (10%)

Experimental Simulation of Closed Timelike Curves,

M. Ringbauer, M.A. Broome, C.R. Myers, A.G. White, T.C. Ralph

Nature Communications **5**, 4145 (2014).

Reproduced in Appendix B.

Contributor	Statement of contribution
M. Ringbauer (Candidate)	Design and construction of experiment (60%) Data acquisition (100%) Data analysis and interpretation (60%) Complete first draft of manuscript (45%) Final draft of manuscript (20%) Referee replies and final manuscript revision (30%) Initial development of concept (25%)
M. A. Broome	Design and construction of experiment (10%) Data analysis and interpretation (20%) Complete first draft of manuscript (35%) Final draft of manuscript (20%) Referee replies and final manuscript revision (30%) Initial development of concept (25%)
C.R. Myers	Design and construction of experiment (10%) Theoretical model (50%) Data analysis and interpretation (10%) Complete first draft of manuscript (10%) Final draft of manuscript (20%) Referee replies and final manuscript revision (10%) Initial development of concept (25%)
A.G. White	Complete first draft of manuscript (10%) Final draft of manuscript (20%) Referee replies and final manuscript revision (20%)
T.C. Ralph	Design and construction of experiment (10%) Theoretical model (50%) Data analysis and interpretation (10%) Complete first draft of manuscript (10%) Final draft of manuscript (20%) Referee replies and final manuscript revision (10%) Initial development of concept (25%)

Reproduced in Appendix C. Preprint at arXiv:1602.05955

Contributor	Statement of contribution
M. Ringbauer (Candidate)	Design and construction of experiment (55%) Data acquisition (100%) Data analysis and interpretation (70%) Complete first draft of manuscript (50%) Final draft of manuscript (25%) Referee replies and final manuscript revision (35%) Initial development of concept (20%)
T.J.Weinhold	Design and construction of experiment (15%) Data analysis and interpretation (10%) Final draft of manuscript (25%) Referee replies and final manuscript revision (20%) Initial development of concept (20%)
A.G. White	Design and construction of experiment (15%) Data analysis and interpretation (10%) Final draft of manuscript (25%) Referee replies and final manuscript revision (20%) Initial development of concept (20%)
M.R. Vanner	Design and construction of experiment (15%) Theoretical model (100%) Data analysis and interpretation (10%) Complete first draft of manuscript (50%) Final draft of manuscript (25%) Referee replies and final manuscript revision (35%) Initial development of concept (40%)

Contributions by others to the thesis

No contributions by others.

Statement of parts of the thesis submitted to
qualify for the award of another degree

None.

Acknowledgements

My first thanks goes to my advisory team. For his unmatched enthusiasm for physics and pretty much anything, I would like to thank Andrew White. Thank you for giving me the independence to pursue my interests and for providing guidance and advice along the way. For being an exceptional advisor, I would like to thank Alessandro Fedrizzi. Your trust, guidance, and advice throughout my time at UQ was invaluable. For introducing me to everything in and around the lab, I thank Till Weinhold, Marcelo Almeida, and Matthew Broome ('The day we understand the beam splitter...').

A special thanks goes to Philip Walther, who set me on the path to quantum foundations. Thank you for the unique opportunities you gave me, and for your enthusiasm and support throughout my undergraduate time.

The work in this thesis would not have been possible without all the talented physicists that I was privileged to work with. I would like to acknowledge all my co-authors, colleagues and fellow PhD students, who contributed to my experience at UQ, especially Juan Loredó for being a great office mate and for countless discussions; Markus Rambach, Sahar Basiri, Nathan Walk, Jacques Pienaar, Saleh Rahimi-Keshari and Farid Shahandeh for interesting and inspiring discussions; Ben Duffus for a great time in the lab; Roberto Muñoz and Sarah Lau for their persistence in the lab; Chris Wood for introducing me to tensor networks so that I can bug everyone else with it; the gremlin hunter, Mike Goggin for a great time in and outside the lab; and last but not least T. Vulpecula for all those little presents still hidden in the lab.

I would also like to acknowledge my theory colleagues Cyril Branciard, Eric Cavalcanti, Fabio Costa, Howard Wiseman, Ivan Kassal, Tim Ralph, and Gerard Milburn, for inspiring discussions. A special thanks to Cyril for making sure I am precise with every detail. I would like to thank all the philosophers I was privileged to work with, especially Sally Shrapnel and Peter Evans, for introducing me to the ways of philosophy of science.

My sincere apologies go to Glen Harris and David McAuslan for dragging you out of bed for our 6:00am climbing sessions. You and all the other friends I made here made my time.

A special thanks goes to Nina for her support throughout my studies, and for dragging me out of the lab every now and then.

Finally, I am extremely grateful for having such a loving and supportive family. My grandparents for their loving nature and countless childhood adventures. My late grandma, unbeatable in her card games, for her unique, sarcastic sense of humour that will always be remembered. Mum and dad, for their unconditional support, great advice, and encouragement. My sisters, who I can always count on, and my cousin, who is always up for a good time.

Keywords

Keywords: quantum foundations, quantum information, single photons, uncertainty relations, interpretations of quantum mechanics, causality, causal models, closed timelike curves, optomechanics, experimental metaphysics

Australian and New Zealand Standard Research Classifications (ANZSRC)

ANZSRC code: 020603 Quantum Information, Computation and Communication 30%

ANZSRC code: 020604 Quantum Optics, 30%

ANZSRC code: 020699 Quantum Physics not elsewhere classified, 40%

Fields of Research (FoR) Classification

FoR code: 0206 Quantum Physics, 100%

Contents

Abstract	i
Declaration by author	ii
Publications during candidature	iii
Publications included in this thesis	iv
Acknowledgements	ix
1 Quantum Information Basics	5
1.1 Quantum States	5
1.1.1 Qubits	5
1.1.2 Beyond Qubits	8
1.1.3 Comparing Quantum States	9
1.1.4 Composite Systems and Entanglement	10
1.2 Quantum Channels	14
1.2.1 Working with Quantum Channels	15
1.2.2 Process Matrix Representation	16
1.3 Quantum Measurements	18
1.3.1 Inconsistent Prescriptions	20
1.3.2 Weak Measurements	21
1.3.3 Quantum Non-Demolition Measurements	22
1.4 Quantum Information in Practice	23
1.4.1 Single Photons	23
1.4.2 Manipulating Polarization Qubits	25
References	30
2 Quantum Tomography	35
2.1 Introduction	35
2.2 Quantum State Tomography	36
2.2.1 Linear Inversion Tomography	37
2.2.2 Maximum Likelihood Estimation	38
2.2.3 Zero Probabilities	40
2.2.4 Error Bars for Quantum Tomography	40
2.2.5 Quantum Process Tomography	41
2.2.6 Caveats and Generalizations	42
2.3 Taming Non-Completely-Positive Maps with Superchannels	43
2.3.1 Constructing the Superchannel \mathcal{M}	44
2.3.2 Superchannel Tomography	46

2.3.3	Superchannels in the Wild	48
2.3.4	Quantifying Initial Correlations	49
2.3.5	Preparation Fidelity	51
2.3.6	Discussion	52
	References	54
3	Introduction to Quantum Foundations	58
3.1	Probability and Randomness	58
3.2	Quantum Correlations	59
3.2.1	EPR Paradox	59
3.2.2	Bell's Theorem	60
3.2.3	A Hierarchy of Correlations	62
3.2.4	Superquantum Correlations	63
3.2.5	Correlation Polytopes	65
3.3	Contextuality	68
3.3.1	Kochen-Specker Contextuality	70
3.3.2	Universal (Operational) Contextuality	71
3.3.3	Noncontextuality and Classicality	72
3.3.4	Wigner Negativity	72
3.4	Paradoxes and Basic Games	73
3.4.1	Quantum Teleportation	74
3.4.2	GHZ Paradox and Multipartite Entanglement	76
3.4.3	Hardy's Paradox	77
3.4.4	Leggett-Garg and Macrorealism	78
3.5	Pre- and Post-Selection Paradoxes	79
3.5.1	Weak Values	80
3.5.2	Classical Anomalous Weak Values	81
	References	84
4	On the Reality of the Wavefunction	93
4.1	Introduction	93
4.1.1	Ontic or Epistemic	95
4.1.2	The Quantum Wavefunction	95
4.1.3	Interpretations of Quantum Mechanics	98
4.2	The Ontological Models Framework	101
4.2.1	The ψ -Ontic/ ψ -Epistemic Distinction	103
4.2.2	Distinguishing Quantum States	104
4.2.3	Known Constraints	107
4.2.4	Examples of Ontological Models	110
4.3	ψ -Ontology Theorems	111
4.3.1	The Pusey-Barrett-Rudolph Theorem	111

4.3.2	Other Theorems, Other Assumptions	114
4.4	Constraining ψ -epistemic models	117
4.4.1	How to Constrain ψ -Epistemic Models	118
4.4.2	Maths Exercise	119
4.4.3	PP-Incompatibility	122
4.4.4	Experimental Robustness	122
4.5	Testing Realist ψ -Epistemic Models	123
4.5.1	Choosing States and Measurements	124
4.5.2	Experimental Setup	124
4.5.3	Experimental Results	130
4.5.4	Experimental Limitations and Extensions	131
4.5.5	Three-Outcome Measurement and Single-Output Decomposition	133
4.6	Discussion and Outlook	133
4.6.1	Where to From Here?	134
4.6.2	Limitations and Possible Extensions of the Method	136
4.6.3	Bell Inequalities and Device Independence	139
	References	141

5 Causality in a Quantum World 147

5.1	Introduction	147
5.2	The Causal Modeling Framework	148
5.2.1	A Causal Model	148
5.2.2	Reading the Causal Graph	152
5.2.3	Causal Discovery	153
5.3	Quantum Correlations: Bell’s Theorem and Beyond	154
5.3.1	Axiomatic Approach	155
5.3.2	Bell’s Assumptions	157
5.3.3	Causal Assumptions	160
5.3.4	No Fine-Tuning	160
5.3.5	Many Roads to Bell’s Theorem	162
5.3.6	Relaxing the Assumptions	163
5.4	Testing Causal Models for Quantum Correlations	167
5.4.1	Experimental Setup for Testing the CHSH Inequality	168
5.4.2	Interventional Approach	169
5.4.3	Observational Approach	171
5.5	Discussion and Outlook	173
5.5.1	Loopholes	174
5.5.2	Relation to Previous Work	175
5.5.3	Beyond Classical Causal Modeling	176
	References	178

6	Pushing Joint-Measurement Uncertainty to the Limit	183
6.1	Introduction	183
6.2	Heisenberg's Uncertainty Principle	184
6.2.1	The Uncertainty Principle	184
6.2.2	Incompatible Observables	186
6.2.3	Preparation Uncertainty	187
6.2.4	Joint-Measurement Uncertainty	189
6.2.5	Measurement-Disturbance	191
6.3	Measuring Measurement Uncertainty	192
6.3.1	Relating $\varepsilon_{\mathcal{A}}$ to Experimental Data	194
6.3.2	The Three-State Method	195
6.3.3	The Weak-Measurement Method	196
6.3.4	Inaccuracies from $\alpha_{\mathcal{M}}$	199
6.4	Testing Joint-Measurement Uncertainty Relations	200
6.4.1	Experimental Configuration	200
6.4.2	Three-State Method	201
6.4.3	Weak-Measurement Method	202
6.5	Discussion and Outlook	204
	References	206
7	Conclusion and Outlook	209
	References	215
A	Information Causality in the Quantum and Post-Quantum Regime	216
B	Experimental Simulation of Closed Timelike Curves	226
C	Generation of Mechanical Interference Fringes by Multi-Photon Quantum Measurement	236

List of Figures

<i>Coverimage</i>	© Martin Ringbauer and Benjamin Duffus	Cover
1.1	Classical and quantum bits.	6
1.2	Entanglement witness.	13
1.3	The action of a quantum channel in the Bloch sphere.	15
1.4	Basics of tensor network notation.	16
1.5	Different representations of a quantum channel.	17
1.6	Representation of measurements on the Bloch sphere.	18
1.7	POVMs in practice.	19
1.8	Two inconsistent prescriptions for quantum measurements.	21
1.9	Waveplate rotation in Bloch sphere.	26
1.10	Experimental implementation of a controlled-NOT gate.	28
1.11	Photons on a beam splitter	29
2.1	Quantum tomography.	36
2.2	The two most relevant choices of tomographically complete measurements for a single qubit.	37
2.3	Quantum tomography with an environment.	45
2.4	Graphical derivation of the superchannel \mathcal{M} for describing the reduced dynamics of a system initially correlated with the environment.	46
2.5	Experimental setup.	48
2.6	Reconstructed superchannel Choi matrix and effective channels from quantum process tomography.	50
2.7	Initial correlation norm.	51
2.8	Optimization of the preparation fidelity.	52
3.1	The scenario considered by Bell.	61
3.2	Hierarchy of correlations.	62
3.3	Quantum Random Access Coding.	64
3.4	Correlation polytopes in the simplest case.	67
3.5	Wigner quasi-probability distribution.	73
3.6	Quantum Teleportation	75
4.1	A classical point particle in position-momentum phase-space.	96
4.2	Classification of interpretations of quantum mechanics with respect to the status they attribute to the wavefunction.	99
4.3	The ontological models framework	102

4.4	The ψ -ontic/ ψ -epistemic distinction.	104
4.5	The ψ -epistemic explanation of imperfect distinguishability of non-orthogonal quantum states.	105
4.6	The experimental protocol for ququarts.	119
4.7	Mapping of the probability measures μ_ψ to sets C_ψ	120
4.8	Experimental setup.	125
4.9	The main sources of systematic imperfections in the setup.	126
4.10	Typical in-situ characterisation curve for a semi-circular waveplate.	128
4.11	Experimental violation of inequality (4.8) for qutrits and ququarts.	131
4.12	Error tolerance of the experimental protocol for testing maximally ψ -epistemic models.	132
4.13	Theoretical values for S using real, and complex-valued states and measurements.	132
4.14	Quantum overlaps of all states used in the experiment.	137
5.1	The causal modeling framework.	148
5.2	Basic notations of causal models.	149
5.3	The d -separation criterion.	153
5.4	Bell-local causal models.	155
5.5	Comparison of various constraints on the causal structure of Bell's theorem.	161
5.6	Many roads to Bell's theorem.	163
5.7	Relaxations of local causality.	164
5.8	Relaxations of measurement independence.	165
5.9	Relation between various explanations for Bell correlations.	166
5.10	The experimental setup.	168
5.11	Observed average causal effect $\text{ACE}_{A \rightarrow B}$ versus measured CHSH-value.	172
5.12	Experimental test of inequality (5.19).	173
6.1	The three faces of Heisenberg's uncertainty principle.	184
6.2	Compatible and incompatible qubit observables.	187
6.3	Preparation Uncertainty Principle.	187
6.4	Preparation Uncertainty Relation.	188
6.5	Joint-measurement uncertainty principle.	189
6.6	Joint-measurement uncertainty relations.	191
6.7	Measurement-disturbance principle.	192
6.8	Measurement-disturbance relations in the $\varepsilon_{\mathcal{A}}/\eta_{\mathcal{B}}$ -plane.	193
6.9	Conceptual idea of determining joint-measurement accuracy.	194
6.10	Estimation of $\alpha_{\mathcal{M}}$ in the presence of experimental imperfections.	196
6.11	Illustration of the experimental configuration.	201
6.12	Experimental setup for the three-state method.	202
6.13	Measurement inaccuracies $\varepsilon_{\mathcal{A}}$ and $\varepsilon_{\mathcal{B}}$ estimated using the three-state method.	203
6.14	Experimental setup for the weak-measurement method.	204

6.15 Measurement inaccuracies, $\varepsilon_{\mathcal{A}}$ vs. $\varepsilon_{\mathcal{B}}$ estimated using the weak-measurement method. 204

Preamble

In 1900 Lord Kelvin delivered a lecture at the Royal Institution of Great Britain at a time where the general sentiment was that physics was mainly understood and the only thing left was to perform more precise measurements. In this lecture, Kelvin pointed out that “the beauty and clearness of the dynamical theory, which asserts heat and light to be modes of motion, is at present obscured by two clouds”. These two clouds referred to the failure of the Michelson-Morley experiment to reveal the luminiferous aether, the hypothetical medium through which light travels; and the fact that the radiation laws of the time made unphysical predictions about the radiation emitted by a black body, known as the UV catastrophe. Just five years later, Kelvin’s two clouds gave rise to the theory of relativity and quantum mechanics, respectively—two theories which caused a radical change in our understanding of the physical world and went on to become the most successful physical theories. Today, quantum mechanics is the basis of all of modern physics with the notable exception of gravity, which remains the territory of relativity and is notoriously difficult to give a quantum description of. Over the last decades the counterintuitive predictions of quantum mechanics held up in every experimental test and were confirmed to unprecedented accuracy. Yet, a century after its inception, the cloud over the interpretation of quantum mechanics is yet to be lifted.

There are two parts to any physical theory: the mathematical formalism, which describes how the theory works, and the physical interpretation, which connects it to the real world. While the former is fixed for a given theory, the latter can, and does in general evolve. In classical mechanics these two go hand-in-hand. There is a parameter x , which is interpreted as the position of a particle, and there is a differential equation which describes how x changes over time, that is, how the particle moves through space and reacts to external forces. As physical theories become more advanced, this one-to-one correspondence breaks down. For example, Maxwell’s theory of electromagnetism was initially interpreted as describing stresses and movements in an elastic medium, the luminiferous aether, which carries light waves. This widely accepted interpretation turned out to be untenable after the famous Michelson-Morley experiment failed to discover the aether and its very existence was difficult to maintain in the light of other results such as the theory of relativity. Today it is understood that the aether does not exist and that electromagnetic waves do not need a medium to propagate. Despite this drastic change in interpretation, however, the mathematical formalism in terms of the Maxwell-Heaviside equations remained unchanged and makes the same experimental predictions.

Quantum mechanics takes this interpretational ambiguity to an embarrassing new level, with not two but over a dozen actively used interpretations. Much of the incentive for this development stems from attempts to finding a satisfactory resolution of the quantum measurement problem. The central object of the theory, the wavefunction, or quantum state, is used

to describe and make predictions about any kind of quantum system. According to quantum mechanics, the wavefunction evolves continuously and deterministically with the Schrödinger equation, which means that knowing the starting conditions, one can perfectly predict the future state of the system. In the course of such an evolution the wavefunction typically ends up in a quantum superposition and the system is, loosely speaking, in multiple states at the same time. Yet, in an experiment we always observe definite outcomes. This is captured by the so-called projection postulate, which asserts that, in a measurement, the wavefunction abruptly collapses probabilistically into one of the possible outcomes. The problem, however, is that quantum mechanics provides no clue as to what counts as a measurement and what does not, neither does it explain why a particular measurement outcome occurs. This leaves plenty of room for interpretations, ranging from emphasizing the role of the observed in creating the measurement outcome, to considering the collapse as a physical process that happens all the time without any measurement, to rejecting the collapse completely and treating every branch of the superposition state as an alternate reality.

Narrowing down the list of interpretations holds the key to a better understanding of quantum mechanics, yet deciding between them is difficult in practice. Besides pure interpretations, there are in fact a few theories, which modify the mathematical structure of quantum mechanics. These theories make predictions that differ from quantum mechanics and are thus, in principle, subject to experimental tests with future experimental capabilities. The vast majority, however, are interpretations which supplement the mathematical formalism of quantum mechanics with a wide range of physical narratives. These interpretations make largely the same predictions as quantum mechanics, but as in the case of the luminiferous aether, this does not necessarily mean they are immune to experimental tests. Besides striving for a better understanding of our best physical theory, research into quantum foundations and interpretations of quantum mechanics has inspired the development of some revolutionary technologies. Quantum computing, for example originated in David Deutsch contemplating over the many worlds interpretation, and quantum cryptography is rooted in John Bell's study of local causal hidden variable theories for quantum mechanics. With applications such as these, foundational research has been pivotal for the development of quantum information theory, which in turn provided a new language and mathematical formalism to rigorously address some of the central questions in quantum foundations. Hence, apart from the desire for our best physical theory to provide more than just predictions, developing a deeper understanding for the theory is crucial for developing it further and harnessing its full potential.

What to expect

The core of this thesis are three experiments, studying the role of reality, causality, and uncertainty in quantum mechanics, which are presented in Chapters 4, 5, and 6, respectively. Each of these chapters is largely self-contained and readers with a strong background in quantum information, quantum foundations and quantum photonics may wish to skip directly to these

chapters. Readers less familiar with these fields will find introductions to the relevant technical concepts, notation, and mathematical background used throughout the thesis in Chapters 1-3. In Appendices A, B, and C three further experiments are included as original publications. These are touching upon the topics of finding physical principles behind quantum mechanics, quantum effects in extreme gravitational fields, and developing a pathway towards tests of macroscopic quantum phenomena, respectively, and are discussed briefly in Chapter 7

In Chapter 1 I introduce the main concepts and notion of quantum information theory, which forms the technical basis for the rest of the thesis. This includes discrete-variable quantum states, processes and measurements, as well as the experimental implementation of these using single photons and linear optics.

In Chapter 2 I discuss quantum tomography, which is a central tool for calibrating, testing and verifying quantum experiments. This chapter also includes the first new results of the present thesis, where we demonstrate a new tomography technique based on quantum superchannels. This technique enables a complete characterization of the evolution of a quantum system even when it is coupled to, and initially correlated with an environment, in which case standard methods fail.

In Chapter 3 I introduce a number of central concepts and ideas relevant in the study of quantum foundations. This includes quantum correlations and how correlation polytopes can be used to characterize them and find Bell-type inequalities. I also briefly touch upon the search for a physical principle that could explain why quantum correlations are not as strong as relativity allows. This is based on an experimental simulation of post-quantum correlations, which is included in full in Appendix A. The rest of the chapter focuses on contextuality—which is widely accepted as the central feature that differentiates quantum theory from classical theory with an epistemic restriction—various quantum paradoxes, and weak values.

Chapter 4 is devoted to the role of the quantum wavefunction within interpretations of quantum mechanics. Despite being the central object of the theory and of crucial importance for making predictions, it remains unclear whether the wavefunction corresponds to physical reality or is rather a representation of our incomplete knowledge of that reality. The core of this chapter then discusses an experiment, where we showed that the latter interpretation cannot fully reproduce quantum predictions.

Chapter 5 discusses the role of causality in quantum mechanics. I first introduce the causal modeling framework as a rigorous mathematical basis for this question. I then introduce a new causally motivated decomposition of the assumptions behind Bell's theorem, and an experiment, where we demonstrated that allowing for superluminal communication of measurement outcomes is not sufficient to explain quantum correlations in terms of classical cause-and-effect relations.

Chapter 6 focuses on a widely overlooked aspect of Heisenberg's uncertainty principle. Contrary to widespread understanding, this central principle not only limits how well two incompatible observables can be prepared on a quantum system, but also how accurately they can be jointly measured. I then present two experiments, where we test the optimal tradeoff in mea-

surement accuracy in a joint approximation of two incompatible measurements on a quantum systems.

Finally, Chapter 7 concludes with a summary of the results presented in this thesis, a discussion of the results included in the appendix, and directions for further research.

CHAPTER 1

Quantum Information Basics

QUANTUM information science views quantum mechanics as a theory that is fundamentally about information and information processing. The simplest quantum system, the qubit, serves as a fundamental quantum of information, and quantum mechanics is built from qubits. This approach thus captures the idea that classical mechanics should be a special case of a fundamentally quantum description, rather than starting from classical mechanics and quantizing it to get quantum mechanics. Working in this simple information theoretic framework has proven very successful for understanding the fundamental properties of quantum mechanics and for developing practical applications.

1.1 Quantum States

1.1.1 Qubits

In classical information science the fundamental unit of information is the *bit*, short for binary digit. A bit is represented by a physical system that has two distinct states, denoted 0 and 1, such as high and low levels of voltage in an electrical circuit, or head and tails of a coin. Quantum information theory is based on the very same premise. The fundamental unit of information is called a *qubit*, short for quantum bit, represented by a quantum system with two distinct states, $|0\rangle$ and $|1\rangle$. The prototypical example of a qubit is a spin-1/2 particle, such as an electron. Spin can be thought of as an intrinsic form of angular momentum, and when measured along some direction will always be found to either point up or down with respect to the measurement direction. The states $|0\rangle$ and $|1\rangle$ correspond to spin up and down, respectively, with respect to the z-direction. In contrast to the classical case, however, a quantum two-level system is not limited to these states, but can be in any state of the form

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle = \alpha \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \beta \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad (1.1)$$

where α and β are complex numbers, known as probability amplitudes, which satisfy $|\alpha|^2 + |\beta|^2 = 1$. The absolute square of the probability amplitude $|\alpha|^2$ is interpreted as the probability of obtaining outcome $|0\rangle$ in a measurement of $\{|0\rangle, |1\rangle\}$. All states of the form Eq. (1.1) lie on the so-called Bloch-sphere, see Fig. 1.1 and are called *pure* quantum states. The poles of the

sphere correspond to the classical bit values $|0\rangle$ and $|1\rangle$, or *computational basis states*, and all other states are *quantum superposition states*. The counter-intuitive nature of such states was famously illustrated by Schrödinger using the example of a cat that is at the same time dead ($|0\rangle$) and alive ($|1\rangle$).




Figure 1.1: **Classical and quantum bits.** (a) A classical bit can be in one of two distinct states, 0 and 1, or in a statistical mixture, represented by the black line. (b) A qubit can be in any of the pure quantum states on the surface of the Bloch-sphere, or in a statistical mixture of them, on the inside of the sphere. Here $|\pm\rangle = (|0\rangle \pm |1\rangle) / \sqrt{2}$ and $|\pm i\rangle = (|0\rangle \pm i|1\rangle) / \sqrt{2}$. Any antipodal pair of states is *orthogonal* and can be distinguished with certainty, just like the classical bit values 0 and 1.

More rigorously, a quantum pure-state $|\psi\rangle$ is a vector in a Hilbert space \mathcal{H} , which is a complex-valued vector space with an inner product $\langle \cdot | \cdot \rangle$. The absolute square of the inner product $|\langle \psi | \phi \rangle|^2$ quantifies the overlap of the states $|\psi\rangle$ and $|\phi\rangle$, and is related to the probability of distinguishing these states. This is one of the crucial differences to the classical case, where any two pure states are either equal or distinct, but not equal with some probability. In the quantum case, the only states that can be distinguished with certainty are *orthogonal states* with $\langle \psi | \phi \rangle = 0$. For a qubit these correspond to antipodal points on the Bloch sphere, such as $|0\rangle$ and $|1\rangle$. Any two nonorthogonal states $|\phi\rangle$ and $|\psi\rangle$, on the other hand, are overlapping and can be distinguished in a single measurement with a probability of at most $\frac{1}{2}(1 + \sqrt{1 - |\langle \psi | \phi \rangle|^2})$ [1]. Whether this is due to insufficient control over the state preparation at the level of the physical properties of the system, or rather due to a fundamental restriction on what can be measured will be discussed in detail in Chapter 4, (see also Heisenberg's uncertainty principle, Chapter 6).

In practice, even a classical bit is not perfectly 0 or 1, but rather mostly 0 with a bit of noise. In general, a bit can be a probabilistic mixture of 0 with probability p and 1 with probability $1 - p$. Think, for example of a bit determined by a coin flip. If the coin is fair, the bit will be completely random, $p = 1/2$, but if the coin is unfair one of the two is more likely and the bit contains some information. In Fig. 1.1 these states of the classical bit correspond to a line segment connecting $|0\rangle$ and $|1\rangle$. Similarly, quantum pure-states are an idealization, and do not exist in practice due to experimental noise. In analogy to the classical case *mixed quantum*

states are within the Bloch-sphere of Fig. 1.1. In order to describe such states mathematically, the concept of the state-vector is generalized to a density matrix ρ , which is represented by the Bloch-vector $\vec{n} = (n_x, n_y, n_z)$

$$\rho = \frac{1}{2} (\mathbb{1} + \vec{n} \cdot \vec{\sigma}) \quad \text{with} \quad \|\vec{n}\| \leq 1, \quad (1.2)$$

where $\vec{\sigma} = \{\sigma_x, \sigma_y, \sigma_z\}$ is the vector consisting of the Pauli-matrices, defined below, which are identified with the x, y, z-directions of the Bloch-sphere. The components of the vector \vec{n} could thus be identified with the average values of the spin along x, y, and z. Formally, the density matrix of a qubit is a 2×2 matrix which is positive semi-definite (i.e. it has no negative eigenvalues) Hermitian (i.e. $\rho^\dagger = \rho$) operator, and has unit trace (i.e. $\text{Tr}[\rho] = \sum_i \rho_{ii} = 1$) for every physical quantum state. Any such operator can be written as a linear combination of the elements of an appropriate operator basis, such as the Pauli basis, $\{\mathbb{1}, \sigma_x, \sigma_y, \sigma_z\}$, which gives rise to the Bloch-sphere representation of Eq. (1.2)

$$\mathbb{1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (1.3)$$

This basis is at the heart of the mathematical mapping of the two-dimensional complex-valued Hilbert space of a single qubit onto the three-dimensional real-valued space of the Bloch-sphere. As a visual consequence of this mapping, angles in the Bloch-sphere are double the ‘‘physical’’ angle between the states (defined by the inner product), such that, for example, the orthogonal states $|0\rangle$ and $|1\rangle$ with $\langle 0|1\rangle = 0$ are not at a right angle, but at antipodal points of the sphere. In the case of a pure state Eq. (1.2) reduces to $\rho = |\psi\rangle\langle\psi|$ and \vec{n} is a unit vector, on the surface of the sphere. This also shows that states defined in this way are invariant under multiplication with a global phase factor $e^{i\varphi}$, reducing the number of free parameters by one. The origin of this nomenclature is the fact that mixed states can be prepared as a statistical mixture of pure states $|\psi_i\rangle$

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i| \quad \text{with} \quad \sum_i p_i = 1, \quad (1.4)$$

where p_i is the probability for the state $|\psi_i\rangle$ in the mixture. Curiously, there are in general infinitely many ways to prepare a given density matrix as a mixture of pure states¹ as per Eq. (1.4). However, as long as two different ways of mixing pure states produces the same density matrix, the linearity of quantum mechanics ensures that there is no observable difference, see Sec. 3.3. To illustrate the difference between pure states and mixed states, consider the following two states

¹Classical probabilistic mixtures in contrast would typically allow for a unique decomposition into pure states.

$$\rho_+ = \frac{1}{2}(|0\rangle + |1\rangle)(\langle 0| + \langle 1|) = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \quad (1.5a)$$

$$\rho_{\mathbb{1}} = \frac{1}{2}(|0\rangle\langle 0| + |1\rangle\langle 1|) = \frac{1}{2} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}. \quad (1.5b)$$

Since these states have identical diagonal elements, they are indistinguishable for a measurement in the computational basis $\{|0\rangle, |1\rangle\}$. What distinguishes these states are the off-diagonal, or *coherence* terms in the density matrix. The state $\rho_{\mathbb{1}} = \frac{1}{2}\mathbb{1}$ is the so-called *maximally mixed state*, which is a 50:50-mixture of $|0\rangle$ and $|1\rangle$ —or in fact any pair of orthogonal states—conveys no information and lies at the centre of the Bloch sphere. The state ρ_+ , on the other hand, is a pure state on the surface of the sphere, which corresponds to a system that is, in some sense, in both states $|0\rangle$ and $|1\rangle$ at the same time. This can be quantified using the so-called *purity* $\mathcal{P} = \text{Tr}[\rho_\psi^2]$. The density matrix of a pure state is a so-called projection operator, for which $\rho_\psi^2 = \rho_\psi$ and thus $\text{Tr}[\rho^2] = 1$. In any other case $\mathcal{P} < 1$ and in the case of a maximally mixed state $\mathcal{P} = 1/2$, which means that measurement outcomes in any basis are completely random. In the case of qubits $\mathcal{P} = \frac{1}{2}(1 + \|\vec{n}\|^2)$, which conforms the intuition that the length of the Bloch-vector is a measure of mixedness of the state.

1.1.2 Beyond Qubits

The concepts introduced above are not restricted to qubits, but apply to any quantum system. A pure d -dimensional quantum state, a *qudit*, is represented by a d -dimensional complex vector with norm 1, which has $2(d-1)$ parameters (one of the d entries is fixed for normalization, and another due to an unobservable global phase). A mixed qudit density matrix has $d^2 - 1$ parameters. The gap in the dimension of these two sets of states implies that the geometric picture gets more complicated than a ball and its bounding sphere. There still exists a map from the set of mixed quantum states to $d^2 - 1$ dimensional vectors in a generalized Bloch ball, but this map is not bijective [2]. Not every point in the ball corresponds to a physical state. Furthermore, the pure state-space is not the full boundary of the ball, but already in the qutrit case “sprinkled” over the $(d^2 - 1)$ -sphere [3].

Continuous Variable Systems

The density matrix also describes so-called *continuous variable systems*, such as coherent states of light, which are becoming increasingly important for quantum information, see Ref. [4] for a comprehensive review. Since such states occupy an infinite dimensional Hilbert space, they cannot be visualized on a generalized Bloch-sphere. Instead, they can be associated with a distribution over a classical phase-space, the so-called *Wigner quasiprobability distribution* [5, 6]

$$W_\rho(q, p) = \frac{1}{\pi\hbar} \int_{-\infty}^{\infty} \langle q-x|\rho|q+x\rangle e^{2ipx/\hbar} dx. \quad (1.6)$$

Here q and p correspond to a pair of (conjugate) phase-space variables, such as position and momentum for a free particle, or amplitude and phase for a state of light. The Wigner-distribution is always normalized to 1 and for states that behave classical reduces to the classical phase-space probability density. The non-classical features of some quantum superposition states, such as Schrödinger’s cat-like states, on the other hand, lead to negative regions in the distribution, see Sec. 3.3.4 and can thus in general not be interpreted as a probability density. The Wigner distribution nonetheless gives a faithful representation of the quantum state and the expectation value of an operator $\hat{G}(q, p)$ in the state ρ can be recovered as a phase-space average of the operator’s Wigner transform $g(q, p)$ using Eq. (1.6) with $W_\rho(q, p)$ acting as the probability density [4]:

$$\langle \hat{G} \rangle = \text{Tr}[\rho \hat{G}] = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} W_\rho(q, p) g(q, p) dq dp .$$

Another distinct feature of the Wigner distribution is that it is bounded above and below by $\frac{1}{\pi\hbar}$. This implies in particular that the distribution cannot be infinitely sharp. In other words, the quantum system cannot have arbitrarily well-defined values for both conjugate properties at the same time—Heisenberg’s preparation uncertainty principle, see Chap. 6. It is worth noting that the Wigner distribution is just one of a whole family of quasi-probability distributions, which all have similar properties [4, 7]. Whether there is a way to avoid negative probabilities and interpret the quantum state as an object that represents limited knowledge, akin to a classical probability distribution in phase space is the topic of Chap. 4.

1.1.3 Comparing Quantum States

In practice, one is often interested in comparing two quantum states, for example to assess how close the experimental state is to the desired ideal state. The two most widely used measures for this purpose are the *trace distance* \mathcal{D} and the *fidelity* \mathcal{F} [8]. The fidelity is defined as [9]

$$\mathcal{F}(\rho, \sigma) := \text{Tr}[\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}]^2 = \|\sqrt{\rho}\sqrt{\sigma}\|_{\text{tr}}^2 , \quad (1.7)$$

where the trace-norm $\|\cdot\|_{\text{tr}}$ is the sum of singular values. Some authors define the fidelity without the final square, but the present definition has the advantage that it is physically more meaningful when generalized to quantum processes [10]. Although the fidelity is not quite a metric (i.e. a distance measure) on the quantum state space, it could easily be turned into one [8], but this is typically not necessary, as it already has most desirable properties for a measure of closeness of quantum states. It is symmetric, bounded to $0 \leq \mathcal{F} \leq 1$ and $\mathcal{F}(\rho, \sigma) = 1$ if and only if the states are equal. In the case where one or both of the states are pure, Eq. (1.7) simplifies to

$$\mathcal{F}(|\psi\rangle, \sigma) = |\langle \psi | \sigma | \psi \rangle| \quad (1.8a)$$

$$\mathcal{F}(|\psi\rangle, |\phi\rangle) = |\langle \psi | \phi \rangle|^2 . \quad (1.8b)$$

In this case the fidelity can thus be interpreted as the probability of finding the quantum system in state $|\psi\rangle$, when it was prepared in state σ (or $|\phi\rangle$).

The trace distance between two quantum states is defined as

$$\mathcal{D}(\rho, \sigma) = \frac{1}{2} \|\rho - \sigma\|_{tr}. \quad (1.9)$$

In contrast to the fidelity, the trace-distance is a proper metric on quantum states. For qubits it corresponds to half the euclidean distance between the states in the Bloch sphere. The trace distance also has an important operational interpretation in the sense that $1/2(1 + \mathcal{D}(\rho, \sigma))$ is the optimal probability of success for distinguishing the two quantum states with a single measurement [8]. Although they have somewhat different properties, fidelity and trace-distance are closely related via [8]

$$1 - \sqrt{\mathcal{F}(\rho, \sigma)} \leq \mathcal{D}(\rho, \sigma) \leq \sqrt{1 - \mathcal{F}(\rho, \sigma)}, \quad (1.10)$$

where the right-hand inequality is satisfied for pure states.

Recalling that in the case where both states are diagonal in the same basis, that is $\rho = \sum_i p_i |i\rangle\langle i|$ and $\sigma = \sum_i q_i |i\rangle\langle i|$, they resemble classical probability distribution. In this case, fidelity and trace-distance reduce to the corresponding classical fidelity \mathcal{F}_c and the variational distance δ , respectively for the probability distributions $P = \{p_i\}$ and $Q = \{q_i\}$ [8]

$$\mathcal{F}_c(P, Q) := \left(\sum_i \sqrt{p_i q_i} \right)^2, \quad (1.11a)$$

$$\delta(P, Q) := \frac{1}{2} \|P - Q\|_1 = \frac{1}{2} \sum_i |p_i - q_i|, \quad (1.11b)$$

1.1.4 Composite Systems and Entanglement

There is of course more to quantum mechanics than single, isolated systems. Classically, a composite system is completely specified by the properties of the individual constituents, and thus lives in the Cartesian product of the individual state spaces. In the quantum case, however, the state space of a composite system is the tensor product of the subsystem's Hilbert spaces [8]. In the case of two qubits, for example, the composite (four-dimensional) Hilbert space is given by $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$, where $\mathcal{H}_A, \mathcal{H}_B$ are the two-dimensional Hilbert spaces of system A and B , respectively, and \otimes denotes the tensor product. Just like in the case of classical systems, if system A is in state $|\psi\rangle_A$ and system B is in state $|\phi\rangle_B$, then the composite system is in the state $|\psi\rangle_A \otimes |\phi\rangle_B$. States of this form are called product states. Taking products of the basis states of the subsystem Hilbert spaces, one can construct a basis for the joint Hilbert space $\{|0\rangle_A, |1\rangle_A\} \otimes \{|0\rangle_B, |1\rangle_B\} = \{|00\rangle_{AB}, |01\rangle_{AB}, |10\rangle_{AB}, |11\rangle_{AB}\}$, where $|01\rangle_{AB}$ is a shorthand for $|0\rangle_A \otimes |1\rangle_B$. Every pure state in this space can be expressed as a linear combination of these basis elements²,

²Note that this space is in fact isomorphic to the state-space of a single ququart with basis $\{|0\rangle, |1\rangle, |2\rangle, |3\rangle\}$, see also Chap. 4. Certain quantum superposition states in the ququart space correspond to entangled states in the bipartite qubit space.

and every mixed state can be expressed as a linear combination of such pure states

$$|\psi\rangle_{AB} = \sum_{i,j} c_{ij} |i\rangle_A \otimes |j\rangle_B. \quad (1.12)$$

Classically, a composite system is completely specified by the properties (e.g. phase-space coordinates) of the individual constituents. Curiously, this does not hold anymore in the tensor-product structure of composite quantum systems. In other words, not every pure state in \mathcal{H}_{AB} (i.e. the states that can be written as Eq. (1.12)) can be expressed as a product of the states of system A and system B . The most prominent example are the four Bell-states:

$$|\Phi^+\rangle = (|00\rangle + |11\rangle)/\sqrt{2} \quad (1.13a)$$

$$|\Phi^-\rangle = (|00\rangle - |11\rangle)/\sqrt{2} \quad (1.13b)$$

$$|\Psi^+\rangle = (|01\rangle + |10\rangle)/\sqrt{2} \quad (1.13c)$$

$$|\Psi^-\rangle = (|01\rangle - |10\rangle)/\sqrt{2}. \quad (1.13d)$$

The states (1.13) are pairwise orthogonal and form a basis of \mathcal{H}_{AB} , such that every state can be decomposed in a linear combination of these states. However, the Bell states cannot be expressed as a product of states of system A and system B . On the contrary, the individual states of the subsystems are completely undefined (i.e. maximally mixed) [11]. More generally, any quantum state in \mathcal{H}_{AB} that can be expressed as a mixture of product states is called a *separable state*

$$\rho_{\text{sep}} = \sum_i p_i \rho_i^A \otimes \rho_i^B \quad \text{with } p_i \geq 0, \sum_i p_i = 1. \quad (1.14)$$

Any state that cannot be written in this form is called *entangled*³. Deciding whether a given density matrix is separable or not, however, is a computationally difficult problem in general (in fact, it is considered NP-hard even for bipartite systems) [13]. Fortunately, in the case of two qubits there is a necessary and sufficient condition given by the positive partial transpose, or Peres-Horodecki criterion. Loosely speaking, the joint state of two qubits is entangled if and only if the matrix obtained by transposing the part pertaining to one subsystem has negative eigenvalues [14]. Entanglement is considered one of the central resources for quantum information, as will be discussed in more detail in Sec. 3.2. As soon as one moves beyond the simple case of pairs of qubits the situation gets very complicated with, for example, multiple qualitatively different forms of entanglement in the case of multipartite systems, see Ref. [14] for more details.

Reduced States

In practice, one is often interested only in a subsystem of a larger composite system. For example, the system of interest might interact with, or be correlated with the environment, which is inaccessible to the experimenter. When performing a measurement on the system,

³Translated from the German term “Verschränkung”, coined by Erwin Schrödinger in 1935 [12]

everything else, including potential correlations, is simply ignored or, in quantum information jargon, *traced over*. To make this more precise, recall that the trace of a matrix is the sum of its diagonal elements, $\text{Tr}[\rho] = \sum_i \rho_{ii}$, which, in the case of a density matrix, corresponds to measuring the system in the computational basis and ignoring the result. In the case of a composite system, one can choose to trace out (i.e. ignore) only specific subsystems [8, 15]. For the so-called *reduced state* of the subsystem A of a bipartite state ρ^{AB} is obtained by applying the *partial trace* to subsystem B

$$\rho^{\text{A}} = \text{Tr}_{\text{B}}[\rho^{\text{AB}}] = \sum_{j=l} (\rho^{\text{AB}})_{ij,kl} |i\rangle\langle k|, \quad (1.15)$$

where ρ^{AB} is written in terms of tensor indices ik for subsystem A and jl for subsystem B . In Sec. 1.2 below, an intuitive tensor-network representation of the partial trace is introduced (c.f. See Fig. 1.4). Unless the two subsystems are in a product state, the partial trace will result in a reduced state that is mixed. The more entangled the initial state, the more mixed the reduced state. In the case of maximally entangled Bell states, the joint properties of the two systems are sharply defined, but the individual properties are completely undefined, and the reduced states are thus completely mixed.

Although preparing mixed states by tracing over part of a pure entangled states seems qualitatively very different from a statistical mixture of pure states, as long as the two methods produce the same density matrix, there is no measurement that can distinguish them. This is also known as the problem of proper (statistical) mixture versus improper (tracing-out) mixture and is a consequence of the linearity of quantum mechanics. Any non-linear evolution beyond standard quantum mechanics could in principle distinguish between the two kinds of mixed states [16, 17], see Appendix B.

Quantifying Non-Classical Correlations

When it comes to quantifying non-classical correlations, the simplest case (two qubits) is well understood, but everything else not so much. A comprehensive review of the concept of entanglement and how to quantify it can be found in Ref. [14]. There are in fact two commonly considered problems: deciding whether a given state is entangled, and quantifying the amount of entanglement.

Since the set of separable quantum states is by definition convex, the Hahn-Banach theorem guarantees that for every entangled state ρ there exists a hyperplane that separates it from the set of separable states, see Fig. 1.2. The Hermitian operator associated with such a hyperplane is called *entanglement witness*, as it has a positive expectation value for all separable states, but a negative value for some entangled states. A witness is called optimal, if the plane is tangent to the space of separable states [14]. Bell inequalities are examples of suboptimal entanglement witnesses, detecting only Bell-nonlocal states. Optimal entanglement witnesses can be further strengthened by taking into account additional information about the quantum state under investigation, which is already available in almost every standard entanglement

witnessing experiment [18].




Figure 1.2: **Entanglement witness.** An entanglement witness W is associated with a hyperplane such that the set of separable states are on one side. Every state in the red shaded region to the right of the witness, such as ρ is correctly identified as an entangled state. However, it is in general not possible to guarantee that a state is separable, since the boundary of the set of separable states is not a simple hyperplane.

This geometric approach of entanglement witnesses can also be generalized to multipartite states [19]. However, the study of entanglement gets significantly more complicated the more parties are involved. In the case of three qubits, for example, there is entanglement of the GHZ-type $|\psi\rangle_{\text{GHZ}} = (|000\rangle + |111\rangle)/\sqrt{2}$ and of the W-type $|\psi\rangle_{\text{W}} = (|001\rangle + |010\rangle + |100\rangle)/\sqrt{3}$ [20]. When tracing out any of the three parties of a GHZ state, the other two are left in a maximally mixed state, while they are left in a maximally entangled state if the initial state was a W-state. Hence, classifying multipartite entangled states requires careful definitions of separability [14]. One could, for example, consider *fully separable* states, in which case states with only bipartite entanglement have the same status as genuinely multipartite entangled states, such as the GHZ-state. On the other hand, measures of partial separability can miss interesting partially entangled states.

Entanglement witnesses are attractive because they do not require complete knowledge of the quantum states, which is experimentally challenging. However, different entangled states require different witnesses to be verified. If, in contrast, the full density matrix of the state is known (e.g. by performing quantum state tomography, see Chap. 2), one can derive general purpose, quantitative entanglement measures. Such a measure should vanish for all separable states and should not increase under local operations and classical communication. Whether a measure is non-zero for all non-separable states depends on the chosen definition of separability [14].

Concurrence

Concurrence \mathcal{C} is a measure of bipartite entanglement first introduced in 1997 [21]. This measure is of particular practical importance since there exists a closed-form expression for its convex roof extension that generalizes it to mixed states [22].

$$C = \max\{0, \lambda_1 - \lambda_2 - \lambda_3 - \lambda_4\}, \quad (1.16)$$

where λ_i are the singular values of $\sqrt{\sqrt{\rho}(\sigma_Y \otimes \sigma_Y)\rho^*(\sigma_Y \otimes \sigma_Y)\sqrt{\rho}}$ in decreasing order. The concurrence can thus be applied to mixed states, as well as pure states. Another often used measure of entanglement is the *tangle* [23], which for a bipartite state is simply the square of the concurrence, but can be generalized to multi-partite states.

Residual Tangle

The residual tangle generalizes the tangle to a measure of *genuine* 3-partite entanglement, i.e. entanglement such that no partition of the system is separable [14, 23].

$$\tau_3(A:B:C) := \tau(A:BC) - \tau(AB) - \tau(AC) . \quad (1.17)$$

For mixed states, it can then be defined via the convex roof extension

$$\tau_3(\rho) = \min \sum_i p_i \tau_3(\pi_i) , \quad (1.18)$$

where the minimum is taken over all pure-state decompositions of $\rho = \sum_i p_i \pi_i$. Since τ_3 is a convex function on the set of density matrices, the optimum can in principle be found easily [24]. The residual tangle has potential to be generalized to an arbitrary number of parties using hyperdeterminants [25].

1.2 Quantum Channels

A *quantum channel* (or process, evolution, dynamics, transformation) describes how the quantum system, and in particular the encoded quantum information, is transformed between state preparation and measurement. One of the central postulates of quantum mechanics is that information can neither be lost, nor created. This implies that the evolution of an isolated quantum system must be described by a unitary operator U (recall unitary means $U^\dagger U = \mathcal{I}$), which corresponds to a generalized rotation, see Fig. 1.3a.

Information only needs to be preserved for isolated systems, but no quantum system is perfectly isolated in practice, and information encoded in fragile quantum states tends to get lost into the environment over time, see also Chap. 2. Hence, although the evolution of system and environment together might be unitary, the *reduced dynamics* of the system alone can in general not be described by a simple rotation. Instead, the quantum state will in general become mixed in the course of the evolution, see Fig. 1.3b.

Formally, a general quantum channel is described by a *completely-positive trace-preserving* (CPTP) map $\mathcal{E}: L(\mathcal{X}) \rightarrow L(\mathcal{Y})$, which maps operators (such as density matrices) on the Hilbert space \mathcal{X} onto operators on the potentially different Hilbert space \mathcal{Y} . In the following we will restrict our attention to the common case $\mathcal{X} = \mathcal{Y}$, but the results can easily be adapted to the general case. The map \mathcal{E} is positive if it maps positive operators onto positive operators, and completely positive (CP) if the same is true for the composite map $\mathcal{E} \otimes \mathcal{I}$, where \mathcal{I} is the identity map on a space at least as large as \mathcal{X} . Completely-positive maps are the most general

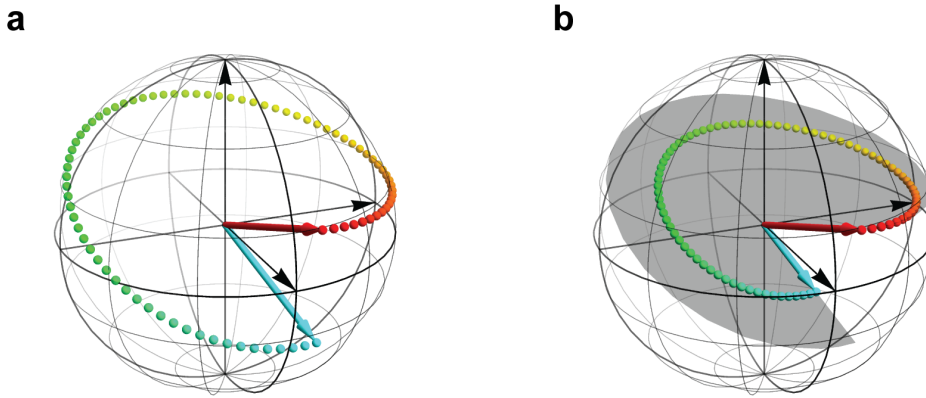


Figure 1.3: **The action of a quantum channel in the Bloch sphere.** (a) The initial state, shown in red, evolves unitarily to the final state, shown in cyan. The intermediate states are coloured accordingly to illustrate the evolution, which occurs completely on the surface of the sphere. (b) A noisy version of the evolution in (a), which is non-unitary and decreases the purity of the state compared to the unitary evolution shown in grey.

transformations that turn a physical quantum state into another physical quantum state [26]. The map is called trace-preserving (TP) if it does not change the trace of an operator.

The requirement that \mathcal{E} is trace-preserving is typically not satisfied in practice due to loss. However, if this loss is unbiased, it is often possible to renormalize the output state (or post-select on those runs of the experiment in which the system made it through) and treat the map as trace-preserving. A more serious problem in practice is that correlations with the environment might make the transformation look non-CP and thus unphysical [27]. There are, however, sophisticated tools to deal with such situations, see Chap. 2, and in the following all quantum channels are assumed to be CPTP.

1.2.1 Working with Quantum Channels

In order to work with a quantum channel it has to be represented in a way that makes clear how it transforms a general quantum state. There are a number of different representations used for different purposes and by different communities, but are all closely related and can be transformed into one another, see Ref. [28] for an accessible discussion. All of these are conveniently represented using a pictorial tensor-network notation [28], see Fig. 1.4.

Since isolated quantum systems evolve unitarily, it should be possible to describe any evolution as unitary by making the system large enough (the “church of the larger Hilbert space”). This is made precise by the *Stinespring dilation theorem* [29], which states that any non-unitary dynamics of the system can be described as a unitary evolution U_{SE} of system S and environment E together, and then ignoring the environment, see Fig. 1.5

$$\mathcal{E}(\rho) = \text{Tr}_E[U_{SE}(\rho \otimes \rho_E)U_{SE}^\dagger]. \quad (1.19)$$

Curiously, it is not actually necessary to include the whole environment, but the effect of an




Figure 1.4: **Basics of tensor network notation.** (a) A state vector $|\Psi\rangle$ (a row-vector) is represented by a triangle with a wire to the right. (b) The conjugate vector $\langle\Psi|$ (a column-vector) is represented by a triangle with a wire to the left. (c) A straight wire corresponds to an identity matrix. (d) A box with two wires corresponds to a matrix (rank-2) tensor, such as a quantum channel or a density matrix. (e) Matrix multiplication (tensor contraction) is represented by connecting the “wires” corresponding to the indices to be summed over. (f) Tensor products are represented by vertical composition. The rank of the resulting tensor is the sum of the tensor ranks. (g) The trace operation for square matrices is represented by contracting the two matrix indices with each other. (h) The partial trace operation for composite square matrices is represented by contracting the two matrix indices of the corresponding subsystems to be traced over. (i) The vectorized identity operator is equivalent to the unnormalized bell state $\sum_i |i\rangle\otimes|i\rangle$. (j) Column vectorization of a matrix is represented by “bending” the right wire upwards to form a vector. “Sliding” a matrix around a bell-state results in the transposition of the matrix.

arbitrarily large environment on a d -dimensional system can always be simulated using an ancillary system of dimension d^2 [30].

Another way to interpret a non-unitary channel is to think of it as a unitary channel with added noise. This motivates the *operator-sum* representation, where the quantum channel \mathcal{E} is decomposed into a sum of linear operators K_i , the so-called *Kraus operators* (or sometimes measurement operators), see Fig. 1.5

$$\mathcal{E}(\rho) = \sum_{i=1}^{d^2} K_i \rho K_i^\dagger \quad \text{where} \quad \sum_{i=1}^{d^2} K_i^\dagger K_i = \mathbf{1} . \quad (1.20)$$

If \mathcal{E} is unitary there is only a single Kraus operator. In general, any completely-positive quantum channel can be decomposed into at most d^2 Kraus operators [8] and any two such decompositions are unitarily equivalent. In the context of quantum error correction the operators K_i are also known as *noise operators*.

1.2.2 Process Matrix Representation

The *Choi-Jamiolkowski isomorphism* establishes a connection between linear maps and linear operators, and thus in particular between quantum channels and quantum states [31, 32]. In the most basic form it refers to what is known as vectorization [28], where an operator U is

turned into a vector $|U\rangle\rangle$,

$$U = \sum_{i,j=0}^{d-1} U_{ij} |i\rangle\langle j| \mapsto |U\rangle\rangle = \sum_{i,j=0}^{d-1} U_{ij} |j\rangle \otimes |i\rangle. \quad (1.21)$$

Here $\{|i\rangle\}_{i=0}^{d-1}$ denotes the computational basis in d dimensions. The set of operators $\{|i\rangle\langle j|\}$ is called the *elementary basis* of the space of operators, since it consists of all $d \times d$ matrices with a single element that is 1, and otherwise zeros. The order $|j\rangle \otimes |i\rangle$ in Eq. (1.21) is a matter of convention, such that the resulting vector consists of the stacked columns of U . The corresponding density matrix is $|U\rangle\rangle\langle\langle U| = \sum_{i,j=0}^{d-1} |i\rangle\langle j| \otimes U|i\rangle\langle j|U^\dagger$. Following this argument, a CPTP map \mathcal{E} is associated with a *Choi-matrix* Λ which is the unnormalized quantum state (with trace d rather than trace 1) obtained from vectorizing \mathcal{E} [33].

$$\Lambda = \sum_{i,j=0}^{d-1} |i\rangle\langle j| \otimes \mathcal{E}(|i\rangle\langle j|). \quad (1.22)$$

This matrix completely describes the evolution of an arbitrary quantum state ρ under the channel \mathcal{E}

$$\mathcal{E}(\rho) = \text{Tr}_1[(\rho^T \otimes \mathbf{1})\Lambda]. \quad (1.23)$$

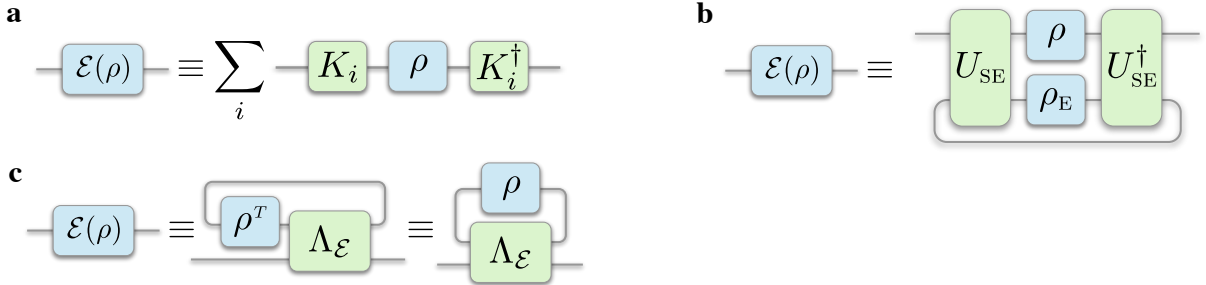


Figure 1.5: **Different representations of a quantum channel.** (a) In the system-environment representation the evolution of the system is interpreted as a joint unitary evolution of system and environment, followed by tracing over the environment. (b) In the operator-sum representation a quantum channel is decomposed into a sum of Kraus operators K_i . (c) In the process matrix representation the channel is associated with a Choi-matrix Λ , which is applied directly to the state ρ .

The Choi matrix is unique for a given basis, but not every basis is equally useful. In particular, the elementary basis used in the definition (1.22) does not correspond to a set of physical quantum states. In practice, the Pauli basis $\{\mathbf{1}, \sigma_x, \sigma_y, \sigma_z\}$ or a tensor product in the case of multiple qubits is most useful, since the matrix elements have a clear interpretation as Pauli operators acting on the state. The Choi matrix written in the Pauli basis is also known as *process matrix*, $\chi = T\Lambda T^\dagger$, where T is a basis-change matrix from elementary to Pauli basis.

The evolution of the state ρ is then given by

$$\mathcal{E}(\rho) = \text{Tr}_1[(\rho \otimes \mathbb{1})(T^\dagger \chi T)] = \sum_{i,j=0}^{d^2-1} \chi_{ij} \sigma_i \rho \sigma_j^\dagger. \quad (1.24)$$

1.3 Quantum Measurements

Recall that $|\langle \phi | \psi \rangle|^2$ is associated with the probability of finding the system in the state $|\phi\rangle$ after it was prepared in the state $|\psi\rangle$. In other words, $|\phi\rangle$ is interpreted as part of a so-called projective or *sharp* measurement on $|\psi\rangle$. In this case, the measurement outcomes correspond to pure states (i.e. unit vectors) on the Bloch-sphere and form an orthonormal basis of the Hilbert space, which ensures that the measurement outcome probabilities sum to 1. Consider, for example, a measurement of $\{|0\rangle, |1\rangle\}$ corresponding to Bloch vectors $(0, 0, \pm 1)$. Assigning the values $a_0 = +1$ and $a_1 = -1$ to the two outcomes, one can define the *observable* $\hat{O} = a_0|0\rangle\langle 0| + a_1|1\rangle\langle 1| = \sigma_z$, associated with the physical property spin in z-direction. The average value of the spin in z-direction is then given by the *expectation value* of \hat{O} on the state ψ

$$\langle \hat{O} \rangle_\psi = \langle \psi | \hat{O} | \psi \rangle = +1|\langle 0 | \psi \rangle|^2 - 1|\langle 1 | \psi \rangle|^2 = P(0|\psi) - P(1|\psi), \quad (1.25)$$

where $P(0/1|\psi)$ is the probability for observing outcome $|0\rangle$ and $|1\rangle$, respectively. In contrast to a classical bit, one measurement is insufficient to fully determine an arbitrary quantum state, see Chap. 2. Just like quantum states are not limited to pure states, measurements are not limited to sharp measurements. Rather, general measurements could be unsharp, represented by non-unit vectors, and even biased towards one outcome.




Figure 1.6: **Representation of measurements on the Bloch sphere.** (a) A sharp measurement on a qubit corresponds to two anti-parallel unit vectors (or orthogonal pure states). A three-outcome qubit POVM, on the other hand, could be composed of (b) noisy projectors, or (c) sub-normalized projectors.

A general quantum measurement is represented by a positive-operator valued measure (POVM) \mathbb{M} , which is a set of Hermitian, positive semi-definite operators $\{E_i\}_1^n$ that sum to the identity ($\sum_{i=1}^n E_i = \mathbb{1}$). The operators $\{E_i\}_{i=1}^n$ are called *POVM elements* (or effects) and correspond to measurement outcomes. In the special case where all E_i are projective (i.e. $E_i = |\phi_i\rangle\langle \phi_i|$), the measurement is called *sharp* and the set $\{|\phi_i\rangle\}$ is an orthonormal basis of

the Hilbert space, thus in particular $n = d$, where d is the Hilbert space dimension. Just like pure states do not exist in practice, the same is true for sharp measurements, and a general measurement is *unsharp*. The probability for the k 'th measurement outcome E_k on the state ρ is given by the *Born rule* [34]

$$P(k | \rho) = \text{Tr}[\rho E_k]. \quad (1.26)$$

The Born rule is one of the key postulates of quantum mechanics crucial for its inherently probabilistic predictions. There is still an ongoing debate whether or not it can be derived from the other postulates of the theory [35, 36]. Interestingly, however, once one accepts that measurement outcomes should be probabilistic, Gleason's theorem [37] gives strong evidence that the Born rule is the unique consistent probability rule that depends only on the quantum state [38].




Figure 1.7: **POVMs in practice.** (a) Using Naimark's theorem every POVM can be implemented by interacting the system with a suitable ancilla and performing a joint projective measurement on system and ancilla. (b) A POVM that is a noisy projective measurement can be implemented as a convex combination of two sharp measurements. (c) A POVM consisting of subnormalized projectors can be implemented using sharp measurements and post-processing.

Notably, the measurement probabilities defined in this way only depend on the measurement operator E_k , and not on the full POVM or any details of how the measurement is implemented⁴. The post-measurement state ρ'_k of the system, on the other hand, does depend on the physical realization of the measurement and cannot be determined from the particular POVM element E_k alone. To this end one needs to specify a (non-unique) *instrument*, which is a set of completely-positive maps $\{\phi_i\}$, where ϕ_k corresponds to the k 'th measurement outcome and $\sum_i \phi_i$ is trace-preserving [39]. The post-measurement state is then given by

$$\rho'_k = \frac{\phi_k(\rho)}{\text{Tr}[\phi_k(\rho)]}. \quad (1.27)$$

Just like mixed quantum states can be considered a subsystem of a larger pure state, POVMs can be interpreted as the effective description of the measurement on a subsystem arising from a projective measurement on a larger system. Formally, Naimark's dilation theorem [40] shows

⁴This feature is reminiscent of the observation that quantum mechanics is impartial about how a certain quantum state was prepared (as a mixture of pure states) or how a certain transformation is implemented (as a mixture of unitary transformations), see Sec. 3.3 for more details.

that a general POVM can be modelled as a projective measurement on system+ancilla on a sufficiently extended Hilbert space. Ignoring the ancilla leads to a non-projective POVM on the system. In practice, implementing POVMs in this way can be quite impractical and, alternatively, one can implement a POVM as a combination of projective measurements and post-processing, see e.g. Fig. 1.7.

In the special case of a single qubit, the most general two-outcome measurement is given by the POVM $\mathbb{M} = \{M_+, M_-\}$. The POVM elements, associated with the outcomes ± 1 , can be written as $M_{\pm} = \frac{1}{2}(\mathbb{1} \pm (\mu\mathbb{1} + \vec{m}\vec{\sigma}))$, with \vec{m} a vector on the Bloch sphere, and $\mu \in \mathbb{R}$ representing the measurement bias, such that $|\mu| + \|\vec{m}\| \leq 1$. One can also define the corresponding observable by $\mathcal{M} = M_+ - M_- = \mu\mathbb{1} + \vec{m}\vec{\sigma}$. An unbiased measurement can be represented as a vector on the Bloch-sphere, and the expectation value of the observable \mathcal{M} on the state $\rho = (\mathbb{1} + \vec{n}\vec{\sigma})/2$ is in general given by

$$\langle \mathcal{M} \rangle_{\rho} = \text{Tr}[\mathcal{M}\rho] = \mu + \vec{m} \cdot \vec{n} , \quad (1.28)$$

where $\vec{m} \cdot \vec{n}$ is the standard scalar product between the real-valued Bloch-vectors \vec{m} and \vec{n} .

1.3.1 Inconsistent Prescriptions

“WHY must I treat the measuring device classically? What will happen to me if I don’t?” - E. Wigner [41]

Consider a system in the initial state $|\psi\rangle = \sum_i c_i |\phi_i\rangle$ with respect to some basis $\{|\phi_i\rangle\}$. When subject to a measurement in the basis $\{|\phi_i\rangle\}$, the prescription outlined above implies that the result k will be obtained with probability $|c_k|^2$ and the system will be left in the post-measurement state $|\phi_k\rangle$, see Fig. 1.8a. The two crucial aspects of this *measurement postulate* are that measurement outcomes occur probabilistically, and the state of the system changes discontinuously (or *collapses*) to the post-measurement state. This, however, is in stark contrast to the usual deterministic and continuous evolution of quantum systems prescribed by the Schrödinger equation.

A “measurement” receives a special (classical) treatment in this prescription, but there is nothing within quantum mechanics that determines whether a given operation should be treated this way or not. Quite the contrary, quantum mechanics applies equally well to the measurement device as it does to the measured system. Denoting the state of the measurement device before the experiment by $|ready\rangle$, a measurement in the basis $\{|\phi_i\rangle\}$ simply corresponds to an interaction between the measurement device and the measured system. After this interaction, system and measurement device are in an entangled state $\sum_i c_i |i\rangle \otimes |\phi_i\rangle$, see Fig. 1.8b.

The inconsistency between these two descriptions of a quantum measurement is often referred to as the *measurement problem* and is one of the biggest open problems in quantum foundations. Finding a prescription for what counts as a measurement and what does not, as well as an explanation for why a specific outcome occurred and not any of the other possibilities are central motivations behind the development of interpretations of quantum mechanics. Suggestions range from “there is no collapse” to “collapse happens all the time even without

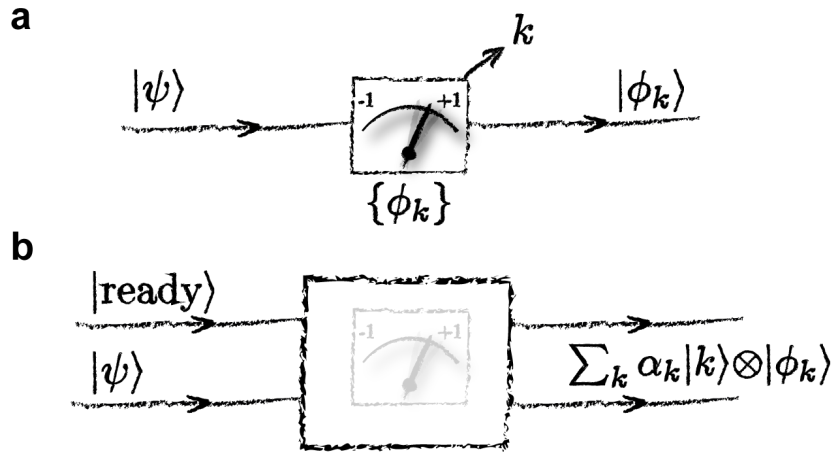


Figure 1.8: **Two inconsistent prescriptions for quantum measurements.** (a) Treating the measurement device as an external device, the projection postulate asserts that, upon measurement of $\{\phi_i\}$, the system randomly collapses to one of the states ϕ_k with probability $P(k) = |\langle \phi_k | \psi \rangle|^2$, producing the measurement outcome k . (b) Alternatively the measurement device could be treated as a quantum system and the measurement process as a physical interaction within quantum mechanics. In this case, the Schrödinger equation implies that the system and the measurement device in general end up in an entangled state.

measurement” to “collapse is just an illusion”, see Chapter 4 for more details.

1.3.2 Weak Measurements

In practice, measurements in quantum optics are typically destructive in the sense that the photon, the carrier of the quantum information, is absorbed by a detector during the measurement process. One way to avoid this is by implementing the measurement in the way illustrated in Fig. 1.8b, known as a *von Neumann measurement*. Here the system is coupled to the meter in such a way that a subsequent measurement of the meter effectively implements a measurement of the desired observable on the system. Such a measurement scheme thus makes it possible to extract information from the system without destroying the physical system. However, the quantum state of the system is nevertheless disturbed, in accordance with the uncertainty principle in the same way as if the measurement was performed directly on the system. In the language above, the quantum state of the system is collapsed to one of the measurement outcomes.

One of the advantages of the von Neumann measurement scheme is that it naturally allows for weak measurements by reducing the interaction strength⁵. The measurement strength can be quantified by a parameter $0 \leq \kappa \leq 1$, where $\kappa = 1$ just corresponds to a standard strong measurement and $\kappa = 0$ means no measurement at all. A weak measurement is everything between these two extremes. In the limit where κ is very small, the disturbance of the system

⁵Traditionally the strength of a von Neumann measurement is determined by the coupling strength g , the interaction time t , and the initial uncertainty in the meter Δ . A measurement is then considered “weak” if $gt \ll \Delta$, see e.g. Ref. [42]

state is greatly reduced—it is on average only slightly biased, rather than projected—but at the same time, the measurement encodes only very little information about the state. However, from a large enough ensemble of identically prepared systems, it is nevertheless possible to infer an *average value* of the measured quantity with high precision. This is quite different from a noisy measurement, which extracts little information while still disturbing the system significantly.

In practice, weak measurements are naturally encountered whenever the measured system has to be preserved. A typical example is gaining information about the behaviour of a system between two strong measurements, which is described in terms of weak values. Such situations can lead to statistically anomalous behaviour, so-called *anomalous weak values*, which has led to some controversy around the concept of weak values, see Sec. 3.5. An example of the use of weak measurements is the reconstruction of the average photon trajectories through a double-slit [43]. The weak measurement here simply produces a (post-selected) conditional average (the average trajectory conditioned on the arrival at a specific point on the screen), which is nothing particularly spooky and agrees well with classical intuition [44]. See Ref. [45] for a review of applications of the concept of weak values and Ref. [46] for the “contextual values” approach to unifying the continuous spectrum of measurement strengths.

1.3.3 Quantum Non-Demolition Measurements

As a consequence of Heisenberg’s uncertainty principle, any quantum measurement that reveals information about the system, must also disturb it, see Chap. 6. A measurement of the position of a particle, for example, typically disturbs the particle’s momentum. In the subsequent evolution this disturbance can couple back to the position, such that a second position measurement might give an inconsistent result.

A quantum non-demolition (QND) (or quantum back-action free) measurement is designed to avoid exactly this by confining the disturbance to the conjugate observables [47, 48]. Importantly, a QND measurement is *not* a novel kind of measurement that avoids disturbance [49], which would violate the uncertainty principle. Rather, it is a measurement that is cleverly engineered such that there is no disturbance on the measured observable and subsequent measurements give the same result. This is achieved by designing a situation where the measured observable is a constant of motion of the free evolution of the system (i.e. commutes with the free Hamiltonian). The momentum of a free particle (in contrast to its position) is an example of such a conserved property. A QND measurement can thus be used to stabilize the measured property, such as photon number in a cavity [50], at the cost of randomizing the conjugate observables. QND measurements should not be confused with non-destructive quantum measurements, which only aim to preserve the physical system, not the quantum state.

1.4 Quantum Information in Practice

There are a large number of quantum systems out there and all have their advantages and disadvantages. These include, for example, spins embedded in pure diamond, superconducting circuits, which are interesting for their similarities to classical electrical circuits, and ions, which have been setting records in precision and number of entangled qubits. A rather unique candidate among all these platforms are single photons.

1.4.1 Single Photons

Whereas most systems, such as ions in a trap or superconducting circuits, are stationary, photons are inherently mobile and in fact quite tricky to hang on to. This makes them a particularly interesting candidate for quantum communication. Moreover, while most architectures struggle to preserve their quantum states for seconds, qubits encoded in the polarization of photons have a *coherence time* beyond measurable⁶. This is a result of the fact that photons virtually do not interact with their environment, which is boon and bane for the architecture. No interaction means that they are easy to work with—without the need for cryogenics and other isolation—but at the same time two-photon interactions become a major challenge. The breakthrough for single photons as a competitive quantum information platform was probably triggered by a seminal paper by Knill, Laflamme, and Milburn [51], showing that measurement induced non-linearity is sufficient to implement the crucial two-photon gates.

Photons are also one of the most flexible platforms with a large number of accessible and controllable degrees of freedom that can be used to encode quantum information and prepare quantum systems of arbitrary finite dimension. The most prominent degree of freedom is polarization, which is equivalent to the spin of a spin-1/2 particle and thus gives the photon a natural qubit structure. Besides polarization, there are other discrete degrees of freedom, such as photon number and transverse spatial mode. On the other hand, there are continuous degrees of freedom, such as frequency (i.e. energy), time-of-arrival or path. These latter degrees of freedom can be binned into discrete sets that can act as a basis for a discrete-variable system.

Polarization

Polarization is the oscillation direction of light as an electromagnetic wave, which naturally parallels the behaviour of a spin-1/2 particle. Specifically, one can define $|0\rangle = |H\rangle$ and $|1\rangle = |V\rangle$, where H and V denote horizontal and vertical polarization, respectively, with respect to some physical reference. Equal superpositions of H and V with phases of 0 and π , respectively define diagonal and anti-diagonal states $|D/A\rangle = (|H\rangle \pm |V\rangle)/\sqrt{2}$, which are the eigenstates of σ_x . Hence, the xz-plane of the Bloch-sphere is identified with linear polarizations. Choosing any other phase-value will in general lead to elliptical polarization with the extreme cases (with phases $\pm\pi/2$) of right- and left- circular polarization $|R/L\rangle = (|H\rangle \pm i|V\rangle)/\sqrt{2}$, the eigenstates of

⁶If the upgraded BICEP detector finds convincing evidence for polarization of the cosmic microwave background that would suggest a coherence time on the order of billions of years.

σ_y . Any polarization state can be decomposed into two orthogonal components, and a rotation from one polarization state to another thus amounts to a phase-shift between the appropriate components.

Path Encoding

When encoding a qubit in path, the computational basis states ($|0\rangle$ and $|1\rangle$ for a qubit) are identified with distinct spatial paths for a single photon. In all other states the photon is in a superposition of these paths with specific amplitudes and phases. In practice, this means that the phase between the two modes has to be stabilized to a fraction of a wavelength, which can be very challenging with bulk optics. Advances in waveguide technology, however, make it now possible to integrate path-encoded circuits into optical chips, which significantly improves phase-stability and reduces their footprint. Current devices still suffer from high losses, however, due to rapid development and electronically controllable phase-shifters, this architecture now achieves comparable quality to bulk optical implementations [52]. When it comes to simultaneously encoding information in multiple degrees of freedom, however, this architecture is more limited, since the waveguide material is typically birefringent and does not support higher-order spatial modes. Nonetheless, they are the most promising candidates for creating high-dimensional quantum systems using single photons.

Transverse Spatial Mode

Using non-Gaussian transverse mode structures is another way to encode a high-dimensional Hilbert space on a single photon. The angular momentum of a photon is composed of a spinorial and an orbital component, neither of which are angular momenta by themselves [53, 54]. As a spin-1 particle, a photon can carry $\pm\hbar$ of spinorial angular momentum, which corresponds to left- and right-circular polarization states (or helicity eigenstates) [54]. Orbital angular momentum, on the other hand, is related to the transverse structure of the field, associated with a helical phase factor $e^{im\varphi}$ for azimuthal angle φ [55]. In general, spin- and orbital- angular momentum are coupled, but at least in the paraxial approximation we can manipulate them independently. While the former couples to optical anisotropies in the material (i.e. birefringence), the latter component couples to inhomogeneities in the form of rotational asymmetries [55].

In particular, orbital angular momentum gives rise to optical vortex beams, which have a zero-intensity “hole” at the centre and have been used for a long time [53]. They offer, for example, a very convenient way to introduce angular momentum to other quantum systems like a Bose-Einstein condensate, or even put these systems in a superposition of rotational states [56]. For quantum information purposes such states offer a (arbitrarily) high-dimensional discrete Hilbert space. Any mode structure can be written as a discrete sum of so-called Ince-Gauss modes [57], which can be obtained using devices such as spatial light modulators. Moreover, states entangled in their transverse mode structure are generated directly from spontaneous parametric downconversion. Although it is technically challenging, arbitrary transformations of these states within any finite dimensional Hilbert space are possible with cylindrical lenses,

inverting prisms, or spatial light modulators.

1.4.2 Manipulating Polarization Qubits

Any operation on a set of qubits can be decomposed into single-qubit rotations and two-qubit entangling operations. In the case of qubits encoded in the polarization of single photons, the former are easy, but the latter are causing some headache.

Single-Qubit Operations

Recall that any two polarization states are related by a phase-shift between two suitably chosen polarization components, which is achieved using so-called *waveplates*. These are made from birefringent material which features different refractive indices (i.e. propagation speeds) for different polarizations, thus delaying one with respect to the other. Birefringence occurs naturally due to anisotropic molecular order in the material, or it can be induced through stress (stress birefringence), electric fields (Kerr effect) or magnetic fields (Faraday effect). The most common birefringent materials are so-called *uniaxial* crystals that have one distinguished crystal axis—the *optic axis*—which has a refractive index that is different from the other two axes. There are also bi-axial materials (the most general case), which are much more complex and have applications elsewhere.

A waveplate is typically made from a uniaxial material, oriented with the optic axis perpendicular to the propagation direction of incoming light. Light is split into the *ordinary* component polarized orthogonal to the optic axis, and the *extraordinary* component polarized parallel to the optic axis. For a positive uniaxial crystal such as quartz, the extraordinary axis experiences a higher refractive index $n_e > n_o$, leading to slower propagation speeds and is thus called the *slow* axis. The retardation $\delta = (n_e - n_o)t/\lambda$ (in multiples of the wavelength λ) is then determined by the difference in refractive index and the material thickness t , which makes the effect sensitive to the angle of incidence, the temperature, and of course the used wavelength.

Although the retardation is fixed by design, arbitrary single-qubit operations can nonetheless be implemented using two kinds of waveplates. The so-called *quarter-waveplate* (QWP) and *half-waveplate* (HWP) impart a delay of a quarter wavelength, $\delta = 1/4$, and a half wavelength $\delta = 1/2$, respectively. In the Bloch sphere this corresponds to a rotation of $\pi/2$ and π , respectively, around an axis in the XZ-plane, determined by the orientation of fast and slow axis, see Fig. 1.9. As a trivial consequence, a waveplate does nothing⁷ to the states along this axis.

A combination of QWP-HWP-QWP is sufficient to go from any point on the Bloch sphere to any other point and thus implement arbitrary unitary transformations [15]. Starting from a fixed point, such as $|H\rangle$ in the case of state preparation (or measurement in reverse for that matter) makes the situation even simpler and two waveplates are sufficient (either QWP-HWP or HWP-QWP). Curiously, informationally complete measurements could even be done with a single *third-waveplate*, with a retardation of $1/3$ [58].

⁷In fact, the waveplate still imparts a global phase of π or $\pi/2$, which can be relevant if it is part of a multi-path interferometer.




Figure 1.9: **Waveplate rotation in Bloch sphere.** (a) Under the action of a half-waveplate the initial quantum state (red) undergoes a π -rotation to the final state (orange) about an axis (black line) in the xz-plane (gray shaded) of the Bloch sphere. The angle θ of the axis with respect to $|H\rangle$ is twice the physical angle of the waveplate. (b) Under the action of a quarter-waveplate the state undergoes a $\pi/2$ rotation around an axis in the xz-plane of the Bloch sphere. In both cases the states aligned with the axis of the waveplate remain unchanged up to a global phase.

Characterization of Waveplates

Although single-qubit manipulation is relatively simple for polarization qubits, great care has to be taken to achieve high precision. The two crucial parameters of a waveplate are the orientation θ_0 of the waveplate's optic axis with respect to horizontal and its retardation value δ . Precise characterization of these parameter is key to high-precision experiments. In practice, this is done by placing the unknown waveplate between two aligned (or crossed) polarization references (e.g. aligned calcite beam displacers), ideally in-situ, that is, at the final place in the setup. Recalling that whatever the actual parameters are, when the one of the waveplate axes are aligned with the reference, any waveplate acts as the identity. Rotating the waveplate results in a characteristic sinusoidal behaviour of the normalized transmitted intensity

$$f(\theta) = 1 - \frac{1}{2}(1 - \cos(2\pi\delta)) \sin^2(2(\theta - \theta_0)) \quad (1.29)$$

$$= 1 - V \sin^2(2(\theta - \theta_0)) , \quad (1.30)$$

where $V = 1/2(1 - \cos(2\pi\delta))$ is the observed visibility. The zero position θ_0 is estimated directly as a fit-parameter⁸, and is independent of the visibility. Hence, it could also be determined between crossed polarizers, while the retardation value δ can only be determined between aligned polarizers. Moreover, the relationship between retardation and visibility is non-linear with a divergent derivative at $\delta = 1/2$. This implies that, in the case of a HWP, the better the

⁸In fact, θ_0 is the angle of either fast or slow axis, and determining which of the two it is requires more complicated methods, such as using Fresnel reflection from a metallic surface [59]. In practice, however, it is almost always sufficient to simply make sure that all axes within an experiment are aligned relative to one another, which merely requires analysing the behaviour of two consecutive waveplates.

waveplate is, the more difficult it is to estimate its performance from the interference visibility. The fidelity of the prepared states with the ideal case, on the other hand, for both QWP and HWP, features the same sinusoidal dependence (shifted for the QWP) on the retardation as the visibility. Hence, the fidelity of a single waveplate operation is also quite resilient to imperfections in the retardation and thus a suboptimal measure for the quality of a waveplate. Significant deviations of the prepared state on the Bloch sphere may produce large fidelities, but lead to rapid error accumulation. Hence, although visibility and fidelity are good indicators for the performance of a waveplate in practice, high-precision experiments could benefit from more involved techniques in the spirit of gate-set tomography or measurement tomography, see Chap. 2.

Two-Qubit Entangling Operations

Not every quantum operation on multiple qubits can be decomposed into local single-qubit operations. They can, however, be decomposed into a series of single-qubit operations and (one kind of) two-qubit entangling gates. The most widely-used operation of this kind is the controlled-NOT (CNOT) gate, which acts on a pair of qubits, such that the operation performed on the second (the *target*) depends on the state of the first (the *control*). Specifically, the CNOT-gate flips the polarization of the target when the control is $|V\rangle$, thus implementing $|VH\rangle \leftrightarrow |VV\rangle$, and acts as the identity otherwise. This can be used to turn a product state $(|H\rangle + |V\rangle)_c \otimes |V\rangle_t$ into an entangled state $(|HH\rangle + |VV\rangle)_{ct}$, where the indices c and t denote control and target qubits, respectively and normalization has been omitted. In the computational basis $\{|HH\rangle, |HV\rangle, |VH\rangle, |VV\rangle\}$ the CNOT gate and the equivalent CZ gate can be represented as transformation matrices

$$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad \text{CZ} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}. \quad (1.31)$$

The CZ is a controlled σ_z -gate and thus related to the CNOT by a simple rotation of the Bloch-sphere of the target qubit, using a Hadamard gate $U_H = (\sigma_x + \sigma_z)/\sqrt{2}$ (a 90° rotation around the Y-direction of the Bloch sphere), such that $\text{CNOT} = (\mathbf{1} \otimes U_H) \cdot \text{CZ} \cdot (\mathbf{1} \otimes U_H)$. In contrast to the CNOT, however, the CZ-gate is symmetric in the two inputs and imparts a π -phase shift on the $|VV\rangle$ -component, which performs the operation $|VD\rangle \leftrightarrow |VA\rangle$, thus creating a maximally entangled state from the input $|DD\rangle$ rather than $|DH\rangle$ for the CNOT. The CZ-gate is relevant because it is closer to experimental implementations in linear optics.

Achieving the required π phaseshift on the $|VV\rangle$ -term using photons which are non-interacting, is quite difficult. The first strategy is to create an extremely non-linear environment, such as via interaction with atoms [60], which is technologically very challenging. The second strategy relies on non-classical interference and measurement-induced nonlinearities [51]. Two identical bosons (such as photons) have a curious property to bunch when incident on a 50:50 beam




Figure 1.10: **Experimental implementation of a controlled-not gate.** (a) The CNOT gate is equivalent to a CZ gate up to local Hadamard rotations on the target qubit. (b) In practice, a CZ gate is implemented using partially polarizing beam splitters (PPBS) and appropriate compensation, see text. The initial state preparation and measurement are performed using a Glan-Taylor polarizer (GT), and a set of HWP and QWP.

splitter from two different ports, such that they both exit at the same port, but never one each at the two outputs [61]. This is known as the *Hong-Ou-Mandel* (HOM) effect [62]. Classically, one would expect that each particle has a 50% chance to come out at either exit port, leading to a 50% chance of them exiting at different ports, see Fig. 1.11.

To exploit this effect, two independent single photons are interfered on a partially-polarizing beam splitter (PPBS), which reflects $2/3$ of $|V\rangle$ and transmits everything else. Denoting the creation operator for $|V\rangle$ in mode i by a_i^\dagger , this implements

$$a_1^\dagger a_2^\dagger \mapsto \frac{(a_1^\dagger + \sqrt{2}a_2^\dagger)(a_2^\dagger - \sqrt{2}a_1^\dagger)}{\sqrt{3}} = \frac{1}{3} \left(a_1^\dagger a_2^\dagger - \sqrt{2}(a_1^\dagger)^2 + \sqrt{2}(a_2^\dagger)^2 - 2(a_1^\dagger a_2^\dagger) \right). \quad (1.32)$$

The two middle terms in the right-hand side expression correspond to both photons exiting in port 1 or 2 and are discarded by post-selecting on having one photon in each arm. The other two terms interfere to give $-\frac{1}{9}|VV\rangle$. In order to balance the amplitudes of $|H\rangle$ and $|V\rangle$ it is then also necessary to discard $2/3$ of $|H\rangle$ in each arm.

The indistinguishability of the initial photons is crucial for the non-classical interference to occur. If the particles are distinguishable, for example because they arrive at slightly different times⁹, no interference is observed, leading to an intensity of $5/9|VV\rangle$. The characteristic drop in intensity of the $|VV\rangle$ -component¹⁰ when going from distinguishable to indistinguishable photons is also known as the HOM dip, see Fig. 1.11d. A CZ-gate implemented in this way is probabilistic and only works with a probability of $1/9$, since in all other cases photons exit at the same port. To filter these erroneous cases, it is necessary to postselect on those events where exactly one photon is found in each output port which implies that these gates cannot be concatenated¹¹. In contrast to most other architectures, however, linear optics has the

⁹At first sight it might seem obvious that photons that do not arrive at the same time cannot interfere, but, in fact, the interference can be restored by erasing the information about which photon was first [63].

¹⁰In fact, the same is true for every input state which has a $|VV\rangle$ -component.

¹¹A series of these gates can, however, be used in an overlapping ladder structure, as long as no two photons

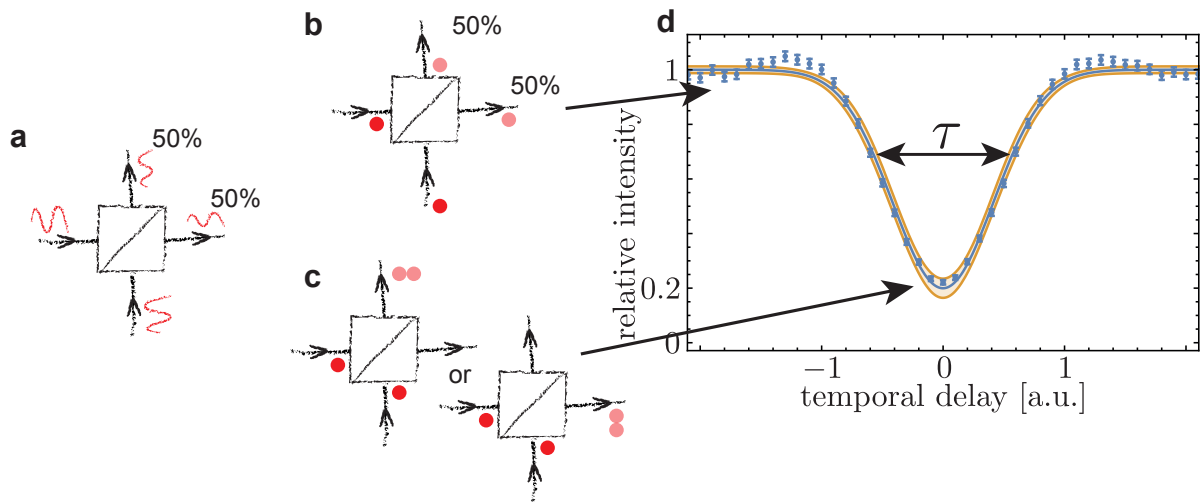


Figure 1.11: **Photons on a beam splitter.** (a) Classical waves on a 50/50 beam-splitter have equal probability for exiting at either port. (b) Two distinguishable single photons (e.g. arriving at different times) behave just as classical waves, but (c) two indistinguishable photons bunch and always exit the same port together. (d) The probability for finding one photon each in the two outputs as a function of their temporal separation—the HOM dip. The width of the dip profile is a quantitative measure of the coherence time τ of the photons, the temporal interval within which they interfere.

advantage that entangled pairs of photons are rather cheap and can be produced directly, without requiring an entangling gate. Using a single entangled pair as a resource, it is possible to implement a CZ-gate that has a success probability of $1/4$, and requires *only* post-selection of the ancillary photons, such that gates of this kind can be concatenated. If entangled pairs of photons are not available (e.g. when using on-demand quantum-dot based photon sources [65]), they can be created with low success probability from indistinguishable, separable photons using so-called fusion gates [66].

interfere twice [64]

References

- [1] Helstrom, C. W. Quantum detection and estimation theory. *J. Stat. Phys.* **1**, 231–252 (1969).
- [2] Bertlmann, R. A. & Krammer, P. Bloch vectors for qudits. *J. Phys. A* **41**, 235303 (2008).
- [3] Goyal, S. K., Simon, B. N., Singh, R. & Simon, S. Geometry of the generalized Bloch sphere for qutrits. *J. Phys. A* **49**, 165203 (2016).
- [4] Braunstein, S. L. & van Loock, P. Quantum information with continuous variables. *Rev. Mod. Phys.* **77**, 513 (2005).
- [5] Wigner, E. On the Quantum Correction For Thermodynamic Equil. *Phys. Rev.* **40**, 749–759 (1932).
- [6] Moyal, J. E. & Bartlett, M. S. Quantum mechanics as a statistical theory. *Mathematical Proceedings of the Cambridge Philosophical Society* **99** (1949).
- [7] Hillery, M., O’Connell, R., Scully, M. & Wigner, E. Distribution functions in physics: Fundamentals. *Phys. Rep.* **106**, 121–167 (1984).
- [8] Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).
- [9] Jozsa, R. Fidelity for Mixed Quantum States. *J. Mod. Opt.* **41**, 2315–2323 (1994).
- [10] Gilchrist, A., Langford, N. K. & Nielsen, M. A. Distance measures to compare real and ideal quantum processes. *Phys. Rev. A* **71**, 062310 (2005).
- [11] Branciard, C., Brunner, N., Gisin, N., Kurtsiefer, C., Lamas-Linares, A., Ling, A. & Scarani, V. Testing quantum correlations versus single-particle properties within Leggett’s model and beyond. *Nat. Phys.* **4**, 681–685 (2008).
- [12] Schrödinger, E. Die gegenwärtige Situation in der Quantenmechanik. *Die Naturwissenschaften* **23**, 807–849 (1935).
- [13] Gurvits, L. Classical deterministic complexity of Edmonds’ Problem and quantum entanglement. In *Proceedings of the thirty-fifth ACM symposium on Theory of computing - STOC ’03* (ACM Press, 2003).
- [14] Horodecki, R., Horodecki, M. & Horodecki, K. Quantum entanglement. *Rev. Mod. Phys.* **81**, 865–942 (2009).

- [15] Langford, N. K. *Encoding , manipulating and measuring quantum information in optics*. Phd, The University of Queensland (2007).
- [16] Cavalcanti, E. G., Menicucci, N. C. & Pienaar, J. L. The preparation problem in nonlinear extensions of quantum theory. *arXiv:1206.2725* (2012).
- [17] Ringbauer, M., Broome, M. a., Myers, C. R., White, A. G. & Ralph, T. C. Experimental simulation of closed timelike curves. *Nat. Commun.* **5**, 4145 (2014).
- [18] Shahandeh, F., Ringbauer, M., Loredo, J. C. & Ralph, T. C. Ultrafine Entanglement Witnessing. *arXiv:1610.04244* (2016).
- [19] Tóth, G., Wieczorek, W., Krischek, R., Kiesel, N., Michelberger, P. & Weinfurter, H. Practical methods for witnessing genuine multi-qubit entanglement in the vicinity of symmetric states. *New J. Phys.* 083002 (2009).
- [20] Dür, W., Vidal, G. & Cirac, J. I. Three qubits can be entangled in two inequivalent ways. *Phys. Rev. A* **62**, 062314 (2000).
- [21] Hill, S. & Wootters, W. K. Entanglement of a Pair of Quantum Bits. *Phys. Rev. Lett.* **78**, 5022–5025 (1997).
- [22] Wootters, W. K. Entanglement of Formation of an Arbitrary State of Two Qubits. *Phys. Rev. Lett.* **80**, 2245–2248 (1998).
- [23] Coffman, V., Kundu, J. & Wootters, W. K. Distributed entanglement. *Phys. Rev. A* **61**, 052306 (2000).
- [24] Cao, K., Zhou, Z.-W., Guo, G.-C. & He, L. Efficient numerical method to calculate the three-tangle of mixed states. *Phys. Rev. A* **81**, 034302 (2010).
- [25] Miyake, A. Classification of multipartite entangled states by multidimensional determinants. *Phys. Rev. A* **67**, 012108 (2003).
- [26] Wood, C. *Non-Completely Positive Maps : Properties and Applications*. Ph.D. thesis, Macquarie University (2009).
- [27] Ringbauer, M., Wood, C. J., Modi, K., Gilchrist, A., White, A. G. & Fedrizzi, A. Characterizing Quantum Dynamics with Initial System-Environment Correlations. *Phys. Rev. Lett.* **114**, 090402 (2015).
- [28] Wood, C. J., Biamonte, J. D. & Cory, D. G. Tensor networks and graphical calculus for open quantum systems. *Quantum Inf. Comput.* **15**, 0759–0811 (2011).
- [29] Stinespring, W. F. Positive functions on C*-algebras. *Proc. Am. Math. Soc.* **6**, 211–211 (1955).

- [30] Schumacher, B. Sending quantum entanglement through noisy channels. *Phys. Rev. A* **54**, 2614 (1996).
- [31] Choi, M.-D. Completely positive linear maps on complex matrices. *Linear Algebra Appl.* **10**, 285–290 (1975).
- [32] Jamiołkowski, A. Linear transformations which preserve trace and positive semidefiniteness of operators. *Rep. Math. Phys.* **3**, 275–278 (1972).
- [33] Wood, C. J. *Initialization and Characterization of Open Quantum Systems*. Ph.D. thesis, University of Waterloo (2015).
- [34] Born, M. On the Quantum Mechanics of Collisions. *Zeitschrift für Physik* **38**, 803–827 (1926).
- [35] Fuchs, C. A. QBism, the Perimeter of Quantum Bayesianism. *arXiv:1003.5209* (2010).
- [36] Aerts, D. A possible explanation for the probabilities of quantum mechanics. *J. Math. Phys.* **27**, 202–210 (1986).
- [37] Gleason, A. Measures on the Closed Subspaces of a Hilbert Space. *Indiana Univ. Math. J.* **6**, 885–893 (1957).
- [38] Carroll, S. M. Why probability in quantum mechanics is given by the wave function squared. Sean Carroll’s Blog (2014).
- [39] Davies, E. B. & J. T. Lewis. An Operational Approach to Quantum Probabillity. *Commun. Math. Phys.* **17**, 239–260 (1970).
- [40] Naimark, M. Spectral functions of a symmetric operator. *Izv. Akad. Nauk SSSR Ser.Mat.* **4**, 277–318 (1940).
- [41] Kwiat, P. G. & Englert, B.-G. Quantum-erasing the nature of reality, or perhaps, the reality of nature? In Barrow, J. D., Davies, P. C. W. & L., H. C. (eds.) *Science and Ultimate Reality: Quantum Theory, Cosmology, and Complexity* (Cambridge University Press, 2004).
- [42] Pusey, M. F. & Leifer, M. S. Logical pre- and post-selection paradoxes are proofs of contextuality. *Electron. Notes Theor. Comput. Sci.* **195**, 295–306 (2015).
- [43] Kocsis, S., Braverman, B., Ravets, S., Stevens, M. J., Mirin, R. P., Shalm, L. K. & Steinberg, A. M. Observing the Average Trajectories of Single Photons in a Two-Slit Interferometer. *Science* **332**, 1170–1173 (2011).
- [44] Bliokh, K. Y., Bekshaev, A. Y., Kofman, A. G. & Nori, F. Photon trajectories, anomalous velocities and weak measurements: a classical interpretation. *New J. Phys.* **15**, 073022 (2013).

- [45] Dressel, J. Weak values as interference phenomena. *Phys. Rev. A* **91**, 032116 (2015).
- [46] Dressel, J. & Jordan, a. N. Contextual-value approach to the generalized measurement of observables. *Phys. Rev. A* **85**, 022123 (2012).
- [47] Grangier, P., Levenson, J. A. & Poizat, J.-P. Characterization of quantum non-demolition measurements in optics. *Nature* **396**, 537–542 (1998).
- [48] Ralph, T. C., Bartlett, S. D., O’Brien, J. L., Pryde, G. J. & Wiseman, H. M. Quantum nondemolition measurements for quantum information. *Phys. Rev. A* **73**, 012113 (2006).
- [49] Monroe, C. Demolishing quantum nondemolition. *Physics Today* **64**, 8 (2011).
- [50] Brune, M., Haroche, S., Raimond, J. M., Davidovich, L. & Zagury, N. Manipulation of photons in a cavity by dispersive atom-field coupling: Quantum-nondemolition measurements and generation of “Schrödinger cat” states. *Phys. Rev. A* **45**, 5193–5214 (1992).
- [51] Knill, E., Laflamme, R. & Milburn, G. J. A scheme for efficient quantum computation with linear optics. *Nature* **409**, 46–52 (2001).
- [52] Carolan, J. *et al.* Universal linear optics. *Science* **349**, 711–716 (2015).
- [53] Molina-Terriza, G., Torres, J. P. & Torner, L. Twisted photons. *Nat. Phys.* **3**, 305–310 (2007).
- [54] Fernandez-Corbaton, I. & Molina-Terriza, G. Role of duality symmetry in transformation optics. *Phys. Rev. B* **88**, 085111 (2013).
- [55] Nagali, E. & Sciarrino, F. Manipulation of Photonic Orbital Angular Momentum for Quantum Information Processing. In *Advanced Photonic Sciences* (InTech, 2012).
- [56] Andersen, M., Ryu, C., Cladé, P., Natarajan, V., Vaziri, A., Helmerson, K. & Phillips, W. Quantized Rotation of Atoms from Photons with Orbital Angular Momentum. *Phys. Rev. Lett.* **97**, 170406 (2006).
- [57] Fickler, R. *Quantum Entanglement of Complex Structures of Photons*. Ph.D. thesis, University of Vienna (2015).
- [58] Hou, Z., Xiang, G., Dong, D., Li, C.-f. & Guo, G.-c. Realization of mutually unbiased bases for a qubit with only one wave plate: theory and experiment. *Opt. Express* **23**, 10018 (2015).
- [59] Logofatu, P. C. Simple method for determining the fast axis of a wave plate. *Optical Engineering* **41**, 3316 (2002).
- [60] Volz, J., Scheucher, M., Junge, C. & Rauschenbeutel, A. Nonlinear π phase shift for single fibre-guided photons interacting with a single resonator-enhanced atom. *Nat. Phot.* **8**, 965–970 (2014).

- [61] Holbrow, C. H., Galvez, E. & Parks, M. E. Photon quantum mechanics and beam splitters. *Am. J. Phys.* **70**, 260 (2002).
- [62] Hong, C. K., Ou, Z. Y. & Mandel, L. Measurement of subpicosecond time intervals between two photons by interference. *Phys. Rev. Lett.* **59**, 2044–2046 (1987).
- [63] Kim, Y.-H. & Grice, W. P. Quantum interference with distinguishable photons through indistinguishable pathways. *arXiv:quant-ph/0304086* 6 (2003).
- [64] Ralph, T. Scaling of multiple postselected quantum gates in optics. *Phys. Rev. A* **70**, 012312 (2004).
- [65] Loredó, J. C. *Enabling Multi-Photon Experiments with Solid-State Emitters: A Farewell to Downconversion*. Ph.D. thesis, The University of Queensland (2016).
- [66] Kok, P. & Lovett, B. W. *Optical Quantum Information Processing* (Cambridge University Press, 2010).

CHAPTER 2

Quantum Tomography

Acknowledgement

The second part of this chapter is based on work that was first published in Ref. [1], and I have incorporated large part of the text of that paper. I would particularly like to acknowledge Christopher Wood, who was chiefly responsible for the theory developed for both, reconstruction and analysis, used in that paper, which also forms part of his PhD thesis submitted to the University of Waterloo. I have included much of this work here, as it is a vital component for the experimental application of the introduced technique. Furthermore, I would like to gratefully acknowledge Christopher Wood for introducing me to the tensor network notation and the quantum tomography methodology that I am using throughout this chapter.

2.1 Introduction

IN practice, the experimenter always has to verify that the three components of a quantum experiment—state preparation, transformation, and measurement—are close to what the theorist wants them to be. The first step towards this is classical characterization of the devices (e.g. testing a waveplate with a laser against a reference), which especially in optics can go a long way to estimating the performance of the device in a quantum experiment [2]. On the other hand, if the goal is to reconstruct which state (or channel, or measurement) the experiment actually implemented, quantum tomography is the gold-standard. Just like tomography in medical imaging, the aim is to reconstruct the unknown state from a large number of snapshots taken at various angles, see Fig. 2.1. In contrast to the classical case, where all measurements are performed on the same patient, a quantum measurement necessarily disturbs the state of the system. As a consequence, it is not possible to reconstruct the exact quantum state of a single particle, and quantum tomography always reconstructs the (average) quantum state of an ensemble of identically prepared particles.

A set of different snapshots that contains enough information to fully determine any quantum state is called an *informationally-complete* set of measurements. For a d -dimensional system, with $d^2 - 1$ free parameters, the minimal number of measurements in such a set is d^2 . One measurement for each parameter and one additional measurement to normalize the experimental frequencies into probabilities. In a very similar fashion, tomography of a measurement uses an informationally complete set of preparations. Up to an additional parameter that can

be interpreted as a bias, the measurement also corresponds to a vector on the Bloch sphere, see also Sec. 1.3.

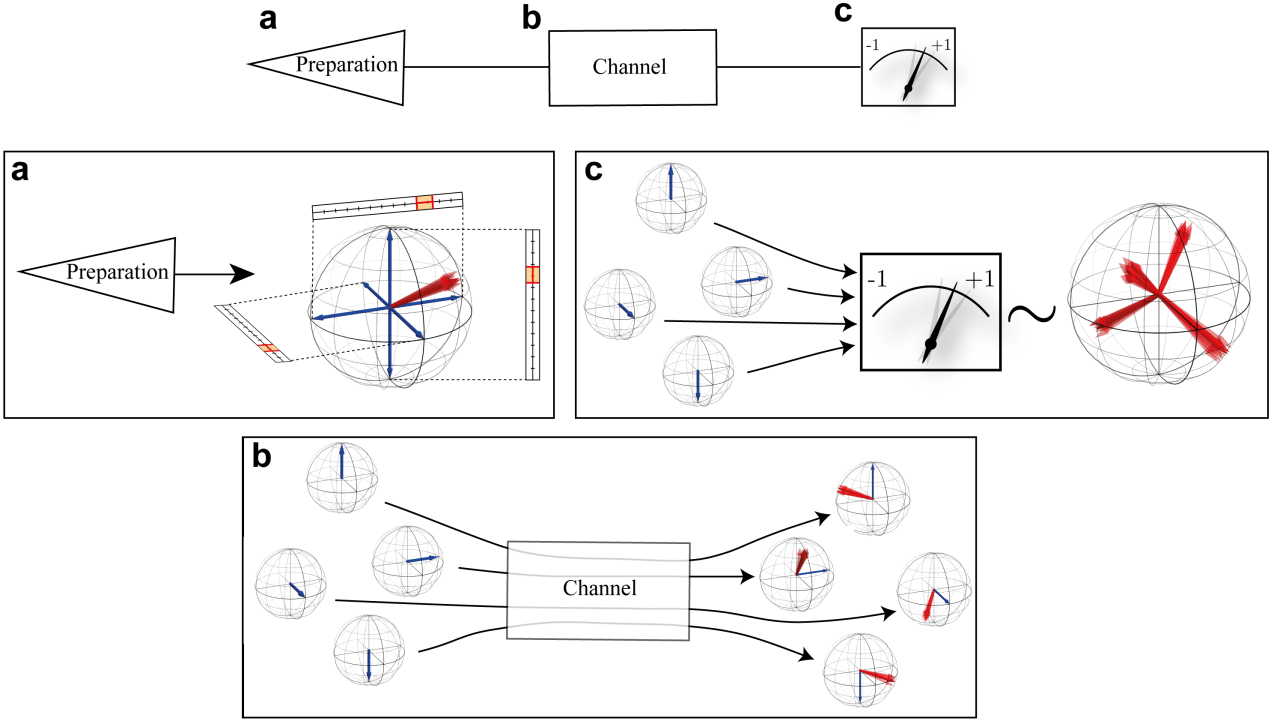


Figure 2.1: **Quantum tomography.** Every quantum experiment can be broken up into state preparation, evolution, and measurement. The aim of quantum tomography is to characterize the experimental implementation of these three steps. **(a)** Quantum state tomography aims to characterize the quantum state produced in the state-preparation step. This is achieved by subjecting this state to a set of well-calibrated measurements. **(b)** Quantum process tomography aims to reconstruct a description of the evolution of arbitrary states through the quantum channel. This is achieved by reconstructing the output states (using quantum state tomography) for a set of well-characterized input states. **(c)** Quantum Measurement Tomography aims to reconstruct a description of the measurement that is implemented by a given measurement device. By probing the measurement device with a set of well-characterized input states, it is possible to reconstruct the measurement operators from the observed statistics.

2.2 Quantum State Tomography

Consider, for example, a single qubit, whose state $\rho = (\mathbb{1} + \vec{n} \cdot \vec{\sigma})/2$ is completely described by the $d^2 - 1 = 3$ real valued parameters that are the components of the Bloch-vector (n_x, n_y, n_z) , see Fig. 2.1. The simplest procedure to find the quantum state is thus by performing measurements along the axes of the Bloch sphere in order to determine the components of the Bloch vector individually. This is an example of a set of *mutually unbiased bases*, which are sets of orthonormal bases, such that the modulus squared of the inner product between elements of different bases is a constant $1/d$, see Fig. 2.2a. In other words, any state from one of these

bases has equal probability for all outcomes when measured in any of the other bases in the set. This latter property is particularly interesting for applications such as quantum cryptography. Although it is unclear whether such sets exist in general, they are known to exist for the most relevant cases of Hilbert spaces with prime-power dimension [3]. This includes multi-qubit spaces, which makes these sets the most widely used measurement for multi-qubit quantum tomography.

In practice, these measurements are often the easiest and most precise. However, they are *over-complete* in the sense that they use more than the d^2 required measurements. As in the classical case, more measurements typically means more reliable results, but at the same time the number of measurements required to characterize a multi-qubit system scales exponentially with the number of qubits. One way to reduce the number of measurements is to only measure one outcome for each axes except one, and assume that the total number of events is the same along every axis. This, however, is a rather asymmetric arrangement of measurements. A more important class of minimal sets of tomography measurements are so-called *symmetric informationally-complete* POVMs [4]. They consist of d^2 subnormalized projectors, which have a fixed inner product and are thus arranged symmetrically in the respective Hilbert space. Such measurements are conjectured to exist in every dimension, but their existence has only been shown numerically for dimensions up to 67. In the case of a single qubit the outcomes of the SIC POVM correspond to the vertices of a regular tetrahedron, see Fig. 2.2.




Figure 2.2: **The two most relevant choices of tomographically complete measurements for a single qubit.** (a) The standard tomography set are the mutually unbiased bases corresponding to the Pauli operators, i.e. the axes of the Bloch-sphere. (b) The minimal set of measurements in the form of a symmetric informationally POVM forms a tetrahedron in the Bloch-sphere. In this case all projectors are sub-normalized to $1/2$ to satisfy the normalization of probabilities required for a POVM.

2.2.1 Linear Inversion Tomography

In practice, the simplest algorithm for quantum tomography is linear inversion. In the following, the notation of Ref. [5] is used and the interested reader is referred to Ref. [5] for a more in-depth discussion. Consider the problem of state-tomography of a d -dimensional system using

a set of informationally complete (potentially over-complete) measurement operators $\{E_i\}_{i=1}^K$, with $K \geq d^2$. Using vectorized notation, the expected probabilities for the various measurement outcomes are then given by $P_i = \text{Tr}[E_i \rho] = \langle\langle E_i | \rho \rangle\rangle$. One can then define a vector $|P\rangle$ of observed probabilities, a matrix S of vectorized measurement operators, and a matrix W of weights [5]

$$\begin{aligned} |P\rangle &= \sum_{i=1}^K P_i |i\rangle \\ S &= \sum_{i=1}^K |i\rangle \langle\langle E_i | \\ W &= \sum_{i=1}^K w_i |i\rangle \langle i|. \end{aligned} \tag{2.1}$$

Here, $|i\rangle$ denotes the unit vector with a single 1 at index i , and $w_i \geq 0$ are weights which specify the relative importance of the various measurements. The matrix S contains the vectorized measurement operators, such that in the ideal case $S|\rho\rangle\rangle = |P\rangle$. The linear inversion problem is finding a $\hat{\rho}_{\text{LIN}}$ that satisfies the equation $S|\rho\rangle\rangle = |P\rangle$. This estimator is given by

$$\hat{\rho}_{\text{LIN}} = (S^\dagger S)^{-1} S^\dagger |P\rangle. \tag{2.2}$$

In the case of a tomographically complete or over-complete set of measurements, the matrix $(S^\dagger S)$ is indeed invertible [5], otherwise the Moore-Penrose pseudo-inverse can be used. The linear inversion estimator is a special case of a weighted least-squares fit [5]

$$\hat{\rho} = \text{argmin}_\rho \|WS|\rho\rangle\rangle - W|P\rangle\|_2, \tag{2.3}$$

where $\|\cdot\|_2$ is the Euclidean vector norm. This has the analytic solution

$$\begin{aligned} |\hat{\rho}\rangle\rangle &= (S^\dagger W^\dagger W S)^{-1} S^\dagger W^\dagger W |P\rangle \\ &= \sum_i w_i^2 P_i \left(\sum_k w_k^2 |E_k\rangle\rangle \langle\langle E_k| \right)^{-1} |E_i\rangle\rangle, \end{aligned} \tag{2.4}$$

which, in the case of uniform weights $w_i = 1$, reduces to Eq. (2.3).

This simple technique, however, does not cope very well with imperfect experimental data. As a result of experimental noise, the obtained estimator $\hat{\rho}_{\text{LIN}}$ is typically not a valid (i.e. positive semi-definite, Hermitian, trace-one) density matrix [6, 7]. The problem is that the procedure does not know how a physical density matrix looks like and simply finds a matrix that works.

2.2.2 Maximum Likelihood Estimation

The main problem of linear inversion, that the result is typically unphysical, can be overcome by only searching the space of valid density matrices for a solution. This solution can in general not be expected to reproduce the observed data perfectly, but it is the physical state that is most likely to have produced the observed data.

To formalize this, note that the probability for jointly observing a set of counts $\{n_1, \dots, n_K\}$ from measurements $\{E_i, \dots, E_K\}$, given that the measured quantum state is ρ , is simply the product of the individual probabilities. Turning this around, the same product describes the *likelihood* L that the measured state was ρ , given that the set of counts $\{n_1, \dots, n_K\}$ was observed

$$L(\rho|\{n_i\}) = P(\{n_i\}|\rho) = \prod_{i=1}^K P(n_i|\rho), \quad (2.5)$$

The goal of *maximum likelihood estimation* [8, 9] is to find the quantum state $\hat{\rho}_{\text{MLE}}$ that maximizes the likelihood function, or equivalently minimizes the negative log-likelihood function (which turns out to be computationally more stable). To make the optimization computationally tractable, it is typically assumed that the probability of observing n_i counts from the measurement E_i is well approximated by a normal distribution centred around the estimator $\text{Tr}[E_i\rho]$

$$P(n_i|\rho) \propto \exp\left[-\frac{(n_i - N_i \text{Tr}[E_i\rho])^2}{2\sigma_i^2}\right]. \quad (2.6)$$

Here, N_i denotes the total number of counts in the basis that contains the measurement E_i . This takes into account that, due to loss or other factors, it cannot be assumed that the total number of counts is the same in every basis. For computational reasons it is also convenient to approximate the estimator-variance σ_i^2 , which a priori depends on ρ , by the constant variance of the observed counts n_i . When conditioned on a fixed N_i the (unconditionally Poisson-distributed) counts n_i follow a binomial distribution with variance $\sigma_i^2 \sim N_i p_i (1 - p_i)$. Hence, the approximation of this distribution by a normal distribution is only valid for large N_i , and when $p_i \neq 0, 1$. Under these conditions maximum likelihood estimation solves the optimization problem [5]:

$$\begin{aligned} \hat{\rho}_{\text{MLE}} &= \arg \min_{\rho} [-\log L(\rho|\{n_i\})] \\ &= \arg \min_{\rho} \sum_{i=1}^K \frac{(n_i - N_i \text{Tr}[E_i\rho])^2}{N_i p_i (1 - p_i)}. \end{aligned} \quad (2.7)$$

As shown in Ref. [5] this can be recast as a constraint least-squares optimization

$$\begin{aligned} &\text{minimize} \quad \|WS|\rho\rangle\rangle - W|p\rangle\|_2 \\ &\text{subject to:} \quad \rho \geq 0, \quad \text{Tr}[\rho] = 1, \end{aligned} \quad (2.8)$$

where the weights are chosen such that

$$w_i = \frac{N_i}{\sigma_i} = \sqrt{\frac{N_i}{p_i(1-p_i)}}, \quad (2.9)$$

The linear inversion method of Eq. (2.3) is a special case where the constraints of Eq. (2.8) are dropped and uniform weights $w_i = 1$ are used. Since the objective function $\|\cdot\|_2$ in Eq. (2.8) is convex and the constraints are semidefinite, the problem can be solved numerically as a *semidef-*

inite program. See Ref. [5] for more details.

2.2.3 Zero Probabilities

If any of the tomography measurements happen to return zero counts, the estimator may assign a probability of zero to it. However, no finite dataset can do justice to a probability of zero¹ [10] since it cannot exclude the possibility that this event is just very unlikely. Moreover, the normal approximation in Eq. (2.9) breaks down and it is not possible to assign any non-zero error to the estimate either [10]. In the case of binary-outcome measurements, a simple way to avoid zero probabilities is to add a small amount of noise [5, 11] to the observed counts, defining the observed probabilities

$$f_i = \frac{n_i + \beta}{N_i + 2\beta}, \quad (2.10)$$

where the factor of 2 in the denominator equals the number of measurements used to normalize the experimental frequencies (i.e. 2 for a two-outcome measurement). The *hedging* parameter β can be chosen freely, but a value of $\beta \sim \frac{1}{2}$ works well in most cases, except almost pure states, where lower values work better [5, 12]. This approach works well to avoid zero-probabilities in the data, but it cannot guarantee that the reconstructed state doesn't *predict* zero probabilities. In particular in situations with low statistics and an unlucky experimenter, it can happen that, even after adding β , the linear inversion solution still lies outside the Bloch-sphere. The maximum likelihood estimator in this case would lie on the surface of the sphere and thus predict probability zero for orthogonal measurements. This can be avoided using *hedged maximum likelihood estimation* [12] where the standard likelihood function of Eq. (2.5) is replaced by product $L(\rho|\{n_i\}) \det(\rho)^\beta$. The optimization problem is then

$$\begin{aligned} & \text{minimize} \quad \frac{1}{2} \|WS|\rho\rangle\rangle - W|p\rangle\|_2^2 - \beta \log \det(\rho) \\ & \text{subject to:} \quad \rho \geq 0, \quad \text{Tr}[\rho] = 1. \end{aligned}$$

This effectively pushes the estimator away from the boundary, but does not change it significantly if beta is chosen as above. In the case of measurements in only a single basis, this actually recovers Eq. (2.10) exactly [12].

2.2.4 Error Bars for Quantum Tomography

Both linear inversion and maximum likelihood estimation produce point estimates $\hat{\rho}$. Since such estimates can never hope to match the true ρ exactly (except in the asymptotic limit), they are only physically meaningful if they have some uncertainty bounds. In the special case of polarization qubits, error estimation for tomography has been discussed in Ref. [9]. More generally, standard approaches typically use the variance of the point estimator to define

¹As discussed in Ref. [10], maximum likelihood estimation is a frequentist tool, which works in the asymptotic limit of infinite sample size, but “applying a frequentist method to relatively small amounts of data is inherently disaster-prone”

ellipsoidal confidence regions. These regions, however, might include unphysical states, even if the point estimator is physical. This can in general be avoided by constructing the regions via bootstrapping methods [13] that perform many iterations of the tomography on either resampled experimental data or experimental data with simulated statistical noise. This method is used in large parts of the material presented later in this thesis.

However, also the latter method quantifies the variance of the point estimator, which, in maximum likelihood methods, is in general biased by the positivity constraint (except in the asymptotic limit) [6, 14]. Indeed, maximum likelihood estimators typically *underestimate* the fidelity with the true state [6, 7], and the uncertainty regions of such a biased estimator are in general suboptimal [14]. In contrast, linear inversion estimators are unbiased but unphysical in a large number of cases and produce larger mean squared errors, which makes them of limited use in practice [7].

A reliable error region for maximum likelihood estimation should contain the true state with high probability and contain no unphysical states. One suggestion that satisfies these requirements are the *likelihood ratio confidence regions* $R_\alpha(D)$ introduced in Ref. [14].

$$R_\alpha(D) = \{\rho \mid \lambda(\rho) < \lambda_\alpha\}, \quad (2.11)$$

where $\lambda(\rho)$ is the log-likelihood ratio $\lambda(\rho) = -2\log[\mathcal{L}(\rho)/\max_{\rho'} \mathcal{L}(\rho')]$ and λ_α is a threshold level, such that region $R_\alpha(D)$ contains the true ρ with probability at least α . The region $R_\alpha(D)$ is a convex region and is proven to be a near-optimal region estimator [14]. Practically, however, these regions are difficult to work with since they generally do not have an explicit description and can have quite arbitrary shapes.

2.2.5 Quantum Process Tomography

Besides quantum state tomography, one of the most important motivations for quantum tomography is to reconstruct the quantum process \mathcal{E} implemented by the experiment. Although the experimenter typically has some prior knowledge what he or she intended to build, this can in principle also be done in a completely black-box fashion. Since a quantum channel maps density matrices to other density matrices one has to reconstruct the output state for each of an informationally complete set of inputs. The linearity of quantum mechanics implies that this is sufficient to know how an arbitrary input state evolves. Intuitively, knowing how the axes of the Bloch-sphere are transformed is enough to know how every point of the sphere is transformed. Exploiting the Choi-Jamiolkowski isomorphism the reconstruction of a quantum process \mathcal{E} acting on a d -dimensional system can be mapped to state-tomography of the d^2 -dimensional Choi matrix $\Lambda_\mathcal{E}$. Using Eq. (1.23) the probability p_{ij} for observing the measurement outcome E_j given that the state ρ_i was subject to the process is

$$p_{ij} = \text{Tr}[E_j^\dagger \text{Tr}_1[(\rho^T \otimes \mathbb{1})\Lambda_\mathcal{E}] = \text{Tr}[(\rho^T \otimes E_j^\dagger)\Lambda_\mathcal{E}] = \langle\langle \rho^* \otimes E_j \mid \Lambda_\mathcal{E} \rangle\rangle.$$

Note that ρ^* is now interpreted as a “measurement” on the Choi matrix. With this identification one can then use the same tools as for quantum state tomography to formulate quantum process tomography as a constrained optimization problem [5]

$$\begin{aligned} & \text{minimize} && \|WS|\Lambda_{\mathcal{E}}\rangle\rangle - W|p\rangle\|_2 \\ & \text{subject to:} && \Lambda_{\mathcal{E}} \geq 0, \quad \text{Tr}[\Lambda_{\mathcal{E}}] = d, \end{aligned} \tag{2.12}$$

where $S = \sum_{i,j=1}^K |i\rangle\langle\rho_i^* \otimes E_j|$ and straight-forward generalizations of $|p\rangle$, and W from Eq. (2.1). In contrast to state tomography, the Choi matrix is normalized such that $\text{Tr}[\Lambda_{\mathcal{E}}] = d$ rather than 1.

Practicality

Since process tomography of an n -qubit process scales as 2^{4n} , it quickly becomes impractical. Hence it would be desirable to use (partially) classical methods to quickly assess the performance of a quantum process without the constraints of process tomography. In a *dual-rail* picture, where the two levels of the qubit correspond to two distinct paths, any n -qubit quantum channel can be described by a $2n$ linear optical network of beam-splitters and phase-shifters. The process is then conveniently represented by a $2n \times 2n$ *transfer matrix*, which describes how every input mode is mapped to every output mode. Characterizing the matrix amounts to measuring the amplitudes and phases of every entry, which can be with a coherent state exciting one or at most two input ports [2].

2.2.6 Caveats and Generalizations

One of the challenges of tomography is to avoid circularity. State tomography relies on well-characterized measurements, but to verify the measurements, well-characterized states are needed. A priori there is no ultimate reference to compare against. As a consequence, the results of quantum tomography can be quite confounded by errors in state preparation and/or measurement. Consider, for example, a single qubit prepared by a polarizer and a set of waveplates. The state tomography might return a purity of 0.9 for this state, which might of course be a result of the system getting entangled with the environment. However, since linear optics is blessed (or cursed) with quantum systems that do not interact much with anything, a more likely explanation is that the tomographic measurements were simply not performing as expected (e.g. unbalanced loss).

Randomized Benchmarking and Gate-Set Tomography

In practice, state preparation uses a fixed element, such as a polarizer or ground-state cooling, followed by a rotation, that is, a quantum channel. Similarly, measurements are implemented using a rotation, followed by a fixed projection, typically in the computational basis. These elements can typically be characterized classically to estimate their performance in the experiment. In linear optics, for example, polarizers based on birefringent crystals reliably split an

incoming beam into two orthogonal polarizations with a practical extinction ratio of at least $1 : 10^4$. Importantly, these elements are fixed throughout the experiment and are typically very repeatable which suggests that the crucial aspect is the characterization of quantum processes and the dynamical elements in the experiment.

Randomized benchmarking [15, 16] addresses this challenge using sequences of quantum gates which are chosen at random, such that they ideally add up to the identity. These sequences are applied to a fixed input state and the output is measured using a fixed two-outcome measurement. Since all sequences are decompositions of the identity, all output states should be the same. However, due to error accumulation the fidelity decreases with the length of the sequence, and the slope of fidelity vs. sequence-length gives an estimate of the average error per gate.

Gate-set tomography [17] goes a step further and aims to provide a complete characterization of the experiment without requiring prior knowledge or perfect references. The experiment is treated as a black-box that can implement a discrete set of quantum gates (buttons on the box). In addition, the box is capable of repeatably preparing an input state and performing a two-outcome measurement in the end. In contrast to other techniques, including randomized benchmarking, there are only very mild assumptions on preparation and measurement, such as repeatability of preparation and measurement, but it is not necessary to know the input state or the measurement that was performed. The collected data for different gate-series is used to reconstruct the gate-set, input state and measurement, up to a gauge freedom. In contrast to randomized benchmarking, which averages over coherent errors such as over-rotations, gate-set tomography is sensitive to these errors and provides a better picture of what is actually happening [18]. Gate-set tomography is also robust against local maxima in the likelihood function [17].

Although these methods are in principle independent of the physical implementation, they are particularly useful in architectures like ion traps, where quantum systems are (mostly) stationary and transformations are induced by lasers, which enables long sequences of many randomly chosen operations. Linear optics, on the other hand, is a bit special in this regard since gates are realized by actual physical elements rather than laser pulses, and the length of the gate-sequence can thus not be changed as dynamically.

2.3 Taming Non-Completely-Positive Maps with Super-channels

This section is based on the publication “Characterizing Quantum Dynamics with Initial System-Environment Correlations” [1].

So far, quantum tomography was concerned with isolated quantum systems and processes. In real-world experiments, however, quantum systems are inevitably coupled to an environment, which typically acts as a source of noise, but may also be harnessed as a resource—for example in initializing quantum states that may be otherwise unobtainable [19–26]. In either case

understanding the joint behaviour of system and environment is essential. Quantum mechanics postulates that the joint system-environment (SE) state evolves unitarily, which need not be true for the system alone. The theory of open quantum systems, nevertheless allows for an operationally complete description of the reduced dynamics of the system, in the case that the initial system-environment state is uncorrelated [27], see Fig. 2.3a. This central assumption is often, however, at best an approximation [28, 29].

In practice, system and environment may be correlated even before the initial state of the system is prepared. Although the state preparation erases all pre-existing correlations, it might leave the environment in a different state, depending on which state of the system was prepared. When these different environment states couple to the evolution again at a later stage, the implemented process is not just non-unitary, but also non-CP [1, 5, 28], see Fig. 2.3. Consider, for example, the extreme case of a maximally entangled initial system-environment state: $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)_{SE}$. A projective preparation of the system into $|0\rangle$ or $|1\rangle$ leaves the environment in orthogonal states. Hence, if the subsequent system evolution is not perfectly isolated, it is coupled to different environment states leading to drastically different reduced dynamics of the system conditional on the used preparation procedure [30]. Standard characterization techniques may in this case return a description of the reduced system dynamics that appears unphysical [28, 31–34]. On the other hand, when complete-positivity is enforced, such as in maximum likelihood quantum process tomography, the reconstructed map does not reliably describe the system dynamics.

While the environment is typically inaccessible to the experimenter, recent results suggest that at least partial information about the initial joint system-environment state can be extracted from measurements of the system alone. Initial correlations can be witnessed through the distinguishability [35–38] and purity [39] of quantum states, which has also been explored experimentally [40–42]. A more operationally complete characterisation can be obtained by using a *quantum superchannel* \mathcal{M} , which explicitly uses the system’s preparation procedure, rather than the prepared state, as an input [34], see Fig. 2.3b. This superchannel approach captures not just the system evolution, but also the dynamical influence of the environment, even in the presence of initial system-environment correlations. Remarkably, it can be experimentally reconstructed using only measurements on the system, and contains quantitative information about the initial correlations and the influence of the environment.

2.3.1 Constructing the Superchannel \mathcal{M}

Consider the situation in Fig. 2.3b, where a system with Hilbert space \mathcal{X}_1 , and an environment with Hilbert space \mathcal{Y}_1 are initially in the joint state $\rho_{SE} \in L(\mathcal{X}_1 \otimes \mathcal{Y}_1)$. In the following $L(\mathcal{H})$ denotes the set of linear operators on the space \mathcal{H} , and $T(\mathcal{H}_1, \mathcal{H}_2)$ denotes the set of linear transformation from the space \mathcal{H}_1 to the space \mathcal{H}_2 . The joint state ρ_{SE} is then subject to a preparation procedure $\mathcal{P} = (\mathcal{P}_S \otimes \mathcal{I}_E) \in T(\mathcal{X}_1 \otimes \mathcal{Y}_1, \mathcal{X}_2 \otimes \mathcal{Y}_2)$, where $\mathcal{Y}_2 = \mathcal{Y}_1$ (due to the identity operation), \mathcal{I}_E is the identity channel on the environment, and $\mathcal{P}_S \in T(\mathcal{X}_1, \mathcal{X}_2)$ acts only on the system to prepare it in a desired input state. This is followed by coupled evolution of the joint




Figure 2.3: **Quantum tomography with an environment.** (a) Quantum process tomography (QPT) is designed for a scenario where an isolated quantum system is prepared in the initial state ρ_S and undergoes a joint unitary evolution U with the environment. Process tomography can reconstruct the reduced dynamics of an informationally complete set of states $\{\rho_{\text{in}}^i\}$ and a tomographically complete set of measurements $\{M_j\}_{j=1}^d$ for each input. (b) The situation depicted in (a) is at best an approximation of the real state of affairs. In any realistic scenario the system is initially correlated with the environment, even before the state preparation procedure \mathcal{P} . Furthermore, the evolution \mathcal{U} is in general also affected by the environment. The quantum superchannel (QSC) takes all these effects into account and treats the experiment as a whole. It can be reconstructed from an informationally complete set of preparation procedures and measurements.

system-environment state, described by a CPTP map $\mathcal{U} \in T(\mathcal{X}_2 \otimes \mathcal{Y}_2, \mathcal{X}_3 \otimes \mathcal{Y}_3)$, see Fig. 2.3b. The output is then given by

$$\begin{aligned}
 \rho'_S &= \text{Tr}_{\mathcal{Y}_3} [\mathcal{U}(\mathcal{P}_S(\rho_{SE}))] \\
 &= \text{Tr}_{\mathcal{X}_2, \mathcal{Y}_2, \mathcal{Y}_3} [(\mathcal{P}(\rho_{SE})^T \otimes \mathbb{1}_{SE})\Lambda_{\mathcal{U}}] \\
 &= \text{Tr}_{\mathcal{X}_2, \mathcal{Y}_2, \mathcal{Y}_3} [(\text{Tr}_{\mathcal{X}_1, \mathcal{Y}_1} [(\rho_{SE}^T \otimes \mathbb{1}_{SE})\Lambda_{\mathcal{P}}]^T \otimes \mathbb{1}_{SE})\Lambda_{\mathcal{U}}].
 \end{aligned} \tag{2.13}$$

Figure 2.4 illustrates the construction of the superchannel \mathcal{M} to describe the evolution of Eq. (2.13), and visualizes how it is related to the Choi matrix of the joint evolution. The equivalent definition of the superchannel Choi matrix $\Lambda_{\mathcal{M}}$ in terms of tensor indices is derived in the supplement of Ref. [1], and given by

$$(\Lambda_{\mathcal{M}})_{i_1 i_2 i_3 | j_1 j_2 j_3} = \sum_{n, m, l} (\rho_{SE})_{i_1 n | j_1 m} (\Lambda_{\mathcal{U}})_{i_1 n i_3 | j_1 m j_3 l}, \tag{2.14}$$

By construction, $\Lambda_{\mathcal{M}} \geq 0$ if \mathcal{U} is CP. However, the same does not hold true for TP. If \mathcal{U} is TP, then

$$\text{Tr}_{\mathcal{X}_3} [\Lambda_{\mathcal{M}}] = \text{Tr}_{\mathcal{Y}_1} [\rho_{SE}] \otimes \mathbb{1}_{\mathcal{X}_2}. \tag{2.15}$$

Hence \mathcal{M} is TP if and only if $\text{Tr}_{\mathcal{Y}_1} [\rho_{SE}] = \mathbb{1}_{\mathcal{X}_1}$, which is the case only for a maximally entangled initial state or if the system is initially in a completely mixed state. In all other cases, different preparation procedures lead to different overall count rates and the superchannel is not trace




Figure 2.4: **Graphical derivation of the superchannel \mathcal{M} for describing the reduced dynamics of a system initially correlated with the environment.** Note that Λ_U and $\Lambda_{\mathcal{P}}$ have each 4 subsystem indices, which correspond to S -input, E -input, S -output and E -output, respectively, and different colors are used to distinguish Choi matrices (green) from density matrices (blue). The expression in Eq. (2.13) corresponds to the third term. Since all tensor wires are contracted the position of ρ_{SE} and $\Lambda_{\mathcal{P}}$ can be exchanged and the preparation procedure can be treated as the effective input state. The Choi matrix $\Lambda_{\mathcal{M}}$ of the superchannel is then defined as the contraction of ρ_{SE} and Λ_U , equivalently to the index contraction given in Eq. (2.14).

preserving. For a TP map \mathcal{U} , the Choi matrix for \mathcal{M} has normalization

$$\text{Tr}[\Lambda_{\mathcal{M}}] = \text{Tr}[\mathcal{U}(\mathbb{1}_{\mathcal{X}_1} \otimes \text{Tr}_{\mathcal{X}_1}[\rho_{SE}])] = d_{\mathcal{X}_1} . \quad (2.16)$$

The quantum superchannel $\mathcal{M} \in T(\mathcal{X}_1 \otimes \mathcal{X}_2, \mathcal{X}_3)$ thus takes the system-preparation procedure $\Lambda_{\mathcal{P}_s} \in L(\mathcal{X}_1 \otimes \mathcal{X}_2)$ as an input and produces an output quantum state $\rho' \in L(\mathcal{X}_3)$ given by

$$\begin{aligned} \rho' &= \mathcal{M}(\Lambda_{\mathcal{P}_s}) \\ &= \text{Tr}_{\mathcal{X}_1, \mathcal{X}_2}[(\Lambda_{\mathcal{P}_s} \otimes \mathbb{1}_{\mathcal{X}_3})\Lambda_U] . \end{aligned}$$

2.3.2 Superchannel Tomography

The crucial property of the superchannel Choi matrix $\Lambda_{\mathcal{M}}$ is that *all* environment indices are contracted and it can thus be tomographically reconstructed from measurements on the system alone. Since the superchannel takes a preparation procedure as an input, this requires an informationally complete set of *preparation procedures*, rather than just a set of prepared states, as well as an informationally complete set of measurements for each output state. A projective preparation procedure $\mathcal{P}_{ij} \in T(\mathcal{X}_1, \mathcal{X}_2)$, consists of an initial projection (or postselection) onto the state ρ_i followed by a rotation to the state ρ_j . The corresponding Choi matrix is thus given by

$$\Lambda_{\mathcal{P}_{ij}} = \rho_i^* \otimes \rho_j ,$$

where $*$ denotes complex conjugation. The probability of observing a click when preparing a state using $\Lambda_{\mathcal{P}_{ij}}$ and then measuring the system by projecting onto a state ρ_k is given by

$$\begin{aligned} p_{ijk} &= \text{Tr} \left[\rho_k^\dagger \text{Tr}_{12} [(\rho_i^\dagger \otimes \rho_j^T \otimes \mathbb{1}_{\mathcal{X}_3}) \Lambda_{\mathcal{M}}] \right] \\ &= \text{Tr} \left[(\rho_i^\dagger \otimes \rho_j^T \otimes \rho_k^\dagger) \Lambda_{\mathcal{M}} \right] \\ &= \langle\langle \rho_i \otimes \rho_j^* \otimes \rho_k | \Lambda_{\mathcal{M}} \rangle\rangle = \langle\langle \Pi_{ijk} | \Lambda_{\mathcal{M}} \rangle\rangle, \end{aligned}$$

where $\Pi_{ijk} \equiv \rho_i \otimes \rho_j^* \otimes \rho_k$. For an informationally complete set of states $\{\rho_i\}_{i=1}^K$, with $K \geq d^2$, one can generalize the vector $|p\rangle$, and matrices S and W introduced in Sec. 2.2.1 to

$$\begin{aligned} f_{ijk} &= \frac{n_{ijk} + \beta}{N_{ijk} + 4\beta} \\ |p\rangle &= \sum_{i,j,k=1}^K f_{ijk} |i, j, k\rangle \\ S &= \sum_{i,j,k=1}^K |i, j, k\rangle \langle\langle \Pi_{ijk} | \\ W &= \sum_{i,j,k=1}^K w_{ijk} |i, j, k\rangle \langle i, j, k|, \end{aligned} \tag{2.17}$$

where f_{ijk} are the hedged experimental frequencies to avoid issues associated with zero probabilities, see Sec. 2.2.3. The total number of events N_{ijk} , which is unknown a priori, is defined by totalling the observed counts for measurement configurations that sum to identity. Since the second index Π_{ijk} corresponds to the rotated state for the initial projective preparation procedure, and thus only the first and third indices correspond to true measurements. The experimental frequencies are thus normalized to $d^2 = 4$ counts, which leads to the factor of 4 in the definition of the hedged frequencies. For the choice of weights w_{ijk} a normal approximation for the distribution of the observed probabilities p_{ijk} is assumed, so that

$$w_{ijk} = \sqrt{\frac{N_{ijk}}{p_{ijk}(1-p_{ijk})}}.$$

As with quantum state- and process- tomography, one can use maximum likelihood estimation, see Sec. 2.2.2, to reconstruct the Choi matrix of the superchannel $\Lambda_{\mathcal{M}}$, by solving the constrained optimization problem

$$\begin{aligned} &\text{minimize} \quad \|WS|\Lambda_{\mathcal{E}}\rangle\rangle - W|p\rangle\|_2 \\ &\text{subject to:} \quad \Lambda_{\mathcal{E}} \geq 0, \quad \text{Tr}[\Lambda_{\mathcal{E}}] = d. \end{aligned} \tag{2.18}$$

Recall that enforcing complete positivity in maximum likelihood tomography is an effective way for combating statistical noise, which might lead to apparent non-CP dynamics. However, in the presence of initial system-environment correlations the reduced evolution of the system would be genuinely non-CP, which can thus not be correctly reconstructed using maximum likelihood quantum process tomography. In contrast, the superchannel \mathcal{M} fully takes the effect of state

preparation into account and is therefore always a completely positive map. As a consequence, the additional constraint of maximum likelihood estimation is justified for overcoming statistical noise, even in the presence of initial correlations.

2.3.3 Superchannels in the Wild

To demonstrate the use of the superchannel in practice, consider the evolution of a single photonic qubit, coupled to, and correlated with an environment, see Fig. 2.5. The experimenter aims to implement the target system evolution described by the unitary operator U_s , chosen as either a Pauli-Z gate ($U_s = Z$), a Hadamard-gate ($U_s = H = R_Y Z R_Y^\dagger$), or a rotation ($U_s = Z R_Y$), where R_Y denotes a $\pi/4$ -rotation around σ_y . Due to coupling to the environment the reduced dynamics of the system will in general deviate from that described by U_s . This can be simulated by replacing the Z operations in the above decomposition of U_s by controlled Z (CZ) operations, switched on and off conditional on the state of the environment, which is modelled as another photonic qubit². In the case of Z and H the environment might thus cause a failure of the system unitary (i.e. the identity operation is implemented), while in the case $U_s = Z R_Y$ it can introduce a phase error.




Figure 2.5: **Experimental setup.** System and environment photons are created in the state ρ_{SE} with controllable degree of entanglement, using the source of Ref. [46]. Arbitrary preparations \mathcal{P}_{ij} on the system and measurements $\{M_j\}$ are implemented by means of polarizers (PBS), quarter- and half-wave plates (QWP, HWP) and single-photon detectors (APD). The joint system-environment evolution U is implemented as a CZ gate between a set of HWPs and QWPs. In the case of no initial correlations this setup implements the target system evolution U_s . The CZ gate is based on non-classical interference at a partially polarizing beam splitter (PPBS) with reflectivities of $r_H = 0$ ($r_V = 2/3$) for horizontally (vertically) polarized light [47].

The initial system-environment state was generated via spontaneous parametric down conversion in the form

$$|\psi\rangle_{SE} = \cos(2\theta)|H\rangle_S|V\rangle_E + \sin(2\theta)|V\rangle_S|H\rangle_E, \quad (2.19)$$

where $|H\rangle, |V\rangle$ correspond to horizontally and vertically polarized photons respectively. In this case the strength of the initial correlations (both quantum and classical) is parametrized by

²This is sufficient to describe a large range of joint system-environment dynamics including common error channels, and to illustrate the technique [40, 43], although a slightly larger environment would be required in the most general case [44, 45].

the tangle $\tau = \sin^2(4\theta)$ and can be tuned from uncorrelated ($\theta = 0$) to maximal correlation ($\theta = \pi/8$) [46]. Specifically, initial states with $\tau = \{0.012, 0.136, 0.423, 0.757, 0.908\}$ were generated, with an average fidelity of $F = 0.96(1)$ with the corresponding ideal state. The system was then subjected to the preparation procedure P_{ij} , which prepared it in the state ρ_j by first projecting onto the state ρ_i followed by a unitary rotation.

Using a set of informationally complete sets $\{|H\rangle, |V\rangle, |D\rangle, |A\rangle, |R\rangle, |L\rangle\}$ the superchannel Choi matrix $\Lambda_{\mathcal{M}}$ can be reconstructed using maximum likelihood estimation, as outlined in Sec. 2.3.2. In the case of vanishing initial correlations, the superchannel factorizes into the density matrix of the effective initial system state and the effective system channel $\Lambda_{\mathcal{M}} = \rho_s \otimes \Lambda_{\mathcal{E}}$ [34]. Hence, to allow for an operational interpretation of $\Lambda_{\mathcal{M}}$, is best written using the state basis for the index corresponding to the effective initial state, and the Pauli basis for the indices corresponding to the effective channel. Figure 2.6b shows the results of maximum likelihood quantum process tomography for the case $U_s = H$, with different choices of preparation procedures that, in the presence of initial correlations, result in vastly different reconstructed channels. The superchannel \mathcal{M} in Fig. 2.6a clearly illustrates the reason for this discrepancy: a term that corresponds to the identity operation and increases with the strength of initial correlations. This is exactly the simulated environment-induced failure mode of the system evolution.

2.3.4 Quantifying Initial Correlations

Initial system-environment correlations reveal themselves through their effect on the system evolution for different preparation procedures. To quantify these effects, one can define an *average initial system state* $\rho_{S,\text{av}} = \text{Tr}_{23}[\Lambda_{\mathcal{M}}]/d$ and an *average effective map* for the evolution of the system as $\Lambda_{\mathcal{E},\text{av}} = \text{Tr}_1[\Lambda_{\mathcal{M}}]$. Recall that for a product initial state ($\rho_{\text{SE}} = \rho_s \otimes \rho_e$) the map \mathcal{M} takes the product form $\Lambda_{\mathcal{M}} = \rho_s \otimes \Lambda_{\mathcal{E}}$. In this case $\rho_{S,\text{av}} = \rho_s$, and $\Lambda_{\mathcal{E},\text{av}} = \Lambda_{\mathcal{E}}$ is the Choi matrix of the channel \mathcal{E} describing the (noisy) evolution of the system alone—the same as would result from conventional quantum process tomography. For a given \mathcal{M} one can now define the corresponding separable superchannel \mathcal{M}_s via $\Lambda_{\mathcal{M}_s} = (\rho_{S,\text{av}} \otimes \Lambda_{\mathcal{E},\text{av}})$. In general $\mathcal{M} \neq \mathcal{M}_s$ and the distance between \mathcal{M} and \mathcal{M}_s can be used to quantify the strength of the initial system-environment correlations. This distance is quantified by the so-called *initial correlation norm*:

$$\|\mathcal{M}\|_{\text{IC}} = \frac{1}{2} \|\mathcal{M} - \mathcal{M}_s\|_{\diamond}. \quad (2.20)$$

The matrix $\mathcal{M} - \mathcal{M}_s$ was introduced as *correlation memory matrix* in Ref. [34] since it describes how the dynamics is affected by initial correlations. The choice of the diamond norm $\|\cdot\|_{\diamond}$ [48] allows for an operational interpretation of the initial correlation norm in terms of channel discrimination [49]. For any two quantum channels $\mathcal{E}_1, \mathcal{E}_2$, the best single shot strategy for deciding if a given channel is \mathcal{E}_1 or \mathcal{E}_2 succeeds with probability $\frac{1}{2} (1 + \frac{1}{2} \|\mathcal{E}_1 - \mathcal{E}_2\|_{\diamond})$. Thus, when $\|\mathcal{M}\|_{\text{IC}} = 0$ there is no operational difference between \mathcal{M} and \mathcal{M}_s , which means that there are no observable system-environment correlations. This can either mean that the initial system-

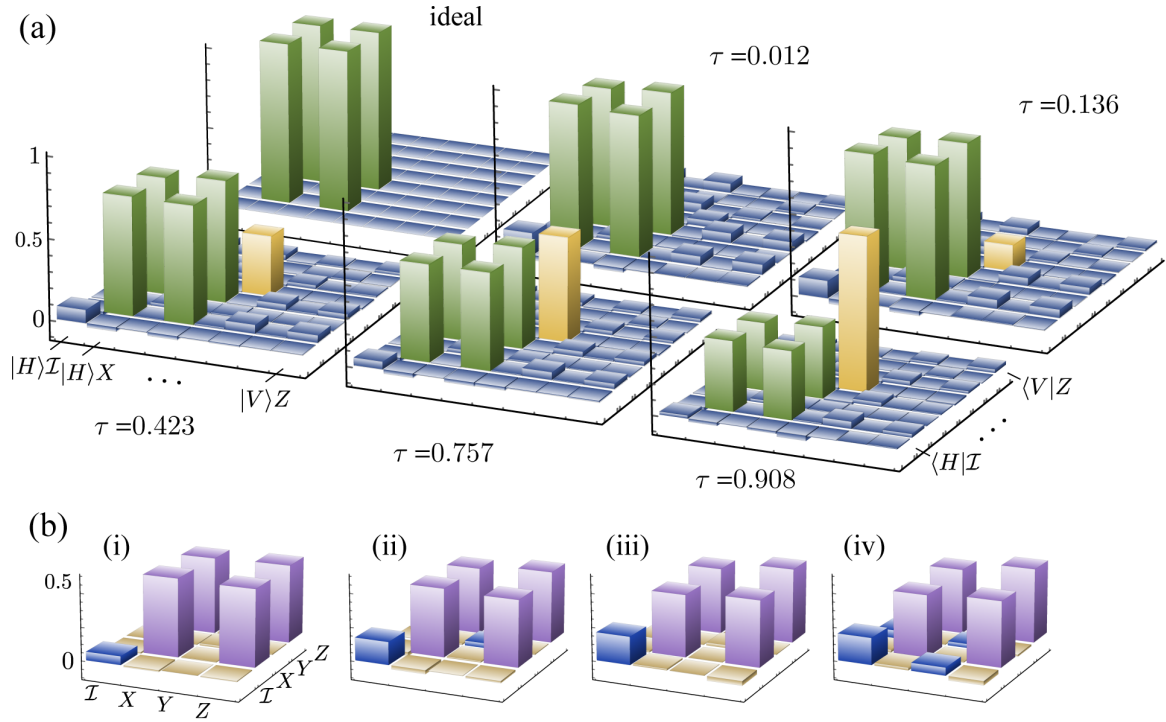


Figure 2.6: **Reconstructed superchannel Choi matrix and effective channels from quantum process tomography.** (a) Real parts of $\Lambda_{\mathcal{M}}$ for $U_s = H$ in the ideal, uncorrelated case and experimental results for increasing strength of initial correlations. The matrices $\Lambda_{\mathcal{M}}$ are shown in a polarization-Pauli basis, with the elements from left to right corresponding to $\{|H\rangle, |V\rangle\} \otimes \{\mathcal{I}, X, Y, Z\}$ and from front to back corresponding to $\{\langle H|, \langle V|\} \otimes \{\mathcal{I}, X, Y, Z\}$. The emergence of a peak corresponding to the identity operation (shown in yellow) is characteristic for the simulated increased tendency of the single-qubit operation U_s (shown in green) to fail in the presence of stronger initial correlations. The negligible imaginary parts are not shown. (b) Real parts of the Choi-matrices (shown in the Pauli basis) for U_s obtained via quantum process tomography for different choices of preparation procedure in the case of low initial correlation $\tau = 0.136$. Cases (i) and (ii) correspond to a fixed ρ_k in Fig. 2.3b, (iii) corresponds to $\rho_k = \rho_i$, and (iv) is the case where $1 \leq k \leq 4$. The information contained in the superchannel \mathcal{M} can be used to identify the optimal preparation procedure.

environment state is indeed uncorrelated, or that the environment is Markovian and initial correlations do not affect the subsequent dynamics. The initial correlation norm thus provides a necessary and sufficient condition for the decoupling of the future state of the system from its past interactions with the environment. When $\|\mathcal{M}\|_{\text{IC}} > 0$ there exists an optimal preparation procedure that can be used as a witness for initial correlations, and the specific value of the norm determines the single shot probability of success for this witness. For a more detailed discussion of the properties of this measure of initial correlation the reader is referred to Ref. [5].

Figure 2.7 shows the measured values of $\|\mathcal{M}\|_{\text{IC}}$ for the three system-environment interactions discussed in Sec. 2.3.3 for a range of simulated initial system-environment states. For all interactions, the maximum obtained value of $\|\mathcal{M}\|_{\text{IC}}$ is ~ 0.5 . This is in agreement with theoretical expectations, since for a maximally correlated initial state the simulated system-

environment coupling would cause a failure of the evolution with probability $1/2$.

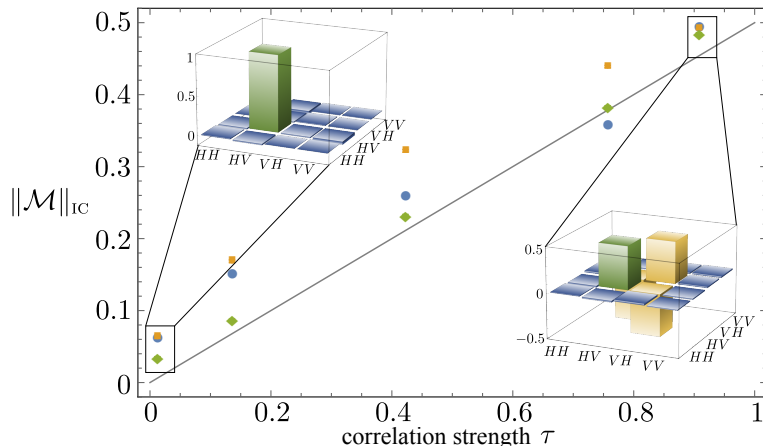


Figure 2.7: **Initial correlation norm.** Shown is the initial correlation norm $\|\mathcal{M}\|_{\text{IC}}$ vs correlation strength τ of ρ_{SE} for $U = \sigma_z$ (blue circles), $U = H$ (yellow squares) and $U = R_Y$ (green diamonds). The values of τ were obtained from state tomography of ρ_{SE} for each experiment. The measured real parts of the states with weakest and strongest initial correlations are shown in the respective insets. The solid line corresponds to the initial correlation norm in the ideal case. Error bars from Poissonian counting statistics are on the order of the symbol size.

2.3.5 Preparation Fidelity

The information contained in \mathcal{M} can also be used to choose a set of preparation procedures that optimize the impact of the environment. Consider a system preparation via initial post-selection on the state ρ_1 . The subsequent evolution is then described by the effective map

$$\Lambda_{\mathcal{E}_{\rho_1}} = \frac{1}{p_{\rho_1}} \text{Tr}_1 [(\rho_1^\dagger \otimes \mathbb{1}_{23}) \Lambda_{\mathcal{M}}] , \quad (2.21)$$

where $p_{\rho_1} = \text{Tr} [(\rho_1^\dagger \otimes \mathbb{1}_{23}) \Lambda_{\mathcal{M}}] / d$ is the probability of success for the post-selection on ρ_1 . Studying the effective maps in Eq.(2.21) for different ρ_1 , one can optimize the preparation procedure for any desired evolution of the system. A figure of merit for this optimization is given by the *preparation fidelity* F_{prep} , which measures the fidelity between the implemented effective map \mathcal{E}_{ρ_1} and the desired target channel \mathcal{U}_s for initial projection onto ρ_1 ,

$$F_{\text{prep}}(\mathcal{M}, \rho_1, U_s) = \frac{1}{d^2} F(\Lambda_{\mathcal{E}_{\rho_1}}, \Lambda_{U_s}) . \quad (2.22)$$

Maximizing F_{prep} over all states ρ_1 for a given target unitary U_s finds the preparation which comes closest to the desired U_s . Curiously, this is, in general, not equivalent to minimizing the impact of the environment, since the optimal preparation might harness some of the environmental correlations to improve the gate performance. Figure 2.8 a and b, show the preparation fidelity with a nominal $U_s = Z$ and $U_s = R_Y Z$, respectively, calculated from an experimentally reconstructed superchannel \mathcal{M} with weak initial correlations of $\|\mathcal{M}\|_{\text{IC}} = 0.062(5)$ and




Figure 2.8: **Optimization of the preparation fidelity.** The preparation fidelity $F_{\text{prep}}(\mathcal{M}, \rho_1, Z)$ for (a) a target $U_s = Z$ and (b) a target $U_s = R_Y Z$ is shown as a density plot on the surface of the Bloch sphere of the initial-projection state ρ_1 . In both cases, the initial correlations are very weak, but the preparation procedure can nonetheless be used to optimize the implementation of the system evolution.

$\|\mathcal{M}\|_{\text{IC}} = 0.034(2)$, respectively. The effect of the environment is minimized for a state that is significantly different to the standard preparation $|H\rangle$, leading to an improvement in fidelity by 0.2%. This demonstrates that, even in the regime of almost uncorrelated initial system-environment states, the information contained in the superchannel can be used to improve the implemented evolution.

Alternatively, one could consider minimizing F_{prep} to find the worst-case preparation, which gives insight into where and why the experimental setup fails. Furthermore, the preparation fidelity defined in Eq. (2.22) is just one example, and one could consider other figures of merit depending on the specific scenario.

2.3.6 Discussion

The superchannel approach offers an operationally motivated and experimentally accessible way of fully characterizing the reduced dynamics of a quantum system that is coupled to an environment, even in the presence of initial correlations. As a direct generalization of quantum process tomography, all the tools developed to improve the efficiency of the latter, such as compressive sensing [50, 51] can also be applied to the reconstruction of \mathcal{M} . Similar to process tomography, however, the reconstruction of the superchannel relies on well-characterized preparations and measurements to produce reliable estimates. In contrast to the reduced system evolution, however, the superchannel is always completely positive, and the use of maximum likelihood reconstruction is therefore justified, even in the presence of initial correlations.

On the practical side, the superchannel contains information about the initial correlations and how the environment couples to the evolution of the system. It can thus be exploited to improve the performance of the experiment, even in the regime of very weak initial correlations, as demonstrated in Sec. 2.3.5. However, since in the limit of vanishing correlations this approach essentially reduced to quantum process tomography, it is most useful in quantum architectures which are strongly coupled to their environment, such as spins in local spin baths. Another application is in quantum control, where control timescales can be much faster than

environmental reset times.

On the fundamental side, the superchannel allows for the study of non-Markovian quantum processes, and two two-point correlation example here could easily be generalized to a multi-point scenario [52]. It has been suggested that such non-Markovianity could be used as a resource [53], and the superchannel approach has also been used to derive the lower bound on entropy production in a generic quantum process [54]. Interestingly, the superchannel is closely related to recent generalizations of the process matrix in the study of the causal structure of quantum mechanics [52]. In particular, experiments such as the one presented in Sec. 2.3.3 are excellent candidates for simulating causally non-separable processes, where two operations are implemented in a superposition of causal order [55].

References

- [1] Ringbauer, M., Wood, C. J., Modi, K., Gilchrist, A., White, A. G. & Fedrizzi, A. Characterizing Quantum Dynamics with Initial System-Environment Correlations. *Phys. Rev. Lett.* **114**, 090402 (2015).
- [2] Rahimi-Keshari, S., Broome, M. A., Fickler, R., Fedrizzi, A., Ralph, T. C. & White, A. G. Direct characterization of linear-optical networks. *Opt. Express* **21**, 13450–13458 (2013).
- [3] Durt, T., Englert, B.-G., Bengtsson, I. & Życzkowski, K. On mutually unbiased bases. *Int. J. Quantum Inf.* **8**, 535–640 (2010).
- [4] Renes, J. M., Blume-Kohout, R., Scott, A. J. & Caves, C. M. Symmetric informationally complete quantum measurements. *J. Math. Phys.* **45**, 2171–2180 (2004).
- [5] Wood, C. J. *Initialization and Characterization of Open Quantum Systems*. Ph.D. thesis, University of Waterloo (2015).
- [6] Schwemmer, C., Knips, L., Richart, D., Weinfurter, H., Moroder, T., Kleinmann, M. & Gühne, O. Systematic Errors in Current Quantum State Tomography Tools. *Phys. Rev. Lett.* **114**, 080403 (2015).
- [7] Shang, J., Ng, H. K. & Englert, B.-G. Quantum state tomography: Mean squared error matters, bias does not. *arXiv:1405.5350* (2014).
- [8] Hradil, Z. Quantum-state estimation. *Phys. Rev. A* **55**, R1561–R1564 (1997).
- [9] James, D. F. V., Kwiat, P. G., Munro, W. J. & White, A. G. Measurement of qubits. *Phys. Rev. A* **64**, 052312 (2001).
- [10] Blume-Kohout, R. Optimal, reliable estimation of quantum states. *New J. Phys.* **12**, 043034 (2010).
- [11] Ferrie, C. & Blume-Kohout, R. Estimating the bias of a noisy coin. *AIP Conf. Proc.* **1443**, 14–21 (2012).
- [12] Blume-Kohout, R. Hedged Maximum Likelihood Quantum State Estimation. *Phys. Rev. Lett.* **105**, 200504 (2010).
- [13] Home, J. P., Hanneke, D., Jost, J. D., Amini, J. M., Leibfried, D. & Wineland, D. J. Complete methods set for scalable ion trap quantum information processing. *Science* **325**, 1227–1230 (2009).

- [14] Blume-Kohout, R. Robust error bars for quantum tomography. *arXiv:1202.5270* (2012).
- [15] Knill, E., Leibfried, D., Reichle, R., Britton, J., Blakestad, R., Jost, J., Langer, C., Ozeri, R., Seidelin, S. & Wineland, D. Randomized benchmarking of quantum gates. *Phys. Rev. A* **77**, 012307 (2008).
- [16] Wallman, J. J. & Flammia, S. T. Randomized benchmarking with confidence. *New J. Phys.* **16**, 103032 (2014).
- [17] Blume-Kohout, R., Gamble, J. K., Nielsen, E., Mizrahi, J., Sterk, J. D. & Maunz, P. Robust, self-consistent, closed-form tomography of quantum logic gates on a trapped ion qubit. *arXiv:1310.4492* (2013).
- [18] Blume-Kohout, R. private communication.
- [19] Plenio, M., Huelga, S., Beige, A. & Knight, P. Cavity-loss-induced generation of entangled atoms. *Phys. Rev. A* **59**, 2468–2475 (1999).
- [20] Bose, S., Knight, P., Plenio, M. & Vedral, V. Proposal for Teleportation of an Atomic State via Cavity Decay. *Phys. Rev. Lett.* **83**, 5158–5161 (1999).
- [21] Beige, A., Braun, D., Tregenna, B. & Knight, P. L. Quantum Computing Using Dissipation to Remain in a Decoherence-Free Subspace. *Phys. Rev. Lett.* **85**, 1762–1765 (2000).
- [22] Diehl, S., Micheli, A., Kantian, A., Kraus, B., Büchler, H. P. & Zoller, P. Quantum states and phases in driven open quantum systems with cold atoms. *Nat. Phys.* **4**, 878–883 (2008).
- [23] Barreiro, J. T., Müller, M., Schindler, P., Nigg, D., Monz, T., Chwalla, M., Hennrich, M., Roos, C. F., Zoller, P. & Blatt, R. An open-system quantum simulator with trapped ions. *Nature* **470**, 486–491 (2011).
- [24] Cormick, C., Bermudez, A., Huelga, S. F. & Plenio, M. B. Dissipative ground-state preparation of a spin chain by a structured environment. *New J. Phys.* **15**, 073027 (2013).
- [25] Lin, Y., Gaebler, J. P., Reiter, F., Tan, T. R., Bowler, R., Sørensen, A. S., Leibfried, D. & Wineland, D. J. Dissipative production of a maximally entangled steady state of two quantum bits. *Nature* **504**, 415–418 (2013).
- [26] Xu, J.-S., Sun, K., Li, C.-F., Xu, X.-Y., Guo, G.-C., Andersson, E., Lo Franco, R. & Compagno, G. Experimental recovery of quantum correlations in absence of system-environment back-action. *Nat. Comms.* **4**, 2851 (2013).
- [27] Petruccione, F. & Breuer, H.-P. *The theory of open quantum systems* (Oxford Univ. Press, 2002).

- [28] Modi, K. & Sudarshan, E. C. G. Role of preparation in quantum process tomography. *Phys. Rev. A* **81**, 052119 (2010).
- [29] Rodríguez-Rosario, C. A., Modi, K., Kuah, A.-M., Shaji, A. & Sudarshan, E. C. G. Completely positive maps and classical correlations. *J. Phys. A* **41**, 205301 (2008).
- [30] Modi, K. Preparation of states in open quantum mechanics. *Open Syst. Inf. Dyn.* **18**, 253–260 (2011).
- [31] Kuah, A.-M., Modi, K., Rodríguez-Rosario, C. A. & Sudarshan, E. C. G. How state preparation can affect a quantum experiment: Quantum process tomography for open systems. *Phys. Rev. A* **76**, 042113 (2007).
- [32] Ziman, M. Quantum process tomography: the role of initial correlations. *arXiv:quant-ph/0603166* (2006).
- [33] Carteret, H. A., Terno, D. R. & Życzkowski, K. Dynamics beyond completely positive maps: Some properties and applications. *Phys. Rev. A* **77** (2008).
- [34] Modi, K. Operational approach to open dynamics and quantifying initial correlations. *Sci. Rep.* **2**, 581 (2012).
- [35] Wißmann, S., Leggio, B. & Breuer, H.-P. Detecting initial system-environment correlations: Performance of various distance measures for quantum states. *Phys. Rev. A* **88**, 022108 (2013).
- [36] Laine, E.-M., Piilo, J. & Breuer, H.-P. Witness for initial system-environment correlations in open-system dynamics. *Europhys. Lett.* **92**, 60010 (2010).
- [37] Rodríguez-Rosario, C. A., Modi, K., Mazzola, L. & Aspuru-Guzik, A. Unification of witnessing initial system-environment correlations and witnessing non-Markovianity. *Europhys. Lett.* **99**, 20010 (2012).
- [38] Gessner, M. & Breuer, H.-P. Detecting Nonclassical System-Environment Correlations by Local Operations. *Phys. Rev. Lett.* **107**, 180402 (2011).
- [39] Rossatto, D. Z., Werlang, T., Castelano, L. K., Villas-Boas, C. J. & Fanchini, F. F. Purity as a witness for initial system-environment correlations in open-system dynamics. *Phys. Rev. A* **84**, 042113 (2011).
- [40] Smirne, A., Brivio, D., Cialdi, S., Vacchini, B. & Paris, M. G. A. Experimental investigation of initial system-environment correlations via trace-distance evolution. *Phys. Rev. A* **84**, 032112 (2011).
- [41] Li, C.-F., Tang, J.-S., Li, Y.-L. & Guo, G.-C. Experimentally witnessing the initial correlation between an open quantum system and its environment. *Phys. Rev. A* **83**, 064102 (2011).

- [42] Gessner, M., Ramm, M., Pruttivarasin, T., Buchleitner, A., Breuer, H.-P. & Häffner, H. Local detection of quantum correlations with a single trapped ion. *Nat. Phys.* **10**, 105–109 (2013).
- [43] Chiuri, A., Greganti, C., Mazzola, L., Paternostro, M. & Mataloni, P. Linear Optics Simulation of Non-Markovian Quantum Dynamics. *Phys. Rev. A* **86**, 010102 (2012).
- [44] Schumacher, B. Sending quantum entanglement through noisy channels. *Phys. Rev. A* **54**, 2614 (1996).
- [45] Narang, G. Simulating a single-qubit channel using a mixed-state environment. *Phys. Rev. A* **75**, 032305 (2007).
- [46] Fedrizzi, A., Herbst, T., Poppe, A., Jennewein, T. & Zeilinger, A. A wavelength-tunable fiber-coupled source of narrowband entangled photons. *Opt. Express* **15**, 15377–15386 (2007).
- [47] Langford, N., Weinhold, T., Prevedel, R., Resch, K., Gilchrist, A., O’Brien, J., Pryde, G. & White, A. Demonstration of a Simple Entangling Optical Gate and Its Use in Bell-State Analysis. *Phys. Rev. Lett.* **95**, 210504 (2005).
- [48] Kitaev, A. Y. Quantum computations: algorithms and error correction. *Russ. Math. Surv.* **52**, 1191 (1997).
- [49] Rosgen, B. & Watrous, J. On the Hardness of Distinguishing Mixed-State Quantum Computations. In *20th Annual IEEE Conference on Computational Complexity (CCC’05)*, 344–354 (IEEE, 2005).
- [50] Shabani, A., Kosut, R. L., Mohseni, M., Rabitz, H., Broome, M. A., Almeida, M. P., Fedrizzi, A. & White, A. G. Efficient Measurement of Quantum Dynamics via Compressive Sensing. *Phys. Rev. Lett.* **106**, 100401 (2011).
- [51] Flammia, S. T., Gross, D., Liu, Y.-K. & Eisert, J. Quantum tomography via compressed sensing: Error bounds, sample complexity, and efficient estimators. *New J. Phys.* **14**, 095022 (2012).
- [52] Pollock, F. a., Rodríguez-Rosario, C., Frauenheim, T., Paternostro, M. & Modi, K. Complete framework for efficient characterisation of non-Markovian processes. *arXiv:1512.00589* (2015).
- [53] Bylicka, B., Chruscinski, D. & Maniscalco, S. Non-markovianity as a resource for quantum technologies. *arXiv:1301.2585* (2013).
- [54] Vinjanampathy, S. & Modi, K. Second law for quantum operations. *arXiv:1405.6140* (2014).
- [55] Oreshkov, O., Costa, F. & Brukner, C. Quantum correlations with no causal order. *Nat. Commun.* **3**, 1092 (2012).

CHAPTER 3

Introduction to Quantum Foundations

THIS section introduces a few concepts that are central not just to quantum foundational research, but also for practical applications in quantum information theory. Many of these concepts will also be important in the later chapters of this thesis. After a short discussion of the probabilistic nature of quantum mechanics, I will focus on the non-classical nature of quantum correlations in Sec. 3.2. After discussing the the famous Einstein-Podolsky-Rosen paradox and Bell’s theorem, I will briefly touch upon the topic of correlations that are stronger than those of quantum mechanics, which includes original results that are included in full in Appendix A. I will also discuss how correlation polytopes can be used to study sets of correlations and derive testable inequalities from them. In Sec. 3.3 I discuss contextuality, which is a distinguishing feature of quantum mechanics over classical physics. Section 3.4 focuses on paradoxes and games that act as primitives for fundamental questions as well as practical quantum information protocols. Finally, Sec. 3.5 discusses pre- and post-selection paradoxes, their foundational importance, and their connection to weak values.

3.1 Probability and Randomness

Classical physics is deterministic such that sufficient knowledge of the current state of the system and its environment allows one to uniquely predict the future state of the system and the outcome of any measurement. The measurement statistics predicted by the Born rule of quantum mechanics, on the other hand, are inherently probabilistic. In this sense, quantum theory is often treated as a generalization of classical probability theory, both of which are special cases in the framework of generalized probabilistic theories [1–4]. Which physical principle singles out quantum mechanics within these theories, however, is still a topic of active research [5], see Appendix A.

The origin of the probabilities in quantum theory itself is part of the so-called *measurement problem* and one of the central motivations for the study of interpretations of quantum mechanics. In classical mechanics, probabilistic behaviour arises from a lack of knowledge about the actual state of the system. Whether the same is true in the quantum case is an important open question, see Chap. 4. However, although modifications such as Bohmian mechanics provide a deterministic description at an underlying, unobservable level, there are nonetheless observable events which are fundamentally unpredictable [6].

This unpredictability is particularly interesting for the computer science community, where random numbers are crucial for security applications. Classically, random numbers are *pseudo random*, which means that they are based on some complicated, but deterministic algorithm. In other words, they are random but not unpredictable, and security based on these numbers could be compromised if an adversary discovers the algorithm behind them. Quantum mechanics, on the other hand, can guarantee unpredictability of the random numbers and thus in principle enables cryptographic security certified by physics.

Curiously, the interpretation of probability itself has an interesting role to play with respect to the interpretation of quantum theory. The two main interpretations of probability are the objective *Frequentist* interpretation, which sees probabilities as a description of the statistical distribution of objective properties of an ensemble of systems, and the subjective *Bayesian* point of view, which treats probabilities as degrees of belief of an agent. Naturally, the Frequentist needs to rely on experiments to determine the inherent probability of an event, while the Bayesian starts with a prior probability distribution that reflects the agent's beliefs before the experiment, and updates this distribution in the light of new experimental evidence. Most importantly, in the Bayesian approach there is no need to assume that quantum systems have objective properties. Giving up this notion of objective reality, however, we are faced with the problem of explaining how two observers can assign consistent probabilities to an experiment, see Chap 4 and Ref. [7]

3.2 Quantum Correlations

As discussed in Sec. 1.1.4, it is in general not possible to decompose the state of a composite system into definite states of the subsystems. Entangled quantum systems remain intimately connected no matter how far apart they are, and feature correlations so strong that they cannot be explained in terms of classical cause-and-effect relations, see Chap. 5 for more details. Such strong correlations seem to be in conflict with relativity, which has led Einstein, Podolsky and Rosen [8] in 1935 to suggest that quantum mechanics must be missing something in its description of the world [8].

3.2.1 EPR Paradox

In a seminal 1935 paper Einstein, Podolsky and Rosen [8] discussed a thought experiment involving a pair of particles correlated in their position degrees of freedom and with vanishing total momentum. In the following we will discuss a version of the experiment involving qubits which was first proposed by Bohm for spin-1/2 particles [9]. Consider two particles in a singlet state, $|\Psi\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$, shared between two spacelike separated¹ parties, Alice and Bob. Alice can then choose to measure her particle in either the $\{|0\rangle, |1\rangle\}$ -basis or the $\{|+\rangle, |-\rangle\}$ -basis.

¹Two events are spacelike separated in relativistic spacetime, if any signal travelling from one event to the other would have to travel faster than light to reach its destination in time. As a consequence, neither event can act as a cause of the other (if we assume that causes propagate at most at the speed of light), see Chap. 5.

The singlet state has the special property that whatever measurement outcome Alice gets, Bob’s system will be projected into the orthogonal state. Hence, if Alice measures in $\{|0\rangle, |1\rangle\}$ and obtains outcome $|0\rangle$, Bob’s system will be in the state $|1\rangle$, and similarly for other measurement choices and outcomes. In summary, for one of Alice’s measurement choices Bob’s state is either $|0\rangle$ or $|1\rangle$, while for the other choice it is either $|+\rangle$ or $|-\rangle$. Hence, Alice can *steer* Bob’s quantum state with her measurement choice [10].

In a model where the quantum state provides a complete description of physical reality (i.e. a ψ -complete ontological model in the language of Chap. 4) this steering phenomenon implies that Alice’s measurement choice directly influences the state of reality (or *ontic* state) of Bob’s system. Einstein famously described this influence as “spooky action at a distance” [11], since it must be instantaneous and not constrained by the fundamental speed limit of relativity. Notably, however, there is no explicit conflict with special relativity, whose central premise is that no *information* can be transmitted faster than light. Since Alice’s measurement outcome is random, so is Bob’s state until he learns about Alice’s measurement outcome. Hence, signal locality is satisfied at the observational level, despite potential superluminal influences at the level of an underlying, unobservable, reality. The conflict between the assumption that the wavefunction provides a complete description of physical reality, and that there is no action at a distance is known as the *EPR paradox*. If Alice can steer Bob’s state, then it is not possible to describe the observed correlations with a so-called *local hidden state model*, which ascribes a definite (hidden) state to Bob’s system in every run of the experiment [12]. This led EPR to suggest that quantum mechanics is incomplete and that the wavefunction must be supplemented by more general additional variables—a local hidden variable model.

Besides its foundational importance, quantum steering has found application in the context of quantum cryptography. Consider a scenario where Alice and Bob want to certify that they share entanglement. Bob trusts that his measurement device operates according to quantum mechanics, but Alice does not know anything about the workings of her device, nor do they trust that the source really produces entangled quantum states. By demonstrating Alice’s ability to steer Bob’s state, that is by violating a steering inequality, they can nonetheless establish that their observations cannot be explained in terms of a definite quantum state of Bob’s system together with appropriately fabricated results of Alice’s device [12]. In other words, Alice and Bob can establish that they share an entangled state and limit the amount of information that a potential eavesdropper could have about their measurement results [13]. This can be used to perform so-called *one-sided device-independent quantum key distribution*, without knowing anything about the workings of Alice’s measurement device or the source of particles.

3.2.2 Bell’s Theorem

The specific correlations arising in the EPR paradox discussed above², where Alice and Bob only measure in either the $\{|0\rangle, |1\rangle\}$ -basis or the $\{|+\rangle, |-\rangle\}$ -basis can indeed be described by a

²The position-momentum correlations observed in the original EPR paradox can also be reproduced by a local hidden variable model, since the state and measurements considered there admit a classical description[14].

local hidden variable model [15–17]. However, John Bell showed that this is not true in general for correlations arising from local measurements on entangled quantum states. Specifically, Bell derived an inequality that must be satisfied by any hidden variable model that obeys some physically motivated assumptions about locality and how classical cause-and-effect relations work [18, 19], see Chap. 5. Quantum correlations, however, violate this inequality and can thus not be reproduced by what have become known as *Bell-local* hidden variable models.

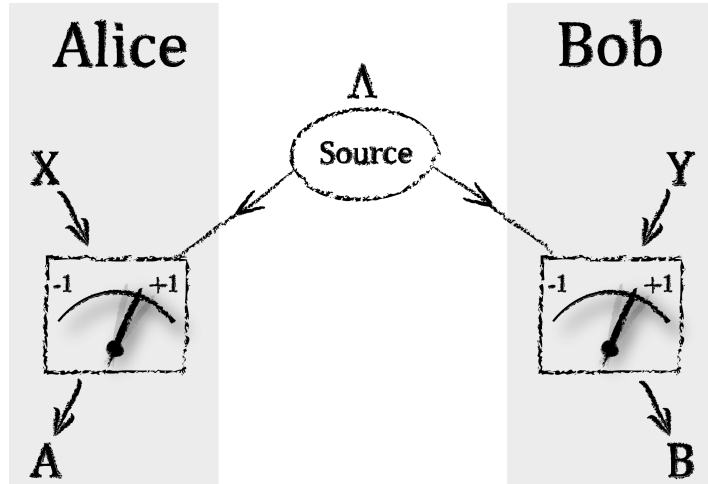


Figure 3.1: **The scenario considered by Bell.** Two parties, Alice and Bob, each perform one of two local measurements on one half of a shared quantum state and obtain one of two possible outcomes. Their measurement choices are represented by the variables X and Y , respectively and the measurement outcomes by A and B , respectively. The measurement outcomes may additionally be influenced by a hidden variable Λ , which need not originate at the source.

In the scenario considered by Bell, Alice and Bob share a quantum state of two qubits. They can each choose to perform one of two measurements on their system, $x \in \{0, 1\}$ for Alice and $y \in \{0, 1\}$ for Bob, and each obtain one of two outcomes, $a \in \{\pm 1\}$ for Alice and $b \in \{\pm 1\}$ for Bob. In a Bell-local hidden variable model Alice and Bob can freely choose their measurement settings, but they are not allowed to communicate, which is typically justified by spacelike separating them. These models are characterized by conditional probability distributions of the form

$$P(a, b|x, y) = \sum_{\lambda} P(a|x, \lambda)P(b|y, \lambda)P(\lambda) , \quad (3.1)$$

where Λ is a hidden variable that is not observable. Loosely speaking, a Bell-local hidden variable model is one in which Alice and Bob try to reproduce the correlations of entangled quantum systems using a pre-arranged strategy and a shared list of random numbers (corresponding to λ), but are not allowed to communicate (represented by the factorization in Eq. (3.1)). All correlations obtained in such a way must satisfy the Clauser-Horne-Shimony-Holt (CHSH) inequality [20]

$$S_2 = \langle A_0B_0 \rangle + \langle A_0B_1 \rangle + \langle A_1B_0 \rangle - \langle A_1B_1 \rangle \leq 2 , \quad (3.2)$$

where $\langle A_xB_y \rangle = \sum_{a,b} abP(a, b|x, y)$ denotes the joint expectation value for the product of outcomes of Alice and Bob for input x and y , respectively. $P(a, b|x, y)$ is the conditional probability

for outcomes a and b , given the settings were x and y , respectively. Correlations obtained from measurements on certain *Bell-nonlocal* entangled quantum states, on the other hand, can violate the inequality up to the so-called *Tsirelson bound* $S_2 = 2\sqrt{2}$.

Similar to steering, Bell inequalities have found application in quantum cryptography. By violating the CHSH inequality Alice and Bob can convince a third party, Charlie, that they share entanglement, even if Charlie does not trust either of them. In other words, a violation of the CHSH inequality allows for a *device-independent* verification of entanglement, without any assumptions about the workings of the used devices. This is very relevant in cryptographic scenarios, where it allows Alice and Bob to establish a secure connection with untrusted devices. In fact, as we show in Chap. 5, security can be established even if Alice’s measurement outcome is leaked.

3.2.3 A Hierarchy of Correlations

When considering mixed states the structure of correlations becomes quite interesting, even for a pair of qubits, see Fig. 3.2. Not every entangled state can be used for steering, and not every steerable state can be used for teleportation [12]. Furthermore, not every un-entangled state is necessarily classical. In fact, just as the classical states of a single qubit lie on a line in the Bloch-sphere, the classical states of two qubits are a negligible part of the set of two-qubit states (a set of measure zero) [21]. All other states exhibit quantum discord which, loosely speaking, means that the state cannot be fully determined without disturbing it. For pure states, discord is equal to entanglement, but this is not true anymore for mixed states, where entangled states are a subset of discordant states, see Fig. 3.2. Discord is a valuable resource by itself and has been linked to advantages in some computational tasks [22], and to remote state preparation, where Alice tries to prepare a state at Bob’s side [23]. The interested reader is referred to Ref. [24] for a comprehensive review on quantum discord.

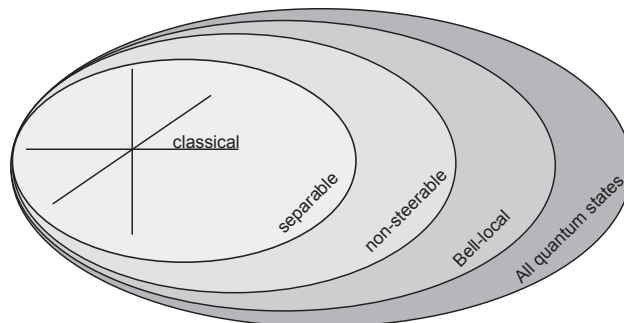


Figure 3.2: **Hierarchy of correlations.** Classically correlated states form a negligible (i.e. measure-zero) subset of the space of quantum states (indicated by black lines). All other states feature quantum discord (indicated by the gray shading) [25]. As the shading gets darker, states get more powerful, with Bell-nonlocal states being the most powerful.

3.2.4 Superquantum Correlations

Curiously, quantum correlations are stronger than any classical correlations, yet they are not the strongest correlations that respect signal locality. Instead, the so-called Popescu-Rohrlich (PR) box, which satisfies $a \oplus b = xy$ (where \oplus denotes addition modulo 2), violates the CHSH inequality (3.2) up to the algebraic maximum of 4 and still remains compatible with signal locality [26]. Moreover, theories which allow for such strong correlations exhibit many features present in quantum mechanics [27]. These include effects that are common to theories with a restriction on knowledge [16], such as no-cloning [26] and uncertainty relations [28], as well as effects that are considered genuinely non-classical, such as monogamy of correlations [28] and nonlocality swapping [29].

This raises the question which physical principle, if not signal locality, singles out quantum mechanics from this family of more general probabilistic theories. A number of such principles have been proposed [30–37], most of which are of information-theoretic nature. A promising example is information causality [30], which strengthens signal locality by requiring that the amount of information that Bob can gain about an unknown to him data set of Alice cannot be more than the amount of information that Alice communicates to him, see Fig. 3.3. Both classical and quantum theory satisfy information causality, but it is violated by many post-quantum theories [30]. Specifically, the principle recovers Tsirelson’s bound for isotropic correlations [30], and also reproduces the full boundary of bipartite quantum correlations for some two-dimensional slices of the no-signaling polytope (see Sec. 3.2.5), but not for all [27]. This short-coming, however, might only be a result of the used *random access coding* protocol, which is not necessarily optimal. Moreover, a sufficient condition for information causality to hold has so far only been found for one fixed protocol [27], but not in general. Using non-locality distillation [38], the set of correlations ruled out by information causality can also be further extended, which suggests again that the protocol is not optimal in every case [39]. Future generalizations of any principle would also have to take into account that no bipartite principle can hope to recover the quantum set of genuinely multipartite correlations, which may appear classical in any bipartition [40].

The Information Causality Game and Random Access Coding

The principle of information causality is typically formulated as a game and is a special case of *random access coding*, which is one of the central primitives in the study of (quantum) communication complexity [41, 42]. The game is played between Alice and Bob, where Alice holds a list of $N = 2^n$ bits $\{a_i\}_{i=0}^{N-1}$ and Bob receives an index $0 \leq b \leq N - 1$, which labels one of Alice’s bits that he should guess. In addition to using arbitrary shared no-signaling resources, Alice is also allowed to send an m -bit message to Bob. The principle of information causality then states that the information that Bob can gain about Alice’s data must be limited to m bits, as measured by the mutual information I or the success probability of the information

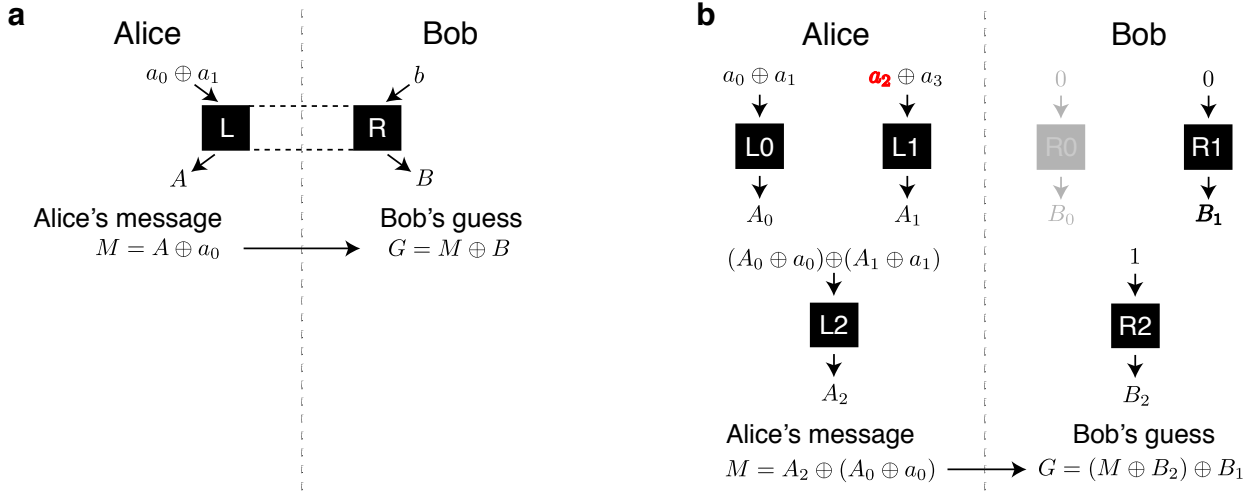


Figure 3.3: Quantum Random Access Coding. A general no-signaling resource is given by a space-like separated (indicated by the dashed line) pair of black boxes producing local outputs A and B for Alice and Bob, when they input a and b , respectively. In the case of a PR-box the outputs of the left (L) and right (R) box would be perfectly correlated according to $A \oplus B = ab$. **(a)** The simplest instance of the information causality game, where Alice has two bits $\{a_0, a_1\}$ and Bob tries to guess the bit a_b . Here Alice inputs $a_0 \oplus a_1$ into her box, obtains output A and sends a message $m = A \oplus a_0$ to Bob. Bob then inputs b into his box and from his output computes his guess $G = M \oplus B$ for Alice's bit. **(b)** Example of the multilevel information causality protocol for $n = 2$. Alice has a list of N bits a_i and Bob tries to guess the bit a_2 (shown in bold, red) using $N-1=3$ pairs of shared black boxes on $n=2$ levels (corresponding boxes labeled $L0/R0, L1/R1, L2/R2$). Alice's inputs at the highest level are similar to the case in (a), while those on the next level(s) contain information about the previous outputs. Bob's inputs b_i and choice of boxes are determined by the binary decomposition $b = \sum_{k=0}^{n-1} b_k 2^k$. From his outputs B_1, B_2 and Alice's 1-bit message M Bob computes a final guess G for Alice's bit a_b . Note that Bob only needs to use one box on each level and ignores the outputs of all the other boxes. Hence, his input to these boxes can be arbitrary and in the experiment presented in Appendix A we chose to use the same input for all boxes on one level.

causality game in Fig. 3.3, see Appendix A for details.

$$I = \sum_{i=0}^{N-1} I_c(a_i; \beta | b = i) \leq m. \quad (3.3)$$

Specifically, consider the case where Alice sends a 1-bit message to Bob, and they are allowed to use arbitrary classical resources (such as a shared list of random numbers), quantum resources (such as shared entangled states), or post-quantum resources (such as shared PR boxes). In the classical case the best possible strategy is “majority vote”, where Alice sends the most frequent bit to Bob [43], which, in the case of random bits, succeeds in 50% of the cases. Using shared entangled states outperforms the classical strategy in a very similar way to the CHSH game [44, 45], but the information gain is still bounded to 1 bit. Using a nested protocol and a shared PR box, however, Alice and Bob can succeed every time, thus violating information

causality. A more detailed experimental study of this principle, using simulated post-quantum correlations can be found in Appendix A.

3.2.5 Correlation Polytopes

The set of classical probability distributions over a set of variables forms a convex polytope, such that any distribution can be written as a weighted sum of the extremal points. Sets of correlations compatible with linear constraints naturally give rise to such polytopes, which turn out to be a powerful tool for the systematic study of the underlying correlations [46]. The aim of this section is to summarize the essential aspects of this approach at the example of Bell-local correlations.

As a simple example, consider a three-dimensional cube with corners $(\pm 1, \pm 1, \pm 1)$, centred around the point $(0, 0, 0)$. The coordinates of any point (x, y, z) within the cube can be decomposed into a weighted sum of the coordinates of the corners. What separates the cube from the rest of three-dimensional space are inequality constraints of the form $-1 \leq x, y, z \leq 1$. In a very similar fashion, the full set of classical probability distributions is a convex polytope, such that any distribution within the set can be written as a weighted sum of the extremal points. The polytope is bounded by the positivity (i.e. $P \geq 0$), and normalization ($P \leq 1$) constraints for classical probabilities. More generally, any boundary of a polytope within a space of equal dimension has the form of an inequality. Boundaries to a space of larger dimension, on the other hand, are formulated as equalities. In the case of the cube, for example, the boundary to the four-dimensional space with coordinates (x, y, z, t) is given by the inequality $t = 0$.

No-Signaling Correlations

In the CHSH scenario, there are two parties, Alice and Bob, with two measurement settings, X and Y , respectively, and two outputs each, A and B , respectively. Hence, full set of conditional probabilities of measurement outcomes given the inputs $P(a, b|x, y)$ forms a 2^4 -dimensional hypercube defined by the constraints $0 \leq P(a, b|x, y) \leq 1$. In addition to these constraints the measurement outcome probabilities must be normalized (i.e. add up to 1) for each combination of measurement settings

$$\sum_{a,b} P(a, b|x, y) = 1 \quad \forall x, y, \quad (3.4)$$

which reduces the number of free parameters by 4. The physical interpretation of the CHSH scenario further requires that the probability distributions satisfy signal locality (or no-signaling), corresponding to the four constraints

$$\begin{aligned} \sum_b P(a, b|x, y) &= \sum_b P(a, b|x, y') & \forall a, x, y, y' \\ \sum_a P(a, b|x, y) &= \sum_a P(a, b|x', y) & \forall b, y, x, x'. \end{aligned} \quad (3.5)$$

This implies that Alice should not be able to gain information about Bob's measurement setting from her local measurement outcome and vice versa. The constraints of Eq. (3.4) and Eq. (3.5)

define the set \mathcal{NS} of no-signaling correlations as a convex 8-dimensional polytope in a 16-dimensional space (since Eq. (3.4) and Eq. (3.5) are 8 equality constraints)

$$\mathcal{NS} = \left\{ \left(\begin{array}{c} P(-1, -1|0, 0) \\ \vdots \\ P(+1, +1|1, 1) \end{array} \right) \mid 0 \leq P(a, b|x, y) \leq 1 \wedge \text{normalization (3.4)} \wedge \text{signal-locality (3.5)} \right\}, \quad (3.6)$$

Each inequality constraint corresponds to a facet of the polytope, and each extremal probability assignment corresponds to a vertex. Crucially, due to the convexity, either facets, or vertices are sufficient to fully specify the polytope. It is typically easier to define the polytope by specifying the facets, which in the case of \mathcal{NS} are simply the 16 positivity constraints. The vertices can then be found by studying the dual polytope, which, loosely speaking, is obtained by turning every vertex into a facet and every facet into a vertex [5, 47]. In practice, this is done using software packages such as PORTA [48]. In the case of the set \mathcal{NS} there are 24 vertices, which consist of 16 local-deterministic assignments and 8 non-local boxes [39], see Fig. 3.4. The local deterministic vertices are such that exactly one outcome occurs for every choice of settings, while the non-local vertices correspond to the PR-box distributions, $a \oplus b = xy$ and the respective symmetries under relabelling of inputs, outputs and parties. The relabelling operation acts as a permutation within the set of non-local vertices and within the set of local vertices.

Parametrization. Instead of dealing with the full probability vector it is often instructive to use a parametrization in terms of the marginal probabilities $\langle A \rangle_{xy} = \sum_b \sum_a aP(a, b|x, y)$ and $\langle B \rangle_{xy} = \sum_a \sum_b bP(a, b|x, y)$, and the correlator $\langle AB \rangle_{ab} = \sum_{a,b} abP(a, b|x, y)$. A priori these quantities are 4-component vectors, however, signal-locality requires that $\langle A \rangle_{xy} = \langle A \rangle_x$ and $\langle B \rangle_{xy} = \langle B \rangle_y$. The remaining 8 parameters fully describe the conditional probabilities

$$P(a, b|x, y) = \frac{1}{4} \left(1 + a \langle A \rangle_x + b \langle B \rangle_y + ab \langle AB \rangle_{ab} \right). \quad (3.7)$$

The set \mathcal{NS} in Eq. (3.6) can thus also be described by the vector $(\langle A \rangle_0, \langle A \rangle_1, \dots, \langle AB \rangle_{11})$.

Bell-local Correlations

The set of Bell-local correlations \mathcal{BL} satisfies signal locality and is thus a subset of the set \mathcal{NS} , with the additional constraint of local causality and measurement independence, given by Eq. (3.1), see Chap. 5 for more details. The 16 vertices of the set Bell-local correlations are exactly the local-deterministic assignments of the no-signaling set. However, the polytope of Bell-local correlations has 8 additional facets, which correspond to the 8 symmetries of the CHSH inequality (3.2) under relabelling of inputs, outcomes and parties, see Fig. 3.4. Also here, relabelling is a symmetry transformation, which does not change the set of correlations, but this is not true in general. An example where the various facets are not related by a simple relabelling is the inequality discussed in Chap. 5, where the outcome B may depend on

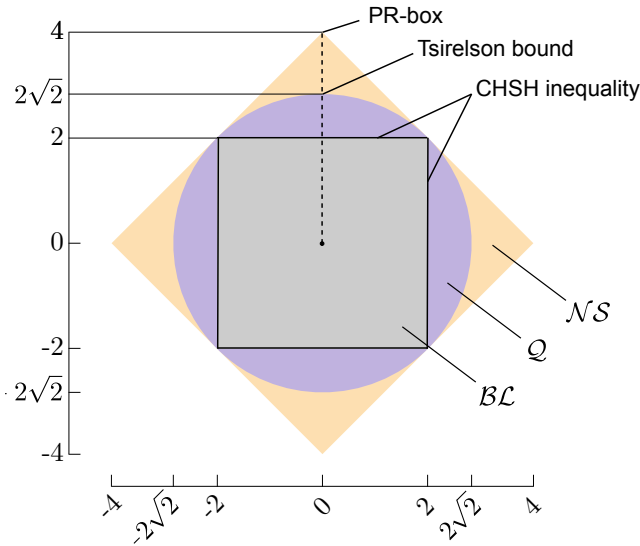


Figure 3.4: **Correlation polytopes in the simplest case.** A sketch of a two-dimensional cut through the structure of the space of possible correlation between two parties, who can each perform one of two dichotomic measurements. The set of Bell-local (\mathcal{BL}) correlations is a polytope, which is bounded by CHSH inequalities. These correlations are a subset of the polytope of correlations compatible with signal locality (\mathcal{NS}). The set of quantum correlations (\mathcal{Q}) is convex, but not a polytope and thus more difficult to characterize. The dashed line indicates isotropic correlations, which include as special cases the maximally mixed state (in the centre), the maximally entangled state (achieving the Tsirelson bound at the boundary between \mathcal{Q} and \mathcal{NS}) and the PR-box, which is the strongest form of correlation compatible with signal locality [5].

outcome A . In either case, studying the polytope corresponding to the probability distributions compatible with a certain constraint (such as local causality) provides an easy way for finding experimentally testable inequalities for these correlations.

The Set of Quantum Correlations

The set of Bell-local correlations is a strict subset of the set \mathcal{Q} of quantum correlations, with the boundary given by the 8 symmetries of the CHSH inequality. At the same time, quantum correlations are a subset of no-signaling correlations, but the exact form of this boundary is, even in the simplest case, still largely unknown. The Tsirelson bound, which marks the maximal violation of the CHSH inequality, is in fact only one point of the continuous boundary between quantum and post-quantum correlations. The rest of the boundary is computationally difficult to study, since the set of quantum correlations, in contrast to Bell-local and no-signaling correlations, is not a polytope, see Fig. 3.4.

However, not every aspect of the problem is intractable and both, numerically deriving, and analytically proving Tsirelson-type bounds for CHSH inequalities with two parties and n binary measurements, can be cast as a semidefinite programming problem and solved efficiently [49]. Building on this result, Navascues et al. [50, 51] developed an infinite hierarchy of conditions

that, in the asymptotic limit, describe the set of quantum correlations. Each level of this hierarchy amounts to testing for the existence of a positive semidefinite matrix with ever more stringent constraints. A failure of this test identifies a candidate correlation as non-quantum, but the hierarchy only converges to the quantum set in the asymptotic limit. Nevertheless, the authors identify situations where a quantum representation can be certified in a finite number of steps [51]. Curiously, the first step of this hierarchy of correlations corresponds to a closed set of correlations that is only slightly larger than quantum correlations and thus called *almost quantum correlations* [52]. This set of correlations was shown to satisfy most physical principles that are meant to single out quantum correlations, and there is numerical evidence suggesting that it might satisfy information causality as well, which makes it an interesting test bed for further studies [52].

3.3 Contextuality

In a noncontextual theory, any two experimental procedures that produce the exact same statistics in every situation, should be represented by the same element in the theory. Consider, for example, the maximally mixed quantum state $\rho = \frac{1}{d}\mathbb{1}$, which can be decomposed into infinitely many different sets of pure states. The linearity of quantum theory implies that any decomposition of this state will return the exact same statistics under any measurement. In a noncontextual ontological model³ of quantum mechanics, the maximally mixed state would therefore have to be represented by the same “state of reality” (or *ontic state*, see Sec. 4.2), regardless of how it was generated. As a consequence of the result presented in Chap. 4 such models are not possible. Following Spekkens [53], there are three relevant notions of contextuality for an operational theory such as quantum mechanics

- i) *Preparation noncontextuality*: If two preparations are operationally indistinguishable (i.e. they produce the same statistics for every measurement), then they should be represented in the same way at the ontic level.
- ii) *Transformation noncontextuality*: If two transformations are operationally indistinguishable (i.e. they produce the same statistics for every preparation and measurement), then they should be represented in the same way at the ontic level.
- iii) *Measurement noncontextuality*: If two measurements are operationally indistinguishable (i.e. they produce the same statistics for every preparation procedure), then they should be represented in the same way at the ontic level.

Formally, one defines equivalence classes of preparations, transformations, and measurements as those elements that are operationally indistinguishable (i.e. produce the same statistics). In the case of quantum theory, these equivalence classes are density matrices, CP maps, and measurement operators, respectively. Everything that distinguishes elements within an equivalence

³An ontological model for an operational theory aims to explain the operational predictions for measurement outcomes in terms of objective physical properties, *ontic states* (or hidden variables), see Sec. 4.2.

class but has no observable effect, such as different pure-state decompositions of a mixed state, is part of the context. For a model to be noncontextual *all the elements* of an equivalence class must be represented in the same way in the model. On the other hand, a model is contextual if it fails for *at least one element*. In this sense the notion of contextuality is more permissive than noncontextuality. The latter, however, holds in classical mechanics, and indeed, the emerging definition of a classical theory is a local non-contextual theory [53, 54].

Preparation Contextuality

The observable statistics of quantum mechanics depend only on the *density matrix* but not the specific pure-state decomposition. However, it seems reasonable to expect that two different ways of preparing the same mixed state—two distinct experimental procedures, using distinct pure states—would lead to some difference at the ontic level. Indeed, any ontological model of quantum mechanics must be preparation contextual for the maximally mixed state⁴ [53], which can (for non-maximally ψ -epistemic models) be extended to all mixed states [55]. These results only rest on the assumption that convex combinations of quantum states are represented by convex sums of epistemic states. This is a fairly uncontroversial assumption if one considers a preparation of mixed states via a classical random choice. Whether preparation contextuality extends to pure states is still an open question.

Transformation Contextuality

The observable statistics of a quantum channel depend only on the *CP map* but not on the specific decomposition into Kraus operators. However, since different Kraus decompositions correspond to different physical implementations, one might again expect that this difference may reveal itself at the level of the transition matrix for ontic states. Indeed, any ontological model of quantum mechanics must be transformation contextual for irreversible transformations, which are the ones with non-trivial Kraus representation [53]. Although there is a close relationship between quantum states and quantum transformations by means of the Choi-Jamiołkowski isomorphism, at the ontic level transformations and preparations play a rather different role. The study of ontological models of quantum mechanics, however, mostly focuses on simple prepare-and-measure experiments, since transformations can often be absorbed into the state preparation, and thus typically play a secondary role.

Measurement Contextuality

A quantum measurement is represented by a POVM, $\{E_i\}$, and the probability for each outcome E_i depends only on this *measurement operator*, but not on the full POVM. In the case of sharp measurements on systems of dimension 3 or larger, there are thus multiple ways to complete a given measurement operator E_i to a proper measurement, which gives rise to a notion of

⁴It is quite ironic that, although many of the strange features of quantum mechanics can be reproduced classically, one of the simplest possible quantum states—the maximally mixed state—shows genuinely non-classical behaviour.

measurement contextuality first discussed by Bell [56], Kochen and Specker [57], see below. An unsharp measurement, on the other hand, can in general be implemented in many different ways as a convex combination of sharp measurements. This gives rise to a notion of measurement contextuality even for qubits that is very similar to that of preparation and transformation contextuality outlined above [53]. Any ontological model for quantum mechanics, which is preparation contextual must also be measurement contextual for unsharp measurements [53].

3.3.1 Kochen-Specker Contextuality

The notion of contextuality was pioneered by Bell [56], Kochen and Specker [57], who showed that, for systems of dimension 3 or larger, there can be no ontological model that is measurement noncontextual, and outcome deterministic. The combination of these two properties is now called *Kochen-Specker noncontextuality*. Outcome determinism requires that the outcome probabilities for a given ontic state can only take the values 0 or 1. This does, however, not imply that the quantum probabilities must be deterministic, since they arise from an average over many ontic states. Curiously, outcome determinism for sharp measurements is a consequence of preparation noncontextuality and the perfect predictability associated with sharp measurements [58].

For systems of dimension 3 or larger⁵, a sharp measurement operator can be part of multiple projective measurements. For example, a measurement along $|0\rangle$ for a qutrit could be part of the contexts $\{|0\rangle, |1\rangle, |2\rangle\}$, as well as $\{|0\rangle, \frac{1}{\sqrt{2}}(|1\rangle + |2\rangle), \frac{1}{\sqrt{2}}(|1\rangle - |2\rangle)\}$. Proofs of the Kochen-Specker theorem are then phrased as graph-theoretic arguments using a number of overlapping measurement bases. In dimension 4, for example, one could consider 18 measurement directions and 9 orthonormal bases, such that each measurement operator is part of exactly two bases. The measurement operators are associated with vertices of a graph, which, in accordance with outcome determinism, can for a fixed ontic state be assigned either probability 0 and coloured white, or probability 1 and coloured black. Normalization of probabilities requires that there is exactly one black vertex in each basis. Such an assignment, however, is not possible in a noncontextual model [59]. The contradiction can be avoided by either allowing values different from 0 and 1 (i.e. giving up outcome determinism), or by allowing the values to depend on the respective basis (i.e. giving up measurement noncontextuality).

In order to make the Kochen-Specker theorem experimentally testable, the logical contradiction must be turned into an inequality [60]. Although this is always possible [61], care has to be taken to not arrive at an inequality that merely rules out models which are logically inconsistent [59]. The simplest case, the Klyachko-Can-Binicioğlu-Shumovsky (KCBS) inequality, involves 5 measurements $\{A_i\}$, each taking values ± 1 , arranged in a pentagram around the state vector of the measured state [62]. These measurements are chosen such that A_i and $A_{i+1 \pmod 5}$ are compatible, which allows for a deterministic value assignment in the inequality

$$\langle A_1 A_2 \rangle + \langle A_2 A_3 \rangle + \langle A_3 A_4 \rangle + \langle A_4 A_5 \rangle + \langle A_5 A_1 \rangle \geq -3. \quad (3.8)$$

⁵Kochen and Specker provided an explicit counterexample to their theorem for two-dimensional systems [57].

Due to the odd number of terms, this is a frustrated network of correlations. Moreover, the inequality is state-dependent, although the notion of contextuality is not. However, this is just a shortcoming of this very simple inequality and, adding more measurement configurations, one can formulate state-independent inequalities.

One of the main drawbacks is that any derivation of Kochen-Specker inequalities relies on the idealized assumptions of exact operational equivalence and sharp, noiseless measurements, which cannot be achieved in practice. Experimental tests (e.g. Ref. [63–65]) are subject to a finite-precision loophole⁶, which has led to a lot of discussion about whether contextuality is experimentally testable at all [58, 66].

3.3.2 Universal (Operational) Contextuality

A violation of a Kochen-Specker inequality can always be attributed to a failure of outcome determinism, rather than measurement noncontextuality. Hence, instead of testing Kochen-Specker noncontextuality, it would be desirable to test noncontextuality directly without idealized assumptions. Recall that, according to Spekkens’ operational definition [53], a model is contextual if it fails to be noncontextual for any preparation, transformation, or measurement. Preparation contextuality is considered the most fundamental since it cannot be avoided in any ontological model of quantum mechanics [55] (except potentially maximally ψ -epistemic models, which are ruled out on other grounds, see Chap. 4). Preparation contextuality, together with perfect predictability for eigenstate measurements, can then be used to derive outcome determinism for sharp measurements [58]. The Kochen-Specker theorem can thus be derived from the assumptions of noncontextuality, operational equivalence and perfect predictability [59]. Although this does not seem to improve the situation much at first sight, perfect predictability is a quality that can be tested experimentally, in contrast to outcome determinism, which acts at the ontic level. Hence, turning the argument around, universal noncontextuality and operational equivalence imply a failure of perfect predictability. One can then derive quantitative bounds on how far from perfect predictability a noncontextual model must be [59]

$$A = \frac{1}{6} \sum_{t \in \{1,2,3\}} \sum_{b \in \{0,1\}} P(X = b | M_t, P_{t,b}) \leq \frac{5}{6}. \quad (3.9)$$

Here, A is the average degree of correlation between the measurement outcome X and the (binary) choice of preparation b , for three measurement choices M_t and three pairs of preparations $P_{t,b}$. If measurements are more predictable than this bound, and operational equivalence holds for measurements and preparations, then universal contextuality is ruled out [59]. Contextuality inequalities derived in this way are robust against noise but still require very low noise to be violated.

⁶In principle, one could extend the Hilbert space in a way that all states can be modelled as pure states of system+ancilla, transformations are unitary and measurements are sharp (the “church of the larger Hilbert space”). This turns a measurement-contextual, but outcome-indeterministic model into a deterministic one, but this is not unique and requires some assumption about extending the ontic state space [58].

The final hurdle is the assumption of exact operational equivalence. Recall that two preparations (measurements) are operationally equivalent if they produce the same statistics for all possible measurements (preparations). Experimentally, this is clearly unfeasible, since not only are no experimental procedures exactly equivalent, but neither can any experiment test all possible measurements. To overcome this in practice, one can rely on the assumption that the chosen sets of preparations and measurements are each informationally complete. Operational equivalence for such sets then implies operational equivalence relative to all possible preparations and measurements, respectively. The statistics of any set of operationally equivalent preparations and measurements can then be inferred from the statistics on those well-characterized informationally complete procedures, see Ref. [67] for details. However, although three binary-outcome measurements are complete for quantum mechanics, this need not be the case for a more general theory that the data must be compared against. Despite some evidence in favour of it⁷, it is unclear if any experiment could ever fully justify the assumption that three measurements are informationally complete for a qubit. Hence, a violation of Eq. (3.9) could still be attributed to a failure of the chosen measurements to be informationally complete, rather than a failure of universal noncontextuality.

3.3.3 Noncontextuality and Classicality

Noncontextuality (in the operational sense defined above) is emerging as the most general definition for classicality [54, 59, 67–72]. Bell-locality, which is another commonly held notion of classicality, can, for example, be considered a special case of contextuality, where Bob’s measurements are Alice’s measurement context and vice versa⁸ [16, 68, 74]. For any preparation noncontextual model there is a Bell-local model, and any proof of Bell’s theorem can be turned into a proof of preparation noncontextuality [71]. The converse implication from preparation noncontextuality to Bell’s theorem works in special cases [75], see Chap. 4. Moreover, negativity of the Wigner quasi-probability distribution, which is a notion of non-classicality for continuous-variable systems, was also found equivalent (when generalized appropriately) to contextuality [70].

3.3.4 Wigner Negativity

When trying to represent the quantum states of continuous variable quantum systems on some kind of phase-space, one is naturally led to quasi-probability distributions, such as the Wigner distribution. These share some features with classical probability distributions and can reproduce quantum measurement statistics. However, in contrast to probability distributions, they are bounded, and thus occupy a minimum area set by the preparation uncertainty principle, see Chap. 6, and they can (and have to be) negative in some areas. Specifically, a quantum

⁷From fitting the data from a (quantum mechanically) over-complete set of four measurements to a theory where three are complete.

⁸In fact, it was shown that the spatial contextuality-inequalities of KCBS, and Bell-inequalities are both special cases of so-called the n -cycle inequalities [68, 73].

state ρ is associated with a quasi-probability distribution $\mu_\rho(\lambda)$ (normalized to $\int d\lambda \mu_\rho(\lambda) = 1$), and a POVM with outcomes $\{E_k\}$ is associated with a response-function $\zeta_{E_k}(\lambda)$ (normalized to $\sum_k \zeta_{E_k}(\lambda) = 1$), such that

$$\text{Tr}[\rho E_k] = \int_\Lambda d\lambda \mu_\rho(\lambda) \zeta_{E_k}(\lambda). \quad (3.10)$$

Here, $\lambda \in \Lambda$ corresponds to a phase space point, or *ontic state* of the system. Traditionally, the ontic state might be associated with a pair of position- and momentum- coordinates, $\lambda = (q, p)$, but any pair of complementary quadratures would do. The appearance of negative values of the distribution in parts of phase space is often considered a signature of non-classical behaviour. However, negativity in one basis does not imply that there doesn't exist a basis in which the representation is non-negative, and Wigner negativity has been found to be neither sufficient, nor necessary for non-classicality [70].

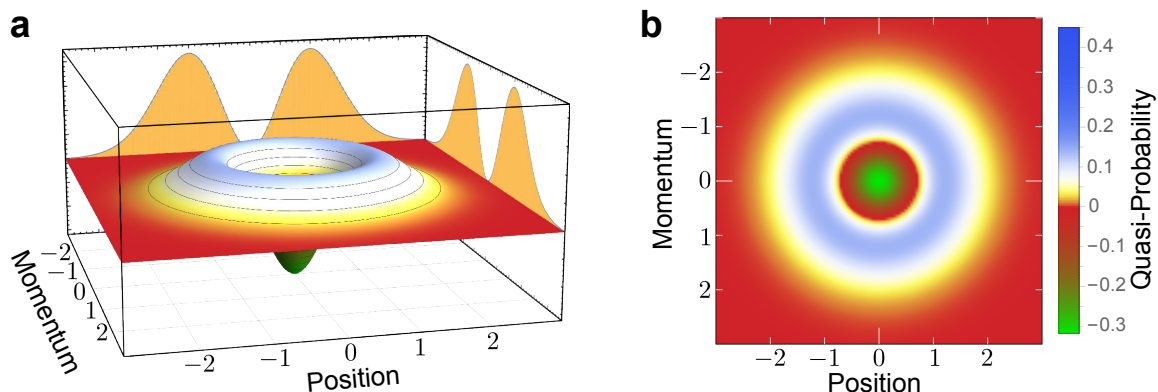


Figure 3.5: **Wigner quasi-probability distribution.** (a) A single-quantum Fock state has a negative dip in the centre of the Wigner-distribution, yet both quadratures remain non-negative and show a double-peaked structure. (b) Two-dimensional projection of the Wigner distribution.

Adopting a more general framework, a quasi-probability representation of (a subset of) quantum theory is non-negative, if the distribution μ is non-negative for all quantum states and the response function ζ is non-negative for all measurements. By definition, a quasi-probability distribution is noncontextual, since μ depends only on the density matrix and ζ only on the POVM element. Noncontextual, positive phase-space distributions, however, cannot reproduce quantum predictions, see Sec. 3.3. Any viable ontological model for quantum mechanics must thus either feature negative probabilities at the ontic level, or be contextual [70].

3.4 Paradoxes and Basic Games

Many of the counterintuitive phenomena observed in quantum systems can be phrased as logical paradoxes—a set of conditions that cannot be satisfied by any classical system. These formulations reveal most clearly the contrast between classical and quantum systems, but they rely on idealized assumptions, such as perfect correlations or exact equivalences. In order to make them experimentally robust, they are used to derive inequalities, which have the form of

games, played by two or more parties, where quantum systems achieve an improvement over the best classical strategy, which represents the bound of the inequality. These games also often act as primitives in the context of quantum cryptography, computation, communication, or for foundational questions.

3.4.1 Quantum Teleportation

Quantum teleportation uses Bell-nonlocal entangled states as a resource for transmitting quantum information without transmitting the physical quantum system that carries the information. What is teleported is *information*, not matter. Consider Alice and Bob sharing a maximally entangled two-qubit quantum state, such as the Bell state $|\Psi^-\rangle_{AB}$, and a third party, Charlie, prepares a state $|\psi\rangle_C$ that he wishes to send to Bob. The joint-state of Alice, Bob and Charlie can be written as

$$\begin{aligned} |\Psi^-\rangle_{AB} \otimes |\psi\rangle_C &= \frac{1}{\sqrt{2}} (|01\rangle - |10\rangle)_{AB} \otimes (\alpha|0\rangle + \beta|1\rangle)_C \\ &= \frac{1}{2} (|\Phi^+\rangle_{AC} (-\beta|0\rangle + \alpha|1\rangle)_B + |\Phi^-\rangle_{AC} (\beta|0\rangle + \alpha|1\rangle)_B + \\ &\quad |\Psi^+\rangle_{AC} (-\alpha|0\rangle + \beta|1\rangle)_B + |\Psi^-\rangle_{AC} (\alpha|0\rangle + \beta|1\rangle)_B). \end{aligned}$$

When Alice performs a joint measurement of AC in the Bell-state basis and obtains outcome $|\Psi^-\rangle$, then Bob's state is projected onto the original state $|\psi\rangle_C$. For all other outcomes of Alice's measurement Bob has to apply a correction, which may be a bit-flip, a sign-flip, or both. To complete the teleportation, Alice must communicate her measurement outcome to Bob, which requires two bits of classical communication, see Fig. 3.6. Without this information, Bob's state is maximally mixed, and thus, although the teleportation happens essentially instantaneous, special relativity is not violated, since the information becomes available only once the classical information arrives, which necessarily travels at most at the speed of light. The classical information itself, however is completely random and thus useless to an eavesdropper, who does not also have access to Bob's particle.

Moreover, Alice cannot even in principle know the state she is teleporting without destroying it. This makes quantum teleportation particularly interesting for cryptographic applications. For example, cluster-state quantum computation uses a chain (or more complex structure) of pairwise entangled qubits and works by teleporting quantum information from one system to the next. The actual computation happens "during" the teleportation by using the error-correction step to implement a quantum operation on the teleported qubit. This can be used for so-called *blind quantum computation* [76], where Charlie prepares an input that he sends together with a set of measurement instructions to Alice to perform a cluster-state quantum computation. Since Alice cannot read the qubits without destroying it, it becomes possible to have a computer perform a computational task without the computer knowing the input, the output, or the kind of computational task that it is performing⁹.

⁹This is a step beyond what is possible with obfuscation algorithms in classical computing, which merely

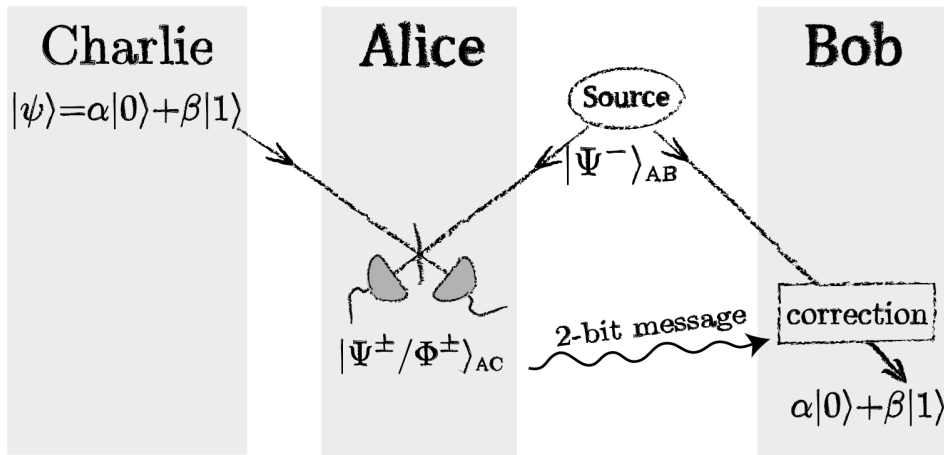


Figure 3.6: **Quantum Teleportation.** Alice and Bob share a maximally entangled state $|\Psi^-\rangle_{AB}$ and Charlie prepares a qubit $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ that he wants to teleport to Bob. Alice performs a bell-state measurement on her part of the shared entangled state together with Charlie's qubit, and sends the outcome as a two-bit message to Bob. Depending on Alice's outcome, Bob applies a correction to his qubit and obtains Charlie's original state. Notably, Alice cannot know the qubit that is teleported, since such a measurement would destroy this state. This protocol is the basis for superdense coding and entanglement swapping, see text.

Superdense Coding

Superdense coding [77] is in some sense the inverse of quantum teleportation. Whereas teleportation is used to transmit a qubit using shared entanglement and two classical bits of communication, superdense coding is used to transmit two classical bits using shared entanglement and one qubit of communication. In the scenario above, Alice can transform the shared state $|\Psi^-\rangle_{AB}$ into any of the four Bell-states by means of bit- and sign-flips on her qubit alone, thereby encoding two classical bits. She then sends her qubit to Bob, who performs a Bell-state measurement to decode the two bits. It has been shown that two bits is the maximum that can be transmitted using a single qubit [77].

Entanglement Swapping

In an important extension of quantum teleportation, Charlie, instead of a pure quantum state, gives Alice one half (C) of an entangled state and keeps the other half (D). The rest of the protocol remains the same and we can write the joint state of Alice (A), Bob (B) and Charlie (CD) as

$$|\Psi^-\rangle_{AB}|\Psi^-\rangle_{CD} = \frac{1}{2} (|\Phi^+\rangle_{AC}|\Phi^+\rangle_{BD} + |\Phi^-\rangle_{AC}|\Phi^-\rangle_{BD} - |\Psi^+\rangle_{AC}|\Psi^+\rangle_{BD} + |\Psi^-\rangle_{AC}|\Psi^-\rangle_{BD}) . \quad (3.11)$$

Hence, Alice's Bell-state measurement on the system AC also projects the system BD onto a Bell state. As before, the protocol requires classical information about Alice's measurement outcome to apply appropriate local error correction. Since the no-cloning theorem prevents

hide the computation, but cannot hide input and output.

noiseless amplification of quantum information entanglement swapping plays a central role in long-distance quantum cryptography [78, 79]. Charlie and Bob are able to share entanglement over a usually prohibitively large distance by breaking it up into smaller segments, and using entanglement swapping and quantum memories. The latter are required to contend with unavoidable losses and inefficiencies.

3.4.2 GHZ Paradox and Multipartite Entanglement

Using multipartite entangled states, it is possible to demonstrate Bell-nonlocality based on a simple logical contradiction, “without inequalities” [80]. Consider three parties who each receive an input bit, denoted X, Y, Z , respectively and each can produce a ± 1 -valued output, denoted A, B, C , respectively, such that

$$\begin{aligned} A_0 B_0 C_0 &= 1 \\ A_1 B_1 C_0 &= -1 \\ A_1 B_0 C_1 &= -1 \\ A_0 B_1 C_1 &= -1 . \end{aligned} \tag{3.12}$$

Here, P_s denotes the outcome for party P for setting s . If the parties are generating their outputs using a pre-arranged strategy and shared classical randomness, but without communicating, it is impossible for them to satisfy all four requirements. Indeed, the best such a Bell-local model can do is to satisfy three of the four. An easy way to see the contradiction is by multiplying all four equations, such that the left-hand side is a product of squares, which must be $+1$, while the right-hand side equals -1 [81]. Using instead a shared a GHZ-state $\frac{1}{2}(|000\rangle + |111\rangle)$, to generate the outputs as the results of measurements of σ_x (for input 0) and σ_y (for input 1), all four requirements ideally satisfied exactly.

The GHZ paradox can be turned into an experimentally testable inequality, known as the Mermin-Ardehali-Belinskii-Klyshko inequality [82]

$$|\langle A_0 B_1 C_1 \rangle + \langle A_1 B_0 C_1 \rangle + \langle A_1 B_1 C_0 \rangle + \langle A_0 B_0 C_0 \rangle| \leq 2 , \tag{3.13}$$

where $\langle A_x B_y C_z \rangle = \sum_{a,b,c} abc P(a, b, c|x, y, z)$ denotes the joint expectation value of A_x, B_y , and C_z . The GHZ-state violates this inequality up to the algebraic maximum of 4. Although a violation of inequality (3.13) demonstrates Bell-nonlocality, it does not demonstrate genuine 3-partite entanglement. A violation could instead be obtained from measurements on a statistical mixture of bipartite entangled states with a classically correlated third particle [83]. To demonstrate genuine tripartite Bell-nonlocality, Svetlichny derived an inequality that cannot

be violated by any states that are only bipartite entangled [81, 83, 84]

$$\begin{aligned}
S_{\text{SVET}} &= \left| \langle A_0 B_0 C_0 \rangle + \langle A_0 B_0 C_1 \rangle + \langle A_0 B_1 C_0 \rangle + \langle A_1 B_0 C_0 \rangle - \right. \\
&\quad \left. \langle A_1 B_1 C_0 \rangle - \langle A_1 B_0 C_1 \rangle - \langle A_0 B_1 C_1 \rangle - \langle A_1 B_1 C_1 \rangle \right| \leq 4 \\
&= S_{\text{CHSH}}(AB)C_0 + S'_{\text{CHSH}}(AB)C_1 \leq 4.
\end{aligned} \tag{3.14}$$

This inequality can be interpreted as a frustrated network of correlations [83]. In the case where Charlie's input is 0, Alice and Bob play the CHSH game, while for 1 they play the CHSH' game (with reversed inputs) [85].

3.4.3 Hardy's Paradox

In the spirit of the GHZ paradox, one can also derive a logical contradiction for only two qubits. Hardy's original presentation [86] relied on two overlapping interferometers, one for an electron and one for a positron, which are arranged such that electron and positron would be expected to annihilate in two of the arms, similar to the Elitzur-Vaidman bomb-test [87]. This narrative, however, rather complicates the paradox which applies to arbitrary two-qubit entangled states (except maximally entangled) [88]. A simplified version in terms of black-boxes has been given by Mermin [89] and, in terms of quantum cakes by Kwiat et al. [90].

Formally, Hardy's paradox can be summarized by the incompatible probability assignment

$$P(-1, +1|0, 1) = 0 \tag{3.15a}$$

$$P(+1, -1|1, 0) = 0 \tag{3.15b}$$

$$P(+1, +1|0, 0) = 0 \tag{3.15c}$$

$$P(+1, +1|1, 1) > 0, \tag{3.15d}$$

where $P(a, b|x, y)$ denotes the probability for output a for Alice and b for Bob, given input x for Alice and y for Bob. Equation (3.15c) implies that, whenever both settings are 0, Alice and Bob never both get outcome +1. Equation (3.15b) and (3.15c), respectively, show that whenever only Alice's setting is 1, she never gets output +1 and whenever only Bob's setting is 1, he never gets output +1. Classically, i.e. in a Bell-local model, this would imply that also when both inputs are 1, the Alice and Bob will never both get outcome +1. Quantum mechanically, however, the probability for this event can be up to $\sim 9\%$, while keeping all other probabilities zero.

As in the GHZ case, this is not yet experimentally testable since it relies on idealized assumptions. Following Mermin [89], however, it can be turned into an inequality

$$P(1, 1|1, 1) - P(0, 1|0, 1) - P(1, 0|1, 0) - P(1, 1|0, 0) \leq 0, \tag{3.16}$$

which is a form of a CH inequality [91]. Hardy's paradox is nonetheless a beautifully clear illustration of the failure of counterfactual reasoning in quantum mechanics.

3.4.4 Leggett-Garg and Macrorealism

With their paper “Quantum mechanics versus macroscopic realism: Is the flux there when nobody looks?”¹⁰, Leggett and Garg [93] aimed to devise an experimental test for whether our world adheres to macrorealism—the view that macroscopic objects have definite properties, independent of, and undisturbed by observation. This is a view held, for example, by *objective collapse models*, which postulate that the quantum description breaks down at some scale, and superpositions of macroscopic objects cannot be sustained. When an object counts as macroscopic is not exactly well-defined but is typically considered a trade-off between the mass [94] or energy [95] of the object, and the timescale over which a superposition can be maintained. Experimental tests of macrorealism thus effectively amount to observing and maintaining superposition states of macroscopic objects, and provide a promising route towards testing such models [96–99].

Specifically, Leggett-Garg’s macrorealism is composed of three assumptions [96]: macrorealism per se, non-invasive measurability and induction. *Macrorealism per se* demands that a macroscopic system, which could be in a number of macroscopically distinct states, always occupies exactly one of these states. What exactly counts as “macroscopically distinct states” is left open and might well be theory-dependent, although there are some arguments pertaining to coarse-grained measurements [100]. *Non-invasive measurability* asserts that it is in principle possible to determine the state of the macroscopic system without disturbing it in any way. *Induction* is the assumption of a temporal order and causal arrow, such that later events cannot influence earlier events. Although non-invasive measurability is closely entwined with macrorealism per se, the former is not a consequence of the latter (with Bohmian mechanics as a counter example) [97]. Leggett and Garg then consider a series of measurements of the same observable Q , performed at three different times on a single quantum system to derive what is now known as *Leggett-Garg inequalities*, which are satisfied by any macrorealistic theory, see Fig. [93].

$$\langle Q_1 Q_2 + Q_2 Q_3 + Q_1 Q_3 \rangle \geq -1 \quad (3.17a)$$

$$\xrightarrow{\text{NIM}} \langle Q_1 Q_2 \rangle_{12} + \langle Q_2 Q_3 \rangle_{23} + \langle Q_1 Q_3 \rangle_{13} \geq -1, \quad (3.17b)$$

where $\langle \cdot \rangle_{ij}$ is the expectation value when only the measurements at time i and j are actually performed. The step from Eq. (3.17a), which is trivially satisfied, to Eq. (3.17b), requires that removing one of the measurements does not change the joint statistics of the other two, which is a consequence of non-invasive measurability (NIM). Focusing on this crucial feature allows for a much stronger argument than Leggett-Garg inequalities. Specifically, Ref. [96] introduced the notion of *no-signaling in time*, which is satisfied if a measurement does not change the statistics of future measurements. No-signaling in time fails *exactly* for the quantum interference terms [101] and, together with the causal arrow, is equivalent to macrorealism [99]. This

¹⁰Inspired by Einstein’s famous question “is the moon there when nobody looks?” [92], which was intended to challenge the idea that there might be no objective reality at all, see Chap. 4 for a detailed discussion.

is in contrast to Leggett-Garg inequalities, which are necessary but not sufficient for macro-realism¹¹. The fact that the system is macroscopic (whatever that means) has not been used anywhere in the argument, and any derived constraints thus hold independent of the system size. This confirms the intuition ruling out macrorealism comes down to observing quantum interference in a system large enough to be considered macroscopic. Hence, experiments such as molecule interferometry [103] or optomechanical interference, see Appendix C serve a very similar purpose to direct violations of no-signaling in time in a macroscopic qubit system, such as a superconducting circuit [104].

Recall that Leggett-Garg compare the joint expectation value of two temporal measurements in the situation where a third measurement is or is not performed, i.e. a change of measurement context. As pointed out in Ref. [97], a macrorealistic model is in fact a Kochen-Specker contextual model for the considered macroscopic observables. This reveals an interesting connection to the Kochen-Specker theorem, which shows that there can be no ontological model that is noncontextual and outcome-deterministic *for all* measurements. Indeed, the KBCS inequality to test Kochen-Specker noncontextuality shares much of the structure of a Leggett-Garg inequality and is based on five measurements of the same observable at different times on an evolving quantum system. There is also a close connection between a violation of no-signaling in time (or a Leggett-Garg inequality) and the appearance of pre-and post-selection paradoxes, see Sec. 3.5. It has been argued that such paradoxes require a measurement between pre- and post-selection that is invasive at the ontic level, but operationally undetectable (i.e. there is no difference in the statistics between the cases where the measurement is performed or not) [105]. Finally, the set of temporal quantum correlations is not equivalent to that of spatial correlations, such that, for example, monogamy of entanglement, which is a central concept for spatial entanglement, does not hold in the temporal analogue.

3.5 Pre- and Post-Selection Paradoxes

When a quantum system is prepared (or pre-selected) in some state $|\psi\rangle$, and then post-selected (i.e. conditioned on a certain outcome of a later measurement) in the state $|\phi\rangle$, quantum theory can give “paradoxical” predictions intermediate measurements. Specifically, in the case of so-called logical pre- and post-selection paradoxes (where probabilities only take values 0 and 1), the paradox is in effect a conflict with Kochen-Specker noncontextual ontological models [69, 106, 107].

Consider the three box paradox [69, 105, 108, 109], where a ball can be in any of three boxes, represented by the states $|1\rangle, |2\rangle$ and $|3\rangle$. At the beginning of the game the system is prepared in the state $(|1\rangle + |2\rangle + |3\rangle)/\sqrt{3}$ (e.g. by placing the ball in box 1 and shuffling the boxes appropriately). An observer then gets to check either if the ball is in box 1 or if it is in box 2.

¹¹The reason for this is that the set of macrorealistic probabilities is a lower-dimensional subset of the set of all temporal quantum correlations (defined by the assumption of the arrow of time). As a consequence, the boundary cannot be formulated in terms of inequalities (i.e. hyperplanes) [102]. In contrast, Bell-local correlations are an equal-dimensional subset of the set of spatial quantum correlations, and the boundary is exactly given by the Bell inequalities [99].

Finally, the system is postselected in the state $(|1\rangle+|2\rangle-|3\rangle)/\sqrt{3}$. The intermediate measurement of box 1 corresponds to $\{|1\rangle\langle 1|, |2\rangle\langle 2|+|3\rangle\langle 3|\}$, and that of box 2 to $\{|2\rangle\langle 2|, |1\rangle\langle 1|+|3\rangle\langle 3|\}$. In either case, the second measurement outcome can be decomposed in a way that is orthogonal to the pre- or post-selected state and thus never occurs [106]. Hence, whenever the observer checks box 1, she finds the ball there, and whenever she checks box 2, she finds the ball there. In other words, the answer she gets depends on the question she asks¹², which already has a very strong contextuality flavour.

However, a priori this does not constitute a proof of the Kochen-Specker theorem, which is concerned with prepare-and-measure experiments, see e.g. Ref. [69, 106] and references therein. On the other hand, logical pre- and post-selection paradoxes are related to proofs of the theorem when considering the intermediate measurements and the final post-selection as counterfactual alternatives [69, 106]. In the paradox scenario itself, however, the intermediate measurement could conceivably disturb the system enough to encode information about the measurement setting, thereby changing the post-selection probability. In fact, any noncontextual model for a pre- and post-selection paradox must involve measurement disturbance [69, 105]. It has been shown, however, that the required form of disturbance is itself incompatible with noncontextuality [69]. In particular, since pre- and post-selection are non-orthogonal, the intermediate measurement must disturb the system in a way that the post-selection changes from being possible to being impossible. Specifically, they rely on transformation contextuality and show that the von Neumann state-update rule corresponds to a transformation that sometimes leaves the state unperturbed [69]. The proof thus crucially relies on the von Neumann state update and that intermediate probabilities are 0/1-valued.

3.5.1 Weak Values

In the case where the intermediate measurement (of observable A) between pre- and post-selection is weak, one can define the so-called *weak value* [110]

$$A_w := \frac{\langle \psi | A | \phi \rangle}{\langle \psi | \phi \rangle}. \quad (3.18)$$

For $|\psi\rangle = |\phi\rangle$, this definition reduces to the expectation value $\langle A \rangle$. Considering a von Neumann-type measurement, the weak value appears as a first-order approximation in the coupling strength, of the *relative* change in the reduced probability distribution of the meter. In contrast to the expectation value, the weak value A_w is in general complex valued, with the real part $\text{Re}[A_w]$ corresponding to the average shift in the meter's probability distribution, and the imaginary part related to the backaction on the system. Consequently weak values have become useful in testing noise-disturbance and joint measurement uncertainty relations, see Chap. 6. They have also been related to the velocities of particles in Bohmian mechanics [111, 112] and used in quantum state tomography and parameter estimation, see Ref. [113–115] for a more detailed discussion of weak values and their applications.

¹²In particular, “the ball is in box 2” is not a possible answer to the first question.

In the case where pre- and post-selection states are almost orthogonal (i.e. $\langle \psi | \phi \rangle \ll 1$), the weak value can get arbitrarily large, while the probability of observing these events becomes arbitrarily small. In particular, this allows for so-called *anomalous* weak values, which are larger than the largest (or smaller than the smallest) eigenvalue of A . Weak values are a quite peculiar statistical construct based on post-selection, which can lead to a variety of strange effects [116]. Hence, there is a lot of controversy about what exactly (if anything) is genuinely non-classical about weak values. The imaginary part, for example, appears also in subsets of quantum mechanics that admit a noncontextual (i.e. classical) model [14, 117]. Similarly, anomalous values, which have been attributed much physical significance [109, 110], occur as statistical anomalies in classical weak value models [118], typically in conjunction with measurement disturbance [114, 118–121].

There is, however, an interesting connection between weak values, contextuality and measurement disturbance. For example, when the intermediate measurement in a logical pre- and post-selection paradox is weak, then the paradox must involve anomalous weak values [69]. Any model that can explain anomalous weak values¹³ must either violate outcome determinism for the projective measurements (e.g. in a ψ -ontic model), and/or be contextual for either the weak measurement or the induced disturbance [122]. Related to this is the observation that an anomalous weak value for A implies, as a consequence of normalization, that at least one projector in the spectral decomposition must be assigned a negative weak value. This suggests an intriguing parallel to negative regions in the Wigner quasi-probability distribution [123], which is also related to the failure of noncontextual models.

3.5.2 Classical Anomalous Weak Values

One of the results that spurred a lot of critical and heated discussion about the status of weak values is an example of anomalous weak values appearing in a pre- and post-selected coin-flipping experiment [120]. In this thought experiment Alice repeatedly hands Bob a coin, heads up (the preselection). Bob then performs a noisy measurement of the coin and, depending on the outcome, might or might not flip the coin before handing it back. Alice then keeps only those cases where the coin was tails up (the postselection) and finds that the conditional average of measurement outcome divided by the measurement strength—the “classical weak value” [120]—can take anomalously large values, just like the quantum counterpart. Of course, a noisy classical measurement is different from a weak quantum measurement (e.g. in the information-disturbance tradeoff), but the purpose of the example is merely to demonstrate that weak values can arise as a statistical artefact in an arguably classical model.

Curiously, just like in the case of the three-box paradox, the model features measurement outcome-dependent disturbance. In a follow-up paper [118], the authors discuss the general conditions for the occurrence of an anomalous post-selection shift, where conditioning on the final state increases the shift of the meter compared to the unconditional case. They show that

¹³With the specific information-disturbance trade-off for weak quantum measurements, and under the assumption that post-selection can be treated as a projective measurement.

in the general 3-variable (ψ, s, ϕ) dichotomous model, where s is the outcome of the intermediate measurement, the conditional independence of ϕ and s given ψ is sufficient to ensure that there is no anomalous shift. In other words, for an anomalous shift to occur the final measurement must be correlated to the intermediate measurement in a way that is not only due to the initial state ψ . In particular, this implies that in an ontological model, there must be some disturbance at the ontic level for an anomalous post-selected shift to occur [118], which agrees with previous results [105]. The coin-flip model is an example for such a correlated model.

Leggett-Garg

Weak values also appear naturally in the context of Leggett-Garg inequality tests using weak measurements (see Sec. 3.4.4). Specifically, Leggett-Garg inequalities are violated *if and only if* the same data features anomalous weak values [124, 125]. Recall that the assumptions of Leggett-Garg are that a macroscopic system (such as a coin) is always in a well-defined state (macroscopic realism), which can be determined by a non-invasive measurement (non-invasive measurability). Although the coin in Ref. [120] is arguably in a definite state at any time, the example features correlated measurement disturbance, which violates the non-invasiveness criterion. This allows for anomalous weak values and at the same time, following Ref. [125], the Leggett-Garg value is $\frac{2\delta}{1-\delta} - 2\delta + 1$, which violates the Leggett-Garg inequality for any value of the disturbance $1 - \delta$.

Weak Value Amplification

When pre- and post-selected states are almost orthogonal the weak value can in principle become arbitrarily large¹⁴, which has led to the suggestion that it might be possible to use this so-called *weak value amplification*— to amplify weak signals. Indeed, many experiments report such an effect [126, 127], while, at the same time, statistical analysis suggests that methods which use only a fraction of the data (the postselected part) cannot outperform methods that use the full dataset [128, 129].

As so often, the disagreement comes down to the definitions and performance metrics used. Simply put, the goal of a weak value experiment is to concentrate most of the information into a small fraction of the events, which are more easily detected, and ignore the rest. Although a large part of the relevant information can be retained in this way [130], the weak value method can only match the performance of a postselection-free method in the limit of an infinitely weak measurement [129]. On the other hand, if the relevant resource is the *measured* particles rather than the *input* particles (such as in optical schemes, since photons are cheap), then the postselection probability is irrelevant and the concentration of information implies that the weak-value method performs better [131]. A rigorous analysis of postselected probabilistic quantum metrology can be found in Ref. [132] and an accessible discussion in Ref. [129].

¹⁴It has been argued that the contributions of higher order terms become significant in the limit where the two states are close to orthogonal such that it may not be possible to reach arbitrarily large values [115].

Another subtle aspect to the comparison is whether the final post-selection is interpreted as simply discarding data, or as an additional measurement. In the latter case, there is an advantage when compared to an experiment without the final measurement. Moreover, if photons that “fail” the post-selection can be recycled, the number of interactions between system and meter can effectively be increased and thus lead to an advantage over schemes without post-selection, but this might again be an unfair comparison. On the more technical side, weak value amplification typically performs worse even in the presence of noise [129]. However, there are also some noise models (such as $1/f$ noise) where it is beneficial [130, 133]. Finally, it might be the case that the optimal estimators (for the non-postselected experiment) cannot be implemented or that the initial state cannot be chosen optimally, in which case the optimal post-selection can still help.

In summary, while there seems to be no fundamental advantage in using weak value amplification, there might be a technical one [113].

References

- [1] Barrett, J. Information processing in generalized probabilistic theories. *Phys. Rev. A* **75**, 032304 (2007).
- [2] Barnum, H. & Wilce, A. Post-Classical Probability Theory. *Quantum Theory: Informational Foundations and Foils* (2012).
- [3] Janotta, P. & Hinrichsen, H. Generalized probability theories: what determines the structure of quantum theory? *J. Phys. A* **47**, 323001 (2014).
- [4] Chiribella, G. & Spekkens, R. W. (eds.) *Quantum Theory: Informational Foundations and Foils* (Springer, 2016).
- [5] Popescu, S. Nonlocality beyond quantum mechanics. *Nat. Phys.* **10**, 264–270 (2014).
- [6] Cavalcanti, E. G. & Wiseman, H. M. Bell nonlocality, signal locality and unpredictability (or What Bohr could have told Einstein at Solvay had he known about Bell experiments). *Found. Phys.* **42**, 1329—1338 (2012).
- [7] Fuchs, C. A., Mermin, N. D. & Schack, R. An introduction to QBism with an application to the locality of quantum mechanics. *Am. J. Phys.* **82**, 749–754 (2014).
- [8] Einstein, A., Podolsky, B. & Rosen, N. Can Quantum-Mechanical Description of Physical Reality Be Considered Complete? *Phys. Rev.* **47**, 777–780 (1935).
- [9] Bohm, D. *Quantum Theory* (Dover Publications, 1989).
- [10] Schrödinger, E. Discussion of Probability Relations between Separated Systems. *Math. Proc. Cambridge Phil. Soc.* **31**, 555–563 (1935).
- [11] Bell, J. S. Bertlmann’s socks and the nature of reality. In *Speakables and Unspeakables in Quantum Mechanics* (Cambridge University Press, 1987).
- [12] Wiseman, H. M., Jones, S. J. & Doherty, A. C. Steering, Entanglement, Nonlocality, and the Einstein-Podolsky-Rosen Paradox. *Phys. Rev. Lett.* **98**, 140402 (2007).
- [13] Branciard, C., Cavalcanti, E. G., Walborn, S. P., Scarani, V. & Wiseman, H. M. One-sided device-independent quantum key distribution: Security, feasibility, and the connection with steering. *Phys. Rev. A* **85**, 010301 (2012).
- [14] Bartlett, S. D., Rudolph, T. & Spekkens, R. W. Reconstruction of Gaussian quantum mechanics from Liouville mechanics with an epistemic restriction. *Phys. Rev. A* **86**, 012103 (2012).

- [15] Werner, R. F. Quantum states with Einstein-Podolsky-Rosen correlations admitting a hidden-variable model. *Phys. Rev. A* **40**, 4277–4281 (1989).
- [16] Spekkens, R. Evidence for the epistemic view of quantum states: A toy theory. *Phys. Rev. A* **75**, 032110 (2007).
- [17] Wood, C. J. & Spekkens, R. W. The lesson of causal discovery algorithms for quantum correlations: Causal explanations of Bell-inequality violations require fine-tuning. *New J. Phys.* **17**, 033002 (2015).
- [18] Bell, J. S. On the Einstein Podolsky Rosen Paradox. *Physics* **1**, 195–200 (1964).
- [19] Bell, J. S. The theory of local beables. *Epistemological Lett.* **9**, 11–24 (1976).
- [20] Clauser, J., Horne, M., Shimony, A. & Holt, R. Proposed Experiment to Test Local Hidden-Variable Theories. *Phys. Rev. Lett.* **23**, 880–884 (1969).
- [21] Ferraro, A., Aolita, L., Cavalcanti, D., Cucchietti, F. M. & Acín, A. Almost all quantum states have nonclassical correlations. *Phys. Rev. A* **81**, 052318 (2010).
- [22] Datta, A., Shaji, A. & Caves, C. M. Quantum Discord and the Power of One Qubit. *Phys. Rev. Lett.* **100**, 050502 (2008).
- [23] Dakić, B. *et al.* Quantum discord as resource for remote state preparation. *Nat. Phys.* **8**, 666–670 (2012).
- [24] Modi, K. A Pedagogical Overview of Quantum Discord. *Open Syst. Inf. Dyn.* **21**, 1440006 (2014).
- [25] Datta, A. Quantum optics: Discord in the ranks. *Nat. Photon.* **6**, 724–725 (2012).
- [26] Popescu, S. & Rohrlich, D. Quantum nonlocality as an axiom. *Found. Phys.* **24**, 379–385 (1994).
- [27] Allcock, J., Brunner, N., Pawłowski, M. & Scarani, V. Recovering part of the boundary between quantum and nonquantum correlations from information causality. *Phys. Rev. A* **80**, 040103 (2009).
- [28] Masanes, L., Acín, A. & Gisin, N. General properties of nonsignaling theories. *Phys. Rev. A* **73**, 012112 (2006).
- [29] Skrzypczyk, P., Brunner, N. & Popescu, S. Emergence of Quantum Correlations from Nonlocality Swapping. *Phys. Rev. Lett.* **102**, 110402 (2009).
- [30] Pawłowski, M., Paterek, T., Kaszlikowski, D., Scarani, V., Winter, A. & Zukowski, M. Information causality as a physical principle. *Nature* **461**, 1101–4 (2009).

- [31] van Dam, W. Implausible Consequences of Superstrong Nonlocality. *arXiv:quant-ph/0501159* (2005).
- [32] Brassard, G., Buhrman, H., Linden, N., Methot, A. A., Tapp, A. & Unger, F. A limit on nonlocality in any world in which communication complexity is not trivial. *Phys. Rev. Lett.* **96**, 250401 (2006).
- [33] Linden, N., Popescu, S., Short, A. & Winter, A. Quantum Nonlocality and Beyond: Limits from Nonlocal Computation. *Phys. Rev. Lett.* **99**, 180502 (2007).
- [34] Navascués, M. & Wunderlich, H. A glance beyond the quantum model. *Proc. Royal Soc. A* **466**, 881–890 (2009).
- [35] Fritz, T., Sainz, A. B., Augusiak, R., Brask, J. B., Chaves, R., Leverrier, A. & Acín, A. Local orthogonality as a multipartite principle for quantum correlations. *Nat. Commun.* **4**, 2263 (2013).
- [36] Barnum, H., Beigi, S., Boixo, S., Elliott, M. B. & Wehner, S. Local Quantum Measurement and No-Signaling Imply Quantum Correlations. *Phys. Rev. Lett.* **104**, 140401 (2010).
- [37] Acín, A., Augusiak, R., Cavalcanti, D., Hadley, C., Korbicz, J. K., Lewenstein, M., Masanes, L. & Piani, M. Unified Framework for Correlations in Terms of Local Quantum Observables. *Phys. Rev. Lett.* **104**, 140404 (2010).
- [38] Brunner, N. & Skrzypczyk, P. Nonlocality Distillation and Postquantum Theories with Trivial Communication Complexity. *Phys. Rev. Lett.* **102**, 160403 (2009).
- [39] Allcock, J., Brunner, N., Linden, N., Popescu, S., Skrzypczyk, P. & Vértesi, T. Closed sets of nonlocal correlations. *Phys. Rev. A* **80**, 062107 (2009).
- [40] Gallego, R., Würflinger, L. E., Acín, A. & Navascués, M. Quantum Correlations Require Multipartite Information Principles. *Phys. Rev. Lett.* **107**, 210403 (2011).
- [41] Ambainis, A., Nayak, A., Ta-Shma, A. & Vazirani, U. Dense Quantum Coding and a Lower Bound for 1-way Quantum Automata. *Proc. 31. STOC* 12 (1998).
- [42] Ambainis, A., Nayak, A., Ta-Shma, A. & Vazirani, U. Dense quantum coding and quantum finite automata. *J. ACM* **49**, 496–511 (2002).
- [43] Ambainis, A., Leung, D., Mancinska, L. & Ozols, M. Quantum Random Access Codes with Shared Randomness. *arXiv:0810.2937* (2008).
- [44] Pawłowski, M. & Scarani, V. Information Causality. *arXiv:1112.1142* (2011).
- [45] Al-Safi, S. W. & Short, A. J. Information causality from an entropic and a probabilistic perspective. *Phys. Rev. A* **84**, 042323 (2011).

- [46] Pitowsky, I. Correlation polytopes: Their geometry and complexity. *Math. Program.* **50**, 395–414 (1991).
- [47] Barrett, J., Linden, N., Massar, S., Pironio, S., Popescu, S. & Roberts, D. Non-local correlations as an information theoretic resource. *Phys. Rev. A* **71**, 022101 (2004).
- [48] Christof, T. & Löbel, A. PORTA – POlyhedron Representation Transformation Algorithm. <http://porta.zib.de/> (2009).
- [49] Wehner, S. Tsirelson bounds for generalized Clauser-Horne-Shimony-Holt inequalities. *Phys. Rev. A* **73**, 022110 (2006).
- [50] Navascués, M., Pironio, S. & Acín, A. Bounding the Set of Quantum Correlations. *Phys. Rev. Lett.* **98**, 010401 (2007).
- [51] Navascués, M., Pironio, S. & Acín, A. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New J. Phys.* **10**, 073013 (2008).
- [52] Navascués, M., Guryanova, Y., Hoban, M. J. & Acín, A. Almost quantum correlations. *Nat. Commun.* **6**, 6288 (2015).
- [53] Spekkens, R. W. Contextuality for preparations, transformations, and unsharp measurements. *Phys. Rev. A* **71**, 052108 (2005).
- [54] Jennings, D. & Leifer, M. No return to classical reality. *Contemp. Phys.* **57**, 60–82 (2016).
- [55] Leifer, M. S. Is the quantum state real ? An extended review of ψ -ontology theorems. *Quanta* **3**, 67–155 (2014).
- [56] Bell, J. S. On the Problem of Hidden Variables in Quantum Mechanics. *Rev. Mod. Phys.* **38**, 447–452 (1966).
- [57] Kochen, S. B. & Specker, E. The Problem of hidden variables in quantum mechanics. *J. Math. Mech.* **17**, 59–87 (1967).
- [58] Spekkens, R. W. The Status of Determinism in Proofs of the Impossibility of a Noncontextual Model of Quantum Theory. *Found. Phys.* **44**, 1125–1155 (2014).
- [59] Kunjwal, R. & Spekkens, R. W. From the Kochen-Specker Theorem to Noncontextuality Inequalities without Assuming Determinism. *Phys. Rev. Lett.* **115**, 110403 (2015).
- [60] Cabello, A. Experimentally Testable State-Independent Quantum Contextuality. *Phys. Rev. Lett.* **101**, 210401 (2008).
- [61] Yu, X.-d., Guo, Y.-q. & Tong, D. M. A proof of the Kochen–Specker theorem can always be converted to a state-independent noncontextuality inequality. *New J. Phys.* **17**, 93001 (2015).

- [62] Klyachko, A. a., Can, M. A., Binicioğlu, S. & Shumovsky, A. S. Simple Test for Hidden Variables in Spin-1 Systems. *Phys. Rev. Lett.* **101**, 020403 (2008).
- [63] Kirchmair, G., Zähringer, F., Gerritsma, R., Kleinmann, M., Gühne, O., Cabello, A., Blatt, R. & Roos, C. F. State-independent experimental test of quantum contextuality. *Nature* **460**, 494–497 (2009).
- [64] Lapkiewicz, R., Li, P., Schaeff, C., Langford, N. K., Ramelow, S., Wieśniak, M. & Zeilinger, A. Experimental non-classicality of an indivisible quantum system. *Nature* **474**, 490–493 (2011).
- [65] Zhang, X., Um, M., Zhang, J., An, S., Wang, Y., Deng, D.-l., Shen, C., Duan, L.-M. & Kim, K. State-Independent Experimental Test of Quantum Contextuality with a Single Trapped Ion. *Phys. Rev. Lett.* **110**, 070401 (2013).
- [66] Barrett, J. & Kent, A. Non-contextuality, finite precision measurement and the Kochen-Specker theorem. *Stud. Hist. Philos. Mod. Phys.* **35**, 151–176 (2004).
- [67] Mazurek, M. D., Pusey, M. F., Kunjwal, R., Resch, K. J. & Spekkens, R. W. An experimental test of noncontextuality without unwarranted idealizations. *arXiv:1505.06244* (2015).
- [68] Araújo, M., Quintino, M. T., Budroni, C., Cunha, M. T. & Cabello, A. All noncontextuality inequalities for the n-cycle scenario. *Phys. Rev. A* **88**, 022118 (2013).
- [69] Pusey, M. F. & Leifer, M. S. Logical pre- and post-selection paradoxes are proofs of contextuality. *Electron. Notes Theor. Comput. Sci.* **195**, 295–306 (2015).
- [70] Spekkens, R. W. Negativity and Contextuality are Equivalent Notions of Nonclassicality. *Phys. Rev. Lett.* **101**, 020401 (2008).
- [71] Pusey, M. F. The robust noncontextuality inequalities in the simplest scenario. *arXiv:1506.04178* (2015).
- [72] Larsson, J. Å. A contextual extension of Spekkens’ toy model. *AIP Conf. Proc.* **1424**, 211–220 (2012).
- [73] Chaves, R. Entropic inequalities as a necessary and sufficient condition to noncontextuality and locality. *Phys. Rev. A* **87**, 022102 (2013).
- [74] Cabello, A., Severini, S. & Winter, A. (Non-)Contextuality of Physical Theories as an Axiom. *arXiv:1010.2163* (2010).
- [75] Leifer, M. S. & Maroney, O. J. E. Maximally Epistemic Interpretations of the Quantum State and Contextuality. *Phys. Rev. Lett.* **110**, 120401 (2013).

- [76] Broadbent, A., Fitzsimons, J. & Kashefi, E. Universal Blind Quantum Computation. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, 517–526 (2009).
- [77] Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).
- [78] Duan, L. M., Lukin, M. D., Cirac, J. I. & Zoller, P. Long-distance quantum communication with atomic ensembles and linear optics. *Nature* **414**, 413–418 (2001).
- [79] Khalique, A., Tittel, W. & Sanders, B. C. Practical long-distance quantum communication using concatenated entanglement swapping. *Phys. Rev. A* **88**, 022336 (2013).
- [80] Greenberger, D. M., Horne, M. A., Shimony, A. & Zeilinger, A. Bell’s theorem without inequalities. *Am. J. Phys.* **58**, 1131 (1990).
- [81] Brunner, N., Cavalcanti, D., Pironio, S., Scarani, V. & Wehner, S. Bell nonlocality. *Rev. Mod. Phys.* **86**, 419–478 (2014).
- [82] Żukowski, M. & Brukner, Č. Bell’s Theorem for General N -Qubit States. *Phys. Rev. Lett.* **88**, 210401 (2002).
- [83] Mitchell, P., Popescu, S. & Roberts, D. Conditions for the confirmation of three-particle nonlocality. *Phys. Rev. A* **70**, 060101 (2004).
- [84] Svetlichny, G. Distinguishing three-body from two-body nonseparability by a Bell-type inequality. *Phys. Rev. D* **35**, 3066–3069 (1987).
- [85] Bancal, J.-D., Brunner, N., Gisin, N. & Liang, Y.-C. Detecting Genuine Multipartite Quantum Nonlocality: A Simple Approach and Generalization to Arbitrary Dimensions. *Phys. Rev. Lett.* **106**, 020405 (2011).
- [86] Hardy, L. Quantum mechanics, local realistic theories, and Lorentz-invariant realistic theories. *Phys. Rev. Lett.* **68**, 2981–2984 (1992).
- [87] Elitzur, A. C. & Vaidman, L. Quantum mechanical interaction-free measurements. *Found. Phys.* **23**, 987–997 (1993).
- [88] Hardy, L. Nonlocality for two particles without inequalities for almost all entangled states. *Phys. Rev. Lett.* **71**, 1665–1668 (1993).
- [89] Mermin, N. D. Quantum mysteries refined. *Am. J. Phys.* **62**, 880 (1994).
- [90] Kwiat, P. G. & Hardy, L. The mystery of the quantum cakes. *Am. J. Phys.* **68**, 33 (2000).
- [91] Rabelo, R., Zhi, L. Y. & Scarani, V. Device-Independent Bounds for Hardy’s Experiment. *Phys. Rev. Lett.* **109**, 180401 (2012).

- [92] Mermin, N. D. Is the moon there when nobody looks? Reality and the quantum theory. *Physics Today* **38**, 38–47 (1985).
- [93] Leggett, A. J. & Garg, A. Quantum mechanics versus macroscopic realism: Is the flux there when nobody looks? *Phys. Rev. Lett.* **54**, 857–860 (1985).
- [94] Ghirardi, G. C., Rimini, A. & Weber, T. Unified dynamics for microscopic and macroscopic systems. *Phys. Rev. D* **34**, 470–491 (1986).
- [95] Milburn, G. J. Intrinsic decoherence in quantum mechanics. *Phys. Rev. A* **44**, 5401–5406 (1991).
- [96] Kofler, J. & Brukner, Č. Condition for macroscopic realism beyond the Leggett-Garg inequalities. *Phys. Rev. A* **87**, 052115 (2013).
- [97] Maroney, O. J. E. & Timpson, C. G. Quantum- vs. Macro- Realism: What does the Leggett-Garg Inequality actually test? *arXiv:1412.6139* (2014).
- [98] Emary, C., Lambert, N. & Nori, F. Leggett-Garg inequalities. *Rep. Prog. Phys.* **77**, 016001 (2014).
- [99] Clemente, L. & Kofler, J. No Fine Theorem for Macrorealism: Limitations of the Leggett-Garg Inequality. *Phys. Rev. Lett.* **116**, 150401 (2016).
- [100] Kofler, J. & Brukner, Č. Classical World Arising out of Quantum Physics under the Restriction of Coarse-Grained Measurements. *Phys. Rev. Lett.* **99**, 180403 (2007).
- [101] Dakić, B., Paterek, T. & Brukner, Č. Density cubes and higher-order interference theories. *New J. Phys.* **16**, 023028 (2014).
- [102] Clemente, L. & Kofler, J. Necessary and sufficient conditions for macroscopic realism from quantum mechanics. *Phys. Rev. A* **91**, 062103 (2015).
- [103] Eibenberger, S., Gerlich, S., Arndt, M., Mayor, M. & Tüxen, J. Matter–wave interference of particles selected from a molecular library with masses exceeding 10 000 amu. *Phys. Chem. Chem. Phys.* **15**, 14696 (2013).
- [104] Knee, G. C., Kakuyanagi, K., Yeh, M.-C., Matsuzaki, Y., Toida, H., Yamaguchi, H., Saito, S., Leggett, A. J. & Munro, W. J. A strict experimental test of macroscopic realism in a superconducting flux qubit. *Nat. Commun.* **7**, 13253 (2016).
- [105] Maroney, O. J. E. Detectability, Invasiveness and the Quantum Three Box Paradox. *arXiv:1207.3114* (2012).
- [106] Leifer, M. S. & Spekkens, R. W. Pre- and Post-Selection Paradoxes and Contextuality in Quantum Mechanics. *Phys. Rev. Lett.* **95**, 200405 (2005).

- [107] Leifer, M. S. & Spekkens, R. W. Logical pre- And post-selection paradoxes, measurement-disturbance contextuality. *Int. J. Theor. Phys.* **44**, 1977–1987 (2005).
- [108] Aharonov, Y. & Vaidman, L. Complete description of a quantum system at a given time. *J. Phys. A* **24**, 2315–2328 (1991).
- [109] Vaidman, L. Weak-measurement elements of reality. *Found. Phys.* **26**, 895–906 (1996).
- [110] Aharonov, Y., Albert, D. & Vaidman, L. How the result of a measurement of a component of the spin of a spin-1/2 particle can turn out to be 100. *Phys. Rev. Lett.* **60**, 1351–1354 (1988).
- [111] Wiseman, H. M. Grounding Bohmian mechanics in weak values and bayesianism. *New J. Phys.* **9**, 165–165 (2007).
- [112] Kocsis, S., Braverman, B., Ravets, S., Stevens, M. J., Mirin, R. P., Shalm, L. K. & Steinberg, A. M. Observing the Average Trajectories of Single Photons in a Two-Slit Interferometer. *Science* **332**, 1170–1173 (2011).
- [113] Dressel, J., Malik, M., Miatto, F. M., Jordan, A. N. & Boyd, R. W. Colloquium: Understanding quantum weak values: Basics and applications. *Rev. Mod. Phys.* **86**, 307–316 (2014).
- [114] Dressel, J. & Jordan, a. N. Contextual-value approach to the generalized measurement of observables. *Phys. Rev. A* **85**, 022123 (2012).
- [115] Wu, S. & Li, Y. Weak measurements beyond the Aharonov-Albert-Vaidman formalism. *Phys. Rev. A* **83**, 052106 (2011).
- [116] Bub, J. & Brown, H. Curious properties of quantum ensembles which have been both preselected and post-selected. *Phys. Rev. Lett.* **56**, 2337–2340 (1986).
- [117] Jozsa, R. Complex weak values in quantum measurement. *Phys. Rev. A* **76**, 044103 (2007).
- [118] Ferrie, C. & Combes, J. Classical correlation alone supplies the anomaly to weak values. *arXiv:1410.8067* (2014).
- [119] Bliokh, K. Y., Bekshaev, A. Y., Kofman, A. G. & Nori, F. Photon trajectories, anomalous velocities and weak measurements: a classical interpretation. *New J. Phys.* **15**, 073022 (2013).
- [120] Ferrie, C. & Combes, J. How the Result of a Single Coin Toss Can Turn Out to be 100 Heads. *Phys. Rev. Lett.* **113**, 120404 (2014).
- [121] Karanjai, A., Cavalcanti, E. G., Bartlett, S. D. & Rudolph, T. Weak values in a classical theory with an epistemic restriction. *New J. Phys.* **17**, 073015 (2015).

- [122] Pusey, M. F. Anomalous Weak Values Are Proofs of Contextuality. *Phys. Rev. Lett.* **113**, 200401 (2014).
- [123] Dressel, J. Weak values as interference phenomena. *Phys. Rev. A* **91**, 032116 (2015).
- [124] Goggin, M. E., Almeida, M. P., Barbieri, M., Lanyon, B. P., O’Brien, J. L., White, a. G. & Pryde, G. J. Violation of the Leggett-Garg inequality with weak measurements of photons. *Proc. Natl. Acad. Sci. USA* **108**, 1256–61 (2011).
- [125] Williams, N. S. & Jordan, A. N. Weak Values and the Leggett-Garg Inequality in Solid-State Qubits. *Phys. Rev. Lett.* **100**, 026804 (2008).
- [126] Hosten, O. & Kwiat, P. Observation of the spin hall effect of light via weak measurements. *Science* **319**, 787–790 (2008).
- [127] Dixon, P. B., Starling, D. J., Jordan, A. N. & Howell, J. C. Ultrasensitive Beam Deflection Measurement via Interferometric Weak Value Amplification. *Phys. Rev. Lett.* **102**, 173601 (2009).
- [128] Ferrie, C. & Combes, J. Weak Value Amplification is Suboptimal for Estimation and Detection. *Phys. Rev. Lett.* **112**, 040406 (2014).
- [129] Knee, G. C., Combes, J., Ferrie, C. & Gauger, E. M. Weak-value amplification: state of play. *Quantum Measurements and Quantum Metrology* **3**, 1–7 (2016).
- [130] Viza, G. I., Martínez-Rincón, J., Alves, G. B., Jordan, A. N. & Howell, J. C. Experimentally quantifying the advantages of weak-value-based metrology. *Phys. Rev. A* **92**, 032127 (2015).
- [131] Zhu, X., Zhang, Y., Pang, S., Qiao, C., Liu, Q. & Wu, S. Quantum measurements with preselection and postselection. *Phys. Rev. A* **84**, 052111 (2011).
- [132] Combes, J., Ferrie, C., Jiang, Z. & Caves, C. M. Quantum limits on postselected, probabilistic quantum metrology. *Phys. Rev. A* **89**, 052117 (2014).
- [133] Jordan, A. N., Martínez-Rincón, J. & Howell, J. C. Technical Advantages for Weak-Value Amplification: When Less Is More. *Phys. Rev. X* **4**, 011031 (2014).

CHAPTER 4

On the Reality of the Wavefunction

Acknowledgement

The work that forms the basis of this chapter was first published in Ref. [1]. I have incorporated text of that paper, where appropriate. The experiments were performed with Benjamin Duffus and also contributed to the research project in his Honours degree at the University of Queensland. The quantum states and measurements required for the experiment were computed by Cyril Branciard.

4.1 Introduction

“Do you really believe the moon exists only when you look at it?” - A. Einstein [2]

EINSTEIN’S famous question highlights how the development of quantum mechanics required a radical rethinking of many beloved concepts of classical physics. In classical physics, systems have objective properties and the moon is there even if we are not looking. The behaviour of quantum systems, on the other hand, challenges this absolute objective reality. Quantum systems do not adhere to traditional cause-and-effect relations (see Chap. 5), they are necessarily disturbed by a measurement (see Chap. 6), and the measurement outcome depends on what other measurements are made as well (see Sec. 3.3). It seems that a measurement is not merely revealing a pre-existing reality anymore, but has a role to play in creating the measurement outcome. In fact, it is not even clear what it means to talk about the reality of a quantum system.

Much of this debate revolves around the quantum wavefunction and the so-called *measurement problem*. The wavefunction, or *quantum state*¹ ψ is the central object of the theory and it is used to describe and make predictions about any kind of quantum system. The wavefunction is ubiquitous as a mathematical tool, and is used across the fields, from quantum chemistry to molecular dynamics in biological processes, yet we don’t know what it actually represents. According to the rules of quantum mechanics, the wavefunction evolves with time in a continuous and deterministic way, until a measurement causes a spontaneous

¹In the following these terms are used interchangeably, as are $|\psi\rangle$ and ψ unless there is potential for confusion.

and probabilistic collapse. Yet, quantum mechanics does not provide a rule for deciding which prescription should be used in a given situation, which is known as the measurement problem and is best illustrated at the example of Wigner’s friend, playing with Schrödinger’s cat.

Schrödinger’s cat is put into a box together with a device that may or may not kill the cat. After a while, the state of the cat will have evolved from alive to a *quantum superposition* of dead and alive at the same time. When Wigner’s friend opens the box, however, she will not find the cat in this state, but instead quantum mechanics predicts that she will have equal chance of finding the cat dead and finding it alive. Wigner, who is outside the laboratory, on the other hand, has not performed a measurement himself, and must therefore describe the joint state of the cat and his friend as a superposition of the cat being dead with the friend seeing a dead cat, and the cat being alive with the friend finding the cat alive. The inconsistency of the description given by Wigner and his friend is what many interpretations of quantum mechanics aim to resolve. The way they interpret the wavefunction itself is the key difference between these interpretations.

The many worlds interpretation of quantum mechanics, for example, always takes the outside perspective, and considers the un-collapsed wavefunction with its intricate superposition and entanglement structure as a literal description of the world. A related interpretation, known as Bohmian mechanics, also postulates that the wavefunction never collapses, but in contrast to many worlds, interprets it as a field that guides particles through the world. On the other end of the spectrum of realities are objective collapse models, which modify quantum mechanics slightly to put a stochastic wavefunction collapse term directly into the Schrödinger equation. This implies that above a certain level of complexity it is not possible to sustain superpositions of macroscopic objects for significant amounts of time. Wigner’s friend would thus never end up in an entangled state in the first place, but rather in a statistical mixture of all possibilities. Somewhat more radical, one could give up the notion of objective reality altogether, and consider the quantum state that Wigner assigns as purely a mathematical tool for making predictions about the outcomes of potential measurements he could make on his friend and the cat, rather than as a description of the actual state of affairs. This dodges the collapse problem, since a wavefunction collapse is not a physical process anymore but rather an update of probabilities in light of new evidence. This “shut up and calculate” doctrine of the Copenhagen interpretation, as Mermin famously put it [3], works very well, but without objective reality one loses the notion of explaining observations, which for many is too large a price to pay.

This chapter is devoted to the nature of the quantum state and to an experimental approach that narrows the wide range of available interpretations a little. Section 4.1.1 motivates this question and provides necessary background, which is formalized in the ontological models framework, introduced in Sec. 4.2. Section 4.3 is dedicated to various no-go theorems for the ψ -epistemic viewpoint, where the wavefunction represents limited knowledge about an underlying objective reality, and why they are not completely no-go after all. Section 4.4 then discusses how it is nonetheless possible to put constraints on all ψ -epistemic models, and

Sec. 4.5 presents an experiment demonstrating that ψ -epistemic models cannot fully reproduce quantum predictions, based on the publication [1]. Section 4.6 finishes with a discussion of the alternatives that are left after the experiment and an outlook for future research. The presentation here necessarily leaves out some of the mathematical details and related research, and the interested reader is referred to an excellent review by Matthew Leifer [4].

4.1.1 Ontic or Epistemic

At the core of Einstein’s question is the desire for a realistic theory of the world, where *realism* in this context means that there is an objective description of the world that is independent of our observation². Within such a theory, one can distinguish *ontic* states (from the Greek word ὄν , ‘being’), or states of reality, and *epistemic* states (from the Greek word ἐπιστήμη , ‘knowledge’), or states of knowledge. Ontic states represent objective, observer-independent properties of a physical system, whereas epistemic states correspond to a probability distribution over a set of ontic states, indicating incomplete knowledge of the actual state of the system [5].

This distinction between states of reality and states of knowledge is ubiquitous and in no way unique to quantum mechanics. Consider, for example a single classical particle moving on a line in one dimension [4, 5]. According to classical mechanics, this particle is completely described by a point in phase space, which is a pair of position and momentum values—the ontic state of the system. At any time t , the particle has a well-defined (objective) ontic state $(x(t), p(t))$, see Fig. 4.1a. Typically, however, complete knowledge about the actual ontic state (or phase-space point) of the system is not available. The system can then be described by a probability distribution over the set of ontic states—an epistemic state, see Fig. 4.1b. This epistemic state can still be used to make predictions about the system. However, the value that the distribution assigns to every ontic state has no physical significance. It merely represents the knowledge that an agent has about the system.

4.1.2 The Quantum Wavefunction

The above discussion about ontic and epistemic states has great significance in the question what kind of object the quantum state is. On the one hand, a large fraction of interpretations treat the wavefunction as an ontic object, on the other hand, treating it as a state of knowledge elegantly avoids some of the conceptual issues of interpretations of quantum mechanics. For example, if the wavefunction was ontic, a measurement would, in the standard formalism, lead to a discontinuous change of objective reality. If the wavefunction was an epistemic state, on the other hand, the measurement merely leads to a discontinuous change in the observer’s knowledge in light of new experimental evidence. Indeed, classical theories which impose restrictions on how much an observer can know (a so-called *epistemic restriction*), can reproduce a range of quantum phenomena, including uncertainty relations, no cloning, teleportation and Gaussian

²The concept of realism has many more facets in the philosophy community, but the definition used here is widely used in foundational physics.




Figure 4.1: **A classical point particle in position-momentum phase-space.** (a) At time t the particle occupies a well-defined ontic state $(x(t), p(t))$ which completely specifies its properties (b) If complete knowledge about the actual ontic state of the system is not available, then the system is described by an epistemic state—a probability distribution over the space of ontic states (here phase-space). The epistemic states can be used to make statistical predictions about the particle’s behaviour.

quantum mechanics [4–7]. Moreover, mixed quantum states are generally viewed as epistemic objects, since they can be prepared as statistical mixtures of pure states (without recording which state was prepared in each run) [5].

The distinction between ontic and epistemic interpretations of the wavefunction is made precise in the ontological models framework, that will be introduced in Sec. 4.2. This allows us to roughly distinguish three main classes of models with respect to the status they attribute to the quantum wavefunction

- *Realist ψ -ontic*: The wavefunction is in direct correspondence with the ontic state of the system, which might additionally include other ontic degrees of freedom.
- *Realist ψ -epistemic*: The quantum state represents knowledge or information about the ontic states of the system. Quantum mechanics is a statistical theory of these epistemic states in the spirit of classical Liouville mechanics with the restriction that complete knowledge of the ontic states is impossible.
- *Operationalist*: The quantum state represents an agent’s knowledge or degree of belief, not about the real state of the quantum system, but only about future measurement outcomes. In these models there is no need for a deeper reality (i.e. no ontic states).

Although this distinction is based on the very general ontological models frameworks, which is formulated with only a minimal set of assumptions, not all interpretations of quantum mechanics fit into this picture. Interpretations with more exotic ontologies that do not fit into the framework will be discussed briefly in Sec. 4.1.3 and Sec. 4.6.1.

ψ -Ontic Interpretations

“No one can understand this theory until he is willing to think of Ψ as a real objective field rather than just a ‘probability amplitude’.” - J.S. Bell [8]

In the ψ -ontic view the quantum state is in direct correspondence with the objective reality of the system. This either means that the wavefunction *is* the ontic state of the system, (i.e. the phase-space point in Fig. 4.1a), or more generally, that it is part of the objective reality, together with other relevant parameters. In either case, changing the wavefunction implies a change of the ontic state (the reality) of the system. Hence the wavefunction itself is attributed a real status. Bell himself proposed a ψ -ontic interpretation with ontic states consisting of ψ together with an additional continuous parameter [9]. While the ψ -ontic view is quite popular, it does entail a number of challenges, such as the measurement problem, Bell-nonlocality and ontological excess baggage, which will be discussed in Sec. 4.2.3

ψ -Epistemic Interpretations

“[...] wave fields were not to be interpreted as a mathematical description of how an event actually takes place [...], rather they are a mathematical description of what we can actually know about the system.” - A. Einstein [10]

Einstein held firm the believe that the wavefunction cannot be a complete description of the reality of a quantum system, and moreover should only represent an experimenter’s limited knowledge about the real state of the system [11]. In the realist ψ -epistemic view the quantum state thus describes the probability of finding the system in a certain ontic state, similar to Liouville mechanics, except that the ontic states are typically not even in principle empirically accessible. The fact that classical statistical mechanics with the right restriction can demonstrate similar behaviour as quantum systems is indeed one of the main motivations for this class of interpretations [7, 12, 13]. The realist ψ -epistemic view (possibly not exactly in the way Einstein imagined it) thus offers very intuitive explanations of a wide range of quantum phenomena [5].

Anti-Realist Interpretations

“It is wrong to think that the task of physics is to find out how nature is. Physics concerns what we can say about nature.” - N. Bohr [14]

At first sight Bohr’s view seems very similar to Einstein’s. Both agree that the wavefunction should not directly describe the reality of a quantum system. In contrast to Einstein, however, Bohr rejects the need for such an objective reality altogether. In this sense, interpretations in this spirit are often called anti-realist³ and they ascribe a key role to agents and their interaction with the world. The purpose of the theory in such an interpretation is generally to help the agent make predictions about future measurement outcomes. Without objective reality, the wavefunction is necessarily an epistemic object, and these interpretations thus resolve conceptual issues such as the wavefunction collapse in a similar way as realist ψ -epistemic

³The term anti-realist has a somewhat negative flavour to it, but it is simply meant to capture the fact that these interpretations reject the notion of objective observer-independent realism. Recently, the term *participatory realism* has been used to emphasize the role that agents play in these interpretations [15]. However, the kind of “realism” featured in these interpretations is certainly non-standard and to what extent it fits into a rigorous definition of realism will be up to philosophy to decide.

interpretations. Moreover, by considering measurement as a primitive, they tend to deny the need for an explanation to the other question of the measurement problem, (what counts as a measurement). However, it seems rather unsatisfactory for our best physical theory to not say anything about the actual physics.

4.1.3 Interpretations of Quantum Mechanics

“Quantum theory is the most useful and powerful theory physicists have ever devised. Yet today, nearly 90 years after its formulation, disagreement about the meaning of the theory is stronger than ever. New interpretations appear every day. None ever disappear” - D. Mermin [16]

Today, there is a large number of active interpretations of quantum mechanics, which enjoy varying popularity in different communities [17–19]. They span a diverse range of physical narratives, yet building on ideas from Refs. [4, 15], they can be roughly classified according to the status they attribute to the quantum wavefunction, see Fig. 4.2. Below is a brief summary of a number of interpretations, which are also to some extent interrelated [20].

Copenhagen Interpretation. The Copenhagen, or *orthodox* interpretation of quantum mechanics is a loosely defined operational interpretation, based on the idea that the theory should only make predictions for future measurement outcomes, and that physical systems do not have well-defined properties before the measurement [21].

Ensemble Interpretation. The quantum state is a purely epistemic description of the physical properties of an ensemble of systems and has no meaning for individual systems [22, 23].

QBism. In Quantum Bayesianism or QBism quantum state represents an agent’s subjective degree of belief about the probabilities of future events [24–26].

Information-Theoretic Interpretation. The central feature of quantum theory is that the information content of a quantum system is fundamentally limited and the quantum state represents a proposition about the system, rather than physical reality [27–29].

Relational Quantum Mechanics. The quantum state only describes the relation between two quantum systems or the information that one system has about another system [30].

Bohmian mechanics. In this theory, also known as the de Broglie-Bohm “pilot-wave” theory, quantum systems are considered to behave as classical-like particles with definite, objective positions, which are deterministically guided by an ontic universal wavefunction [31, 32].

Collapse Models. Collapse models are ψ -ontic and consider the collapse a real, physical process and postulate a modification of the dynamics of quantum mechanics by introducing a

Realist		ψ -ontic	ψ -epistemic
		Dirac-von Neumann Bohmian mechanics Collapse models Modal interpretations Bell	Ensemble Interpretation (Einstein, Ballentine) Spekkens Toy Model
Ontological Models	Other	Transactional Two-state vector Many Worlds	Consistent Histories*
		Anti-Realist	
		Copenhagen (Bohr) Information-theoretic QBism Relational interpretation Brukner	

Figure 4.2: **Classification of interpretations of quantum mechanics with respect to the status they attribute to the wavefunction.** The ontological models framework captures all interpretations which feature an objective, observer-independent reality, where measurement outcomes are unique single events that are not relative to anything and there is backwards in time causation. Realist interpretations outside this framework thus either feature retrocausal influences (such as the transactional interpretation and the two-state vector formalism), or propose that more than one measurement outcomes is realized in a single experiment (such as many worlds). Anti-realist interpretations on the other hand are necessarily ψ -epistemic and interpret the wavefunction as an observer-dependent object that allows an agent to make predictions about future measurement outcomes. *The consistent histories formalism is somewhat open to interpretation itself and could be seen as a retrocausal ψ -epistemic interpretation, as a version of many worlds, or as an operational approach.

stochastic collapse term to the Schrödinger evolution, which becomes relevant only for large systems [33–35].

Everett “Many-Worlds” Interpretation. The many-worlds interpretation takes the wavefunction and with it the superposition principle as a literal description of the world, where in some sense every possible measurement outcome is instantiated [36].

Consistent Histories. The quantum state is an epistemic object that is used to assign probabilities to various alternate histories of a quantum system [37, 38].

Modal Interpretations. This name refers to a class of realist ψ -ontic models where the quantum state is regarded as a dynamical state that contains the full range of possible physical properties of a quantum system. The dynamical state always evolves according to the

Schrödinger equation (i.e. does not collapse), and is supplemented by a value state that represents the actual values of these physical properties, which are well-defined at any time [39, 40].

Transactional interpretation. Building on the Wheeler-Feynman absorber theory this interpretation identifies the evolution of a quantum system with a standing wave composed of a retarded wave from the source and an advanced wave from the detector, both of which are given an ontic status [41, 42].

Two-State Vector Formalism. Similarly to the transactional interpretation, the two-state vector formalism considers an ontology consisting of a state vector for the pre-selection (preparation) and another state vector for the post-selection (measurement) [43].

Retrocausal Interpretations. Besides the transactional interpretation and the two-state vector formalism, other retrocausal, or causally symmetric interpretations, motivated by time symmetry of the underlying equations, have been proposed, mainly as a reaction to Bell's theorem [44–46].

Many Interacting Worlds. Quantum mechanics arises through repulsive interaction between a large number of classical worlds where systems have real objective properties, but the wavefunction does not play a primary role [47].

Brukner. Brukner recently suggested an interpretation which is based on the premise that there are no objective facts about the world, but any record is only meaningful relative to some observer. This interpretation combines elements from a variety of generally anti-realist interpretations [20].

Quantum Darwinism. Strictly speaking this is not an interpretation, but rather a mechanism that explains how classical physics emerges out of a fundamentally quantum description through decoherence and redundant information encoding in the environment. However, Zurek emphasizes that the quantum state should be viewed as “mostly information”, placing this approach close to the Copenhagen camp [48].

Just a Matter of Taste?

The debate about interpretations of quantum mechanics is often dismissed as a matter of philosophy, without practical consequences. A notable exception, of course, are collapse models, which are alternative theories that make predictions significantly different from quantum mechanics, and indeed have already been experimentally constrained [49]. However, even those aside, the remaining contenders are far from untestable. As discussed in more detail below, already Einstein pointed out that interpretations in which the wavefunction is the complete description of physical reality are in conflict with locality assumptions. This was later strengthened by Bell, who showed that any realist interpretation of quantum mechanics was in conflict

with local causality, see Chap. 5. Bell’s theorem had a large impact in quantum information theory and is now the basis for quantum cryptography.

For interpretations of quantum mechanics, however, Bell’s theorem is a very general result that poses a challenge, but can be accommodated in most interpretations. These interpretations feature a vast variety of physical narratives and conceptually different approaches. This shapes the ideas derived from them. ψ -ontic and ψ -epistemic interpretations, for example, suggest quite different behaviour in the classical limit [4]. All of this makes it seem rather unlikely that there should be no way to distinguish them in practice [15]. Indeed, the recent years have seen a revival of interest in the interpretational question, with a number of results on classes of realist ψ -epistemic interpretations [50–54], classes of realist interpretations in general [55] and “single-world” interpretations [56]. Although all of these results rely on a number of assumptions and address only subclasses of the targeted interpretations they clearly show that interpretations of quantum mechanics are certainly more “experimental metaphysics” [57] than pure philosophy.

4.2 The Ontological Models Framework

The formal basis for discussing realist interpretations of quantum mechanics in the traditional sense is the framework of ontological models first introduced in Ref. [11]. This framework is based on a minimal set of assumptions about what can be considered objectively real. These basic assumptions are also the basis of Bell’s theorem, which, using a few additional assumptions, ruled out “local hidden variable” models (cf. Chap. 5), which are an important special case of what is captured by the ontological models framework. The term “hidden variable”, however, is rather unfortunate and overly suggestive, since the ontic states in general need not be hidden or might not even be any different from the wavefunction itself. The following discussion of the framework focuses on the essential aspects, using pure states and projective measurements, referring the reader to Ref. [4, 11, 58] for the general treatment (including mixed states and POVMs) and mathematical subtleties.

The central assumption of the framework, realism, is captured in the existence of a set of ontic states that completely specify the system. The quantum state merely represents limited information about the ontic state occupied by the system. The ingredients can be summarized as follows, see Fig. 4.3:

- i) There exist a set Λ of ontic states λ , which completely specify the system.
- ii) Every preparation of a quantum state $|\psi\rangle$ actually prepares the quantum system in an ontic state λ , sampled from a (classical) probability density $\mu_\psi(\lambda)$ over the set of ontic states—the epistemic state
- iii) For each measurement M with outcomes $\{m\}$ the observed measurement statistics depend *only* on the ontic state via a *response function* $\xi_M(m | \lambda)$, but not on the quantum state

$|\psi\rangle$. The conditional probability for outcome m when measuring M on $|\psi\rangle$ is given by

$$P_M(m | \psi) = \int_{\Lambda} \xi_M(m | \lambda) \mu_{\psi}(\lambda) d\lambda . \quad (4.1)$$

iv) Additionally the model should reproduce quantum predictions (i.e. Born rule probabilities) for the measurement M with eigenbasis $\{|m\rangle\}$

$$P_M(m | \psi) = |\langle m | \psi \rangle|^2 . \quad (4.2)$$

While the above requirements are formulated for models of quantum mechanics, the ontological models framework can also be used to describe ontological models for general operational theories. Besides realism the ontological models framework implicitly assumes a temporal order of events and a causal arrow with respect to this order. Otherwise the ontic state could not screen off the measurement outcome from the quantum state as in Eq. (4.1). These assumptions are also shared with Bell's theorem, which, however, relies on additional assumptions about the causal relation of two separate quantum systems, see Sec. 5.3.2. The ontological models framework, on the other hand, is primarily concerned with prepare-and-measure experiments on single quantum systems, where no locality questions arise. Everything more complicated requires additional assumptions, such as rules for state-update upon measurement, when considering state-evolution within this framework.

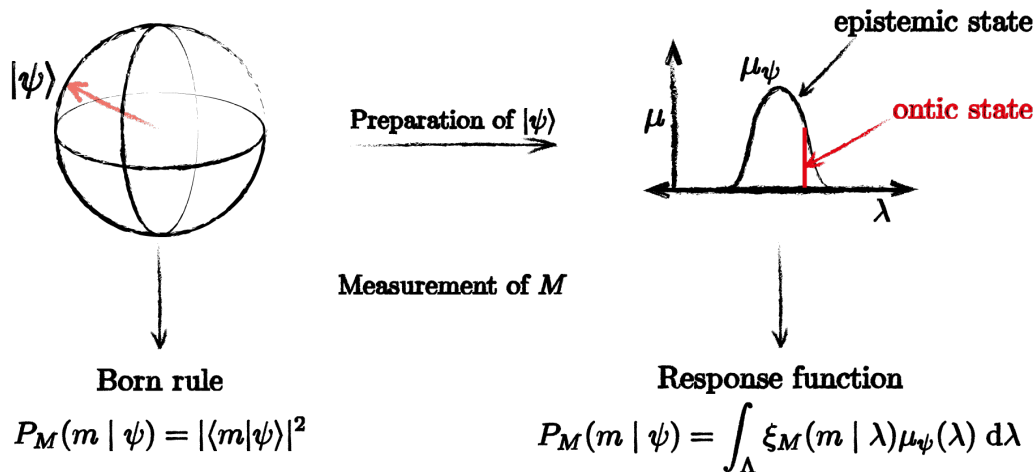


Figure 4.3: **The ontological models framework.** Every preparation of the quantum state $|\psi\rangle$ (here shown in the simplest form of a qubit on the Bloch sphere) actually prepares the system in an ontic state λ , chosen at random from a classical probability distribution μ_{ψ} in the ontic state space—the epistemic state. The probabilities for each outcome $\{m\}$ of the measurement M are specified by the response function ξ_M and depend only on λ . For a model to be viable, the response function must reproduce Born-rule probabilities when averaging over lambda.

The ontological models framework is thus general enough to allow for

- **Measurement contextuality.** Ontological models can be measurement contextual, since the response function can depend on the full POVM M , rather than just the measurement outcome m .
- **Preparation contextuality.** Ontological models can be preparation contextual, and do not require that the preparation of a mixed state ρ corresponds to a unique measure μ_ρ , see Ref. [4] for details.
- **Indeterminacy.** Ontological models do not assume predetermination at the ontic level, and measurement outcomes can result stochastically from the response function.
- **Arbitrary ontic state space.** There are no assumptions on the form and structure of the ontic states and the ontic state space, other than that it is a measurable space.

Rigorous definition and subtleties Leifer, in Ref. [4], rigorously defines an ontological model as a set $(\Lambda, \Sigma, \Delta, \xi)$, where Λ is a measurable set with σ -algebra Σ . The function Δ maps every quantum state ρ to a set Δ_ρ of σ -additive probability measures $\mu: \Sigma \rightarrow [0, 1]$ that are non-negative ($\mu_\psi(\lambda) \geq 0$) and normalized ($\mu(\Lambda) = 1$). The function ξ maps every outcome m of a POVM M to a measurable response function $\xi_M(m | \cdot): \Lambda \rightarrow [0, 1]$ that is non-negative ($\xi_M(m | \lambda) \geq 0$) and normalized ($\sum_m \xi_M(m | \lambda) = 1 \quad \forall M, \forall \lambda$). The model reproduces quantum predictions if for every quantum state ρ

$$\int_{\Lambda} \xi_M(m | \lambda) d\mu(\lambda) = \text{Tr}[m\rho] \quad \forall M, \forall m \in M, \forall \mu \in \Delta_\rho. \quad (4.3)$$

Notably, this definition allows for both measurement contextuality (via the dependence of ξ on M , not just m) and preparation contextuality (via the set Δ_ρ).

4.2.1 The ψ -Ontic/ ψ -Epistemic Distinction

In any model within the ontological models framework (ψ -ontic as well as ψ -epistemic), the quantum state is a representation of limited information about the underlying ontic state of the system; a coarse-grained picture of the real state of the system, which potentially lacks some of the detail contained in the ontic states. The difference between ψ -ontic and ψ -epistemic is that these coarse-grained patches are non-overlapping for a ψ -ontic model, but they might overlap for a ψ -epistemic one, see Fig. 4.4.

More specifically, in a ψ -ontic model every ontic state belongs to a unique quantum state. Consequently, any pair of non-identical (pure) quantum states $|\psi\rangle$ and $|\phi\rangle$ is associated with a pair of *disjoint* epistemic states μ_ψ and μ_ϕ (in the general case every pair of measures from the sets Δ_ψ and Δ_ϕ are disjoint), see Fig. 4.4b. In a ψ -ontic model, a change in the quantum state thus necessarily comes with a change in the ontic state, and the quantum state is therefore said to be “in direct correspondence with reality”. The converse⁴, however, need not be true

⁴This is the “completeness” question raised by Einstein, whether a quantum state is associated to a unique “element of reality” (i.e. ontic state) [59]. Models which feature such a one-to-one correspondence between




Figure 4.4: **The ψ -ontic/ ψ -epistemic distinction.** (a) Two non-orthogonal quantum states $|\psi\rangle$ and $|\phi\rangle$ on the Bloch sphere. (b) In a ψ -ontic interpretation the epistemic states are disjoint for any pair of non-orthogonal quantum states. The quantum state is uniquely determined by the ontic state. (c) In a ψ -epistemic model the probability distributions for two non-orthogonal states may overlap and a single ontic state may be compatible with multiple quantum states.

in general and a single quantum state can be compatible with multiple ontic states. Moreover, these ontic states can have a complicated internal structure, and the number of ontic states per quantum state need not be constant either.

On the other hand, a model is ψ -epistemic if there is at least one pair of non-orthogonal quantum states, whose epistemic states overlap, see Fig. 4.4c. This definition is much less stringent than that of a ψ -ontic model and there are thus some important sub-classes. Most notably, *maximally ψ -epistemic* models, where the overlap of epistemic states for any two non-orthogonal states is large enough to provide a mechanism that fully explains why such quantum states cannot be perfectly distinguished, as a result of a lack of information about the underlying ontic states.

4.2.2 Distinguishing Quantum States

Consider two non-orthogonal quantum states $|\psi\rangle$ and $|\phi\rangle$, which are assigned distinct state-vectors, yet no single measurement can perfectly distinguish them in practice⁵. This curious feature of quantum mechanics has an elegant explanation in a ψ -epistemic model, which was nicely illustrated in Ref. [60] using a classical example, which is reproduced in Fig. 4.5. When the two epistemic states μ_ψ and μ_ϕ overlap, then a preparation of either state sometimes leaves the system in the same ontic state. In all such cases there is no measurement that could distinguish the two preparations.

The intuitive explanation of the indistinguishability of quantum states in terms of epistemic overlap is indeed one of the main motivations for ψ -epistemic models. Moreover, it gives an operational interpretation to the epistemic overlap as the residual indistinguishability when

quantum states and ontic states are thus called *ψ -complete*, and are a subclass of ψ -ontic models.

⁵To avoid too cluttered writing, this imperfect distinguishability of non-orthogonal quantum states is often referred to as “indistinguishability”.

all available (accessible and hidden) information is taken into account. To make this more precise, and quantify *how much* of the observed indistinguishability can indeed be explained in this way it is necessary to define measures of distinguishability of quantum states and classical probability distributions. The *quantum overlap* $\omega_Q(|\psi\rangle, |\phi\rangle)$ of the states $|\psi\rangle$ and $|\phi\rangle$ can be defined as [58, 60, 61]:

$$\begin{aligned}\omega_Q(|\psi\rangle, |\phi\rangle) &:= 1 - \delta_Q(|\psi\rangle, |\phi\rangle) \\ &= 1 - \sqrt{1 - |\langle\psi|\phi\rangle|^2},\end{aligned}\tag{4.4}$$

where $\delta_Q(|\psi\rangle, |\phi\rangle) = \sqrt{1 - |\langle\psi|\phi\rangle|^2}$ denotes the quantum trace-distance, a distance measure for quantum states, see Sec. 1.1.3. Since the trace distance can be directly generalized to mixed states, the measure ω_Q can be too. For two classical probability distributions one can use the classical equivalent of the trace-distance, the variational distance $\delta_C(\mu_\psi, \mu_\phi) = \frac{1}{2} \int_\Lambda |\mu_\psi(\lambda) - \mu_\phi(\lambda)|$ to define the *classical overlap*⁶ of the corresponding probability distributions or epistemic states:

$$\begin{aligned}\omega_C(\mu_\psi, \mu_\phi) &:= 1 - \delta_C(\mu_\psi, \mu_\phi) \\ &= \int_\Lambda \min[\mu_\psi(\lambda), \mu_\phi(\lambda)] d\lambda.\end{aligned}\tag{4.5}$$




Figure 4.5: **The ψ -epistemic explanation of imperfect distinguishability of non-orthogonal quantum states.** This example is reproduced from Barrett et al. [60]. **(a)** Consider a red and a blue deck of cards, two distinct “states”, where the red deck only contains cards of hearts and the blue deck only contains aces. Drawing a card from either deck at random corresponds to preparing the respective state and the face of the card represents the ontic state of the system after the preparation. **(b)** The probability distributions associated with the preparation of the two states overlap at the ace of hearts. Whenever a preparation results in this ontic state it is not possible to tell which deck it was drawn from. Hence, the two decks cannot be distinguished perfectly from just the ontic states.

These measures are operationally motivated. The probability of success for distinguishing the two quantum states $|\psi\rangle$ and $|\phi\rangle$ using the optimal quantum measurement is $1 - \omega_Q(|\psi\rangle, |\phi\rangle)/2$.

⁶This nomenclature might be somewhat confusing, but recall that the epistemic states are for all practical purposes just classical probability distributions over some ontic state space Λ .

On the other hand, the probability of success for distinguishing the two states using the optimal measurement on the ontic level (i.e. using information about the actual ontic state, rather than just the quantum state) is given by $1 - \omega_{\text{C}}(\mu_{\psi}, \mu_{\phi})/2$. The classical and quantum measures are thus compatible in terms of their operational interpretation, and must satisfy

$$0 \leq \omega_{\text{C}}(\mu_{\psi}, \mu_{\phi}) \leq \omega_{\text{Q}}(|\psi\rangle, |\phi\rangle) \leq 1. \quad (4.6)$$

The first and last inequalities merely correspond to the normalization $0 \leq \omega_{\text{C}/\text{Q}} \leq 1$, while the middle inequality implies that the classical overlap cannot be larger than the quantum overlap. In other words, distinguishing two states based only on their quantum states cannot be easier than based on the additional information about the actual ontic states. If this inequality was violated then the two states would be less distinguishable in the model than according to quantum theory and the model could not reproduce quantum predictions. Using the relation between these overlap measures ontological models can be classified as

(i) ψ -ontic if (up to sets of measure zero)

$$\omega_{\text{C}}(\mu_{\psi}, \mu_{\phi}) = 0 \quad \forall |\psi\rangle \neq |\phi\rangle.$$

(ii) ψ -epistemic if

$$\exists |\psi\rangle \neq |\phi\rangle : \omega_{\text{C}}(\mu_{\psi}, \mu_{\phi}) \neq 0.$$

(iii) maximally ψ -epistemic if

$$\omega_{\text{C}}(\mu_{\psi}, \mu_{\phi}) = \omega_{\text{Q}}(|\psi\rangle, |\phi\rangle) \quad \forall |\psi\rangle, |\phi\rangle.$$

Maximally ψ -epistemic models saturate the middle inequality in Eq. (4.6), which means that in such models *all* the indistinguishability of non-orthogonal states can be explained in terms of overlapping classical probability distributions. At the other end of the spectrum are ψ -ontic models, which feature absolutely no overlap of the probability distributions. These models thus require a different mechanism to explain the imperfect distinguishability, commonly referred to as coarse-grained measurements. Just like quantum states might correspond to a coarse-grained preparation of ontic states, quantum measurements might only be able to reveal coarse-grained information about these ontic states. Indeed, coarse-grained measurements seems to be an common feature, albeit not necessarily involved in explaining indistinguishability, in most ontological models, even including maximally ψ -epistemic ones such as Spekkens' toy model [5].

All ontological models that do not fall in the two extreme categories are non-maximally ψ -epistemic, and can only explain a fraction of the indistinguishability in terms of overlapping probability distributions. As a figure of merit to quantify the relative amount of indistinguishability that a model can explain one can use the ratio of classical-to-quantum overlap

$$\kappa(\psi, \phi) = \frac{\omega_{\text{C}}(\mu_{\phi}, \mu_{\psi})}{\omega_{\text{Q}}(|\phi\rangle, |\psi\rangle)}, \quad (4.7)$$

for a pair of non-orthogonal quantum states (such that $\omega_Q(|\phi\rangle, |\psi\rangle) > 0$). In the limiting case of a ψ -ontic model, the overlap ratio $\kappa = 0$ for every pair of states, while for the other extreme, a maximally ψ -epistemic model $\kappa = 1$ for every pair of states. In any other case there is at least one pair of states such that $0 < \kappa < 1$. However, for such non-maximally ψ -epistemic models there need not be any clear relation between ω_C and ω_Q , other than the physical constraint Eq. (4.6).

Rethinking ψ -Epistemicity

The notion of ψ -epistemicity, defined as everything that is not ψ -ontic, is very permissive, and allows for some rather contrived models. Hence, a number of finer distinctions were introduced and discussed in detail Ref. [4].

Sometimes ψ -ontic models are models where each pure state has some “region of the ontic state space all to itself” [4]. Specifically, the support of each μ_ψ contains at least one λ that is not in the support of any other μ_ϕ . This notion is much stronger, excludes maximally ψ -epistemic models and, curiously, proving that ontological models must be sometimes ψ -ontic would imply much the same results as a proof that they must be ψ -ontic [4].

Pairwise ψ -epistemic models are ψ -epistemic for *any* pair of non-orthogonal quantum states. While this gets rid of some of the more contrived models it is also a very strong notion.

Symmetric ψ -epistemic models were introduced in Ref. [62], and are such that the overlap between two epistemic states depends only on the inner product of the corresponding quantum states. In other words, states with similar inner product have similar classical overlap in an appeal to rotational invariance.

4.2.3 Known Constraints

Although the formal ontological models framework is rather young, the study of realist interpretations of quantum mechanics has a long history. Going back to Einstein famously challenging whether ψ -complete interpretations could reproduce quantum predictions [11, 59], a number of individual results have ruled out specific subclasses of ontological models. Some well-known results for regarding ontological models include

- i) EPR [59] / Einstein [11] incompleteness argument.
- ii) All ontological models are Bell-nonlocal [63].
- iii) All ontological models are Kochen-Specker contextual (Bell-Kochen-Specker theorem [9, 64] and generalizations thereof [65]).
- iv) The size of the ontic state space for sometimes ψ -ontic models is continuously infinite, already for a single qubit [66–68].

All of these results were derived within the ontological models framework and are in fact closely related, see Ref. [4]. For example, Bell’s theorem generalizes the EPR argument, and Kochen-Specker contextuality generalizes Bell’s theorem. Moreover, many results implied by ψ -ontology

already follow from the impossibility of maximally ψ -epistemic models [4], which motivates the importance of this class of models for providing an intuitive classical explanation for quantum phenomena.

Incompleteness and the EPR Argument

EPR considered two parties, Alice and Bob, who share a maximally entangled quantum state. Depending on the measurement that Alice performs on her part of the shared state, she steers Bob's system to a state from different ensembles. Assuming separability of the ontic state space together with the assumption that the ontic state of Bob cannot depend on the choice of measurement at Alice (a locality assumption) [69, 70] Einstein concluded that there must be multiple quantum states compatible with a single ontic state.

Hence, there is a conflict between ψ -complete models and EPR's notion of locality. The latter, also referred to as the *no-telepathy* assumption [70] asserts that Bob's ontic state is independent of Alice's measurement setting. This is different from Bell's locality assumption, which is concerned instead with the independence of Bob's measurement outcome from Alice's measurement setting [71], see also Chap. 5. EPR's original assumptions were relatively strong, and phrased rather vaguely, see Ref. [70] for a detailed discussion. A cleaner, and also stronger, version of the argument was given by Einstein in his correspondence with Schrödinger, as discussed and formalized in Ref. [11]. This argument indeed suffices to establish a conflict with all ψ -ontic, rather than just ψ -complete models.

Bell-Nonlocality

Bell's theorem builds on the EPR argument, and demonstrates that *any* ontological model, ψ -ontic or ψ -epistemic, that reproduces quantum predictions must be Bell-nonlocal [63, 72]. Note that the notion of Bell-locality incorporates a series of assumptions, beyond the ontological models framework, which capture classical notions of locality and causality, see Chap. 5 or Ref. [73]. These additional assumptions are necessary to specify how two or more quantum systems interact, and they are formulated in terms of statistical or causal constraints. Crucially, however, Bell's theorem makes no assumptions on the structure of the ontic state space and thus applies to arbitrary ontological models.

Experimental violations of Bell inequalities imply that in order for ontological models to reproduce quantum predictions, they must violate the classical notion of locality and causality. Going beyond Bell-local models, it is shown in Chap. 5 that not even non-local ontological models, which allow superluminal communication of measurement outcomes, can reproduce quantum correlations.

Curiously, Ref. [4] showed that, using preparation contextuality of the maximally mixed state, an argument similar to that of Einstein [11, 59], is in fact sufficient to prove that no Bell-local ontological model can reproduce the predictions of quantum mechanics. The idea is to use a steering argument to reinterpret a Bell-experiment as a prepare-and-measure experiment for Bob's quantum system, where the preparation procedure is given by Alice's measurement.

A version of Bell’s local causality assumption then implies that the response function separates into a function that depends only on Alice’s measurement and one that depends only on Bob’s. The preparation results in a maximally mixed state in case Alice “forgets” her measurement outcome. Bell-locality now implies that the model must be preparation contextual for this state, which is a contradiction [4].

Contextuality

The premise of noncontextuality is that if two experimental procedures (preparations, transformations or measurements) are operationally equivalent (i.e. produce the same statistics in every situation) then they should be represented in the same way at the ontic level [65]. The discussion of contextuality in Sec. 3.3 can thus be interpreted as constraints on the possible ontological models.

One of the earliest results in this space is the (Bell-)Kochen-Specker theorem [9, 64], which showed that there can be no outcome-deterministic and measurement noncontextual models for $d \geq 3$. Viable models do, however, exist if either of these conditions is violated [4]. Using the operational definition of contextuality in Ref. [65] it was further shown (in theory and experiment) that quantum mechanics is also incompatible with the combination of measurement noncontextuality and preparation noncontextuality [74]. Preparation noncontextuality in fact implies outcome determinism for sharp measurements [65] and seems to be a more fundamental feature, since there are preparation contextual models that are measurement-noncontextual (by virtue of being outcome-indeterministic), but not vice versa [65].

Curiously, preparation noncontextual models must be maximally ψ -epistemic. This is based on the simple argument that in a preparation noncontextual model the epistemic state of a proper mixture of quantum states must be independent of the pure-state decomposition. Decomposing the maximally mixed state into two sets that involve a pair of states that are non-maximally ψ -epistemic thus implies that the corresponding epistemic states cannot be identical and the model is preparation contextual. Furthermore, a maximally ψ -epistemic model must be outcome-deterministic and measurement noncontextual, thus Kochen-Specker noncontextual. The latter conclusion follows from the fact that the response-function for a measurement of ϕ must be equal to 1 for every ontic state corresponding to ϕ , for every measurement containing ϕ , and thus vanish everywhere else, see Ref. [75] for details. Turning the above argument around, Kochen-Specker contextuality implies that there can be no maximally ψ -epistemic models, which in turn implies preparation contextuality for the maximally mixed state. The Kochen-Specker theorem, however is not enough to rule out non-maximally ψ -epistemic models [4].

Ontological Excess Baggage

The term *ontological excess baggage* was coined by Hardy to describe the observation that the size of the ontic state space must be infinite even for a qubit [66]. Since a single qubit $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ is specified by two continuous, complex-valued parameters α, β , there are uncountably infinitely many distinct quantum states for a qubit. Any ψ -ontic model that

reproduces the quantum predictions exactly⁷ thus requires continuously infinitely many ontic states, since there must be at least one ontic state per quantum state [66]. Indeed, it is enough for a model to be sometimes ψ -ontic to reach this conclusion [4]. For ψ -epistemic models, on the other hand, this might not be a serious problem, since the underlying ontic state space might be finite, even if specifying the exact probability distribution over this space requires continuous parameters [4].

Montina further showed that the size of the ontic state space must also grow exponentially with the number of systems [67]. Specifically, for a Hilbert space of dimension d , any ontological model with Markovian evolution must contain at least $2d - 2$ real parameters, and contain something that looks and behaves like a quantum state [68]. If the Markov assumption is not made, the ontic state space might be smaller, which they demonstrate by constructing a one dimensional model for a qubit.

4.2.4 Examples of Ontological Models

Part of the motivation for studying ontological models of quantum mechanics comes from the fact that there are numerous examples of such models in the literature. These were typically developed as a proof-of-principle to demonstrate that such models can reproduce certain aspects of quantum mechanics. Such examples are very valuable for understanding the consequences of giving up various assumptions in no-go theorems for ontological models, such as Bell’s theorem.

One of the earliest examples is Bell’s ψ -ontic model that he introduced to demonstrate that deterministic hidden-variable models can indeed reproduce quantum predictions [9]. The ontic state in this model consists of the quantum state and an additional continuous parameter, which essentially encodes the measurement outcome probabilities. The model is thus ψ -complete and, although originally formulated as a measurement-contextual model for orthogonal measurements on qubits, it has since been generalized to higher-dimensional systems, and can also be turned into a Kochen-Specker noncontextual model [4].

An indeterministic, ψ -complete version of Bell’s model is the Beltrametti-Bugajski model [76]. In contrast to Bell’s model the measurement outcome probabilities are not encoded in an ontic state using a deterministic response function, but instead obtained directly from an indeterministic response-function that directly implements the trace-operator. The model is thus measurement noncontextual, and since the ontic states are simply the pure quantum states, it is preparation contextual.

One of the most influential ontological toy models was introduced by Spekkens to demonstrate that a very simple (maximally) ψ -epistemic model for qubits can give intuitive explanations for a range of quantum phenomena [5]. The model reproduces the behaviour of a qubit for $\sigma_x, \sigma_y, \sigma_z$ -preparations and measurements using an ontic state space of only four states. The crucial aspect is a kind of uncertainty principle, called the *knowledge-balance principle*, which asserts that only “half” of the information about the ontic state of the system can be known

⁷If the model is only assumed to reproduce quantum predictions to within ε , it might be possible to relax this to a countable subset.

at any time. The model can be generalized to larger dimensions and, in odd dimensions, was found to correspond to the stabilizer formalism of quantum mechanics [4]. It has also been generalized to continuous variables, where it recovers Gaussian quantum mechanics [7].

Another interesting example is the Kochen-Specker model for orthonormal measurements on a single qubit, originally introduced as a counterexample to the Kochen-Specker theorem in two dimensions [64]. Again this model is maximally ψ -epistemic [4] and the ontic state-space is simply the Bloch-sphere. A pure quantum state ψ is associated with a probability measure in the hemisphere that has ψ at its pole, with a density proportional to the cosine of the angle between ψ and the corresponding ontic state λ . The response functions are deterministic functions of the angle between ontic state and measurement direction, 1 for all ontic states in the hemisphere with the measurement direction as its pole, and 0 otherwise, see Ref. [77] for more details.

4.3 ψ -Ontology Theorems

“The great tragedy of Science—the slaying of a beautiful hypothesis by an ugly fact [...]” - T. H. Huxley [78]

Realist ψ -epistemic models provide some of the most intuitive explanations for a range of quantum phenomena. Although there are no fully worked-out interpretations of this kind, the existence of toy models for Pauli measurements on qubits and other fractions of quantum mechanics motivates the question whether such models can be extended to all of quantum mechanics. For the longest time it was generally believed that questions about interpretations of quantum mechanics were of metaphysical nature and could not be rigorously tested. This view was challenged by a paper entitled “On the reality of the quantum state” by Pusey Barrett and Rudolph (hereafter PBR), proving a no-go theorem which showed that realist ψ -epistemic models are in conflict with quantum predictions [50].

However, in a follow-up work some of the same authors proved that their previous result could be avoided and that it is always possible to construct ψ -epistemic models [79]. This latter conclusion exploits the fact that the PBR theorem as well as all other no-go theorems for ψ -epistemic models [51–53, 62] necessarily rely on assumption beyond the ontological models framework, typically concerning the structure of the ontic state space. Hence, all current no-go theorems actually rule out various subclasses of ψ -epistemic models that satisfy their specific assumption. What was shown in Ref. [79] and also in Ref. [62] is that in the most general case, without additional assumptions, it is impossible to completely rule out ψ -epistemic models.

4.3.1 The Pusey-Barrett-Rudolph Theorem

PBR showed, with a deceptively simple argument that ψ -epistemic models make predictions that are in conflict with those of quantum mechanics. Although their result hinges on a few critical assumptions and idealizations, it spurred a burst of interest in the study of ontological

models for quantum mechanics. This section summarizes the PBR argument and the key points of the underlying assumptions. For a detailed discussion of these assumptions, and how they could be weakened the reader is referred to Ref. [4].

In a Nutshell

PBR consider two non-orthogonal quantum states $|\psi_1\rangle, |\psi_2\rangle$ and a two-outcome measurement $\{m_1, m_2\}$, such that $\langle m_i | \psi_i \rangle = 0$. Any viable ontological model must reproduce these zero probabilities, which means that the response function for measurement outcome m_i must vanish for all ontic states corresponding to ψ_i , that is $\xi(m_1 | \text{supp}(\mu_{\psi_2})) = \xi(m_2 | \text{supp}(\mu_{\psi_1})) = 0$. Since the response function ξ must be normalized for every ontic state, there can be no ontic state in the support of both ψ_1 and ψ_2 and thus the epistemic states μ_{ψ_1} and μ_{ψ_2} must be disjoint. This observation by itself is not particularly surprising, given that this kind of measurement (this is essentially the unambiguous state discrimination problem) only exists for orthogonal states, which by definition have to have disjoint epistemic states.

The important step is to now use two copies of the previous states, building the set $\{|\phi_1\rangle = |\psi_1\rangle \otimes |\psi_1\rangle, |\phi_2\rangle = |\psi_1\rangle \otimes |\psi_2\rangle, |\phi_3\rangle = |\psi_2\rangle \otimes |\psi_1\rangle, |\phi_4\rangle = |\psi_2\rangle \otimes |\psi_2\rangle\}$. These product states are formed under the *preparation independence* assumption. This assumption implies that the epistemic states of a product quantum state should factorize into the epistemic states of the individual systems, $\mu_{\psi_1 \otimes \psi_2}(\lambda) = \mu_{\psi_1}(\lambda_1)\mu_{\psi_2}(\lambda_2)$. Now consider a four-outcome measurement $\{m_1, m_2, m_3, m_4\}$, such that $\langle m_i | \phi_i \rangle = 0$. Using the same argument as before, if there is an ontic state in the support of both ψ_1 and ψ_2 , it must be in the support of all four ϕ_i , thus violating the normalization of the response function ξ . In the case where $|\psi_1\rangle = |0\rangle$ and $|\psi_2\rangle = |+\rangle$, the required measurement is simply a projection onto four orthogonal maximally entangled states, which can be implemented experimentally.

The remaining step is to show that this result holds for arbitrary non-orthogonal states $\psi_1/\psi_2 = \cos\theta|0\rangle \pm \sin\theta|1\rangle$. Since the proof only requires two different states it is sufficient to consider qubits. In general the authors then consider composite system of n copies, each prepared either in ψ_1 or ψ_2 , subject to a 2^n -outcome measurement M_n , such that each outcome occurs with zero probability for at least one of the 2^n possible joint states⁸. If the epistemic states μ_{ψ_1} and μ_{ψ_2} overlap, there is a certain probability that the preparation of the joint state yields an ontic states that is compatible with all 2^n quantum states, which leads to the same contradiction as before. Using this argument, epistemic overlap can be ruled out for arbitrarily close ψ_1 and ψ_2 , by choosing n arbitrarily large, and the authors construct an explicit measurement circuit for all n .

Experimental Prospects

Experimentally the rapid increase in n (and especially the asymptotic limit) is rather prohibitive. Limiting the the number of copies n is equivalent to restricting the result to pairs of states with sufficiently small quantum overlap. A more serious problem is that experimental

⁸If such a measurement exists the states are called *Post-Peierls*(PP)-incompatible [80].

noise means that the outcome probabilities never vanish exactly. Hence, one has to study the case where the model only reproduces quantum predictions to within ε . Based on the preparation independence assumption PBR showed that $\omega_C(\mu_{\psi_1}, \mu_{\psi_2}) \leq 2\sqrt[n]{\varepsilon}$ holds in this case [50]. In other words, in the presence of non-zero experimental uncertainty, it is not possible anymore to completely rule out ψ -epistemic models for the tested states. Instead the experiment can merely establish bounds on the absolute value of the classical overlap in ψ -epistemic models that satisfy preparation independence. This argument has allowed to experimentally rule out maximally ψ -epistemic models under the assumption of preparation independence [81].

A Closer Look

It has been suggested that PBR’s preparation independence assumption has a similarity to Bell’s “measurement independence”-assumption [54]. It has subsequently undergone some scrutiny and, as discussed in detail by Leifer [4] can be decomposed into two separate (well-known) assumptions that are used at different places in the proof.

- i) Cartesian product assumption or *separability for product states*: All properties of a product state are fully determined by the properties of the individual systems. There are no “holistic” ontic states [4, 82].
- ii) No-correlation assumption or *factorization for product states*: product quantum states are represented by product epistemic states [4, 82].

Separability. This notion goes back to Einstein and implies that spatially separate systems should have independent ontic states (or “elements of reality”). In this sense the properties of a composite system are fully determined by the properties of the individual constituents, just as in the case of classical mechanics. This implies that the ontic state space for n particles should have a Cartesian product structure $\Lambda = \Lambda_1 \times \dots \times \Lambda_n$. The epistemic states nonetheless live in the corresponding tensor product space due to a simple closure argument⁹ [82].

Experimental tests of Bell-local and Leggett’s crypto-local models suggest that separability is problematic for entangled states. Also in ψ -ontic models separability cannot hold in general, since it would imply a Cartesian product structure of the quantum state space and not allow for entangled states [82]. Hence, PBR only makes this assumption for product states, which is physically motivated by the idea that states which can actually be prepared spacelike separated from one another should not have any holistic properties. On the ontic level this means that only the set of ontic states in the support of product state distributions is assumed to be a subset of the Cartesian product of the corresponding ontic spaces. The PBR result could be interpreted as a demonstration of the failure of separability for product states in ψ -epistemic models.

⁹This would also make for a good motivation of ψ -epistemic models, since the tensor product structure of quantum mechanics would arise naturally [82].

Factorization. This is the assumption that composite quantum states should be represented by product epistemic states, $\mu(\lambda) = \mu_1(\lambda_1)\mu_2(\lambda_2)\cdots\mu_n(\lambda_n)$. Although closely related, this assumption is different from separability, which implies that $\lambda = (\lambda_1, \dots, \lambda_n)$. As with separability, PBR only assume factorization for product states, which is motivated by the idea that two states which are prepared in an entirely uncorrelated fashion, such that no measurement can reveal any correlations between them, then there should also be no correlation at the ontic level [4]. In the case where separability for product states holds, this essentially reduces to preparation noncontextuality for product states [82].

Weakening the Assumptions

The PBR theorem highlights a conflict between empirical observations and the assumptions of ψ -epistemicity, separability for product states and factorization for product states. In the face of an experimental violation at least one of these three assumptions must be given up or relaxed. It is thus an interesting question to what extent the various assumptions could be relaxed, while still leading to the same conflict. Two such efforts are briefly discussed below.

Compactness or *compatibility* is a notion introduced in Refs. [4, 83, 84]. It states that, if there is an ontic state λ in the overlap of μ_{ψ_1} and μ_{ψ_2} , then there must be an ontic state λ' which is in the support of the epistemic state for all 4 product states $|\psi_1\rangle \otimes |\psi_1\rangle, |\psi_1\rangle \otimes |\psi_2\rangle, |\psi_2\rangle \otimes |\psi_1\rangle, |\psi_2\rangle \otimes |\psi_2\rangle$. While this is probably the weakest assumption required for the PBR proof to work, it lacks much of the physical motivation of separability and factorization.

The weak preparation independence postulate or *local independence* as termed by Ref. [54] separates the ontic states of the joint system into local and non-local ontic states, where the latter are inaccessible under local measurements. Separability for product states is then weakened to $\Lambda = \Lambda_1 \times \Lambda_2 \times \Lambda_{NL}$ and factorization for product states is weakened to $\mu_{\psi_1 \otimes \psi_2}(\Lambda_{12} \times \Lambda_{NL}) = (\mu_{\psi_1} \times \mu_{\psi_2})(\Lambda_{12})$. Hence, under local measurements there is no difference between a model that satisfies this assumption and one that satisfies preparation independence; the additional ontic variable is only revealed under entangling measurements, which is the final step of the PBR proof. In Ref. [54] a toy model is described that evades the PBR conclusion, but satisfies weak preparation independence. It is, however, unclear whether this can be generalized to a model for all of quantum theory [4, 54]. Other explicit ψ -epistemic models, which are derived from Bell's model [62, 79] all violate PIP and its weaker version.

4.3.2 Other Theorems, Other Assumptions

The PBR result was followed by a series of ψ -ontology theorems, using a variety of different assumptions [51–53, 62]. The PBR theorem is concerned with multiple copies of a system, whose preparations are assumed to be independent [50]. Others consider multiple systems under the assumption of free choice and a notion of locality [51], or single systems, assuming notions of continuity [52], invariance [53], or symmetry of the ontic state space [62].

Notably, the counterexamples of Ref. [79] are explicitly constructed for single systems and would allow for superluminal influence from the measurement choices to the ontic states. These

models therefore do not directly apply to the Bell-like scenario in Ref. [51], which is by assumption free of superluminal influences [85, 86]. However, the latter assumes a notion of locality and has thus been challenged in the light of Bell’s theorem.

Colbeck, Renner

The first of a series of papers by Colbeck and Renner was in fact published *before* PBR in 2011, and focused on a slightly different aspect of ontological models than the ψ -ontic/ ψ -epistemic discussion [87]. The authors argue that knowledge of the ontic state can, under assumptions of free will and locality, never allow for more precise predictions than knowledge of only the quantum state. As a by-product their work also includes a ψ -ontology theorem, which was published separately in a simplified form [51]. The central assumption in this work is a strong notion of “free choice”, which is meant to capture the idea that the experimenter is free to choose measurement settings independently of everything, including ontic variables and pre-existing information, that is not in its future light cone in any reference frame, such that it is only correlated with events that could have been caused by it [51]. Fundamentally this incorporates Bell’s measurement independence (i.e. measurement choices are independent of the ontic variables), as well as notions of relativistic causality, temporal order and causal arrow, see Sec. 5.3.1. Hence, it implicitly assumes parameter independence [4], stating that the outcome of one freely chosen measurement should not depend on the choice for another space-like separated measurement.

Under these assumptions they choose a bipartite Bell-like scenario that satisfies the conditions that (i) measurement is chosen before the outcome is obtained (causal arrow), (ii) ontic variables cannot influence the measurement choice (measurement-independence), (iii) measurements on one side cannot influence the outcomes on the other side (parameter independence). They show that this implies signal locality in the ontological model, and using chained Bell inequalities they bound the amount of correlation between the ontic state and the local measurement outcomes. In the limit of infinitely many measurements in the chain it follows that the ontic state carries the same information about the local measurement outcomes as the quantum state [51]. Assuming further that appending ancillas preserves ontological indistinctness, this argument can be turned into a ψ -ontology theorem, see Ref. [4] for details.

In fact, the Colbeck-Renner result establishes that any ontological model which reproduces quantum predictions and is consistent with their free choice assumption must be ψ -complete, rather than just ψ -ontic, which is stronger than the PBR theorem. However, imposing locality assumptions on ontological models has received some criticism [4] and it is known that Bohmian Mechanics, for example, is a ψ -ontic (but not ψ -complete) model that does not satisfy parameter independence and is therefore not ruled out by this result. Furthermore locality assumptions require a background causal structure, which in the light of recent developments towards understanding the causal structure of quantum mechanics is rather undesirable, see Chap. 5.

Hardy

Hardy showed that there can be no ψ -epistemic model that satisfies (*restricted*) *ontic indifference* [53]. This is a kind of symmetry assumption which implies that there exists at least one pure quantum state ψ , for which any unitary that leaves it unchanged can be implemented in a way that also leaves the underlying ontic states unchanged. Although ontic indifference has been criticised as an assumption for a ψ -ontology theorem [4], Hardy's theorem is relevant as one of the few results that study dynamics in ontological models. In contrast to PBR, Hardy's theorem applies to a single system and, as pointed out in Ref. [4], makes precise the hand-waving argument in favour of ψ -ontic models based on interference. Strictly speaking Hardy's result only holds for infinite-dimensional Hilbert spaces. In order to extend the result to finite dimensions it is necessary to add an ancillary system of arbitrary dimension (which could be "the rest of the universe") under the assumption that this preserves ontic indistinctness. This is motivated in a similar fashion as PBR's separability for product states, see Ref. [53] and [4] for a more detailed discussion.

Patra et al.

Patra et al. [52] derived a no-go theorem for ψ -epistemic models that satisfy a weak *separability* assumption and a strong *continuity* assumption. The first requires that if a preparation of ψ has non-zero probability for the ontic state λ , then a preparation of n independent copies of ψ must have non-zero probability for the joint ontic state $(\lambda, \dots, \lambda)$. The continuity assumption is meant to translate the continuous quantum state-space to the ontic space in the sense that a slight variation of ψ should only cause a slight change of μ_ψ . The authors then classify the possible ψ -epistemic models by two parameters: their continuity δ and their epistemicity [88] ε . Here δ is the radius of a ball around the state ψ , such that there exists an ontic state in μ_ψ that is also in μ_ϕ for all states ϕ within this δ -ball. The quantity $\varepsilon = \sum_\lambda \min_k \mu_{\psi_k}(\lambda)$ is a measure of the overlap of epistemic states for a choice of d quantum states ψ_k (where d is also the Hilbert space dimension) in the δ -ball around ψ .

They then prove two no-go theorems. The first rules out δ -continuous models for $\delta \geq 1 - \sqrt{(d-1)/d}$. The second makes use of the separability assumption, proving that there exists no ψ -epistemic model that is both δ -continuous and satisfies the separability assumption. In Ref. [88] a generalized, experimentally testable version of the theorem is presented, using a redefinition of δ and ε to contend with imperfect detection efficiency. The predictions of this theorem have been confirmed experimentally using time-bin encoding on an optical coherent state [88]. Finally, the notion of δ -continuity has been criticised for being much too strong for ψ -ontology theorems. The particular problems are nicely illustrated with a classical example in Ref. [4].

Aaronson et al.

Aaronson et al. [62] showed that there exists no ψ -epistemic model that is *maximally non-trivial* and *symmetric under unitary transformations*. Maximally non-trivial, or *pairwise ψ -epistemic* [4] models are such that the epistemic states for any pair of non-orthogonal quantum states, independent of their inner product, have non-zero overlap. The second assumption, symmetry under unitary transformations, fixes the ontic state space to either the space of pure states or the space of unitary operators, and requires that the epistemic state is a function f_ψ only of the overlap between ψ and λ (i.e. $\mu_\psi(\lambda) = f_\psi(|\langle\psi|\lambda\rangle|)$). As a consequence, μ_ψ is invariant under all unitary operations that keep ψ fixed. The motivation for this assumption comes from a stronger notion of complete unitary invariance, which the authors call *strong symmetry*, where the epistemic state is a fixed function of the overlap between ψ and λ for all quantum states. In this case, applying a unitary U to a quantum state is equivalent to applying it to the ontic states. Hence, in such theories one can simply use the Schrödinger equation to describe time evolution on the ontic level. In contrast to Hardy's ontic indifference, however, these symmetry assumptions do not require that the ontic states remain fixed under unitary transformations. While this no-go theorem is very aesthetic, the restriction of the ontic state space, to be almost equivalent to the quantum state space, is very restrictive. It seems not too surprising that in this case the quantum state would turn out to be ontic.

Curiously, the authors also show that without the symmetry assumption it is always possible to construct pairwise ψ -epistemic models for any dimension. Loosely speaking, this entails first constructing a ψ -epistemic model with overlapping epistemic states for an arbitrary pair of pure states and then combining a dense set of these models via a form of convex combination. This gives by construction pairwise ψ -epistemic model. However, the ontic state space is $\mathbb{C}\mathbb{P}^{d-1} \times [0, 1] \times \mathbb{N}$, since the construction requires to keep one copy of each original ontic space [62], thus violating symmetry under unitary transformations. Furthermore, the overlap for any pair of states is very small: $\omega_c(\psi, \phi) \sim (|\langle\psi|\phi\rangle|/d)^{\mathcal{O}(d)}$.

4.4 Constraining ψ -epistemic models

The moral of the previous section is that ruling out ψ -epistemic models is not possible without additional assumptions beyond the ontological models framework. The present section explores an alternative approach without additional assumptions, but with more moderate demands. Instead of completely ruling out ψ -epistemic models, one can aim to bound the extent to which overlapping probability distribution can explain observations, in particular the imperfect distinguishability of quantum states. In a pioneering work Barrett et al. [60] showed that no ψ -epistemic model can fully explain this latter phenomenon. Specifically, there are sets of states for which the ratio κ of classical-to-quantum overlaps must be smaller than one, for systems of dimension $d \geq 3$, and that this bound decreases linearly with the dimension. This result was subsequently improved to an exponential scaling with dimension [61], and Ref. [58] showed that an arbitrarily low bound on κ can be achieved for any dimension $d \geq 4$. Whether the same is

true for $d = 3$ is still an open question.

Common to all three works is that they *do not* rely on any auxiliary assumptions beyond the ontological models framework. They all rule out the most compelling class of ψ -epistemic models, those that fully explain quantum indistinguishability in terms of overlapping probability distributions. The rest of this section is devoted to a sketch of the derivation of an experimentally testable inequality for maximally ψ -epistemic models, derived in Ref. [1, 58] and the corresponding experimental protocol to test it.

4.4.1 How to Constrain ψ -Epistemic Models

Recall that in ψ -epistemic models the epistemic states μ_ψ and μ_ϕ corresponding to two non-orthogonal quantum states $|\psi\rangle$ and $|\phi\rangle$ can overlap. This offers an intuitive qualitative explanation for why these states cannot be perfectly distinguished, but it is unclear whether this explanation reproduces the phenomenon quantitatively. A good figure of merit to assess how much of the overlap a model can explain is the ratio $\kappa(\psi, \phi)$ of classical-to-quantum overlap. Any value of $\kappa < 1$ (for any pair of states) implies that the overlap of the epistemic states is insufficient to explain quantum indistinguishability and thus rules out maximally ψ -epistemic models. The problem is that the epistemic overlap is defined at the level of ontic state and cannot be measured directly. The fascinating result of Ref. [60] is that with a bit of measure theory it is nonetheless possible to use experimental data to impose upper bounds on the amount of overlap that is possible in any ψ -epistemic model, see Ref. [58, 60] for details. The experimental protocol is illustrated on the example of ququarts in the real subspace in Fig. 4.6, using the following steps

- (i) Pick a reference state ψ_0 and a set of $n \geq 2$ quantum states: $\{\psi_j\}_{j=1}^n$.
- (ii) For each triplet of states $\{\psi_0, \psi_{j_1}, \psi_{j_2}\}$ (with $j_1 < j_2$) which includes the reference pick a measurement $M_{j_1 j_2}$ with 3 outcomes (m_0, m_1, m_2) and measure the probabilities $P_{M_{j_1 j_2}}(m_i|\psi_{j_i})$ for outcome m_i on the state ψ_{j_i} (with $j_0 = 0$) of the triplet.
- (iii) From the experimental data estimate $1 + \sum_{1 \leq j_1 < j_2 \leq n} \sum_{i=0}^2 P_{M_{j_1 j_2}}(m_i|\psi_{j_i})$, which is an upper bound on the classical overlap, see below.
- (iv) Compute the sum of pairwise quantum overlaps $\sum_{1 \leq j \leq n} \omega_Q(|\psi_0\rangle, |\psi_j\rangle)$ for the used states.

Any maximally ψ -epistemic model must satisfy

$$S(\{\psi_j\}, \{M_{j_1 j_2}\}) = \frac{1 + \sum_{1 \leq j_1 < j_2 \leq n} \sum_{i=0}^2 P_{M_{j_1 j_2}}(m_i|\psi_{j_i})}{\sum_{1 \leq j \leq n} \omega_Q(|\psi_0\rangle, |\psi_j\rangle)} \geq 1. \quad (4.8)$$

The quantity $S(\{\psi_j\}, \{M_{j_1 j_2}\})$ is an upper bound on the minimal overlap ratio for the set of states $\{|\psi_j\rangle\}_{j=1}^n$. This implies that within the sets of states measured, there is at least one pair

of states with an overlap of at most $S(\{\psi_j\}, \{M_{j_1 j_2}\})$. In particular, $S < 1$ is inconsistent with maximally ψ -epistemic models.

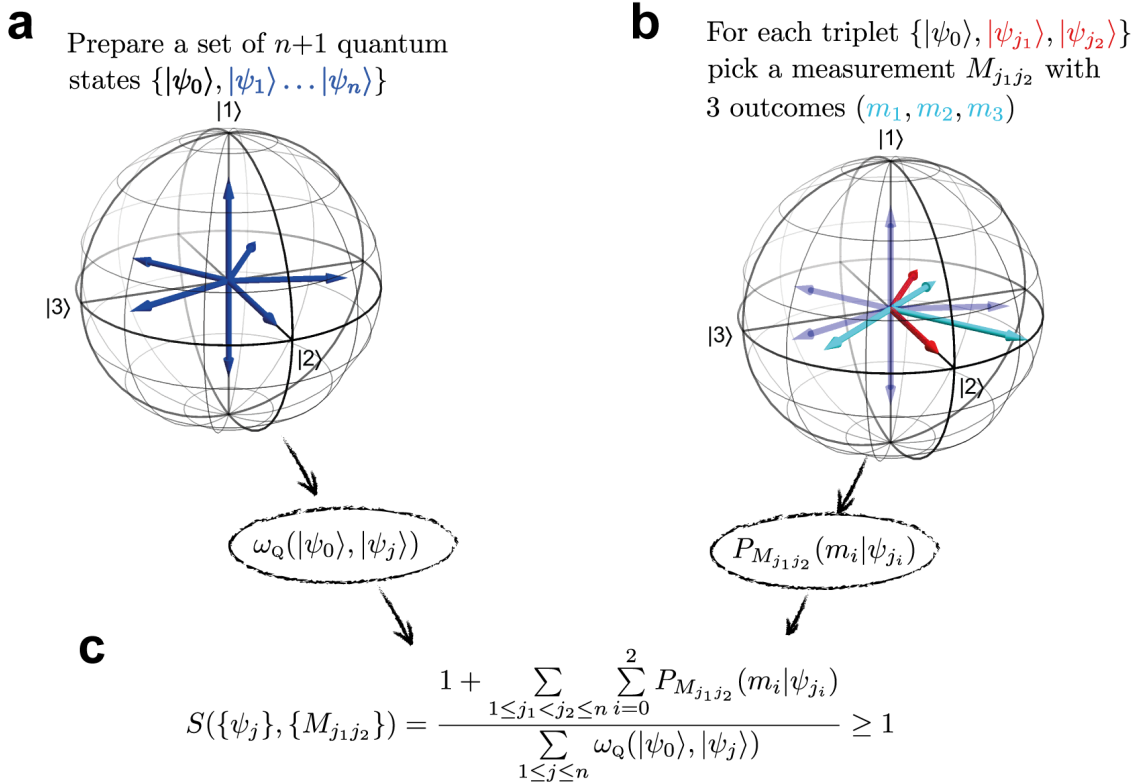


Figure 4.6: **The experimental protocol for ququarts.** (a) Prepare $n + 1$ quantum states for $n \geq 2$, shown in blue. The states are shown in the subspace orthogonal to the reference state $|\psi_0\rangle = |0\rangle = (1, 0, 0, 0)$, spanned by the computational basis states $|1\rangle, |2\rangle, |3\rangle$. The pairwise quantum overlaps $\omega_Q(|\psi_0\rangle, |\psi_j\rangle)$ can be computed directly for the used states. (b) For each triplet of states $\{\psi_0, \psi_{j_1}, \psi_{j_2}\}$ (with $j_1 < j_2$), which includes the reference, pick a measurement $M_{j_1 j_2}$ with 3 outcomes (m_0, m_1, m_2) . The outcome probabilities for outcome m_i on state $|\psi_{j_i}\rangle$, denoted $P_{M_{j_1 j_2}}(m_i | \psi_{j_i})$, can be measured experimentally. (c) The value $S(\{\psi_j\}, \{M_{j_1 j_2}\})$ can be computed from the experimental data and is an upper bound on minimal overlap present in any ψ -epistemic model that reproduces the experimental statistics for the chosen states and measurements.

4.4.2 Maths Exercise

Following Ref. [58] this section sketches how inequality (4.8) can be derived with just a bit of elementary measure theory. Those readers for whom it is obvious, please skip ahead to Sec. 4.4.3, those who want the full details are referred to Ref. [4, 58].

Since $\mu_\psi: \Lambda \rightarrow [0, 1]$ is a positive, measurable function, it defines a region, the “area under the curve”, denoted $C_\psi \subset \Lambda \times \mathbb{R}^+$, see Fig. 4.7

$$C_\psi := \{(\lambda, x): 0 \leq x \leq \mu_\psi(\lambda)\} .$$

The symmetric overlap of multiple such regions C_{ψ_i} is simply given by the intersection

$$\bigcap_i C_{\psi_i} = \{(\lambda, x): 0 \leq x \leq \min_i [\mu_{\psi_i}(\lambda)]\} . \quad (4.9)$$

The volume $\nu(C)$ of a region C can be determined by integrating over $\Lambda \times \mathbb{R}^+$ using the measure $d\nu = d\lambda \times dx$. In the case of one set C_ψ this recovers the normalization condition

$$\nu(C_\psi) = \int_{C_\psi} d\nu = \int_\Lambda d\lambda \int_0^{\mu_\psi(\lambda)} dx = \int_\Lambda \mu_\psi(\lambda) d\lambda = 1 .$$

In the case of multiple sets the volume generalizes the classical overlap defined in Eq. (4.5),

$$\begin{aligned} \nu(C_\psi \cap C_\phi) &= \int_\Lambda \min[\mu_\psi(\lambda), \mu_\phi(\lambda)] d\lambda = \omega_C(\mu_\psi, \mu_\phi) \\ \nu(\bigcap_i C_{\psi_i}) &= \int_\Lambda \min_i [\mu_{\psi_i}(\lambda)] d\lambda . \end{aligned}$$

Different (e.g. asymmetric) definitions are possible with different overlap measures, but the present definition is widely used for its operational interpretation in terms of single-shot distinguishability.




Figure 4.7: **Mapping of the probability measures μ_ψ to sets C_ψ .** (a) The set C_ψ literally represents the area under the curve of μ_ψ . (b) The symmetric overlap is defined as the intersection of the two sets C_ψ and C_ϕ

For the triplets $(\psi_0, \psi_{j_1}, \psi_{j_2})$ it follows that

$$\begin{aligned} \sum_{i=0}^2 \xi_{M_{j_1, j_2}}(m_i | \lambda) = 1 \quad \forall \lambda \\ \nu(C_{\psi_0} \cap C_{\psi_{j_1}} \cap C_{\psi_{j_2}}) &\stackrel{\downarrow}{=} \int_\Lambda \left(\sum_{i=0}^2 \xi_{M_{j_1, j_2}}(m_i | \lambda) \right) \min_i [\mu_{\psi_{j_i}}(\lambda)] d\lambda \\ &\leq \sum_{i=0}^2 \int_\Lambda \xi_{M_{j_1, j_2}}(m_i | \lambda) \mu_{\psi_{j_i}}(\lambda) d\lambda \\ &= \sum_{i=0}^2 P_{M_{j_1, j_2}}(m_i | \psi_{j_i}) . \end{aligned} \quad (4.10)$$

This establishes a connection between the overlap of the triplets $(\psi_0, \psi_{j_1}, \psi_{j_2})$ and the observed measurement probabilities¹⁰ $P_{M_{j_1, j_2}}(m_i | \psi_{j_i})$. The next step is to relate the left-hand side of

¹⁰At this point it is not yet assumed that the measurement probabilities agree with the Born rule. In the following, however, it will be implicitly assumed that the considered models reproduce quantum statistics, i.e.

inequality (4.10) to the classical overlap, which is defined for pairs of states. Denoting the intersections with the reference set C_{ψ_0} by $A_j = C_{\psi_0} \cap C_{\psi_j}$ for $j = 1, \dots, n$ as in Ref. [58] one can use a standard counting argument, the so-called inclusion-exclusion principle to obtain

$$\nu\left(\bigcup_{j=1}^n A_j\right) = \sum_{k=1}^n (-1)^{k-1} \sum_{1 \leq j_1 < \dots < j_k \leq n} \nu(A_{j_1} \cap \dots \cap A_{j_k}) \quad (4.11)$$

$$\geq \sum_{j=1}^n \nu(A_j) - \sum_{1 \leq j_1 < j_2 \leq n} \nu(A_{j_1} \cap A_{j_2}). \quad (4.12)$$

Truncating the sum in Eq. (4.11) yields alternating upper and lower bounds, known as Bonferroni inequalities. Since, by definition, every A_j contains the set C_{ψ_0} the union of all A_j cannot be larger than C_{ψ_0} , and the left-hand side of the inequality (4.12) is therefore upper-bounded by 1. Some reshuffling implies:

$$\begin{aligned} \sum_{j=1}^n \underbrace{\nu(A_j)}_{\substack{=\nu(C_{\psi_0} \cap C_{\psi_j}) \\ =\omega_C(\mu_{\psi_0}, \mu_{\psi_j})}} &\leq \underbrace{\nu\left(\bigcup_{j=2}^n A_j\right)}_{\leq \nu(C_{\psi_0})=1} + \sum_{1 \leq j_1 < j_2 \leq n} \underbrace{\nu(A_{j_1} \cap A_{j_2})}_{\substack{=\nu(C_{\psi_0} \cap C_{\psi_{j_1}} \cap C_{\psi_{j_2}}) \\ \leq \sum_{i=0}^2 P_{M_{j_1, j_2}}(m_i | \psi_{j_i})}} \\ \sum_{j=1}^n \omega_C(\mu_{\psi_0}, \mu_{\psi_j}) &\leq 1 + \sum_{1 \leq j_1 < j_2 \leq n} \sum_{i=0}^2 P_{M_{j_1, j_2}}(m_i | \psi_{j_i}). \end{aligned}$$

Hence, the sum of pairwise classical overlaps ω_C , between the reference state ψ_0 and the states $\{\psi_j\}_{j=1}^n$, can be upper-bounded from experimental data. Crucially, this is a bound on the *sum* of overlaps and does not directly constrain the overlap of any specific pair of states, other than by the consistency condition Eq. (4.6), that the classical overlap can never exceed the quantum overlap. This, however, implies that whenever the ratio between the sums of classical overlaps and the sums of quantum overlaps is bounded below 1, there must be at least one pair of states with an overlap ratio below 1.

$$\begin{aligned} \frac{1 + \sum_{1 \leq j_1 < j_2 \leq n} \sum_{i=0}^2 P_{M_{j_1, j_2}}(m_i | \psi_{j_i})}{\sum_{j=1}^n \omega_Q(|\psi_0\rangle, |\psi_j\rangle)} &\geq \frac{\sum_{j=1}^n \omega_C(\mu_{\psi_0}, \mu_{\psi_j})}{\sum_{j=1}^n \omega_Q(|\psi_0\rangle, |\psi_j\rangle)} \\ &= \frac{\sum_{j=1}^n \kappa(\psi_0, \psi_j) \omega_Q(|\psi_0\rangle, |\psi_j\rangle)}{\sum_{j=1}^n \omega_Q(|\psi_0\rangle, |\psi_j\rangle)} \\ &\geq \min_j [\kappa(\psi_0, \psi_j)] \frac{\sum_{j=1}^n \omega_Q(|\psi_0\rangle, |\psi_j\rangle)}{\sum_{j=1}^n \omega_Q(|\psi_0\rangle, |\psi_j\rangle)}. \end{aligned}$$

For maximally ψ -epistemic models it must hold that $\kappa(\psi_0, \psi_j) = 1$ for all states ψ_j , which yields inequality (4.8), where the left-hand side is an experimentally accessible figure of merit that distinguishes maximally from non-maximally ψ -epistemic models. In the special case where all quantum overlaps are equal, a similar argument can be used to obtain a bound on the average overlap ratio

satisfy Eq. (4.2).

$$\begin{aligned}
\frac{1 + \sum_{1 \leq j_1 < j_2 \leq n} \sum_{i=0}^2 P_{M_{j_1, j_2}}(m_i | \psi_{j_i})}{\sum_{j=1}^n \omega_{\mathcal{Q}}(|\psi_0\rangle, |\psi_j\rangle)} &\geq \frac{\sum_{j=1}^n \kappa(\psi_0, \psi_j) \omega_{\mathcal{Q}}(|\psi_0\rangle, |\psi_j\rangle)}{\sum_{j=1}^n \omega_{\mathcal{Q}}(|\psi_0\rangle, |\psi_j\rangle)} \\
&= \frac{\omega_{\mathcal{Q}}(|\psi_0\rangle, |\psi_1\rangle) \sum_{j=1}^n \kappa(\psi_0, \psi_j)}{n \omega_{\mathcal{Q}}(|\psi_0\rangle, |\psi_1\rangle)} \\
&= \frac{1}{n} \sum_{j=1}^n \kappa(\psi_0, \psi_j) .
\end{aligned}$$

Under this condition there exist states and measurements [58] for which the average ratio scales as

$$\frac{1}{n} \sum_{j=1}^n \kappa(\psi_0, \psi_j) < 8/n^{(d-3)/(d-2)} . \quad (4.13)$$

Hence, for systems of dimension $d > 3$ arbitrary low bounds can be imposed on the average (and thus also minimal) classical-to-quantum overlap ratio, showing that the explanation of quantum indistinguishability in terms of overlapping probability distributions is arbitrarily bad. Whether this result holds for $d = 3$ as well is still an open question. Arbitrary low bounds, of course, come at a price: they require arbitrary large sets of states, which, experimentally, is rather unfortunate. Nevertheless, nontrivial bounds can already be established for any system of dimension $d \geq 3$, as soon as the number of states is $n \geq 3$.

4.4.3 PP-Incompatibility

Although the approach of Ref. [58] builds on the ideas of Ref. [50, 60], there is a crucial difference in the choice of states. Barrett et al. [60] (mostly for convenience) focus on so-called PP-incompatible states, which are such that for each triplet $(\psi_0, \psi_{j_1}, \psi_{j_2})$ a measurement M_{j_1, j_2} exists (in the subspace spanned by the states), such that $P_{M_{j_1, j_2}}(m_i | \psi_{j_i}) = 0$ for $0 < i < 2$ (with $j_0 = 0$). Their argument is then similar to PBR, that, when quantum theory assigns a probability of zero, so must the ontological model. In other words $C_{\psi_0} \cap C_{\psi_{j_1}} \cap C_{\psi_{j_2}}$ must be of measure zero, which simplified the derivation. Such a configuration can be achieved by picking states from mutually unbiased bases, but this fixes the number of states for given dimension. The derivation of Ref. [58] is more flexible in that it does not make any such assumption. In fact, the states constructed for the analytical bound also turn out to be PP-incompatible for any n , but this bound might not be optimal and considering non-PP incompatible states might yield better bounds.

4.4.4 Experimental Robustness

Inequality (4.8) allows, in principle, for an experimental test of ψ -epistemic models without any auxiliary assumptions beyond the ontological models framework. For such a test to be practical, however, it has to be sufficiently robust against experimental imperfections.

Recall that the number of triplets of states required for the test scales as $\frac{n(n-1)}{2}$, and for each triplet a 3-outcome has to be implemented to estimate the probabilities $P_{M_{j_1, j_2}}(m_i | \psi_{j_i})$

with $0 < i < 2$. Hence, the number of terms in the numerator of inequality (4.8) is $\frac{3n(n-1)}{2}$. To estimate the error tolerance of the inequality one can add a uniform error ε to each term and compute the largest value of ε such that the inequality is still violated [58]

$$\varepsilon < \frac{1}{\frac{3}{2}n(n-1)} \left(\frac{1 + \sum_{1 \leq j_1 < j_2 \leq n} \sum_{i=0}^2 P_{M_{j_1 j_2}}(m_i | \psi_{j_i})}{\sum_{1 \leq j \leq n} \omega_Q(|\psi_0\rangle, |\psi_j\rangle)} - 1 \right). \quad (4.14)$$

For the states used in Ref. [58] to obtain the analytical bound Eq. (4.13) the maximally permissible error per measurement scales roughly as $\varepsilon \lesssim 1/(12n^{(d-1)/(d-2)})$, which is on the order of 10^{-3} per measurement. Increasing the number of states n thus not only requires a quadratic increase in the number of measurements, but also leads to a *less* noise-robust experiment. This implies a tradeoff between large n and sufficiently low noise to observe a strong violation. On the positive side, the bounds are obtained for projective measurements, which can in general be implemented with higher precision than arbitrary POVMs.

Since the states used in the experiment are obtained from numerical optimization they are in general *not* PP-incompatible, which means that the probabilities in the numerator of inequality (4.8) are not all expected to vanish. Instead, some vanish and others are just very close to zero. In either case, however, measuring a bounded quantity (such as a probability) very close to the boundary leads to a very sensitive experiment and asymmetric error distributions that are skewed away from the boundary. In fact, no realistic experiment with finite data can hope to reproduce a value on the boundary, such as an ideal probability of zero. Finally, the denominator of inequality (4.8) depends on the prepared quantum states, which makes accurate characterization of these states crucial for reliable estimates.

4.5 Testing Realist ψ -Epistemic Models

This section focuses on the “easiest experiment possible” [89]: a prepare-and-measure experiment to test (maximally) ψ -epistemic models and impose experimental constraints on the possible class of models using the ideas developed in the previous section. In particular, these experiments test inequality (4.8) and use the parameter S to classify the models that remain compatible with the experimental results. As discussed above, there are two important ingredients to such an experimental test.

Firstly, a quantum system of dimension 3 (qutrit) or larger is required. Higher-dimensional systems would lead to exponentially stronger bounds in inequality (4.8), but at the same time, achieving the required precision in state preparation and measurement becomes increasingly challenging. Practically this implies a tradeoff, which, with current technology, rather points towards the low-dimensional end. A conceptually interesting aspect is the divisibility of quantum systems. The ψ -ontology theorems of Ref. [50, 52], for example, use independence assumptions and composite Hilbert spaces, while the derivations presented in Sec. 4.4, following Ref. [58, 60, 61] are concerned with single quantum systems. Keeping with this spirit it is conceptually desirable to study high-dimensional Hilbert spaces on individual systems, rather than

constructing them by combining lower-dimensional spaces (e.g. using single photons rather than pairs). Finally, any low-dimensional system can always be embedded in a higher-dimensional Hilbert space, which means that a violation of inequality (4.8) for dimension d implies a violation for any larger d . For this reason it might be desirable to use the lowest-dimensional system, a qutrit, which also happens to be naturally indivisible.

Secondly, a non-trivial bound requires the preparation of at least 4 (i.e. at least $n = 3$) different quantum states [1, 58]. While increasing the number of states n allows, in theory, for a larger violation of inequality (4.8), this comes at the cost of an increase in the number of measurements that is quadratic in n , and with an overall decrease in the error-tolerance of the experiment. Hence, a tradeoff between violation and experimental imperfections is expected to occur, where at some point the accuracy of the experiment will be insufficient to achieve an improvement by adding another state.

4.5.1 Choosing States and Measurements

The choice of states and measurements is critical for achieving strong bounds on the minimal ratio of classical-to-quantum overlaps. For every dimension $d \geq 4$, Branciard [58] showed that there are families of $n + 1$ states, achieving arbitrarily low bounds as $n \rightarrow \infty$. The states used to obtain these analytic bounds, however are not necessarily optimal for an experimental scenario. However, they can be further subject to numerical optimization to achieve a stronger violation of inequality (4.8). In the present experiment this optimization was performed over the real-valued subspace of the respective Hilbert-space. This reduces the achievable violation for larger n , but allows for a simpler and more precise optical setup. The obtained solutions are given in the supplementary information of Ref. [1], but although they achieve an improvement over the analytical bound, they cannot be guaranteed to be optimal, since the optimization is non-convex.

Instead of optimizing the states and measurements to achieve the largest violation of inequality (4.8), one might instead optimize for higher error tolerance, which tends to give slightly different results, or for any other figure of merit, see Sec. 4.6.2. Figure 4.6a shows an example set of states used in the experiment for $d = 4, n = 7$ projected onto the real-valued subspace orthogonal to $|\psi_0\rangle$. The measurement $M_{2,4}$ for this case is illustrated in Figure 4.6b, projected onto the same subspace.

4.5.2 Experimental Setup

The requirements for a high-dimensional system on a single particle that can be controlled to extreme precision and accuracy can be met using hyper-encoding in both the polarization and path degrees of freedom on a single photon, see Fig. 4.8. With two separate spatial paths this encoding naturally represents a ququart with computational basis states $|0\rangle = |H\rangle_1, |1\rangle = |V\rangle_1, |2\rangle = |H\rangle_2, |3\rangle = |V\rangle_2$, where $|p\rangle_m$ denotes a state of polarisation p in path m . A qutrit can be realised in this setup by simply leaving the state $|3\rangle$ unpopulated.

In contrast to fully path-encoded designs, which would enable encoding arbitrarily high dimensions, the present approach does not require any nested interferometers. Hence, it can deliver much greater stability and precision. Furthermore, reducing all state-preparation and measurement to polarization rotations with a fixed two-path interferometer is what enables the required precision and measurement accuracy (see Sec. 4.4.4), at least in 3 and 4 dimensions. While this setup could in principle be extended to arbitrary dimension by adding additional path degrees of freedom, this would require nested interferometers, which is detrimental for accuracy.

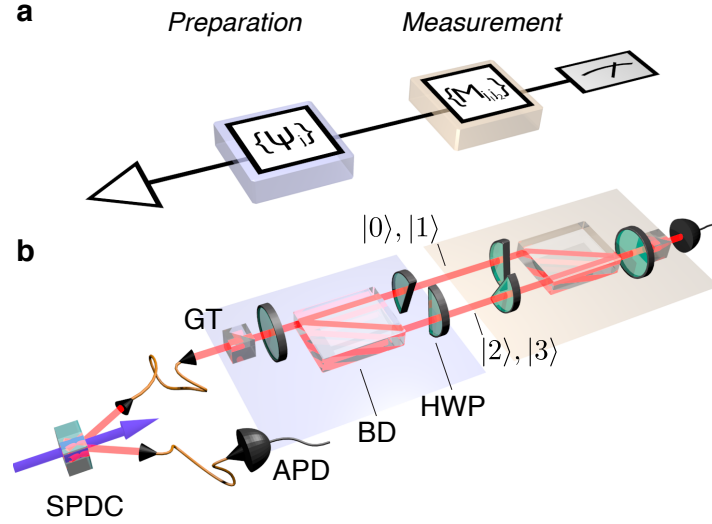


Figure 4.8: **Experimental setup.** (a) A single photon is prepared in the initial state $|H\rangle$ by means of a Glan-Taylor polariser (GT). The subsequent half-wave plate HWP defines the amplitudes of the *path*-encoded modes—i.e. the relationship between $\{|0\rangle, |1\rangle\}$ and $\{|2\rangle, |3\rangle\}$. A calcite beam displacer spatially separates these modes and another HWP in each arm sets the amplitude in this spatial mode. (b) Arbitrary projective measurements onto ququart states can be implemented using the reverse of the setup used for state preparation. An additional 45° offset is introduced to the measurement HWPs to achieve equal path-length of the two arms in the Jamin-Lebedev interferometer. Using only one output port of the final analysing polariser ensures maximal fidelity of the measurement process.

Calibration of the Experiment

On paper the setup of Fig. 4.8 looks very promising, but whether it can reach the stringent accuracy requirements depends on careful calibration and analysis of experimental imperfections. Any single-photon experiment necessarily suffers from statistical noise due to the Poissonian counting statistics. This affects all the experimentally estimated probabilities $P_M(m | \psi)$ in inequality (4.8) and can be taken into account using a Monte Carlo routine. Besides statistical noise, there are also a few points where the experiment could be affected by systematic imperfections, see Fig. 4.9.

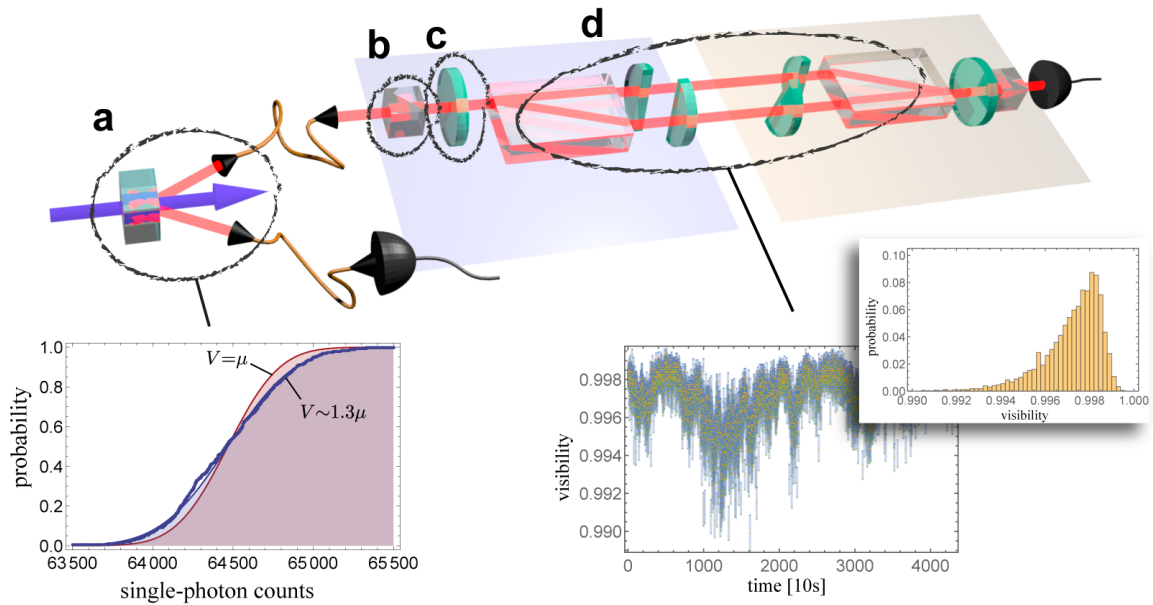


Figure 4.9: **The main sources of systematic imperfections in the setup.** (a) **Counting Statistics.** The main source of statistical noise in the experiment is the single photon counting statistics. Ideally, a random series of discrete events should be Poisson distributed, with equal variance and mean count rate. In practice, however, added noise drives the source into a slightly super-Poissonian regime with a larger variance of $V \sim 1.3\mu$. The counting statistics of the source are the dominating factor for statistical noise. (b) **Polarization Reference.** Everything has to be aligned to a single reference in the beginning to avoid systematic imperfections. (c) **Waveplates Inaccuracies** in the calibration and imperfections in the waveplate affect the accuracy of the experiment, while the precision is determined by the precision of the rotation stage. Detailed plots are given in Fig. 4.10. (d) **Interferometer.** Phase fluctuations in the Jamin-Lebedev polarisation-interferometer are critical for the precision of the optical setup. The inset shows a trace of the interferometer visibility for 10s measurements over a period of 12 hours, and a histogram of the visibilities. Including some long-term drifts the interferometer proved relatively stable with a median contrast of $\sim 400:1$.

Counting Statistics. The number of discrete events occurring within a fixed time-frame follows a Poisson distribution. This is true if the events occur at a fixed average rate and independently of the time since the last event. Under realistic experimental conditions these requirements do not hold exactly, and, due to added noise, the variance of the single-photon source turns out to be slightly larger than Poissonian, see Fig. 4.9a. For a measurement time of 10s per point we observe a variance of $V \sim 1.3\mu$. This corresponds to the typical measurement time in the experiment, and different noise timescales mean that the distribution depends on the chosen measurement time. Besides these statistical errors the source also exhibits systematic drifts due to fluctuations in the laser power. While the characteristic time-scale for these is on the order of tens of minutes and thus longer than typical measurement times, they can nonetheless affect longer measurement runs. This effect can be minimized by arranging measurements in time such that those used for normalization of a set of probabilities

are taken in close succession. Furthermore, laser drift can be monitored on the trigger port, which is detected directly at the source. Another source of fluctuations is polarization drift in the fibre from the source to the experiment, which only affects the overall experimental count-rate due to the initial Glan-Taylor polarizer. In contrast to laser drift, this effect cannot easily be corrected for, but it can be detected as a disproportional drop in count-rates between experiment and trigger. In the present experiment this effect was negligible on the relevant time-scales.

Reference. One of the more obvious sources of systematic error is imperfect calibration of the polarization optics. Since every optical element in the experimental setup is birefringent with some preferred axis, it is crucial to ensure all these axes are aligned with respect to a common reference. In the present case this reference is provided by the Glan-Taylor polarizer, which prepares a reference polarization state (here horizontal) with a nominal contrast of $10^5:1$. All other components are then calibrated in-situ, one by one, to this reference. Although such a method in general offers advantages over external calibration, it can suffer from accumulation of errors. In the present experimental design, however, this can largely be avoided by calibrating the elements as the setup is built, starting from the detector end.

Waveplates. The waveplates are the only “moving parts” in the experimental setup of Fig. 4.8, and all the state-preparation and measurement is reduced to polarization rotation. As a consequence, the waveplates are crucial for the accuracy of the experiment through imperfect calibration, steering, and manufacturing imperfections (i.e. imperfect retardance). Furthermore, the repeatability of the motorized rotation stages on the order of $\sim 1/100^\circ$ limit the precision of the experiment. Of the above, the dominating sources of imperfections are systematic offsets in the calibrated zero position, and, most importantly, non-ideal retardance. Recall that the latter can be determined from a calibration of the waveplate between aligned polarizers. Since the retardance is very sensitive to the angle-of-incidence, it is crucial that this step is performed in-situ. Table 4.1 summarizes the calibration results, with the waveplates numbered in the order in which they appear in the setup, see Fig. 4.8.

Element	$\delta\phi_0$	V
HWP1	0.04°	0.999 ± 0.002
HWP2	0.05°	1.000 ± 0.002
HWP3	0.06°	1.000 ± 0.003
HWP4	0.06°	0.997 ± 0.002
HWP5	0.06°	1.000 ± 0.002
HWP6	0.04°	0.999 ± 0.002

Table 4.1: **Calibration data for the used half-waveplates.** $\delta\phi_0$ is the standard error of the fitted zero position of the waveplate, and V is the visibility (with standard error) obtained from the fit.

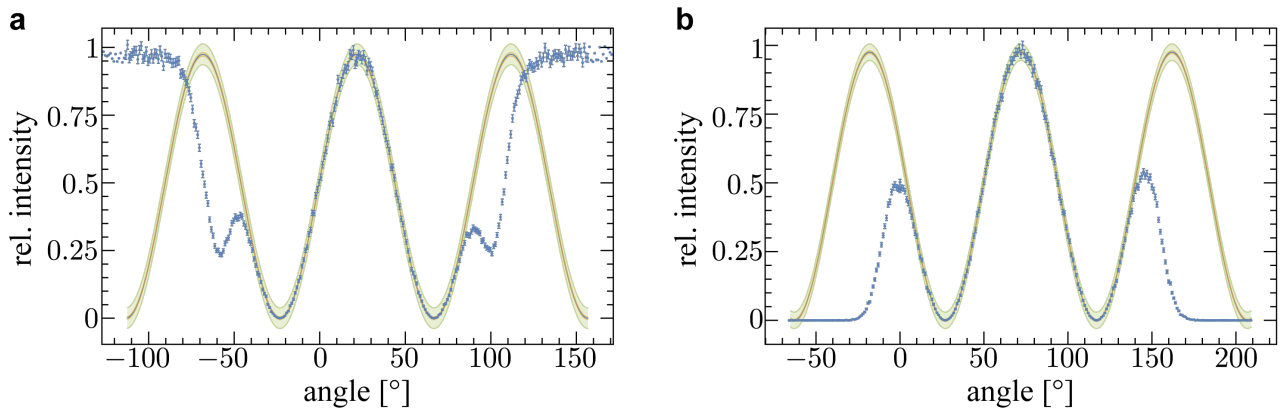


Figure 4.10: **Typical in-situ characterisation curve for a semi-circular waveplate.** Measured calibration data is shown in blue with error bars indicating statistical uncertainties. The fit-curve is shown in blue, with the 3σ confidence intervals in orange and 3σ prediction intervals in green. The fit describes the data well within a region of about $\pm 50^\circ$ around the upright position. Outside this region the waveplate fails to cover the whole beam, which leads to scattering. **(a)** In-situ calibration of a waveplate between aligned polarizers **(b)** In-situ calibration between crossed polarizers.

Another important aspect is that waveplates are not perfectly flat, which results in a slight deflection of the beam, that is further amplified by the distance of the waveplate from the output of the setup. This causes a variation in the single-photon coupling at the output of the setup, which translates into incorrect estimates of the experimental probabilities. This effect could in principle be suppressed by using multi-mode fibres, which perform much wider spatial filtering than single-mode fibres. However, this treats the symptom, not the cause, and the loss of tight spatial filtering results in additional noise. Furthermore, steering within an interferometer affects the overlap of the resulting beams and thus the quality of the interference. It is therefore preferable to keep the spatial selectivity of the single-mode fibre to ensure proper overlap of the interfering paths, and instead minimize steering through careful selection and positioning of the waveplates.

The interferometer. Since measurements of path-encoded quantum systems require interference of multiple spatial modes (except for the special case of computational-basis measurements), the quality of this interference (here the visibility of the Jamin-Lebedev interferometer) is crucial for the precision of the experiment. The two main sources of imperfect interference are a difference in path-length between the two arms, and imperfect beam overlap. The former is typically a result of drifts in the phase-alignment of the interferometer (here the yaw of the second beam displacer), which causes a slowly varying phase error in the prepared states that can be well approximated by a static unitary operation on the experimental timescale. Imperfect beam overlap, on the other hand results in de-phasing of the produced states and is caused by imperfect alignment, waveplate steering, length-unmatched beam displacers, or optical imperfections. In the present setup the optical quality of the interferometer limits the contrast to $\sim 400:1$.

The treatment of interferometer errors is somewhat complicated by the fact that every interferometer in a path-encoded scheme is part of (at least) two operations. In the present case of a single interferometer, errors affect the combination of state preparation and measurement and, in the absence of an intermediate evolution, can be attributed to either. The experimental protocol, however, is robust against such phase errors. Coherent errors cannot change the quantum overlap of a pair of states, while the choice $|\psi_0\rangle = |0\rangle$ ensures that incoherent errors are not a problem either. On the measurement side, recall that the exact form of the implemented POVM is irrelevant in the experiment. Hence, interferometer errors only affect the numerical value of the probabilities $P_{M_{j_1 j_2}}(m_i | \psi_i)$, thus weaken the achievable violation of inequality (4.8).

Performance Estimation

The calibration data gathered above can be used to estimate the performance of the optical setup to develop an idea for the expected values of S , and the uncertainties associated with the remaining systematic errors. Recall that the error budget is very tight, and, since the measured quantity is bounded and expected to be close to the boundary, most imperfections contribute in an unfavourable way. In particular, this step can provide an estimation for the tradeoff between number of states and accumulation of errors. This can be done using a Monte Carlo routine, which samples simulated experimental values assuming normal distributed variations in the calibration parameters. In each run the simulation picks a random value from the distribution of calibrated waveplate uncertainties, Poissonian noise and interferometer visibility.

The denominator of inequality (4.8) is the sum of the quantum overlaps $\sum_{j=1}^n \omega_Q(\psi_0, \psi_j)$ and thus only affected by state preparation. Moreover, each term in the sum is phase-independent and thus only affected by amplitude errors, which are dominated by waveplate imperfections and can be estimated from classical calibration data. Alternatively one might consider quantum state tomography to reconstruct the experimentally generated state and calculate the overlaps from these. This approach, however, conflates preparation and measurement errors and cannot provide a clear estimate of the quality of the state preparation alone. Moreover, state tomography always returns mixed states, which would require a generalization of the overlap measure ω_Q to mixed states. Hence, for the present purpose only coherent errors are modelled.

The numerator of inequality (4.8) depends on both the preparation and measurement parts of the setup. Classical calibration data for the waveplates (fit data with normal distributed standard errors) and the interferometer phase (sampled from zero-centred normal distribution with width matched to the average contrast) can be used to estimate the range of values.

Monte Carlo Approach. Under the approximation that numerator and denominator are independent the Monte Carlo routine can be performed in two steps. First the distribution of the denominator is estimated using wave-plate calibration data. This distribution is then used as an independent variable, together with the observed uncertainty distribution for the numerator to estimate the distribution of S . The resulting 1σ regions of expected values for S for ququarts

are shown as shaded areas in in Fig. 4.11. In practice this assumption is a good approximation, but not exactly satisfied. One could consider a more involved Monte Carlo routine, which does not use any information about the experimental distribution of uncertainties and instead samples the full parameter S from the classical calibration data. This would be computationally intense and is not necessary for the present purpose, where the Monte Carlo routine should only provide an estimate for the expected experimental performance.

4.5.3 Experimental Results

The theoretical predictions indicate that a violation of inequality (4.8) is possible for any dimension $d \geq 3$, using $n + 1 \geq 3 + 1$ quantum states are used. In practice the predicted violation for $n = 3$ is below the experimental precision and the first conclusive results are obtained for $n \geq 4$. The experimental data shown in Fig. 4.11 agrees well with the trend predicted by quantum mechanics when taking into account the calibration of the experiment, and conclusively rules out maximally ψ -epistemic models for any dimension $d \geq 3$. Since it is always possible to embed a low-dimensional quantum system into a higher-dimensional Hilbert space, the results for dimension d translate to any higher-dimensional Hilbert space. The result does, however, not hold in the case $d = 2$, where an explicit maximally ψ -epistemic model exists [5].

These experimental results were achieved with an average error per measurement of $\varepsilon \sim 0.0005_{-0.0004}^{+0.0012}$, see Fig. 4.12, which leads to the expected tradeoff seen in Fig. 4.11a between a larger violation of inequality (4.8) and error accumulation due to a large number of measurements, with saturation around $n \sim 12$. Achieving further improvements beyond this number of states would require higher experimental precision and interferometric stability.

Quantum State Tomography

An independent assessment of the quality of the experimental state preparation can be obtained from quantum state tomography. In principle, the tomography results could be used in the denominator of inequality (4.8). However, since the results of state tomography are conflated with coherent and incoherent errors in the interferometer and in the measurement, it significantly underestimates the quality of the state preparation alone. It can nonetheless provide an interesting overall figure of merit for the experiment. We performed quantum state tomography on all states used in the experiment, for all used values of n . For ququarts we found a median fidelity of $\overline{\mathcal{F}} = 0.998_{-0.002}^{+0.002}$ and median purity of $\overline{\mathcal{P}} = 0.999_{-0.003}^{+0.001}$, where the error bounds correspond to the first quantile of experimental values (i.e. $\sim 68\%$ of values).

Experimental Errors

Since the states used in the experiment are obtained from numerical optimization they are in general *not* PP-incompatible, which means that the probabilities $P_{M_{j_1 j_2}}(m_i | \psi_i)$ in the numerator of inequality (4.8) are not all expected to vanish. Instead, some vanish and others are just

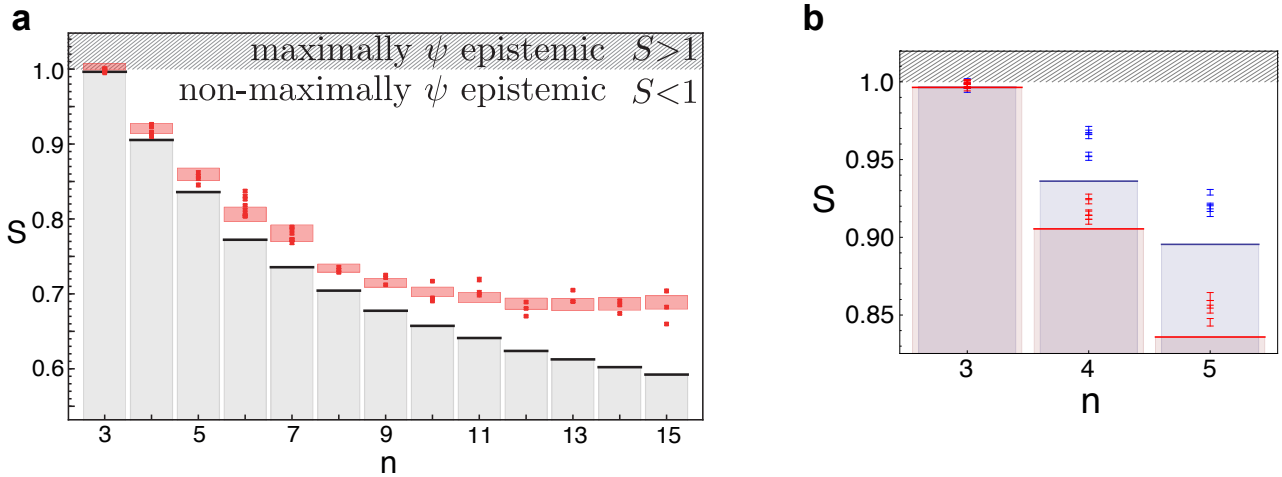


Figure 4.11: **Experimental violation of inequality (4.8) for qutrits and ququarts.** (a) Values of S using real-valued ququart states, $d=4$, as a function of the number of states n . The thick black line on top of the shaded bars represent theoretical values for the used states and measurements for a given value of n , and the hatched region corresponds to $S \geq 1$, which is compatible with maximally ψ -epistemic models. Any value within the white area of the graph is incompatible with maximally ψ -epistemic models. Each data point contains 1-sigma error bars indicating statistical errors due to Poissonian counting statistics. The red-shaded areas indicate the 1σ range of expected values from the above performance estimation of the optical setup, not including systematic long-term drifts. To take the latter into account at least three data points were measured for each value of n . The spread between these individual runs is in general larger than the error bars and increases with the number of states n . (b) The comparison between the $d=3$ (blue) and $d=4$ (red) cases highlights the advantage of higher-dimensional systems as soon as $n > d$.

very close to zero. In either case, however, measuring a bounded quantity (such as a probability) very close to the boundary leads to a very sensitive experiment and asymmetric error distributions that are skewed away from the boundary. In fact, no realistic experiment with finite data can hope to reproduce a value that ideally lies on the boundary, such as a probability of zero. Finally, the denominator of inequality (4.8) depends on the prepared quantum states, which makes accurate characterization of these states crucial for reliable estimates.

4.5.4 Experimental Limitations and Extensions

The present dual-encoding approach combines the high precision of (two-dimensional) polarization encoding with the high-dimensional nature of path encoding without the need for complicated and imprecise interferometry. In principle the Hilbert-space dimension accessible to this scheme could be increased by factors of two. However, this would require adding an additional interferometer in each arm of the current design, which significantly complicates the alignment and would be detrimental for experimental precision. Other possibilities include a fully path-encoded approach using integrated optics for increased interferometric stability, but also such an approach grows in complexity very quickly as the Hilbert space dimension is increased. It




Figure 4.12: **Error tolerance of the experimental protocol for testing maximally ψ -epistemic models.** The maximal permissible error per measurement ε is shown for a qutrit (orange) and ququart (blue) for $3 \leq n \leq 15$. The error tolerance decreases with dimension and with the number of states (i.e. with anything that allows for a larger violation). The shaded gray area corresponds to the 1σ range of absolute errors per measurement achieved in the experiment, with the solid gray line representing the median of the experimental distribution, which is asymmetric, since the measured quantity is bounded.

might thus be preferable to pursue a naturally high-dimensional system, such as orbital angular momentum of photons (which, however cannot currently reach the required precision), d -level atomic systems, or spin-orbit coupled electrons.

A direct extension of the current approach would be to include quarter-waveplates to access to the full ququart state-space, rather than just the real-valued subspace. This will allow for stronger bounds on S , as soon as $n \geq 7$, as shown in Fig. 4.13. Finally, improving the experimental performance will allow to use a larger number of states before reaching the error-accumulation tradeoff.




Figure 4.13: **Theoretical values for S using real, and complex-valued states and measurements.** For $n \geq 7$, using complex valued states can give a significant advantage over the real-valued subspace. These values are based on numerical optimization and cannot be guaranteed to be optimal.

4.5.5 Three-Outcome Measurement and Single-Output Decomposition

Irrespective of the dimension of the system, inequality (4.8) uses triplets of states and three-outcome measurements. To construct such a measurement for $d > 3$, consider a d -dimensional measurement, where all outcomes, except the first three, are chosen orthogonal to the subspace spanned by the triplet of states under consideration. Quantum theory then assigns vanishing probabilities to these outcomes, and a three-outcome POVM can be constructed by grouping them with one of the first outcomes, e.g. $m'_0 = m_0 + \sum_{i=3}^d m_i$. In practice, however, these probabilities do not vanish exactly and this step might add some experimental noise. Moreover, it is rather challenging to directly project on a large subspace, so instead, all outcomes are measured, and combined in post-processing. The method is nonetheless valid, since the derivation of inequality (4.8) makes no assumption about the POVM.

In order to obtain reliable estimates of the outcome probabilities of a d -outcome measurement it is crucial to ensure that there is no systematic bias among the outcomes. In a direct implementation of the d -outcome measurement this would require balancing of losses, coupling- and detection-efficiencies for all outcomes to within 10^{-3} , which is extremely challenging. In the present implementation the d measurement outcomes are instead measured sequentially in close succession. Thus only one fibre-coupling assembly and detector is used, which eliminates the main sources of systematic bias. However, since such a successive measurement design requires different waveplate settings for each measurement outcome, it might be susceptible to errors introduced by steering effects or other imperfections. In order to ensure that the implemented operation is indeed a proper 3-outcome measurement (which need not be the ideal measurement) one has to ensure that the normalised sum of the measurement operators for the three outcomes is as close as possible to the identity operator. Using the calibration data for the optical elements the average drop in measurement fidelity caused by these imperfections can be bounded to 0.0007 ± 0.0002 .

4.6 Discussion and Outlook

A violation of inequality (4.8) demonstrates that *no ψ -epistemic model that reproduces the experimental observations can fully explain the limited distinguishability of non-orthogonal quantum states in terms of overlapping probability distributions*. In other words, it rules out maximally ψ -epistemic interpretations of quantum mechanics. Crucially, observing a value $S(\{\psi_j\}, \{M_{j_1 j_2}\}) < 1$ does not mean that a ψ -epistemic model can at most explain this amount of indistinguishability. Instead, it means that within the chosen set of states there is at least pair of states with a ratio of classical-to-quantum overlap of at most $S(\{\psi_j\}, \{M_{j_1 j_2}\})$.

For this reason the choice of states and measurements is crucial (and the name of the S -parameter very clumsy). As a result, the S -parameter is also of limited use for classifying ψ -epistemic models between the two extremes of maximally ψ -epistemic and ψ -ontic. Without

any auxiliary assumptions it remains possible to consider models where different mechanisms explain the indistinguishability of different sets of states. Since inequality (4.8) is symmetric under unitary transformations the particular parameter of interest is the inner product (or quantum overlap) of the used states. Although inequality (4.8) was derived for pure states, it can straightforwardly be generalized to mixed states [4], which has been found, at least in some cases, to allow for stronger bounds on the overlap ratio [90].

Comparison to PBR

It is worth spending a few moments to compare the present approach with the seminal PBR result. In the noiseless case, PBR establishes an overlap ratio of 0 (i.e. the classical overlap vanishes) for any pair of non-orthogonal quantum states with small enough inner product, using two copies of the state. In order to extend this result to states with arbitrary quantum overlap, they have to use n copies of the states, with $n \rightarrow \infty$. Taking into account experimental noise, the PBR argument is not able to completely rule out ψ -epistemic models anymore. Instead, it only places bounds on the absolute value of the classical overlaps. Once the PBR result is established for a family of states it follows that all states with smaller quantum overlap also have vanishing classical overlap, but not for states with larger overlap. However, even in the noisy case the experiment uses composite systems and has to rely on the preparation independence assumption.

In contrast, the present result uses a single copy of each state and establishes an upper bound on the minimal pairwise overlap ratio within a set of states. It does in general not give information about the individual overlaps, and a bound of 0 can only be reached asymptotically, as the number of states in the set gets arbitrarily large. An observed value of $S < 1$ implies that there is at least one pair of states with a classical-to-quantum overlap-ratio of at most S , while all other states could in principle have a ratio of 1. Without additional assumptions this cannot be extended to arbitrary pairs of states.

4.6.1 Where to From Here?

With maximally ψ -epistemic models ruled out, we are left with a number of options, which all come with their own problems.

Maintaining the notion of objective, observer-independent reality in the sense of the ontological model framework, one can either subscribe to non-maximally ψ -epistemic models or adopt the ψ -ontic viewpoint. While the former are not completely ruled out, their main feature, overlapping epistemic states, is insufficient to fully explain the observed indistinguishability and must be supplemented with coarse-grained measurements. They are also constrained by the existing experimental results and by no-go theorems which rule out many appealing sub-classes of these models. The ψ -ontic view, on the other hand, faces the same old issues, such as the measurement problem, where the projection postulate in a ψ -ontic interpretation implies an instantaneous, probabilistic change in the system's objective reality, which was criticized in Ref. [55]. Moreover, realist interpretations must provide an explanation for how the ontology

can be contextual, while measurement probabilities only depend on the specific measurement operator, or why the underlying reality seems to allow for superluminal effects, while all observations obey signal locality. A more detailed discussion of the consequences of a ψ -ontic interpretation can be found in Ref. [4].

A popular alternative option is to give up the notion of observer-independent reality and consider agent-centric interpretations, such as QBism [24, 25, 80]. Without objective reality, they circumvent many of the explanatory issue that ψ -ontic interpretations have to face up to. The collapse in the measurement problem is simply interpreted as a Bayesian update of the agent’s probability assignment in light of new evidence. The consequences of Bell’s theorem, on the other hand can be avoided by either giving up the idea that correlations need to be explained [73], or by assuming a subjective agent-centric perspective that denies the existence of correlations over spacelike separated regions. Nonetheless, such interpretations typically lose the notion of “explaining” observations in favour of simply accommodating them.

Finally, if the existing ontological models are not appealing, or the idea of an interpretation not providing any explanation for its predictions seems unsatisfactory, one could consider more exotic ontologies, see e.g. Ref. [4]. One such option are retrocausal interpretations, which are typically ψ -ontic, although consistent histories could be read as a ψ -epistemic retrocausal interpretation. Another option is to consider ontologies with multiple realities, such as the ψ -ontic many worlds, or relational interpretations. The latter, however, venture into the territory of anti-realism, since they feature frame-dependent reality, which makes it difficult for observers to agree on events.

Coarse Graining

In a ψ -ontic models the imperfect distinguishability of non-orthogonal quantum states is explained by a restriction that measurements can only reveal coarse-grained information about the ontic states, reminiscent of the uncertainty principle, see Chap. 6. Indeed, in deterministic ontological models the uncertainty principle requires the measurements to be coarse-grained, independent of the status of the quantum state in these models [91]. Related to this it has been argued that no ontological model, which satisfies a notion of free choice and locality, can allow for more precise predictions than quantum mechanics [87].

Curiously, coarse-grained measurements seem to be a common feature—albeit not necessarily involved in explaining indistinguishability—in most ontological models, including maximally ψ -epistemic ones such as Spekkens’ toy model [5] or the Kochen-Specker model [64]. Spekkens’ toy model, for example, is based on the so-called *knowledge-balance* principle, which asserts that only half the information needed to determine the ontic state can be known. To satisfy this principle the model features overlapping probability distributions, coarse-grained measurements, and measurement disturbance. Since the model is maximally ψ -epistemic, however, the coarse-graining is not necessary to explain the indistinguishability of non-orthogonal states.

In non-maximally ψ -epistemic models, on the other hand, coarse-graining would have to partially explain the imperfect distinguishability. Although it might seem excessive to have two

mechanism to explain one phenomenon, the study of worked-out ontological models suggests that it is not entirely implausible to expect both to play a role [5, 91]. This possibility is also interesting in light of the observation that the overlap ratio only goes to zero as the states become more distinguishable. Hence, all current theory and experimental results are indeed compatible with models where coarse-graining explains most of the overlap for states that are far from indistinguishable and epistemic overlap only gains importance as the states come closed together. In order to rule out this possibility it is important to establish bounds for states with large quantum overlaps.

Dynamics in Ontological Models

So far a lot of attention has been on prepare-and-measure experiments, which is a promising direction that allows for very general results without additional assumptions. Going beyond this simplest case, it will be interesting to study state update upon measurement, such as in Spekkens' toy model, or dynamics in ontological models. The latter is, for example, considered explicitly in the works of Hardy [53], and Colbeck-Renner [51], and implicitly by Aaronson [62] via restricting to an ontic state-space where dynamics at the ontic level are as in quantum mechanics. In fact, Colbeck-Renner only require that adding ancillas preserves ontological indistinctness, which is satisfied if the action of a unitary is described as a stochastic map. Hardy goes a step further by imposing ontic indifference, which assumes that unitary dynamics that keep the quantum state fixed correspond to dynamics that keep the ontic state fixed.

4.6.2 Limitations and Possible Extensions of the Method

An interesting observation is that in all current works on the topic, stronger bounds on overlap ratios are only achieved for states with decreasing quantum overlap, see Fig. 4.14. In other words, in cases where overlapping probability distributions can only explain a small fraction of the overlap, there is also almost no overlap to be explained in the first place. The states used in Ref. [58], for example, have an inner product that scales as $|\langle \psi_0^{(n)} | \psi_j^{(n)} \rangle|^2 = \frac{1}{4} n^{-1/(d-2)}$. This allows for an alternative interpretation of Branciard's result as showing that there exist states $|\psi\rangle$ and $|\phi\rangle$, such that

$$\kappa(\psi, \phi) \leq \frac{4^d}{8} |\langle \psi_0^{(n)} | \psi_j^{(n)} \rangle|^{2(d-3)} \quad \text{with} \quad |\langle \psi | \phi \rangle| \rightarrow 0$$

This behaviour is common for all current results and would indeed be expected from models where coarse-graining and epistemic overlap both play a role in explaining quantum indistinguishability. Such models are compatible with current experimental results and could only be ruled out by showing that overlapping probability distributions contribute very little to explaining the indistinguishability, independent of the quantum overlap.




Figure 4.14: **Quantum overlaps of all states used in the experiment.** The solid line is the average quantum overlap as a function of n . Since the states are numerically optimized, they are in general not symmetrically distributed, although in some cases, such as $n = 12$, the optimization routine actually settles on a symmetric distribution with equal overlaps.

Overlap Measures and Figures of Merit

The overlap measures ω_c and ω_q used here have a clear operational interpretation in terms of single-shot distinguishability with optimal measurements. However, this choice, which leads to symmetric overlap regions, is not unique. Interestingly, using other measures can produce qualitatively different results [4]. Using, for example, asymmetric overlap region associated with unambiguous state discrimination [92], Spekkens’ toy model and the Kochen-Specker model turn out to be non-maximally ψ -epistemic [91], while they are maximally ψ -epistemic with respect to the present measure. This dependence implies that “epistemicity” can be defined with respect to a variety of measures, which is an interesting area of further research.

Another point of discussion is the use of the overlap ratio κ as a figure of merit for studying ψ -epistemic models. As a relative measure, κ quantifies the fraction of the indistinguishability explained by a model, but fails to capture how much indistinguishability there is to be explained in the first place. This can be somewhat deceiving, since $S \rightarrow 0$ implies that ψ -epistemic models can only explain an arbitrarily small fraction, but in all current work, this limit is only achieved for states with an arbitrarily small amount of quantum overlap to be explained.

As an alternative it has been suggested to measure the absolute difference between the quantum and classical overlaps, the *overlap deficit* [91].

$$\delta_\omega(\psi_0, \psi_j) = \omega_q(\psi_0, \psi_j) - \omega_c(\psi_0, \psi_j) . \quad (4.15)$$

A bound on the maximal overlap deficit for the chosen set of states can directly be obtained from that on the overlap ratio, as follows: for a given value of S , there exists a pair of states (ψ_0, ψ_j) such that $\kappa(\psi_0, \psi_j) \leq S$. For those two states,

$$\begin{aligned}
\delta_\omega(\psi_0, \psi_j) &\equiv \omega_Q(|\psi_0\rangle, |\psi_j\rangle) - \omega_C(\mu_{\psi_0}, \mu_{\psi_j}) \\
&= \omega_Q(|\psi_0\rangle, |\psi_j\rangle) (1 - \kappa(\psi_0, \psi_j)) \\
&\geq \min_{j'} [\omega_Q(|\psi_0\rangle, |\psi_{j'}\rangle)] (1 - S) ,
\end{aligned}$$

and hence

$$\max_j [\delta_\omega(\psi_0, \psi_j)] \geq \min_j [\omega_Q(|\psi_0\rangle, |\psi_j\rangle)] (1 - S) . \quad (4.16)$$

The states and measurements used in the experiment in Sec. 4.5 can thus also be used directly to derive bounds on the overlap deficit. However, since these states and measurements were optimized for S , they are not optimal for the overlap deficit as a figure of merit.

The main reason for considering the overlap deficit is that it has a more instructive limiting behaviour, since $\delta_\omega \rightarrow 1$ is only achieved for quantum states that are completely overlapping, while the epistemic states are completely disjoint. However, for experimentally accessible values away from the asymptotic limits, both measures have similar shortcomings. A value of $S = 0.7$ implies that the model can explain 70% of the quantum overlap, but gives no indication if that is a lot of overlap or not. On the other hand, a value of $\delta_\omega = 0.3$ implies an absolute difference of 0.3 between quantum and classical overlap, but does not indicate what this says about the explanatory power of the model. The most informative would thus be a combination of either quantity and a list of quantum overlaps for the prepared states, or a combination of both, overlap deficit and overlap ratio.

Fair-sampling Assumption

Any photonic experiment suffers from optical losses and imperfect detection efficiency. To address this most experiments rely on a fair-sampling assumption, which asserts that the set of detected events (e.g. detected photons) is representative for the full ensemble. Notable exceptions are the loophole-free Bell-inequality tests of 2015 [93, 94], see Chap. 5. It is often under-appreciated, that there are actually two aspects to this assumption. On the practical side, fair-sampling must be justified by making sure that any kind of loss or inefficiency is unbiased and affects every photon equally (e.g. there are no polarization-dependent effects). Fundamentally, fair-sampling assumes that the tested ontological model is fair-sampled in the sense that it is not the case that some ontic states are more likely to be detected than others, thus giving a biased result. The latter cannot be justified experimentally, and experiments therefore typically test the fair-sampled version of an ontological model. Similarly, the experiment presented here relies on a fair-sampling assumption.

To estimate the required detection efficiency to avoid the fair-sampling requirement, one can take into account non-detection events by assigning a fixed or random outcome to them. Suppose for simplicity that the detection efficiencies for all three outcomes and for all states and measurements are the same, η , and that the experimenter chooses to output m_0 whenever none of the detectors click. Hence, the probabilities for m_0 are given by $P_{M_{j_1 j_2}}^{(\eta)}(m_0|\psi_0) =$

$\eta P_{M_{j_1 j_2}}^{(\eta=1)}(m_0|\psi_0) + (1-\eta)$ and for m_i with $i \geq 1$ by $P_{M_{j_1 j_2}}^{(\eta)}(m_i|\psi_{j_i}) = \eta P_{M_{j_1 j_2}}^{(\eta=1)}(m_i|\psi_{j_i})$. The value of S as defined in Eq. (4.8) is then given by

$$S^{(\eta)}(\{\psi_j\}, \{M_{j_1 j_2}\}) = \frac{1 + \sum_{j_1 < j_2} \left[(1-\eta) + \eta \sum_{i=0}^2 P_{M_{j_1 j_2}}^{(\eta=1)}(m_i|\psi_{j_i}) \right]}{\sum_j \omega_Q(|\psi_0\rangle, |\psi_j\rangle)}. \quad (4.17)$$

Note that the same value would be obtained here if the no-detection events were replaced by random outcomes. Ruling out maximally ψ -epistemic models requires $S^{(\eta)} < 1$, and thus

$$\eta > \frac{1 + \frac{n(n-1)}{2} - \sum_j \omega_Q(|\psi_0\rangle, |\psi_j\rangle)}{\frac{n(n-1)}{2} - \sum_{j_1 < j_2} \sum_{i=0}^2 P_{M_{j_1 j_2}}^{(\eta=1)}(m_i|\psi_{j_i})}. \quad (4.18)$$

As a consequence of the tight error threshold for a violation of inequality (4.8) the efficiency requirements for a test without the fair-sampling assumption are naturally much stricter than, for example, in the case of Bell inequalities. Specifically, using the states and measurements of the present experiment, the best scenario, $d = 4, n = 5$ requires an overall efficiency of at least ~ 0.976 . While this is currently out of reach for photonic implementations, other architectures routinely achieve such efficiencies and would be more suitable for such an experiment.

4.6.3 Bell Inequalities and Device Independence

When it comes to tests of ontological models, the first thing that comes to mind is usually Bell's theorem. In contrast to the prepare-and-measure experiments discussed in this chapter, Bell's theorem is concerned with the correlations between two spatially separated quantum systems. Under some additional assumptions one can derive Bell inequalities, which must be satisfied by the correlations produced by any Bell-local ontological model, see Chap. 5 for details. Quantum correlations, on the other hand violate this inequality. So far this is very similar to the present case with inequality (4.8).

However, Bell inequality have the remarkable property of being device independent. This means that they do not make any assumptions about the mechanism that generated the observed correlations and could thus be tested using a pair of black boxes. This powerful feature allows Bell-inequality violations to be used for certifying that two untrusted parties share entanglement, which is crucial for security purposes, see Chap. 3. From a foundational perspective this means that Bell inequalities can be used to rule out local-causal models simply by finding correlations that violate the inequality, without having to commit to quantum mechanics for describing these correlations.

In the case of testing maximally ψ -epistemic models, however, the situation is different. Instead of asking whether a model could have produced some experimental observations, the experimental observations are used to constrain whether a model can quantitatively reproduce quantum predictions. This approach is thus necessarily device dependent, in that quantum mechanics is required to compute the reference to compare against. Here, this reference is the

sum of quantum overlaps in the denominator of (4.8). Whether device-independent tests in this generality are possible thus remains an open question. However, some models, which are often counted towards the ψ -ontic camp are actually alternative theories and are known to make predictions that deviate slightly from quantum mechanics. One example is gravitational collapse models, which predict that quantum superpositions of objects above a certain mass scale cannot be maintained for a significant amount of time. Such models make predictions that are explicitly different from quantum mechanics and thus lend themselves more naturally to device-independent tests [49].

References

- [1] Ringbauer, M., Duffus, B., Branciard, C., Cavalcanti, E. G., White, A. G. & Fedrizzi, A. Measurements on the reality of the wavefunction. *Nat. Phys.* **11**, 249–254 (2015).
- [2] Pais, A. Einstein and the quantum theory. *Rev. Mod. Phys.* **51**, 863–914 (1979).
- [3] Mermin, N. D. Could feynman have said this? *Physics Today* **57**, 10–11 (2004).
- [4] Leifer, M. S. Is the quantum state real ? An extended review of ψ -ontology theorems. *Quanta* **3**, 67–155 (2014).
- [5] Spekkens, R. Evidence for the epistemic view of quantum states: A toy theory. *Phys. Rev. A* **75**, 032110 (2007).
- [6] Jennings, D. & Leifer, M. No return to classical reality. *Contemp. Phys.* **57**, 60–82 (2016).
- [7] Bartlett, S. D., Rudolph, T. & Spekkens, R. W. Reconstruction of Gaussian quantum mechanics from Liouville mechanics with an epistemic restriction. *Phys. Rev. A* **86**, 012103 (2012).
- [8] Bell, J. S. *Speakable and Unspeakable in Quantum Mechanics* (Cambridge University Press, 1987).
- [9] Bell, J. S. On the Problem of Hidden Variables in Quantum Mechanics. *Rev. Mod. Phys.* **38**, 447–452 (1966).
- [10] Einstein, A. *Out of My Later Years* (Citadel Press, 1956).
- [11] Harrigan, N. & Spekkens, R. W. Einstein, Incompleteness, and the Epistemic View of Quantum States. *Found. Phys.* **40**, 125–157 (2010).
- [12] Caves, C. M. & Fuchs, C. A. Quantum information: How much information in a state vector? *arXiv:quant-ph/9601025* (1996).
- [13] Daffertshofer, A., Plastino, A. R. & Plastino, A. Classical No-Cloning Theorem. *Phys. Rev. Lett.* **88**, 210601 (2002).
- [14] Oppenheimer, R. The Philosophy of Niels Bohr. *Bulletin of the Atomic Scientists* **19** (1963).
- [15] Cabello, A. Interpretations of quantum theory: A map of madness. *arXiv:1509.04711* (2015).

- [16] Mermin, N. D. Quantum mechanics: Fixing the shifty split. *Physics Today* **65**, 8 (2012).
- [17] Schlosshauer, M., Kofler, J. & Zeilinger, A. A snapshot of foundational attitudes toward quantum mechanics. *Stud. Hist. Philos. Mod. Phys.* **44**, 222–230 (2013).
- [18] Sommer, C. Another Survey of Foundational Attitudes Towards Quantum Mechanics. *arXiv:1303.2719* (2013).
- [19] Norsen, T. & Nelson, S. Yet Another Snapshot of Foundational Attitudes Toward Quantum Mechanics. *arXiv:1306.4646* (2013).
- [20] Brukner, C. On the quantum measurement problem. *arXiv:1507.05255* 19–22 (2015).
- [21] Bohr, N. *The Philosophical Writings of Niels Bohr* (Ox Bow Press, 1987).
- [22] Einstein, A. Physik und realität. *J. Franklin Inst.* **221**, 313–347 (1936).
- [23] Ballentine, L. The Statistical Interpretation of Quantum Mechanics. *Rev. Mod. Phys.* **42**, 358–381 (1970).
- [24] Fuchs, C. A. QBism, the Perimeter of Quantum Bayesianism. *arXiv:1003.5209* (2010).
- [25] Mermin, N. D. QBism puts the scientist back into science. *Nature* **507**, 421–423 (2014).
- [26] Fuchs, C. A., Mermin, N. D. & Schack, R. An introduction to QBism with an application to the locality of quantum mechanics. *Am. J. Phys.* **82**, 749–754 (2014).
- [27] Zeilinger, A. A foundational principle for quantum mechanics. *Found. Phys.* **29**, 631–643 (1999).
- [28] Paterek, T., Dakić, B. & Brukner, Č. Theories of systems with limited information content. *New J. Phys.* **12** (2010).
- [29] Brukner, C. & Zeilinger, A. Information and fundamental elements of the structure of quantum theory. *arXiv:quant-ph/0212084* (2002).
- [30] Rovelli, C. Relational quantum mechanics. *Int. J. Theor. Phys.* **35**, 1637–1678 (1996).
- [31] Bohm, D. A Suggested Interpretation of the Quantum Theory in Terms of "Hidden" Variables. I & II. *Phys. Rev.* **85**, 166–193 (1952).
- [32] Passon, O. What you always wanted to know about Bohmian mechanics but were afraid to ask. *arXiv:quant-ph/0611032* (2006).
- [33] Ghirardi, G. C., Rimini, A. & Weber, T. Unified dynamics for microscopic and macroscopic systems. *Phys. Rev. D* **34**, 470–491 (1986).
- [34] Weinberg, S. Collapse of the state vector. *Phys. Rev. A* **85**, 062116 (2012).

- [35] Milburn, G. J. Intrinsic decoherence in quantum mechanics. *Phys. Rev. A* **44**, 5401–5406 (1991).
- [36] Everett, H. I. “Relative State” Formulation of Quantum Mechanics. *Rev. Mod. Phys.* **29**, 454–462 (1957).
- [37] Griffiths, R. B. A consistent quantum ontology. *Stud. Hist. Philos. Mod. Phys.* **44**, 93–114 (2013).
- [38] Hohenberg, P. C. Colloquium : An introduction to consistent quantum theory. *Rev. Mod. Phys.* **82**, 2835–2844 (2010).
- [39] Lombardi, O. & Dieks, D. Modal interpretations of quantum mechanics. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (2016), spring 2016 edn.
- [40] Dieks, D. & Vermaas, P. E. (eds.) *The Modal Interpretation of Quantum Mechanics*, vol. 60 of *The Western Ontario Series in Philosophy of Science* (Springer Netherlands, 1998).
- [41] Cramer, J. G. The transactional interpretation of quantum mechanics. *Rev. Mod. Phys.* **58**, 647–687 (1986).
- [42] Kastner, R. E. *The Transactional Interpretation of Quantum Mechanics* (Cambridge University Press, 2013).
- [43] Aharonov, Y. & Vaidman, L. The Two-State Vector Formalism of Quantum Mechanics: an Updated Review. *arXiv:quant-ph/0105101* (2001).
- [44] Price, H. Toy models for retrocausality. *Stud. Hist. Philos. Mod. Phys.* **39**, 752–761 (2008).
- [45] Evans, P. W. Retrocausality at no extra cost. *Synthese* **192**, 1139–1155 (2015).
- [46] Evans, P. W. Quantum causal models, faithfulness and retrocausality. *arXiv:1506.08925* (2015).
- [47] Hall, M. J., Deckert, D.-A. & Wiseman, H. M. Quantum Phenomena Modeled by Interactions between Many Classical Worlds. *Phys. Rev. X* **4**, 041013 (2014).
- [48] Zurek, W. H. Quantum Darwinism. *Nat. Phys.* **5**, 181–188 (2009).
- [49] Bassi, A. Models of spontaneous wave function collapse: what they are, and how they can be tested. *J. Phys.* **701**, 012012 (2016).
- [50] Pusey, M. F., Barrett, J. & Rudolph, T. On the reality of the quantum state. *Nat. Phys.* **8**, 476–479 (2012).
- [51] Colbeck, R. & Renner, R. Is a system’s wave function in one-to-one correspondence with Its elements of reality? *Phys. Rev. Lett.* **108**, 150402 (2012).

- [52] Patra, M. K., Pironio, S. & Massar, S. No-Go Theorems for ψ -Epistemic Models Based on a Continuity Assumption. *Phys. Rev. Lett.* **111**, 090402 (2013).
- [53] Hardy, L. Are Quantum States Real? *Int. J. Mod. Phys. B* **27**, 1345012 (2013).
- [54] Emerson, J., Serbin, D., Sutherland, C. & Veitch, V. The whole is greater than the sum of the parts: on the possibility of purely statistical interpretations of quantum theory. *arXiv:1312.1345* (2013).
- [55] Cabello, A., Gu, M., Gühne, O., Larsson, J.-Å. & Wiesner, K. Thermodynamical cost of some interpretations of quantum theory. *arXiv:1509.03641* (2015).
- [56] Frauchiger, D. & Renner, R. Single-world interpretations of quantum theory cannot be self-consistent. *arXiv:1604.07422* (2016).
- [57] Shimony, A. Search for a worldview which can accommodate our knowledge of microphysics. In *Philosophical Consequences of Quantum Theory* (University of Notre Dame Press, 1989).
- [58] Branciard, C. How ψ -Epistemic Models Fail at Explaining the Indistinguishability of Quantum States. *Phys. Rev. Lett.* **113**, 020409 (2014).
- [59] Einstein, A., Podolsky, B. & Rosen, N. Can Quantum-Mechanical Description of Physical Reality Be Considered Complete? *Phys. Rev.* **47**, 777–780 (1935).
- [60] Barrett, J., Cavalcanti, E. G., Lal, R. & Maroney, O. J. E. No ψ -Epistemic Model Can Fully Explain the Indistinguishability of Quantum States. *Phys. Rev. Lett.* **112**, 250403 (2014).
- [61] Leifer, M. S. ψ -Epistemic Models are Exponentially Bad at Explaining the Distinguishability of Quantum States. *Phys. Rev. Lett.* **112**, 160404 (2014).
- [62] Aaronson, S., Bouland, A., Chua, L. & Lowther, G. ψ -epistemic theories: The role of symmetry. *Phys. Rev. A* **88**, 032111 (2013).
- [63] Bell, J. S. The theory of local beables. *Epistemological Lett.* **9**, 11–24 (1976).
- [64] Kochen, S. B. & Specker, E. The Problem of hidden variables in quantum mechanics. *J. Math. Mech.* **17**, 59–87 (1967).
- [65] Spekkens, R. W. Contextuality for preparations, transformations, and unsharp measurements. *Phys. Rev. A* **71**, 052108 (2005).
- [66] Hardy, L. Quantum ontological excess baggage. *Stud. Hist. Philos. Mod. Phys.* **35**, 267–276 (2004).
- [67] Montina, A. Exponential complexity and ontological theories of quantum mechanics. *Phys. Rev. A* **77**, 022104 (2008).

- [68] Montina, A. State-space dimensionality in short-memory hidden-variable theories. *Phys. Rev. A* **83**, 032107 (2011).
- [69] Howard, D. Einstein on locality and separability. *Stud. Hist. Philos. Sci.* **16**, 171–201 (1985).
- [70] Wiseman, H. M. The two Bell’s theorems of John Bell. *J. Phys. A* **47**, 424001 (2014).
- [71] Wiseman, H. M. & Rieffel, E. G. Reply to Norsen’s paper ”Are there really two different Bell’s theorems?”. *arXiv:1503.06978* (2015).
- [72] Bell, J. S. On the Einstein Podolsky Rosen Paradox. *Physics* **1**, 195–200 (1964).
- [73] Wiseman, H. M. & Cavalcanti, E. G. Causarum Investigatio and the Two Bell’s Theorems of John Bell. *arXiv:1503.06413* (2015).
- [74] Mazurek, M. D., Pusey, M. F., Kunjwal, R., Resch, K. J. & Spekkens, R. W. An experimental test of noncontextuality without unphysical idealizations. *Nat. Commun.* **7**, 11780 (2016).
- [75] Leifer, M. S. & Maroney, O. J. E. Maximally Epistemic Interpretations of the Quantum State and Contextuality. *Phys. Rev. Lett.* **110**, 120401 (2013).
- [76] Beltrametti, E. G. & Bugajski, S. A classical extension of quantum mechanics. *J. Phys. A* **28**, 3329–3343 (1995).
- [77] Harrigan, N. & Rudolph, T. Ontological models and the interpretation of contextuality. *arXiv:0709.4266* (2007).
- [78] Huxley, T. H. *Collected Essays VIII* (1894).
- [79] Lewis, P. G., Jennings, D., Barrett, J. & Rudolph, T. Distinct Quantum States Can Be Compatible with a Single State of Reality. *Phys. Rev. Lett.* **109**, 150404 (2012).
- [80] Caves, C., Fuchs, C. & Schack, R. Conditions for compatibility of quantum-state assignments. *Phys. Rev. A* **66**, 062111 (2002).
- [81] Nigg, D., Monz, T., Schindler, P., Martinez, E. A., Hennrich, M., Blatt, R., Pusey, M. F., Rudolph, T. & Barrett, J. Can different quantum state vectors correspond to the same physical state? An experimental test. *New J. Phys.* **18**, 013007 (2015).
- [82] Spekkens, R. Why I Am Not a Psi-ontologist. *Lecture at Perimeter Institute, PIRSA:12050021* (2012).
- [83] Schlosshauer, M. & Fine, A. Implications of the Pusey-Barrett-Rudolph Quantum No-Go Theorem. *Phys. Rev. Lett.* **108**, 260404 (2012).

- [84] Hall, M. J. W. Generalisations of the recent Pusey-Barrett-Rudolph theorem for statistical models of quantum phenomena. *arXiv:1111.6304* (2011).
- [85] Colbeck, R. private communication.
- [86] Branciard, C. private communication.
- [87] Colbeck, R. & Renner, R. No extension of quantum theory can have improved predictive power. *Nat. Commun.* **2**, 411 (2011).
- [88] Patra, M. K., Ollslager, L., Duport, F., Safioui, J., Pironio, S. & Massar, S. Experimental refutation of a class of ψ -epistemic models. *Phys. Rev. A* **88**, 032112 (2013).
- [89] Leifer, M. The reality of the quantum state from Kochen-Specker contextuality. *Talk at EmQM15, Vienna* (2015).
- [90] Knee, G. C. Towards optimal experimental tests on the reality of the quantum state. *arXiv:1609.01558* (2016).
- [91] Leifer, M. private communication.
- [92] Herzog, U. & Bergou, J. Distinguishing mixed quantum states: Minimum-error discrimination versus optimum unambiguous discrimination. *Phys. Rev. A* **70**, 022302 (2004).
- [93] Shalm, L. K. *et al.* Strong Loophole-Free Test of Local Realism. *Phys. Rev. Lett.* **115**, 250402 (2015).
- [94] Giustina, M. *et al.* Significant-Loophole-Free Test of Bell's Theorem with Entangled Photons. *Phys. Rev. Lett.* **115**, 250401 (2015).

CHAPTER 5

Causality in a Quantum World

Acknowledgement

This chapter is based on work that was first published in Ref. [1], and, where appropriate, I have incorporated text of that paper. The experiments were performed with Christina Giarmatzi and the theory was largely developed by Rafael Chaves.

5.1 Introduction

THE previous chapter outlined the classification of various interpretations of quantum mechanics with respect to the status of the quantum wavefunction. The results discussed there suggest that, if one wants to maintain realism, then the wavefunction cannot be treated as a state of knowledge. However, this is not the only constraint on realist interpretations of quantum mechanics. One of the earliest such results was formulated by John Bell in 1964 [2], and again, but differently in 1976 [3]. In modern language Bell’s result implies that the correlations between entangled quantum systems cannot be explained in terms of, possibly hidden, cause and effect relations.

After four decades of experimental effort, Bell’s prediction has been confirmed in an unambiguous way [4–6]. These experiments, however, cannot reveal where exactly quantum mechanics breaks with our classical notion of cause and effect. One way to reconcile the conflict is by introducing nonlocal causal influences, which has coined the term “quantum nonlocality”, but this is not the only possibility [7–9]. Consequently the question of how much the various assumptions have to give way to recover a causal explanation has become an exciting area of research [10–18]. The framework of causal modeling has proven an excellent platform for the study of stronger-than-classical correlation [16, 19–25], and non-classical causality [26, 27]. This chapter focuses on using the theory of causal modeling to study relaxations of Bell’s assumptions, in particular relaxations of the local causality assumption. This approach leads to testable inequalities, which is used in Sec. 5.4 to rule out a class of causal models that aim to explain quantum correlations by means of a nonlocal causal influence from one measurement outcome to the other.

5.2 The Causal Modeling Framework

“Correlation doesn’t imply causation, but it does waggle its eyebrows suggestively and gesture furtively while mouthing ‘look over there’.” - R. Munroe [28]

Cause-and-effect relations are ubiquitous for describing observed correlations in empirical science. So ubiquitous in fact, that it was not until recently that causation was realized as an implicit assumption and put on firm mathematical grounds in the form of the causal modeling framework [29–31]. The notion of causality is in this framework based on the fundamental idea that if a variable acts as the cause for another one, actively changing or *intervening* on the first should result in a change in the second. This notion of intervention is the crucial aspect that enables a distinction between correlation and causation. Causal modeling then provides a powerful language to translate, in both directions, between intuitive networks of cause-effect relations and abstract conditional probability distributions for a set of events. The purpose of this is two-fold: a given structure of causal dependences and independences imposes constraints on the allowed joint probabilities, and given a joint probability distribution it is possible to discover the underlying causal structure.




Figure 5.1: **The causal modeling framework.** The causal modeling framework translates between **(a)** joint probability distributions and **(b)** an intuitive structure of causal dependences and independences.

5.2.1 A Causal Model

Formally, a causal model consists of a causal structure, represented by a *directed acyclic graph*, see Fig. 5.1b, and a set of causal-statistical parameters. The nodes of the graph (circles in Fig. 5.1b), correspond to distinct events such as measurements, and each event is associated with a random variable that encodes the potential values, such as potential measurement outcomes or potential values of a physical property. Two nodes can be connected by a directed edge (arrows in Fig. 5.1b), which represents a causal influence from one node to the other. Note that, although the flow of causal influences is often associated with the flow of physical time, it is not necessary to make this association and there is no a priori notion of time in a causal

model. Furthermore, the graph is “acyclic”, meaning that there are no closed paths or cycles since otherwise events may cause themselves, which defies the idea of causal explanation. The causal relations between causally connected nodes are described in intuitive family terms such as ancestors and children, see Fig. 5.2a. Any causal graph can be described in terms of three main building blocks: chain, collider and fork, see Fig. 5.2b.

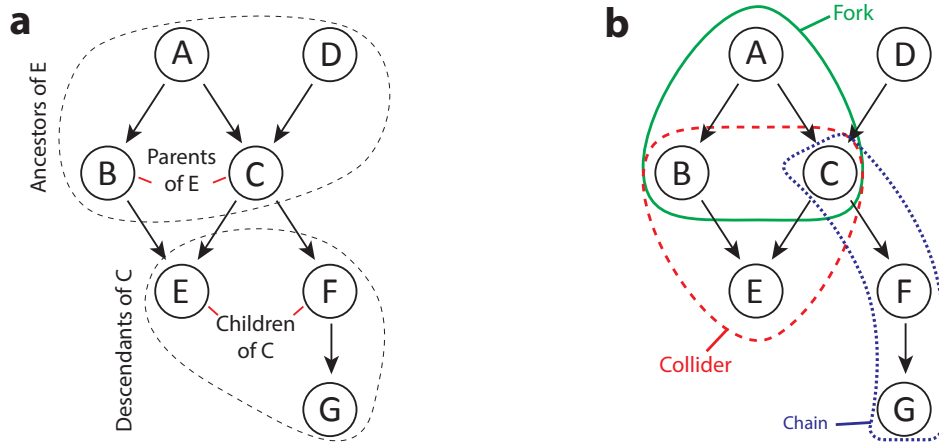


Figure 5.2: **Basic notations of causal models.** (a) Relationships between nodes of the graph are denoted using family terms. Direct causes of a variable are called its parents and variables are called children of their direct causes. The set of all generations of parents (children) of a variable are called its ancestors (descendants). (b) Each causal graph is composed of a combination of the three basic elements: forks, colliders and chains.

The causal parameters quantify the causal relations in the model by specifying the conditional probabilities $P(A|\text{Pa}(A))$ for every variable A given its graph-theoretic parents $\text{Pa}(A)$. In the case of so-called *exogenous* (as opposed to endogenous) variables—nodes without parent nodes—this amounts to specifying the initial distribution.

The dependence between endogenous variables is in general probabilistic, but it can be modelled as deterministic functions with added noise. In particular, explicitly adding an uncorrelated *noise variable* to every node turns any probabilistic model into a deterministic one. Conversely, exogenous nodes that only affect a single child-node can be absorbed into the child-node, turning the deterministic dependences of that node into probabilistic ones. In other words, any randomness in a causal model can be traced back to ignorance over some variables. A crucial assumption at this point, the so-called *Markov* assumption, is that all exogenous variables (including noise variables) are independently distributed. This ensures that the causal model is complete in the sense that there are no un-modelled common causes.

The assumption that there exist free variables, which are not caused by anything, is a crucial ingredient of causal models, and the very notion of a “cause” is defined in terms of freely chosen interventions. An intervention is an independent external influence that modifies the causal mechanisms of a given model [32]. In the strongest case, it sets a variable X to one of its values x , while leaving all other variables unperturbed. This operation, which is typically denoted $\text{do}(x)$, thereby overrides the mechanisms that would otherwise determine the value of X . If there is a causal influence from X to Y , then intervening on X (i.e. changing X)

should cause a change in Y . The maximal shift caused by the intervention can then be used to quantify the strength of the causal influences

$$\mathcal{C}_{A \rightarrow B} = \sup_{x, x', y} |P(y|\text{do}(x)) - P(y|\text{do}(x'))|. \quad (5.1)$$

Causal Markov Condition

A Markovian causal model (i.e. one that satisfies the Markov assumption of uncorrelated noise variables) can be fully specified by the conditional probability distributions of all variables, given their parents. Together with the law of total probability this definition implies that the full joint probability distribution can be written in the product form:

$$P(X_1, \dots, X_n) = \prod_i^n P(X_i | Pa(X_i)), \quad (5.2)$$

which is known as the causal Markov condition¹. An equivalent formulation is that “For each pair of distinct X and Y , where X does not cause Y , $P(X|Y, Pa(X)) = P(X|Pa(X))$.” [33]. In other words, every variable is independent of its non-descendants, given its parents ($X \perp\!\!\!\perp Nd(X) | Pa(X)$). The parent nodes, which are typically associated with direct causes, thus “screen off” the effect of indirect causes, which are ancestors further away in the graph.

In a Markovian causal model all correlations are accounted for by the causal influences in the model. Hence, the Markov condition, could be read as implying that for every statistical dependence there should be a causal dependence [34]. Non-Markovian models on the other hand contain some correlations that are not explained by the causal structure, which can be due to *unmeasured* common causes. Whether a model is Markovian or not can thus depend on the level of grain, or the detail of the causal description. When the model cannot fully describe the observed correlations, one can include *latent* or unobserved (maybe even unobservable) variables into the causal structure until the extended model satisfies the Markov condition. This is a dangerous path, since adding enough latent variables everything could be explained in a superdeterministic way, which defies the purpose. However, it might be that there are indeed unmeasured common causes, which could be modelled as latent variables, or as in the case of hidden variable models of quantum mechanics, there might be fundamental restrictions on which variables can and cannot be observed.

Reichenbach’s Common Cause Principle

A central motivation for the causal modeling program is Reichenbach’s common cause principle. Essentially, it states that if two events are correlated, but neither is a direct cause of the other, then there must be a common cause for both, which explains the correlation. More recently Reichenbach’s principle has been decomposed into two separate assumptions [7, 22, 23]: the existence of a common cause and the notion of “explaining” correlations.

¹Strictly speaking, Eq. (5.2) is the Markov condition, and the causal interpretation requires the implicit assumption that interventions on the various variables are possible.

Principle of common cause If two variables are correlated and neither is a direct cause of the other, then there must be a common cause for both.

Factorization of probabilities A common cause Z completely explains the correlations between X and Y (assuming neither is a cause of the other), if and only if conditioning on Z removes the correlations between X and Y . The common cause “screens-off” X and Y , such that their joint distribution factorizes

$$P(X, Y|Z) = P(X|Z)P(Y|Z)$$

Reichenbach’s principle is a guiding idea of classical causal models, and both its parts are inherent features of the definition of a Markovian causal model. In such a model all correlations are accounted for by causal influences from parent-nodes to children-nodes, which implies the principle of common cause. The Markov condition, which again follows from the definition of a Markovian causal model, then implies factorization.

Faithfulness

A causal model is called a *faithful* representation of a probability distribution if all conditional independences in the distribution are a consequence of the causal structure alone, and not due to a specific choice of model parameters. Faithfulness is thus in some sense the converse to the Markov condition and can be read as stating that every observed statistical independence should imply a causal independence [34]. Together the Markov and faithfulness conditions thus establish a one-to-one correspondence between causal and statistical (in-)dependences. However they do not fully determine the causal structure and additional information (e.g. interventional data) is required to establish the direction of causal influences.

Causal models that do not satisfy faithfulness are called *fine-tuned*, since they contain causal influences which are hidden from the statistics by virtue of specifically chosen causal parameters. This could occur, for example, if two variables are causally connected via two different paths, which cancel each other exactly, see Ref. [34] for a detailed discussion. To illustrate the essential features of fine-tuning, consider a simple quantum² example of a source of particles S which produces qubits in ± 1 -eigenstates of σ_Y , and a binary measurement in σ_X or σ_Z , with the setting represented by X and the outcome by A . Due to the specific choice of values for the source, the variables X and A appear statistically independent (A is completely random for any choice of X). However, any disturbance of the quantum states produced by the source will break the apparent independence of X and A .

Fine-tuned causal models are often considered unnatural, since the causal parameters are chosen from a set of measure zero in parameter space. Hence, there is a vanishing probability for a model to be fine-tuned, unless there is a natural mechanism that drives the model to this point in parameter space. One prominent example is Bohmian mechanics, where equilibration leads

²Although one has to be careful when modeling quantum systems with classical causal models, properties such as faithfulness still apply [32].

to fine-tuned parameters. The question of naturalness and fine-tuning is also relevant in other fields, such as cosmology or particle physics. The mass of the Higgs boson, for example, appears to be fine-tuned and physicists are aiming to find a *natural* mechanism that could explain this value in a theory that eventually supersedes the standard model, such as supersymmetry [35]. Finally, although minimal variations of the fine-tuned parameters can completely change the pattern of conditional independences in a causal model, there are intrinsic forms of fine-tuning where such disturbances are not possible [34, 36].

5.2.2 Reading the Causal Graph

The pattern of causal influences in the causal graph is a representation of the conditional independence relations that characterize the causal model. These relations are of the form $(X \perp\!\!\!\perp Y|Z)$ which reads X is independent of Y given Z , and means that knowing Z one cannot gain information about X by learning Y . In terms of probability distributions this is equivalent to the condition $P(X, Y|Z) = P(X|Z)P(Y|Z)$ (“when Z is known then X and Y are independent”) or $P(X|Y, Z) = P(X|Z)$ (“the probability of X does not depend on Y if Z is known”). In terms of conditional independence relations the Markov condition states that $(X \perp\!\!\!\perp nd(X)|Pa(X)) \forall X$; “every variable is independent of its non-descendants $nd(X)$ given its parents $Pa(X)$ ”.

Since the conditional independence relations satisfy the semi-graphoid axioms it is sufficient to specify a generating set, from which all relations of the model can be calculated [23]. Letting X, Y, Z denote sets of events, and $X \cup Y$ the union of such sets the semi-graphoid axioms are [23]

1. *symmetry*: $(X \perp\!\!\!\perp Y|Z) \Leftrightarrow (Y \perp\!\!\!\perp X|Z)$.

Intuitively, if, knowing Z , no knowledge about X can be gained from learning Y , then it is also not possible to gain knowledge about Y from learning X .

2. *decomposition*: $(X \perp\!\!\!\perp Y \cup W|Z) \Rightarrow (X \perp\!\!\!\perp Y|Z) \& (X \perp\!\!\!\perp W|Z)$.

Intuitively, if learning Y and W does not provide information about X , given Z , then neither Y nor W independently provide information about X given Z .

3. *weak union*: $(X \perp\!\!\!\perp Y \cup W|Z) \Rightarrow (X \perp\!\!\!\perp Y|W \cup Z)$.

Intuitively, if learning Y and W does not provide information about X given Z , then knowing W does not change the fact that learning Y does not provide information about X .

4. *contraction*: $(X \perp\!\!\!\perp Y|Z) \& (X \perp\!\!\!\perp W|Z \cup Y) \Rightarrow (X \perp\!\!\!\perp Y \cup W|Z)$.

Intuitively, if Y is irrelevant to X given Z then conditioning on Y cannot change whether W is relevant to X or not.

The d -separation criterion offers an intuitive graphical tool for extracting all conditional independence relations that are implied by a causal graph (and only those). Two variables (or disjoint sets of variables) X and Y are d -separated by the set of variables Z (i.e. conditionally

independent given Z), if and only if removing Z renders every undirected path (i.e. ignoring the direction of arrows) between X and Y inactive and does not create new active paths [23], see Fig. 5.3.




Figure 5.3: **The d -separation criterion.** In order to establish whether X and Y are d -separated by Z one needs to check the three primitives. In the case of a direct path which contains a (a) chain or (b) fork with the middle element in Z , removing this element breaks the path. (c) In the case of a collider (which is an inactive path to start with) the middle element should *not* be in Z (indicated by $\notin Z$) nor connected to an element in Z via a directed path, otherwise removing it creates an active path, since conditioning on a common effect leaves the causes correlated.

5.2.3 Causal Discovery

One of the main applications and central motivations of causal modeling is the idea of discovering causal structure underlying some empirical data. In practice, this task is very difficult, even in the classical domain, where a causal model always exists, since systems can be assigned objective properties [37–40], see Ref. [16] for an accessible introduction.

Causal discovery algorithms traditionally try to find a causal structure that is compatible with the set of conditional independences of a given empirical probability distribution. In general this approach can only hope to return equivalence classes of candidate structures, since e.g. a chain and fork are equivalent as far as the conditional independences are concerned. Additional information, such as interventional data, or space-time relations between events can help to distinguish between the different structures in a class. Another major challenge for causal discovery algorithms is that the empirical data may not contain all relevant variables, and unobserved common causes result in a non-Markovian model. To address such problems, and make the task computationally tractable causal discovery algorithms rely on two central principles [16]: faithfulness, and minimality. Since any correlation can be explained with enough latent variables, these principles are used to restrict the number of latent variables and causal links introduced by the algorithms. Minimality in particular is an appeal to Occam’s razor to prefer the most specific model. If one model can simulate another but not vice versa, then the

second is preferred.

Interestingly, the ambiguity in conditional independence relations together with the principle of minimality can lead traditional algorithms into situation where they return a causal structure that is compatible with the conditional independences, but not with the actual observed probability distribution. Such a case is in fact the Bell scenario, which satisfies exactly the same conditional independences as the EPR scenario. However, the case considered by Bell (recall, they differ in the choice of measurement direction) features stronger correlations than EPR, and cannot be explained causally. Furthermore conditional independences can never be satisfied exactly with finite empirical data, even if they hold for the underlying mechanism. Hence, it has been suggested that causal discovery algorithms should take into account the full joint probability distribution, which contains more information than the conditional independences alone [16, 41]. In Ref. [16] it was pointed out that causal discovery algorithms which satisfy minimality and faithfulness fail to produce a causal structure that can reproduce Bell-correlations. All causal models that can explain the observed correlations must therefore be fine-tuned.

5.3 Quantum Correlations: Bell’s Theorem and Beyond

In the scenario envisioned by Bell a source produces pairs of quantum systems, one of which is sent to Alice, and the other to Bob. Alice (Bob) then has the choice between two measurement settings, represented by the values of the random variable X (Y) and obtains one of two outcomes, represented by the variable A (B), see Fig. 5.4a. In addition to the measurement settings, the outcomes also depend on the properties of the particles. These are in part determined by known source parameters C and in part by unknown (or unknowable) “hidden” variables Λ [7]. For simplicity both can be described by a single variable Λ , which might have a complicated internal structure.

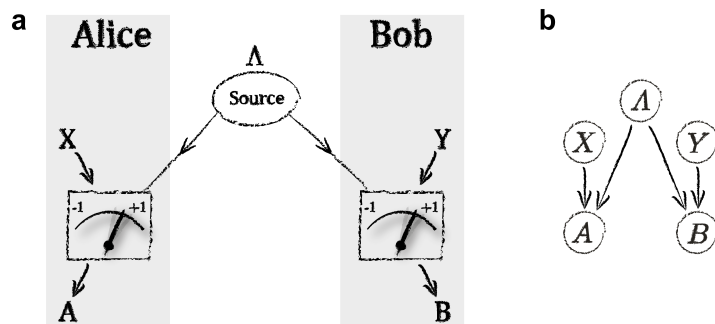


Figure 5.4: **Bell-local causal models.** Two parties, Alice and Bob, each perform one of two local measurements on one half of a shared quantum state and obtain one of two possible outcomes. Their measurement choices are represented by the variables X and Y , respectively and the measurement outcomes by A and B , respectively. **(a)** The relevant variables in a space-time diagram. The hidden variable Λ need not originate at the source of the particles. **(b)** The bare minimum: five random variables (circles) and their causal dependencies (arrows).

Abstracting the physics out of the situation, there are five relevant variables X, Y, A, B , and Λ . Following classical intuition, the measurement outcome on each particle should only depend on the properties of the local particle, and the chosen measurement setting. This is captured by the causal structure in Fig. 5.4b, and by the conditional joint probability distributions compatible with this structure, which are Bell's local hidden variable models:

$$p(a, b|x, y) = \sum_{\lambda} p(a|x, \lambda)p(b|y, \lambda)p(\lambda) \quad (5.3)$$

This set of conditional probability distributions forms a polytope. The facets of this polytope represent constraints of positivity of probabilities and the 8 symmetries of the Clauser-Horne-Shimony-Holt [42] (CHSH) Bell inequality under relabelling of the inputs, outputs and parties.

$$S_2 = \langle A_0 B_0 \rangle + \langle A_0 B_1 \rangle + \langle A_1 B_0 \rangle - \langle A_1 B_1 \rangle \leq 2, \quad (5.4)$$

where $\langle A_x B_y \rangle = \sum_{a,b=\pm 1} ab p(a, b|x, y)$ is the joint expectation value of A_x and B_y .

Fundamentally, Bell's theorem and the models Eq. (5.3) are based on the factorization of the conditional probability distribution $P(A, B|X, Y, \Lambda)$ into $P(A|X, \Lambda)P(B|Y, \Lambda)$. An experimental violation [5, 6] of a CHSH (or other Bell-type) inequality implies that at least one of the assumptions underlying this factorization is at variance with quantum correlations. However, there are a number of different sets of assumptions that lead to the above factorization, which caused some confusion about the interpretation of the theorem and experimental violations of Bell-inequalities.

5.3.1 Axiomatic Approach

Wiseman and Cavalcanti [7] proposed an axiomatic approach, which describes the various assumptions underlying Bell's theorem in terms of a set of fundamental axioms. This approach clarifies some of the persisting misunderstandings and conflicting interpretations of Bell's theorem. It is furthermore very interesting as a guideline for future efforts in the study of the merits and cost of various relaxations of these assumptions in candidate interpretations of quantum correlations. They show that Bell's theorem follows from four axioms (denoted A1-A4), which are considered uncontroversial, and four postulates (denoted P1-P4), which are more likely to be questioned. The four axioms are [7]

(A1) **Macroreality**

An event observed by any observer is a real single event and not relative to anything or anyone.

(A2) **Reasonable spacetime**

Experiments take place in a well-behaved spacetime, which is time-orientable, such as Minkowski space.

(A3) **Temporal order of events**

There is a notion of past and future with respect to every event (not necessarily the one of special relativity).

(A4) **Causal arrow**

Causes are in the past of their effects, excluding retrocausal interpretations of quantum mechanics.

Macroreality is a notion of realism at the observational level, which ensures that different observers give the same account of an experiment. This allows for some operational theories, but excludes, for example, QBism or interpretations which do not feature an absolute reality, such as relational interpretations. This notion is thus different from the notion of realism used in Chap. 4, which assumes that there is an observer-independent objective reality, rather than just operational predictions. Macroreality is a minimal assumption in both, the ontological models framework and the framework of classical causal modeling. The other three axioms are rather uncontroversial for earth-based experiments. However, interesting effects can occur in situations where they might fail, such as in the presence of closed timelike curves, see Appendix B, or for superpositions of massive objects, where no clear order of events exists [43]. The four postulates are [7]

(P1) **Free Choice**

There are events in the model which have no relevant causes.

(P2) **Relativistic causality**

Past and future are defined via past and future light cone. Causes cannot propagate faster than the speed of light

(P3) **Common cause principle**

If two variables are correlated but not direct causes of one another then there must be a common cause for both.

(P4) **Decorrelating Explanation**

A variable acts as a common cause for the correlation between two variables (which are not direct causes of one another), if conditioning on the variable removes the correlation.

Free choice is a crucial primitive in the causal modeling framework, used to define the concept of cause and effect via (freely chosen) interventions. Notably, this definition does not make any reference to the philosophically problematic concept of free will. Although giving up free choice would make it difficult to maintain a notion of causality, it is possible to relax this notion, allowing for superdeterministic models, where measurement settings, and measurement outcomes are both determined by a common cause. While this might seem “about as plausible, and appealing, as, belief in ubiquitous alien mind-control.” [7], the study of models with restricted free choice is nonetheless interesting from a conceptual and practical point of view [10, 12–14, 44].

In its abstract form, the causal modeling framework relies only on A1, P1, and Reichenbach’s principle (P3 and P4). Axioms A2 and A3 are, strictly speaking, not necessary since causal

models do not require a background spacetime and do not have an a priori notion of time. In the physical interpretation of the framework, however, physical time is typically associated with the direction of causation and axioms A2-A4 then provide a physical justification for the acyclicity of causal graphs (i.e. no causal loops).

Curiously, this approach highlights that Bell’s theorem follows from the causal modeling framework, (P1, P3, and P4), together with relativistic causality (P2). Hence, in light of experimental violations of Bell inequalities one either has to give up relativistic causality, or the classical way of thinking about cause and effect. Ref. [7] suggests that advocates of realism like Bell himself tend to give up relativistic causality, in order to maintain a notion of explaining correlations. On the other hand, it is argued that advocates of an operational interpretation of quantum mechanics are not particularly attached to the factorization part of Reichenbach’s principle

5.3.2 Bell’s Assumptions

Bell’s theorem and the factorization of Eq. (5.3) can be obtained from a number of different sets of assumptions about the specific causal structure. All of these in turn follow from subsets of the fundamental axioms and postulates discussed in Ref. [7]. In this section the most commonly used “higher-level” assumptions for deriving Bell’s theorem are introduced and discussed. These are typically formulated directly in terms of constraints on the conditional probability distribution and do not necessarily have a causal interpretation. Section 5.3.3 then presents a novel decomposition of the central assumption of local causality into causally meaningful sub-assumptions, and Sec. 5.3.5 discusses how all of these assumptions come together to imply Bell’s theorem.

Measurement Independence requires that the measurement settings X, Y are uncorrelated with the hidden variable Λ . In terms of conditional probability distributions this implies that $P(x, y|\lambda) = P(x, y)$ or equivalently $P(\lambda|x, y) = P(\lambda)$. This condition is often referred to as the “free choice assumption”, and motivated on the grounds of apparent experimental free will, which, however, is a problematic association [45]. It is important to distinguish between the fundamental axiom of free choice, which states that there exist variables without relevant causes, and measurement independence, which states that the measurement settings are such variables. Assuming that measurement settings are freely chosen variables further implies marginal independence of Alice’s and Bob’s settings. In practice, measurement independence is justified by spacelike separation of the measurement settings from the source, and appealing to relativistic causality, see Sec. 5.5.1.

Predetermination is the assumption that the measurement-outcome probabilities are deterministic functions of the measurement settings and the hidden variable:

$$P(a, b|x, y, \lambda) \in \{0, 1\} \quad \forall x, y, \lambda \quad (5.5)$$

Importantly, this does not imply that the measurement outcomes are predictable, since they are obtained from averaging over the unknown hidden variable. Bell used this assumption together with parameter independence in his 1964 derivation, see Ref. [7] for details. Predetermination also implies the strictly weaker condition outcome independence, which Bell subsequently used in the 1976 version of his theorem. Although these assumptions are related, only the latter can be derived from the causal assumptions directly [7]. Nevertheless, the two versions of Bell's theorem can be shown to be equivalent [46].

Predictability is the operational pendant to predetermination:

$$P(a, b|x, y, c) \in \{0, 1\} \quad \forall x, y, c \quad (5.6)$$

Predictability implies that the measurement outcomes are fully determined by the measurement settings X, Y and the source parameters C . This property holds, in principle, for any classical system, where sufficient knowledge of the starting configuration (which might be exceedingly complicated) allows one to calculate future configurations. It was shown that Bell's theorem together with signal locality implies that some quantum phenomena are fundamentally unpredictable [8, 47]. This is particularly interesting in a cryptographic context, as these phenomena can be used to generate perfect randomness.

Signal locality or no-signalling captures the idea that signals cannot travel faster than the speed of light, and follows directly from special relativity (or relativistic causality in the language above). In the Bell scenario it implies that the local measurement outcome should not carry information about the remote (space-like separated) measurement setting. This is captured by the statistical independence relations

$$P(a|x, y) = P(a|x) \quad \text{and} \quad P(b|x, y) = P(b|y), \quad (5.7)$$

or in the language of causal models by the conditional independences ($A \perp\!\!\!\perp Y|X$) and ($B \perp\!\!\!\perp X|Y$). Since signal locality is an empirically observed, operational concept that involves only observable probability distributions it can be treated as an observed independence, rather than an assumption.

Local causality is the assumption used in Bell's 1976 paper [3] and captures the idea that there should be no causal influence from one side of the experiment to the spacelike separated other side. Formally, it is given by the constraints

$$P(a|b, x, y, \lambda) = P(a|x, \lambda) \quad \text{and} \quad P(b|a, x, y, \lambda) = P(b|y, \lambda). \quad (5.8)$$

Local causality is much stronger than the operational principle of signal locality and directly assumes the factorization of the joint conditional probability distribution. In terms of conditional independence relations it reads ($A \perp\!\!\!\perp BY|X, \Lambda$) and ($B \perp\!\!\!\perp AX|Y, \Lambda$). Local causality follows

from the joint assumption of parameter independence (or locality) and outcome independence.

Parameter independence generalizes the statistical constraints of signal locality to include the hidden variable. It is defined as the joint assumption of the two statistical constraints:

$$P(a|x, y, \lambda) = P(a|x, \lambda) \quad \text{and} \quad P(b|x, y, \lambda) = P(b|y, \lambda). \quad (5.9)$$

This is the *locality* assumption that Bell used together with predetermination in his 1964 derivation of the theorem [2, 7]. The two statistical constraints in Eq. (5.9) individually correspond to the conditional independences $(A \perp\!\!\!\perp Y|X, \Lambda)$ and $(B \perp\!\!\!\perp X|Y, \Lambda)$, respectively. The first constraint corresponds to a causal model as in Fig. 5.4b with additional links $X \rightarrow B$ and $A \rightarrow B$, but no link $Y \rightarrow A$, see Fig. 5.5, and similarly for the second constraint. Importantly, parameter independence, which is the joint assumption of both statistical constraints, *does not* correspond to a causal model³. As a consequence, parameter independence *does not* imply causal independence (in the sense of causal modeling) of either outcome from the remote setting as was first noted in Ref. [48], see Fig. 5.5. A failure of parameter independence (as is the case in collapse models and Bohmian Mechanics) also does not necessarily imply signaling [8], since the dependence of an outcome on the remote setting can disappear when averaging over the hidden variable.

Outcome independence is one way to supplement parameter independence to reach local causality, and is given by the statistical constraints

$$P(a|b, x, y, \lambda) = P(a|x, y, \lambda) \quad \text{and} \quad P(b|a, x, y, \lambda) = P(b|x, y, \lambda). \quad (5.10)$$

This condition can be interpreted as precluding any direct causal influence from Alice’s measurement outcome to Bob’s or vice versa. In other words, the measurement outcomes are independent, given knowledge of everything else: $(A \perp\!\!\!\perp B|X, Y, \Lambda)$. This assumption is also known as *statistical completeness*, owing to the observation that if Bob’s outcome is statistically independent of Alice’s outcome, all correlations must be due to ignorance of Λ [45].

5.3.3 Causal Assumptions

Outcome independence and parameter independence together imply local causality, which has a clear interpretation as forbidding any causal link from either side of the experiment to the other. The individual sub-assumptions, however, do not have a causal interpretation. In the spirit of causal modeling we introduce a novel decomposition of local causality into *causal parameter independence* and *causal outcome independence*, which have a clear causal interpretation as excluding various causal links. Since there cannot be a bidirectional causal link between A and

³The sets of probability distributions compatible with both models is in the intersection of the sets that are compatible with either model. In the case of parameter independence these models are mutually exclusive, because one contains a link $A \rightarrow B$ and the other a link $B \rightarrow A$.

B , each of these assumptions correspond to two sets of statistical constraints, one for models with a potential link $A \rightarrow B$ and one for models with a link $B \rightarrow A$. Formulated in this way, the assumptions apply to any linear combination of the models (i.e. any convex combination of conditional probability distributions) with either causal link.

Causal parameter independence captures the idea that there should be no causal link from the measurement settings to the remote measurement outcomes. For models with a link $A \rightarrow B$, the corresponding statistical constraints are

$$P(a|x, y, \lambda) = P(a|x, \lambda) \quad \text{and} \quad P(b|a, x, y, \lambda) = P(b|a, y, \lambda), \quad (5.11)$$

The respective conditions for the model with the link $B \rightarrow A$ are obtained by swapping a with b , and x with y in Eq. (5.11). In contrast to ordinary parameter independence either constraint ensures that there is no direct causal link from either measurement setting to the remote outcome.

Causal outcome independence excludes the causal link from Alice's measurement outcome A to Bob's measurement outcome B . In terms of probability distributions, this corresponds to

$$P(b|a, x, y, \lambda) = P(b|x, y, \lambda). \quad (5.12)$$

Again, this condition is for the causal graph with a link $A \rightarrow B$ and the condition for the graph with this reversed causal link can be obtained by swapping a with b in Eq. (5.12). In this case, the statistical constraints starting from either model turn out to be equivalent and imply the factorization $P(a, b, x, y, \lambda) = P(a|x, y, \lambda)P(b|x, y, \lambda)$. blabla

5.3.4 No Fine-Tuning

As discussed above, one of the common assumptions (but not a central principle) of the causal modeling framework is causal faithfulness or no fine-tuning. This implies that all observed conditional independences should be a consequence of the causal structure, and not just of a particular, fine-tuned, choice of model parameters. In the Bell-scenario there are two (approximate) empirical conditional independence relations

$$\begin{aligned} (X \perp\!\!\!\perp Y) & \quad (\text{marginal independence of settings}) \\ (A \perp\!\!\!\perp Y | X), (B \perp\!\!\!\perp X | Y) & \quad (\text{signal locality}). \end{aligned} \quad (5.13)$$

Using a causal discovery approach one can now find all causal structures that respect these observed conditional independence relations. Under the mild assumption that measurement settings X, Y precede the measurement outcomes A, B it was shown in Ref. [16] that the only causal structure for the variables A, B, X, Y (with or without hidden variables), which faithfully represents Eq. (5.13) is that of Bell's local causal models in Fig. 5.4b. This implies that




Figure 5.5: **Comparison of various constraints on the causal structure of Bell’s theorem.** The causal links forbidden by the respective assumption are shown in dashed green. In the case of parameter independence only the causal model for the first constraint is shown; the joint assumption of both does not correspond to a causal model. The statistical constraints corresponding to causal outcome independence and causal parameter independence are asymmetric in a and b , and swapping them would result in a causal structure where the arrow between A and B is reversed.

even without appealing to physical notions such as local causality or measurement independence, Bell-local causal models are obtained as the only faithful representation of the observed conditional independences.

All other causal models (i.e. all explanations of Bell correlations that are in accordance with Reichenbach’s principle and free choice) must then be fine-tuned. This fine-tuning can take various forms, see Ref. [16] for a physical, and Ref. [34] for a philosophical discussion. One such mechanism is based on *canceling paths*, where two variables are causally connected via multiple

paths that balance each other exactly such that the variables appear uncorrelated. This form of fine-tuning appears, for example in models which abandon local causality (cf. Fig. 5.7c), where it ensures that signal locality is satisfied. Another possibility is *marking*, where a causal link is hidden by another causal influence from an independent variable. In Fig. 5.7b, for example, the direct causal link from A to B can be masked by the link from Λ to B via a specific structure of the hidden variable [16].

Fine-tuned explanations are sometimes considered unnatural, since they are very sensitive to disturbances of the causal parameters. However, fine-tuned causal parameters can arise quite naturally through processes such as equilibration [16, 49]. Bohmian mechanics is a prominent example, where signal locality and the Born rule are fine-tuned features of the equilibrium state. Furthermore, Ref. [34] proposed a fine-tuning mechanism that is similar to cancelling paths, but acts on the level of values, rather than variables, which is interesting for spontaneous collapse interpretations and retro-causal models [36]. This *internal cancelling paths* mechanism effectively introduces a probabilistic dependence between three variables in a chain, such that there are multiple paths to arrive at each value of the final variable, resulting in statistical independence of the first and last variable. If the probabilities are chosen right (according to the born rule) then this reproduces quantum predictions. They argue that disturbances (specifically, unitary operations) might disturb the causal structure at the variable level, but not at the value level. Thus this kind of fine-tuning would not be detectable.

5.3.5 Many Roads to Bell's Theorem

Figure 5.6 brings some order to all the causal assumptions introduced in the preceding section, by depicting how they can be combined to derive Bell's theorem.

The derivation of Bell's theorem from signal locality and predictability follows that correlations which violate a Bell inequality, while satisfying signal locality must contain fundamentally unpredictable events [8]. Bell's 1964 theorem is based on the assumption of parameter independence (or locality) together with predetermination. In contrast to before, neither of these assumptions are operational. Hence they cannot be tested independently, and a Bell-inequality violation does not contain any information about their independent validity [8]. Bell's 1976 theorem is based on local causality, which can be decomposed into either the statistical constraints of parameter independence and outcome independence, or into the causal constraints of causal parameter independence and causal outcome independence. Either combination leads to Bell's theorem, but only the latter has a causal interpretation. Finally, the structure of Bell's local causal models can be obtained as the only causal structure that can faithfully reproduce the observed conditional independences [16].

5.3.6 Relaxing the Assumptions

“Bell's theorem, in all its forms, tells us not what quantum mechanics is, but what quantum mechanics is not” - Zukowski, Brukner (2014) [9]

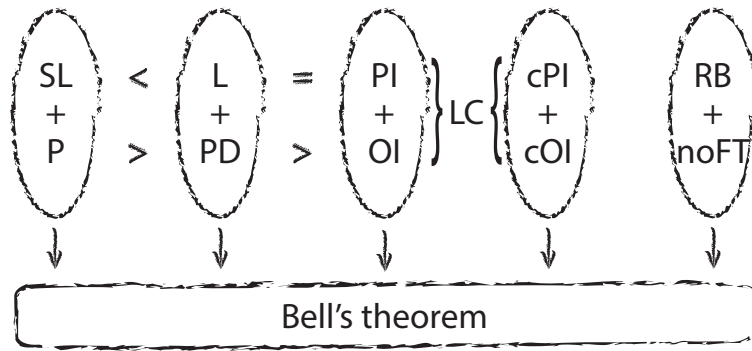


Figure 5.6: **Many roads to Bell's theorem.** Deriving Bell's theorem from empirically verified signal-locality (SL) and the strong assumption of predictability (P) implies that some quantum phenomena are fundamentally unpredictable [8, 47]. Bell's 1964 derivation of his theorem is based on locality (L) and predetermination (PD) [2], while the 1976 version is based on local causality (LC) [3], which can be seen as the conjunction of parameter independence (PI) (another name for locality) and outcome independence (OI). Alternatively, local causality can be obtained from causal parameter independence (cPI) and causal outcome independence (cOI) [1]. All these derivations also assume measurement independence (MI). Ref. [16] shows that Bell's local causal models also follow from Reichenbach's principle (RB), that is prepositions 3 and 4, as the only causal model that can satisfy signal locality and marginal independence of settings without violating the no fine-tuning assumption (noFT). Symbols <, >, = indicate that the assumption to the right is stronger, weaker or equal to the assumption to the left of the symbol.

A violation of a CHSH inequality reveals a conflict between experimental observations and the predictions of Bell-local causal models. Although this is often attributed to some kind of nonlocal causal influence, which coined the term “quantum nonlocality”, a failure of any of the assumptions might allow for a causal explanation of the observed correlations. Consequently, there is a significant body of literature examining to what degree the various assumptions have to be relaxed in order to recover a causal model for quantum correlations [10–17, 45]. The causal modeling framework offers a unique platform for a unified treatment and an intuitive interpretation of these relaxations [15, 16]. A large number of proposed causal models can be classified by studying the corresponding causal structure with respect to Bell's original local causal structure.

Local Causality

Local causality can be relaxed by either relaxing causal outcome independence (Fig. 5.7a), causal parameter independence (Fig. 5.7b), or both (Fig. 5.7c). Since Alice and Bob are assumed to be spacelike separated from each other, models that violate local causality involve *superluminal* causal influences⁴. Such models can nonetheless be in accordance with signal lo-

⁴Relativity merely implies that there can be no superluminal information transfer (i.e. signal locality), but not that causal influences cannot travel faster than the speed of light. For example, both phase- and group-velocity of a wavepacket can be superluminal, but the signal velocity (identified with the velocity of an edge) is always bounded by the speed of light in vacuum [50]. In other words, causal influences can conceivably travel

cality by virtue of fine-tuned causal parameters. A violation of causal parameter independence implies that one or both measurement settings causally influence the measurement outcome on the respective other side of the experiment. For a violation of causal outcome independence the most general case is a convex combination of models that contain either the link $A \rightarrow B$ or $B \rightarrow A$. A model with a simultaneous link from A to B and B to A cannot be given a causal interpretation, since it involves a causal loop, such that A could cause itself.




Figure 5.7: **Relaxations of local causality.** Local causality can be violated by (a) a violation of causal parameter independence, (b) a violation of causal outcome independence, or (c) a violation of both.

Models featuring nonlocal causal influences are among the most popular causal explanations of quantum correlations. An early example of how a violation of causal parameter independence can be exploited to simulate singlet-state correlations was introduced by Toner and Bacon in 2003 [51]. In their protocol Alice and Bob share two random Bloch-vectors for each run of the experiment (the hidden variable). Alice uses these two vectors to partition the Bloch-sphere into two regions and sends the location of her measurement choice with respect to these regions (one bit of information) to Bob. Bob similarly partitions the Bloch-sphere into two different regions and computes his output based on the relative orientation of his and Alice’s measurement, see Ref. [51] for details. Curiously, the hidden variable in this model must contain an infinite amount of information, which is a necessary feature of ontological models of quantum mechanics, known as the ontological excess baggage theorem, see Chap. 4 for details. They demonstrate that a 1-bit message is sufficient to simulate all singlet state correlations, which under appropriate circumstances can be compressed to ~ 0.85 bits. More generally, the minimum entropy of a binary message in this scenario is directly proportional to the binary entropy of the CHSH violation (i.e. $S - 2$) and thus at most ~ 0.736 [15].

Measurement Independence

Measurement independence assumes that there is no correlation between the measurement settings and the hidden variable. This assumption could be violated by a causal influence from the hidden variable to the measurement setting, see Fig. 5.8a. Such models are sometimes referred to as *super-deterministic*, since they seem to imply that measurement settings cannot be chosen freely. Appealing to experimental free will to justify measurement independence, however, is questionable, not the least because the concept of free will itself is problematic [52] and not well-defined [45].

faster than light, as long as they do not permit superluminal signaling.

On the other hand the hidden variable might be causally influenced by the measurement settings. Since the latter are typically chosen after the particles leave the source, such an influence would in general be to the past, one speaks of a *retro-causal* model, see Fig. 5.8b. Such models, however, are fine-tuned to satisfy signal locality, and the causal influences may even propagate at subluminal speeds within the past lightcone. However, not every model, which allows for causal influences to the past can be treated within the causal modeling framework. In general such models may lead to causal loops, such as feedback from the measurement outcomes to the hidden variable. Examples of models involving such causal loops Cramer’s transactional interpretation [53], the two-state vector formalism [54] or Huw Price’s interpretation [55]. Such models are not captured by the causal modeling framework and it is difficult to even establish a clear notion of causation in such theories.




Figure 5.8: **Relaxations of measurement independence.** **(a)** In a superdeterministic model the variable Λ may have a direct causal influence on Alice’s or Bob’s measurement setting. **(b)** In a retrocausal model the measurement settings may have a direct causal influence on the hidden variable.

Explanations of quantum correlations via relaxed measurement independence has continuously attracted attention [10, 12–15, 44]. As discussed above, this does not necessarily mean superdeterministic models where everything is predetermined, but rather models with some degree of correlation between the measurement settings and the hidden variable. This correlation can be measured in terms of the mutual information $I(X, Y: \Lambda)$, which quantifies how much information can be gained about the measurement settings from learning Λ or vice versa. In Ref. [13] it was shown that a mutual information of ~ 0.0463 is sufficient to simulate CHSH correlations, and ~ 0.0663 is enough for any set of projective measurement on a singlet state. These numbers were confirmed to be tight in the case of bipartite correlations with binary inputs and outputs, and have been generalized in various directions [15].

On the other hand, Ref. [18] showed that quantum correlations cannot be reproduced even with arbitrary (though not complete) relaxations of measurement independence. The apparent contradiction with Ref. [13] is resolved in the details. Specifically, in Ref. [18] a relaxation of measurement independence is defined as a deviation from the case where all conditional probabilities $P(x, y|\lambda)$ are balanced. Their result crucially relies on Hardy’s paradox and the assumption that none of the probabilities $P(x, y|\lambda)$ are zero for models with arbitrary relaxation of measurement independence. Ref. [13, 15] on the other hand follow a more general approach that does not make any assumptions on the underlying distribution of the inputs, defining measurement independence via the variational distance between the distributions $P(x, y)$ and

$P(x, y|\lambda)$.

On Common Grounds

The causal modeling approach provides a common framework for many of the discussed alternative explanations of quantum correlations and reveals their close connection, see Ref. [16] for a detailed discussion. Although all the models discussed above are obtained by adding one or more causal arrows to the causal structure of the Bell-local models, they can have very different physical interpretations. For example, a model with an additional hidden variable that correlates X and B could be super-deterministic, retro-causal or superluminal, depending on the direction of the causal influences and the position of this additional variable in space-time [16], see Fig. 5.9. From the point of view of causal modeling, these three models are identical. Causal models are formulated without any reference to a spacetime structure, and are thus completely defined by the structure of causal dependences and the causal-statistical parameters.




Figure 5.9: **Relation between various explanations for Bell correlations.** Consider a model where an additional hidden variable μ causally connects Alice’s setting choice X and Bob’s measurement outcome B . **a** If μ is placed in the common past of X and B the model is superdeterministic (or retro-causal if the direction of the arrow between X and μ is reversed). **b** If μ is placed in the future of X , but in the past of B then it mediates a superluminal influence from X to B , thus violating local causality. **c**) If μ is placed in the common future of X and B the model is retro-causal but not necessarily nonlocal. (Figure adapted from Ref. [16])

Another common feature of all models discussed above, which can reproduce Bell correlations is that they are unfaithful representations of the empirical conditional independence relations. In order to satisfy the empirical constraints of marginal independence of the measurement settings and signal-locality, the causal parameters of these models must be fine-tuned [16]. Indeed, the Bell-local causal structure is the only faithful representation of these constraints, but it cannot reproduce the strength of the observed correlations.

5.4 Testing Causal Models for Quantum Correlations

Experimental tests of causal models for quantum correlations go back to the early tests of the CHSH inequality by Freedman and Clauser [56] in 1972. These experiments demonstrated that Bell-local causal models cannot reproduce quantum predictions, but they had to rely

on additional assumptions to contend with experimental imperfections and a lack of space-like separation. Since then, much experimental effort has been devoted to overcoming the requirement for such assumptions and Bell-local causal models have now been ruled out in a conclusive, loophole-free fashion [4–6]. This opens the path to study in more detail which of the assumptions behind Bell-local models must give way to recover a causal picture of quantum correlations.

A violation of the CHSH inequality is often attributed to a nonlocal causal influence, or “quantum nonlocality”, and in particular to a failure of outcome independence⁵, see e.g. Ref. [18]. As discussed in Sec. 5.6 the conventional statistical notions of parameter independence and outcome independence do not have a causal interpretation. Instead, one can decompose local causality into the causally meaningful assumptions of causal parameter independence and causal outcome independence. The class of models which satisfy the former, but may violate the latter are a direct generalization of Bell-local causal models. This class contains all models where the measurement outcome A may influence the outcome B (see Fig. 5.7a), or vice versa, or any linear combination of these. The set of conditional probability distributions compatible with a linear combination of two causal models is the convex combination of the conditional probability distributions compatible with either model. Hence, the conditional probability distributions compatible with the models studied here are of the form

$$p(a, b|x, y) = \sum_{\lambda} p(a|x, \lambda)p(b|a, y, \lambda)p(\lambda) + \sum_{\mu} p(a|b, x, \mu)p(b|y, \mu)p(\mu), \quad (5.14)$$

with $\sum_{\lambda} p(\lambda) + \sum_{\mu} p(\mu) = 1$.

One approach to testing these models, is to directly measure the strength of the causal link between A and B in Fig. 5.7a. This can be done using controlled interventions, where the variable A is changed independent of the other variables in the model, and measuring the associated change in the distribution of B . This approach can be used in any scenario, but is device-dependent and requires detailed knowledge of the underlying causal mechanism, including the hidden variables. Another approach, much in the spirit of quantum information, is to study the polytope formed by the conditional joint probability distributions of Eq. (5.14). The facets of this polytope are Bell-type inequalities, which allow for a device-independent test of the models. If any of these inequalities is violated by quantum systems, then the model cannot reproduce all quantum predictions.

5.4.1 Experimental Setup for Testing the CHSH Inequality

To test bipartite nonlocal causal models, pairs of photons are created in the state $\cos \gamma|HV\rangle + \sin \gamma|VH\rangle$, where the parameter γ (the polarization angle of the pump beam) continuously controls the degree of entanglement as measured by the concurrence $\mathcal{C} = |\sin 2\gamma|$, see Sec. 3.2. Alice and Bob can each perform arbitrary local measurements on one of these photons using a

⁵One motivation for this might be that parameter independence, when suitably generalized holds in orthodox quantum mechanics [7].

combination of HWP and QWP, followed by a polarizer, see Fig. 5.10. In the case of the CHSH inequality it is in fact sufficient to choose all measurement directions on a great circle of the Bloch-sphere, which can be implemented using only a HWP [1].




Figure 5.10: **The experimental setup.** (a) Photons pairs are generated via spontaneous parametric downconversion in a periodically poled KTP (ppKTP) crystal, using the Sagnac design of Ref. [57]. Alice and Bob perform local measurement on their photons in the equatorial plane of the Bloch sphere using a half-waveplate (HWP) and a polarizing beam splitter (PBS). Additional quarter-wave plates (QWP) can be used for quantum state tomography of the initial entangled state. In the interventional experiment an additional combination of QWP and polarizer (POL) are used between Alice’s basis choice and her measurement. The causal variables are indicated using the notation of Fig. b, and the hidden variable Λ need only be in the causal past of the measurement outcomes, but not necessarily originate at the source. (b) Sketch of the effect of the intervention in the casual model of Fig. 5.7a.

The measurement settings X , and Y are generated using the online quantum random number generator of the Australian National University, based on Ref. [58]. In contrast to the classical counterpart, which typically produces random numbers from some underlying deterministic algorithm, a quantum random number generator exploits the inherent unpredictability of quantum mechanics and are thus capable, in principle, to produce truly unpredictable random numbers [59]. This can be relevant for security purposes and to ensure that the measurement settings are statistically independent from the hidden variable Λ (except in a superdeterministic model). In the present experiment, the use of a quantum random number generator is not strictly necessary, as we only require that the measurement settings satisfy marginal independence, that is $P(X, Y) = P(X)P(Y)$, which can be quantified using the maximal variation between the joint and product distribution:

$$V_{\text{MI}} = \sup_{x,y} |P(x, y) - P(x)P(y)|. \quad (5.15)$$

Besides marginal independence, the conditional probabilities also have to satisfy signal locality, which is a consequence of special relativity. In practice, however, no finite dataset can perfectly satisfy this condition, and residual signaling has to be taken into account when

interpreting the results of a Bell-test experiment. Again, this can be quantified by the maximal shift in the conditional probability of A given X due to a change of Y and vice versa for Bob⁶

$$\begin{aligned} V_{\text{SL}}^A &= \sup_{a,x,y,y'} |P(a|x,y) - P(a|x,y')| \\ V_{\text{SL}}^B &= \sup_{b,y,x,x'} |P(b|x,y) - P(b|x',y)|. \end{aligned} \tag{5.16}$$

5.4.2 Interventional Approach

In the CHSH-scenario, passive observations alone are not enough to determine whether correlations between A and B are due to direct causation or a common cause Λ . This is because a message sent from A to B could in principle contain information about the settings, rather than the outcomes, which is known to allow for simulation of the correlations [48]. An intervention on the variable A , however, would break all incoming causal influences, and the remaining correlation between A and B must stem from direct causation. The maximal shift of the probability distribution of B upon intervention on A can be used to quantify the strength of this causal link, using the so-called *average causal effect* [15, 29, 60]

$$\text{ACE}_{A \rightarrow B} = \sup_{b,y,a,a'} |p(b|\text{do}(a),y) - p(b|\text{do}(a'),y)|. \tag{5.17}$$

Here $\text{do}(a)$ indicates that the variable A is set to the value a by means of an intervention, see Sec. 5.2.1. In contrast to the measure used in Ref. [15], the average causal effect does not require knowledge of the hidden variable⁷ and is thus experimentally accessible, and satisfies [1]

$$\min \text{ACE}_{A \rightarrow B} = \max [0, (S_{\text{CHSH}} - 2)/2]. \tag{5.18}$$

Here the maximum is taken over the eight symmetries of the CHSH quantity, and $\min \text{ACE}_{A \rightarrow B}$ refers to the minimal strength of the causal link $A \rightarrow B$ required for a causal explanation of a given set of quantum correlations, which is directly proportional to S_{CHSH} .

Experimentally, a valid intervention must modify the causal mechanism in a way that sets the value of the variable A independent of all other variables⁸, thus breaking all incoming causal arrows on A . Modifying a causal structure in this way necessarily requires some knowledge of the underlying causal mechanisms. Consider the textbook example of interventions in a medical setting, trying to establish a causal connection between hypertension and coronary heart disease. An intervention with blood-pressure lowering medication requires knowledge of the physiology of hypertension (i.e. the variable that is intervened upon). Moreover it has to be assumed that the specific way of dividing patients into treatment and control groups is uncor-

⁶Different tests could be designed using the equivalent conditions $P(A|X,Y) = P(A|X)$ or $P(A,Y|X) = P(A|X)P(Y|X)$, but these are only reliable in the case where marginal independence of the settings holds exactly.

⁷The term ‘‘average’’ is used because the marginal distributions over the observed variables are defined as an average over the unobserved ones.

⁸Strictly speaking, any modification of the ‘‘natural’’ causal mechanism is an intervention, even if it only partially breaks the causal links [32, 61], but for the present purpose the strong definition is most appropriate.

related from genetic factors or any other relevant causes. Similarly in other cases, interventions always require detailed knowledge about potential confounding factors, and cannot exist in a theoretical vacuum.

In the present case A corresponds to the outcome of a polarization measurement in the linear-polarization plane, which is taken to instantiate at Alice’s PBS. The requirements of an intervention on A can be met under the assumption that the local degrees of freedom behave according to quantum mechanics. Such assumptions are common in quantum steering scenarios and semi-device independent quantum cryptography, where it is assumed that the devices of at least one of the labs can be trusted and work according to quantum mechanics. However, this assumption does restrict the class of causal models that can be tested with an interventional approach. Randomly projecting Alice’s photon onto circular polarization states—which are mutually unbiased with respect to to all CHSH measurements—then erases all relevant information, and re-preparation in the eigenstates of the PBS forces one of the two outcomes $A = \pm 1$. This corresponds to the operators $\{|H\rangle\langle R/L|, |V\rangle\langle R/L|\}$, which are implemented experimentally using a quarter-waveplate and a polarizer.

This kind of intervention relies on the symmetry of the problem, where all information is encoded in linear polarization states. If this was not the case, an alternative approach could be to replace Alice’s photon after the setting choice with a new photon. This is technically more challenging, since it requires careful synchronization of the new photon with the incoming photon of Alice and might potentially open additional loopholes. The simpler implementation of preserving the photon but erasing the relevant polarization information is much less invasive, but involves an unbiased loss of 50% due to the use of a polarizer. This could potentially be overcome by using a polarizing beamsplitter instead and combining both beams at Alice’s detectors. However, to avoid any bias the settings for the intervention elements are chosen randomly using quantum random numbers from the Australian National University’s online quantum random number generator based on Ref.[58]. This loss can thus be absorbed into the fair-sampling assumption.

A crucial aspect of the experiment is to establish multi-point correlations between all involved variables, not just statistical correlations between the outcomes. Each individual click event was registered and correlated using an AIT-TTM8000 time-tagging module with a temporal resolution of 82 ps. Outcome probabilities, used to estimate $\text{ACE}_{A \rightarrow B}$, were then computed from a total of 48,000 coincidence counts, with no more than one event for each set of random choices for X, Y and the two elements of I .

Experimental Results

Using this interventional approach the average causal effect can be bounded to below $\text{ACE}_{A \rightarrow B} = 0.02^{+0.02}_{-0.02}$ for six different partially entangled states, and is largely independent of the observed CHSH value see Fig. 5.11. As shown in Fig. 5.11 a slight increase of the observed values of $\text{ACE}_{A \rightarrow B}$ with the entanglement of the used state is observed. This is in accordance with the expectation that states with higher entanglement are more susceptible to Poissonian noise due

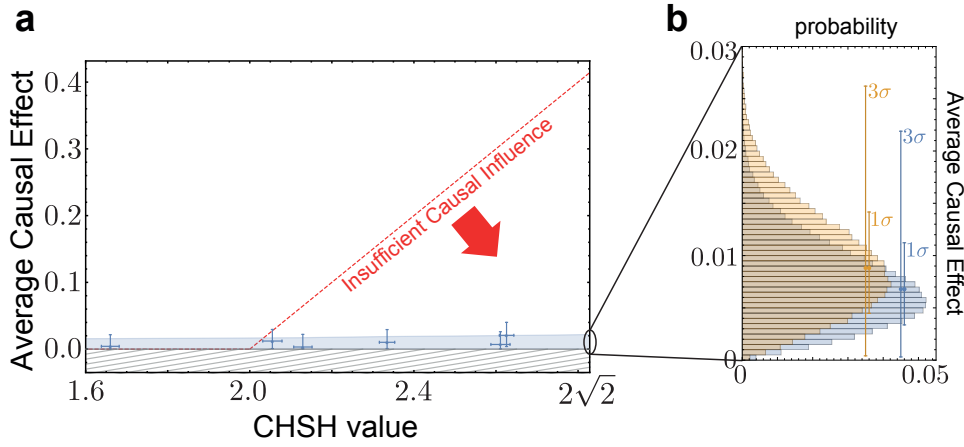


Figure 5.11: **Observed average causal effect $ACE_{A \rightarrow B}$ versus measured CHSH-value.** (a) Any value below the dashed red line, given by Eq. (5.18), is not sufficient to explain the observed CHSH-violation. All errors represent 3σ statistical confidence intervals obtained from a Monte-Carlo simulation of the Poissonian counting statistics. The blue shaded area represent the 3σ region of Poissonian noise compatible with $ACE_{A \rightarrow B} = 0$. Since the quantity $ACE_{A \rightarrow B}$ is bounded from below by 0 (the hatched area), the distribution of statistical errors is expected to be asymmetric. (b) Distribution of expected values for a nominal $ACE_{A \rightarrow B} = 0$ due to statistical noise (blue) or statistical noise and systematic imperfections (orange). 1σ and 3σ statistical confidence intervals are indicated.

to a lower average count-rate per measurement setting. Moreover, since the quantity $ACE_{A \rightarrow B}$ is bounded from below by 0, every systematic imperfection and statistical noise tend to push the observed values away from this boundary. Monte-Carlo simulations of the Poissonian noise alone suggest expected values between $ACE_{A \rightarrow B} = 0.005^{+0.011}_{-0.005}$ for the least entangled and $ACE_{A \rightarrow B} = 0.007^{+0.015}_{-0.006}$ for the most entangled state in the experiment. These median and variance of these values increase slightly when additionally taking into account known systematic imperfections in the used waveplates.

As shown in Fig.5.11, all observed values are indeed compatible with $ACE_{A \rightarrow B} = 0$ at the 3σ -level of statistical noise (blue shaded region). Within current experimental capabilities this demonstrate that CHSH violations above a value of $S_{\text{CHSH}} = 2.05 \pm 0.02$ cannot be fully explained by means of a direct causal influence from one outcome to the other. In other words, the potential causal influence between Alice's and Bob's measurement—the green arrow in Fig. 5.7a—is not sufficiently strong.

Marginal independence of the measurement settings was tested using Eq. (5.15), obtaining a median of $V_{\text{MI}} \sim 0.0014$ and sample variance of $\sim 10^{-6}$ across the six datapoints in Fig. 5.11. Signal locality for Alice (Bob) was tested using Eq. (5.16) and found to be satisfied to within a median variation of $V_{\text{SLA}} \sim 0.0047$ ($V_{\text{SLB}} \sim 0.0092$) with sample variance of $\sim 7 \times 10^{-6}$ ($\sim 7 \times 10^{-5}$).

5.4.3 Observational Approach

Interventions are an interesting conceptual tool and central to the study of causal models. In principle they can be used on any causal model (not on any variable of the model though) to quantify the strength of causal links. However, since interventions aim to modify the causal mechanism which determines the value of a variable (here A) they require detailed knowledge about the physical representation of this variable and are thus necessarily device-dependent.

Instead of trying to quantify a potential causal link, however, one could instead study the set of probability distributions that are compatible with the model. These form a polytope, which is bounded by Bell-type inequalities, such as the CHSH inequalities for the model in Fig. 5.4b. In the case of a causal influence either from A to B or from B to A , the corresponding probability distributions are given by Eq. (5.14). In the CHSH case with two inputs and two outputs for each party, this set covers all quantum distributions and this method can make no nontrivial statement about such models. Curiously, however, the situation changes when considering a situation with more settings than outcomes. This suggests communication of the measurement settings is more powerful than communication of the measurement outcomes, in accordance with a similar result of Ref. [48]. It would be interesting to study this behaviour in more detail. Specifically in the case of three settings and two outcomes, one of the facets of the polytope is given by the inequality [15]

$$S_3 = \langle E_{00} \rangle - \langle E_{02} \rangle - \langle E_{11} \rangle + \langle E_{12} \rangle - \langle E_{20} \rangle + \langle E_{21} \rangle \leq 4. \quad (5.19)$$

This inequality is symmetric in A and B , and thus satisfied by any model with a causal influence from either party to the other (i.e. any convex combination of the model in Fig. 5.7a and the model with the $A \rightarrow B$ arrow reversed), but can be violated by entangled quantum states. Testing this inequality thus amounts to testing the model in Fig. 5.7a in a device-independent fashion and without committing to any particular temporal ordering of A and B . Curiously, Eq. (5.19) seems to be related to a chained Bell inequality of Ref. [62] under relabelling of the settings. However, such relabellings are not symmetry transformations in a model where B may statistically depend on A , and there is no logical relationship between the inequalities.

To test inequality (5.19) it is sufficient to consider measurements in the equatorial plane of the Bloch-sphere, see Fig. 5.12a. These can be implemented using the setup in Fig. 5.10 with the intervention elements I removed. The maximal violation of $S_3 = 3\sqrt{3}$ is achieved for a maximally entangled state (corresponding to $\gamma = 45^\circ$) and measurement angles with respect to $|H\rangle$ of $\{-\pi/6, 7\pi/6, \pi/2\}$ for Alice and $\{-\pi/3, \pi/3, \pi\}$ for Bob. As in the case of the CHSH inequality, it is possible to optimize these measurement settings for the specific state used in the experiment. In this case a violation is theoretically possible for every non-separable quantum state, see Fig. 5.12b. Experimentally, a maximal value of $S_3 = 5.16^{+0.02}_{-0.02}$ was achieved, correspond to a violation of Eq. (5.19) by more than 170 standard deviations.

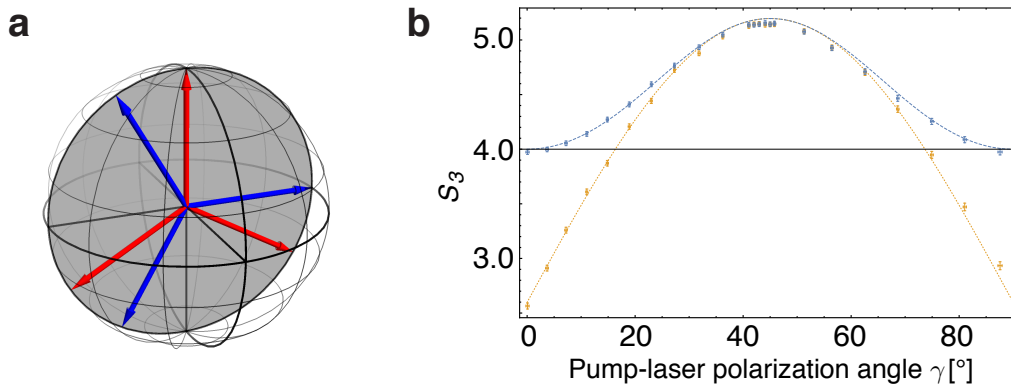


Figure 5.12: **Experimental test of inequality** (5.19). **(a)** The measurement settings for Alice (red) and Bob (blue) achieving maximal violation of the inequality can be chosen in the equatorial plane of the Bloch sphere. **(b)** Observed values S_3 for a variety of quantum states of the form $\cos(\gamma)|HV\rangle + \sin(\gamma)|VH\rangle$. Data from a fixed measurement scheme (optimal for the maximally entangled state, $\gamma = 45^\circ$) are shown in orange, with the dotted, orange line representing the corresponding theory prediction. The blue data and blue dashed theory line corresponds to the case where measurement settings were optimized for the prepared states. The black line represents the bound of inequality (5.19); any point above this line cannot be explained causally by a model of the form in Fig. 5.7a. Error-bars correspond to 3σ statistical confidence intervals.

5.5 Discussion and Outlook

The local causal models ruled out by Bell’s theorem have a natural representation within the causal modeling framework by the causal structure of Fig. 5.4b. Most relaxations of Bell’s assumptions can similarly be given a causal interpretation, which corresponds to structures with additional causal arrows. Such model directly generalize the original structure and contain Bell’s local causal models as a special case. This makes the causal modeling framework an excellent, and intuitive platform for studying such relaxations. In contrast to Bell’s original argument, causal models are formulated without any reference to a spacetime structure. Curiously, however, the assumptions of classical causality alone are sufficient to derive Bell’s local causal model as the only faithful representation of the experimental observations in a CHSH experiment [16].

When it comes to testing candidate causal models, interventions offer a direct way of quantifying the strength of potential causal links, that is very close to the spirit of causal modeling. However, this approach is necessarily device- and theory-dependent, which limits its applicability to testing ontological models for quantum correlations with a specific form of the hidden variables. Moreover, not every variable can be intervened upon, and models such as those that relax measurement independence [17, 18, 63], or feature objective wavefunction collapse [34] require different methods. On the other hand, studying the set of conditional joint probability distributions compatible with a given causal models produces device-independent inequalities. A violation of such an inequality constitutes a general test of arbitrary hidden variable models

that respect the corresponding causal structure.

Curiously, the CHSH scenario with two parties, two inputs and two outcomes, does not lead to nontrivial device-independent inequalities for a causal structure with additional arrows with respect to the Bell-local one. However, once the number of possible measurement settings is larger than the number of outcomes, a model that only allows communication of measurement outcomes cannot reproduce quantum predictions any more. This suggests that communication of the measurement settings is more powerful than communication of the measurement outcomes, which is in accordance with a similar result of Ref. [48]. It would be interesting to study this behaviour in more detail and for more complicated scenarios.

The results in Sec. 5.4 rule out any causal model which explains CHSH correlations in terms of a causal link from either measurement outcome to the other. In other words, a violation of causal outcome independence is insufficient to recover a causal explanation of quantum correlations. This adds another piece to the puzzle of ontological models for quantum correlations, see Chap. 4. It is now of great interest to extend these methods to other classes of causal models for CHSH correlations [12–18], and beyond.

5.5.1 Loopholes

Early experimental tests of the CHSH inequality had to rely on auxiliary assumptions to contend with experimental imperfections and a lack of spacelike separation. This allows for loopholes, where a violation of the inequality could be attributed to a failure of one or more of the auxiliary assumptions, rather than a failure of Bell-locality.

Fair Sampling

Most photonic experiments, for example, suffer from low detection efficiency and thus rely on *fair-sampling*. A model is fair-sampled, if the loss is unbiased at the hidden variable level. On the other hand, unfair-sampled causal models might be such that events which lead to a violation of the tested inequality are more likely to be detected than those that don't. Such an argument can, however, only be made up to a certain threshold efficiency. In the case of the standard CHSH inequality this threshold is $\sim 83\%$ [64]. However, by settling for a smaller possible violation of the inequality, efficiencies as low as $2/3$ can be sufficient to rule out Bell-locality [65]. In the present case the required efficiency for an experiment without the fair-sampling assumptions is $\sim 88\%$, but it is possible that a similar argument as before can be used to lower this bound [1].

Locality

In standard CHSH tests, Bell's assumption of local causality is typically justified by ensuring spacelike separation of the two sides of the experiment and appealing to special relativity. Using fast electronics this requirement can be met with photons using a separation on the order of merely 100m [5, 6]. Importantly, however, this requirement is used to *justify* why

one should adopt a Bell-local model. The causal model itself, does not contain any notion of spacelike separation or even a notion of physical time and speed of causal influences. In particular when testing causal models that relax local causality, and are thus explicitly nonlocal, spacelike separation cannot be used to justify any causal independence. The causal and statistical independences are implied by the model, independent of any spacetime narrative. If the inequalities implied by the model are violated, then the model cannot explain the experimental observations.

Measurement Independence

As with local causality, the measurement independence assumption, that the choices of measurement settings are statistically independent of the hidden variable, is typically justified by spacelike separation of the setting choices from the source [6]. Curiously, this assumes that the hidden variable originates at, or close to the source, which need not necessarily be the case. As with the locality loophole, this step is merely a justification by appealing to special relativity, rather than a strict requirement. A violation of the CHSH inequality shows that Bell’s local causal model cannot reproduce the experimental observations. If the measurement settings are not spacelike separated, then a possible explanation is subluminal influences from the measurement settings to the hidden variable. If they are spacelike separated, then such an explanation would either require superluminal influences, or would require the hidden variable to originate further in the past. To push the latter case some billions of years into the past, it has been suggested to use light from distant quasars to determine measurement settings [63].

5.5.2 Relation to Previous Work

Leggett models

Previous work on models beyond Bell-locality focused mainly on Leggett’s *crypto nonlocality* [66–68], which refers to a class of nonlocal hidden variable models where the hidden variables are product quantum states. In other words, in each run of the experiment the subsystems are assumed to behave as if they were in definite pure quantum states, and thus satisfy local quantum mechanics, while there are no restrictions on the correlations [11, 69]. Assuming additionally measurement independence, signal locality and non-negativity of probabilities, such a model nonetheless imposes constraints on the correlations of two quantum systems. The constraints from crypto-locality, however, are logically unrelated to Bell-locality and either may be satisfied or violated independent of the other [70]. Rather than focusing on nonlocality, Leggett’s model is “an attempt at keeping the correlations and reintroducing sharp properties at the individual level as well” [11]. A violation of generalized Leggett inequalities implies that the constituents of a maximally-entangled singlet-state cannot have even partially defined individual properties, in the sense of local mixed states, rather than pure states.

Finite Speed Influences

Other experiments have studied nonlocal causal models by establishing lower bounds on the speed of potential superluminal causal influences in CHSH experiments [71, 72]. In a multipartite case it was further shown that for any model based on finite-speed superluminal causal influences, it is possible to construct a multipartite experiment which would reveal a violation of signal locality for this model [72]. The results presented in Sec. 5.4 are consistent with this theorem, since causal models are formulated without any reference to a spacetime structure, allowing causal influences to propagate at sub- or superluminal, instantaneous, or even to the past, as long as it does not create any causal loop.

5.5.3 Beyond Classical Causal Modeling

Causal models are a very useful tool for a systematic study of relaxations of the causal assumptions behind Bell-locality. This covers a large range of models, which despite being fine-tuned, all feature some compromise of how we think about causality, including measurement dependence, retro-causality, and nonlocal influences.

However, not every model fits within the causal modeling framework. One example is to treat the pair of measurement outcomes in a CHSH-inequality test as one entity with multiple causes (X, Y, Λ) . In this case the causal Markov condition does not apply, since both outcomes are a single event, but this seems to avoid the problem, rather than solving it [33]. This argument is akin to the observation that the Born rule seems to be in conflict with the factorization property of classical causal modeling [23]. Related to this idea, one could consider “perspectival” dependences, which are undirected edges that can adopt either direction depending on the perspective and epistemic restrictions of the agent that assigns the causal model [36]. In the CHSH case, for example, Alice might imagine a model where her measurement outcome causes Bob’s via a path through the hidden variable, $A \rightarrow \Lambda \rightarrow B$, while Bob might assign the reverse $B \rightarrow \Lambda \rightarrow A$. Neither of these models are objective, but rather represent the direction in which information is gained about the system [36]. Similar perspectival problems arise in cosmology, where we attribute objective status to quantities like the average density of the universe, which are not empirically accessible to us as observers [73].

Finally the causal modeling framework might need to be generalized to appropriately treat quantum systems. This might entail introducing “non-classical” nodes as a generalization of latent variables [25], which are considered unobservable, thus defying the possibility of interventions. Alternatively one may consider all nodes in principle observable and instead modify the criterion by which independence is decided, such as giving up the factorization part of Reichenbach’s principle [23]. Other work has considered replacing conditional probabilities with quantum conditional states in an effort to generalize Bayes theorem and Bayesian conditioning to the quantum case [74]. An alternative approach is to redefine causal models from bottom up, using the quantum formalism, where nodes are local operations and causal influences are represented by quantum channels [32]. Such a framework accommodates a general notion of

intervention (as the choice of an instrument, or a set of quantum operations), which means that causal discovery remains possible. This framework also recovers classical causal models in the appropriate limit (where all operations are diagonal in a fixed basis), without any changes to the causal structure [32].

References

- [1] Ringbauer, M., Giarmatzi, C., Chaves, R., Costa, F., White, A. G. & Fedrizzi, A. Experimental test of nonlocal causality. *Sci. Adv.* **2**, e1600162–e1600162 (2016).
- [2] Bell, J. S. On the Einstein–Podolsky–Rosen Paradox. *Physics* **1**, 195 (1964).
- [3] Bell, J. S. The theory of local beables. *Epistemological Letters* **9**, 11–24 (1976).
- [4] Hensen, B. *et al.* Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature* **526**, 682–686 (2015).
- [5] Shalm, L. K. *et al.* Strong Loophole-Free Test of Local Realism. *Phys. Rev. Lett.* **115**, 250402 (2015).
- [6] Giustina, M. *et al.* Significant-Loophole-Free Test of Bell’s Theorem with Entangled Photons. *Phys. Rev. Lett.* **115**, 250401 (2015).
- [7] Wiseman, H. M. & Cavalcanti, E. G. *Causarum Investigatio* and the Two Bell’s Theorems of John Bell. *arXiv:1503.06413* (2015).
- [8] Cavalcanti, E. G. & Wiseman, H. M. Bell nonlocality, signal locality and unpredictability (or What Bohr could have told Einstein at Solvay had he known about Bell experiments). *Found. Phys.* **42**, 1329—1338 (2012).
- [9] Zukowski, M. & Brukner, Č. Quantum non-locality—it ain’t necessarily so... *J. Phys. A: Math. Theor* **47**, 424009 (2014).
- [10] Brans, C. H. Bell’s Theorem Does Not Eliminate Fully Causal Hidden Variables. *Int. J. Theor. Phys.* **27**, 219–226 (1988).
- [11] Branciard, C., Brunner, N., Gisin, N., Kurtsiefer, C., Lamas-Linares, A., Ling, A. & Scarani, V. Testing quantum correlations versus single-particle properties within Leggett’s model and beyond. *Nat. Phys.* **4**, 681–685 (2008).
- [12] Hall, M. J. W. Local Deterministic Model of Singlet State Correlations Based on Relaxing Measurement Independence. *Phys. Rev. Lett.* **105**, 250404 (2010).
- [13] Hall, M. J. W. Relaxed Bell inequalities and Kochen-Specker theorems. *Phys. Rev. A* **84**, 022102 (2011).
- [14] Barrett, J. & Gisin, N. How Much Measurement Independence Is Needed to Demonstrate Nonlocality? *Phys. Rev. Lett.* **106**, 100406 (2011).

- [15] Chaves, R., Kueng, R., Brask, J. B. & Gross, D. Unifying Framework for Relaxations of the Causal Assumptions in Bell’s Theorem. *Phys. Rev. Lett.* **114**, 140403 (2015).
- [16] Wood, C. J. & Spekkens, R. W. The lesson of causal discovery algorithms for quantum correlations: causal explanations of Bell-inequality violations require fine-tuning. *New J. Phys.* **17**, 033002 (2015).
- [17] Aktas, D., Tanzilli, S., Martin, A., Pütz, G., Thew, R. & Gisin, N. Demonstration of Quantum Nonlocality in the Presence of Measurement Dependence. *Phys. Rev. Lett.* **114**, 220404 (2015).
- [18] Pütz, G. & Gisin, N. Measurement dependent locality. *New J. Phys.* **18**, 055006 (2016).
- [19] Chaves, R., Majenz, C. & Gross, D. Information–theoretic implications of quantum causal structures. *Nat. Commun.* **6**, 5766 (2015).
- [20] Fritz, T. Beyond Bell’s theorem: correlation scenarios. *New J. Phys.* **14**, 103001 (2012).
- [21] Fritz, T. Beyond Bell’s Theorem II: Scenarios with Arbitrary Causal Structure. *Comm. Math. Phys.* 1–45 (2015).
- [22] Cavalcanti, E. G. & Lal, R. On modifications of Reichenbach’s principle of common cause in light of Bell’s theorem. *J. Phys. A: Math. Theor.* **47**, 424018 (2014).
- [23] Pienaar, J. & Brukner, Č. A graph-separation theorem for quantum causal models. *New J. Phys.* **17**, 073020 (2015).
- [24] Leifer, M. S. & Spekkens, R. W. Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference. *Phys. Rev. A* **88**, 052130 (2013).
- [25] Henson, J., Lal, R. & Pusey, M. F. Theory-independent limits on correlations from generalized Bayesian networks. *New J. Phys.* **16**, 113043 (2014).
- [26] Brukner, Č. Quantum causality. *Nat. Phys.* **10**, 259–263 (2014).
- [27] Oreshkov, O. & Giarmatzi, C. Causal and causally separable processes. *New J. Phys.* **18**, 093020 (2016).
- [28] Munroe, R. XKCD webcomic at <https://xkcd.com/552/> (2009).
- [29] Pearl, J. *Causality*. (Cambridge University Press, 2009).
- [30] Spirtes, P., Glymour, N. & Scheines, R. *Causation, Prediction, and Search, 2nd ed* (The MIT Press, 2001).
- [31] Woodward, J. *Making Things Happen: A Theory of Causal Explanation* (Oxford University Press, 2005).

- [32] Costa, F. & Shrapnel, S. Quantum causal modelling. *New J. Phys.* **18**, 063032 (2016).
- [33] Hausman, D. M. & Woodward, J. Independence, invariance and the causal Markov condition. *Br. J. Philos. Sci.* **50**, 521–583 (1999).
- [34] Näger, P. M. The causal problem of entanglement. *Synthese* **193**, 1127–1155 (2016).
- [35] Clarke, J. private communication.
- [36] Evans, P. W. Quantum causal models, faithfulness and retrocausality. *arXiv:1506.08925* (2015).
- [37] Geiger, D. & Meek, C. Quantifier Elimination for Statistical Problems. In *Proc. 15th Conf. Uncertainty in Artificial Intell.*, 226–235 (1999).
- [38] Tian, J. & Pearl, J. On the Testable Implications of Causal Models with Hidden Variables. In *Proc. 18th Conf. Uncertainty in Artificial Intell.*, 519–527 (Morgan Kaufmann Publishers Inc., 2002).
- [39] Chaves, R., Luft, L., Maciel, T. O., Gross, D., Janzing, D. & Schölkopf, B. Inferring latent structures via information inequalities. *Proc. 30th Conf. Uncertainty in Artificial Intell.* 112–121 (2014).
- [40] Mooij, J. M., Peters, J., Janzing, D., Zscheischler, J. & Schölkopf, B. Distinguishing cause from effect using observational data: methods and benchmarks. *J. Mach. Learn. Res.* **17**, 1–102 (2016).
- [41] Ferrie, C. private communication.
- [42] Clauser, J. F., Horne, M. A., Shimony, A. & Holt, R. A. Proposed Experiment to Test Local Hidden-Variable Theories. *Phys. Rev. Lett.* **23**, 880–884 (1969).
- [43] Zych, M. private communication.
- [44] Pütz, G., Rosset, D., Barnea, T. J., Liang, Y.-C. & Gisin, N. Arbitrarily Small Amount of Measurement Independence Is Sufficient to Manifest Quantum Nonlocality. *Phys. Rev. Lett.* **113**, 190402 (2014).
- [45] Hall, M. J. W. *The Significance of Measurement Independence for Bell Inequalities and Locality*, 189–204 (Springer International Publishing, 2016).
- [46] Fine, A. Hidden Variables, Joint Probability, and the Bell Inequalities. *Phys. Rev. Lett.* **48**, 291–295 (1982).
- [47] Masanes, L., Acín, A. & Gisin, N. General properties of nonsignaling theories. *Phys. Rev. A* **73**, 012112 (2006).

- [48] Pawłowski, M., Kofler, J., Paterek, T., Seevinck, M. & Brukner, Č. Non-local setting and outcome information for violation of Bell’s inequality. *New J. Phys.* **12**, 083051 (2010).
- [49] Dash, D. Restructuring dynamic causal systems in equilibrium. *Proc. 10th Conf. Artificial Intell. and Stat.* (2005).
- [50] Stenner, M. D., Gauthier, D. J. & Neifeld, M. A. The speed of information in a ‘fast-light’ optical medium. *Nature* **425**, 695–698 (2003).
- [51] Toner, B. F. & Bacon, D. Communication cost of simulating Bell correlations. *Phys. Rev. Lett* **91**, 187904 (2003).
- [52] Hossenfelder, S. Free will is dead, let’s bury it. Backreaction (Blog) (2016).
- [53] Cramer, J. G. An Overview of the Transactional Interpretation. *Int. J. Theor. Phys.* **27** (1988).
- [54] Aharonov, Y. & Vaidman, L. The Two-State Vector Formalism of Quantum Mechanics: an Updated Review. *arXiv:quant-ph/0105101* (2001).
- [55] Price, H. & Wharton, K. Disentangling the Quantum World. *Entropy* **17**, 7752–7767 (2015).
- [56] Freedman, S. J. & Clauser, J. F. Experimental Test of Local Hidden-Variable Theories. *Phys. Rev. Lett.* **28**, 938–941 (1972).
- [57] Fedrizzi, A., Herbst, T., Poppe, A. & Zeilinger, A. A wavelength-tunable fiber-coupled source of narrowband entangled photons. *Opt. Express* **15**, 15377–15386 (2007).
- [58] Symul, T., Assad, S. M. & Lam, P. K. Real time demonstration of high bitrate quantum random number generation with coherent laser light. *App. Phys. Lett.* **98** (2011).
- [59] Herrero-Collantes, M. & Garcia-Escartin, J. C. Quantum Random Number Generators. *arXiv:1604.03304* (2016).
- [60] Janzing, D., Balduzzi, D., Grosse-Wentrup, M. & Schölkopf, B. Quantifying causal influences. *Ann. Stat.* **41**, 2324–2358 (2013).
- [61] Korb, K. B., Hope, L. R., Nicholson, A. E. & Axnick, K. Varieties of Causal Intervention. In *PRICAI 2004: Trends in Artificial Intell.*, 322–331 (Springer, 2004).
- [62] Braunstein, S. L. & Caves, C. M. Wringing out better Bell inequalities. *Nucl. Phys. B* **6**, 211–221 (1989).
- [63] Gallicchio, J., Friedman, A. S. & Kaiser, D. I. Testing Bell’s Inequality with Cosmic Photons: Closing the Setting-Independence Loophole. *Phys. Rev. Lett.* **112**, 110405 (2014).

- [64] Garg, A. & Mermin, N. D. Detector inefficiencies in the Einstein-Podolsky-Rosen experiment. *Phys. Rev. D* **35**, 3831–3835 (1987).
- [65] Eberhard, P. H. Background level and counter efficiencies required for a loophole-free Einstein-Podolsky-Rosen experiment. *Phys. Rev. A* **47**, R747–R750 (1993).
- [66] Leggett, A. J. Nonlocal hidden-variable theories and quantum mechanics: An incompatibility theorem. *Found. Phys.* **33**, 1469–1493 (2003).
- [67] Gröblacher, S., Paterek, T., Kaltenbaek, R., Brukner, Č., Zukowski, M., Aspelmeyer, M. & Zeilinger, A. An experimental test of non-local realism. *Nature* **446**, 871–875 (2007).
- [68] Paterek, T., Fedrizzi, A., Gröblacher, S., Jennewein, T., Zukowski, M., Aspelmeyer, M. & Zeilinger, A. Experimental Test of Nonlocal Realistic Theories Without the Rotational Symmetry Assumption. *Phys. Rev. Lett.* **99**, 210406 (2007).
- [69] Colbeck, R. & Renner, R. Hidden variable models for quantum theory cannot have any local part. *Phys. Rev. Lett.* **101**, 050403 (2008).
- [70] Branciard, C. Bell’s local causality, Leggett’s crypto-nonlocality, and quantum separability are genuinely different concepts. *Phys. Rev. A* **88**, 042113 (2013).
- [71] Salart, D., Baas, A., Branciard, C., Gisin, N. & Zbinden, H. Testing the speed of ‘spooky action at a distance’. *Nature* **454**, 861–864 (2008).
- [72] Bancal, J.-D., Pironio, S., Acín, A., Liang, Y.-C., Scarani, V. & Gisin, N. Quantum non-locality based on finite-speed causal influences leads to superluminal signalling. *Nat. Phys.* **8**, 867–870 (2012).
- [73] Davies, T. private communication.
- [74] Leifer, M. S. & Spekkens, R. W. Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference. *Phys. Rev. A* 052130 (2013).

CHAPTER 6

Pushing Joint-Measurement Uncertainty to the Limit

Acknowledgement

This chapter is based on work that was first published in Ref. [1], and the body of this chapter incorporates text of that paper. The experiments were performed with Devon Biggerstaff and the theory work was done by Cyril Branciard.

6.1 Introduction

THE state of a classical physical system is captured by a point in phase space, for example, a set of position and momentum values for a mechanical system. These properties are in principle exactly defined and with a suitable measurement device they can be measured to arbitrarily high precision. In the quantum realm, however, this dream of perfect measurement hits a fundamental limit, the Heisenberg uncertainty principle.

Heisenberg famously illustrated this principle at the example of a γ -ray microscope, used for measuring position and momentum of an electron. The position can be measured by exposing the electron to a flash of light and taking a picture of it, and the momentum can be determined from two such pictures. As with every optical microscope the resolution of such a picture is limited by the wavelength of the light that is used. In order to increase the precision of the position measurement shorter wavelength light must be used. This, however, means that the light is more energetic and thus imparts a larger momentum kick to the electron, which disturbs the momentum measurement by causing the electron to move in a random direction. Heisenberg postulated that this is not a shortcoming of the specific apparatus, but a fundamental feature of nature.

The Heisenberg microscope has become the archetypal illustration of Heisenberg's uncertainty principle and the famous uncertainty relation $\Delta x \Delta p \geq \hbar/2$. What is less well-known is that there is much more to the story than this relation, which was neither derived by Heisenberg, nor does it capture the microscope example. Heisenberg's seminal 1927 paper [2] does not contain any rigorous quantitative statements, but instead describes three *qualitative* features of the uncertainty associated with incompatible observables, see Fig. 6.1:

- (I) Two incompatible observables cannot be arbitrarily well defined on a quantum state.
- (II) Two incompatible observables cannot be jointly measured with arbitrary accuracy.
- (III) The measurement of one such observable disturbs the subsequent measurement of the other.




Figure 6.1: **The three faces of Heisenberg’s uncertainty principle.** (a) **Preparation Uncertainty.** A quantum system cannot be prepared in a state ρ , such that two incompatible observables A and B are arbitrarily well defined on ρ . (b) **Joint-Measurement Uncertainty.** Two incompatible observables A and B cannot be jointly measured on a state ρ with arbitrary accuracy. A joint-measurement device \mathcal{M} can only approximate them by compatible observables A and B , respectively. (c) **Measurement-Disturbance.** The approximate measurement \mathcal{A} of the observable A disturbs the subsequent measurement of the incompatible observable B .

6.2 Heisenberg’s Uncertainty Principle

“The uncertainty principle refers to the degree of indeterminateness in the possible present knowledge of the simultaneous values of various quantities with which the quantum theory deals; it does not restrict, for example, the exactness of a position measurement alone or a momentum measurement alone” - W. Heisenberg, 1930

6.2.1 The Uncertainty Principle

Heisenberg’s seminal 1927 paper “Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik” [2] provided much of the intuition that still guides the study of quantum uncertainty. The purpose of the paper was to motivate and provide intuition¹ for the *uncertainty principle*, which describes qualitative uncertainty trade-offs. The quantitative aspect of the uncertainty principle is captured by so-called *uncertainty relations*, which Heisenberg did not derive in his paper. He did, however, conjecture that a relation of the following form should hold

$$q_1 p_1 \sim h, \tag{6.1}$$

¹The key word in the title is “anschaulich”, which in this context is best translated to “intuitive”.

where q_1 and p_1 represent, depending on the context, various forms of position and momentum uncertainties, respectively. Heisenberg did not prove this relation—which he saw as a physical manifestation of the commutation relations of incompatible observables—but he justified it for a number of explicit physical examples. Heisenberg’s conjecture, Eq. (6.1), captures two central insights

- The scale at which quantum uncertainty becomes important is on the order of the fundamental constant h .
- Uncertainty relations describe trade-offs, rather than absolute limits; any one measurement can, in principle, be made arbitrarily precise at the cost of making the other increasingly imprecise.

Heisenberg realized the status of the uncertainty principle as a fundamental feature of quantum theory: “if there existed experiments, which allowed simultaneously a “sharper” determination of p and q than equation (6.1) permits, then quantum mechanics would be impossible” [2]. Heisenberg’s paper contains at least three distinct ideas which he illustrates with practical examples.

Measurement and disturbance The tradeoff between measurement accuracy and disturbance (III above) is one of the central topics of Heisenberg’s work. Based on his famous γ -ray microscope, he argues that knowing the position of an electron (i.e. measuring it) with (in)accuracy q_1 must cause a disturbance of the electron’s momentum, such that it can only be determined up to the accuracy p_1 [2]. Based on the Compton-effect he concludes that p_1 and q_1 must obey Eq. (6.1).

Joint measurements Heisenberg also discusses an error-trade-off scenario (II above), stating that Eq. (6.1) shows that an “exact energy measurement can only be achieved with according inaccuracy in time” [2].

Preparation uncertainty Although Heisenberg focuses mostly on measurement-based arguments, he typically associates uncertainty with the width of the wavefunction. In particular, he argues that a Gaussian wavepacket with a width of q_1 in position must have a width of at least p_1 in momentum, satisfying Eq. (6.1), and sketches a proof of a preparation uncertainty relation. He also discusses a “fundamental uncertainty in the initial conditions”, which today we would call preparation uncertainty.

An epistemic view point Heisenberg’s discussion is strongly influenced by his Copenhagenist view of the quantum state. For example, in his discussion of the atomic orbit of an electron, he suggests that the orbit “comes into existence only due to it’s observation” [2]. Following this theme, he argues that terms like “position” only make sense in conjunction with an experimental prescription for measuring it. As a consequence of this measurement-centred discussion

Heisenberg generally associates the uncertainty in the knowledge of the position (which he uses synonymously with measurement accuracy) with the width of the (post-measurement) wavefunction. The lack of rigorous distinction between Heisenberg's different conceptual ideas, together with Heisenberg's endorsement of the Kennard relation [3] $\Delta x \Delta p \geq \hbar/2$ might have contributed to the still ongoing debates about the precise formulation of measurement uncertainty relations.

6.2.2 Incompatible Observables

Two observables A and B are compatible, or *jointly measurable*, if there is a measurement \mathcal{M} such that the statistics of A and B can be obtained as the marginal distributions of \mathcal{M} . As an example consider a system consisting of two qubits and two-outcome observables $A = \sigma_z$ and $B = \sigma_x$. The measurement of A on the first qubit is trivially compatible with a measurement of B on the second qubit. The four-outcome measurement \mathcal{M} defined by the POVM elements $\{|0\rangle\langle 0| \otimes |+\rangle\langle +|, |0\rangle\langle 0| \otimes |-\rangle\langle -|, |1\rangle\langle 1| \otimes |+\rangle\langle +|, |1\rangle\langle 1| \otimes |-\rangle\langle -|\}$ reproduces the statistics of both measurements by summing over the respective other outcome. In contrast, when A and B are measured on the same system they are incompatible and there is no single joint measurement that can reproduce the statistics of both for all input states.

The crucial observation is that A and B are non-commuting observables $[A, B] = [\sigma_z \sigma_x - \sigma_x \sigma_z] = 2i\sigma_y \neq 0$, but they do commute when measured on different systems. Any pair of commuting observables can be jointly measured, but the converse is only true for sharp observables (i.e. projective measurements) [4]. In this case, a commonly used measure of compatibility is thus the effective commutator of the two observables

$$C_{AB} = \frac{1}{2i} \langle [A, B] \rangle_\rho \quad (6.2)$$

An interesting property of this quantity is that it depends on the state ρ on which the measurement is performed. Hence, C_{AB} can be zero even for two non-commuting sharp observables, such as $A = \sigma_z$, $B = \sigma_x$, and ρ is in the xz-plane of the Bloch-sphere [5].

In the general case commutativity is not necessary for observables to be compatible. Necessary and sufficient conditions for joint measurability for general POVMs have been derived [6–8] and reviewed in Ref. [4]. For example, sufficiently unsharp versions of a pair of non-commuting unbiased projective measurements, see Fig 6.2, can be jointly measured if and only if the corresponding Bloch vectors satisfy [6]

$$\|\vec{a} + \vec{b}\| + \|\vec{a} - \vec{b}\| \leq 2 \quad (6.3)$$

Besides, or perhaps related to, their role in quantum uncertainty, incompatible observables also play a crucial role in other areas of quantum foundations, such as the study of non-classical correlations. Specifically, joint measurability is equivalent to quantum steering in the sense that steering is possible if and only if the steering party uses incompatible measurements [9, 10]. In the case where Alice uses a pair of incompatible two-outcome POVMs it is even possible




Figure 6.2: **Compatible and incompatible qubit observables.** (a) Two projective observables along orthogonal directions on the Bloch-sphere are incompatible and cannot be jointly measured. There is no physical joint measurement (i.e. vector within the Bloch-sphere) that could reproduce the statistics of both. (b) Sufficiently unsharp versions of the observables in a) can be jointly measured by means of the measurement \mathcal{M}

to demonstrate Bell-nonlocality with a suitable bipartite state and observables for the other party [11]. Hence, there is no local-causal model that can reproduce the correlations arising from pairs of incompatible two-outcome measurements on a shared quantum system. For more general sets of POVMs, however, the latter result does not hold up [9, 12].

6.2.3 Preparation Uncertainty

A quantum system cannot be prepared in a state ρ such that two incompatible observables A and B are arbitrarily well defined on ρ .

The *preparation uncertainty principle* is motivated by Heisenberg's argument that a Gaussian wavepacket with width q_1 in position must have a corresponding width p_1 in momentum, such that $q_1 p_1 \sim \hbar$. The width of the wavepacket is a measure of the statistical spread of the corresponding properties in the state ρ , rather than our ability to measure it. Heisenberg's argument thus implies that it is impossible to *prepare* a quantum state that is very well localized in position and at the same time has a well-defined momentum. Formally, the statistical spread




Figure 6.3: **Preparation Uncertainty Principle.** It is not possible to prepare a quantum state ρ which simultaneously has well-defined values for two incompatible observables A and B . Preparation-uncertainty relations show that a gain in precision for A comes with a loss in precision for B .

(or width of the wavepacket) is quantified by the standard deviation $\Delta A = \sqrt{\text{Tr}[A^2\rho] - \text{Tr}[A\rho]^2}$.

Using this definition, Kennard in 1927 derived the famous uncertainty relation for position and momentum, which was soon generalized by Robertson to arbitrary sharp observables [3, 13]

$$\Delta A \Delta B \geq \frac{1}{2} |\langle [A, B] \rangle_\rho|. \quad (6.4)$$

This “textbook” uncertainty relation is most well-known in its canonical position-momentum form, $\Delta x \Delta p \geq \hbar/2$, where it implies that the product of position uncertainty and momentum uncertainty must be larger than the constant $\hbar/2$.

For other observables, in particular bounded qubit observables, the relation, however, has two rather undesirable properties, which have received more attention in recent rigorous discussions of the uncertainty principle, see e.g. Ref. [5]. It is state-dependent, and it imposes bounds that can in general not be satisfied. In the case of a qubit, for example, whenever ρ lies in the plane spanned by A and B , but is not aligned with either A or B , the right-hand side of Eq. (6.4) vanishes, although both ΔA and ΔB are strictly positive [5]. The state-dependence of the bound itself means that evaluating it requires knowledge of the state, but once A, B and ρ are all specified, there is no freedom for ΔA and ΔB anymore. The relation thus becomes somewhat superfluous, since the relevant uncertainties can be calculated directly [14].




Figure 6.4: **Preparation Uncertainty Relation.** Abbott et al’s tight state-independent preparation uncertainty relation in the $\Delta A/\Delta B$ -plane for $\vec{a} \cdot \vec{b} = 0$ (blue), $\vec{a} \cdot \vec{b} = 1/2$ (orange), and $\vec{a} \cdot \vec{b} = 3/4$ (green) [5]. The gray area is forbidden by the respective uncertainty relation.

When the goal of a preparation uncertainty relation is to describe the minimum-uncertainty states with respect to two incompatible observables, a formulation that depends only on the observables, but not on the state is desirable. In the case of qubits, simple geometric arguments can be used to provide tight state-independent uncertainty relations for arbitrary observables [5]. Tight here means that these relations describe the optimal trade-off between the uncertainties and can always be saturated by some quantum state. For two sharp measurements the optimal trade-off is given by [5]

$$(\Delta A)^2 + (\Delta B)^2 + 2|\vec{a} \cdot \vec{b}| \sqrt{1 - (\Delta A)^2} \sqrt{1 - (\Delta B)^2} \geq 1 + (\vec{a} \cdot \vec{b})^2. \quad (6.5)$$

Curiously, geometric arguments naturally lead to a relation which does not have the product form of the Robertson relation, Eq. (6.4), but rather involve sums of variances. This ensures that as one uncertainty goes to zero, the other does not necessarily diverge, as one would expect, at least for bounded observables.

6.2.4 Joint-Measurement Uncertainty

Two incompatible observables A and B cannot be jointly measured on a quantum state ρ with arbitrary accuracy.

In contrast to preparation uncertainty, this is not a statement about the quantum state, but about our ability to *measure* the two incompatible observables. When A and B are incompatible, there is no joint measurement that reproduces the statistics of both observables for any input state. However, the joint measurement can still be approximated² by measuring an observable \mathcal{M} and defining approximations $\mathcal{A}=f(\mathcal{M})$ and $\mathcal{B}=g(\mathcal{M})$, see Fig. 6.5. Specifically, for an outcome m of \mathcal{M} , the joint-measurement apparatus outputs $f(m)$ to approximate the measurement of A and $g(m)$ to approximate the measurement of B . In fact, every measurement can be used as an approximate joint measurement for any pair of observables, some are just more useful than others [15].




Figure 6.5: **Joint-measurement uncertainty principle.** Two incompatible observables A and B cannot be jointly measured on a quantum state ρ with arbitrary accuracy. Any joint-measurement performed on the state ρ , can only measure approximate observables \mathcal{A} and \mathcal{B} . Measurement uncertainty relations aim to quantify the optimal trade-off between the approximation errors for A and B .

Joint-measurement uncertainty relations aim to quantify the optimal trade-off between the approximation errors for A and B . There are actually at least two questions that could be asked

1. What is the best performance of an approximate joint measurement device over all possible input states?
2. How well can the joint measurement be approximated on a given quantum state ρ .

In the first case, a state-independent uncertainty relation is most appropriate, since it should hold for every quantum state, see e.g. Ref. [16]. In the second case, which we will discuss here,

²Recall that, for example, sufficiently unsharp versions of an incompatible pair of projective measurements are jointly measurable, see Sec. 6.2.2.

the trade-off in approximation errors should be optimized for a specific, known quantum state ρ . Since in general this might work better for some states than for others, we expect this case to be described by a state-dependent relation. In contrast to the preparation uncertainty case, once A, B and ρ are fixed, there is still freedom in the approximation errors, since they also depend on the implementation of the joint measurement.

Formally the measurement inaccuracy or approximation error can be quantified by the root-mean-square error between the approximate and ideal observables [17–20]

$$\begin{aligned}\varepsilon_{\mathcal{A}} &= \text{Tr}[(\mathcal{A}-A \otimes \mathbb{1})^2 \rho \otimes |\zeta\rangle\langle\zeta|]^{1/2}, \\ \varepsilon_{\mathcal{B}} &= \text{Tr}[(\mathcal{B}-B \otimes \mathbb{1})^2 \rho \otimes |\zeta\rangle\langle\zeta|]^{1/2},\end{aligned}\tag{6.6}$$

where $|\zeta\rangle\langle\zeta|$ is the initial state of an ancillary system used for the approximate measurement. Heisenberg’s intuition was that the measurement errors should also be restricted by a relation of the form $\varepsilon_{\mathcal{A}}\varepsilon_{\mathcal{B}} \sim h$. Hence, it was generally assumed that Robertson’s relation, Eq. (6.4) would also hold for joint-measurement uncertainty, with the measurement inaccuracies $\varepsilon_{\mathcal{A}}$ and $\varepsilon_{\mathcal{B}}$ replacing the standard deviations ΔA and ΔB in Eq. (6.4)

$$\varepsilon_{\mathcal{A}}\varepsilon_{\mathcal{B}} \geq |C_{AB}|.\tag{6.7}$$

It was not until 1965, that this relation was formally proven for position and momentum measurements by Arthurs and Kelly [21] and by Arthurs and Goodman in 1988 for the general case [22]. The derivation of this relation, however, relied on the restrictive assumption that the mean errors are independent of the measured state, and can be violated when all possible approximate measurements are considered³ [17, 19, 23], see Fig. 6.6.

To overcome this problem, Ozawa derived a universally valid joint-measurement uncertainty relation [18] (and a similar relation was derived by Hall [15] with $\Delta\mathcal{A}$ and $\Delta\mathcal{B}$ instead of ΔA and ΔB)

$$\varepsilon_{\mathcal{A}}\varepsilon_{\mathcal{B}} + \varepsilon_{\mathcal{A}}\Delta B + \varepsilon_{\mathcal{B}}\Delta A \geq |C_{AB}|.\tag{6.8}$$

Ozawa’s relation, however, although it cannot be violated by any quantum state, can also not be saturated, except in special cases [19]. It was not until 2013 that Branciard [19, 24] derived a joint-measurement uncertainty relation, which is universally valid and can be saturated.

$$\Delta B^2\varepsilon_{\mathcal{A}}^2 + \Delta A^2\varepsilon_{\mathcal{B}}^2 + 2\sqrt{\Delta A^2\Delta B^2 - C_{AB}^2}\varepsilon_{\mathcal{A}}\varepsilon_{\mathcal{B}} \geq C_{AB}^2,\tag{6.9}$$

where ΔA and ΔB are the standard deviations of A and B on the state ρ . In the case where ρ is a pure state, the relation is *tight* [19], which means that it describes the optimal trade-off between the inaccuracies of the approximate measurements \mathcal{A} and \mathcal{B} , see Fig. 6.6. Interestingly, achieving this optimal trade-off and saturating Eq. (6.9) may require the approximate observables \mathcal{A} and \mathcal{B} to have different spectra than A and B —i.e. the optimal output values $f(m)$ and $g(m)$ may

³In particular the Arthur-Kelly relation implies that as one estimation error goes to zero, the other must go to infinity. For bounded operators, such as qubit observables this cannot happen. In the case of the Robertson relation this is not a problem, since there the right-hand side goes to zero as one of $\Delta A, \Delta B$ goes to zero.




Figure 6.6: **Joint-measurement uncertainty relations.** Joint-measurement uncertainty relations are manifest as forbidden regions in the $\varepsilon_{\mathcal{A}}/\varepsilon_{\mathcal{B}}$ -plane, here shown for two orthogonal, projective qubit measurements and a pure state that is orthogonal to both, see Fig. 6.11. Values below either curve are not allowed by the respective uncertainty relation. The naive Arthurs-Kelly relation (6.7) (dashed line, top-right) imposes a very loose bound that can be violated by general approximate measurements. Ozawa’s relation (6.8) (dashed line, bottom-left) is universally valid, but cannot be saturated in practice. Branciard’s joint-measurement relation (6.9) (solid, blue) is tight for pure states and quantifies the optimal trade-off between the inaccuracies of the approximate measurements \mathcal{A} and \mathcal{B} .

not be eigenvalues of A and B .

6.2.5 Measurement-Disturbance

A measurement of one observable disturbs a subsequent measurement of an incompatible observable.

The measurement-disturbance scenario is the motive of the famous Heisenberg microscope. In a sequential measurement of two incompatible observables, the more accurate the first measurement is, the more it disturbs the subsequent measurement. The aim of measurement-disturbance relations is to find the optimal trade-off between the inaccuracy $\varepsilon_{\mathcal{A}}$ of the first measurement and the disturbance $\eta_{\mathcal{B}}$ imparted onto the second, see Fig. 6.7. As in the previous case, one could be interested in a state-independent assessment of the sequential-measurement device. This would quantify the minimal disturbance that such a device must impart on some quantum state, given that it can perform the first to a certain accuracy (on some potentially different state) [16, 25].

The measurement-disturbance scenario can in fact be seen as a special of the joint-measurement scenario discussed above. To see this, simply draw a box around the sequential measurement scheme and call it a joint-measurement device, see Fig. 6.7b. The $\varepsilon_{\mathcal{B}}$ of the second measurement is then interpreted as the disturbance $\eta_{\mathcal{B}}$ caused by the first measurement. Mathematically




Figure 6.7: **Measurement-disturbance principle.** (a) An approximate measurement \mathcal{A} of the observable A disturbs the subsequent measurement of the incompatible observable B , thus turning it into a disturbed (or approximate) measurement \mathcal{B} . The more accurate the first measurement is, the more it disturbs the second. (b) The measurement-disturbance scenario is a special case of the joint-measurement scenario, where the approximate measurement \mathcal{B} is actually a disturbed measurement of the ideal observable B .

this identification requires the disturbed observable \mathcal{B} to have the same spectrum as the ideal observable B (e.g. if B is ± 1 -valued, then \mathcal{B} must be too).

The non-optimal relations (6.7) and (6.8) are not affected by this restriction and can equally be formulated for the measurement-disturbance case with $\varepsilon_{\mathcal{B}} = \eta_{\mathcal{B}}$, see Fig. 6.8. Branciard’s relation, on the other hand quantifies the optimal trade-off, which is typically achieved with approximate observables that do not have the same spectrum as the ideal observables. Hence, Branciard’s relation is in general stronger (i.e. more restrictive) in the measurement-disturbance case than in the joint-measurement case [24]. In particular in the case of ± 1 -valued observables (such that $A^2=B^2=\mathbb{1}$), when also imposing $\mathcal{A}^2=A^2=\mathbb{1}$ and/or $\mathcal{B}^2=B^2=\mathbb{1}$, relation (6.9) is modified to [1, 24]:

$$\text{Eq. (6.9)} \begin{cases} \varepsilon_{\mathcal{A}} \rightarrow \sqrt{1 - (1 - \varepsilon_{\mathcal{A}}^2/2)^2} \\ \varepsilon_{\mathcal{B}} \rightarrow \sqrt{1 - (1 - \varepsilon_{\mathcal{B}}^2/2)^2} \end{cases}, \quad (6.10)$$

where the replacement is made for the observable(s) on which the same-spectrum assumption is imposed. As shown in Fig. 6.8, this relation also imposes an upper bound on the inaccuracy and disturbance for the observables that are assumed to satisfy the same-spectrum assumption.

6.3 Measuring Measurement Uncertainty

Although Heisenberg discussed several examples that would fall into the joint-measurement and measurement-disturbance scenarios, he mostly focused on the properties of the post-measurement quantum state, rather than the measurement. Kennard in deriving the famous preparation uncertainty relation concluded that measurement uncertainty cannot be isolated from uncertainty of the state, since the experimenter does not have access to a “real” value of the measured observable to compare against [3]. This has led to the long-standing belief that measurement uncertainty relations are experimentally untestable [26].

The game changed when Ozawa proposed an method for experimentally estimating approx-




Figure 6.8: **Measurement-disturbance relations in the $\varepsilon_{\mathcal{A}}/\eta_{\mathcal{B}}$ -plane.** In the case where the approximate measurement of \mathcal{A} disturbs a subsequent measurement of \mathcal{B} , measurement-disturbance relations illustrate the forbidden regions in the $\varepsilon_{\mathcal{A}}/\eta_{\mathcal{B}}$ -plane, here shown for two orthogonal, projective qubit measurements and a pure state that is orthogonal to both, see Fig. 6.11. Values below either curve are not allowed by the respective uncertainty relation. The naive Arthurs-Kelly relation (6.7) (dashed line, top-right) and Ozawa’s universally valid relation (6.8) (dashed line, bottom-left) are the same as in the joint-measurement case of Fig. 6.6. As a consequence of the same spectrum assumption Branciard’s measurement-disturbance relation (6.9) (solid, green) is stricter than the corresponding joint-measurement uncertainty relation (solid, blue), and also imposes an upper bound on the disturbance. Branciard’s relation is tight for pure states and quantifies the optimal trade-off between the inaccuracies of \mathcal{A} and the disturbance imposed on \mathcal{B} .

imation errors—the *three-state method* [27]. This method results directly from a decomposition of $\varepsilon_{\mathcal{A}}$, see Sec. 6.3.2, and is reminiscent of measurement tomography where $\varepsilon_{\mathcal{A}}$ is estimated from the measurement statistics of \mathcal{M} on the states ρ , $A\rho A$ and $(\mathbf{1}+A)\rho(\mathbf{1}+A)/\|\cdot\|$, and similarly for $\varepsilon_{\mathcal{B}}$. As in the case of quantum tomography, the three-state method relies on exact preparation of the additional states, and the estimates will thus be somewhat confounded by preparation errors.

This requirement for exact state-preparation can be avoided using the conceptually very different *weak measurement method* [28], where the state is subject to a weak pre-measurement of A or B before the approximating measurement \mathcal{M} . The inaccuracy $\varepsilon_{\mathcal{A}}$ is then estimated from the joint statistics of \mathcal{M} and the pre-measurement of A , and similarly for B . While this method might be conceptually more appealing than the three-state method, since all measurements are performed on the state ρ , the weak measurement necessarily disturbs the state. This method can thus only hope to satisfy Eq. (6.9) in the limit of vanishing measurement-strength of the pre-measurement.

The rest of this section is devoted to formalizing these methods in a way that they are robust to experimental imperfections, which will be somewhat technical. The general idea is




Figure 6.9: **Conceptual idea of determining joint-measurement accuracy.** (a) In the three-state method the approximation error $\varepsilon_{\mathcal{A}}$ is estimated from the measurement statistics of the approximate joint measurement \mathcal{M} on three different quantum states. (b) The weak measurement method uses a weak pre-measurement of the observable A followed by the approximating measurement \mathcal{M} to determine the inaccuracy of the approximation \mathcal{A} . The hadamard gates (dashed) are only needed when measuring $A = \sigma_x$, and are omitted for the weak measurement of $B = \sigma_z$. Note that the state ρ entering the approximate joint-measurement device is disturbed by the semi-weak measurement and thus partially mixed.

captured by Fig. 6.10 and Eqs. (6.15),(6.16),(6.12).

6.3.1 Relating $\varepsilon_{\mathcal{A}}$ to Experimental Data

Consider the joint approximation of two incompatible ± 1 -valued, sharp qubit observables $A = \vec{a} \cdot \vec{\sigma}$ and $B = \vec{b} \cdot \vec{\sigma}$ on a single-qubit state ρ . The measurement apparatus used for this joint approximation is treated as a black box, which acts on qubits and outputs binary values $m = \pm 1$. In full generality, it is a 2-outcome POVM $\mathbb{M} = \{M_+, M_-\}$, with elements $M_{\pm} = \frac{1}{2}(\mathbb{1} \pm (\mu\mathbb{1} + \vec{m} \cdot \vec{\sigma}))$, where $\mu \in \mathbb{R}$ and $\vec{m} = (m_x, m_y, m_z)$ a vector in the Bloch sphere, such that $|\mu| + \|\vec{m}\| \leq 1$. We also define the observable

$$\mathcal{M} = M_+ - M_- = \mu\mathbb{1} + \vec{m} \cdot \vec{\sigma}. \quad (6.11)$$

The outputs m are used to define the estimates $f(m)$ and $g(m)$ for A and B , respectively. Generalizing Eq. (6.6) the root-mean-square error $\varepsilon_{\mathcal{A}}$ (and equivalently $\varepsilon_{\mathcal{B}}$) can be decomposed as [15, 19, 27]:

$$\begin{aligned} \varepsilon_{\mathcal{A}}^2 &= \sum_m \text{Tr}[(A - f(m)\mathbb{1})M_m(A - f(m)\mathbb{1})\rho] \\ &= 1 + \sum_m f(m)^2 \text{Tr}[M_m\rho] - 2\sum_m f(m) \text{Re}[\text{Tr}[M_m A\rho]], \end{aligned} \quad (6.12)$$

using $A^2 = \mathbb{1}$ and $\text{Tr}[A^2\rho] = 1$ for a ± 1 -valued observable A and a pure-state ρ . The second term $\text{Tr}[M_m\rho]$ is simply the probability of outcome m , which is measured directly. Hence, estimating the inaccuracy $\varepsilon_{\mathcal{A}}$ comes down to estimating the last term in the expansion (6.12) from experimental data, which is what the three-state method and the weak-measurement method were developed for. However, as we will see below, both methods rely on idealized assumptions, which are not robust against experimental imperfections [1].

In the case where \mathbb{M} is a two-outcome measurement we can use the definition of \mathcal{M} in

Eq. (6.11) to define

$$\alpha_{\mathcal{M}} := \text{Re} \left[\text{Tr}[\mathcal{M}A\rho] \right] = \mu \langle A \rangle_{\rho} + \vec{m} \cdot \vec{a}, \quad (6.13)$$

such that $\text{Re} \left[\text{Tr}[\mathcal{M}_{\pm}A\rho] \right] = \frac{1}{2} (\langle A \rangle_{\rho} \pm \alpha_{\mathcal{M}})$ (equivalently $\beta_{\mathcal{M}}$ for the approximate/disturbed measurement \mathcal{B}). Equation (6.13) shows, that estimating the term $\text{Re} \left[\text{Tr}[\mathcal{M}_m A\rho] \right]$ can be reduced to the problem of estimating the unknown parameters μ and \vec{m} of the black-box measurement⁴.

6.3.2 The Three-State Method

The three-state method is motivated by decomposing the term $\alpha_{\mathcal{M}}$ in Eq. (6.13) as [27]

$$\alpha_{\mathcal{M}} = \text{Re} \left[\text{Tr}[\mathcal{M}A\rho] \right] = \frac{1}{2} \left[\text{Tr} \left[\mathcal{M}(\mathbb{1}+A)\rho(\mathbb{1}+A) \right] - \text{Tr} \left[\mathcal{M}A\rho A \right] - \text{Tr} \left[\mathcal{M}\rho \right] \right]. \quad (6.14)$$

In other words, $\alpha_{\mathcal{M}}$ can be estimated from the expectation values of M on the three states ρ , $A\rho A$ and $(\mathbb{1}+A)\rho(\mathbb{1}+A)/\|(\mathbb{1}+A)\rho(\mathbb{1}+A)\|$. This decomposition, however, relies on perfect state-preparations, and is not robust against experimental imperfections.

The idealized assumptions behind Eq. (6.14) can be avoided by directly estimating the parameters μ and \vec{m} in the definition of $\alpha_{\mathcal{M}}$. Specifically, the expectation value of \mathcal{M} on the state ρ with Bloch-vector $\vec{\rho}$ is given by $\langle \mathcal{M} \rangle_{\rho} = \mu + \vec{m} \cdot \vec{\rho}$, and similarly for the other two states ρ_1 and ρ_2 . To be compatible with the experimental observations, (μ, \vec{m}) must then be in the set

$$S_{\rho, \rho_1, \rho_2}^{exp} = \left\{ (\mu, \vec{m}) \left| \begin{array}{l} |\mu| + \|\vec{m}\| \leq 1 \quad \text{and} \\ \left\{ \begin{array}{l} \mu + \vec{m} \cdot \vec{\rho} = \langle \mathcal{M} \rangle_{\rho}^{exp} \\ \mu + \vec{m} \cdot \vec{\rho}_1 = \langle \mathcal{M} \rangle_{\rho_1}^{exp} \\ \mu + \vec{m} \cdot \vec{\rho}_2 = \langle \mathcal{M} \rangle_{\rho_2}^{exp} \end{array} \right. \end{array} \right. \right\}. \quad (6.15)$$

As illustrated in Fig. 6.10, the first constraint ensures that the measurement is physical (in particular that the vector \vec{m} is within the Bloch sphere), while the latter 3 have the form of hyperplanes. The values of (μ, \vec{m}) compatible with this set define bounds on $\alpha_{\mathcal{M}}$,

$$\alpha_{\mathcal{M}}^{min(max)} = \min(\max)_{(\mu, \vec{m}) \in S_{\rho, \rho_1, \rho_2}^{exp}} \left[\mu \langle A \rangle_{\rho} + \vec{m} \cdot \vec{a} \right], \quad (6.16)$$

such that $\alpha_{\mathcal{M}} \in [\alpha_{\mathcal{M}}^{min}, \alpha_{\mathcal{M}}^{max}]$. Using, for example, Monte Carlo sampling, the experimental uncertainties on $\langle M \rangle_{\rho}^{exp}$ and $\vec{\rho}$ can be explicitly taken into account in the set $S_{\rho, \rho_1, \rho_2}^{exp}$. In contrast to the original three-state method, there are no assumptions on the state preparation, other than that it is well-characterised such that the constraints on (μ, \vec{m}) from Eq. (6.15) to be sharply defined. The choice $\rho_1 \simeq A\rho A$ and $\rho_2 \simeq (\mathbb{1}+A)\rho(\mathbb{1}+A)/\|(\mathbb{1}+A)\rho(\mathbb{1}+A)\|$, motivated by Eq. (6.14), ensures that the range of possible values for $\alpha_{\mathcal{M}}$ is small (i.e. $\alpha_{\mathcal{M}}^{min} \simeq \alpha_{\mathcal{M}}^{max}$), but

⁴This could in principle be done using measurement tomography, but as we will see below, the three-state and weak-measurement methods allow for a more direct and physically meaningful estimation that avoids some of the practical problems of tomography.




Figure 6.10: **Estimation of $\alpha_{\mathcal{M}}$ in the presence of experimental imperfections.** The constraints on (μ, \vec{m}) defining the set $S_{\rho, \rho_1, \rho_2}^{exp}$ can be interpreted as hypersurfaces in four dimensions. In the case where $\mu = 0$ (i.e. where the measurement is unbiased), these can be visualized as hyperplanes in the (m_x, m_y, m_z) -space. The constraint $|\mu| + \|\vec{m}\| \leq 1$ ensures that the value of \vec{m} lies within the Bloch-sphere. The red ball at the intersection of the planes indicates the position of \vec{m} in the specific example here (the size of this ball has no significance).

virtually any set of states ρ, ρ_1, ρ_2 would allow for non-trivial constraints on $\alpha_{\mathcal{M}}$.

Adding a fourth state, such that the set of four states is informationally complete (e.g. the state ρ'_2 , used for the estimation of $\varepsilon_{\mathcal{B}}$) this technique amounts to tomography of the measurement apparatus. In the presence of experimental imperfections this would in general require a maximum likelihood approach to ensure the result is physical, which is significantly more complicated and might also lead to biased estimates, see Chap. 2.

6.3.3 The Weak-Measurement Method

The idea behind the weak measurement method is to actually measure the term $\alpha_{\mathcal{M}} = \text{Re} [\text{Tr}[\mathcal{M}A\rho]]$ by identifying it with a so-called *weak-valued average* [28]. Specifically it is the average of the product ma over the weak-valued joint probability distribution $P_{\rho}^{wv}(m, a)$, which represents the probability that an initial weak measurement of A yields result a and the final measurement \mathcal{M} on ρ yields outcome m [29], such that

$$P_{\rho}^{wv}(m, a) = \text{Re} [\text{Tr}[M_m \Pi_a^A \rho]] \quad (6.17)$$

$$\alpha_{\mathcal{M}} = \sum_{m, a} m a P_{\rho}^{wv}(m, a). \quad (6.18)$$

Here Π_a^A denotes the projector corresponding to the eigenvalue a of A (such that $A = \Pi_{+1}^A - \Pi_{-1}^A$ in our case of a dichotomic observable, with $a = \pm 1$). Note that weak-valued joint probabilities can be negative, but are always normalized: $\sum_{m, a} P_{\rho}^{wv}(m, a) = 1$ [29].

Since the weak measurement cannot be infinitely weak in practice, it will disturb the state. Hence, the state ρ , which enters the measurement apparatus and on which the inaccuracies $\varepsilon_{\mathcal{A}}$ and $\varepsilon_{\mathcal{B}}$ are defined, must be taken to be the perturbed state after the weak measurement. The

method itself works for any measurement strength [1, 30]. Increasing the strength, however, results in a more mixed state ρ , which in turn makes it harder to approach and saturate the bounds imposed by the measurement uncertainty relations (6.9), and (6.10)

In the present case the joint measurement of $A = \sigma_x$ and $B = \sigma_z$ is approximated on the state $\rho_0 = | + y \rangle \langle + y |$. A weak measurement of the system with measurement strength $|\kappa|$ is implemented by coupling the system to a meter qubit in the state $|\kappa\rangle = \sqrt{\frac{1+\kappa}{2}}|0\rangle + \sqrt{\frac{1-\kappa}{2}}|1\rangle$, with $\kappa \in [-1, 1]$, using a controlled-NOT gate [31, 32], and subsequently measuring the meter in σ_z . This setup natively performs a weak σ_z -measurement on the system, which can be turned into a σ_x -measurement by means of additional Hadamard gates (U_H) on the system before and after the CNOT. The state after the unitary transformation $U_A = (U_H \otimes \mathbb{1}) \cdot U_{\text{CNOT}} \cdot (U_H \otimes \mathbb{1})$ is then given by

$$\begin{aligned} \rho_{12}^{A,th} &= U_A \cdot (| + y \rangle \langle + y | \otimes |\kappa\rangle \langle \kappa |) \cdot U_A^\dagger \\ &= \frac{1}{4} \left(\mathbb{1} + \sqrt{1 - \kappa^2} \sigma_Y \otimes \mathbb{1} + \sqrt{1 - \kappa^2} \mathbb{1} \otimes \sigma_X + \sigma_Y \otimes \sigma_X + \kappa \sigma_X \otimes \sigma_Z - \kappa \sigma_Z \otimes \sigma_Y \right), \end{aligned} \quad (6.19)$$

with subscripts 1 and 2 denoting system and ancillary qubit, respectively, while the superscript “*th*” indicates that these are ideal, theoretical states. After the interaction, the first qubit is in the state

$$\rho_1^{A,th} = \text{Tr}_2 \rho_{12}^{A,th} = \frac{1}{2} \left(\mathbb{1} + \sqrt{1 - \kappa^2} \sigma_Y \right). \quad (6.20)$$

The projective σ_z -measurement of the ancilla with initial state $|\kappa\rangle$, after the interaction U_A effectively amounts to performing a POVM with elements $\text{Tr}_2[(U_A^\dagger \cdot (\mathbb{1} \otimes \frac{\mathbb{1} \pm \sigma_Z}{2}) \cdot U_A) \cdot (\mathbb{1} \otimes |\kappa\rangle \langle \kappa |)] = \frac{\mathbb{1} \pm \kappa \sigma_X}{2}$, which is a weak measurement of $A = \sigma_x$ on the first qubit. Recalling that $\mathcal{M} = \mu \mathbb{1} + \vec{m} \cdot \vec{\sigma}$, the joint expectation value of this weak measurement and the final strong measurement of \mathcal{M} on the state $\rho_{12}^{A,th}$ of Eq. (6.19) is thus given by

$$\langle \mathcal{M} \otimes \sigma_Z \rangle_{\rho_{12}^{A,th}} = \kappa m_X. \quad (6.21)$$

Similarly, for a weak pre-measurement of B using $U_B = U_{\text{CNOT}}$ the state of the first qubit after the interaction is

$$\rho_1^{B,th} = \text{Tr}_2 \rho_{12}^{B,th} = \frac{1}{2} \left(\mathbb{1} + \sqrt{1 - \kappa^2} \sigma_Y \right) = \rho_1^{A,th}. \quad (6.22)$$

The projective σ_z -measurement of the second qubit of $\rho_{12}^{B,th}$ (with outcome $b = \pm 1$) effectively implements a POVM $\{ \frac{\mathbb{1} \pm \kappa \sigma_Z}{2} \}$ —a weak measurement of $B = \sigma_z$ —on the first qubit, such that the joint expectation value with \mathcal{M} is given by

$$\langle \mathcal{M} \otimes \sigma_Z \rangle_{\rho_{12}^{B,th}} = \kappa m_Z. \quad (6.23)$$

For $A = \sigma_x$, and $B = \sigma_z$, using (6.11) (with $\langle A \rangle_{\rho^{th}} = \langle B \rangle_{\rho^{th}} = 0$), together with (6.21) and (6.23),

implies

$$\alpha_{\mathcal{M}} = m_x = \frac{\langle \mathcal{M} \otimes \sigma_z \rangle_{\rho_{12}^{A,th}}}{\kappa}, \quad (6.24)$$

$$\beta_{\mathcal{M}} = m_z = \frac{\langle \mathcal{M} \otimes \sigma_z \rangle_{\rho_{12}^{B,th}}}{\kappa}. \quad (6.25)$$

Hence, the values $\alpha_{\mathcal{M}}$ and $\beta_{\mathcal{M}}$ can in principle be estimated directly from the experimentally accessible quantities $\langle \mathcal{M} \otimes \sigma_z \rangle_{\rho_{12}^{A,th}}$ and $\langle \mathcal{M} \otimes \sigma_z \rangle_{\rho_{12}^{B,th}}$. In practice, however, the state preparation and the 2-qubit interactions will necessarily be imperfect, resulting in approximate states $\rho_{12}^A \simeq \rho_{12}^{A,th}$ and $\rho_{12}^B \simeq \rho_{12}^{B,th}$. These imperfect states—provided they are carefully characterized—can nevertheless be used to restrict the possible values of $\alpha_{\mathcal{M}}$ and $\beta_{\mathcal{M}}$ to small intervals, in a manner similar to that used in the three-state method above. Furthermore, ρ_1^A and ρ_1^B (with $\rho_1^{A/B} = \text{Tr}_2 \rho_{12}^{A/B}$) will in practice be slightly different. Therefore the state ρ —on which the joint measurement of A and B is approximated, and which enters the definition of the inaccuracies $\varepsilon_{\mathcal{A}}$ and $\varepsilon_{\mathcal{B}}$, of $\alpha_{\mathcal{M}}$ and $\beta_{\mathcal{M}}$, and of $\langle A \rangle_{\rho}$ and $\langle B \rangle_{\rho}$ in particular—should be taken to be the average state $\rho = \frac{1}{2}(\rho_1^A + \rho_1^B)$.

The value of $\alpha_{\mathcal{M}}$ can now be bounded using the experimentally accessible expectation values $\langle \mathcal{M} \otimes \sigma_z \rangle_{\rho_{12}^A}^{exp}$, and $\langle \mathcal{M} \otimes \mathbb{1} \rangle_{\rho_1^A}^{exp} = \langle \mathcal{M} \rangle_{\rho_1^A}^{exp}$. Specifically

$$\langle \mathcal{M} \otimes \mathbb{1} \rangle_{\rho_{12}^A} = \mu + \vec{m} \cdot \vec{\rho}_1^A \quad (6.26a)$$

$$\langle \mathcal{M} \otimes \sigma_z \rangle_{\rho_{12}^A} = \mu \rho_{12}^{A,1z} + \vec{m} \cdot \vec{\rho}_{12}^{A,\cdot z} \quad (6.26b)$$

$$\text{with } \rho_{12}^{A,1z} = \langle \mathbb{1} \otimes \sigma_z \rangle_{\rho_{12}^A}$$

$$\text{and } \vec{\rho}_{12}^{A,\cdot z} = (\langle \sigma_x \otimes \sigma_z \rangle_{\rho_{12}^A}, \langle \sigma_y \otimes \sigma_z \rangle_{\rho_{12}^A}, \langle \sigma_z \otimes \sigma_z \rangle_{\rho_{12}^A}),$$

To be compatible with the experimental observations, (μ, \vec{m}) must therefore be in the set

$$T_{\rho_{12}^A}^{exp} = \left\{ (\mu, \vec{m}) \left| \begin{array}{l} |\mu| + \|\vec{m}\| \leq 1 \quad \text{and} \\ \left\{ \begin{array}{l} \mu \rho_{12}^{A,1z} + \vec{m} \cdot \vec{\rho}_{12}^{A,\cdot z} = \langle \mathcal{M} \otimes \sigma_z \rangle_{\rho_{12}^A}^{exp} \\ \mu + \vec{m} \cdot \vec{\rho}_1^A = \langle \mathcal{M} \rangle_{\rho_1^A}^{exp} \end{array} \right. \right. \right\},$$

which implies that $\alpha_{\mathcal{M}} \in [\alpha_{\mathcal{M}}^{min}, \alpha_{\mathcal{M}}^{max}]$ with

$$\alpha_{\mathcal{M}}^{min(max)} = \min(\max)_{(\mu, \vec{m}) \in T_{\rho_{12}^A}^{exp}} [\mu \langle A \rangle_{\rho} + \vec{m} \cdot \hat{a}]. \quad (6.27)$$

As with the three-state method, careful characterization of the state ρ_{12}^A is essential for the constraints on (μ, \vec{m}) , describing the set $T_{\rho_{12}^A}^{exp}$ in (6.27) to be sharply defined. An analogous analysis applies to $\beta_{\mathcal{M}}$, using the state ρ_{12}^B .

Since the set (6.27) is defined by only two hyperplane constraints, rather than three as in the case of the three-state method, the bounds are expected to be relatively weak. The

constraints can, however, be strengthened significantly by taking into account experimental data from both the weak measurements of A and B instead of treating them separately. The measurement parameters (μ, \vec{m}) are then, for both $\varepsilon_{\mathcal{A}}$ and $\varepsilon_{\mathcal{B}}$, estimated from the set

$$T^{exp} = \left\{ (\mu, \vec{m}) \left| \begin{array}{l} |\mu| + \|\vec{m}\| \leq 1 \quad \text{and} \\ \mu \rho_{12}^{A,1z} + \vec{m} \cdot \vec{\rho}_{12}^{A,\cdot z} = \langle \mathcal{M} \otimes \sigma_z \rangle_{\rho_{12}^A}^{exp} \\ \mu \rho_{12}^{B,1z} + \vec{m} \cdot \vec{\rho}_{12}^{B,\cdot z} = \langle \mathcal{M} \otimes \sigma_z \rangle_{\rho_{12}^B}^{exp} \\ \mu + \vec{m} \cdot \vec{\rho} = \langle \mathcal{M} \rangle_{\rho}^{exp} \end{array} \right. \right\}. \quad (6.28)$$

6.3.4 Inaccuracies from $\alpha_{\mathcal{M}}$

Once $\alpha_{\mathcal{M}}$ has been bounded to some interval $[\alpha_{\mathcal{M}}^{min}, \alpha_{\mathcal{M}}^{max}]$, Eq. (6.12) can be used to estimate $\varepsilon_{\mathcal{M}}$. For this, however, it is also necessary to specify the approximating function $f(m)$. This is what distinguishes the joint-measurement from the measurement-disturbance case, since the latter restricts the approximating function for B to have the same spectrum as the ideal observable.

With the Same-Spectrum Assumption

In the case where the same-spectrum assumption is imposed the values of $f(m)$ are constrained to eigenvalues of the ideal observable, here $f(m) = \pm 1$. Hence f is either constant, such that $f(+1) = f(-1) = \pm 1$, or balanced, such that $f(+1) = -f(-1) = \tau$, with $\tau = \pm 1$. Eq. (6.12) then reduces to

$$\begin{aligned} f \text{ constant} &\rightarrow \varepsilon_{\mathcal{A}}^2 = 2 \mp 2 \langle A \rangle_{\rho} \\ f \text{ balanced} &\rightarrow \varepsilon_{\mathcal{A}}^2 = 2 - 2\tau \alpha_{\mathcal{M}} \end{aligned}$$

In the case where f is constant, $\varepsilon_{\mathcal{A}}^2$ is independent of \mathcal{M} and the role of f as a useful approximation is rather questionable. In the balanced case, the bounds on $\alpha_{\mathcal{M}}$ directly and linearly transform into bounds on $\varepsilon_{\mathcal{A}}$, such that

$$(\varepsilon_{\mathcal{A}}^{min(max)})^2 = \begin{cases} 2 - 2\alpha_{\mathcal{M}}^{max(min)} & \text{if } \tau = +1 \\ 2 + 2\alpha_{\mathcal{M}}^{min(max)} & \text{if } \tau = -1 . \end{cases} \quad (6.29)$$

Without the Same-Spectrum Assumption

If the same-spectrum assumption is not imposed f can be an arbitrary real-valued function. Substituting the definition of $\alpha_{\mathcal{M}}$ into Eq. (6.12) implies

$$(\varepsilon_{\mathcal{A}}^{\min(\max)})^2 = \begin{cases} 1 + \sum_m f(m)^2 \operatorname{Tr}[M_m \rho] - \langle A \rangle_{\rho} \sum_m f(m) \\ \quad - [f(+1) - f(-1)] \alpha_{\mathcal{M}}^{\max(\min)} & \text{if } f(+1) - f(-1) \geq 0 \\ 1 + \sum_m f(m)^2 \operatorname{Tr}[M_m \rho] - \langle A \rangle_{\rho} \sum_m f(m) \\ \quad - [f(+1) - f(-1)] \alpha_{\mathcal{M}}^{\max(\min)} & \text{if } f(+1) - f(-1) \leq 0 . \end{cases} \quad (6.30)$$

The output values $f(m)$ are then chosen to minimize the inaccuracy $\varepsilon_{\mathcal{A}}$ for each fixed configuration of the measurement apparatus (i.e. each fixed \mathcal{M}) [15, 19]

$$f_{\text{opt}}(m = \pm 1) = \frac{\operatorname{Re}[\operatorname{Tr}[M_{\pm} A \rho]]}{\operatorname{Tr}[M_{\pm} \rho]} = \frac{\langle A \rangle_{\rho} \pm \alpha_{\mathcal{M}}}{2 \operatorname{Tr}[M_{\pm} \rho]}. \quad (6.31)$$

However, if $\alpha_{\mathcal{M}}$ is not known precisely, but only bounded by $\alpha_{\mathcal{M}}^{\min}$ and $\alpha_{\mathcal{M}}^{\max}$, it is not possible to unambiguously define $f(m) = f_{\text{opt}}(m)$. Instead, the range of possible values for $\varepsilon_{\mathcal{A}}$ can be optimized by defining either $f(+1) = \frac{\langle A \rangle_{\rho} + \alpha_{\mathcal{M}}^{\min}}{2 \operatorname{Tr}[M_{+} \rho]}$ or $f(+1) = \frac{\langle A \rangle_{\rho} + \alpha_{\mathcal{M}}^{\max}}{2 \operatorname{Tr}[M_{+} \rho]}$, and either $f(-1) = \frac{\langle A \rangle_{\rho} - \alpha_{\mathcal{M}}^{\min}}{2 \operatorname{Tr}[M_{-} \rho]}$ or $f(-1) = \frac{\langle A \rangle_{\rho} - \alpha_{\mathcal{M}}^{\max}}{2 \operatorname{Tr}[M_{-} \rho]}$, whichever combination yields the smallest range of possible values for $\varepsilon_{\mathcal{A}}$ from (6.30).

6.4 Testing Joint-Measurement Uncertainty Relations

6.4.1 Experimental Configuration

To test the joint-measurement uncertainty relation (6.9) and measurement-disturbance relation (6.10) experimentally, we use polarization-encoded qubit to approximate the joint measurement of the incompatible observables $A = \sigma_x$ and $B = \sigma_z$ on the state $|y+\rangle\langle y+| = (\mathbb{1} + \sigma_y)/2$. In this configuration ρ, A, B are all mutually orthogonal, which in the ideal case implies $C_{AB}^2 = 1$. In the case of the weak measurement method with a semi-weak measurement of strength $0 < |\kappa| < 1$, the actual state is perturbed to $\rho = (\mathbb{1} + \sqrt{1 - \kappa^2} \sigma_y)/2$, such that $C_{AB}^2 = 1 - \kappa^2$. In either case, the measurement apparatus that performs the joint approximation of A and B is chosen to implement a projective measurement $\mathcal{M} = \cos \theta \sigma_z + \sin \theta \sigma_x$ onto a direction in the xz-plane of the Bloch sphere, see Fig. 6.11. The outputs $m = \pm 1$ of the measurement \mathcal{M} are used to compute approximations $f(m)$ and $g(m)$ for A and B , respectively. In the case where the same spectrum assumption is imposed, these values are restricted to ± 1 , otherwise they can be optimized to minimize the inaccuracies $\varepsilon_{\mathcal{A}}$ and $\varepsilon_{\mathcal{B}}$.




Figure 6.11: **Illustration of the experimental configuration.** (a) The joint measurement of $A = \sigma_x$ and $B = \sigma_z$ is approximated on the state $\rho = (\mathbb{1} + \sigma_y)/2$, shown in red on the Bloch-sphere. Since A, B and ρ are all pairwise orthogonal the measurements are non-commuting and maximally incompatible on ρ . \mathcal{M} is a projective measurement in the xz -plane of the Bloch-sphere, providing joint approximations \mathcal{A} and \mathcal{B} (visualization highly simplified) for A and B , respectively. (b) Treated as a black box, the measurement device \mathcal{M} performs an approximate joint-measurement of A and B . (c) Experimentally, the measurement \mathcal{M} is implemented as a actual measurement of B after a half-waveplate (which encodes information about A). The half-waveplate rotation encodes information about A in the final measurement, while disturbing the measurement of B . In the case where \mathcal{B} is taken to be the disturbed measurement this corresponds to the measurement-disturbance scenario.

6.4.2 Three-State Method

Recall that the three-state method—in the experimentally robust variant introduced in Sec. 6.3.2—allows the inaccuracies $\varepsilon_{\mathcal{A}}$, and $\varepsilon_{\mathcal{B}}$ to be constrained to small intervals, using the measurement statistics of \mathcal{M} on three different quantum states. In practice, $\rho \simeq |+y\rangle\langle +y|$, and the other states are chosen as $\rho_1 \simeq A\rho A \simeq B\rho B \simeq |-y\rangle\langle -y|$ and $\rho_2 \simeq (\mathbb{1} + A)\rho(\mathbb{1} + A)/\|\cdot\| \simeq |+x\rangle\langle +x|$ (respectively $\rho'_2 \simeq (\mathbb{1} + B)\rho(\mathbb{1} + B)/\|\cdot\| \simeq |+z\rangle\langle +z|$), which is motivated by the original three-state method and tends to produce the strongest constraints on $(\varepsilon_{\mathcal{A}}, \varepsilon_{\mathcal{B}})$. The experimental setup that implements these preparations and measurements is shown in Fig. 6.12.

Using well-characterized state preparations, the weak measurement method produces sharply defined intervals of compatible values for the inaccuracies $\varepsilon_{\mathcal{A}}$ and $\varepsilon_{\mathcal{B}}$, shown in Fig. 6.13. It is crucial to note that these intervals are *not* error regions around a point estimator. Instead the method produces intervals of possible values that are all equally compatible with the experimental data⁵. The effect of statistical errors from the Poissonian counting statistics is to smear out the boundaries, as indicated in Fig. 6.13b. Hence, although there are “error-bars” on the boundary values of each interval, the method will always produce finite intervals.

⁵In order to obtain a point estimate the method must be modified to resemble measurement tomography (using an informationally complete set of states), which, however, would require sophisticated methods to contend with statistical uncertainties and biased estimators, see Chap. 2.




Figure 6.12: **Experimental setup for the three-state method.** (a) Conceptually, in the three state method the measurement inaccuracy $\varepsilon_{\mathcal{A}}$ of the approximate measurement \mathcal{A} is estimated from the measurement statistics of \mathcal{M} on the three quantum states ρ, ρ_1, ρ_2 (or ρ, ρ_1, ρ'_2 for $\varepsilon_{\mathcal{B}}$). (b) In the experiment, quantum states are prepared on single photons at a wavelength of $\lambda=820nm$ from a triggered spontaneous parametric down-conversion (SPDC) source, using a Glan-Taylor polarizer (GT), a quarter- and a half-waveplate (QWP, HWP). A HWP and a polarizer implement the measurement \mathcal{M} . The additional QWP between these elements is used for state tomography, with avalanche photodiodes (APD) used for detection.

As a consequence of the state-dependence of relation (6.9), and (6.10) the bounds they impose are sensitive to the state ρ on which the joint measurement is approximated. In the case presented here, ρ was prepared with a fidelity of $\mathcal{F} = 0.999172(7)$ and purity of $\mathcal{P} = 0.99917(2)$. The decreased purity implies that the relations are not exactly tight anymore, which can in principle be fixed with slightly modified inequalities [24]. More importantly, however, the incompatibility parameter C_{AB} is sensitive to imperfections in the state (not just impurity). The used state here achieved $C_{AB}^2 = 0.99669(3)$, which becomes visible as a small correction of the bound in Fig. 6.13 (solid line), away from the optimal trade-off (dot-dashed line) for the ideal case $C_{AB}^2 = 1$.

6.4.3 Weak-Measurement Method

Recall that the weak-measurement method—in the experimentally robust variant introduced in Sec. 6.3.3—allows the inaccuracies $\varepsilon_{\mathcal{A}}$, and $\varepsilon_{\mathcal{B}}$ to be constrained to small intervals, using the joint statistics of a semi-weak pre-measurement A and B , respectively, and a final measurement of \mathcal{M} . In practice this is achieved using a non-deterministic controlled-NOT (CNOT) gate [33], as discussed in Sec. 6.3.3, and shown in Fig. 6.14

The averaged disturbed state after the weak measurement with an average value of $\kappa = -0.262(4)$ had a reduced purity of $\mathcal{P}=0.964(1)$ (and fidelity of $\mathcal{F}=0.99998(6)$). This corresponds to $C_{AB}^2=0.928(2)$, which results in a significant deviation of the bounds imposed by relation (6.9) and (6.10) from the ideal case, see Fig. 6.15. Moreover, the imperfect process fidelity of the CNOT gate of $\mathcal{F}_p=0.964(1)$ results in larger systematic errors, which appear as displacements of the estimated intervals away from the optimal trade-off. Moreover, Eq. (6.26) becomes increasingly sensitive to experimental imperfections as the measurement strength $|\kappa|$ decreases, which forces a trade-off between small intervals and tight uncertainty bounds.

The intervals obtained from the weak-measurement method are significantly larger than those obtained from the three-state method. One of the two main reasons for this is that the

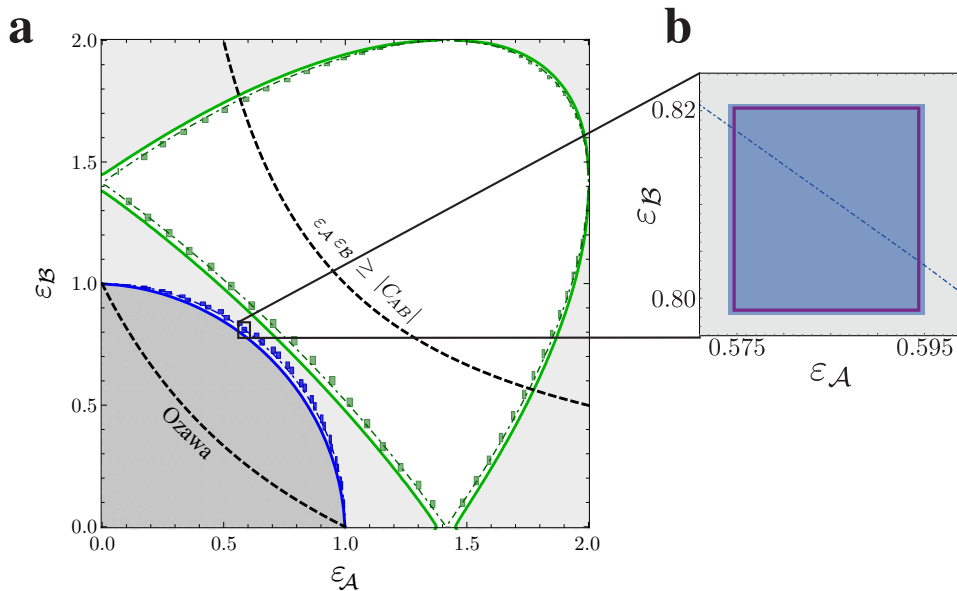


Figure 6.13: **Measurement inaccuracies ε_A and ε_B estimated using the three-state method.** (a) The blue rectangles represent the intervals of compatible values of ε_A and ε_B without the same-spectrum assumption. The solid blue curve corresponds to the bound imposed by relation (6.9), for the experimental values of $\Delta A, \Delta B$ and C_{AB} ; the values below this bound are forbidden by quantum theory. The dot-dashed blue line is the bound imposed by (6.9) for the ideal case $\Delta A = \Delta B = C_{AB} = 1$. The green rectangles and curves represent the corresponding data when the same-spectrum assumption is imposed on both \mathcal{A} and \mathcal{B} , now invoking relation (6.10), which also *upper-bounds* the values of ε_A and ε_B . For comparison, the black dashed curves indicate the bounds imposed by the suboptimal Arthurs-Kelly relation (6.7)—which is violated by our data—and by Ozawa’s relation (6.8)—which is indeed satisfied, but cannot be saturated. The shown intervals include 1σ statistical errors due to Poissonian photon-counting statistics. (b) The blue-shaded region corresponds to the intervals for $(\varepsilon_A, \varepsilon_B)$ including 1σ statistical uncertainties. The solid purple line represents the size of the interval without taking into account statistical uncertainty.

count-rates have to be kept low to avoid errors from higher-order emissions of the downconversion source. As a consequence statistical uncertainty significantly contributes to the size of the intervals for $(\varepsilon_A, \varepsilon_B)$, as shown in Fig. 6.15. Secondly, the weak measurement method only imposes two constraints on the set of possible values of the measurement parameters that are used to estimate $(\varepsilon_A, \varepsilon_B)$. This naturally leads to weaker constraints and thus larger intervals than in the case of the three-state method, which provides three constraints. If, however, instead of treating them separately, one takes into account the experimental data from both the weak measurement of A and of B , then the size of the obtained intervals can be significantly reduced. The corresponding smaller intervals, shown as dark rectangles in Fig. 6.15, are then typically dominated by statistical errors.




Figure 6.14: **Experimental setup for the weak-measurement method.** (a) Conceptually, in the weak-measurement method the inaccuracy ε_A of the approximate measurement \mathcal{A} is estimated from the joint statistics of a weak pre-measurement of A , and the strong approximating measurement \mathcal{M} . (b) Experimentally, the weak measurement is implemented using a non-deterministic controlled-not gate with an appropriately chosen meter qubit, which is subsequently measured in σ_z . The state ρ after the weak measurement (slightly mixed) is then subjected to the strong measurement \mathcal{M} , implemented as in Fig. 6.12.

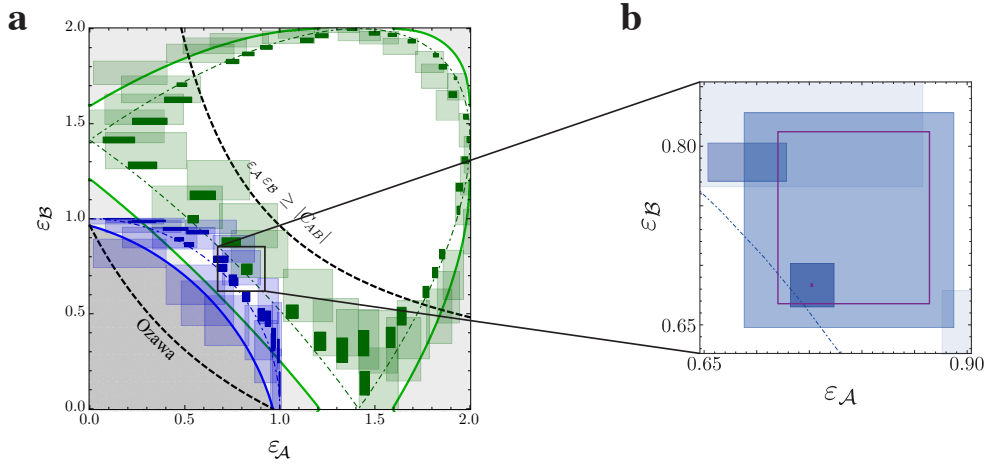


Figure 6.15: **Measurement inaccuracies, ε_A vs. ε_B estimated using the weak-measurement method.** (a) The results are presented as in Fig. 6.13. The darker rectangles represent the smaller intervals of compatible values that can be obtained by using the data from *both* semi-weak measurements. Due to the reduced purity of the input state ρ on which the joint measurement is approximated the bounds imposed by the relations (6.9), and (6.10) for the experimental values of ΔA , ΔB and C_{AB} (solid lines) deviate significantly from the ideal bounds (dot-dashed lines). (b) Contribution of statistical errors to the size of the intervals, as in Fig. 6.13b. In contrast to the three-state method, the intervals are dominated by statistical errors.

6.5 Discussion and Outlook

Although the alleged violations of “Heisenberg’s uncertainty relation” [34–37] ultimately turned out to be perfectly compatible with the uncertainty principle, they highlighted that it is far from a solved problem. The experimental demonstration that the Arthurs-Kelly reformulation

of the famous Robertson relation cannot be used to quantify joint-measurement uncertainty has led to a lot of debate and disagreement. Much of the confusion in fact can be traced back to the exact definitions of measurement inaccuracy and disturbance used by various authors and the different problems that are addressed [5, 15, 17, 19, 38–40].

Specifically, one might be interested in quantifying the worst-case performance of a joint-measurement device for two incompatible observables over all possible quantum states. In this case a state-independent uncertainty relation would be desirable. However, for any fixed quantum state it is in general possible to do much better than this state-independent bound. Branciard’s joint-measurement relation (6.9) and measurement-disturbance relation (6.10) address the latter problem. For pure states they quantify the optimal trade-off between the inaccuracies of the approximate measurements \mathcal{A} and \mathcal{B} and thus describe exactly what is and what is not possible in quantum mechanics. In the case of mixed states these relations are not tight anymore, but there are ways to generalize them appropriately [20, 24].

Curiously, the recent discussion about the measurement uncertainty principle revealed that the problem of preparation uncertainty was not completely solved either. As pointed out in Ref [5], the Robertson relation provides a state-dependent lower bound, which often provides only a rather weak constraint, and can even vanish, although both uncertainties are non-zero. Moreover, the state-dependence itself means that A, B and ρ must be specified to evaluate the relation, which leaves no freedom for ΔA and ΔB . Hence, such a state-dependent relation cannot be used to answer questions about the optimal trade-off between the uncertainties of A and B for any quantum state. In the case of state-dependent measurement uncertainty relations this is not the case, since $\varepsilon_{\mathcal{A}}$ and $\varepsilon_{\mathcal{B}}$ depend on the measurement device as well, and not only on A, B and ρ . For qubits (or more generally finite-dimensional systems) geometric arguments are very powerful for finding uncertainty relations, and have already been used to fully characterize preparation uncertainty for qubits. For the joint-measurement and measurement-disturbance case some progress has been made [40], but the general state-independent case is still open [14, 40].

Quantum information theory inspired an alternative approach to uncertainty relations based on bounding the information content of a measurement, quantified by the entropy. This has the interesting operational interpretation in the sense that the more predictable the results of one observable are, the more random must the results of another, incompatible observable be. In the preparation uncertainty scenario, the relations of Ref. [5] directly translate into tight state-independent entropic uncertainty relations. For joint measurements this approach has recently been used to derive both state-independent [41] and state-dependent [42, 43] entropic joint-measurement uncertainty relations. It would be interesting to study the connections between Ozawa’s mean errors and the entropic approach in more detail.

References

- [1] Ringbauer, M., Biggerstaff, D. N., Broome, M. A., Fedrizzi, A., Branciard, C. & White, A. G. Experimental Joint Quantum Measurements with Minimum Uncertainty. *Phys. Rev. Lett.* **112**, 020401 (2014).
- [2] Heisenberg, W. Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik. *Z. Phys.* **43**, 172–198 (1927).
- [3] Kennard, E. H. Zur Quantenmechanik einfacher Bewegungstypen. *Z. Phys.* **44**, 326–352 (1927).
- [4] Heinosaari, T., Reitzner, D. & Stano, P. Notes on Joint Measurability of Quantum Observables. *Found. Phys.* **38**, 1133–1147 (2008).
- [5] Abbott, A., Alzieu, P.-L., Hall, M. & Branciard, C. Tight State-Independent Uncertainty Relations for Qubits. *Mathematics* **4** (2016).
- [6] Busch, P. & Schmidt, H.-J. Coexistence of qubit effects. *Quantum Inf. Process.* **9**, 143–169 (2008).
- [7] Stano, P., Reitzner, D. & Heinosaari, T. Coexistence of qubit effects. *Phys. Rev. A* **78**, 012315 (2008).
- [8] Yu, S., Liu, N.-l., Li, L. & Oh, C. H. Joint measurement of two unsharp observables of a qubit. *Phys. Rev. A* **81**, 062116 (2010).
- [9] Quintino, M. T., Vértesi, T. & Brunner, N. Joint Measurability, Einstein-Podolsky-Rosen Steering, and Bell Nonlocality. *Phys. Rev. Lett.* **113**, 160402 (2014).
- [10] Uola, R., Budroni, C., Gühne, O. & Pellonpää, J.-P. One-to-One Mapping between Steering and Joint Measurability Problems. *Phys. Rev. Lett.* **115**, 230402 (2015).
- [11] Wolf, M. M., Perez-garcia, D. & Fernandez, C. Measurements incompatible in Quantum Theory cannot be measured jointly in any other local theory. *Phys. Rev. Lett.* **103**, 230402 (2009).
- [12] Quintino, M. T., Bowles, J., Hirsch, F. & Brunner, N. Incompatible quantum measurements admitting a local-hidden-variable model. *Phys. Rev. A* **93**, 052115 (2016).
- [13] Robertson, H. The Uncertainty Principle. *Phys. Rev.* **34**, 163–164 (1929).
- [14] Branciard, C. private communication.

- [15] Hall, M. Prior information: How to circumvent the standard joint-measurement uncertainty relation. *Phys. Rev. A* **69**, 052113 (2004).
- [16] Busch, P., Lahti, P. & Werner, R. F. Heisenberg uncertainty for qubit measurements. *Phys. Rev. A* **89**, 012129 (2014).
- [17] Ozawa, M. Universally valid reformulation of the Heisenberg uncertainty principle on noise and disturbance in measurement. *Phys. Rev. A* **67**, 042105 (2003).
- [18] Ozawa, M. Uncertainty relations for joint measurements of noncommuting observables. *Phys. Lett. A* **320**, 367–374 (2004).
- [19] Branciard, C. Error-tradeoff and error-disturbance relations for incompatible quantum measurements. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 6742 (2013).
- [20] Ozawa, M. Error-disturbance relations in mixed states. *arXiv:1404.3388* (2014).
- [21] Arthurs, E. & Kelly Jr, J. L. On the simultaneous measurement of a pair of conjugate observables. *Bell Syst. Tech. J* **44**, 725–729 (1965).
- [22] Arthurs, E. & Goodman, M. Quantum Correlations: A Generalized Heisenberg Uncertainty Relation. *Phys. Rev. Lett.* **60**, 2447–2449 (1988).
- [23] Ozawa, M. Position measuring interactions and the Heisenberg uncertainty principle. *Phys. Lett. A* **299**, 1–7 (2002).
- [24] Branciard, C. Deriving tight error-trade-off relations for approximate joint measurements of incompatible quantum observables. *Phys. Rev. A* **89**, 022124 (2014).
- [25] Rozema, L. A., Mahler, D. H., Hayat, A. & Steinberg, A. M. A Note on Different Definitions of Momentum Disturbance. *arXiv:1307.3604* (2013).
- [26] Werner, R. F. The uncertainty relation for joint measurement of position and momentum. *Quantum Info. Comput.* **4**, 546–562 (2004).
- [27] Ozawa, M. Uncertainty relations for noise and disturbance in generalized quantum measurements. *Ann. Phys.* **311**, 350–416 (2004).
- [28] Lund, A. P. & Wiseman, H. M. Measuring measurement–disturbance relationships with weak values. *New J. Phys.* **12**, 093011 (2010).
- [29] Wiseman, H. Directly observing momentum transfer in twin-slit “which-way” experiments. *Phys. Lett. A* **311**, 285 – 291 (2003).
- [30] Weston, M. M., Hall, M. J. W., Palsson, M. S., Wiseman, H. M. & Pryde, G. J. Experimental Test of Universal Complementarity Relations. *Phys. Rev. Lett.* **110**, 220402 (2013).

- [31] Nielsen, M. & Chuang, I. *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).
- [32] Pryde, G. J., O’Brien, J. L., White, A. G., Ralph, T. C. & Wiseman, H. M. Measurement of quantum weak values of photon polarization. *Phys. Rev. Lett.* **94**, 220405 (2005).
- [33] Langford, N. K., Weinhold, T. J., Prevedel, R., Resch, K. J., Gilchrist, A., O’Brien, J. L., Pryde, G. J. & White, A. G. Demonstration of a simple entangling optical gate and its use in bell-state analysis. *Phys. Rev. Lett.* **95**, 210504 (2005).
- [34] Rozema, L. A., Darabi, A., Mahler, D. H., Hayat, A., Soudagar, Y. & Steinberg, A. M. Violation of Heisenberg’s Measurement-Disturbance Relationship by Weak Measurements. *Phys. Rev. Lett.* **109**, 100404 (2012).
- [35] Erhart, J., Sponar, S., Sulyok, G., Badurek, G., Ozawa, M. & Hasegawa, Y. Experimental demonstration of a universally valid error–disturbance uncertainty relation in spin measurements. *Nat. Phys.* **8**, 185–189 (2012).
- [36] Baek, S.-Y., Kaneda, F., Ozawa, M. & Edamatsu, K. Experimental violation and reformulation of the Heisenberg’s error-disturbance uncertainty relation. *Sci. Rep.* **3**, 2221 (2013).
- [37] Sulyok, G., Sponar, S., Erhart, J., Badurek, G., Ozawa, M. & Hasegawa, Y. Violation of Heisenberg’s error-disturbance uncertainty relation in neutron spin measurements. *Phys. Rev. A* **88** (2013).
- [38] Busch, P., Lahti, P. & Werner, R. F. Proof of Heisenberg’s Error-Disturbance Relation. *Phys. Rev. Lett.* **111**, 160405 (2013).
- [39] Ozawa, M. Disproving Heisenberg’s error-disturbance relation. *arXiv:1308.3540* (2013).
- [40] Busch, P., Lahti, P. & Werner, R. F. Colloquium : Quantum root-mean-square error and measurement uncertainty relations. *Rev. Mod. Phys.* **86**, 1261–1281 (2014).
- [41] Buscemi, F., Hall, M. J. W., Ozawa, M. & Wilde, M. M. Noise and Disturbance in Quantum Measurements: An Information-Theoretic Approach. *Phys. Rev. Lett.* **112**, 050401 (2014).
- [42] Kaniewski, J., Tomamichel, M. & Wehner, S. Entropic uncertainty from effective anticommutators. *Phys. Rev. A* **90**, 012332 (2014).
- [43] Coles, P. J. & Furrer, F. State-dependent approach to entropic measurement–disturbance relations. *Phys. Lett. A* **379**, 105–112 (2015).

CHAPTER 7

Conclusion and Outlook

THE work presented in this thesis addresses some of the key foundational questions on the path towards a better understanding of quantum mechanics. In summary, we have presented an experiment that rules out interpretations of quantum mechanics which consider the wavefunction as purely a representation of incomplete knowledge of an underlying objective reality. We have demonstrated the experimental application of causal modelling techniques to quantum correlations, showing that not even nonlocal causal influences from one measurement outcome to the other can explain quantum correlations. Finally, we have experimentally studied a largely overlooked aspect of Heisenberg's uncertainty principle, namely the constraints it imposes on the joint measurement of two incompatible observables. All three aspects are currently very active fields of research in both theory and experiment and contribute significantly to our understanding of the physical content of quantum mechanics. Beyond the central scope of the thesis we present three additional works in the appendix, which investigate the limits of quantum theory. The first focuses on an experimental simulation of super-quantum correlations and information causality as a candidate principle for defining the boundary between quantum mechanics and more powerful theories. Introducing relativistic quantum information, we have simulated a model of closed timelike curves, giving insight into how quantum systems may behave under extreme gravitational conditions. Finally, we have introduced and demonstrated a method for preparing non-Gaussian states of motion of a macroscopic mechanical oscillator, which will be a key ingredient in future tests of the limits of quantum mechanics in the macroscopic regime. In the following we will discuss in more detail the main results, their relevance and connection, as well as interesting directions for future research.

The nature of the quantum wavefunction is the central issue in the debate over the interpretation of quantum mechanics, which is almost as old as the theory itself. While probably everyone working in quantum mechanics thought about this question at some point, it was generally considered a topic of philosophical interest with no practical impact and not open to empirical exploration. This changed with the development of the ontological models framework that formalized the problem in mathematical rigour and enabled to the seminal PBR theorem. This, and many subsequent results showed that interpretations wherein the wavefunction is a representation of incomplete knowledge of an underlying reality are untenable because they cannot reproduce one of the crucial features of quantum theory, the imperfect distinguishability of non-orthogonal quantum states. This strong result, however, only holds under additional

assumptions on the structure of the ontic state space. In chapter 4 we present an experiment that follows a different path, inspired by the PBR result. Without making additional assumptions, one cannot completely rule out realist ψ -epistemic models, but can still constrain their explanatory power. Specifically we showed that such models can at most explain $\sim 70\%$ of the quantum overlap observed in the experiment, which required the preparation and measurement of four-dimensional quantum systems with unprecedented accuracy.

Our results demonstrate that the epistemic overlap is insufficient for explaining explain quantum indistinguishable, yet it remains unclear whether it plays any role at all. A common feature of all known experimental protocols is that they establish more stringent bounds on the overlap ratio only as the amount of quantum overlap goes to zero. In other words the closer we are to showing that a ψ -epistemic model cannot explain anything, the less there is to explain. It is an important open question whether similar bounds can be achieved with a fixed inner product. Experimentally, imposing more stringent bounds on the overlap ratio will require improvements in the preparation and measurement precision for higher-dimensional quantum systems. Finally, for practical reasons our current results rely on a fair sampling assumption and repeating the experiment in an architecture which achieves the required (almost perfect) detection efficiency will allow to also test unfair-sampled maximally ψ -epistemic models.

One of the most important implications of ruling out maximally ψ -epistemic models, however, is that it demonstrates that the discussion about interpretations of quantum mechanics is open to empirical exploration. This opens the route towards tests of other interpretations, which can be roughly grouped into ψ -ontic models (within or outside the ontological models framework), operational interpretations, and more exotic ontologies. The class of ψ -ontic interpretations hosts many popular candidates such as many-worlds and Bohmian Mechanics. A few of these, most notably, collapse models, are actually alternative theories which are known to make predictions that differ from quantum mechanics and will be tested with future macroscopic experiments, such as the optomechanical system discussed in Appendix C. More generally, there are arguments against larger classes of realist [1] and single-world [2] interpretations, based on heating due to information erasure and self-consistency, respectively. However, also these results hinge on a number of questionable assumptions. With regards to operational interpretations the experimental prospects are largely unknown, but gedanken experiments such as that of Wigner’s friend can go a long way to studying these interpretations. In the end, it is quite likely that a consistent interpretation of quantum mechanics will feature aspects of a range of the current interpretations [3]

The study of ontological models (or hidden variable models) actually has a much longer history than the ontological models framework. It was the work of John Bell, which showed that these models can indeed be constrained experimentally, which initiated half a century of “experimental metaphysics” [4], culminating in a series of loophole-free Bell-inequality violations in 2015 [5–7]. Today Bell-inequality violations are at the heart of device-independent entanglement verification and, in principle, unbreakable quantum cryptography, which highlights their

practical importance beyond the metaphysics. These applications, of course came long after Bell, who was concerned with the causal structure of quantum mechanics. This is a central aspect for realist interpretations, which, in contrast to operational interpretations, are actually trying to *explain* observed correlations.

Roughly speaking, Bell-inequality violations show that quantum correlations cannot be explained in terms of classical cause-and-effect relations under the assumption of relativistic causality. In chapter 5 we discuss how the mathematical theory of causal modeling provides a natural framework for the study of non-classical correlations. Within this framework we introduce a decomposition of Bell's local causality assumption into the causally motivated assumptions of causal parameter independence and causal outcome independence. We then show experimentally that giving up causal outcome independence, i.e. allowing for a causal influence from one measurement outcome to the other, is not enough to explain the observed correlations causally. Since causal models are formulated without reference to a background spacetime, the tested causal influence could be sub- or super-luminal, or even to the past, as long as it does not create causal loops. These results rule out a large class of nonlocal causal models of which Bell's models are a special case and highlight the strength of the causal modeling approach for a systematic study of the various assumptions behind causal explanations of quantum correlations. It would be interesting to apply similar methods to other scenarios and relaxations of other causal assumptions to study the causal structure of quantum mechanics in more detail.

Another interesting aspect of our result is that the considered set of joint probability distributions is only bounded by non-trivial Bell-type inequalities in a scenario with three settings and two outcomes. This suggests that setting information is more important to simulating quantum correlations and could in the standard two-setting, two-outcome case be encoded in the measurement outcome. It would be interesting to study this feature in more complicated scenarios and in cases with a very large number of settings and outcomes. However, this leads to an exponential growth of the dimension of the corresponding probability polytopes, which makes their study intractable with standard methods. Finally, all causal model that can reproduce Bell correlations must be fine-tuned in the sense that the causal parameters must be chosen carefully such that whatever causal link the model adds compared to Bell-local models does not lead to superluminal signaling [8]. This brings in an interesting alternative perspective, and although not every fine-tuned model is unstable, this suggests that we might have to reconsider our notion of causality in the quantum world.

Besides reality and causality, one of the most fundamental aspects of quantum mechanics is Heisenberg's uncertainty principle, which provides physical meaning to the notion of incompatible observables. It is well-known that no quantum state can simultaneously have arbitrarily well-defined position and momentum. Contrary to common belief, however, Heisenberg's uncertainty principle encompasses much more than that and, in fact, contains three distinct statements about nature. These, respectively, quantify how well two incompatible observables

can be prepared on a quantum system, how well they can be jointly measured, and how much a measurement of one disturbs the measurement of the other. Heisenberg himself has not quantified any of these principles and until about a decade ago only preparation uncertainty had a quantitative uncertainty relation, although it was often incorrectly used for the other two cases. In chapter 6 we present the first experimental test of tight joint-measurement uncertainty relations derived in Ref. [9]. These relations quantify the optimal tradeoff in measurement accuracy for the joint approximation of two incompatible observables on a quantum state. The measurement-disturbance case is in fact closely related and quantified by similar relations, and in both cases we achieved close-to saturation of these bounds using high-precision photonics. Practically, understanding the optimal error-tradeoff is crucial for developing high-precision joint measurements, or can be exploited for trading off accuracy in one measurement for uncertainty in another.

Curiously, despite its fundamental and practical importance, the quantitative aspect of Heisenberg’s uncertainty principle is still poorly understood, and even worse, is often treated as a solved problem. In fact, only the preparation uncertainty principle has received considerable attention and has been quantified by a state-dependent relation. However, such relations are only meaningful in restricted cases, as pointed out in Ref. [10], who provided preparation uncertainty relations for the general case. In the measurement-uncertainty case the situation is a bit more difficult, and both, state-dependent, and state-independent relations are meaningful, but for different purposes. State-dependent relations, tested here, quantify the optimal tradeoff in measurement accuracy for a given quantum state, while state-independent relations would quantify the worst-case performance of a measurement device. A general family of uncertainty relations that is optimal for the latter case is still an open problem. Experimentally, it would be interesting to perform a detailed analysis of state-independent preparation uncertainty relations and also compare the state-dependent and state-independent cases in the measurement uncertainty scenario.

The work included in Appendices A-C go beyond fundamental questions within quantum mechanics and explore the boundaries of the theory. Despite all their success in their respective regimes, our current two major physical theories, quantum mechanics and general relativity, are fundamentally incompatible and can thus not be the end of the story. A good understanding not only of the physical content of quantum theory, but also of its limits will be key in the development of a theory that bridges the gap and possibly supersedes quantum mechanics and relativity.

One of the curious features in this respect is that quantum correlations, despite being stronger than any classical correlations, are not as strong as relativity would permit. Moreover, there are theories with stronger correlations than quantum mechanics that are still consistent with relativity and display many features that are considered characteristic for quantum mechanics. In Sec. 3.2.4 and Appendix A we demonstrate how non-unitary effects, specifically polarization-dependent loss, can lead to apparent correlations that violate the CHSH inequality.

ity beyond the quantum bound. We use these simulated correlations to study information causality, which is a promising candidate for a physical principle to recover quantum correlations. Information causality strengthens signal locality and implies that the information that Bob can gain about an unknown to him data set of Alice, should be limited to the amount of information communicated to him by Alice. This principle can explain the strength of quantum correlations in a number of scenarios, but not always. Exploring this and other physical principles, or a combination thereof further would be of great interest for building quantum theory on a framework of physical principles, rather than mathematical axioms.

In Appendix B we simulate the behaviour of a quantum system that traverses a closed time-like curve and thus effectively travels back in time to interact with itself. Such closed timelike curves are puzzling phenomena that appear to be compatible with the equations of general relativity under extreme conditions. In the classical case these are usually disregarded because they lead to paradoxes. However, in the quantum case these paradoxes are resolved due to the larger state-space, which, in particular, allows for mixed states. We simulate the behaviour of a quantum system in such a situation, revealing effects such as perfect distinguishability of quantum states and nonlinear evolution. Curiously the latter phenomenon makes it possible to distinguish between proper mixtures of quantum states, where different pure states are combined to a mixed state, and improper mixtures, where part of a higher-dimensional system is traced out, although in standard quantum mechanics there is no observable difference between the two. Exploring the behaviour of quantum systems under such extreme conditions can provide important clues to where quantum mechanics clashes with relativity, which is one of the key challenges of modern physics.

Finally, an important open question is how large one can make a system before quantum mechanics breaks down. On the one hand, quantum theory should be universal and apply to any physical system, on the other hand quantum effects are not observed at the everyday, macroscopic scale. A prominent suggestion is that quantum mechanics is in fact just a linear approximation to an alternative theory, which features a stochastic “collapse” term in the Schrödinger evolution. Such theories are compatible with all present observations, but imply that quantum superposition states of macroscopic objects cannot be maintained over extended periods of time. A very promising platform to test such theories is the emerging field of optomechanics, which studies the behaviour of mechanical resonators, interfaced through single photons. However, this field has for a long time been facing two major roadblocks. The mechanical systems must be cooled very close to their motional ground state, and the interaction with the photons must be very strong. In Appendix C we develop and demonstrate a novel technique for preparing quantum states of motion of a mechanical resonator, which sidesteps these major challenges. Inspired by the KLM scheme for preparing non-classical states of light using measurement induced nonlinearities, our method relies on photon counting in an optical interferometer of photons reflected from the mechanical resonator. We showed theoretically that this method produces non-classical states with arbitrary low single-photon coupling strength and for arbitrary thermal occupation of the system (although could systems work better, of

course). We demonstrate one of the main features of the method, that it generates interference fringes in the position quadrature distribution, on a mechanical resonator in a thermal regime. This method opens a feasible route for tests of collapse models of quantum mechanics and the exploration of macroscopic quantum phenomena.

Quantum mechanics is our most successful physical theory, yet we do not understand what it says about the physical world. In this thesis I have aimed to make significant contributions to developing such a physical understanding. My work demonstrates that the quantum wavefunction, the central object of the theory, cannot be interpreted as a representation of imperfect knowledge about an underlying reality. This is part of a decade-old problem that was long believed to not be open for empirical tests. I explored the causal structure of quantum mechanics, generalizing Bell's seminal theorem to rule out a large class of nonlocal causal models and pave the way for further exploration of quantum causality in more complicated scenarios. This can have important implications for public-outcome quantum cryptography schemes. I have also studied the role of quantum uncertainty, despite being widely regarded a solved problem, has many more facets expected. The measurement uncertainty principle is crucial for high-precision joint measurements and quantum metrology. Finally, my results push the boundaries of quantum mechanics by exploring the behaviour of quantum systems under extreme relativistic conditions, post-quantum theories and demonstrating the first generation of interference in the position quadrature distribution of a macroscopic mechanical resonator. These contributions address a range of key questions within quantum foundations and set a number of directions for future research towards a better understanding of our best physical theory.

References

- [1] Cabello, A., Gu, M., Gühne, O., Larsson, J.-Å. & Wiesner, K. Thermodynamical cost of some interpretations of quantum theory. *arXiv:1509.03641* (2015).
- [2] Frauchiger, D. & Renner, R. Single-world interpretations of quantum theory cannot be self-consistent. *arXiv:1604.07422* (2016).
- [3] Brukner, C. On the quantum measurement problem. *arXiv:1507.05255* (2015).
- [4] Shimony, A. Search for a worldview which can accommodate our knowledge of microphysics. In *Philosophical Consequences of Quantum Theory*.
- [5] Hensen, B. *et al.* Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature* **526**, 682–686 (2015).
- [6] Giustina, M. *et al.* Significant-Loophole-Free Test of Bell’s Theorem with Entangled Photons. *Phys. Rev. Lett.* **115**, 250401 (2015).
- [7] Shalm, L. K. *et al.* Strong Loophole-Free Test of Local Realism. *Phys. Rev. Lett.* **115**, 250402 (2015).
- [8] Wood, C. J. & Spekkens, R. W. The lesson of causal discovery algorithms for quantum correlations: Causal explanations of Bell-inequality violations require fine-tuning. *New J. Phys.* **17**, 033002 (2015).
- [9] Branciard, C. Error-tradeoff and error-disturbance relations for incompatible quantum measurements. *Proc. Nat. Acad. Sci. USA* **110**, 6742–6747 (2013).
- [10] Abbott, A., Alzieu, P.-L., Hall, M. & Branciard, C. Tight State-Independent Uncertainty Relations for Qubits. *Mathematics* **4**, 8 (2016).

APPENDIX A

Information Causality in the Quantum and Post-Quantum Regime

This appendix reproduces a paper of which I am the lead author and which was first published in *Scientific Reports* **4**, 6955 (2014).

Information Causality in the Quantum and Post-Quantum Regime

Martin Ringbauer^{1,2}, Alessandro Fedrizzi^{1,2}, Dominic W. Berry³ and Andrew G. White^{1,2}

¹*Centre for Engineered Quantum Systems, ²Centre for Quantum Computation and Communication Technology, School of Mathematics and Physics, University of Queensland, Brisbane, QLD 4072, Australia.*

³*Department of Physics and Astronomy, Macquarie University, Sydney, NSW 2109, Australia*

Quantum correlations can be stronger than anything achieved by classical systems, yet they are not reaching the limit imposed by relativity. The principle of information causality offers a possible explanation for why the world is quantum and why there appear to be no even stronger correlations. Generalizing the no-signaling condition it suggests that the amount of accessible information must not be larger than the amount of transmitted information. Here we study this principle experimentally in the classical, quantum and post-quantum regimes. We simulate correlations that are stronger than allowed by quantum mechanics by exploiting the effect of polarization-dependent loss in a photonic Bell-test experiment. Our method also applies to other fundamental principles and our results highlight the special importance of anisotropic regions of the no-signalling polytope in the study of fundamental principles.

I. INTRODUCTION

Quantum mechanics is one in a large class of theories which are consistent with relativity in the sense that they do not allow signals to be sent faster than the speed of light. Many of these theories exhibit strong non-local correlations between distant particles that cannot be explained by the properties of the individual particles alone. Surprisingly, quantum mechanics is not the most non-local among them, which raises the question about the physical principle that singles out quantum mechanics and sets the limit on the possible strength of correlations in nature.

Here we experimentally address this fundamental question by testing the principle of information causality in the classical, quantum and post-quantum regime. While the no-signaling principle limits the speed with which distant parties can communicate, information causality states that the accessible information cannot be more than the information content of a communicated message, no matter what other shared resources are used. Both classical and quantum mechanics satisfy this principle, while it is violated by most post-quantum theories [1].

We experimentally emulate correlations of various strengths from classical to almost maximally non-local and demonstrate a violation of the principle of information causality in the case where the simulated correlations are beyond the quantum regime. Apparent super-quantum correlations are, in our approach, a consequence of the non-unitary evolution of quantum states when subjected to polarization-dependent loss with post-selection [2]. For moderate loss, we find that initially entangled states can result in super-quantum correlations, while unentangled states still appear classical. For higher loss on the other hand we observe super-quantum correlations even for classical input states.

II. THEORETICAL FRAMEWORK

No-signaling resources can formally be treated as pairs of black boxes shared between arbitrarily separated Alice and Bob [3], see Fig. 1a). Each box has a single input and output and the correlation between them is only restricted by the no-signaling principle. This means that the local outcome only depends on the local input, such that Alice cannot learn anything about Bob's input from only her output.

A typical quantum example of such a resource is a pair of entangled particles, shared between Alice and Bob, where inputs correspond to measurement settings and outputs to measurement outcomes. Since the work of John Bell—and numerous subsequent confirming experiments—it is now widely accepted that these particles exhibit non-local correlations, which have no classical explanation. Under the no-signaling constraint alone, however, there are even stronger non-local correlations than quantum entanglement [4]. The maximum that is compatible with relativity is achieved by the so-called Popescu-Rohrlich (PR)-box [4], characterized by perfect correlations of the form $A \oplus B = ab$, between Alice's and Bob's inputs a and b and outputs A and B , respectively. Here \oplus denotes addition modulo 2, equivalent to the logical XOR, where $A \oplus B = 0$ when $A = B$ and 1 otherwise.

A convenient operational way of quantifying non-locality is the Clauser-Horne-Shimony-Holt (CHSH) inequality [5]. This experimentally testable reformulation of Bell's inequality is satisfied by any correlation that can be described by a local hidden variable model. Such models are a description of correlations that can arise in classical systems, but cannot describe non-local correlations obtained from e.g. entangled quantum states. Written in terms of correlations of the form $A \oplus B = ab$

the inequality takes the form

$$S = \sum_{a=0}^1 \sum_{b=0}^1 P(A \oplus B = ab \mid a, b) \leq 3. \quad (1)$$

Here $P(A \oplus B = ab \mid a, b)$ denotes the probability for obtaining outputs A, B , which satisfy $A \oplus B = ab$ given the inputs a for Alice and b for Bob. While this inequality is satisfied by any classical correlations, it can be violated in the quantum case. This violation, however, is bounded to a value of $2 + \sqrt{2} \approx 3.41$, known as Tsirelson's bound [6]. Note, that inequality (1) is presented here in a slightly different form than conventionally [5], where the classical bound is 2 and Tsirelson's bound is $2\sqrt{2}$. They are, however, linearly related and the difference is a simple rescaling of S .

Despite being a simple consequence of the mathematical formalism of quantum mechanics, it is unclear what the physical motivation is for this seemingly sub-optimal limit on the strength of quantum correlations. In fact even the algebraic maximum $S = 4$ can be achieved (by the PR-box) without violating the no-signaling principle.

III. THE PRINCIPLE OF INFORMATION CAUSALITY

This principle is physically motivated by the fact that, according to special relativity, faster-than-light information transfer would allow information to be sent backwards in time and thus violate causality. Nevertheless, it does not explain why super-quantum correlations such as the PR-box are incompatible with quantum mechanics and seem not to exist in nature. A possible explanation is offered by the principle of information causality—a generalization of no-signaling—which states that there cannot be more information available than was transmitted [7].

This can be understood on the basis of the following elementary information-theoretic protocol: Bob tries to gain information from a set of data that is only known to Alice. The parties are allowed to use an arbitrary amount of shared no-signaling resources, but may not communicate more than m classical bits. In this case, the information causality principle states that the amount of information accessible to Bob should be limited to m classical bits [7].

In the simplest instance, Alice has a set of two bits $\{a_0, a_1\}$ and Bob wants to guess one of them, which we denote a_b [8], see Fig. 1a). Alice and Bob then input $a_0 \oplus a_1$ and b into their respective black box and obtain outputs A and B . From this output Alice computes an $m = 1$ -bit message $M = A \oplus a_0$ and sends it to Bob, who calculates his guess for Alice's bit as $G = M \oplus B = a_0 \oplus A \oplus B$. In the case of a shared PR-box, Bob can guess either one of Alice's bits perfectly, since in that case $A \oplus B = ab$ and thus $G = a_0 \oplus b(a_0 \oplus a_1)$.

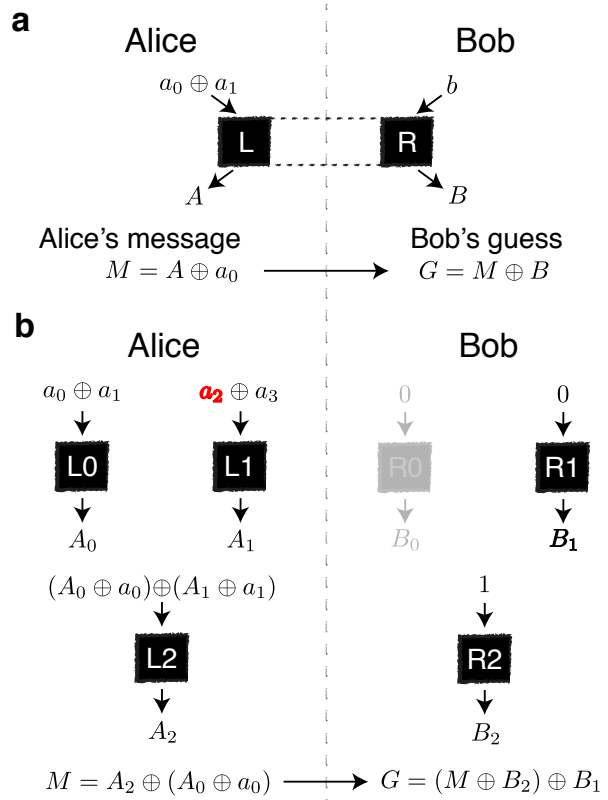


FIG. 1. **Illustration of the information causality protocol.** **a)** A general no-signaling resource is given by a space-separated (indicated by the dashed line) pair of black boxes producing local outputs A and B for Alice and Bob, when they input a and b , respectively. In the case of a PR-box the outputs of the left (L) and right (R) box would be perfectly correlated according to $A \oplus B = ab$. The inputs and outputs depicted here correspond to the simplest instance of the information causality protocol. **b)** Example of the multilevel information causality protocol for $n = 2$. Alice has a list of N bits a_i and Bob tries to guess the bit a_3 (shown in bold, red) using $N-1=3$ pairs of shared black boxes on $n=2$ levels (corresponding boxes labeled $L0/R0, L1/R1, L2/R2$). Bob's inputs b_i and choice of boxes are determined by the binary decomposition $b = \sum_{k=0}^{n-1} b_k 2^k$. From his outputs B_1, B_2 and Alice's 1-bit message M Bob computes a final guess G for Alice's bit a_b . Note that Bob only needs to use one box on each level and ignores the outputs of all the other boxes. Hence, his input to these boxes can be arbitrary and in the experiment we chose to use the same input for all boxes on one level.

In the more general case considered here, Alice has a dataset $\{a_0, \dots, a_{N-1}\}$ of $N = 2^n$ bits and Bob wants to guess the bit with index $b = \sum_{k=0}^{n-1} b_k 2^k$. As discussed in Ref. [7], Alice and Bob can achieve this task by using a nested version of the protocol outlined above, with $N-1$ black boxes on n levels and 1 bit of classical communication.

The protocol is illustrated in Fig. 1b) for the case $n = 2$. From every output Alice computes a temporary

message $M_{k,i}$, where k denotes the level and i the number of the box on that level. Since she is only allowed 1 bit of communication, she uses these temporary messages as the inputs for the boxes on the next-lower level and only sends the final message to Bob. Depending on b_n Bob then decodes either $M_{n-1,1}$ or $M_{n-1,2}$ and then moves on to the next-higher level until he reaches the bit of interest.

Bob's success can then be quantified by

$$I = \sum_{k=0}^{N-1} I(a_k : G | b = k), \quad (2)$$

where $I(a_k : G | b = k)$ is the Shannon mutual information between the k 'th bit of Alice's list and Bob's guess for it [7]. This quantity can further be bounded as

$$I \geq \sum_{k=0}^{N-1} 1 - h(P_k), \quad (3)$$

where $h(P_k)$ is the binary entropy of the success probability P_k for guessing the k 'th bit.

IV. EXPERIMENTAL IMPLEMENTATION

Experimentally, we generate apparent super-quantum correlations based on the effect of polarization-dependent loss in a post-selected Bell-test experiment [2], see Fig. 2a). We use photon pairs created from a continuous-wave pumped spontaneous parametric down-conversion source in a polarization Sagnac design [9, 10], as illustrated in Fig. 2b). Using this approach we obtain photon pairs with very high efficiency and in a continuously tunable fashion that enables us to produce any bipartite quantum state [11].

In particular, we used the maximally entangled state $|\psi^+\rangle = (|H\rangle|V\rangle + |V\rangle|H\rangle)/\sqrt{2}$ as the initial state, where $|H/V\rangle$ represent horizontal and vertical polarization, respectively. For comparison, we also considered the corresponding fully decohered and thus separable state $\rho_{\text{sep}} = (|HV\rangle\langle HV| + |VH\rangle\langle VH|)/2$. This state was produced as a mixture of the two pure state components $|HV\rangle$ and $|VH\rangle$ by probabilistically mixing the respective coincidence counts.

The initial state is then subjected to polarization-dependent loss, introduced to the system by means of a Jamin-Lebedev polarization-interferometer, which allows individual control of the degree of loss for each polarization mode for both Alice and Bob, see Fig. 2c). In the symmetric case considered here the loss was parametrized by a single parameter κ , where $\kappa=0$ corresponds to the loss-free scenario and $\kappa=1$ means complete loss of one polarization. With this setup we simulated correlations of increasing strength, ranging from classical to quantum and close to maximal non-signaling as discussed in detail in the methods section.

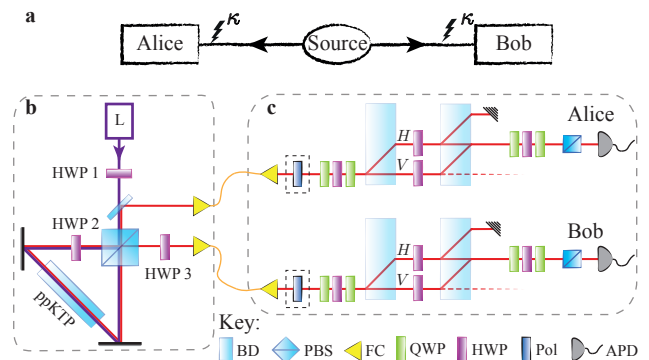


FIG. 2. **The experimental approach.** a) Pairs of single photons are created at the source and are subjected to polarization-dependent loss before Alice and Bob perform their measurements. b) The photon-source used in the experiment is spontaneous parametric down-conversion in a 10 mm long periodically poled KTiOPO_4 (ppKTP) crystal inside a polarization Sagnac interferometer using a grating stabilized continuous wave pump laser (L) at a wavelength of $\lambda = 410$ nm. By controlling the phase and polarization of this laser and adjusting the additional half-wave-plate in Bob's arm of the source, HWP3, any two-qubit states can be produced. c) Polarization-dependent loss is introduced to the system in a controllable way using an interferometer based on calcite beam displacers (BD), which split the horizontal and vertical polarization components into two spatial modes. The two HWPs in the interferometer are set to rotate the polarization by 90° , which ensures equal path-length of the two spatial modes upon recombination at the second set of BD. The degree of loss for each polarization is then proportional to the offset of the corresponding HWP from this setting. Finally, a series of quarter-wave plates (QWP), HWP and polarizing beam splitter (PBS) is used to perform the Bell measurements. Note: additional polarizers may be introduced before the interferometer to produce high quality separable states.

Using these correlations we investigated the information causality protocol on up to four levels (corresponding to a 16-bit data-set for Alice) with 1-bit of communication. Crucially, we implemented the protocol in Fig. 1b) on a shot-by-shot basis, rather than estimating the performance from coincidence probabilities. For this we used an AIT-TTM8000 time-tagging module with a temporal resolution of 82 ps to register the single photon counts for all possible outcomes. From this data, using passive feed-forward, i.e. at the processing stage, we were able to reconstruct over 10^5 individual trials of the protocol for each of the 21 settings of uniformly increasing κ .

V. RESULTS

At a correlation strength of $S=3.874(5)$, the information available to Bob is at least $I \geq 1.86(2)$ bits, despite only receiving 1 bit from Alice. For four nesting lev-

els of the protocol we establish lower bounds as high as $I \geq 7.47(11)$ bit, which violates the information causality inequality $I \leq 1$ by almost 60 standard deviations. Similarly for weaker correlations, Bob has more information available than contained in Alice’s message for all nesting levels as soon as the correlation strength surpasses $S \approx 3.5$. The fact that this value is significantly higher than Tsirelson’s bound of $S_Q \approx 3.41$ emphasizes that the quantity I only recovers this bound in the asymptotic limit $n \rightarrow \infty$.

In the following we therefore consider an alternative figure of merit, motivated by identifying the protocol in Fig. 1b) as a special case of a so-called random access code [12]. Using similar ideas as in Ref. [7], the efficiency of this task can be bounded by

$$\eta = \sum_{k=0}^{N-1} (2P_k - 1)^2 \leq 1, \quad (4)$$

which thus also encompasses the principle of information causality [12]. This bound, however, can indeed be saturated by quantum states for any size of Alice’s dataset, as illustrated in Fig. 3. Note that our data violates the bound before the correlations surpass Tsirelson’s bound. This is a result of a slight anisotropy in the simulated correlations due to experimental imperfections and a resulting bias for certain data-sets. It is not present when considering isotropic correlations, see Fig. 3b). Crucially, this highlights the dependence of both figures of merit (3) and (4) to the specific random choice of Alice’s data-set.

In particular, the separable state used in the simulation produces entanglement-like correlations for one measurement choice of Alice and uncorrelated outputs for the other, see Fig. S1. Hence, depending on the choice of data-set the figures of merit η and I might resemble the behavior expected for an entangled state, for a completely mixed state or, for higher nesting level, anything in-between. Only when averaging over all possible datasets, $\{a_i\}$, for a given level or employing the “depolarization” protocol introduced in Ref. [13] to make the correlations isotropic without changing the CHSH value, can the quantities (3) and (4) be used as reliable figures of merit, see Fig. S1 and S2. Note, however, that anisotropic super-quantum correlations (averaged over all datasets) do not necessary violate Tsirelson’s bound. In this case the principle of information causality cannot be probed using the depolarization approach, since it would result in isotropic correlations and information causality would not be violated.

VI. DISCUSSION

In contrast to the full set of no-signaling correlations, and the set of classical correlations, which both have the form of a well-characterized polytope, much less is known about the quantum set [3, 14]. Understanding the set of quantum correlations theoretically and characterizing




FIG. 3. Experimental results for the efficiency in the information causality protocol. **a)** Shown is the efficiency of the protocol for increasing strength of correlation, see methods section. The data points represent $n = 1$ (blue circles), $n = 2$ (red squares), $n = 3$ (yellow diamonds) and $n = 4$ (green triangles) levels in the protocol, where at each level a random dataset $\{a_i\}$ was used. Error-bars represent the standard deviation of 5 individual runs of every protocol. The lines correspond to theoretical expectations for the given correlation strength. **b)** A zoom into the region where our data violates Tsirelson’s bound (indicated by the grey, vertical line). Our data violates the bound of $\eta \leq 1$ already before the correlation strength surpasses Tsirelson’s bound, which is a result of a finite sample size and the particular choice of random dataset, see Sec SI. In the right panel, the same plot for isotropic correlations obtained from using the protocol of Ref. [13] shows very good agreement with the theoretical predictions.

it experimentally should thus be a primary aim from a practical as well as a fundamental perspective. Information causality, which has been proposed as a physical principle to reconstruct the set of quantum correlations, has already proven successful in recovering the famous Tsirelson bound. This limit of quantum correlations, however, is only one extremal point on the continuous boundary and there exist correlations below it, which nevertheless do not admit a quantum description [1]. Information causality also rules out such correlations for some 2-dimensional slices of the full (8-dimensional) no-signaling polytope, while it does not for other slices [1]. This shortcoming, nevertheless is not definite and might just be a result of a suboptimal protocol in Fig. 1b).

A violation of information causality would in particular imply that the tested theory does not admit a suitable measure of one of the most elementary information theoretic quantities: entropy [12, 15]. Such a measure

is assumed to be consistent with the classical limit and such that the entropy change ΔH of a composite system XY satisfies $\Delta H(XY) \geq \Delta H(X) + \Delta H(Y)$ under local evolution of the subsystems X and Y . Hence, a failure of these requirements could be interpreted as allowing for the generation of non-local correlations via local transformations. Similar consequences might also arise from the violation of alternatives to information causality, which are more or less successful in recovering part of the quantum boundary. Examples include the principles of local orthogonality [16], the requirement that the theory has a suitable classical limit [17] or that certain communication [18–20] or computational tasks [21] are non-trivial.

Our method of simulating super-quantum correlations could be adapted to explore some of these alternative principles as well. Of particular interest, however, would be a test of information causality in the multipartite case, since most of the above principles are formulated in the bipartite setting, which is bound to fail in recovering the full quantum boundary due to the existence of multipartite super-quantum correlations, which obey every bipartite principle [22, 23]. While there are studies of information causality for higher-dimensional systems, which strengthen its position as a physical principle that determines quantum correlations [24], a suitable generalization to the multipartite case is still an open problem.

As highlighted by our experiment, special focus has to be put on anisotropic regions of the no-signaling polytope. Specifically we find that the introduced figures of merit are not valid in a single instance of the protocol and have to be averaged over all possible datasets or estimated from the depolarized, isotropic data. This subtle, but very important detail is clearly highlighted by our experimental results, where we show how even a small amount of imbalance can result in a violation of the principle by quantum states for a specific choice of parameters, while obeying the principle on average.

VII. METHODS

Examining the results of a CHSH-inequality test make it clear where our data crosses the boundary of the quantum set. In our investigation we focused on the scenario of a fixed maximally entangled state $|\psi_+\rangle$ in situations with different amounts of loss, as shown in Fig. 4. We further considered the state ρ_{sep} , which resembles the state $|\psi_+\rangle$ after full decoherence as might happen during propagation between Alice and Bob. This allows for an intuitive comparison between the entangled and unentangled case.

The tested inequality has the form of a CHSH-inequality with measurements in the YZ-plane of the Bloch sphere. For the lossless case $\kappa = 0$, Alice’s and Bob’s measurements can be viewed as the application of appropriate basis-rotations (around the X-axis) followed by projective measurements in the $|H/V\rangle$ -basis. These rotations can also be seen as phase-gates in the diag-

onal polarization basis $|\pm\rangle = \frac{1}{\sqrt{2}}(|H\rangle \pm |V\rangle)$. In the case where polarization dependent loss is present, these phase-gates become non-unitary. They act as the identity on the state $|u\rangle = (\sqrt{1+\kappa}|H\rangle + \sqrt{1-\kappa}|V\rangle)/\sqrt{2}$ and impose a phase on the non-orthogonal state $|w\rangle = (\sqrt{1+\kappa}|H\rangle - \sqrt{1-\kappa}|V\rangle)/\sqrt{2}$, where $\kappa = \langle u|w\rangle$. The precise relation between κ and the degree of loss is discussed in Ref. [2]. As non-unitary operations can only be performed non-deterministically, postselection on success is required, which results in the observation of apparent super-quantum correlations. Finally, we use the first step of the depolarization procedure in Ref. [13] to symmetrize the simulated correlations, while preserving their possible anisotropy.

Curiously, we note that moderate polarization-dependent loss can lead to super-quantum correlations for entangled states without invalidating the CHSH inequality for separable states, as suggested in Ref. [2]. This observation even holds when optimizing the separable state for maximal CHSH-value, for each degree of loss [2]. Note, however, that these results were obtained using the same measurements for both separable and entangled states, whereas arbitrary hidden variable theories would allow arbitrary measurements.




FIG. 4. **Experimental results.** Shown are the experimentally obtained values for the CHSH-parameter S for both the entangled state $|\psi^+\rangle$ (blue circles) and the separable state ρ_{sep} (red squares), together with the theoretical predictions (blue and red lines, respectively) for these states, versus the amount of polarization-dependent loss as parametrized by κ . The gray dashed line represents the theoretical expectation for the optimal separable state for a given amount of loss. In the experiment we observe a violation of Tsirelson’s bound for $\kappa \geq 0.3$. Interestingly, we identify a region ($0.3 \leq \kappa \leq 0.372$) where the quantum bound of the inequality is violated, while the classical bound still holds for all separable states. With the chosen, fixed, separable state ρ_{sep} we observe a first violation at $\kappa = 0.5$. Errors from a Monte-Carlo sampling of the Poissonian counting statistics are not visible on the scale of this plot.

Figure 4 illustrates the obtained values of the CHSH parameter S and compares them to the ideal case, which,

for the initially entangled state, is described by

$$S_{|\psi_+\rangle}(\kappa) = 3 \frac{\kappa^2 - \cos \frac{\Theta}{2}}{2(1 - \kappa^2 \cos \frac{\Theta}{2})} - \frac{\kappa^2 - \cos \frac{3\Theta}{2}}{2(1 - \kappa^2 \cos \frac{3\Theta}{2})} + 2. \quad (5)$$

Here Θ is a function of κ , which can be analytically approximated by $\Theta = \pi(17 + \cos(\pi\kappa))/12$, as discussed in Ref. [2].

We experimentally violate Tsirelson's bound by more than 7 standard deviations, $S=3.423(1)$ at a loss parameter of $\kappa=0.3$. At this point, the achieved value for the unentangled state, $S=2.821(2)$, is indeed well below the classical bound of 3 and even the optimal unentangled state does not violate the inequality until $\kappa \approx 0.37$. In the region $0.3 \leq \kappa \leq 0.37$ it is therefore possible to exploit super-quantum correlations from entangled states while unentangled states still appear classical. With increasing loss, both states eventually violate Tsirelson's bound and approach the numerical maximum of $S = 4$, with experimental values of $S = 3.9341(6)$ and $S = 3.929(2)$ for the entangled and separable state, respectively. The increasing deviation from the theoretical predictions in Fig. 4 is a result of the decreasing signal-to-noise ratio in the single-photon detectors for high-loss settings.

Related experiments have observed apparent violations

of Tsirelson's bound as a consequence of explicit violations of the detection loophole [25] or the fair-sampling assumption [26]. The latter is in fact typically violated if the quantum system of interest has more (possibly) correlated degrees of freedom than those tested in the Bell-inequality [27]. Violation of Tsirelson's bound has also been considered as an intermediate step in deriving three-qubit inequalities [28].

Acknowledgments

We thank Tim Ralph and Cyril Branciard for helpful discussions. We also thank the team from the Austrian Institute of Technology for kindly providing the time-tagging modules for this experiment. We acknowledge financial support from the ARC Centres of Excellence for Engineered Quantum Systems (CE110001013) and Quantum Computation and Communication Technology (CE110001027). D.W.B. is funded by an Australian Research Council Future Fellowship (FT100100761), and A.F. by an Australian Research Council Discovery Early Career Researcher Award (DE130100240). A.G.W. acknowledges support from a UQ Vice-Chancellor's Senior Research Fellowship.

-
- [1] Allcock, J., Brunner, N., Pawłowski, M. & Scarani, V. Recovering part of the boundary between quantum and nonquantum correlations from information causality. *Phys. Rev. A* **80**, 040103 (2009).
- [2] Berry, D. W., Jeong, H., Stobińska, M. & Ralph, T. C. Fair-sampling assumption is not necessary for testing local realism. *Phys. Rev. A* **81**, 012109 (2010).
- [3] Barrett, J. *et al.* Non-local correlations as an information theoretic resource. *Phys. Rev. A* **71**, 022101 (2005).
- [4] Popescu, S. & Rohrlich, D. Quantum nonlocality as an axiom. *Found. Phys.* **24**, 379–385 (1994).
- [5] Clauser, J., Horne, M., Shimony, A. & Holt, R. Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett.* **23**, 880–884 (1969).
- [6] Cirel'son, B. S. Quantum generalizations of Bell's inequality. *Lett. Mat. Phys.* **4**, 93–100 (1980).
- [7] Pawłowski, M. *et al.* Information causality as a physical principle. *Nature* **461**, 1101–1104 (2009).
- [8] van Dam, W. *Nonlocality & Communication Complexity*. Phd thesis, University of Oxford (2000).
- [9] Kim, T., Fiorentino, M. & Wong, F. N. C. Phase-stable source of polarization-entangled photons using a polarization Sagnac interferometer. *Phys. Rev. A* **73**, 012316 (2006).
- [10] Fedrizzi, A., Herbst, T., Poppe, A. & Zeilinger, A. A wavelength-tunable fiber-coupled source of narrowband entangled photons. *Opt. Exp.* **15**, 15377–15386 (2007).
- [11] Smith, D. H. *et al.* Conclusive quantum steering with superconducting transition-edge sensors. *Nature Comm.* **3**, 625 (2012).
- [12] Al-Safi, S. W. & Short, A. J. Information causality from an entropic and a probabilistic perspective. *Phys. Rev. A* **84**, 042323 (2011).
- [13] Masanes, L., Acín, A. & Gisin, N. General properties of nonsignaling theories. *Phys. Rev. A* **73**, 012112 (2006).
- [14] Brunner, N., Cavalcanti, D., Pironio, S., Scarani, V. & Wehner, S. Bell nonlocality. *Rev. Mod. Phys.* **86**, 419–478 (2014).
- [15] Dahlsten, O. C. O., Lercher, D. & Renner, R. Tsirelson's bound from a generalized data processing inequality. *New J. Phys.* **14**, 063024 (2012).
- [16] Fritz, T. *et al.* Local orthogonality as a multipartite principle for quantum correlations. *Nature Comm.* **4**, 2263 (2013).
- [17] Navascués, M. & Wunderlich, H. A glance beyond the quantum model. *Proc. R. Soc. A* **466**, 881–890 (2009).
- [18] van Dam, W. Implausible consequences of superstrong nonlocality (2005). arXiv:quant-ph/0501159.
- [19] Brassard, G. *et al.* A limit on nonlocality in any world in which communication complexity is not trivial. *Phys. Rev. Lett.* **96**, 250401 (2006).
- [20] Brunner, N. & Skrzypczyk, P. Nonlocality distillation and postquantum theories with trivial communication complexity. *Phys. Rev. Lett.* **102**, 160403 (2009).
- [21] Linden, N., Popescu, S., Short, A. & Winter, A. Quantum nonlocality and beyond: limits from nonlocal computation. *Phys. Rev. Lett.* **99**, 180502 (2007).
- [22] Gallego, R., Würflinger, L. E., Acín, A. & Navascués, M. Quantum correlations require multipartite information principles. *Phys. Rev. Lett.* **107**, 210403 (2011).
- [23] Yang, T. H., Cavalcanti, D., Almeida, M. L., Teo, C. & Scarani, V. Information-causality and extremal tripartite correlations. *New J. Phys.* **14**, 013061 (2012).

- [24] Cavalcanti, D., Salles, A. & Scarani, V. Macroscopically local correlations can violate information causality. *Nature Comm.* **1**, 136 (2010).
- [25] Tasca, D., Walborn, S., Toscano, F. & Souto Ribeiro, P. Observation of tunable Popescu-Rohrlich correlations through postselection of a Gaussian state. *Phys. Rev. A* **80**, 030101 (2009).
- [26] Romero, J., Giovannini, D., Tasca, D., Barnett, S. & Padgett, M. Tailored two-photon correlation and fair-sampling: a cautionary tale. *New J. Phys.* **15**, 083047 (2013).
- [27] Dada, A. C. & Andersson, E. On Bell inequality violations with high-dimensional systems. *IJQI* **09**, 1807–1823 (2011).
- [28] Cabello, A. Violating Bell's inequality beyond Cirel'son's bound. *Phys. Rev. Lett.* **88**, 060403 (2002).

Supplemental Material

Here we discuss in more detail the subtleties of the information causality protocol for individual choices of the data-set for Alice. In particular we will discuss the effect of anisotropic correlations of decohered entangled states and present results for the case where these are transformed to isotropic correlations by incorporating the protocol proposed in Ref. [13] into our experiment.

SI. SEPARABLE STATE CORRELATIONS AND THE CHOICE OF DATA-SET

As discussed previously, the experiment was performed for both entangled and separable initial states. While the entangled state allows for generation of super-quantum correlations in the range $3.379(1) \leq S \leq 3.9341(6)$, the separable state covers a larger range of $2.698(2) \leq S \leq 3.929(2)$. This range in particular includes a large part of the set of classical ($S \leq 3$) and quantum correlations ($S \leq 3.41$). Here it is important to reiterate the form of the used state. Since it has been chosen as the fully decohered version of $|\psi_+\rangle$ it still retains correlations of the same strength in one basis, while correlations in any orthogonal basis are lost—a classically correlated state.

a		Outputs: AB				b		Outputs: AB			
		00	01	10	11			00	01	10	11
Inputs: ab	00	0.43	0.07	0.07	0.43	Inputs: ab	00	0.25	0.25	0.25	0.25
	01	0.43	0.07	0.07	0.43		01	0.25	0.25	0.25	0.25
	10	0.43	0.07	0.07	0.43		10	0.43	0.07	0.07	0.43
	11	0.07	0.43	0.43	0.07		11	0.07	0.43	0.43	0.07

FIG. S1. **Measurement probabilities for the lossless case $\kappa = 0$.** Shown are the theoretically expected measurement probabilities for both the **a)** entangled and **b)** separable state for a lossless Bell-test experiment. Note that the correlations are the same for the two states when Alice chooses to measure 1, but there are no correlations for the separable state when she measures 0.

An important consequence of this feature, as discussed in the main text, is that the success probability in the information causality protocol then depends on the specific choice of data-set for Alice. As an example consider the simplest instance, where Alice has two bits a_0 and a_1 . Alice uses $a_0 \oplus a_1$ as her input and Bob uses $b \in \{0, 1\}$. The probability of success in this scenario will be the same as for the entangled state whenever $a_0 \oplus a_1 = 1$ (that is for the data-sets $\{0, 1\}$ and $\{1, 0\}$) and the same as for random guessing ($1/2$) in the other two cases, where $a_0 \oplus a_1 = 0$. Although Ref. [7] discussed the related case where the probability of success may depend on Bob’s choice, the feature revealed here has important practical implications, since the calculation of the figures of merit always requires to consider all of Bob’s possible choices, while only one data-set for Alice has to be considered. Depending on the choice of data-set, any degree of effi-

ciency can be achieved with classical states.

There are several ways to circumvent these problems. Clearly, averaging over all possible data-sets will recover the performance expected from the CHSH-value corresponding to the respective state. However, as the possible choices for Alice’s data-set grow doubly-exponentially, this approach is typically unfeasible, in particular when testing the protocol on a shot-by-shot basis. Here we employed the “depolarization” protocol introduced by Ref. [13], which takes any set of correlations to an isotropic one without changing the CHSH-value, using 3 bits of shared randomness. The drawback of this method, however, is that it precludes a test of information causality in the anisotropic regime of super-quantum correlations below Tsirelson’s bound, see Ref. [1].

SII. RESULTS FOR ISOTROPIC CORRELATIONS

Figure S2 shows the experimental lower bounds on the mutual information measure I for the isotropic correlations obtained from the depolarization protocol applied to the initially separable state. Figure S3 shows the random access code efficiency η for the same data. In both cases we observe excellent agreement with the theoretical predictions.




FIG. S2. **Experimental results for the mutual information gain in the information causality protocol.** Shown is the lower bound on the mutual information gain in the protocol for increasing strength of isotropic correlation. The data points represent $n = 1$ (blue circles), $n = 2$ (red squares), $n = 3$ (yellow diamonds) and $n = 4$ (green triangles) levels in the protocol, where at each level a random dataset $\{a_i\}$ was used. Error-bars represent the standard deviation of 5 individual runs of every protocol. The lines correspond to theoretical expectations for the given correlation strength.




FIG. S3. **Experimental results for the efficiency in the information causality protocol.** Shown is the efficiency of the protocol for increasing strength of isotropic correlation. The data points represent $n = 1$ (blue circles), $n = 2$ (red squares), $n = 3$ (yellow diamonds) and $n = 4$ (green triangles) levels in the protocol, where at each level a random dataset $\{a_i\}$ was used. Error-bars represent the standard deviation of 5 individual runs of every protocol. The lines correspond to theoretical expectations for the given correlation strength.

APPENDIX B

Experimental Simulation of Closed Timelike Curves

This appendix reproduces a paper of which I am the lead author and which was first published in *Nature Communications* **5**, 4145 (2014).

Experimental Simulation of Closed Timelike Curves

Martin Ringbauer^{1,2*}, Matthew A. Broome^{1,2}, Casey R. Myers¹, Andrew G. White^{1,2} and Timothy C. Ralph²

¹*Centre for Engineered Quantum Systems,* ²*Centre for Quantum Computer and Communication Technology, School of Mathematics and Physics, University of Queensland, Brisbane, QLD 4072, Australia*

Closed timelike curves are among the most controversial features of modern physics. As legitimate solutions to Einstein’s field equations, they allow for time travel, which instinctively seems paradoxical. However, in the quantum regime these paradoxes can be resolved leaving closed timelike curves consistent with relativity. The study of these systems therefore provides valuable insight into non-linearities and the emergence of causal structures in quantum mechanics—essential for any formulation of a quantum theory of gravity. Here we experimentally simulate the non-linear behaviour of a qubit interacting unitarily with an older version of itself, addressing some of the fascinating effects that arise in systems traversing a closed timelike curve. These include perfect discrimination of non-orthogonal states and, most intriguingly, the ability to distinguish nominally equivalent ways of preparing pure quantum states. Finally, we examine the dependence of these effects on the initial qubit state, the form of the unitary interaction, and the influence of decoherence.

INTRODUCTION

One aspect of general relativity that has long intrigued physicists is the relative ease with which one can find solutions to Einstein’s field equations that contain closed timelike curves (CTCs)—causal loops in space-time that return to the same point in space and time [1–3]. Driven by apparent inconsistencies—like the grandfather paradox—there have been numerous efforts, such as Novikov’s self-consistency principle [4] to reconcile them or Hawking’s chronology protection conjecture [5], to disprove the existence of CTCs. While none of these classical hypotheses could be verified so far, the situation is particularly interesting in the quantum realm. In his seminal 1991 paper Deutsch showed for quantum systems traversing CTCs there always exist unique solutions, which do not allow superluminal signalling [6, 7]. Quantum mechanics therefore allows for causality violation without paradoxes whilst remaining consistent with relativity.

Advances in the field of Deutsch CTCs have shown some very surprising and counter-intuitive results, such as the solution of NP-complete problems in polynomial time [8], unambiguous discrimination of any set of non-orthogonal states [9], perfect universal quantum state cloning [10, 11] and the violation of Heisenberg’s uncertainty principle [12]. The extraordinary claims of what one could achieve given access to a quantum system traversing a CTC have been disputed in the literature, with critics pointing out apparent inconsistencies in the theory such as the information paradox or the linearity trap [13, 14]. However, it has been shown that the theory can be formulated in such a way that these inconsistencies are resolved [7, 15].

Modern experimental quantum simulation allows one to ask meaningful questions that provide insights into the behaviour of complex quantum systems. Initial results have been obtained in various areas of quantum mechanics [16–18] and in particular in the field of relativistic quantum information [19–23]. This recent experimental success, coupled with the growing interest for the study of non-linear extensions to quantum mechanics, motivates the question of whether the fundamentally non-linear dynamics and the unique behaviour arising from CTCs can be simulated experimentally.

In this article we use photonic systems to simulate the quantum evolution through a Deutsch CTC. We demonstrate how the CTC-traversing qubit adapts to changes in the input state $|\psi\rangle$, and unitary interaction U to ensure physical consistency according to Deutsch’s consistency relation [6]. We observe non-linear evolution in the circuit suggested by Bacon [8] and enhanced distinguishability of two non-orthogonal states after the action of an optimised version of a circuit proposed by Brun et al. [9]. Using the self-consistent formulation of Ref. [7] we then move beyond the simplest implementations and find a striking difference in the behaviour of the system for direct as opposed to entanglement-assisted state preparation. Finally, we explore the system’s sensitivity to decoherence.




FIG. 1: **Model of a quantum state $|\psi\rangle$ interacting with an older version of itself.** This situation can equivalently be interpreted as a chronology-respecting qubit interacting with a qubit trapped in a CTC. The CTC in general consists of a causal worldline with its past and future ends connected via a wormhole (indicated by black triangles).

*Electronic address: ringbauer.martin@gmail.com

RESULTS

The Deutsch model. While there has been some recent success on alternative models of CTCs based on postselection [23–25], we focus on the most prominent model for describing quantum mechanics in the presence of CTCs, introduced by Deutsch [6]. Here a quantum state $|\psi\rangle$ interacts unitarily with an older version of itself, Fig. 1. With the inclusion of an additional SWAP gate, this can equivalently be treated as a two-qubit system, where a chronology-respecting qubit interacts with a qubit ρ_{CTC} trapped in a closed timelike curve. The quantum state of ρ_{CTC} in this picture is determined by Deutsch’s consistency relation:

$$\rho_{\text{CTC}} = \text{Tr}_1 [U' (|\psi\rangle\langle\psi| \otimes \rho_{\text{CTC}}) U'^{\dagger}], \quad (1)$$

where U' is the unitary U followed by a SWAP gate, Fig. 1. This condition ensures physical consistency—in the sense that the quantum state may not change inside the wormhole—and gives rise to the non-linear evolution of the quantum state $|\psi\rangle$. The state after this evolution is consequently given by $\rho_{\text{OUT}} = \text{Tr}_2 [U' (|\psi\rangle\langle\psi| \otimes \rho_{\text{CTC}}) U'^{\dagger}]$. The illustration in Fig. 1 further shows that the requirement of physical consistency forces ρ_{CTC} to adapt instantly to any changes in the surroundings, such as a different interaction unitary U or input state $|\psi\rangle$. While Eq. (1) is formulated in terms of a pure input state $|\psi\rangle$ it can be directly generalised to mixed inputs [7].

Simulating CTCs. Our experimental simulation of a qubit in the (pure) state $|\psi\rangle$ traversing a CTC relies on the circuit diagram shown in Fig. 2a). A combination of single qubit unitary gates before and after a controlled-Z gate allows for the implementation of a large set of controlled-unitary gates U . Using polarisation-encoded single photons, arbitrary single qubit unitaries can be realised using a combination of quarter-wave (QWP) and half-wave plates (HWP); additional SWAP gates before or after U are implemented as a physical mode-swap. The controlled-Z gate is based on non-classical (Hong-Ou-Mandel) interference of two single photons at a single partially polarising beam-splitter (PPBS) that has different transmittivities $\eta_V=1/3$ for vertical (V) and $\eta_H=1$ for horizontal (H) polarisation [26]—a more detailed description of the implementation of the gate can be found in Ref. [27]. Conditioned on post-selection it induces a π phase-shift when the two interfering single-photon modes are vertically polarised, such that $|VV\rangle \rightarrow -|VV\rangle$ with respect to all other input states.

One of the key features of a CTC is the inherently non-linear evolution that a quantum state $|\psi\rangle$ undergoes when traversing it. This is a result of Deutsch’s consistency relation, which makes ρ_{CTC} dependent on the input state $|\psi\rangle$. In order to simulate this non-linear behaviour using linear quantum mechanics we make use of the

effective non-linearity obtained from feeding extra information into the system. In our case we use the classical information about the preparation of the state $|\psi\rangle$ and the unitary U to prepare the CTC qubit in the appropriate state ρ_{CTC} as required by the consistency relation Eq. (1). After the evolution we perform full quantum state tomography on the CTC qubit in order to verify that the consistency relation is satisfied.

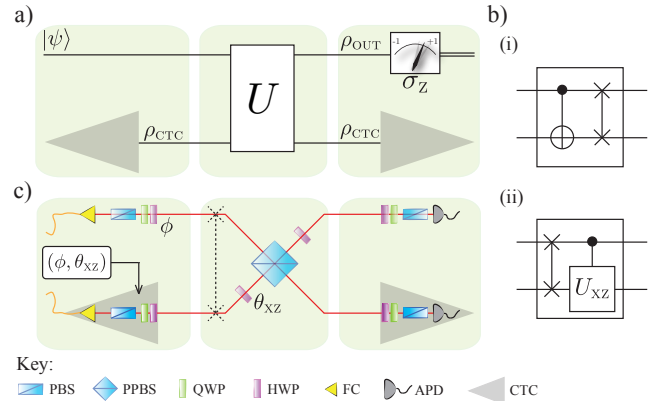


FIG. 2: Experimental details. **a)** The circuit diagram for a general unitary interaction U between the state $|\psi\rangle$ and the CTC system. **b)** The specific choice of unitary in the demonstration of the **(i)** non-linear evolution and **(ii)** perfect discrimination of non-orthogonal states. **c)** Experimental setup for case (ii). Two single photons, generated via spontaneous parametric down-conversion in a nonlinear β -barium-borate crystal, are coupled into two optical fibres (FC) and injected into the optical circuit. Arbitrary polarisation states are prepared using a Glan-Taylor polariser (POL), a quarter-wave (QWP) and a half wave-plate (HWP). Non-classical interference occurs at the central partially-polarising beam-splitter (PPBS) with reflectivities $\eta_H=0$ and $\eta_V=2/3$. Two avalanche photo-diodes (APD) detect the single photons at the outputs. The states $|\psi\rangle$ are chosen in the xz -plane of the Bloch sphere, parametrised by ϕ and CU_{xz} is the corresponding controlled unitary, characterised by the angle θ_{xz} . The SWAP gate was realized via relabeling of the input modes.

Non-linear evolution. As a first experiment we investigate the non-linearity by considering a Deutsch CTC with a CNOT interaction followed by a SWAP gate as illustrated in Fig. 2b)(i). This circuit is well-known for the specific form of non-linear evolution:

$$\alpha|H\rangle + e^{i\varphi}\beta|V\rangle \rightarrow (\alpha^4 + \beta^4)|H\rangle\langle H| + 2\alpha^2\beta^2|V\rangle\langle V|, \quad (2)$$

which has been shown to have important implications for complexity theory, allowing for the solution of NP-complete problems with polynomial resources [8]. According to Deutsch’s consistency relation, Eq. (1) the state of the CTC-qubit for this interaction is given by

$$\rho_{\text{CTC}} = \alpha^2|H\rangle\langle H| + \beta^2|V\rangle\langle V|. \quad (3)$$

We investigate the non-linear behaviour experimentally for 14 different quantum states $|\psi\rangle = \cos(\frac{\phi}{2})|H\rangle +$

$e^{i\varphi} \sin(\frac{\phi}{2})|V\rangle$, with $\phi \in \{0, \frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4}, \pi\}$ and a variety of phases $\varphi \in \{0, 2\pi\}$, where the locally available information ϕ and φ is used to prepare ρ_{CTC} . In standard (linear) quantum mechanics no unitary evolution can introduce additional distinguishability between quantum states. To illustrate the non-linearity in the system we employ two different distinguishability measures: the trace-distance $\mathcal{D}(\rho_1, \rho_2) = \frac{1}{2} \text{Tr}[|\rho_1 - \rho_2|]$, where $|\rho| = \sqrt{\rho^\dagger \rho}$ and a single projective measurement with outcomes “+” and “-”:

$$\mathcal{L}(\rho_1, \rho_2) = \langle +|\rho_1|+\rangle \langle -|\rho_2|-\rangle + \langle -|\rho_1|-\rangle \langle +|\rho_2|+\rangle. \quad (4)$$

While the metric \mathcal{D} is a commonly used distance measure it does not have an operational interpretation and requires full quantum state tomography in order to be calculated experimentally. The measure \mathcal{L} in contrast is easily understood as the probability of obtaining different outcomes in minimum-error discrimination of the two states using a single projective measurement on each system. The operational interpretation and significance of \mathcal{L} is discussed in more detail in the section. Both \mathcal{D} and \mathcal{L} are calculated between the state $|\psi\rangle$ and the fixed reference state $|H\rangle$ after being evolved through the circuit shown in Fig. 2b(i). The results are plotted and compared to standard quantum mechanics in Fig. 3. If the state $|\psi\rangle$ is not known then, based only on the knowledge of the reference state $|H\rangle$ and the evolution in Eq. (2) it is natural and optimal to use the measure \mathcal{L} with a σ_z -measurement.

We observe enhanced distinguishability for all states with an initial trace-distance to $|H\rangle$ smaller than $1/\sqrt{2}$ (i.e. $\phi \leq \frac{\pi}{2}$), as clearly demonstrated by the σ_z -based measure, see Fig. 3. Note, however, that this advantage over standard quantum mechanics is not captured by the metric $\mathcal{D}(\rho_1, \rho_2)$ unless the non-linearity is amplified by iterating the circuit on the respective output at least 3 times, see inset of Fig. 3. This shows that the non-linearity is not directly related to the distance between two quantum states. By testing states with various polar angles for each azimuthal angle on the Bloch sphere, we confirm that any phase information is erased during the evolution and that the evolved state ρ_{OUT} is indeed independent of φ , up to experimental error. We further confirm, with an average quantum state fidelity of $\mathcal{F} = 0.998(2)$ between the input and output state of ρ_{CTC} in Eq. (3), that the consistency relation (1) is satisfied for all tested scenarios.

Non-orthogonal state discrimination. While it is the crucial feature, non-linear state evolution is not unique to the SWAP.CNOT interaction, but rather a central property of all non-trivial CTC interactions. Similar circuits have been found to allow for perfect distinguishability of non-orthogonal quantum states [9], leading to disconcerting possibilities such as breaking of quantum cryptography [9], perfect cloning of quantum states [10, 11], and violation of Heisenberg’s uncertainty




FIG. 3: Non-linear evolution in a Deutsch CTC with SWAP.CNOT interaction. Both the trace distance \mathcal{D} , and the σ_z -based distinguishability measure \mathcal{L} (equal to within experimental error in this case) of the evolved states ρ_{OUT} after the interaction with the CTC are shown as yellow diamonds. The blue circles (red squares) represent the measure \mathcal{D} (\mathcal{L}) between the input states $|\psi\rangle$ and $|H\rangle$ in the case of standard quantum mechanics. Note that due to the phase-independence of the evolution in Eq. (2) states that only differ by a phase collapse to a single data point. Crucially, the metric \mathcal{D} does not capture the effect of the non-linearity, while \mathcal{L} does, indicated by the red shaded region. Error bars obtained from a Monte Carlo routine simulating the Poissonian counting statistics are too small to be visible on the scale of this plot. **Inset:** The dashed black lines with decreasing thickness represent theoretical expectations for \mathcal{D} and \mathcal{L} from 2, 3, 4 and 5 iterations of the circuit.

principle [12]. In particular it has been shown that a set $\{|\psi_j\rangle\}_{j=0}^{N-1}$ of N distinct quantum states in a space of dimension N can be perfectly distinguished using an N -dimensional CTC-system. The algorithm proposed by Brun et al. [9] relies on an initial SWAP operation between the input and the CTC-system, followed by a series of N controlled unitary operations, transforming the input states to an orthogonal set, which can then be distinguished.

In our simulation of this effect we consider the qubit case $N=2$, which consequently would require two controlled unitary operations between the input state and the CTC system. We note, however, that without loss of generality the set of states to be discriminated can be rotated to the xz-plane of the Bloch sphere, such that $|\psi_0\rangle = |H\rangle$ and $|\psi_1\rangle = \cos(\frac{\phi}{2})|H\rangle + \sin(\frac{\phi}{2})|V\rangle$ for some angle ϕ . In this case, the first controlled unitary is the identity operation \mathcal{I} , while the second performs a controlled rotation of $|\psi_1\rangle$ to $|V\rangle$ as illustrated in Fig. 4a). In detail, the gate CU_{xz} applies a π rotation to the target qubit conditional on the state of the control qubit, about an axis in the xz-plane defined by the angle θ_{xz} . For the optimal choice $\theta_{\text{xz}} = \frac{\phi - \pi}{2}$ the gate rotates the state $|\psi_1\rangle$ to $|V\rangle$, orthogonal to $|\psi_0\rangle$, enabling perfect distinguishability by means of a projective σ_z measurement, see Fig. 4a).

In practice the gate CU_{xz} is decomposed into a

controlled-Z gate between appropriate single qubit rotations, defining the axis θ_{xz} . The latter are realised by half-wave plates before and after the PPBS, set to an angle of $\theta_{xz}/4$ with respect to their optic axis, see Fig. 2c):

$$CU_{xz}(\theta_{xz}) = (\mathcal{I} \otimes \text{HWP}(\theta_{xz}/4)) \cdot \text{CZ} \cdot (\mathcal{I} \otimes \text{HWP}(\theta_{xz}/4))$$

$$= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cos(\theta_{xz}) & \sin(\theta_{xz}) \\ 0 & 0 & \sin(\theta_{xz}) & -\cos(\theta_{xz}) \end{pmatrix}. \quad (5)$$

Note that relation (1) requires that $\rho_{\text{CTC}} = |H\rangle\langle H|$, whenever the input state is $|H\rangle$, independent of the gate CU_{xz} . Crucially, this consistency relation ensures that any physical CTC-system adapts instantly to changes in ϕ and θ_{xz} , parametrising the input state and gate, respectively. In our simulation these two parameters are used to prepare the corresponding state ρ_{CTC} , as shown in Fig. 2c).

In a valid experimental simulation the input and output states ρ_{CTC} have to match, i.e. ρ_{CTC} has to satisfy relation (1). This has been verified for all following experiments with an average quantum state fidelity of $\mathcal{F} = 0.996(7)$.

In the experiment, we prepared near-pure quantum states directly on single photons using a Glan-Taylor polariser followed by a combination of a HWP and a QWP. We simulated CTC-aided discrimination of non-orthogonal states for 32 distinct quantum states $|\psi_1\rangle$ with $\phi \in [0, 2\pi)$. For each state we implemented CU_{xz} with the optimal choice of $\theta_{xz} = \frac{\phi - \pi}{2}$. Furthermore we tested the ability of this system to distinguish the set $\{|\psi_0\rangle, |\psi_1\rangle\}$ given non-optimal combinations of ϕ and θ_{xz} . For this we either chose $\phi = 3\pi/2$ and varied the gate over the full range of $\theta_{xz} \in [-\frac{\pi}{2}, \frac{\pi}{2})$, or chose CU_{xz} as a controlled Hadamard (optimal for $\phi = 3\pi/2$) and varied the state $|\psi_1\rangle$ over the full range of $\phi \in [0, 2\pi)$. The output state is characterized by quantum state tomography, which provides sufficient data to obtain \mathcal{L} for arbitrary measurement directions as well as for the calculation of the trace-distance.

Figure 5a) illustrates the observed distinguishability \mathcal{L} for the above experiments and compares it to the expectation from standard quantum mechanics. In the latter case the measure \mathcal{L} is maximized by choosing the optimal projective measurement, based on the available information about the states $|\psi_0\rangle$ and $|\psi_1\rangle$. Crucially, the optimized \mathcal{L} is directly related to the trace-distance \mathcal{D} as $\mathcal{L} = \frac{1}{2}(1 + \mathcal{D}^2)$ and therefore captures the same qualitative picture, without the requirement for full quantum state tomography. In the CTC case a σ_z -measurement is chosen, which is optimal when $\theta_{xz} = \frac{\phi - \pi}{2}$. Otherwise further optimization is possible based on the knowledge of θ_{xz} (see section and Fig. S1 for more details). Furthermore, we note that the above scenario can also be interpreted in a state-identification rather than state-

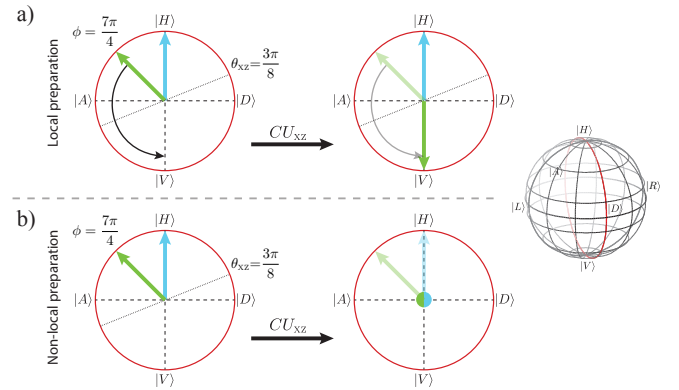


FIG. 4: **Bloch-sphere evolution of states traversing a CTC.** In the case of (a) local state preparation, the state $|\psi_0\rangle = |H\rangle$ (blue) is unaffected by CU_{xz} , while $|\psi_1\rangle$ (green) undergoes a π rotation about the axis defined by θ_{xz} . The axis is chosen as $\theta_{xz} = \frac{\phi - \pi}{2}$ such that $|\psi_1\rangle \mapsto |V\rangle$ which can then be perfectly distinguished from $|\psi_0\rangle$. (b) For non-local preparation of the initial states and the same choice of θ_{xz} the controlled unitary maps both initial states to the maximally mixed state $\frac{1}{2}(|H\rangle\langle H| + |V\rangle\langle V|)$. The probability of distinguishing the two states is therefore $1/2$ —as good as randomly guessing.

discrimination picture, which is discussed in more detail in section and illustrated in Fig. S2.

Local vs. non-local state preparation. Due to the inherent non-linearity in our simulated system, care must be taken when describing mixed input states ρ_{in} . In particular a distinction between proper and improper mixtures can arise which is unobservable in standard (linear) quantum mechanics [28]. This ambiguity is resolved in Ref. [7] by requiring the consistency condition to act shot-by-shot—i.e. independently in every run of the experiment—on the reduced density operator of the input state. For proper mixtures this means that ρ_{in} is always taken as a pure state, albeit a different one shot-by-shot. For improper mixtures in contrast, ρ_{in} will always be mixed. A similar, but much more subtle and fascinating feature, which has received less attention in the literature so far occurs with respect to preparation of pure states [29]. While in standard quantum mechanics a pure state prepared directly (locally) on a single qubit is equivalent to one that has been prepared non-locally through space-like separated post-selection of an entangled resource state, this is not true under the influence of a CTC. The origin of this effect is not the non-linear evolution, but rather the local absence of classical information about the post-selection outcome. The role of locally available classical information in entanglement-based preparation schemes is a matter of current debate and still to be clarified.

A possible resource state for alternatively preparing $|\psi_0\rangle$ and $|\psi_1\rangle$ could be of the form $|\Psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle \otimes |\psi_0\rangle +$

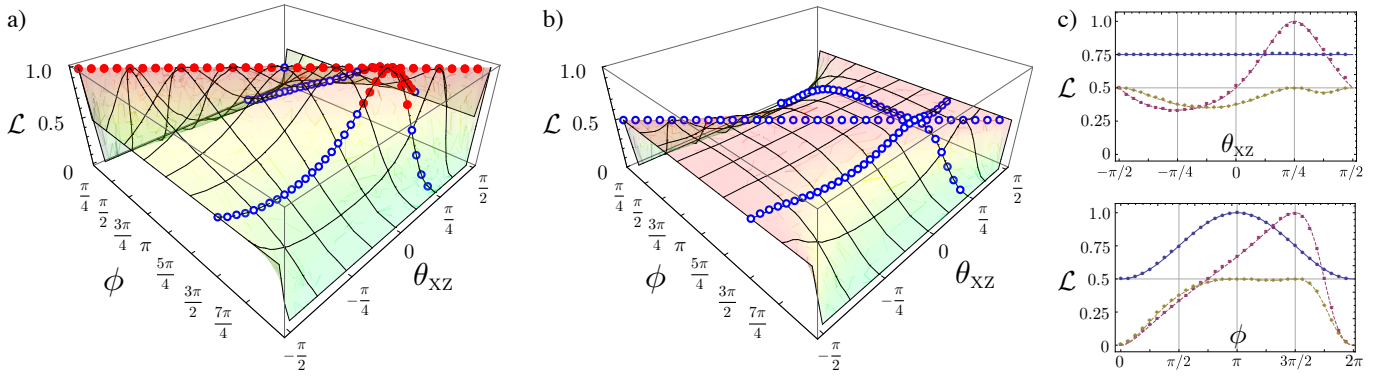


FIG. 5: **Experimental results.** Probability of state discrimination for **a)** locally prepared and **b)** non-locally prepared states $|\psi_0\rangle=|H\rangle$ and $|\psi_1\rangle = \cos(\frac{\phi}{2})|H\rangle + \sin(\frac{\phi}{2})|V\rangle$ as measured by \mathcal{L} . The surface represents the theoretically predicted probability depending on the state and gate parameters ϕ and θ_{XZ} , respectively. Solid, red (open, blue) data-points indicate better (worse) performance than standard quantum mechanics. **c)** Cross-sectional views of the combined plots a) and b) reveal the rich structure in the dependencies on the initial parameters for **(top)** a fixed state ($\phi=3\pi/2$) and **(bottom)** a fixed gate ($\theta_{XZ}=\pi/4$). Here red squares (yellow diamonds) correspond to the CTC case with local (non-local) preparation and blue circles represents standard quantum mechanics. Error bars obtained from a Monte Carlo routine simulating the Poissonian counting statistics are too small to be visible on the scale of this plot.

$|1\rangle \otimes |\psi_1\rangle$), where projection of the first qubit onto the state $|0\rangle$ and $|1\rangle$ leaves the second qubit in the state $|\psi_0\rangle$ and $|\psi_1\rangle$, respectively. From the point of view of ρ_{CTC} , however, there exists no information about the outcome of this projective measurement. Hence it “sees” and adapts to the mixed state $\rho_{in} = \text{Tr}_1[|\Psi\rangle\langle\Psi|] = \frac{1}{2}(|\psi_0\rangle\langle\psi_0| + |\psi_1\rangle\langle\psi_1|)$. The state of the CTC qubit is therefore different for local and non-local preparation. If this was not the case, it would enable superluminal signalling, which is in conflict with relativity [29].

Figure 4b) illustrates the evolution induced by CU_{XZ} , when the input states $|H\rangle$ and $|\psi_1\rangle$ are prepared using an entangled resource $|\Psi\rangle$, rather than directly. The results of the previously discussed distinguishability experiments for this case are shown in Fig. 5b). In Fig. 5c) they are compared to the case of local preparation and to standard quantum mechanics for a fixed input state and a fixed gate, respectively. Again, consistency of our simulation is ensured by a quantum state fidelity of $\mathcal{F} = 0.9996(3)$ between the input and output states of ρ_{CTC} .

In our simulation we find that the CTC-system can indeed achieve perfect distinguishability of the (directly prepared) states $|\psi_0\rangle$ and $|\psi_1\rangle$ even for arbitrarily close states if the appropriate gate is implemented, see Fig. 5a). Furthermore we show that the advantage over standard quantum mechanics persists for a wide range of non-optimal gate-state combinations, outside of which, however, the CTC-system performs worse (blue points). Notably, we find that for non-locally prepared input states CTC-assisted state discrimination never performs better than random guessing—a probability of 0.5—as shown in Fig. 5b). The predictions for standard quantum mechanics, in contrast are independent of the way the states $|\psi_0\rangle$ and $|\psi_1\rangle$ are prepared.

Decoherence. We further investigated the effect of two important decoherence mechanisms on the simulated CTC-system, shown in Fig. 2a). The first is a single qubit depolarising channel acting on the input state $|\psi\rangle$, which can be modelled as

$$\rho \mapsto (1 - \frac{3p}{4})\rho + \frac{p}{4}(\sigma_x\rho\sigma_x + \sigma_y\rho\sigma_y + \sigma_z\rho\sigma_z), \quad (6)$$

where $(\sigma_x, \sigma_y, \sigma_z)$ are the 3 Pauli matrices and $p \in [0, 1]$ quantifies the amount of decoherence.

The second form of decoherence concerns the controlled unitary CU_{XZ} and is described as

$$\rho \mapsto (1 - \varepsilon)CU_{XZ} \rho CU_{XZ}^\dagger + \varepsilon\rho, \quad (7)$$

where $\varepsilon \in [0, 1]$ is the probability of the gate to fail, describing the amount of decoherence that is present. For $\varepsilon = 0$ the gate acts as an ideal controlled rotation CU_{XZ} , while it performs the identity operation for $\varepsilon = 1$.

We tested the robustness of the state-discrimination circuit in Fig. 2b)(ii) against both forms of decoherence. For this test we chose CU_{XZ} as a controlled Hadamard (i.e. $\theta_{XZ}=\pi/4$) and the initial states $|\psi_0\rangle=|H\rangle$ and $|\psi_1\rangle=\frac{1}{\sqrt{2}}(|H\rangle-|V\rangle)$ (i.e. $\phi=3\pi/2$). Figure 6 shows the distinguishability \mathcal{L} of the evolved states as a function of both decoherence mechanisms over the whole range of parameters $p \in [0, 1]$ and $\varepsilon \in [0, 1]$. Note, that the decoherence channel in Eq. (7) does not have an analogue in the standard quantum mechanics case (i.e. without a CTC), hence only the channel in Eq. (6) is considered for comparison. It is further naturally assumed, that the experimenter has no knowledge of the specific details of the decoherence and therefore implements the optimal measurements for the decoherence-free case. The physical validity of the simulation is ensured by consistency of

ρ_{CTC} across the boundary of the wormhole with an average fidelity of $\mathcal{F} = 0.997(4)$.




FIG. 6: **State discrimination as a function of gate and qubit decoherence for locally prepared states.** Here ε quantifies the decoherence of the unitary interaction CU_{xz} (with $\theta_{xz}=\pi/4$), which has no analogue in the standard quantum mechanics case and p the single qubit depolarisation of the input qubits $|H\rangle$ and $|\psi_1\rangle$ (with $\phi=3\pi/2$). The system demonstrates robustness against both forms of decoherence and the CTC-advantage persists up to $p=\sqrt{2}-1$ and $\varepsilon=\frac{1}{3}$, respectively. The semi-transparent blue surface represents the optimum in standard quantum mechanics. Error bars obtained from a Monte Carlo routine simulating the Poissonian counting statistics are too small to be visible on the scale of this plot.

It is worth noting, that the interpretation of decoherence effects in the circuit in Fig. 2a) is very different from the linear scenario without a CTC. In the case of single-qubit depolarisation the initially pure input state becomes mixed. In contrast to the linear case now an important distinction has to be made with respect to the origin of the decoherence. If it results from an interaction with the environment, which is the case considered here, then ρ_{CTC} “sees” an improper mixture and adjusts to the mixed density matrix of the input state. If, however, the origin of the mixture is classical fluctuations in the preparation apparatus, then shot-by-shot pure states enter the circuit and the consistency relation holds accordingly shot-by-shot, resulting in a proper mixture at the output. This shows that in the presence of a CTC it would be possible to identify the origin of the decoherence in an experimental setup.

Furthermore, careful analysis of the decoherence of the unitary gate U reveals parallels to effects seen in non-local state preparation. The decoherence is assumed to arise from non-local coupling to the environment. Again, due to a lack of classical knowledge of the outcome of an eventual measurement of the environment, ρ_{CTC} “sees” the mixed process in Eq. (7) in every run of the experiment. In the case of full decoherence the distinguishabil-

ity is reduced to 0.5 as in standard quantum mechanics. The differences between local and non-local decoherence in their interpretation and effect is one of the key insights from our simulation.

DISCUSSION

Quantum simulation is a versatile and powerful tool for investigating quantum systems that are hard or even impossible to access in practice [20]. Although no CTCs have been discovered to date, quantum simulation nonetheless enables us to study their unique properties and behaviour. Here we simulated the immediate adaptation of ρ_{CTC} to changes in the CTC’s environment and in particular the effect of different forms of decoherence. We also show that the non-linearity inherent in the system is in fact not uniform as shown in Fig. 3, suggesting that non-linear effects only become apparent in certain scenarios and for a specific set of measurements.

Moreover, we find intriguing differences with respect to nominally equivalent ways of pure state preparation. Although acknowledged in Ref. [29] this feature has not been further investigated in the present literature. Importantly this effect arises due to consistency with relativity, in contrast to the similar effect for mixed quantum states discussed earlier, which is a direct result of the non-linearity of the system [7].

Our study of the Deutsch model provides insights into the role of causal structures and non-linearities in quantum mechanics, which is essential for an eventual reconciliation with general relativity.

Acknowledgments

We thank Nathan Walk and Nicolas Menicucci for insightful discussions. We acknowledge financial support from the ARC Centres of Excellence for Engineered Quantum Systems (CE110001013) and Quantum Computation and Communication Technology (CE110001027). A.G.W. and T.C.R. acknowledge support from a UQ Vice-Chancellor’s Senior Research Fellowship.

Author contributions

M.R., M.A.B., C.R.M. and T.C.R. developed the concepts, designed the experiment, analysed the results and wrote the paper. M.R. performed the experiments and analysed data. T.C.R. and A.G.W. supervised the project and edited the manuscript.

Supplemental Material

DISTINGUISHING (MIXED) QUANTUM STATES

The measure \mathcal{L} introduced in Eq. (4) has an operational interpretation as the probability of obtaining the outcome “different” when comparing two quantum states by means of a single projective measurement on each system. Notably, the minimum-error measurement for discriminating two quantum states is indeed a projective measurement in a direction that depends on the two states [30–32]. Hence, considering only projective measurements is not a restriction and the measure is optimal with the right choice of measurement direction. This result in particular also holds for mixed quantum states, which will become very relevant in the next section.

The situation in the main text can be recast as a game where Alice prepares two quantum systems, one in state $|\psi_0\rangle$ and one in state $|\psi_1\rangle$ and sends them to Bob, whose task is to determine whether they are different or not. If Alice indeed sends two different states, then the measure \mathcal{L} is understood as Bob’s probability of either guessing both states correctly or both incorrectly, which are the two cases where he successfully distinguishes the states. Hence, \mathcal{L} is a natural measure for this task and—given that Bob uses the knowledge about the states to be distinguished—also optimal. In fact, given an optimal choice of measurement direction, \mathcal{L} is directly related to the trace-distance metric \mathcal{D} and therefore a similarly suitable measure of distinguishability:

$$\mathcal{L} = p_{\text{correct}}^2 + p_{\text{error}}^2 = 1 - \frac{1}{2}|\langle\psi_0|\psi_1\rangle|^2 = \frac{1}{2}(1 + \mathcal{D}^2)$$

Optimal CTC implementation

In the main text we investigated the case where the controlled unitary CU_{XZ} is chosen non-optimally for the state $|\psi_1\rangle$. Notably, this is considered a conscious choice of the experimenter, in contrast to decoherence of the gate, which is beyond their control. Hence, the knowledge about θ_{XZ} is available and can be used to optimize the measurement direction of the final projective measurement as is done in the case of standard quantum mechanics. Although in the non-optimal CTC-case the output states are mixed, this does not change the fact that the optimal measurement is projective. Hence, the quantity \mathcal{L} can be optimized depending on the states $|\psi_0\rangle, |\psi_1\rangle$ and the gate CU_{XZ} . The corresponding results are shown in Fig. S1, which differs from Fig. 5 in that the CTC case is now optimal for the chosen gate.

Note that now the CTC-circuit always achieves the best distinguishability of 1/2 for non-local state preparation. In the local case the advantage over standard

quantum mechanics is extended to a wider range of non-optimal combinations. Furthermore, we observe recovery of distinguishability for combinations far from optimal, see Fig. S1.

State identification

As an alternative approach we consider a scenario where Alice prepares two known quantum states $|\psi_0\rangle$ and $|\psi_1\rangle$ at random and sends them—one at a time—to Bob, who is given the task of identifying each of the states. Similarly to the state-discrimination case, the optimal measurement is a projective measurement in a direction that depends on the two states. The figure of merit that is intrinsically related to this task is the probability of success,

$$p_{\text{succ}} = p_{|\psi_0\rangle}p(\psi_0 | |\psi_0\rangle) + p_{|\psi_1\rangle}p(\psi_1 | |\psi_1\rangle),$$

where $p_{|\psi\rangle}$ is the probability for the state $|\psi\rangle$ to be sent and $p(\phi | |\psi\rangle)$ is Bob’s probability for guessing $|\phi\rangle$ in the case where he received the state $|\psi\rangle$. The optimal measurement direction can again be chosen based on knowledge of the two states to be identified. In the scenario considered here, this information is available to Bob and the states are prepared with equal probability. The optimal probability of success is then given by

$$p_{\text{succ}} = \frac{1}{2}(1 + \sqrt{1 - |\langle\psi_0|\psi_1\rangle|^2}) = \frac{1}{2}(1 + \mathcal{D}).$$

Hence, in this case p_{succ} is also directly related to \mathcal{D} , making it an equivalent measure of distinguishability. We have analyzed our experiment from this point of view and find the same qualitative behavior: the CTC circuit outperforms standard quantum mechanics for the optimally chosen unitary interaction, as well as for a range of non-optimal choices. The results are shown in Fig. S2, which parallels Fig. 5.

Note that in contrast to Fig. 5, the CTC-circuit always achieves, but never surpasses 1/2 in the case of non-locally prepared states.

References

- [1] M. S. Morris, K. S. Thorne, and U. Yurtsever, Wormholes, Time Machines, and the Weak Energy Condition, *Physical Review Letters*, **61**, 1446-1449, (1988).
- [2] M. S. Morris and K. S. Thorne, Wormholes in spacetime and their use for interstellar travel: A tool for teaching general relativity, *American Journal of Physics*, **56**, 395, (1988).
- [3] J. R. Gott, Closed timelike curves produced by pairs of moving cosmic strings: Exact solutions, *Physical Review Letters*, **66**, 1126-1129, (1991).




FIG. S1: **Experimental results for the state discrimination scenario with an optimal CTC implementation.** Probability of state discrimination for **a)** locally prepared and **b)** non-locally prepared states $|\psi_0\rangle=|H\rangle$ and $|\psi_1\rangle = \cos(\frac{\phi}{2})|H\rangle + \sin(\frac{\phi}{2})|V\rangle$ as measured by \mathcal{L} . The surface represents the theoretically predicted probability for the optimal CTC implementation, depending on the state and gate parameters ϕ and θ_{XZ} , respectively. Solid, red (open, blue) data-points indicate better (worse) performance than standard quantum mechanics, which also implements the optimal measurement, making use of all the available information. **c)** Cross-sectional views of the combined plots **a)** and **b)** as in Fig. 5. Error bars obtained from a Monte Carlo routine simulating the Poissonian counting statistics are too small to be visible on the scale of this plot.

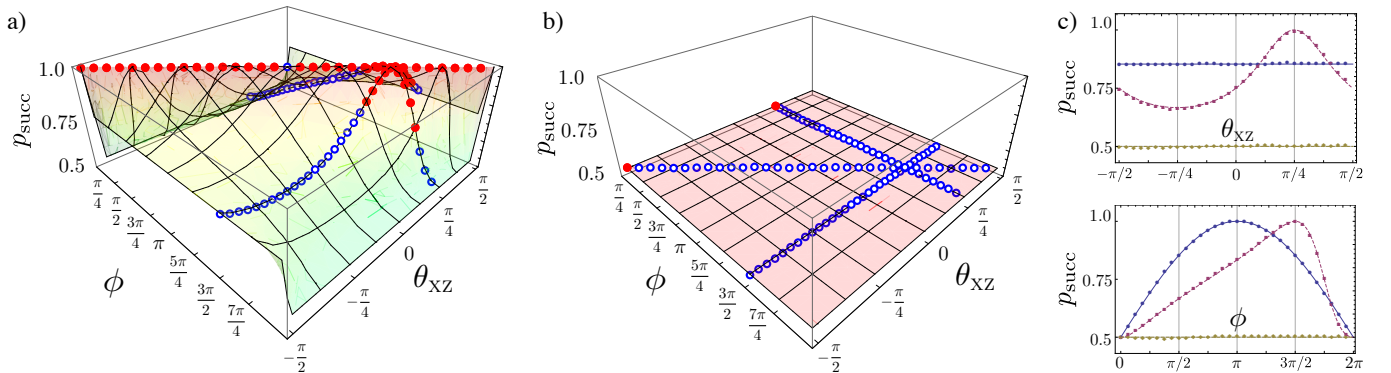


FIG. S2: **Experimental results for the state identification scenario.** Probability of state identification p_{succ} for **a)** locally prepared and **b)** non-locally prepared states $|\psi_0\rangle=|H\rangle$ and $|\psi_1\rangle = \cos(\frac{\phi}{2})|H\rangle + \sin(\frac{\phi}{2})|V\rangle$. The surface represents the theoretically predicted probability depending on the state and gate parameters ϕ and θ_{XZ} , respectively. Solid, red (open, blue) data-points indicate better (worse) performance than standard quantum mechanics, which implements the optimal measurement, making use of all the available information. **c)** Cross-sectional views of the combined plots **a)** and **b)** as in Fig. 5. Error bars obtained from a Monte Carlo routine simulating the Poissonian counting statistics are too small to be visible on the scale of this plot.

- [4] I. Novikov, *Evolution of the Universe* (Cambridge University Press, Cambridge, England, 1983).
- [5] S. W. Hawing, Chronology protection conjecture, *Physical Review D*, **46**, 603-611, (1992).
- [6] D. Deutsch, Quantum mechanics near closed timelike lines, *Physical Review D*, **44**, 3197-3217, (1991).
- [7] T. C. Ralph and C. R. Myers, Information Flow of quantum states interacting with closed timelike curves, *Physical Review A*, **82**, 062330, (2010).
- [8] D. Bacon, Quantum computational complexity in the presence of closed timelike curves, *Physical Review A*, **70**, 032309, (2004).
- [9] T. A. Brun, J. Harrington, and M. M. Wilde, Localized Closed Timelike Curves Can Perfectly Distinguish Quantum States, *Physical Review Letters*, **102**, 210402, (2009).
- [10] D. Ahn, C. R. Myers, T. C. Ralph, and R. B. Mann, Quantum state cloning in the presence of a closed time-like curve, *Physical Review A*, **88**, 022332, (2013).
- [11] T. A. Brun, M. M. Wilde, and A. Winter, Quantum state cloning using Deutschian closed timelike curves, *Physical Review Letters*, **111**, 190401, (2013).
- [12] J. L. Pienaar, T. C. Ralph, and C. R. Myers, Open Timelike Curves Violate Heisenberg's Uncertainty Principle, *Physical Review Letters*, **110**, 060501, (2013).
- [13] C. Bennett, D. Leung, G. Smith, and J. A. Smolin, Can Closed Timelike Curves or Nonlinear Quantum Mechanics Improve Quantum State Discrimination or Help Solve Hard Problems?, *Physical Review Letters*, **103**, 170502, (2009).
- [14] W. Kłobus, A. Grudka, and A. Wójcik, Comment on "Information flow of quantum states interacting with closed timelike curves", *Physical Review A*, **84**, 056301, (2011).
- [15] T. C. Ralph and C. R. Myers, Reply to "Comment on 'Information flow of quantum states interacting with closed timelike curves'", *Physical Review A*, **84**, 056302, (2011).

- (2011).
- [16] T. Kitagawa, M. A. Broome, A. Fedrizzi, M. S. Rudner, E. Berg, I. Kassal, A. Aspuru-Guzik, E. Demler, and A. G. White, Observation of topologically protected bound states in photonic quantum walks, *Nature Communications*, **3**, 882, (2012).
- [17] X.-S. Ma, B. Dakic, S. Kropatschek, W. Naylor, Y.-H. Chan, Z.-S. Gong, L.-M. Duan, A. Zeilinger, and P. Walther, Quantum simulation of the wavefunction to probe frustrated Heisenberg spin systems, *Nature Physics*, **7**, 399-405, (2011).
- [18] J. Simon, W. S. Bakr, R. Ma, M. E. Tai, P. M. Preiss, and M. Greiner, Quantum simulation of antiferromagnetic spin chains in an optical lattice, *Nature*, **472**, 307-312, (2011).
- [19] R. Gerritsma, G. Kirchmair, F. Zähringer, E. Solano, R. Blatt, and C. F. Roos, Quantum simulation of the Dirac equation, *Nature*, **463**, 68-71, (2010).
- [20] J. Casanova, C. Sabín, J. León, I. L. Egusquiza, R. Gerritsma, C. F. Roos, J. J. García-Ripoll, and E. Solano, Quantum Simulation of the Majorana Equation and Unphysical Operations, *Physical Review X*, **1**, 021018, (2011).
- [21] T. G. Philbin, C. Kuklewicz, S. Robertson, S. Hill, F. König, and U. Leonhardt, Fiber-Optical Analog of the Event Horizon, *Science*, **319**, 1367-1370, (2008).
- [22] N. C. Menicucci, S. Jay Olson, and G. J. Milburn, Simulating quantum effects of cosmological expansion using a static ion trap, *New Journal of Physics*, **12**, 095019, (2010).
- [23] S. Lloyd, L. Maccone, R. Garcia-Patron, V. Giovannetti, Y. Shikano, S. Pirandola, L. Rozema, A. Darabi, Y. Soudagar, L. Shalm, et al., Closed Timelike Curves via Postselection: Theory and Experimental Test of Consistency, *Physical Review Letters*, **106**, 040403, (2011).
- [24] S. Lloyd, L. Maccone, R. Garcia-Patron, V. Giovannetti, Y. Shikano, Quantum mechanics of time travel through post-selected teleportation, *Physical Review D*, **84**, 025007, (2011).
- [25] T. A. Brun, M. M. Wilde, Perfect State Distinguishability and Computational Speedups with Postselected Closed Timelike Curves, *Foundations of Physics*, **42**, 341-361, (2011).
- [26] T. C. Ralph, N. Langford, T. Bell, and A. G. White, Linear optical controlled-NOT gate in the coincidence basis, *Physical Review A*, **65**, 062324, (2002).
- [27] N. K. Langford, T. J. Weinhold, R. Prevedel, K. J. Resch, A. Gilchrist, J. L. O'Brien, G. J. Pryde, and A. G. White, Demonstration of a Simple Entangling Optical Gate and Its Use in Bell-State Analysis, *Physical Review Letters*, **95**, 210504, (2005).
- [28] B. d'Espagnat, *Conceptual Foundations of Quantum Mechanics*, 2nd ed. (Addison Wesley, 1976).
- [29] E. G. Cavalcanti, N. C. Menicucci, and J. L. Pienaar, The preparation problem in nonlinear extensions of quantum theory, Preprint at <http://arxiv.org/abs/1206.2725> (2012).
- [30] C. W. Helstrom, Quantum detection and estimation theory, *Journal of Statistical Physics*, **1**, 231-252, (1969).
- [31] U. Herzog and J. Bergou, Distinguishing mixed quantum states: Minimum-error discrimination versus optimum unambiguous discrimination, *Physical Review A*, **70**, 022302, (2004).
- [32] G. Jaeger and A. Shimony, Optimal distinction between two non-orthogonal quantum states, *Physics Letters A*, **197**, 83-87, (1995).

APPENDIX C

Generation of Mechanical Interference Fringes by Multi-Photon Quantum Measurement

This appendix reproduces a paper of which I am the lead author, and which is available at [arXiv:1602.05955](https://arxiv.org/abs/1602.05955).

Generation of Mechanical Interference Fringes by Multi-Photon Quantum Measurement

M. Ringbauer,^{1,2} T. J. Weinhold,^{1,2} A. G. White,^{1,2} M. R. Vanner^{1,2,3}

¹*Centre for Engineered Quantum Systems, ²Centre for Quantum Computer and Communication Technology, School of Mathematics and Physics, University of Queensland, Brisbane, QLD 4072, Australia,*
³*Clarendon Laboratory, Department of Physics, University of Oxford, OX1 3PU, United Kingdom*

The exploration of wave phenomena and quantum properties of massive systems offers an intriguing pathway to study the foundations of physics and to develop a suite of quantum-enhanced technologies. Here we present an optomechanical scheme to prepare non-Gaussian quantum states of motion of a mechanical resonator using photonic quantum measurements. Our method is capable of generating non-classical mechanical states without the need for strong single-photon coupling, and is resilient against optical loss and initial mechanical thermal occupation. Additionally, our approach provides a route to generate larger mechanical superposition states using effective interactions with multi-photon quantum states. We experimentally demonstrate this technique on a mechanical thermal state in the classical limit and observe interference fringes in the mechanical position distribution that show phase super-resolution. This opens a feasible route to explore and exploit quantum phenomena at a macroscopic scale.

INTRODUCTION

Generating quantum superposition states in macroscopic systems is an important goal in experimental quantum science. Studying such states will allow us to probe the limits of applicability of quantum mechanics and to harness quantum physics for new technologies. Early evidence for quantum phenomena with massive systems was provided by electron diffraction experiments [1] and, through the efforts of the last nine decades, quantum-matter-wave behaviour has been observed for neutrons [2], trapped ion systems [3], ultracold atoms [4], and even molecules comprising many hundreds of atoms [5]. A promising route to explore quantum behaviour on an even more macroscopic scale is provided by quantum optomechanics [6] where a mechanical oscillator interacts with an optical field via radiation pressure. This versatile quantum-optical platform enables tests of fundamental physics [7–9], the development of microwave-to-optical interfaces [10], and high-precision weak-force sensors [11]. Indeed, the LIGO gravitation-wave antenna can be considered a large-scale optomechanical system [12, 13]. Recently, impressive progress has been made using both opto- and electro-mechanical systems with examples including single-phonon-level operations [14–17], quantum coherent coupling [18], mechanically-induced squeezing of light [19, 20], and even opto-mechanical entanglement [21]. Experimental efforts continue in a diverse set of directions, however, progress is hindered by three main factors: weak single-photon coupling, sensitivity to optical loss, and mechanical decoherence. The approach we introduce here exploits measurement-induced non-linearities [22] to overcome the challenges of weak coupling and optical loss. In addition, our approach offers resilience against initial thermal occupation and can also be readily employed in cryogenic systems to overcome the challenge of thermal decoherence.

To understand radiation pressure consider the reflection of a single photon from a mechanical resonator. The reflection imparts a momentum—inversely proportional to the wavelength—onto the resonator and concurrently the optical field acquires a phase-shift—proportional to the mechanical displacement. This momentum, however, is typically very small compared to the quantum noise of the resonator due to weak optomechanical coupling. In order to enhance the strength of this interaction, experimental efforts often employ an optical cavity to increase the number of reflections [6]. Utilising such cavity enhancement, the seminal works of Bose, Jacobs, and Knight [7]; and Marshall *et al.* [8], proposed using a superposition of the optical vacuum and a single photon to generate optomechanical entanglement with the motion of a mechanically oscillating mirror forming part of a Fabry-Perot cavity. In this case, the mechanical resonator is subject to a quantum superposition of the identity operation (no photon present) and a displacement operation (single photon present), thus generating a mechanical Schrödinger-cat state.

Our method enables the generation of mechanical superposition states without the need for non-classical optical input states such as single photons. Rather, we use an optomechanical interaction for a time much shorter than the mechanical period [23] with a weak optical coherent state and then project the reflected field, via photon counting, onto a superposition of zero and one photon to conditionally generate the superposition. In contrast to Refs. [7, 8], the measurement used in our scheme leaves the mechanics in a single-mode superposition state, which is achieved by projection from a more easily prepared form of optomechanical correlations. A single-mode quantum superposition state—with components separated in the momentum quadrature—will show interference fringes in the position probability distribution with a frequency inversely proportional to the superposition separation. When the photon detection registers

the required event to prepare such a state an independent readout beam is then used to verify and characterise the mechanical state. The process of our scheme is non-classical and generates quantum states of motion independent of the coupling strength and for finite initial thermal occupation. Moreover, the size of the generated mechanical superposition state can be made larger by projection onto an optical $N00N$ -state [24–26], which can be implemented using coherent state inputs, multi-port interferometry, and multi-photon coincidence measurements [27]. Unlike preparation of large $N00N$ states, projection onto these states is experimentally simple. We implement this process experimentally and observe the first mechanical interference fringe pattern and the predicted phase-superresolution for a two-photon coincidence measurement. In this proof-of-concept experiment we observe a high-visibility fringe pattern in the mechanical position but we do not observe any non-classical features due to quantum decoherence and low readout sensitivity. Generating and observing single-mode non-classicality, such as Wigner negativity, of a massive mechanical oscillator remains an outstanding challenge. Our measurement-based-scheme offers a promising route to achieving this important goal and can be readily applied to a number of optomechanical systems beyond the system used here. This experimental technique provides a powerful platform to empirically explore open-quantum-system dynamics, test potential collapse models of the wavefunction [28–30], and enable the development of quantum-enhanced weak force sensors.

RESULTS

Mechanical state preparation via photon detection— For simplicity we first describe our scheme for two-port photon counting before generalizing to multi-port quantum measurement. The two-port case describes our experimental results and can be implemented with Mach-Zehnder-type interferometers. Figure 1A shows a conceptual model of our experiment, where a weak coherent state is injected and interacts with a mechanical resonator in one arm of an interferometer. The two optical fields inside the interferometer then interfere on a beam-splitter and photon counting is performed on the two output ports. A single-photon click on one of the detectors effectively projects the optical state that was inside the interferometer onto the path-entangled number state, i.e. $(|10\rangle + |01\rangle)/\sqrt{2}$, which is the state used in Refs. [7, 8]. If, instead, both detectors register a single photon then, due to second-order quantum interference, the optical state that interacted with the mechanical resonator within the interferometer must have been a 2-photon $N00N$ -state $(|20\rangle - |02\rangle)/\sqrt{2}$. In this case, the mechanical oscillator was subject to a superposition of the identity operation and a two-photon radiation-pressure displacement, thus enhancing the size of the superposition by a factor of two compared with the single photon detection case.

This operation can be conveniently described using the measurement- (or Kraus-) operator approach, which allows us to conveniently compute the mechanical state after a measurement as well as the measurement outcome probabilities. Consider a mechanical resonator in the pure initial state $|\psi_{\text{in}}\rangle_{\text{M}}$ and coherent states $|\alpha\rangle_1|\alpha\rangle_2$ in the interferometer arms 1 and 2, respectively (i.e. the state after the initial 50/50 beamsplitter in Fig. 1A). The state immediately after the interaction with mode 1 of the interferometer is then given by $e^{i\mu a_1^\dagger a_1 X_{\text{M}}}|\alpha\rangle_1|\alpha\rangle_2|\psi_{\text{in}}\rangle_{\text{M}}$, where, $\mu = 4\pi x_0/\lambda$ is the momentum transfer per photon in units of the mechanical quantum noise. Here, $x_0 = \sqrt{\hbar/m\omega}$ is the mechanical ground state size (m ; effective mass, ω ; mechanical angular frequency), $X_{\text{M}} = x/x_0$ is the mechanical position operator in units of x_0 , and $a_{1,2}$ are the annihilation operators for the interferometer arms 1 and 2. The two optical fields then interfere on a beam-splitter, via $B^\dagger a_1 B = t a_1 + r a_2$ and $B^\dagger a_2 B = r a_1 - t a_2$, where r^2 and t^2 denote the intensity reflectivity and transmittivity, respectively. The mechanical state after this interaction and photon measurement is $|\psi_{\text{out}}\rangle_{\text{M}} \propto \sum_n \langle n|_1 \langle m|_2 B e^{i\mu a_1^\dagger a_1 X_{\text{M}}} |\alpha\rangle_1 |\alpha\rangle_2 |\psi_{\text{in}}\rangle_{\text{M}}$, where m and n denote the number of photons detected in mode 1 and 2 after the beam-splitter, respectively. Assuming a 50:50 beam-splitter and including a static phase-shift ϕ in mode 2 of the interferometer we can introduce the measurement operator,

$$\Upsilon = \frac{e^{-|\alpha|^2}}{\sqrt{m!n!}} \frac{\alpha^{m+n}}{(\sqrt{2})^{m+n}} (e^{i\mu X_{\text{M}}} + e^{i\phi})^m (e^{i\mu X_{\text{M}}} - e^{i\phi})^n. \quad (1)$$

The mechanical output state can now be written $|\psi_{\text{out}}\rangle_{\text{M}} \propto \Upsilon |\psi_{\text{in}}\rangle_{\text{M}}$. This operation corresponds to a superposition of a mechanical displacement $e^{i\mu X_{\text{M}}}$ and the identity operation with a controllable phase $e^{i\phi}$. More generally we may compute the mechanical state via $\rho_{\text{M}}^{\text{out}} = \Upsilon \rho_{\text{M}}^{\text{in}} \Upsilon^\dagger / \mathcal{P}$, where $\rho_{\text{M}}^{\text{in},\text{out}}$ are the input and output mechanical density matrices, respectively, and $\mathcal{P} = \int_{-\infty}^{\infty} dX_{\text{M}} \Upsilon^\dagger \Upsilon \langle X_{\text{M}} | \rho_{\text{M}}^{\text{in}} | X_{\text{M}} \rangle$ is the probability for obtaining the photon counting outcomes m and n , see Supplementary Materials for further details.

To examine the mechanical position probability distribution of the state after the interaction and measurement we may write $\langle X_{\text{M}} | \rho_{\text{M}}^{\text{out}} | X_{\text{M}} \rangle \propto \Upsilon^\dagger \Upsilon \langle X_{\text{M}} | \rho_{\text{M}}^{\text{in}} | X_{\text{M}} \rangle$. The function $\Upsilon^\dagger \Upsilon$ is oscillatory in X_{M} , as obtaining a click in our interferometer gives periodic information about X_{M} , and can be interpreted as a filter acting on the initial mechanical position probability distribution. This oscillatory behaviour is intrinsically linked with the cubic nature of the full optomechanical radiation-pressure interaction $a_1^\dagger a_1 X_{\text{M}}$ and the non-linearity of photon counting. This allows our scheme to generate non-Gaussian states of motion, which would not be possible in the more commonly considered linearized regime with quadratic interactions and linear measurements.

Experimental Setup— In our experiment (Fig. 1B) we use a high-stress 1.7×1.7 mm Si_3N_4 membrane [32] em-

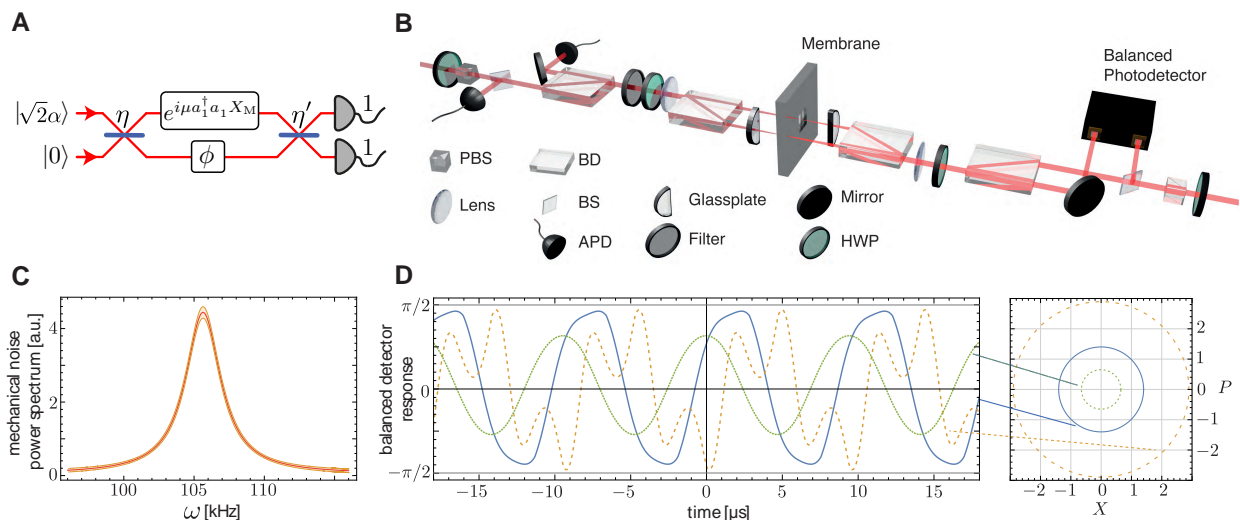


FIG. 1. **Experimental scheme.** (A) A mechanical resonator interacts with a weak optical coherent state inside an interferometer with static phase-shift ϕ formed by two beam splitters with reflectivities η and η' . The optomechanical radiation-pressure interaction is described via the unitary operation $e^{i\mu a_1^\dagger a_1 X_M}$, where a_1 describes the optical field operator, and X_M describes the mechanical position. Mechanical interference fringes are generated via photon counting on the two interferometer outputs, which projects the field inside the interferometer onto a path-entangled photon number state. (B) In our experiment the two interferometer paths in A are represented by orthogonal polarizations of a weak coherent-state, and a half-wave plate (HWP) acts as a tunable beam splitter. For the interaction with the mechanical resonator—a SiN membrane—the polarization modes are split into distinct optical paths using a calcite beam displacer (BD). One of the beams reflects off the membrane, while the other one reflects off the adjacent frame and acquires a static phase shift ϕ , controlled by the yaw degree of freedom of the BD. The modes are then interfered, separated, and detected by two single-photon detectors (APD). A position measurement of the membrane is performed from the other side of the membrane using a similar setup and an independent readout beam. Lenses focus the preparation and readout beams to a spot size of $\sim 50 \mu\text{m}$ on the membrane and glass plates are used to compensate for shifts in the foci due to the birefringence in the BDs. (C) Noise power spectrum of the fundamental vibrational drum mode of the mechanical resonator with resonance frequency of $\omega_M/2\pi = 105.64 \text{ kHz}$. (D) Balanced-detector time-traces used to measure the mechanical state for three example phase space points (shown on the right) with displacements of 103 nm (dotted, green), 222 nm (solid blue) and 458 nm (dashed, orange), respectively. Here X and P are the mechanical position and momentum quadratures, respectively, in units of the interferometer readout range $\lambda_R/4 = 158.2 \text{ nm}$ (corresponding to the turning-points of the interferometer).

bedded in a $10 \times 10 \text{ mm}$ Si-frame. The membrane has a thickness of $50 \pm 2.5 \text{ nm}$ and at our state-preparation-field wavelength of 795 nm has a measured reflectivity of $23.0 \pm 0.5 \%$, while the frame has a reflectivity of $20.5 \pm 0.2 \%$. The noise-power spectrum of the fundamental drum mode at $\omega_M/2\pi = 105.64 \pm 0.02 \text{ kHz}$ is shown in Fig. 1C. At room temperature and at atmospheric pressure the mechanical line-width (FWHM) was measured to be $\delta\omega_M/2\pi = 3.10 \pm 0.05 \text{ kHz}$ and the effective mass is on the order of 100 ng , which comprises approximately 10^{16} atoms. In order to probe the regime where the optomechanical phase shifts are large, i.e. $\mu^2 \langle X_M^2 \rangle \gtrsim 1$, we use a ring-piezo to drive the membrane motion.

The membrane is mounted in a way that allows optical access from both sides, and forms the central part of two folded Mach-Zehnder interferometers (MZI), see Fig. 1B. One interferometer is used for mechanical state preparation with photon counting, as illustrated in Fig. 1A, while the other is used for mechanical position readout using a balanced detector and a $\sim 100 \mu\text{W}$ laser at a wavelength

of 632.8 nm . Our setup employs a compact polarization interferometer design that does not require active phase stabilisation [33]. The two arms of the MZI in Fig. 1A are represented by orthogonal polarizations and the role of the beamsplitter is achieved by a half-wave plate (HWP), which allows for precise control of the splitting ratio. For the interaction with the membrane the two polarizations are separated using a calcite beam-displacer and recombined after reflection from the mechanical device.

An APD click—either one of the detectors or a coincidence event within a 7.8 ns window—triggers the balanced detector and recording of a $50 \mu\text{s}$ long trace at a sampling rate of 100 MS/s , see Fig. 1D. The mechanical quadratures X and P , defined in units of the interferometer readout range ($\lambda_R/4$), are then extracted from a fit to this time trace. For small mechanical displacements the time trace is almost sinusoidal, but becomes overmodulated as the resonator displacement surpasses $\sim 100 \text{ nm}$. In addition, for larger mechanical displacements we observe a mechanical-position-dependent am-

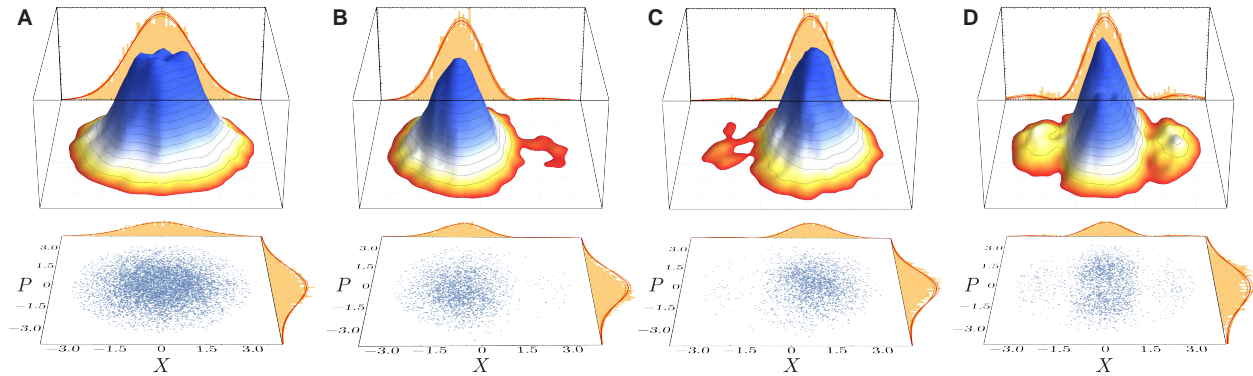


FIG. 2. **Observed mechanical states of motion showing interference fringes.** Each plot shows the measured phase-space points (bottom), together with the corresponding mechanical position and momentum quadrature histograms, and a normalized 3D probability-density histogram of these points (top), with the mechanical position quadrature histogram reproduced again for easy visualization. The quadrature histograms contain a fit (solid, red line) using the theory model and corresponding 3σ mean prediction bands (shaded orange). All axes are in units of the interferometer readout range. (A) Measured initial Gaussian thermal state. (The truncation at the origin is due to finite resolution of our data acquisition.) (B) Conditional motional state prepared via $\{m, n\} = \{0, 1\}$ detection, which is shifted in phase-space along $-X$ and shows the start of a second peak in the right tail of the Gaussian envelope. (C) Conditional state for $\{m, n\} = \{1, 0\}$ detection, which is as B, but shifted along $+X$. Again, note the second peak, now on the left. (D) Two photon ($\{m, n\} = \{1, 1\}$) detection generates a mechanical fringe pattern in X with twice the frequency of the single photon detection cases due to super-resolution of the measurement. Three maxima of the fringe pattern are observed. In all of these cases the interferometer phase ϕ was set to $\pi/2$. Also note that the momentum quadrature remains close to the initial distribution for all of these measurements.

plitude modulation of the interferometer signal. We have taken the first order corrections due to this amplitude modulation into account when fitting the time traces, see Supplementary Material. For each type of click event we record ~ 3000 such time traces to create a phase-space histogram of the mechanical motion. By using a combination of spectral and polarization filters the read-out beam transmitted through the membrane is suppressed below the dark-count level of our single-photon detectors of ~ 150 Hz.

Large phase-shift regime— Figure 2 shows the measured mechanical phase-space distributions prepared via one- and two-photon measurements on a piezo-driven initial Gaussian thermal-state (Fig. 2A) with RMS position fluctuations of 198 ± 2 nm. This corresponds to the regime of large optical phase-shifts, i.e. $\mu^2 \langle X_M^2 \rangle \gtrsim 1$. For single-photon detection, low frequency fringes are observed with a π phase-shift between the detection events $\{m, n\} = \{0, 1\}$ (Fig. 2B) and $\{1, 0\}$ (Fig. 2C). Moreover, we observe the start of the second fringe peak in the tails of the Gaussian envelope, which is on the right in Fig. 2B and the left in Fig. 2C. In the case of a two-photon detection event $\{m, n\} = \{1, 1\}$ (Fig. 2D), the mechanical resonator interacted with an effective two-photon $N00N$ state. Consequently, we observe phase super-resolution in the mechanical interference pattern at twice the fringe-frequency of the single photon cases. We would like to highlight that this work goes significantly beyond previous all-optical schemes observing phase super-resolution via multi-photon projection [27, 31] as our scheme maps the fringe pattern into the state of another *bosonic mode* rather than being a modulation to the photon count rate

with a *scalar* phase quantity. Thus, the super-resolution achieved by such projections is used here as a resource for state preparation. In a quantum regime, the fringe pattern observed can then be interpreted as either the quantum interference in the superposition state or from the filter of the quantum measurement. This measurement-based technique provides a considerable advantage for ultimately generating non-classical states of mechanical motion and can be employed in other quantum optical systems. Note that all conditional states (Fig. 2B-D) feature interference fringes in the position distribution, while the momentum quadrature remains as the initial Gaussian distribution. The conditional mechanical states shown here were prepared with the phase set to $\phi = \pi/2$, which gives a fringe maxima in the centre of the distribution for the $\{1, 1\}$ event.

We would like to additionally note here that our method can be utilised to determine the optomechanical coupling strength μ by fitting to the fringe pattern observed. This technique requires a well calibrated position axis in units of the mechanical ground state size and can be performed for any mechanical thermal occupation.

Small phase-shift regime— Our scheme can also generate non-Gaussian states of motion in the regime of small optomechanical phase shifts, i.e. $\mu^2 \langle X_M^2 \rangle \ll 1$. Indeed, mechanical non-classicality, in the form of Wigner negativity, can be generated independent of the coupling strength, providing a promising route to generate and explore macroscopic mechanical quantum states even for systems with weak single-photon-coupling. For small μ , applying our scheme to the mechanical ground state for $\{m, n\} = \{0, 1\}$ with $\phi = 0$ yields $(e^{i\mu X_M} - 1)|0\rangle \simeq$

$(i\mu X_M)|0\rangle = i\mu 2^{-1/2}|1\rangle$. Here, $|1\rangle$ denotes a single phonon Fock state and the mechanical position operator in terms of the phonon annihilation (b) and creation (b^\dagger) operators is $X_M = 2^{-1/2}(b + b^\dagger)$. This conditional state is a result of quantum interference and has no classical description [34–36]. Note that in this regime the detection event $\{m, n\} = \{1, 1\}$ generates a similar state. The Wigner function of this mechanical single-phonon Fock state and its two conjugate quadrature distributions are shown in Fig. 3B. A mechanical Fock state has a rotationally invariant distribution in phase-space and has the quadrature distribution $\text{Pr}(X_M) = 2\pi^{-1/2}X_M^2 e^{-X_M^2}$ for all quadrature angles. For $\mu^2 \langle X_M^2 \rangle \ll 1$ the filter for these detection events is $\Upsilon^\dagger \Upsilon \propto \mu^2 X_M^2$. Thus, we note that the position probability distribution after this operation has the same form as a Fock state for any thermal occupation, not just the ground state.

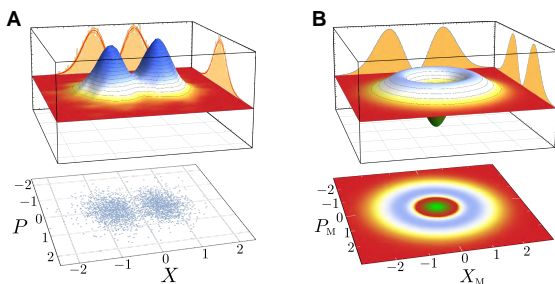


FIG. 3. **Non-Gaussian mechanical states prepared with weak single-photon coupling.** (A) Experimentally observed mechanical phase-space distribution for the detection event $\{m, n\} = \{1, 1\}$ with $\phi = 0$ in the weak coupling and weak drive regime, in units of the interferometer read-out range. The position probability distribution here has the same form as that of a single phonon Fock state. (B) Theoretical Wigner function of single phonon Fock state, plotted in units of the mechanical quantum noise. This state would be obtained when applying our scheme to the mechanical ground state for the same parameters as in Fig. A. Note that the Fock state has a non-Gaussian momentum distribution whereas our data has a Gaussian momentum distribution due to the thermal initial state.

To observe this type of fringe pattern we use the detection event $\{m, n\} = \{1, 1\}$ on a mechanical thermal state with RMS position fluctuations of 91 ± 1 nm. The observed mechanical phase-space distribution and quadrature distributions are shown in Fig. 3A. The measured position probability distribution is in good agreement with the theoretical prediction. Note that, for this large thermal excitation, however, the momentum quadrature probability distribution remains Gaussian.

Generalisation to larger N00N states— Our scheme can be extended to generate mechanical superposition states with increasing separation size by only changing the optical measurement. This extension would require two coherent states, one of which interacts with the mechanical system, while the other acquires a static phase

shift, together with $N - 2$ vacuum ancilla modes injected into an optical N -port interferometer, see Fig. 4A. N -fold single-photon detection at the output of the N -port then projects the optical field so that the mechanical resonator effectively interacted with an optical $N00N$ -state. The mechanical resonator is thus subject to a superposition of a radiation-pressure force with 0 or N photons thus enhancing the separation in the superposition state (cf. Fig. 4B). See the Supplementary Material for a mathematical description and the heralding probability. Here, the high-frequency fringes of the optical $N00N$ -state are mapped onto the mechanical position probability distribution. Importantly, the states generated by our scheme can exhibit strong negativity of the Wigner quasi-probability distribution independent of the optomechanical coupling strength μ . As shown in Fig. 4C, this negativity approaches zero from below asymptotically with increasing \bar{n} thus the scheme is resilient against initial thermal occupation. Furthermore, we would like to highlight that when using large $N00N$ states this negativity scales with \bar{n}^{-1} . This scaling should be contrasted to the scaling available using single quanta addition [34], which goes with \bar{n}^{-2} , see the supplementary for more details. Moreover, for low amplitude optical coherent states, our scheme is robust against optical loss and inefficiency due to the single photon conditioning.

DISCUSSION

We have introduced a technique that exploits the quantum nature of multi-photon measurements to generate non-classical states of motion of a mechanical resonator. Additionally, we have performed a proof-of-concept experimental demonstration of this technique, which allowed for the first observation of mechanical interference fringes within a thermal distribution. These fringes, albeit at a classical level, have the same qualitative form as those expected from a canonical quantum superposition state, which would be the result of applying our method to a low entropy initial state. This highlights that the appearance of fringes in the position quadrature distribution is not a sufficient condition for non-classical behaviour. Moreover, fringes can be generated by our method independent of the initial thermal occupation, which emphasizes the importance of precise calibration of the position quadrature measurements and characterization of the initial state. While our continuous-wave read-out technique is suitable for the regime experimentally tested here, this technique cannot resolve displacements below the size of the ground state, which is required to observe mechanical non-classicality. Different techniques, such as quantum non-demolition pulsed quadrature measurements [23, 37] will allow this limit to be surpassed and perform mechanical quantum state reconstruction [38]. Looking ahead, one possible experimental approach to generating significant non-classicality with our scheme is to use MHz-frequency oscillators, which

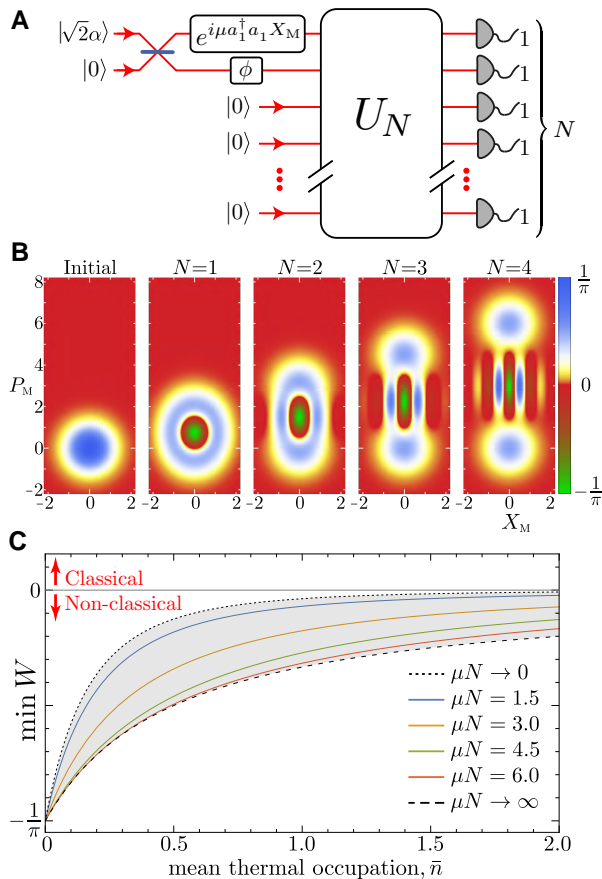


FIG. 4. **Generation of larger superposition states.** (A) Optical setup requiring two input coherent states, a linear optical network U_N , and single photon counters. One of the two input coherent states interacts with the mechanical resonator and then the inputs are projected onto an optical $N00N$ state, which generates a mechanical superposition state with a separation that increases with N , see text. (B) Simulated Wigner functions for an initial mechanical ground state ($\bar{n} = 0$) and superposition states prepared via multiphoton detection. From left to right: initial ground state, and mechanical states prepared using 1-, 2-, 3-, 4-photon $N00N$ -state projections for $\mu = 1.5$ and $\phi = 0$. Note that these mechanical states are generated via superposition of identity and displacement and are thus not symmetric around the origin. (C) Minimum of the Wigner distribution ($\min W$) as a function of the initial thermal occupation \bar{n} for various values of μN . Note that our scheme generates Wigner-negativity even in the limit of weak coupling $\mu N \rightarrow 0$ (dotted, black line) and saturates for large coupling $\mu N \rightarrow \infty$ (dashed, black line). All experimental configurations (i.e. arbitrary values of μ, N and \bar{n}) result in states within the grey shaded area, and always feature Wigner-negativity. Moreover, for arbitrary single-photon coupling strength μ and arbitrary $N00N$ -state size N the generated states achieve the maximum possible negativity of $-1/\pi$, as $\bar{n} \rightarrow 0$. The coloured solid lines correspond to $N = 1, 2, 3, 4$ multi-photon coincidence events, for $\mu = 1.5$, as in Fig.B.

have been experimentally cooled to a thermal occupation of $\bar{n} \simeq 0.34$ [39]. For electro-mechanical systems, one could employ a secondary coupling outside the resolved-sideband regime to a weak optical field or to a microwave field together with superconducting qubits, to implement our protocol. At this thermal occupation, the Wigner negativity generated by our scheme for weak coupling is -0.07 , and in the limit of strong coupling is -0.19 , which should be compared with the maximum value achievable of $-\pi^{-1} \simeq -0.32$.

Our scheme offers four main advantages: (i) Our process generates non-classical mechanical states—signified by negativity in the Wigner quasiprobability distribution—without the need for strong single photon coupling and (ii) with favourable scaling against the initial thermal occupation. (iii) Our scheme is resilient against optical loss, as photon counting with low amplitude input coherent states can be used. And (iv), larger superposition states can be prepared by changing only the measurement and projecting onto optical $N00N$ states. This combination of advantages thus dramatically improves the feasibility to generate and observe mechanical single-mode non-classicality—a key outstanding goal of the field—and the high-visibility mechanical fringes observed here are a key step towards achieving this goal.

ACKNOWLEDGMENTS

We would like to thank C. M. Caves, M. S. Kim, G. J. Milburn, I. Pikovski, J. Schmiedmeyer, and F. Shahandeh for helpful discussions and B. Duffus and T. Vulpecula for experimental assistance. This work was supported by an ARC Discovery Project (DP140101638) and in part by the ARC Centres for Engineered Quantum Systems (CE110001013) and Quantum Computation and Communication Technology (CE110001027). AGW acknowledges support through a UQ Vice-Chancellor’s Research and Teaching Fellowship, and MRV through a UQ Early Career Researcher Grant.

Supplementary Information: Generation of Mechanical Interference Fringes by Multi-Photon Quantum Measurement

Here we provide further details of the theoretical model, the experimental setup, and the data analysis techniques used.

SI. THEORETICAL MODEL

In this section we give a mathematical model to describe mechanical state preparation via interaction with an optical field followed by photon counting as introduced in this work. We first discuss the details of the two-port case, which can be achieved using a standard Mach-Zehnder interferometer configuration. We then generalise this to multi-port interferometry, which can be used to generate larger mechanical superposition states by projection onto an optical $N00N$ state.

A. Two-port model

For optomechanical systems where the interaction time is much shorter than the mechanical period the mechanical free evolution can be neglected, and the radiation-pressure interaction is in this case described by the unitary

$$U = e^{i\mu a_1^\dagger a_1 X_M}. \quad (\text{S1})$$

Here, a_1 is optical annihilation operator in mode 1 ($[a_1, a_1^\dagger] = 1$), and $X_M = (b + b^\dagger)/\sqrt{2}$ is the mechanical position operator in units of the mechanical quantum noise for mechanical annihilation operator b . Defining $X_M = x/x_0$, and $P_M = p/p_0$ (with mechanical position x , and momentum p), we have $[X_M, P_M] = i$ for $[x, p] = i\hbar$ and $x_0 = \sqrt{\hbar/(m\omega_M)}$ and $p_0 = \sqrt{\hbar m\omega_M}$ where m is the effective mass and ω_M is the mechanical angular frequency. The unitary in Eq. (S1) accurately describes our cavity-free experiment, which uses a simple reflection, as well as cavity-based experiments in the regime where the cavity amplitude decay rate κ is much larger than the mechanical resonance frequency.

The optomechanical coupling strength, μ , quantifies the momentum kick per photon in units of p_0 . For the simple case of a single reflection this strength is

$$\mu = 4\pi x_0/\lambda, \quad (\text{S2})$$

where λ is the optical wavelength. This can easily be seen by noting that: (i) the momentum transfer per single photon $p_0\mu = 2\hbar k = 4\pi\hbar/\lambda$, and (ii) that phase shift $\mu X_M = 2kx$, where k is the optical wavenumber. For a cavity-optomechanical system with interaction Hamiltonian $H/\hbar = -g_0 a_1^\dagger a_1 (b + b^\dagger)$ the momentum kick per photon is enhanced linearly by the cavity finesse \mathcal{F}

$$\mu \propto \frac{g_0}{\kappa} = \frac{4\mathcal{F}x_0}{\lambda}. \quad (\text{S3})$$

We will use a Kraus operator approach to determine the state of the mechanical resonator conditioned on a ‘click’-event in the interferometer. The Kraus- or measurement operator also allows the heralding probabilities for these events to be easily computed.

For our two-port case, the Kraus operator is

$$\Upsilon_{m,n} = {}_2\langle n| {}_1\langle m| B_{12} e^{i\mu a_1^\dagger a_1 X_M} |\alpha\rangle_1 |\alpha e^{i\phi}\rangle_2, \quad (\text{S4})$$

where $\alpha \in \mathbb{R}$ denotes the coherent-state size, B_{12} is the beam-splitter operator, and we have included the static phase shift ϕ into the coherent state inside the interferometer. Using a matrix representation, we write the action of a 50:50 beam-splitter as

$$\begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \rightarrow \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \end{pmatrix}. \quad (\text{S5})$$

The measurement operator then becomes

$$\Upsilon_{m,n} = \frac{e^{-\alpha^2}}{\sqrt{m!n!}} \frac{\alpha^{m+n}}{(\sqrt{2})^{m+n}} (e^{i\mu X_M} + e^{i\phi})^m (e^{i\mu X_M} - e^{i\phi})^n. \quad (\text{S6})$$

The state of the mechanical resonator after the interaction and ‘click’-event is determined by

$$\rho_M^{\text{out}} = \Upsilon \rho_M^{\text{in}} \Upsilon^\dagger / \mathcal{P}, \quad (\text{S7})$$

where \mathcal{P} is the heralding probability, which ensures that $\text{Tr}\{\rho_M^{\text{out}}\} = 1$.

Important to this work is observing the mechanical position probability distribution after the interaction and measurement. The operator Υ depends only on the operator X_M and we can thus write

$$\langle X_M | \rho_M^{\text{out}} | X_M \rangle = \frac{1}{\mathcal{P}} \Upsilon^\dagger(X_M) \Upsilon(X_M) \langle X_M | \rho_M^{\text{in}} | X_M \rangle. \quad (\text{S8})$$

Hence, $\Upsilon^\dagger \Upsilon$ can be interpreted as a filter, acting on the position distribution of the resonator

$$\Upsilon^\dagger \Upsilon = \frac{1}{m!n!} e^{-2\alpha^2} \alpha^{2m} \alpha^{2n} \times (1 + \cos(\mu X_M - \phi))^m (1 - \cos(\mu X_M - \phi))^n. \quad (\text{S9})$$

Note that α will only affect the heralding probability and does not change the form of the conditional mechanical state. See Figs. 2 and 3 of the main text for experimental observations of such mechanical interference fringes. The heralding probability \mathcal{P} is given by

$$\begin{aligned} \mathcal{P} &= \text{Tr}\{\Upsilon^\dagger \Upsilon \rho_M^{\text{in}}\} \\ &= \int_{-\infty}^{\infty} dX_M \Upsilon^\dagger \Upsilon \langle X_M | \rho_M^{\text{in}} | X_M \rangle. \end{aligned} \quad (\text{S10})$$

For a mechanical thermal state with position distribution

$$\langle X_M | \rho_M^{\text{in}} | X_M \rangle = \frac{1}{\sqrt{\pi(1+2\bar{n})}} \exp \left[\frac{-X_M^2}{1+2\bar{n}} \right], \quad (\text{S11})$$

where \bar{n} is the mean thermal occupation, the heralding probabilities $\mathcal{P}_{m,n}$ for the click events $\{m,n\}$ take the form

$$\begin{aligned} \mathcal{P}_{0,0} &= e^{-2\alpha^2}, \\ \mathcal{P}_{0,1} &= e^{-2\alpha^2} \alpha^2 (1 - e^{-\mu^2(1+2\bar{n})/4} \cos \phi), \\ \mathcal{P}_{1,0} &= e^{-2\alpha^2} \alpha^2 (1 + e^{-\mu^2(1+2\bar{n})/4} \cos \phi), \\ \mathcal{P}_{1,1} &= \frac{1}{2} e^{-2\alpha^2} \alpha^4 (1 - e^{-\mu^2(1+2\bar{n})} \cos(2\phi)). \end{aligned} \quad (\text{S12})$$

Note that for finite α and $\mu^2(1+2\bar{n}) > 0$ the heralding probability is non-zero for all ϕ .




FIG. S1. **Unnormalized theoretical mechanical position probability distributions for various click-events.** An initial Gaussian state ($\{0,0\}$) is subject to $\Upsilon_{m,n}$, with $\mu = 1, \bar{n} = 4, \phi = \pi$ for single-photon ($\{0,1\}, \{1,0\}$) and two-photon ($\{1,1\}$) click-events. Note that the case $\{1,1\}$ generates a fringe-pattern with twice the frequency compared to the single-photon cases.

B. Multi-port model

Here we discuss how larger mechanical superposition states can be generated using single photon detection at each of the outputs of a multi-port interferometer. See Fig. 4A for an optical schematic. A coherent state interacts with a mechanical resonator and then (together with a second reference coherent state) is projected onto an optical $N00N$ state where N is the size of multi-port. The mechanics will have then undergone a superposition of the identity operation (zero photons) and a displacement corresponding to the N -photon radiation pressure kick. The corresponding unitary U_N for a real-bordered symmetric canonical multi-port, can be represented by a matrix with elements

$$M_N^{(k,l)} = \frac{1}{\sqrt{N}} e^{i2\pi kl/N}, \quad (\text{S13})$$

where $k, l \in [0, N-1]$. This form is a multi-port generalisation of the matrix used in Eq. (S5). As an example, the three-port case is

$$\mathbf{M}_3 = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 & 1 & 1 \\ 1 & e^{i2\pi/3} & e^{i4\pi/3} \\ 1 & e^{i4\pi/3} & e^{i2\pi/3} \end{pmatrix}. \quad (\text{S14})$$

The output field operators are given by $a_{\text{out}} = \mathbf{M} a_{\text{in}}$, where $a_{\text{in/out}}$ are N -dimensional column vectors. The state of light before the unitary U_N for an N -fold coincidence event is

$$\begin{aligned} \langle 000\dots 0 | a_N \dots a_3 a_2 a_1 U_N^\dagger = \\ \langle 000\dots 0 | U_N a_N U_N^\dagger \dots U_N a_3 U_N^\dagger U_N a_2 U_N^\dagger U_N a_1 U_N^\dagger. \end{aligned} \quad (\text{S15})$$

Here $U_N a_j U_N^\dagger$ is readily computed from the matrix expression, Eq. (S13) above.

In our scheme vacuum is injected into modes 3 to N , which leaves an optical $N00N$ state in modes 1 and 2. Vacuum inputs in modes 3 to N implies that only the first two columns of \mathbf{M} (Eq. (S13)) are being used. That is, there is no excitation in modes 3 to N . We can then write the un-normalised state for modes 1 and 2 as

$$\langle \varphi_{1,2} | \propto \langle 00 | \frac{1}{(\sqrt{N})^N} \prod_{m=0}^{N-1} (a_1 + e^{-i2\pi m/N} a_2). \quad (\text{S16})$$

All the cross terms in the product vanish and we have

$$\begin{aligned} \langle \varphi_{1,2} | \propto \frac{1}{(\sqrt{N})^N} \langle 00 | (a_1^N - (-1)^N a_2^N), \\ \propto \frac{\sqrt{N!}}{(\sqrt{N})^N} (\langle 0N | - (-1)^N \langle N0 |). \end{aligned} \quad (\text{S17})$$

The measurement operator that acts on the mechanical resonator for this multi-port case is then

$$\Upsilon_N = \frac{1}{(\sqrt{N})^N} e^{-|\alpha|^2} \alpha^N (e^{iN\mu X_M} - (-1)^N e^{iN\phi}). \quad (\text{S18})$$

Choosing $\phi = \pi$ as in the main text, this operator has the same form as $\Upsilon_{1,0}$. The displacements, however, are increased by a factor of N , compared to Eq.(S9). As a consequence, the frequency of the cosine in the filter-function $\Upsilon_N^\dagger \Upsilon_N$ is increased by a factor of N and thus exhibit phase-super resolution.

$$\Upsilon_N^\dagger \Upsilon_N = \frac{2}{N^N} e^{-2|\alpha|^2} |\alpha|^{2N} (1 - (-1)^N \cos(N\mu X_M - N\phi)). \quad (\text{S19})$$

Using this expression we can also compute the heralding probability $\mathcal{P}_N = \text{Tr} \left\{ \Upsilon_N^\dagger \Upsilon_N \rho_M \right\}$, which, for an initial thermal state with \bar{n} is given by

$$\begin{aligned} \mathcal{P}_N(\bar{n}) &= \frac{2}{N^N} e^{-2|\alpha|^2} |\alpha|^{2N} \times \\ &(1 - (-1)^N \exp[-\frac{1}{4}\mu^2(1+2\bar{n})N^2] \cos(N\phi)). \end{aligned} \quad (\text{S20})$$

Importantly, this expression is always positive when the mechanical position distribution has non-zero spread, $\mu^2(1+2\bar{n}) > 0$.

Using $\Upsilon_N^\dagger \Upsilon_N$ we can quantify how much Wigner-negativity can be generated by our scheme for given parameters N , \bar{n} and μ . We find that the minimum of the Wigner-function is given by

$$\min W = \frac{-1}{\pi(1+2\bar{n})} \frac{1 - \exp[-(1/4)\mu^2 N^2 / (1+2\bar{n})]}{1 - \exp[-(1/4)\mu^2 N^2 (1+2\bar{n})]}. \quad (\text{S21})$$

Note that for $\bar{n} = 0$, $\min W = -1/\pi$ independent of the other parameters, reaching the lowest possible value for a Wigner function. In the opposite limit, where $\bar{n} \rightarrow \infty$, $\min W$ approaches zero from below, see Fig. S2A. Note also, that projecting onto larger optical $N00N$ -states increases the generated Wigner-negativity for fixed thermal occupation \bar{n} , as illustrated in Fig. S2B.

It is also instructive to study the limits of weak ($\mu N \rightarrow 0$) and strong ($\mu N \rightarrow \infty$) optomechanical coupling

$$\min W \xrightarrow{\mu N \rightarrow 0} \frac{-1}{\pi} \frac{1}{(1+2\bar{n})^3}, \quad (\text{S22})$$

$$\min W \xrightarrow{\mu N \rightarrow \infty} \frac{-1}{\pi} \frac{1}{1+2\bar{n}}. \quad (\text{S23})$$

Thus, Wigner negativity is generated by our scheme even for weak optomechanical coupling. Moreover, the amount of negativity that can be generated is bounded by the two power-laws Eqs. (S22),(S23), as shown in Fig. 4C and Fig. S2A

SII. EXPERIMENTAL DETAILS

Our experiment makes use of an inherently stable polarization interferometer design, which can be intuitively understood in analogy to a Mach-Zehnder interferometer (MZI), see Fig. S3. The interferometer has two modes, which are represented by orthogonal polarizations $|H\rangle$ and $|V\rangle$. A half-wave plate takes the role of the beam splitter in the MZI, with the advantage of a very precise and continuously tunable splitting ratio. The two polarization modes are then separated into different spatial modes for the interaction with the mechanical resonator and recombined afterwards, using calcite beam displacers. This design achieves good phase stability due to common-mode rejection of all non-rotational mechanical noise, without requiring any active locking.

With our implementation the two beam splitters in the MZI are the same optical element and thus constrained to equal splitting ratios, which, however, is not a restriction for our experiment. Furthermore, our experimental arrangement allows to probe the position of the membrane from the other side, which significantly simplifies the experimental setup. Independent of losses in the setup, the optimal signal-to-noise ratio (SNR) is achieved with a splitting ratio of $1/2$.

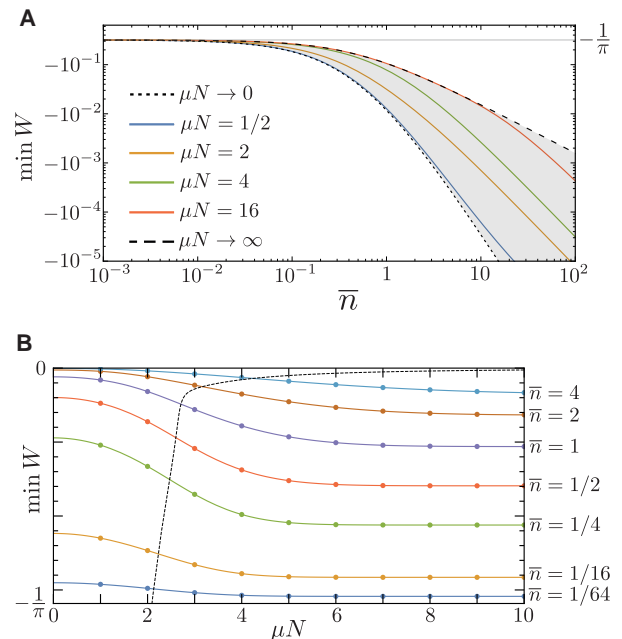


FIG. S2. **Wigner-negativity generated by our scheme.** (A) Loglog-plot of the Wigner-negativity as a function of the initial thermal occupation for various values of μN . For small \bar{n} the generated negativity is close to the maximal value of $-1/\pi$, and experiences power-law behaviour for large \bar{n} . (B) Wigner-negativity as a function of μN for various values of initial thermal occupation \bar{n} . The black-dashed line traces the inflection-points of the negativity as a function of μN and is given to aid observation of the scaling of these curves.

In order to achieve a large initial thermal state, the membrane was mounted on a Steminc SM412 ring-piezo with a capacitance of 1.8 nF and a nominal resonance frequency of 1.7 MHz. The piezo was driven with noise across the mechanical resonance frequency at 106 kHz. Note, however, that the drive couples to $\sqrt{X^2 + P^2}$, such that the drive voltage must be Chi-distributed in order to achieve Gaussian initial states in X and P . The piezo was driven with a discretized version of this distribution, sampled at 3.2 MS/s. Figure S4 shows this distribution, together with the measured distributions of X , P and $\sqrt{X^2 + P^2}$.

Synchronization of the two APDs for coincidence detection and of the APDs and the balanced detector for the position measurement was achieved using a pulsed laser diode. The diode pulses with a FWHM of 12 ns at a variable repetition rate were used as an input on the position readout side. The beam was reflected off the membrane and detected by the balanced photodetector, with one input blocked. Part of the beam is transmitted through the membrane, attenuated to the single-photon level and detected by the APDs. This technique allowed for synchronization of the two APDs to within 1 ns and between single-photon detection and position readout to within 7 ns, which is 4 orders of magnitude below the




FIG. S3. **Reflective polarization Mach-Zehnder interferometer.** (A) In a standard Mach-Zehnder interferometer, an input beam in mode a or b is split on a beamsplitter (BS) with reflectivity r_p into two spatial modes c and d . One mode obtains a static phase-shift ϕ , while the other interacts with the mechanical resonator via $e^{i\mu a_1^\dagger a_1 X}$. The two beams then interfere on a second beamsplitter with reflectivity r_p' and are split into the output modes e and f . (B) In the polarization interferometer the modes a, d, e correspond to horizontal polarization and b, c, f to vertical polarization. The BS are replaced by a half-waveplate (HWP), which allows for precise control of the splitting ratio. In the reflective design here, the polarization modes are spatially separated on a calcite beam displacer (BD). One of the beams reflects off the mechanical resonator, the other of the static frame of the resonator. The beams are recombined into the same spatial mode and interfere in the HWP. Finally they can be separated using another BD. The arrow over the mode-labels indicates the propagation direction.

time of a mechanical period.

SIII. DATA ANALYSIS

Upon an appropriate trigger signal (either APD1-click, APD2-click or coincidence click within 7.8 ns) a trace of the homodyne signal is recorded. This trace consists of 5000 points, sampled at a rate of 100 MS/s, thus resulting in a window of $\pm 25\mu\text{s}$ around the trigger event. The X - and P -values for each trace were obtained from a fit of the mechanical response function and the phase-space distribution was reconstructed from 3000 such measurements. The expected mechanical response function is of the form

$$A * \cos(X \cos[\omega_M t] + P \sin[\omega_M t] + \phi_R) + c, \quad (\text{S24})$$

where A is the full amplitude of the homodyne signal, ω_M is the mechanical resonance frequency, ϕ_R is the static phase of the readout interferometer and c is the residual DC-component in the signal due to asymmetric loss in the two arms of the homodyne interferometer. In our experiment we used a balanced photodetector with a gain




FIG. S4. **Drive spectrum and response of the mechanical resonator.** (A) Probability density histogram of the measured voltage distribution used to drive the piezo. The orange line corresponds to a Chi-distribution fit with shape-parameter $\nu = 1.98 \pm 0.03$ and the blue shaded area represents the 3σ confidence region of the fit. (B) The measured distributions of X (blue) and P (orange) using the drive in Fig. A, together with fitted Normal-distributions (blue and orange lines, respectively) with shape-parameters $\sigma_X = 0.740 \pm 0.009$ and $\sigma_P = 0.735 \pm 0.008$, respectively. The corresponding 3σ confidence regions are shown in complementary colors (orange and blue, respectively). (C) The phase-space norm $\sqrt{X^2 + P^2}$ is linearly proportional to the drive-voltage and is thus well-described by a Rayleigh distribution with shape-parameter $\sigma = 0.727 \pm 0.008$ (solid line with shaded 3σ regions).

of 10^5 V/A and a bandwidth of 4 MHz, and the DC-component c was compensated to zero by adding a tunable loss element in front of one detector. Since ω_M , A and c were measured independently and remained constant throughout the experiment, the only free variables in the fit were X , P and the read-out phase ϕ_R . These variables have distinct effects on the shape of the response, thus allowing for unique and stable fitting.

While Eq. (S24) describes the response very well in the low drive regime $\sqrt{X^2 + P^2} \lesssim 1$, we observe a mechanical position dependent optical amplitude modulation in the readout signal for larger drive. We attribute this amplitude modulation to a mechanical position dependent change in the reflectivity which is consistent with additional measurements of the optical transmission with time. The reflectivity of the mechanical device is dependent on thin-film interference and thus any small change in the refractive index of the can significantly affect the reflectivity. We modelled this modulation using a multi-

plicative function of the form

$$1 - d * \left| \cos \left[\omega_M t + \arctan(X/P) - \frac{\pi}{4} \right] \right|, \quad (\text{S25})$$

where d describes the relative strength of the amplitude modulation. The parameter d was used as an additional free parameter in the fit and was found to scale linearly with the drive strength. We used the function $1 - d|x|$ as the first-order approximation, which gave stable fits in the regime used. The arctan term transforms to the rotating-frame picture, while the $\pi/4$ phase-shift results in zero modulation when the membrane position is zero.

Note that the absolute value implies that the modulation depends on the displacement of the membrane, but not the direction. This model describes the observed response very well for $\sqrt{X^2 + P^2} \lesssim 3.2$, which is the regime where all data was taken. For larger drive we observe an additional asymmetry in the amplitude modulation that might be due to anisotropic stress in the material. Finally, for very low drive strength, the signal-to-noise ratio is limited by the resolution of the oscilloscope, resulting in a truncation in the center of the phase-space plot in Fig. 2A.

-
- [1] C. J. Davisson, L. H. Germer, *Nature* **119**, 558-560 (1927).
- [2] H. Rauch, W. Treimer, U. Bonse, *Phys. Lett. A* **47**, 369–371 (1974).
- [3] C. Monroe, D. M. Meekhof, B. E. King, D. J. Wineland, *Science* **272**, 1131–1136 (1996).
- [4] Y. Shin, M. Saba, T. A. Pasquini, W. Ketterle, D. E. Pritchard, A. E. Leanhardt, *Phys. Rev. Lett.* **92**, 050405 (1992).
- [5] S. Eibenberger, S. Gerlich, M. Arndt, M. Mayor, J. Tuxen, *Phys. Chem. Chem. Phys.* **15**, 14696–14700 (2013).
- [6] M. Aspelmeyer, T. J. Kippenberg, F. Marquardt, *Rev. Mod. Phys.* **86**, 1391–1452 (2014).
- [7] S. Bose, K. Jacobs, P. Knight, *Phys. Rev. A* **59**, 3204–3210 (1999).
- [8] W. Marshall, C. Simon, R. Penrose, D. Bouwmeester, *Phys. Rev. Lett.* **91**, 130401 (2003).
- [9] I. Pikovski, M. R. Vanner, M. Aspelmeyer, M. S. Kim, C. Brukner, *Nature Physics* **8**, 393–397 (2012).
- [10] R. W. Andrews, *et al. Nature Physics* **10**, 321–326 (2014).
- [11] D. Rugar, R. Budakian, H. J. Mamin, B. W. Chui, *Nature* **430**, 329–332 (2004).
- [12] C. M. Caves, *Phys. Rev. D* **23**, 1693–1708 (1981).
- [13] LIGO Scientific Collaboration and Virgo Collaboration, *Phys. Rev. Lett.* **116**, 061102 (2016).
- [14] A. D. O’Connell, *et al. Nature* **464**, 697–703 (2010).
- [15] K. C. Lee, *et al. Nature Photonics* **6**, 41–44 (2012).
- [16] M. R. Vanner, M. Aspelmeyer, and M. S. Kim, *Phys. Rev. Lett.* **110**, 010504 (2013).
- [17] R. Riedinger, *et al. Nature* **530**, 313–316 (2016).
- [18] E. Verhagen, *et al. Nature* **482**, 63–67 (2012).
- [19] T. P. Purdy, P. L. Yu, R. W. Peterson, N. S. Kampel, C. A. Regal, *Phys. Rev. X* **3**, 031012 (2013).
- [20] A. Safavi-Naeini, *et al. Nature* **500**, 185–189 (2013).
- [21] T. A. Palomaki, J. D. Teufel, R. W. Simmonds, K. W. Lehnert, *Science* **342**, 710–713 (2013).
- [22] E. Knill, R. Laflamme, and G. J. Milburn, *Nature* **409**, 46–52 (2001).
- [23] M. R. Vanner, *et al. Proc. Natl. Acad. Sci. USA* **108**, 16182–16187 (2011).
- [24] A. N. Boto, *et al. Phys. Rev. Lett.* **85**, 2733 (2000).
- [25] P. Walther, *et al., Nature* **429**, 158 (2004).
- [26] M. W. Mitchell, J. S. Lundeen, A. M. Steinberg, *Nature* **429**, 161 (2004).
- [27] K. J. Resch, *et al. Phys. Rev. Lett.* **98**, 223601 (2007).
- [28] G. C. Ghirardi, A. Rimini, T. Weber, *Phys. Rev. D* **34**, 470–491 (1986).
- [29] L. Diosi, *Phys. Rev. A* **40**, 1165–1174 (1989).
- [30] R. Penrose, *Gen. Relat. Gravit.* **28**, 581–600 (1996).
- [31] G. Khoury, H. S. Eisenberg, E. J. S. Fonseca, D. Bouwmeester, *Phys. Rev. Lett.* **96**, 203601 (2006).
- [32] A. B. Zwickl, *et al. Appl. Phys. Lett.* **92**, 103125 (2008).
- [33] M. Ringbauer, B. Duffus, C. Branciard, E. G. Cavalcanti, A. G. White, A. Fedrizzi, *Nature Physics* **11**, 249–254 (2013).
- [34] A. Zavatta, V. Parigi, M. Bellini, *Phys. Rev. A* **75**, 052106 (2007).
- [35] D. Kleckner, I. Pikovski, *et al. New J. Phys.* **10**, 095020 (2008).
- [36] B. Pepper, R. Ghobadi, E. Jeffrey, C. Simon, D. Bouwmeester, *Phys. Rev. Lett.* **109**, 023601 (2012).
- [37] M. R. Vanner, J. Hofer, G. D. Cole, M. Aspelmeyer, *Nat. Commun.* **4**, 2295 (2013).
- [38] M. R. Vanner, I. Pikovski, M. S. Kim, *Ann. Phys. (Berlin)* **527**, 15–26 (2015).
- [39] J. D. Teufel, *et al. Nature* **475**, 359–363 (2011).