

Quantum theories with alternative measurement postulates

Vincenzo Fiorentino

PHD

UNIVERSITY OF YORK
MATHEMATICS

September 2023

Abstract

Two measurements of an observable always yield identical outcomes when implemented in quick succession on a quantum system. In the standard formulation of non-relativistic quantum theory, this phenomenon is accounted for by the collapse of the state induced by a non-destructive measurement. In the first part of this thesis, we investigate the role of the collapse in quantum theory in a systematic way. We introduce the framework of *Alternative-Measurement Theories* (AMTs) defined as no-signalling foils of quantum mechanics sharing with it all standard postulates, except for the *update rule* assigning post-measurement states. By embedding quantum mechanics in a set of structurally similar theories, we are able to identify properties that depend on the collapse and show which of them are unique to the Lüders projection rule. We show that update rules can affect the ontological status of quantum states, or not allow for non-local correlations. The operational equivalence of different experimental strategies to measure product observables is found to rely on the chosen update rule, as well as the feasibility of protocols such as local tomography. The comparison with individual AMTs, such as “passive quantum theory” characterised by measurements which cause *no* state update, allows us to identify operational assumptions from which the Lüders rule can be derived. We show that the repeatability of measurement outcomes is insufficient; a stronger assumption on degenerate observables proves sufficient, at least for non-composite systems. In the second part, we focus on support uncertainty relations in spaces of prime dimensions. Tao derived a state-independent inequality which holds for the support sizes of a pure qudit state in two bases related by a discrete Fourier transform. We generalise Tao’s uncertainty relation to complete sets of $(d + 1)$ mutually unbiased bases in prime-dimensional spaces and investigate the sharpness of the obtained lower bounds.

Contents

Abstract	2
Contents	3
List of Tables	7
List of Figures	8
Acknowledgments	9
Author's declaration	10
Preface	11
I Quantum theories with alternative measurement postulates	15
Introduction	16
1 Quantum theory with collapsing measurements	22
1.1 The standard postulates of quantum theory	22
1.2 Quantum measurements	24
1.2.1 The measurement problem	24
1.2.2 Lüders' and von Neumann's projection postulates	25
1.2.3 Generalised measurements and POVMs	26
1.2.4 Measurements on subsystems	27

1.3	Quantum operations	29
1.3.1	Channels	30
1.3.2	Instruments	30
1.4	Alternative axiomatisations	32
1.5	Beyond quantum theory	33
2	A quantum theory with non-collapsing measurements	35
2.1	The postulates of passive quantum theory	35
2.2	Unconventional measurements: a review	36
2.3	Passive measurements	41
2.3.1	Single systems	41
2.3.2	Composite systems	45
2.3.3	Generalised p-measurements	47
2.4	Proper and improper mixtures	48
2.5	Channels and instruments in pQT	51
2.5.1	p-Channels	51
2.5.2	p-Instruments	52
2.5.3	Complete positivity	54
2.5.4	Generalised p-instruments	55
2.6	Simulating quantum theory	57
2.7	No-signalling and linearity	58
2.8	Local realism	60
2.9	Contextuality	62
2.10	An ontological model for pQT	65
2.11	Quantum information with p-measurements	72
2.12	Summary and discussion	73
3	Quantum theories with alternative measurements	76
3.1	The AMT framework	76
3.1.1	Construction	76
3.1.2	Assumptions	81
3.2	The role of update rules	82
3.2.1	Multiple extensions to composite systems	82
3.2.2	The limitations of complete positivity	83
3.2.3	Composition compatibility	85

3.2.4	Context-dependence: von Neumann's postulate	86
3.2.5	Signalling: partial repeatability	88
3.3	Generalised instruments and AMT simulation	92
3.3.1	Generalised instruments and observables	92
3.3.2	Subtheories	94
3.4	Properties of AMTs	97
3.4.1	Preparation indistinguishability	97
	Mono- and multi-partite procedures	101
	Closure under sequential composition	102
	Mutual simulability	105
3.4.2	Ideality	107
3.4.3	Deterministic repeatability	109
3.4.4	Information-disturbance trade-off	111
3.5	Summary and future work	113

II Support uncertainty relations 118

Introduction 119

4 Support inequalities for complete sets of mutually unbiased bases 122

4.1	Preliminaries	122
4.1.1	The support size of a quantum state	122
4.1.2	Support inequalities for a Fourier pair of bases	123
4.1.3	Mutually unbiased bases in prime dimensions	125
4.2	Support inequalities...	127
4.2.1	...for arbitrary pairs of MU bases	127
4.2.2	...for complete sets of $(d + 1)$ MU bases	130
4.3	Saturating support inequalities for MU bases	131
4.3.1	Constraints on saturating states	131
4.3.2	Dimension $d = 3$	135
4.3.3	Dimensions $d = 5$ and $d = 7$	136
4.3.4	Numerical results for $5 \leq d \leq 19$	136
4.3.5	Dimensions $d > 19$	137
4.4	Sharp lower bounds	138
4.4.1	Dimension $d = 3$	138

4.4.2	Dimension $d = 5$	138
4.4.3	Dimension $d = 7$	140
4.5	Summary and conclusions	142
A	Appendix to Part I	144
A.1	Definitions of operator sets	144
A.2	Proofs of Lemmata 3.1 and 3.2	144
A.3	Comparison with the work of Wilson and Ormrod	148
B	Appendix to Part II	150
B.1	Proofs of Lemmata 4.1, 4.3, 4.4 and Corollary 4.2	150
	References	157

List of Tables

3.1	Summary of the operational properties of the update rules featured in Chap. 3.	116
4.1	Summary of known lower and sharp bounds for the generalised support inequality involving the standard complete set of $(d + 1)$ MU bases, for small prime dimensions.	142

List of Figures

2.1	Schematic depiction of Gisin’s argument.	59
2.2	The star-shaped pattern of ten observables on \mathcal{H}_8 leading to a proof of the Kochen-Specker theorem.	63
2.3	Example of contextuality in the local hidden variable model compatible with pQT.	69
2.4	Example of the lack of correlations in the mono-partite scenario in the local hidden variable model compatible with pQT.	70
3.1	Example of partially repeatable measurements.	90
3.2	Circuit diagrammatic representation of Eq. (3.50).	95
3.3	Pictorial proof that Eq. (3.50) implies that $\text{AMT}(\beta)$ can simulate all instruments realisable in $\text{AMT}(\alpha)$	96
3.4	Schematic depiction of the arguments of Secs. 3.4.1 and 3.4.2 to derive the Lüders rule in the context of the AMT framework.	109
3.5	Schematic depiction of the alternative argument of Sec. 3.4.4 to derive the Lüders rule in the context of the AMT framework.	114

Acknowledgments

Non sarebbe stato possibile realizzare questa tesi senza il supporto delle persone a me vicine. In primis, un immenso grazie ai miei genitori e a mio fratello, il cui amore e affetto trascende qualsivoglia distanza fisica.

I would like to thank my supervisor, Stefan, for his guidance over the past four years. His insights, attention to detail and calm disposition helped me grow as a researcher and a writer. I express my gratitude to the members of the TAP, Prof Colbeck and Dr Pusey, for their helpful advice on the research. I also acknowledge funding from WW Smith fund.

Throughout this journey, I've been fortunate to cross paths with exceptional individuals. To my peers at the Department of Mathematics, who managed to make the long office hours more enjoyable and entertaining. Special thanks to Rutvij, Dingjia, Vasilis, Ali, Peiyun, Max, Vilasini, Hamid, Jamie, Andrew, Lewis, Kuntal and Shashaank for the exciting conversations, the crazy badminton games and the laugh tales over multiple pints of beer.

Last but not least, to Laura. Words cannot express my thanks, but they're yours anyway. Like everything else.

Author's declaration

I declare that the work presented in this thesis, except where otherwise stated, is based on my own research carried out at the University of York and has not been submitted previously for any degree at this or any other university. Sources are acknowledged by explicit references.

Chapter 2 expands on:

- Fiorentino, V. & Weigert, S. A Quantum Theory with Non-collapsing Measurements. *arXiv:quant-ph/2303.13411* (2023).

Chapter 4 has been published as:

- Fiorentino, V. and Weigert, S. Uncertainty relations for the support of quantum states. *J. Phys. A: Math. Theor.* **55**, 495305 (2022).

Preface

Quantum mechanics emerged a century ago from the joint effort of many physicists, driven by the goal to explain certain experimental observations that could not be justified within the *classical* framework of physics. To address phenomena such as the black-body radiation spectrum [68], the photoelectric effect [103, 126] and the discrete frequencies in the spectrum of hydrogen [11, 155], it was necessary to challenge long-established concepts about the nature of the Universe. Intuitions such as the quantization of energy [69, 146] and the wave-like behaviour of particles [53] provided the conceptual groundwork for the development of Heisenberg, Born, and Jordan’s matrix mechanics [30, 31, 98], as well as Schrödinger’s wave mechanics [158]. However, it was von Neumann’s subsequent contributions that provided the theory with a rigorous mathematical foundation. In his book [177], von Neumann outlines a set of assumptions that, for the most part, make up the “standard” formulation of quantum theory, as commonly taught in academic curricula.

Despite the experimental success of quantum theory, a fundamental debate persists regarding its interpretation. There is still no unanimous agreement on what its mathematical objects represent of the world. For instance, the role of quantum states is a subject of contention: some assert that states are *ontic*, objective attributes of the systems they describe, while others suggest that they merely encode our *epistemic* knowledge about them [92].

The conceptual ambiguities within quantum mechanics can partly be attributed to the abstract nature of its standard postulates. Rather than offering explicit physical principles, these postulates associate quantum systems with separable Hilbert spaces and define their dynamical and observable properties in terms of operators acting on these spaces. Integral to the standard formulation is the *projection postulate*,

also known as the *collapse* of the quantum state, *state-reduction*, or *quantum jump*. This postulate introduces a seemingly abrupt and nonlinear update of the state of a quantum system following a measurement. The so-called *measurement problem* refers to the challenge of understanding the mechanism of the collapse (or whether there is one and when it occurs), which seems to suggest a hard “cut” between the quantum and classical realms [9]. The ontic and epistemic interpretations of states offer distinct perspectives on the problem. If the state corresponds to an objective property of the system, then the collapse must be a physical process, and the resolution of the measurement problem hinges on comprehending how this process takes place. Conversely, if the state encodes our knowledge about properties of that system, then the collapse amounts to an update of our description of the system, conditioned on the new information gained from the measurement.

The measurement problem is a central subject of research within the field of *quantum foundations*, which seeks to clarify the structural properties of quantum theory and address the conceptual challenges stemming from its non-classical character. Notable achievements within this domain include proofs of quantum nonlocality and contextuality [20, 21, 122]. These results have also found practical application in the emerging domain of *quantum information*. Here, the computational power of quantum theory and its potential for (secret) communication are studied to devise algorithms and protocols that outperform their classical counterparts.

This thesis presents results from two research projects within the field of quantum foundations. The following paragraphs will introduce the topics and provide a concise summary of our contributions. For a more comprehensive understanding of the motivations and results, readers are directed to the introductions at the beginning of each part.

In Part I, we investigate the role of the projection postulate in quantum theory. To do so in a systematic way, we construct *foils* of quantum theory that retain all the standard postulates, except for replacing the collapse with alternative state update rules. The “method” of juxtaposing quantum mechanics with structurally similar models or embedding it in a set of foil theories has been successfully employed in the past. Examples include *Generalised Probabilistic Theories* (GPTs) [16, 91], *ontological models* replicating certain features of quantum theory [167] and modifications of quantum theory with *non-linear time evolutions* [26, 84] or with *additional types of measurement devices* [3, 118–120]. In a similar way, by comparing quantum theory to toy models with different post-measurement states, we aim to better understand

the conceptual and practical consequences of the collapse in the standard postulates of quantum theory. Furthermore, we look for operational principles that uniquely isolate the Lüders collapse of quantum mechanics from all feasible alternatives.

Chap. 1 provides an overview and discussion of the standard postulates of quantum theory. In Chap. 2, we explore *passive quantum theory* (pQT), which postulates that measurements *do not* disturb the state of a system. *Passive* measurements can be used to implement an *individual state determination* (ISD) procedure [33] and were initially introduced in [3] as a hypothetical computational resource. Later, they also appeared in [120] as *stochastic eigenvalue readout devices* (SERDs). The connection between pQT and previous or concurrent independent work on hypothetical measurement devices [3, 118–120] which enable the implementation of an ISD procedure, is examined. By analysing the properties of pQT, one can explore how the projection postulate shapes quantum theory. The collapse is shown to be necessary for the operational equivalence of two different experimental procedures for measuring observables on composite systems, thereby contributing to the emergence of correlations between distant systems—and thus leading to nonlocality and allowing for local tomography. The projection postulate also affects the ontology of quantum states, as “passive measurements” render the state of a *single* system observable [33, 119], which has implications for the interpretation of mixed states and the computational power of the theory [3, 120].

Chap. 3 generalises the idea of pQT and introduces *Alternative-Measurement Theories* (AMTs) through the concept of an *update rule*. We require update rules to satisfy two constraints: (non-relativistic) no-signalling and context-independence (which asserts that the post-measurement state is independent of the outcomes that are *not* observed). In contrast to no-signalling, the reason for demanding context-independence is not based on physical consistency. Instead, it is a property abstracted from quantum theory, which restricts the possible behaviours of the foil theories. We show that quantum theory can be isolated within the set of AMTs by imposing the following two operational assumptions: (i) the outcome probability distribution of an observable is not affected by a prior measurement of a coarse-graining of the same observable (this effectively encodes a form of trade-off between the information gained from measurements and the disturbance applied to states); (ii) a measurement on some subsystem ‘A’ with outcome represented by Π_x is operationally equivalent to a measurement on the composite system ‘AB’ with outcome represented by $\Pi_x \otimes \mathbb{I}_B$ (a property that we call “composition compatibility”). Alternatively, the Lüders rule

can be derived by the combined assumptions of (i) “strong indistinguishability”, stating that a sequence of measurements performed on one or more subsystems cannot distinguish between different preparations of the same joint state, and (ii) “ideality”, according to which a measurement does not change the state of the system if the outcome is certain.

In Part II, our focus will be on *uncertainty relations*. Preparational uncertainty relations describe the impossibility of preparing a quantum system in such a way that the uncertainties relating to two or more *incompatible* observables take arbitrarily small values. Heisenberg’s relation [99] concerning the product of the variances of position and momentum distributions is a notable example of preparational uncertainty relation. Others include *entropic* inequalities [55, 178] and *support* inequalities [64, 173]. In support inequalities, the uncertainty of a state with respect to a basis is quantified by the *support size*, i.e. the number of non-zero expansion coefficients of a Hilbert space vector expressed in that orthonormal basis.

Chap. 4 deals with an additive support uncertainty relation introduced by Tao [173] which holds for the computational and Fourier bases in prime-dimensional Hilbert spaces. Tao’s uncertainty relation finds applications in quantum foundations [7, 51, 52] but also in signal processing and recovery [38, 39]. Our work establishes that this inequality remains valid for any pair of bases selected from the standard complete set of $(d + 1)$ *mutually unbiased* (MU) bases existing in a Hilbert space of prime dimension d . This result allows us to construct a generalised uncertainty relation that encompasses the supports across all $(d + 1)$ MU bases. The bound we obtain appears to be sharp for dimension three only. Analytic and numerical results for prime dimensions up to nineteen suggest that the bound cannot be saturated in general. For prime dimensions two to seven, we construct sharp bounds on the support sizes in $(d + 1)$ MU bases and identify some of the states achieving them.

Part I

Quantum theories with alternative measurement postulates

Introduction

Experimental evidence for the state change of a quantum system induced by measurements has been available since 1925. Using a cloud chamber, Compton and Simon [44] studied the scattering of “x-ray quanta” by electrons. They discovered that the angle characterising the path of the recoiling electron and the angle of the photon scattering direction are strongly correlated. Knowing one of them is sufficient to determine where the particles interacted. According to von Neumann [177], this experiment implements two subsequent measurements of one single observable, which outputs the spatial coordinate of the interaction locus. The measurements of the angles can be carried out in arbitrary temporal order and in quick succession, leading to identical results. The resulting *deterministic repeatability* is argued by von Neumann to be equivalent to assuming the *projection postulate*: immediately after measuring an observable with non-degenerate eigenvalues, a quantum system will reside in the unique eigenstate associated with the observed outcome. This is in stark contrast to classical mechanics, where measurements reveal pre-existent properties of the probed systems without affecting their state, given by a point in phase space.

The incompatibility between the linear, deterministic time evolution of states and the non-linear, stochastic nature of the collapse induced by measurements is commonly referred to as the *measurement problem*. Different interpretations of quantum theory have been developed to try and resolve or “explain away” the measurement problem. These include interpretations that do not require objective “quantum jumps” (e.g. Many Worlds [57]), and others that try to describe them as manifestations of a fundamental stochastic dynamical law that operates at all times (e.g. GRW [83]).

The work presented in this part of the thesis does not directly delve into the interpretational issues surrounding measurements in quantum mechanics. Instead, we adopt an operational approach to address three main questions:

1. What is the role of the projection postulate within quantum theory? In other words, what are the conceptual and practical consequences of incorporating the collapse in the set of postulates?
2. Do alternatives to the projection postulate exist that give rise to self-consistent physical theories?
3. Are there collapse-dependent properties unique to quantum theory? In other words, can we identify operational principles that distinguish quantum mechanics from alternative theories with different post-measurement states?

To explore these questions, we will construct *foils* of quantum theory in a precise way: all standard postulates will remain unchanged, except for Lüders' projection postulate, which will be replaced by different rules for state updates.

The idea of constructing foil theories by modifying one or more axioms of quantum theory has been applied in various contexts. For example, previous studies have explored quantum-like theories with non-linear dynamical rules [5, 26, 71, 119, 151, 152, 179, 188], in contrast to the linear Schrödinger equation, or with a different method for calculating outcome probabilities [79, 80, 82, 101]. The resulting deviations from quantum mechanics have helped to illustrate the extent to which quantum theory relies on the original postulate. One example is the work of Gisin et al. [18, 85, 147, 163] who showed that nonlinear variants of Schrödinger's equation may enable superluminal communication between distant parties. However, the fact that Gisin's argument does not apply to *all* nonlinear transformations [45, 46, 71, 72, 101, 118–120, 151, 152] sparked further research on the potential inclusion of nonlinear deterministic dynamics into quantum theory. Recently, Wilson and Ormrod [182] demonstrated how the linear and unitary deterministic dynamics of quantum systems can be derived from the assumption of “local applicability”, provided the other standard postulates of quantum theory are in place.

In the more specific context of modifications to the projection postulate, we acknowledge Kent's *causal quantum theory* [116, 117] which posits that the Lüders collapse is a well-defined physical process that satisfies strict local causality. Namely, only measurements carried out in the causal past of a quantum system can affect its state. Notably, the demand for self-consistency in causal quantum theory requires that the

notion of observable be also adjusted: “sharp” observables cannot be probed, implying that all measurements are inherently fuzzy and imprecise. Furthermore, Kent’s hypothetical readout devices [118–120] represent an alternative class of measurements, introduced alongside standard quantum measurements. They are postulated to reveal (complete or partial) information about the *local state* of a system without disturbing its quantum state. The local state is defined to be the density operator obtained after accounting only for the measurement events which have taken place in the causal past.

In analogy with causal quantum theory, the foil theories presented in this document make distinct predictions from standard quantum mechanics due to a fundamental revision of the rule for state updates. However, they serve different purposes. Causal quantum theory aims to reconcile locality with the observed violation of Bell’s inequality in a way that can be potentially tested experimentally, and proposes a specific solution to the measurement problem, involving the collapse as a physical reality. In contrast, our toy models are not intended as post-quantum theories but rather function as hypothetical tools to explore the role and uniqueness of the collapse, as expressed by the three questions outlined earlier.

Conceptually, our approach is inspired by the success of applying a similar strategy to foundational topics in quantum theory using *Generalised Probabilistic Theories* (GPTs) [16, 91]. They contain both quantum and classical mechanics as special cases, among other hypothetical but structurally similar theories. Features such as the no-cloning theorem [14, 16], steering [170] or uncertainty relations [141] can also be derived for these no-signalling foil theories. Interestingly, GPTs may exhibit correlations stronger than those produced by quantum theory [148], and constructions other than the traditional tensor product can be used to describe composite systems [110, 181]. By understanding the specific options realised in quantum theory, one hopes to explain why nature seems to “prefer” those options over others.

To address the first of the three questions, we will discuss a toy model called *passive quantum theory* (pQT) which assumes that measuring any observable causes *no* state update. In this model, non-collapsing measurement devices, as introduced in [3], *replace* standard quantum measurements rather than *augmenting* them. Consequently, repeated measurements of the same observable yield outcomes according to the same probability distribution. Although pQT shares many features with standard quantum theory, it is manifestly different from it. As a foil theory, pQT does not aim to

reproduce quantum theory, unlike unitary models that try to eliminate the need for a projection postulate. Passive measurements have far-reaching consequences, both from a conceptual and an applied point of view. Suspending the collapse highlights the subtle ways in which the projection postulate shapes quantum theory. In fact, some concepts of quantum theory turn out to be equivalent as a result of the non-trivial state update described by the Lüders projection. For instance, “proper” and “improper” mixtures can be distinguished by non-collapsing measurements [33, 119]. As a consequence, density operators do not always provide a complete description of the observable properties of a system. Furthermore, there exist distinct experimental procedures to measure local observables which are equivalent in quantum theory but lead to different joint outcome probabilities in pQT. In particular, without the projective state-update after measurements, spatially separated parties may not register *nonlocal* correlations, making passive quantum theory a fully local theory. However, it still cannot be reproduced by a non-contextual hidden variable model.

The passive measurements of pQT represent an alternative to the projection postulate that results in a self-consistent, no-signalling quantum-like theory [119, 120]. To address the second question and explore what other postulated measurement behaviours give rise to operationally valid alternatives to quantum theory, we will formalise the concept of an *update rule* based on a small set of operational assumptions. An update rule serves as a map that describes how states transform following a measurement on any given system. We will use this concept, along with the remaining postulates of quantum theory, to construct the framework of *Alternative-Measurement Theories* (AMTs). Both pQT and quantum theory—despite not being an ‘alternative’—are examples of AMTs.

To address the third question, we will consider the entire framework of AMTs, rather than individual toy models. Our investigation aims to determine whether the properties identified through the comparison with pQT, which rely on the collapse, are unique to quantum theory (and therefore equivalent to assuming the standard projection postulate) or more commonly found within AMTs. We will show, in accordance with results in [37, 50], that the combination of two assumptions, namely the (*strong*) *indistinguishability of preparations* and *ideality*, precisely define quantum theory within the AMT set. A comparable argument cannot be made for *deterministic repeatability* alone, as indicated by the differing predictions between the Lüders rule of quantum theory and von Neumann’s original projection postulate. However, we will prove that the Lüders rule for *single* systems can be derived uniquely

by imposing a trade-off between the information gained from measurements and the resulting disturbance applied to the systems.

Besides identifying operational principles that can distinguish the quantum mechanical update rule from all possible alternatives, our approach suggests exploring toy models capable of faithfully reproducing quantum measurements. We will show that “linear” AMTs constitute a class of “mutually simulable” theories. That is, any AMT characterised by an update rule giving rise to linear instruments can reproduce the experimental predictions of quantum theory, despite being manifestly different from it.

Part I of the thesis is organised as follows. Chap. 1 provides a brief discussion on the standard axiomatisation of quantum theory and introduces the concept of instruments, which will play a key role in the subsequent two chapters. Additionally, we examine the apparent incompleteness of the standard set of axioms and review some of the main criticisms that have been raised in the past.

In Chap. 2, we describe passive quantum theory (pQT) and conduct a comprehensive comparison with quantum theory to elucidate the extent to which its properties rely on the standard measurement postulate. A key distinction between the two theories is the possibility in pQT to observe the state of an individual system. The idea of states being observable properties has been explored in the past; therefore, we begin the chapter by reviewing earlier work on the topic and its connection to pQT. We corroborate that passive quantum theory provides a consistent no-signalling foil of quantum theory that is not locally tomographic. However, we find that it is able to “simulate” standard quantum measurements, provided that some operational restrictions—such as a finite time delay for the “projection” to occur—are satisfied. Additionally, we present a deterministic and local hidden variable model which is compatible with the predictions of pQT, and we review how the modification of the axiom changes the computational capabilities of the theory.

In Chap. 3, we formalise the concept of an update rule and use it to define Alternative-Measurement Theories. We examine and justify the basic assumptions underlying the AMT framework, namely *context-independence* and *non-relativistic no-signalling*, and we motivate why *complete positivity* is unfit to construct a framework of theories where the state updates can be nonlinear. Within each AMT, we define generalised instruments and observables, allowing us to describe how one AMT can simulate measurements of another AMT. We then show that ideality, completeness of the density matrix description and information-disturbance trade-offs can serve

as operational properties from which the quantum mechanical update rule can be derived. We also discuss how our work aligns with—and differs from—earlier works deriving the mathematical description of deterministic transformations of quantum states [182] or of quantum measurements [76, 132, 168] from a small set of assumptions. The AMT framework remains largely unexplored, and Chap. 3 concludes by discussing potential avenues for future research. These include the search for a “nonlinear” AMTs (or proving its non-existence) that can simulate quantum mechanics, exploration of how local tomography and restrictions on multipartite correlations impact the set of toy models, and the generalization of core concepts to different mathematical structures, such as that of GPTs.

Quantum theory with collapsing measurements

The goal of this chapter is to offer a concise overview of the standard postulates of quantum theory. After presenting their content (Sec. 1.1), we will take a closer look at the role of measurements (Sec. 1.2). In particular, we will compare the Lüders collapse with von Neumann's original projection postulate and we will address the formulation of the theory in terms of generalised measurements (or POVMs). Channels and instruments will then be introduced as tools to describe the deterministic and probabilistic transformations of states, respectively (Sec. 1.3). Moreover, we will survey some common criticisms of the standard Copenhagen-type framework and explore efforts to establish alternative axiomatisations (Sec. 1.4). The chapter will conclude with a section on the role of foil theories in investigating the unique properties of quantum mechanics (Sec. 1.5).

Throughout this document, we will be using standard concepts such as bounded operators $L \in \mathcal{L}(\mathcal{H})$ on a complex, finite-dimensional Hilbert space \mathcal{H} , self-adjoint operators $M \in \mathcal{L}_s(\mathcal{H})$, projectors $\Pi \in \mathcal{P}(\mathcal{H})$, unitaries $U \in \mathcal{U}(\mathcal{H})$ and density matrices $\rho \in \mathcal{S}(\mathcal{H})$. A list of definitions can be found in Appendix A.1.

1.1 The standard postulates of quantum theory

The mathematical description of a physical system typically starts with specifying the *states* it might reside in as well as their *time evolution*. It is also necessary to indicate how *observable quantities* are represented in the theory and how *composition* of systems is described. Finally, the theory must be connected with observations made

by *measurements* so that theoretical predictions can be compared to experimental data.

The set of postulates presented in this section, often referred to as the *standard* formulation of quantum theory in finite dimensions, builds upon von Neumann's axiomatisation of 1932 [177]. While the precise number of postulates may differ across sources [112, 140, 149], they all share the common framework of describing the theory in terms of complex Hilbert spaces and the associated operators. The first four axioms establish the mathematical stage, while the last two regulate the probabilistic occurrence of measurement outcomes and the conditional post-measurement states.

- (S) To every physical system there corresponds a complex, separable and finite-dimensional Hilbert space, $\mathcal{H} = \mathbb{C}^d$, and the *states* correspond to rays $|\psi\rangle \in \mathcal{H}$ or, more generally, to density operators $\rho \in \mathcal{S}(\mathcal{H})$.
- (T) Reversible transformations, including the *time evolution* of quantum states, are described by unitary maps of the form $\rho \mapsto U\rho U^\dagger$, with $U \in \mathcal{U}(\mathcal{H})$.
- (C) The state space of a *composite* system is obtained from tensoring the Hilbert spaces describing its constituent parts.
- (O) *Observable* quantities are represented by self-adjoint operators $M \in \mathcal{L}_s(\mathcal{H})$. The eigenvalues m_x of M represent the possible outcomes of a measurement of the corresponding observable.
- (P) The *probability* that a measurement of the observable represented by M returns the outcome m_x is given by the *Born rule*

$$p(m_x) = \text{Tr}(\Pi_x \rho) \tag{1.1}$$

where $\Pi_x \in \mathcal{P}(\mathcal{H})$ projects onto the eigenspace of M with eigenvalue m_x and ρ is the state of the system when the measurement happens.

- (M[†]) If a measurement outputs the outcome m_x , then the pre-measurement state $\rho \in \mathcal{S}(\mathcal{H})$ is updated to the normalised post-measurement state according to the *Lüders rule* \mathbf{w}^\dagger :

$$\rho \xrightarrow{m_x} \mathbf{w}^\dagger(\Pi_x, \rho) = \frac{\Pi_x \rho \Pi_x}{\text{Tr}(\Pi_x \rho)}. \tag{1.2}$$

1.2 Quantum measurements

1.2.1 The measurement problem

When viewed as a physical theory describing an objective reality independent of the observer, quantum mechanics appears to be incomplete in the following sense. A realist perspective suggests that measurements should be regarded as an “emergent” phenomenon, resulting from the chain of interactions between a “measured” system and the physical constituents making up a “measuring” apparatus. To simply assume the probabilistic nature of measurements and their influence on states—Axioms (O), (P) and (M⁺)—, without providing an explanation of how these features emerge from more fundamental interactions, is an indication of this incompleteness. One can argue that the issue of incompleteness does not arise if the formalism is viewed operationally, i.e. as a mere toolset for obtaining (objective or subjective) probabilities about future experiments.

The so-called measurement problem captures, at least in part, this enigma. It refers to the question of how to reconcile the linear time evolution of states with the observed definiteness of measurement outcomes. According to Axiom (M⁺), a measurement causes a superposition of states to collapse to a single definite state. However, assuming an irreversible collapse merely shifts the problem to how to reconcile the unitary dynamics of the Schrödinger equation with the seemingly instantaneous projection occurring during the measurement process.

Numerous attempts have been made to address the measurement problem. They usually fall into two categories: *interpretational* attempts, which propose alternative interpretations of quantum theory without altering its predictions, or *modifications* of quantum theory, which suggest changes or additions to the mathematical formalism. Among the proposed solutions, we acknowledge the *Many-Worlds interpretation* [57], which posits that the wave function never collapses, and instead, the observer and the observed become entangled in a larger system. Other approaches, such as *quantum decoherence* [157], argue that interactions between a quantum system and its environment give the appearance of wave function collapse. *Objective collapse theories* [17, 19, 83] suggest modifying quantum mechanics in such a way that it includes a mechanism for spontaneous wave function collapse. However, opinions about the proposed solutions remain divided, and the measurement problem continues to be an active area of research.

The work presented in this document does not delve into the complexities of the measurement problem. Instead, our objective is to explore the implications and distinctive features of the projection postulate. To achieve this, we will construct foil theories by substituting Axiom (M^L) with (M^P) and (M^A) , introduced in Chap. 2 and 3, respectively. Therefore, measurements will continue to be considered a primitive concept in the toy theories generated through our approach.

1.2.2 Lüders' and von Neumann's projection postulates

Modern formalisations of quantum theory generally employ the projection postulate as presented by Lüders in 1950 [128], rather than the version introduced by von Neumann in [177]. Von Neumann's postulate, in fact, failed to identify post-measurement states following the measurement of a degenerate observable, i.e. where at least one eigenvalue has multiplicity greater than one. The original collapse postulate can be formalised as follows:

(M^{vN}) If a measurement of a *non-degenerate* observable represented by M outputs the outcome m_x , then the pre-measurement state is updated with the unique eigenstate of M associated to m_x .

In [177], von Neumann describes the outcome degeneracy of quantum measurements as resulting from classical post-processing of more fundamental non-degenerate measurements. In other words, according to von Neumann, one does not directly implement a measurement of a degenerate observable M ; instead, a measurement of a non-degenerate *refinement* M' commuting with M is carried out, and the outcomes are coarse-grained. However, this approach presents a challenge: for any given degenerate observable, there exist infinitely many possible refinements, leading to an infinite number of possible post-measurement mixtures. For example, given any orthonormal basis $\{|e_i\rangle\}_{i=1}^d$ of \mathcal{H} , the observable $M' = \sum_{i=1}^d \lambda_i |e_i\rangle\langle e_i|$, with $\lambda_i \in \mathbb{R}$ for all i and $\lambda_i \neq \lambda_j$ for $i \neq j$, is a refinement of the identity, $M = \mathbb{I}$.

The lack of an unambiguous assignment of output states in the degenerate case was the primary reason that led Lüders to introduce his revised postulate. With Axiom (M^L) , Lüders preserves the observed repeatability of measurement outcomes while introducing a trade-off between the information acquired from the measurement and the irreversible “disturbance” imposed on the system. Lüders justifies this replacement by considering the special case of a trivial measurement represented by the identity \mathbb{I} , which, according to von Neumann, significantly disturbs the system

to the extent that it can result in *any* state, while simultaneously providing no information about its initial preparation. In contrast, Lüders' axiom ensures that measuring \mathbb{I} is operationally equivalent to *not* performing any measurement.

1.2.3 Generalised measurements and POVMs

Some versions of the standard postulates prefer to adopt the broader notion of *generalised measurements*—or *generalised observables*—to characterise observable quantities of quantum systems. The ensuing axiom, taken from [140], would then read

(gO) Observable quantities are represented by collections $\mathbf{M} = \{M_x\}_x$ of *measurement operators* $M_x \in \mathcal{L}(\mathcal{H})$ satisfying the completeness equation $\sum_x M_x^\dagger M_x = \mathbb{I}$. Each M_x corresponds to a possible outcome of a measurement of \mathbf{M} .

The Hermitian operators appearing in (O) form the strict subset of *sharp* observables [37]. We will often denote sharp observables by the corresponding set of orthogonal projectors, rather than by the Hermitian operator.

There exists a close relationship between generalised measurements and *Positive Operator-Valued Measures (POVMs)*, i.e. collections $\{E_x\}_x$ of *effects*, positive semi-definite operators $E_x \geq O$ that sum to the identity $\sum_x E_x = \mathbb{I}_d$. To each generalised measurement, in fact, there corresponds a unique POVM via the relation $E_x = M_x^\dagger M_x$ for all x .

For generalised measurements, the Born rule reads

$$p(x) = \text{Tr} \left(M_x^\dagger M_x \rho \right), \quad (1.3)$$

whereas the Lüders rule is written as

$$\rho \xrightarrow{x} \frac{M_x \rho M_x^\dagger}{\text{Tr}(M_x^\dagger M_x \rho)}. \quad (1.4)$$

If instead POVMs are postulated, then the Born rule is modified accordingly,

$$p(x) = \text{Tr} (E_x \rho), \quad (1.5)$$

but no equivalent of Lüders rule can be specified only using the POVM elements. POVMs are often preferred in scenarios where the post-measurement state is of no interest.

Naimark's dilation theorem [138, 144] shows that any generalised measurement can be modelled by performing a sharp measurement on an ancilla after it has been

coupled to the original system by some unitary operator. Therefore, assuming (gO) is not necessary. Axiom (O) is sufficient to capture the richness of the quantum mechanical measurement process.

The possibility to define observable quantities in different ways plays a crucial role in Gleason-type theorems. Gleason’s theorem [86] from 1957 represents an attempt to derive some of the standard postulates of quantum theory from the others, along with more intuitive assumptions.¹ The theorem shows that the state space of quantum theory (S) and the Born rule for outcome probabilities (P) can be derived for dimensions $d > 2$ from Axiom (O) and a simple operational definition of states in terms of probability assignments. Gleason-type theorems aim to generalise this result by including qubits. However, all known generalisations rely on stronger assumptions concerning observables, such as (gO) or other related assumptions that include specific classes of *unsharp* observables [34, 41, 76, 186].

1.2.4 Measurements on subsystems

The Lüders postulate (\mathbf{M}^L) introduced in Sec. 1.1 describes the state change of a single quantum system \mathcal{H}_A as a result of a measurement performed on it,

$$\mathbf{w}_A^L(\Pi_x, \rho) = \frac{\Pi_x \rho \Pi_x}{\text{Tr}(\Pi_x \rho)}. \quad (1.6)$$

We will use subscripts to indicate the system for which the transformation yields the output state. In the case of a generic single system, we will use the label ‘A’.

Eq. (1.6) does not address, however, what happens to a *composite system* following a measurement on one of its constituents. In other words, the map does not describe the state change of a quantum system as a result of a measurement performed on a possibly distant, entangled system. Normally, one “extends” \mathbf{w}_A^L to a rule \mathbf{w}_{AB}^L for bipartite systems $\mathcal{H}_A \otimes \mathcal{H}_B$ in the following way. Given an initial joint state ρ_{AB} , a measurement performed on subsystem ‘A’ with outcome represented by $\Pi_x \in \mathcal{P}(\mathcal{H}_A)$ induces a state-update of the form

$$\rho_{AB} \xrightarrow{x_A} \mathbf{w}_{AB}^L(\Pi_x, \rho_{AB}) = \frac{(\Pi_x \otimes \mathbb{I}_B) \rho_{AB} (\Pi_x \otimes \mathbb{I}_B)}{\text{Tr}(\Pi_x \otimes \mathbb{I}_B \rho_{AB})}. \quad (1.7)$$

¹Naimark’s theorem can be regarded as an example of such effort, allowing (gO) to be replaced by (O). In a similar spirit, a recent paper by Wilson and Ormrod [182] shows that (T) can be replaced by an assumption of “local applicability” for deterministic transformations; although it remains debatable whether “local applicability” is a more intuitive assumption.

Similar expressions account for measurements on \mathcal{H}_B and, more generally, on any component of an arbitrary composite system.

In our notation, $w_{AB}^L : \mathcal{P}(\mathcal{H}_A) \times \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B) \rightarrow \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$ represents the complete *Lüders rule* that describes how states in quantum theory update following measurements on a subsystem. The subscript denotes the composite system for which the output state is returned, while the constituent on which the measurement is performed can be inferred from the Hilbert space on which the projector $\Pi_x \in \mathcal{P}(\mathcal{H}_A)$ representing the outcome is defined.

The single-system projection of Eq. (1.6) is implied by Eq. (1.7). Let $\rho_A = \text{Tr}_B(\rho_{AB})$, then

$$w_A^L(\Pi_x, \rho_A) = \text{Tr}_B \left[w_{AB}^L(\Pi_x, \rho_{AB}) \right]. \quad (1.8)$$

Furthermore, notice the following property of the Lüders rule: for arbitrary $\Pi_x \in \mathcal{P}(\mathcal{H}_A)$ and joint states ρ_{AB} , we can write

$$w_{AB}^L(\Pi_x, \rho_{AB}) = w_{AB}^L(\Pi_x \otimes \mathbb{I}, \rho_{AB}), \quad (1.9)$$

where ‘AB’ on the right-hand side is regarded a single system, since $\Pi_x \otimes \mathbb{I} \in \mathcal{P}(\mathcal{H}_A \otimes \mathcal{H}_B)$. In other words, the maps on the left- and right-hand sides of Eq. (1.9) are not the same: the first is defined on $\mathcal{P}(\mathcal{H}_A) \times \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$ and describes a measurement on subsystem ‘A’, whereas the second acts on $\mathcal{P}(\mathcal{H}_A \otimes \mathcal{H}_B) \times \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$ and describes a measurement on ‘AB’. Therefore, in quantum theory, a measurement on subsystem ‘A’ with outcome represented by Π_x is operationally equivalent to a measurement on the composite system ‘AB’ with outcome represented by $\Pi_x \otimes \mathbb{I}$. Both scenarios lead to the same update of the joint state.

Eq. (1.7) provides a natural but not unique way to extend Eq. (1.6) to composite systems. For example, the mapping

$$\tilde{w}_{AB}^L(\Pi_x, \rho_{AB}) = w_A^L(\Pi_x, \text{Tr}_B(\rho_{AB})) \otimes \text{Tr}_A(\rho_{AB}) \quad (1.10)$$

assigns different post-measurement states to composite systems while still being compatible with the Lüders projection for single systems,

$$w_A^L(\Pi_x, \rho_A) = \text{Tr}_B \left[\tilde{w}_{AB}^L(\Pi_x, \rho_{AB}) \right]. \quad (1.11)$$

We conclude that the standard axioms of quantum theory, in the form presented in Sec. 1.1 which can be found in [140, 149], are incomplete from an operational standpoint. To describe quantum mechanical state updates in both single *and* composite systems, Axiom (\mathbf{M}^L) needs to be replaced by a more general axiom:

(M $_{\otimes}^L$) If a measurement on system ‘A’ outputs the outcome x , then the joint pre-measurement state ρ_{AB} of the composite system ‘AB’ is updated to the normalised joint post-measurement state according to the *Lüders rule* w^L :

$$\rho_{AB} \xrightarrow{x_A} w_{AB}^L(\Pi_x, \rho_{AB}) = \frac{(\Pi_x \otimes \mathbb{I}_B) \rho_{AB} (\Pi_x \otimes \mathbb{I}_B)}{\text{Tr}(\Pi_x \otimes \mathbb{I}_B \rho_{AB})}. \quad (1.12)$$

By extending the Lüders projection to composite systems, we can more easily formulate certain properties of quantum theory, such as (non-relativistic) *no-signalling*. For all $\rho_{AB} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$ and all observables² $\{\Pi_x\}_x$ on \mathcal{H}_A ,

$$\sum_x p(x) \text{Tr}_A [w_{AB}^L(\Pi_x, \rho_{AB})] = \text{Tr}_A(\rho_{AB}) \quad (1.13)$$

describes the fact that un-conditional measurements on subsystem ‘A’ do not alter the reduced state of (the possibly distant) subsystem ‘B’. A similar equation accounts for measurements on subsystem ‘B’. Consequently, “local” measurements on quantum systems cannot be used to send signals to distant parties.

1.3 Quantum operations

In quantum theory, a state will change either due to unitary evolution or to measurements. The overall effect of any sequence composed of deterministic and probabilistic transformations, possibly executed with the mediation of ancillary systems, can be modelled by a *quantum operation*. It is a mapping from the set of density operators $\mathcal{S}(\mathcal{H})$ to the set of sub-normalised states

$$\bar{\mathcal{S}}(\mathcal{H}) = \{\rho \in \mathcal{L}_s(\mathcal{H}) : \rho \geq O, 0 \leq \text{Tr}(\rho) \leq 1\}. \quad (1.14)$$

Different realisations of the same state ρ cannot be distinguished by unitary transformations or quantum measurements. Any operation \mathcal{N} , in fact, preserves convex combinations,

$$\mathcal{N}\left(\sum_{i=1}^N p_i \rho_i\right) = \sum_{i=1}^N p_i \mathcal{N}(\rho_i) \quad (1.15)$$

for all $\rho_i \in \mathcal{S}(\mathcal{H})$ and $0 \leq p_i \leq 1$, $\sum_i p_i = 1$. A mapping \mathcal{N} satisfying Eq. (1.15) has a unique linear extension to $\mathcal{L}(\mathcal{H})$ and, in fact, operations are usually defined on this larger space [97].

²As mentioned in Sec. 1.2.3, we will often denote sharp observables by the corresponding set of orthogonal projectors, rather than by the associated Hermitian operator.

Definition 1. A mapping \mathcal{N} on $\mathcal{L}(\mathcal{H})$ is a *quantum operation* if it is

- *linear*: $\mathcal{N}(c_1 L_1 + c_2 L_2) = c_1 \mathcal{N}(L_1) + c_2 \mathcal{N}(L_2)$, for all $c_1, c_2 \in \mathbb{C}$, $L_1, L_2 \in \mathcal{L}(\mathcal{H})$;
- *completely positive*: the map $\mathcal{N} \otimes \mathcal{I}$ on $\mathcal{L}(\mathcal{H} \otimes \mathcal{H}')$ is positive for all finite dimensional extensions \mathcal{H}' ;
- *trace non-increasing*: $\text{Tr}[\mathcal{N}(L)] \leq \text{Tr}(L)$ for all $L \in \mathcal{L}(\mathcal{H})$.

1.3.1 Channels

The set of *quantum channels* is defined as the strict subset of quantum operations that are trace-preserving, $\text{Tr}[\mathcal{N}(L)] = \text{Tr}(L)$. They represent the most general deterministic transformations of quantum states. Throughout the document, we will denote channels with the letter η . Stinespring's dilation theorem [144, 171] asserts that for any channel η there exists a Hilbert space \mathcal{H}_E , a (pure) state $\xi \in \mathcal{S}(\mathcal{H}_E)$ and a unitary operator $U \in \mathcal{U}(\mathcal{H} \otimes \mathcal{H}_E)$ such that

$$\eta(\rho) = \text{Tr}_E [U \rho \otimes \xi U^\dagger] . \quad (1.16)$$

This means that every channel can be understood as the result of an interaction between the system and a suitable environment. Furthermore, there are multiple experimental setups that can give rise to the same channel, as the corresponding *dilation* $\langle \mathcal{H}_E, U, \xi \rangle$ is not unique.

1.3.2 Instruments

Instruments serve as a valuable tool for describing the effect on the state of different strategies to measure an observable.

Definition 2. Let $M = \sum_{x \in X} m_x \Pi_x$ be an observable of \mathcal{H} . A M -compatible *quantum instrument* is a collection of maps $\{\omega_x\}_{x \in X}$ such that

- For each $x \in X$, ω_x is a quantum operation;
- $\text{Tr}[\omega_x(\rho)] = \text{Tr}(\Pi_x \rho)$ for all $\rho \in \mathcal{S}(\mathcal{H})$.

Note that an instrument is usually defined as a mapping from an outcome space (X, Σ) to the set of quantum operations [97]. However, since we will only deal with discrete observables, a compatible instrument is completely determined by the finite set of operations described in Def. 2.

If a measurement of M is performed on the quantum state ρ , the resulting un-normalized post-measurement state conditioned on the outcome m_x is represented by $\omega_x(\rho)$. The un-conditional output state is instead given by $\omega_X(\rho) = \sum_{x \in X} \omega_x(\rho)$, where ω_X defines a quantum channel. If a measurement of $N = \sum_{y \in Y} n_y \Pi_y$ is performed immediately after the measurement of M , the joint sequential probabilities can be expressed as

$$p(o_1 = m_x, o_2 = n_y) = \text{Tr}[\Pi_y \omega_x(\rho)], \quad (1.17)$$

where o_1 and o_2 denote the random variables corresponding to the first and second outcomes of the time-ordered sequence, respectively. For each observable M , there are infinitely many M -compatible quantum instruments that describe different post-measurement states for the system.

The *Lüders instrument* holds a significant place in quantum theory. For any (not necessarily normalised) state ρ and any outcome Π_x of a sharp measurement, it is composed of maps of the form

$$\omega_x^\perp(\rho) = \Pi_x \rho \Pi_x. \quad (1.18)$$

When convenient, we will use the operator Π_x as a subscript, rather than the label x , i.e. $\omega_{\Pi_x}^\perp(\rho)$. The Lüders instrument is related to the Lüders rule for single systems \mathbf{w}^\perp appearing in (\mathbf{M}^\perp) —cf. (1.2)—via

$$\omega_x^\perp(\rho) = \text{Tr}(\Pi_x \rho) \mathbf{w}_x^\perp \left(\frac{\rho}{\text{Tr}(\rho)} \right), \quad (1.19)$$

where we set $\omega_x^\perp(O) = O$, with O being the zero operator. Note that the single-system Lüders rule \mathbf{w}^\perp was only defined on normalised states, thus the need to divide by the trace of $\rho \in \bar{\mathcal{S}}(\mathcal{H})$. While \mathbf{w}^\perp is nonlinear over $\bar{\mathcal{S}}(\mathcal{H})$, the associated instrument map ω^\perp in Eq. (1.18) is obviously linear.

A theorem by Ozawa [97, 142], which follows from Stinespring's dilation theorem, justifies thinking of the Lüders instrument as the *fundamental* description of the operational effects of quantum measurements, in line with its status of axiom of the theory. Before stating the result, we need to define the notion of a *measurement model* [97].

Definition 3. Let $\mathbf{M} = \{M_x\}_x$ be a generalised observable on a system associated with the Hilbert space \mathcal{H} . Let \mathcal{H}_E be a Hilbert space (associated to an ancillary

system), $\xi \in \mathcal{S}(\mathcal{H}_E)$ a pure state, $U \in \mathcal{U}(\mathcal{H} \otimes \mathcal{H}_E)$ and $N = \sum_x n_x |n_x\rangle\langle n_x| \in \mathcal{L}_s(\mathcal{H}_E)$. The quadruple $\mathcal{M} = \langle \mathcal{H}_E, \xi, U, N \rangle$ is a *measurement model* of \mathbf{M} if it satisfies the *probability reproducibility condition*,

$$\mathrm{Tr} \left[(\mathbb{I} \otimes |n_x\rangle\langle n_x|) U (\rho \otimes \xi) U^\dagger \right] = \mathrm{Tr} (M_x^\dagger M_x \rho), \quad (1.20)$$

for all $\rho \in \mathcal{S}(\mathcal{H})$ and $M_x \in \mathbf{M}$.

Ozawa showed that for any quantum instrument $\{\omega_x\}_x$ compatible with some generalised observable $\mathbf{M} = \{M_x\}_x$, there exists a measurement model $\mathcal{M} = \langle \mathcal{H}_E, \xi, U, N \rangle$ of \mathbf{M} such that

$$\omega_x(\rho) = \mathrm{Tr}_E \left[(\mathbb{I} \otimes |n_x\rangle\langle n_x|) U (\rho \otimes \xi) U^\dagger (\mathbb{I} \otimes |n_x\rangle\langle n_x|) \right] \quad (1.21a)$$

$$= \mathrm{Tr}_E \left[\omega_{\mathbb{I} \otimes |n_x\rangle\langle n_x|}^\mathbf{L} (U \rho \otimes \xi U^\dagger) \right]. \quad (1.21b)$$

Eq. (1.21b) implies that all quantum instruments, whether they describe sharp or unsharp measurements, can be obtained by carrying out a sharp Lüders measurement on an ancilla that is suitably coupled with the system. Therefore, not only is it sufficient to consider sharp observables as fundamental (see Sec. 1.2.3), but also the Lüders projection ($\mathbf{M}_\otimes^\mathbf{L}$) is sufficient to capture the full range of quantum instruments. A theorem due to Hayashi [94, 97] reinforces this idea, showing that every quantum instrument compatible with a discrete (generalised) observable $\mathbf{M} = \{M_x\}_x$ can be expressed as a post-processing of the state obtained from a Lüders measurement of the same observable,

$$\omega_x = \eta_x \circ \omega_x^\mathbf{L}, \quad (1.22)$$

where $\{\eta_x\}_x$ is a set of outcome-dependent quantum channels.

1.4 Alternative axiomatisations

Quantum mechanics has been remarkably successful in explaining experimental observations at small scales. However, the standard postulates do not provide a physical explanation for its empirical success. They associate mathematical objects with physical quantities without identifying underlying physical principles that would justify this connection.

Numerous attempts have been made to formulate quantum theory in different terms. Von Neumann and Birkhoff developed “*quantum logic*” as a logical foundation

for quantum theory [28]. Their approach, however, did not eliminate—or justify—the underlying Hilbert space structure. More recently, researchers have turned to information-theoretic [42, 49, 91, 133] or operational [129, 131] approaches to identify fundamental principles that could serve as the basis for quantum theory. None of these efforts has gained sufficient traction to replace the established Hilbert space-based formalism as the *standard* formulation of quantum mechanics.

The separate treatment of space and time in quantum mechanics has also been the subject of criticism. The Sum-Over-Histories approach [73], which is central to the interpretation of quantum mechanics known as *quantum measure theory* [43, 65, 165], is based on concepts such as events and histories that are defined on the four-dimensional manifold of spacetime. By taking special relativity into account from the outset, quantum measure theory seeks to reframe quantum theory in the language of logic, without relying on the Hilbert space setting and its resulting view of a state as a description of a system in three-dimensional space evolving with respect to absolute time via Schrödinger’s equation. Despite these efforts, a complete characterization of quantum theory in measure-theoretical terms has not yet been achieved.

1.5 Beyond quantum theory

It remains unclear which properties of quantum theory can be deemed genuinely non-classical, and why Nature appears to exclude “superquantum” features. To try and investigate the distinctive characteristics of quantum theory, it is useful to embed it in a larger family of theories that are structurally similar.

The framework of *Generalised Probabilistic Theories* (GPTs) introduces *operational* theories defined by their own sets of states (S), observables (O), and rules for system composition (C). Properties such as the no-cloning theorem, uncertainty relations or teleportation were shown to hold for some GPTs as well. Many models, importantly, exhibit superquantum correlations, e.g. Boxworld [89], and one can define GPTs that share the same set of correlations with quantum theory [185].

Another approach consists of *ontological models* which assume that all properties of a system at any given time are determined by an objective, observer-independent “ontic state” (encoded in “hidden variables”). Outcome probabilities typically emerge from the inaccessible nature of the ontic state, in contrast with operational theories

in which probabilities represent a primitive concept. Thus, the quantum state of a system corresponds to a probability distribution over the underlying ontic space. Spekkens’ toy model [167] is a prominent example of a hidden variable theory that replicates many features that are generally regarded as quantum, including the complementarity of observables, remote steering, teleportation, superdense coding and interference phenomena [40, 93]. It does not, however, replicate nonlocal correlations, contextuality or the computational speedup over classical algorithms [124]. Spekkens’ theory continues to be a subject of active research.

Constructing foil theories by modifying specific postulates of quantum theory has been instructive to reveal which features of the theory depend on which of its underlying assumptions. Examples include modifications of the Born rule (P) [79, 80, 82, 101], substitution of the standard tensor product used in (C) with minimum or maximum tensor products [15, 110] and restriction of the observables in (O) to PT-symmetric [23] or normal [154] operators. Other examples involve formulating quantum theory over different number fields, such as the real numbers [153, 172] or quaternions [74]. Kent’s causal quantum theory [116, 117] represents an attempt to modify the projection postulate (M^L) in a way that rules out any “spooky action at a distance” while maintaining the same local predictions as standard quantum theory. Additionally, readout devices [118–120] exemplify types of measurements that can be consistently added to (or replace) the standard collapse-inducing quantum measurements. In particular, supplementing standard quantum measurements with non-collapsing ones has been shown to provide computational and communication advantages [2–4, 137].

Nonlinear generalisations of Schrödinger’s equation (cf. (T)) have been suggested, most notably in Weinberg’s work [179] which, according to Gisin [85] and Polchinski [147], allows for superluminal communication. Gisin argues [18, 163] that, in order to avoid superluminal signalling, the dynamics of quantum systems must necessarily be described by completely positive linear maps on density matrices. However, *nonlinear* transformations do exist [45, 46, 71, 72, 101, 119, 120, 151, 152] that do not lead to signalling, suggesting a re-examination of Gisin’s argument.

A quantum theory with non-collapsing measurements

In this chapter, we will introduce and explore “passive quantum theory” (pQT), a foil of non-relativistic quantum theory defined by measurements that do not alter the state of a measured system (Sec. 2.1). Before describing the modified measurement process in more detail (Sec. 2.3), we embed pQT in the context of earlier work introducing hypothetical measurements allowing one to *observe* the state of an individual system (Sec. 2.2). Next, we will turn to exploring differences and similarities of pQT with quantum theory. The comparison aims at highlighting aspects of quantum theory that depend on the standard projection postulate. Specifically, we investigate the operational equivalence of two different strategies to measure observables on composite systems (Sec. 2.3.2), the modified ontology of states (Sec. 2.4), nonlinear state transformations (Secs. 2.5 and 2.7), locality and contextuality (Secs. 2.8 to 2.10) and the computational power of the foil theory (Sec. 2.11).

2.1 The postulates of passive quantum theory

Passive quantum theory (pQT) is defined by the same set of postulates as quantum theory (see Sec. 1.1), namely (S), (T), (C), (O) and (P), *except* for the Lüders rule (M^L). We replace (M^L) by the postulate (M^P) which defines passive measurements, denoting pure states in pQT (or *p-states*) by $|\widehat{\psi}\rangle$.

(M^P) A system resides in the same p-state $|\widehat{\psi}\rangle$ *before and after* measuring an observable.

It is sufficient to consider pure states to account for the modified measurement behaviour. The generalisation to density operators is implied and will be addressed in Sec. 2.4.

Measurements in pQT provide information about the state of a system in the form of a single outcome, as they do in quantum theory. They are, however, fundamentally *passive* in the sense that the trivial update rule (M^P) causes no disturbance to the system.¹ With only one type of time evolution, the toy model does not exhibit the tension between the unitary dynamics of a quantum system and its probabilistic evolution due to measurements.² Nevertheless, a substantial part of the measurement problem persists in pQT: Axiom (T) cannot explain the emergence of measurement outcomes as postulated in (O); it remains unclear when and why a measurement actually happens.

The foil theory pQT does not represent an *interpretation* of quantum theory, as its predictions differ. Proposals of unitary quantum mechanics, such as modal interpretations [60], Bohm’s theory [29] or the relative-state approach [57], aim to *eliminate* the projection postulate but do not change the predictions of quantum theory.

Replacing (M^L) by (M^P) has the striking consequence that p-states become *observable* (see Sec. 2.3.1). We will now review earlier contributions where, in one way or another, both the hypothetical means to observe a quantum state and the implications thereof are discussed.

2.2 Unconventional measurements: a review

In his 1997 paper [33], Busch imagines the existence of a procedure called *individual state determination* (ISD) which allows one to directly *observe* the state (i.e. the density operator) of an *individual* quantum system. Drawing on earlier work on state discrimination [59, 106, 109, 145], the consequences of an ISD procedure are explored, highlighting the clashes with standard quantum theory that result.

In quantum theory, it is impossible to reliably discriminate between non-orthogonal states when a single system is provided, and the state-update rule prevents infor-

¹The possibility of non-collapsing measurements was listed by von Neumann [177] as early as 1932 as one of three possible reactions of a physical system to a measurement.

²Replacing (M^L) with (M^P) differs from *augmenting* quantum theory with (M^P), which would lead to yet a different theory with *three* competing “time evolutions”: (T), (M^L) and (M^P).

mation about a system being gained without affecting its state. In contrast, an ISD procedure would allow one to unambiguously identify any state in a set of unknown non-orthogonal pure states. *Proper* mixtures, resulting from incomplete knowledge about the preparation of a system, would no longer be fundamental: any prior uncertainty about the ‘real’ state of an individual system can be eliminated using an ISD procedure. The possibility to discern between two preparations of the same mixed state ρ implies that density operators provide only an *incomplete* description of a system [33]. Carrying out an ISD procedure on a subsystem of an entangled system would reveal the mixed reduced state. Therefore, *improper* mixtures, which arise from entanglement with other systems, are observable just as pure states are [33]. Busch also points out that combining the ISD procedure with standard collapse-inducing measurements allows distant parties to communicate superluminally fast.

Throughout, Busch only considers the case of quantum theory *augmented* by an ISD procedure. It is instructive to consider an “ISD-only” version of the theory, in which standard quantum measurements are absent. In such a theory, signalling between distant parties is not possible, but one can still distinguish between proper and improper mixtures and discriminate between unknown non-orthogonal states.

The passive measurements of pQT can be used to implement an ISD procedure. Busch briefly mentions measurements “with no state changes” [33, p. 7] as a possible method to observe a quantum state, though this is unfeasible in standard quantum theory. In line with Busch’s argument, passive measurements can be introduced at the expense of fundamentally changing quantum theory. Consequently, the implications he describes, e.g. the revised role of mixed states, also apply to pQT—except for the fact that distant parties cannot use passive measurements to communicate (see Sec. 2.7). In other words, pQT can be regarded as a “realisation” of the “ISD-only” version of quantum theory. That being said, the goal of exploring pQT in preparation for the more general treatment of AMTs differs from the thrust of Busch’s paper, which is to show how a hypothetical (unspecified) ISD procedure violates accepted properties of standard, unmodified quantum theory.

Given access to a source of classical randomness, any ISD procedure allows one to effectively *simulate* passive measurements, albeit with an inevitable time delay (see Sec. 2.6). To do so, one first observes the state of an individual system using ISD. Next, one calculates the associated probability distribution and randomly selects one outcome based on that distribution. This procedure takes finite time, even if

done for a second run when the probability distribution is already known. Similarly, an ISD procedure can be simulated in pQT to arbitrary precision in a *finite* time interval since a finite sequence of tomographically complete passive measurements will reveal the (approximate) state of a system. However, an *infinite* sequence of measurements is necessary for *exact* state reconstruction (Sec. 2.3.1).

In 2005, Kent showed that nonlinear transformations of quantum states do not necessarily imply superluminal signalling [119], in contrast to *Gisin’s argument* [18, 85, 147, 163]. The argument is based on a hypothetical *state readout device* (SRD), i.e. a measurement device capable to extract the *local* state of an individual system without disturbing it. As outlined in the Introduction, the local state denotes the reduced state of the subsystem accessible to one party subsequent to all (standard, collapse-inducing) measurements conducted on it or on any system entangled with it within its causal past. In essence, a measurement performed on another space-like separated part of the system updates the quantum state of the accessible subsystem—as prescribed by the projection postulate—but it does not update the local state *instantly*. The update only happens once sufficient time has elapsed for a light-speed signal to propagate between the systems. Consequently, the quantum state and the local state of a system may temporarily diverge. Thus, the SRD represents one way among others [45, 46, 71, 72, 101, 151, 152] to implement a *nonlinear* time evolution that is not ruled out by Gisin’s argument. In fact, by letting the known local state determine which unitary gate to apply, an observer can implement *any* locally varying nonlinear transformation on $\mathcal{S}(\mathcal{H})$.

Kent’s readout device can be thought of as a specific ISD procedure which reveals the *local* state of the measured system rather than its *quantum* state. Consequently, quantum theory *augmented* by an SRD exhibits many of the properties outlined by Busch in [33], without violating the no-signalling requirement. For example, different preparations of the same mixed state can be distinguished by invoking the SRD just *once*. It also becomes possible to effectively *clone* an initially unknown quantum state, without violating the (dynamical) no-cloning theorem—cf. Sec. 2.3.1.

Both an SRD and measurements in pQT do not disturb the measured system, nor do they enable signalling between distant parties. Nevertheless, it seems appropriate to consider “SRD-augmented” quantum theory and pQT as *distinct* theories. In contrast to pQT, in the “SRD-augmented” quantum theory one can still perform standard measurement-based protocols such as teleportation or entanglement swapping, or violate Bell’s inequalities—cf. Sec. 2.8 and Sec. 2.11. The motivation to

define the local state (and the reason it differs from the quantum state) rests on the possibility to collapse quantum states with standard quantum measurements, which do not exist in pQT. While not explicitly considered in [119], an “SRD-only” quantum theory—i.e. a theory where standard quantum measurements are absent (or never performed) but a readout device is assumed to exist—would represent another specific example of an “ISD-only” theory, using Busch’s terminology. The “SRD-only” theory will exhibit the same limitations regarding measurement-based protocols and correlation strength that we will encounter in pQT. Just as a generic ISD procedure, a state readout device can effectively simulate passive measurements (modulo a time delay), and *vice versa*.

In 2016, Aaronson et al. introduced non-collapsing (or passive) measurements to explore their potential computational advantages [3]. Specifically, they considered a version of quantum theory *augmented* by *non-adaptive* non-collapsing measurements. Here, *non-adaptive* means that the outcomes of non-collapsing measurements cannot be used to determine the unitaries applied later in the circuit. The toy model was shown to enable efficient solutions to the Graph Isomorphism problem (the task of determining whether two finite graphs are isomorphic) and to offer improvements over Grover’s algorithm. However, non-adaptive passive measurements do not allow NP-hard problems to be solved efficiently. This contrasts with other modifications of quantum theory—such as nonlinear dynamics, postselection, and changes to the Born rule—which had been shown to extend the power of quantum computation to NP or beyond [1, 5]. The results in [3] catalysed further research into non-adaptive non-collapsing measurements [2, 4, 137] while the computational power of *adaptive* non-collapsing measurements remains largely unexplored.

Similar conclusions were drawn in 2021 when Kent also introduced non-collapsing measurements via the so-called *stochastic eigenvalue readout device* (SERD). SERDs were presented alongside several other hypothetical devices to explore alternative methods for obtaining information about quantum states in theories that involve localised collapse [120]. A *single* use of the SRD can simulate a SERD or, equivalently, a passive measurement, except for an unavoidable finite time delay, while repeated uses of the SERD allow one to reconstruct the local state [120]. In addition, invoking an infinite-precision SRD once would enable the solutions of NP problems in polynomial time, while a SERD only provides partial information about the state, hence will not be as efficient as an SRD—cf. Sec. 2.11. The foil theory considered by Aaronson et al. in [3] coincides with “SERD-augmented” quantum theory as

presented in [120], and it differs from passive quantum theory for the same reasons as the “SRD-augmented” quantum theory. Importantly, “SERD-only” quantum theory—where standard quantum measurements are absent—is *identical* to pQT.

Responding to the claim that the measurement postulates of quantum mechanics are operationally redundant [132], Kent introduced new hypothetical readout devices in 2023 [118]. A *stochastic positive operator readout device* (SPOD) is defined, akin to a SERD but associated with POVMs instead of PVMs. Generalised measurements in pQT, discussed in Secs. 2.3.3 and 2.5.4, will correspond to SPODs. In this paper, Kent also considers the “SERD-only” (and the “SPOD-only”) theory which is equivalent to pQT, introduced independently³ in [75].

In summary, the idea of extracting information about an unknown quantum state, a key feature of measurements in pQT, has been addressed in a number of ways dating back to at least 1997. We are aware of three proposals that are based on unconventional procedures or devices, with different motivations provided. Firstly, Busch is interested in exploring the consequences of adding hypothetical procedures that make the state of quantum systems observable [33]. Any such procedure would—in the absence of further modifications—clash with fundamental properties of quantum theory, such as the relativistically motivated no-signalling. Secondly, Kent’s readout devices [118–120] are meant (*i*) to illustrate that, upon suitable modifications, quantum states may evolve nonlinearly without leading to superluminal signalling; and (*ii*) to show that alternatives to the standard projection postulate are compatible with quantum theory, thereby refuting the idea that the standard state update rule could be derived from the other postulates. Thirdly, Aaronson et al. introduced non-collapsing measurements alongside collapsing measurements to investigate how acquiring information about the state of the register at key points in the quantum circuit could enhance the computational power of the theory [3].

Our own motivation for passive quantum theory, where the state update in a quantum measurement is suppressed, is to identify those properties of quantum theory which depend on the standard update, and in which way. Such an investigation has not yet been carried out systematically.

When discussing pQT, we will revisit and expand upon some established properties of ISD- or SERD-augmented theories, including, for instance, an argument for the incompleteness of quantum channels as descriptions of p-state transformations

³Ref. [118] was released in July 2023 (during the final stages of editing of the thesis), i.e. after the pre-print [75] that introduced passive quantum theory in March 2023.

(Sec. 2.5.1), an examination of the operational constraints on simulating quantum measurements within pQT (Sec. 2.6), and a critique of the claim in [163] that Gisin’s argument (as presented in the paper) does not rely on the projection postulate (Sec. 2.7).

In addition, we establish further properties of pQT, such as the equivalence (under a reasonable assumption) of postulate (M^P) for passive measurements with *sequential commutativity* (Sec. 2.3), a discussion on the role of *complete positivity* in foils such as pQT with nonlinear instruments (Sec. 2.5.3), and the distinction between *direct* and *indirect* classes of p-instruments (Sec. 2.5.4) which does not exist in standard quantum theory. We also review contextuality in pQT (Sec. 2.9), and present a deterministic ontological model compatible with pQT, based on the so-called “Bell model” [125] (Sec. 2.10). We will indicate throughout how our findings relate to the earlier work on unconventional measurements summarised above.

2.3 Passive measurements

We begin by discussing passive measurements on single systems (Sec. 2.3.1), then examine the case of composite systems (Sec. 2.3.2). We conclude with a brief section on generalised measurements in pQT (Sec. 2.3.3).

2.3.1 Single systems

The predictions of standard quantum theory and passive quantum theory actually agree as long as post-measurement states are not involved. For instance, the expectation value of an observable $M = \sum_x m_x \Pi_x$ can be determined in pQT just as in quantum mechanics: the eigenvalue m_x of the observable M will occur with probability $p(m_x)$, according to Axiom (P), upon measuring it repeatedly on an *ensemble* of systems each of which resides in the p-state $|\widehat{\psi}\rangle$. Thus, preparational uncertainty relations [99, 115, 159] hold for the variances of non-commuting observables in pQT. Similarly, entropic [107, 156] and support [64, 173] inequalities carry over from standard quantum theory since the Shannon entropy of a p-state and its support-size (see Sec. 4.1.1) can be evaluated in the same way as for quantum states. Consequently, Heisenberg’s original plausibility argument—i.e. measuring the position of an electron will cause an uncertainty of its momentum, due to an

uncontrollable state change—is invalid, as already noticed by Busch [36]. Preparational uncertainty relations describe fundamental limitations on the ability to prepare states and they are not related to irreversible measurement-based disturbances.

However, the expectation value of an observable may be obtained in pQT from a *single* copy of a p-state $|\widehat{\psi}\rangle$, in contrast to the ensemble needed in quantum theory. Since a (non-destructive) passive measurement of the observable M does not update a p-state, it is possible to repeat the measurement *on one and the same* system as often as is necessary to determine the outcome probabilities $p(m_x)$ given by the Born rule to arbitrary precision. As a result, and in line with the previous work discussed in Sec. 2.2, the collapse-free theory allows us to reconstruct an unknown p-state $|\widehat{\psi}\rangle$ from a *single* system by simply repeating p-measurements of an informationally complete set of observables [35, 108, 150]. *A fortiori*, an experimenter can tell apart any two distinct non-orthogonal p-states $|\widehat{\psi}\rangle$ and $|\widehat{\phi}\rangle$ *with certainty*, even when being presented with a single copy only, as already observed in [33, 119]. To discriminate between two p-states, full tomography is, in general, not needed: repeated measurements of a single observable with different expectation values in these states would be sufficient.

Successful single-copy state reconstruction means that p-states should be thought of as *observable* quantities assigned to individual quantum systems, akin to states of classical objects. The flip side of the lack of collapse is that an experimenter can no longer make use of measurements to *prepare* states. A desired p-state can only be prepared by suitably evolving some known p-state in time, i.e. dynamically.

The Lüders rule of quantum theory entails *deterministic repeatability* as observed in actual experiments.

Definition 4 (*Deterministic repeatability* (DR)). Consecutive measurements of the same observable, performed on the same system without intervening unitary evolution, yield identical outcomes.

Consider a system on which two successive measurements are performed (without intervening evolution). Let o_1 and o_2 be the random variables denoting the outcomes of the first and second measurement, respectively. Deterministic repeatability ensures that for any observable $M \in \mathcal{L}_s(\mathcal{H})$ with set of outcomes $\{m_x\}$, the conditional probability $p(o_2 = m_x | o_1 = m_x)$ of observing outcome m_x , after the first measurement

of M has returned m_y , is given by

$$p(o_2 = m_x | o_1 = m_y) = \delta_{m_x m_y} \quad \text{for all } m_x, m_y, \quad (2.1)$$

for any initial state of the system with $p(o_1 = m_y) \neq 0$.

Clearly, measurements in pQT are not deterministically repeatable. Instead, repeated p-measurements are consistent with *probabilistic repeatability*.

Definition 5 (*Probabilistic repeatability* (PR)). Consecutive measurements of the same observable, performed on the same system without intervening unitary evolution, yield identically distributed outcomes.

In other words, the conditional probabilities obey

$$p(o_2 = m_x | o_1 = m_y) = p(o_1 = m_x) \quad \text{for all } m_x, m_y, \quad (2.2)$$

and for any p-state.

Passive measurements are also consistent with the requirement of *sequential commutativity*.

Definition 6 (*Sequential commutativity* (SC)). Given any two observables M and N , the *joint sequential probability* of observing any pair of outcomes m_x and n_y is independent of the order in which the measurements are carried out.

Mathematically, sequential commutativity holds if

$$p(o_1 = m_x, o_2 = n_y) = p(o_1 = n_y, o_2 = m_x), \quad (2.3)$$

hence

$$p(o_1 = m_x) p(o_2 = n_y | o_1 = m_x) = p(o_1 = n_y) p(o_2 = m_x | o_1 = n_y), \quad (2.4)$$

for all m_x, n_y . For a given Hilbert space, combining the Born rule with sequential commutativity implies postulate (M^P), provided we make the reasonable assumption that later outcomes do not affect the state obtained as a result of an earlier measurement. This suggests that passive measurements are more closely related to classical probability theory than quantum measurements. In the context of the general framework of Alternative-Measurement Theories (AMTs) presented in Chap. 3, Lemma 2.1 shows that sequential commutativity (SC) represents an operational definition of the update rule of pQT—cf. Def. 9.

Lemma 2.1. *Sequential commutativity (SC) implies passive measurements (\mathbf{M}^P) if we assume*

- (i) *the standard postulates for states (S), observables (O) and outcome probabilities (P);*
- (ii) *state-updates do not depend on the outcomes of future measurements; i.e. if $|\psi\rangle \xrightarrow{o_1} |\phi\rangle \xrightarrow{o_2} \dots$ denotes a sequence of state-updates for the observed outcomes o_1, o_2, \dots , then $\phi(\psi, o_1, o_2, \dots) = \phi(\psi, o_1)$.*

Proof. Let the system reside in the state $|\psi\rangle \in \mathcal{H}_x^M$, where $\mathcal{H}_x^M \equiv \Pi_x^M \mathcal{H}$ is the eigenspace of M with (possibly degenerate) eigenvalue m_x . Let $N = M$, then Eq. (2.4) implies that $p(o_2 = m_y | o_1 = m_x) = 0$ if $x \neq y$. Since $\sum_y p(o_2 = m_y | o_1 = m_x) = 1$, it must follow that $p(o_2 = m_x | o_1 = m_x) = 1$. Hence, the post-measurement state after recording outcome m_x must still lie in \mathcal{H}_x^M , $|\psi\rangle \xrightarrow{m_x} |\psi_x^M\rangle \in \mathcal{H}_x^M$. More generally (and making use of assumption (ii)), if a measurement on a pure state returns an outcome with probability 1, then the post-measurement state must remain in the eigenspace of the observable associated with that outcome.

Suppose $|\psi_x^M\rangle \neq |\psi\rangle$ (this is only meaningful if $\dim \mathcal{H}_x^M > 1$), hence observing m_x leads to a non-trivial update of the state. Let N be an observable different from M such that $\Pi_y^N = |\psi\rangle\langle\psi|$ for some eigenvalue n_y . Since $p(o_1 = n_y) = 1$, it follows from the argument in the previous paragraph that $p(o_2 = n_y | o_1 = n_y) = 1$, hence $|\psi\rangle \xrightarrow{n_y} |\psi\rangle \in \mathcal{H}_y^N$ (the post-measurement state must coincide with the pre-measurement state since $\dim \mathcal{H}_y^N = 1$). Therefore, for the initial state $|\psi\rangle \in \mathcal{H}_y^N \subset \mathcal{H}_x^M$, sequential commutativity ensures that

$$p(o_1 = m_x) p(o_2 = n_y | o_1 = m_x) = p(o_1 = n_y) p(o_2 = m_x | o_1 = n_y) \quad (2.5a)$$

$$\implies p(o_2 = n_y | o_1 = m_x) = 1 \quad (2.5b)$$

since $p(o_2 = m_x | o_1 = n_y) = p(o_1 = m_x) = 1$. But Eq. (2.5b) implies that $\langle\psi|\psi_x^M\rangle = 1$, leading to a contradiction. We conclude that $|\psi_x^M\rangle = |\psi\rangle$ when $|\psi\rangle \in \mathcal{H}_x^M$. This property is known as *ideality* (cf. Sec. 3.4.2).

At this point, we know that if a measurement returns an outcome with probability 1, then the post-measurement state must coincide with the pre-measurement state. It remains to consider the more general case, wherein the state is *not* an eigenstate of the measured observable. Let N now be any observable such that $p(o_1 = n_y) \neq 0$ for some n_y , and let M be any observable such that $p(o_1 = m_x) = 1$. From ideality we have that $p(o_2 = n_y | o_1 = m_x) = p(o_1 = n_y)$, hence Eq. (2.4) becomes

$p(o_1 = n_y) = p(o_1 = n_y)p(o_2 = m_x|o_1 = n_y)$, or $p(o_2 = m_x|o_1 = n_y) = 1$. Therefore, after recording outcome n_y , the state must still lie in \mathcal{H}_x^M . This result holds for *any* \mathcal{H}_x^M that contains $|\psi\rangle$, in particular the one-dimensional subspace with projector $\Pi_x^M = |\psi\rangle\langle\psi|$. Since we assume that the state-update induced by $o_1 = n_y$ cannot depend on $o_2 = m_x$, we exclude the possibility of a ‘conspiracy’ wherein $|\psi\rangle$ is transformed into a different state, still inside the required \mathcal{H}_x^M . We can thus conclude that the measurement of such arbitrary observable N must leave the initial state unchanged, i.e. $|\psi\rangle \xrightarrow{n_y} |\psi\rangle$ holds for any $|\psi\rangle$ such that $p(o_1 = n_y) \neq 0$. \square

The no-cloning theorem [58, 184] expresses the impossibility to produce copies of unknown, non-orthogonal states *dynamically*, i.e. through the application of unitary evolution of a system. The theorem is a consequence of the linearity of the quantum time evolution (T) of states (S) of composite systems (C). Since the proof does not involve the measurement postulate, the no-go result also holds in passive quantum theory. However, in pQT an alternative cloning procedure exists, based on single-copy state-reconstruction followed by the preparation of a new system in the observed p-state. More generally, quantum state cloning is an immediate application of *any* ISD procedure [33], such as the readout devices introduced in [118–120] which encompass passive measurements [3]. In quantum theory, the equivalent of this measurement-based procedure would require an ensemble of identically prepared systems.

2.3.2 Composite systems

In line with the discussion of Sec. 1.2.4, Axiom (M^P) is insufficient to describe the impact on *joint* states of measurements carried out on a subsystem. For a complete characterisation of passive measurements, the more general Axiom (M_⊗^P) is introduced.

(M_⊗^P) A composite system resides in the same p-state $|\Phi\rangle$ *before and after* measuring an observable of one of its subsystems.

For simplicity, we will focus on bipartite systems, i.e. $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$. Consider product observables such as $M = A \otimes B$. In quantum theory, M can be measured by a single apparatus \mathcal{D}_{AB} that spans both parts of the system in state $|\Phi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$; this method will be called the *multi-partite* scenario. Alternatively, M can be measured using two *mono-partite* devices \mathcal{D}_A and \mathcal{D}_B that measure $A \in \mathcal{L}_s(\mathcal{H}_A)$

and $B \in \mathcal{L}_s(\mathcal{H}_B)$, respectively, followed by some classical communication. If the constituents of the system are at different locations, only the method using separate devices can be implemented.

In the multi-partite scenario, one obtains the eigenvalues $m_{xy} = a_x b_y$ of M with probabilities $p(m_{xy})$ and the system ends up in the state $\Pi_x^A \otimes \Pi_y^B |\Phi\rangle / \sqrt{p(m_{xy})}$ (for simplicity, we consider non-degenerate eigenvalues m_{xy}). In the mono-partite scenario, the experimenters can, by comparing their records, determine the probabilities $p(o_1 = a_x)$ and $p(o_2 = b_y | o_1 = a_x)$ (as well as $p(o_1 = b_y)$ and $p(o_2 = a_x | o_1 = b_y)$), where o_1 and o_2 denote the outcomes of the measurement sequence, which govern the occurrence of the eigenvalues a_x and b_y . The identities

$$p(m_{xy}) = \text{Tr} \left(\Pi_x^A \otimes \Pi_y^B |\Phi\rangle\langle\Phi| \right) \quad (2.6a)$$

$$= \text{Tr} \left(\Pi_x^A \otimes \mathbb{I} |\Phi\rangle\langle\Phi| \right) \text{Tr} \left(\mathbb{I} \otimes \Pi_y^B \frac{\Pi_x^A \otimes \mathbb{I} |\Phi\rangle\langle\Phi| \Pi_x^A \otimes \mathbb{I}}{\text{Tr}(\Pi_x^A \otimes \mathbb{I} |\Phi\rangle\langle\Phi|)} \right) \quad (2.6b)$$

$$= p(o_1 = a_x) p(o_2 = b_y | o_1 = a_x) \quad (2.6c)$$

and

$$p(m_{xy}) = \text{Tr} \left(\Pi_x^A \otimes \Pi_y^B |\Phi\rangle\langle\Phi| \right) \quad (2.7a)$$

$$= \text{Tr} \left(\mathbb{I} \otimes \Pi_y^B |\Phi\rangle\langle\Phi| \right) \text{Tr} \left(\Pi_x^A \otimes \mathbb{I} \frac{\mathbb{I} \otimes \Pi_y^B |\Phi\rangle\langle\Phi| \mathbb{I} \otimes \Pi_y^B}{\text{Tr}(\mathbb{I} \otimes \Pi_y^B |\Phi\rangle\langle\Phi|)} \right) \quad (2.7b)$$

$$= p(o_1 = b_y) p(o_2 = a_x | o_2 = b_y) \quad (2.7c)$$

entail that, in quantum theory, the distribution of the outcomes of M is independent of the chosen method of implementation. The final state of the composite system will also coincide,

$$\begin{aligned} \frac{\mathbb{I} \otimes \Pi_y^B}{\sqrt{p(o_2 = b_y | o_1 = a_x)}} \left(\frac{\Pi_x^A \otimes \mathbb{I} |\Phi\rangle}{\sqrt{p(o_1 = a_x)}} \right) &= \frac{\Pi_x^A \otimes \mathbb{I}}{\sqrt{p(o_2 = a_x | o_1 = b_y)}} \left(\frac{\mathbb{I} \otimes \Pi_y^B |\Phi\rangle}{\sqrt{p(o_1 = b_y)}} \right) \\ &= \frac{\Pi_x^A \otimes \Pi_y^B |\Phi\rangle}{\sqrt{p(m_{xy})}}. \end{aligned} \quad (2.8)$$

Therefore, the Lüders rule (\mathbf{M}_{\otimes}^L) ensures the operational equivalence between the two experimental strategies for measuring product observables in quantum theory.

The measurement axiom (\mathbf{M}_{\otimes}^P) breaks the equivalence between the procedures in pQT. When a passive measurement of M is performed on the p-state $|\widehat{\Phi}\rangle$ by

the multi-partite device \mathcal{D}_{AB} , the outcome probabilities are identical to those of quantum theory,

$$p(m_{xy}) = \langle \widehat{\Phi} | \Pi_x^A \otimes \Pi_y^B | \widehat{\Phi} \rangle = \langle \Phi | \Pi_x^A \otimes \Pi_y^B | \Phi \rangle, \quad (2.9)$$

and the process will not disturb the joint state.

In contrast, in the mono-partite scenario, the two devices \mathcal{D}_A and \mathcal{D}_B will each perform a measurement that does not update the joint state [3, 120]. As a result, they will both operate on the same p-state $|\widehat{\Phi}\rangle$, regardless of their order of implementation. The lack of collapse trivialises the conditional probabilities, $p(o_2 = b_y | o_1 = a_x) = p(o_1 = b_y) \equiv p(b_y)$ and $p(o_2 = a_x | o_1 = b_y) = p(o_1 = a_x) \equiv p(a_x)$, hence the measurement outcomes produced by \mathcal{D}_A and \mathcal{D}_B will not be correlated. Therefore, in pQT, the two-device strategy does not generally constitute a measurement of M ,

$$p(m_{xy}) \neq p(o_1 = a_x) p(o_2 = b_y | o_1 = a_x) = p(a_x) p(b_y), \quad (2.10)$$

$$p(m_{xy}) \neq p(o_1 = b_y) p(o_2 = a_x | o_1 = b_y) = p(b_y) p(a_x). \quad (2.11)$$

The two scenarios remain equivalent only when the initial state is separable, $|\widehat{\Phi}\rangle = |\widehat{\phi}_A\rangle \otimes |\widehat{\phi}_B\rangle$.

If the two subsystems are spatially separated, there is no way to measure M without first bringing them together. Thus, passive quantum theory is not *locally tomographic*⁴. In other words, an informationally complete set of observables of the form $A \otimes B$ exists, but their expectation values can only be obtained by carrying out p-measurements using devices \mathcal{D}_{AB} spanning both constituents. Observables that cannot be expressed in the form $A \otimes B$ can only be implemented by multi-partite devices, as in standard quantum theory.

2.3.3 Generalised p-measurements

Naimark's dilation theorem [138, 144] does not involve post-measurement states, hence it applies to pQT. For any generalised observable $\mathbf{M} = \{M_x\}_x$ on \mathcal{H} , there exists some unitary operator $U \in \mathcal{U}(\mathcal{H} \otimes \mathcal{H}_E)$ such that, for any p-state $|\widehat{\psi}\rangle \in \mathcal{H}$ and a fixed ancilla p-state $|\widehat{\xi}\rangle \in \mathcal{H}_E$, we have

$$U(|\widehat{\psi}\rangle \otimes |\widehat{\xi}\rangle) = \sum_x M_x |\widehat{\psi}\rangle \otimes |\widehat{n}_x\rangle, \quad (2.12)$$

⁴Note that quantum theory *augmented* with passive measurements [3] or other readout devices [118–120] remains locally tomographic, as opposed to quantum theory in which standard quantum measurements are *replaced* accordingly.

where $\{|\widehat{n}_x\rangle\}_x$ is an orthonormal basis of \mathcal{H}_E . The probability reproducibility condition—cf. Eq. (1.20) in Sec. 1.3.2—confirms that the outcome distribution of a sharp ancilla measurement of $N = \sum_x n_x |\widehat{n}_x\rangle\langle\widehat{n}_x|$ agrees with the predictions of the Born rule for a measurement of \mathbf{M} on $|\widehat{\psi}\rangle$,

$$\mathrm{Tr} \left[\left(\mathbb{I} \otimes |\widehat{n}_x\rangle\langle\widehat{n}_x| \right) \left(\sum_{x'x''} M_{x'} |\widehat{\psi}\rangle\langle\widehat{\psi}| M_{x''}^\dagger \otimes |\widehat{n}_{x'}\rangle\langle\widehat{n}_{x''}| \right) \right] = \langle\widehat{\psi}| M_x^\dagger M_x |\widehat{\psi}\rangle \quad (2.13)$$

Therefore, even in pQT, the *measurement model* $\mathcal{M} = \langle\mathcal{H}_E, \xi, U, N\rangle$, where $\xi = |\widehat{\xi}\rangle\langle\widehat{\xi}|$, can be regarded as defining a strategy to implement a measurement of \mathbf{M} . However, in contrast with quantum theory, the lack of a state update means that the protocol leaves the system entangled with the ancilla. To decouple them, the inverse unitary U^\dagger must be applied.

2.4 Proper and improper mixtures

Up to this point, our focus has been on pure states described by rays in Hilbert space. In quantum theory, probabilistic preparations of pure states give rise to mixed states. We describe probabilistic preparations mathematically by means of *Gemenge* (German for ‘mixture’), i.e. collections $\{(p_1, \psi_1), (p_2, \psi_2), \dots, (p_n, \psi_n)\}$ where $p_i \geq 0$, $\sum_i p_i = 1$ and $|\psi_i\rangle \in \mathcal{H}$. For example, the Gemenge $\mathcal{G} = \{(q, \psi_1), (1 - q, \psi_2)\}$, describes a system that is prepared in state $|\psi_1\rangle$ with (classical) probability q or in state $|\psi_2\rangle$ with probability $(1 - q)$ [37]. Alternatively, it can describe an ensemble of systems, each residing in either $|\psi_1\rangle$ or $|\psi_2\rangle$ with probabilities q and $(1 - q)$, respectively. An element of such ensemble is then completely described by the density matrix

$$\rho_{\mathcal{G}} = q |\psi_1\rangle\langle\psi_1| + (1 - q) |\psi_2\rangle\langle\psi_2|, \quad (2.14)$$

which returns the outcome probabilities and post-measurement states prescribed by the Born rule (P) and Lüders rule (\mathbf{M}^L), respectively. Eq. (2.14) denotes a mixed state if $q \neq 0, 1$.

To any Gemenge there corresponds a unique density operator; however, to each mixed density operator there exist infinitely many compatible Gemenge. Despite requiring different experimental set-ups to prepare, in quantum theory these mixtures are operationally indistinguishable. No quantum measurement (or sequence thereof) can reveal whether a system was prepared according to \mathcal{G}_1 or \mathcal{G}_2 , if $\rho_{\mathcal{G}_1} = \rho_{\mathcal{G}_2}$. This

follows from the linearity of both the Born rule and the Lüders instrument on the space of density matrices [97].

Mixed states also arise in quantum theory when an experimenter can access only part of a composite system. Given an entangled bipartite state ρ_{AB} , the reduced state of each subsystem obtained with the partial trace, $\rho_A = \text{Tr}_B(\rho_{AB})$ and $\rho_B = \text{Tr}_A(\rho_{AB})$, will be mixed. Mixed states with this origin are often referred to as *improper*, while the term *proper*—or *epistemic*—is used for probabilistic Gemenge of states [13, 48, 174]. The indistinguishability of different realisations of the same density operator implies that no quantum measurement can determine whether ρ refers to a proper or an improper mixture.

In pQT, mixed states play a different role, as mentioned already in Sec. 2.2. Upon receiving a system prepared according to $\mathcal{G} = \{(q_1, \psi_1), (q_2, \psi_2), \dots, (q_n, \psi_n)\}$, an experimenter can choose to perform single-copy reconstruction and identify the pure p-state $|\widehat{\psi}_i\rangle$ of the system. Thus, as pointed out already in [33, 119], a *proper* mixture represents an incomplete and disposable—hence not fundamental—description of a p-system, that can be improved by extracting more information via passive measurements. In principle, any ignorance reflected in the classical “mixing” probability can be removed, enabling an experimenter to retrieve the correct Gemenge describing the preparation of an ensemble of p-systems.

However, improper mixtures still play a role in pQT. If an observer has access only to half of a pair residing in an entangled p-state $|\widehat{\Psi}\rangle \in \mathcal{H}_{AB}$, then single-copy reconstruction on, say, subsystem ‘A’ will identify the improper mixed state $\rho_A = \text{Tr}_B(|\widehat{\Psi}\rangle\langle\widehat{\Psi}|)$ associated with it (the *local* state in the terminology of [119, 120]). Unless the observer gets access to the entire composite system, a mixed state remains the most accurate description of the subsystem. Improper mixtures can therefore be “observed” in pQT in the same way as pure p-states. In particular, single-copy tomography defines a protocol to *witness* entanglement, as the detection of “mixedness” guarantees the presence of entanglement with other systems. However, it is impossible to identify the joint state $|\widehat{\Psi}\rangle$ uniquely, since many states lead to the same mixture at ‘A’,

$$\text{Tr}_B \left[(\mathbb{I}_A \otimes U_B) |\widehat{\Psi}\rangle\langle\widehat{\Psi}| (\mathbb{I}_A \otimes U_B^\dagger) \right] = \text{Tr}_B (|\widehat{\Psi}\rangle\langle\widehat{\Psi}|) = \rho_A \quad (2.15)$$

for any $U_B \in \mathcal{U}(\mathcal{H}_B)$.

We conclude that, in general, density operators in pQT do not represent a complete description of the observable properties of a system, as inferred from [33,

119, 120]. A proper mixture represented by ρ can be distinguished from an improper mixture described by the same operator. Additionally, passive measurements allow us to distinguish between different preparations of an ensemble corresponding to the same proper mixed state. For improper mixtures, however, the density matrix represents the best possible description of the system, as it does in quantum theory.

Having extended the discussion to density operators, we can define, in analogy with the Lüders rule w^L for single systems of Eq. (1.2), the *passive update rule* for single systems, $w^P : \mathcal{P}(\mathcal{H}) \times \mathcal{S}(\mathcal{H}) \rightarrow \mathcal{S}(\mathcal{H})$,

$$w^P(\Pi_x, \rho) = \rho. \quad (2.16)$$

Notice that, according to both w^L and w^P , the output state depends (at most) on the pre-measurement state and the observed outcome. In other words, the post-measurement state is independent of the *context* of Π_x , i.e. the measured observable $\{\Pi_x\}_x$.

We can use Eq. (2.16) to reformulate Axiom (M^P) in a way that resembles more closely the structure of (M^L) .

(M^P) If a measurement outputs the outcome x , then the pre-measurement state $\rho \in \mathcal{S}(\mathcal{H})$ is updated to the normalised post-measurement state according to the *passive rule* w^P as follows:

(i) If ρ is a pure state, then

$$\rho \xrightarrow{x} w^P(\Pi_x, \rho) = \rho. \quad (2.17)$$

(ii) If ρ is a proper mixture with Gemenge⁵ $\mathcal{G} = \{(p_1, \rho_1), \dots, (p_n, \rho_n)\}$, then

$$\rho = \sum_{i=1}^n p_i \rho_i \xrightarrow{x} \sum_{i=1}^n p(i|x) w^P(\Pi_x, \rho_i) = \sum_{i=1}^n p(i|x) \rho_i, \quad (2.18)$$

where $p(i|x) = p_i p(x|i)/p(x)$. In terms of Gemenge, $\mathcal{G} \xrightarrow{x} \mathcal{G}_x = \{(p(1|x), \rho_1), \dots, (p(n|x), \rho_n)\}$.

As improper mixtures arise from entanglement in composite systems, they do not feature in Axiom (M^P) which pertains to non-composite systems.

⁵We can express Gemenge in terms of density matrices, rather than state-vectors. This will be more useful later, as we will also want to consider proper mixtures consisting of improper mixed states.

2.5 Channels and instruments in pQT

2.5.1 p-Channels

The proof of Stinespring’s dilation theorem [144, 171] does not involve measurements, hence the result holds in pQT as well. Consequently, channels—linear, completely positive, trace-preserving maps on $\bar{\mathcal{S}}(\mathcal{H})$ —provide a representation for deterministic transformations of p-states. In particular, every channel η can be implemented by letting the system interact with an environment which is then traced out, and there exist multiple strategies to execute the same channel.

However, as a result of the different ontology of p-states (cf. Sec. 2.4), it is possible to distinguish between different experimental procedures giving rise to the same channel. For example, consider the entangling unitary operator

$$U(|\widehat{\psi}\rangle \otimes |\widehat{0}\rangle) = \sqrt{1-p} |\widehat{\psi}\rangle \otimes |\widehat{0}\rangle + \sqrt{p/3} (\sigma_x |\widehat{\psi}\rangle \otimes |\widehat{1}\rangle + \sigma_y |\widehat{\psi}\rangle \otimes |\widehat{2}\rangle + \sigma_z |\widehat{\psi}\rangle \otimes |\widehat{3}\rangle), \quad (2.19)$$

which, for $\rho = |\widehat{\psi}\rangle\langle\widehat{\psi}|$ and $\xi = |\widehat{0}\rangle\langle\widehat{0}|$, gives rise to the *depolarising channel* when substituted in Eq. (1.16),

$$\eta_{\text{dep}}(\rho) = (1-p)\rho + p/3 (\sigma_x \rho \sigma_x + \sigma_y \rho \sigma_y + \sigma_z \rho \sigma_z). \quad (2.20)$$

In quantum theory, the transformation can be given an alternative yet equivalent interpretation: the state is left undisturbed with probability $(1-p)$, or one of the three operations of bit flip, phase flip or their combination is performed, each with probability $p/3$. However, this interpretation does not hold in pQT. If η_{dep} is implemented by applying the transformation of Eq. (2.19), then reconstruction of the resulting p-state via sequential passive measurements will result in a mixed state. On the other hand, if a classical mixture of the qubit unitaries \mathbb{I} , σ_x , σ_y , σ_z is implemented, then the same procedure will point to a pure state, revealing which of the four gates was actually applied. The fact that either the unitary or the measurement-based view can be consistently upheld is another example of conflation of concepts within quantum theory which results from the particular form of Lüders rule.

We conclude that, due to the distinguishability between proper and improper mixtures of p-states, quantum channels η do not represent complete descriptions of the (deterministic) transformations of states in pQT. Information regarding the

particular method of implementation of a given channel η can in fact be obtained by performing p-measurements on the system residing in $\eta(\rho)$.

2.5.2 p-Instruments

In Sec. 1.3.2, we discussed how quantum instruments serve as a useful tool for describing the impact of different experimental strategies to measure a given observable. Specifically, we explained how the Lüders instrument, derived by combining Axiom (P) with Axiom (M^L),

$$\omega_x^L(\rho) = \text{Tr}(\Pi_x \rho) w_x^L\left(\frac{\rho}{\text{Tr}(\rho)}\right) = \Pi_x \rho \Pi_x, \quad (2.21)$$

can be considered as the fundamental quantum instrument compatible with an arbitrary observable M . Other M -compatible quantum instruments are derivatives of it, as indicated by Eq. (1.21b) and Eq. (1.22).

In a similar manner, we can construct the *fundamental p-instrument* $\{\omega_x^P\}$ by combining Axiom (P) with Axiom (M^P). Then, by considering post-processing or the inclusion of ancillary systems, we can define *generalised p-instruments* in terms of $\{\omega_x^P\}$, cf. Sec. 2.5.4. Recalling w^P from Eq. (2.16), for an arbitrary sharp observable with outcomes denoted by $x \in X$, the corresponding *fundamental p-instrument* $\{\omega_x^P\}_x$ consists of maps on $\bar{\mathcal{S}}(\mathcal{H})$ of the form

$$\omega_x^P(\rho) = \text{Tr}(\Pi_x \rho) w_x^P\left(\frac{\rho}{\text{Tr}(\rho)}\right) = \text{Tr}(\Pi_x \rho) \frac{\rho}{\text{Tr}(\rho)}. \quad (2.22)$$

Note that we must set $\omega_x^P(\rho) = O$ when $\text{Tr}(\rho) = 0$, i.e. when ρ is the zero operator. The operator $\omega_x^P(\rho)$ represents the un-normalised conditional post-measurement state, while $\omega_X^P(\rho) = \sum_{x \in X} \omega_x^P(\rho)$ denotes the un-conditional state. In analogy with the Lüders instrument, the fundamental p-instrument map ω_x^P is 1-homogeneous, i.e. $\omega_x^P(\lambda\rho) = \lambda \omega_x^P(\rho)$ for $\lambda \in [0, 1]$. This property enables us to account for sequences of measurements. In fact, $\omega_y^P(\omega_x^P(|\widehat{\psi}\rangle\langle\widehat{\psi}|))$ describes the state at the end of a time-ordered sequence of passive measurements with outcomes x and y , respectively. The joint sequential probability is then given by $\text{Tr}[\omega_y^P(\omega_x^P(|\widehat{\psi}\rangle\langle\widehat{\psi}|))]$.

The incompleteness of density operators as descriptions of the observable properties of p-systems is a consequence of the fact that ω_x^P does not define a quantum operation (except in the trivial cases where $\Pi_x \in \{O, \mathbb{I}\}$). In fact, ω_x^P cannot be

implemented in quantum theory as it fails to preserve convex combinations, so that no linear extension to $\mathcal{L}(\mathcal{H})$ exists,

$$\omega_x^{\text{P}} \left(\sum_i p_i |\widehat{\psi}_i\rangle\langle\widehat{\psi}_i| \right) \neq \sum_i p_i \omega_x^{\text{P}} (|\widehat{\psi}_i\rangle\langle\widehat{\psi}_i|) . \quad (2.23)$$

Eq. (2.23) encodes the distinguishability between proper and improper mixed states. Specifically, the left-hand side of Eq. (2.23) describes the absence of disturbance on an *improper* mixture described by $\rho = \sum_i p_i |\widehat{\psi}_i\rangle\langle\widehat{\psi}_i|$, whereas the right-hand side represents the information update resulting from a passive measurement on the *proper* mixture described by the Gemenge $\mathcal{G} = \{(p_1, \psi_1), (p_2, \psi_2), \dots\}$.

To illustrate how passive measurements impose a classical update on proper mixtures, let us consider the example of a single qubit prepared according to $\mathcal{G} = \{(1/2, |\widehat{0}\rangle\langle\widehat{0}|), (1/2, |\widehat{\mp}\rangle\langle\widehat{\mp}|)\}$, hence described by $\rho = (|\widehat{0}\rangle\langle\widehat{0}| + |\widehat{\mp}\rangle\langle\widehat{\mp}|) / 2$. A passive measurement of σ_z modeled by the p-instrument $\{\omega_0^{\text{P}}, \omega_1^{\text{P}}\}$ is performed, and the outcome “0” (eigenvalue +1) is obtained. Then, the updated normalised description of the system is given by the proper mixture

$$\rho_0 = \frac{1}{p(0)} \left[\frac{1}{2} \omega_0^{\text{P}} (|\widehat{0}\rangle\langle\widehat{0}|) + \frac{1}{2} \omega_0^{\text{P}} (|\widehat{\mp}\rangle\langle\widehat{\mp}|) \right] = \frac{2}{3} |\widehat{0}\rangle\langle\widehat{0}| + \frac{1}{3} |\widehat{\mp}\rangle\langle\widehat{\mp}| . \quad (2.24)$$

The new operator reflects the fact that, among the two available options, the observed outcome suggests that it is more likely that the qubit was prepared in the p-state $|\widehat{0}\rangle$. If this was indeed the case, then the mixture would converge to $|\widehat{0}\rangle\langle\widehat{0}|$ in the limit of infinite repetitions of σ_z p-measurements. Importantly, for this procedure to be effective, the initial decomposition \mathcal{G} must include the correct pure state. If the experimenter is unaware whether ρ denotes a proper or improper mixture and assigns equal probability to each scenario, she would initially describe the system with the Gemenge

$$\mathcal{G}' = \{(1/4, |\widehat{0}\rangle\langle\widehat{0}|), (1/4, |\widehat{\mp}\rangle\langle\widehat{\mp}|), (1/2, \rho)\} , \quad (2.25)$$

which is compatible with the same mixed state, $\rho_{\mathcal{G}'} = \rho_{\mathcal{G}} = \rho$. Moreover, if the Gemenge \mathcal{G} is also unknown and no specific Gemenge is assigned a higher probability, she would use an equal mixture of all proper mixtures compatible with ρ . By applying the p-instrument maps to individual elements of the sum, the experimenter can then systematically reconstruct the correct p-state associated with the single system.

2.5.3 Complete positivity

In quantum theory, instrument maps ω_{AB}^L describing subsystem measurements can be constructed, unsurprisingly, by combining the Born rule of Axiom (P) with the generalised Lüders rule w_{AB}^L of Axiom (M_{\otimes}^L) (cf. Sec. 1.2.4),

$$\omega_{AB}^L(\Pi_x, \rho_{AB}) = \text{Tr}[\Pi_x \text{Tr}_B(\rho_{AB})] w_{AB}^L\left(\Pi_x, \frac{\rho_{AB}}{\text{Tr}(\rho_{AB})}\right) \quad (2.26)$$

$$= (\Pi_x \otimes \mathbb{I}_B) \rho_{AB} (\Pi_x \otimes \mathbb{I}_B) \quad (2.27)$$

$$= (\omega_x^L \otimes \mathcal{I})(\rho_{AB}) . \quad (2.28)$$

Here, $\Pi_x \in \mathcal{P}(\mathcal{H}_A)$ denotes an arbitrary outcome of subsystem ‘A’.

The requirement of complete positivity of ω_x^L ensures that ω_{AB}^L always returns valid joint states of $\mathcal{H}_A \otimes \mathcal{H}_B$. The fact that the mapping $\omega_x^L \otimes \mathcal{I}$ can be interpreted as describing the effect on the larger system of a *local* measurement on \mathcal{H}_A is ensured by the following consistency condition,

$$\text{Tr}_B \left[(\omega_x^L \otimes \mathcal{I})(\rho_{AB}) \right] = \omega_x^L[\text{Tr}_B(\rho_{AB})] , \quad (2.29)$$

which holds for all $\rho_{AB} \in \bar{\mathcal{S}}(\mathcal{H}_A \otimes \mathcal{H}_B)$.

Complete positivity can be defined for both linear and nonlinear transformations [6]. However, for nonlinear maps this requirement does not generally encode the same desired operational properties. As pointed out in [47], complete positivity is “physically unfitting” for nonlinear dynamics. Taking the example of passive measurements, the fundamental p-instrument map ω^P —cf. Eq. (2.22)—is *not* completely positive, since a completely positive and 1-homogeneous map is necessarily linear [6]. However, this does not imply that it is impossible to consistently “extend” ω^P to describe the impact of p-measurements carried out on part of a composite system. In fact, the generalised Axiom (M_{\otimes}^P) introduced in Sec. 2.3.2 serves precisely that purpose.

In analogy with quantum theory, Axiom (M_{\otimes}^P) of pQT defines the generalised passive update rule w_{AB}^P which outputs *normalised* states of composite systems:

$$w_{AB}^P(\Pi_x, \rho_{AB}) = \rho_{AB} . \quad (2.30)$$

The transformations of Eq. (2.30) is defined for any finite-dimensional extension \mathcal{H}_B . Furthermore, for arbitrary $\Pi_x \in \mathcal{P}(\mathcal{H}_A)$ and ρ_{AB} , it holds that

$$w_{AB}^P(\Pi_x, \rho_{AB}) = w_{AB}^P(\Pi_x \otimes \mathbb{I}, \rho_{AB}) , \quad (2.31)$$

where ‘AB’ on the right-hand side is regarded a single system, since $\Pi_x \otimes \mathbb{I} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$. This property is also valid in standard quantum theory—cf. Eq. (1.9).

The p-instrument maps ω_{AB}^P , therefore, can be defined as follows,

$$\omega_{AB}^P(\Pi_x, \rho_{AB}) = \text{Tr}[\Pi_x \text{Tr}_B(\rho_{AB})] w_{AB}^P\left(\Pi_x, \frac{\rho_{AB}}{\text{Tr}(\rho_{AB})}\right) \quad (2.32)$$

$$= \text{Tr}(\Pi_x \otimes \mathbb{I} \rho_{AB}) \frac{\rho_{AB}}{\text{Tr}(\rho_{AB})}. \quad (2.33)$$

The interpretation of $\omega_{AB}^P(\Pi_x, \rho_{AB})$ as the *un-normalised* state of a bipartite system $\mathcal{H}_A \otimes \mathcal{H}_B$ after a passive measurement on subsystem \mathcal{H}_A yields outcome Π_x is ensured by the same consistency condition used for the Lüders instrument—Eq. (2.29)—and, more generally, for quantum operations,

$$\text{Tr}_B[\omega_{AB}^P(\Pi_x, \rho_{AB})] = \omega_A^P[\Pi_x, \text{Tr}_B(\rho_{AB})], \quad (2.34)$$

which holds for all joint states ρ_{AB} .

A more detailed discussion on the role of complete positivity in the context of more general *update rules* is provided in Sec. 3.2.2 of Chap. 3.

2.5.4 Generalised p-instruments

Having introduced the fundamental p-instrument for both single and composite systems, we can define *generalised p-instruments* by accounting for post-processing of passive measurements and the inclusion of ancillary systems. Although ω_S^P is only defined for sharp measurement outcomes of system \mathcal{H}_S , we can extend it to measurement operators (or, equivalently, to effects) in the following way. Let $\mathcal{M} = \langle \mathcal{H}_E, \xi, U, N = \{\Pi_x^N\}_x \rangle$ be a measurement model compatible with the generalised observable $\mathbf{M} = \{M_x\}_x$, we define

$$\omega_{M_x}^P(\rho) = \text{Tr}_E\left\{U^\dagger \left[\omega_{\mathbb{I} \otimes \Pi_x^N}^P(U \rho \otimes \xi U^\dagger)\right] U\right\}, \quad (2.35)$$

where $\{\omega_{\mathbb{I} \otimes \Pi_x^N}^P\}$ describes the effect on the composite system of the passive sharp measurement on the ancilla \mathcal{H}_E —cf. Eq. (2.31). Then, the collection $\{\omega_{M_x}^P\}_x$ describes a passive measurement of \mathcal{M} which, as already mentioned in Sec. 2.2, corresponds⁶ to the action of the SPOD introduced in [118]. In particular, Eq. (2.35) captures

⁶Technically, a SPOD implements a POVM rather than a generalised measurement, but this distinction is not significant since the device does not alter the state of the measured system.

the fact that, in order to bring the system back to its initial state ρ , the inverse coupling unitary U^\dagger must be applied after measuring \mathbf{M} via the strategy defined by \mathcal{M} —cf. Sec. 2.5.1. For example, consider a passive measurement of $M = \{\Pi_x^M\}$, which may indeed be performed using an ancilla. Let $\rho = |\widehat{\psi}\rangle\langle\widehat{\psi}|$ and $\xi = |\widehat{0}\rangle\langle\widehat{0}|$, where $|\widehat{\psi}\rangle = \sum_i c_i |\widehat{i}\rangle$. We can then set $U(|\widehat{\psi}\rangle \otimes |\widehat{0}\rangle) = \sum_i c_i |\widehat{i}\rangle \otimes |\widehat{i}\rangle = |\widehat{\Psi}\rangle$. A measurement of $N = M$ on the ancilla will yield the same outcome probability distribution as a measurement of M on $|\widehat{\psi}\rangle$. However, the p-measurement on the ancilla will not alter the joint state, i.e. $\mathbf{w}_{\mathbb{I} \otimes \Pi_x^M}^{\mathbf{P}}(U\rho \otimes \xi U^\dagger) = |\widehat{\Psi}\rangle\langle\widehat{\Psi}|$, hence to complete the protocol and return the system to its original state, the unitary U^\dagger must be applied.

The generalisation of $\omega^{\mathbf{P}}$ to unsharp measurements allows us to define generalised p-instruments for single systems. A formal definition of generalised instruments applying to any theory in the AMT framework will be given in Sec. 3.3.1 of Chap. 3. In the special case of pQT, they can be expressed more simply by a formula resembling the quantum case, see Eq. (1.22).

Definition 7. Given a generalised observable $\mathbf{M} = \{M_x\}_x$, an \mathbf{M} -compatible generalised p-instrument is any collection $\{\omega_{M_x}\}_x$ of maps such that

$$\omega_{M_x} = \eta_x \circ \omega_{M_x}^{\mathbf{P}}, \quad (2.36)$$

where $\{\omega_{M_x}^{\mathbf{P}}\}_x$ describes a passive measurement of \mathbf{M} and $\{\eta_x\}_x$ is a set of outcome-dependent channels.

In quantum theory, Ozawa’s theorem [97, 142]—cf. Sec. 1.3.2—demonstrates that a physical implementation of the channel in Eq. (1.22) is not necessary. There always exists a measurement model—defining a strategy terminating with an ancilla measurement—that realises the corresponding quantum instrument. This is not true in pQT, where there exist p-instruments for which the time-ordered sequence shown in Eq. (2.36) represents the only way to implement the transformations. For example, Eq. (2.35) describes a measurement strategy for unsharp \mathbf{M} that cannot be realised without applying a unitary channel *after* the ancilla measurement. Given that time is required to physically implement a quantum channel, these strategies inevitably take longer to execute, and can thus be distinguished from those where post-measurement channels are not needed. In Sec. 3.3.1 of Chap. 3, this distinction in the experimental implementation will define the difference between *direct* and *indirect* generalised instruments in arbitrary AMTs. The possibility to realise any instrument “directly”,

i.e. without the need for further evolution after the measurement, should not be expected to hold *a priori*. In quantum theory, this is a consequence of the specific form of the projection postulate.

2.6 Simulating quantum theory

Measurements on passive quantum systems can be made to look as if they were performed on a quantum system, modulo a time delay. For non-degenerate Hermitian operators M , a Lüders measurement of the associated observable returns a post-measurement state $|m_x\rangle$ that is independent of the initial one $|\psi\rangle$ (except that $|m_x\rangle$ must have a non-zero expansion coefficient in $|\psi\rangle$). One can replicate the process in pQT by performing a passive measurement of the same observable and substituting $|\widehat{\psi}\rangle$ for $|m_x\rangle$. The “replacement” time would be subject to “quantum speed limits” [54] if the set $\{|m_x\rangle\}$ was unavailable and the experimenter had to generate the required p-state by unitary evolution. For degenerate Hermitian operators, the collapsed state depends also on the input state. Single-copy reconstruction of ρ is therefore necessary to determine $w_x^{\downarrow}(\rho)$. However, if the measured system is, say, half of an entangled pair residing in $|\widehat{\Psi}\rangle$, then the simulation protocol fails, unless the experimenter has access to the whole composite system. In fact, the improper mixture $\rho_A = \text{Tr}_B(|\widehat{\Psi}\rangle\langle\widehat{\Psi}|)$ is insufficient to determine the collapsed product state $|m_x\rangle \otimes |\widehat{\phi}\rangle$. The pure p-state $|\widehat{\Psi}\rangle$ must instead be reconstructed; however, as mentioned in Sec. 2.3.2, this cannot be done with mono-partite measurements. Therefore, other than the time delay due to p-state replacement and the need for an infinite sequence of p-measurements, the protocol requires that all entangled systems be at the same location or be experimentally accessible in some other way. If these restrictions are satisfied, then the Lüders instrument—hence any quantum instrument—of an arbitrary observable M can be simulated in pQT.

One class of quantum instruments that can be simulated without resorting to multiple passive measurements is *conditional state preparators*.

Definition 8. Given an observable $M = \{M_x\}$, an M -compatible instrument is called a *conditional state preparator* if it is composed of maps of the form

$$\omega_x(\rho) = \text{Tr}(M_x^\dagger M_x \rho) \xi_x \quad (2.37)$$

where $\xi_x \in \mathcal{S}(\mathcal{H})$.

They are equivalent to discarding the system and preparing a new one in the state ξ_x independent of the original. Substituting complete state contractions $\eta_x(\rho) = \xi_x$ in the formula for generalised p-instruments of Eq. (2.36), in fact, makes the transformation linear—more specifically, a quantum operation. It is known that if M is a rank-1 observable, i.e. for any x , $M_x^\dagger M_x = e(x) \Pi_x$ for some projector Π_x and $0 < e(x) < 1$, then all M -compatible quantum instruments are of the form of Eq. (2.37) [97]. Therefore, pQT can simulate any quantum measurement of a rank-1 observable with a single passive device, i.e. without first reconstructing the state. In contrast, quantum theory cannot simulate the passive measurements of non-trivial observables featured in pQT.

2.7 No-signalling and linearity

If local measurements do not alter the state of a composite system, then they cannot be used by two distant parties, say Alice and Bob, to communicate. Given a bipartite system $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$ and any observable $A = \sum_x a_x \Pi_x$ on \mathcal{H}_A , then

$$\sum_x \text{Tr}_A \left[\omega_{AB}^P(\Pi_x, \rho_{AB}) \right] = \text{Tr}_A(\rho_{AB}) \quad (2.38)$$

holds for all $\rho_{AB} \in \mathcal{S}(\mathcal{H}_{AB})$. Similarly, for any observable $B = \sum_y b_y \Pi_y$ on \mathcal{H}_B ,

$$\sum_y \text{Tr}_B \left[\omega_{AB}^P(\Pi_y, \rho_{AB}) \right] = \text{Tr}_B(\rho_{AB}) . \quad (2.39)$$

Eq. (2.38) and (2.39) establish (quantum) *no-signalling* (NS) in pQT from Alice to Bob and from Bob to Alice, respectively. Clearly, the relations hold for arbitrary p-instruments implemented locally. That is, appending channels of the form $\eta_x \otimes \mathcal{I}_B$ to $\omega_{AB}^P(\Pi_x, \rho_{AB})$ in Eq. (2.38), or $\mathcal{I}_A \otimes \eta_y$ to $\omega_{AB}^P(\Pi_y, \rho_{AB})$ in Eq. (2.39), preserves the equalities.

A violation of Eq. (2.38) and (2.39), along with the assumption that measurement disturbances propagate instantaneously (in some reference frame), amounts to a violation of the *relativistic* no-signalling principle, according to which information cannot travel superluminally. Typically, nonlinear transformations are *inconsistent* with relativistic no-signalling, but exceptions indeed exist, such as those implemented using Kent's readout devices [118–120] (cf. end of the section). In [85], Gisin showed that a nonlinear modification of the Schrödinger equation presented by Weinberg

[179] led to arbitrarily fast communication between observers. *Gisin's argument*, later improved in [163], affirms that, by assuming the Hilbert space setting (S), the Born rule (P) and that information cannot propagate faster than light, it follows that the time evolution of quantum states must be described by linear, completely positive maps. This result depends on the fact that different convex combinations of quantum states can be used to describe one and the same mixed state. Imagine to remotely prepare a mixed state in one of two distinct convex combinations (Gemenge) by performing local measurements on one part of a bipartite system. If quantum states were to evolve nonlinearly in time, a space-like separated observer could subsequently distinguish these decompositions, leading to signalling—see Fig.2.1.




Figure 2.1: Schematic depiction of Gisin's argument. By choosing to measure A' or A'' , Alice can prepare one of two equivalent decompositions of the same mixed state, $\rho'_B(t_1)$ or $\rho''_B(t_1)$. If density matrices evolved nonlinearly, Bob could be able to distinguish between the time-evolved $\rho'_B(t_2)$ and $\rho''_B(t_2)$, thus revealing whether A' or A'' was chosen. In particular, this mechanism allows superluminal signalling if Bob's measurement takes place outside the causal future of Alice's measurement.

The authors of [163] claim that the argument just given (as Gisin's argument) does *not* rely on the projection postulate. But if this was the case, then it would also apply to pQT, ruling out any deterministic nonlinear evolution of p-states. However, in analogy to the argument in [119], an experimenter can effectively implement nonlinear deterministic transformations by reconstructing the reduced p-state of a quantum system and applying unitaries based on the result.

The authors justify their claim by stating that the conditional post-measurement state can be inferred from the Born rule alone, which specifies both the single and

joint outcome probabilities,

$$p(o_2 = b_y | o_1 = a_x) = \frac{\text{Tr}(\Pi_x \otimes \Pi_y \rho_{AB})}{\text{Tr}(\Pi_x \otimes \mathbb{I} \rho_{AB})} = \text{Tr} \left\{ \mathbb{I} \otimes \Pi_y \underbrace{\frac{\Pi_x \otimes \mathbb{I} \rho_{AB} \Pi_x \otimes \mathbb{I}}{\text{Tr}(\Pi_x \otimes \mathbb{I} \rho_{AB})}}_{w_{AB}^L(\Pi_x, \rho_{AB})} \right\}. \quad (2.40)$$

The issue with their argument lies in the expectation that the outcome probability distribution of a *single* measurement on $\mathcal{H}_A \otimes \mathcal{H}_B$ (with outcome $\Pi_x \otimes \Pi_y$) coincides with that of *two* local measurements on \mathcal{H}_A (with outcome Π_x) and \mathcal{H}_B (with outcome Π_y), each performed with a separate device. In fact, the use in Eq. (2.40) of

$$p(o_1 = a_x, o_2 = b_y) = \text{Tr}(\Pi_x \otimes \Pi_y \rho_{AB}) \quad (2.41)$$

is tantamount to assuming the operational equivalence between the mono- and multi-partite procedures described in Sec. 2.3.2. The equivalence, however, is not guaranteed by the Born rule alone, as it relies on the update rule for post-measurement states. Therefore, as already pointed out by Holman [105], Gisin's argument is *not* independent of the projection postulate. The supposedly instantaneous and nonlocal collapse, which causes the transition from an improper mixture to a proper mixture, is essential for the argument to apply.

It should be noted that Gisin's argument does not rule out all nonlinear time evolutions [45, 46, 71, 72, 101, 119, 120, 151, 152]. Alternative state-update rules, rather than alternative time evolutions, may also result in nonlinear transformations of the joint and reduced states, and possibly enable signalling. We already mentioned in Sec. 2.2 that the readout devices presented in [118–120] represent an example of measurement-induced transformations *not* dismissed by Gisin-type arguments, that is, consistent with no-signalling. Passive quantum theory, where all measurements are non-collapsing [3], is one such example.

2.8 Local realism

Bell's theorem demonstrates that no physical theory of local hidden variables can reproduce all of the predictions of quantum mechanics [20]. The widely accepted conclusion is that quantum theory is incompatible with the combined notions of *realism* and *locality*. Realism assumes that measurement outcomes are not created during the measurement but reveal pre-existing properties possessed by the system

and independent of whether the measurement is performed. Mathematically, this is modelled by introducing hidden variables specifying all observable properties of the system. Locality assumes that outcome probabilities cannot depend on events happening arbitrarily far away. It is usually regarded as the combination of two assumptions, *(i) parameter independence* and *(ii) outcome independence* [111, 160, 176]. Given two space-like separated parties measuring their own systems, the outcome probability of one measurement is not influenced by the other party's *(i)* choice of measurement or *(ii)* observed outcome. *Free choice* also features among the assumptions of Bell's theorem: it stipulates that the choices of local measurements are independent of the hidden variables.

In pQT, the restriction of space-like separation means that no correlations between entangled systems will be observed. Therefore, Bell's inequalities will not be violated. More generally, this is true for quantum theory with readout devices [118–120] *replacing* standard quantum measurements, albeit not for quantum theory *augmented* by them, which remains nonlocal. Furthermore, causal quantum theory [116, 117] is another example of a local foil obtained by modifying the projection postulate.

Consider the simplest 2-input 2-output CHSH inequality

$$|\langle AB \rangle + \langle A'B \rangle + \langle AB' \rangle - \langle A'B' \rangle| \leq 2, \quad (2.42)$$

where one party measures either $A \otimes \mathbb{I}$ or $A' \otimes \mathbb{I}$ while the other party, with their own separate device, measures either $\mathbb{I} \otimes B$ or $\mathbb{I} \otimes B'$. Implementing mono-partite p-measurements leads to factorised joint outcome probabilities,

$$p(a_x, b_y) = p(a_x)p(b_y), \quad (2.43)$$

for all bipartite p-states ρ_{AB} . Hence the correlations appearing in (2.42) take the form

$$\langle AB \rangle_{\text{pQT}} = \sum_{xy} a_x b_y p(a_x, b_y) = \langle A \otimes \mathbb{I} \rangle \langle \mathbb{I} \otimes B \rangle, \quad (2.44)$$

differing from those predicted by quantum theory,

$$\langle AB \rangle_{\text{QT}} = \langle A \otimes B \rangle. \quad (2.45)$$

Eq. (2.43) implies the existence of a joint probability distribution for all four outcomes,

$$p(a_w, a'_x, b_y, b'_z) = p(a_w) p(a'_x) p(b_y) p(b'_z), \quad (2.46)$$

and the marginals return the correct probabilities for each choice of measurements. The existence of a joint probability distribution entails the constraint of Eq. (2.42), making the toy model compatible with local realism. A similar argument holds for any other Bell-type inequality.

Replacing the Lüders collapse with a passive rule deprives the theory of its nonlocal features, even though all outcome distributions featured in quantum theory can be observed in pQT too. In fact, a passive measurement of $A \otimes B$ implemented with a *single* device \mathcal{D}_{AB} spanning both constituents will return correlations of the form (2.45). However, since this procedure requires access to both subsystems, the obtained violation of Bell's inequality does not imply a rejection of locality.

2.9 Contextuality

A hidden variable model for pQT can be local but cannot be non-contextual, as per the theorems by Kochen and Specker [122] and by Bell [21]. A *context* can be regarded as a collection of observables $\{A, B, C, \dots\}$ that can be jointly measured. For sharp observables, joint measurability coincides with the commutativity of their corresponding self-adjoint operators [95, 96]. An observable can belong to more than one context, e.g. $\mathcal{C}_1 = \{A, B_1, B_2, \dots\}$ and $\mathcal{C}_2 = \{A, B'_1, B'_2, \dots\}$, but the set union $\mathcal{C}_1 \cup \mathcal{C}_2$ may not be jointly measurable. A hidden-variable model is *non-contextual* (à la Kochen-Specker) if the value assigned by the ontic state (or hidden variable) λ to any observable is independent of the context in which it is measured, e.g.

$$v_\lambda(A|\mathcal{C}_1) = v_\lambda(A|\mathcal{C}_2) = v_\lambda(A) . \quad (2.47)$$

For contextuality proofs, it suffices to find a set of observables A, B, C, \dots for which we can show that it is impossible to associate to each operator one of its eigenvalues, $v(A)v(B)v(C), \dots$, such that all functional relationships between mutually commuting subsets of the observables are also satisfied by the associated values. These arguments do not require the notion of a post-measurement state, hence they apply to pQT. They highlight the inherent inconsistency between the assumption of realism supported by Eq. (2.47) and the requirement that observables correspond to self-adjoint operators on a Hilbert space. The coexistence of contextuality and locality in a realist model of pQT can be exemplified using a state-independent version of the Kochen-Specker theorem provided by Mermin [135]. The star pattern




Figure 2.2: The star-shaped pattern of ten observables on \mathcal{H}_8 leading to a proof of the Kochen-Specker theorem. Notice the use of simplified notation, e.g. σ_x^1 instead of $\sigma_x^1 \otimes \mathbb{I}_2 \otimes \mathbb{I}_2$. Each straight line defines a context, and each operator belongs to two contexts.

depicted in Fig. 2.2 shows ten dichotomic observables (with eigenvalues ± 1) on a eight-dimensional qudit, \mathcal{H}_8 , in such a way that any straight line of four operators defines a context. In particular, each observable belongs to exactly two contexts. Notice the use of simplified notation, e.g. σ_x^1 instead of $\sigma_x^1 \otimes \mathbb{I}_2 \otimes \mathbb{I}_2$. Denote the operators by A_i , $i = 1, \dots, 10$ and call C_j , $j = 1, \dots, 5$ the products of the operators on the j -th line. Then $C_j = \mathbb{I}$ for all contexts $j \neq 5$, where $j = 5$ denotes the horizontal line, for which $C_5 = -\mathbb{I}$. If we assume the existence of a hidden variable λ that assigns all measurement outcomes, then a measurement of C_i reveals the product of the eigenvalues assigned to the four operators in the context. In particular, if we also assume that the assignment happens non-contextually, then the following must hold,

$$\left[\prod_{i=1}^{10} v_\lambda(A_i) \right]^2 = \prod_{j=1}^5 v_\lambda(C_j). \quad (2.48)$$

However, $(\pm)\mathbb{I}$ represents the trivial observable that always returns the outcome $(\pm)1$ regardless of the p-state $\rho \in \mathcal{S}(\mathcal{H}_8)$. Hence, $v_\lambda(\mathbb{I}) = +1$ and $v_\lambda(-\mathbb{I}) = -1$ for any λ . This means that the right-hand side of (2.48) evaluates to -1 , leading to a contradiction. We conclude that, as with quantum theory, pQT cannot be reproduced by a hidden variable model that assigns values to observables independently of what

other observables are jointly measured.

In [135], Mermin shows that, by an appropriate change of the assumptions, the same scenario can be recasted into the inequality-free proof of Bell's theorem provided by Greenberg, Horne and Zeilinger [87, 134]. Instead of considering the ten operators as observables on an eight-dimensional qudit, they should be viewed as pertaining to measurements performed on three space-like separated qubits. Each context is then specified by three independent devices that register the value assigned to an observable of the form σ_j^i , $i \in \{1, 2, 3\}$ and $j = \{x, y\}$. Without communication between parties, no one can know which of the two contexts including their chosen measurement describes the scenario.

Instead of assuming non-contextuality for all observables, we only do so when it can be justified by locality. In other words, we impose Eq. (2.47) only to the six observables of the form σ_j^i , for which a dependency on the context implies a nonlocal influence. Dropping non-contextuality for the product-observables in the horizontal line breaks the chain of relations that leads to the contradiction. However, in quantum theory the argument can be restored by making it state-dependent. Suppose the three qubits reside in a GHZ state, i.e. an eigenstate of the four commuting product-operators. Then the values assigned to the product-observables cannot differ from one context to the other, $v_\lambda(\sigma_{j_1}^1 \sigma_{j_2}^2 \sigma_{j_3}^3 | \mathcal{C}_j) = v_\lambda(\sigma_{j_1}^1 \sigma_{j_2}^2 \sigma_{j_3}^3 | \mathcal{C}_{j'})$. If the three parties were to share notes after performing their measurements, they would always find that the products of their three outcomes equal the eigenvalue corresponding to the GHZ state. This is a consequence of the fact that local measurements of $\sigma_{j_1}^1$, $\sigma_{j_2}^2$ and $\sigma_{j_3}^3$, followed by classical communication, are, as a whole, equivalent to a measurement of $\sigma_{j_1}^1 \sigma_{j_2}^2 \sigma_{j_3}^3$. Therefore, by restricting to such states, we can derive Eq. (2.48) using only the assumption of locality. Then, the contradiction demonstrates the incompatibility of quantum theory with local realism without the use of inequalities.

In pQT, the use of GHZ states ensures that the values assigned to the product-observables are independent of the context. However, in contrast with standard quantum theory, the use of three devices, each registering one outcome, followed by classical communication between the parties, *cannot* be interpreted as a passive measurement of $\sigma_{j_1}^1 \sigma_{j_2}^2 \sigma_{j_3}^3$. This means that, in general,

$$v_\lambda(\sigma_{j_1}^1) v_\lambda(\sigma_{j_2}^2) v_\lambda(\sigma_{j_3}^3) \neq v_\lambda(\sigma_{j_1}^1 \sigma_{j_2}^2 \sigma_{j_3}^3). \quad (2.49)$$

As a result, the chain of relations leading to the contradiction breaks, invalidating the locality-based argument and in agreement with our earlier discussion of Sec. 2.8.

To observe the correlations, the whole context must be measured jointly by a *single* device, which is only possible when the subsystems are brought together in the same place. Our conclusions are summarised in the following remark.

Remark. In pQT, the assignment of observable quantities by the hidden variable λ must be context-dependent if the operators composing the context are measured *jointly* by a single device. However, when the context is specified by multiple devices, the predictions of pQT are consistent with a non-contextual assignment.

2.10 An ontological model for pQT

We propose an explicit local and contextual *ontological model* that is consistent with the predictions of passive quantum theory. It is based on the so-called ‘‘Bell model’’ [125], generalising the method to explain qubit measurements deterministically [21]. In this section, we use the term ‘‘local’’ in a relativistic sense, meaning that changes of the ontic state of a system that is spread across multiple regions do not occur superluminally. Another model for pQT, which is local but not deterministic, can be obtained by identifying the *ontic state* of a system with its *local state*, along the lines described in [116].

We employ a slightly different notion of an ontological model compared to that in, say, [125], which includes only (generalised versions of) the first four items from the list below. In the present context, an ontological model consists of:

- a measurable space (Λ, Σ) called the *ontic space*, where Λ is the set of *ontic states* and Σ a σ -algebra on Λ ;
- a probability measure μ^ψ over (Λ, Σ) for each p-state $\psi \in \mathcal{H}_d$;
- the probability $\Pr(x|\lambda, \mathcal{D})$ that a measurement performed with device \mathcal{D} on one or more constituents of a system described by the ontic state $\lambda \in \Lambda$ returns the outcome x ;
- a probability measure $T(U, \lambda)$ on (Λ, Σ) describing the update of the ontic state following the implementation of a unitary gate U on one or more constituents of a system described by $\lambda \in \Lambda$;
- a probability measure $W(x, \lambda, \mathcal{D})$ on (Λ, Σ) describing the (possibly outcome-dependent) update of the ontic state following a measurement with device \mathcal{D} .

The ontic space. Consider a d -dimensional p-system, where $d = p_1^{n_1} p_2^{n_2} \dots$ is the unique prime factorisation of d . Let

$$P_d = (\underbrace{p_1, p_1, \dots, p_1}_{n_1 \text{ times}}, \underbrace{p_2, p_2, \dots, p_2}_{n_2 \text{ times}}, \dots) \quad (2.50)$$

be the finite tuple of prime numbers appearing in the decomposition, including repetitions. Denoting the elements of P_d with s_i , $i = 1, \dots, |P_d|$, we can write

$$\mathcal{H}_d = \bigotimes_{i=1}^{|P_d|} \mathcal{H}_{s_i}. \quad (2.51)$$

The ontic space of the system is assumed to be

$$\Lambda = \Lambda_1 \times \Lambda_2 \quad \text{where} \quad \Lambda_1 = P(\mathcal{H}_d)^{\times |P_d|}, \quad \Lambda_2 = [0, 1]^{\times |P_d|}, \quad (2.52)$$

where $P(\mathcal{H}_d)$ is the set of rays in \mathcal{H}_d and $|P_d|$ is the cardinality of P_d . An equivalent and useful factorisation is

$$\Lambda = \bigotimes_{i=1}^{|P_d|} (\Lambda_1^{s_i} \times \Lambda_2^{s_i}) \quad \text{where} \quad \Lambda_1^{s_i} = P(\mathcal{H}_d), \quad \Lambda_2^{s_i} = [0, 1]. \quad (2.53)$$

Then, the ontic state of a single d -dimensional p-system takes the form

$$\lambda = ((\lambda_1^{s_1}, \lambda_2^{s_1}), (\lambda_1^{s_2}, \lambda_2^{s_2}), \dots, (\lambda_1^{s_{|P_d|}}, \lambda_2^{s_{|P_d|}})) \quad (2.54)$$

or, equivalently, $\lambda = (\lambda_1, \lambda_2)$ where

$$\lambda_1 = (\lambda_1^{s_1}, \lambda_1^{s_2}, \dots, \lambda_1^{s_{|P_d|}}) \in \Lambda_1, \quad \lambda_2 = (\lambda_2^{s_1}, \lambda_2^{s_2}, \dots, \lambda_2^{s_{|P_d|}}) \in \Lambda_2. \quad (2.55)$$

In other words, a two-variable ontic state $(\lambda_1^{s_i}, \lambda_2^{s_i})$ is assigned to each subsystem (or degree of freedom) and the ontic state of the entire system is the Cartesian product of the ontic states of all subsystems. The variable $\lambda_1^{s_i}$ can be regarded as subsystem s_i 's own description of the composite system it is a part of, whereas $\lambda_2^{s_i}$ is s_i 's “hidden” variable that is inaccessible to observers. Anticipating what will be discussed in a later paragraph, the reason we assign to each subsystem their own description of the composite system is to model the time evolution of states, i.e. the mapping T of the ontological model, in a local way, in the sense that implementing a unitary gate on a subsystem does not superluminally update the ontic state of a distant subsystem. Call $\Sigma_1^{s_i}$ and $\Sigma_2^{s_i}$ the Borel σ -algebras of $\Lambda_1^{s_i}$ and $\Lambda_2^{s_i}$, respectively. Then $\Sigma = \Sigma_1 \otimes \Sigma_2$, where $\Sigma_1 = \bigotimes_{i=1}^{|P_d|} \Sigma_1^{s_i}$ and $\Sigma_2 = \bigotimes_{i=1}^{|P_d|} \Sigma_2^{s_i}$, is the σ -algebra of Λ ,

generated by elements of the form $\Omega = \Omega_1 \times \Omega_2$, $\Omega_1 \in \Sigma_1$, $\Omega_2 \in \Sigma_2$ [125]. The tuple (Λ, Σ) defines a measurable space.

The p-state $\psi \in \mathcal{H}_d$ is represented by the product measure

$$\mu^\psi(\Omega) = \int_{\Lambda_2} \mu_1^\psi(\Omega_{\lambda_2}) d\mu_2^\psi(\lambda_2) , \quad (2.56)$$

where μ_i^ψ is a probability measure on Λ_i and

$$\Omega_{\lambda_2} = \{\lambda_1 \in \Lambda_1 \mid (\lambda_1, \lambda_2) \in \Omega\} . \quad (2.57)$$

Suppose a procedure involving the entire system—i.e. a global unitary—prepares the p-state ψ . At the end of it, all constituents will share the same description of the joint state, $\lambda_1 = (\psi, \psi, \dots)$. This can be achieved by setting

$$\mu_1^\psi(\lambda_1) = \prod_{i=1}^{|P_d|} \delta(\lambda_1^{s_i} - \psi) . \quad (2.58)$$

The “hidden” variable λ_2 , on the other hand, is uniformly distributed, i.e. $\mu_2^\psi = \mu_2$ is the uniform measure on Λ_2 . The underlying mechanism producing λ_2 —whether intrinsically probabilistic or deterministic—is irrelevant to our analysis. What matters is that the model is *realist*, i.e. all observable properties of the system are uniquely determined by (λ_1, λ_2) which take on definite values at any point in time and regardless of the measurements we may or may not perform.

The measurement process. The behaviour of p-measurements (and sequences thereof) is modelled as follows. For the moment, we will assume that all subsystems agree in their description of the composite system they belong to, i.e. $\lambda_1^{s_i} = \psi$ for all i . Let $\mathcal{D}_{(1,r)}$ denote a device implementing a measurement on the first $r \leq |P_d|$ subsystems in (2.51) of the observable represented by $\{\Pi_x\}_{x=1}^N$, $\Pi_x \in \mathcal{P}(\otimes_{i=1}^r \mathcal{H}_{s_i})$ with $N \leq \dim(\otimes_{i=1}^r \mathcal{H}_{s_i})$. We can consider this scenario without loss of generality, since any device acts like $\mathcal{D}_{(1,r)}$ for a suitable ordering of the elements of (2.50). A *global* measurement, with the entire system as input, is one where $r = |P_d|$.

Consider the unit r -cube $[0, 1]^r$. The device $\mathcal{D}_{(1,r)}$ splits the hypercube into exactly N non-overlapping regions $\{R_x\}_{x=1}^N$, each corresponding to a possible measurement outcome x . Each region has volume

$$V(R_x) = \text{Tr}(\Pi_x \rho_{(1,r)}) , \quad (2.59)$$

where $\rho_{(1,r)}$ is the reduced state of the first r subsystems,

$$\rho_{(1,r)} = \text{Tr}_{s_j, j>r} (|\psi\rangle\langle\psi|) . \quad (2.60)$$

Notice that the locally available entries of λ_1 , $\lambda_1^{(1,r)} = (\lambda_1^{s_1}, \dots, \lambda_1^{s_r})$ are sufficient to determine $\rho_{(1,r)}$. The arrangement of the regions in the unit r -cube and their shape do not matter and can be assumed to be determined by the device. The result of the measurement will then be the value x labelling the region where the point $\lambda_2^{(1,r)} = (\lambda_2^{s_1}, \dots, \lambda_2^{s_r}) \in [0, 1]^r$ resides. In other words, the function Pr of our ontological model is defined by the formula

$$\text{Pr} \left(x | \lambda_1^{(1,r)}, \lambda_2^{(1,r)}, \mathcal{D}_{(1,r)} \right) = \begin{cases} 1 & \text{if } \lambda_2^{(1,r)} \in R_x \\ 0 & \text{otherwise} \end{cases} . \quad (2.61)$$

Compliance with the Born rule is ensured by Eq. (2.59). The measurement also leads to an update of the ontic state which is outcome-independent,

$$\lambda_1^{s_i} \mapsto \lambda_1^{s_i} \quad i \in \{1, \dots, |P_d|\} \quad (2.62a)$$

$$\lambda_2^{s_i} \mapsto \begin{cases} \widetilde{\lambda}_2^{s_i} & i \leq r \\ \lambda_2^{s_i} & i > r \end{cases} \quad (2.62b)$$

where each $\widetilde{\lambda}_2^{s_i}$ is drawn uniformly from the unit interval $[0, 1]$. Formally, Eqs. (2.62) characterise the update map W of our ontological model. In other words, the measurement does not affect the description of the composite p-state associated to *any* subsystem, $\lambda_1 \mapsto \lambda_1$, and does not update the hidden variables corresponding to the subsystems that did not interact with the device.⁷ Assuming uniform measure for $\widetilde{\lambda}_2^{s_i}$ ensures that repetitions of the same measurement will agree with the probabilistic repeatability of pQT. Single-copy state reconstruction will return the reduced state $\rho_{(1,r)}$, as expected.

To update the hidden variables $\lambda_2^{s_i}$ of all the subsystems, a device must necessarily act on the entire system, $r = |P_d|$. Furthermore, for any $r_1 < |P_d|$, a measurement by $\mathcal{D}_{(1,r_1)}$ of $\{\Pi_x\}$ is operationally indistinguishable from a measurement of $\{\Pi_x \otimes \mathbb{I}_m\}$ performed by $\mathcal{D}_{(1,r_2)}$ on $r_2 > r_1$ subsystems, where $m = \prod_{r_1 < i \leq r_2} s_i$. This follows from the use of the Born rule in (2.59) which ensures that the two procedures return

⁷Notice that the mechanism can be made consistent with quantum measurements if one sets $\lambda_1^{s_i} \mapsto \Pi_x \otimes \mathbb{I} \lambda_1^{s_i} / \sqrt{p(x)}$ for all i .

the same outcome statistics for any initial p-state $\psi \in \mathcal{H}_d$. Nevertheless, the two measurements are not identical at the ontic level, as the number of hidden variables updated by the two devices is different, r_1 for \mathcal{D}_1 and r_2 for \mathcal{D}_2 . However, such discrepancy cannot be observed.

Contextuality and locality. The model complies with the contextual behaviour of pQT. Consider the following example in $d = 3$, hence $|P_3| = 1$, with the two qutrit observables $\mathcal{C}_1 = \{\Pi_a, \Pi_b, \Pi_c\}$ and $\mathcal{C}_2 = \{\Pi_x, \Pi_b, \Pi_y\}$. The sets define two different contexts for the operator Π_b . Let the system reside in $\lambda = (\lambda_1, \lambda_2)$ with λ_1 not an eigenvector of any of the five projectors above. Then, contextual behaviour can arise from the different ways \mathcal{C}_1 and \mathcal{C}_2 can split the unit hypercube (here the unit interval). Fig. 2.3 shows an example where $v_\lambda(\Pi_b|\mathcal{C}_1) = 1$ but $v_\lambda(\Pi_b|\mathcal{C}_2) = 0$. Notice that there will be values of λ_2 for which $v_\lambda(\Pi_b|\mathcal{C}_1) = v_\lambda(\Pi_b|\mathcal{C}_2)$ and these will generally depend on how the regions are arranged—in Fig. 2.3 the order follows the appearance of the operators in the sets.




Figure 2.3: An example of how the unit interval can be split according to the elements of $\mathcal{C}_1 = \{\Pi_a, \Pi_b, \Pi_c\}$ (above) and $\mathcal{C}_2 = \{\Pi_x, \Pi_b, \Pi_y\}$ (below), given the ontic state (λ_1, λ_2) . The shaded areas highlight the values of λ_2 for which $v_\lambda(\Pi_b|\mathcal{C}_1) \neq v_\lambda(\Pi_b|\mathcal{C}_2)$.

The machinery behind a measurement relies exclusively on the properties of the *single* device implementing it and on the ontic states of the measured subsystems. As a result, no contextual behaviour will arise whenever the context is defined by *multiple* devices, in accordance with the results of Sec. 2.9. In particular, if we assume that the un-measured subsystems are at space-like separation, then the use of locally available information, along with the restriction to only disturb local hidden variables, ensure that measurements in the model are an inherently local process. The model is also compatible with the assumption that measurement settings are independent of λ .

Lack of correlations. The following example aims to clarify how the model works for mono- and multi-partite measurements on a pair of qubits. Let $d = 4$, hence $P_4 = \{2, 2\}$ and consider two qubit observables $A = \{\Pi_0^A, \Pi_1^A\}$ and $B = \{\Pi_0^B, \Pi_1^B\}$, along with the two-qubit observable $M_{AB} = \{\Pi_0^A \otimes \Pi_0^B, \Pi_0^A \otimes \Pi_1^B, \Pi_1^A \otimes \Pi_0^B, \Pi_1^A \otimes \Pi_1^B\}$. A passive measurement of M_{AB} can only be realised with a single measurement device. Let $\lambda = (\lambda_1, \lambda_2)$ be the ontic state with $\lambda_1^A = \lambda_1^B = \alpha|0_A 0_B\rangle + \beta|1_A 1_B\rangle$ where $\alpha, \beta \neq 0$. Fig. 2.4(a) illustrates one way to arrange the two non-zero volume regions in the unit square, along with the measurement outcome. If A and B are instead measured with two separate devices acting on subsystem s_A and s_B , respectively, then Fig. 2.4(b) illustrates how the modulo-2 sum of the outcomes can happen to be 1, in agreement with the predictions of pQT.




Figure 2.4: (a). Example of a measurement of M_{AB} with outcome ‘00’ on the ontic state $((\lambda_1^A, \lambda_2^A), (\lambda_1^B, \lambda_2^B))$, where $\lambda_1^i = \alpha|0_A 0_B\rangle + \beta|1_A 1_B\rangle$. (b). Example of two distinct p-measurements of A and B performed on $(\lambda_1^A, \lambda_2^A)$ and $(\lambda_1^B, \lambda_2^B)$, respectively, for which $v_\lambda(A) + v_\lambda(B) = 1 \pmod{2}$.

Unitary evolution. So far, we have assumed that λ_1 is always composed of $|P_d|$ copies of the prepared p-state ψ . The reason for assigning one vector $\lambda_1^{s_i} \in P(\mathcal{H}_d)$ to each subsystem lies in the need to model unitary evolution in a local way, as done for measurements. In fact, had we set $\Lambda_1 = P(\mathcal{H}_d)$, then a unitary on, say, \mathcal{H}_{s_2} alone, U_{s_2} , would change the p-state of the entire system to $U_{s_2}\psi$. A measurement on a space-like separated subsystem, say \mathcal{H}_{s_1} , would then use the updated $U_{s_2}\psi$ in Eq. (2.60) to calculate the volumes via Eq. (2.59). This would effectively amount

to a nonlocal influence between the two subsystems as a result of the ψ -ontology of the model: if the p-state is part of the ontic state, then any change of it can potentially establish an instantaneous interaction between systems. Assigning to each subsystem their own description of the composite system, $\lambda_1^{s_i}$, is a way to avoid such non-classical behaviour. By doing so, we can establish the mapping T within our ontological model, whereby a local unitary on \mathcal{H}_{s_2} only updates $\lambda_1^{s_2}$, leaving unchanged the other entries of λ_1 , at least until enough time has passed for a signal to reach the corresponding subsystems. Therefore, in the case of space-like separation, a device measuring on \mathcal{H}_{s_1} will still compute the volumes using $\lambda_1^{s_1} = \psi$. The correct statistics will be obtained because applying U_{s_2} does not affect the reduced state of s_1 . We have thus introduced a mechanism resembling Kent's local-state update⁸ [119], which accounts for modifications of $\lambda_1^{s_i}$ due to unitary gates applied in the causal past of s_i . This way, by the time s_1 and s_2 are brought to the same location, their corresponding entries in λ_1 will be identical, bringing us back to the previously examined scenario. In addition, applying a joint unitary $U_{s_1 s_2}$ on s_1 and s_2 will cause a new update of the locally available variables, $\lambda_1^{s_1} = \lambda_1^{s_2} = U_{s_1 s_2} U_{s_2} \psi$, which will in turn extend to the remaining entries of λ_1 without resorting to superluminal influences.

Taking into account unitary evolution, we provide a slight modification of the formula for reduced states. Since, in order to perform a passive measurement on $\bigotimes_{i=1}^r \mathcal{H}_{s_i}$, the r systems must not be space-like separated⁹, we conclude that at the time of the measurement, $\lambda_1^{(1,r)} = (\lambda_1^{s_1}, \dots, \lambda_1^{s_r})$ will be composed of equal entries, hence

$$\rho_{(1,r)} = \text{Tr}_{s_j, j>r} (|\lambda_1^{s_i}\rangle\langle\lambda_1^{s_i}|) \quad \text{for any } i \leq r. \quad (2.63)$$

Eq. (2.60) constitutes a special case where no unitary gate is applied on any subsystem in the causal past of the measurement event.

The model is ψ -ontic but not ψ -complete, as the ontic state includes the quantum state as well as other variables [125]. Nevertheless, it may be the case that the

⁸Notice that, while we use it to model unitary evolution of the local ontic states, the original purpose of Kent's mechanism was to propagate measurement-induced disturbances.

⁹It is known that, in quantum theory, measurements of nonlocal observables can in principle be performed on spatially separated systems via local operations and classical communication. However, existing implementations [88, 175] rely on quantum state teleportation, which necessitates the collapse of the state of the composite system as a result of local measurements. Consequently, these protocols cannot be implemented in pQT.

“true” p-state of the system does not appear in the entries of λ_1 for some time. It should be mentioned that we have only focused on pure states of the composite system. To account for entanglement with additional systems, one could further assume that the mechanism for updating λ_1 can replace each entry with vectors from higher-dimensional Hilbert spaces, effectively allowing the ontic space Λ to change with time.

The model raises several unanswered questions, such as how each degree of freedom can carry such a large amount of information, how the update mechanism à la Kent can work, how the device arranges regions in the hypercube, and how the hidden variables are generated. However, what we can conclude is that suspending the collapse enables a fully classical explanation of all the predictions of the theory.

2.11 Quantum information with p-measurements

The state update induced by quantum measurements is essential for many protocols in quantum information. Teleportation [25] and entanglement swapping [143], for example, rely on system-wide state changes as a result of local measurements. Thus, they will no longer work in pQT. The impossibility to “steer” the state of a distant subsystem means that quantum key distribution protocols based on entangled states [70] are also ruled out. At the same time, single-copy reconstruction would allow for perfect eavesdropping on p-states, i.e. without leaving a trace. The security of protocols like BB84 [24] would therefore be compromised.

Both adding collapse-free measurements to standard quantum theory and replacing collapsing measurements by them will modify its computational power, as shown in [3, 120]. In the following paragraphs, we review some key implications, with further details available in these papers.

“*Quantum parallelism*” may be exploited in full since the p-state encoding the result of a quantum computation is observable, at least in principle [3]. The algorithms by Deutsch and Jozsa [56], Grover [90] and Simon [164], for example, involve “oracles” which “evaluate” a function $f(x)$ by means of a unitary operator, viz. $U_f : |x, 0\rangle \mapsto |x, f(x)\rangle$. Letting the linear operator U_f act on the symmetric superposition $|s\rangle = 2^{-n/2} \sum_{x=0}^{2^n-1} |x, 0\rangle$, the output state carries information about the function $f(x)$ for all its values. A projective quantum measurement on the final state $U_f|s\rangle$ will, however, reveal at most one value of $f(x)$, necessitating for

further calls of the oracle. As mentioned in Sec. 2.2, a single implementation of an infinite-precision local-state readout device (SRD) can reveal all the values of $f(x)$ [119, 120], making it possible to solve NP problems in polynomial time [5]. In pQT, however, passive measurements only provide partial information about the state of a system. Nevertheless, all values $f(x)$ can be extracted from the final p-state $U_f|\widehat{s}\rangle$ with at least 2^n measurements. Hence, only a single call to the oracle is necessary within pQT which represents a substantial reduction in computational cost, but a large increase in *measurement complexity*. The improvement does not violate arguments of optimality such that for Grover’s algorithm [189]: by restricting to a single p-measurement, it is still impossible to produce a more reliable solution using fewer oracle calls (e.g. $O(\sqrt{2^n})$ for Grover’s search algorithm).

In quantum algorithm complexity, the cost of a single measurement is generally ignored due to the collapse making measurement sequences irrelevant. However, for p-algorithms it is essential to estimate the resource cost of each measurement to determine their feasibility [3]. The advantages provided by p-measurements are more evident when the structure of the specific problem is leveraged. For instance, without requiring additional measurements beyond the ones outlined in Simon’s algorithm, in pQT a single implementation of Simon’s routine is sufficient. This amounts to a linear improvement in terms of query complexity with no increase in measurement complexity. Similarly, with even a single additional p-measurement at the end of Shor’s factorisation algorithm [161], one eliminates the need for continued fractions. This step is required in standard quantum theory to extract powers of prime factors from the only available measurement outcome.

Although shared entanglement has no use for communication in pQT, a single p-qubit can potentially carry an unlimited amount of information [3]. In fact, an arbitrary n-bit message $m = b_1b_2\dots b_n$, $b_i \in \{0, 1\}$, can be encoded in, say, the (real) amplitude of its p-state, such as $\alpha = \sum_{i=1}^n b_i 2^{-i}$ in $\alpha|\widehat{0}\rangle + \beta e^{i\phi}|\widehat{1}\rangle$. After the system is transmitted to a receiver, it can *in principle* be decoded using a sequence of p-measurements.

2.12 Summary and discussion

We have reviewed a collapse-free foil theory of quantum mechanics by assuming that measurements do not cause states to update, cf. Sec. 2.1. Such “passive”

measurements were first introduced in [3] and correspond to a class of hypothetical devices presented in [120]. The predictions of this model, pQT, agree with those of quantum theory as long as its post-measurement states are discarded systematically. Any non-quantum feature of pQT can be traced back to states not collapsing when measuring generic observables. Being manifestly different from quantum theory, the model represents a tool to study the role of the collapse rather than suggesting an alternative interpretation aiming to circumvent the projection postulate.

The possibility of single-copy state reconstruction turns p-states into observable quantities [33, 119], albeit only by virtue of infinitely long sequences of p-measurements, cf. Sec. 2.2, Sec. 2.3.1. As any time-evolved state can in principle be accessed directly in pQT, the cost and computational power of known quantum algorithms must be evaluated anew [3], cf. Sec. 2.11. Assuming that the measured state has been reconstructed—a procedure that, for entangled states, requires access to all subsystems—pQT can simulate quantum theory *including* the collapse if one accepts a time delay in state updates, cf. Sec. 2.6. In contrast, quantum theory cannot simulate pQT.

In standard quantum theory, projective measurements can be used to prepare specific states. In pQT, a desired state can only be prepared dynamically, i.e. by suitably evolving a known state in time.

The comparison with pQT shows that some concepts of quantum theory are equivalent because of the non-trivial state update described by the standard projection postulate. As is well-known, proper and improper mixtures of quantum states are indistinguishable. This is no longer true in pQT where passive measurements can be used to reveal each of the individual states forming a proper mixture, cf. Sec. 2.4. This, in turn, affects the valid interpretations for the physical realisation of quantum channels, cf. Sec. 2.5.1. Similarly, in quantum theory the observable $A \otimes B$ can be measured by either a single global device or by two local devices implementing A and B separately, if supplemented by classical communication. In pQT, these two scenarios lead to entirely different outcome statistics, cf. Sec. 2.3.2. A consequence of the distinguishability between these two operational procedures is the impossibility to perform local tomography. Furthermore, Bell's inequalities can no longer be violated in the usual setting, cf. Sec. 2.8, meaning that the predictions of pQT are consistent with local realism. However, since the model shares the same state and observable spaces of quantum theory, it necessarily exhibits contextual behavior, cf. Sec. 2.9. In Sec. 2.10, we presented a fully-fledged deterministic hidden variable

model for pQT consistent with both relativistic locality and contextuality.

It is not difficult to see that passive measurements represent just one specific case of possible alternative state-update rules leading to quantum-like theories. However, if consistent modifications of the projection postulate exist, an intriguing question arises: can we identify a physical principle which singles out the quantum mechanical update rule among these alternatives?

Quantum theories with alternative measurements

3.1 The AMT framework

3.1.1 Construction

The analysis of pQT given in Chap. 2 shows that modifications of the projection postulate can give rise to consistent quantum-like theories. Comparing them with quantum mechanics allows us to identify the subtle ways in which the projection postulate shapes quantum theory. The ontology of quantum states, the incompatibility with “local realism” and the operational equivalence between different procedures for measuring local observables were all shown to depend on the specific form of the standard update rule w^\perp . Other features of the theory, such as contextuality and preparational uncertainty relations, are independent of it.

Passive measurements offer an alternative to the collapse, but they are not the only option. In this chapter, we will define a wide variety of *Alternative-Measurement Theories* (AMTs), which are identical to quantum mechanics except for the assigned post-measurement states. The AMT *framework* allows us to explore questions that go beyond the scope of Chap. 2. For example, we can search for principles that single out the quantum mechanical update rule amongst other operationally meaningful alternatives. Our approach also suggests looking for toy models that successfully simulate quantum measurements, despite being manifestly different from quantum mechanics.

To define AMTs, we introduce the concept of an *update rule* w^A which characterises

the transformation of states resulting from a measurement performed on a system (or a subsystem)¹. Recall that all other axioms of quantum theory will remain unchanged—cf. Sec. 1.1. We begin by considering measurements on a *single* system, generically labeled with ‘A’ in a serif font, as done throughout Chap. 1 and 2. Abstracting from standard and passive quantum theories—i.e. from (\mathbf{M}^L) and (\mathbf{M}^P) , respectively—we require that every pair of a pre-measurement state of the system and a sharp measurement outcome (associated with a projector) uniquely determines a post-measurement state, i.e. \mathbf{w}_A^A is a mapping from $\mathcal{P}(\mathcal{H}_A) \times \mathcal{S}(\mathcal{H}_A)$ to $\mathcal{S}(\mathcal{H}_A)$. From an abstract point of view, we assume two properties: (R1) *completeness*, stating that every measurement that might be performed must specify a valid post-measurement state of the system; and (R2) *context-independence*, according to which the post-measurement state does not depend on the context of Π_x , i.e. it may have been measured along with any other commuting set of projectors.

We formalise these features in the alternative postulate (\mathbf{M}^A) .

(\mathbf{M}^A) If a measurement outputs the outcome x , then a pre-measurement state ρ is updated to the normalised post-measurement state as follows:

(i) If ρ is a pure state, then

$$\rho \xrightarrow{x} \mathbf{w}_A^A(\Pi_x, \rho). \quad (3.1)$$

(ii) If ρ is a proper mixture with Gemenge $\mathcal{G} = \{(p_1, \rho_1), \dots, (p_n, \rho_n)\}$, then

$$\rho = \sum_{i=1}^n p_i \rho_i \xrightarrow{x} \sum_{i=1}^n p(i|x) \mathbf{w}_A^A(\Pi_x, \rho_i), \quad (3.2)$$

where $p(i|x) = p_i p(x|i)/p(x)$.

As improper mixtures arise from measurements on entangled states of composite systems, they do not feature in (\mathbf{M}^A) , as will be explained in the following paragraphs. To simplify the notation in the single-system case, we define $\mathbf{w}_x^A(\cdot) \equiv \mathbf{w}_A^A(\Pi_x, \cdot)$ as the mapping from $\mathcal{S}(\mathcal{H}_A)$ to itself representing the assignment of states conditioned on observing outcome x .

Having defined the effect of measurements on single systems, we now need to describe their action on composite systems, since a system entering a measurement device may be entangled with other systems. In other words, an update rule must also account for the effect on the states of *composite* systems when measurements

¹We will utilise the sans-serif superscript A to denote a generic AMT.

are carried out only on some subsystems. We realise that, in analogy with quantum theory and pQT (cf. Sec. 2.3.2), Axiom (M^A) is operationally incomplete and a generalised postulate (M_{\otimes}^A) must be introduced. To do so, we need to define what constitutes a valid update rule in the AMT framework.

Definition 9 (*AMT update rule*). Let \mathcal{H}_A and \mathcal{H}_B be finite-dimensional Hilbert spaces. An *update rule* w^A is a family of functions $\{w_{AB}^A\}_{AB}$ (one for each dimension of \mathcal{H}_A and \mathcal{H}_B) satisfying (R1) *completeness*, (R2) *context-independence*—i.e. $w_{AB}^A : \mathcal{P}(\mathcal{H}_A) \times \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B) \rightarrow \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$ —and two additional conditions:

(R3) *Self-consistency*: let $\mathcal{H}_B = \mathcal{H}_{B'} \otimes \mathcal{H}_{B''}$ (hence $w_{AB}^A \equiv w_{AB'B''}^A$), then for all $\mathcal{H}_{B''}$, $\rho_{AB'B''} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_{B'} \otimes \mathcal{H}_{B''})$ and $\Pi_x \in \mathcal{P}(\mathcal{H}_A)$, subsystem update rules must satisfy the relation

$$\mathrm{Tr}_{B''} \left[w_{AB'B''}^A (\Pi_x, \rho_{AB'B''}) \right] = w_{AB'}^A (\Pi_x, \mathrm{Tr}_{B''} (\rho_{AB'B''})) ; \quad (3.3)$$

(R4) (*Quantum*) *No-signalling*: for all $\rho_{AB} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$ and all observables $\{\Pi_x\}$ on \mathcal{H}_A ,

$$\sum_x p(x) \mathrm{Tr}_A \left[w_{AB}^A (\Pi_x, \rho_{AB}) \right] = \mathrm{Tr}_A (\rho_{AB}) ; \quad (3.4)$$

Given these conditions, we can interpret $w_{AB}^A (\Pi_x, \rho_{AB})$ as the normalised state of the composite system ‘AB’, initially residing in ρ_{AB} , after a measurement has been performed on subsystem ‘A’ resulting in outcome x . The subscript of w^A indicates the (possibly composite) system for which the output state is returned, while the subsystem on which the measurement is performed can be inferred from the Hilbert space on which the projector Π_x is defined. Def. 9 considers measurements on subsystem ‘A’, but update rules are of course assumed to account for measurements on *any* constituent, e.g. $\Pi_y \in \mathcal{P}(\mathcal{H}_B)$.

In line with our previous discussion on single systems, we require post-measurement joint states to be (R1) always well-defined and (R2) uniquely determined by the observed local outcome Π_x and the pre-measurement state ρ_{AB} . In addition, condition (R3) ensures that the assignment of states is unambiguous. In particular, for any outcome Π_x registered by a device acting on ‘A’, the output state assigned to the composite system ‘AB’ can also be obtained from the output state assigned to the larger composite system ‘AB’B’’. The assignment of post-measurement states is required to be independent of the observer’s choice of how to partition the inaccessible part of the system. Condition (R4), on the other hand, assures that local

measurements, in the sense of measurements on distant subsystems, cannot be used to transmit information—cf. Sec. 1.2.4.

The single-system mapping of Eq. (3.1) is now obtained as a special case. Let $\rho_A = \text{Tr}_B(\rho_{AB})$, using (R3) we set

$$\mathbf{w}_A^A(\Pi_x, \rho_A) = \text{Tr}_B \left[\mathbf{w}_{AB}^A(\Pi_x, \rho_{AB}) \right]. \quad (3.5)$$

Eq. (3.5) describes the update of the state of ‘A’ when initially entangled to ‘B’. Therefore, it specifies how *improper mixtures* update following measurements.

An AMT is defined by replacing the Lüders rule \mathbf{w}^L with an update rule satisfying Def. 9.

(M $_{\otimes}^A$) If a measurement on system ‘A’ outputs the outcome x , then the joint pre-measurement state ρ_{AB} of the composite system ‘AB’ is updated to the normalised joint post-measurement state according to the *update rule* \mathbf{w}^A as follows:

(i) If ρ_{AB} is a pure or improper mixed state, then

$$\rho_{AB} \xrightarrow{x_A} \mathbf{w}_{AB}^A(\Pi_x, \rho_{AB}); \quad (3.6)$$

(ii) If ρ_{AB} is a proper mixture with Gemenge $\mathcal{G} = \{(p_1, \rho_{AB}^1), \dots, (p_n, \rho_{AB}^n)\}$, then

$$\rho_{AB} = \sum_{i=1}^n p_i \rho_{AB}^i \xrightarrow{x_A} \sum_{i=1}^n p(i|x) \mathbf{w}_{AB}^A(\Pi_x, \rho_{AB}^i), \quad (3.7)$$

where $p(i|x) = p_i p(x|i)/p(x)$.

By considering proper and improper mixtures separately, we allow the AMT framework to include model theories such as pQT where sequential measurements can be used to distinguish between the two classes.

Definition 10. An *Alternative-Measurement Theory* (AMT) is a theory defined by the quantum postulates—cf. Sec. 1.1—for states (S), time evolution (T), system composition (C), observables (O), probabilistic outcome production (P) and Axiom (M $_{\otimes}^A$) assigning post-measurement states.

Both quantum theory and pQT are examples of AMTs. Quantum theory is characterised by the Lüders rule,

$$\mathbf{w}_{AB}^L(\Pi_x, \rho_{AB}) = \frac{(\Pi_x \otimes \mathbb{I}) \rho_{AB} (\Pi_x \otimes \mathbb{I})}{\text{Tr}(\Pi_x \otimes \mathbb{I} \rho_{AB})}, \quad (3.8)$$

whereas passive quantum theory, i.e. quantum theory with non-collapsing measurements [3] (or SERDs [120]) replacing standard measurements, is characterised by the passive update rule,

$$\mathbf{w}_{AB}^P(\Pi_x, \rho_{AB}) = \rho_{AB}. \quad (3.9)$$

We note that the other types of readout devices explored in [118–120] do not, at least in an obvious way, provide alternative update rules as defined in Def. 9. Update rules, in the context of Def. 9, refer to assignments of post-measurement states for measurements of *quantum observables*, which are represented by self-adjoint operators. Devices such as SRDs or VNEMs fall outside our framework as they yield the *local state* of a system or its *von Neumann entropy*, respectively [118] (although it may be conceivable to simulate such devices within valid AMTs, e.g. pQT). Nonetheless, implementing (possibly local-state-dependent, thus nonlinear) evolution rules after the action of a SERD—mentioned explicitly in [118, 120]—indeed provides different update rules, hence AMTs. Our framework, however, is not confined to these examples: it also encompasses update rules capable of modifying (as in standard quantum theory) the state of arbitrarily distant entangled systems, something that cannot be achieved by applying unitaries based on the observed local state of a system. In particular, anticipating the discussion of Sec. 3.3.2, the AMTs obtained by assuming post-processing of SERDs constitute *subtheories* of pQT.

By combining Axioms (\mathbf{M}_{\otimes}^A) and (\mathbf{P}) we can construct *instrument maps* $\omega_{AB}^A : \mathcal{P}(\mathcal{H}_A) \times \bar{\mathcal{S}}(\mathcal{H}_A \otimes \mathcal{H}_B) \rightarrow \bar{\mathcal{S}}(\mathcal{H}_A \otimes \mathcal{H}_B)$, which allow us to jointly describe both the obtained outcome and the post-measurement state. In the following sections, we will often work with ω^A and sub-normalised states rather than directly with update rules \mathbf{w}^A and normalised states. Let us define

$$\omega_{AB}^A(\Pi_x, \rho_{AB}) = \text{Tr}[\Pi_x \text{Tr}_B(\rho_{AB})] \mathbf{w}_{AB}^A\left(\Pi_x, \frac{\rho_{AB}}{\text{Tr}(\rho_{AB})}\right), \quad (3.10)$$

then $\omega_{AB}^A(\Pi_x, \rho_{AB})$ denotes the *un-normalised* post-measurement state of ‘AB’ conditioned on observing outcome x on system ‘A’. As with standard instrument maps, the trace of $\omega_{AB}^A(\Pi_x, \rho_{AB})$ returns the probability of observing x ,

$$\text{Tr}[\omega_{AB}^A(\Pi_x, \rho_{AB})] = \text{Tr}(\Pi_x \text{Tr}_B(\rho_{AB})) = \text{Tr}(\Pi_x \otimes \mathbb{I} \rho_{AB}). \quad (3.11)$$

We can concatenate instrument maps meaningfully. Given any initial joint state ρ_{AB} (not necessarily normalised), the operator $\omega_{AB}^A(\Pi_y, \omega_{AB}^A(\Pi_x, \rho_{AB}))$ describes the joint state at the end of a time-ordered sequence of two measurements with

outcomes Π_x and Π_y , respectively. Note that the measurements do not have to be performed on the same subsystem. This property is ensured by the requirement that ω_{AB}^A be *1-homogeneous* over $\bar{\mathcal{S}}(\mathcal{H}_A \otimes \mathcal{H}_B)$: for all $\Pi_x \in \mathcal{P}(\mathcal{H}_A)$, $\rho_{AB} \in \bar{\mathcal{S}}(\mathcal{H}_A \otimes \mathcal{H}_B)$ and $\lambda \in [0, 1]$,

$$\omega_{AB}^A(\Pi_x, \lambda \rho_{AB}) = \lambda \omega_{AB}^A(\Pi_x, \rho_{AB}) . \quad (3.12)$$

For proper mixtures, the conditional un-normalised post-measurement state is obtained by applying ω_{AB}^A to each element of the sum,

$$\rho_{AB} = \sum_{i=1}^n p_i \rho_{AB}^i \xrightarrow{x_A} \sum_{i=1}^n p_i \omega_{AB}^A(\Pi_x, \rho_{AB}^i) . \quad (3.13)$$

3.1.2 Assumptions

To recap, the AMT framework rests on four assumptions:

- (R1) *Completeness*: Every system must be assigned a post-measurement state as a result of any measurement that may be performed on some of its components.
- (R2) *Context-independence*: The initial state of the system and the observed outcome uniquely determine the post-measurement state.
- (R3) *Self-consistency*: The assignment of post-measurement states does not depend on whether or not the probed system is regarded as part of a larger composite system.
- (R4) (*Quantum*) *No-signalling*: Measurements on separated subsystems do not provide a signalling mechanism.

Assumptions (R1) and (R3) ensure that the behaviour of a system is coherent under measurements. In contrast, (R2) and (R4) encode operational principles that will restrict the set of mathematically-consistent measurement behaviours in a non-trivial way.

In Sec. 3.2, we will continue our discussion on the definition of update rules. In particular, we will examine how different AMTs can share the same measurement behaviour for single systems and we will explore the limitations of “complete positivity” in constructing the AMT framework. Then, in Sec. 3.2.4 and 3.2.5, we will analyse more closely assumptions (R2) and (R4) by providing examples of measurement behaviours that are ruled out exclusively by either context-independence or no-signalling.

3.2 The role of update rules

3.2.1 Multiple extensions to composite systems

As discussed in Sec. 1.2.4, Axiom (M^L) fails to describe how a measurement on one component affects entangled systems in quantum theory. Axiom (M_{\otimes}^L) was therefore introduced to describe any measurement scenario in the theory. However, Axiom (M_{\otimes}^L) is not implied uniquely from Axiom (M^L) .

More generally, different update rules may be compatible with the same behaviour observed for single systems. Therefore, operational conditions on the behaviour of single systems are usually insufficient to isolate a specific update rule w^A from the rest.

For example, it is easy to verify that

$$w_{AB}^A(\Pi_x, \rho_{AB}) = w_A^A(\Pi_x, \text{Tr}_B(\rho_{AB})) \otimes \text{Tr}_A(\rho_{AB}) \quad (3.14)$$

defines a valid update rule compatible with any single-system rule w_A^A obeying (R1) completeness and (R2) context-independence. In fact, since the reduced state of subsystem ‘B’ is not affected, the map in Eq. (3.14) satisfies (R3) self-consistency and (R4) no-signalling.

Substituting w_A^L —cf. Eq. (1.6)—into Eq. (3.14) leads to an alternative update rule, \tilde{w}^L , still compatible with the standard collapse, cf. Eq. (1.10). Therefore, there are at least two AMTs featuring the Lüders projection for single systems: standard quantum theory, identified by w^L , and *correlation-free* quantum theory, identified by \tilde{w}^L . The predictions of the two models agree for single-system measurements, but differ when entangled systems are measured independently. In particular, the behaviour described by Eq. (3.14) is equivalent to cloning the reduced state $\text{Tr}_B(\rho_{AB})$ and performing a measurement on the copy while the original is discarded². As in pQT, correlations between distant entangled systems cannot be observed but, in contrast to pQT, they cannot even be retrieved by subsequent multi-partite measurements.

²Note that the ‘correlation-free’ theory cannot be obtained by instead changing the state space to exclude entanglement. Such a modification would, in fact, also affect the behaviour of measurements performed on the entire composite system, which one does not obtain by replacing the standard update rule with \tilde{w}^L .

3.2.2 The limitations of complete positivity

A map \mathcal{N} defined on the space of bounded operators $\mathcal{L}(\mathcal{H}_A)$ is said to be *completely positive* if, for any finite-dimensional space \mathcal{H}_B with $m = \dim(\mathcal{H}_B)$, and all positive operators $\sigma \in \mathcal{L}(\mathcal{H}_A \otimes \mathcal{H}_B)$, where $\sigma = \sum_{ijkl} c_{ijkl} |i\rangle\langle j| \otimes |k\rangle\langle l| = \sum_{kl} \sigma_{kl} \otimes |k\rangle\langle l|$, we have

$$\sigma' = \sum_{k,l=1}^m \mathcal{N}(\sigma_{kl}) \otimes |k\rangle\langle l| \geq O. \quad (3.15)$$

This definition, taken from [6], was introduced to apply to linear as well as nonlinear maps. If \mathcal{N} is linear, then $\sigma' \equiv (\mathcal{N} \otimes \mathcal{I})(\sigma)$.

Quantum mechanics is an example of an AMT in which instrument maps associated with a measurement on subsystem ‘A’ with outcome Π_x obey the relation

$$\omega_{AB}^A(\Pi_x, \rho_{AB}) = \left(\omega_x^A \otimes \mathcal{I}_B \right) (\rho_{AB}), \quad (3.16)$$

where $\omega_x^A(\cdot) \equiv \omega_A^A(\Pi_x, \cdot)$ is the single-partite map acting on $\bar{\mathcal{S}}(\mathcal{H}_A)$ conditioned on Π_x . Assuming ω_x^A has a unique linear extension to $\mathcal{L}(\mathcal{H}_A)$, the *complete positivity* of ω_x^A is a necessary and sufficient condition for Eq. (3.16) to define a valid update rule w^A via the relation given in Eq. (3.10). In fact, if ω_x^A is completely positive, then it is a quantum operation, hence ω_{AB}^A is both (R3) self-consistent and (R4) no-signalling—if this was *not* the case, quantum theory would violate at least one of these two requirements. Conversely, if ω_x^A is not completely positive, then ω_{AB}^A does not map onto $\bar{\mathcal{S}}(\mathcal{H}_A \otimes \mathcal{H}_B)$, hence Eq. (3.16) does not always represent a valid state of the composite system. Therefore, complete positivity represents a way to consistently “extend” single-system *linear* maps ω_A^A to composite systems. In other words, it provides a way to define (M_{\otimes}^A) from (M^A) in the case of single-system rules w_A^A that give rise to linear instruments.

We could have defined arbitrary update rules constructively using complete positivity of the single-system mappings ω_A^A , rather than in terms of the conditions (R1)-(R4) of Def. 9. Clearly, a more restricted framework of theories would have resulted. This is an immediate consequence of Eq. (3.14) which provides an update rule consistent with Def. 9 regardless of whether ω_A^A is completely positive or not—a feature that distinguishes our approach from others, such as [162], where complete positivity of transformations is assumed from the outset. Nevertheless, can Eq. (3.15) serve as an effective means to extend *any* single system behaviour into a fully fledged update rule?

It turns out that this is not possible, as complete positivity defined by Eq. (3.15) is “physically unfit” for nonlinear dynamics. The authors of [47] observed that a nonlinear version of \mathcal{N} generally maps different decompositions of the same state ρ_{AB} to different operators. In the AMT framework, this translates into an ambiguous—hence unphysical—assignment of joint post-measurement states. Let us illustrate this feature using passive measurements. Firstly, notice that complete positivity, as defined in [6], requires the map to be defined on the space of bounded operators, whereas the single-system rule ω_A^P and its corresponding p-instrument map ω_A^P are defined on $\mathcal{S}(\mathcal{H}_A)$ and $\bar{\mathcal{S}}(\mathcal{H}_A)$, respectively. Therefore, one needs to modify their definitions so that they also apply to non-Hermitian operators, i.e. $\omega_A^P(\Pi_x, L) = \text{Tr}(\Pi_x L) L$, where $L \in \mathcal{L}(\mathcal{H}_A)$. Now suppose we have a qubit pair in the state $|\Psi_{AB}\rangle = c_0|00\rangle + c_1|11\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$, and we perform a computational basis measurement on subsystem ‘A’ with outcome $\Pi_0 = |0\rangle\langle 0|$. By choosing two different bases, say $\{|0\rangle, |1\rangle\}$ and $\{|+\rangle, |-\rangle\}$, for \mathcal{H}_B , we can express the operator $\rho_{AB} = |\Psi_{AB}\rangle\langle\Psi_{AB}|$ in two different ways:

$$\rho_{AB} = \sum_{k,l \in \{0,1\}} \sigma_{kl} \otimes c_k c_l^* |k\rangle\langle l| = \sum_{\tilde{k}, \tilde{l} \in \{+, -\}} \tilde{\sigma}_{\tilde{k}\tilde{l}} \otimes |\tilde{k}\rangle\langle\tilde{l}|/2 \quad (3.17)$$

where we have set $\sigma_{kl} = |k\rangle\langle l|$, $\tilde{\sigma}_{\tilde{k}\tilde{l}} = |\phi_{\tilde{k}}\rangle\langle\phi_{\tilde{l}}|$ and $|\phi_{\pm}\rangle = c_0|0\rangle \pm c_1|1\rangle$. Applying the transformation defined in Eq. (3.15) to these decompositions leads to two different outputs:

$$\sum_{k,l \in \{0,1\}} \omega_A^P(\Pi_0, \sigma_{kl}) \otimes c_k c_l^* |k\rangle\langle l| = |c_0|^2 |00\rangle\langle 00|, \quad (3.18)$$

$$\sum_{\tilde{k}, \tilde{l} \in \{+, -\}} \omega_A^P(\Pi_0, \tilde{\sigma}_{\tilde{k}\tilde{l}}) \otimes |\tilde{k}\rangle\langle\tilde{l}|/2 = |c_0|^2 \rho_{AB}. \quad (3.19)$$

Both are valid states of $\mathcal{H}_A \otimes \mathcal{H}_B$, hence the argument against complete positivity applies regardless of whether Eq. (3.15) holds in all cases. In particular, the output given in Eq. (3.18) is inconsistent with the action of ω_A^P on the reduced state of qubit ‘A’, $\omega_A^P(\Pi_0, \rho_A) = |c_0|^2 \rho_A$. Using a third basis for \mathcal{H}_B would likely result in yet another final state.

This suggests that complete positivity is not a sensible criterion to construct instrument maps for composite systems given arbitrary instrument maps for single systems. The requirement does not entail for nonlinear maps the same desirable operational properties that it does for linear maps.

3.2.3 Composition compatibility

In quantum theory, a measurement with outcome Π_x performed by a device \mathcal{D}_A on subsystem \mathcal{H}_A affects the joint state in the same way as a measurement with outcome $\Pi_x \otimes \mathbb{I}_B$ performed by a global device \mathcal{D}_{AB} on the entire system $\mathcal{H}_A \otimes \mathcal{H}_B$. Formally, this property reads

$$\mathbf{w}_{AB}^L(\Pi_x, \rho_{AB}) = \mathbf{w}_{AB}^L(\Pi_x \otimes \mathbb{I}_B, \rho_{AB}) , \quad (3.20)$$

where on the right-hand side ‘AB’ is regarded as a single system, since $\Pi_x \otimes \mathbb{I} \in \mathcal{P}(\mathcal{H}_A \otimes \mathcal{H}_B)$. In other words, the two expressions in Eq. (3.20) are different: that on the left-hand side is defined on $\mathcal{P}(\mathcal{H}_A) \times \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$, while the other on $\mathcal{P}(\mathcal{H}_A \otimes \mathcal{H}_B) \times \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$. Let us define this operational feature for arbitrary AMTs.

Definition 11 (*Composition Compatibility (CC)*). An update rule \mathbf{w}^A satisfies *composition compatibility (CC)* if, for every $\mathcal{H}_A, \mathcal{H}_B$, joint state $\rho_{AB} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$ and local outcome $\Pi_x \in \mathcal{P}(\mathcal{H}_A)$,

$$\mathbf{w}_{AB}^A(\Pi_x, \rho_{AB}) = \mathbf{w}_{AB}^A(\Pi_x \otimes \mathbb{I}_B, \rho_{AB}) . \quad (3.21)$$

Again, although Eq. (3.21) specifies measurements on ‘A’ only, the property is assumed to hold for measurements on any component.

In analogy with complete positivity, composition compatibility can be used to *define* an update rule \mathbf{w}^A given only the behaviour for single systems, \mathbf{w}_A^A . However, in contrast with complete positivity, (CC) can be also used for maps \mathbf{w}_A^A giving rise to nonlinear instruments. Passive quantum theory is one such example, in line with Eq. (2.31). Furthermore, composition compatibility singles out quantum theory within the set of AMTs consistent with the Lüders projection for single systems \mathbf{w}_A^L —cf. Sec. 3.2.1.

However, this strategy of constructing update rules does not work in general. For instance, consider the mapping

$$\mathbf{w}_A^{\text{mix}}(\Pi_x, \rho) = \frac{\mathbb{I}_A}{d_A} \quad (3.22)$$

where $d_A = \dim(\mathcal{H}_A)$, which completely depolarises any measured state. If the outcome represented by $|0\rangle\langle 0| \otimes \mathbb{I}_2$ is observed from a measurement by a global device

\mathcal{D}_{AB} on a 4-dimensional qudit (labelled ‘AB’) in the product state $|00\rangle$, then the post-measurement state would be

$$\mathbf{w}_{AB}^{\text{mix}}(|0\rangle\langle 0| \otimes \mathbb{I}_2, |00\rangle\langle 00|) = \frac{\mathbb{I}_4}{4}. \quad (3.23)$$

If we imposed composition compatibility, then Eq. (3.23) would also describe the joint state of two (possibly distant) qubits after a measurement on \mathcal{H}_A performed with a local device \mathcal{D}_A . Consequently, the reduced state of qubit ‘B’ would transition from a pure state to the maximally mixed state. Clearly, such transformation would enable signalling between the parties, thus violating (R4).

3.2.4 Context-dependence: von Neumann’s postulate

The assumption (R2) of context-independence ensures that all identical experimental runs (i.e. the same state ρ is measured and the same outcome Π_x is observed) lead to the same post-measurement state. Conceptually, it resembles the standard idea of non-contextuality in quantum foundations [166] which stipulates that the *context*—i.e. any information distinguishing two operationally equivalent processes—does not result in any difference at the ontic level. Similarly, here we choose to assume that the post-measurement state must not depend on any information beyond the observed outcome and the input state. Specifically, the state assignment cannot depend on “hidden” variables with probability distributions which vary from one experimental run to the next.

Many possible measurement behaviours are dismissed by this constraint. For example, consider a map acting non-trivially whenever Π_x is obtained from a measurement of the observable M but behaving passively when it is obtained from a measurement of a different observable N . The output would then depend on the experimenter’s choice of observable, hence the map does not define a valid update rule³. Our assumption does not exclude the possibility that \mathbf{w}^A may act trivially whenever the outcome Π_x is observed and non-trivially whenever some other outcome $\Pi_y \neq \Pi_x$ is registered. Single-system maps of the form

$$\mathbf{w}_A^A(\Pi_x, \rho, \lambda) = \begin{cases} \rho_0 & \text{if } \lambda = 0 \\ \rho_1 & \text{if } \lambda = 1 \end{cases}, \quad (3.24)$$

³Of course, one might choose to define update rules differently to include such possibility.

however, are ruled out by context-independence. In this case, the parameter λ might represent some environmental variable, noise internal to the device, or a switch that flips between different update rules.

However, if instead we assumed that the probability that the hidden variable λ takes on a specific value is the same across all experimental runs, i.e. $p(\lambda = 0) = q$, then it is possible to define an update rule that accounts for such probabilistic behavior. For single systems, this would be

$$\mathbf{w}_A^{\Lambda}(\Pi_x, \rho) = q \rho_0 + (1 - q) \rho_1. \quad (3.25)$$

In contrast with Eq. (3.24), here the post-measurement state is a proper mixture and identical in all runs. One might still consider the “real” post-measurement state to be either ρ_0 or ρ_1 based on the value of some hidden λ , however the distribution of λ is independent of the “context” of the particular experiment. Specifically, Eq. (3.25) might depict a scenario where, with probability q (equal for all measurements), the state is projected to the subspace \mathcal{H}_x (i.e. a standard quantum measurement device is implemented), while with probability $1 - p$ it does not change (i.e. a SERD or passive measurement device is implemented). In other words, the assumption of context-independence does not rule out *all* possible versions of quantum theory augmented by readout devices [120].

The projection postulate introduced by von Neumann is a notable example for which a similar argument can be made. As discussed in Sec. 1.2.2, according to von Neumann [177], outcome degeneracy of quantum measurements results from classical post-processing of non-degenerate measurements. That is, one does not directly implement a measurement of a degenerate observable M , but a measurement of a non-degenerate *refinement* M' commuting with M is carried out, and the outcomes are coarse-grained. Given a degenerate observable, there exist, however, infinitely many refinements. For example, both qutrit observables $\{|0\rangle\langle 0|, |1\rangle\langle 1|, |2\rangle\langle 2|\}$ and $\{|+\rangle\langle +|, |-\rangle\langle -|, |2\rangle\langle 2|\}$ are valid refinements of the degenerate observable $\{\mathbb{I} - |2\rangle\langle 2|, |2\rangle\langle 2|\}$. As a result, the post-measurement state will not only depend on ρ and Π_x but also on the chosen orthonormal basis $\mathcal{B}_x = \{|x_i\rangle\}_{i=1}^{\text{Tr}(\Pi_x)}$ of $\mathcal{H}_x \equiv \Pi_x \mathcal{H}$,

$$\mathbf{w}_A^{\text{vN}}(\Pi_x, \rho, \mathcal{B}_x) = \frac{1}{p(x)} \sum_i p(x_i) \mathbf{w}_A^{\downarrow}(|x_i\rangle\langle x_i|, \rho), \quad (3.26)$$

where $p(x) = \text{Tr}(\Pi_x \rho)$ and $p(x_i) = \text{Tr}(|x_i\rangle\langle x_i| \rho)$. Therefore, von Neumann’s postulate—which, as Lüders postulate, only focuses on the state-updates of the

measured system—does not give rise to a valid update rule, unless the probability of using any given basis \mathcal{B}_x is fixed for each experiment resulting in the outcome Π_x . Notice that the corresponding instrument map ω_A^{vN} is linear and completely positive, hence Eq. (3.16) allows to extend the rule to composite systems in a way that satisfies no-signalling. We can thus conclude that “context-independence” (R2) is the only assumption responsible for ruling out von Neumann’s update rule.

3.2.5 Signalling: partial repeatability

An example of measurement behavior that is ruled out by the quantum no-signalling condition was already provided in Eqs. (3.22)-(3.23). Assuming that measurement disturbances propagate instantaneously, such toy models would enable superluminal signalling between parties, thus also violating the relativistic no-signalling principle. In other words, the possibility to construct toy models via our approach that violate (R4) (cf. (3.4)) indicates that the Hilbert space structure of quantum theory on its own does not guarantee compatibility with special relativity. In this section, we investigate other toy models that violate the quantum no-signalling condition.

Let us examine the following measurement behaviours that can be considered to “lie between” the standard collapse of quantum theory and the passive rule of pQT. Given $\lambda \in [0, 1]$, we introduce a one-parameter family of mappings w_A^λ for a single system ‘A’. Let $\Pi_x \in \mathcal{P}(\mathcal{H}_A)$ and $\rho \in \mathcal{S}(\mathcal{H}_A)$, define

$$w_A^\lambda(\Pi_x, \rho) = \frac{G_\lambda^{1/2}(\Pi_x) \rho G_\lambda^{1/2}(\Pi_x)}{\text{Tr}(G_\lambda(\Pi_x) \rho)} \quad (3.27)$$

where

$$G_\lambda(\Pi_x) = (1 - \lambda) \Pi_x + \lambda \mathbb{I} \quad (3.28)$$

and

$$G_\lambda^{1/2}(\Pi_x) = \left(1 - \sqrt{\lambda}\right) \Pi_x + \sqrt{\lambda} \mathbb{I} \quad (3.29)$$

is its unique positive square root. The Lüders rule w_A^L and the passive rule w_A^P for single systems are recovered by setting $\lambda = 0$ and $\lambda = 1$, respectively. In other words, the family $\{w_A^\lambda\}$ smoothly interpolates between quantum theory and pQT.

The induced instrument maps ω_A^λ ,

$$\omega_A^\lambda(\Pi_x, \rho) = \text{Tr}(\Pi_x \rho) \frac{G_\lambda^{1/2}(\Pi_x) \rho G_\lambda^{1/2}(\Pi_x)}{\text{Tr}(G_\lambda(\Pi_x) \rho)}, \quad (3.30)$$

are nonlinear in the state ρ for $\lambda \neq 0$. Therefore, as pQT, these new models do not preserve the indistinguishability of different preparations of the same mixed state.

Consider the vector state $|\psi\rangle = \sum_{i=0}^{d-1} c_i |i\rangle \in \mathcal{H}_A$ and a measurement outcome $\Pi_X = \sum_{i \in X} |i\rangle\langle i|$ for $X \subseteq \{0, \dots, d-1\}$. The output state is then

$$\mathbf{w}_A^\lambda(\Pi_X, |\psi\rangle\langle\psi|) = |\psi_X\rangle\langle\psi_X|, \quad (3.31)$$

where

$$|\psi_X\rangle = \frac{1}{N} \left(\sum_{j \in X} c_j |j\rangle + \sqrt{\lambda} \sum_{i \notin X} c_i |i\rangle \right), \quad (3.32)$$

with N being the normalisation factor. Therefore, the parameter λ , fixed for all measurements in the corresponding toy model, determines how likely a second measurement, performed immediately after the first one, will yield the same outcome as the first. The closer λ is to 0, the higher the boost in moduli of the amplitudes corresponding to the “observed” components of the superposition. We call this phenomenon *partial repeatability*.

Definition 12 (*Partial repeatability (PA)*). For any outcome x of a measurement repeated on the same system without intervening unitary evolution, the probability of observing x for the *second* time is higher than the probability to observe it for the *first* time,

$$p(o_2 = x | o_1 = x) > p(o_1 = x), \quad (3.33)$$

where $o_1(o_2)$ is the variable corresponding to the outcome of the first (second) measurement.

The transformations of Eq. (3.27) satisfy partial repeatability, except when $\lambda = 1$. In fact, partial repeatability does not include the *probabilistic* repeatability (PR) known from pQT, where $p(o_2 = x | o_1 = x) = p(o_1 = x)$, cf. Sec. 2.3.1. Furthermore, Eq. (3.32) shows that the disturbance to the state does not affect the relative phases, and that the ratio between post- and pre-measurement amplitudes is the same even for “unobserved” components, i.e. where $i \notin X$.

In order to extend the update rule to composite systems, we assume composition compatibility (CC), following the example of quantum theory and pQT, see Sec. 3.2.3. That is, subsystem measurements induce joint state updates according to

$$\mathbf{w}_{AB}^\lambda(\Pi_x, \rho_{AB}) = \mathbf{w}_{AB}^\lambda(\Pi_x \otimes \mathbb{I}_B, \rho_{AB}). \quad (3.34)$$

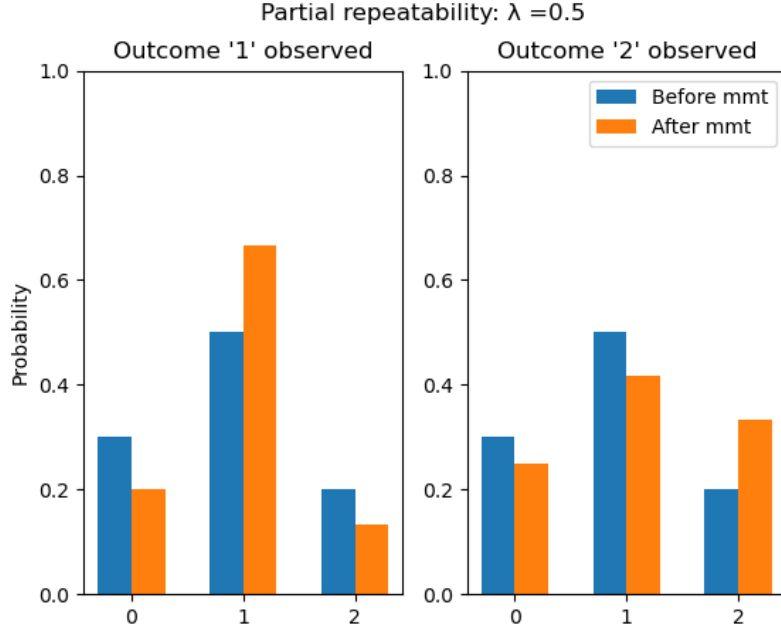


Figure 3.1: Example of partial repeatability resulting from the single-system transformation w_A^λ . Taking $\lambda = 0.5$ and the initial qutrit state $|\psi\rangle = \sqrt{0.3}|0\rangle + \sqrt{0.5}|1\rangle + \sqrt{0.2}|2\rangle$, the plots show the outcome probabilities before and immediately after a measurement in the computational basis, where outcome ‘1’ (left) or ‘2’ (right) is observed.

To show that Eq. (3.34) is self-consistent (R3), notice that

$$G_\lambda^{1/2}(\Pi_x \otimes \mathbb{I}_B) = G_\lambda^{1/2}(\Pi_x) \otimes \mathbb{I}_B, \quad (3.35)$$

hence, as expected,

$$\text{Tr}_{B''} \left[w_{AB'B''}^\lambda(\Pi_x, \rho_{AB'B''}) \right] = \text{Tr}_{B''} \left[\frac{G_\lambda^{1/2}(\Pi_x) \otimes \mathbb{I}_{B'B''} \rho_{AB'B''} G_\lambda^{1/2}(\Pi_x) \otimes \mathbb{I}_{B'B''}}{\text{Tr}(G_\lambda(\Pi_x \otimes \mathbb{I}_{B'B''}) \rho_{AB'B''})} \right] \quad (3.36)$$

$$= \frac{G_\lambda^{1/2}(\Pi_x) \otimes \mathbb{I}_{B'} \text{Tr}_{B''}(\rho_{AB'B''}) G_\lambda^{1/2}(\Pi_x) \otimes \mathbb{I}_{B'}}{\text{Tr}(G_\lambda(\Pi_x \otimes \mathbb{I}_{B'}) \text{Tr}_{B''}(\rho_{AB'B''}))} \quad (3.37)$$

$$= w_{AB'}^\lambda(\Pi_x, \text{Tr}_{B''}(\rho_{AB'B''})). \quad (3.38)$$

However, Eq. (3.34) does *not* define a valid update rule for $\lambda \notin \{0, 1\}$ because it violates quantum no-signalling (R4). Let Alice and Bob share a pair of qubits in the entangled state

$$|\psi\rangle = \alpha|00\rangle + e^{i\theta}\beta|11\rangle, \quad \alpha, \beta \in \mathbb{R}, \quad (3.39)$$

and consider a computational basis measurement performed by Alice. Then

$$p(0_A) \mathbf{w}_{AB}^\lambda(\Pi_0, |\psi\rangle\langle\psi|) + p(1_A) \mathbf{w}_{AB}^\lambda(\Pi_1, |\psi\rangle\langle\psi|) = \alpha^2 |\psi_{0_A}\rangle\langle\psi_{0_A}| + \beta^2 |\psi_{1_A}\rangle\langle\psi_{1_A}|, \quad (3.40)$$

where $\beta^2 = 1 - \alpha^2$. The normalised post-measurement states are

$$|\psi_{0_A}\rangle = \frac{\alpha|00\rangle + e^{i\theta}\sqrt{\lambda}\beta|11\rangle}{\sqrt{\alpha^2 + \lambda\beta^2}}, \quad |\psi_{1_A}\rangle = \frac{\sqrt{\lambda}\alpha|00\rangle + e^{i\theta}\beta|11\rangle}{\sqrt{\lambda\alpha^2 + \beta^2}}. \quad (3.41)$$

Alice can signal to Bob since

$$\begin{aligned} p(0_A) \text{Tr}_A [\mathbf{w}_{AB}^\lambda(\Pi_0, |\psi\rangle\langle\psi|)] + p(1_A) \text{Tr}_A [\mathbf{w}_{AB}^\lambda(\Pi_1, |\psi\rangle\langle\psi|)] &= \\ &= \left(\frac{\alpha^4}{\alpha^2 + \lambda\beta^2} + \frac{\lambda\alpha^2\beta^2}{\lambda\alpha^2 + \beta^2} \right) |0\rangle\langle 0| + \left(\frac{\lambda\alpha^2\beta^2}{\alpha^2 + \lambda\beta^2} + \frac{\beta^4}{\lambda\alpha^2 + \beta^2} \right) |1\rangle\langle 1| \\ &\neq \alpha^2|0\rangle\langle 0| + \beta^2|1\rangle\langle 1| = \text{Tr}_A(|\psi\rangle\langle\psi|) \end{aligned} \quad (3.42)$$

for all $0 < \lambda < 1$ and provided $\alpha^2 \notin \{0, 1/2, 1\}$. The probability that Bob will observe the outcome “0” depends on whether or not Alice performs her measurement:

$$p(0_B|\text{no A mmt}) = \alpha^2 \quad (3.43)$$

$$p(0_B|\text{A mmt}) = \frac{\alpha^4}{\alpha^2 + \lambda\beta^2} + \frac{\lambda\alpha^2\beta^2}{\lambda\alpha^2 + \beta^2} \quad (3.44)$$

The two probabilities match when qubits are not entangled or, surprisingly, when they are maximally entangled. Consequently, by sharing an ensemble of (not maximally) entangled systems, Alice and Bob can violate Bell’s inequalities up to the algebraic maximum.

Despite the failure of \mathbf{w}^λ of Eq. (3.34) to define an update rule, it is indeed possible to construct valid AMTs featuring partial measurements. Eq. (3.14), for instance, provides a valid extension $\tilde{\mathbf{w}}_{AB}^\lambda$ of \mathbf{w}_A^λ to composite systems. Alternatively, we can interpolate between \mathbf{w}^L and \mathbf{w}^P in the following way. Let $\mu \in (0, 1)$ and $\Pi_x \in \mathcal{P}(\mathcal{H}_A)$, and consider the one-parameter family of update rules defined by

$$\mathbf{w}_{AB}^\mu(\Pi_x, \rho_{AB}) = (1 - \mu) \mathbf{w}_{AB}^L(\Pi_x, \rho_{AB}) + \mu \mathbf{w}_{AB}^P(\Pi_x, \rho_{AB}). \quad (3.45)$$

The transformations of Eq. (3.45) satisfy partial repeatability. Unlike the previous family of transformations, here the higher probability of re-observing x comes from the non-zero chance—equal for all measurements—of the state collapsing via the Lüders rule. Eq. (3.45) does not lead to signalling since it is the convex combination of no-signalling update rules.

3.3 Generalised instruments and AMT simulation

3.3.1 Generalised instruments and observables

In any AMT, measuring an observable $M = \sum_{x \in X} m_x \Pi_x^M$ defined on a system leads to a post-measurement state as prescribed by the update rule. However, there exist alternative strategies to measure the same observable. For instance, a device that implements a measurement of $N = U M U^\dagger$ can be used, provided the system first undergoes the unitary evolution $\rho \mapsto U \rho U^\dagger$. Though a different experimental set-up is required, the outcome statistics will agree with the Born rule for M , $p(x) = \text{Tr}(\Pi_x^M \rho)$. Another strategy is to couple the system with an ancilla, which is then suitably measured. Furthermore, adding a channel at the end of any of these strategies will define yet a different scheme for measuring M . This holds true for all AMTs because channels can be defined in terms of axioms which do not involve the update rule.

Although every measurement is *ultimately* modeled by the update rule of the theory, different schemes for measuring M will generally result in different post-measurement states of the original system. We account for this flexibility by defining *generalised M -compatible instruments* for any model theory in the framework. This will also allow us to include *unsharp* measurements in our formalism—Sec. 1.2.3.

It follows from Naimark’s dilation theorem [144, 171] that we can interpret a *measurement model* $\mathcal{M} = \langle \mathcal{H}_E, \xi, U, N \rangle$, where \mathcal{H}_E denotes an ancillary system, as describing one way to implement a measurement of $M = \{M_x\}_{x \in X}$ on \mathcal{H}_S . For this to hold, letting $N = \{\Pi_x^N\}_{x \in X}$, the probability reproducibility condition—cf. Eq. (1.20)—must be satisfied,

$$\text{Tr}\left[\left(\mathbb{I} \otimes \Pi_x^N\right) U \rho \otimes \xi U^\dagger\right] = \text{Tr}\left(M_x^\dagger M_x \rho\right). \quad (3.46)$$

In the language of AMTs, given a theory characterised by the update rule \mathbf{w}^A —with associated fundamental instrument map ω^A given by Eq. (3.10)—, the strategy described by \mathcal{M} induces the following transformation of system \mathcal{H}_S ,

$$\omega_{\mathcal{M}}(M_x, \rho) = \text{Tr}_E \left[\omega_{SE}^A \left(\Pi_x^N, U \rho \otimes \xi U^\dagger \right) \right]. \quad (3.47)$$

Therefore, with Eq. (3.47) we can describe the impact of unsharp measurements, despite the fact that the update rule \mathbf{w}^A and the fundamental instrument map ω^A are only defined for projectors.

However, Eq. (3.47) only covers a subset of possible strategies for measuring M . Instead of operating on the ancilla, a sharp measurement could in principle be carried out on \mathcal{H}_S directly or on the whole composite system $\mathcal{H}_S \otimes \mathcal{H}_E$. We want to be able to account for these options in order to include, for example, the ancilla-free strategy mentioned at the beginning of the section, where the unitary gate in \mathcal{M} would have the form $U \otimes \mathbb{I}_E$. In addition, we also would like to include cases of AMTs that are *not* “compositionally compatible” (CC), i.e. where measuring $\{\Pi_x\}$ on \mathcal{H}_E is not equivalent to measuring $\{\mathbb{I}_S \otimes \Pi_x\}$ on $\mathcal{H}_S \otimes \mathcal{H}_E$. To achieve this level of generality, we will work with a generalised notion of measurement model $\mathcal{M} = \langle \mathcal{H}_E, \xi, U, N \rangle$, in which N can refer to a (sharp) observable defined on *any* subspace $\mathcal{H}_{\tilde{E}} \subseteq \mathcal{H}_S \otimes \mathcal{H}_E$. The reproducibility condition (3.46) can then be written in the following way,

$$\mathrm{Tr} \left[\omega_{SE}^A \left(\Pi_x^N, U \rho \otimes \xi U^\dagger \right) \right] = \mathrm{Tr} \left(M_x^\dagger M_x \rho \right), \quad (3.48)$$

where $\Pi_x^N \in \mathcal{P}(\mathcal{H}_{\tilde{E}})$, while Eq. (3.47) continues to hold.

We now define, for any theory in the framework, the M -compatible generalised instruments describing all possible single-apparatus strategies implementing a measurement of M and their effect on the system. In essence, these strategies follow the simple paradigm of first executing an M -compatible measurement model \mathcal{M} and then applying an outcome-dependent channel to the composite system. In particular, the channels are assumed to be implemented dynamically, i.e. via unitary evolution involving additional ancillary systems. The alternative way to implement channels via unconditional measurements (see discussion in Sec. 2.5.1) is ruled out by the restriction to a single apparatus.

Definition 13. The collection of maps $\{\omega_{M_x}\}_x$ is an M -compatible generalised instrument realisable in the AMT with update rule \mathbf{w}^A if

$$\omega_{M_x}(\rho) \equiv \omega(M_x, \rho) = \mathrm{Tr}_E \left[\eta_{SE}^x \circ \omega_{SE}^A \left(\Pi_x^N, U \rho \otimes \xi U^\dagger \right) \right] \quad (3.49)$$

where $\mathcal{M} = \langle \mathcal{H}_E, \xi, U, N = \{\Pi_x^N \in \mathcal{P}(\mathcal{H}_{\tilde{E}})\} \rangle$ is an M -compatible measurement model and $\{\eta_{SE}^x\}_x$ is a collection of channels on $\mathcal{S}(\mathcal{H} \otimes \mathcal{H}_E)$.

For a given AMT, the *fundamental instrument* $\{\omega_x^A\}_x$ for a sharp observable M is recovered by setting $\eta_{SE}^x = \mathcal{I}_{SE}$, $U = \mathbb{I}_{SE}$ and letting $N = M$ be an observable of \mathcal{H}_S . In the special case of quantum theory, the generalised instruments compatible with M coincide with its quantum instruments defined in Def. 2.

A theorem by Ozawa [142]—already mentioned in Sec. 1.3.2—shows that for every \mathbf{M} -compatible *quantum* instrument $\{\omega_{M_x}\}_x$, there exists a measurement model \mathcal{M} such that, for all states ρ , $\omega_{M_x}(\rho) = \omega_{\mathcal{M}}(M_x, \rho)$. In other words, all quantum instruments can be realised without physically implementing channels η_{SE}^x after the measurement on the ancilla. This property, however, does not hold for arbitrary AMTs. In pQT, for instance, as discussed in Sec. 2.5.4, there exists p-instruments that cannot be realised without applying some channels after the measurement. An example of such “indirect” p-instrument is provided by the map $\omega(M_x, \rho) = \text{Tr}(M_x \rho) \rho / \text{Tr}(\rho)$ for some non-trivial unsharp \mathbf{M} (i.e. the M_x are not all multiples of the identity). Despite being identical to the fundamental p-instrument ω_A^P (which however is only defined for sharp observables), the transformation requires that a unitary gate is applied after the ancilla measurement. In fact, any \mathbf{M} -compatible measurement model \mathcal{M} generates entanglement between the system and the ancilla, which is not broken by a passive measurement of the observable N . Hence, to bring the system back to its initial state ρ , the de-coupling unitary U^\dagger must be applied to the composite $\mathcal{H}_S \otimes \mathcal{H}_E$.

We thus identify two classes of generalised instruments. On one hand, *direct* instruments describe ways of measuring \mathbf{M} that do not require the implementation of channels η_{SE}^x . On the other hand, *indirect* instruments describe those strategies where channels, hence the time delay needed for their implementation, are unavoidable.

Definition 14. The \mathbf{M} -compatible generalised instrument $\{\omega_{M_x}\}_x$, realisable in the AMT with update rule \mathbf{w}^A , is said to be *direct* if there exists an \mathbf{M} -compatible measurement model \mathcal{M} such that $\omega_{M_x}(\rho) = \omega_{\mathcal{M}}(M_x, \rho)$ for all $\rho \in \bar{\mathcal{S}}(\mathcal{H}_S)$. Otherwise, it is said to be *indirect*.

3.3.2 Subtheories

In Sec. 2.6, we discussed the possibility to simulate quantum theory with passive measurements. We now generalise the idea to arbitrary AMTs with different update rules. Specifically, if $\text{AMT}(\beta)$ can reproduce the update rule of $\text{AMT}(\alpha)$, we say that $\text{AMT}(\alpha)$ is a *subtheory* of $\text{AMT}(\beta)$.

Definition 15. $\text{AMT}(\alpha)$, with update rule \mathbf{w}^α , is a *subtheory* of $\text{AMT}(\beta)$, with update rule \mathbf{w}^β , if, for any \mathcal{H}_A , \mathcal{H}_B and any sharp measurement $M = \{\Pi_x\}_x$ on \mathcal{H}_A , there exists an M -compatible measurement model $\mathcal{M}_M = \langle \mathcal{H}_C, \xi, U_{CA}, N \rangle$ — with

$N = \{\Pi_x^N\}_x$ a sharp observable of $\mathcal{H}_{\tilde{E}} \subseteq \mathcal{H}_C \otimes \mathcal{H}_A$ — and a collection of channels $\{\eta_{CA}^x\}_x$ on $\mathcal{S}(\mathcal{H}_C \otimes \mathcal{H}_A)$ such that

$$\omega_{AB}^\alpha(\Pi_x, \rho_{AB}) = \text{Tr}_C \left[(\eta_{CA}^x \otimes \mathbb{I}_B) \circ \omega_{CAB}^\beta \left(\Pi_x^N, (U_{CA} \otimes \mathbb{I}_B) (\xi \otimes \rho_{AB}) (U_{CA}^\dagger \otimes \mathbb{I}_B) \right) \right] \quad (3.50)$$

for all joint states $\rho_{AB} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$.

A diagrammatic illustration of the condition of Eq. (3.50), which is articulated in terms of fundamental instrument maps rather than the corresponding update rules, is depicted in Fig. 3.2. In short, we require that the impact on the joint state of any sharp subsystem measurement predicted by $\text{AMT}(\alpha)$ can be simulated in $\text{AMT}(\beta)$ by a procedure that involves a measurement of the *same* subsystem, and possibly an ancilla. In particular, we impose the restriction that the simulation protocol cannot involve subsystem \mathcal{H}_B , which is not directly measured in the “target” procedure of $\text{AMT}(\alpha)$. This guarantees the exclusion of simulation strategies that require the experimenter to have access to both systems \mathcal{H}_A and \mathcal{H}_B , even in cases where these systems are space-like separated.




Figure 3.2: Circuit diagrammatic representation of Eq. (3.50). The operator $U_M = U_{CA}$ denotes the unitary included in the M -compatible measurement model \mathcal{M}_M , whereas U_x , which acts on \mathcal{H}_{CA} and an additional ancilla \mathcal{H}_D (initialised in some state ξ_D), realises the outcome-dependent channel η_{CA}^x —i.e. $\langle \mathcal{H}_D, U_x, \xi_D \rangle$ is a *dilation* of η_{CA}^x . We use dashed lines around the meter to indicate that the measurement may be performed only on part of the system(s) of interest, here $\mathcal{H}_{\tilde{E}} \subseteq \mathcal{H}_C \otimes \mathcal{H}_A$.

A direct consequence of Eq. (3.50) is that, for any generalised observable M , all M -compatible generalised instruments realisable in $\text{AMT}(\alpha)$ are also realisable in $\text{AMT}(\beta)$. This is illustrated schematically in Fig. 3.3. Therefore, if (3.50) holds,

then $\text{AMT}(\beta)$ can successfully simulate *any* measurement process that appears in $\text{AMT}(\alpha)$. The simulability relation defined by the notion of a subtheory does not define a partial order within the set of AMTs. In fact, in Sec. 3.4.1, we will explore a subset of different AMTs that are *mutually simulable*.

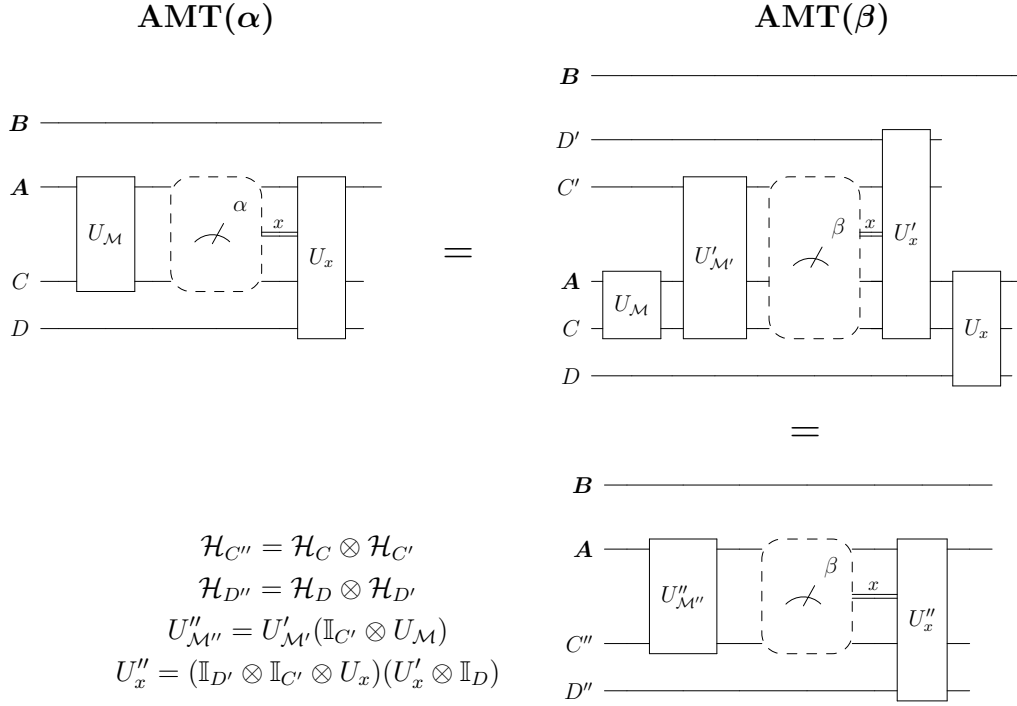


Figure 3.3: Pictorial proof that Eq. (3.50) implies that $\text{AMT}(\beta)$ can simulate all instruments realisable in $\text{AMT}(\alpha)$. For any generalised observable \mathcal{M} of \mathcal{H}_A , any \mathcal{M} -compatible generalised instrument in $\text{AMT}(\alpha)$ can be realised by a circuit like the one on the left, cf. (3.49). The horizontal equivalence is ensured by Eq. (3.50), while the vertical equivalence follows by suitable relabeling of systems and operators.

When simulation was addressed in Sec. 2.6 in the specific case of pQT, it was argued that all quantum measurements can be simulated, at least in principle, with passive measurements. However, the protocol described required both an infinite sequence of passive measurements and access to entangled subsystems. According to the more restrictive definition in terms of Eq. (3.50), it is clear that standard quantum theory is not a subtheory of passive quantum theory, and vice versa. In fact, no passive measurement can induce the collapse of a distant entangled system, and at the same time no non-trivial quantum measurement leaves the joint state unaltered.

3.4 Properties of AMTs

Having introduced the AMT framework and illustrated it by providing examples of acceptable and excluded measurement behaviors, our attention now turns to exploring the properties of the theories allowed within the framework. The objective of the upcoming sections is to characterise physically reasonable update rules which respect additional operationally-inspired properties. Ultimately, we will search for features that distinguish standard quantum theory from all other AMTs. Table 3.1 in Sec. 3.5 summarises the main update rules discussed in this chapter along with some of their features.

3.4.1 Preparation indistinguishability

In all AMTs, a *single* measurement cannot distinguish between different preparations of the same density matrix. By performing one measurement on each element of an infinite ensemble described by the mixed state ρ , an experimenter cannot decide whether the ensemble was prepared according to Gemenge \mathcal{G}_1 or \mathcal{G}_2 , where $\rho_{\mathcal{G}_1} = \rho_{\mathcal{G}_2} = \rho$. This follows from the linearity of the Born rule over $\mathcal{S}(\mathcal{H})$.

Suppose $\mathcal{G}_1 = \{(1/2, |0\rangle), (1/2, |1\rangle)\}$ and $\mathcal{G}_2 = \{(1/2, |+\rangle), (1/2, |-\rangle)\}$, hence $\rho_{\mathcal{G}_1} = \rho_{\mathcal{G}_2} = \mathbb{I}/2 \in \mathcal{S}(\mathcal{H}_2)$. In quantum theory, the two preparations are indistinguishable even if the experimenter is allowed to perform a *sequence* of measurements on each element of the ensemble. However, this equivalence does not hold in all AMTs. In pQT, for example, two measurements in the computational basis are sufficient to discriminate between the two scenarios. If successive measurements on the same system always yield identical outcomes, i.e. $p(o_1 = x_1, o_2 = x_2) = \delta_{x_1 x_2}/2$, then the ensemble was prepared according to \mathcal{G}_1 ; otherwise, \mathcal{G}_2 was employed. The distinguishability, as discussed in Sec. 2.5.2, is a consequence of the nonlinearity of the fundamental p-instrument map ω^P over $\mathcal{S}(\mathcal{H})$.

We now formalise the idea of “preparation indistinguishability” for sequential measurements. In analogy with the extension (M_{\otimes}^A) of Axiom (M^A) , we will start by considering non-composite systems only (Def. 16) and then provide a stronger definition that accounts for subsystem measurements (Def. 17).

Definition 16 (*Weak Indistinguishability (WI)*). Physically different preparations of one and the same state of a composite system cannot be distinguished by a sequence of measurements performed on it as a whole.

Given an arbitrary AMT, this property is guaranteed by the convex-linearity over $\mathcal{S}(\mathcal{H}_A)$ of the single-system instrument map ω_A^A induced by the defining update rule w^A . That is, (WI) holds if and only if

$$\omega_A^A(\Pi_x, \lambda\rho + (1 - \lambda)\sigma) = \lambda\omega_A^A(\Pi_x, \rho) + (1 - \lambda)\omega_A^A(\Pi_x, \sigma) \quad (3.51)$$

for all $\Pi_x \in \mathcal{P}(\mathcal{H}_A)$, $\rho, \sigma \in \mathcal{S}(\mathcal{H}_A)$ and $\lambda \in [0, 1]$. Convex-linearity over density matrices implies the existence of a unique linear extension over the larger space of bounded operators $\mathcal{L}(\mathcal{H}_A)$ [97]. For simplicity, we will speak of “linearity” rather than “convex-linearity” and denote the extension to $\mathcal{L}(\mathcal{H}_A)$ with the same symbol as the map originally defined on $\mathcal{S}(\mathcal{H}_A)$, namely ω_A^A .

As shown in Sec. 3.2.1, it is possible to devise update rules such that ω_A^A is linear over $\mathcal{S}(\mathcal{H}_A)$ but ω_{AB}^A is nonlinear over the larger, composite system $\mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$. In this case, even though the AMT satisfies (WI), different preparations of the same *joint* state may be distinguished by exploiting the composite nature of the system. As an example, we consider the update rule \tilde{w}^L defined by Eq. (3.14), that is compatible with the Lüders projection for single systems w_A^L . The corresponding instrument map for composite systems $\tilde{\omega}_{AB}^L$ is nonlinear over density matrices. To see this, suppose that the two-qubit state ρ_{AB} has been prepared either via $\mathcal{G}_1 = \{(1/2, |\Psi^+\rangle), (1/2, |\Psi^-\rangle)\}$ or via $\mathcal{G}_2 = \{(1/2, |00\rangle), (1/2, |11\rangle)\}$ ⁴. In such an AMT, measurements in the computational basis of each qubit can reveal which Gemenge was employed. In fact, if the outcomes obtained from the two mono-partite measurements are *not* identical, then the qubits must have been prepared according to \mathcal{G}_1 , see Sec. 3.2.1. Therefore, measurement outcomes from *different* constituents may be used to gain knowledge about the preparation of the composite system. This possibility neither implies a violation of quantum no-signalling nor can it be achieved with measurements performed on only one constituent.

In order to exclude such a scenario, we define the following stronger notion of preparation indistinguishability.

Definition 17 (*Strong Indistinguishability* (SI)). Physically different preparations of one and the same state of a composite system cannot be distinguished by a sequence of measurements performed on it as a whole or on its constituents.

Since a linear map ω_{AB}^A cannot induce a nonlinear map ω_A^A , the condition is equivalent to the convex-linearity over density matrices of the instrument map for

⁴Here, $|\Psi^\pm\rangle = (|00\rangle \pm |11\rangle)/\sqrt{2}$.

composite systems induced by the defining update rule w^A . That is, (SI) holds if and only if

$$\begin{aligned} \omega_{AB}^A(\Pi_x, \lambda\rho_{AB} + (1-\lambda)\sigma_{AB}) &= \\ &= \lambda\omega_{AB}^A(\Pi_x, \rho_{AB}) + (1-\lambda)\omega_{AB}^A(\Pi_x, \sigma_{AB}) \end{aligned} \quad (3.52)$$

for all \mathcal{H}_B , $\Pi_x \in \mathcal{P}(\mathcal{H}_A)$, $\rho_{AB}, \sigma_{AB} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$ and $\lambda \in [0, 1]$. In other words, the map ω_{AB}^A is linear over $\mathcal{L}(\mathcal{H}_A \otimes \mathcal{H}_B)$.

If the map ω_{AB}^A is linear, then all generalised instruments—cf. Eq. (3.49)—of the corresponding theory will be linear. The converse is also true, hence linearity of all generalised instruments is equivalent to “strong indistinguishability”.

Lemma 3.1. *An update rule w^A is (SI) if and only if, for any generalised observable M on \mathcal{H}_A , all M -compatible generalised instruments realisable in the corresponding AMT are composed of linear maps over $\bar{\mathcal{S}}(\mathcal{H}_A)$.*

Proof. See Appendix A.2. □

We will now show that assuming (SI) is equivalent to fixing a particular extension of instrument maps to composite systems. The proof of this result (cf. Theorem 3.1 below) relies on the following lemma.

Lemma 3.2. *Let $M, N : \mathcal{L}(\mathcal{H}_A \otimes \mathcal{H}_B) \rightarrow \mathcal{L}(\mathcal{H}_A \otimes \mathcal{H}_B)$ be linear mappings satisfying $M(\rho_A \otimes \rho_B) = N(\rho_A \otimes \rho_B)$ for all $\rho_A \in \mathcal{S}(\mathcal{H}_A)$ and $\rho_B \in \mathcal{S}(\mathcal{H}_B)$. Then $M = N$.*

Proof. See Appendix A.2. □

Theorem 3.1. *An update rule w^A is (SI) if and only if the corresponding instrument maps—defined by Eq. (3.10)—have the form*

$$\omega_{AB}^A(\Pi_x, \rho_{AB}) = \left(\omega_x^A \otimes \mathcal{I}_B \right) (\rho_{AB}) . \quad (3.53)$$

where $\omega_x^A(\cdot) \equiv \omega_A^A(\Pi_x, \cdot)$.

Proof. In order for Eq. (3.53) to (i) be well-defined and (ii) assign post-measurement states unambiguously, ω_A^A must be (i) linear over $\mathcal{L}(\mathcal{H}_A)$ and (ii) completely positive. But if ω_A^A is linear over $\mathcal{L}(\mathcal{H}_A)$, then the map ω_{AB}^A is linear over $\mathcal{L}(\mathcal{H}_A \otimes \mathcal{H}_B)$. Hence, an update rule w^A for which Eq. (3.53) holds is necessarily (SI).

To show the converse, namely that (SI) implies Eq. (3.53), consider an arbitrary pure product state $\rho_\psi \otimes \rho_\phi = |\psi\rangle\langle\psi| \otimes |\phi\rangle\langle\phi| \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$. Given any update rule w^A , the requirement of no-signalling (R4) implies that

$$\mathrm{Tr}_A \left[\omega_{AB}^A (\Pi_x, \rho_\psi \otimes \rho_\phi) \right] = \rho_\phi \quad (3.54)$$

whereas from self-consistency (R3) we have that

$$\mathrm{Tr}_B \left[\omega_{AB}^A (\Pi_x, \rho_\psi \otimes \rho_\phi) \right] = \omega_x^A (\rho_\psi) . \quad (3.55)$$

The post-measurement state of the joint system must therefore be given by

$$\omega_{AB}^A (\Pi_x, \rho_\psi \otimes \rho_\phi) = \omega_x^A (\rho_\psi) \otimes \rho_\phi . \quad (3.56)$$

Consider now an arbitrary product state $\rho_A \otimes \rho_B = \sum_{ij} p_i q_j \rho_A^i \otimes \rho_B^j$, where $0 \leq p_i, q_j \leq 1$, $\sum_i p_i = \sum_j q_j = 1$ and ρ_A^i, ρ_B^j are pure states. If w^A is (SI), it follows from the linearity of ω_{AB}^A over joint states and Eq. (3.56) that

$$\omega_{AB}^A (\Pi_x, \rho_A \otimes \rho_B) = \sum_{ij} p_i q_j \omega_{AB}^A (\Pi_x, \rho_A^i \otimes \rho_B^j) \quad (3.57)$$

$$= \sum_{ij} p_i q_j \omega_x^A (\rho_A^i) \otimes \rho_B^j \quad (3.58)$$

$$= \sum_{ij} p_i q_j \omega_x^A \otimes \mathcal{I} (\rho_A^i \otimes \rho_B^j) \quad (3.59)$$

$$= \omega_x^A \otimes \mathcal{I} (\rho_A \otimes \rho_B) , \quad (3.60)$$

where we used the definition $\mathcal{N} \otimes \mathcal{I} (L_A \otimes L_B) = \mathcal{N} (L_A) \otimes L_B$ for any $L_A \in \mathcal{L}(\mathcal{H}_A)$ and $L_B \in \mathcal{L}(\mathcal{H}_B)$. Hence, for any fixed outcome $\Pi_x \in \mathcal{P}(\mathcal{H}_A)$, the two linear mappings $M(\cdot) = \omega_{AB}^A (\Pi_x, \cdot)$ and $N(\cdot) = \omega_x^A \otimes \mathcal{I}(\cdot)$ defined on $\mathcal{L}(\mathcal{H}_A \otimes \mathcal{H}_B)$ agree on all product operators $\rho_A \otimes \rho_B$. Lemma 3.2 applies: the two maps must be equal. Therefore,

$$\omega_{AB}^A (\Pi_x, \rho_{AB}) = \left(\omega_x^A \otimes \mathcal{I}_B \right) (\rho_{AB}) \quad (3.61)$$

holds for any $\Pi_x \in \mathcal{P}(\mathcal{H}_A)$ and joint state $\rho_{AB} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$. \square

We have shown that assuming ‘‘strong indistinguishability’’ implies that (i) the single-system instrument map ω_A^A is linear and completely positive and (ii) the composite-system instrument map ω_{AB}^A is given by the well-defined relation (3.53). In other words, the assumption singles out quantum instruments. Since the Luders instrument is not the only quantum instrument compatible with an

arbitrary sharp observable M , it follows that the set of (SI) AMTs includes standard quantum theory, though not exclusively. Adding an outcome-dependent channel to the Lüders rule, in fact, defines a different (SI) update rule. Nevertheless, from the theorems by Ozawa [142] and Hayashi [94, 97], mentioned in Sec. 1.3.2, it follows that quantum theory can simulate any (SI) AMT.

Corollary 3.1. *Every (SI) AMT is a subtheory of standard quantum theory.*

Mono- and multi-partite procedures

Another consequence of assuming strong indistinguishability is the equivalence between mono- and multi-partite procedures for measuring observables of the form $A \otimes B$, which we encountered in Sec. 2.3.2.

Definition 18 (*Equivalence of mono- and multi-partite procedures (MM)*). The outcome probabilities $\{p(x, y)\}_{xy}$ from a measurement of $\{\Pi_x \otimes \Pi_y\}_{xy}$ using a global device \mathcal{D}_{AB} on $\mathcal{H}_A \otimes \mathcal{H}_B$ coincide with the joint probabilities from measurements of $\{\Pi_x\}_x$ on \mathcal{H}_A and $\{\Pi_y\}_y$ on \mathcal{H}_B performed by *local* devices \mathcal{D}_A and \mathcal{D}_B , respectively.

Note that classical communication between the parties is required to obtain the joint probabilities when performing mono-partite measurements. Given an arbitrary AMT with update rule w^A , the requirement for (MM) can be expressed as

$$\mathrm{Tr}(\Pi_x \otimes \Pi_y \rho_{AB}) = \mathrm{Tr}\left[(\mathbb{I} \otimes \Pi_y) \omega_{AB}^A(\Pi_x, \rho_{AB})\right], \quad (3.62)$$

holding for all ρ_{AB} and local outcomes $\Pi_x \in \mathcal{P}(\mathcal{H}_A)$ and $\Pi_y \in \mathcal{P}(\mathcal{H}_B)$. For the reverse order of the local measurements, we have

$$\mathrm{Tr}(\Pi_x \otimes \Pi_y \rho_{AB}) = \mathrm{Tr}\left[(\Pi_x \otimes \mathbb{I}) \omega_{AB}^A(\Pi_y, \rho_{AB})\right]. \quad (3.63)$$

Notice that the identities (3.62) and (3.63) only require agreement of the outcome probabilities and not necessarily of the output state of the composite system.

Lemma 3.3. *If the update rule w^A is (SI), then the corresponding AMT satisfies (MM).*

Proof. If the update rule is (SI), then the instrument map for single systems ω_A^A is linear and completely positive. Therefore, it has an operator-sum (Kraus) representation $\omega_A^A(\Pi_x, L) = \sum_r K_r L K_r^\dagger$, where $L \in \mathcal{L}(\mathcal{H}_A)$ and $\sum_r K_r^\dagger K_r = \Pi_x$ [97]. Then

we can write

$$\mathrm{Tr}\left[(\mathbb{I} \otimes \Pi_y) \left(\omega_A^{\mathbb{A}}(\Pi_x) \otimes \mathcal{I}\right) (\rho_{AB})\right] = \mathrm{Tr}\left[\sum_r (\mathbb{I} \otimes \Pi_y) (K_r \otimes \mathbb{I}) \rho_{AB} (K_r^\dagger \otimes \mathbb{I})\right] \quad (3.64)$$

$$= \mathrm{Tr}\left[\sum_r (K_r^\dagger K_r \otimes \mathbb{I}) (\mathbb{I} \otimes \Pi_y) \rho_{AB}\right] \quad (3.65)$$

$$= \mathrm{Tr}(\Pi_x \otimes \Pi_y \rho_{AB}) \quad (3.66)$$

since the operators $K_r \otimes \mathbb{I}$ and $\mathbb{I} \otimes \Pi_y$ commute for all values of r and y . \square

Closure under sequential composition

Another way to ensure linearity of all instrument maps is by assuming the following principle, which, to the best of our knowledge, was first introduced in [76].

Definition 19 (*Closure under sequential composition (CS)*). For any sequence of generalised measurements performed on a system, there exists a *single* generalised measurement that is equivalent to it.

In other words, given an arbitrary sequence of measurements, it is possible to identify a single (generalised) measurement such that the outcome probabilities obtained from performing it on each element of an infinite ensemble coincide with the joint probabilities observed from performing the sequence of measurements on each element. The equivalence must hold for any state ρ .

The equivalence between (CS) and (SI) can be understood as follows: if we require sequences of measurements to have the same mathematical representation as single measurements, then preparations that cannot be distinguished based on the outcome statistics of single measurements also cannot be distinguished based on the statistics of many measurements.

In the language of our framework, an AMT is “sequentially closed” (CS) if, for any pair of generalised observables $\mathbf{F} = \{F_i\}$ and $\mathbf{G} = \{G_j\}$ of \mathcal{H}_A and any \mathbf{F} -compatible instrument $\{\omega_{F_i}\}$ realisable in the AMT, there exists a generalised observable $\mathbf{H} = \{H_{ij}\}$ such that

$$\mathrm{Tr}(H_{ij}^\dagger H_{ij} \rho) = \mathrm{Tr}[G_j^\dagger G_j \omega_{F_i}(\rho)], \quad (3.67)$$

for any state $\rho \in \bar{\mathcal{S}}(\mathcal{H}_A)$. The generalisation to longer sequences follows by induction.

Lemma 3.4. *An update rule w^A is (CS) if and only if it is (SI).*

Proof. Suppose w^A is (CS). Let $\rho = p\rho_1 + q\rho_2$ with $\rho_1 \neq \rho_2$ and non-negative numbers p, q with $p + q = 1$. By linearity of the trace, we write

$$\mathrm{Tr}[H_{ij}^\dagger H_{ij} (p\rho_1 + q\rho_2)] = p \mathrm{Tr}(H_{ij}^\dagger H_{ij} \rho_1) + q \mathrm{Tr}(H_{ij}^\dagger H_{ij} \rho_2). \quad (3.68)$$

Using (3.67) leads to

$$\mathrm{Tr}[G_j^\dagger G_j \omega_{F_i} (p\rho_1 + q\rho_2)] = p \mathrm{Tr}[G_j^\dagger G_j \omega_{F_i} (\rho_1)] + q \mathrm{Tr}[G_j^\dagger G_j \omega_{F_i} (\rho_2)] \quad (3.69)$$

$$= \mathrm{Tr}[G_j^\dagger G_j (p\omega_{F_i} (\rho_1) + q\omega_{F_i} (\rho_2))] . \quad (3.70)$$

According to Assumption (CS), Eq. (3.69) must hold for any $\mathbf{G} = \{G_j\}$. But since two quantum states that assign the same probabilities to all outcomes of all measurements are necessarily equal, it follows that

$$\omega_{F_i} (p\rho_1 + q\rho_2) = p\omega_{F_i} (\rho_1) + q\omega_{F_i} (\rho_2) \quad (3.71)$$

Therefore, for any generalised measurement $\mathbf{F} = \{F_i\}$, all \mathbf{F} -compatible generalised instruments realisable in the AMT are composed of linear maps. According to Lemma 3.1, this implies that the update rule w^A satisfies (SI).

To show that (SI) implies (CS), let $\{\omega_{F_i}\}$ be any \mathbf{F} -compatible generalised instrument in the AMT. Using Lemma 3.1 again, we have that ω_{F_i} is linear for every i . We can thus define the dual map $\omega_{F_i}^* : \mathcal{L}(\mathcal{H}) \rightarrow \mathcal{L}(\mathcal{H})$ in the following way (see [97]),

$$\mathrm{Tr}[\omega_{F_i}^* (A) B] = \mathrm{Tr}[A \omega_{F_i} (B)] \quad \forall A, B \in \mathcal{L}(\mathcal{H}). \quad (3.72)$$

Since $\omega_{\mathbf{F}} = \sum_i \omega_{F_i}$ is trace-preserving, $\omega_{\mathbf{F}}^*$ is unital, $\omega_{\mathbf{F}}^* (\mathbb{I}) = \mathbb{I}$. Let $\mathbf{G}_j = G_j^\dagger G_j$ denote the POVM elements of $\mathbf{G} = \{G_j\}$, then $\mathbf{H}_{ij} = \omega_{F_i}^* (\mathbf{G}_j)$ defines a POVM,

$$\sum_{ij} \omega_{F_i}^* (\mathbf{G}_j) = \sum_j \omega_{\mathbf{F}}^* (\mathbf{G}_j) = \omega_{\mathbf{F}}^* (\underbrace{\sum_j \mathbf{G}_j}_{\mathbb{I}}) = \mathbb{I} \quad (3.73)$$

where the second step follows from the linearity of $\omega_{\mathbf{F}}^*$. Hence, for arbitrary \mathbf{F} and \mathbf{G} , there exists a generalised observable $\mathbf{H} = \{H_{ij}\}$ with POVM elements $\mathbf{H}_{ij} = H_{ij}^\dagger H_{ij}$ such that

$$\mathrm{Tr}[\mathbf{G}_j \omega_{F_i} (\rho)] = \mathrm{Tr}[\mathbf{H}_{ij} \rho] \quad (3.74)$$

for all states $\rho \in \bar{\mathcal{S}}(\mathcal{H})$. We conclude that w^A satisfies (CS). \square

Closure under sequential composition—cf. Def. 19—was introduced in 2017 by Flatt et al. in [76] to extend the Gleason-Busch theorem [34] to sequential measurements. The authors justify this principle by an assumption of non-contextuality at the level of the description of measurements.

In 2019, Masanes et al. adopt (CS) and argue that quantum instruments represent the only consistent measurement-induced transformations of states [132]. However, a recent objection by Stacey [168] maintained that linearity was assumed in their argument, rather than derived (the paper [132] was also critically reviewed in [118], where objections were raised using readout devices). In their response [81] to [168], the authors clarified how (CS) encapsulates their operational definition of a *measurement* as

“any experiment that takes a quantum system [...] as its input and generates one of several possible outcomes” ([81], p. 3).

With this definition, the process involving the implementation of a sequence of devices is considered a single measurement in its own right. Therefore, it must be represented by the same type of mathematical objects, which they assume to be POVMs.

In contrast to both [132] and [76], the AMT framework employs a different notion of measurement, based on the standard axiomatization of quantum theory. We regard a measurement as

any experiment that returns the value of one of the observable quantities of a system.

In particular, according to Axiom (O), the observable quantities are represented by self-adjoint operators. Rather than using sharp observables instead of POVMs, the key difference from the definition adopted in [132] is the specification of what measurement outcomes represent, namely values of observables, which have a precise mathematical representation. Now, a sequence of measurements might not, as a whole, return the value of one of the observables of the system, so the entire process does not generally constitute a measurement in its own right. Therefore, measurement sequences are conceptually distinct from the single measurements that compose them, and do not need to be described by the same mathematical objects. Although the lack of specification in the definition employed in [132] might suggest that it is weaker, it is really the AMT definition that provides greater generality. The AMT framework, in fact, includes model theories that are inconsistent with the

definition in [132], e.g. pQT where the sequential measurements of two observables cannot be described as a measurement of a single observable, as well as theories that are consistent with it, i.e. (SI) AMTs.

Mutual simulability

The principle of strong indistinguishability significantly limits the range of feasible measurement behaviours. Corollary 3.1 asserts that any AMT that conforms to this restriction can be simulated by quantum theory. Likewise, every (SI) AMT has the ability to replicate the standard collapse.

Theorem 3.2. *Standard quantum theory is a subtheory of every (SI) AMT.*

Proof. Consider the Lüders instrument $\{\omega_{\Pi_x}^L\}_x$ of a sharp observable $\{\Pi_x\}_x$ on \mathcal{H}_A . Ozawa's theorem (cf. Sec. 1.3.2) applies: there exists a measurement model $\mathcal{M} = \langle \mathcal{H}_C, \xi, U_{CA}, N = \{\Pi_x^N\}_x \rangle$ with $\Pi_x^N \in \mathcal{P}(\mathcal{H}_C)$ such that

$$\omega_{\Pi_x}^L(\rho_A) = \text{Tr}_C \left[\left(\Pi_x^N \otimes \mathbb{I}_A \right) U_{CA} (\xi \otimes \rho_A) U_{CA}^\dagger \right] \quad (3.75)$$

for all $\rho_A \in \mathcal{S}(\mathcal{H}_A)$. Recall that ξ is the initial state of the ancilla \mathcal{H}_C , U_{CA} is the coupling unitary and N is the sharp ancilla observable. We can use the same measurement model \mathcal{M} for a measurement on a subsystem. Let $\rho_{AB} \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$, then

$$\omega_{AB}^L(\Pi_x, \rho_{AB}) = \left(\omega_{\Pi_x}^L \otimes \mathcal{I}_B \right) (\rho_{AB}) \quad (3.76)$$

$$= \text{Tr}_C \left[\left(\Pi_x^N \otimes \mathbb{I}_A \otimes \mathbb{I}_B \right) \sigma_{CAB} \right] \quad (3.77)$$

where $\sigma_{CAB} = (U_{CA} \otimes \mathbb{I}_B) (\xi \otimes \rho_{AB}) (U_{CA}^\dagger \otimes \mathbb{I}_B)$. We can rewrite this equation in the following way,

$$\omega_{AB}^L(\Pi_x, \rho_{AB}) = \text{Tr}_C \left[\left(\Pi_x^N \otimes \mathbb{I}_A \otimes \mathbb{I}_B \right) \sigma_{CAB} \left(\Pi_x^N \otimes \mathbb{I}_A \otimes \mathbb{I}_B \right) \right] \quad (3.78)$$

$$= \text{Tr}_C \left[\left(\omega_{\Pi_x^N}^L \otimes \mathcal{I}_{AB} \right) (\sigma_{CAB}) \right]. \quad (3.79)$$

Since appending channels to the ancilla \mathcal{H}_C does not alter the reduced state of $\mathcal{H}_A \otimes \mathcal{H}_B$ (violation of this would amount to signalling), we have that

$$\omega_{AB}^L(\Pi_x, \rho_{AB}) = \text{Tr}_C \left\{ \left[\left(\eta_x \circ \omega_{\Pi_x^N}^L \right) \otimes \mathcal{I}_{AB} \right] (\sigma_{CAB}) \right\} \quad (3.80)$$

for any channel η_x on $\mathcal{S}(\mathcal{H}_C)$. Hayashi's theorem (cf. Sec. 1.3.2) asserts that, given any generalised observable $\mathbf{M} = \{M_x\}_x$, every \mathbf{M} -compatible quantum instrument

$\{\omega_{M_x}\}_x$ is composed of maps of the form $\omega_{M_x} = \eta_x \circ \omega_{M_x}^{\text{L}}$, where $\{\eta_x\}_x$ is some set of outcome-dependent channels. We know from Theorem 3.1 that the fundamental instrument of any (SI) AMT, $\{\omega_{\Pi_x}^{(\text{SI})}\}_x$, is a quantum instrument. Therefore, it follows from Eq. (3.80) that in any (SI) AMT we can write

$$\omega_{AB}^{\text{L}}(\Pi_x, \rho_{AB}) = \text{Tr}_C \left[\left(\omega_{\Pi_x}^{(\text{SI})} \otimes \mathcal{I}_{AB} \right) (\sigma_{CAB}) \right], \quad (3.81)$$

which holds for all joint states ρ_{AB} . We use Eq. (3.53) to rewrite the expression in the form of Eq. (3.50) appearing in the definition of a subtheory:

$$\omega_{AB}^{\text{L}}(\Pi_x, \rho_{AB}) = \text{Tr}_C \left[\omega_{CAB}^{(\text{SI})} \left(\Pi_x^N, (U_{CA} \otimes \mathbb{I}_B) (\xi \otimes \rho_{AB}) (U_{CA}^\dagger \otimes \mathbb{I}_B) \right) \right]. \quad (3.82)$$

Therefore, quantum theory is a subtheory of any (SI) AMT. \square

In particular, since no non-trivial channel follows the measurement on \mathcal{H}_C , Eq. (3.82) shows that the Lüders instrument is a *direct* generalised instrument in all models that obey (SI)—cf. Def. 14. Combining Corollary 3.1 and Theorem 3.2, we can conclude that all AMTs with update rules $\mathbf{w}^{(\text{SI})}$ which give rise to *linear* fundamental instruments are *mutually simulable*.

Corollary 3.2. *Every (SI) AMT is a subtheory of every other (SI) AMT.*

Corollary 3.2 suggests that the Lüders' rule \mathbf{w}^{L} is not the only update rule that is able to model the experimental results predicted by quantum mechanics. Any (SI) update rule can replicate the standard quantum collapse by suitably measuring an ancilla after having coupled it with the original system. For example, consider the (SI) AMT with linear fundamental instrument map $\omega_{AB}^{\text{mix}} = \omega_A^{\text{mix}} \otimes \mathcal{I}_B$, where $\omega_A^{\text{mix}}(\Pi_x, \rho) = \text{Tr}(\Pi_x \rho) \mathbb{I}_A / d_A$. In other words, the update rule for non-composite systems $\mathbf{w}_A^{\text{mix}}$ simply replaces the pre-measurement state with the maximally mixed state (see Def. 3.22). Clearly, $\mathbf{w}_A^{\text{mix}}$ and \mathbf{w}_A^{L} assign different post-measurement states to the system interacting with the device. However, despite the fundamental difference in measurement behaviour, in such AMT it is possible to induce the standard collapse of the state by following the measurement strategy outlined in Eq. 3.82, in which the device interacts with an ancilla. Repeating the same procedure with a new ancilla, immediately after the first measurement, will yield identical outcomes, in line with the deterministic repeatability of quantum mechanics—see Sec. 3.4.3. In general, since AMTs differ exclusively by their update rules, it follows that (SI) AMTs can capture all features of quantum theory, such as deterministic repeatability,

nonlocal correlations, measurement-based protocols and algorithms. Therefore, the choice to use the Lüders rule or any other (SI) rule is a matter of convention and convenience. Considering the fact that our formalism has the simple goal of modelling the operational effects of measurements, there is no *a priori* reason to prefer one (SI) update rule over another.

The results of this section also allow us to conclude that neither is pQT a subtheory of (SI) AMTs, nor do any of them feature as subtheories of pQT. This is due to the fact that passive measurements cannot disturb the states of distant entangled systems.

3.4.2 Ideality

In order to distinguish quantum theory from all other (SI) AMTs, an additional condition must be imposed. It was shown in [37, 50] that the Lüders instrument is the only *ideal* quantum instrument for sharp observables.

Definition 20 (*Ideality (ID)*). A measurement does not change the state of the system if the outcome is certain.

Just as for preparation indistinguishability, we can distinguish a weaker and a stronger notion of ideality. “Weak ideality” is defined for non-composite systems only. Given an update rule w_A^A , weak ideality amounts to the condition

$$\mathrm{Tr}(\Pi_x \rho) = 1 \implies w_A^A(\Pi_x, \rho) = \rho. \quad (3.83)$$

In contrast, “strong ideality” applies to the joint state of the composite system, thus preserving correlations with any entangled subsystem that was not measured. Mathematically, strong ideality is equivalent to the condition

$$\mathrm{Tr}(\Pi_x \otimes \mathbb{I} \rho_{AB}) = 1 \implies w_{AB}^A(\Pi_x, \rho_{AB}) = \rho_{AB}. \quad (3.84)$$

For (SI) theories, the two notions are equivalent, since the update rule on the joint system can be expressed in terms of the single-system rule, according to Eq. (3.53) (Theorem 3.1). However, the equivalence does not hold in general, as demonstrated by the AMT defined by \tilde{w}^L —cf. Sec. 3.2.1—which only satisfies weak ideality.

We now provide a different proof of a fundamental theorem shown in [37, 50].

Theorem 3.3. *If an update rule has the properties (SI) and (ID), it must be equal to the Lüders rule w^L .*

Proof. Let $\{|i\rangle, i = 0, \dots, d-1\}$ be some orthonormal basis of \mathcal{H}_d . Consider the outcome represented by the projector $\Pi_X = \sum_{i \in X} |i\rangle\langle i|$ where $X \subseteq \{0, \dots, d-1\}$. Assuming that the update rule w^A is (ID) implies that $\omega_X^A(|\psi\rangle\langle\psi|) = |\psi\rangle\langle\psi|$ for all $|\psi\rangle \in \mathcal{H}_X = \Pi_X \mathcal{H}_d$. Additionally, all update rules must satisfy the following: $\omega_X^A(|\phi\rangle\langle\phi|) = O$ for all $|\phi\rangle \in \mathcal{H}_X^\perp$, where $\mathcal{H}_X^\perp = (\mathbb{I}_d - \Pi_X)\mathcal{H}_d$. For example, $\omega_X^A(|j\rangle\langle j|) = O$ for all $j \notin X$. But since we assume (SI), the map ω_X^A is linear over $\mathcal{L}(\mathcal{H}_d)$ and completely positive, with Kraus representation $\omega_X^A(L) = \sum_r K_r L K_r^\dagger$, where $L \in \mathcal{L}(\mathcal{H}_d)$. It follows that $\omega_X^A(|j\rangle\langle j|) = \sum_r K_r |j\rangle\langle j| K_r^\dagger = O$, hence $K_r |j\rangle = 0$ for all r and all $j \notin X$. This implies that $K_r |\phi\rangle = 0$ for all r and all $|\phi\rangle = \sum_{j \notin X} c_j |j\rangle \in \mathcal{H}_X^\perp$.

Let $|\xi\rangle$ be an arbitrary state of \mathcal{H}_d , which can be decomposed as follows,

$$|\xi\rangle = \underbrace{\Pi_X |\xi\rangle}_{=|\psi\rangle} + \underbrace{(\mathbb{I}_d - \Pi_X) |\xi\rangle}_{=|\phi\rangle} = |\psi\rangle + |\phi\rangle. \quad (3.85)$$

Then, due to the linearity of ω_X^A , we can write

$$\omega_X^A(|\xi\rangle\langle\xi|) = \omega_X^A(|\psi\rangle\langle\psi|) + \omega_X^A(|\psi\rangle\langle\phi|) + \omega_X^A(|\phi\rangle\langle\psi|) + \omega_X^A(|\phi\rangle\langle\phi|) \quad (3.86)$$

$$= |\psi\rangle\langle\psi| + \sum_r K_r |\psi\rangle\langle\phi| K_r^\dagger + \sum_r K_r |\phi\rangle\langle\psi| K_r^\dagger + \sum_r K_r |\phi\rangle\langle\phi| K_r^\dagger \quad (3.87)$$

$$= |\psi\rangle\langle\psi|. \quad (3.88)$$

Hence, $\omega_X^A(|\xi\rangle\langle\xi|) = \Pi_X |\xi\rangle\langle\xi| \Pi_X = \omega_X^L(|\xi\rangle\langle\xi|)$ for any $|\xi\rangle \in \mathcal{H}_d$. The extension to composite systems ω_{AB}^L is fixed by Theorem 3.1. Having isolated the Lüders instrument map ω^L , the corresponding Lüders rule w^L follows from Eq. (3.10). \square

A schematic summary of the assumptions involved in the derivation of w^L within the AMT framework is shown in Fig. 3.4.

Notice that ideality alone is not sufficient to imply the standard quantum update rule w^L , as there exist ideal “nonlinear” AMTs such as pQT. Other cases can be envisaged: let w_A^A be any valid update rule for single systems (not necessarily ideal), then

$$w_A^{(\text{ID})}(\Pi_x, \rho) = \text{Tr}(\Pi_x \rho) w_A^L(\Pi_x, \rho) + [1 - \text{Tr}(\Pi_x \rho)] w_A^A(\Pi_x, \rho) \quad (3.89)$$

defines another single-system rule that complies with the weaker notion of ideality.

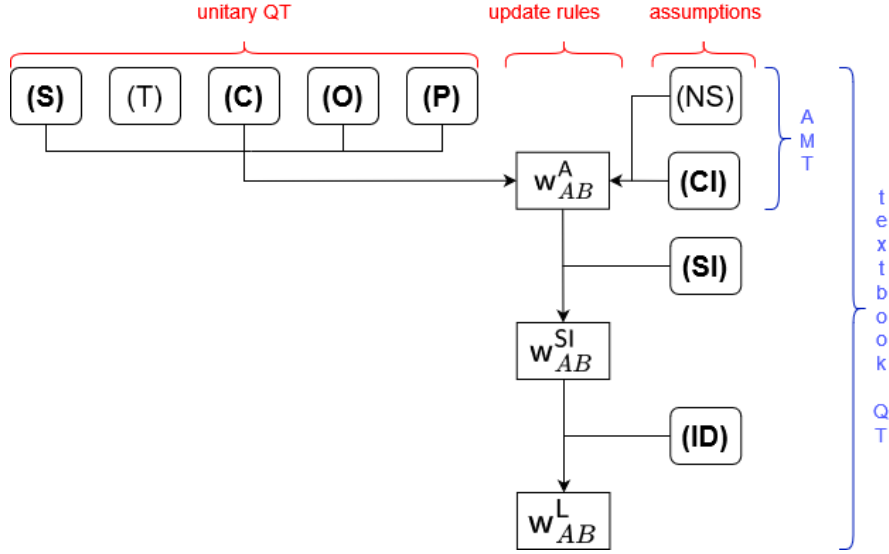


Figure 3.4: Schematic depiction of the arguments of Secs. 3.4.1 and 3.4.2 to derive the Lüders rule in the context of the AMT framework. Refer to Sec. 1.1 for the standard postulates of unitary quantum theory, and recall that (NS) = “no-signalling” ((R4) in Sec. 3.1.1), (CI) = “context-independence” ((R2) in Sec. 3.1.1), (SI) = “strong indistinguishability” and (ID) = “ideality”. The Lüders rule is uniquely implied by the assumptions in bold. Notably, while (NS) is required to define generic update rules, it does not feature in the derivation culminating with Theorem 3.3. In fact, quantum no-signalling is ensured by the requirement of (SI).

3.4.3 Deterministic repeatability

The original projection postulate was introduced by von Neumann to model the deterministic repeatability observed in the scattering experiment carried out by Compton and Simon [44]. We provide the definition once again.

Definition 4 (*Deterministic repeatability* (DR)). Consecutive measurements of the same observable, performed on the same system without intervening unitary evolution, yield identical outcomes.

In order to satisfy deterministic repeatability (DR), the post-measurement state of the probed system must have full support in the linear subspace $\mathcal{H}_x \equiv \Pi_x \mathcal{H}$ of \mathcal{H} for any outcome Π_x ,

$$\Pi_x w_A^A(\Pi_x, \rho) \Pi_x = w_A^A(\Pi_x, \rho) . \tag{3.90}$$

For non-degenerate outcomes, where $\Pi_x = |x\rangle\langle x|$, the only option is to have

$$\mathbf{w}_A^A(\Pi_x, \rho) = |x\rangle\langle x| = \mathbf{w}_A^L(\Pi_x, \rho) . \quad (3.91)$$

However, for degenerate outcomes, infinitely many quantum states have full support in \mathcal{H}_x . We present three examples of non-Lüders deterministically repeatable update rules for non-composite systems:

- (i) following the discussion of Sec. 3.2.4, von Neumann's projection \mathbf{w}_A^{vN} does not define an update rule because it violates context-independence (R3). However, by fixing an orthonormal basis $\mathcal{B}_x = \{|x_i\rangle\}_{i=1}^{\text{Tr}(\Pi_x)}$ for each subspace \mathcal{H}_x , we obtain a valid update rule $\tilde{\mathbf{w}}_A^{\text{vN}}$ which satisfies (DR),

$$\tilde{\mathbf{w}}_A^{\text{vN}}(\Pi_x, \rho) = \frac{1}{p(x)} \sum_i p(x_i) \mathbf{w}_A^L(|x_i\rangle\langle x_i|, \rho) , \quad (3.92)$$

where $p(x) = \text{Tr}(\Pi_x \rho)$ and $p(x_i) = \text{Tr}(|x_i\rangle\langle x_i| \rho)$.

- (ii) the map

$$\mathbf{w}_A^A(\Pi_x, \rho) = \frac{1}{\text{Tr}(\Pi_x)} \sum_{ij} |x_i\rangle\langle x_j| \quad (3.93)$$

reduces to Eq. (3.91) when Π_x is rank-1 but, in case of outcome degeneracy, collapses any initial state to an equal superposition of the elements of a preferred basis \mathcal{B}_x of \mathcal{H}_x .

- (iii) another option is to update ρ to the maximal mixed state \mathbb{I}_x/d_x of \mathcal{H}_x , which does not require the specification of any basis for the subspace.

Notably, all proposed examples define linear instruments for non-composite systems. Nevertheless, we can conclude that deterministic repeatability (DR), even when supported by strong indistinguishability (SI), is insufficient to isolate quantum theory within the broad framework of AMTs. In [100], the authors show that the Lüders rule for single systems can be derived via an information-theoretic approach by imposing, alongside (DR), that the post-measurement state minimises the relative information with respect to the pre-measurement state ρ . In other words, they prove that $\mathbf{w}_A^L(\Pi_x, \rho)$ is the least distinguishable state from ρ among those compatible with the repeatability of measurement outcomes. A similar result, framed in terms of trace distances, already appeared in [102].

3.4.4 Information-disturbance trade-off

As discussed in Sec. 1.2.2, Lüders presents two main objections to von Neumann's projection postulate [128]. Firstly, he argues that the post-measurement state should depend only on the outcome and the initial state. Secondly, he suggests that

“the measurement of a highly degenerate quantity permits only relatively weak assertions regarding the considered ensemble. For that reason, the resulting change in state should likewise be small” (Eng. trans. [127] of [128], p. 665).

In the limit of the trivial observable $\{\mathbb{I}, O\}$, a measurement reveals no information about the original state and it should have no effect on it. This idea of a trade-off between applied disturbance and acquired information is encoded in the following principle.

Definition 21 (*Information-disturbance trade-off (TO)*). The outcome probability distribution of an observable M is not affected by a prior measurement of a coarse-graining of M .

Consider an infinite ensemble prepared in an arbitrary quantum state, and examine the following two scenarios: (*i*) a measurement of M is performed on each element of the ensemble; (*ii*) a measurement of a coarse-graining of M followed by a measurement of M are performed on each element of the ensemble. Condition (TO) stipulates that the outcome probabilities $\{p(x)\}$ observed from the measurements of M in scenarios (*i*) and (*ii*) must be identical. In other words, for all quantum states and outcomes Π_x and Π_y such that $\Pi_x \Pi_y = \Pi_y \Pi_x = \Pi_x$, the relation $p(o_1 = y) p(o_2 = x | o_1 = y) = p(x)$ should hold (where o_1 and o_2 denote the first and second outcomes of two consecutive measurements, respectively). In terms of single-system update rules, the condition can be expressed as

$$\mathrm{Tr}(\Pi_y \rho) \mathrm{Tr}[\Pi_x \omega_A^\Lambda(\Pi_y, \rho)] = \mathrm{Tr}(\Pi_x \rho) , \quad (3.94)$$

or, in terms of the associated instrument maps, as

$$\mathrm{Tr}[\Pi_x \omega_A^\Lambda(\Pi_y, \rho)] = \mathrm{Tr}(\Pi_x \rho) . \quad (3.95)$$

Here, x denotes the outcome of the M measurement, while y denotes the outcome of the coarse-graining. If $\Pi_x \neq \Pi_y$, then $\mathcal{H}_x \subset \mathcal{H}_y$, meaning that y provides less

information about the original state than x . As a result, the disturbance applied to the state cannot be arbitrarily strong but must preserve coherence in the subspace \mathcal{H}_y , in the sense that $p(o_2 = x | o_1 = y) = \text{const.} \times p(x)$ for all $\Pi_x \in \mathcal{P}(\mathcal{H}_y)$. In the limiting case of $\Pi_y = \mathbb{I}$, one obtains $\mathbf{w}_A^\Lambda(\mathbb{I}, \rho) = \rho$, in agreement with Lüders' remark.

The trade-off implies both *deterministic repeatability* and the weak version of *ideality*.

Lemma 3.5. *If an update rule for non-composite systems \mathbf{w}_A^Λ is (TO), then it is (DR).*

Proof. Consider an arbitrary projector Π_x and a system residing in some state $\rho \in \mathcal{S}(\mathcal{H}_A)$ such that $\text{Tr}(\Pi_x \rho) \neq 0$. Letting $\Pi_y = \Pi_x$ in Eq. (3.94) leads to

$$\text{Tr} \left[\Pi_x \mathbf{w}_A^\Lambda(\Pi_x, \rho) \right] = 1. \quad (3.96)$$

This implies that the post-measurement state $\mathbf{w}_A^\Lambda(\Pi_x, \rho)$ has full support in the subspace $\mathcal{H}_x = \Pi_x \mathcal{H}_A$, i.e.

$$\Pi_x \mathbf{w}_A^\Lambda(\Pi_x, \rho) \Pi_x = \mathbf{w}_A^\Lambda(\Pi_x, \rho), \quad (3.97)$$

which is the defining condition of (DR)—cf. Eq. (3.90). \square

Lemma 3.6. *If an update rule for single systems \mathbf{w}_A^Λ is (TO), then it is weakly (ID).*

Proof. We need to show that Eq. (3.94) implies condition (3.83). From Lemma 3.5, we know that $\mathbf{w}_A^\Lambda(\Pi_y, \rho)$ has full support \mathcal{H}_y . Let the initial state ρ also have full support in \mathcal{H}_y , i.e. $\text{Tr}(\Pi_y \rho) = 1$, and consider an informationally complete set of observables on \mathcal{H}_y . If $\mathbf{w}_A^\Lambda(\Pi_y, \rho) \neq \rho$, then for at least one of the outcomes of an observable in the set, say $\Pi_x \in \mathcal{P}(\mathcal{H}_y)$, we have $\text{Tr} \left[\Pi_x \mathbf{w}_A^\Lambda(\Pi_y, \rho) \right] \neq \text{Tr}(\Pi_x \rho)$, in contradiction with (TO). Therefore, the transformation must *not* disturb the states of measured systems when the outcome is certain, i.e. it satisfies the weaker notion of ideality. \square

The following theorem demonstrates the uniqueness of Axiom (M^L) in relation to (TO), a result briefly discussed in [22], albeit without investigating its validity for composite systems.

Theorem 3.4. *The Lüders rule for single systems \mathbf{w}_A^L is the only one consistent with (TO).*

Proof. Substituting w_A^L in (3.94) leads to

$$\mathrm{Tr}(\Pi_x \Pi_y \rho \Pi_y) = \mathrm{Tr}(\Pi_x \rho), \quad (3.98)$$

which holds for all ρ and $\Pi_x \Pi_y = \Pi_y \Pi_x = \Pi_x$. Hence the Lüders projection for single systems satisfies (TO).

To show the converse, suppose w^A is (TO) but $w_A^A \neq w_A^L$. Then, there must exist some $\Pi_y \in \mathcal{P}(\mathcal{H}_A)$ and $\rho \in \mathcal{S}(\mathcal{H}_A)$ such that $w_A^A(\Pi_y, \rho) \neq w_A^L(\Pi_y, \rho)$. Since both rules are (TO), hence (DR), it holds that $w_A^A(\Pi_y, \rho), w_A^L(\Pi_y, \rho) \in \mathcal{H}_y$. We employ the same argument used in the proof of Lemma 3.6. Consider a set of observables of \mathcal{H}_y which allow to reconstruct any quantum state of \mathcal{H}_y tomographically. As $w_A^A(\Pi_y, \rho)$ and $w_A^L(\Pi_y, \rho)$ are different states, they must yield different predictions for at least one outcome $\Pi_x \in \mathcal{P}(\mathcal{H}_y)$ in the set,

$$\mathrm{Tr}[\Pi_x w_A^A(\Pi_y, \rho)] \neq \mathrm{Tr}[\Pi_x w_A^L(\Pi_y, \rho)]. \quad (3.99)$$

Substituting $w_A^L(\Pi_y, \rho) = \Pi_y \rho \Pi_y / \mathrm{Tr}(\Pi_y \rho)$ and using the fact that $\Pi_x \Pi_y = \Pi_y \Pi_x = \Pi_x$, we obtain

$$\mathrm{Tr}(\Pi_y \rho) \mathrm{Tr}[\Pi_x w_A^A(\Pi_y, \rho)] \neq \mathrm{Tr}(\Pi_x \rho), \quad (3.100)$$

which represents a contradiction with our assumption of (TO). We thus conclude that $w_A^A = w_A^L$. \square

Therefore, the trade-off principle (TO) captures the distinctive behaviour of quantum measurements for single systems. To recover the full quantum mechanical update rule w^L , however, one must add another assumption that pertains to composite systems, such as “composition compatibility” (CC) or “strong indistinguishability” (SI). Fig. 3.5 presents a scheme of the assumptions featured in this alternative derivation of the Lüders rule.

3.5 Summary and future work

In this chapter, we have introduced the framework of Alternative-Measurement Theories and initial findings that show the type of questions it raises. The AMT framework is motivated by the methodology of research conducted in Generalised Probabilistic Theories (GPTs) and ontological models such as Spekken’s toy theory.

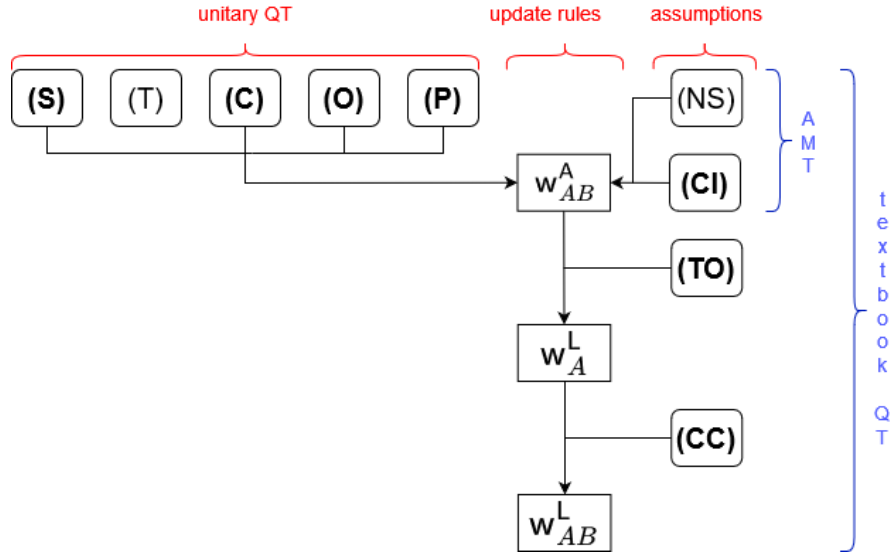


Figure 3.5: Schematic depiction of the alternative argument of Sec. 3.4.4 to derive the Lüders rule in the context of the AMT framework. As in Fig. 3.4, “no-signalling” is not used in the derivation but is implied by the requirement of “composition compatibility” (CC). As mentioned in the text, the derivation also works with (SI) replacing (CC).

AMTs share the Hilbert space structure with quantum mechanics but are characterized by different rules governing the assignment of post-measurement states. The main aims of this framework are (i) to examine the possible alternatives to the Lüders projection that give rise to self-consistent physical theories, and (ii) to explore which features distinguish it from the rest. Our motivations align well with work by Wilson and Ormrod [182], who identify an operational principle that, when supported with the remaining axioms of quantum theory, allows them to derive the unitary deterministic transformations of quantum states, Axiom (T). A more detailed comparison between our work and [182] can be found in Sec. A.3 of Appendix A.

The AMT framework is built on the concept of an *update rule*, see Sec. 3.1.1, which assigns post-measurement states to systems conditioned on the outcomes observed. In particular, to define update rules we do not assume the linearity of the corresponding instrument maps—cf. Sec. 3.1.2. Nonlinear instruments can therefore arise in our framework, which enables us to reconsider the operational role of complete positivity, see Sec. 3.2.1. A summary of the main update rules discussed in this chapter and their operational properties can be found in Table 3.1.

The operational assumptions of “context-independence” and (quantum) “no-signalling” are discussed in Sec. 3.2.4 and 3.2.5, respectively, where we present examples of measurement behaviours ruled out by these requirements. In Sec. 3.3.1, we illustrate how to account for unsharp observables in any AMT and we define generalised instruments describing the impact on states of different experimental strategies to measure the same observables. Generalised instruments are then used in Sec. 3.3.2 to define subtheories, i.e. AMTs that can be faithfully simulated by another AMT.

The linearity over the space of density operators of all instrument maps encodes the impossibility to distinguish between different realisations of the same mixed states. In Sec. 3.4.1, we present different ways to characterise the set of update rules giving rise to linear instruments. Our discussion on linearity connects with a recent debate [81, 168] on the derivation of quantum instruments in [76, 132]. We examine how the two sides of the argument employ different operational definitions of a measurement. Then, we show that the set of AMTs with linear instruments, which includes but is not limited to quantum mechanics, is *mutually simulable*. In fact, any update rule in the set can simulate the deterministic repeatability of quantum measurements by introducing an ancillary system on which to perform measurements.

To uniquely recover the textbook version of the collapse, linearity needs to be supplemented by other assumptions such as ideality—cf. Sec. 3.4.2. We also find, in Sec. 3.4.3, that repeatability alone does not single out the Lüders rule, as shown by suitable modifications of von Neumann’s original postulate. However, the trade-off between the disturbance applied by measurements and information gained—see Assumption (TO)—represents a unique property of the Lüders rule for single systems, Sec. 3.4.4. The argument does not extend to composite systems; additional requirements seem necessary to fully describe measurements in quantum theory.

Looking ahead, we intend to continue searching for operational statements that are equivalent to the projection postulate. An important step would be to characterise all update rules compatible with the standard projection for *single systems*. Presently, only two examples are known, namely w^\perp and \tilde{w}^\perp —cf. Sec. 3.2.1. Other examples may help us to better understand the structure of the quantum mechanical update rule and their relationship to the assumptions underlying the AMT framework.

Another direction is to explore whether “nonlinear” AMTs can exist that have

	(CI)	(NS)	(CC)	(WI)	(SI)	(MM)	(DR)	(PA)	(PR)	(SC)	(ID)	(TO)
w^A :	$\frac{\omega_A^A}{\omega_{AB}^A}$	✓	✓									
w^I :	$\frac{\omega_A^I(\Pi_x, \rho_A) = \Pi_x \rho_A \Pi_x}{\omega_{AB}^I(\Pi_x, \rho_{AB}) = (\omega_x^I \otimes \mathcal{I}_B)(\rho_{AB})}$	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
w^P :	$\frac{\omega_A^P(\Pi_x, \rho_A) = \text{Tr}(\Pi_x \rho_A) \rho_A}{\omega_{AB}^P(\Pi_x, \rho_{AB}) = \text{Tr}(\Pi_x \otimes \mathbb{I}_B \rho_{AB}) \rho_{AB}}$	✓	✓	✓					✓	✓		✓
\tilde{w}^I :	$\frac{\omega_A^I(\Pi_x, \rho_A) = \Pi_x \rho_A \Pi_x}{\tilde{\omega}_{AB}^I(\Pi_x, \rho_{AB}) = \left[\omega_x^I(\text{Tr}_B(\rho_{AB})) \otimes \text{Tr}_A(\rho_{AB}) \right] / \text{Tr}(\rho_{AB})}$	✓	✓	✓	✓	✓	✓	✓	✓	✓	(✓)	✓
w^{mix} :	$\frac{\omega_A^{\text{mix}}(\Pi_x, \rho_A) = \text{Tr}(\Pi_x \rho_A) \mathbb{I}_A / d_A}{\omega_{AB}^{\text{mix}}(\Pi_x, \rho_{AB}) = (\omega_x^{\text{mix}} \otimes \mathcal{I}_B)(\rho_{AB})}$	✓	✓	✓	✓	✓						
\tilde{w}^λ :	$\frac{\omega_A^\lambda(\Pi_x, \rho_A) = \text{Tr}(\Pi_x \rho_A) G_\lambda^{1/2}(\Pi_x) \rho_A G_\lambda^{1/2}(\Pi_x) / \text{Tr}(G_\lambda(\Pi_x) \rho_A)}{\tilde{\omega}_{AB}^\lambda(\Pi_x, \rho_{AB}) = \left[\omega_x^\lambda(\text{Tr}_B(\rho_{AB})) \otimes \text{Tr}_A(\rho_{AB}) \right] / \text{Tr}(\rho_{AB})}$	✓	✓	✓				✓				(✓)
w^μ :	$\frac{\omega_A^\mu(\Pi_x, \rho_A) = \text{Tr}(\Pi_x \rho_A) \left[(1 - \mu) w_A^I(\Pi_x, \rho_A) + \mu w_A^P(\Pi_x, \rho_A) \right]}{\omega_{AB}^\mu(\Pi_x, \rho_{AB}) = \text{Tr}(\Pi_x \otimes \mathbb{I}_B \rho_{AB}) \left[(1 - \mu) w_{AB}^I(\Pi_x, \rho_{AB}) + \mu w_{AB}^P(\Pi_x, \rho_{AB}) \right]}$	✓	✓	✓				✓	✓			✓
\tilde{w}^{VN} :	$\frac{\tilde{\omega}_A^{\text{VN}}(\Pi_x, \rho_A) = \sum_{i=1}^{\text{Tr}(\Pi_x)} \langle x_i \rho_A x_i \rangle \langle x_i }{\tilde{\omega}_{AB}^{\text{VN}}(\Pi_x, \rho_{AB}) = (\tilde{\omega}_x^{\text{VN}} \otimes \mathcal{I}_B)(\rho_{AB})}$	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 3.1: Summary of the operational properties of some of the update rules discussed in Chap. 3. For each update rule w , the state update is described by the induced instrument maps ω for single and composite systems, cf. Eq. (3.10). The parentheses in the (ID) column indicate that the corresponding rule satisfies only the weak notion of ideality. Legend: (CI): context-independence, (NS): quantum no-signalling (Sec. 3.1.2), (CC): composition compatibility (Sec. 3.2.3), (WI): weak indistinguishability, (SI): strong indistinguishability, (MM): equivalence of mono- and multi-partite procedures (Sec. 3.4.1), (DR): deterministic repeatability (Sec. 3.4.3), (PA): partial repeatability (Sec. 3.2.5), (PR): probabilistic repeatability, (SC): sequential commutativity (Sec. 2.3.1), (ID): ideality (Sec. 3.4.2), (TO): information-disturbance trade-off (Sec. 3.4.4).

quantum theory as a subtheory, or to prove the impossibility to do so. We also do not know whether or not there exist AMTs that exhibit super-quantum correlations without violating non-relativistic no-signalling. We also propose to investigate the effect of imposing local tomography on the set of toy models. The consequences of using different update rules on the computational capabilities of the theory will also need to be examined, in analogy with the case of pQT of Chap. 2.

Furthermore, it is worth investigating the potential for generalising the AMT framework beyond the Hilbert space structure. One may consider using our definition of update rule as a template to construct more general operational theories. This would then allow to explore the operational role of sequential measurements if different states, observables and rules for the composition of systems and for outcome probabilities are postulated.

Part II

Support uncertainty relations

Introduction

No quantum particle can reside in a state with both its position and momentum distributions being localised arbitrarily well. For these incompatible observables, Heisenberg’s uncertainty relation [99, 115] establishes a finite lower bound for the product of their variances. This result relies on a fundamental property of Fourier theory: a real (or complex) function with *finite* support on the real line has a Fourier transform which must be non-zero almost everywhere [77]. It is, however, difficult to quantify the support of functions on unbounded intervals. Using variances instead of the supports of probability distributions circumvents this difficulty.

The situation is different for quantum systems with finite-dimensional Hilbert spaces since the *support (size)* of a pure state—defined as the number of non-zero components in a given orthonormal basis—is always finite. A computational basis state in \mathbb{C}^d , say, has support equal to one, and the support of its (discrete) Fourier transform equals d since all basis states contribute. Thus, the *product* of the support sizes equals d which turns out to be its smallest possible value [64].

The underlying product inequality has been generalised in a number of directions [136, 180]. Tao derived an *additive* inequality [173] which is valid in spaces \mathbb{C}^d of prime dimensions $d = p$: the *sum* of the supports of a state and its discrete Fourier transform is bounded from below by the value $(p + 1)$. This bound is sharp since a computational basis state and its Fourier transform saturate it.

Support inequalities and their generalisations have found applications in signal processing [38, 39], for example, and they can be used to identify non-classical quantum states [52] for which the Kirkwood-Dirac quasiprobability distribution [61, 121] is *not* a probability distribution. Such states provide an advantage in quantum metrology [8] and play a role in weak measurements [7, 66, 104] and contextuality

[123].

Using the support size of a quantum state as a measure of uncertainty has an unexpected—and previously unnoticed—operational advantage. Quantum supports take a *finite* set of integer values only, in stark contrast to other measures. Variances of observables in a given state or its von Neumann entropy take *real numbers* as values which demands many measurements to determine experimentally. However, a *finite* number of measurements may already suffice to determine the *exact* support size of a quantum state. This situation occurs whenever the state at hand has (i) “full” support in the basis considered and (ii) each outcome has been registered at least once. Conformity with a given support inequality may be verified by a *finite* number of measurements as small as the bound itself. This property depends, of course, on the assumption that the measuring device never registers incorrect outcomes; limited detection efficiency does not invalidate the argument, however.

Variance-based uncertainty relations also exist for more than two observables associated with multiple orthonormal bases [62, 63, 113, 114]: position and momentum may be supplemented by a third continuous variable which is canonical to each of them. The eigenbases of these three observables are mutually unbiased and related by fractional Fourier transforms. The product of their variances satisfies a *triple uncertainty relation* [113]. Importantly, the best known lower bound of this inequality does *not* follow from the pair uncertainty relations but must be determined independently. No quantum state exists that satisfies all three pair uncertainty relations simultaneously. In a similar vein, entropic uncertainty relations capture the incompatibility of up to $(d + 1)$ observables in finite-dimensional systems, linked to a complete set of mutually unbiased bases known to exist in prime-power dimension [10, 107, 156].

The main goal of this chapter is to extend Tao’s additive support uncertainty relation to the case of more than two bases, inspired by the triple uncertainty relation for continuous variables. The focus will be on prime-dimensional spaces where $(p + 1)$ mutually unbiased bases exist, known as *complete sets*. The support sizes of a state in any pair of mutually unbiased bases from such a set are expected to satisfy Tao’s bound but they might not saturate all pair bounds simultaneously.

Chap. 4 is structured as follows. Sec. 4.1 sets up notation by briefly describing known product and sum inequalities for the support of a vector, and the properties of complete sets of mutually unbiased bases are summarised. In Sec. 4.2, Tao’s additive uncertainty relation for the support of quantum states is shown to hold for

any pair of mutually unbiased bases in a complete set, and the generalised additive support inequality involving all $(p + 1)$ mutually unbiased bases is established as a direct consequence. According to Sec. 4.3, the bounds provided by the generalised support inequality cannot be saturated for prime dimensions $2 \leq d \leq 19$, except $d = 3$. Higher *achievable* bounds are derived in Sec. 4.4 for prime numbers up to $d = 7$. In Sec. 4.5, we summarise and discuss the results obtained. The proofs of some lemmata are relegated to Appendix B.1.

Support inequalities for complete sets of mutually unbiased bases

4.1 Preliminaries

4.1.1 The support size of a quantum state

The *support* (*size*) of a Hilbert-space vector $\psi \in \mathcal{H}_d$ is given by the number of its non-zero expansion coefficients $\psi_v = \langle v | \psi \rangle$ in an orthonormal basis $\mathcal{B} = \{|v\rangle, v = 0, 1, \dots, d-1\}$,

$$|\text{supp}(\psi, \mathcal{B})| = \#(\psi_v \neq 0, v = 0 \dots d-1) \in \{0 \dots d\}. \quad (4.1)$$

The only vector with vanishing support is the zero vector. Due to normalisation, the support of a quantum state must be at least one, and the maximum is achieved whenever the state ψ is a linear combination of all d basis states. The support size of a state clearly depends on the chosen basis.

Thinking of the support size as the (improper) L^0 -“norm”—given $\psi = (\psi_1, \dots, \psi_d)$, it is defined as $\|\psi\|_0 = |\{i : \psi_i \neq 0\}|$ —we will use the notation

$$|\text{supp}(\psi, \mathcal{B})| = \|\psi\|_{\mathcal{B}}. \quad (4.2)$$

The set of expansion coefficients $\{\psi_v, v = 0 \dots d-1\}$ has three obvious *support-conserving* symmetries. The support size is invariant (*i*) under rephasing each expansion coefficient separately,

$$\|\psi\|_{\mathcal{B}} = \|R\psi\|_{\mathcal{B}}, \quad R = \text{diag}(e^{i\tau_0}, e^{i\tau_1}, \dots, e^{i\tau_{d-1}}), \quad (4.3)$$

with real numbers $\tau_v, v = 0, \dots, d-1$; (ii) under permuting the components of any state among themselves

$$\|\psi\|_{\mathcal{B}} = \|P\psi\|_{\mathcal{B}}, \quad P \in S_d, \quad (4.4)$$

where S_d is the permutation group acting on sets of d elements; and (iii) under the complex conjugation of some (or all) of its components,

$$\|\psi\|_{\mathcal{B}} = \|K\psi\|_{\mathcal{B}}, \quad K = \prod_{\text{some } v \in \{0 \dots d-1\}} K_v, \quad (4.5)$$

where each operator $K_v, v = 0, \dots, d-1$, maps one expansion coefficient of the state ψ in the basis \mathcal{B} to its complex conjugate, $K_v\psi_v = \psi_v^*$, and does not change the others. In the basis \mathcal{B} , the permutations P are represented by a matrix of order d containing exactly one unit entry in each row and column; hence the unitary invariances of rephasing and permuting coefficients are conveniently combined into *monomial* matrices $M \equiv RP$. The third invariance described by the operators K_v will play no role.

4.1.2 Support inequalities for a Fourier pair of bases

Given two distinct orthonormal bases \mathcal{B} and \mathcal{B}' of \mathcal{H}_d , one may ask to which extent a state can be “localised” in both of them. Clearly, the product of its support sizes in \mathcal{B} and \mathcal{B}' may take values between one and d^2 . If the bases are related by $\mathcal{B}' = F\mathcal{B}$, where F is the discrete Fourier transform with matrix elements (in the \mathcal{B} -basis)

$$F_{vv'} = \frac{1}{\sqrt{d}} e^{-2\pi i v v' / d} \quad v, v' \in \{0 \dots d-1\}, \quad (4.6)$$

then the *product* of the support sizes of a state ψ and its Fourier transform $\tilde{\psi} = F^\dagger \psi$ is bounded from below [64],

$$\|\psi\|_{\mathcal{B}} \|\tilde{\psi}\|_{\mathcal{B}'} \geq d, \quad (4.7)$$

where we use the fact that the support size of the Fourier transformed state $\tilde{\psi}$ in the basis \mathcal{B} is equal to the support size of the state ψ in the basis \mathcal{B}' , i.e.

$$\|\tilde{\psi}\|_{\mathcal{B}} = \|F^\dagger \psi\|_{\mathcal{B}} = \|\psi\|_{\mathcal{B}'}. \quad (4.8)$$

The inequality (4.7) represents a finite-dimensional analogue of Heisenberg’s uncertainty relation for position and momentum observables of a quantum particle:

quantum states localised in position, say, necessarily come with a broad variance in momentum, the Fourier-transformed position observable.

For spaces \mathcal{H}_d with *prime* dimensions d , an *additive* inequality for the supports of a quantum state in a pair of Fourier-related bases is known [173],

$$\|\psi\|_{\mathcal{B}} + \|\psi\|_{\mathcal{B}'} \geq d + 1, \quad (4.9)$$

which is stronger than the multiplicative relation (4.7), as follows from the inequality $d + 1 - x \geq d/x$, for $x \in [1, d]$. In the terminology of [52], any two bases \mathcal{B} and \mathcal{B}' are said to be *completely incompatible* if and only if the support sizes of the expansion coefficients of any (non-zero) vector $\psi \in \mathcal{H}_d$ satisfy this bound.

The inequality (4.9) is a special case of a theorem valid for finite additive abelian groups G with only trivial subgroups, a condition which implies that the cardinality $|G|$ must be a prime number [173]. Consider a complex-valued function $f : G \rightarrow \mathbb{C}$ and its transform $\tilde{f} : G \rightarrow \mathbb{C}$, defined by

$$\tilde{f}(v') = \frac{1}{\sqrt{|G|}} \sum_{v \in G} f(v) \overline{e(v, v')}, \quad (4.10)$$

where $e(v, v')$ is a “bi-character” of G satisfying $e(v_1 + v_2, v') = e(v_1, v')e(v_2, v')$ and an analogous relation for its second argument. In the particular case of $e(v, v') = e^{-2\pi i v v' / d}$, one obtains an inequality for the supports of f and $\tilde{f} \equiv F^\dagger f$.

Theorem 4.1 (*Tao’s theorem*). *If $f : G \rightarrow \mathbb{C}$ is a non-zero function and the cardinality $|G|$ of the group G is prime, then*

$$|\text{supp}(f)| + |\text{supp}(\tilde{f})| \geq |G| + 1. \quad (4.11)$$

Upon identifying $f(v)$ with $\langle v | \psi \rangle$ and $\tilde{f}(v')$ with $\langle v' | \psi \rangle$, respectively, we obtain the inequality (4.9) relative to the bases \mathcal{B} and \mathcal{B}' introduced via F in Eq. (4.6).

The main ingredient of Tao’s proof is a fundamental property of the Fourier matrix in prime dimensions [78, 169, 173] which dates back to the 1920s: all its square submatrices are invertible.

Theorem 4.2 (*Chebotařev’s theorem*). *If d is prime, then all minors of the Fourier matrix F in Eq. (4.6) are non-zero.*

The inequalities (4.7) and (4.9) involve a *pair* of mutually unbiased bases of \mathcal{H}_d , namely the computational basis \mathcal{B} and its Fourier transform. We will now introduce

larger sets of mutually unbiased bases to formulate more general support inequalities. Not surprisingly, Chebotarëv's theorem must be generalised to other matrices which emerge when establishing bounds on support sizes of quantum states in multiple bases (cf. Sec. 4.2.1).

4.1.3 Mutually unbiased bases in prime dimensions

Two orthonormal bases of the space $\mathcal{H}_d = \mathbb{C}^d$ are said to be *mutually unbiased* (MU) if the inner products between any two states (not of the same basis) have modulus $1/\sqrt{d}$. Then, to know the outcome of a projective measurement performed in one basis implies complete uncertainty about the outcome of a subsequent projective measurement performed in the other.

When d is a power of a prime number p , i.e. $d = p^n$, sets of $(d + 1)$ mutually unbiased bases have been constructed [12, 67, 108, 183]. Such complete sets are both *maximal*—in the sense that no additional MU basis can be added to it—and *tomographically complete*: the probability distributions of outcomes in the $(d + 1)$ bases uniquely encode an unknown quantum state. It is an open problem whether complete sets of MU bases exist in composite dimensions, $d \neq p^n$.

For $d = 2$, the eigenstates of the Pauli operators Z_2 , X_2 and $X_2Z_2 = -iY_2$ form a complete set which has a simple structure. Representing the computational basis \mathcal{B}_0 by the identity matrix $H_0 = I_{(2 \times 2)}$, the following two (2×2) *complex Hadamard matrices*—i.e. $(n \times n)$ unitary matrices with all entries of modulus $1/\sqrt{n}$ —encode the bases which are MU to \mathcal{B}_0 ,

$$H_1 = F = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \quad H_2 = DF \quad \text{where } D = \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}. \quad (4.12)$$

If d is an odd prime, the eigenstates of the $(d + 1)$ *generalised Pauli operators* Z_d , X_d , X_dZ_d , $X_dZ_d^2$, ..., $X_dZ_d^{d-1}$, represent a maximal set of MU bases. The phase and shift operators Z_d and X_d are defined as follows: $Z_d|x\rangle = \omega^x|x\rangle$ and $X_d|x\rangle = |x + 1 \pmod{d}\rangle$, where the states $\{|x\rangle \equiv |\phi_x^0\rangle, x = 0 \dots d - 1\}$ form the computational basis \mathcal{B}_0 and $\omega \equiv e^{2i\pi/d}$ is a d -th root of the number 1 [67]. The k -th state of the j -th basis (following the given order, where $j = 0 : Z_d, j = 1 : X_d, \dots, j = d : X_dZ_d^{d-1}$) is given by

$$|\phi_k^j\rangle = \frac{1}{\sqrt{d}} \sum_{x=0}^{d-1} \omega^{-kx} \omega^{(j-1)x^2} |\phi_x^0\rangle, \quad j \in \{1 \dots d\}, \quad k \in \{0 \dots d - 1\}. \quad (4.13)$$

For each value of j , the equimodular expansion coefficients

$$[H_j]_{xk} = \langle x | \phi_k^j \rangle = \frac{1}{\sqrt{d}} \omega^{-kx} \omega^{(j-1)x^2}, \quad j \in \{1 \dots d\}, x, k \in \{0 \dots d-1\}, \quad (4.14)$$

define a *complex-valued Hadamard matrix*. These are unitary matrices since their columns are given by the components (in the computational basis \mathcal{B}_0) of d orthogonal vectors. When combined with the computational basis \mathcal{B}_0 , the states given in Eq. (4.13) form a complete set of MU bases for Hilbert spaces of prime dimensions, which we will refer to as the *standard set*. In this chapter, all MU bases will be taken from the standard set.

The Hadamard matrix H_1 in (4.14) coincides with the discrete Fourier matrix F given in Eq. (4.6). The remaining Hadamard matrices H_j map the computational basis \mathcal{B}_0 of \mathcal{H}_d to other orthonormal bases denoted by \mathcal{B}_j . Adopting an active view of these transformations, the state ψ is mapped to the state $H_j^\dagger \psi$. The relation between the supports of the state ψ in \mathcal{B}_0 and the j -th MU basis \mathcal{B}_j reads,

$$\|H_j^\dagger \psi\|_0 = \|\psi\|_j, \quad (4.15)$$

abbreviating the notation introduced in (4.8), i.e. $\|\psi\|_{\mathcal{B}_j} \equiv \|\psi\|_j$, $j \in \{0 \dots d\}$.

The columns of the d Hadamard matrices H_j in (4.14) are related in a simple way to each other, namely by

$$|\phi_k^j\rangle = D^{j-1} B^k |\phi_0^1\rangle, \quad j \in \{1 \dots d\}, k \in \{0 \dots d-1\}, \quad (4.16)$$

with two diagonal ($d \times d$) matrices B and D ; in other words, all states of the complete set of MU bases can be generated easily from any given state such as $|\phi_0^1\rangle$ —except for the states of the computational basis \mathcal{B}_0 . Within each Hadamard matrix, the matrix B cyclically shifts a given column to the right,

$$B|\phi_k^j\rangle = \begin{cases} |\phi_{k+1}^j\rangle, & k = 0, \dots, d-2, \\ |\phi_0^j\rangle, & k = d-1; \end{cases} \quad (4.17)$$

its entries are given by the components of the second column of the Fourier matrix $F = H_1$,

$$B = \text{diag} \left(1, \omega^{-1}, \dots, \omega^{-(d-1)} \right), \quad (4.18)$$

except for the factor \sqrt{d} .

The matrix D is given by the components of the first column of the second Hadamard matrix H_2 , i.e.

$$D = \text{diag} \left(1, \omega^1, \dots, \omega^{(d-1)^2} \right), \quad (4.19)$$

cyclically mapping a state of the j -th MU basis to the corresponding one of the MU basis with label $(j + 1)$,

$$D|\phi_k^j\rangle = \begin{cases} |\phi_k^{j+1}\rangle, & j = 1, \dots, d-1, \\ |\phi_k^1\rangle, & j = d. \end{cases} \quad (4.20)$$

In terms of Hadamard matrices, this property reads

$$DH_j = \begin{cases} H_{j+1}, & j = 1, \dots, d-1, \\ H_1, & j = d. \end{cases} \quad (4.21)$$

Writing $H_j = D^{j-1}H_1 \equiv D^{j-1}F$, Chebotarëv's theorem is seen to imply that the minors of all Hadamard matrices H_j , $j = 1 \dots d$, are non-zero: the ranks of the minors of F do not change upon multiplying their rows with non-zero scalars. In view of Eq. (4.12), this generalisation is also valid for the case of $d = 2$.

4.2 Support inequalities...

How well can one localise quantum states simultaneously in $(d + 1)$ MU bases? To answer this question, we need to minimise the support sizes of a state relative to the complete set. In a first step, we now show that the support sizes of a quantum state relative to any *pair* of standard MU bases also satisfy Tao's bound (4.9). Second, by combining the resulting pair inequalities, we establish a state-independent lower bound.

4.2.1 ...for arbitrary pairs of MU bases

Tao's result establishes—for a space of prime dimension d supporting a cyclic abelian group—a sharp inequality for the support sizes of a quantum state and its Fourier transform. Most pairs of the MU bases introduced in Eq. (4.13) are, however, not related by a Fourier transform. Nevertheless, Tao's bound also holds for the supports of the images of any quantum state generated by two Hadamard matrices as we will show now.

Theorem 4.3. *Given any pair of distinct standard MU bases associated with matrices H_j and H_k , $j, k \in \{0 \dots d\}$, the support sizes of a quantum state $|\psi\rangle \in \mathcal{H}_d$ satisfy the state-independent sharp bound,*

$$\|\psi\|_j + \|\psi\|_k \geq d + 1, \quad (4.22)$$

where d is any prime number.

It is important to realise that Theorem 4.3 does not cover *arbitrary* pairs of MU bases in prime dimensions but only those defined in Eq. (4.14). Nevertheless, *all* pairs of MU bases in dimensions $d = 2, 3$ and 5 are found to be completely incompatible since in these dimensions *all* Hadamard matrices are equivalent to the Fourier matrix. Already for the next prime, $d = 7$, other types of Hadamard matrices exist [32].

Proof. The case of dimension $d = 2$ is straightforward. If a state $|\psi\rangle \in \mathcal{H}_2$ has support one in one MU basis, it must have support two in both other bases, due to being MU to their members. Thus, the sum of the supports of any state in two bases must be at least three.

For odd primes d , we will consider two cases separately: either (i) one of the bases in Eq. (4.22) is the computational basis, so that $j = 0$, say, or (ii) neither of them.

(i) Defining the vector $\phi = D^{1-k}\psi$, we obtain

$$\|\psi\|_0 + \|\psi\|_k = \|\psi\|_0 + \|H_k^\dagger \psi\|_0 = \|\psi\|_0 + \|F^\dagger D^{1-k} \psi\|_0 = \|D^{k-1} \phi\|_0 + \|F^\dagger \phi\|_0, \quad (4.23)$$

recalling that $H_k = D^{k-1}F$ holds according to Eq. (4.21). Since D is a diagonal unitary hence a support-conserving unitary matrix (cf. (4.3)), we obtain

$$\|\psi\|_0 + \|\psi\|_k = \|\phi\|_0 + \|F^\dagger \phi\|_0 \geq d + 1, \quad (4.24)$$

where Tao's theorem was used in the last step.

(ii) Now consider the case where the non-zero labels j and k differ from each other. Defining the vector $\phi = H_j^\dagger \psi$, the sum of the support sizes can be written as

$$\|\psi\|_j + \|\psi\|_k = \|H_j^\dagger \psi\|_0 + \|H_k^\dagger \psi\|_0 = \|\phi\|_0 + \|H_k^\dagger H_j \phi\|_0. \quad (4.25)$$

The product $H_k^\dagger H_j$ of two distinct Hadamard matrices is, in fact, always equal to another Hadamard matrix H_t^\dagger , $t \neq 0$, up to a monomial matrix $M(j, k)$; this result

is the content of Lemma 4.1 stated directly after the proof. As the matrix $M(j, k)$ is support-conserving (cf. Eqs. (4.3) and (4.4)) for all states of the space \mathcal{H}_d , we find

$$\|\phi\|_0 + \|M(j, k)H_t^\dagger\phi\|_0 = \|\phi\|_0 + \|H_t^\dagger\phi\|_0 \geq d + 1, \quad (4.26)$$

where (4.24) of Part (i) has been used in the final step. \square

The proof just relies on dissolving products of the Hadamard matrices H_j which encode a complete set of MU bases. Clearly, products of the form $H_k^\dagger H_j$, $j \neq k$, are Hadamard matrices since their matrix elements, being overlaps of MU vectors, have modulus $1/\sqrt{d}$,

$$\left[H_k^\dagger H_j \right]_{\ell\ell'} = \langle \phi_\ell^k | \phi_{\ell'}^j \rangle. \quad (4.27)$$

When $d = 2$, one finds explicitly that $H_1^\dagger H_2 = MH_2^\dagger$ and $H_2^\dagger H_1 = M'H_2^\dagger$, with monomial matrices M and M' . In other words, the phases of the matrix elements (4.27) coincide with those of the adjoint of another transition matrix after permuting and rephasing its rows. This property actually holds for any prime dimension.

Lemma 4.1. *Let d be an odd prime and $j, k \in \{1 \dots d\}$ with $j \neq k$. Then*

$$H_k^\dagger H_j = M(j, k)H_t^\dagger \quad (4.28)$$

for a monomial matrix $M(j, k)$ if and only if $t = 1 + \chi \in \{1, \dots, d\}$ where the integer χ satisfies $4(j - k)\chi = 1 \pmod{d}$.

Proof. See Appendix B.1. \square

Furthermore, Lemma 4.1 allows us to generalise Chebotarëv's theorem (Theorem 4.2) to the product matrices $H_k^\dagger H_j$, for distinct indices j and k .

Corollary 4.1. *If d is prime, then all minors of the Hadamard matrices $H_k^\dagger H_j$, $j, k \in \{0 \dots d\}$ and $j \neq k$, are non-zero.*

Proof. Let d be an odd prime. Any H_t , $t \in \{1 \dots d\}$, has only non-zero minors, as was mentioned after Eq. (4.21), as do the adjoints H_t^\dagger . Therefore, the claim holds if one of the labels j, k , is zero. For both $j, k \neq 0$, Lemma 4.1 applies. Since rephasing and permuting the rows of a matrix do not change the rank of any submatrix, we can conclude that the matrices $H_k^\dagger H_j$, $j, k \neq 0$ and $j \neq k$ also have non-zero minors only.

In dimension $d = 2$, the result follows from inspecting the products $H_1^\dagger H_2 = MH_2^\dagger$ and $H_2^\dagger H_1 = M'H_2^\dagger$. \square

According to Corollary 4.1 the vectors formed by the columns (or rows) of all square submatrices of the Hadamard matrices $H_k^\dagger H_j$, $j \neq k$, are linearly independent. What is more, any set of d (or fewer) vectors taken from *any two* MU bases are linearly independent.

Corollary 4.2. *Given a complete standard set of MU bases in the space \mathcal{H}_d of prime dimension d , any set of d (or fewer) vectors taken from any two MU bases are linearly independent.*

Proof. See Appendix B.1. □

Theorem 4.3 also demonstrates that all pairs of MU bases taken from the complete standard set in prime dimension are *completely incompatible*, in the sense of Sec. 4.1.2. This statement is stronger than the results of [52, 173] since the bases we consider are not necessarily related by the Fourier matrix F .

4.2.2 ...for complete sets of $(d + 1)$ MU bases

Let us denote the sum of the numbers of non-zero expansion coefficients of a state $\psi \in \mathcal{H}_d$ in a complete standard set of MU bases by

$$\mathcal{S}(d) = \|\psi\|_0 + \|\psi\|_1 + \cdots + \|\psi\|_d. \quad (4.29)$$

Then, the inequalities (4.22) imply that the overall support size $\mathcal{S}(d)$ cannot fall below a certain threshold.

Theorem 4.4. *For any prime d , the overall support $\mathcal{S}(d)$ of a quantum state $|\psi\rangle \in \mathcal{H}_d$ in a complete standard set of MU bases satisfies the additive state-independent bound*

$$\mathcal{S}(d) \geq \frac{(d+1)^2}{2} \equiv T(d). \quad (4.30)$$

Proof. Write down $(d+1)$ copies of the support inequality (4.22) with indices $(j, j+1)$, $j = 0, \dots, d-1$, and $(d, 0)$, respectively. Adding them up, the right-hand-sides of (4.22) give $(d+1)^2$. Since each term $\|\psi\|_j \equiv \|H_j^\dagger \psi\|_0$, $j = 0, \dots, d$, occurs twice in the sum, we divide by two and obtain the inequality (4.30). □

An alternative proof treats all pair supports equally: write down (4.22) for all $d(d+1)$ distinct pairs of indices (j, k) and consider the sum of the supports. After removing common factors, the bound (4.30) on $\mathcal{S}(d)$ follows.

Support inequalities other than Eqs. (4.22) and (4.30) exist. They may involve any number between two and $(d + 1)$ MU bases. For example, picking the first three MU bases and combining the associated pair inequalities from (4.22) leads to the additive “triple support inequality”,

$$\mathcal{S}(d; 3) \equiv \|\psi\|_0 + \|F^\dagger\psi\|_0 + \|H_2^\dagger\psi\|_0 \geq \frac{3}{2}(d + 1). \quad (4.31)$$

Clearly, this inequality cannot be saturated for dimension $d = 2$ because the overall support $\mathcal{S}(2)$ of a state is always an integer number. Taking only two possible values, the smallest achievable value of the triple support size $\mathcal{S}(2; 3) \equiv \mathcal{S}(2)$ equals $T_s(2) = 5$; here and in the following, *achievable*—or *sharp*—bounds of $\mathcal{S}(d)$ are denoted by $T_s(d)$.

The lower bound on the triple uncertainty relation for continuous variables [113], derived similarly by combining pair uncertainty relations, can also not be reached. Theorem 4.4 is not constructive, hence it is not obvious whether the case of $d = 2$ represents an exception or whether the inequalities (4.30) are never sharp. In the next section we will first derive some general results about multiple-support inequalities, followed by a closer look at dimensions $3 \leq d \leq 19$.

4.3 Saturating support inequalities for MU bases

To saturate the bound of the inequality (4.30) means to identify states that minimise all support pair relations simultaneously. We present a number of rigorous results for prime dimensions $d \leq 7$. Numerical methods are then used to determine whether the generalised inequality can be saturated for dimensions up to $d = 19$.

4.3.1 Constraints on saturating states

Our first general result is a necessary and sufficient condition that the support inequality (4.30) involving a complete set of $(d + 1)$ MU bases be saturated.

Theorem 4.5 (Equal support sizes). *The additive support inequality for a complete standard set of MU bases (4.30) is saturated by a state $\psi \in \mathcal{H}_d$ if and only if it has the same support in all $(d + 1)$ MU bases, i.e.*

$$\|\psi\|_j = \frac{d + 1}{2}, \quad j \in \{0 \dots d\}, \quad (4.32)$$

where d is an odd prime.

Proof. Substituting the values (4.32) into (4.30) directly produces the lower bound.

For the converse, we show that the supports must have the values given in (4.32) if equality holds in Eq. (4.30). Noting that the support of any state $\psi \in \mathcal{H}_d$ ranges from 1 to d , i.e.

$$\|\psi\|_0 = \frac{d+1}{2} \pm \Delta, \quad \Delta \in \left\{0, 1, \dots, \frac{1}{2}(d-1)\right\}, \quad (4.33)$$

we will proceed by exhausting all its values in the computational basis \mathcal{B}_0 . It turns out that the the minimum in (4.30) cannot be reached if the support is either (i) smaller or (ii) larger than $(d+1)/2$, leaving (iii) the values in (4.32) as the only option.

(i) If $\|\psi\|_0 = (d+1)/2 - \Delta$, $\Delta > 0$, then (4.22) implies that $\|\psi\|_j \geq (d+1)/2 + \Delta$, $j = \{1 \dots d\}$. Hence, the sum of the supports in all $(d+1)$ MU bases equals

$$\begin{aligned} \mathcal{S}(d) &= \sum_{j=0}^d \|\psi\|_j \geq \frac{d+1}{2} - \Delta + d \left(\frac{d+1}{2} + \Delta \right) \\ &\geq \frac{(d+1)^2}{2} + (d-1)\Delta > \frac{(d+1)^2}{2}. \end{aligned} \quad (4.34)$$

Therefore, the inequality cannot be saturated by a state which has support smaller than $(d+1)/2$ in the basis \mathcal{B}_0 .

(ii) Assume that $\|\psi\|_0 = (d+1)/2 + \Delta$, $\Delta > 0$. Clearly, the lower bound of the sum in Eq. (4.30) can only be reached if the support of the state ψ is smaller than $(d+1)/2$ in at least one of the MU bases, $\|\psi\|_{j^*} < (d+1)/2$, $j^* \in \{1 \dots d\}$, say. Repeating the argument from (i) relative to the MU basis \mathcal{B}_{j^*} instead of \mathcal{B}_0 implies that the inequality (4.30) cannot be saturated.

(iii) If $\|\psi\|_0 = (d+1)/2$ then (4.22) implies that $\|\psi\|_j \geq (d+1)/2$, $j = \{1 \dots d\}$. However, given these bounds, the minimum of $\mathcal{S}(d)$ in (4.30) can be achieved only if the support of the state $\psi \in \mathcal{H}_d$ takes the value $(d+1)/2$ in all other MU bases as well. \square

The second general result states that a specific d -th root of unity can appear at most *twice* in the columns of the Hadamard matrices H_j , $j = 2 \dots d$, given in (4.14). The proof of another necessary—but not sufficient—condition for saturating the generalised inequality (4.30) will rely on this limit of the occurrences of roots.

Lemma 4.2 (*Frequency of roots*). *Let d be prime and consider the states $|\phi_k^j\rangle$, $j = 2, \dots, d$, in Eq. (4.13) forming the bases \mathcal{B}_j which are MU to both the identity*

and the Fourier matrix. Any d -th root ω^n , $n \in \{0 \dots d\}$, figures at most twice among the numbers $\sqrt{d}\langle x|\phi_k^j\rangle$, $x \in \{0 \dots d-1\}$.

Proof. We need to determine the number of solutions of the equation $\omega^{-kx+(j-1)x^2} = \omega^n$ which becomes $(j-1)x^2 - kx - n \pmod{d} = 0$ upon taking the logarithm and rearranging. Since $j \neq 1$, the equation is quadratic for each n and there can be at most two integer solutions for the unknown x . The extension to the special case of $d = 2$ is trivial. \square

According to Theorem 4.5, a state saturating (4.30) must have $(d-1)/2$ vanishing expansion coefficients in each MU basis of the standard set, in any odd prime dimension. A third general result is that there are constraints on the distributions of these zeroes when expanded in the MU bases of a complete set.

To spell out these constraints, let us introduce the *zero distributions* \mathcal{Z}^j of a state $\psi \in \mathcal{H}_d$ which list the indices of the vanishing expansion coefficients in the $(d+1)$ bases of the complete set,

$$\mathcal{Z}^j = \left\{ \kappa \in \{0 \dots d-1\} : \langle \phi_\kappa^j | \psi \rangle = 0 \right\}, \quad j = 0, \dots, d. \quad (4.35)$$

Using the relation $\langle \phi_\kappa^j | \psi \rangle = \langle \phi_\kappa^0 | H_j^\dagger | \psi \rangle$, one can also think of \mathcal{Z}^j as the set of vanishing coefficients of the state $H_j^\dagger | \psi \rangle$ in the computational basis.

Two zero distributions of vectors in the same Hilbert space are said to be *compatible*, $\mathcal{Z} \sim \mathcal{Z}'$, if they are equal up to a cyclic shift. In other words, two compatible distributions $\mathcal{Z} = \{\kappa_1, \kappa_2, \dots, \kappa_\delta\}$ and $\mathcal{Z}' = \{\kappa'_1, \kappa'_2, \dots, \kappa'_\delta\}$ must have the same number δ of elements and the mapping $\kappa_i \mapsto \kappa_i + \mu \pmod{d}$ for some fixed integer μ must be a bijection from \mathcal{Z} to \mathcal{Z}' . Compatibility of zero distributions is an equivalence relation between classes of d elements.

The extension of Chebotarëv's Theorem shown in Sec. 4.2.1 and Lemma 4.2 imply a constraint on zero distributions for all prime dimensions $d > 3$. This property will be used in Sec. 4.3.3 to prove that the support inequality (4.30) *cannot* be saturated in dimensions $d = 5$ and $d = 7$.

Theorem 4.6. *Let $d > 3$ be prime and $\psi \in \mathcal{H}_d$ be a state with $(d-1)/2$ expansion coefficients vanishing in the computational basis and in two more standard MU bases, i.e.*

$$\|\psi\|_0 = \|\psi\|_{j_1} = \|\psi\|_{j_2} = \frac{d+1}{2}, \quad j_1 > j_2 \neq 0. \quad (4.36)$$

Then the zero distributions associated with the vectors $H_{j_1}^\dagger|\psi\rangle$ and $H_{j_2}^\dagger|\psi\rangle$, respectively, are incompatible.

Proof. Since the state ψ has $d_- \equiv (d-1)/2$ vanishing components in three bases with labels $j=0, j_1, j_2$, it satisfies $3d_-$ conditions,

$$\begin{aligned} \langle \phi_{\kappa_1^0}^0 | \psi \rangle = \langle \phi_{\kappa_2^0}^0 | \psi \rangle = \dots = 0, & \quad \mathcal{Z}^0 = \{ \kappa_1^0, \kappa_2^0, \dots, \kappa_{d_-}^0 \}, \\ \langle \phi_{\kappa_1^1}^{j_1} | \psi \rangle = \langle \phi_{\kappa_2^1}^{j_1} | \psi \rangle = \dots = 0, & \quad \mathcal{Z}^{j_1} = \{ \kappa_1^1, \kappa_2^1, \dots, \kappa_{d_-}^1 \}, \\ \langle \phi_{\kappa_1^2}^{j_2} | \psi \rangle = \langle \phi_{\kappa_2^2}^{j_2} | \psi \rangle = \dots = 0, & \quad \mathcal{Z}^{j_2} = \{ \kappa_1^2, \kappa_2^2, \dots, \kappa_{d_-}^2 \}. \end{aligned} \quad (4.37)$$

We proceed by contradiction. To assume that the zero distributions \mathcal{Z}^{j_1} and \mathcal{Z}^{j_2} are compatible means that they are related by a cyclic shift by some integer $\mu \in \{0 \dots d-1\}$. In particular, we can arrange the elements in the two sets such that

$$\kappa_i^2 = \kappa_i^1 + \mu \pmod{d} \quad \text{for all } i \in \{1 \dots d_-\}. \quad (4.38)$$

Then, according to Eq. (4.16), the corresponding states must be related by powers of the matrices D and B ,

$$|\phi_{\kappa_i^2}^{j_2}\rangle = D^{j_2-1} B^{\kappa_i^2} |\phi_0^1\rangle = D^{j_2-1} B^{\kappa_i^2} D^{-j_1+1} B^{-\kappa_i^1} |\phi_{\kappa_i^1}^{j_1}\rangle = D^{j_2-j_1} B^\mu |\phi_{\kappa_i^1}^{j_1}\rangle, \quad (4.39)$$

where we have used the fact that D and B commute. Defining $V_\mu^\dagger = D^{j_2-j_1} B^\mu$, the third set of conditions in (4.37) turns into

$$\langle \phi_{\kappa_i^2}^{j_2} | \psi \rangle = \langle \phi_{\kappa_i^1}^{j_1} | V_\mu \psi \rangle = 0 \quad \text{for all } i \in \{1 \dots d_-\}. \quad (4.40)$$

Since V_μ is diagonal in the computational basis, we have

$$\langle \phi_{\kappa_i^0}^0 | \psi \rangle = \langle \phi_{\kappa_i^0}^0 | V_\mu \psi \rangle = 0 \quad \text{for all } i \in \{1 \dots d_-\}, \quad (4.41)$$

which means that \mathcal{Z}^0 and \mathcal{Z}^{j_1} are zero distributions for the pair of vectors ψ and $V_\mu \psi$. In other words, these two states are both orthogonal to the same set of $2d_- = (d-1)$ vectors

$$\left\{ \phi_{\kappa_1^0}^0, \dots, \phi_{\kappa_{d_-}^0}^0, \phi_{\kappa_1^1}^{j_1}, \dots, \phi_{\kappa_{d_-}^1}^{j_1} \right\}, \quad (4.42)$$

stemming from the computational basis \mathcal{B}_0 and the basis \mathcal{B}_{j_1} . According to Corollary 4.2, this is a set of $(d-1)$ linearly independent vectors so that only one unique ray in \mathcal{H}_d can exist that is orthogonal to all of them. Therefore, the vectors ψ and $V_\mu \psi$ must be collinear, i.e. $V_\mu \psi = \lambda \psi$ for some non-zero scalar $\lambda \in \mathbb{C}$.

Since $V_\mu = D^{j_1-j_2}B^{-\mu}$ is diagonal in \mathcal{B}_0 , the computational basis states are eigenvectors of V_μ . By assumption, the state ψ has $d_+ \equiv (d+1)/2$ non-zero coefficients in this basis. Thus, the state ψ will be an eigenvector of the unitary V_μ only if λ is an eigenvalue with multiplicity of d_+ (at least). However, this is impossible for prime dimensions $d > 3$: the non-zero matrix elements on the diagonal of V_μ coincide with the components of the vector $\sqrt{d}|\phi_{-\mu}^{j_1-j_2+1}\rangle$ in the computational basis but for $j_1 > j_2 \neq 0 \pmod{d}$ no more than two of the components may coincide according to Lemma 4.2. Thus, at most two of the eigenvalues of V_μ can coincide. No contradiction arises for dimension $d = 3$ where ψ has exactly two non-vanishing coefficients in the computational basis. \square

4.3.2 Dimension $d = 3$

To prove that the bound (4.30) can be achieved in the space \mathcal{H}_3 , we exhibit the states which minimise the support inequality.

Theorem 4.7. *The state ψ saturates the generalised support inequality (4.30) in dimension $d = 3$ if and only if it is one of the following nine (non-normalised) qutrit states,*

$$\begin{pmatrix} 1 \\ -\omega^m \\ 0 \end{pmatrix}, \quad \begin{pmatrix} 1 \\ 0 \\ -\omega^m \end{pmatrix}, \quad \begin{pmatrix} 0 \\ 1 \\ -\omega^m \end{pmatrix}, \quad m \in \{0, 1, 2\}, \quad (4.43)$$

with $\omega \equiv e^{2i\pi/3}$ being a third root of unity.

Proof. Theorem 4.5 implies that a state ψ saturates Eq. (4.30) with respect to a complete standard set of MU bases if and only if it has support two in each of them, i.e. $\|\psi\|_j \equiv \|H_j^\dagger \psi\|_0 = 2$, $j = 0 \dots 3$. First, we assume that the third component of a candidate state vanishes in the computational basis, i.e. $\psi = (a, b, 0)^T$, with non-zero complex numbers a and b . Applying the matrices H_j^\dagger , $j = 1, 2, 3$, to it, we find four vectors,

$$\begin{pmatrix} a \\ b \\ 0 \end{pmatrix}, \quad \begin{pmatrix} a+b \\ a+\omega b \\ a+\omega^2 b \end{pmatrix}, \quad \begin{pmatrix} a+\omega^2 b \\ a+b \\ a+\omega b \end{pmatrix}, \quad \begin{pmatrix} a+\omega b \\ a+\omega^2 b \\ a+b \end{pmatrix}. \quad (4.44)$$

The components of the last three vectors agree, except for permutations. Hence, support size two can occur in three different ways: one component of each vector

vanishes if

$$b = -\omega^m a, \quad m \in \{0, 1, 2\}, \quad (4.45)$$

holds for some value of m . After removing an irrelevant phase, we obtain the first three vectors given in Eq. (4.43). Second, repeating this argument for initial vectors of the form $\psi = (a, 0, b)^T$ and $\psi = (0, a, b)^T$, respectively, leads to the remaining six vectors in (4.43).

Having exhausted all three-component vectors in the computational basis with support two, we have shown that the nine vectors in (4.43) are the only states saturating the support inequality (4.30) for $d = 3$. \square

4.3.3 Dimensions $d = 5$ and $d = 7$

We will show that it is impossible to reach the lower bound of the support inequality (4.30) in dimensions $d = 5$ and $d = 7$. The proof relies on a property of the zero distributions of the vectors $H_j^\dagger \psi$, $j = 0 \dots d$, which were introduced in Sec. 4.3.1.

Theorem 4.8. *The additive support uncertainty relation (4.30) cannot be saturated in dimensions $d = 5$ and $d = 7$.*

Proof. Let \mathcal{Z}_n^d be the set of the zero distributions with n zeroes among the computational-basis coefficients of qudit states in the Hilbert space \mathcal{H}_d . These distributions are determined by choosing n out of d indices; hence, there are $|\mathcal{Z}_n^d| = \binom{d}{n}$ such sets. Recalling that *compatible* sets of zero distributions form equivalence classes, obtained from rigidly shifting a given one, only $|\mathcal{Z}_n^d / \sim| = \binom{d}{n} / d$ *incompatible* zero distributions exist.

According to Theorem 4.5, a state $|\psi\rangle \in \mathcal{H}_d$ saturating (4.30) for $d > 3$, must have $n = (d - 1)/2$ zeroes in each basis. In addition, a saturating state requires the existence of at least d incompatible zero distributions as Theorem 4.6 does not allow compatible zero distributions for more than two bases. In other words, the inequality $|\mathcal{Z}_{(d-1)/2}^d / \sim| \geq d$ must hold. Clearly, this does not happen for $d = 5$ and $d = 7$ since $|\mathcal{Z}_2^5 / \sim| = 2 < 5$ and $|\mathcal{Z}_3^7 / \sim| = 5 < 7$, respectively. When $d \geq 11$, however, the inequality is satisfied, with $|\mathcal{Z}_5^{11} / \sim| = 42 > 11$, for example. \square

4.3.4 Numerical results for $5 \leq d \leq 19$

For prime numbers d greater than seven, more than d incompatible zero distributions exist which removes the bottleneck we exploited to prove Theorem 4.8. In the

absence of an analytic handle on the problem, we will use numerical means to check whether the bound imposed by (4.30) can be reached for dimensions larger than $d = 7$.

A saturating state necessarily has $(d - 1)/2$ zeroes in each MU basis. Thus, if one picks two distinct MU bases with labels $j_1, j_2 \in \{0 \dots d\}$, say, with corresponding zero distributions \mathcal{Z}^{j_1} and \mathcal{Z}^{j_2} , the state will have vanishing scalar products with a total of $(d - 1)$ states which—in view of Corollary 4.2—are known to be linearly independent. Consequently, there is a *unique* ray $\psi^\perp \in \mathcal{H}_d$ associated with any two zero distributions of the type considered. If the support size of the states ψ^\perp generated in this way (i.e. for all possible choices of initial zero distributions \mathcal{Z}^{j_1} and \mathcal{Z}^{j_2}) is always larger than $(d + 1)/2$ in some third MU basis, then the support inequality (4.30) cannot be saturated: if no state with support size $(d + 1)/2$ in *three* MU bases exists, then no state with support size $(d + 1)/2$ in $(d + 1)$ MU bases will exist. Since only a finite number of zero distributions needs to be checked for a given dimension d , this approach actually represents an *algorithm* to check whether the lower bound can be reached.

Running the program for prime numbers with $5 \leq d \leq 19$ means to check an exponentially increasing number of cases. On a standard PC, the program ran about a second for $d = 5$ and $d = 7$ while it took about a week for $d = 17$. No state has been found which would display $(d - 1)/2$ zeroes in three MU bases. For dimensions $d = 5$ and $d = 7$, this result is stronger than that of Sec. 4.3.3 since the non-existence of a state with two and three zeroes, respectively, is sufficient to derive Theorem 4.8, but not *vice versa*. Due to the exponential increase in the number of zero distributions, dimensions larger than $d = 19$ were out of reach.

4.3.5 Dimensions $d > 19$

To satisfy the additive support inequality (4.30) relative to $(d + 1)$ MU bases, a state needs to satisfy more than one pair relation (4.22) simultaneously which seems unlikely. It is all the more surprising that for dimension $d = 3$ the bound $T(3) = 8$ is actually sharp, i.e. $T_s(3) = T(3)$. Our results for prime dimensions up to $d = 19$ suggest that this case is exceptional.

We conjecture that *the generalised uncertainty relation (4.30) in prime dimensions can only be saturated when $d = 3$* . Here is a plausibility argument to support this view. Assume that a saturating state $\psi \in \mathcal{H}_d$ exists for some prime dimension $d \geq 3$.

According to Theorem 4.5, the state must be orthogonal to exactly $(d - 1)/2$ vectors from each of the $(d + 1)$ MU bases. Corollary 4.2 implies that orthogonality with respect to just two such sets—i.e. $(d - 1)$ vectors—already determines a unique state. Therefore, the remaining $(d - 1)^2/2$ vectors (one set of $(d - 1)/2$ vectors is associated with each of the $(d - 1)$ MU bases not yet considered) must all lie in the same $(d - 1)$ -dimensional subspace orthogonal to the state ψ . This is known to happen for $d = 3$ but seems hard to satisfy for larger dimensions.

4.4 Sharp lower bounds

According to the results presented in Sec. 4.3, no states exist which would saturate the lower bounds (4.30) for the support sizes in dimensions up to $d = 19$, with the exception of $d = 3$. The focus of this section will be on identifying achievable bounds.

4.4.1 Dimension $d = 3$

Theorem 4.7 in Sec. 4.3.2 displays the states which achieve the lower bound (4.30) in dimension $d = 3$. In other words, the bound for the overall support of qutrit states ψ is sharp, $\mathcal{S}(3) \geq 8$, where $\mathcal{S}(d) \equiv \sum_{j=0}^d \|\psi\|_j$ for $\psi \in \mathcal{H}_d$.

4.4.2 Dimension $d = 5$

Theorem 4.8 shows that, for any state $\psi \in \mathcal{H}_5$, the overall support of the states $H_j^\dagger \psi$, $j = \{0 \dots d\}$, must satisfy $\mathcal{S}(5) > 18$. In this section, we will prove a sharp lower bound, namely $\mathcal{S}(5) \geq 22 \equiv T_s(5)$.

To begin, we generalise Lemma 4.2 which will be necessary for the proof of Lemma 4.4.

Lemma 4.3. *Let d be an odd prime and $\omega \equiv e^{\frac{2i\pi}{d}}$. Consider two states $|\phi_{k_1}^{j_1}\rangle, |\phi_{k_2}^{j_2}\rangle \in \mathcal{H}_d$ taken from different standard MU bases, $j_1, j_2 \neq 0$, and let $\{|x\rangle\}$ be the computational basis. Then there can be at most two values of $x \in \{0 \dots d - 1\}$ such that*

$$\langle x | \phi_{k_1}^{j_1} \rangle = \omega^n \langle x | \phi_{k_2}^{j_2} \rangle \quad (4.46)$$

for the same value of $n \in \{0 \dots d - 1\}$. If two different states are taken from the same basis, $j_1 = j_2$, then the equation has exactly one solution for each value of n .

Proof. See Appendix B.1. \square

Now consider a state with two vanishing expansion coefficients in both the computational basis and a second basis of the complete set. It turns out that such a state can have only non-zero coefficients in the remaining four bases, resulting in a total support size of 26.

Lemma 4.4. *If the support of a state $\psi \in \mathcal{H}_5$ equals three in both the computational basis and another standard MU basis with label $j \neq 0$, i.e. $\|\psi\|_0 = \|\psi\|_j = 3$, then its support size in each of the remaining four bases equals five, $\|\psi\|_{j'} = 5$, with $j' \neq 0, j$.*

Proof. The proof, given in Appendix B.1, uses Corollary 4.1 and Lemma 4.3. \square

This result allows us to determine a sharp bound $T_s(5)$ for the support size $\mathcal{S}(5)$.

Theorem 4.9. *Given a state $\psi \in \mathcal{H}_5$, the sharp bound on its overall support size $\mathcal{S}(5)$ in a complete standard set of MU bases is given by $T_s(5) = 22$.*

Proof. To construct the bound, we go through all possible values of the support size of the state ψ in the computational basis, i.e. $\|\psi\|_0 \in \{1 \dots 5\}$.

For $\|\psi\|_0 = 1$, the pair inequalities (4.22) imply that the state ψ must have full support in all other five MU bases, i.e. $\|\psi\|_{j \neq 0} = 5$. Hence, the overall support of a computational basis state is given by $\mathcal{S}(5) = 26$.

For $\|\psi\|_0 = 2$, the pair inequalities (4.22) imply that the state ψ can have at most one zero in each of the other five MU bases, i.e. $\|\psi\|_{j \neq 0} = 4$. Hence, the overall support of ψ is given by $\mathcal{S}(5) = 22$. All 300 states of the form

$$|\psi\rangle = \frac{1}{\sqrt{2}} \left(|\phi_{k_1}^j\rangle - \omega^n |\phi_{k_2}^j\rangle \right), \quad j \in \{0 \dots 5\}, \quad k_1, k_2, n \in \{0 \dots 4\}, \quad k_1 \neq k_2. \quad (4.47)$$

achieve this bound. More generally, for primes $d > 3$, there are

$$(d+1)d \binom{2}{d} = \frac{1}{2} (d^2 - 1) d^2 \quad (4.48)$$

such states as j takes $(d+1)$ values, n takes d values and there are $\binom{2}{d}$ different pairs of k_1 and k_2 . The three-dimensional case is an exception, as demonstrated by Theorem 4.7.

For $\|\psi\|_0 = 3$, the pair inequalities (4.22) rule out a support size lower than three in any basis from the set. We apply Lemma 4.4: the support size of ψ can equal

three in only one other MU basis while the state must have full support in the others, leading to $\mathcal{S}(5) = 26$. It is also possible to have $\|\psi\|_j = 4$ in all bases but the first one, i.e. for $j \neq 0$. In this case seven expansion coefficients would vanish over the complete set, resulting in an overall support size of $\mathcal{S}(5) = 23$. This bound is larger than the one already obtained for the case of $\|\psi\|_0 = 2$.

If $\|\psi\|_0 = 4$ and all other support sizes are also equal to four, the resulting overall support of $\mathcal{S}(5) = 24$ is again larger than the previous bound of $\mathcal{S}(5) = 22$ obtained for $\|\psi\|_0 = 2$. To improve on the value of $\mathcal{S}(5) = 24$, at least one of the other norms must fall below four, i.e. $1 \leq \|\psi\|_{j^*} \leq 3$ for some $j^* \neq 0$. This assumption, however, sends us back to one of the cases already discussed: we formally map $j^* \mapsto 0$ and repeat the arguments given for $1 \leq \|\psi\|_0 \leq 3$.

Similarly, full support in all six MU bases cannot beat any of the bounds given so far. Improving on the value of $\mathcal{S}(5) = 30$ is only possible by decreasing some of the support sizes, so that we will end up in one of the previously discussed cases. Having considered all support sizes of a state in a basis, we have exhausted all possibilities and conclude that the bound on the overall support of a state $\psi \in \mathcal{H}_5$ in six MU bases is indeed given by $T_s(5) = 22$. \square

4.4.3 Dimension $d = 7$

Our aim is to identify states which minimise the overall support $\mathcal{S}(7) = \sum_{j=0}^7 \|\psi\|_j$. To determine the sharp bound for $d = 7$, we will proceed as in the previous section. However, since no equivalent to Lemma 4.4 is known, we will partly rely on numerical results.

For $\|\psi\|_0 = 1$, the pair inequalities (4.22) imply that the state ψ must have full support in all other seven MU bases, $\|\psi\|_{j \neq 0} = 7$. Hence, the overall support of ψ is given by $\mathcal{S}(7) = 50$.

For $\|\psi\|_0 = 2$, the pair inequalities (4.22) imply that the state ψ can have at most one zero in each of the other seven MU bases, i.e. $\|\psi\|_{j \neq 0} = 6$. Hence, the overall support of ψ is given by $\mathcal{S}(7) = 44$, achieved by states of the form

$$|\psi\rangle = \frac{1}{\sqrt{2}} \left(|\phi_{k_1}^j\rangle - \omega^n |\phi_{k_2}^j\rangle \right), \quad j \in \{0 \dots 7\}, \quad k_1, k_2, n \in \{0 \dots 6\}, \quad k_1 \neq k_2. \quad (4.49)$$

According to Eq. (4.48), there are 1176 such states.

For $\|\psi\|_0 = 3$, the smallest possible value of $\mathcal{S}(7)$ compatible with the pair inequalities is $\mathcal{S}(7) = 38$, as the states in the other bases must have support size at

least five each, i.e. $\|\psi\|_{j \neq 0} = 5$. However, no state achieving this bound has been found (numerically). The computations show that a state with support sizes three and five in two MU bases must have full support in the remaining six MU bases so that $\mathcal{S}(7) = 50$. Assuming support size six in all but the first MU basis, the overall support would be $\mathcal{S}(7) = 45$ which is higher than the bound of $\mathcal{S}(7) = 44$ achievable for $\|\psi\|_0 = 2$.

Given a support size of four in the first MU basis, $\|\psi\|_0 = 4$, not all other support sizes can be equal to four according to Theorem 4.8. One case corresponds a state having support size four in the first and one other MU basis. It is possible to (numerically) construct states for which the remaining six supports sizes must be equal to six, leading to $\mathcal{S}(7) = 44$. We neither know analytic expressions for these states nor their total number. The other scenario compatible with $\|\psi\|_0 = 4$ corresponds to the remaining seven support sizes each equalling five, i.e. $\|\psi\|_{j \neq 0} = 5$, leading to $\mathcal{S}(7) = 39$. However, this case cannot be realised: our numerical investigations show that any set composed of three vectors from \mathcal{B}_0 and two vectors from each $\mathcal{B}_{j \neq 0}$ spans the full space.

Assume now that $\|\psi\|_0 = 5$ and that the support of ψ in the other MU bases is also at least five (we exclude all cases with $\|\psi\|_{j^*} < 5$ for some $j^* \neq 0$ since—upon relabeling the MU bases—they have effectively already been considered). An overall support of $\mathcal{S}(7) = 40$ results, below the previously obtained value of $\mathcal{S}(7) = 44$ for $\|\psi\|_0 = 2$. Numerically searching for states achieving this bound, we find that pairs of states with support size five in two MU bases exist but no triples, ruling out the value $\mathcal{S}(7) = 40$. Assuming support size five in two bases and at least six in the remaining six MU bases leads to a higher support size, $\mathcal{S}(7) = 46$.

Starting out with a support size of $\|\psi\|_0 > 5$, no smaller lower bound will exist if all other support sizes take a value of at least six as $\mathcal{S}(7) \geq 48$ follows immediately. If not all support sizes take a value of at least six we are being sent back to a previously discussed case. Thus, we have established the sharp bound of $T_s(7) = 44$ on the overall support of seven-component vectors in eight MU bases, partly relying on numerics.

4.5 Summary and conclusions

Tao's uncertainty relation provides a lower bound on the sum of the support sizes of a state $\psi \in \mathcal{H}_d$ in the standard basis and its Fourier transform, for prime dimensions d . By generalising the bound to arbitrary pairs of mutually unbiased bases (cf. Theorem 4.3), we show in Theorem 4.4 that the sum of the support sizes of a state ψ in a complete standard set of $(d + 1)$ MU bases cannot fall below $T(d) \equiv (d + 1)^2/2$. The bound is found to be sharp for $d = 3$, and proofs were given that it cannot be saturated for dimensions $d = 2, 5$ and 7 . Numerical results indicate that no states exist which achieve the bound for prime numbers up to $d \leq 19$. Table 4.1 summarises these results. We conjecture that the inequality is saturated in dimension $d = 3$ only.

d	2	3	5	7	11	13	17	19
$T(d)$	9/2	8	18	32	72	98	162	200
$T(d)$ achievable?	×	✓	×	×	(×)	(×)	(×)	(×)
$T_s(d)$	5	8	22	(44)	?	?	?	?

Table 4.1: Lower and sharp bounds $T(d)$ and $T_s(d)$, respectively, on the support sizes of states $\psi \in \mathcal{H}_d$ when expanded in the complete standard set of $(d + 1)$ MU bases, for small prime dimensions (numerical results in parentheses).

Tao's pair support inequality has been used to identify *KD-nonclassical* states, i.e. states for which the Kirkwood-Dirac quasiprobability distribution has negative or complex contributions [52]. Given two orthonormal bases of a finite-dimensional space $\mathcal{H}_d, d \in \mathbb{N}$, with no common elements, a state ψ is found to be KD-nonclassical if the sum of its support sizes in these bases is greater than $(d + 1)$. KD-classicality is readily generalised to complete sets of MU bases instead of pairs only. In this context, the results of Sec. 4.3 mean that no states exist which are KD-classical with respect to the standard set of $(d + 1)$ MU bases in small prime dimensions. When $d = 3$, the claim follows by directly computing the complex KD distributions of the nine minimal uncertainty states of Eq. (4.43).

The uncertainty of quantum states involving more than two MU bases has been studied before. Building on a result for a pair of mutually unbiased observables [130], entropic uncertainty relations have been found which involve $(d + 1)$ MU bases [107, 156]. Similarly, Heisenberg's uncertainty relation for continuous variables has a counterpart based on *three* observables satisfying the canonical commutation relation pairwise [113]. Often, the generalisations are straightforward but the resulting

inequalities tend not to be achievable. Sharp bounds and the minimising states are usually difficult to find (see e.g. [114, 187] and the review [178]). In this respect, the additive inequality proposed here is no exception.

Support uncertainty relations for multiple MU bases have many interesting features. As for the pair inequalities, a *finite* number of measurements can be sufficient to confirm that a quantum state satisfies a specific bound. The minimum number of required measurements is simply given by the value of the relevant bound, be it sharp or not: it is sufficient that $T(d)$ different outcomes be registered when measurements in the MU bases are performed on the state ψ . This property also ensures that KD-nonclassicality may sometimes be detected with a finite number of measurements.

Furthermore, the lower bounds of support inequalities neither depend on the state considered nor on the value of Planck's constant. The absence of \hbar as a parameter suggests that no support inequalities for continuous variables will emerge in the limit of systems with ever larger dimensions d . The maximal support size of a quantum state grows without bound and, therefore, does not approach a well-defined quantitative measure for uncertainty. Finally, we would like to point out that determining bounds on support sizes is experimentally difficult since they are basis-dependent quantities.

Establishing sharp bounds for dimensions $d \geq 11$ remains an open question which will require new insights since numerical approaches become unfeasible with increasing dimensions. Other directions of future work will be to study support uncertainty relations for smaller sets of MU bases such as triples, for example. The simplification stems from the considerably smaller number of parameters in comparison to complete MU sets. Preliminary analytical and numerical results for small prime dimensions $3 \leq d \leq 19$ suggest that no state can saturate the bound $T(d; 3)$ on the triple uncertainty relation (4.31) for $d \neq 3$.

— A —

Appendix to Part I

A.1 Definitions of operator sets

In the following list, let \mathcal{H} be a complex, finite-dimensional Hilbert space.

- *Bounded operators*: $\mathcal{L}(\mathcal{H}) : \{L : \mathcal{H} \rightarrow \mathcal{H} : L(c\psi + \phi) = cL(\psi) + L(\phi) \forall \psi, \phi \in \mathcal{H}, c \in \mathbb{C}; \exists t \geq 0 : \|L\psi\| \leq t\|\psi\| \forall \psi \in \mathcal{H}\}$;
- *Self-adjoint operators (Hermitian)*: $\mathcal{L}_s(\mathcal{H}) = \{M \in \mathcal{L}(\mathcal{H}) : M^\dagger = M\}$;
- *Projectors*: $\mathcal{P}(\mathcal{H}) = \{\Pi \in \mathcal{L}_s(\mathcal{H}) : \Pi^2 = \Pi\}$;
- *Unitary operators*: $\mathcal{U}(\mathcal{H}) = \{U \in \mathcal{L}(\mathcal{H}) : U^\dagger U = U U^\dagger = \mathbb{I}\}$;
- *Density operators (normalised)*: $\mathcal{S}(\mathcal{H}) = \{\rho \in \mathcal{L}_s(\mathcal{H}) : \rho \geq 0, \text{Tr}(\rho) = 1\}$.

A.2 Proofs of Lemmata 3.1 and 3.2

Lemma 3.1. *An update rule w^A is (SI) if and only if, for any generalised observable M on \mathcal{H}_A , all M -compatible generalised instruments realisable in the corresponding AMT are composed of linear maps over $\bar{\mathcal{S}}(\mathcal{H}_A)$.*

Proof. It follows from Eq. (3.49) that, if the update rule w^A is (SI), hence the map ω_{AB}^A is linear over $\mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$, then all generalised instruments are composed of linear maps over $\mathcal{S}(\mathcal{H}_A)$. It remains to show that any AMT that is not (SI) will necessarily feature some nonlinear instruments. This is trivial whenever ω_A^A is nonlinear. Suppose now w^A is (WI) but not (SI), i.e. ω_A^A is linear but ω_{AB}^A is not. There will exist at least two (normalised) joint states, ρ_{AB} and σ_{AB} , and a local

outcome Π_x of some observable $N = \{\Pi_x\}_x$ of \mathcal{H}_A , such that

$$\omega_{AB}^A(\Pi_x, \lambda\rho_{AB} + (1-\lambda)\sigma_{AB}) \neq \lambda\omega_{AB}^A(\Pi_x, \rho_{AB}) + (1-\lambda)\omega_{AB}^A(\Pi_x, \sigma_{AB}) \quad (\text{A.1})$$

for some $0 < \lambda < 1$. Tracing out subsystem ‘B’ on both sides must lead to an equality, since we are assuming (WI). However, for any pair of different joint states with same reduced state for a subsystem (here: the left- and right-hand sides of Eq. (A.1)), there exists a channel mapping them to other joint states with different reduced states¹. Denoting this (possibly outcome-dependent) channel by η_{AB}^x , we have that, for this specific value of λ and pair of states (ρ_{AB}, σ_{AB}) ,

$$\begin{aligned} \text{Tr}_B \left[\eta_{AB}^x \circ \omega_{AB}^A(\Pi_x, \lambda\rho_{AB} + (1-\lambda)\sigma_{AB}) \right] &\neq \\ \lambda \text{Tr}_B \left[\eta_{AB}^x \circ \omega_{AB}^A(\Pi_x, \rho_{AB}) \right] + (1-\lambda) \text{Tr}_B \left[\eta_{AB}^x \circ \omega_{AB}^A(\Pi_x, \sigma_{AB}) \right] &. \end{aligned} \quad (\text{A.2})$$

Note that we do not require Eq. (A.2) to hold for arbitrary mixtures of arbitrary states.

In the remaining steps of the proof, we will show that Eq. (A.2) implies the existence of a nonlinear generalised instrument. Let $\mathcal{H}_E = \mathcal{H}_B \otimes \mathcal{H}_{B'}$ denote a composite system, initialised in state $\xi = |0\rangle\langle 0| \otimes \xi_{B'}$, and let $U \in \mathcal{U}(\mathcal{H}_A \otimes \mathcal{H}_E)$ be the unitary inducing the following channel on $\mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$,

$$\eta_U(\varphi_{AB}) = \text{Tr}(|00\rangle\langle 00| \varphi_{AB}) \rho_{AB} + \text{Tr}((\mathbb{I}_{AB} - |00\rangle\langle 00|) \varphi_{AB}) \sigma_{AB}. \quad (\text{A.3})$$

In other words, $\langle \mathcal{H}_{B'}, U, \xi_{B'} \rangle$ defines a dilation of the channel η_U , i.e. $\eta_U(\varphi_{AB}) = \text{Tr}_{B'}(U \varphi_{AB} \otimes \xi_{B'} U^\dagger)$. Setting $\rho_A = |0\rangle\langle 0|$ and $\sigma_A = |1\rangle\langle 1|$, it follows that

$$\eta_U(\rho_A \otimes |0\rangle\langle 0|) = \rho_{AB}, \quad \eta_U(\sigma_A \otimes |0\rangle\langle 0|) = \sigma_{AB}, \quad (\text{A.4})$$

and that

$$\eta_U([\lambda\rho_A + (1-\lambda)\sigma_A] \otimes |0\rangle\langle 0|) = \lambda\rho_{AB} + (1-\lambda)\sigma_{AB}. \quad (\text{A.5})$$

¹For any $\varphi_1, \varphi_2 \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$ such that $\varphi_1 \neq \varphi_2$ but $\text{Tr}_B(\varphi_1) = \text{Tr}_B(\varphi_2)$, there must exist some $\Pi_y \in \mathcal{P}(\mathcal{H}_A \otimes \mathcal{H}_B)$ such that $p(y|\varphi_1) = \text{Tr}(\Pi_y \varphi_1) \neq \text{Tr}(\Pi_y \varphi_2) = p(y|\varphi_2)$. Then, fixing some states $\xi_1, \xi_2 \in \mathcal{S}(\mathcal{H}_A \otimes \mathcal{H}_B)$ with $\text{Tr}_B(\xi_1) \neq \text{Tr}_B(\xi_2)$, the channel $\eta(\varphi) = p(y|\varphi)\xi_1 + (1-p(y|\varphi))\xi_2$ is such that $\text{Tr}_B[\eta(\varphi_1)] \neq \text{Tr}_B[\eta(\varphi_2)]$. The channel can be implemented unitarily on a larger space $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$, where \mathcal{H}_C is an ancilla.

Therefore, using the property of self-consistency (R3) of update rules (cf. Def. 9), we can write

$$\omega_{AB}^A(\Pi_x, \rho_{AB}) = \omega_{AB}^A(\Pi_x, \eta_U(\rho_A \otimes |0\rangle\langle 0|)) \quad (\text{A.6})$$

$$= \omega_{AB}^A(\Pi_x, \text{Tr}_{B'}(U \rho_A \otimes |0\rangle\langle 0| \otimes \xi_{B'} U^\dagger)) \quad (\text{A.7})$$

$$= \text{Tr}_{B'}[\omega_{ABB'}^A(\Pi_x, U \rho_A \otimes |0\rangle\langle 0| \otimes \xi_{B'} U^\dagger)] . \quad (\text{A.8})$$

Similar expressions can be written for σ_{AB} and the mixture $\lambda\rho_{AB} + (1-\lambda)\sigma_{AB}$. Substituting these expressions in Eq. (A.2), and recalling that $\xi = |0\rangle\langle 0| \otimes \xi_{B'}$, leads to the following relation,

$$\begin{aligned} & \text{Tr}_B \left\{ \eta_{AB}^x \circ \text{Tr}_{B'} \left[\omega_{ABB'}^A(\Pi_x, U(\lambda\rho_A + (1-\lambda)\sigma_A) \otimes \xi U^\dagger) \right] \right\} \neq \\ & \quad \lambda \text{Tr}_B \left\{ \eta_{AB}^x \circ \text{Tr}_{B'} \left[\omega_{ABB'}^A(\Pi_x, U\rho_A \otimes \xi U^\dagger) \right] \right\} + \\ & \quad (1-\lambda) \left\{ \eta_{AB}^x \circ \text{Tr}_{B'} \left[\omega_{ABB'}^A(\Pi_x, U\sigma_A \otimes \xi U^\dagger) \right] \right\} . \quad (\text{A.9}) \end{aligned}$$

Since the channel η_{AB}^x does not act on ‘B’’, we can take the partial trace over ‘B’ outside the curly bracket,

$$\begin{aligned} & \text{Tr}_{BB'} \left[\eta_{AB}^x \otimes \mathcal{I}_{B'} \circ \omega_{ABB'}^A(\Pi_x, U(\lambda\rho_A + (1-\lambda)\sigma_A) \otimes \xi U^\dagger) \right] \neq \\ & \quad \lambda \text{Tr}_{BB'} \left[\eta_{AB}^x \otimes \mathcal{I}_{B'} \circ \omega_{ABB'}^A(\Pi_x, U\rho_A \otimes \xi U^\dagger) \right] + \\ & \quad (1-\lambda) \text{Tr}_{BB'} \left[\eta_{AB}^x \otimes \mathcal{I}_{B'} \circ \omega_{ABB'}^A(\Pi_x, U\sigma_A \otimes \xi U^\dagger) \right] . \quad (\text{A.10}) \end{aligned}$$

Recalling that $\mathcal{H}_E = \mathcal{H}_B \otimes \mathcal{H}_{B'}$, and relabeling ‘A’ with ‘S’, we recast Eq. (A.10) in the notation of Eq. (3.49),

$$\begin{aligned} & \text{Tr}_E \left[\eta_{SE}^x \circ \omega_{SE}^A(\Pi_x, U(\lambda\rho_S + (1-\lambda)\sigma_S) \otimes \xi U^\dagger) \right] \neq \\ & \quad \lambda \text{Tr}_E \left[\eta_{SE}^x \circ \omega_{SE}^A(\Pi_x, U\rho_S \otimes \xi U^\dagger) \right] + \\ & \quad (1-\lambda) \text{Tr}_E \left[\eta_{SE}^x \circ \omega_{SE}^A(\Pi_x, U\sigma_S \otimes \xi U^\dagger) \right] , \quad (\text{A.11}) \end{aligned}$$

where we set $\eta_{SE}^x = \eta_{SB}^x \otimes \mathcal{I}_{B'}$. Applying Eq. (3.49) to both sides leads to

$$\omega_{M_x}(\lambda\rho_S + (1-\lambda)\sigma_S) \neq \lambda\omega_{M_x}(\rho_S) + (1-\lambda)\omega_{M_x}(\sigma_S) . \quad (\text{A.12})$$

Therefore, for any (WI) update rule, there exists some generalised observable $\mathbf{M} = \{M_x\}_x$ on \mathcal{H}_S —with measurement model $\mathcal{M} = \langle \mathcal{H}_E, \xi, U, N = \{\Pi_x \in \mathcal{P}(\mathcal{H}_S)\}_x \rangle$ —such that an \mathbf{M} -compatible generalised instrument, $\{\omega_{M_x}\}_x$, is not composed of linear maps over $\bar{\mathcal{S}}(\mathcal{H}_S)$. Specifically, ω_{M_x} does not preserve arbitrary convex combinations of $\rho_S = |0\rangle\langle 0|$ and $\sigma_S = |1\rangle\langle 1|$. It follows that the requirement of (SI) implies the linearity of all generalised instruments in the corresponding AMT. \square

Lemma 3.2. *Let $M, N : \mathcal{L}(\mathcal{H}_A \otimes \mathcal{H}_B) \rightarrow \mathcal{L}(\mathcal{H}_A \otimes \mathcal{H}_B)$ be linear mappings satisfying $M(\rho_A \otimes \rho_B) = N(\rho_A \otimes \rho_B)$ for all $\rho_A \in \mathcal{S}(\mathcal{H}_A)$ and $\rho_B \in \mathcal{S}(\mathcal{H}_B)$. Then $M = N$.*

Proof. Let X and Y be arbitrary elements of $\mathcal{L}(\mathcal{H}_A)$ and $\mathcal{L}(\mathcal{H}_B)$, respectively. They can be expressed as a complex linear combination of self-adjoint operators,

$$X = H + iK, \quad H, K \in \mathcal{L}_s(\mathcal{H}_A), \quad (\text{A.13})$$

$$Y = W + iZ, \quad W, Z \in \mathcal{L}_s(\mathcal{H}_B). \quad (\text{A.14})$$

Moreover,

$$H = H_+ - H_-, \quad H_+, H_- \geq O, \quad (\text{A.15})$$

$$K = K_+ - K_-, \quad K_+, K_- \geq O, \quad (\text{A.16})$$

$$W = W_+ - W_-, \quad W_+, W_- \geq O, \quad (\text{A.17})$$

$$Z = Z_+ - Z_-, \quad Z_+, Z_- \geq O. \quad (\text{A.18})$$

This allows us to express X as follows,

$$\begin{aligned} X = \text{Tr}(H_+) \underbrace{\frac{H_+}{\text{Tr}(H_+)}}_{\rho_H^+} - \text{Tr}(H_-) \underbrace{\frac{H_-}{\text{Tr}(H_-)}}_{\rho_H^-} \\ + i \text{Tr}(K_+) \underbrace{\frac{K_+}{\text{Tr}(K_+)}}_{\rho_K^+} - i \text{Tr}(K_-) \underbrace{\frac{K_-}{\text{Tr}(K_-)}}_{\rho_K^-}, \end{aligned} \quad (\text{A.19})$$

where $\rho_i^j \in \mathcal{S}(\mathcal{H}_A)$ for $i \in \{H, K\}$ and $j \in \{+, -\}$. To avoid issues of ill-definedness, we set $\rho_H^+ = \rho_*$ if $H_+ = O$, where $\rho_* \in \mathcal{S}(\mathcal{H}_A)$ is arbitrary, with analogous conditions for the other states. A similar expression for Y can be written,

$$Y = \text{Tr}(W_+) \rho_W^+ - \text{Tr}(W_-) \rho_W^- + i \text{Tr}(Z_+) \rho_Z^+ - i \text{Tr}(Z_-) \rho_Z^-, \quad (\text{A.20})$$

where $\rho_m^n \in \mathcal{S}(\mathcal{H}_B)$ for $m \in \{W, Z\}$ and $n \in \{+, -\}$. Therefore, $X \otimes Y$ can be expressed as a linear combination (over \mathbb{C}) of product density matrices,

$$X \otimes Y = \sum_{ijmn} c_{im}^{jn} \rho_i^j \otimes \rho_m^n. \quad (\text{A.21})$$

From the linearity of M and N , and the condition that $M(\rho_i^j \otimes \rho_m^n) = N(\rho_i^j \otimes \rho_m^n)$ for all i, j, m, n , it follows that $M(X \otimes Y) = N(X \otimes Y)$ for arbitrary $X \otimes Y \in \mathcal{L}(\mathcal{H}_A \otimes \mathcal{H}_B)$. The two maps are therefore equal, $M = N$. \square

A.3 Comparison with the work of Wilson and Ormrod

In a recent pre-print [182], Wilson and Ormrod adopt a similar approach to the one presented in this chapter, contributing to showcasing the types of arguments and results that can be achieved. Their objective is to identify an operational principle that allows for the unique recovery of the unitary, *deterministic* transformations of quantum states—i.e. Axiom (T)—while retaining all other standard postulates of quantum theory, including the Lüders projection for single systems, Axiom (M^L).

The authors show that imposing unitary dynamics of quantum states is equivalent to assuming the “*local applicability*” of deterministic transformations. Roughly speaking, this refers to the possibility of applying a transformation to a system without having it act on a (possibly far away) environment. In particular, the transformation is independent of any measurement performed on the environment. Interestingly, we observe that their definition of local applicability and our definition of update rule, cf. Def. 9, overlap in certain aspects. Both concepts are defined in terms of a number of basic operational conditions, none of which depend on any prior notion of linearity. Specifically, a condition of self-consistency (R3) also appears in the definition of local applicability, and neither update rules nor locally applicable transformations allow for violations of quantum no-signalling (R4).

In line with with our discussion in Sec. 3.2.1, the authors of [182] notice that, in order to “extend” the action of a transformation to the larger state space of a composite system without assuming linearity, families of functions need to be introduced. In the case of AMTs, fixing \mathcal{H}_A , these are the collections $\{w_{AB}^A\}$ defined for any finite-dimensional extension \mathcal{H}_B , describing the effect of a measurement of ‘A’ on any composite system that includes it. Similarly, families of functions $\{\mathcal{L}_X\}$, defined for any finite-dimensional extension \mathcal{H}_X , are used to define locally applicable transformations.

The authors of [182] also introduce the concept of a *state-measurement theory*, i.e. a “bare” operational theory without notion of system composition and time evolution. In particular, their definition includes a measurement-induced state update mapping u_A . This map resembles the single-system update rule w_A^A featured in this document for the authors assume, albeit only implicitly, context-independence (R2). However, they impose an additional condition on “null measurements”, stating that when the observed outcome provides no information about the pre-measurement state, then the measurement cannot disturb the system. This requirement does

not appear in our Def. 9 of update rules. In fact, measurements in AMTs may disturb the state of a system regardless of the amount of information revealed. In other words, in constructing the AMT framework, we do not assume *a priori* that measuring a “trivial” observable is operationally equivalent to *not* measuring any observable. While the former might require the implementation of a physical device or involve post-processing, the latter, being the absence of measurement, is not formally captured by an update rule. If no measurement takes place, the update rule w^A is simply not applied. In Sec. 3.4.4, we showed that the information-disturbance trade-off (TO), which implies the described behaviour for null measurements, is a rather strong condition, able to uniquely identify the Lüders projection for single systems.

Considering these observations, we notice that a more general definition of a *spatial state-measurement theory* (i.e. a state-measurement theory with a notion of composite system) than the one presented in [182] can be provided by replacing the partial function u_A with a suitable generalisation of our notion of an update rule w^A that does not refer to specific axioms of quantum theory.

Gisin’s argument is also reviewed by the authors of [182]. They emphasise the need of additional assumptions to those of the original argument [84] in order to isolate the Schrödinger equation within the possible convex-linear, deterministic time evolutions compatible with the no-superluminal-signalling principle. In our work, as discussed in Sec. 2.7, we address a different aspect of Gisin’s argument. Specifically, we show that, contrary to the claims made in [163], the argument indeed relies on the projection postulate. Furthermore, in line with numerous examples appeared in the literature [45, 46, 71, 72, 101, 118–120, 151, 152], all nonlinear instruments featured in the AMT framework represent instances of nonlinear dynamical transformations that do not lead to signalling. Therefore, they are not dismissed by Gisin-type arguments.

— B —

Appendix to Part II

B.1 Proofs of Lemmata 4.1, 4.3, 4.4 and Corollary 4.2

We present proofs of Lemma 4.1, Corollary 4.2 and Lemmata 4.3 and 4.4, in this order.

Lemma 4.1. *Let d be an odd prime and $j, k \in \{1 \dots d\}$ with $j \neq k$. Then*

$$H_k^\dagger H_j = M(j, k) H_t^\dagger \quad (4.28)$$

for a monomial matrix $M(j, k)$ if and only if $t = 1 + \chi \in \{1, \dots, d\}$ where the integer χ satisfies $4(j - k)\chi = 1 \pmod{d}$.

Proof. Using Eqs. (4.14), we calculate the matrix elements of the product $H_k^\dagger H_j$, with $j, k \neq 0$ and $j \neq k$,

$$\left[H_k^\dagger H_j \right]_{\ell\ell'} = \langle \phi_\ell^k | \phi_{\ell'}^j \rangle = \frac{1}{\sqrt{d}} G_d(j - k, \ell - \ell'), \quad (B.1)$$

with the generalised Gauss sum [27]

$$G_d(j, \ell) = \frac{1}{\sqrt{d}} \sum_{x=0}^{d-1} \omega^{jx^2 + \ell x}. \quad (B.2)$$

Using $1 = \omega^{(\ell - \ell')^2 \chi} \omega^{-(\ell - \ell')^2 \chi}$ in (B.1) and letting χ be an integer satisfying $4(j - k)\chi \equiv 1 \pmod{d}$, we obtain a standard Gauss sum $G_d(j - k, 0)$ with known closed form.

Explicitly, for $j \neq k$, we obtain

$$\begin{aligned}
 G_d(j-k, \ell-\ell') &= \omega^{-(\ell-\ell')^2\chi} \frac{1}{\sqrt{d}} \sum_x \omega^{(j-k)x^2+(\ell-\ell')x} \omega^{(\ell-\ell')^2\chi} \\
 &= \omega^{-(\ell-\ell')^2\chi} \frac{1}{\sqrt{d}} \sum_x \omega^{(j-k)[x+\bar{2}(j-k)(\ell-\ell')]^2} \\
 &= \omega^{-(\ell-\ell')^2\chi} \left[\frac{1}{\sqrt{d}} \sum_x \omega^{(j-k)x^2} \right] = \omega^{-(\ell-\ell')^2\chi} G_d(j-k, 0) \\
 &= \omega^{-(\ell-\ell')^2\chi} \left(\frac{j-k}{d} \right) \varepsilon_d,
 \end{aligned} \tag{B.3}$$

where \bar{a} denotes the multiplicative inverse of a , $\bar{a}a \equiv 1 \pmod{d}$, while $\left(\frac{a}{b}\right)$ denotes the Jacobi symbol of the integers a and b , and

$$\varepsilon_d = \begin{cases} 1 & \text{if } d \equiv 1 \pmod{4}, \\ i & \text{if } d \equiv 3 \pmod{4}. \end{cases} \tag{B.4}$$

The sum $G_d(j-k, \ell-\ell')$ in (B.3) reduces to a phase factor as it should since the components of the matrix $H_k^\dagger H_j$ are given by the overlap of states stemming from different MU bases.

Combining (B.1) and (4.14), we now determine the elements of the matrix $V \equiv H_k^\dagger H_j H_t$ for arbitrary $t \neq 0$:

$$V_{\ell\ell'} = \sum_{\ell''=0}^{d-1} \langle \phi_\ell^k | \phi_{\ell''}^j \rangle \langle \ell'' | \phi_{\ell'}^t \rangle = \frac{1}{d} \sum_{\ell''=0}^{d-1} G_d(j-k, \ell-\ell'') \omega^{-\ell'\ell''+(t-1)\ell''^2}. \tag{B.5}$$

We can simplify this expression by substituting (B.3) into it, to find

$$\begin{aligned}
 V_{\ell\ell'} &= \frac{1}{d} \left(\frac{j-k}{d} \right) \varepsilon_d \sum_{\ell''=0}^{d-1} \omega^{-(\ell-\ell'')^2\chi} \omega^{-\ell'\ell''+(t-1)\ell''^2} \\
 &= \frac{1}{d} \left(\frac{j-k}{d} \right) \varepsilon_d \omega^{-\ell^2\chi} \sum_{\ell''=0}^{d-1} \omega^{(t-1-\chi)\ell''^2+(2\ell\chi-\ell')\ell''}.
 \end{aligned} \tag{B.6}$$

Letting $t = 1 + \chi$, we obtain sums over all d -th roots of one which vanish unless the exponents of ω vanish,

$$\sum_{\ell''=0}^{d-1} \omega^{(2\ell\chi-\ell')\ell''} = \begin{cases} d & \text{if } \ell' = 2\ell\chi \pmod{d}, \\ 0 & \text{otherwise.} \end{cases} \tag{B.7}$$

Thus, for this value of t , the matrix elements of V take the form

$$V_{\ell\ell'} = \begin{cases} \left(\frac{j-k}{d}\right) \varepsilon_d \omega^{-\ell^2\chi} & \text{if } \ell' = 2\ell\chi \pmod{d}, \\ 0 & \text{otherwise,} \end{cases} \tag{B.8}$$

so that the only non-zero elements of the matrix V are those with indices $(\ell, 2\ell\chi \bmod d)$. Each row ℓ has exactly one non-zero entry and the map $\ell \mapsto 2\ell\chi \bmod d$ constitutes a permutation of the elements of $\{0 \dots d-1\}$ since d is a prime number and $2\chi \not\equiv 0 \pmod d$. (Assume this was not the case, i.e. $2\chi x \bmod d = 2\chi y \bmod d$ for some $x, y \in \{0 \dots d-1\}$, $x \neq y$. Then $2\chi(x-y) = nd$ for some integer n which is never the case whenever d is prime and $\chi \not\equiv 0 \pmod d$.) As a consequence, each column will also display exactly one non-zero entry.

Therefore, the product V of three Hadamard matrices is equal to a monomial matrix $M(j, k)$ if $t = 1 + \chi$, i.e.

$$H_k^\dagger H_j = M(j, k) H_{1+\chi}^\dagger. \quad (\text{B.9})$$

We complete the proof by showing that the matrix $V = H_k^\dagger H_j H_t$ is *not* monomial for any other value of t . For $t \neq 1 + \chi$, the sum on the right-hand side of (B.6) represents another generalised Gauss sum so that

$$V_{\ell\ell'} = \frac{1}{\sqrt{d}} \binom{j-k}{d} \varepsilon_d \omega^{-\ell^2\chi} G_d(t-1-\chi, 2\ell\chi - \ell') \quad (\text{B.10})$$

with $G_d(t-1-\chi, 2\ell\chi - \ell') = \sqrt{d} \langle \phi_{2\ell\chi}^{1+\chi} | \phi_{\ell'}^t \rangle$. We can now substitute the expression (B.3) and obtain

$$V_{\ell\ell'} = \frac{1}{\sqrt{d}} \varepsilon_d^2 \binom{j-k}{d} \left(\frac{t-1-\chi}{d} \right) \omega^{-\ell^2\chi} \omega^{-(2\ell\chi - \ell')^2 \tilde{\chi}} \quad (\text{B.11})$$

where $\tilde{\chi} \in \{1 \dots d\}$ is an integer satisfying $4(t-1-\chi)\tilde{\chi} \equiv 1 \pmod d$. Hence, the matrix elements $V_{\ell\ell'}$ are all non-zero confirming that the matrix V is not monomial unless $t = 1 + \chi$.

An alternative, shorter proof of Lemma 4.1 can be given by representing the Hadamard matrices H_j as 2×2 matrices in $\mathbf{SL}(2, \mathbb{Z}/d\mathbb{Z})$ (cf. [139]). \square

Corollary 4.2. *Given a complete standard set of MU bases in the space \mathcal{H}_d of prime dimension d , any set of d (or fewer) vectors taken from any two MU bases are linearly independent.*

Proof. Construct a matrix M of order $d \times (d_1 + d_2)$ from any $(d_1 + d_2) \leq d$ column vectors—expressed in the computational basis—from the two MU bases \mathcal{B}_{j_1} and \mathcal{B}_{j_2} . Then left-multiply M by $H_{j_1}^\dagger$. Since $\langle x | H_{j_1}^\dagger | \phi_k^{j_1} \rangle = \langle x | k \rangle$, the first d_1 columns will be elements of the computational basis, while the remaining d_2 columns will be taken

from $H_{j_1}^\dagger H_{j_2}$ since $\langle x | H_{j_1}^\dagger | \phi_k^{j_2} \rangle = \langle x | H_{j_1}^\dagger H_{j_2} | k \rangle$. By swapping rows appropriately via a permutation operator P which does not change linear independence of column vectors, the top left square can be mapped to the d_1 -dimensional identity. For example, if we consider $d = 5$ and $d_1 = d_2 = 2$, we obtain a 5×4 matrix,

$$PH_{j_1}^\dagger M = \left(\begin{array}{cc|cc} 1 & 0 & * & * \\ 0 & 1 & * & * \\ \hline 0 & 0 & * & * \\ 0 & 0 & * & * \\ 0 & 0 & * & * \end{array} \right), \quad (\text{B.12})$$

where the asterisks refer to the elements of $H_{j_1}^\dagger H_{j_2}$.

The $(d_1 + d_2) \leq d$ vectors are linearly dependent only if M does *not* have full rank, i.e. $\text{rank}(M) < d_1 + d_2$. Since $PH_{j_1}^\dagger$ is unitary, it follows that $\text{rank}(PH_{j_1}^\dagger M) < d_1 + d_2$. Given the form of the matrix (B.12), the bottom-right part of $PH_{j_1}^\dagger M$ must contain a $(d_2 \times d_2)$ submatrix with vanishing determinant. However, this is prohibited by Corollary 4.1 which ensures for all prime numbers d that $H_{j_1}^\dagger H_{j_2}$ has non-vanishing minors if $j_1 \neq j_2$. Thus, all $(d_1 + d_2)$ column vectors of M must be linearly independent. \square

Lemma 4.3. *Let d be an odd prime and $\omega \equiv e^{\frac{2i\pi}{d}}$. Consider two states $|\phi_{k_1}^{j_1}\rangle, |\phi_{k_2}^{j_2}\rangle \in \mathcal{H}_d$ taken from different standard MU bases, $j_1, j_2 \neq 0$, and let $\{|x\rangle\}$ be the computational basis. Then there can be at most two values of $x \in \{0 \dots d-1\}$ such that*

$$\langle x | \phi_{k_1}^{j_1} \rangle = \omega^n \langle x | \phi_{k_2}^{j_2} \rangle \quad (4.46)$$

for the same value of $n \in \{0 \dots d-1\}$. If two different states are taken from the same basis, $j_1 = j_2$, then the equation has exactly one solution for each value of n .

Proof. By substituting (4.13) into (4.46) and taking the logarithm, one obtains $ax^2 + bx - n = 0 \pmod{d}$ where $a = j_1 - j_2$ and $b = k_2 - k_1$. This quadratic equation can have no more than two integer solutions. Thus, at most two components of the states $|\phi_{k_1}^{j_1}\rangle$ and $|\phi_{k_2}^{j_2}\rangle$ can be identical in the computational basis, up to multiplication by ω^n . If $j_1 = j_2$, then $a = 0$ and the equation is linear with a single solution for each value of n . \square

Lemma 4.4. *If the support of a state $\psi \in \mathcal{H}_5$ equals three in both the computational basis and another standard MU basis with label $j \neq 0$, i.e. $\|\psi\|_0 = \|\psi\|_j = 3$, then its support size in each of the remaining four bases equals five, $\|\psi\|_{j'} = 5$, with $j' \neq 0, j$.*

Proof. Four scalar products with the state ψ vanish,

$$\langle x_1 | \psi \rangle = \langle y_1 | \psi \rangle = \langle \phi_{x_2}^j | \psi \rangle = \langle \phi_{y_2}^j | \psi \rangle = 0, \quad (\text{B.13})$$

two for each of the bases. Hence, the zero distributions of the states ψ and $H_j^\dagger \psi$, are given by $\mathcal{Z}^0 = \{x_1, y_1\}$ and $\mathcal{Z}^j = \{x_2, y_2\}$ respectively, with four integer numbers $x_1, \dots, y_2 \in \{0 \dots 4\}$. Now suppose that there is a third MU basis $\mathcal{B}_{j'}$, different from both \mathcal{B}_0 and \mathcal{B}_j , in which the state ψ does *not* have full support. In other words, there is at least one vanishing scalar product, $\langle \phi_{x_3}^{j'} | \psi \rangle = 0$, say, where $x_3 \in \{0 \dots 4\}$. Expressing the components of the five vectors $|x_1\rangle, |y_1\rangle, |\phi_{x_2}^j\rangle, |\phi_{y_2}^j\rangle, |\phi_{x_3}^{j'}\rangle$ with respect to the computational basis and arranging them into a 5×5 matrix, we find, after permuting the rows and rephasing the last three vectors,

$$M = \frac{1}{\sqrt{5}} \begin{pmatrix} \sqrt{5} & 0 & * & * & * \\ 0 & \sqrt{5} & * & * & * \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & \omega^a & \omega^c & \omega^e \\ 0 & 0 & \omega^b & \omega^d & \omega^f \end{pmatrix}, \quad a, \dots, f \in \{0 \dots 4\}. \quad (\text{B.14})$$

Being elements of the Hadamard matrices H_j and $H_{j'}$, the entries of the last three columns are powers of ω , a fifth root of one. Corollary 4.2 ensures the linear independence of the first four vectors.

If the determinant of M does not vanish, $\det M \neq 0$, then the five column vectors forming it are linearly *independent*, thus spanning \mathcal{H}_5 . However, the only state being orthogonal to all of \mathcal{H}_5 is $\psi = 0$ which does not represent a quantum state. Thus, for an acceptable state ψ producing the given five vanishing expansion coefficients, the five vectors involved must be linearly *dependent*, i.e. $\det M = 0$. Consequently, the determinant of the bottom right 3×3 matrix of M must vanish,

$$\Delta \equiv \det \begin{pmatrix} 1 & 1 & 1 \\ \omega^a & \omega^c & \omega^e \\ \omega^b & \omega^d & \omega^f \end{pmatrix} = \omega^{a+d} + \omega^{e+b} + \omega^{c+f} - \omega^{c+b} - \omega^{e+d} - \omega^{a+f} = 0. \quad (\text{B.15})$$

Each of the six terms in this expression is a power of a fifth root ω of unity, hence non-zero. It is well known that the set $\{\omega^n \mid n = 0, \dots, 3\}$ is linearly independent over the rational numbers \mathbb{Q} . As a consequence, every non-zero complex number that is expressible as a linear combination (over \mathbb{Q}) of these roots of unity has a unique expression. Since $\omega^4 = -1 - \omega - \omega^2 - \omega^3$, it must follow that the only decomposition

of zero over \mathbb{Q} in terms of fifth roots of 1 is $0 = q(1 + \omega + \omega^2 + \omega^3 + \omega^4)$, with some rational number $q \in \mathbb{Q}$. In other words, for the sum to vanish all five roots must be multiplied to the same rational coefficient.

We distinguish two cases: either $q \neq 0$ or $q = 0$. Since (B.15) involves six terms with coefficients ± 1 , we conclude that the case of $q \neq 0$ cannot be realised: it is impossible to get all five roots to appear with the same non-zero coefficient. For example, let $(a + d) = (e + b) \pmod{5}$, then Eq. (B.15) reduces to

$$\Delta = 2\omega^{a+d} + \omega^{c+f} - \omega^{c+b} - \omega^{e+d} - \omega^{a+f} \quad (\text{B.16})$$

Since all roots must appear, the exponents in (B.16) are all different. However, the coefficients are not equal throughout and the sum cannot vanish. A similar argument holds for any other equality between exponents.

The case of $q = 0$ must therefore apply: the determinant Δ vanishes if and only if the six terms in Eq. (B.15) cancel each other in pairs, i.e. the powers of ω must occur an even number of times, and with an equal number of positive and negative coefficients. Hence, the first term in (B.15) is necessarily paired up with one of the powers with a negative coefficient leading. Three cases arise which we will consider separately.

(i) For the first and the fourth term to cancel, we must have $(a + d) = (c + b) \pmod{5}$, or

$$(a - b) = (c - d) \pmod{5}, \quad (\text{B.17})$$

relating the expansion coefficients of two vectors of the *same* basis, namely $|\phi_{x_2}^j\rangle$ and $|\phi_{y_2}^j\rangle$. Consequently, the third and fourth column vectors in the matrix M in (B.14) have (at least) two equal entries in identical positions, up to an irrelevant common phase factor. This would result in a vanishing 2×2 submatrix of H_j contradicting Corollary 4.1 (and Lemma 4.3). Thus the determinant Δ cannot vanish in this case.

(ii) For the first and the fifth term to cancel, we must have $(a + d) = (e + d) \pmod{5}$, or

$$a = e \pmod{5}, \quad (\text{B.18})$$

relating the expansion coefficients of two vectors of *different* bases, namely $|\phi_{x_2}^j\rangle$ and $|\phi_{x_3}^{j'}\rangle$. Corollary 4.1 does not apply to this case. We do know, however, that the fourth term in the sum (B.15) must pair up with either the second or the third term of the sum in (B.15). In the first case, we find $(c + b) = (e + b) \pmod{5}$, or

$$c = e \pmod{5}. \quad (\text{B.19})$$

Given the constraint (B.18), we we obtain the identity

$$a = c \pmod{5}, \quad (\text{B.20})$$

again relating the expansion coefficients of two vectors of the *same* basis, namely $|\phi_{x_2}^j\rangle$ and $|\phi_{y_2}^j\rangle$. As in the Case (i), a contradiction to Corollary 4.1 arises.

In the second case, we pair up terms three and four of the sum (B.15), leading to the identity $(c + b) = (c + f) \pmod{5}$, or

$$b = f \pmod{5}. \quad (\text{B.21})$$

Together with Eq. (B.18), it follows that the last three elements of the third and fifth columns of M are identical. However, according to Lemma 4.3, two vectors stemming from two different bases MU to the computational basis can have at most two identical components.

(iii): Assuming that the first and the sixth term of the sum (B.15) cancel again leads to a contradiction along the lines of the argument considered in Case (ii).

Thus, we are forced to conclude that the determinant Δ cannot not vanish for any $j' \neq 0, j$ and any x_3 , which implies that the state ψ must have full support. \square

References

1. Aaronson, S. Quantum computing, postselection, and probabilistic polynomial-time. *Proc. R. Soc. A* **461**, 3473–3482 (2005).
2. Aaronson, S. PDQP/qpoly= ALL. *arXiv*: [quant-ph/1805.08577](https://arxiv.org/abs/1805.08577) (2018).
3. Aaronson, S., Bouland, A., Fitzsimons, J. & Lee, M. The space "just above" BQP. *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, 271–280 (2016).
4. Aaronson, S., Grewal, S., Iyer, V., Marshall, S. & Ramachandran, R. PDQMA=DQMA= NEXP: QMA With Hidden Variables and Non-collapsing Measurements. *arXiv*: [quant-ph/2403.02543](https://arxiv.org/abs/2403.02543) (2024).
5. Abrams, D. S. & Lloyd, S. Nonlinear Quantum Mechanics Implies Polynomial-Time Solution for NP-Complete and #P Problems. *Phys. Rev. Lett.* **81**, 3992–3995 (1998).
6. Ando, T. & Choi, M.-D. Non-Linear Completely Positive Maps. *North-Holland Math. Stud.* **122**, 3–13 (1986).
7. Arvidsson-Shukur, D. R. M., Drori, J. C. & Halpern, N. Y. Conditions tighter than noncommutation needed for nonclassicality. *J. Phys. A: Math. Theor.* **54**, 284001 (2021).
8. Arvidsson-Shukur, D. R. M., Yunger Halpern, N., Lepage, H. V., Lasek, A. A., Barnes, C. H. W. & Lloyd, S. Quantum advantage in postselected metrology. *Nat. Commun.* **11**, 3775 (2020).
9. Bacciagaluppi, G. & Crull, E. Heisenberg (and Schrödinger, and Pauli) on hidden variables. *Stud. Hist. Philos. Sci. B* **40**, 374–382 (2009).

10. Ballester, M. A. & Wehner, S. Entropic uncertainty relations and locking: Tight bounds for mutually unbiased bases. *Phys. Rev. A* **75**, 022319 (2007).
11. Balmer, J. J. Notiz über die Spectrallinien des Wasserstoffs. *Ann. Phys.* **261**, 80–87 (1885).
12. Bandyopadhyay, S., Boykin, P., Roychowdhury, V. & Vatan, F. A New Proof for the Existence of Mutually Unbiased Bases. *Algorithmica* **34**, 512–528 (2002).
13. Baracca, A., Bergia, S., Bigoni, R. & Cecchini, A. Statistics of observations for «Proper» and « Improper » mixtures in quantum mechanics. *Riv. Nuovo Cim.* **4**, 169–188 (1974).
14. Barnum, H., Barrett, J., Leifer, M. & Wilce, A. Cloning and Broadcasting in Generic Probabilistic Theories. *arXiv: [quant-ph/0611295](https://arxiv.org/abs/quant-ph/0611295)* (2006).
15. Barnum, H., Barrett, J., Leifer, M. & Wilce, A. Generalized No-Broadcasting Theorem. *Phys. Rev. Lett.* **99**, 240501 (2007).
16. Barrett, J. Information processing in generalized probabilistic theories. *Phys. Rev. A* **75**, 032304 (2007).
17. Bassi, A. & Ghirardi, G. Dynamical reduction models. *Phys. Rep.* **379**, 257–426 (2003).
18. Bassi, A. & Hejazi, K. No-faster-than-light-signaling implies linear evolution. A re-derivation. *Eur. J. Phys.* **36**, 055027 (2015).
19. Bassi, A., Lochan, K., Satin, S., Singh, T. P. & Ulbricht, H. Models of wavefunction collapse, underlying theories, and experimental tests. *Rev. Mod. Phys.* **85**, 471–527 (2013).
20. Bell, J. S. On the Einstein Podolsky Rosen paradox. *Physics* **1**, 195–200 (1964).
21. Bell, J. S. On the Problem of Hidden Variables in Quantum Mechanics. *Rev. Mod. Phys.* **38**, 447–452 (1966).
22. Bell, J. S. & Nauenberg, M. The moral aspect of quantum mechanics. In *Preludes in theoretical physics*, 278–286 (1966).
23. Bender, C. M., Boettcher, S. & Meisinger, P. N. \mathcal{PT} -symmetric quantum mechanics. *J. Math. Phys.* **40**, 2201–2229 (1999).

24. Bennett, C. H. & Brassard, G. Quantum cryptography: public key distribution and coin tossing. *Theor. Comput. Sci.* **560**, 7–11 (2014).
25. Bennett, C. H., Brassard, G., Crépeau, C., Jozsa, R., Peres, A. & Wootters, W. K. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Phys. Rev. Lett.* **70**, 1895–1899 (1993).
26. Beretta, G. P., Gyftopoulos, E. P., Park, J. L. & Hatsopoulos, G. N. Quantum thermodynamics. A New Equation of Motion for a Single Constituent of Matter. *Nuovo Cim. B* **82**, 169–191 (1984).
27. Berndt, B. C., Williams, K. S. & Evans, R. J. *Gauss and Jacobi sums* (Wiley, 1998).
28. Birkhoff, G. & Von Neumann, J. The Logic of Quantum Mechanics. *Ann. Math.* **37**, 823–843 (1936).
29. Bohm, D. A Suggested Interpretation of the Quantum Theory in Terms of "Hidden" Variables. I. *Phys. Rev.* **85**, 166–179 (1952).
30. Born, M., Heisenberg, W. & Jordan, P. Zur Quantenmechanik II. *Z. Phys.* **35**, 557–615 (1926).
31. Born, M. & Jordan, P. Zur quantenmechanik. *Z. Phys.* **34**, 858–888 (1925).
32. Bruzda, W., Tadej, W. & Życzkowski, K. *Catalogue of Complex Hadamard Matrices* <https://chaos.if.uj.edu.pl/~karol/hadamard/>.
33. Busch, P. Is the Quantum State (an) Observable? *Potentiality, Entanglement and Passion-at-a-Distance: Quantum Mechanical Studies for Abner Shimony Volume Two*, 61–70 (1997).
34. Busch, P. Quantum States and Generalized Observables: A Simple Proof of Gleason's Theorem. *Phys. Rev. Lett.* **91**, 120403 (2003).
35. Busch, P. Informationally complete sets of physical quantities. *Int. J. Theor. Phys.* **30**, 1217–1227 (1991).
36. Busch, P., Lahti, P. & Werner, R. F. Measurement uncertainty relations. *J. Math. Phys.* **55**, 042111 (2014).
37. Busch, P., Lahti, P. J. & Mittelstaedt, P. *The Quantum Theory of Measurement* (Springer-Verlag, 1996).

38. Candes, E., Romberg, J. & Tao, T. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theory*. **52**, 489–509 (2006).
39. Candes, E. J. & Romberg, J. Quantitative Robust Uncertainty Principles and Optimally Sparse Decompositions. *Found. Comput. Math.* **6**, 227–254 (2006).
40. Catani, L., Leifer, M., Schmid, D. & Spekkens, R. W. Why interference phenomena do not capture the essence of quantum theory. *arXiv: quant-ph/2111.13727* (2022).
41. Caves, C. M., Fuchs, C. A., Manne, K. K. & Renes, J. M. Gleason-Type Derivations of the Quantum Probability Rule for Generalized Measurements. *Found. Phys.* **34**, 193–209 (2004).
42. Chiribella, G., D’Ariano, G. M. & Perinotti, P. Informational derivation of quantum theory. *Phys. Rev. A* **84**, 012311 (2011).
43. Clements, K., Dowker, F. & Wallden, P. Physical Logic. *The Incomputable: Journeys Beyond the Turing Barrier*, 47–61 (2017).
44. Compton, A. H. & Simon, A. W. Directed Quanta of Scattered X-Rays. *Phys. Rev.* **26**, 289–299 (1925).
45. Czachor, M. Nonlocal-looking equations can make nonlinear quantum dynamics local. *Phys. Rev. A* **57**, 4122 (1998).
46. Czachor, M. & Doebner, H.-D. Correlation experiments in nonlinear quantum mechanics. *Phys. Lett. A* **301**, 139–152 (2002).
47. Czachor, M. & Kuna, M. Complete positivity of nonlinear evolution: A case study. *Phys. Rev. A* **58**, 128–134 (1998).
48. d’Espagnat, B. Towards a separable “empirical reality”? *Found. Phys.* **20**, 1147–1172 (1990).
49. Dakic, B. & Brukner, C. Quantum Theory and Beyond: Is Entanglement Special? *arXiv: quant-ph/0911.0695* (2009).
50. Davies, E. B. *Quantum theory of open systems* (Academic Press, 1976).
51. De Bièvre, S. Relating incompatibility, noncommutativity, uncertainty, and Kirkwood–Dirac nonclassicality. *J. Math. Phys.* **64** (2023).

52. De Bièvre, S. Complete Incompatibility, Support Uncertainty, and Kirkwood-Dirac Nonclassicality. *Phys. Rev. Lett.* **127**, 190404 (2021).
53. De Broglie, L. *Recherches sur la théorie des quanta* PhD thesis (Migration-université en cours d'affectation, 1924).
54. Deffner, S. & Campbell, S. Quantum speed limits: from Heisenberg's uncertainty principle to optimal quantum control. *J. Phys. A: Math. Theor.* **50**, 453001 (2017).
55. Deutsch, D. Uncertainty in Quantum Measurements. *Phys. Rev. Lett.* **50**, 631–633 (1983).
56. Deutsch, D. & Jozsa, R. Rapid solution of problems by quantum computation. *Proc. Math. Phys. Eng. Sci.* **439**, 553–558 (1992).
57. Dewitt, B. S. & Graham, N. *The Many-Worlds Interpretation of Quantum Mechanics* (Princeton University Press, 2015).
58. Dieks, D. Communication by EPR devices. *Phys. Lett. A* **92**, 271–272 (1982).
59. Dieks, D. Overlap and distinguishability of quantum states. *Phys. Lett. A* **126**, 303–306 (1988).
60. Dieks, D. & Vermaas, P. E. *The Modal Interpretation of Quantum Mechanics* (Springer Science & Business Media, 1998).
61. Dirac, P. A. M. On the Analogy Between Classical and Quantum Mechanics. *Rev. Mod. Phys.* **17**, 195–199 (1945).
62. Dodonov, V. V., Kurmyshev, E. V. & Man'ko, V. I. Generalized uncertainty relation and correlated coherent states. *Phys. Lett. A* **79**, 150–152 (1980).
63. Dodonov, V. V. Variance uncertainty relations without covariances for three and four observables. *Phys. Rev. A* **97**, 022105 (2018).
64. Donoho, D. L. & Stark, P. B. Uncertainty Principles and Signal Recovery. *SIAM J. Appl. Math.* **49**, 906–931 (1989).
65. Dowker, F. & Ghazi-Tabatabai, Y. The Kochen–Specker theorem revisited in quantum measure theory. *J. Phys. A: Math. Theor.* **41**, 105301 (2008).
66. Dressel, J. & Jordan, A. N. Significance of the imaginary part of the weak value. *Phys. Rev. A* **85**, 012107 (2012).

67. Durt, T. About mutually unbiased bases in even and odd prime power dimensions. *J. Phys. A: Math. Gen.* **38**, 5267–5283 (2005).
68. Ehrenfest, P. Welche Züge der Lichtquantenhypothese spielen in der Theorie der Wärmestrahlung eine wesentliche Rolle? *Ann. Phys.* **341**, 91–118 (1911).
69. Einstein, A. Zur theorie der lichterzeugung und lichtabsorption. *Ann. Phys.* **325**, 199–206 (1906).
70. Ekert, A. K. Quantum cryptography based on Bell’s theorem. *Phys. Rev. Lett.* **67**, 661 (1991).
71. Ferrero, M., Salgado, D. & Sánchez-Gómez, J. L. Nonlinear quantum evolution does not imply supraluminal communication. *Phys. Rev. A* **70**, 014101 (2004).
72. Ferrero, M., Salgado, D. & Sánchez-Gómez, J. L. On nonlinear evolution and supraluminal communication between finite quantum systems. *Int. J. Quantum Inf.* **3**, 257–261 (2005).
73. Feynman, R. P. Space-Time Approach to Non-Relativistic Quantum Mechanics. *Rev. Mod. Phys.* **20**, 367–387 (1948).
74. Finkelstein, D., Jauch, J. M., Schiminovich, S. & Speiser, D. Foundations of Quaternion Quantum Mechanics. *J. Math. Phys.* **3**, 207–220 (1962).
75. Fiorentino, V. & Weigert, S. A quantum theory with non-collapsing measurements. *arXiv: [quant-ph/2303.13411](https://arxiv.org/abs/quant-ph/2303.13411)* (2023).
76. Flatt, K., Barnett, S. & Croke, S. Gleason-Busch theorem for sequential measurements. *Phys. Rev. A* **96**, 062125 (2017).
77. Folland, G. B. & Sitaram, A. The uncertainty principle: A mathematical survey. *J. Fourier Anal. Appl.* **3**, 207–238 (1997).
78. Frenkel, P. E. Simple proof of Chebotarev’s theorem on roots of unity. *arXiv: [math/0312398](https://arxiv.org/abs/math/0312398)* (2004).
79. Galley, T. D. & Masanes, L. Any modification of the Born rule leads to a violation of the purification and local tomography principles. *Quantum* **2**, 104 (2018).
80. Galley, T. D. & Masanes, L. Classification of all alternatives to the Born rule in terms of informational properties. *Quantum* **1**, 15 (2017).

81. Galley, T. D., Masanes, L. & Müller, M. P. Reply to "Masanes-Galley-Müller and the State-Update Postulate". *arXiv: [quant-ph/2212.03629](https://arxiv.org/abs/quant-ph/2212.03629)* (2022).
82. Galvan, B. Generalization of the Born rule. *Phys. Rev. A* **78**, 042113 (2008).
83. Ghirardi, G. C., Rimini, A. & Weber, T. Unified dynamics for microscopic and macroscopic systems. *Phys. Rev. D* **34**, 470–491 (1986).
84. Gisin, N. Stochastic quantum dynamics and relativity. *Helv. Phys. Acta* **62**, 363–371 (1989).
85. Gisin, N. Weinberg's non-linear quantum mechanics and supraluminal communications. *Phys. Lett. A* **143**, 1–2 (1990).
86. Gleason, A. M. Measures on the Closed Subspaces of a Hilbert Space. *J. Math. Mech.* **6**, 885–893 (1957).
87. Greenberger, D. M., Horne, M. A. & Zeilinger, A. Going Beyond Bell's Theorem. *Bell's Theorem, Quantum Theory and Conceptions of the Universe*, 69–72 (1989).
88. Groisman, B. & Reznik, B. Measurements of semilocal and nonmaximally entangled states. *Phys. Rev. A* **66**, 022110 (2002).
89. Gross, D., Müller, M., Colbeck, R. & Dahlsten, O. C. O. All Reversible Dynamics in Maximally Nonlocal Theories are Trivial. *Phys. Rev. Lett.* **104**, 080402 (2010).
90. Grover, L. K. A fast quantum mechanical algorithm for database search. *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, 212–219 (1996).
91. Hardy, L. Quantum theory from five reasonable axioms. *arXiv: [quant-ph/0101012](https://arxiv.org/abs/quant-ph/0101012)* (2001).
92. Harrigan, N. & Spekkens, R. Einstein, incompleteness, and the epistemic view of quantum states. *Found. Phys.* **40**, 125–157 (2010).
93. Hausmann, L., Nurgalieva, N. & del Rio, L. A consolidating review of Spekkens' toy theory. *arXiv: [quant-ph/2105.03277](https://arxiv.org/abs/quant-ph/2105.03277)* (2021).
94. Hayashi, M. *Quantum Information Theory: Mathematical Foundation* (Springer Berlin Heidelberg, 2017).

95. Heinosaari, T., Miyadera, T. & Ziman, M. An invitation to quantum incompatibility. *J. Phys. A: Math. Theor.* **49**, 123001 (2016).
96. Heinosaari, T. & Wolf, M. M. Nondisturbing quantum measurements. *J. Math. Phys.* **51**, 092201 (2010).
97. Heinosaari, T. & Ziman, M. *The Mathematical Language of Quantum Theory: From Uncertainty to Entanglement* (Cambridge University Press, 2011).
98. Heisenberg, W. Über quantentheoretische Umdeutung kinematischer und mechanischer Beziehungen. *Z. Phys.* **33**, 879–893 (1925).
99. Heisenberg, W. Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik. *Z. Phys.* **43**, 172–198 (1927).
100. Hellmann, F., Kamiński, W. & Kostecki, R. P. Quantum collapse rules from the maximum relative entropy principle. *New J. Phys.* **18**, 013022 (2016).
101. Helou, B. & Chen, Y. Extensions of Born’s rule to non-linear quantum mechanics, some of which do not imply superluminal communication. *J. Phys.: Conf. Ser.* **880**, 012021 (2017).
102. Herbut, F. Derivation of the change of state in measurement from the concept of minimal measurement. *Ann. Phys.* **55**, 271–300 (1969).
103. Hertz, H. Ueber einen Einfluss des ultravioletten Lichtes auf die elektrische Entladung. *Ann. Phys.* **267**, 983–1000 (1887).
104. Hofmann, H. F. On the role of complex phases in the quantum statistics of weak measurements. *New J. Phys.* **13**, 103009 (2011).
105. Holman, M. On Arguments for Linear Quantum Dynamics. *arXiv: quant-ph/0612209* (2006).
106. Ivanovic, I. D. How to differentiate between non-orthogonal states. *Phys. Lett. A* **123**, 257–259 (1987).
107. Ivanovic, I. D. An inequality for the sum of entropies of unbiased quantum measurements. *J. Phys. A: Math. Gen.* **25**, L363–L364 (1992).
108. Ivanović, I. D. Geometrical description of quantal state determination. *J. Phys. A: Math. Gen.* **14**, 3241–3245 (1981).
109. Jaeger, G. & Shimony, A. Optimal distinction between two non-orthogonal quantum states. *Phys. Lett. A* **197**, 83–87 (1995).

110. Janotta, P. & Lal, R. Generalized probabilistic theories without the no-restriction hypothesis. *Phys. Rev. A* **87**, 052131 (2013).
111. Jarrett, J. P. On the Physical Significance of the Locality Conditions in the Bell Arguments. *Noûs* **18**, 569–589 (1984).
112. Kaye, P., Laflamme, R. & Mosca, M. *An Introduction to Quantum Computing* (OUP Oxford, 2006).
113. Kechrimparis, S. & Weigert, S. Heisenberg uncertainty relation for three canonical observables. *Phys. Rev. A* **90**, 062118 (2014).
114. Kechrimparis, S. & Weigert, S. Geometry of uncertainty relations for linear combinations of position and momentum. *J. Phys. A: Math. Theor.* **51**, 025303 (2017).
115. Kennard, E. H. Zur Quantenmechanik einfacher Bewegungstypen. *Z. Phys.* **44**, 326–352 (1927).
116. Kent, A. Causal quantum theory and the collapse locality loophole. *Phys. Rev. A* **72**, 012107 (2005).
117. Kent, A. Testing causal quantum theory. *Proc. R. Soc. A* **474**, 20180501 (2018).
118. Kent, A. The measurement postulates of quantum mechanics are not redundant. *arXiv: [quant-ph/2307.06191](https://arxiv.org/abs/quant-ph/2307.06191)* (2023).
119. Kent, A. Nonlinearity without superluminality. *Phys. Rev. A* **72**, 012108 (2005).
120. Kent, A. Quantum state readout, collapses, probes, and signals. *Phys. Rev. D* **103**, 064061 (2021).
121. Kirkwood, J. G. Quantum Statistics of Almost Classical Assemblies. *Phys. Rev.* **44**, 31–37 (1933).
122. Kochen, S. & Specker, E. P. The Problem of Hidden Variables in Quantum Mechanics. *J. Math. Mech.* **17**, 59–87 (1967).
123. Kunjwal, R., Lostaglio, M. & Pusey, M. F. Anomalous weak values and contextuality: Robustness, tightness, and imaginary parts. *Phys. Rev. A* **100**, 042116 (2019).

124. Larsson, J.-Å. A contextual extension of Spekkens' toy model. *AIP Conf. Proc.* **1424**, 211–220 (2012).
125. Leifer, M. S. Is the quantum state real? An extended review of ψ -ontology theorems. *Quanta* **3**, 67–155 (2014).
126. Lenard, P. Ueber die lichtelektrische Wirkung. *Ann. Phys.* **313**, 149–198 (1902).
127. Lüders, G. Concerning the state-change due to the measurement process. *Ann. Phys.* **15**, 663–670 (2006).
128. Lüders, G. Über die Zustandsänderung durch den Meßprozeß. *Ann. Phys.* **443**, 322–328 (1950).
129. Ludwig, G. *Foundations of Quantum Mechanics I* (Springer-Verlag, 1983).
130. Maassen, H. & Uffink, J. B. M. Generalized entropic uncertainty relations. *Phys. Rev. Lett.* **60**, 1103–1106 (1988).
131. Mackey, G. W. *Mathematical Foundations of Quantum Mechanics* (Addison-Wesley, 1963).
132. Masanes, L., Galley, T. D. & Müller, M. P. The measurement postulates of quantum mechanics are operationally redundant. *Nat. Commun.* **10**, 1361 (2019).
133. Masanes, L. & Müller, M. P. A derivation of quantum theory from physical requirements. *New J. Phys.* **13**, 063001 (2011).
134. Mermin, N. D. Simple unified form for the major no-hidden-variables theorems. *Phys. Rev. Lett.* **65**, 3373–3376 (1990).
135. Mermin, N. D. Hidden variables and the two theorems of John Bell. *Rev. Mod. Phys.* **65**, 803–815 (1993).
136. Meshulam, R. An uncertainty inequality for groups of order pq . *Eur. J. Comb.* **13**, 401–407 (1992).
137. Miloschewsky, D. & Podder, S. Revisiting BQP with Non-Collapsing Measurements. *arXiv: [quant-ph/2411.04085](https://arxiv.org/abs/quant-ph/2411.04085)* (2024).
138. Naimark, M. About second-kind self-adjoint extensions of symmetrical operator. *Izv. Akad. Nauk USSR, Ser. Mat* **4**, 53–104 (1940).

139. Neuhauser, M. An explicit construction of the metaplectic representation over a finite field. *J. Lie Theory* **12**, 15–30 (2002).
140. Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, 2010).
141. Oppenheim, J. & Wehner, S. The Uncertainty Principle Determines the Nonlocality of Quantum Mechanics. *Science* **330**, 1072–1074 (2010).
142. Ozawa, M. Quantum measuring processes of continuous observables. *J. Math. Phys.* **25**, 79–87 (1984).
143. Pan, J.-W., Bouwmeester, D., Weinfurter, H. & Zeilinger, A. Experimental Entanglement Swapping: Entangling Photons That Never Interacted. *Phys. Rev. Lett.* **80**, 3891–3894 (1998).
144. Paulsen, V. *Completely Bounded Maps and Operator Algebras* (Cambridge University Press, 2002).
145. Peres, A. How to differentiate between non-orthogonal states. *Phys. Lett. A* **128**, 19 (1988).
146. Planck, M. Über irreversible Strahlungsvorgänge. *Ann. Phys.* **306** (1900).
147. Polchinski, J. Weinberg’s nonlinear quantum mechanics and the Einstein-Podolsky-Rosen paradox. *Phys. Rev. Lett.* **66**, 397–400 (1991).
148. Popescu, S. & Rohrlich, D. Quantum nonlocality as an axiom. *Found. Phys.* **24**, 379–385 (1994).
149. Preskill, J. *Course Information for Physics 219/Computer Science 219 - Quantum Computation* (2021).
150. Prugovečki, E. Information-theoretical aspects of quantum measurement. *Int. J. Theor. Phys.* **16**, 321–331 (1977).
151. Ray, R. K. & Beretta, G. P. No-signaling in Nonlinear Extensions of Quantum Mechanics. *arXiv: [quant-ph/2301.11548](https://arxiv.org/abs/quant-ph/2301.11548)* (2023).
152. Rembieliński, J. & Caban, P. Nonlinear evolution and signaling. *Phys. Rev. Research.* **2**, 012027 (2020).
153. Renou, M.-O., Trillo, D., Weilenmann, M., Le, T. P., Tavakoli, A., Gisin, N., Acín, A. & Navascués, M. Quantum theory based on real numbers can be experimentally falsified. *Nature* **600**, 625–629 (2021).

154. Roberts, B. W. Observables, disassembled. *Stud. Hist. Philos. Sci. B - Stud. Hist. Philos. Mod. Phys.* **63**, 150–162 (2018).
155. Rydberg, J. R. Recherches sur la constitution des spectres d'émission des éléments chimiques. *K. Sven. vetensk.akad. handl.* **23** (1890).
156. Sánchez, J. Entropic uncertainty and certainty relations for complementary observables. *Phys. Lett. A* **173**, 233–239 (1993).
157. Schlosshauer, M. Quantum decoherence. *Phys. Rep.* **831**, 1–57 (2019).
158. Schrödinger, E. Quantisierung als Eigenwertproblem. *Ann. Phys.* **385** (1926).
159. Schrödinger, E. Zum Heisenbergschen Unschärfeprinzip. *Sitzungsber. Preuss. Akad. Wiss. (Phys.-Math. Klasse)* **19**, 296–303 (1930).
160. Shimony, A. Events and Processes in the Quantum World. *Quantum Concepts in Space and Time*, 182–203 (1986).
161. Shor, P. Algorithms for quantum computation: discrete logarithms and factoring. *Proceedings 35th Annual Symposium on Foundations of Computer Science*, 124–134 (1994).
162. Shrapnel, S., Costa, F. & Milburn, G. Updating the Born rule. *New J. Phys.* **20**, 053010 (2018).
163. Simon, C., Bužek, V. & Gisin, N. No-Signaling Condition and Quantum Dynamics. *Phys. Rev. Lett.* **87**, 170405 (2001).
164. Simon, D. R. On the Power of Quantum Computation. *SIAM J. Comput.* **26**, 1474–1483 (1997).
165. Sorkin, R. D. Quantum mechanics as quantum measure theory. *Mod. Phys. Lett. A* **9**, 3119–3127 (1994).
166. Spekkens, R. W. Contextuality for preparations, transformations, and unsharp measurements. *Phys. Rev. A* **71**, 052108 (2005).
167. Spekkens, R. W. Evidence for the epistemic view of quantum states: A toy theory. *Phys. Rev. A* **75**, 032110 (2007).
168. Stacey, B. C. Masanes-Galley-Müller and the State-Update Postulate. *arXiv: quant-ph/2211.03299* (2022).
169. Steinhagen, P. & Lenstra, H. W. Chebotarëv and his density theorem. *Math. Intell.* **18**, 26–37 (1996).

170. Stevens, N. & Busch, P. Steering, incompatibility, and Bell-inequality violations in a class of probabilistic theories. *Phys. Rev. A* **89**, 022123 (2014).
171. Stinespring, W. F. Positive Functions on C*-Algebras. *Proc. Am. Math. Soc.* **6**, 211–216 (1955).
172. Stueckelberg, E. C. Quantum theory in real Hilbert space. *Helv. Phys. Acta* **33**, 458 (1960).
173. Tao, T. An uncertainty principle for cyclic groups of prime order. *Math. Res. Lett.* **12**, 121–127 (2005).
174. Timpson, C. G. & Brown, H. R. Proper and improper separability. *Int. J. Quantum Inf.* **3**, 679–690 (2005).
175. Vaidman, L. Instantaneous measurement of nonlocal variables. *Phys. Rev. Lett.* **90**, 010402 (2003).
176. Valdenebro, A. G. Assumptions underlying Bell’s inequalities. *Eur. J. Phys.* **23**, 569–577 (2002).
177. von Neumann, J. *Mathematische Grundlagen der Quantenmechanik* (Springer Berlin, 1932).
178. Wehner, S. & Winter, A. Entropic uncertainty relations—a survey. *New J. Phys.* **12**, 025009 (2010).
179. Weinberg, S. Testing quantum mechanics. *Ann. Phys.* **194**, 336–386 (1989).
180. Wigderson, A. & Wigderson, Y. The uncertainty principle: Variations on a theme. *Bull. Amer. Math. Soc.* **58**, 225–261 (2021).
181. Wilce, A. Tensor products in generalized measure theory. *Int. J. Theor. Phys.* **31**, 1915–1928 (1992).
182. Wilson, M. & Ormrod, N. On the Origin of Linearity and Unitarity in Quantum Theory. *arXiv: [quant-ph/2305.20063](https://arxiv.org/abs/quant-ph/2305.20063)* (2023).
183. Wootters, W. K. & Fields, B. D. Optimal state-determination by mutually unbiased measurements. *Ann. Phys.* **191**, 363–381 (1989).
184. Wootters, W. K. & Zurek, W. H. A single quantum cannot be cloned. *Nature* **299**, 802–803 (1982).
185. Wright, V. J. *Gleason-type theorems and general probabilistic theories* PhD thesis (University of York, 2019).

186. Wright, V. J. & Weigert, S. A Gleason-type theorem for qubits based on mixtures of projective measurements. *J. Phys. A: Math. Theor.* **52**, 055301 (2019).
187. Wu, S., Yu, S. & Mølmer, K. Entropic uncertainty relation for mutually unbiased bases. *Phys. Rev. A* **79**, 022104 (2009).
188. Yalabik, M. C. Nonlinear Schrödinger equation for quantum computation. *Mod. Phys. Lett. B* **20**, 1099–1106 (2006).
189. Zalka, C. Grover's quantum searching algorithm is optimal. *Phys. Rev. A* **60**, 2746–2751 (1999).